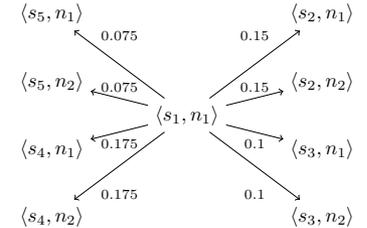
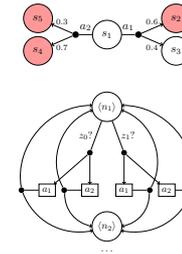
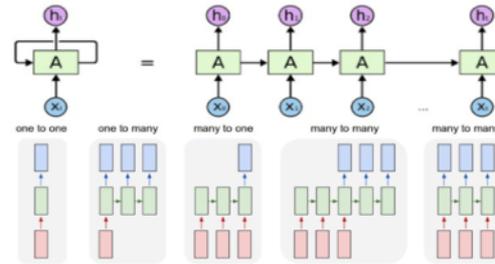
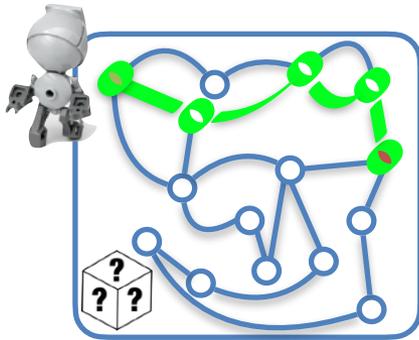


# Planning under Uncertainty

## AI Safety via Formal Verification



Nils Jansen

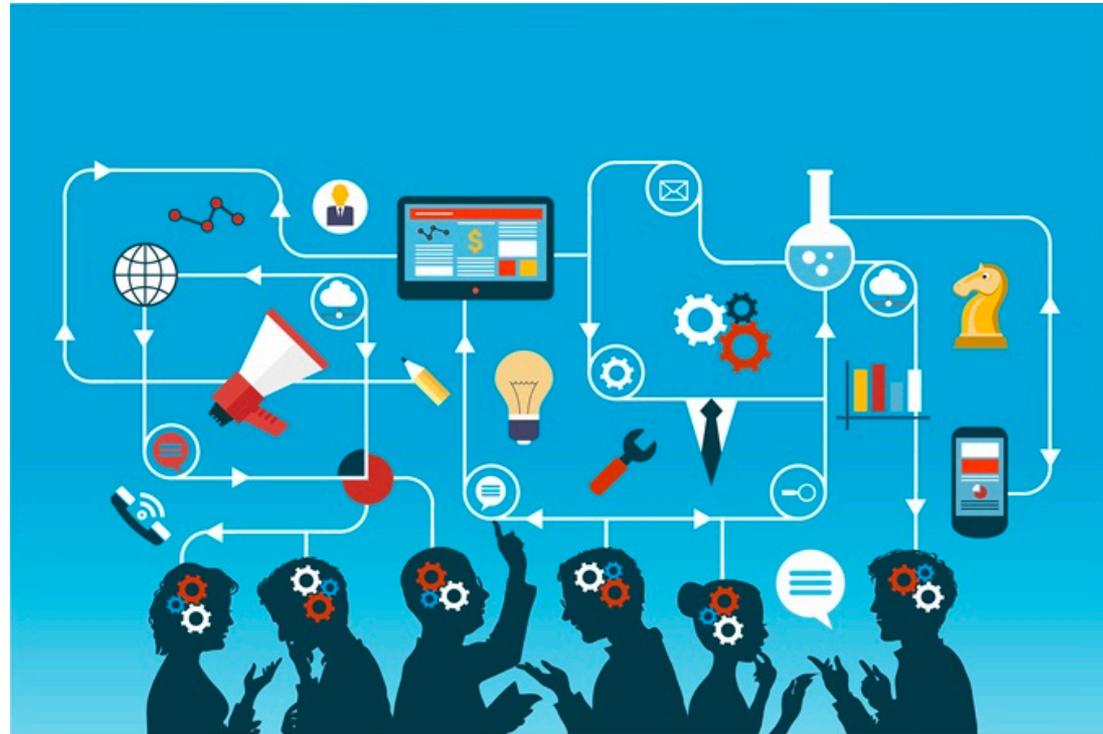
NFM AI-SAFETY, May 11, 2020

# joint work with:

Bernd Becker, Steve Carr, Murat Cubuktepe,  
Alexandru Serban, Marnix Suilen, Ufuk Topcu, Ralf Wimmer,

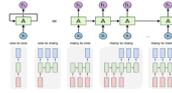


Radboud University



# AI Safety - A Restricted View

## Neural Network Controllers



- how to ensure **safe** decision-making?
- what is **safety**?

## Partial Information



- how to operate if the system state is **not fully observable**?
- what **guarantees** can be given?

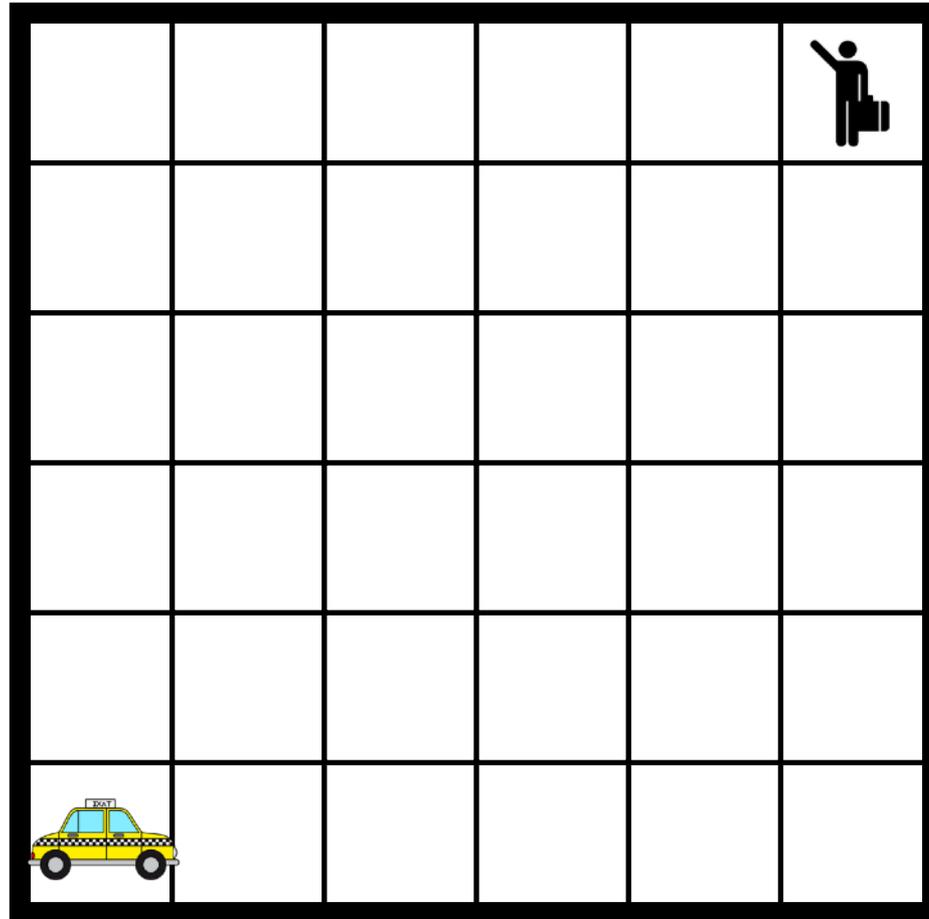
## Uncertainty



- how to robustly account for imperfect sensor information?
- what is robustness?

# Help the Taxi

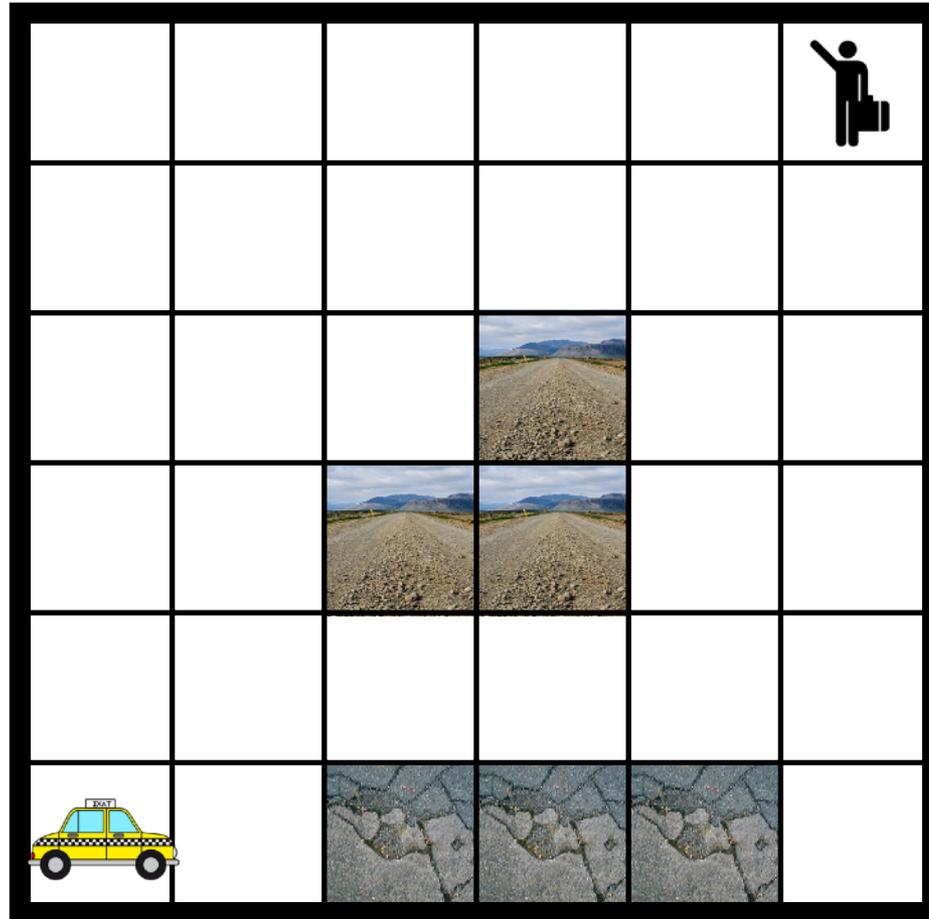
Find the best way to  
the passenger



# Help the Taxi

Find the best way to  
the passenger

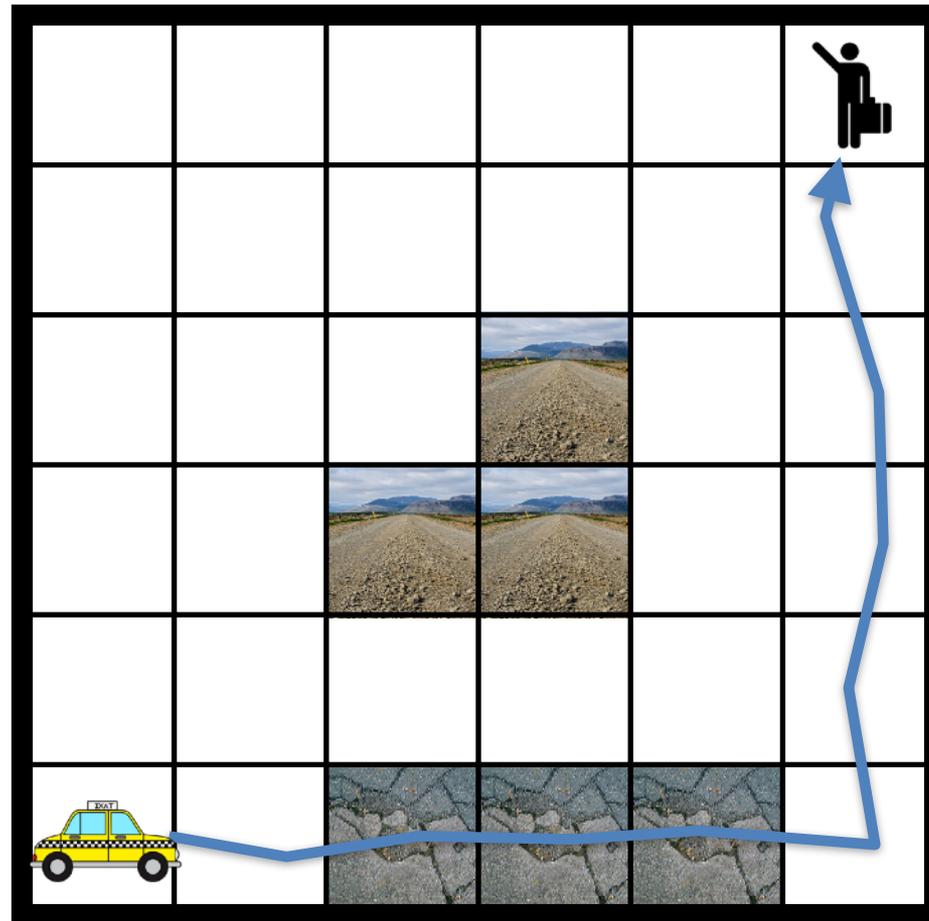
Take **bad roads** into  
account



# Help the Taxi

Find the best way to  
the passenger

Take **bad roads** into  
account

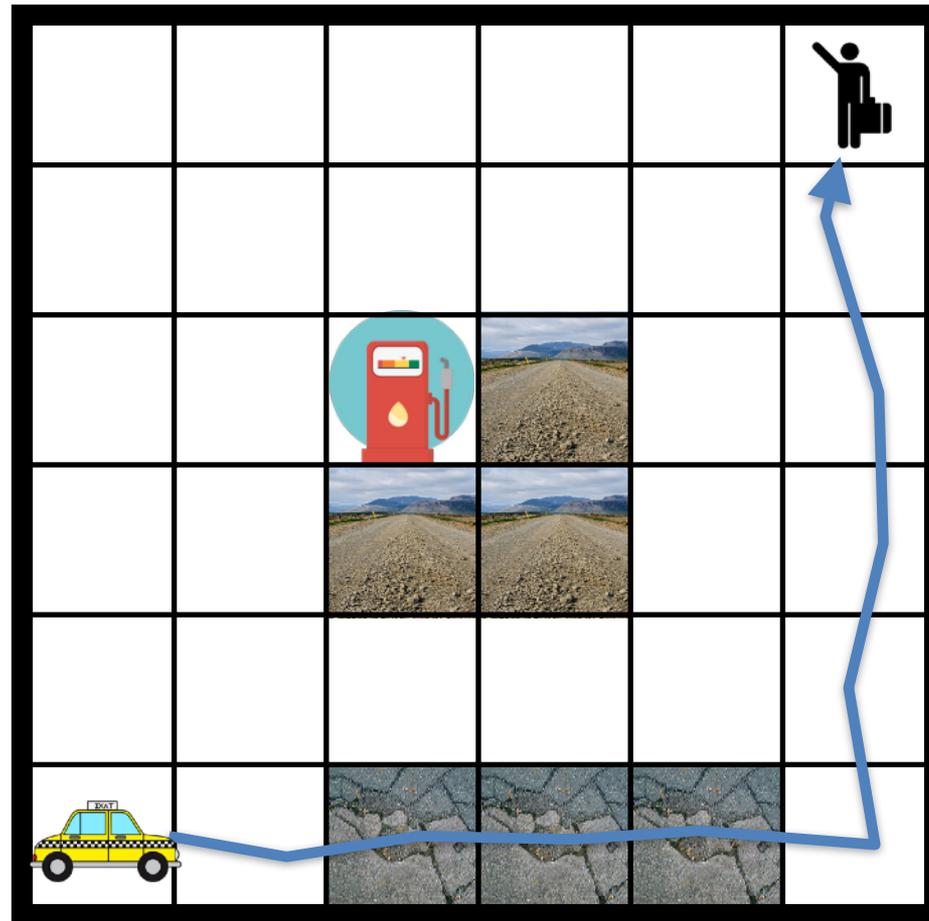


# Help the Taxi

Find the best way to  
the passenger

Take **bad roads** into  
account

**Refuel** on the way



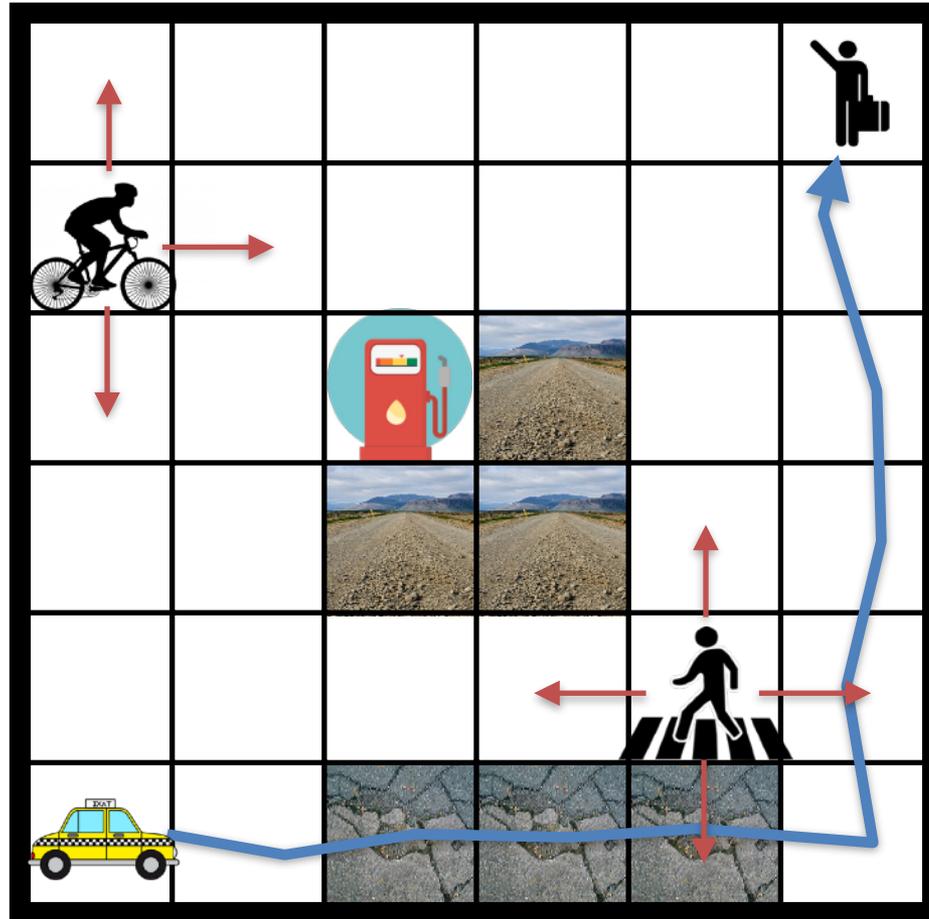
# Help the Taxi

Find the best way to the passenger

Take **bad roads** into account

**Refuel** on the way

Avoid randomly moving **pedestrians and cyclists**





# Help the Taxi

Find the best way to the passenger

Take **bad roads** into account

**Refuel** on the way

Avoid randomly moving **pedestrians and cyclists**

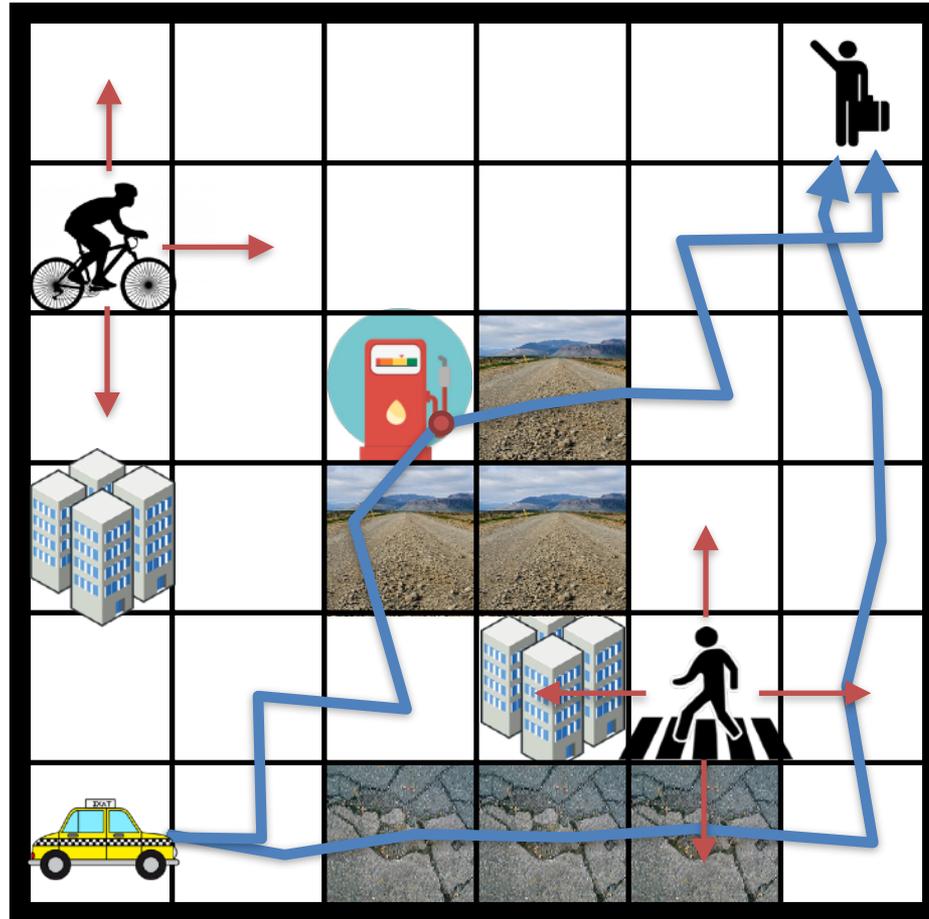
Account for **restricted view**

$Pr_{max}(\neg(\text{ped} \vee \text{cyc}) \cup (\diamond(\text{fuel} \wedge \bigcirc \diamond \text{pass})))$

**temporal logic**

$EC_{min}(\diamond \text{passenger})$

**expected cost**



Find safe and/or cost-optimal policy to get to the passenger

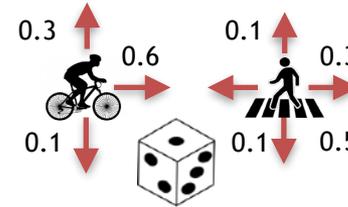
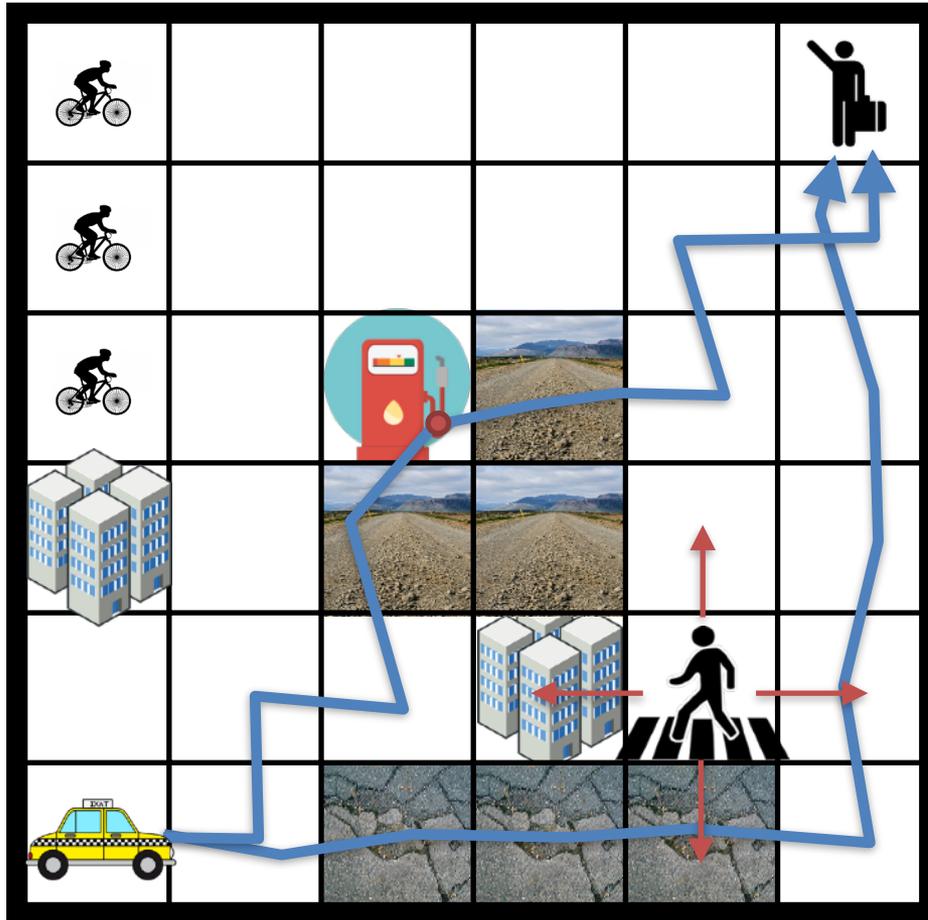
Account for randomness and partial information



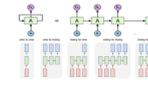




# Help the Taxi: A Closer Look



probability distributions on movements



efficient control via neural network

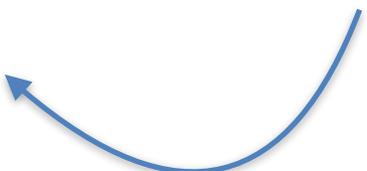


partial observability induces unknown system state



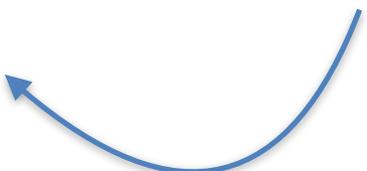
# Two Notoriously Hard Problems

Underlying formal problem:  
Find a **provably correct**  
**Neural Network controller**  
for a **partially observable**  
**Markov Decision Process**

1. Verification of Neural Networks
  2. POMDP synthesis
- 

# Two Notoriously Hard Problems

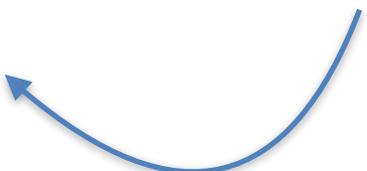
Underlying formal problem:  
Find a **provably correct**  
**Neural Network controller**  
for a **partially observable**  
**Markov Decision Process**

1. Verification of Neural Networks
  2. POMDP synthesis
- 

- Randomized with **infinite memory**: **undecidable**, optimal results.

