Robust Probabilistic Bisimilarity for Labelled Markov Chains

ZAINAB FATMI

Joint-research with Franck van Breugel, Stefan Kiefer, and David Parker

CAV July 2025

Labelled Markov Chains

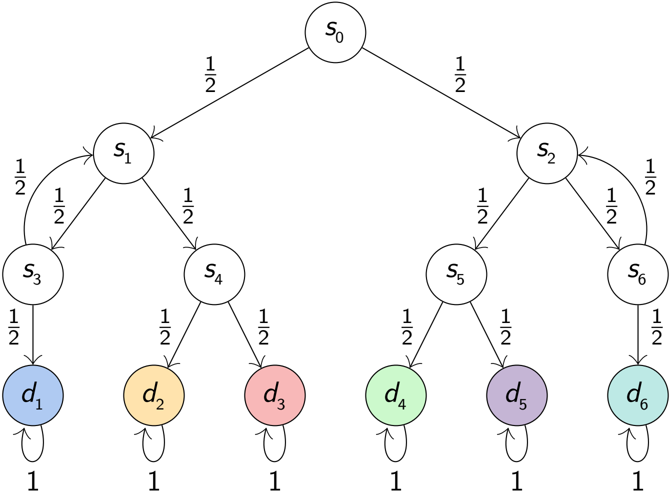

Probabilistic Bisimilarity

One of the major challenges of model checking is that the number of states in a model is often too large to check non-trivial properties of the system, called the state space explosion problem.

Larsen and Skou (1989) introduced the concept of probabilistic bisimulation, which is now implemented in probabilistic model checking tools such as PRISM and Storm.

Probabilistic Bisimilarity

Probabilistic bisimilarity is an equivalence relation on states.

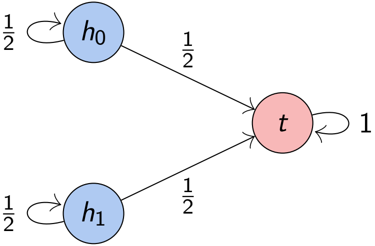

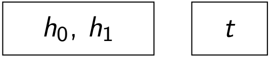

Example 1

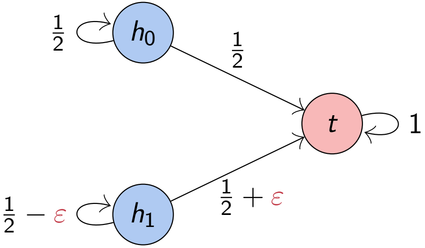

$$\varepsilon = \frac{1}{8}$$

h_0	h_1	t

Probabilistic Bisimilarity Distance

Probabilistic bisimilarity distances, introduced by Desharnais et al. (1999), are a generalization of probabilistic bisimilarity. The distances between states captures the similarity of their behaviour.

The pseudometric has the following properties:

- two states have distance 0 iff they are probabilistic bisimilar, and
- if two states have different labels then they have distance 1.

Example 1

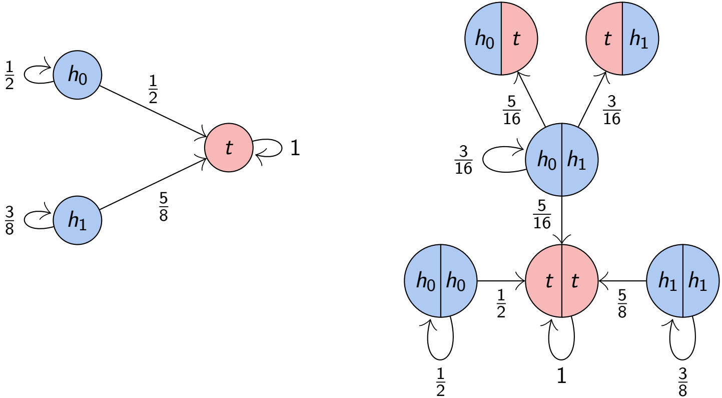

1/2
h0
1/2
t
1
3/8
h1
5/8
h0 t
t h1
5/16
3/16
3/16
h0 h1
5/16
h0 h0
1/2
t t
5/8
h1 h1
1/2
1
3/8

