Robotic Collaboration Through Scalable Reactive Synthesis

NSF NRI FND: 1830549
PIs: Lydia E. Kavraki and Moshe Y. Vardi

RICE UNIVERSITY

Robotic Collaboration Through Scalable Reactive Synthesis

Motivation: We want humans and robots to safely and effectively collaborate on complex tasks

Overall Approach: We formally model the human-robot ensemble as a game in order to provide formal guarantees such as:

- Safety
- Optimality
- Correctness

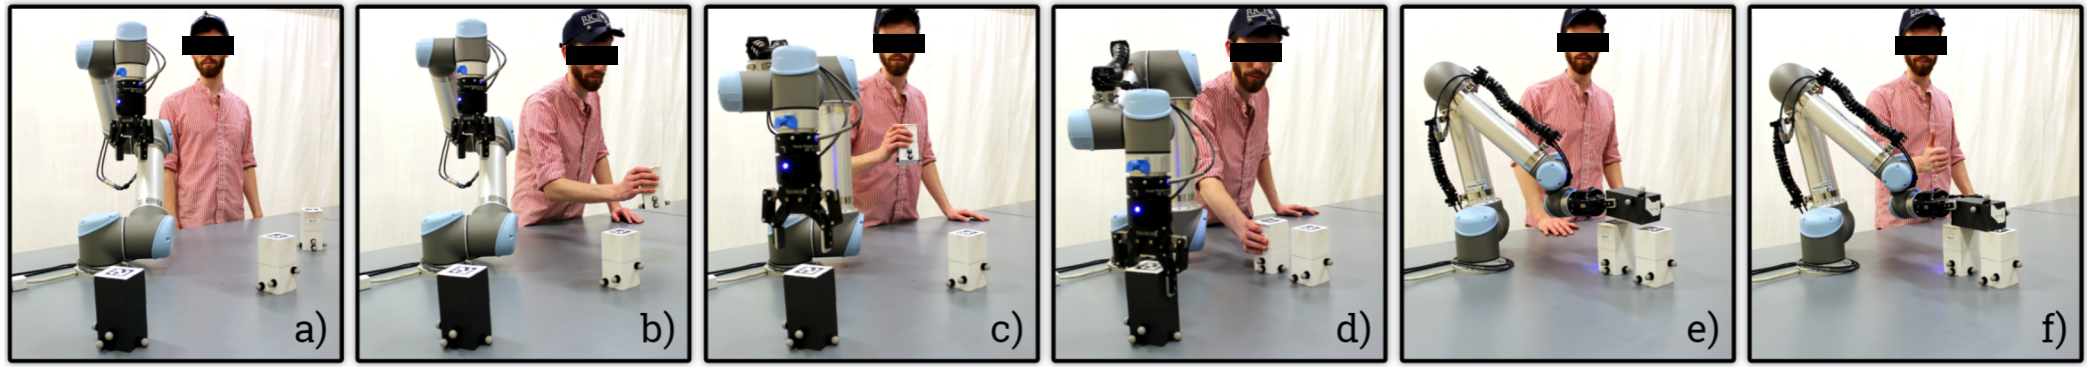

Worst-case Synthesis

Problem: Given a finite-horizon temporal specification and a model of possible human-robot actions, synthesize a policy to guarantee task completion

We combine a logical specification with a domain abstraction

- This becomes a 2-player game

Significant speedup by using a symbolic approach (Binary Decision Diagrams[1]) developed to deal with state-space explosion in model-checking community[2]

