

Robust Probabilistic Bisimilarity for Labelled Markov Chains

Syyeda Zainab Fatmi¹, Stefan Kiefer¹, David Parker¹, and Franck van Breugel²



¹ University of Oxford, Oxford, UK
{syyeda.fatmi, stefan.kiefer, david.parker}@cs.ox.ac.uk

² York University, Toronto, Canada
franck.van.breugel@lassonde.yorku.ca



Abstract. Despite its prevalence, probabilistic bisimilarity suffers from a lack of robustness under minuscule perturbations of the transition probabilities. This can lead to discontinuities in the probabilistic bisimilarity distance function, undermining its reliability in practical applications where transition probabilities are often approximations derived from experimental data. Motivated by this limitation, we introduce the notion of robust probabilistic bisimilarity for labelled Markov chains, which ensures the continuity of the probabilistic bisimilarity distance function. We also propose an efficient algorithm for computing robust probabilistic bisimilarity and show that it performs well in practice, as evidenced by our experimental results.

Keywords: (probabilistic) model checking · labelled Markov chain · probabilistic bisimilarity · behavioural pseudometric.

1 Introduction

In the analysis and verification of probabilistic systems, one of the foundational concepts is identifying and merging system states that are behaviourally indistinguishable. Kemeny and Snell [22] introduced the notion of lumpability for Markov chains and it was adapted to the setting of labelled Markov chains by Larsen and Skou [25], known as probabilistic bisimulation. State-of-the-art probabilistic verification tools [24,18] implement a variety of methods for minimizing the state space of the system by collapsing probabilistically bisimilar states. This can significantly improve verification efficiency in some cases [21].

However, due to the sensitivity of behavioural equivalences to small changes in the transition probabilities, Giacalone et al. [16] proposed using behavioural pseudometrics to capture the behavioural similarity of states in a probabilistic system. Instead of classifying states as either equivalent or inequivalent, the pseudometric maps each pair of states to a real value in the unit interval, thus also quantifying the behavioral difference between non-equivalent states. Behavioural pseudometrics have been studied in the context of systems biology [32], games [5], planning [8] and security [4], among others.

In probabilistic verification, the most widely studied example of such a behavioural pseudometric is the *probabilistic bisimilarity distance*. It generalizes probabilistic bisimilarity quantitatively; in particular, the distance between two states is zero if and only if they are probabilistically bisimilar. The probabilistic bisimilarity distance was introduced by Desharnais et al. [10], based on a real-valued semantics for Larsen and Skou’s probabilistic modal logic [25]. A formula φ of this logic maps any state s to a number $\llbracket\varphi\rrbracket(s) \in [0, 1]$. The probabilistic bisimilarity distance between two states s, t can be characterized as $\delta(s, t) = \sup_{\varphi} |\llbracket\varphi\rrbracket(s) - \llbracket\varphi\rrbracket(t)| \in [0, 1]$, where φ ranges over all formulas. The lower the distance between two states, the more similar their behaviour. As shown by Van Breugel and Worrell [3], the probabilistic bisimilarity distance can also be characterized as a fixed point of a function (we use this definition in this paper).

However, as pointed out by Jaeger et al. [19], probabilistic bisimilarity distances are sometimes not continuous, leading to unexpected and abrupt changes in behaviour between two states when the transition probabilities are perturbed. Since the probabilities of the labelled Markov chain are usually obtained experimentally and, therefore, are often an approximation [31,13,27,28,30], the lack of robustness of probabilistic bisimilarity is a serious drawback. This inconsistency undermines the reliability of probabilistic bisimilarity as a measure of system equivalence and can be particularly problematic when used in practical applications where approximate models are prevalent.

Example 1. Consider Figure 1a on page 6. When $\varepsilon = 0$, states h_0 and h_1 are probabilistically bisimilar; i.e., their distance $\delta_0(h_0, h_1)$ equals 0 (the subscript of δ indicates ε). For $\varepsilon > 0$ we have $\delta_\varepsilon(h_0, h_1) > 0$; i.e., h_0 and h_1 are no longer bisimilar. However, when ε is small then $\delta_\varepsilon(h_0, h_1)$ is small. In fact, one can show that $\delta_\varepsilon(h_0, h_1) \leq 2\varepsilon$, which implies that $\lim_{\varepsilon \rightarrow 0} \delta_\varepsilon(h_0, h_1) = 0$. This means that in this example, the distance is continuous in ε . One may say that states h_0, h_1 are not only probabilistically bisimilar, but also robustly so, in that they remain “almost” bisimilar when the transition probabilities are perturbed. Intuitively, states h_0 and h_1 behave similarly even for small positive ε : both states carry a blue label and perform a geometrically distributed number of self-loops (about two in expectation) before transitioning to state t .

Example 2. Consider Figure 1b on page 6. When $\varepsilon = 0$, states h_2 and h_3 are probabilistically bisimilar; i.e., their distance $\delta_0(h_2, h_3)$ equals 0. But for any $\varepsilon > 0$ we have $\delta_\varepsilon(h_2, h_3) = 1$; i.e., h_2 and h_3 behave “maximally” differently in terms of the probabilistic bisimilarity distance. We have $\lim_{\varepsilon \rightarrow 0} \delta_\varepsilon(h_2, h_3) = 1$; so, in this example, the distance is discontinuous in ε . One may say that although states h_2, h_3 are probabilistically bisimilar, they are not robustly so, because upon perturbing the transition probabilities the behaviour of h_3 changes completely. For any positive ε , state h_2 remains in a self-loop forever, whereas h_3 eventually reaches the (red-labelled) state t_3 with probability 1. Since reachability properties are at the heart of probabilistic model checking, it may be unsafe to merge states h_2 and h_3 if the transition probabilities are not known precisely.

In this paper, we address this issue by introducing the notion of *robust probabilistic bisimilarity* for labelled Markov chains. Robust probabilistic bisimilarity is a particular probabilistic bisimulation, implying that robust probabilistic bisimilarity is a subset of probabilistic bisimilarity. Crucially, we show that our definition ensures the continuity of the probabilistic bisimilarity distance function. This means that for any two states that are robustly probabilistically bisimilar, their probabilistic bisimilarity distance remains small even after small perturbations of any transition probabilities. Note that it is easy to see that the distance from [12] is robust in this sense; on the other hand, states with very small distance in terms of [12] can have very different long-term behaviour, as in Example 2.

Secondly, we develop a polynomial-time algorithm for computing robust probabilistic bisimilarity. It is suitable for large-scale verification tasks, opening the door to checking probabilistic models from the literature for (lack of) robustness of their probabilistic bisimilarity relation. Thus, one can identify pairs of states that may be dangerous to merge if the transition probabilities are not known precisely. We present experimental results on the applicability and efficiency of an implementation of our algorithm on models from the Quantitative Verification Benchmark Set (QVBS) [17] and the examples included in the Java PathFinder extension jpf-probabilistic [14].

The rest of the paper is structured as follows. Section 2 introduces the model of interest, namely a labelled Markov chain, and probabilistic bisimilarity. In Section 3, we formally define probabilistic bisimilarity distances and further examine how the bisimilarity distance changes when the transition function is varied. Section 4 describes robust probabilistic bisimilarity and demonstrates that it ensures the continuity of the bisimilarity distance function. In Section 5, we present a polynomial-time algorithm for computing robust probabilistic bisimilarity. Section 6 reports experimental results on the algorithm's implementation. Finally, Section 7 concludes the paper and discusses directions for future research. The full version of this paper, found in [15], includes omitted proofs and further details.

2 Labelled Markov Chains and Probabilistic Bisimilarity

In this section, we present some fundamental concepts that underpin this paper.

Let X be a nonempty finite set. A function $\mu : X \rightarrow [0, 1]$ is a *subprobability distribution* on X if $\sum_{x \in X} \mu(x) \leq 1$. We denote the set of subprobability distributions on X by $\mathcal{S}(X)$. For $\mu \in \mathcal{S}(X)$ and $A \subseteq X$, we often write $\mu(A)$ instead of $\sum_{x \in A} \mu(x)$. For a distribution $\mu \in \mathcal{S}(X)$ we define the *support* of μ by $\text{support}(\mu) = \{x \in X \mid \mu(x) > 0\}$. A subprobability distribution μ on X is a *probability distribution* if $\mu(X) = 1$. We denote the set of probability distributions on X by $\mathcal{D}(X)$.

