Learning Verified Monitors for Hidden Markov Models

Luko van der Maas, Sebastian Junges

October 30, 2025

Radboud Universiteit


150
180
120
210
90
240
km/h
60
270
30
300
0
330
1/2
0
1
4
3
5
2
6
1/min x 1000
1
7
0
8
120
°C
70
170


150
180
120
210
90
km/h
240
60
270
30
300
0
330
1/2
0
1
4
3
5
2
1/min x 1000
6
1
7
0
8
120
°C
70
170

When should the light turn on?

Learning a monitor

- Monitor is a classifier

Learning a monitor

- Monitor is a classifier
- Classic learning approaches do not guarantee safety

Learning a monitor

- Monitor is a classifier
- Classic learning approaches do not guarantee safety
- We introduce a model

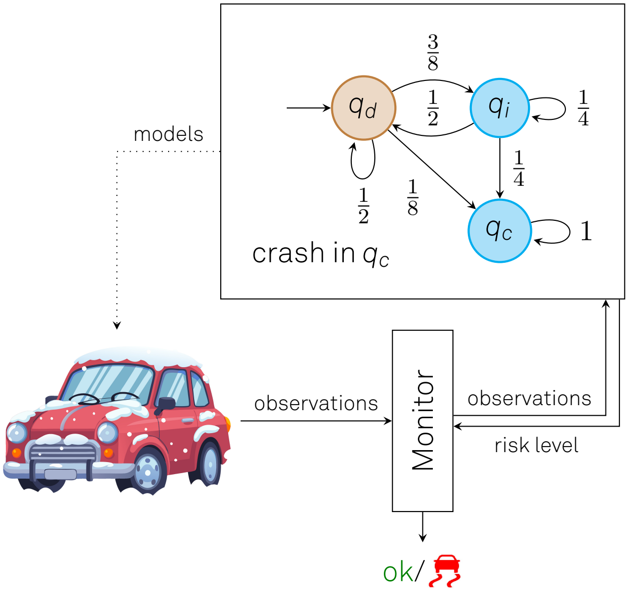

Learning a monitor

- Monitor is a classifier
- Classic learning approaches do not guarantee safety
- We introduce a model
- Verify the monitor is safe on the model

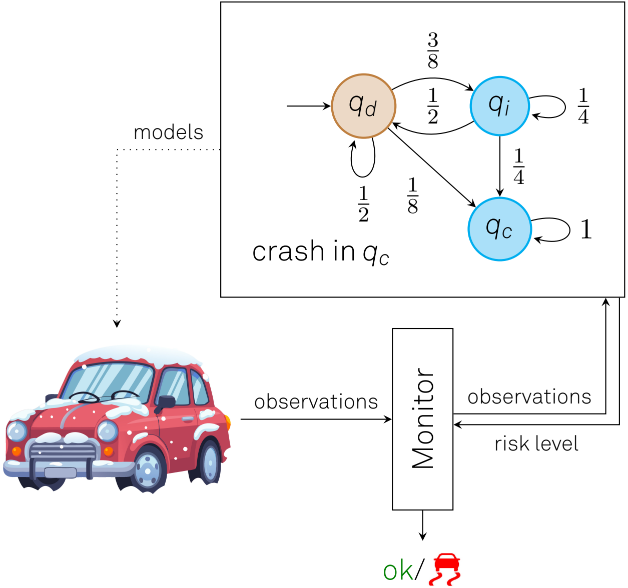

Overview

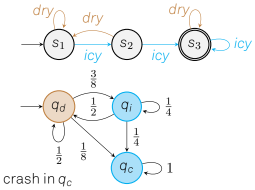

Problem statement

Verification approach

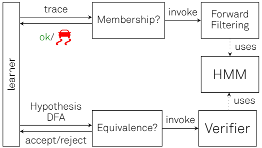

Learning approach

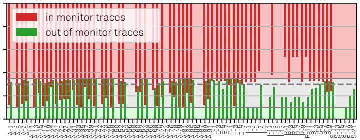

Correctness results

What is our System Model?

Definition (Hidden Markov Model)

- States: S
- Transition function: $\mathbf{P} : S \to \Delta S$
- Observations: Z
- Observation function: $obs : S \to Z$

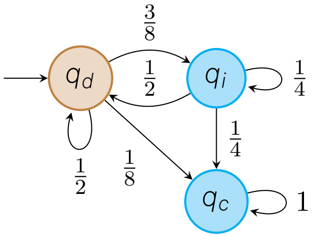

$S = \{q_d, q_i, q_c\}$

$Z = \{dry : \bigcirc, icy : \bigcirc\}$

Monitoring Observations

Question

Having observed the observation sequence τ, what is the probability of being in q_c?