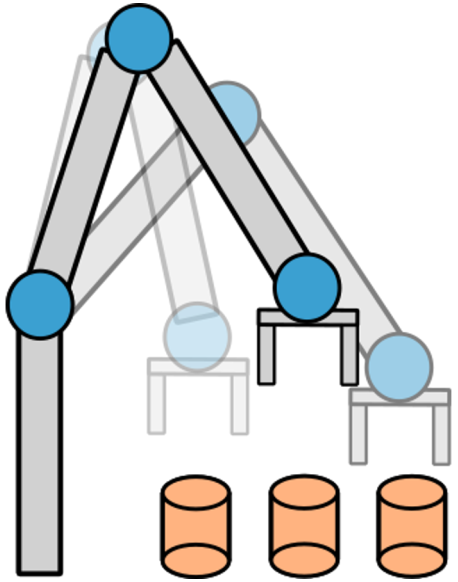

BDDs implicitly group "equivalent" states together

1. R. E. Bryant, "Graph-Based Algorithms for Boolean Function Manipulation," in IEEE Transactions on Computers, Vol. 100, pgs. 677-691, IEEE, 1986.
2. K. He,A. M. Wells, L. E. Kavraki, and M. Y. Vardi, "Efficient Symbolic Reactive Synthesis for Finite-Horizon Tasks," in IEEE Intl. Conf. on Robotics and Automation, 2019.

What about Stochastic Actions?

Our prior work: Given a finite-horizon temporal specification and a model of possible human-robot actions, synthesize a policy to guarantee task completion

Our current work: Given a finite-horizon temporal specification and a **stochastic** model of possible human-robot actions, synthesize a policy to **maximize the probability** of task completion

Expected-case Synthesis

- Markov Decision Processes (MDPs) are a popular model for stochastic systems
- MDP synthesis is doubly-exponential in the length of the specification
- We address scalability by again using a symbolic approach
- Allows us to deal with expected-case rather than worst-case

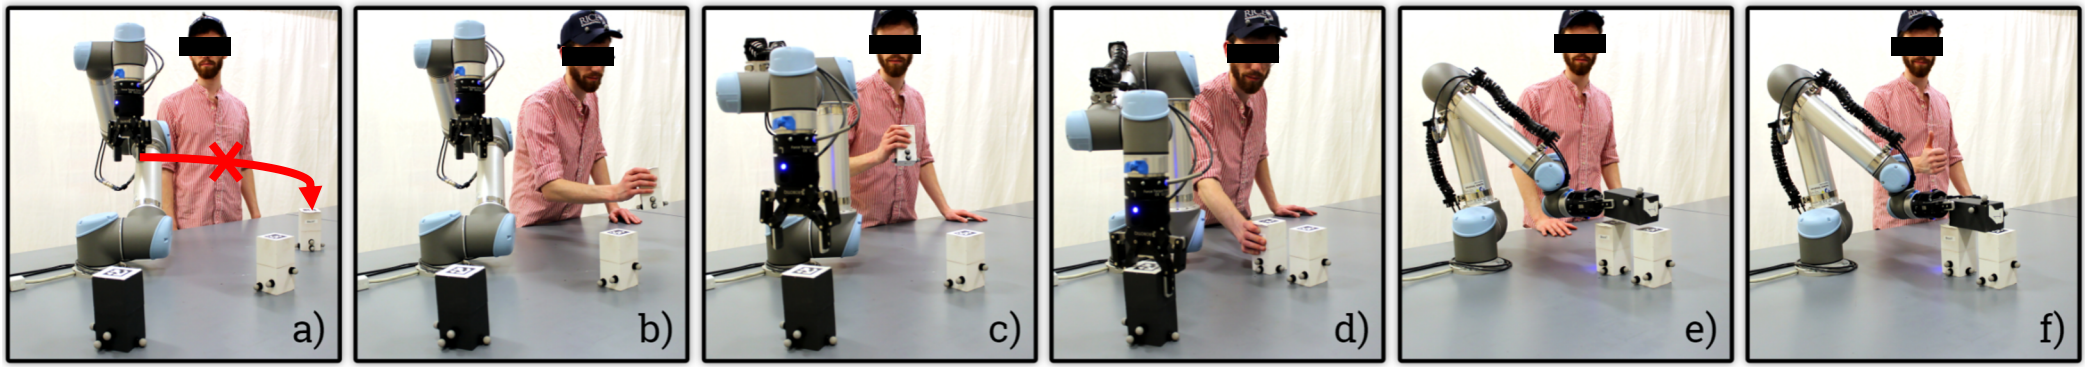

At state (a) we anticipate human cooperation (b-d), leading to faster expected task completion (e-f).

