

Constrained and Robust Policy Synthesis with Satisfiability-Modulo-Probabilistic-Model-Checking

Linus Heck, Filip Macák, Milan Češka, Sebastian Junges


AAAI 2026

Presented at NWO ICT.Open2026

Radboud Universiteit



Introduction



Grid where robot slips to the left with probability 0.1

Goal: Find a robot policy that **provably** collects all rocks with probability ≥ 0.9 :

- ▶ Wherever the rocks are (robust)
- ▶ With a simple set of rules (constrained)

Previous work has been unable to find a robot policy that always works.

At the end of this talk, we will see how we can solve this!

Challenges and Approach

Challenge: Combinatorial search space **and** uncertainty


- ▶ Monolithic encodings (e.g. MILPs): Very slow at evaluating MDP policies
- ▶ Branch & Bound: Cannot reason about arbitrary constraints

Our approach: **SMT Solving + Probabilistic Model Checking** (SMPMC)

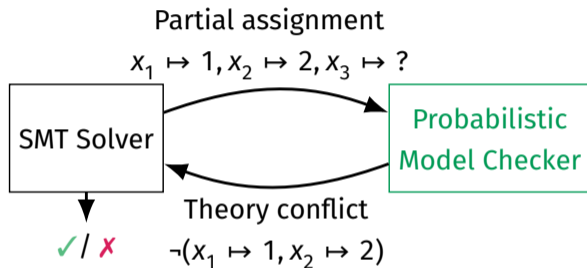
- ▶ Search space and constraints: handled by the SMT solver
- ▶ Solving underlying MDPs: handled by the probabilistic model checker (e.g. using value iteration)

SMT Problem Formulation

Input formula to SMT solver:



SMT Solving Meets Probabilistic Model Checking



- ▶ **Smarter than Guess-And-Check:** Check partial assignments, prune early
- ▶ **More flexible than Branch & Bound:** SMT solver handles search space

Experiments


- ▶ 372 benchmarks from different domains: Sets of MCs, POMDPs with FSCs, Multi-Environment MDPs
- ▶ Comparisons: PAYNT (state-of-the-art for policy synthesis), monolithic SMT (baseline)

Q1: Is SMPMC competitive with the comparisons at policy synthesis?

Q2: Can SMPMC handle both robustness and arbitrary constraints?

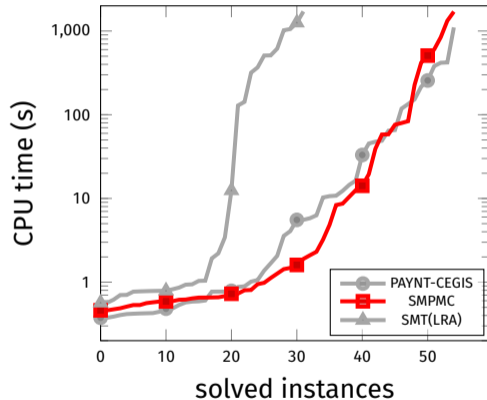
Results: Comparisons Against Baselines

Unconstrained synthesis (single env.):



⇒ competitive with Branch & Bound!

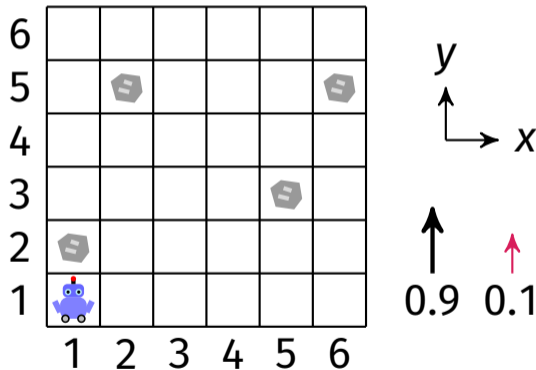

Constrained synthesis (single env.):



⇒ competitive with Guess-And-Check!



Solving Rocks with Constrained and Robust Synthesis

SMPMC finds a five-node decision tree that picks up all rocks in a spiral:



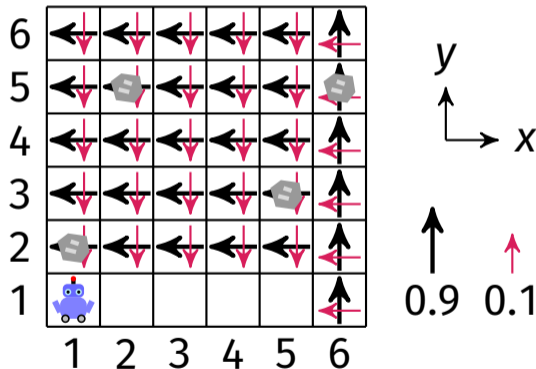

Solving Rocks with Constrained and Robust Synthesis

SMPMC finds a five-node decision tree that picks up all rocks in a spiral:





Solving Rocks with Constrained and Robust Synthesis

SMPMC finds a five-node decision tree that picks up all rocks in a spiral:



Solving Rocks with Constrained and Robust Synthesis

SMPMC finds a five-node decision tree that picks up all rocks in a spiral:



Contributions

- ▶ SMPMC: A novel, versatile, and efficient approach that tightly integrates satisfiability solvers with probabilistic model checking
- ▶ First effective approach that constructs robust decision tree policies
- ▶ Competitive with state-of-the-art heuristics, solves new problems