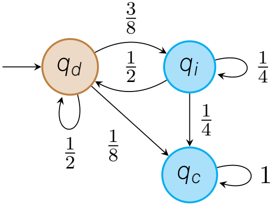

Monitoring Observations

Question

Having observed the observation sequence τ, what is the probability of being in q_c?

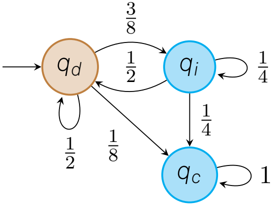

$\tau_1 =$ 

Monitoring Observations

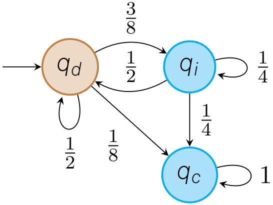

Question

Having observed the observation sequence τ, what is the probability of being in q_c?

$\tau_1 =$  0

Monitoring Observations

Question

Having observed the observation sequence τ, what is the probability of being in q_c?

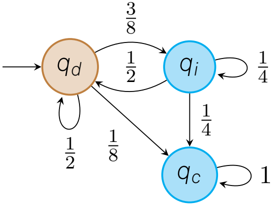

$\tau_1 =$  0

$\tau_2 =$ 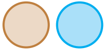

Monitoring Observations

Question

Having observed the observation sequence τ, what is the probability of being in q_c?

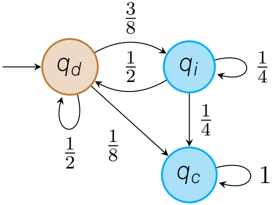

$\tau_1 =$  0

$\tau_2 =$ 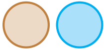 $^1/_4$

Monitoring Observations

Question

Having observed the observation sequence τ, what is the probability of being in q_c?

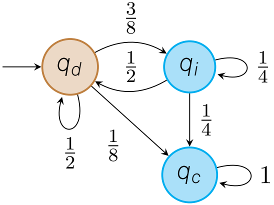

$\tau_1 =$  0

$\tau_2 =$ 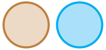 $1/4$

$\tau_3 =$ 

Monitoring Observations

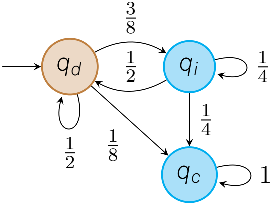

Question

Having observed the observation sequence τ, what is the probability of being in q_c?

$\tau_1 =$  0

$\tau_2 =$ 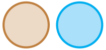 $^1/_4$

$\tau_3 =$  $^5/_8$

Monitoring Observations

Question

Having observed the observation sequence τ, what is the probability of being in q_c?

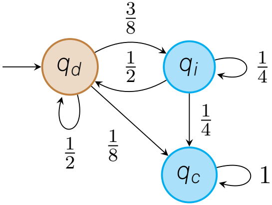

$\tau_1 =$  0

$\tau_2 =$ 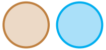 $1/4$

$\tau_3 =$  $5/8$

$\tau_4 =$ 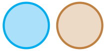

Monitoring Observations

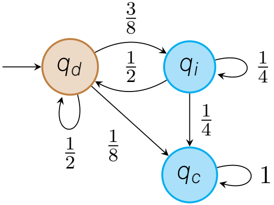

Question

Having observed the observation sequence τ, what is the probability of being in q_c?

$\tau_1 =$  0

$\tau_2 =$ 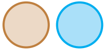 $1/4$

$\tau_3 =$  $5/8$

$\tau_4 =$ 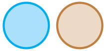 $?$

Monitoring Observations

Question

Probability above $\lambda = 0.3$ is **unsafe.**
Should the warning light go on?

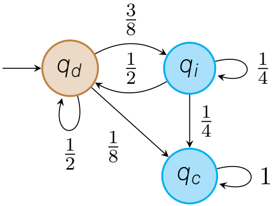

$\tau_1 =$  ok

$\tau_2 =$ 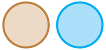 ok

$\tau_3 =$  

$\tau_4 =$ 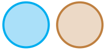 ?

Monitoring Observations

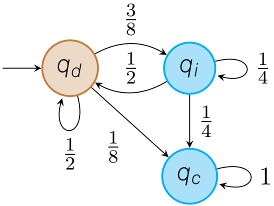

Question

Probability above $\lambda = 0.3$ is **unsafe.**
Should the warning light go on?

