


Controller Synthesis under Model Uncertainty and Structural Constraints

Milan Češka




Applications & Recent results

Partial observability





Decentralised agents with a **hyperproperty** spec




Synthesis of finite state controllers for POMDPs (UAI'22, CAV'23) and for decentralised planning (AAMAS'25)

Uncertainty about the model topology




Synthesis of policy trees for multiple-environment MDPs (ATVA'24)


DT found by dtControl for an optimal policy




Our smaller DT achieves 96% of the optimal value



Synthesis of small almost optimal decision trees for MDPs (CAV'25)



synthesiser



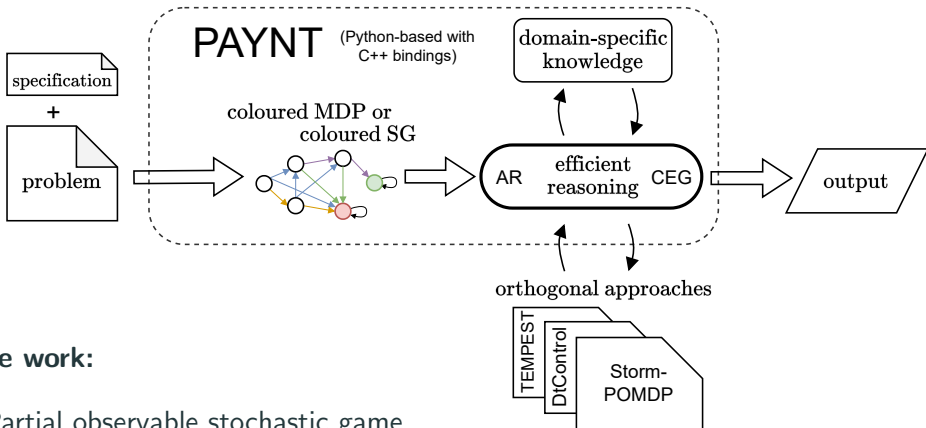
optimal solution (MC)

- probability of each outcome is 1/6
- minimise the expected number of steps

specification

Synthesis of finite-state probabilistic programs from sketches (TACAS'21, JAIR'25)

Available via our synthesis framework PAYNT [CAV'21, JAIR'25]




Future work:

- Partial observable stochastic game
- Robust MDPs and POMDPs
- Robust monitoring/shielding for safe RL

Team

Roman Andriushchenko



Synthesis, Robustness

Filip Macák



POMDPs, POSGs

David Hudák



Safe RL

Cooperation with Sebastian Junges, Joost-Pieter Katoen, and Nils Jansen

Policies Grow on Trees: Model Checking Families of MDPs

Roman Andriushchenko Milan Češka Sebastian Junges **Filip Macák**


Distinguished Paper at ATVA'24

Motivation

Previous work: exploring families of discrete-time Markov chains (DTMCs)

Increased interest in robustness in nondeterministic models


- obtain strategy which works well in multiple environments
- what if we don't know the exact model?



Family of MDPs

Family $\{M_i\}_{i \in \mathcal{I}}$ of MDPs = MDP with parameters

- parameters affect MDP topology
- $i \in \mathcal{I}$ is a parameter assignment, $|\mathcal{I}| < \infty$
- choice of parameter assignment $i \in \mathcal{I}$ represents uncontrollable nondeterminism (adversary, environment)
- choice of action $\alpha \in \text{Act}$ represents controllable nondeterminism



- parameters: $OX = \{2, 3, 4, 5\}$ and $OY = \{2, 3, 4\}$


Robustness problem

input: family $\{M_i\}_{i \in \mathcal{I}}$ of MDPs

input: PCTL reachability property $P(F T) \bowtie \lambda$

output: **robust controller** σ s.t. $\forall i \in \mathcal{I}: P(M_i^\sigma \models F T) \bowtie \lambda$

- requires non-memoryless controllers
- related to solving POMDPs




Problem statement

input: family $\{M_i\}_{i \in \mathcal{I}}$ of MDPs

input: PCTL reachability property $P(\text{F } T) \bowtie \lambda$



output: for each parameter assignment $i \in \mathcal{I}$ a controller σ_i s.t. $P(M_i^{\sigma_i} \models \text{F } T) \bowtie \lambda$
(if such σ_i exists)

6						goal
5						
4		σ_1	σ_2	σ_3	σ_4	
3		σ_5	σ_6	σ_7	σ_8	
2		\emptyset	σ_9	σ_{10}	σ_{11}	
1	 start					
y/x	1	2	3	4	5	6

Problem statement


Additional requirement: produce a **decision tree** of controllers

- nodes of the tree reason about a single parameter
- leaves of the tree (describing sub-families) contain controllers (or \emptyset)
- space-efficient, fast lookup, more understandable for engineers



Stochastic game abstraction

- **Player 1** picks an action, **Player 2** picks a parameter assignment



the above is an over-approximation since Player 2 is too powerful:

- Player 2 can pick parameter assignments inconsistently
 - consistent abstraction would mimic the all-in-one abstraction
- Player 2 acts second
 - this order avoids the abstraction blow-up

