Deep Reinforcement Learning in a Handful of Trials using Probabilistic Dynamics Models

Kurtland Chua, Roberto Calandra, Rowan McAllister, Sergey Levine

University of California, Berkeley

How Long Does Learning Take?

~50 million frames

[Mnih et al. 2015]

~21 million games

[Silver et al. 2017]

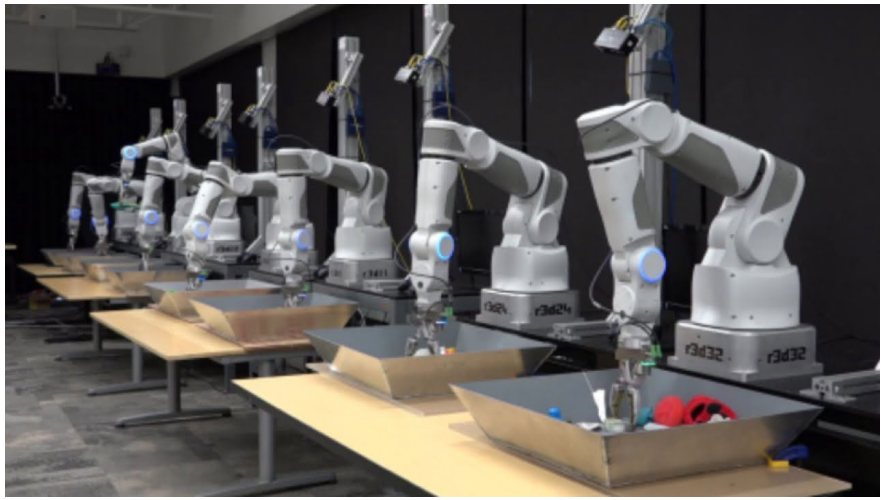

[Levine et al. 2017]

~800,000 grasp attempts

Can we speed this up?