$\tau_1 =$  ok

$\tau_2 =$ 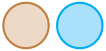 ok

$\tau_3 =$ 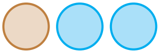 

$\tau_4 =$ 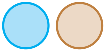 ?

$$\mathbb{U}_\lambda = \{\tau_3, \ldots\}$$

Monitoring Observations

Question

Probability below $\lambda_s = 0.1$ is **safe.**
Probability above $\lambda_u = 0.3$ is **unsafe.**
Should the warning light go on?

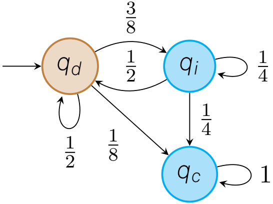

$\tau_1 =$  ok

$\tau_2 =$ 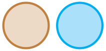 ?

$\tau_3 =$  

$\tau_4 =$ 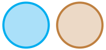 ?

Monitoring Observations

Question

Probability below $\lambda_s = 0.1$ is **safe.**
Probability above $\lambda_u = 0.3$ is **unsafe.**
Should the warning light go on?

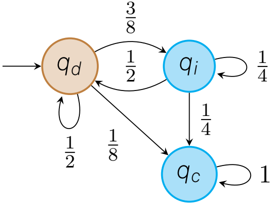

$\tau_1 =$  ok

$\tau_2 =$ 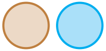 ?

$\tau_3 =$  

$\tau_4 =$ 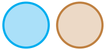 ?

$$\mathbb{U}_{\lambda_u} = \{\tau_3, \ldots\} \qquad \mathbb{S}_{\lambda_s} = \{\tau_1, \ldots\}$$

Monitoring Observations

Question

Probability below $\lambda_s = 0.1$ is **safe.**
Probability above $\lambda_u = 0.3$ is **unsafe.**
Horizon of 3 observations.
Should the warning light go on?

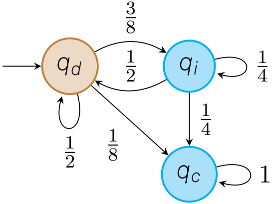

$\tau_1 =$  ok

$\tau_2 =$ 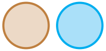 ?

$\tau_3 =$  

$\tau_4 =$ 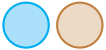 ?

$$\mathbb{U}^{\leq 3}_{\lambda_u} = \{\tau_3\}$$

$$\mathbb{S}^{\leq 3}_{\lambda_s} = \{\tau_1, \ldots\}$$

What is a Monitor?

$$\mathbb{U}^{\leq 3}_{\lambda_u} = \{\bigcirc\ \bigcirc\ \bigcirc\} \qquad \mathbb{S}^{\leq 3}_{\lambda_s} = \{\bigcirc, \ldots\}$$

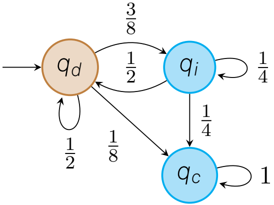

What is a Monitor?

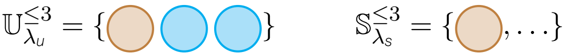

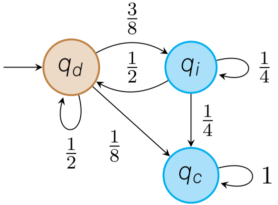

$\mathcal{A}:$

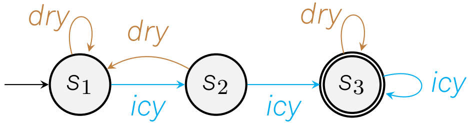

What is a Monitor?

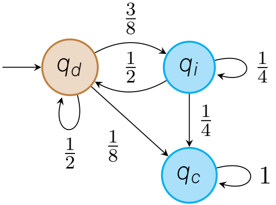

$\mathcal{A}$:

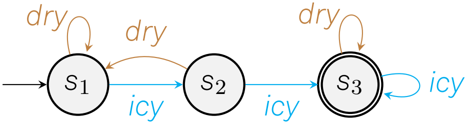

Monitor Correctness

$$\mathbb{U}^{\leq 3}_{\lambda_u} \subseteq \mathcal{L}(\mathcal{A}) \subseteq \Sigma^* \setminus \mathbb{S}^{\leq 3}_{\lambda_s}$$

What is a Monitor?