A *Markov chain* is a pair $\langle S, \tau \rangle$ consisting of a finite set S of states and a transition probability function $\tau : S \rightarrow \mathcal{D}(S)$. A *labelled Markov chain* is a tuple $\langle S, L, \tau, \ell \rangle$ where $\langle S, \tau \rangle$ is a Markov chain, L is a finite set of labels and $\ell : S \rightarrow L$

is a labelling function. A *path* in a Markov chain $\langle S, \tau \rangle$ is a sequence of states $s_0, s_1, s_2 \dots$ such that $s_i \in S$ and $\tau(s_i)(s_{i+1}) > 0$ for all $i \geq 0$.

For the remainder, we fix a labelled Markov chain $\langle S, L, \tau, \ell \rangle$, and we will study perturbations of the transition probability function τ .

For all $\mu, \nu \in \mathcal{D}(X)$, the set $\Omega(\mu, \nu)$ of *couplings* of μ and ν is defined by

$$\Omega(\mu, \nu) = \{ \omega \in \mathcal{D}(X \times X) \mid \forall x \in X : \omega(x, X) = \mu(x) \wedge \omega(X, x) = \nu(x) \}.$$

We write $\omega(x, X)$ for $\sum_{y \in X} \omega(x, y)$.

Definition 1. *An equivalence relation $R \subseteq S \times S$ is a probabilistic bisimulation (or just bisimulation) if for all $(s, t) \in R$, $\ell(s) = \ell(t)$ and there exists $\omega \in \Omega(\tau(s), \tau(t))$ such that $\text{support}(\omega) \subseteq R$. States s and t are bisimilar, denoted $s \sim t$, if $(s, t) \in R$ for some bisimulation R .*

If $|\ell(S)| = 1$ then $\sim = S \times S$. In the remainder, we assume that the labelled Markov chain contains states with different labels, that is, $|\ell(S)| \geq 2$. Hence, we also have that $|S| \geq 2$.

Definition 1 [20, Definition 4.3] differs from the standard definition [25, Definition 6.3] which defines a bisimulation as an equivalence relation $R \subseteq S \times S$ such that for all $(s, t) \in R$, $\ell(s) = \ell(t)$ and for all R -equivalence classes C , $\tau(s)(C) = \tau(t)(C)$, where $\tau(s)(C) = \sum_{t \in C} \tau(s)(t)$. Nevertheless, an equivalence relation R is a bisimulation by Definition 1 if and only if it is a bisimulation as per the standard definition (see [20, Theorem 4.6]).

3 Probabilistic Bisimilarity Distances

Definition 2. *The probabilistic bisimilarity distance (or just bisimilarity distance), $\delta_\tau : S \times S \rightarrow [0, 1]$, is the least fixed point of the function $\Delta_\tau : (S \times S \rightarrow [0, 1]) \rightarrow (S \times S \rightarrow [0, 1])$ defined by*

$$\Delta_\tau(d)(s, t) = \begin{cases} 1 & \text{if } \ell(s) \neq \ell(t) \\ \inf_{\omega \in \Omega(\tau(s), \tau(t))} \sum_{u, v \in S} \omega(u, v) d(u, v) & \text{otherwise.} \end{cases}$$

Intuitively, the smaller the distance between two states, the more similar they behave.

Theorem 1 ([11, Theorem 4.10]). *For all $s, t \in S$, $s \sim t$ if and only if $\delta_\tau(s, t) = 0$.*

Quantitative μ -calculus [23,1,26] is an expressive modal logic that uses fixed point operators to define properties of transition systems. It supports the concise representation of a wide range of properties, including reachability, safety, and the probability of satisfying a general ω -regular specification. We use the syntax described in [5], except that we use the operator next instead of pre_1 and pre_2 . Let $q\mu$ denote the set of quantitative μ -calculus formulae, then a formula $\varphi \in q\mu$ maps states to a numerical value within $[0, 1]$, that is, $\llbracket \varphi \rrbracket : S \rightarrow [0, 1]$. The bisimilarity distances can be characterized in terms of the quantitative μ -calculus [5, Equation 2.3] as $\delta_\tau(s, t) = \sup_{\varphi \in q\mu} |\llbracket \varphi \rrbracket(s) - \llbracket \varphi \rrbracket(t)|$.

3.1 Examples

We now investigate how the bisimilarity distance changes when the transition function varies. In the following, let $\varepsilon \in [0, \frac{1}{2}]$. Define τ_ε as shown in Figure 1. For example, $\tau_{1/6}(h_5)(t_5) = \frac{2}{3}$. Then $\tau_- : [0, \frac{1}{2}] \rightarrow (S \rightarrow \mathcal{D}(S))$ and $\delta_{\tau_-} : [0, \frac{1}{2}] \rightarrow (S \times S \rightarrow [0, 1])$.

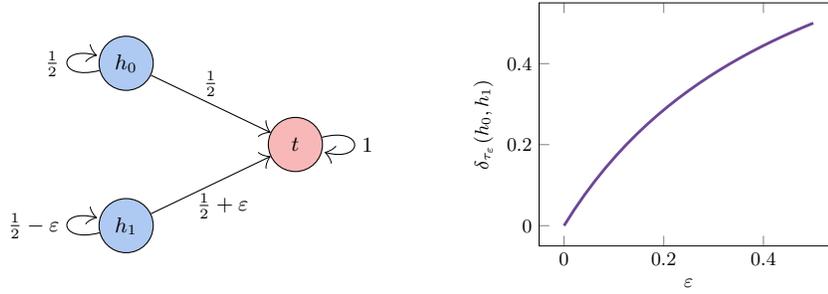
Example 3. Consider Figure 1a. As ε increases, h_1 becomes more biased and the distance between h_0 and h_1 increases proportionally. One can show that $\delta_{\tau_\varepsilon}(h_0, h_1) = \frac{\varepsilon}{0.5+\varepsilon} \leq 2\varepsilon$. Note that if ε is small then the distance is also small and $\lim_{\varepsilon \rightarrow 0} \delta_{\tau_\varepsilon}(h_0, h_1) = 0$.

The formula $\varphi = \mu V. \text{next}(\text{tails} \vee V) \ominus 0.5$ distinguishes the states h_0 and h_1 the most, that is $\delta_{\tau_\varepsilon}(h_0, h_1) = \frac{\varepsilon}{0.5+\varepsilon} = |\llbracket \varphi \rrbracket(h_0) - \llbracket \varphi \rrbracket(h_1)|$. The quantifier μ denotes the least fixed point of the recursive formula involving the variable V . Intuitively, a state satisfies V if the next state is *tails* or satisfies V with probability greater than a half. More precisely, considering state h_1 , $\llbracket \varphi \rrbracket(h_1)$ is the expected value of $\max(\llbracket \varphi \rrbracket(s), \llbracket \text{tails} \rrbracket(s)) - \frac{1}{2}$, where s denotes the random successor state of h_1 . Then, $\llbracket \varphi \rrbracket(h_1)$ evaluates to $\sum_{n=0}^{\infty} \varepsilon (\frac{1}{2} - \varepsilon)^n = \frac{\varepsilon}{0.5+\varepsilon}$. Each summand in the series represents the probability of reaching state t in $n+1$ steps, starting from state h_1 , with 0.5 subtracted at each step. On the other hand, $\llbracket \varphi \rrbracket(h_0) = 0$.

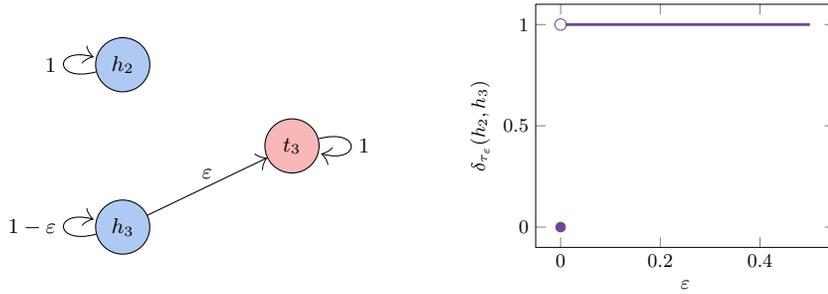
Example 4. In Figure 1b, when $\varepsilon = 0$, the states h_2 and h_3 are bisimilar with $\delta_{\tau_0}(h_2, h_3) = 0$. However, if $\varepsilon > 0$, then $\delta_{\tau_\varepsilon}(h_2, h_3) = 1$. This difference is evident when considering the probability of eventually reaching a state labelled with *tails* when starting in h_2 compared to h_3 . In the first Markov chain, $\llbracket \diamond \text{tails} \rrbracket = 0$, while in the second Markov chain, $\llbracket \diamond \text{tails} \rrbracket = 1$. This property can be expressed as the quantitative μ -calculus formula $\mu V. \text{next}(\text{tails} \vee V)$. This example was also presented in [19].

