

Scalable Methods for Computing State Similarity in Deterministic Markov Decision Processes

Pablo Samuel Castro

Google Brain
psc@google.com

Abstract

We present new algorithms for computing and approximating bisimulation metrics in Markov Decision Processes (MDPs). Bisimulation metrics are an elegant formalism that capture behavioral equivalence between states and provide strong theoretical guarantees on differences in optimal behaviour. Unfortunately, their computation is expensive and requires a tabular representation of the states, which has thus far rendered them impractical for large problems. In this paper we present a new version of the metric that is tied to a behavior policy in an MDP, along with an analysis of its theoretical properties. We then present two new algorithms for approximating bisimulation metrics in large, deterministic MDPs. The first does so via sampling and is guaranteed to converge to the true metric. The second is a differentiable loss which allows us to learn an approximation even for continuous state MDPs, which prior to this work had not been possible.

Introduction

A **finite Markov Decision Process (MDP)** is defined as a 5-tuple $\mathcal{M} = \langle \mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \gamma \rangle$, where \mathcal{S} is a finite set of states, \mathcal{A} is a finite set of actions, $\mathcal{P} : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$ is the next state transition function (where $\Delta(X)$ is the probability simplex over the set X), $\mathcal{R} : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is the reward function (assumed to be bounded by R_{max}), and $\gamma \in [0, 1]$ is a discount factor. An MDP is the standard formalism for expressing sequential decision problems, typically in the context of planning or reinforcement learning (RL). The set of states \mathcal{S} is one of the central components of this formalism, where each state $s \in \mathcal{S}$ is meant to encode sufficient information about the environment such that an agent can learn how to behave in a *consistent* manner. Figure 1 illustrates a simple MDP where each cell represents a state.

There is no canonical way of defining the set of states for a problem. Indeed, improperly designed state spaces can have drastic effects on the algorithm used. Consider the grid MDP in the bottom of Figure 1, where an agent must learn how to navigate to the green cells, and imagine we create an exact replica of the MDP such that the agent randomly transitions between the two layers for each move. By doing so we have




Figure 1: A grid MDP (bottom) with a copy of itself (top). The goal of an agent is to find the shortest path to the green cells. At each iteration, the agent has a 50% chance of jumping to the other level.

doubled the number of states and the complexity of the problem. However, from a planning perspective the two copies of each state should be indistinguishable. A stronger notion of *state identity* is needed that goes beyond the labeling of states and which is able to capture *behavioral indistinguishability*.

In this paper we explore notions of behavioral similarity via state pseudometrics¹ $d : \mathcal{S} \times \mathcal{S} \rightarrow \mathbb{R}$, and in particular those which assign a distance of 0 to states that are behaviorally indistinguishable. Pseudometrics further allow us to reason about states based on what we may know about other similar states. This is a common use-case in fields such as formal verification, concurrency theory, and in safe RL, where one may want to provide (non-)reachability guarantees. In the context of planning and reinforcement learning, these can be useful for state aggregation and abstraction.

Our work builds on bisimulation metrics (Ferns, Panangaden, and Precup 2004) which provide us with theoretical guarantees such as states that are close to each other (with respect to the metric) will have similar optimal value functions. These theoretical properties render them appealing for planning and learning, and they have previously been used for state aggregation (Ferns, Panangaden, and Precup 2004; Ferns et al. 2006), policy transfer (Castro and Precup 2010), representation discovery (Ruan et al. 2015), and exploration (Santara et al. 2019).

Unfortunately, these metrics are expensive to compute (Chen, van Breugel, and Worrell 2012) and require fully enu-

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹A pseudometric is a metric d where $\forall s, t \in \mathcal{S}. s = t \implies d(s, t) = 0$, but not the converse.

merating the states, even when using sampling-based approximants (Ferns et al. 2006), on-the-fly approximants (Comanici, Panangaden, and Precup 2012; Bacci et al. 2013b), or approximants which exploit structure in the state space (Bacci et al. 2013a). The full-state enumeration requirement has thus far rendered bisimulation metrics impractical in problems with large state spaces, and completely incompatible with continuous state spaces. Additionally, bisimulation metrics can be overly *pessimistic* in the sense that they consider worst-case differences between states. Although desirable for certain applications, such as in guaranteeing safe behaviors, it can prove overly restrictive for many practical problems of interest.