$\mathbb{U}^{\leq 3}_{\lambda_u} = \{$  $\}$ $\mathbb{S}^{\leq 3}_{\lambda_s} = \{$  $, \ldots\}$

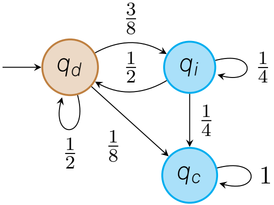

$\mathcal{A}$:

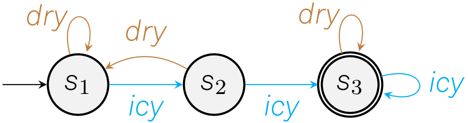

Monitor Correctness

$$\mathbb{U}^{\leq 3}_{\lambda_u} \subseteq \mathcal{L}(\mathcal{A}) \subseteq \Sigma^* \setminus \mathbb{S}^{\leq 3}_{\lambda_s}$$

Verifying Monitors (this paper)

No Missed Alarms Problem

Given a HMM generating a set of traces $\mathbb{U}_{\lambda_u}^{\leq h}$, and a monitor $\mathcal{A}$, verify that

$$\forall \tau \in \mathbb{U}_{\lambda_u}^{\leq h}.\ \tau \in \mathcal{L}(\mathcal{A})$$

Verifying Monitors (this paper)

No Missed Alarms Problem

Given a HMM generating a set of traces $\mathbb{U}_{\lambda_u}^{\leq h}$, and a monitor $\mathcal{A}$, verify that

$$\forall \tau \in \mathbb{U}_{\lambda_u}^{\leq h}.\ \tau \in \mathcal{L}(\mathcal{A})$$

$\Downarrow$ Find a counter example

Find Missed Alarm Problem

Given a HMM generating a set of traces $\mathbb{U}_{\lambda_u}^{\leq h}$, and a monitor $\mathcal{A}$,

$$\exists \tau \in \mathbb{U}_{\lambda_u}^{\leq h}.\ \tau \notin \mathcal{L}(\mathcal{A})$$

Verifying Monitors (this paper)

No Missed Alarms Problem

Given a HMM generating a set of traces $\mathbb{U}_{\lambda_u}^{\leq h}$, and a monitor $\mathcal{A}$, verify that

$$\forall \tau \in \mathbb{U}_{\lambda_u}^{\leq h}.\ \tau \in \mathcal{L}(\mathcal{A})$$

$\Downarrow$ Find a counter example

Find Missed Alarm Problem

Given a HMM generating a set of traces $\mathbb{U}_{\lambda_u}^{\leq h}$, and a monitor $\mathcal{A}$,

$$\exists \tau \in \mathbb{U}_{\lambda_u}^{\leq h}.\ \tau \notin \mathcal{L}(\mathcal{A})$$

Complexity

Finding a missed alarm is NP-complete (proof in the paper).

Searching for Missed Alarms

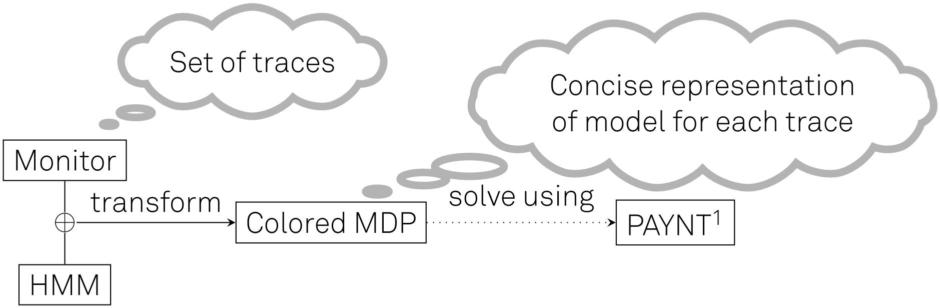

- Writing conditional probability properties using reachability, by Baier et al.[2].
- Equate traces in the HMM to policies in the colored MDP, by Badings et al. [3].

[1] R. Andriushchenko et al., “PAYNT: A tool for inductive synthesis of probabilistic programs,”, 2021.
[2] C. Baier et al., “Computing conditional probabilities in markovian models efficiently,”, 2014
[3] T. S. Badings et al., “Ctmcs with imprecisely timed observations,”, 2024