Example 5. The first Markov chain in Figure 1c represents fair coin flips, while the second Markov chain represents potentially biased coin flips. When $\varepsilon = 0$, the states h_4 and h_5 are bisimilar with $\delta_{\tau_0}(h_4, h_5) = 0$. However, if $\varepsilon > 0$, one can show that $\delta_{\tau_\varepsilon}(h_4, h_5) = 1$. Intuitively, this is because small differences in probabilities can compound and lead to qualitative differences in the long-run behaviour.

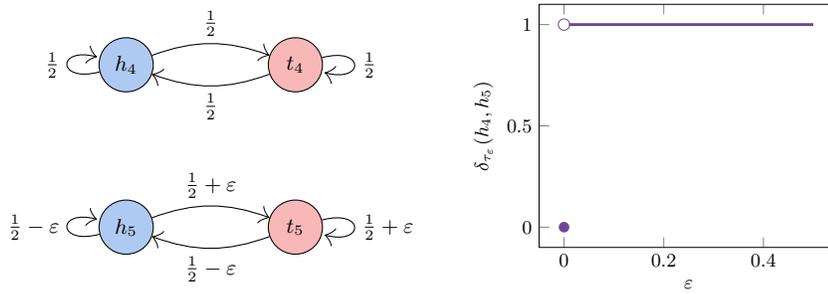
Let us illustrate this. Assume that a point is awarded each time the coin lands on *tails* and a point is deducted each time it lands on *heads*. Let us examine the limit behaviour of the Markov chains. Observe that the Markov chains behave like a random walk on the integer number line, \mathbb{Z} , starting at 0. At each step, the first Markov chain goes up by one with probability $\frac{1}{2}$ and down by one with probability $\frac{1}{2}$. On the other hand, at each step, the second Markov chain goes up by one with probability $\frac{1}{2} + \varepsilon$ and down by one with probability $\frac{1}{2} - \varepsilon$. Let Y_1, Y_2, Y_3, \dots be the sequence of independent random variables, where Y_i denotes the i^{th} step taken by the random walk, with $Y_i = 1$ for a step up and $Y_i = -1$ for a step down. Define $S_n = \sum_{i=1}^n Y_i$. In the first Markov chain, $P(\liminf_{n \rightarrow \infty} S_n = -\infty) = 1$ and $P(\limsup_{n \rightarrow \infty} S_n = \infty) = 1$, by the Hewitt-Savage zero-one law [7, Example 5.19]. In contrast, in the second Markov chain,



(a) Repeated tosses of a fair coin (top) and a biased coin (bottom) until each lands on tails.



(b) Single toss of a rigged coin (top) and repeated tosses of an extremely biased coin until it lands on tails (bottom).



(c) Repeated tosses of a fair coin (top) and a biased coin (bottom).

Fig. 1: Various examples featuring fair and biased coins. States labeled with *heads* are shown in blue, while states labeled with *tails* are shown in red.

we have $P(\lim_{n \rightarrow \infty} S_n = \infty) = 1$, by the law of large numbers [29]. Thus, in the first Markov chain, with equal chances of gaining or losing points at each step, the random walk almost surely oscillates infinitely. In contrast, in the second Markov chain, the upward bias introduced by $\varepsilon > 0$ guarantees that the total number of points will eventually diverge to $+\infty$.

We see that small changes in the transition probabilities can lead to significant changes in the behaviour and, thus, in the distances between states. This example is similar to the one presented in [19].

In the remainder, we conservatively assume that the transition function can be varied arbitrarily, that is, changes to the transition function are not restricted to specific transitions with constrained variables as in the examples. Also, different from [19], the changes might “add transitions.” Therefore, we are interested in the continuity of the function $\delta_{\tau}(s, t) : (S \rightarrow \mathcal{D}(S)) \rightarrow [0, 1]$. See [15, Appendix A] for the metric on distributions $S \rightarrow \mathcal{D}(S)$ used here.

The bisimilarity distance function $\delta_{\tau}(s, t)$ is *lower semi-continuous* at τ if for any sequence $(\tau_n)_n$ converging to τ we have $\liminf_n \delta_{\tau_n}(s, t) \geq \delta_{\tau}(s, t)$ and *upper semi-continuous* at τ if we have $\limsup_n \delta_{\tau_n}(s, t) \leq \delta_{\tau}(s, t)$. Lastly, $\delta_{\tau}(s, t)$ is *continuous* at τ if it is both upper semi-continuous and lower semi-continuous at τ .

The examples in Figure 1 suggest that the bisimilarity distance function δ_{τ} is lower semi-continuous at τ . Indeed the following proposition shows that this holds in general, even allowing for arbitrary modifications of τ .

Proposition 1. *For all $s, t \in S$, the function $\delta_{\tau}(s, t) : (S \rightarrow \mathcal{D}(S)) \rightarrow [0, 1]$ is lower semi-continuous at τ , that is, if $(\tau_n)_n$ converges to τ then $\liminf_n \delta_{\tau_n}(s, t) \geq \delta_{\tau}(s, t)$.*

In Figure 1c, the bisimilarity distance function is not upper semi-continuous. Specifically, $\limsup_{\epsilon \rightarrow 0} \delta_{\tau_{\epsilon}}(h_4, h_5) = 1$, while $\delta_{\tau_0}(h_4, h_5) = 0$. As a result, small perturbations of τ cause a jump in the distance from 0 to 1. The main goal of this paper is to characterize and identify the continuity of the bisimilarity distance function for bisimilar pairs of states.

The following subsets of $S \times S$ play a key role in the subsequent discussion.

Definition 3. *The sets S_{Δ}^2 , $S_{0,\tau}^2$, S_1^2 , $S_{?,\tau}^2$, and $S_{0?}^2$ are defined by*

$$\begin{aligned} S_{\Delta}^2 &= \{ (s, s) \mid s \in S \} \\ S_{0,\tau}^2 &= \{ (s, t) \in S \times S \mid s \neq t \wedge s \sim t \} \\ S_1^2 &= \{ (s, t) \in S \times S \mid \ell(s) \neq \ell(t) \} \\ S_{?,\tau}^2 &= (S \times S) \setminus (S_{\Delta}^2 \cup S_{0,\tau}^2 \cup S_1^2) \\ S_{0?}^2 &= S_{0,\tau}^2 \cup S_{?,\tau}^2 \end{aligned}$$

The first four sets form a partition of $S \times S$. Observe that the sets $S_{0,\tau}^2$ and $S_{?,\tau}^2$ depend on τ and may, therefore, change when we perturb τ , whereas the sets S_{Δ}^2 and S_1^2 stay the same. Note that $S_{0?}^2 = (S \times S) \setminus (S_{\Delta}^2 \cup S_1^2)$. Hence, this set

also stays the same if we perturb τ . Furthermore, note that $\sim = S_{\Delta}^2 \cup S_{0,\tau}^2$ and for all $(s, t) \in S_1^2$, we have $\delta_{\tau}(s, t) = 1$.

Definition 4. Let $\tau : S \rightarrow \mathcal{D}(S)$. The set \mathcal{P}_{τ} of policies for τ is defined by

$$\mathcal{P}_{\tau} = \left\{ P : S \times S \rightarrow \mathcal{D}(S \times S) \left| \begin{array}{l} \forall (s, t) \in S_{\Delta}^2 \cup S_{0,\tau}^2 : P(s, t) \in \Omega(\tau(s), \tau(t)) \\ \forall (s, t) \in S_1^2 : \text{support}(P(s, t)) = \{(s, t)\} \end{array} \right. \right\}.$$

Note that a policy $P \in \mathcal{P}_{\tau}$ induces a Markov chain $\langle S \times S, P \rangle$. The subscript τ is omitted when clear from the context. The following proposition characterizes δ_{τ} in terms of policies.