In this paper we address these impracticalities with the following key contributions:

1. An *on-policy* variant of bisimulation which focuses only on the behavior of interest, rather than worst-case scenarios, along with an analysis of its theoretical properties.
2. A new sampling-based online algorithm for exact computation of the original and on-policy bisimulation metrics with guaranteed convergence.
3. A differentiable loss function for learning an approximation of the two bisimulation metrics using neural networks. We provide empirical evidence of this learning algorithm on MDPs with large and continuous state spaces. To the best of the author’s knowledge, this is the first work proposing a mechanism for approximating bisimulation metrics with neural networks.

Background

Given an MDP \mathcal{M} , a policy $\pi : \mathcal{S} \rightarrow \Delta(\mathcal{A})$ induces a corresponding state-value function $V^\pi : \mathcal{S} \rightarrow \mathbb{R}$ (Puterman 1994): $V^\pi(s) = \mathbb{E}_{a \sim \pi(s)} [\mathcal{R}(s, a) + \gamma \mathbb{E}_{s' \sim \mathcal{P}(s, a)} V^\pi(s')]$. In the control setting, we are typically in search of the optimal value function (Bellman 1957):

$$V^*(s) = \max_{a \in \mathcal{A}} [\mathcal{R}(s, a) + \gamma \mathbb{E}_{s' \sim \mathcal{P}(s, a)} V^*(s')].$$

Bisimulation relations, originally introduced in the field of concurrency theory, were adapted for MDPs by Givan, Dean, and Greig (2003), capture a strong form of behavioral equivalence: if $s, t \in \mathcal{S}$ are bisimilar, then $V^*(s) = V^*(t)$.

Definition 1. *Given an MDP \mathcal{M} , an equivalence relation $E \subseteq \mathcal{S} \times \mathcal{S}$ is a **bisimulation relation** if whenever $(s, t) \in E$ the following properties hold, where \mathcal{S}_E is the state space \mathcal{S} partitioned into equivalence classes defined by E :*

1. $\forall a \in \mathcal{A}, \mathcal{R}(s, a) = \mathcal{R}(t, a)$
2. $\forall a \in \mathcal{A}, \forall c \in \mathcal{S}_E, \mathcal{P}(s, a)(c) = \mathcal{P}(t, a)(c)$, where $\mathcal{P}(s, a)(c) = \sum_{s' \in c} \mathcal{P}(s, a)(s')$,

Two states $s, t \in \mathcal{S}$ are **bisimilar** if there exists a bisimulation relation E such that $(s, t) \in E$. We denote the largest² bisimulation relation as \sim .

Equivalence relations can be brittle: they require exact equivalence under probabilistic transitions. This can be espe-

²Note that there can be a number of equivalence relations satisfying these properties. The smallest is the identity relation, which is vacuously a bisimulation relation.

cially problematic if we are estimating transition probabilities from data, as it is highly unlikely they will match exactly.

Extending the work of Desharnais et al. (1999) for labeled Markov processes, Ferns, Panangaden, and Precup (2004) generalized the notion of MDP bisimulation relations to metrics, yielding a smoother notion of similarity than equivalence relations. Let \mathbb{M} be the set of all pseudometrics on \mathcal{S} . A pseudometric $d \in \mathbb{M}$ induces an equivalence relation $E_d := \{(s, t) | d(s, t) = 0\}$. That is, any two states with distance 0 will be collapsed onto the same equivalence class.