Transformation 1/4

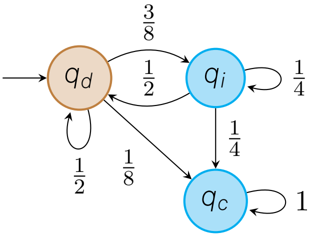

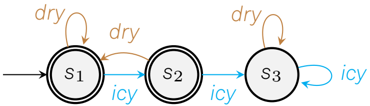

Transformation 1/4

Find an unsafe trace which is not in the monitor

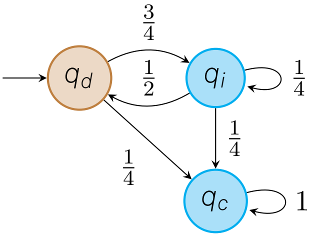

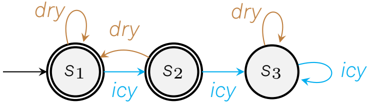

Transformation 1/4

Find an unsafe trace which is ~~not~~ in the monitor

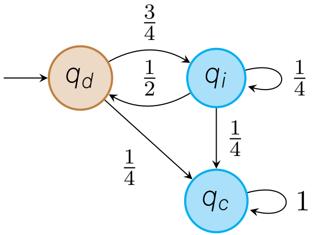

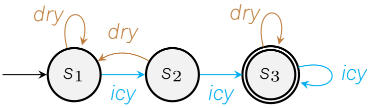

Transformation 1/4

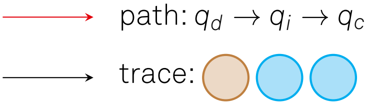

Find an unsafe trace which is ~~not~~ in the monitor

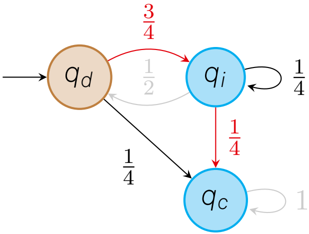

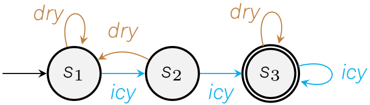

Transformation 2/4

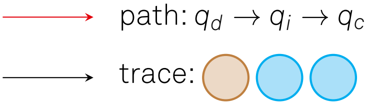

Find an unsafe trace which is ~~not~~ in the monitor

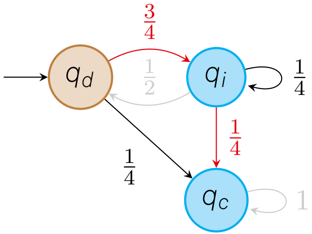

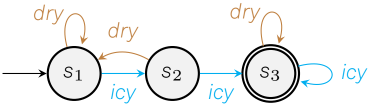

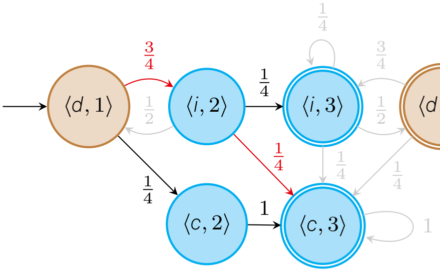

Transformation 3/4

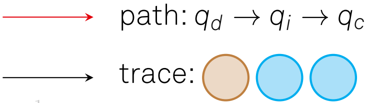

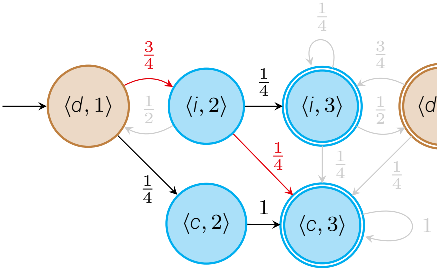

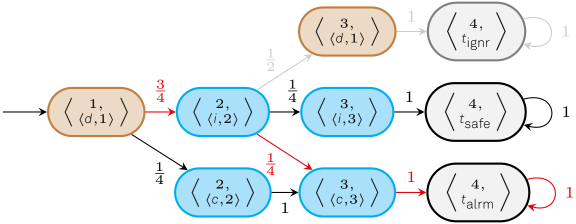

Find a trace which

- does not end in t_{ignr},
- has probability $> \lambda_u$ to reach $\left\langle {4, \atop t_{\text{alrm}}} \right\rangle$.

Transformation 4/4

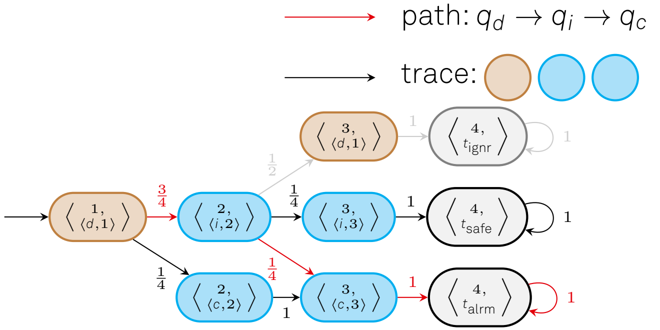

Find a trace which

- does not end in t_{ignr},
- has probability $> \lambda_u$ to reach $\left\langle \begin{matrix} 4, \\ t_{alrm} \end{matrix} \right\rangle$.