Model-Based Reinforcement Learning

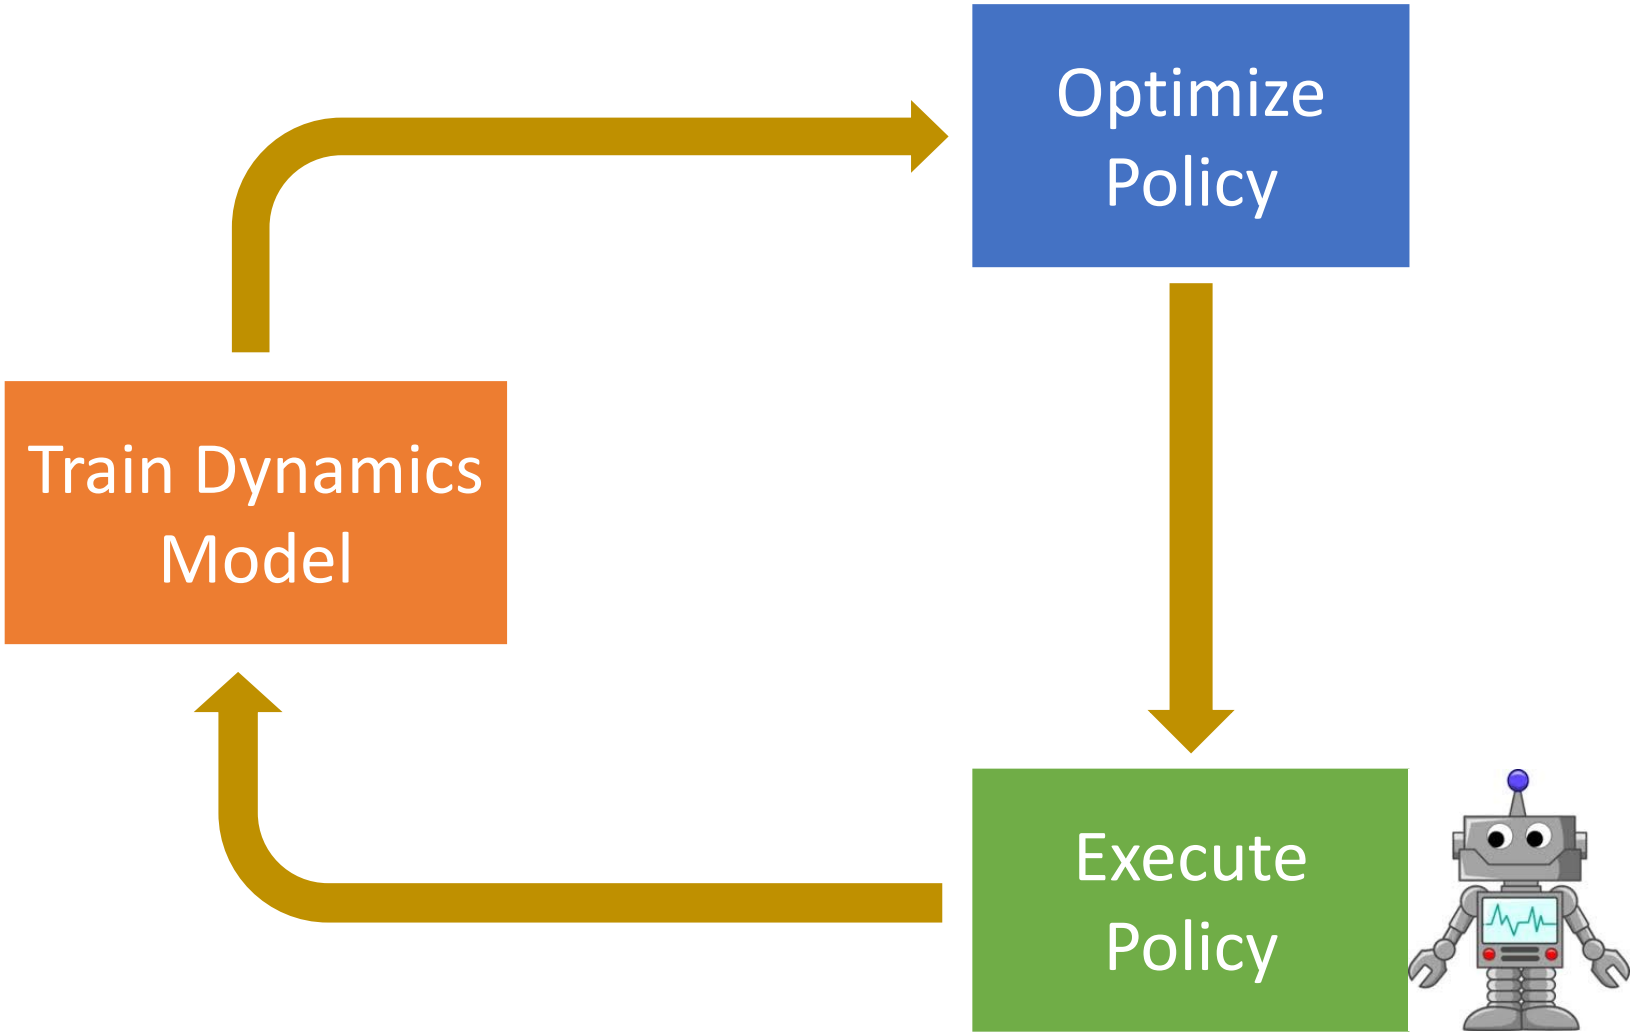

Comparative Performance on HalfCheetah

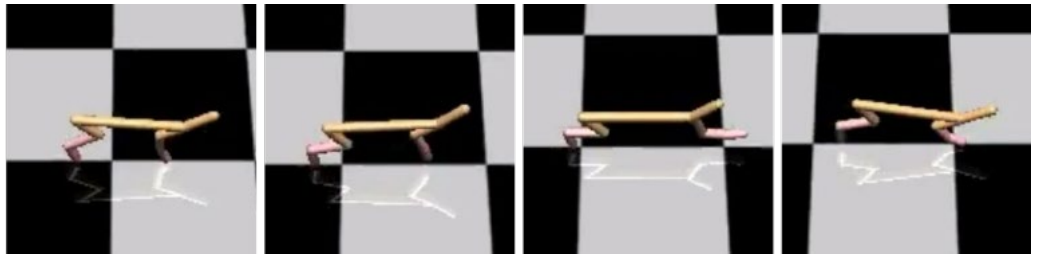

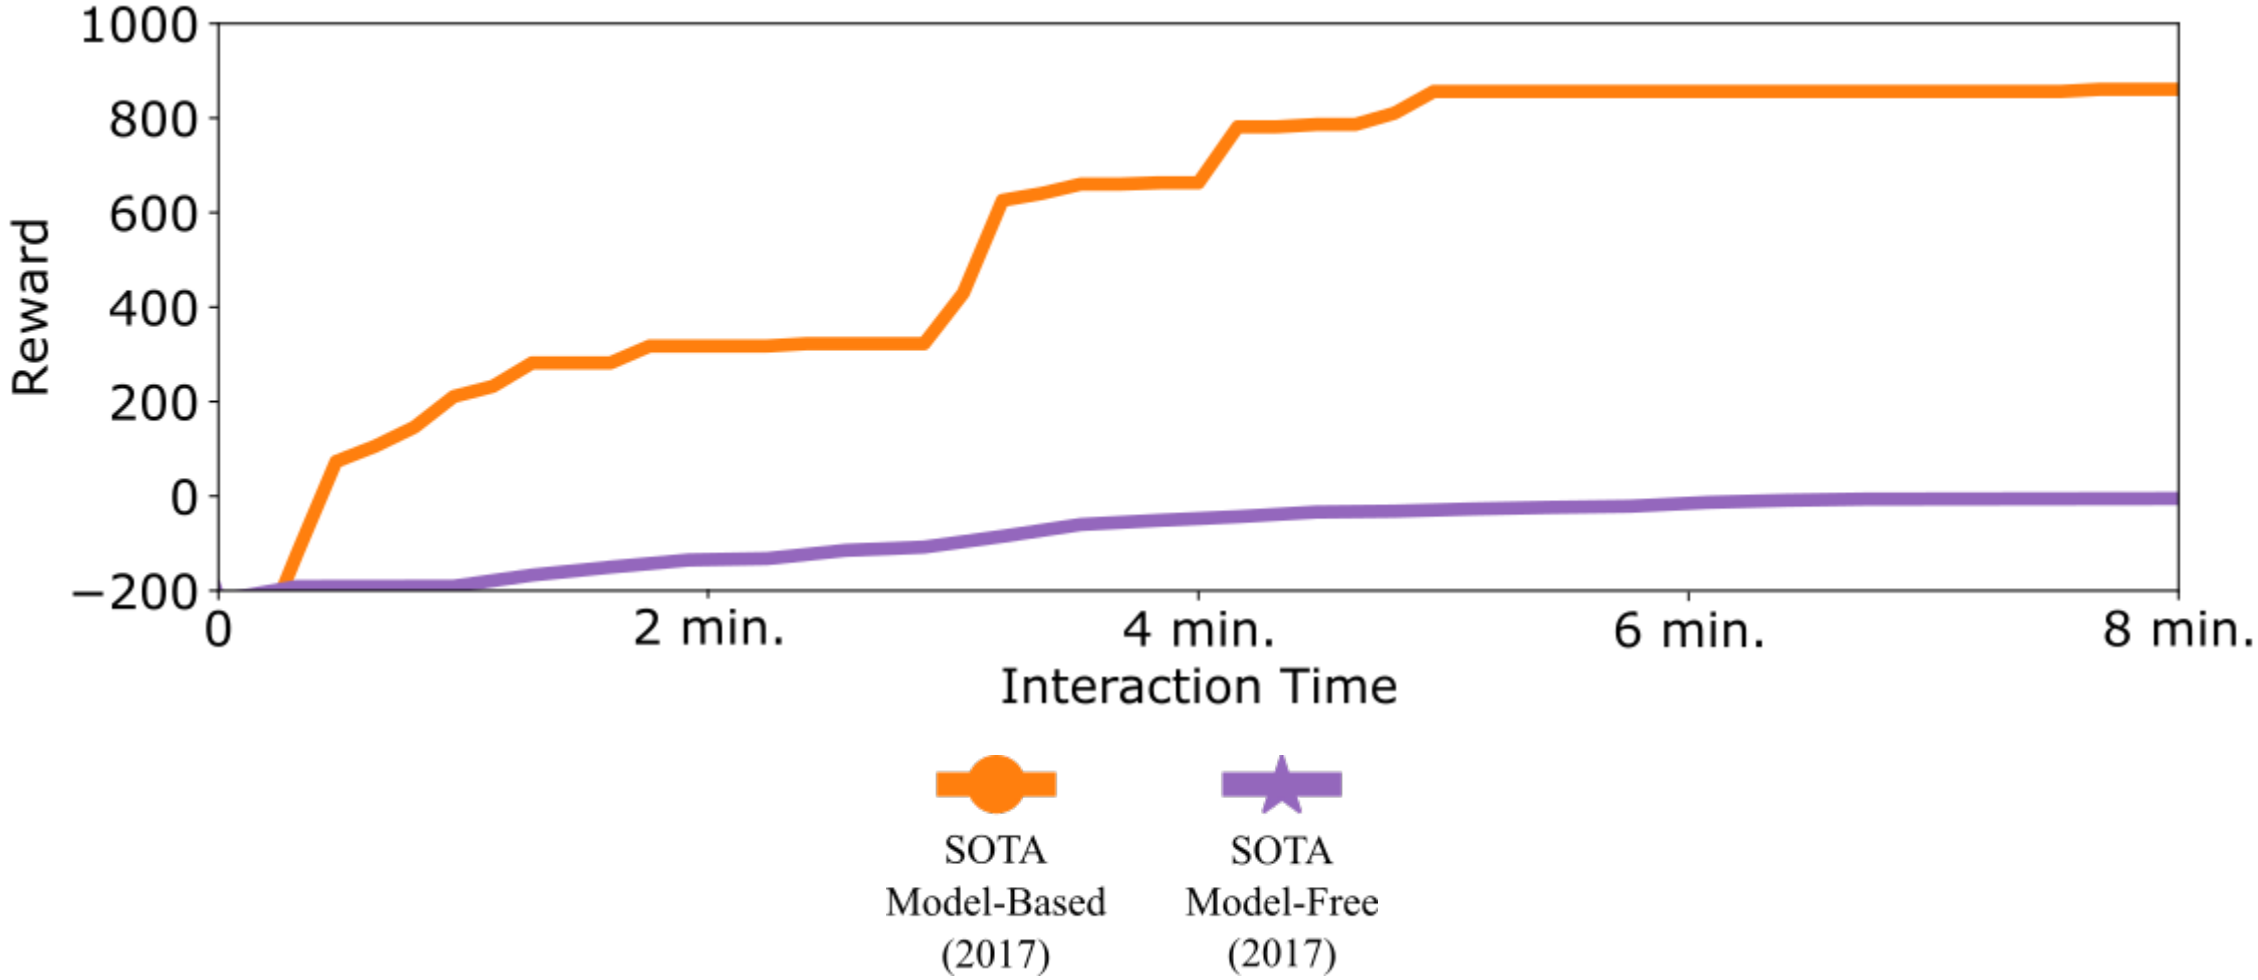

Comparative Performance on HalfCheetah

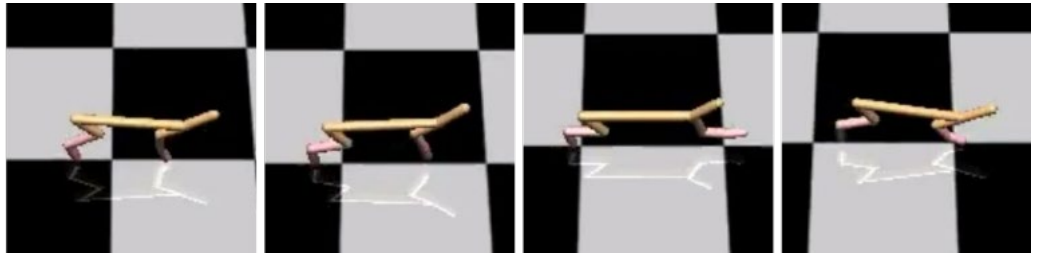

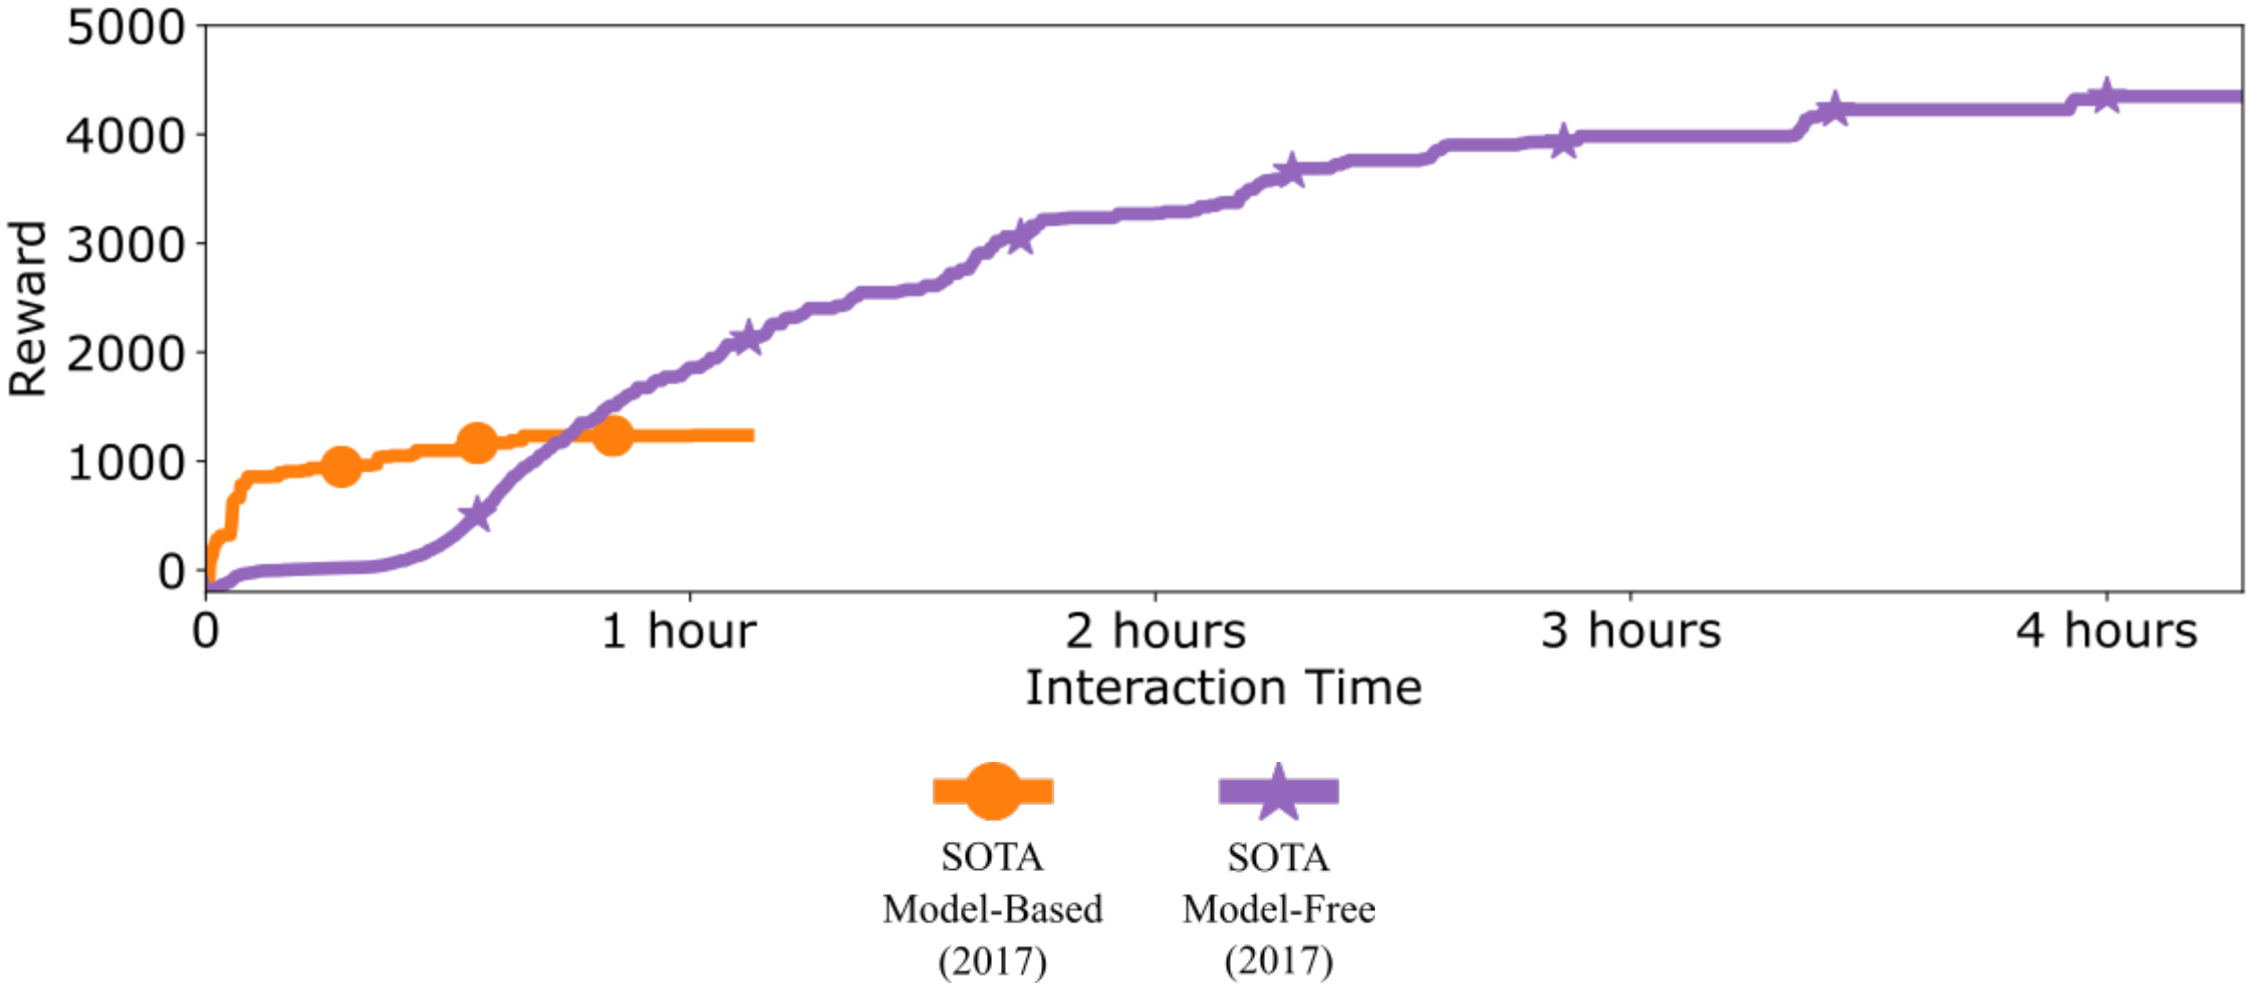