Definition 2. (Ferns, Panangaden, and Precup 2004) *A pseudometric $d \in \mathbb{M}$ is a **bisimulation metric** if E_d is \sim .*

Bisimulation metrics use the 1-Wasserstein metric $\mathcal{W}_1 : \mathbb{M} \rightarrow \mathbb{P}$, where \mathbb{P} is the set of all metrics between probability distributions over \mathcal{S} . Given two state distributions $X, Y \in \Delta(\mathcal{S})$ and a pseudometric $d \in \mathbb{M}$, the Wasserstein $\mathcal{W}_1(d)(X, Y)$ can be expressed by the following (primal) linear program (LP), which “lifts” a pseudometric $d \in \mathbb{M}$ onto one in \mathbb{P} (Villani 2008):

$$\begin{aligned} \max_{\mathbf{u}} \quad & \mathbb{R}^{|\mathcal{S}|} \sum_{s \in \mathcal{S}} (X(s) - Y(s)) u_s & (1) \\ \forall s, s' \in \mathcal{S}, \quad & u_s - u_{s'} \leq d(s, s') \\ & 0 \leq \mathbf{u} \leq 1 \end{aligned}$$

Theorem 1. (Ferns, Panangaden, and Precup 2004): *Define $\mathcal{F} : \mathbb{M} \rightarrow \mathbb{M}$ by*

$$\begin{aligned} \mathcal{F}(d)(s, t) = & & (2) \\ \max_{a \in \mathcal{A}} & (|\mathcal{R}(s, a) - \mathcal{R}(t, a)| + \gamma \mathcal{W}_1(d)(\mathcal{P}(s, a), \mathcal{P}(t, a))) \end{aligned}$$

then \mathcal{F} has a unique fixed point, d_\sim , and d_\sim is a bisimulation metric.

The operator \mathcal{F} can be used to iteratively compute a bisimulation metric as follows. Starting from an initial estimate d_0 , we can compute $d_{n+1} = \mathcal{F}(d_n) = \mathcal{F}^{n+1}(d_0)$. By iteratively applying \mathcal{F} $\lceil \frac{\ln \delta}{\ln \gamma} \rceil$ times, one can compute d_\sim up to an accuracy δ , with an overall complexity of $O(|\mathcal{A}| |\mathcal{S}|^4 \log |\mathcal{S}| \frac{\ln \delta}{\ln \gamma})$.

On-policy bisimulation

The strong theoretical guarantees of bisimulation relations and metrics are largely due to their inherent “pessimism”: they consider equivalence under all actions, even pathologically bad ones (i.e. actions that never lead to positive outcomes for the agent). Indeed, there exist systems where $V^*(s) = V^*(t)$, but $d_\sim(s, t)$ can be arbitrarily large, providing no useful bounds on the optimal behaviour from s and t (see Figure 2). Castro and Precup (2010) also demonstrated that this pessimism yields poor results when using bisimulation metrics for policy transfer.

Another disadvantage of bisimulation relations and metrics is that they are computed via exact action matching between states; however, actions with the same label may induce very different behaviours from different states, resulting in an improper behavioral comparison when using bisimulation. In the system in Figure 2, s and t have equal optimal values,




Figure 2: Edge labels indicate action ($\{a, b\}$) and non-zero rewards ($[K]$). When $\gamma = 0.9$, $V^*(s) = V^*(t) = 10K$, while $d_{\sim}(s, t) = 10K$. Lax bisimulation assigns distance 0 between s and t (action a from s would be matched with action b from t).

but their optimal action is different (a from s , b from t). Taylor, Precup, and Panagaden (2009) overcame this problem by the introduction of lax-bisimulation metrics (definition and theoretical results provided in the supplemental). We note, however, that their method is still susceptible to the pessimism discussed above.

It is often the case that one is interested in behaviours relative to a particular policy π . In reinforcement learning, for example, many algorithms maintain a behaviour policy which is improved iteratively as the agent interacts with the environment. In these situations the pessimism of bisimulation can become a hindrance: if the action maximizing the distance between two states is never chosen by π , we should not include it in the computation!