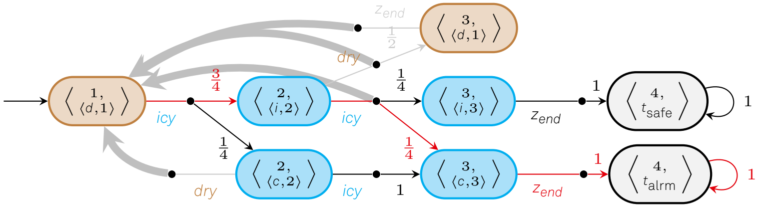

Find a policy $\sigma : \mathbb{N}^{\leq 4} \to Z$ s.t.

- we reach an end state,
- reach t_{alrm} with prob. $\geq \lambda_u$.

Transformation 4/4

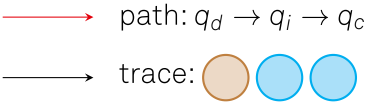

Find a policy
$\sigma : \mathbb{N}^{\leq 4} \to Z$ s.t.

- we reach an end state,
- reach t_{alrm} with prob. $\geq \lambda_u$.

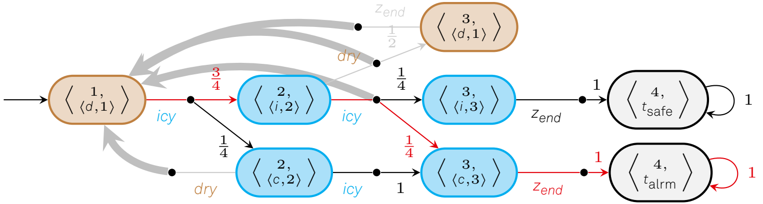

Solvable by PAYNT

Transformation 4/4

path: $q_d \to q_i \to q_c$

trace:

Find a policy $\sigma : \mathbb{N}^{\leq 4} \to Z$ s.t.

- we reach an end state,
- reach t_{alrm} with prob. $\geq \lambda_u$.

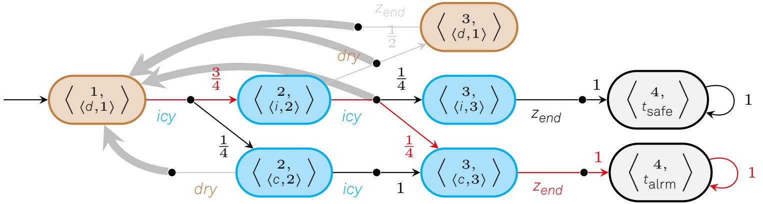

Solvable by PAYNT

Theorem 1

"The transformation is correct"

Learning Verified Monitors for Hidden Markov Models

Learning a monitor

- Active automata learning using L*.

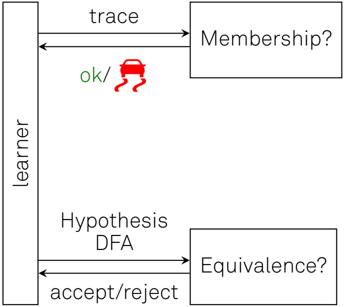

Learning a monitor

- Active automata learning using L*.
- MQ: Forward Filtering implemented by Premise[4] on the HMM with threshold λ_l.

$$\lambda_s \leq \lambda_l \leq \lambda_u$$

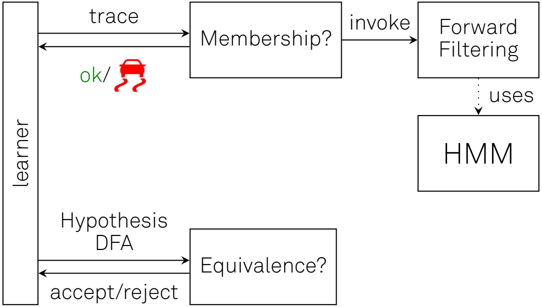

[4] S. Junges et al., “Runtime monitors for markov decision processes,”, 2021

Learning a monitor

- Active automata learning using L*.
- MQ: Forward Filtering implemented by Premise[4] on the HMM with threshold λ_l.

$$\lambda_s \leq \lambda_l \leq \lambda_u$$

- EQ: is a candidate monitor correct.