# Two Notoriously Hard Problems

Underlying formal problem:  
Find a **provably correct**  
**Neural Network controller**  
for a **partially observable**  
**Markov Decision Process**

1. Verification of Neural Networks
  2. POMDP synthesis
- 

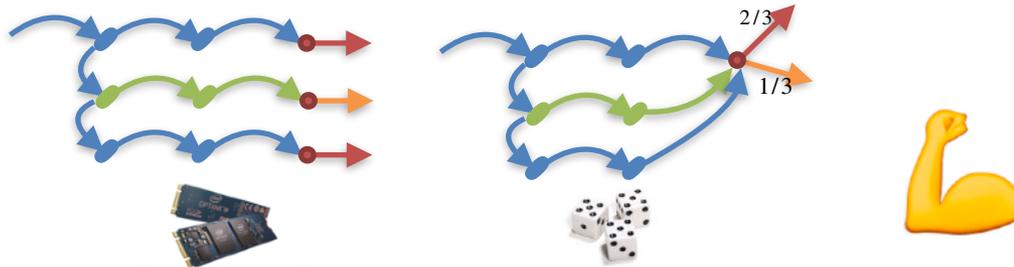
- **Randomized** with **infinite memory**: **undecidable**, optimal results.
- **Randomized** with **finite memory**: **NP-hard**, **SQRT-SUM-hard**, in **PSPACE**, not optimal in general, but sufficient for many applications.

# Two Notoriously Hard Problems

Underlying formal problem:  
Find a **provably correct**  
**Neural Network controller**  
for a **partially observable**  
**Markov Decision Process**

1. Verification of Neural Networks
2. POMDP synthesis

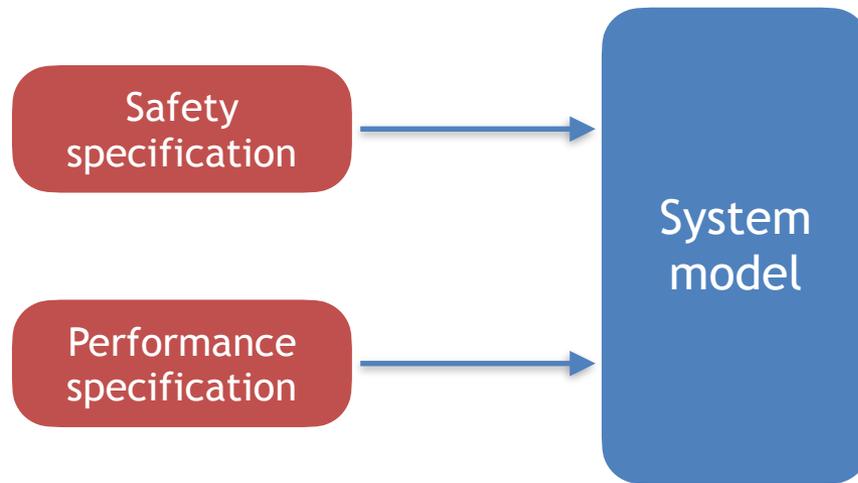
- Randomized with **infinite memory**: **undecidable**, optimal results.
- Randomized with **finite memory**: **NP-hard**, **SQRT-SUM-hard**, in **PSPACE**, not optimal in general, but sufficient for many applications.
- Intuitively: **Randomization** can often **trade off** memory.



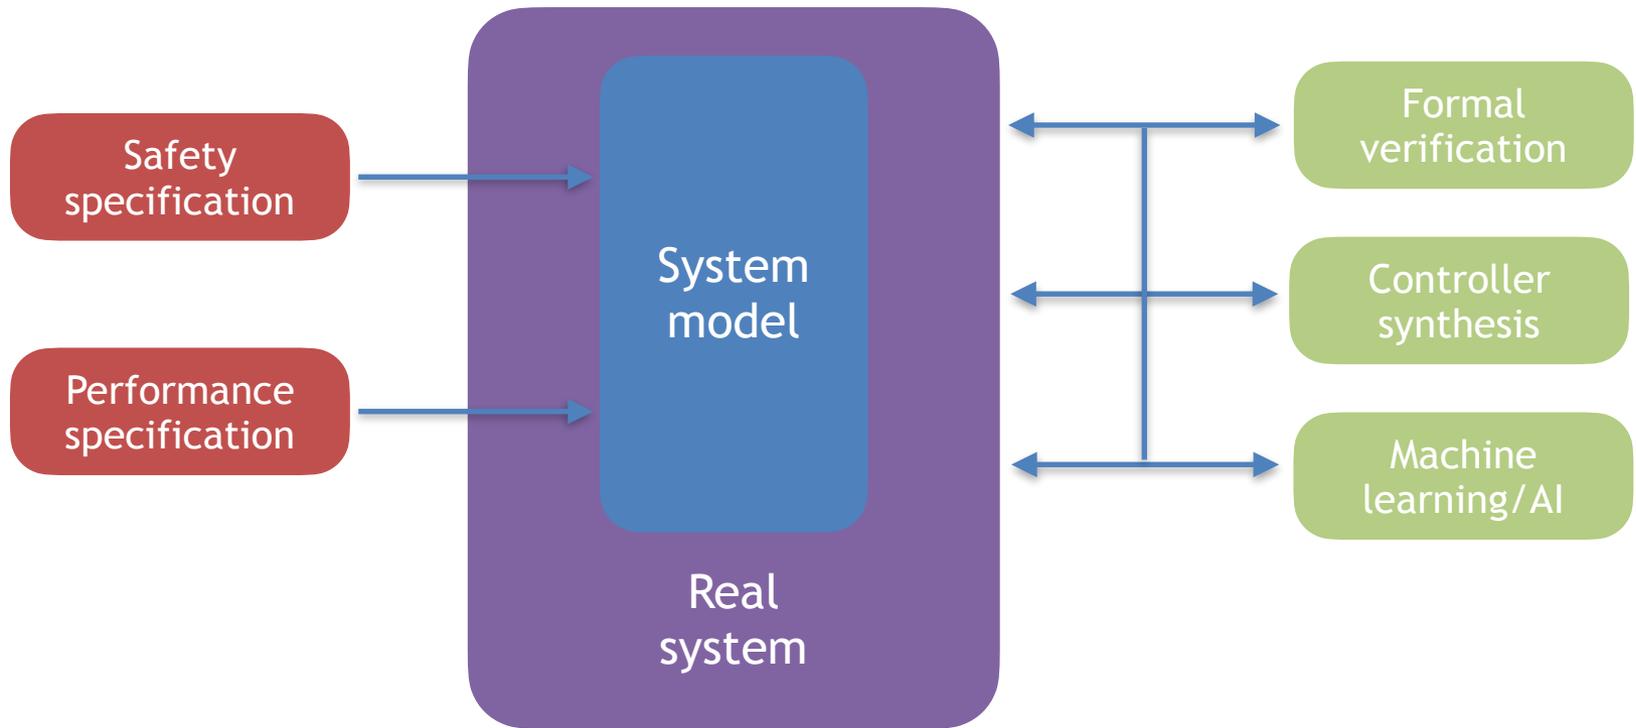
$$\sigma: \text{ObsSeq}_{fin} \rightarrow \text{Distr}(\text{Act})$$