Results

Different expressivity compared to our prior work in reactive synthesis

We test on gridworld (and other) domains using PRISM model checker

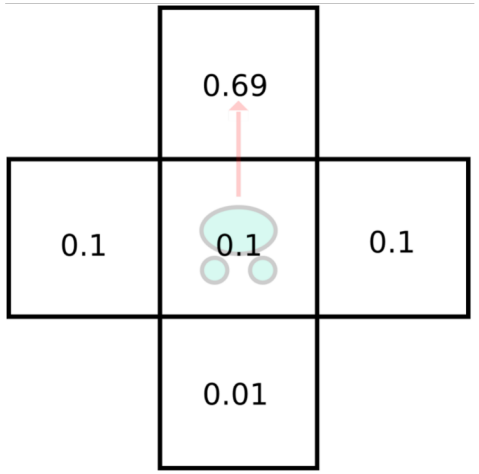

Robot transition probability distribution

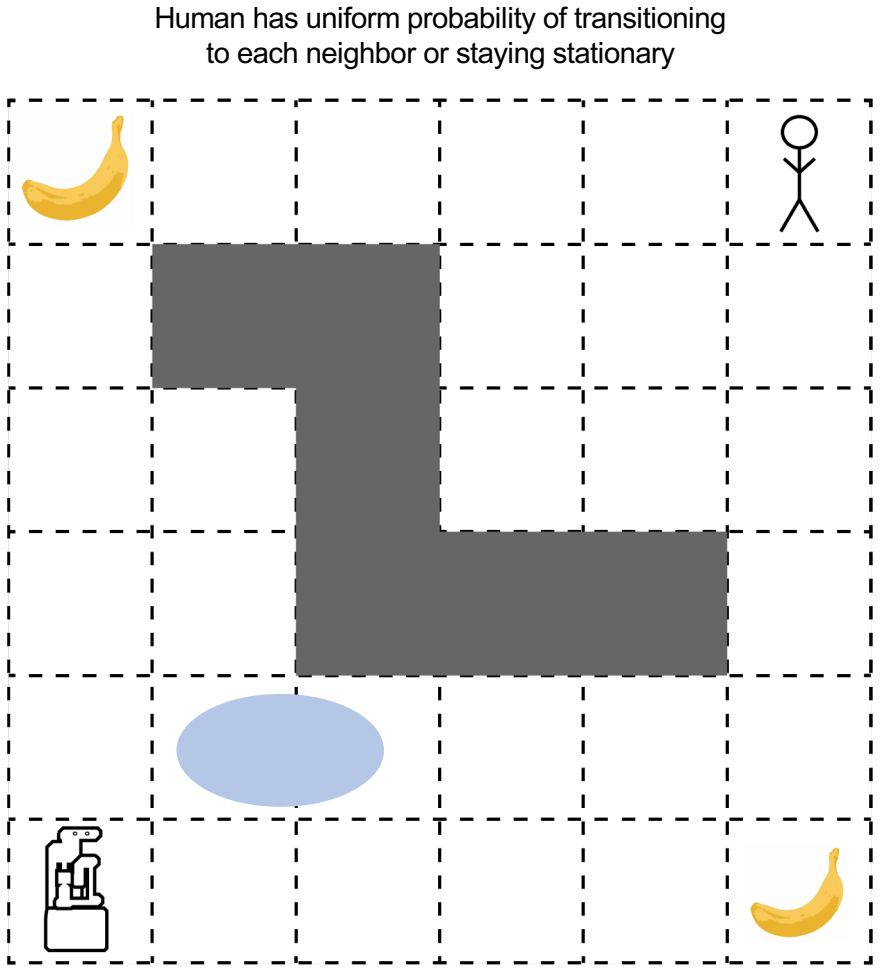

Gridworld with a human, two subgoals and a wet region to avoid

M. Kwiatkowska, Gethin Norman and David Parker. PRISM 4.0: Verification of Probabilistic Real-time Systems. In Proc. 23rd International Conference on Computer Aided Verification (CAV'11), volume 6806 of LNCS, pages 585-591, Springer, 2011.