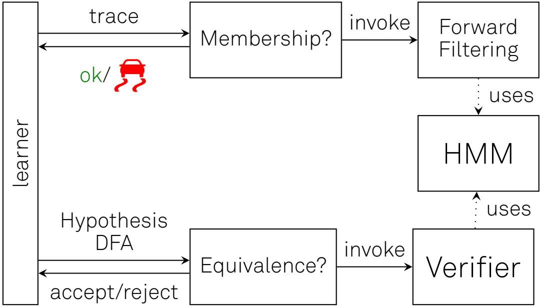

[4] S. Junges et al., "Runtime monitors for markov decision processes,", 2021

Learning a monitor

- Active automata learning using L*.
- MQ: Forward Filtering implemented by Premise[4] on the HMM with threshold λ_l.

$$\lambda_s \leq \lambda_l \leq \lambda_u$$

- EQ: is a candidate monitor correct.

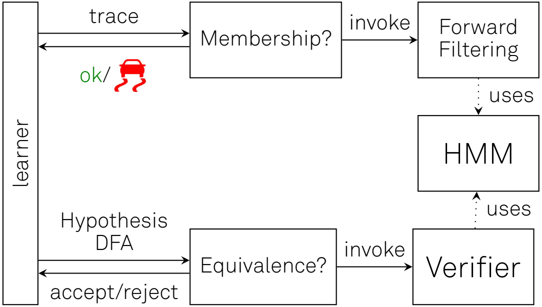

Theorem 2

"Monitors learned using our verification algorithm are correct."

[4] S. Junges et al., "Runtime monitors for markov decision processes,", 2021

Correctness Experiments

Benchmark A-63/64

290 states, 1258 transitions,
50 observations, horizon of 10
$\lambda_s = \lambda_l = \lambda_u = 0.3$

Correctness Experiments

Benchmark A-63/64

290 states, 1258 transitions,
50 observations, horizon of 10
$\lambda_s = \lambda_l = \lambda_u = 0.3$

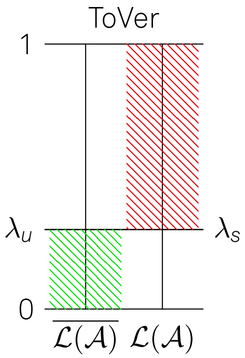

Correctness Experiments

Benchmark A-63/64

290 states, 1258 transitions,
50 observations, horizon of 10
$\lambda_s = \lambda_l = \lambda_u = 0.3$

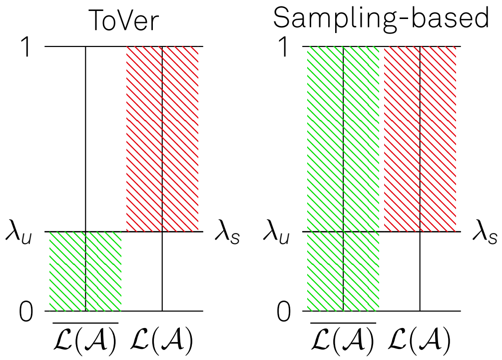

Correctness Experiments

Benchmark A-63/64

290 states, 1258 transitions,
50 observations, horizon of 10
$\lambda_s = 0.1 \quad \lambda_l = 0.3 \quad \lambda_u = 0.35$

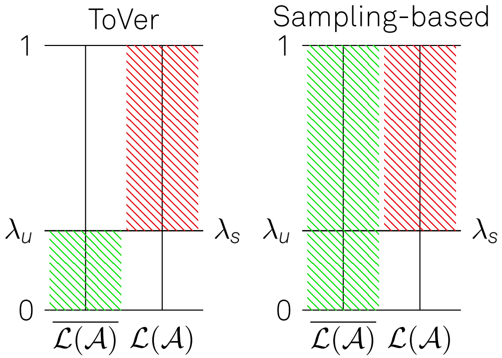

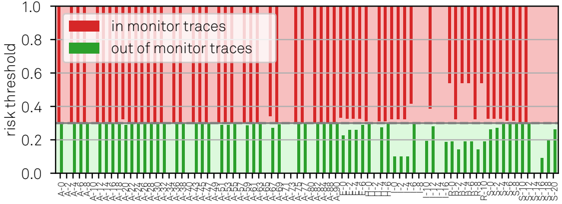

(a) Learning with verification, $\lambda_s = \lambda_l = \lambda_u$

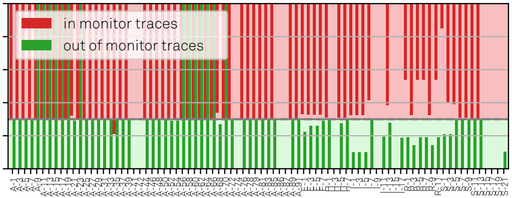

(b) Learning with sampling-based verification

Correctness Experiments

Benchmark A-63/64

290 states, 1258 transitions,
50 observations, horizon of 10
$\lambda_s = 0.1 \quad \lambda_l = 0.3 \quad \lambda_u = 0.35$