We introduce a new notion of bisimulation, *on-policy bisimulation*, defined relative to a policy π . This new notion also removes the requirement of matching on action labels by considering the dynamics induced by π , rather than the dynamics induced by each action. We first define:

$$\mathcal{R}_s^\pi := \sum_a \pi(a|s) \mathcal{R}(s, a)$$

$$\forall C \in \mathcal{S}_{E^\pi}, \mathcal{P}_s^\pi(C) := \sum_a \pi(a|s) \sum_{s' \in C} P(s, a)(s')$$

Definition 3. Given an MDP \mathcal{M} , an equivalence relation $E^\pi \subseteq \mathcal{S} \times \mathcal{S}$ is a π -bisimulation relation if whenever $(s, t) \in E^\pi$ the following properties hold:

1. $\mathcal{R}_s^\pi = \mathcal{R}_t^\pi$
2. $\forall C \in \mathcal{S}_{E^\pi}, \mathcal{P}_s^\pi(C) = \mathcal{P}_t^\pi(C)$

Two states $s, t \in \mathcal{S}$ are π -bisimilar if there exists a π -bisimulation relation E^π such that $(s, t) \in E^\pi$. Denoting the largest bisimulation relation as \sim_π , $d \in \mathbb{M}$ is a π -bisimulation metric if E_d is \sim_π .

Theorem 2. Define $\mathcal{F}^\pi : \mathcal{M} \rightarrow \mathcal{M}$ by $\mathcal{F}^\pi(d)(s, t) = |\mathcal{R}_s^\pi - \mathcal{R}_t^\pi| + \gamma \mathcal{W}_1(d)(\mathcal{P}_s^\pi, \mathcal{P}_t^\pi)$, then \mathcal{F}^π has a least fixed point d_{\sim}^π , and d_{\sim}^π is a π -bisimulation metric.

Proof. (Sketch) This proof mimics the proof of Theorem 4.5 from (Ferns, Panangaden, and Precup 2004). All complete proofs are provided in the supplemental material. \square

The following result demonstrates that π -bisimulation metrics provide similar theoretical guarantees as regular bisimulation metrics, but with respect to the value function induced by π .

Theorem 3. Given any two states $s, t \in \mathcal{S}$ in an MDP \mathcal{M} , $|V^\pi(s) - V^\pi(t)| \leq d_{\sim}^\pi(s, t)$.

Proof. (Sketch) This is proved by induction. The result follows by expanding V^π , taking the absolute value of each term separately, and noticing that V^π is a feasible solution to the primal LP in Equation 1, so is upper-bounded by $\mathcal{W}_1(d_{\sim}^\pi)$. \square

Under a fixed policy π , an MDP reduces to a Markov chain. Bisimulation relations for Markov chains have previously been studied in concurrency theory (Baier et al. 2006; Katoen et al. 2007). Further, π -bisimulation can be used to define a notion of *weak*-bisimulation for MDPs (Baier et al. 2006; Ferrer Fioriti et al. 2016).

Bisimulation metrics for deterministic MDPs

In this section we investigate the properties of deterministic MDPs, which in concurrency theory are known as transition systems (Sangiorgi 2011).

Definition 4. A deterministic MDP \mathcal{M} is one where for all $s \in \mathcal{S}, a \in \mathcal{A}$, there exists a unique $\mathcal{N}(s, a) \in \mathcal{S}$ such that $\mathcal{P}(s, a)(\mathcal{N}(s, a)) = 1$.

As the next lemma shows, under a system with deterministic transitions, computing the Wasserstein metric (and approximants) is no longer necessary.

Lemma 1. Given a deterministic MDP \mathcal{M} , for any two states $s, t \in \mathcal{S}$, action $a \in \mathcal{A}$, and pseudometric $d \in \mathbb{M}$, $\mathcal{W}_1(d)(\mathcal{P}(s, a), \mathcal{P}(t, a)) = d(\mathcal{N}(s, a), \mathcal{N}(t, a))$.