Example 1

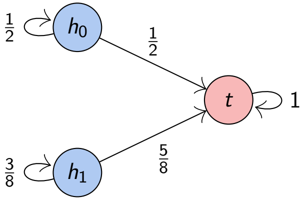

$$\delta_{\tau_\varepsilon}(h_0, h_1) = \frac{1}{5}$$

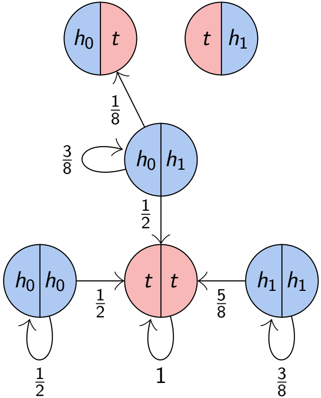

Example 1

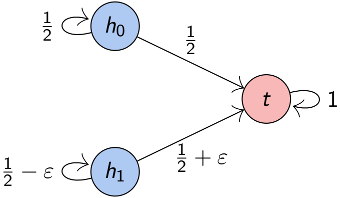

$\frac{1}{2}$
h_0
$\frac{1}{2}$
t
1
$\frac{1}{2} - \varepsilon$
h_1
$\frac{1}{2} + \varepsilon$

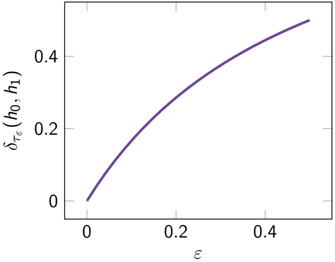

$\delta_{\tau_\varepsilon}(h_0, h_1)$
0.4
0.2
0
0
0.2
0.4
ε

Example 2

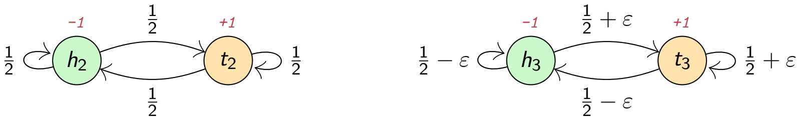

$$\varepsilon = \frac{1}{100}$$

Assume that you earn a point each time the coin lands on *tails* and lose a point each time it lands on *heads*.

Example 2

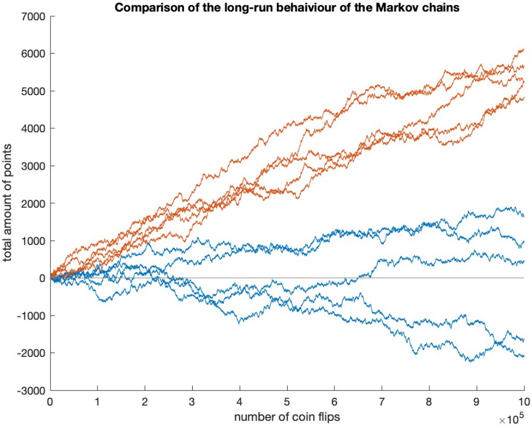

Comparison of the long-run behaviour of the Markov chains
7000
6000
5000
4000
3000
2000
1000
0
-1000
-2000
-3000
total amount of points
0
1
2
3
4
5
6
7
8
9
10
number of coin flips
×10^5

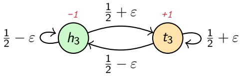

−1
+1
h_3
t_3
$\frac{1}{2}+\varepsilon$
$\frac{1}{2}-\varepsilon$
$\frac{1}{2}-\varepsilon$
$\frac{1}{2}+\varepsilon$

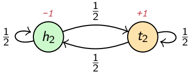

−1
+1
h_2
t_2
$\frac{1}{2}$
$\frac{1}{2}$
$\frac{1}{2}$
$\frac{1}{2}$

Example 2

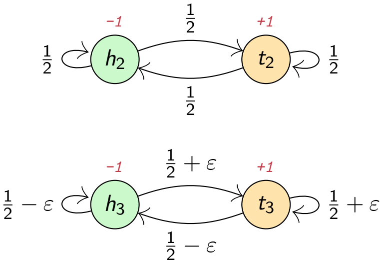

−1
+1
$\frac{1}{2}$
h_2
t_2
$\frac{1}{2}$
$\frac{1}{2}$
$\frac{1}{2}$
−1
+1
$\frac{1}{2}+\varepsilon$
h_3
t_3
$\frac{1}{2}-\varepsilon$
$\frac{1}{2}+\varepsilon$
$\frac{1}{2}-\varepsilon$

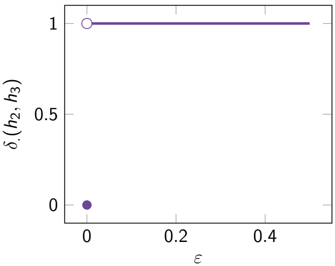

$\delta_{\cdot}(h_2, h_3)$
1
0.5
0
0
0.2
0.4
ε

Why is this a Problem?