Proposition 2. For all $s, t \in S$, $\delta_{\tau}(s, t) = \min_{P \in \mathcal{P}} \gamma_P$, where γ_P is the probability with which (s, t) reaches S_1^2 in $\langle S \times S, P \rangle$.

Proof Sketch. The proof follows from [2, Theorem 10.15] and [6, Theorem 8]. \square

Example 6. Consider the labelled Markov chain in Figure 1a when $\varepsilon = \frac{1}{8}$. Then the probability with which (h_0, h_1) reaches S_1^2 for any policy $P \in \mathcal{P}$ is $\geq \frac{1}{5}$. Any policy P such that $P(h_0, h_1) = \{(h_0, h_1) \mapsto \frac{3}{8}, (h_0, t) \mapsto \frac{1}{8}, (t, t) \mapsto \frac{1}{2}\}$ achieves the minimum probability of $\frac{1}{5}$. The Markov chain induced by such a policy P is illustrated in Figure 2. Thus, $\delta_{\tau_{\varepsilon}}(h_0, h_1) = \frac{1}{5}$.

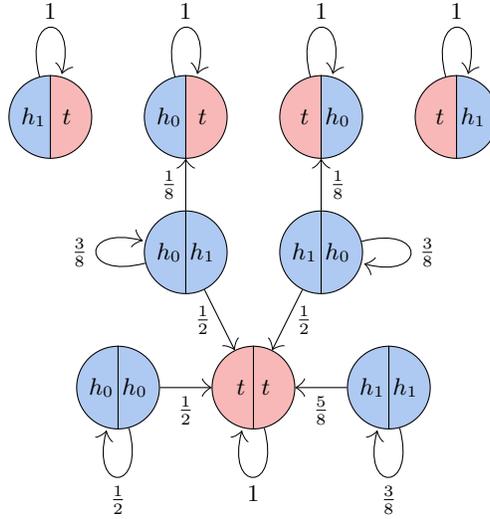


Fig. 2: The Markov chain $\langle S \times S, P \rangle$ induced by the policy P such that (h_0, h_1) reaches S_1^2 with probability $\frac{1}{5}$.

4 Robust Probabilistic Bisimilarity

We aim to define a notion of robust bisimilarity which is a bisimulation that is robust against perturbations of the transition function τ . As we will see in Theorem 2 below, the following definition fulfills this requirement.

Definition 5. Robust probabilistic bisimilarity (or just robust bisimilarity), denoted \simeq , is defined for $s, t \in S$ as $s \simeq t$ if there exists a policy $P \in \mathcal{P}$ such that (s, t) reaches S_{Δ}^2 with probability 1 in $\langle S \times S, P \rangle$.

Lemma 1. Robust bisimilarity, \simeq , is a bisimulation.

Proof Sketch. Clearly, \simeq is reflexive and symmetric. We prove in [15, Appendix] that \simeq is transitive as well and, therefore, an equivalence relation.

Let $s, t \in S$ such that $s \simeq t$. Let $P \in \mathcal{P}$ be the policy such that (s, t) reaches S_{Δ}^2 with probability 1 in $\langle S \times S, P \rangle$. Then, it follows from the definition of \mathcal{P} that $(s, t) \notin S_1^2$. Thus, $\ell(s) = \ell(t)$.

Let $\omega = P(s, t)$, $u, v \in S$ and $(u, v) \in \text{support}(\omega)$. Hence, $\omega(u, v) > 0$ and (u, v) is reachable from (s, t) . Therefore, (u, v) must reach S_{Δ}^2 with probability 1 in $\langle S \times S, P \rangle$. Consequently, $u \simeq v$. As a result, $\text{support}(\omega) \subseteq \simeq$. \square

Therefore, $\simeq \subseteq \sim$ and, by Theorem 1, for any $s, t \in S$ such that $s \simeq t$ we have $\delta_{\tau}(s, t) = 0$.

Example 7. In Figure 1a, when $\varepsilon = 0$, then $h_0 \simeq h_1$, since there exists a policy $P \in \mathcal{P}$ such that (h_0, h_1) reaches $(t, t) \in S_{\Delta}^2$ with probability 1 in $\langle S \times S, P \rangle$. Indeed, take $P(h_0, h_1) = \{(h_0, h_1) \mapsto \frac{1}{2}, (t, t) \mapsto \frac{1}{2}\}$ as shown in Figure 3. Hence, $h_0 \simeq h_1$. Note, however, that $h_2 \not\simeq h_3$ and $h_4 \not\simeq h_5$.

The following theorem provides the rationale behind the term robust bisimilarity. It establishes that for all robust bisimilar pairs of states, small perturbations of τ result in a correspondingly small change in the distance between them.

Theorem 2. For all $s, t \in S$, if $s \simeq t$ then the function $\delta_{-}(s, t) : (S \rightarrow \mathcal{D}(S)) \rightarrow [0, 1]$ is continuous at τ , that is, for any sequence $(\tau_n)_n$ converging to τ we have $\lim_n \delta_{\tau_n}(s, t) = 0$.

Proof Sketch. To build some intuition behind this theorem, we first outline the underlying idea. Let $P \in \mathcal{P}$ be the policy such that (s, t) reaches S_{Δ}^2 with probability 1 in $\langle S \times S, P \rangle$. Then, for some k , the probability of (s, t) reaching S_{Δ}^2 within k steps is almost one, say $1 - x$, where $x > 0$ is a small value. When the transition function τ is perturbed by a small ε , there is a policy P' such that the transitions in $\langle S \times S, P' \rangle$ differ from those in $\langle S \times S, P \rangle$ only by a small $\varepsilon' > 0$. Therefore, (s, t) still reaches S_{Δ}^2 with high probability in $\langle S \times S, P' \rangle$.

To argue the last point in slightly more detail, observe that if $\varepsilon > 0$ is small enough so that $(1 - \varepsilon')^k \geq 1 - x$ then the probability, say p , of any individual path of length at most k from (s, t) to S_{Δ}^2 remains at least $(1 - x) \cdot p$ after the

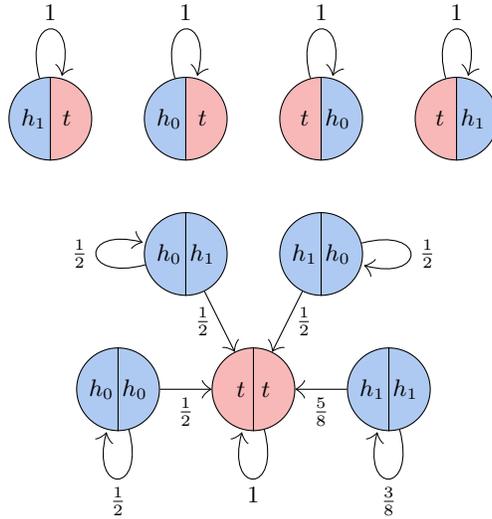


Fig. 3: The Markov chain $\langle S \times S, P \rangle$ induced by the policy P such that (h_0, h_1) reaches S_{Δ}^2 with probability 1.

perturbation. It follows that the probability of *all* paths of length at most k from (s, t) to S_{Δ}^2 remains at least $(1 - x) \cdot (1 - x) \geq 1 - 2x$ after the perturbation.

In [15, Appendix], we provide a different, formal proof using matrix norms. There we construct a graph consisting of the closed communication classes of $\langle S \times S, P \rangle$ that are reachable from (s, t) . Let $P_n \in \mathcal{P}_{\tau_n}$. We then show that for all closed communication classes C reachable from (s, t) and for all pairs $(u, v) \in C$, it holds that $\lim_n \gamma_{P_n}(u, v) = \gamma_P(u, v) = 0$, by induction on the length of a longest path from C .

By Proposition 2, we have $\lim_n \delta_{\tau_n}(s, t) \leq \lim_n \gamma_{P_n}(s, t)$. Using the above result, we conclude that $\lim_n \delta_{\tau_n}(s, t) \leq \gamma_P(s, t) = 0$. \square

Towards an algorithm for computing \simeq , let us develop another characterization of robust bisimilarity. Given a policy $P \in \mathcal{P}$, we say that a set $R \subseteq S \times S$ *supports* a path $(u_1, v_1) \dots (u_n, v_n)$ in $\langle S \times S, P \rangle$ if for all $1 \leq i \leq n$ we have $(u_i, v_i) \in R$ and $\text{support}(P(u_i, v_i)) \subseteq R$.