Deterministic Neural Nets as Models

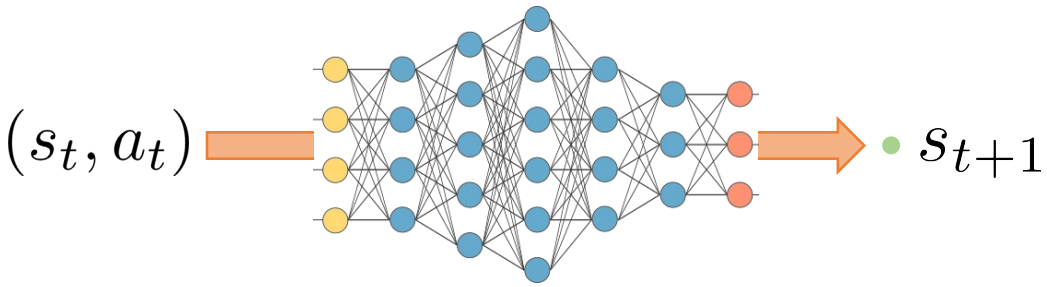

Deterministic Neural Nets as Models

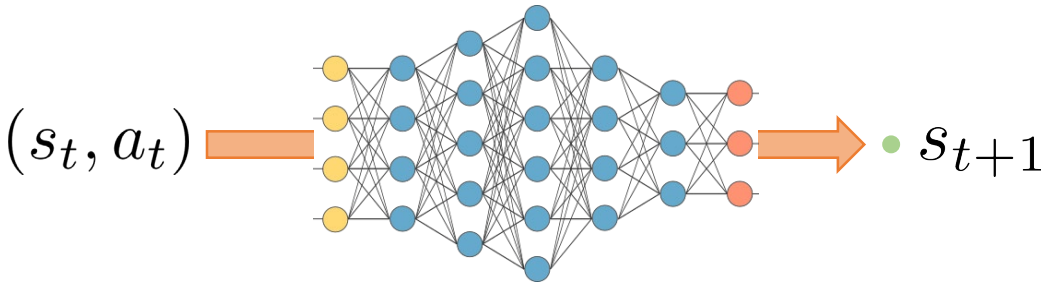

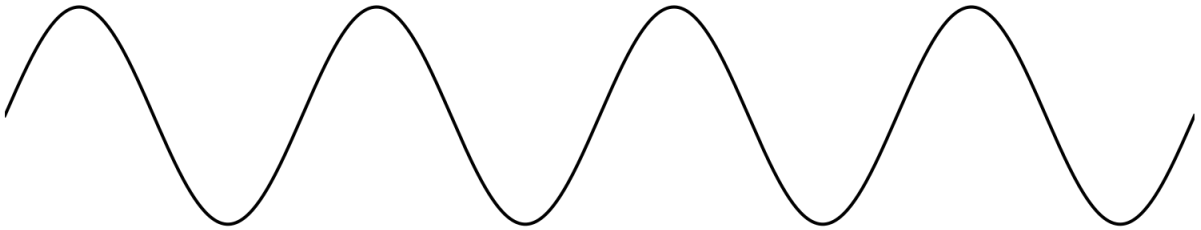

Deterministic Neural Nets as Models

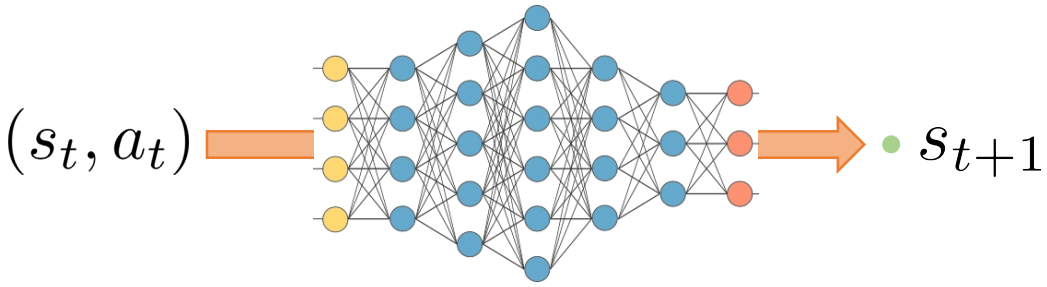

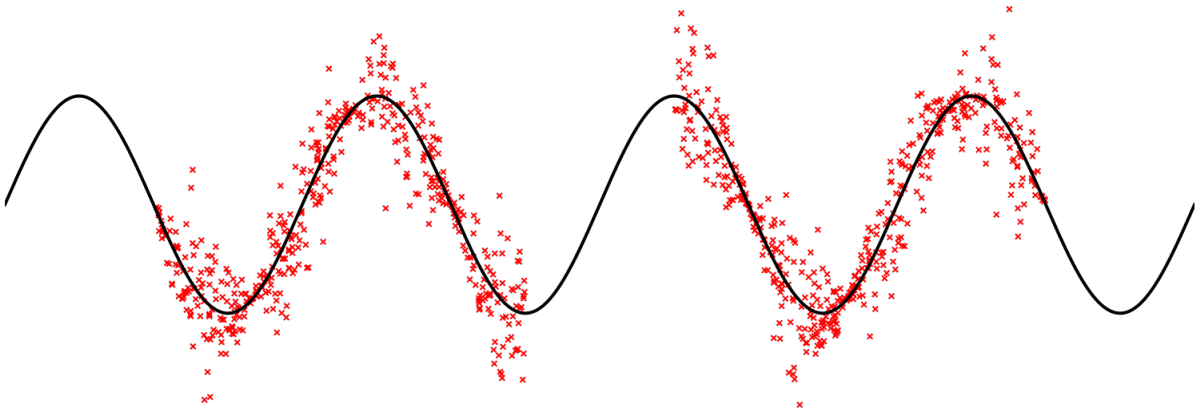

Deterministic Neural Nets as Models

