# Bisimulation Metric Computation for Markov Decision Processes

Sherry Shanshan Ruan

*School of Computer Science*

*McGill University*

August 20, 2014

# Overview

1. Background
2. Motivating Example
   - An example of probabilistic bisimulations
3. Problem and Motivation
   - Bisimulation metrics
   - Related work on bisimulation metric computation
4. Approach and Contributions
   - First contribution: metrics over partitions
   - Second contribution: Kantorovich metric over partitions
   - Third contribution: theoretical guarantees
   - Fourth contribution: experiments on synthetic data
5. Future Work

# Background

- Markov Decision Process (MDPs) are a powerful mathematical model in planning and learning under uncertainty
- A major task in the study of MDPs is to compute the optimal value function
- The state space of MDPs is often too large to apply dynamic programming algorithms
- Use *probabilistic bisimulation* to aggregate states and reduce the size of the state space

# Probabilistic bisimulation relations

## Definition (probabilistic bisimulation relations)

Given a labeled Markov process $(S, \mathcal{L}, R, P)$, a *probabilistic bisimulation relation* is an equivalence relation $\sim$ on $S$ such that if $s \sim t$ then

$$(1)\ \forall a \in \mathcal{L}, R_{sa} = R_{ta}$$

$$(2)\ \forall a \in \mathcal{L}, \forall c \in S/\sim, P^a_{sc} = P^a_{tc}$$

# An example of probabilistic bisimulations

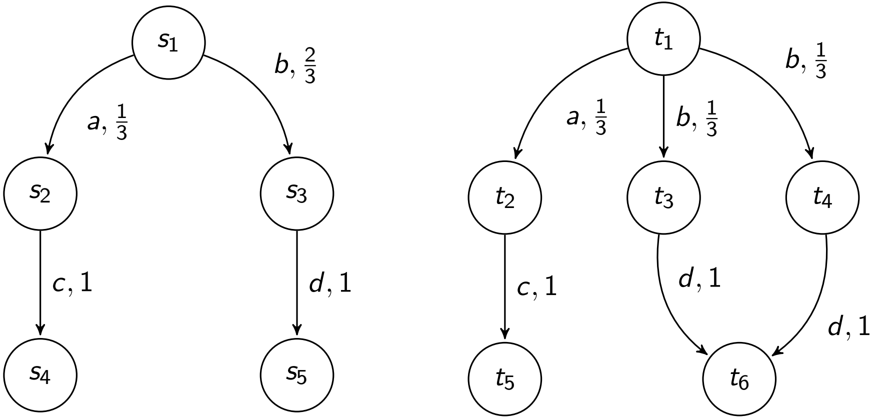


Quick Question:

Which states are bisimilar?