Definition 6. A robust bisimulation is a bisimulation $R \subseteq S \times S$ such that for all $(s, t) \in R$, there exists a policy $P \in \mathcal{P}$ such that R supports a path from (s, t) to S_{Δ}^2 in $\langle S \times S, P \rangle$.

Proposition 3. Robust bisimilarity, \simeq is a robust bisimulation.

Proof Sketch. By Lemma 1, \simeq is a bisimulation. Let $P \in \mathcal{P}$ be the policy such that (s, t) reaches S_{Δ}^2 with probability 1 in $\langle S \times S, P \rangle$. Observe that for all (u, v) reachable from (s, t) , (u, v) must reach S_{Δ}^2 with probability 1 in $\langle S \times S, P \rangle$.

Consequently, $\text{support}(P(s, t)) \subseteq \simeq$. In fact, \simeq supports a path from (s, t) to S_{Δ}^2 in $\langle S \times S, P \rangle$, and we can conclude that \simeq is a robust bisimulation. \square

Proposition 4. *For any robust bisimulation $R \subseteq S \times S$, we have $R \subseteq \simeq$.*

Proof Sketch. We construct a policy $P \in \mathcal{P}$ such that for every $(s, t) \in R$, R supports a path from (s, t) to S_{Δ}^2 in $\langle S \times S, P \rangle$ and for all $(s, t) \in S_{\Delta}^2$, $\text{support}(P(s, t)) \subseteq S_{\Delta}^2$. Note that P is designed to simultaneously ensure that all pairs in R have an R -supported path to S_{Δ}^2 in $\langle S \times S, P \rangle$. It follows from a standard result in Markov chain theory that all $(s, t) \in R$ reach S_{Δ}^2 with probability 1 in $\langle S \times S, P \rangle$. \square

It follows from Propositions 3 and 4 that \simeq , that is, robust bisimilarity, is the greatest robust bisimulation. This is analogous to ordinary bisimulation, where bisimilarity is the greatest bisimulation.

5 Algorithm

In this section, we present an efficient algorithm to compute robust bisimilarity; see Algorithm 1. The algorithm relies on the following properties of any robust bisimulation R :

1. for every $(s, t) \in R$ there exists a policy P such that R supports a path from (s, t) to S_{Δ}^2 in $\langle S \times S, P \rangle$,
2. R is an equivalence relation, and
3. R is a bisimulation.

Robust bisimilarity is the greatest relation with these properties. More formally, we define a function, Refine, such that robust bisimilarity is the greatest fixed point of Refine.

Algorithm 1: Computing robust bisimilarity for labelled Markov chains

Input: A labelled Markov chain with a finite set S of states and a transition probability function $\tau : S \rightarrow \mathcal{D}(S)$, and the set of pairs of bisimilar states $\sim = S_{0, \tau}^2 \cup S_{\Delta}^2$

Output: The set of pairs of robustly bisimilar states $R = \simeq$

```

1  $R \leftarrow \sim$ 
2 repeat
3    $R_{\text{old}} \leftarrow R$ 
4    $R \leftarrow \text{Refine}(R)$  /* see Algorithm 2 */
5 until  $R = R_{\text{old}}$ 
6 return  $R$ 

```

For any L, U with $L \subseteq U \subseteq S \times S$, write $[L, U] = \{ R \subseteq S \times S \mid L \subseteq R \subseteq U \}$ and $[L, U]_{\mathcal{B}} = \{ R \in [L, U] \mid R \text{ is a bisimulation} \}$.

- The function $\text{Filter} : [S_{\Delta}^2, \sim]_{\mathcal{B}} \rightarrow [S_{\Delta}^2, \sim]$ is defined as
 $\text{Filter}(R) = \{ (s, t) \in R \mid \exists P \in \mathcal{P} \text{ such that } R \text{ supports a path from } (s, t) \text{ to } S_{\Delta}^2 \text{ in } \langle S \times S, P \rangle \}$.
- The function $\text{Prune} : [S_{\Delta}^2, \sim] \rightarrow [S_{\Delta}^2, \sim]$ is defined as
 $\text{Prune}(R) = \{ (s, t) \in R \mid \forall (t, u) \in R : (s, u) \in R \wedge \forall (u, s) \in R : (u, t) \in R \}$.
- The function $\text{Bisim} : [S_{\Delta}^2, \sim] \rightarrow [S_{\Delta}^2, \sim]_{\mathcal{B}}$ is defined as
 $\text{Bisim}(R)$ is the largest bisimulation R' with $R' \subseteq R$.
 Given an equivalence relation R , $\text{Bisim}(R)$ can be computed in polynomial time (see [15, Proposition 24]).
- Lastly, the function $\text{Refine} : [S_{\Delta}^2, \sim]_{\mathcal{B}} \rightarrow [S_{\Delta}^2, \sim]_{\mathcal{B}}$ is defined as
 $\text{Refine}(R) = \text{Bisim}(\text{Prune}(\text{Filter}(R)))$.

Proposition 5. *Bisim and Filter are monotone with respect to \subseteq . However, Prune is not.*

Proof Sketch. A counterexample for Prune is as follows. Let $S = \{s, t, u\}$, $A = \{(s, s), (t, t), (u, u), (s, t), (t, s)\}$ and $B = \{(s, s), (t, t), (u, u), (s, t), (t, s), (t, u), (u, t)\}$. A and B are symmetric and reflexive and, thus, can be visualized as an undirected graph as shown in Figure 4. Observe that $A \subseteq B$, however, $\text{Prune}(A) = A \not\subseteq \text{Prune}(B) = \{(s, s), (t, t), (u, u)\}$. \square

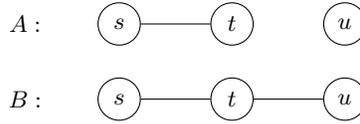


Fig. 4: Graph visualization of the relations A and B used in the proof of Proposition 5.

Note that Algorithm 1 is not a typical fixed point iteration, since we do not know whether Refine is monotone.

Algorithm 2: Refine

Input: A set $R \in [S_{\Delta}^2, \sim]_{\mathcal{B}}$
Output: $\text{Refine}(R)$

```

1  $R \leftarrow \text{Filter}(R)$  /* see Algorithm 3 */
2  $R \leftarrow \text{Prune}(R)$  /* see Algorithm 4 */
3  $R \leftarrow \text{Bisim}(R)$ 
4 return  $R$ 

```

Proposition 6. *Any relation $R \subseteq S \times S$ is a robust bisimulation if and only if it is a fixed point of Refine.*

Proof Sketch. Let $R \subseteq S \times S$. Assume that R is a robust bisimulation. By definition, $\text{Refine}(R) \subseteq R$. Since R is a robust bisimulation, $R \subseteq \text{Refine}(R)$.

Assume that R is a fixed point of Refine, then R is a bisimulation and for every $(s, t) \in R$ there exists a policy P such that R supports a path from (s, t) to S_{Δ}^2 in $\langle S \times S, P \rangle$. Therefore, R is a robust bisimulation. \square

It follows from Propositions 4 and 6 that every fixed point of Refine is a subset of \simeq . Furthermore, by Propositions 3 and 6, \simeq is a fixed point of Refine. Therefore, \simeq is the greatest fixed point of Refine.

Let $Q \subseteq S \times S$ and $s, t, u, v \in S$. We use the following notation below: $\text{Post}((s, t)) = \text{support}(\tau(s)) \times \text{support}(\tau(t))$ and $\text{Pre}(Q) = \{(s, t) \in S \times S \mid \text{Post}((s, t)) \cap Q \neq \emptyset\}$.

Algorithm 3: Filter

Input: A set $R \in [S_{\Delta}^2, \sim]_{\mathcal{B}}$
Output: Filter(R)

- 1 $Q \leftarrow S_{\Delta}^2$
- 2 $n \leftarrow 0$
- 3 **repeat**
- 4 $Q_{\text{old}} \leftarrow Q$
- 5 **foreach** $(s, t) \in (R \cap \text{Pre}(Q_{\text{old}})) \setminus Q_{\text{old}}$ **do**
- 6 $Q \leftarrow Q \cup \{(s, t)\}$
- 7 **end**
- 8 $n \leftarrow n + 1$
- 9 **until** $Q = Q_{\text{old}}$
- 10 **return** Q

Proposition 7. *Given $R \in [\simeq, \sim]_{\mathcal{B}}$, for all $(s, t), (t, u) \in \text{Filter}(R)$, if $s \simeq t$ or $t \simeq u$ then $(s, u) \in \text{Filter}(R)$.*

Proof Sketch. We show that if $t \simeq u$ then $(s, u) \in \text{Filter}(R)$. The case $s \simeq t$ is similar. Write $s_1 = s$ and $t_1 = t$ and $u_1 = u$.