$$\sigma: \text{Obs} \rightarrow \text{Distr}(\text{Act})$$

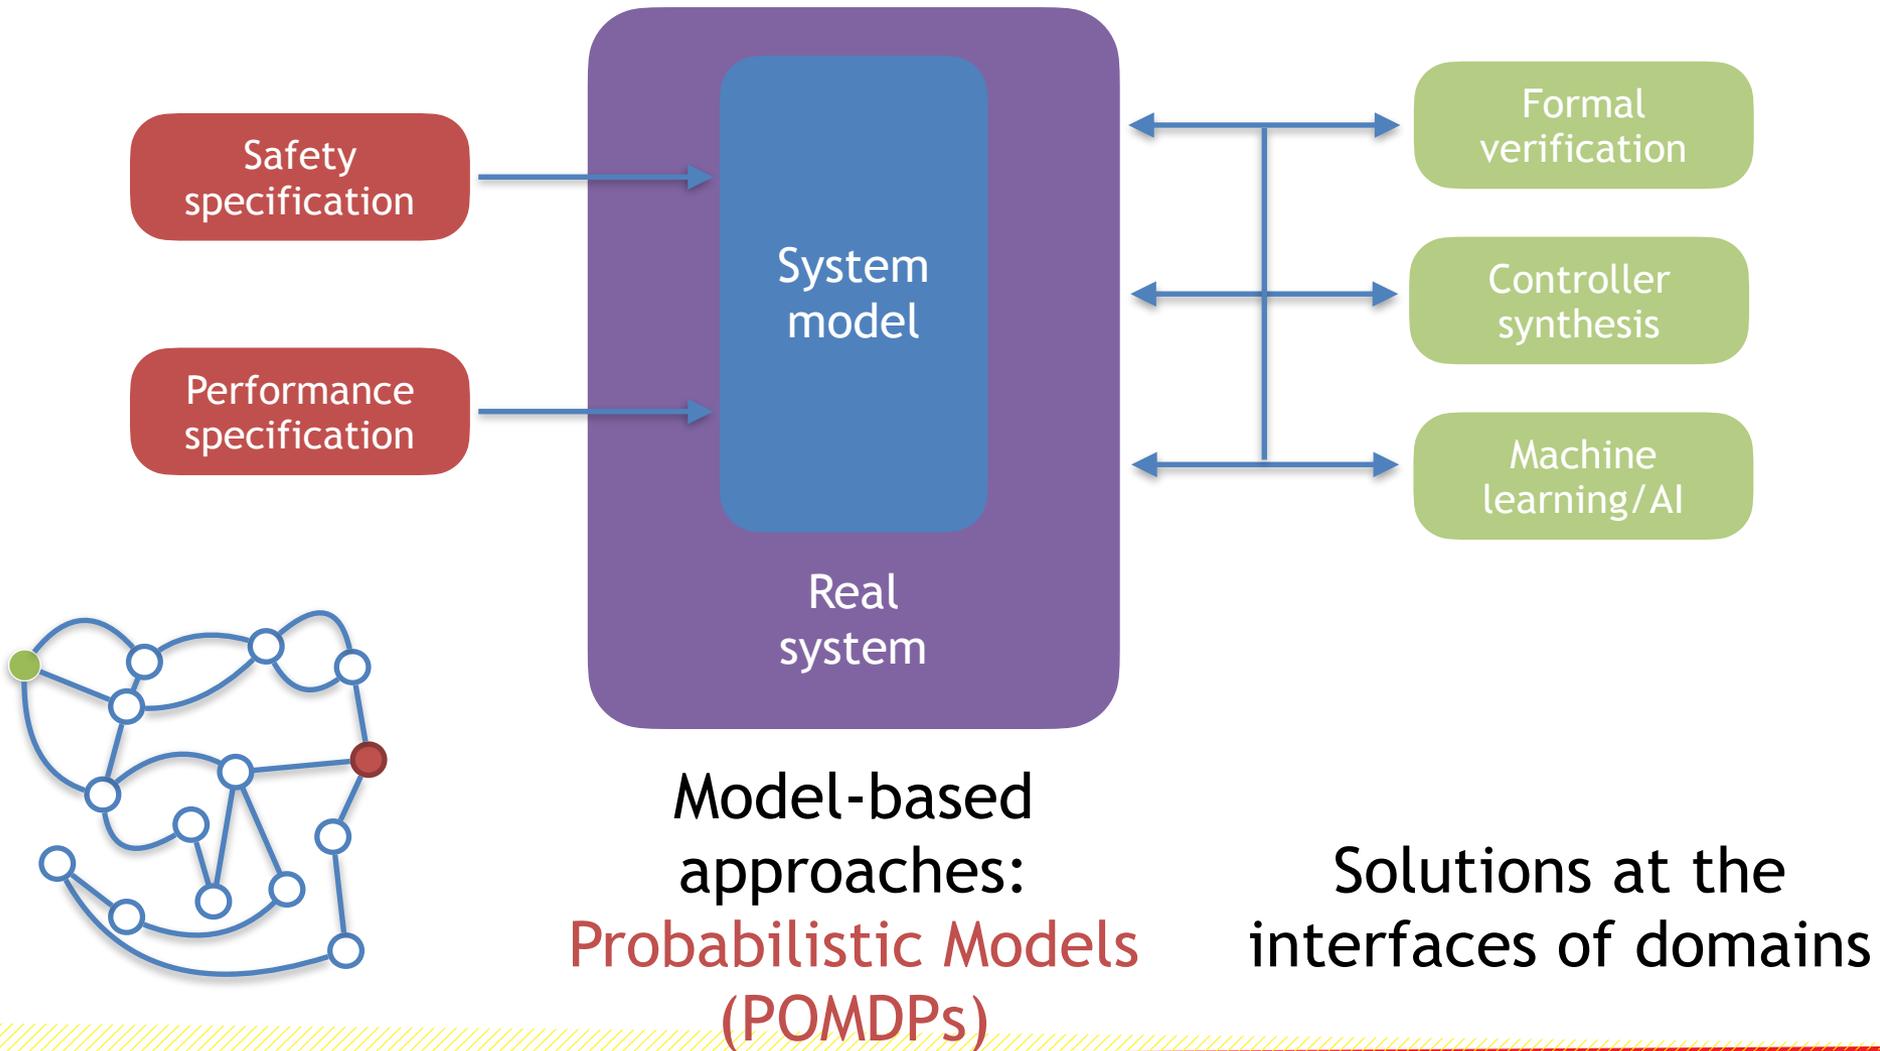
# Decision-Making under Uncertainty



# Decision-Making under Uncertainty



# Decision-Making under Uncertainty



# POMDPs - Applications



Stock Market



Surveying Threatened Species



Health Care



Wireless Sensor Networks



Autonomous Systems



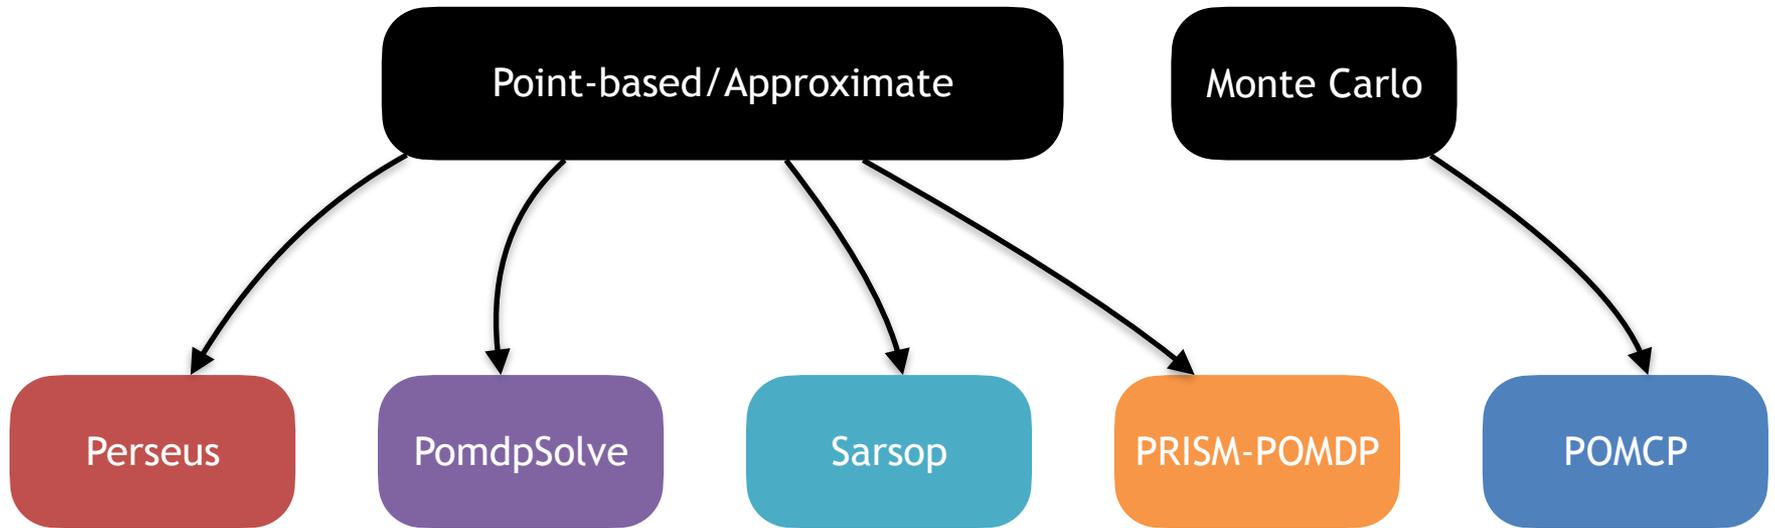
Machine Vision

# POMDP Solving - State of The Art

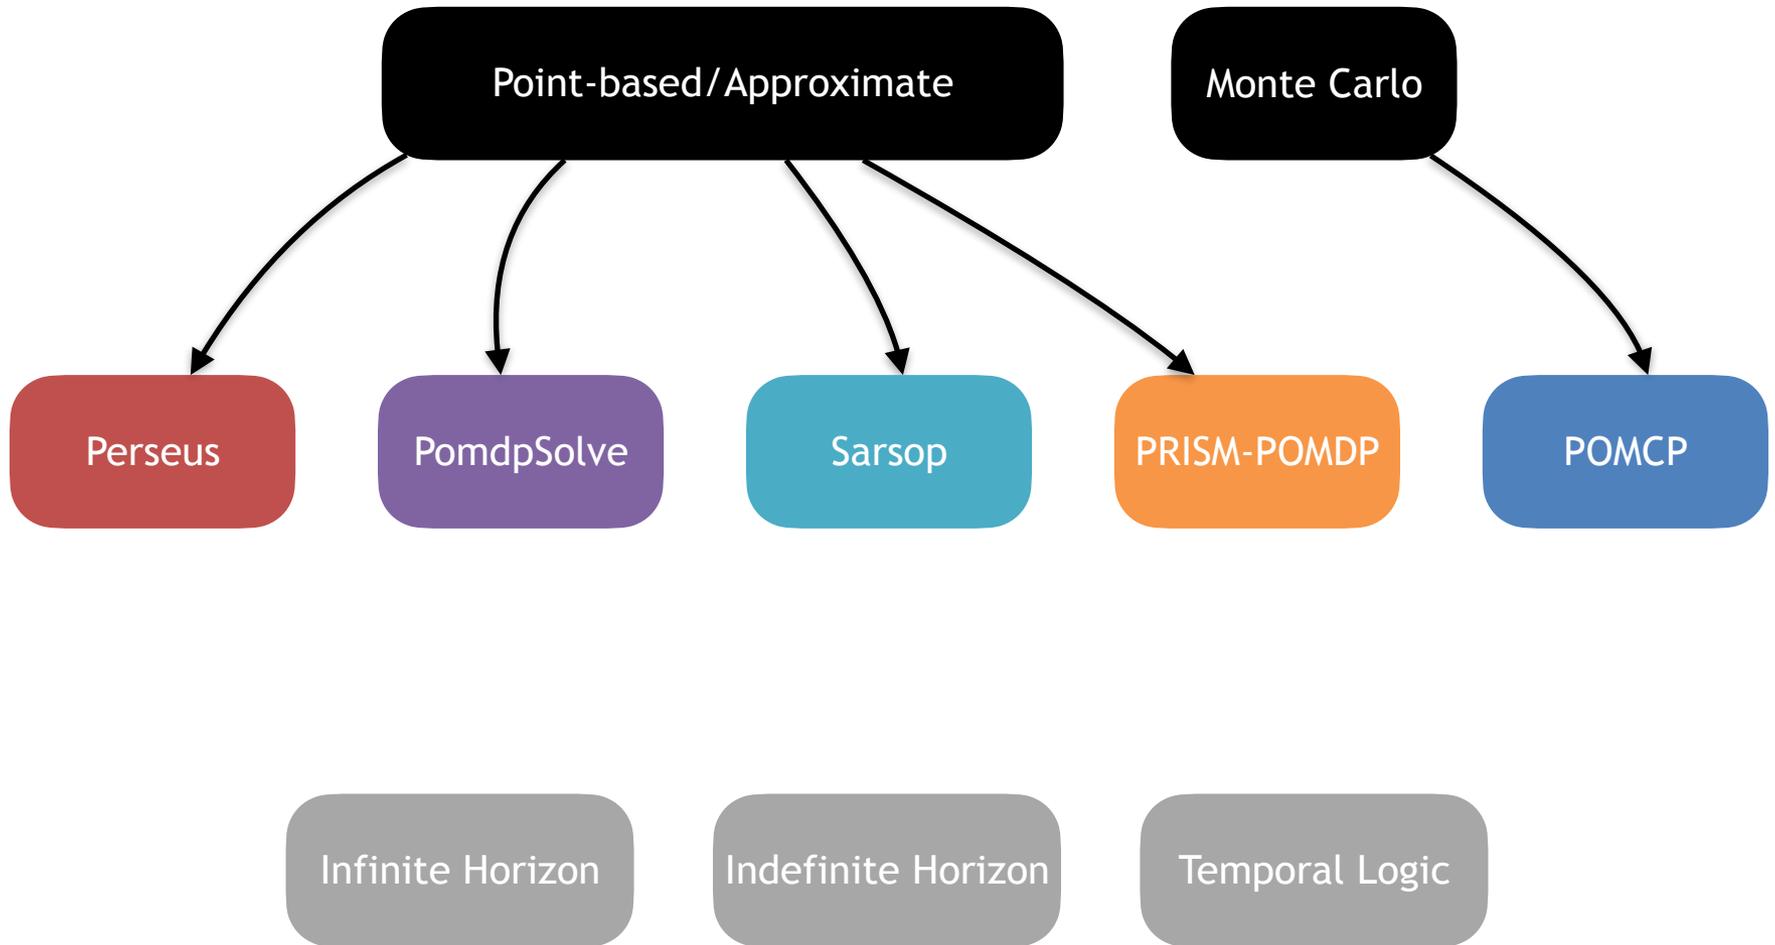
Point-based / Approximate

Monte Carlo

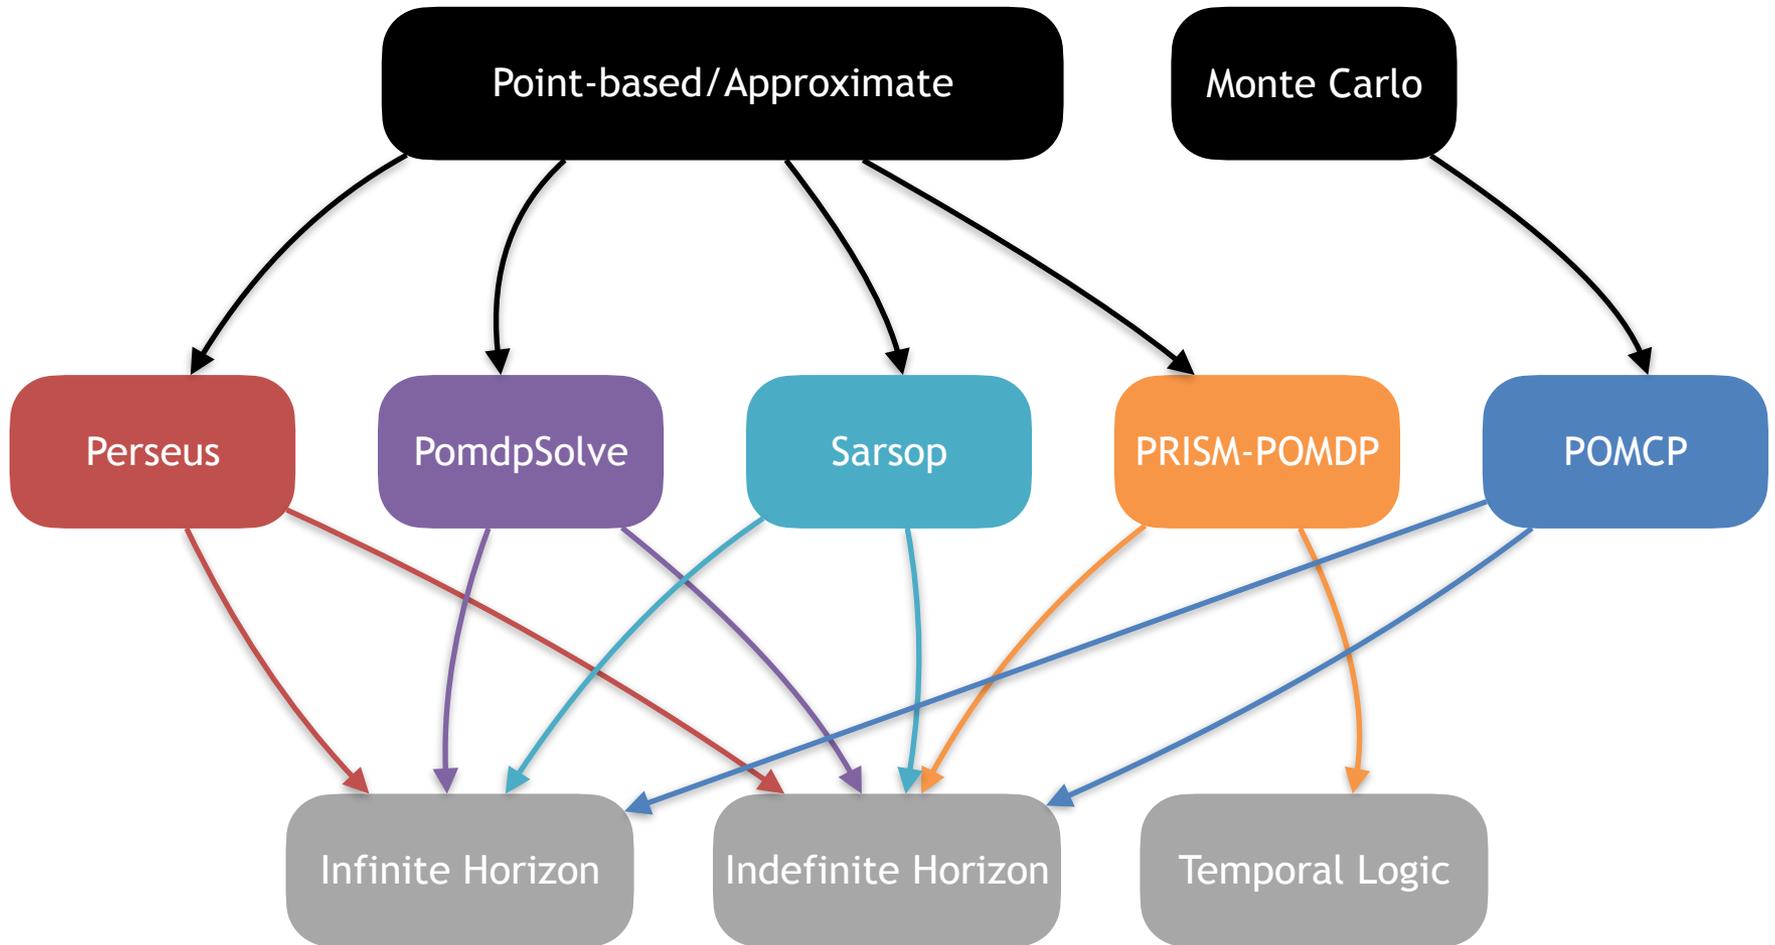
# POMDP Solving - State of The Art



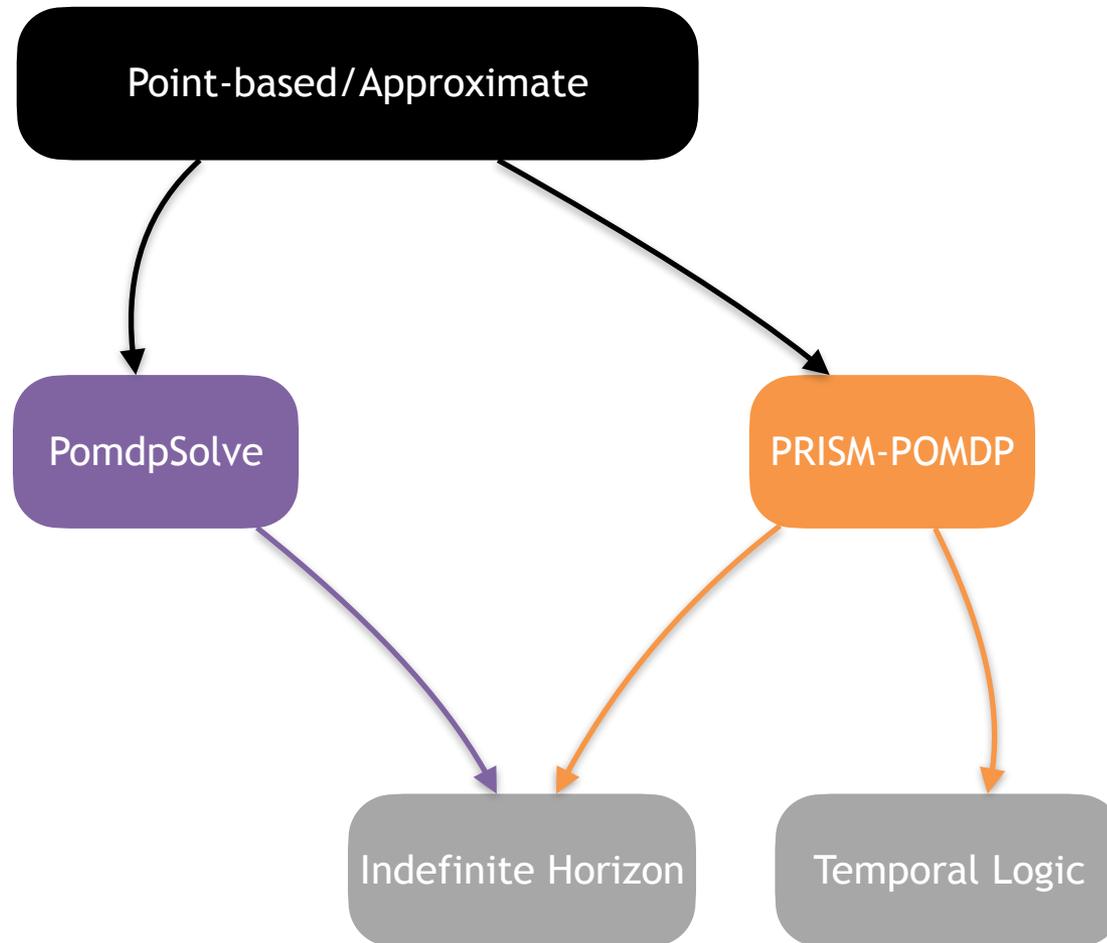
# POMDP Solving - State of The Art



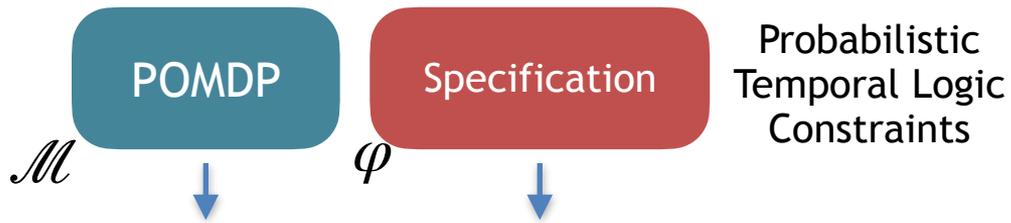
# POMDP Solving - State of The Art



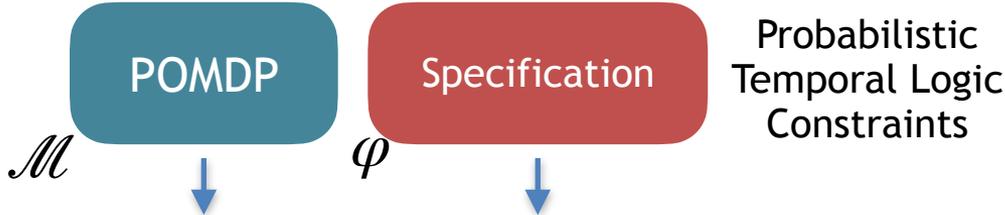
# POMDP Solving - State of The Art



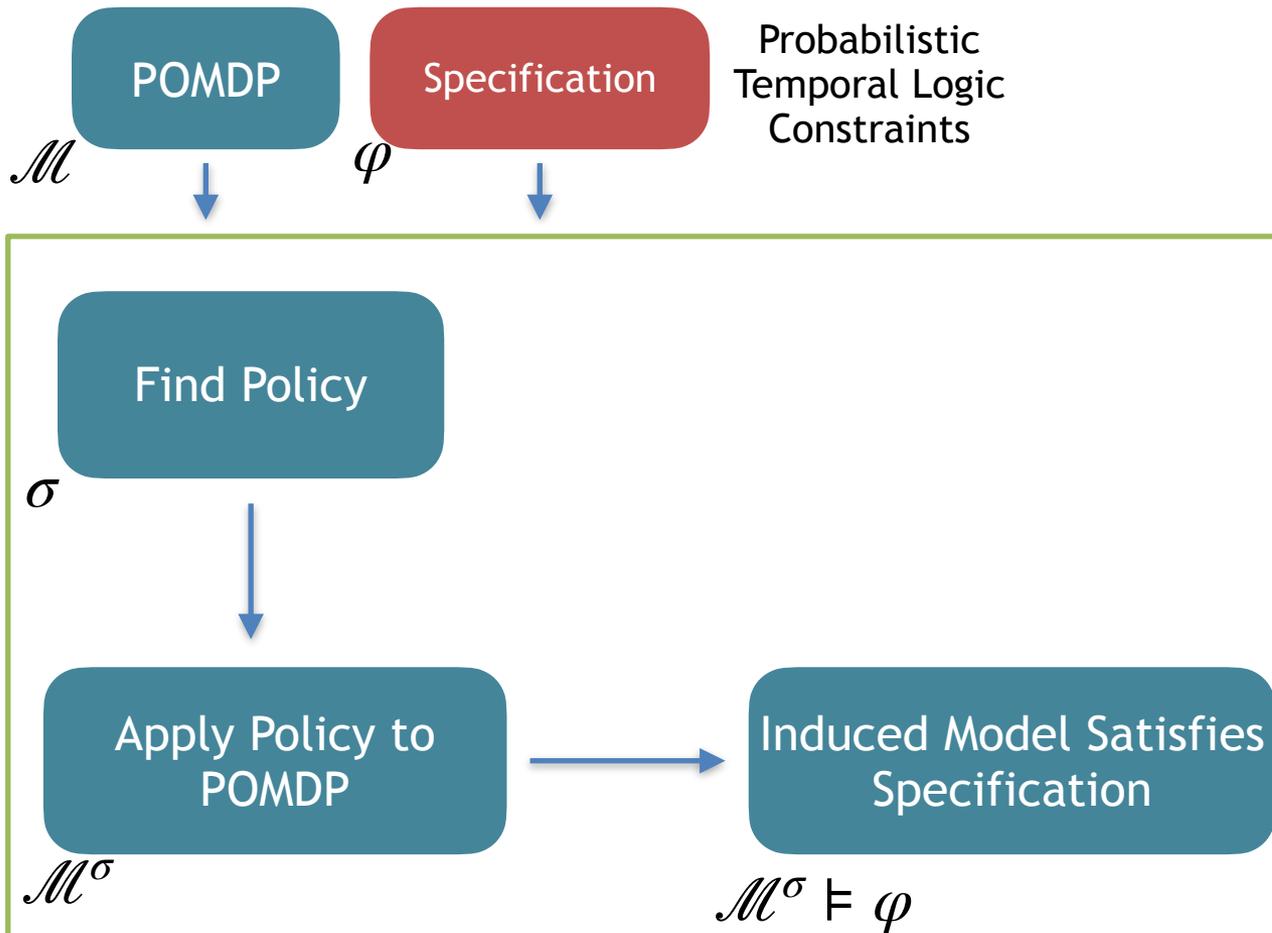
# The Problem: Policy Synthesis



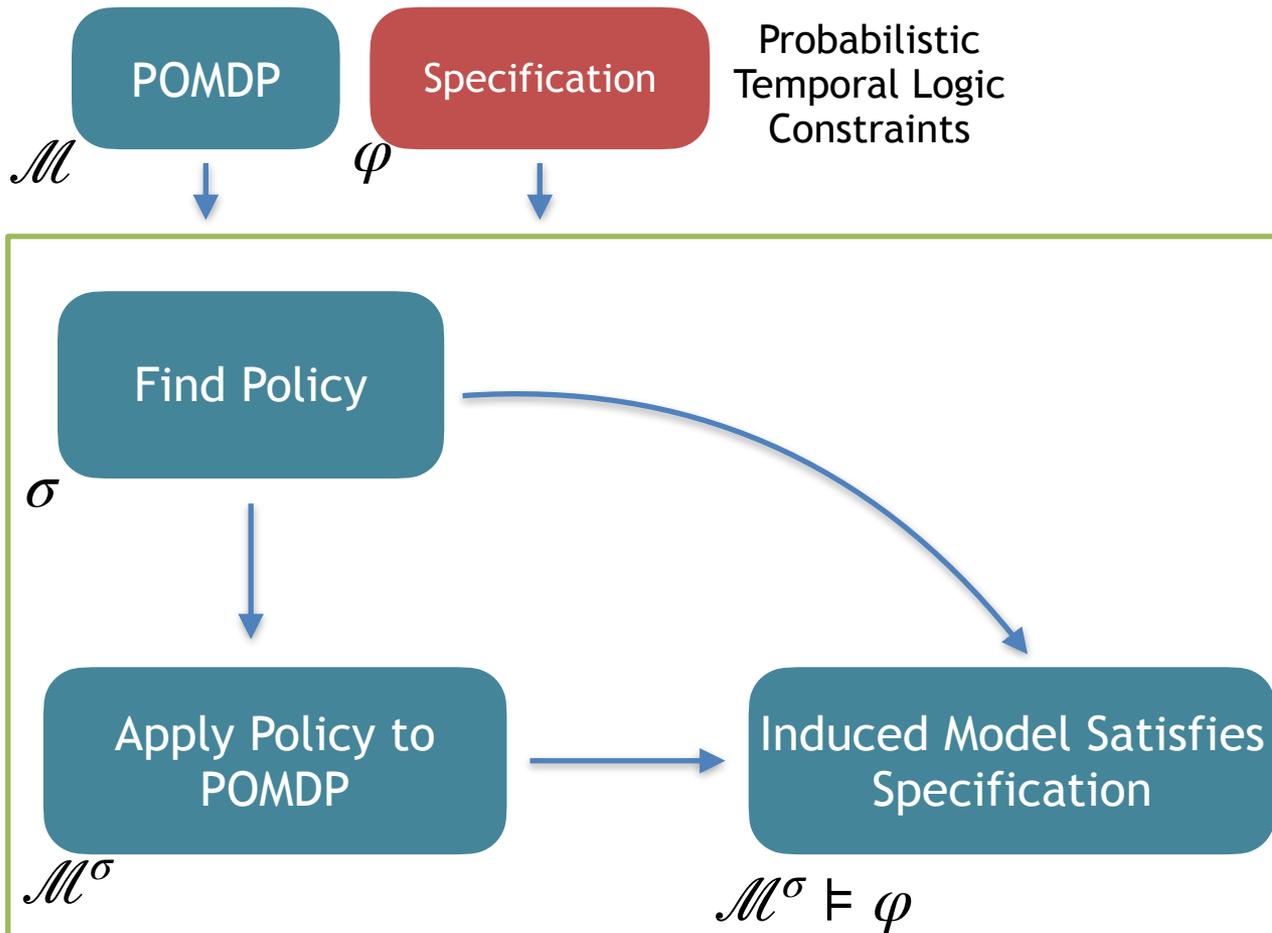
# The Problem: Policy Synthesis



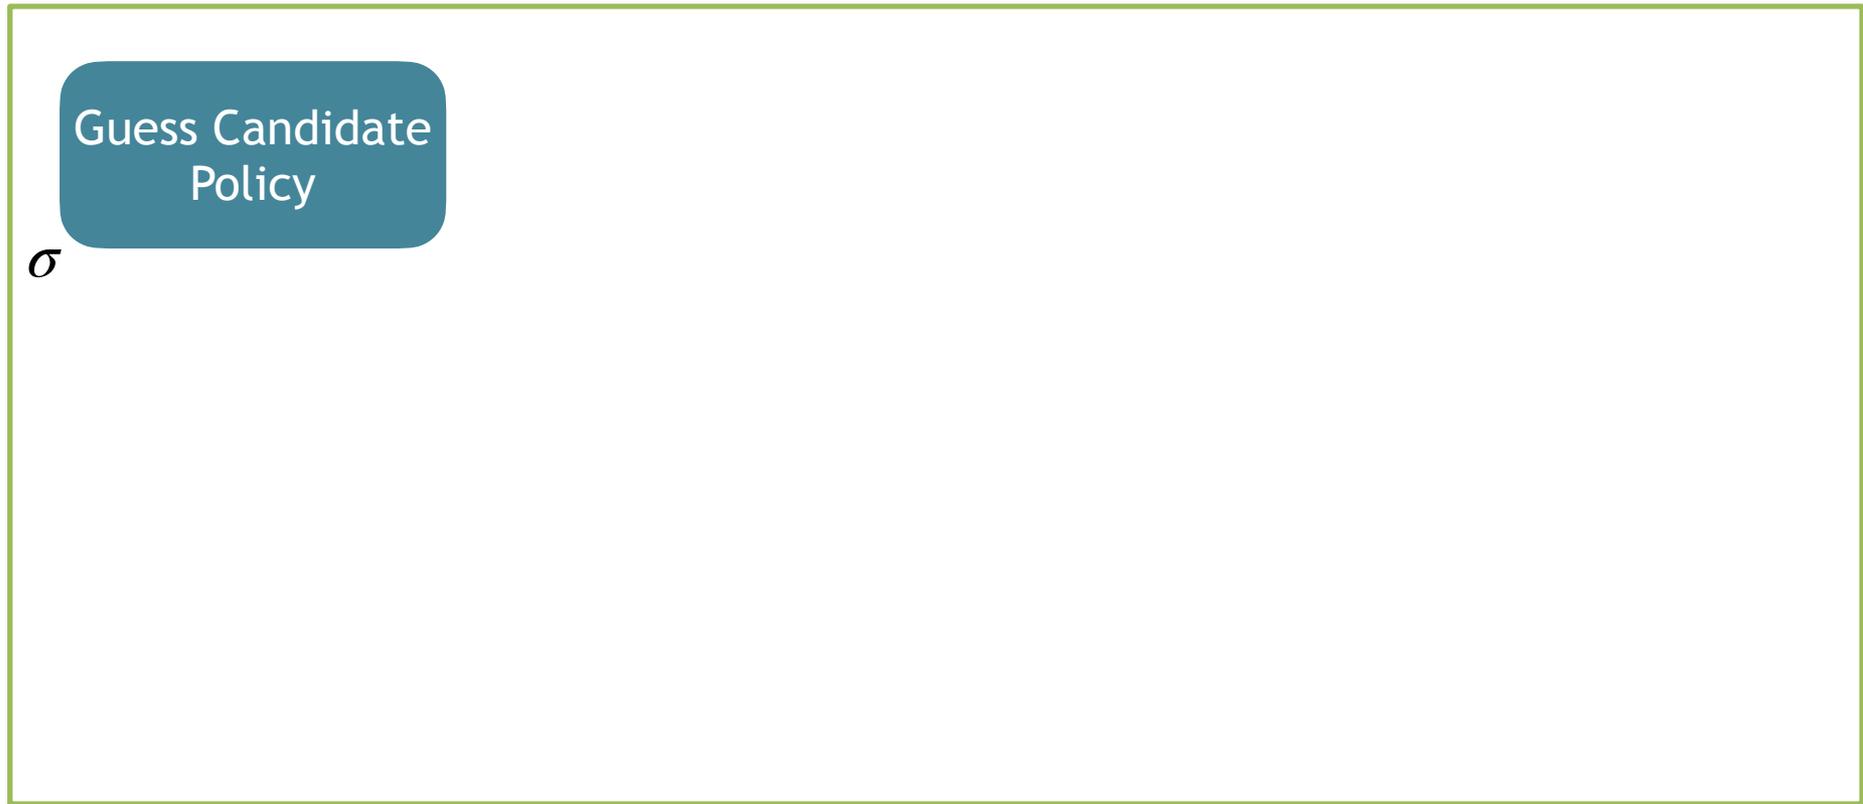
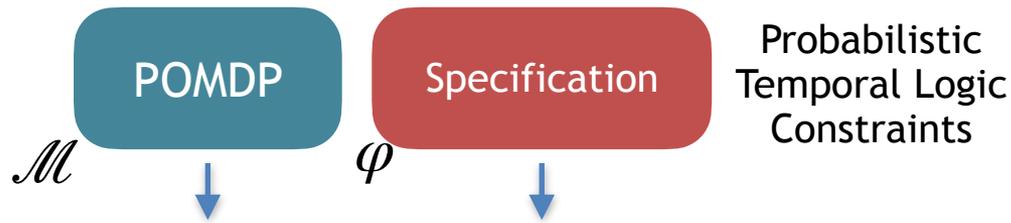
# The Problem: Policy Synthesis