The probabilities of the labelled Markov chain are usually obtained experimentally and, therefore, are often an approximation.

This undermines the reliability of probabilistic bisimilarity as a measure of system equivalence.

Example 1 Revisited

Robust probabilistic bisimilarity := there exists a coupling that reaches a "diagonal" pair (s, s) almost surely.

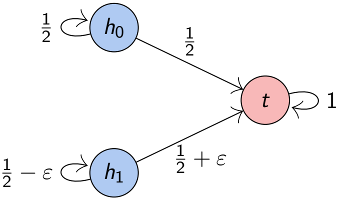

$\varepsilon = 0$
$\delta_{\tau_\varepsilon}(h_0, h_1) = 0$

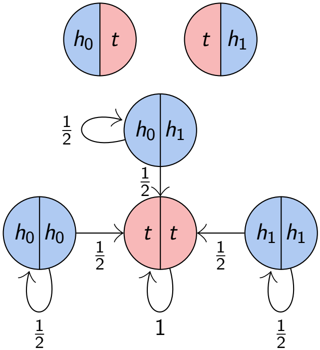

Robust probabilistic bisimilarity is an equivalence relation on states; it ensures that the bisimilarity distance is continuous.

Algorithm (partition-refinement based)

In the product Markov chain, we say that a set $R \subseteq S \times S$ stabilizes a pair of states (u_0, v_0) if there exists a path

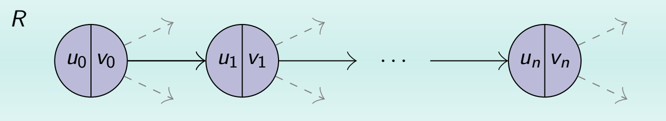

such that

1. $u_n = v_n$
2. for all $0 \leq i \leq n$ we have: $(u_i, v_i) \in R$ and $\mathrm{support}((u_i, v_i)) \subseteq R$.

Algorithm (partition-refinement based)

Robust probabilistic bisimilarity is the greatest relation R with the following properties:

1. for every $(s, t) \in R$ there exists a coupling such that R stabilizes (s, t)
2. R is an equivalence relation
3. R is a bisimulation

It can be computed in polynomial time $\mathcal{O}(n^6)$, with a partition refinement algorithm.

Implementation

We implemented our algorithm in PRISM. We ran experiments on 170 models of up to 116,000 states from

- the Quantitative Verification Benchmark Set (QVBS) and
- the Java PathFinder (JPF) extension jpf-probabilistic.

Experimental Results

Outcome

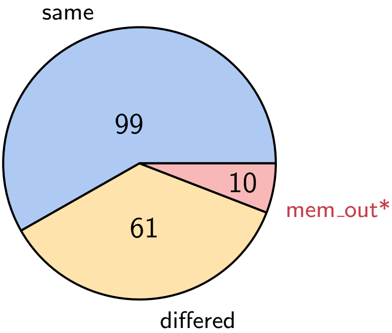

Runtime

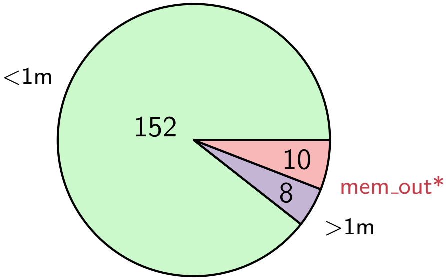

*JVM limited to 8GB

Implementation

UNIVERSITY OF OXFORD DEPARTMENT OF COMPUTER SCIENCE

We implemented our algorithm in PRISM. We ran experiments on 170 models of up to 116,000 states from

- the Quantitative Verification Benchmark Set (QVBS) and
- the Java PathFinder (JPF) extension jpf-probabilistic.

Thus, we can identify states that are bisimilar, but not robustly so, in feasible time.

CAV Artifact Evaluation Available

CAV Artifact Evaluation Reusable

Summary

We introduced the concept of robust probabilistic bisimilarity for labelled Markov chains, which is an equivalence relation that:

1. is a subset of probabilistic bisimilarity, and
2. ensures the continuity of the probabilistic bisimilarity distance function.

We presented a polynomial-time algorithm and showed that it performs well in practice.