The idea behind the proof is that since $\text{Filter}(R) \subseteq R$, we have $(s, t), (t, u) \in R$. Since R is an equivalence relation, $(s, u) \in R$. We define a policy $P \in \mathcal{P}$ such that for all $(s, t) \in R \cap S_{0?}^2$, $\text{support}(P(s, t)) = \text{Post}((s, t)) \cap R$. We then show that since $(s, t) \in \text{Filter}(R)$, there exists a path $(s_1, t_1), \dots, (s_n, t_n)$ in $\langle S \times S, P \rangle$, where $s_n = t_n$.

Assume that $(t, u) \in \simeq$. Recall that \simeq is a bisimulation. Since t_1, \dots, t_n is a path in the original Markov chain $\langle S, \tau \rangle$, there is also a path u_1, \dots, u_n in $\langle S, \tau \rangle$ such that $(t_i, u_i) \in \simeq$ for all $1 \leq i \leq n$. Since $\simeq \subseteq R$, there exists a path

Algorithm 4: Prune

Input: A set $Q \in [S_\Delta^2, \sim]$
Output: Prune(Q)

- 1 $E \leftarrow Q$
- 2 **foreach** $(s, t) \in Q$ **do**
- 3 **foreach** $u \in S : (t, u) \in Q$ **do**
- 4 **if** $(s, u) \notin Q$ **then**
- 5 $E \leftarrow E \setminus \{(s, t), (t, u)\}$
- 6 **end**
- 7 **end**
- 8 **end**
- 9 **return** E

$(t_1, u_1), \dots, (t_n, u_n)$ in $\langle S \times S, P \rangle$. Note that $(s_i, u_i) \in R$ for all $1 \leq i \leq n$. Hence, there exists a path $(s_1, u_1), \dots, (s_n, u_n) = (t_n, u_n)$ in $\langle S \times S, P \rangle$. See Figure 5.

Since $(t_n, u_n) \in \simeq$, we know that (t_n, u_n) reaches S_Δ^2 with probability 1. Therefore, there is a path $(t_n, u_n), \dots, (t_m, u_m)$, with $t_m = u_m$ in $\langle S \times S, P \rangle$ and $(t_i, u_i) \in \simeq$ for all $n \leq i \leq m$. Thus, there exists paths $(s_1, u_1), \dots, (s_n, u_n)$ and $(t_n, u_n), \dots, (t_m, u_m)$ in $\langle S \times S, P \rangle$, with $(s_n, u_n) = (t_n, u_n)$. By the definition of P , R supports the same path in $\langle S \times S, P \rangle$. Hence, $(s, u) \in \text{Filter}(R)$. \square

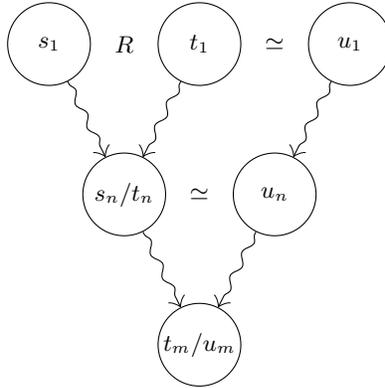


Fig. 5: Illustration of the proof of Proposition 7.

Proposition 7 allows us to prove the following proposition.

Proposition 8. $R \in [\simeq, \sim]_{\mathcal{B}}$ is a loop invariant of Algorithm 1.

Proof Sketch. R is initialized to \sim , so the loop invariant holds before the loop.

Assume that the loop invariant holds before an iteration of the loop. Since $\simeq \subseteq R$, Filter is monotone and \simeq is a fixed point of Refine, we have that \simeq is a subset of Filter(R).

If $s \simeq t$, then $(s, t) \in \text{Filter}(R)$. Then, by Proposition 7, for all $(t, u) \in \text{Filter}(R)$ we have $(s, u) \in \text{Filter}(R)$ and for all $(u, s) \in \text{Filter}(R)$ we have $(u, t) \in \text{Filter}(R)$. Hence, $(s, t) \in \text{Prune}(\text{Filter}(R))$, and we have that \simeq is a subset of Prune(Filter(R)).

Bisim is monotone, therefore, \simeq is a subset of Bisim(Prune(Filter(R))) and Refine(R). By the definition of Bisim, Refine(R) $\in [\simeq, \sim]_{\mathcal{B}}$. Thus, the loop invariant is maintained in each iteration of the loop. \square

Using the loop invariant established in Proposition 8, we can now prove the correctness of Algorithm 1.

Theorem 3. *Algorithm 1 computes the set \simeq .*

Proof Sketch. It is immediate from the definitions of Bisim, Filter and Prune that Refine(R) $\subseteq R$ holds for all $R \subseteq S \times S$. By Proposition 8, $\simeq \subseteq R$, thus, it computes a fixed point of Refine greater than or equal to \simeq . Since \simeq is the greatest fixed point of Refine, we can conclude that Algorithm 1 computes \simeq . \square

In [15, Proposition 27], we show that Algorithm 1 has a time complexity of $\mathcal{O}(n^6)$, where $n = |S|$. The computational bottleneck is the function Filter.

6 Experiments

To evaluate the efficiency and usefulness of our robust bisimilarity algorithm, we implemented it in the widely used probabilistic model checker PRISM [24], an open-source tool providing quantitative verification and analysis of several types of probabilistic models, including labelled Markov chains.

6.1 Implementation

PRISM's implementation of the traditional (i.e., non-robust) bisimilarity algorithm, Bisim, is a standard partition-refinement approach which uses the signature-based method of Derisavi [9]. The initial partition is based on the labelling of the states. Let Π be the current partition and E_{Π} be the set of equivalence classes in Π . Then the new partition is computed as $\{(s, t) \in \Pi \mid \forall B \in E_{\Pi} : \tau(s)(B) = \tau(t)(B)\}$.

We implemented Algorithm 1 in Java as part of PRISM's explicit-state model checking engine. Each state and equivalence class (referred to as a block) is represented by an integer ID. The current partition of the state space is tracked by an array that is indexed by state IDs and contains the corresponding block IDs. To store the list of successors for each state, we use a map. Bisim is run on the input Markov chain to obtain the set of bisimilar states.

The function Filter first constructs R from the current partition and initializes Q to S_{Δ}^2 . In our approach, R is implemented as an array indexed by block

IDs, with each block containing a list of states. Conversely, Q is implemented as an array indexed by state IDs, with each state storing the set of states related to it. Predecessors of Q in R are added to Q until a fixed point is reached. A pair of states $(s, t) \in R \setminus Q$ is a predecessor of Q if they have some successors that are related in Q . Specifically, there must exist successors s' and t' of s and t , respectively, such that $t' \in Q[s']$ and vice versa.

Prune constructs a new partition of the state space by grouping states in the same block if they have the same neighbourhood in Q , that is, they are related to the same states. In other words, s and t are placed in the same block if $Q[s] = Q[t]$ holds. Bisim is then called with the current partition passed as the initial partition. This process continues until no further refinement is possible, resulting in the set of robustly bisimilar states. Finally, the minimized Markov chain is constructed.

6.2 Experimental Setup

We evaluated our algorithm by applying it to all (discrete-time) labelled Markov chains from the Quantitative Verification Benchmark Set (QVBS) [17], a comprehensive collection of probabilistic models which is designed as a benchmark suite for quantitative verification and analysis tools and is the foundation of the Quantitative Verification Competition (QComp), which compares the performance, versatility, and usability of such tools.

For an additional source of models, we also use jpf-probabilistic [14]. Java PathFinder (JPF) [33] is the most popular model checker for Java code, and the JPF extension jpf-probabilistic provides Java implementations of sixty randomized algorithms [14]. As shown in [14], JPF, extended by jp-probabilistic and jpf-label, can be used in tandem with PRISM to check properties of these algorithms and supplement JPF’s qualitative results with quantitative information. A description of the subset of these algorithms utilized in our study is provided in [15, Appendix J].