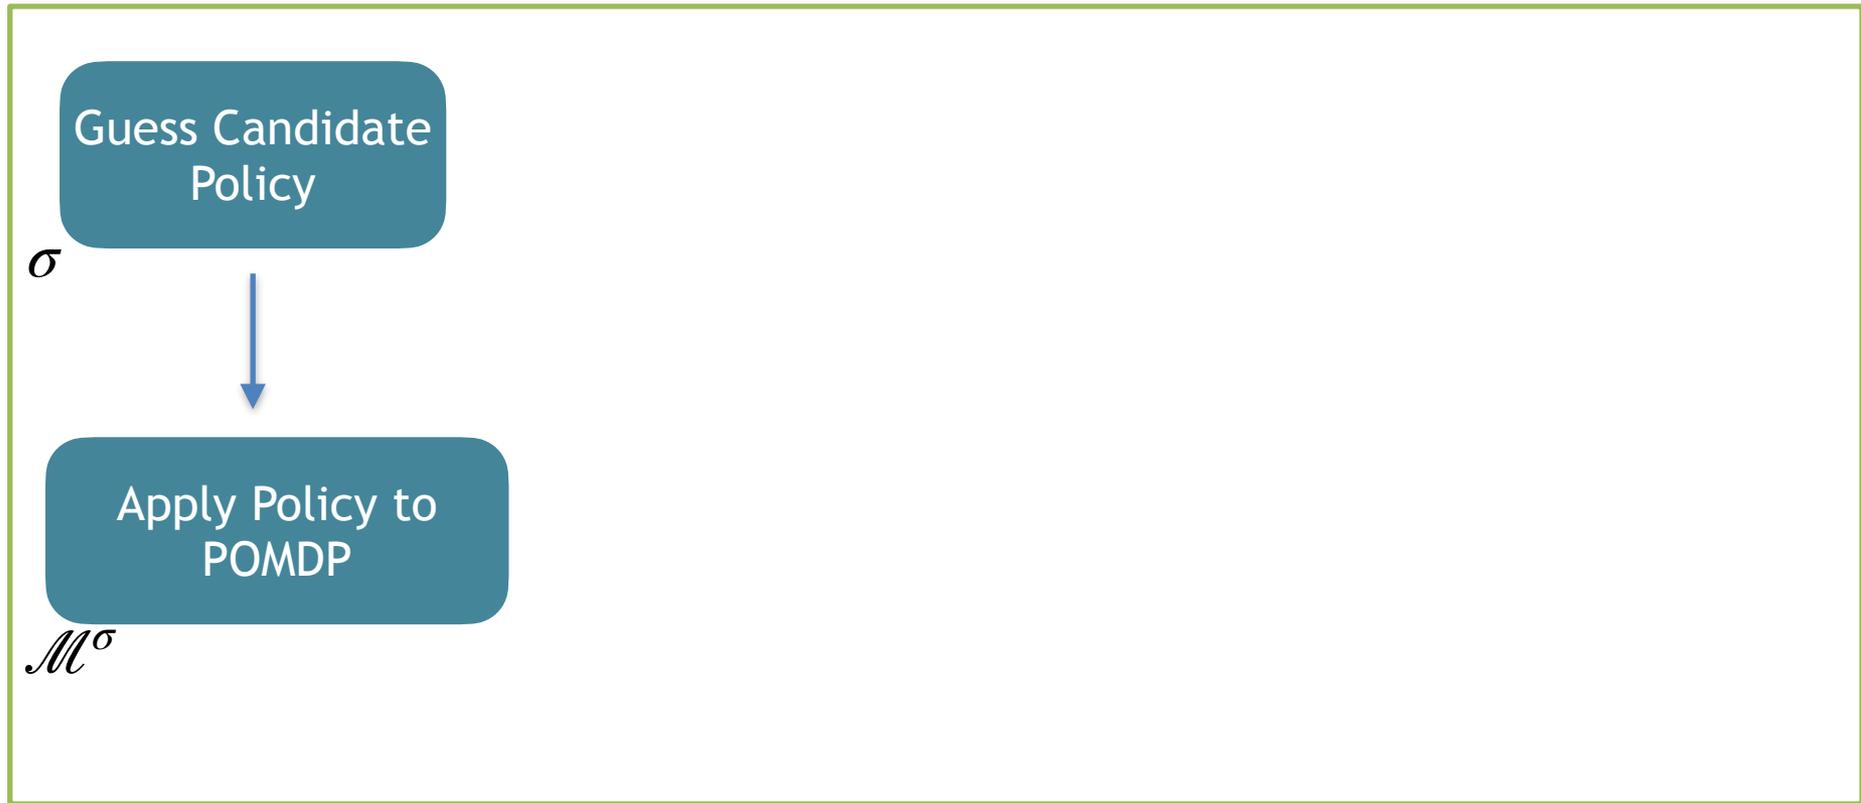
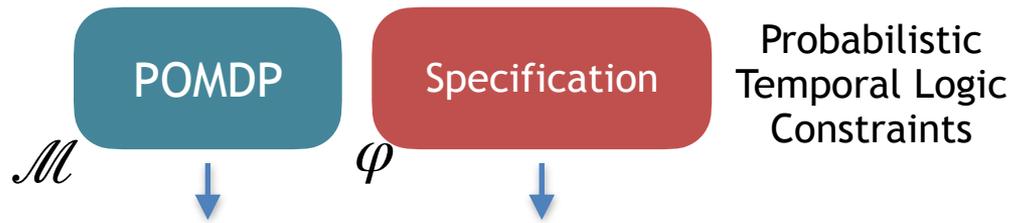
# The Problem: Policy Synthesis



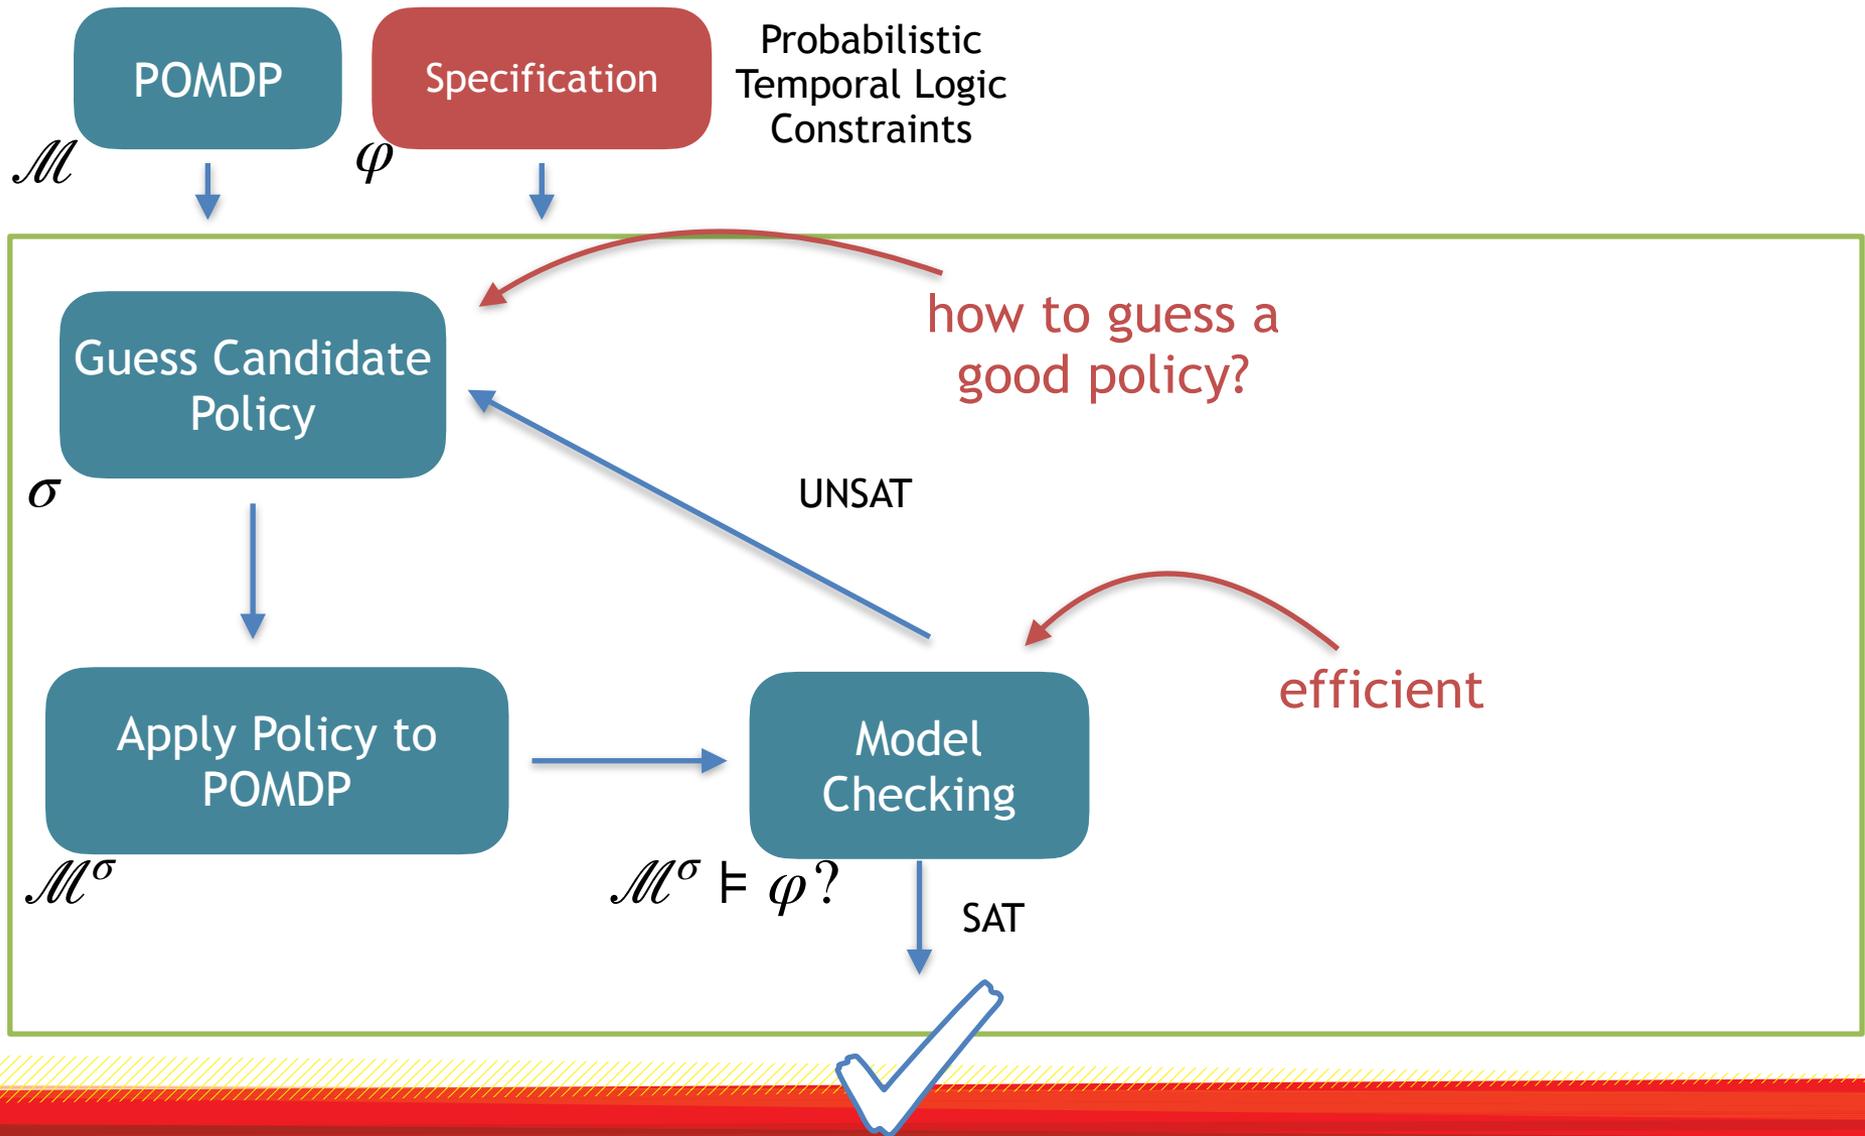
# Be Lazy: Guess a Policy and Verify!



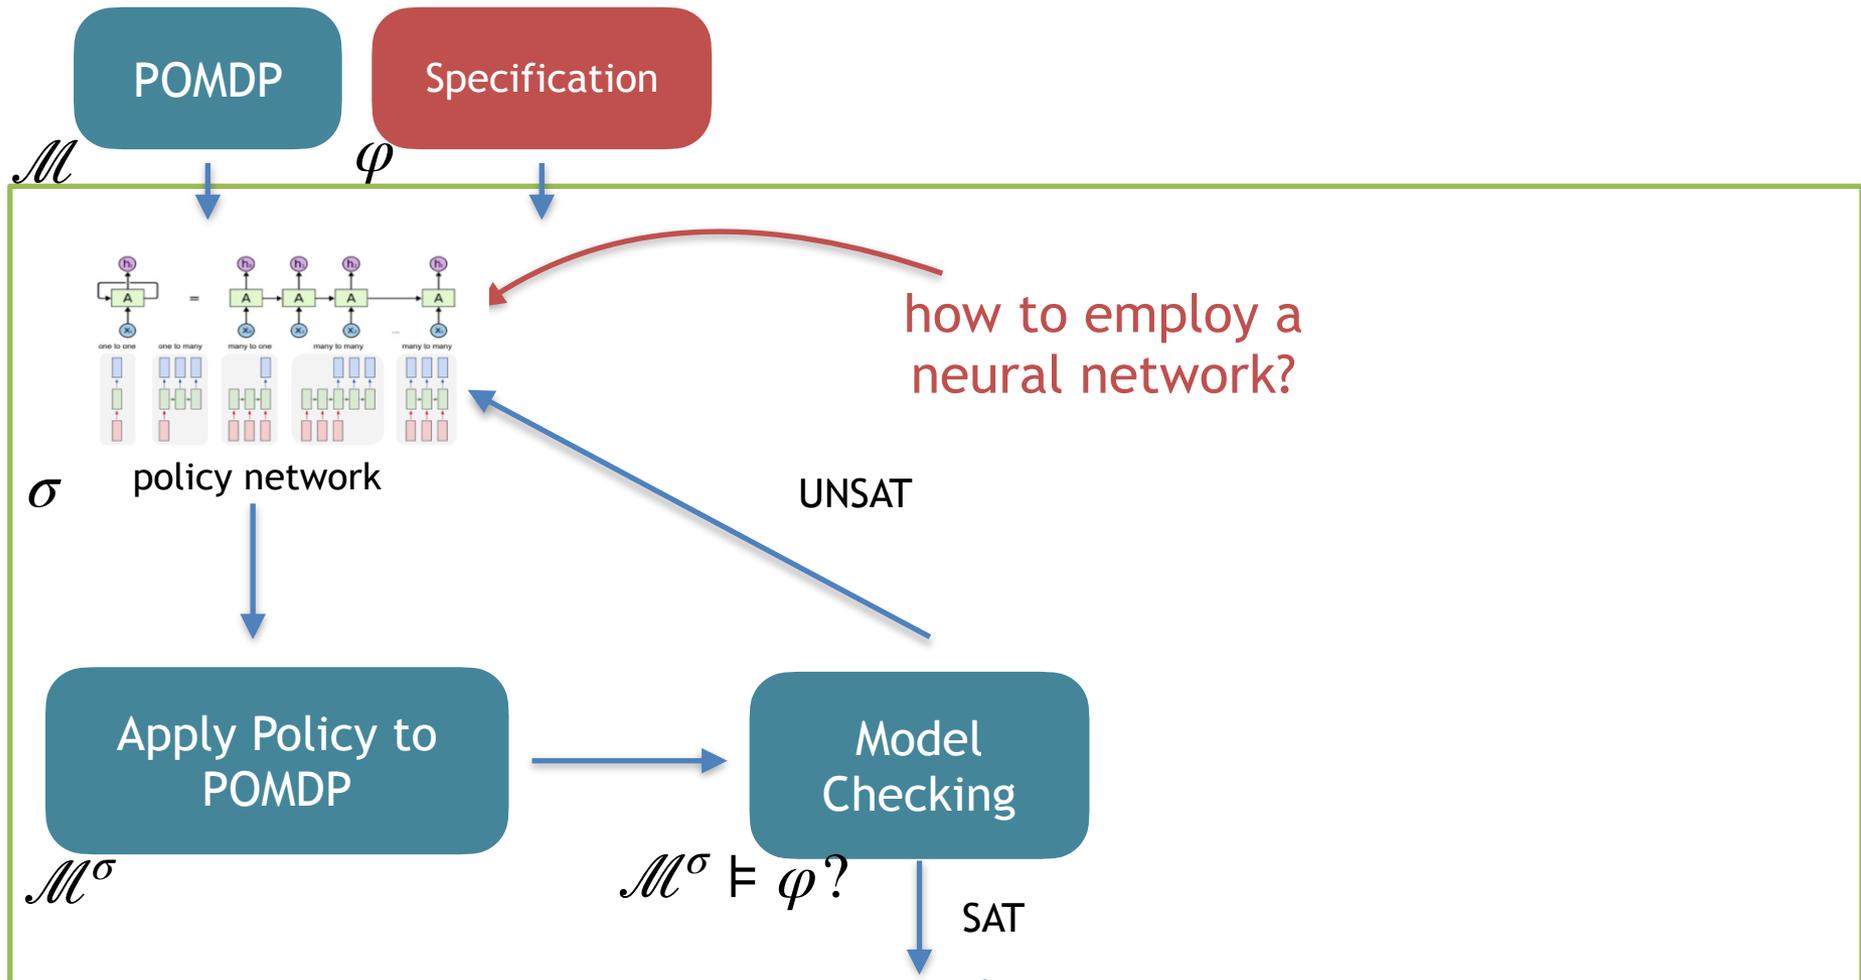
# Be Lazy: Guess a Policy and Verify!



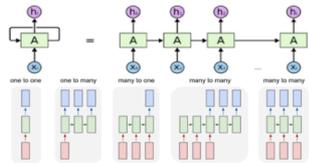
# Be Lazy: Guess a Policy and Verify!