Proof. (Sketch) The result follows by considering the dual formulation of the primal LP in Equation 1, which implies the dual variables $\lambda_{s,t}$ must all be either 1 or 0, by virtue of determinism. \square

By considering only deterministic policies (e.g. policies that assign probability 1 to a single action) in the on-policy case, Lemma 1 allows us to rewrite the operator $\mathcal{F}(d)(s, t)$ in Theorem 1 and $\mathcal{F}^\pi(d)(s, t)$ in Theorem 2 as:

$$\max_{a \in \mathcal{A}} (|\mathcal{R}(s, a) - \mathcal{R}(t, a)| + \gamma d(\mathcal{N}(s, a), \mathcal{N}(t, a)))$$

and

$$|\mathcal{R}(s, \pi(s)) - \mathcal{R}(t, \pi(t))| + \gamma d(\mathcal{N}(s, \pi(s)), \mathcal{N}(t, \pi(t))),$$

respectively. Note the close resemblance to value functions, there is in fact a strong connection between the two: Ferns and Precup (2014) proved that d_{\sim} is the optimal value function of an optimal coupling of two copies of the original MDP.

Even in the deterministic setting, the computation of bisimulation metrics can be intractable in MDPs with very large or continuous state spaces. In the next sections we will leverage the results just presented to introduce new algorithms that are able to scale to large state spaces and learn an approximant for continuous state spaces.

Computing bisimulation metrics with sampled transitions

We present the algorithm and results in this section for the original bisimulation metric, d_\sim , but all the results presented here hold for the on-policy variant d_\sim^π ; the main difference is that actions in the trajectory are given by π and thus, may differ between states being compared.

The update operator \mathcal{F} is generally applied in a dynamic-programming fashion: all state-pairs are updated in each iteration by considering all possible actions. However, requiring access to all state-pairs and actions in each iteration is often not possible, especially when data is concurrently being collected by an agent interacting with an environment. In this section we present an algorithm for computing the bisimulation metric via access to *transition samples*. Specifically, assume we are able to sample pairs of transitions $\{(s, a, \mathcal{R}(s, a), \mathcal{N}(s, a)), (t, a, \mathcal{R}(t, a), \mathcal{N}(t, a))\}$ from an underlying distribution \mathcal{D} (note the action is the same for both). This can be, for instance, a uniform distribution over all transitions in a replay memory (Mnih et al. 2015) or some other sampling procedure. Let \mathcal{T} be the set of all pairs of valid transitions; for legibility we will use the shorthand $\tau_{s,t,a} \in \mathcal{T}$ to denote a pair of transitions from states $s, t \in \mathcal{S}$ under action $a \in \mathcal{A}$. We assume that $\mathcal{D}(\tau) > 0$ for all $\tau \in \mathcal{T}$.

We first define an iterative procedure for computing d_\sim by sampling from \mathcal{D} . Let $d_0 \equiv 0$ be the everywhere-zero metric. At step n , let $\tau_{s_n, t_n, a_n} \in \mathcal{T}$ be a sample from \mathcal{D} and define d_n as:

$$d_n(s, t) = d_{n-1}(s, t), \quad \forall s \neq s_n, t \neq t_n$$

$$d_n(s_n, t_n) = \max \left[\begin{array}{c} d_{n-1}(s_n, t_n), \\ |\mathcal{R}(s_n, a_n) - \mathcal{R}(t_n, a_n)| + \\ \gamma d_{n-1}(\mathcal{N}(s_n, a_n), \mathcal{N}(t_n, a_n)) \end{array} \right] \quad (3)$$

In words, when we sample a pair of states, we only update the distance estimate for these two states if applying the \mathcal{F} operator gives us a larger estimate. Otherwise, our estimate remains unchanged.