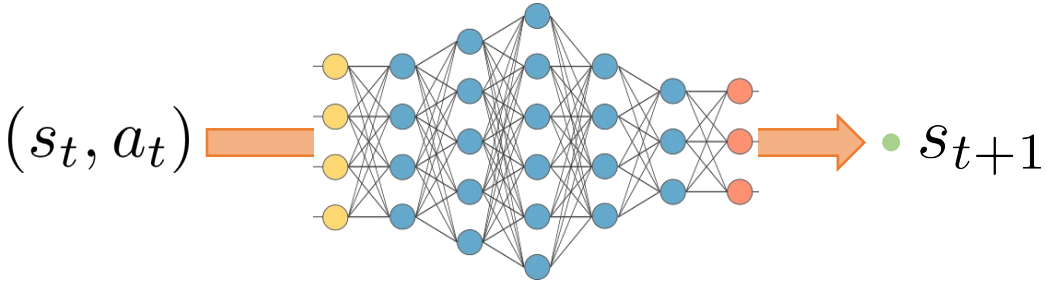

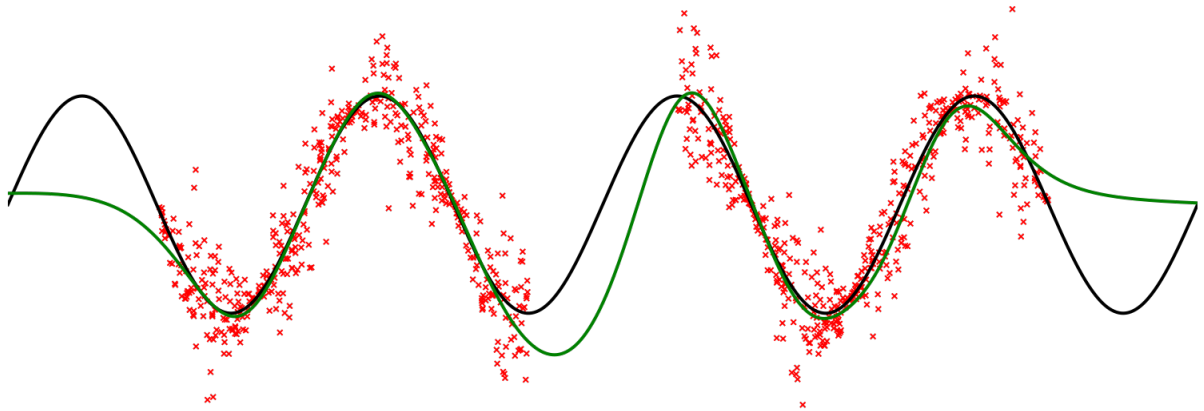

Deterministic Neural Nets as Models

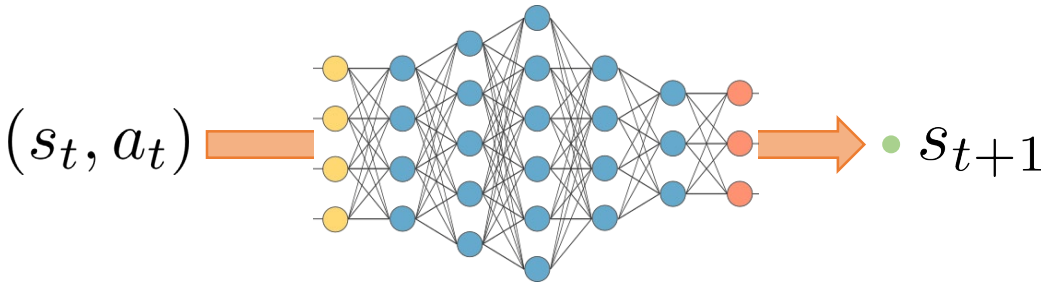

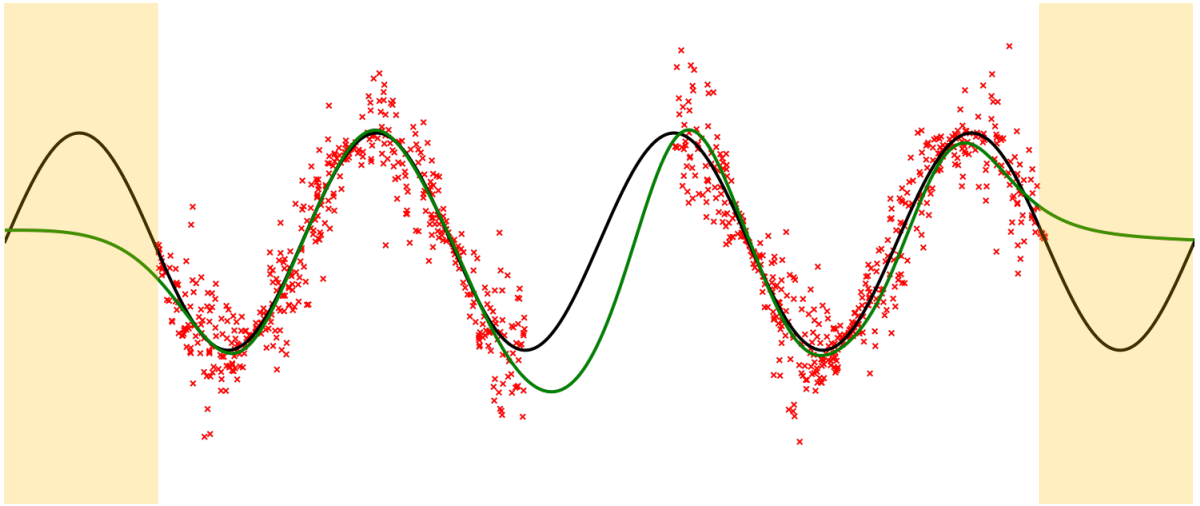

Probabilistic Neural Nets as Models

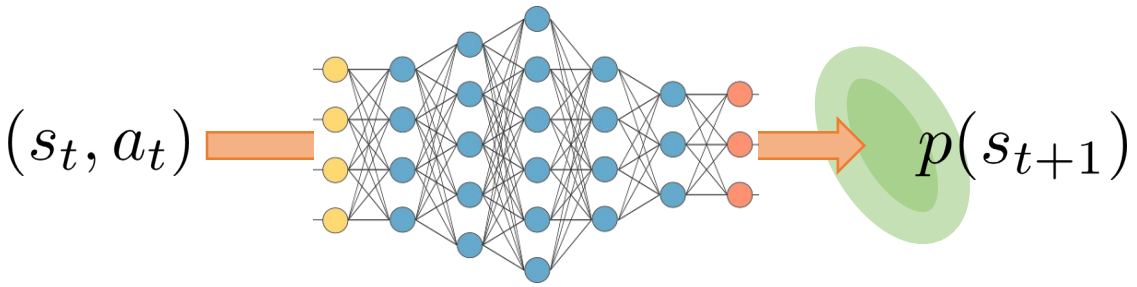

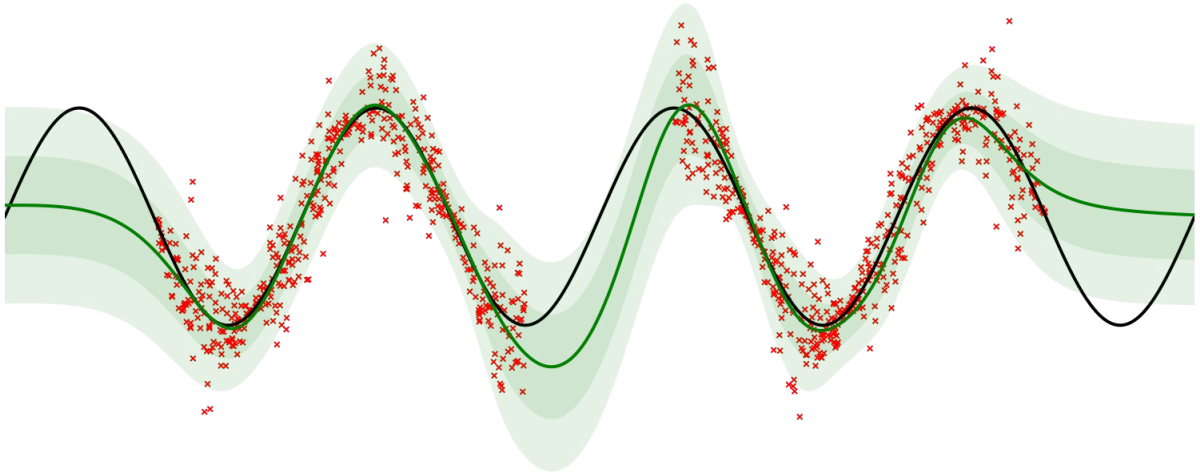

Probabilistic Ensembles as Models