# Let Machine Learning do the Guessing?

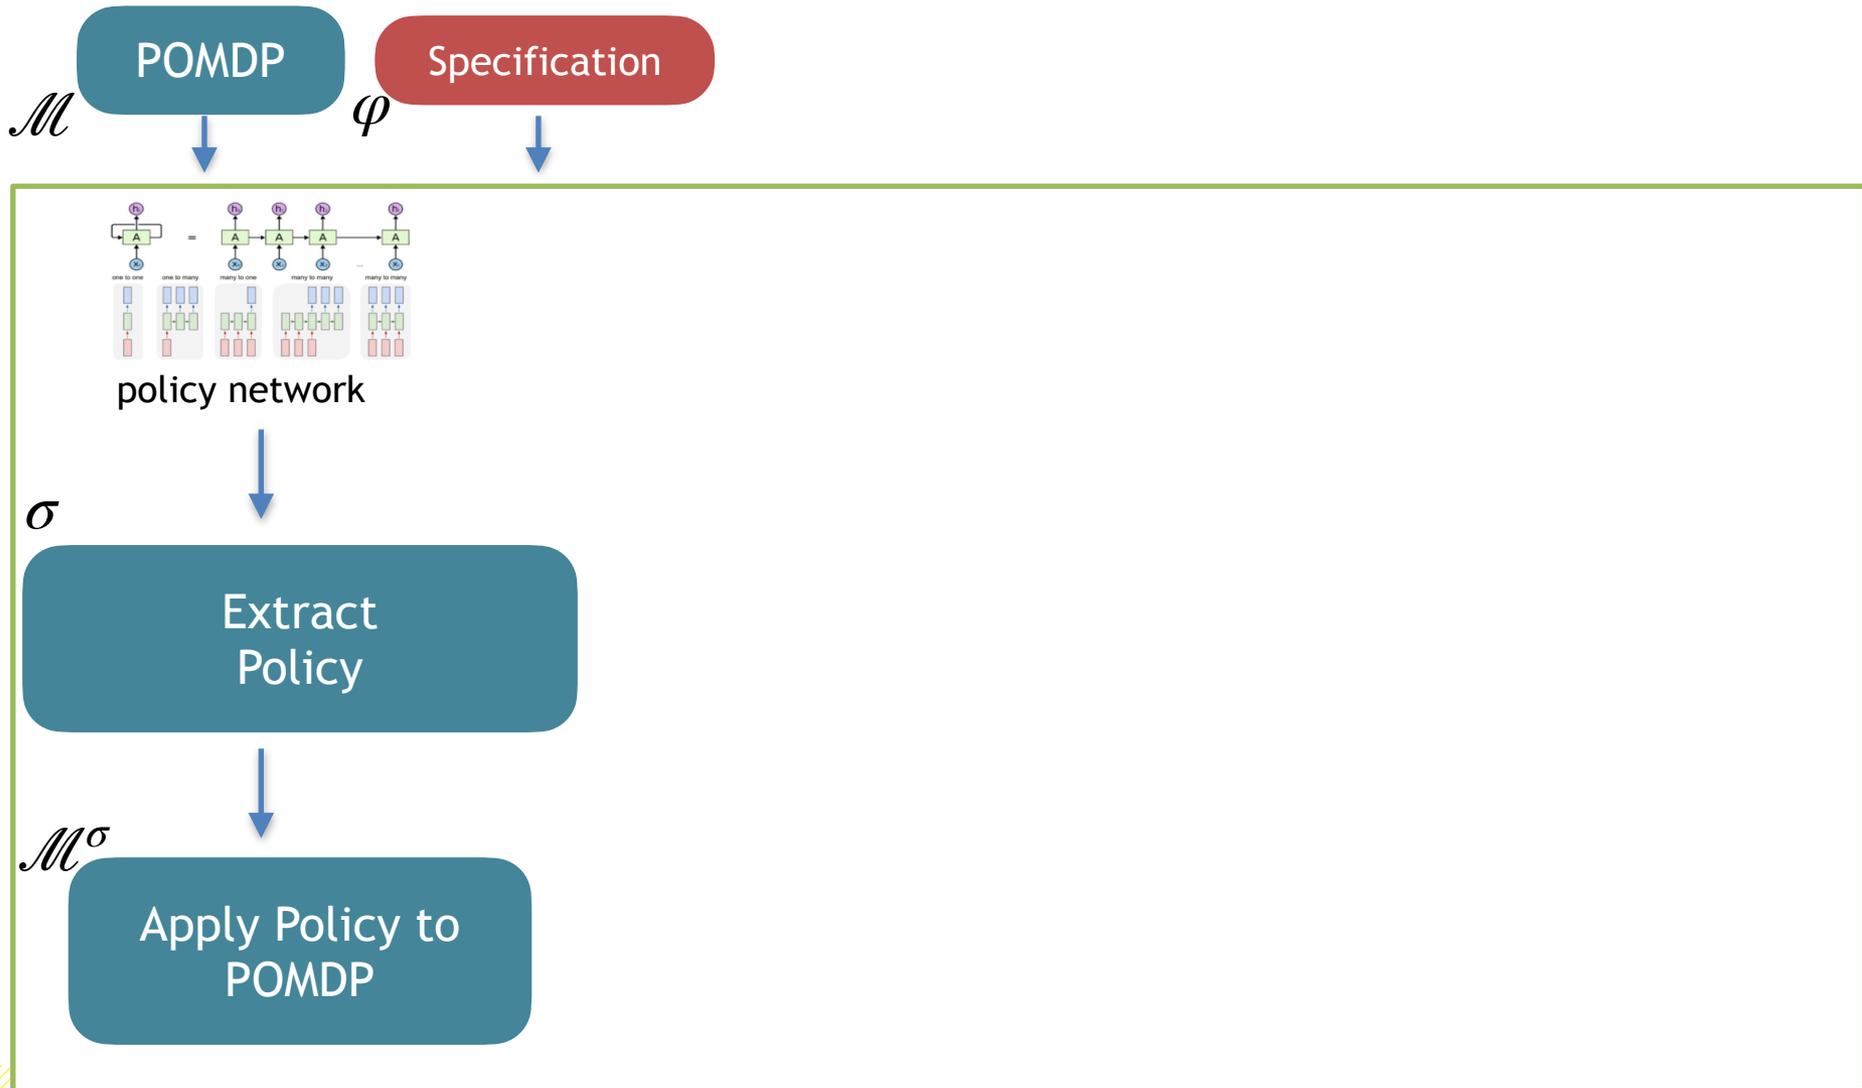


# RNN Policy Improvement

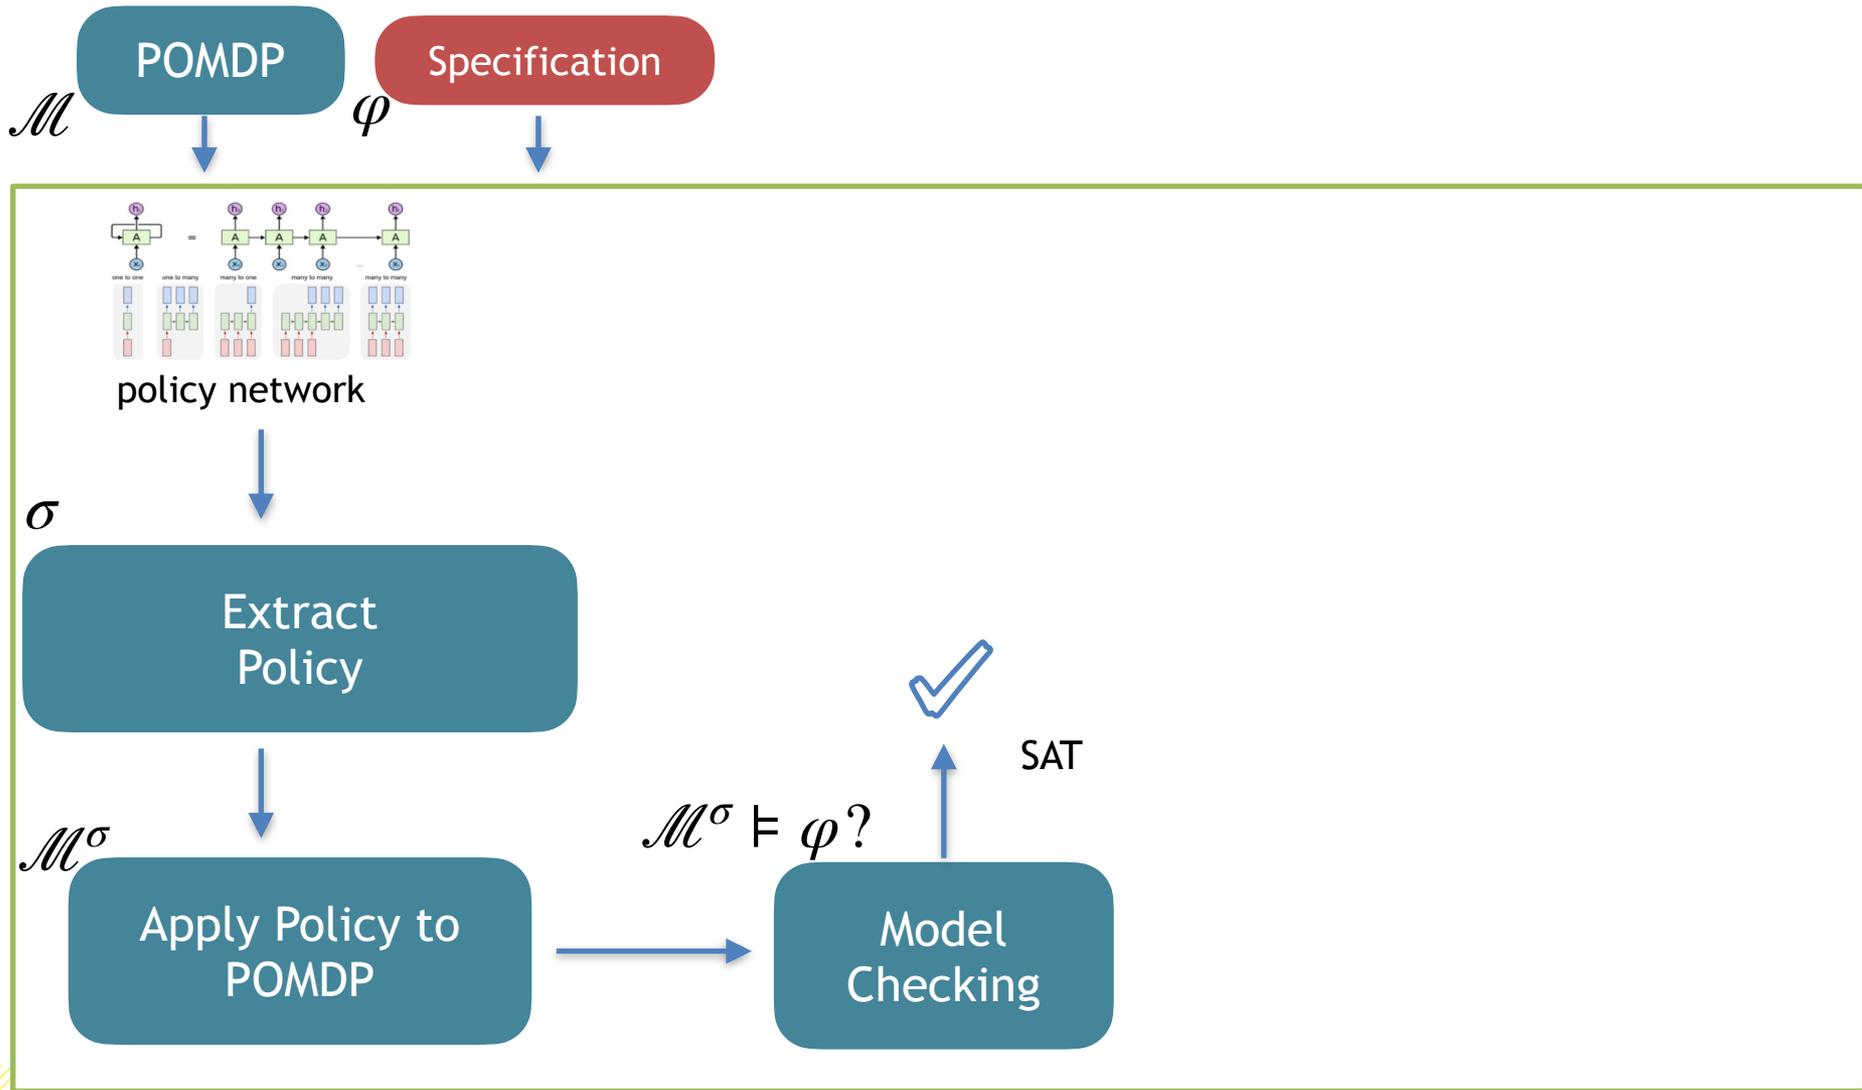


policy network

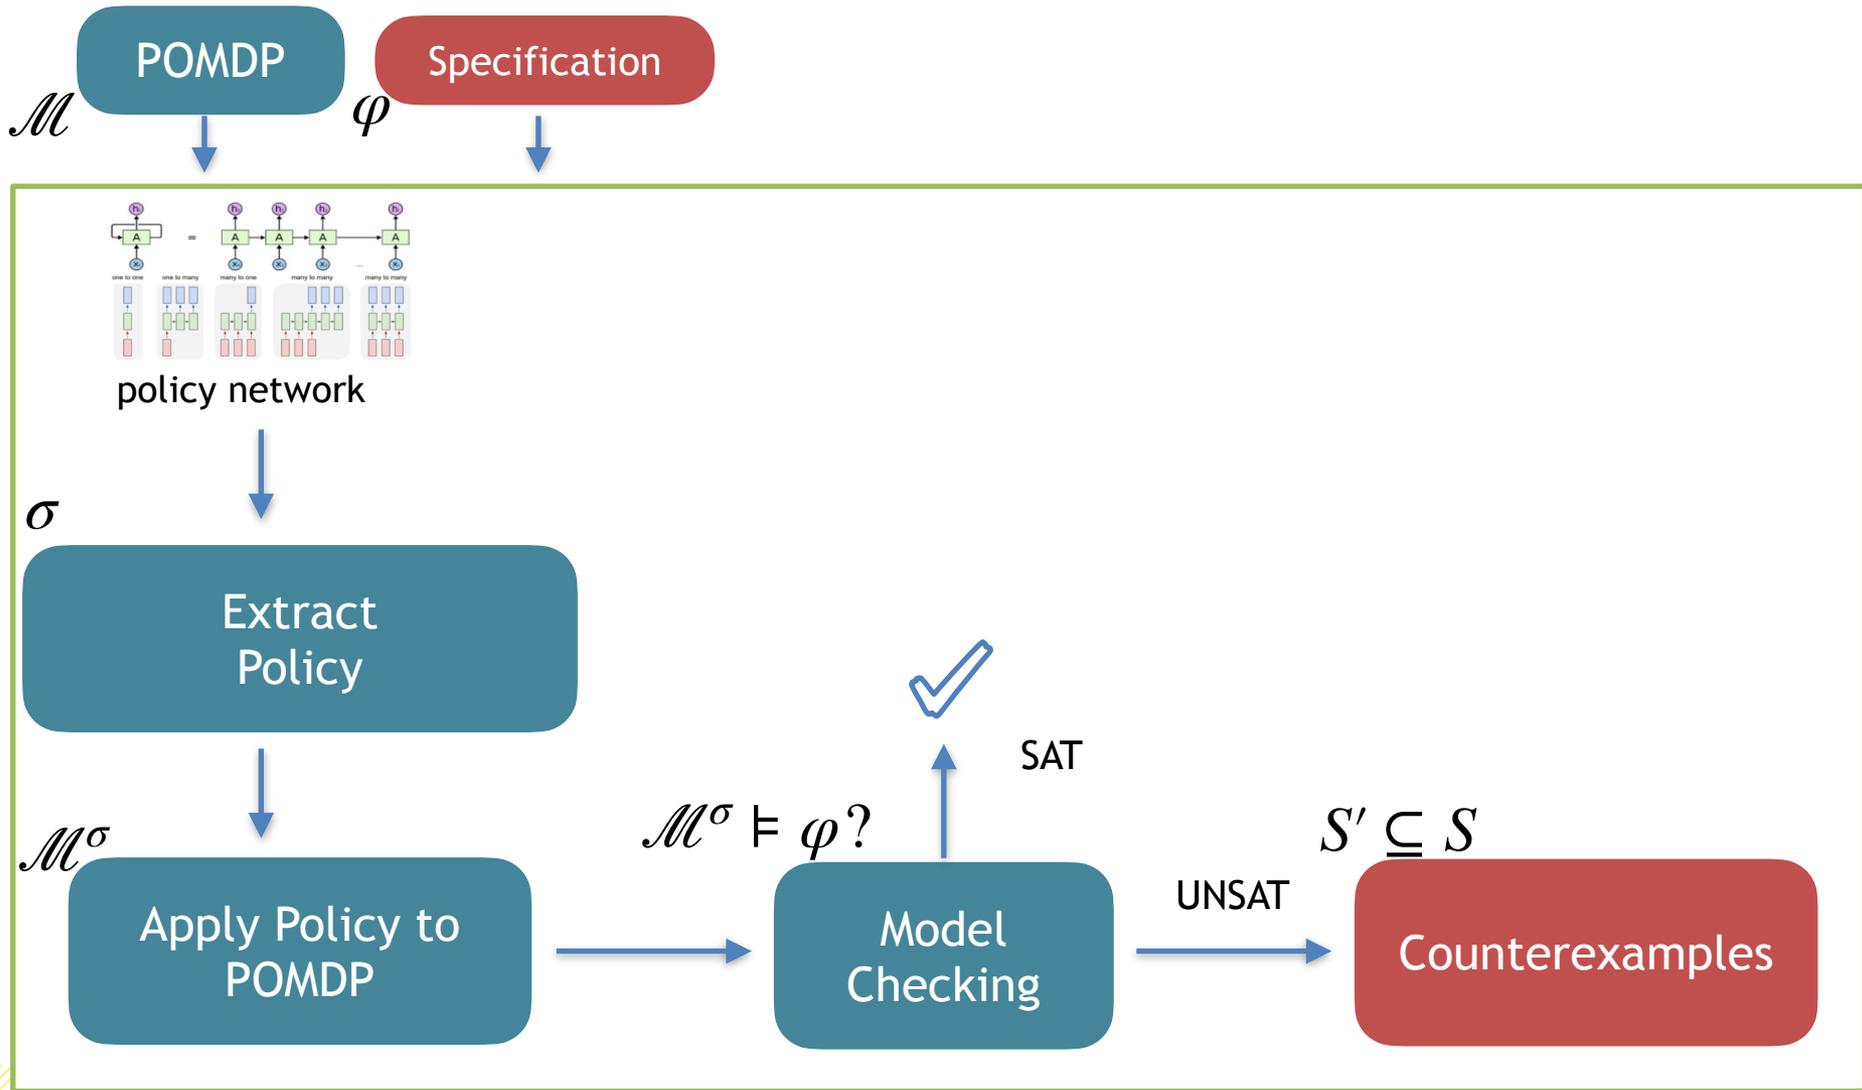
# RNN Policy Improvement



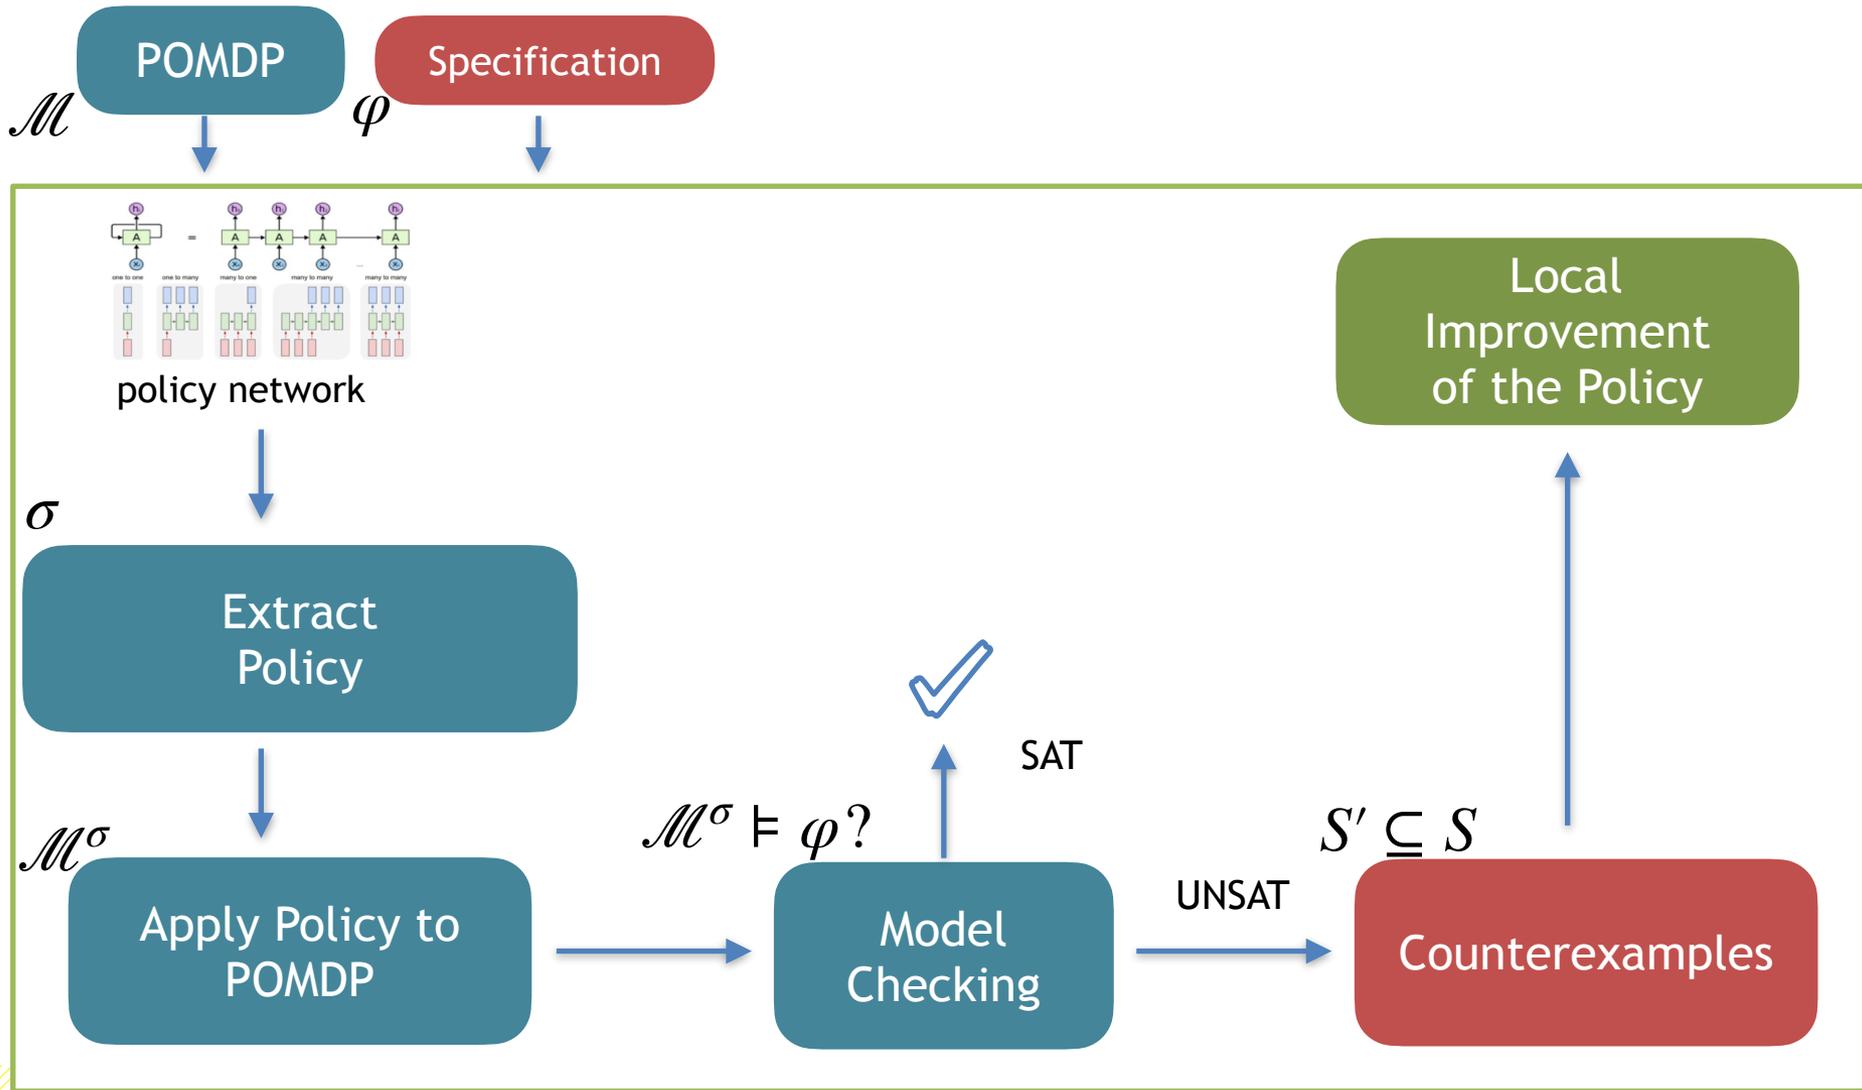
# RNN Policy Improvement



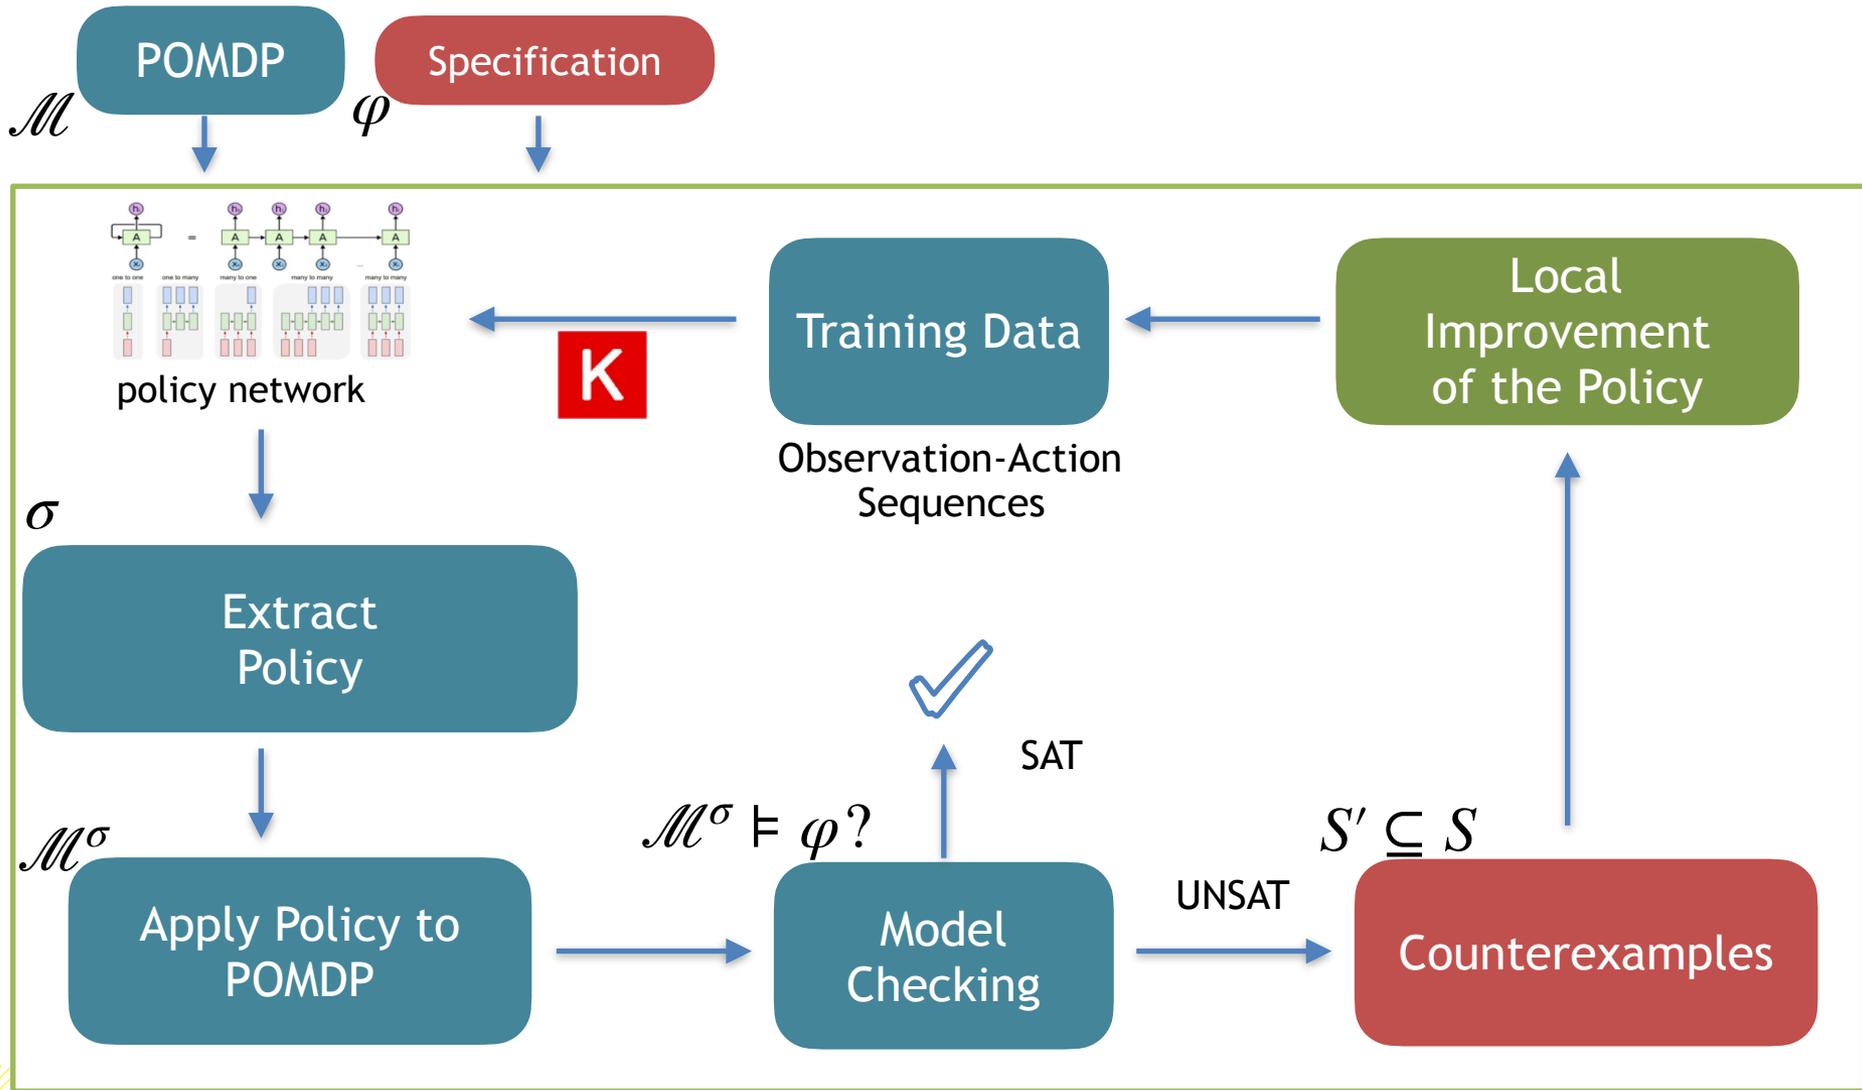
# RNN Policy Improvement



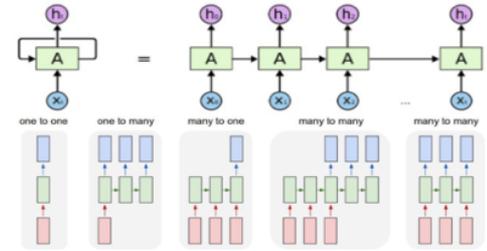
# RNN Policy Improvement



# RNN Policy Improvement



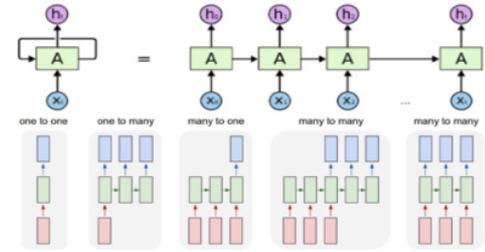
# Learning Strategies with RNNs



# Learning Strategies with RNNs

## Recurrent Neural Network

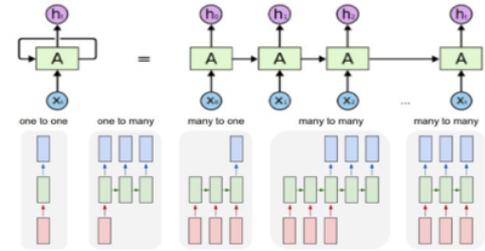
- long short-term memory (LSTM) architecture to learn dependencies in sequential data
- trained with observation-action sequences  $ObsSeq_{fin}$
- policy network  $\sigma: ObsSeq_{fin} \rightarrow Distr(Act)$
- observations are input labels, actions are output labels



# Learning Strategies with RNNs

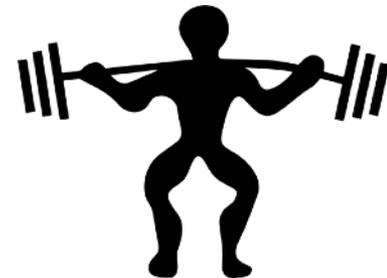
## Recurrent Neural Network

- long short-term memory (LSTM) architecture to learn dependencies in sequential data
- trained with observation-action sequences  $ObsSeq_{fin}$
- policy network  $\sigma: ObsSeq_{fin} \rightarrow Distr(Act)$
- observations are input labels, actions are output labels

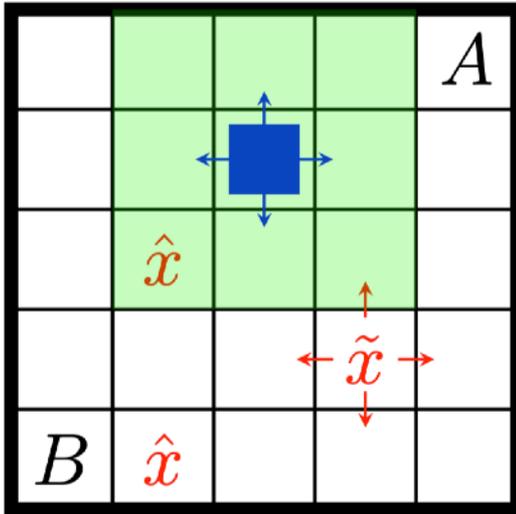


## Initial Training

- compute optimal MDP policy
- generate (possible) observation-action sequences



# Experiments - LTL



Problem	$ S $	$ Act $	$ Z $
Navigation ( $c$ )	$c^4$	4	256
Delivery ( $c$ )	$c^2$	4	256
Slippery ( $c$ )	$c^2$	4	256
Maze( $c$ )	$3c + 8$	4	7
Grid( $c$ )	$c^2$	4	2
RockSample[4, 4]	257	9	2
RockSample[5, 5]	801	10	2
RockSample[7, 8]	12545	13	2

Problem	States	Type, $\varphi$	RNN-based Synthesis		PRISM-POMDP	
			Res.	Time (s)	Res.	Time (s)
Navigation (3)	333	$\mathbb{P}_{\max}^{\mathcal{M}}, \varphi_1$	0.74	<b>14.16</b>	<b>0.84</b>	73.88
Navigation (4)	1088	$\mathbb{P}_{\max}^{\mathcal{M}}, \varphi_1$	0.82	<b>22.67</b>	<b>0.93</b>	1034.64
Navigation (4) [2-FSC]	13373	$\mathbb{P}_{\max}^{\mathcal{M}}, \varphi_1$	0.91	47.26	–	–
Navigation (4) [4-FSC]	26741	$\mathbb{P}_{\max}^{\mathcal{M}}, \varphi_1$	0.92	59.42	–	–
Navigation (4) [8-FSC]	53477	$\mathbb{P}_{\max}^{\mathcal{M}}, \varphi_1$	<b>0.92</b>	85.26	–	–
Navigation (5)	2725	$\mathbb{P}_{\max}^{\mathcal{M}}, \varphi_1$	0.91	<b>34.34</b>	MO	MO
Navigation (5) [2-FSC]	33357	$\mathbb{P}_{\max}^{\mathcal{M}}, \varphi_1$	0.92	115.16	–	–
Navigation (5) [4-FSC]	66709	$\mathbb{P}_{\max}^{\mathcal{M}}, \varphi_1$	0.92	159.61	–	–
Navigation (5) [8-FSC]	133413	$\mathbb{P}_{\max}^{\mathcal{M}}, \varphi_1$	<b>0.92</b>	250.91	–	–
Navigation (10)	49060	$\mathbb{P}_{\max}^{\mathcal{M}}, \varphi_1$	0.79	<b>822.87</b>	MO	MO
Navigation (10) [2-FSC]	475053	$\mathbb{P}_{\max}^{\mathcal{M}}, \varphi_1$	0.83	1185.41	–	–
Navigation (10) [4-FSC]	950101	$\mathbb{P}_{\max}^{\mathcal{M}}, \varphi_1$	<b>0.85</b>	1488.77	–	–
Navigation (10) [8-FSC]	1900197	$\mathbb{P}_{\max}^{\mathcal{M}}, \varphi_1$	0.81	1805.22	–	–
Navigation (15)	251965	$\mathbb{P}_{\max}^{\mathcal{M}}, \varphi_1$	<b>0.91</b>	<b>1271.80*</b>	MO	MO
Navigation (20)	798040	$\mathbb{P}_{\max}^{\mathcal{M}}, \varphi_1$	<b>0.96</b>	<b>4712.25*</b>	MO	MO
Navigation (30)	4045840	$\mathbb{P}_{\max}^{\mathcal{M}}, \varphi_1$	<b>0.95</b>	<b>25191.05*</b>	MO	MO
Navigation (40)	–	$\mathbb{P}_{\max}^{\mathcal{M}}, \varphi_1$	TO	TO	MO	MO
Delivery (4) [2-FSC]	80	$\mathbb{E}_{\min}^{\mathcal{M}}, \varphi_2$	6.02	35.35	<b>6.0</b>	<b>28.53</b>
Delivery (5) [2-FSC]	125	$\mathbb{E}_{\min}^{\mathcal{M}}, \varphi_2$	8.11	<b>78.32</b>	<b>8.0</b>	102.41
Delivery (10) [2-FSC]	500	$\mathbb{E}_{\min}^{\mathcal{M}}, \varphi_2$	<b>18.13</b>	<b>120.34</b>	MO	MO
Slippery (4) [2-FSC]	460	$\mathbb{P}_{\max}^{\mathcal{M}}, \varphi_3$	0.78	67.51	<b>0.90</b>	<b>5.10</b>
Slippery (5) [2-FSC]	730	$\mathbb{P}_{\max}^{\mathcal{M}}, \varphi_3$	0.89	84.32	<b>0.93</b>	<b>83.24</b>