In order to explore both the benefits and the efficiency of our algorithm, we run both the robust and traditional bisimilarity algorithms on all models. For the latter, we use PRISM’s existing implementation, in order to provide a comparable implementation. Our experiments were run on a MacBook with an M1 chip and 16GB memory, and with the Java virtual machine limited to 8GB.

6.3 Results

Table 1 shows results for all benchmarks where the minimized models obtained by traditional bisimilarity and robust bisimilarity differ. These are of particular interest because they are instances where our algorithm identifies that a model minimized in traditional fashion may not be robust. In fact, in all benchmarks we have checked, we have observed that the distance between pairs of states that are not robustly bisimilar is discontinuous. This leads us to the conjecture that for bisimilar states, robust bisimilarity is also a necessary condition for continuity. The property used for each benchmark dictates the labelling used for the model.

Table 1: The results of the benchmarks for which the minimized models differ.

Benchmark			Bisimilarity		Robust Bisimilarity		
Name (prop.)	Parameters	States	Min	Time	Min	Time	
brp (p1)	N=32	MAX=2	1349	646	0.036	901	0.054
		MAX=3	1766	871	0.043	1127	0.062
	N=64	MAX=4	2183	1096	0.051	1353	0.068
		MAX=5	2600	1321	0.058	1579	0.075
		MAX=2	2693	1286	0.080	1797	0.105
		MAX=3	3526	1735	0.084	2247	0.119
		MAX=4	4359	2184	0.103	2697	0.130
brp (p4)	N=32	MAX=5	5192	2633	0.132	3147	0.167
		MAX=2	1349	10	0.012	711	3.690
		MAX=3	1766	12	0.013	937	6.291
		MAX=4	2183	14	0.018	1163	9.331
	N=64	MAX=5	2600	16	0.021	1389	13.952
		MAX=2	2693	10	0.017	1415	27.299
		MAX=3	3526	12	0.015	1865	45.031
crowds (positive)	CS=5	MAX=4	4359	14	0.016	2315	69.949
		MAX=5	5192	16	0.018	2765	102.941
		TR=3	1198	41	0.018	505	0.231
		TR=4	3515	61	0.021	1484	1.304
	CS=10	TR=5	8653	81	0.038	3659	7.575
TR=6		18817	101	0.071	7969	34.765	
TR=3		6563	41	0.024	2320	8.296	
oscillators (power)	T=8	TR=4	30070	61	0.078	10524	196.233
		TR=5	111294	81	0.190	38770	2946.840
		N=3	57	28	0.007	38	0.009
oscillators (time)	T=6	N=6	1717	1254	0.037	1255	0.037
		N=8	6436	5148	0.100	5149	0.122
		N=3	57	28	0.007	38	0.008
set isolation (good sample)	U=13	N=6	1717	1254	0.032	1255	0.036
		N=8	6436	5148	0.111	5149	0.115
		ST=3	8196	19	0.029	27	21.885
		ST=4	8196	20	0.029	26	24.325
		ST=5	8196	21	0.032	25	24.330
		ST=6	8196	22	0.031	24	25.162

In the table, *Min* denotes the number of states in the minimized model and *Time* denotes the amount of time taken (in seconds) to compute bisimilarity.

The results are promising, since robust bisimilarity, although (unsurprisingly) slower than traditional bisimilarity, remains practical across a wide range of standard benchmarks. Table 2 displays some of the largest models per benchmark along with the time required to compute robust bisimilarity. The longest time recorded is about 50 minutes for the *crowds* benchmark. This may be due to

Table 2: Models with the maximum state space per benchmark.

Benchmark			Robust Bisimilarity	
Name	Property	States	Min States	Time
crowds	positive	111294	38770	2946.84
egl	messages	115710	131	153.01
herman	steps	32768	612	25.29
oscillators	power	24311	17877	0.42

Table 3: Summary of all benchmarks with the change due to robust bisimilarity.

Benchmark			Average % Increase	
Name	Property	Instances	States	Time
brp	p1	12	27.93	28.95
	p2	12	7.76	80.54
	p4	12	9193.43	142193.57
crowds	positive	7	12306.72	273258.24
egl	messages	6	-	20693.83
erdős-rényi model	connected	18	-	799.07
fair biased coin	heads	9	-	0.00
has majority element	incorrect	24	-	16.08
herman	steps	7	-	297.99
leader-sync	elected & time	18	-	518.98
haddad-monmege	target	3	-	0.00
oscillators	power & time	14	5.12	11.73
pollards factorization	input	8	-	0.00
queens	success	6	-	1193.93
set isolation	good sample	4	25.06	79035.95
Total		160	1231.68	25589.22

Table 4: Models for which robust bisimilarity results in an `OutOfMemoryError`.

Benchmark				Bisimilarity		
Name	Property	Parameters		States	Min States	Time
crowds	positive	CS=10	TR=6	352535	101	0.57
egl	messages	N=5	L=8	156670	171	0.87
		unfair	N=5	L=2	33790	229
	L=4			74750	469	0.26
	L=6		115710	709	0.40	
nand	reliable	N=20	L=8	156670	949	0.70
			K=1	78332	39982	0.81
			K=2	154942	102012	1.89
			K=3	231552	164042	3.67
queens	success		N=10	23492	527	0.08

the fact that the *crowds* benchmark has many non-robustly bisimilar pairs of states, which we believe makes the benchmark harder. Other benchmarks, e.g. *brp (p4)*, point in a similar direction.

The complete set of experiments includes 170 models, of which 160 are aggregated in Table 3. This table presents the average percentage increase in both the state space of the minimized model and the computation time for robust bisimilarity compared to traditional bisimilarity. The *crowds* benchmark exhibits the largest average percentage increase for both metrics. The reported values may seem large, however, it is important to note that the traditional bisimilarity algorithm required a maximum of 2.14 seconds per model in this table.

Furthermore, robust bisimilarity was successfully computed in less than a minute for 152 models (over 89%). Of the total set of models, the remaining 10 (approximately 6%), listed in Table 4, could only be minimized using traditional bisimilarity, as the robust bisimilarity computation ran out of available memory before completion. This issue occurred with all instances of the *nand* benchmark and half of the instances of the *egl* benchmark.

Ultimately, robust bisimilarity proves feasible for large models, despite needing more resources than traditional bisimilarity. Furthermore, it offers a more reliable method of determining equivalence, which can be particularly beneficial in mission-critical applications, which require a higher level of precision.

7 Conclusions and Future Work

To address the lack of robustness of probabilistic bisimilarity, we have introduced the concept of robust bisimilarity for labelled Markov chains. Robust bisimilarity ensures that the distance function remains continuous even under perturbations of transition probabilities. Additionally, we have presented a computationally efficient algorithm, with experimental results demonstrating that robust bisimilarity is plausible for large-scale verification tasks.

Our work opens new avenues for future exploration. First, a logical characterization of robust bisimilarity could provide deeper insights. Second, while we have established in Theorem 2 that robust bisimilarity is a sufficient condition for continuity, we conjecture that for bisimilar states, robust bisimilarity is in fact also a necessary condition for continuity. We also aim to define continuity for non-bisimilar state pairs, to complete the theoretical characterization of robustness. Thirdly, in [19] it was shown that when the distances are discounted (i.e., differences that manifest themselves later count less), the distance function becomes continuous. This raises the question: can we identify the properties for which the discontinuity is relevant? The examples suggest that these are long-term properties. Finally, we plan to investigate specific types of perturbations of the transition probabilities, such as those that do not introduce new transitions, preserving the graph structure, as seen in Figures 1a and 1c, unlike the perturbation shown in Figure 1b which adds a new transition.