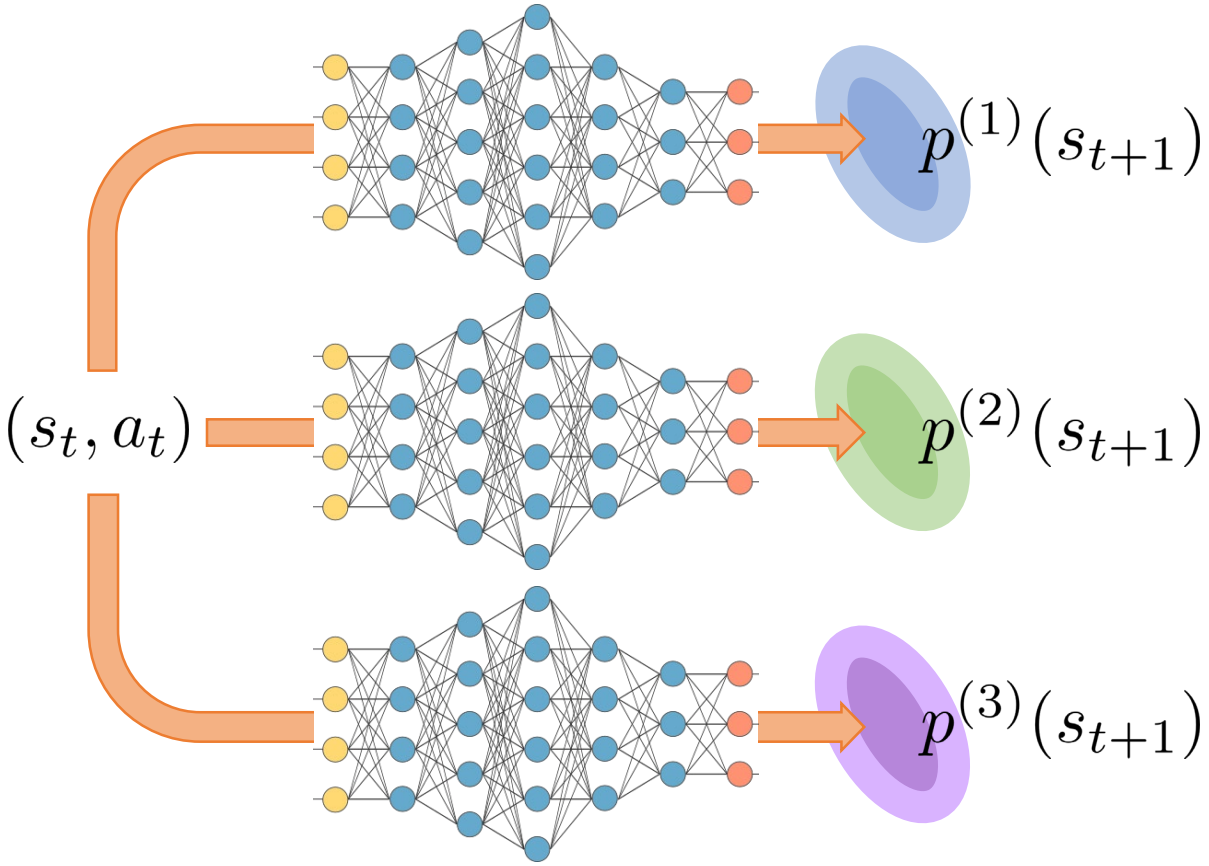

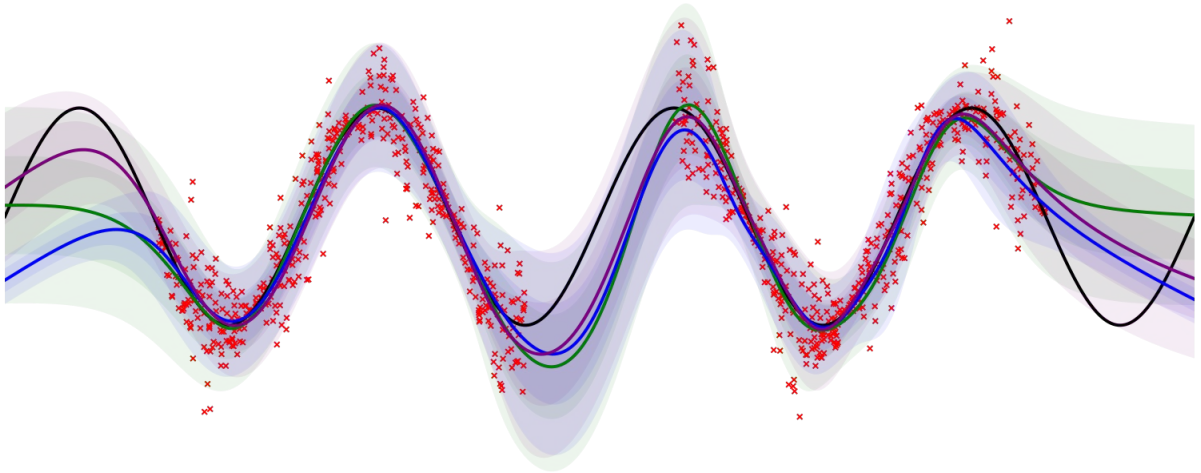

Probabilistic Ensembles as Models

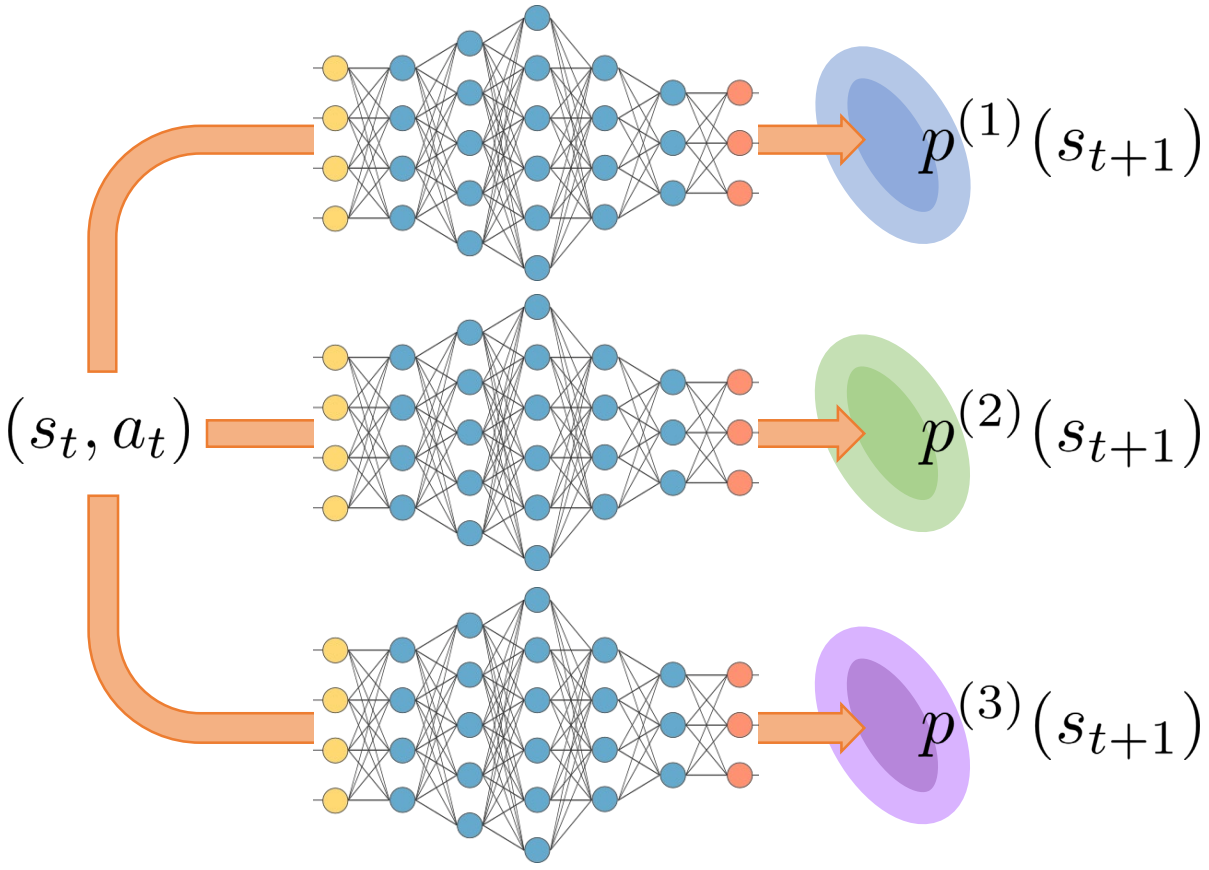

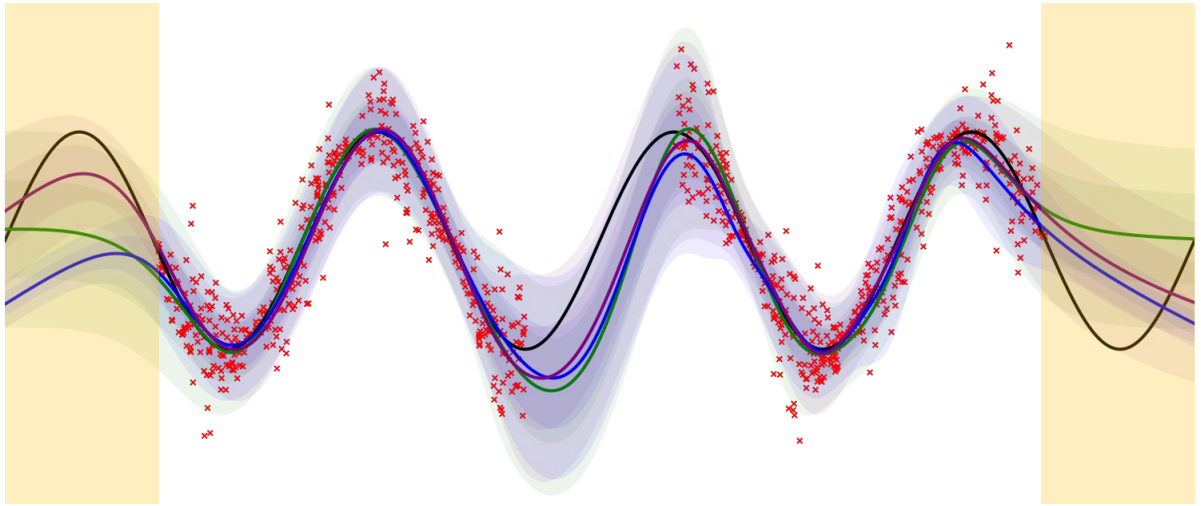

Trajectory Sampling for State Propagation

Trajectory Sampling for State Propagation

s_t

Trajectory Sampling for State Propagation