Slippery (10) [2-FSC]	2980	$\mathbb{P}_{\max}^{\mathcal{M}}, \varphi_3$	<b>0.98</b>	<b>119.14</b>	MO	MO
Slippery (20) [2-FSC]	11980	$\mathbb{P}_{\max}^{\mathcal{M}}, \varphi_3$	<b>0.99</b>	<b>1580.42</b>	MO	MO

# Experiments - Standard POMDPs

Problem	Type	RNN-based Synthesis			PRISM-POMDP		pomdpSolve	
		States	Res	Time (s)	Res	Time (s)	Res	Time (s)
Maze (1)	$E_{\min}^{\mathcal{M}}$	68	4.31	31.70	<b>4.30</b>	<b>0.09</b>	4.30	0.30
Maze (2)	$E_{\min}^{\mathcal{M}}$	83	5.31	46.65	5.23	2.176	<b>5.23</b>	<b>0.67</b>
Maze (3)	$E_{\min}^{\mathcal{M}}$	98	8.10	58.75	7.13	38.82	<b>7.13</b>	<b>2.39</b>
Maze (4)	$E_{\min}^{\mathcal{M}}$	113	11.53	58.09	8.58	543.06	<b>8.58</b>	<b>7.15</b>
Maze (5)	$E_{\min}^{\mathcal{M}}$	128	14.40	<b>68.09</b>	13.00	4110.50	<b>12.04</b>	132.12
Maze (6)	$E_{\min}^{\mathcal{M}}$	143	22.34	<b>71.89</b>	MO	MO	<b>18.52</b>	1546.02
Maze (10)	$E_{\min}^{\mathcal{M}}$	203	100.21	<b>158.33</b>	MO	MO	MO	MO
Grid (3)	$E_{\min}^{\mathcal{M}}$	165	2.90	38.94	2.88	2.332	<b>2.88</b>	<b>0.07</b>
Grid (4)	$E_{\min}^{\mathcal{M}}$	381	4.32	79.99	4.13	1032.53	<b>4.13</b>	<b>0.77</b>
Grid (5)	$E_{\min}^{\mathcal{M}}$	727	6.623	91.42	MO	MO	<b>5.42</b>	<b>1.94</b>
Grid (10)	$E_{\min}^{\mathcal{M}}$	5457	<b>13.630</b>	<b>268.40</b>	MO	MO	MO	MO
RockSample[4, 4]	$E_{\max}^{\mathcal{M}}$	2432	17.71	35.35	N/A	N/A	<b>18.04</b>	<b>0.43</b>
RockSample[5, 5]	$E_{\max}^{\mathcal{M}}$	8320	18.40	<b>43.74</b>	N/A	N/A	<b>19.23</b>	621.28
RockSample[7, 8]	$E_{\max}^{\mathcal{M}}$	166656	20.32	<b>860.53</b>	N/A	N/A	<b>21.64</b>	20458.41

# Improving the Policy

# Improving the Policy

- Identify **critical decisions** that lead to states with high probability of **violating** the specification.

# Improving the Policy

- Identify **critical decisions** that lead to states with high probability of **violating** the specification.
- For each observation with critical decision, **minimize** the number of different critical actions.

Local linear program

$$\begin{aligned} & \max_{\gamma(z)(a), a \in Act} \min_{s \in S} p_s & (1) \\ \text{subject to} & \\ \forall s \in O^{-1}(z). & \quad p_s = \sum_{a \in Act} \gamma(z)(a) \cdot \sum_{s' \in S} \mathcal{P}(s, a, s') \cdot p^*(s') \end{aligned}$$

# Improving the Policy

- Identify **critical decisions** that lead to states with high probability of **violating** the specification.
- For each observation with critical decision, **minimize** the number of different critical actions.
- **Retrain** with the new (locally improved) policy.

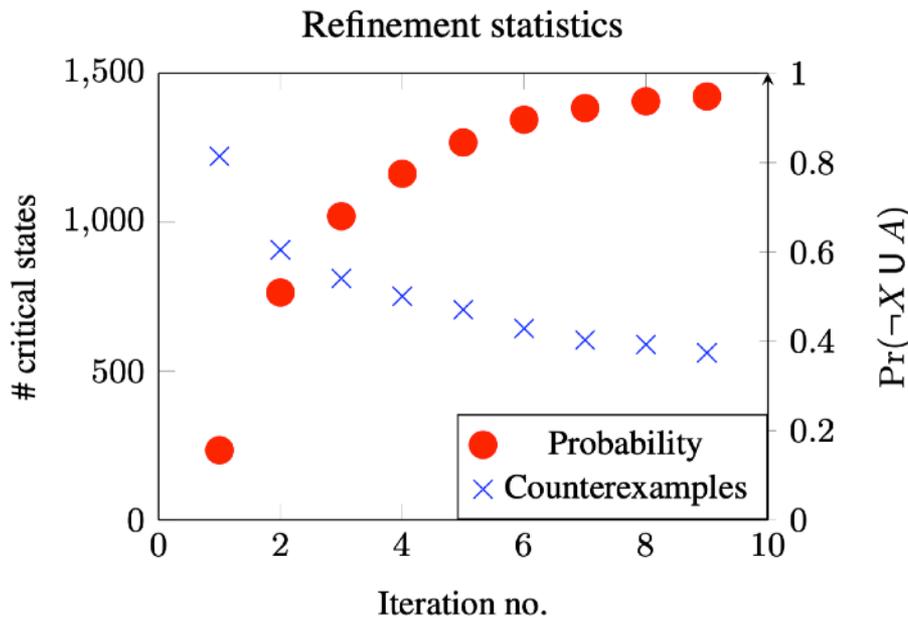
Local linear program

$$\begin{aligned} & \max_{\gamma(z)(a), a \in Act} \min_{s \in S} p_s & (1) \\ \text{subject to} & \\ \forall s \in O^{-1}(z). & \quad p_s = \sum_{a \in Act} \gamma(z)(a) \cdot \sum_{s' \in S} \mathcal{P}(s, a, s') \cdot p^*(s') \end{aligned}$$

# Improving the Policy

- Identify **critical decisions** that lead to states with high probability of **violating** the specification.
- For each observation with critical decision, **minimize** the number of different critical actions.
- **Retrain** with the new (locally improved) policy.

Local linear program

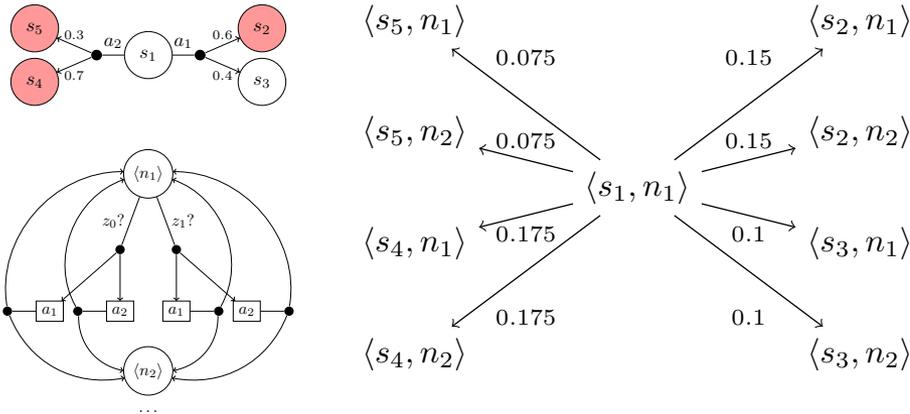


$$\begin{aligned}
 & \max_{\gamma(z)(a), a \in Act} \min_{s \in S} p_s \\
 & \text{subject to} \\
 & \forall s \in O^{-1}(z). \quad p_s = \sum_{a \in Act} \gamma(z)(a) \cdot \sum_{s' \in S} \mathcal{P}(s, a, s') \cdot p^*(s')
 \end{aligned} \tag{1}$$

Even if specification is satisfied, there may be critical states and decisions!

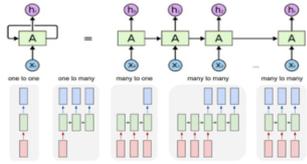
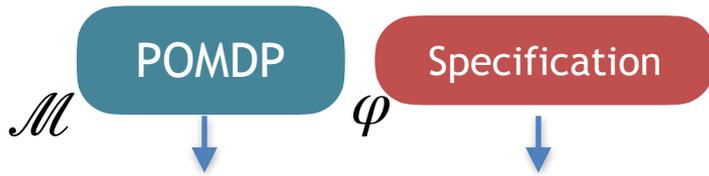
# Finite-memory Strategies (FSC)

- Encode **finite memory** directly into the state space:



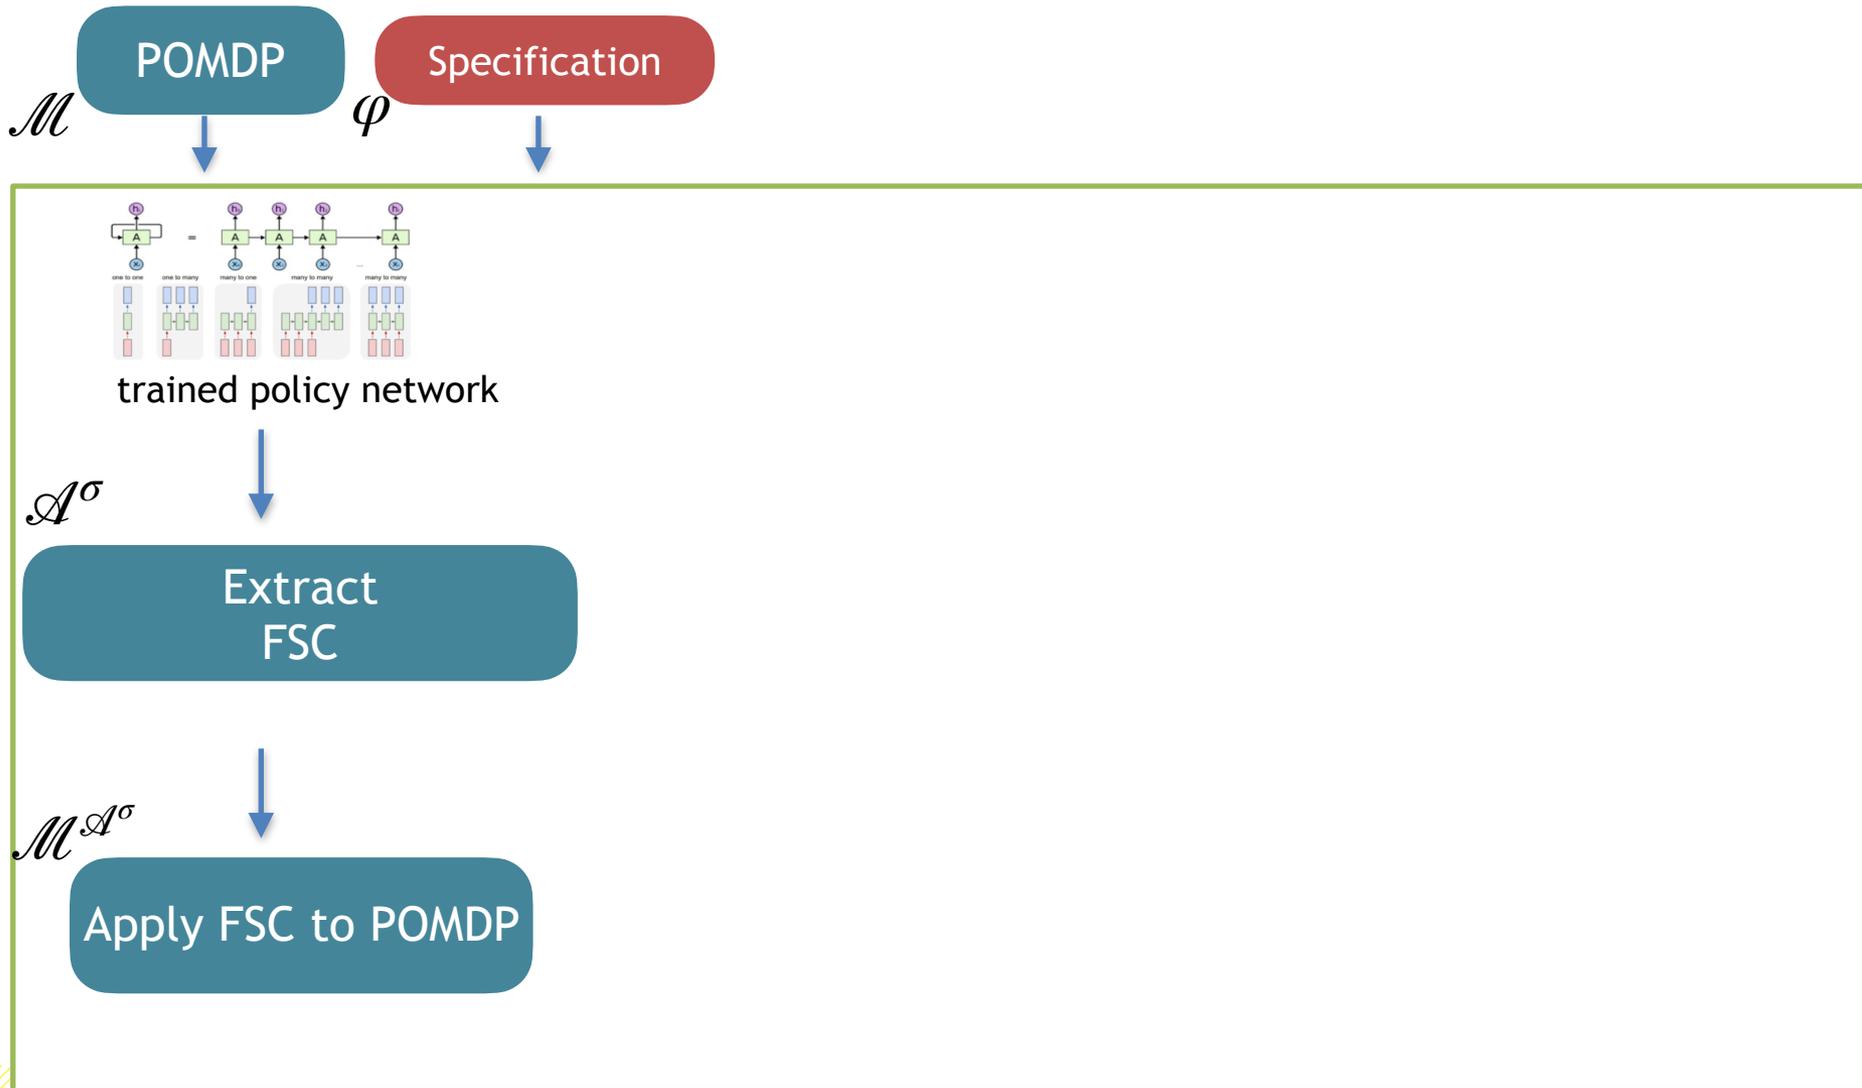
- So far: Predefined memory update, internal memory state of the RNN largely ignored
- How to infer a **memory-update function** to construct an FSC from the RNN?