# An example of probabilistic bisimulations

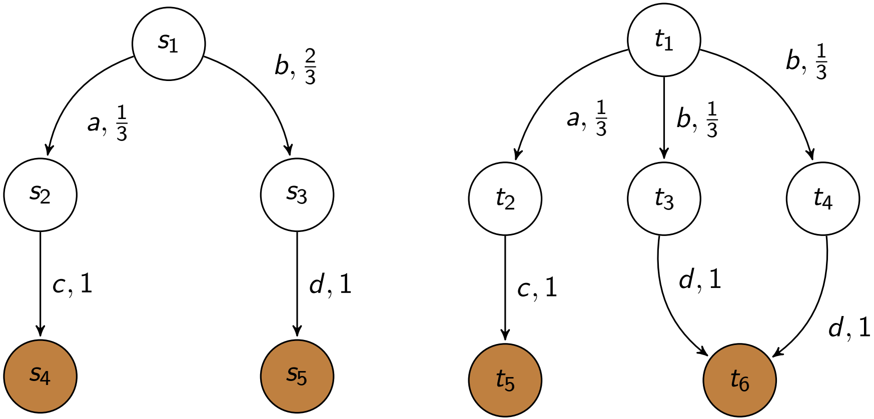


$$s_4 \sim s_5 \sim t_5 \sim t_6$$

# An example of probabilistic bisimulations

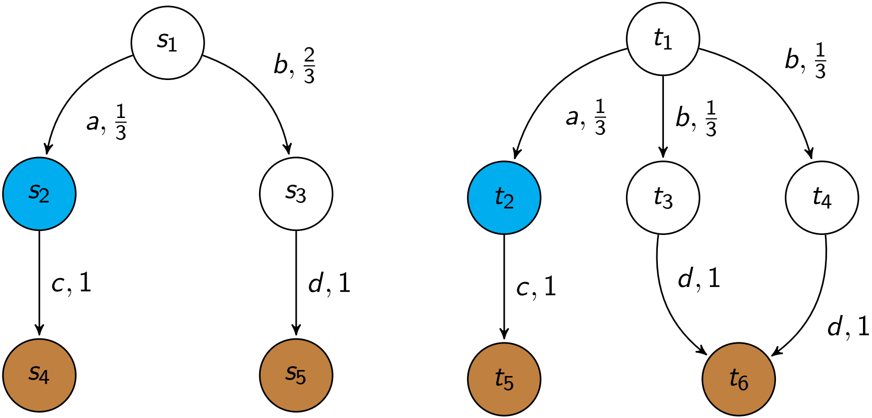


$$s_4 \sim s_5 \sim t_5 \sim t_6$$

$$s_2 \sim t_2$$

# An example of probabilistic bisimulations

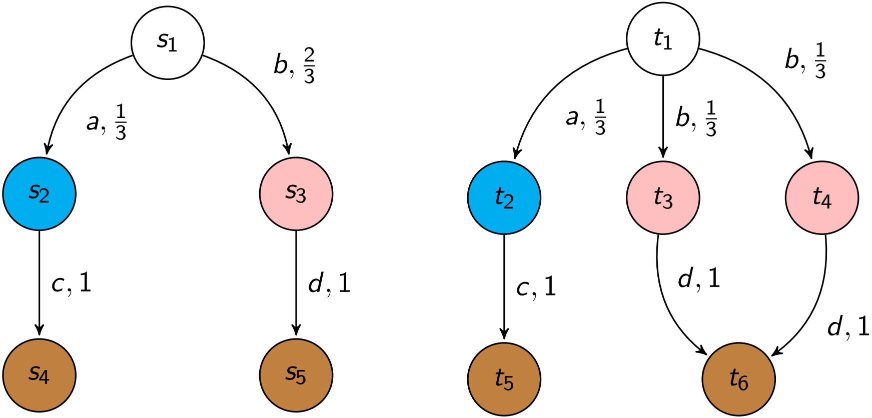


$$s_4 \sim s_5 \sim t_5 \sim t_6$$

$$s_2 \sim t_2 \quad s_3 \sim t_3 \sim t_4$$

# An example of probabilistic bisimulations

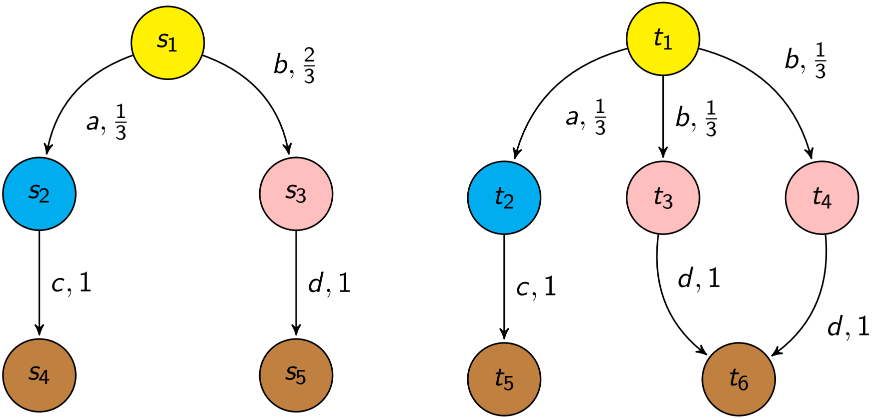


$$s_4 \sim s_5 \sim t_5 \sim t_6$$

$$s_2 \sim t_2 \quad s_3 \sim t_3 \sim t_4 \quad s_1 \sim t_1$$

# Bisimulation metrics

- *Bisimulation metrics* are a quantitative analogue of *bisimulation relations*
- Metrics are smooth with respect to changes on the transition probabilities over the state space
- *Bisimulation metrics* can also be used to aggregate states ($\epsilon$-neighborhood)

# Bisimulation metrics

## Definition (bisimulation metrics)

Let $s$ and $s'$ be two states, the bisimulation metric is a fixed point of of $F$ defined as follows

$$F(d)(s, s') := \max_a (|R_{sa} - R_{s'a}| + \gamma \mathcal{T}(d, P_{sa}, P_{s'a}))$$

- To compute *bisimulation metrics*, we need to measure the distance between probability distributions
- We can use different metrics to compare probability distributions

# Related work on bisimulation metric computation

- Ferns et al. (2004) propose an iteratively improving approximation algorithm
  - They prove convergence
  - **Restrictive**: one has to compare all pairs of states at every iteration
- Comanici et al. (2012) incorporate an on-the-fly technique focusing on a partial set of pairs of states
  - They adopt *heuristics* to select certain pairs of states and only update distances between them
  - **Restrictive**: the metric is still over the entire state space

# First contribution: metrics over partitions

- At each time step, we partition the state space into **blocks** based on the *approximate bisimulation*
- The metric is over partitions instead of the entire state space
- We adopt *heuristics* to select certain pairs of **blocks** at each iteration and only update distances between them
- The *heuristics* provides more flexibility in designing strategies for a variety of MDPs

# Second contribution: Kantorovich metric over partitions

- We use the *Kantorovich metric* to compare probability distributions
- The metric finds the **optimal** coupling of two probability distributions

### Definition (discrete Kantorovich metric)

Let $\Omega$ be a discrete state space , $\mu, \nu$ be two probability measures on $\Omega$, $d$ be a given metric on $\Omega$

$$\mathcal{T}(d, \mu, \nu) := \max \left\{ \left| \sum_{x \in \Omega} h(x)\mu(x) - \sum_{x \in \Omega} h(x)\nu(x) \right| : \|h\|_L \leq 1 \right\}$$

The max is taken over all $h$ s.t. $|h(x) - h(y)| \leq d(x, y)$

# Second contribution: Kantorovich metric over partitions

- *Kantorovich metric* can be transformed into transportation problems by Kantorovich-Rubinstein Duality (max $\rightarrow$ min)
- The minimal cost for transforming one histogram into the other is defined as *the Earth Mover's Distance (EMD)*
- We adopt *the Earth Mover's Distance (EMD)* algorithms (Rubner et al., 2000) to find the best coupling of probability distributions
- By using our partition strategy, the algorithm only needs to search couplings over partitions instead of the entire state space

# Third contribution: theoretical guarantees

- We formalize the asynchronous partition algorithm in a mathematical framework
- We establish the correspondence between the theoretical characterization and the algorithm
- We prove that our asynchronous partition algorithm converges to the same bisimulation metric

# Fourth contribution: experiments on synthetic data

- We implement the iterative asynchronous partition algorithm described previously
- We run experiments on synthetic data

# Future work

- Extend experiments on synthetic data to real data with large state space
- Design elaborate heuristics based on different data to further optimize the algorithm

# Thank You