Trajectory Sampling for State Propagation

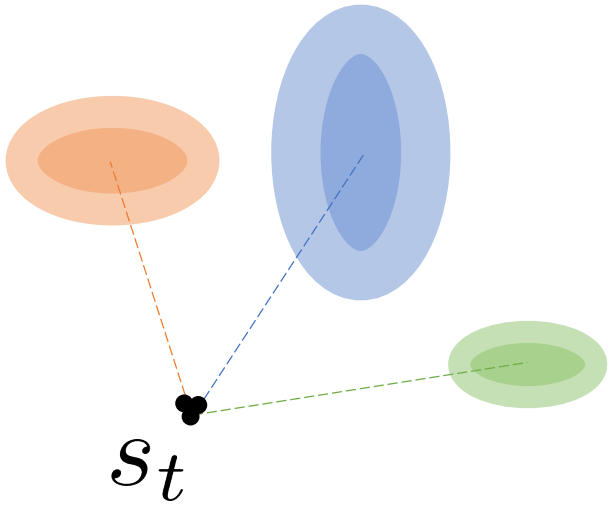

Trajectory Sampling for State Propagation

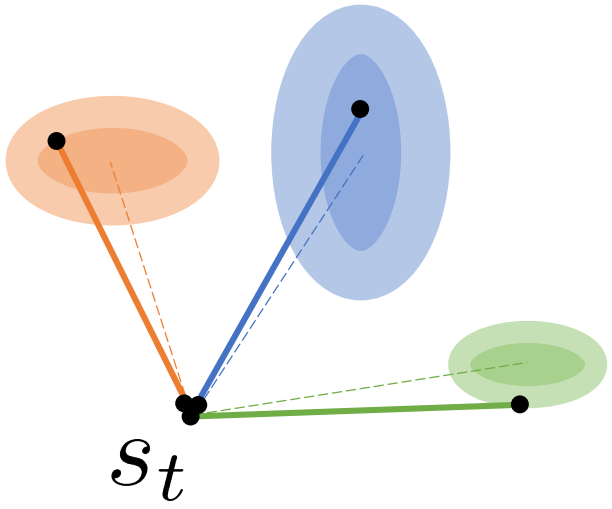

Trajectory Sampling for State Propagation

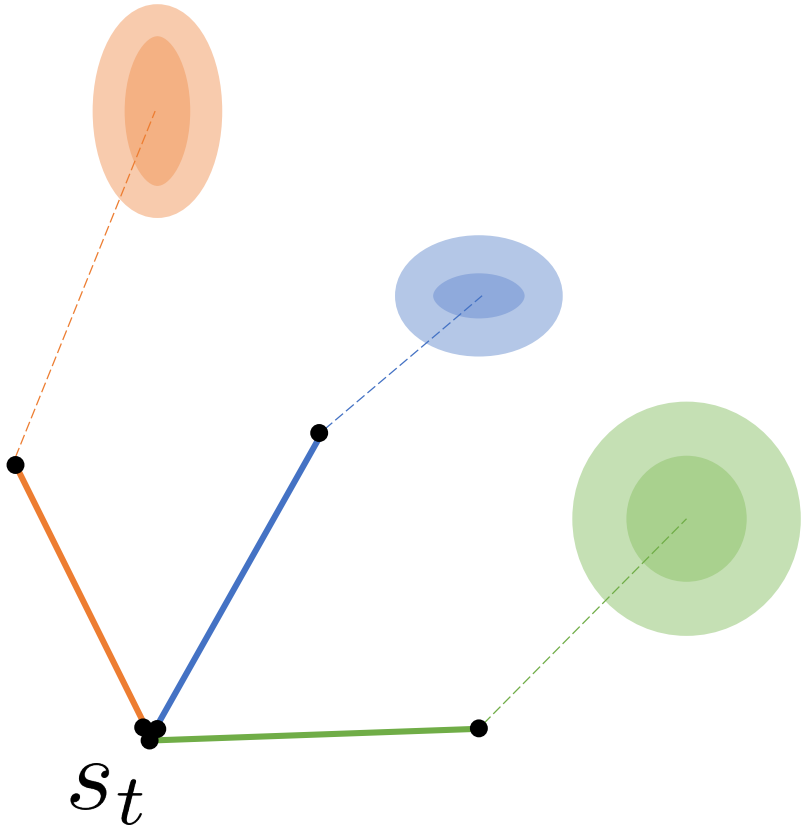

Trajectory Sampling for State Propagation

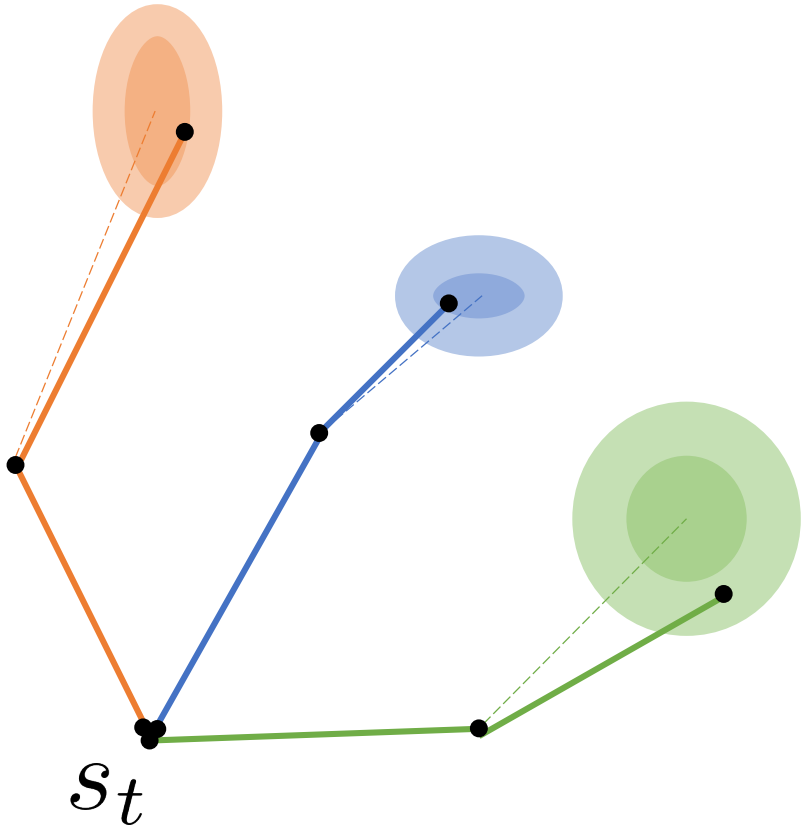