Results

Our symbolic approach improves scalability

- Significant runtime improvement (4x)
- Significant memory improvement (up to 17 vs 9 subgoals)

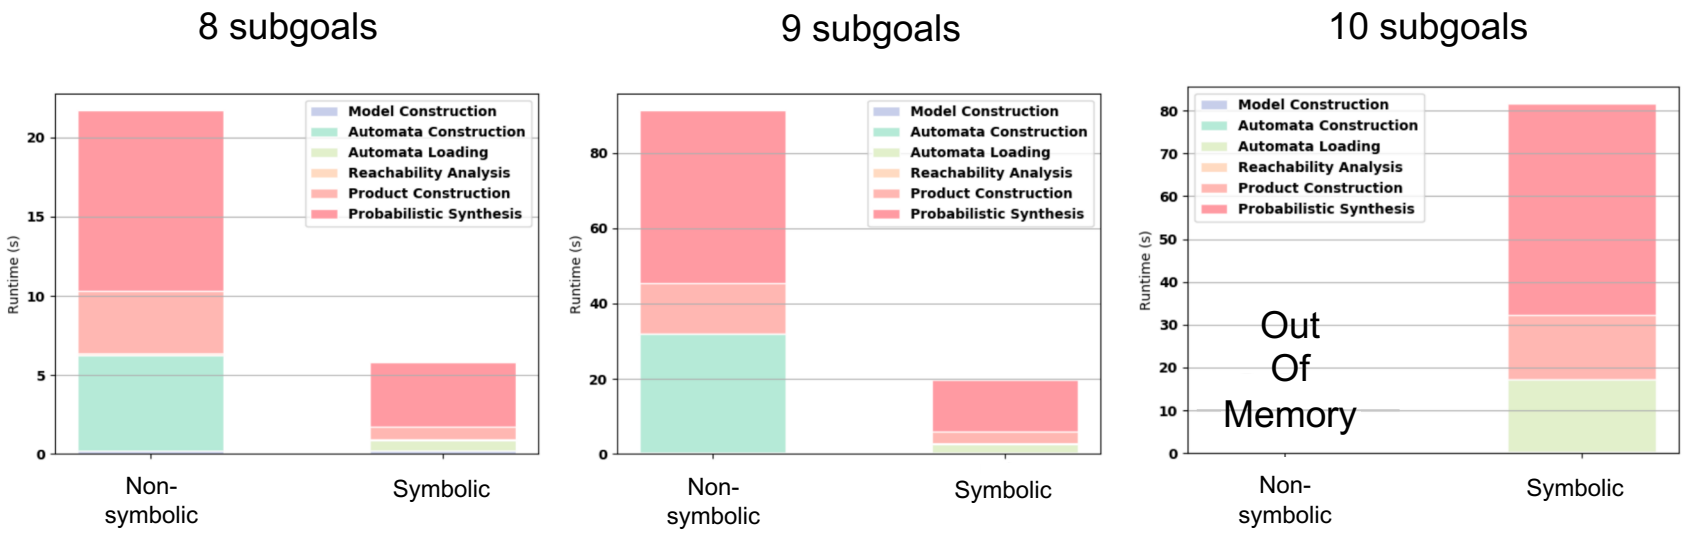

A. M. Wells, M. Lahijanian, L. E. Kavraki, and M. Y. Vardi, “LTLf Synthesis on Probabilistic Systems,” Electronic Proceedings in Theoretical Computer Science, vol. 326, pp. 166–181, 2020.

Future Work

Stochastic games (combine reactive and probabilistic synthesis)

Improve scalability

- Fully symbolic approach
- Factoring
- Decomposition