Theorem 4. *If d_n is updated as in Equation 3 and $d_0 \equiv 0$, $\lim_{n \rightarrow \infty} d_n = d_\sim$ almost surely.*

Proof. (Sketch) We first show that since we are sampling state pairs and actions infinitely often, all state pairs will receive a non-vacuous update at least once (Maximizing action lemma); then show that $d_n \leq d_\sim$ for all $n \in \mathbb{N}$ (Monotonicity lemma). We then use these two results to show that the difference $\|d_\sim - d_n\|_\infty$ is a contraction and the result follows by the Banach fixed-point theorem. Note that the maximizing action lemma as presented here is for the original bisimulation metric; for the on-policy variant, the equivalent result is that all states in the Markov chain induced by π are updated infinitely often. \square

Learning an approximation

We leverage the sampling approach from the previous section to devise a learning algorithm for approximating d_\sim and

d_\sim^π for MDPs with large (or continuous) state spaces, using function approximators in the form of neural networks. Let $\phi : \mathcal{S} \rightarrow \mathbb{R}^k$ be a k -dimensional representation of the state space and let $\psi_\theta : \mathbb{R}^{2k} \rightarrow \mathbb{R}$ be a neural network parameterized by θ , that receives a concatenation of two state representations such that $\psi_\theta([\phi(s), \phi(t)]) \approx d_\sim(s, t)$ (see Figure 3).

Following the practice introduced by (Mnih et al. 2015) we make use of online parameters θ and target parameters θ^- , where the online parameters are updated at each iteration while the target parameters are updated every C iterations. Given a pair of states $s \neq t$ and action $a \in \mathcal{A}$, at iteration i we define the target objective $\mathbf{T}_{\theta_i^-}(s, t, a)$ for d_\sim as:

$$\max \left[\begin{array}{c} |\mathcal{R}(s, a) - \mathcal{R}(t, a)| + \gamma \psi_{\theta_i^-}([\phi(\mathcal{N}(s, a)), \phi(\mathcal{N}(t, a))]), \\ \psi_{\theta_i^-}([\phi(s), \phi(t)]) \end{array} \right]$$

and equal to 0 whenever $s = t$. The target objective $\mathbf{T}_{\theta_i^\pi}^\pi(s, t)$ for d_\sim^π is:

$$|\mathcal{R}(s, \pi(s)) - \mathcal{R}(t, \pi(t))| + \gamma \psi_{\theta_i^\pi}([\phi(\mathcal{N}(s, \pi(s))), \phi(\mathcal{N}(t, \pi(t)))])$$

We can then define our loss $\mathcal{L}_{s,t,a}^{(\pi)}$ as: $\mathbb{E}_{\mathcal{D}} \left(\mathbf{T}_{\theta_i^\pi}^\pi(s, t, a) - \psi_{\theta_i^\pi}([\phi(s), \phi(t)]) \right)^2$. This loss is specified for a single pair of transitions, but we can define an analogous target and loss with mini-batches, which allows us to train our approximant more efficiently using specialized hardware such as GPUs:

$$\mathbf{T} = (1 - \mathbf{I}) \otimes \max \left(\mathbf{R}^2 + \gamma \beta \psi_{\theta_i^-}(\mathbf{N}^2), \beta \psi_{\theta_i^-}(\mathbf{S}^2) \right)$$

$$\mathcal{L}_i(\theta_i) = \mathbb{E}_{\mathcal{D}} \left[\mathbf{W} \otimes (\psi_{\theta_i}(\mathbf{S}^2) - \mathbf{T})^2 \right]$$

$$\mathbf{T}^\pi = (1 - \mathbf{I}) \left(\mathbf{R}^2 + \gamma \beta \psi_{\theta_i^\pi}(\mathbf{N}^2) \right)$$

$$\mathcal{L}_i^\pi(\theta_i) = \mathbb{E}_{\mathcal{D}} \left[(\psi_{\theta_i^\pi}(\mathbf{S}^2) - \mathbf{T}^\pi)^2 \right]$$

\mathbf{R}^2 , \mathbf{S}^2 , and \mathbf{N}^2 are batches of rewards, states, and next-states, respectively, \mathbf{W} is a mask used to enforce action matching when approximating d_\sim , and \mathbf{I} is the identity matrix. $\psi(\mathbf{X})$ indicates applying ψ to a matrix \mathbf{X} elementwise, and \otimes stands for the Hadamard product. We multiply by $(1 - \mathbf{I})$ to zero out the diagonals, since those represent approximations to $d_\sim^\pi(s, s) \equiv 0$. The parameter β is a stability




Figure 3: Using a neural network for learning ψ as an approximant to d_\sim or d_\sim^π .