# RNN Policy Improvement

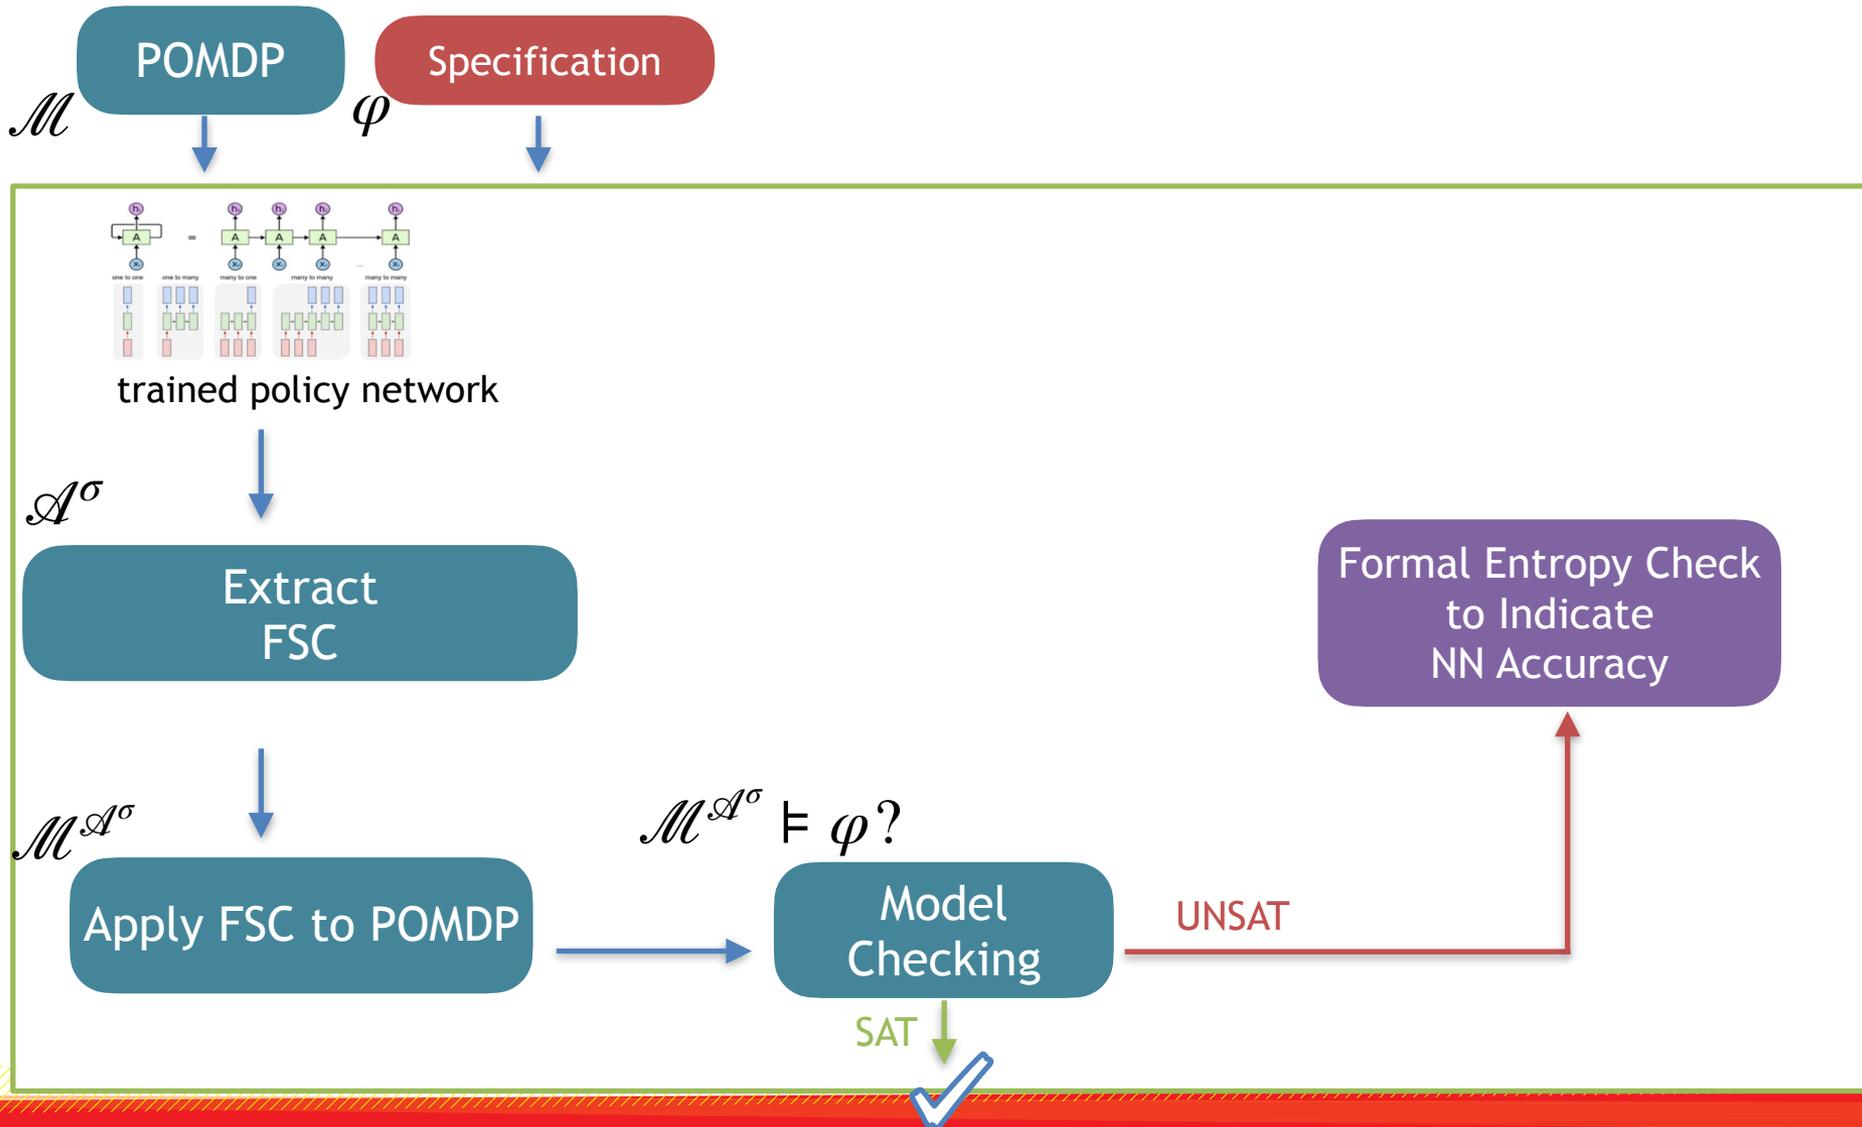


trained policy network

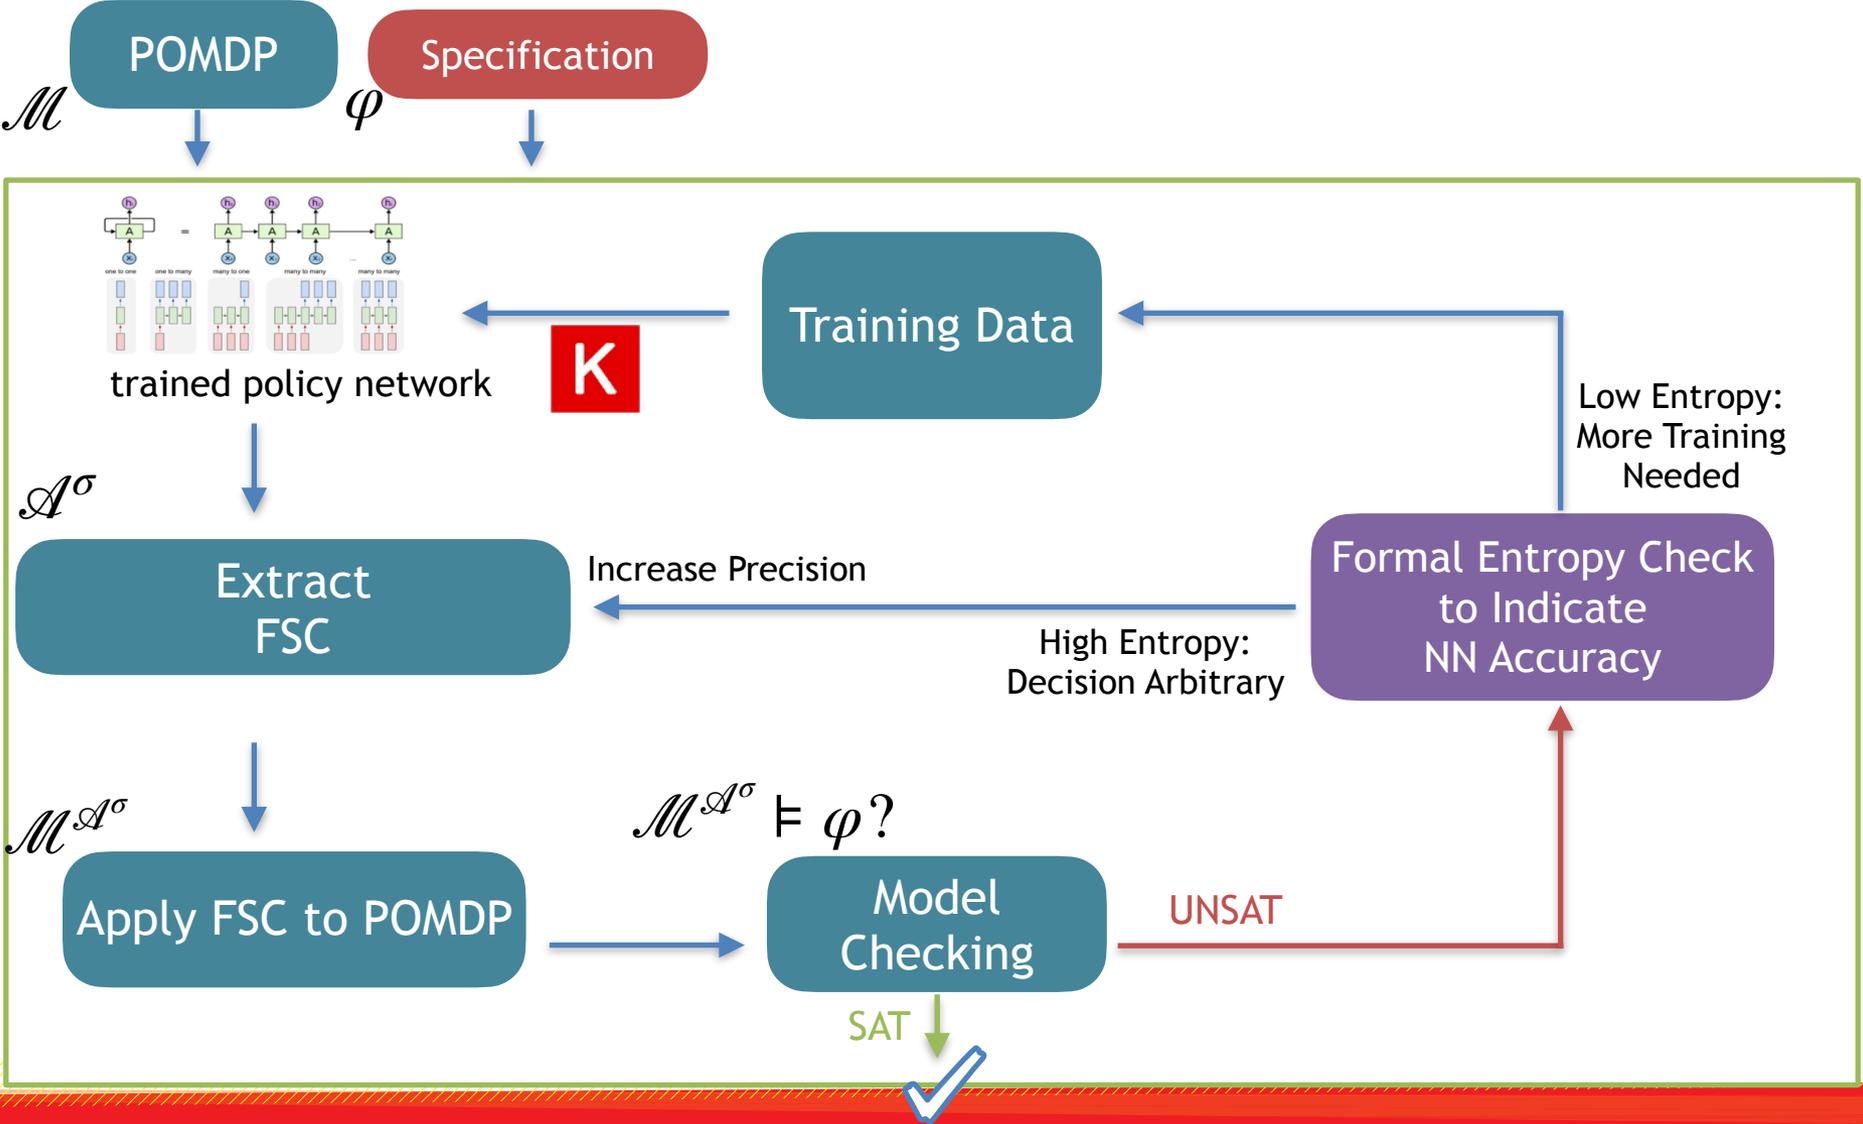
# RNN Policy Improvement



# RNN Policy Improvement

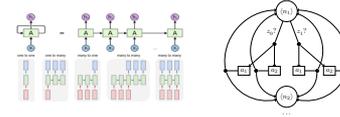


# RNN Policy Improvement



# Extracting and Improving FSCs

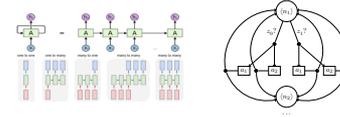
## Quantized Bottleneck Insertion [1]



- Discretization of the continuous memory structure of the RNN
- Mapping to a **pre-defined** number of finite memory state

[1] Anurag Koul, Alan Fern, and Sam Greydanus. Learning finite state representations of recurrent policy networks. In ICLR, 2019

# Extracting and Improving FSCs



## Quantized Bottleneck Insertion [1]

- Discretization of the continuous memory structure of the RNN
- Mapping to a **pre-defined** number of finite memory state

## Entropy Check

$$H(\text{Crit}_{\mathcal{A}}^{\mathcal{M}}) = \frac{1}{|\text{Crit}_{\mathcal{A}}^{\mathcal{M}}|} \sum_{(n,s) \in \text{Crit}_{\mathcal{A}}^{\mathcal{M}}} H(n, s)$$

- Localized to critical states
- **Entropy high**: NN is likely extrapolating
- **Entropy low**: NN (may be) well trained, increase precision of FSC extraction

[1] Anurag Koul, Alan Fern, and Sam Greydanus. Learning finite state representations of recurrent policy networks. In ICLR, 2019

# Experiments

Problem	S	Z	Type	Extraction Approach			Handcrafted			PRISM-POMDP		SolvePOMDP	
				Memory	Value	Time (s)	Memory	Value	Time (s)	Value	Time (s)	Value	Time (s)
Maze(1)	11	7	Min	2	4.33	80.31	2	4.31	30.70	<b>4.30</b>	<b>0.09</b>	4.30	0.30
Maze(2)	14	7	Min	3	5.34	114.23	3	5.31	46.65	5.23	2.176	<b>5.23</b>	<b>0.67</b>
Maze(5)	23	7	Min	3	13.29	160.12	6	14.40	68.09	13.00*	4110.50	<b>12.04</b>	<b>134.46</b>
Maze(10)	38	7	Min	5	<b>23.02</b>	210.01	11	100.21	<b>158.33</b>	MO	MO	MO	MO
Grid(3)	9	2	Min	3	2.90	87.31	2	2.90	38.94	2.88	2.332	<b>2.88</b>	<b>0.06</b>
Grid(4)	16	2	Min	7	4.20	124.31	3	4.32	79.99	4.13	1032.53	<b>4.13</b>	<b>0.73</b>
Grid(5)	25	2	Min	9	5.91	250.14	4	6.623	91.42	MO	MO	<b>5.42</b>	<b>1.97</b>
Grid(10)	100	2	Min	9	<b>12.92</b>	1031.21	9	13.63	<b>268.40</b>	MO	MO	MO	MO
Grid(25)	625	2	Min	16	<b>35.32</b>	6514.30	24	531.05	<b>622.31</b>	MO	MO	MO	MO
Navigation (4)	256	256	Max	8	0.92	160.32	8	0.92	80.26	<b>0.93*</b>	<b>1034.64</b>	NA	NA
Navigation (5)	625	256	Max	8	<b>0.95</b>	311.65	8	0.92	<b>253.11</b>	MO	MO	NA	NA
Navigation (10)	10 <sup>4</sup>	256	Max	8	<b>0.90</b>	2561.02	4	0.85	<b>1471.17</b>	MO	MO	NA	NA
Navigation (20)	1.6 × 10 <sup>5</sup>	256	Max	9	<b>0.98</b>	8173.03	4	0.96	<b>7068.24</b>	MO	MO	NA	NA

- scales to larger examples than other solvers (but not optimal!)
- higher-quality than handcrafted memory structure

# Experiments

Problem	S	Z	Type	Extraction Approach			Handcrafted			PRISM-POMDP		SolvePOMDP	
				Memory	Value	Time (s)	Memory	Value	Time (s)	Value	Time (s)	Value	Time (s)
Maze(1)	11	7	Min	2	4.33	80.31	2	4.31	30.70	<b>4.30</b>	<b>0.09</b>	4.30	0.30
Maze(2)	14	7	Min	3	5.34	114.23	3	5.31	46.65	5.23	2.176	<b>5.23</b>	<b>0.67</b>
Maze(5)	23	7	Min	3	13.29	160.12	6	14.40	68.09	13.00*	4110.50	<b>12.04</b>	<b>134.46</b>
Maze(10)	38	7	Min	5	<b>23.02</b>	210.01	11	100.21	<b>158.33</b>	MO	MO	MO	MO
Grid(3)	9	2	Min	3	2.90	87.31	2	2.90	38.94	2.88	2.332	<b>2.88</b>	<b>0.06</b>
Grid(4)	16	2	Min	7	4.20	124.31	3	4.32	79.99	4.13	1032.53	<b>4.13</b>	<b>0.73</b>
Grid(5)	25	2	Min	9	5.91	250.14	4	6.623	91.42	MO	MO	<b>5.42</b>	<b>1.97</b>
Grid(10)	100	2	Min	9	<b>12.92</b>	1031.21	9	13.63	<b>268.40</b>	MO	MO	MO	MO
Grid(25)	625	2	Min	16	<b>35.32</b>	6514.30	24	531.05	<b>622.31</b>	MO	MO	MO	MO
Navigation (4)	256	256	Max	8	0.92	160.32	8	0.92	80.26	<b>0.93*</b>	<b>1034.64</b>	NA	NA
Navigation (5)	625	256	Max	8	<b>0.95</b>	311.65	8	0.92	<b>253.11</b>	MO	MO	NA	NA
Navigation (10)	10 <sup>4</sup>	256	Max	8	<b>0.90</b>	2561.02	4	0.85	<b>1471.17</b>	MO	MO	NA	NA
Navigation (20)	1.6 × 10 <sup>5</sup>	256	Max	9	<b>0.98</b>	8173.03	4	0.96	<b>7068.24</b>	MO	MO	NA	NA

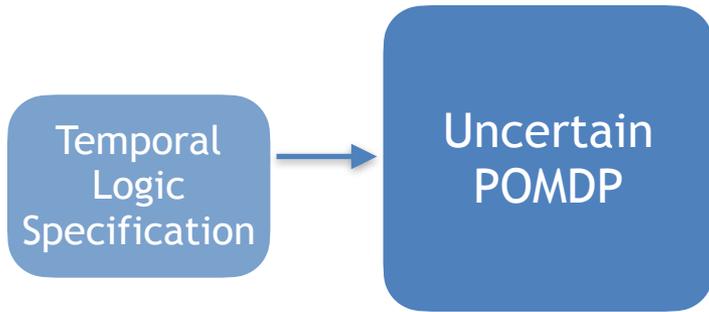
- scales to larger examples than other solvers (but not optimal!)
- higher-quality than handcrafted memory structure