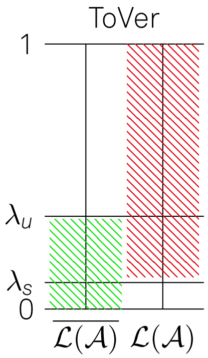

Correctness Experiments

Benchmark A-63/64

290 states, 1258 transitions,
50 observations, horizon of 10

$\lambda_s = 0.1 \quad \lambda_l = 0.3 \quad \lambda_u = 0.35$

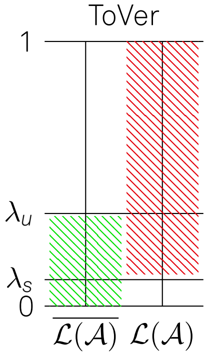

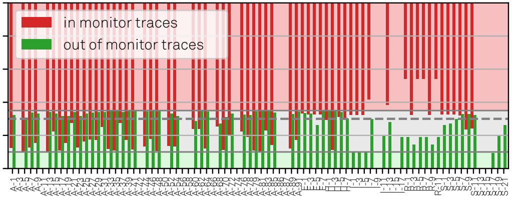

(c) Learning with verification, $\lambda_s < \lambda_l < \lambda_u$

Conclusion

Summary

- We present a verification algorithm for HMM monitors.
- We prove the verification problem is coNP-complete.
- We integrate it with active automata learning to learn correct monitors.
- We learn monitors with up to 1500 states in 11 hours on models with 100s of states.

Future interests

- Ideas to adapt AAL more to our specific problem.
- Adapt colored MDP model checking more to our specific problem of conditional probabilities.
- Learn models from data such that they are useful for monitoring.

Email: `luko.vandermaas@ru.nl`

Results: Runtime

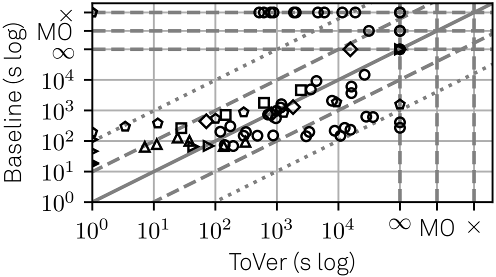

(a) $\lambda_s < \lambda_l < \lambda_u$, Runtime

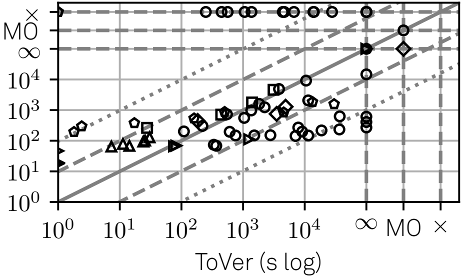

(b) $\lambda_s = \lambda_l = \lambda_u$, Runtime

Results: Monitor Size

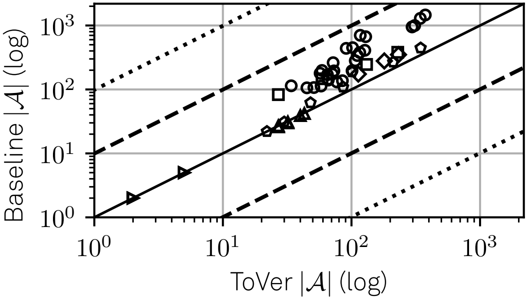

(a) $\lambda_s < \lambda_l < \lambda_u$, Size of $\mathcal{A}$

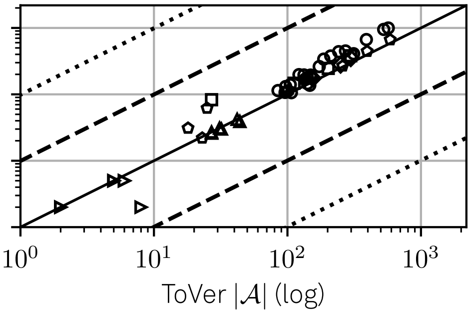

(b) $\lambda_s = \lambda_l = \lambda_u$, Size of $\mathcal{A}$