Trajectory Sampling for State Propagation

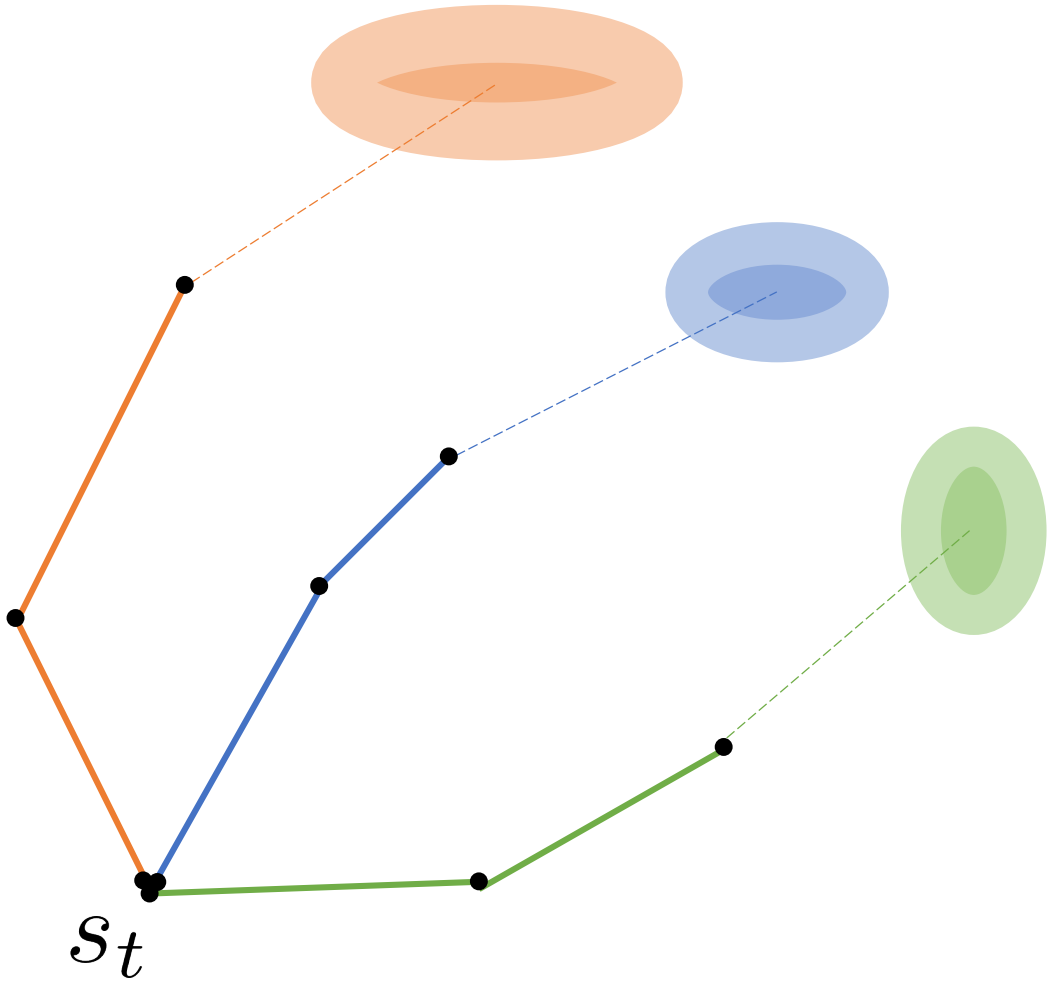

Trajectory Sampling for State Propagation

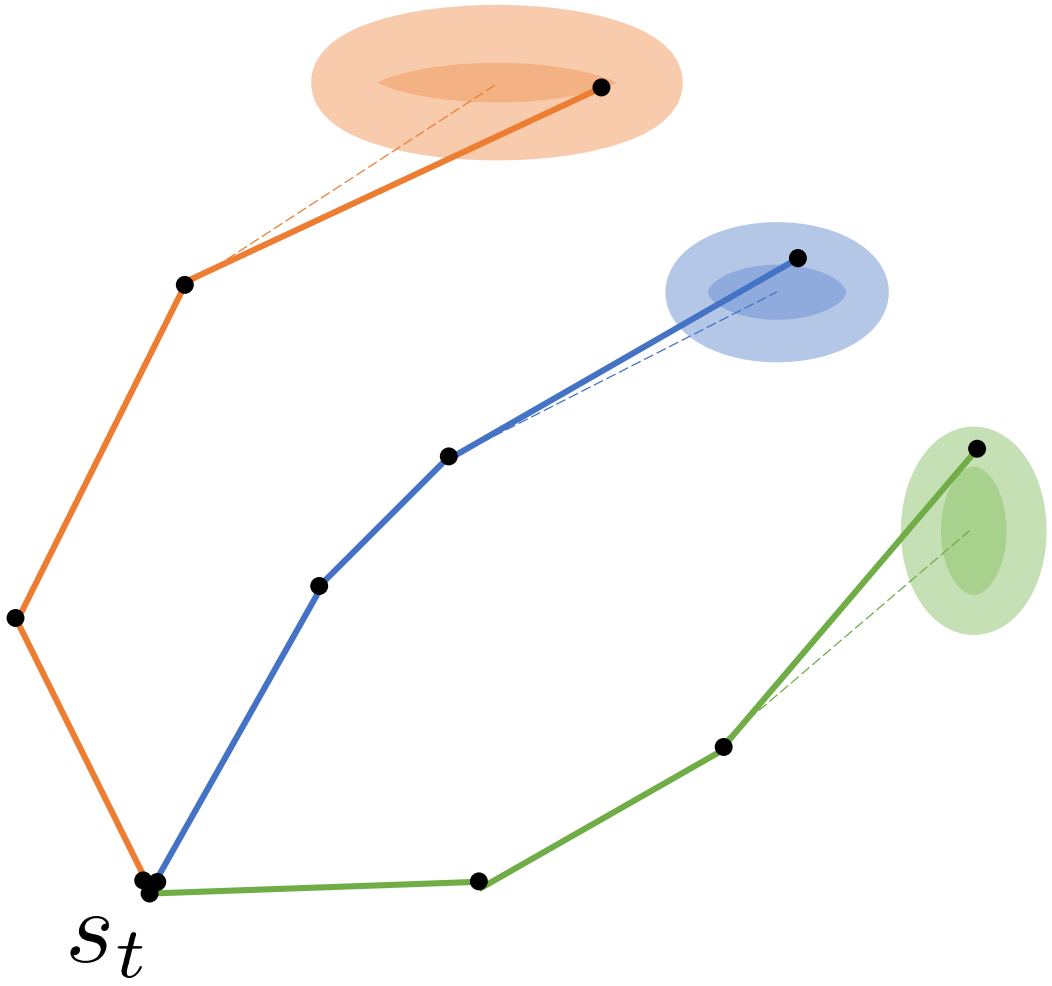

Trajectory Sampling for State Propagation

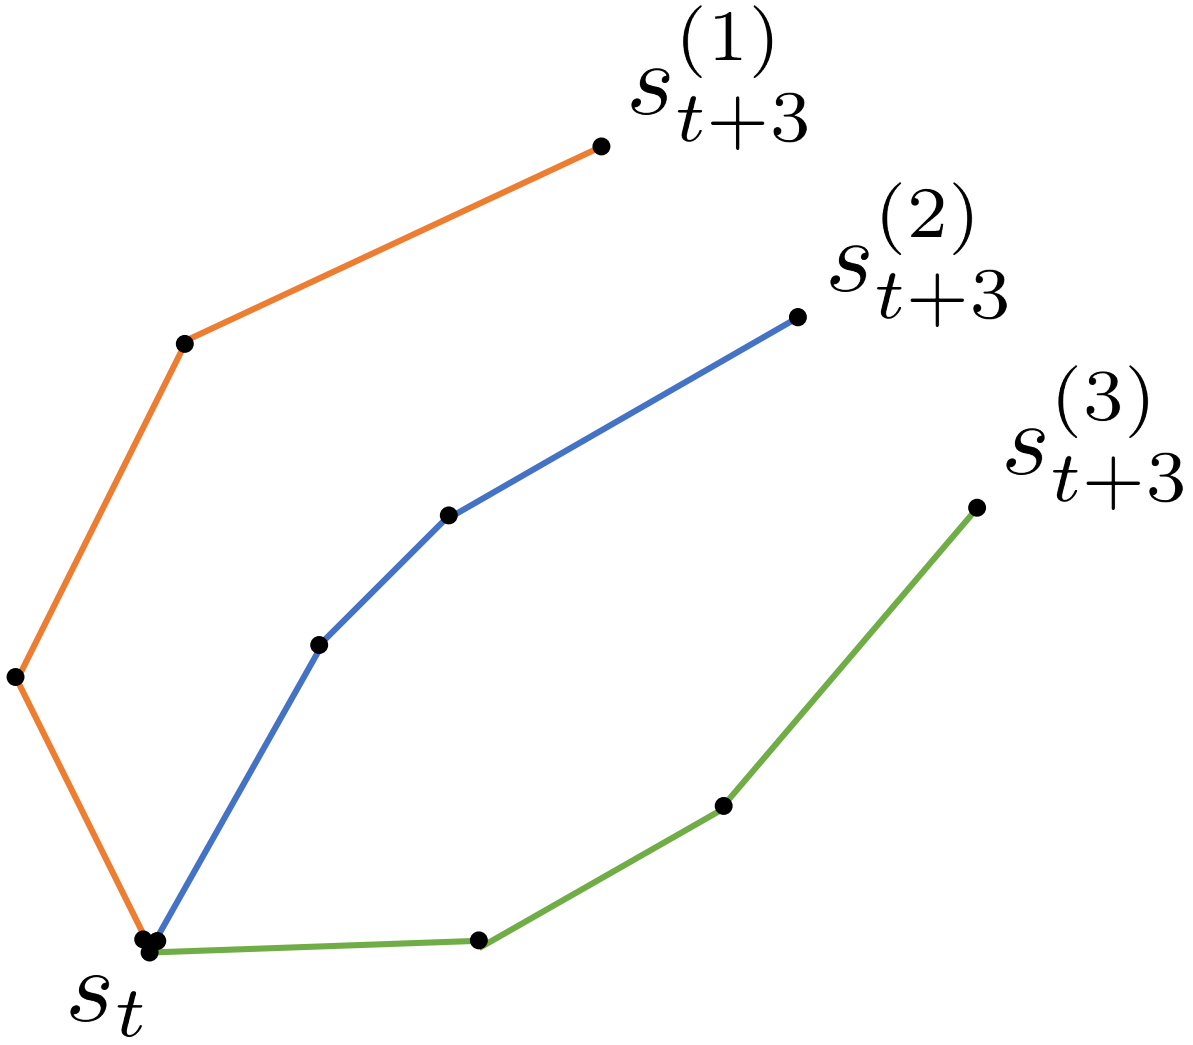

Experimental Results

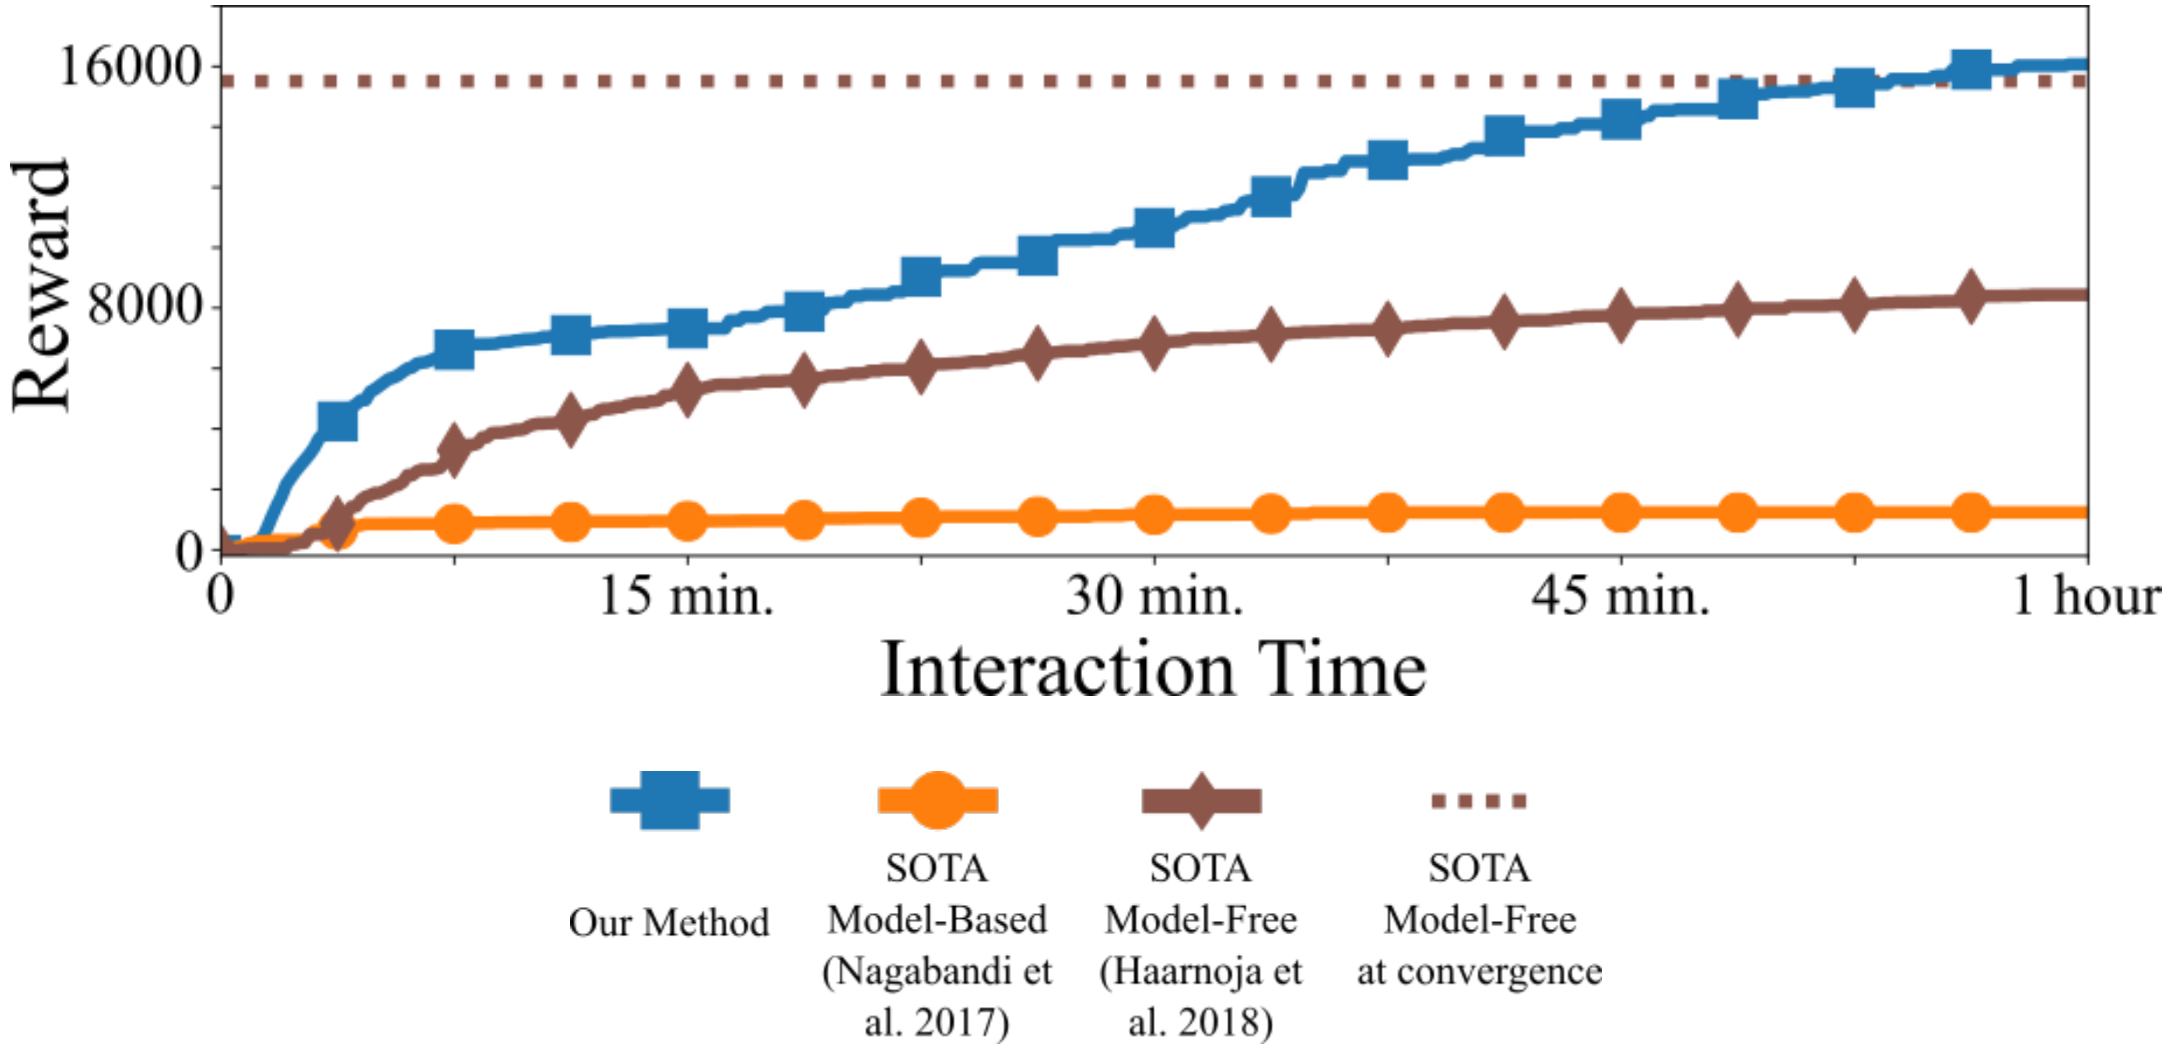

Deep Reinforcement Learning in a Handful of Trials using Probabilistic Dynamics Models

Poster #165

Code: https://github.com/kchua/handful-of-trials

Website: https://sites.google.com/view/drl-in-a-handful-of-trials

- ✓Data efficient
- ✓Competitive asymptotic performance
- ✓Easy to implement


Kurtland Chua


Roberto Calandra


Rowan McAllister


Sergey Levine