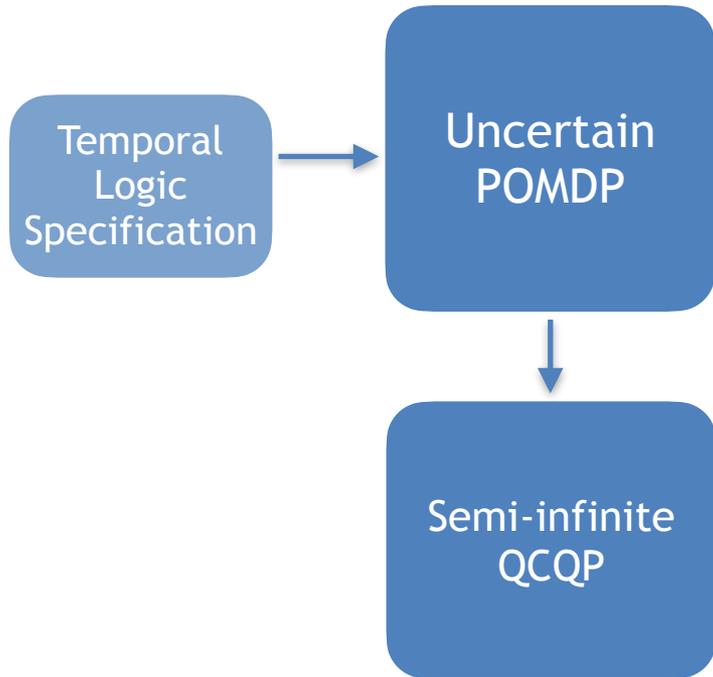




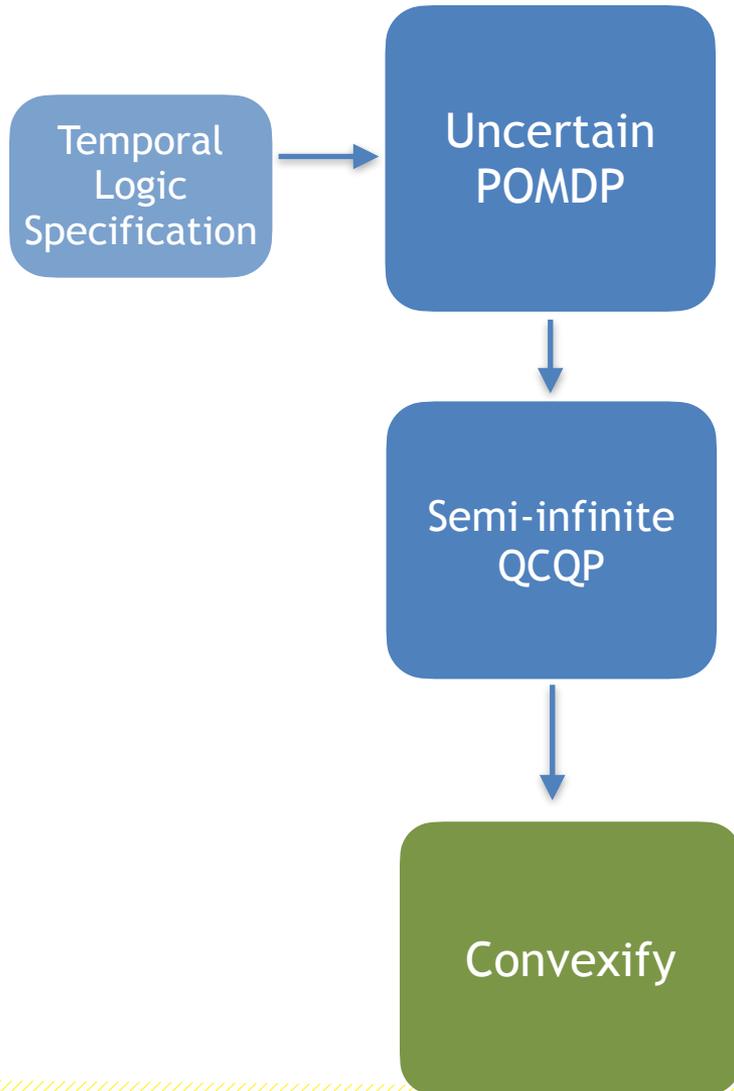
# Story - Convex-concave Procedure



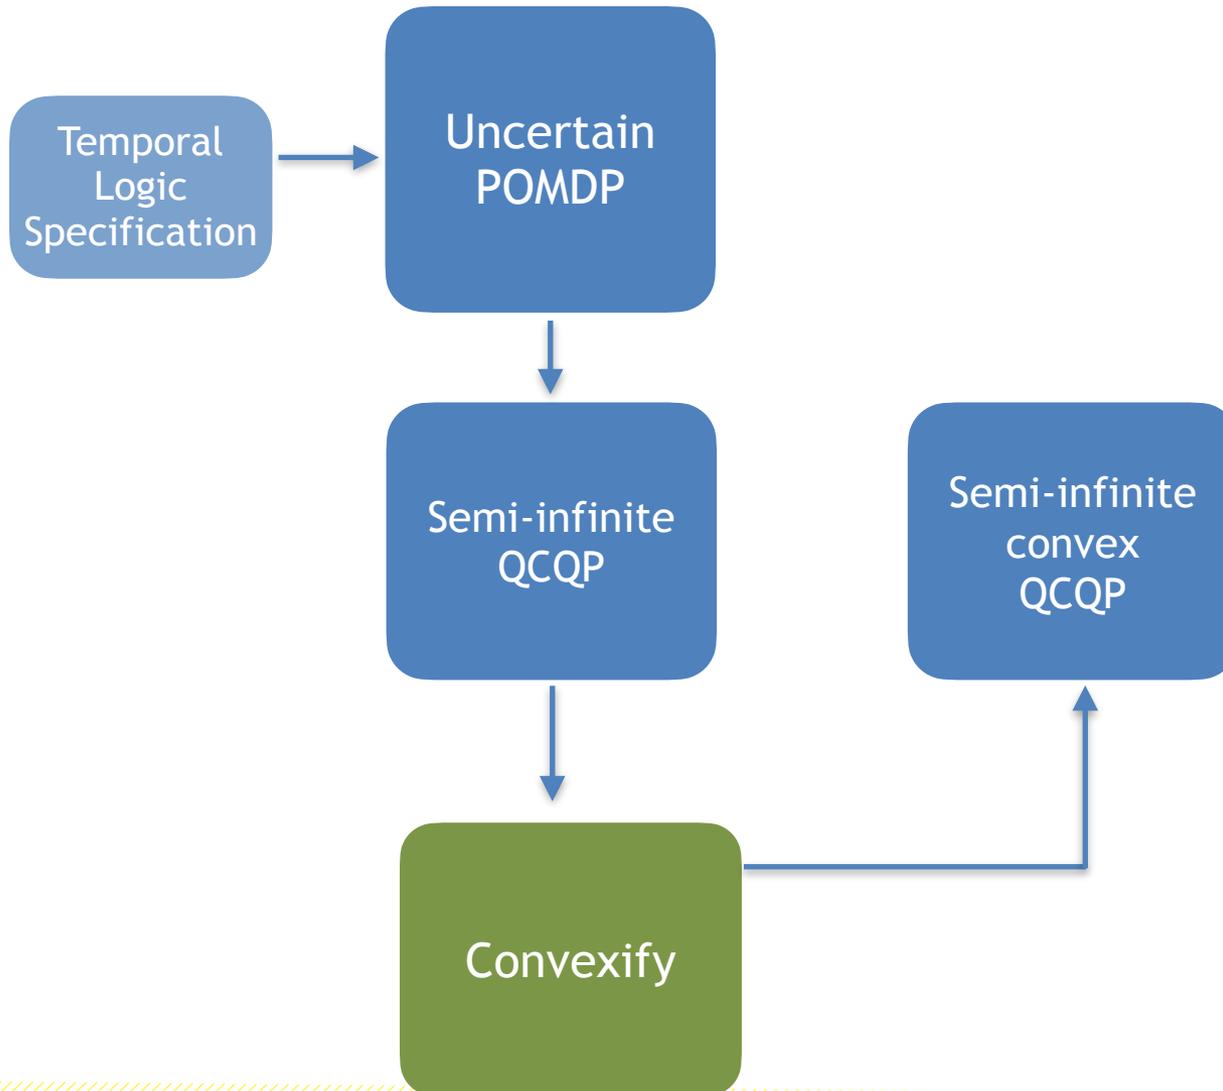
# Story - Convex-concave Procedure



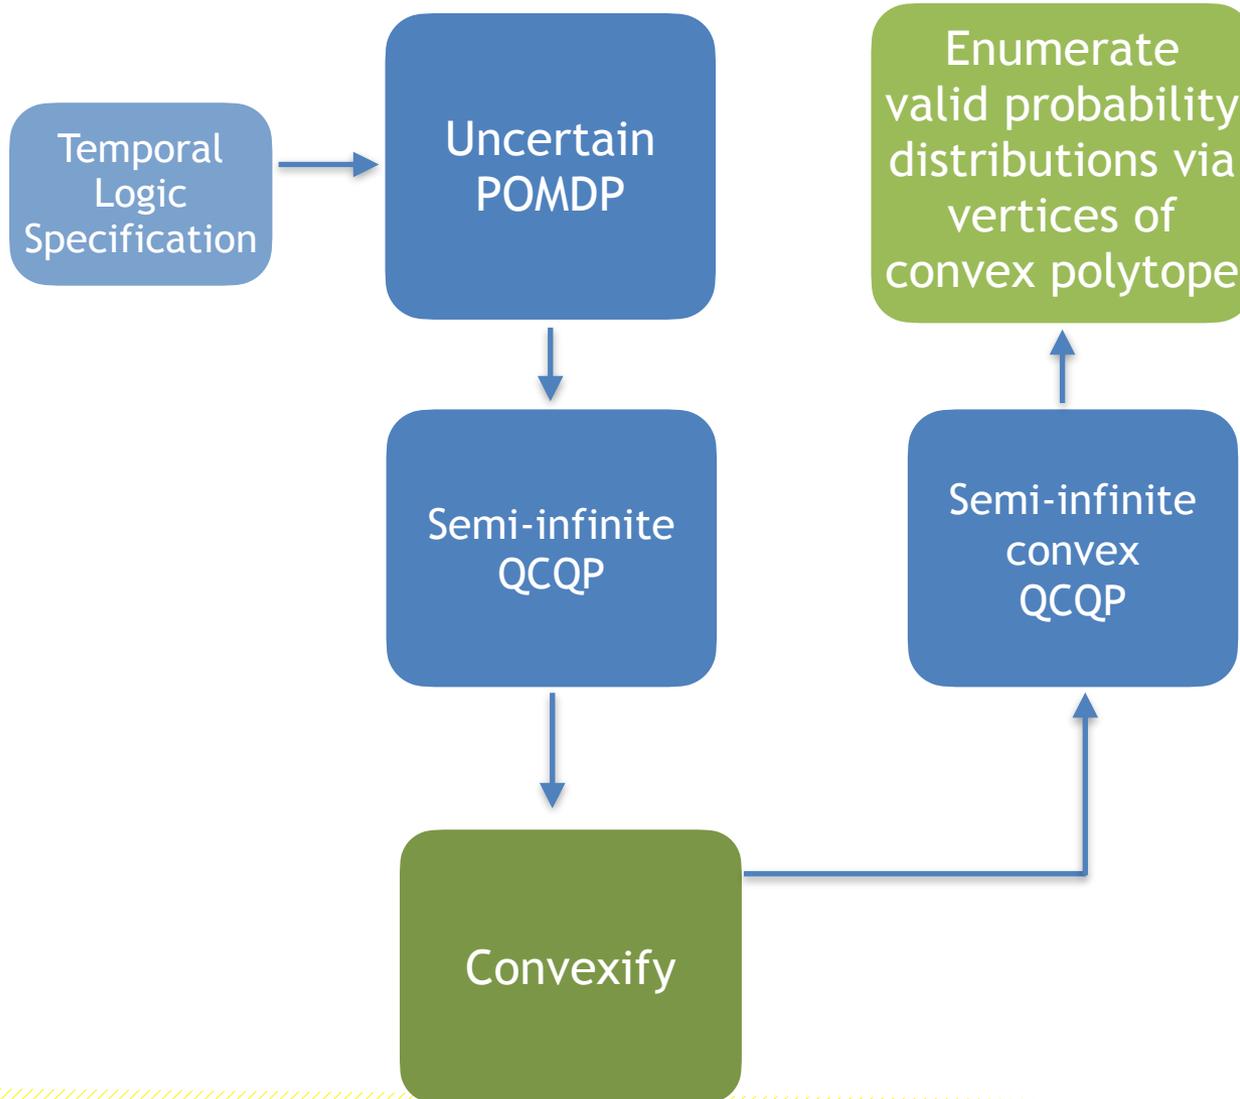
# Story - Convex-concave Procedure



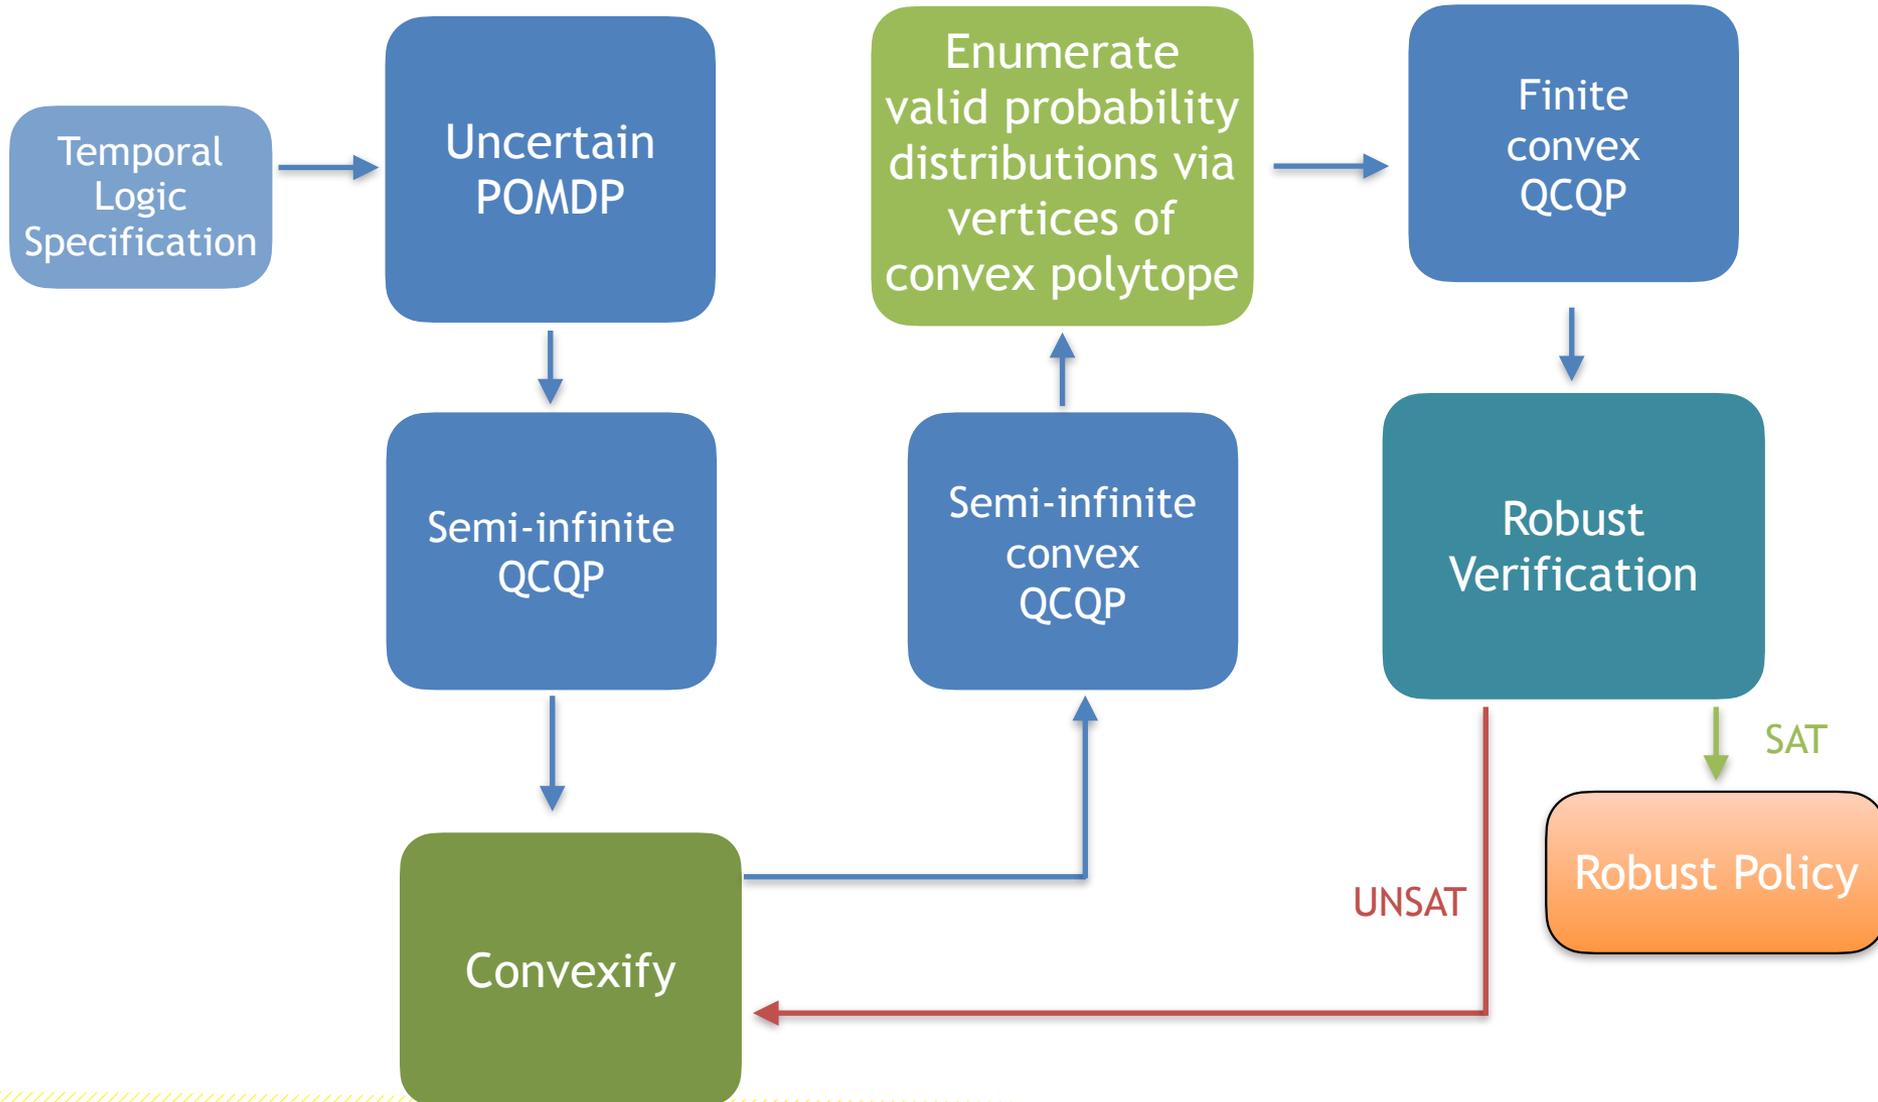
# Story - Convex-concave Procedure



# Story - Convex-concave Procedure



# Story - Convex-concave Procedure



# Interested?

Steven Carr, Nils Jansen, Ralf Wimmer, Alexandru C. Serban, Bernd Becker, Ufuk Topcu. **Counterexample-Guided Strategy Improvement for POMDPs Using Recurrent Neural Networks**. In IJCAI 2019, pages 5532-5539.

Steven Carr, Nils Jansen, Ufuk Topcu. **Verifiable RNN-Based Policies for POMDPs Under Temporal Logic Constraints**. In IJCAI 2020, to appear.

Marnix Suilen, Nils Jansen, Murat Cubuktepe, Ufuk Topcu. **Robust Policy Synthesis for Uncertain POMDPs via Convex Optimization**. In IJCAI 2020, to appear.

Leonore Winterer, Sebastian Junges, Ralf Wimmer, Nils Jansen, Ufuk Topcu, Joost-Pieter Katoen, Bernd Becker. **Strategy Synthesis for POMDPs in Robot Planning via Game-Based Abstractions**. In Transactions of Automatic Control, 2020, to appear.

Leonore Winterer, Ralf Wimmer, Nils Jansen, Bernd Becker. **Strengthening Deterministic Policies for POMDPs**. In NFM, 2020, to appear.

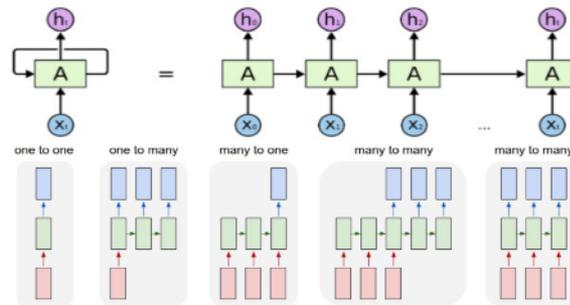
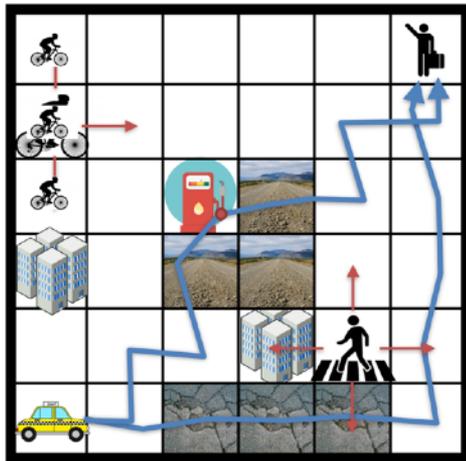
Murat Cubuktepe, Nils Jansen, Sebastian Junges, Joost-Pieter Katoen, Ufuk Topcu. **Scenario-Based Verification of Uncertain MDPs**. In TACAS, 2020, to appear.

Sebastian Junges, Nils Jansen, Ralf Wimmer, Tim Quatmann, Leonore Winterer, Joost-Pieter Katoen, Bernd Becker. **Finite-state Controllers of POMDPs via Parameter Synthesis**. In Conference on Uncertainty in Artificial Intelligence (UAI), pages 519-529, 2018.

<http://nilsjansen.org>

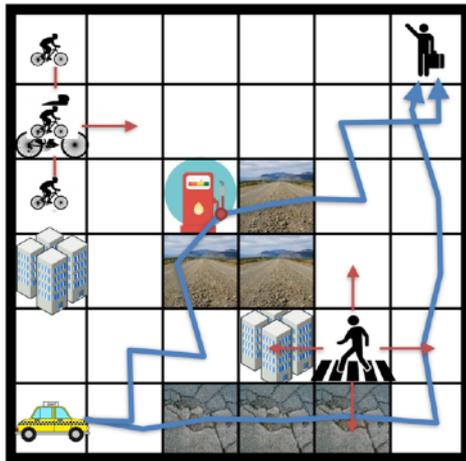
# Conclusion

- Novel way to generate **provably correct** POMDP strategies
- **Good** scalability, **not optimal**
- Results **transferrable**
- First result on robust policies for uncertain POMDPs for temporal logic constraints

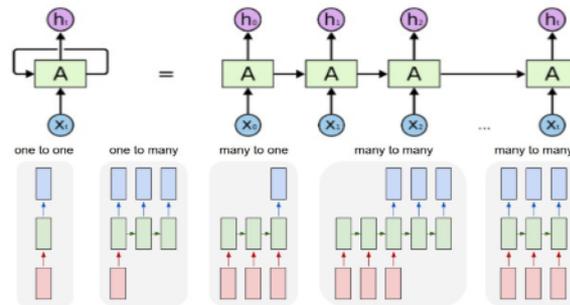


# Conclusion

- Novel way to generate **provably correct** POMDP strategies
- **Good** scalability, **not optimal**
- Results **transferrable**
- First result on robust policies for uncertain POMDPs for temporal logic constraints



$$Pr_{max}(\neg(\text{ped} \vee \text{cyc}) \text{ U } (\diamond(\text{fuel} \wedge \bigcirc \diamond \text{pass})))$$



$$EC_{min}(\diamond \text{passenger})$$