Acknowledgments. This research was supported in part by the Clarendon Fund and by the Natural Sciences and Engineering Research Council of Canada.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. de Alfaro, L., Majumdar, R.: Quantitative solution of omega-regular games. In: Vitter, J.S., Spirakis, P.G., Yannakakis, M. (eds.) *Proceedings of the 33rd Annual Symposium on Theory of Computing*. pp. 675–683. ACM, Heraklion, Crete, Greece (Jul 2001)
2. Baier, C., Katoen, J.P.: *Principles of model checking*. The MIT Press, Cambridge, MA, USA (2008)
3. van Breugel, F., Worrell, J.: Towards quantitative verification of probabilistic transition systems. In: Orejas, F., Spirakis, P.G., van Leeuwen, J. (eds.) *Proceedings of the 28th International Colloquium on Automata, Languages and Programming*. *Lecture Notes in Computer Science*, vol. 2076, pp. 421–432. Springer, Crete, Greece (Jul 2001)
4. Cai, X., Gu, Y.: Measuring anonymity. In: Bao, F., Li, H., Wang, G. (eds.) *Proceedings of the 5th International Conference on Information Security Practice and Experience*. *Lecture Notes in Computer Science*, vol. 5451, pp. 183–194. Springer, Xi'an, China (Apr 2009)
5. Chatterjee, K., de Alfaro, L., Majumdar, R., Raman, V.: Algorithms for game metrics (full version). *Logical Methods in Computer Science* **6**(3) (2010)
6. Chen, D., van Breugel, F., Worrell, J.: On the complexity of computing probabilistic bisimilarity. In: Birkedal, L. (ed.) *Proceedings of the 15th International Conference on Foundations of Software Science and Computational Structures*. *Lecture Notes in Computer Science*, vol. 7213, pp. 437–451. Springer-Verlag, Tallinn, Estonia (Mar/Apr 2012)
7. Çinlar, E.: *Probability and stochastics*, Graduate Texts in Mathematics, vol. 261. Springer-Verlag, New York, NY, US (2011)
8. Comanici, G., Precup, D.: Basis function discovery using spectral clustering and bisimulation metrics. In: Burgard, W., Roth, D. (eds.) *Proceedings of the 25th AAAI Conference on Artificial Intelligence*. pp. 325–330. AAAI Press, San Francisco, California, USA (Aug 2011)
9. Derisavi, S.: Signature-based symbolic algorithm for optimal Markov chain lumping. In: *Proceedings of the 4th International Conference on the Quantitative Evaluation of Systems*. pp. 141–150. IEEE Computer Society, Edinburgh, Scotland, UK (Sep 2007)
10. Desharnais, J., Gupta, V., Jagadeesan, R., Panangaden, P.: Metrics for labeled Markov systems. In: Baeten, J.C.M., Mauw, S. (eds.) *Proceedings of the 10th International Conference on Concurrency Theory*. *Lecture Notes in Computer Science*, vol. 1664, pp. 258–273. Springer-Verlag, Eindhoven, The Netherlands (Aug 1999)
11. Desharnais, J., Gupta, V., Jagadeesan, R., Panangaden, P.: Metrics for labelled Markov processes. *Theoretical Computer Science* **318**(3), 323–354 (Jun 2004)
12. Desharnais, J., Laviolette, F., Tracol, M.: Approximate analysis of probabilistic processes: Logic, simulation and games. In: *Proceedings of the 5th International Conference on the Quantitative Evaluation of Systems*. pp. 264–273. IEEE Computer Society, Saint-Malo, France (Sep 2008)
13. Eastman, J.R., He, J.: A regression-based procedure for Markov transition probability estimation in land change modeling. *Land* **9**(11) (Oct 2020)
14. Fatmi, S.Z., Chen, X., Dhamija, Y., Wildes, M., Tang, Q., van Breugel, F.: Probabilistic model checking of randomized Java code. In: Laarman, A., Sokolova, A.

- (eds.) Proceedings of the 27th International Symposium on Model Checking Software, SPIN. Lecture Notes in Computer Science, vol. 12864, pp. 157–174. Springer (Jul 2021)
15. Fatmi, S.Z., Kiefer, S., Parker, D., van Breugel, F.: Robust probabilistic bisimilarity for labelled Markov chains. CoRR (May 2025). <https://doi.org/10.48550/arXiv.2505.15290>
 16. Giacalone, A., Jou, C., Smolka, S.A.: Algebraic reasoning for probabilistic concurrent systems. In: Broy, M., Jones, C.B. (eds.) Proceedings of the Working Conference on Programming Concepts and Methods. pp. 443–458. North-Holland, Sea of Galilee, Israel (Apr 1990)
 17. Hartmanns, A., Klauck, M., Parker, D., Quatmann, T., Ruijters, E.: The quantitative verification benchmark set. In: Vojnar, T., Zhang, L. (eds.) Proceedings of the 25th International Conference on Tools and Algorithms for the Construction and Analysis of Systems. Lecture Notes in Computer Science, vol. 11427, pp. 344–350. Springer, Prague, Czech Republic (Apr 2019)
 18. Hensel, C., Junges, S., Katoen, J., Quatmann, T., Volk, M.: The probabilistic model checker storm. International Journal on Software Tools for Technology Transfer **24**(4), 589–610 (2022)
 19. Jaeger, M., Mao, H., Larsen, K.G., Mardare, R.: Continuity properties of distances for Markov processes. In: Norman, G., Sanders, W.H. (eds.) Proceedings of the 11th International Conference on Quantitative Evaluation of Systems. Lecture Notes in Computer Science, vol. 8657, pp. 297–312. Springer-Verlag, Florence, Italy (Sep 2014)
 20. Jonsson, B., Larsen, K.: Specification and refinement of probabilistic processes. In: Proceedings of the 6th Annual Symposium on Logic in Computer Science. pp. 266–277. IEEE, Amsterdam, The Netherlands (Jul 1991)
 21. Katoen, J., Kemna, T., Zapreev, I.S., Jansen, D.N.: Bisimulation minimisation mostly speeds up probabilistic model checking. In: Proceedings of the 13th International Conference on Tools and Algorithms for the Construction and Analysis of Systems. Lecture Notes in Computer Science, vol. 4424, pp. 87–101. Springer (2007)
 22. Kemeny, J.G., Snell, J.L.: Finite Markov chains. Springer-Verlag, Heidelberg, Germany (1960)
 23. Kozen, D.: A probabilistic PDL. In: Johnson, D.S., Fagin, R., Fredman, M.L., Harel, D., Karp, R.M., Lynch, N.A., Papadimitriou, C.H., Rivest, R.L., Ruzzo, W.L., Seiferas, J.I. (eds.) Proceedings of the 15th Annual Symposium on Theory of Computing. pp. 291–297. ACM, Boston, Massachusetts, USA (Apr 1983)
 24. Kwiatkowska, M., Norman, G., Parker, D.: PRISM 4.0: Verification of probabilistic real-time systems. In: Gopalakrishnan, G., Qadeer, S. (eds.) Proceedings of the 23rd International Conference on Computer Aided Verification. Lecture Notes in Computer Science, vol. 6806, pp. 585–591. Springer-Verlag, Snowbird, Utah, USA (Jul 2011)
 25. Larsen, K., Skou, A.: Bisimulation through probabilistic testing. In: Proceedings of the 16th Annual ACM Symposium on Principles of Programming Languages. pp. 344–352. ACM, Austin, TX, USA (Jan 1989)
 26. McIver, A., Morgan, C.: Abstraction, Refinement and Proof for Probabilistic Systems. Monographs in Computer Science, Springer (2004)
 27. Mizutani, D., Lethanh, N., Adey, B.T., Kaito, K.: Improving the estimation of Markov transition probabilities using mechanistic-empirical models. Frontiers in Built Environment **3**, 58 (Oct 2017)

28. Olariu, E., Cadwell, K.K., Hancock, E., Trueman, D., Chevrou-Severac, H.: Current recommendations on the estimation of transition probabilities in Markov cohort models for use in health care decision-making: a targeted literature review. *ClinicoEconomics and Outcomes Research* **9**, 537–546 (Sep 2017)
29. Spitzer, F.: Principles of random walk. Graduate Texts in Mathematics, Springer-Verlag, New York, NY, USA (1964)
30. Srivastava, T., Latimer, N.R., Tappenden, P.: Estimation of transition probabilities for state-transition models: A review of NICE appraisals. *PharmacoEconomics* **39**(8), 869–878 (Aug 2021)
31. Tang, Q., van Breugel, F.: Algorithms to compute probabilistic bisimilarity distances for labelled Markov chains. In: Meyer, R., Nestmann, U. (eds.) Proceedings of the 28th International Conference on Concurrency Theory. LIPIcs, vol. 85, pp. 27:1–27:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, Berlin, Germany (Sep 2017)
32. Thorsley, D., Klavins, E.: Approximating stochastic biochemical processes with Wasserstein pseudometrics. *IET Systems Biology* **4**, 193–211 (2010)
33. Visser, W., Havelund, K., Brat, G., Park, S., Lerda, F.: Model checking programs. *Automated Software Engineering* **10**(2), 203–232 (Apr 2003)