Robust policy heuristic

- assume a family \mathcal{M} of MDPs and a specification $P(\text{F } T) \geq 0.9$
- construct game abstraction $\mathcal{G}(\mathcal{M})$
- the following is a **sufficient** (but not necessary) condition for σ_1 to be a robust controller for \mathcal{M} :

$$\max_{\sigma_1} \min_{\sigma_2} P(\mathcal{G}(\mathcal{M})^{\sigma_1 \sigma_2} \models \text{F } T) \geq 0.9$$

- if the above condition does *not* hold and σ_2 is consistent in its parameter assignment, then this assignment is **unsatisfiable**

Proving unsatisfiability heuristic

- assume a family \mathcal{M} of MDPs and a specification $P(F \ T) \geq 0.9$
- the following is a **sufficient** (but not necessary) condition for no MDP in \mathcal{M} being satisfiable:

$$\max_{\sigma_1} \max_{\sigma_2} P(\mathcal{G}(\mathcal{M})^{\sigma_1 \sigma_2} \models F \ T) < 0.9$$

- such “game” abstraction is simply an MDP
- if the above condition does *not* hold and σ_2 is consistent in its parameter assignment, then this assignment is **satisfiable**

Abstraction refinement step: if neither of the tests was successful, we split family \mathcal{M} into smaller subfamilies based on the controller (σ_1, σ_2) for the game abstraction $\mathcal{G}(\mathcal{M})$

- if σ_2 is not consistent i.e. in parameter X , we split wrt. X to disallow such an inconsistency in the subfamilies
- if σ_2 is consistent, representing some satisfiable assignment i , we try to separate i (and other assignments in which σ_2 is consistent) into a smaller subfamily

Proposed Algorithm

Algorithm 1 Policy tree synthesis

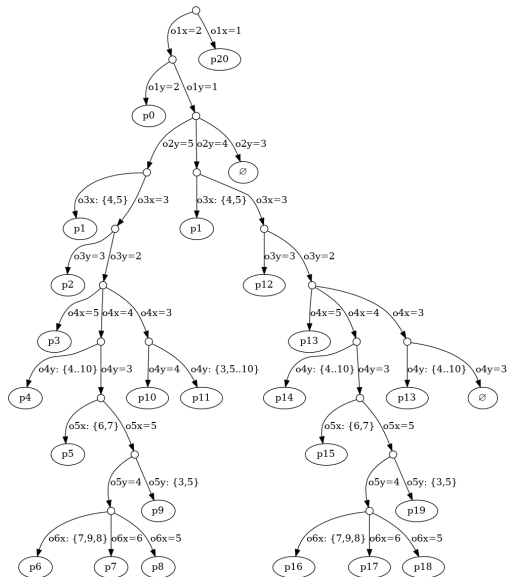
Input: family $\mathcal{M} = \{M_i\}_{i \in \mathcal{I}}$ of MDPs, PCTL property φ

Output: policy tree for \mathcal{M} wrt. φ

```
1: function BUILDTREE( $\mathcal{M}, \varphi$ )
2:    $\sigma \leftarrow$  try to find a robust controller for  $\mathcal{M}$  wrt.  $\varphi$ 
3:   if succeeded then
4:     return LEAFNODE( $\mathcal{M}, \sigma$ )
5:   try to prove that no  $M_i \in \mathcal{M}$  can satisfy  $\varphi$ 
6:   if succeeded then
7:     return LEAFNODE( $\mathcal{M}, \emptyset$ )
8:    $\mathcal{M}', \mathcal{M}'' \leftarrow \text{split}(\mathcal{M})$ 
9:   return INNERNODE( $\mathcal{M}, \text{BUILDTREE}(\mathcal{M}', \varphi), \text{BUILDTREE}(\mathcal{M}'', \varphi)$ )
```

Main idea: given a family of MDPs, try to find a robust controller or try to prove that no satisfying MDP exists, split the family if a conclusive result was not obtained

Decision tree example



Experimental results

model	model info			our approach		speedup wrt.	
	$ S_{\mathcal{M}} $	$ \mathcal{M} $	SAT %	P/SAT %	time	1-by-1	all-in-1
dodge-2	2e5	3e4	100	0.1	122	8	1.1
dodge-3	2e5	9e7	100	<0.01	1445	†1764	MO
dpm-10-b	9e3	1e5	22	0.02	74	21	TO
obs-8-6	5e2	5e4	90	0.6	6	4	1.5
obs-10-6	8e2	3e6	98	<0.01	5	412	MO
obs-10-9	1e3	4e8	100	<0.01	259	†1661	MO
rov-1000	2e4	4e6	99	0.03	1402	†65	TO
uav-work	9e3	2e6	99	<0.01	113	55	TO
virus	2e3	7e4	83	0.9	50	0.8	TO
rocks-6-4	3e3	7e3	100	34	102	0.2	0.1

Conclusion

Main contributions:

1. We contribute a scalable approach to policy synthesis for sets of MDPs
2. The key technique is a game-based abstraction with abstraction refinement
3. The resulting algorithm finds policies for millions of MDPs and provides a compact representation of them

Future work:

- Investigate the robustness problem further
- Incorporate the compact representation of policies (e.g. as decision trees)
- Extend the framework to families of POMDPs

Thank you for attention!