Figure 4: GridWorld and its (x, y) representation. Black densities illustrate the sampled states when adding noise.

parameter that begins at 0 and is incremented towards 1 every C iterations. Its purpose is to gradually “grow” the effective horizon of the bisimulation backup and maximization. This is helpful since the approximant ψ_θ can have some variance initially, depending on how θ is initialized. Further, Jiang et al. (2015) demonstrate that using shorter horizons during planning can often be better than using the true horizon, especially when using a model estimated from data. Note that, in general, the approximant ψ is not guaranteed to be a proper pseudometric. A lengthier discussion, including the derivation of these matrices, is provided in the supplemental material.

Empirical evaluation

In this section we provide empirical evidence for the effectiveness of our bisimulation approximants³. We begin with a simple 31-state GridWorld, on which we can compute the bisimulation metric exactly, and use a “noisy” representation which yields a continuous-state MDP. Having the exact metric for the 31-state MDP allows us to quantitatively measure the quality of our learned approximant.

We then learn a π -bisimulation approximant over policies generated by reinforcement learning agents trained on Atari 2600 games. In the supplemental material we provide an extensive discussion of the hyperparameter search we performed to find the settings used for both experiments. Training was done on a Tesla P100 GPU.

GridWorld

We first evaluate our learning algorithms on the 31-state GridWorld environment illustrated in Figure 4. There are 4 actions (up, down, left, right) with deterministic transitions, and where an action driving the agent towards a wall keeps the agent in the same cell. There is a single reward of +1.0 received upon entering either of the green cells, and a reward of -1.0 for taking an action towards a wall. We display the

³Code available at https://github.com/google-research/google-research/tree/master/bisimulation_aai2020

bisimulation distances between all states in the supplemental, computed using the sampling approach.

We represent each state by its coordinates $(x, y) \in \mathbb{R}^2$, as illustrated in Figure 4. To estimate d_\sim we use a network with an input layer of dimension 4, one fully connected hidden layer with 729 hidden units, and an output of length 1. The input is the concatenation of two state representations, normalized to be in $[-1, 1]$, while the output value is the estimate to d_\sim . We sampled state pairs and actions uniformly randomly, and ran our experiments with $\gamma = 0.99$, $C = 500$, $b = 256$, and increased β from 0 to 1 by a factor of 0.9 every time the target network was updated; we used the Adam optimizer (Kingma and Ba 2015) with a learning rate of 0.01. Because of the maximization term in the target, these networks can have a tendency to overshoot (although the combination of target networks and the β term helps stabilize this); we ran the training process for 2500 steps, which, for this problem, was long enough to converge to a reasonable solution before overshooting. The full hyperparameter settings are provided in the supplemental.

To evaluate the learning process, we measure the absolute error: $\|d_\sim - \psi\|_\infty$ using the true underlying state space for which we know the value of d_\sim . Note that because there is no fixed learning target, absolute errors are not guaranteed to be bounded. For this reason we also report the normalized error: $\|\frac{d_\sim}{\|d_\sim\|_2} - \frac{\psi}{\|\psi\|_2}\|_\infty$ as in practice one is mostly interested in relative, rather than absolute, distances. Figure 5a and Figure 5b display the results of our experiments over 10 independent runs; the shaded areas represent the 95% confidence interval.

In addition to training on the 31 state-MDP, we constructed a continuous variant by adding Gaussian noise to the state representations; this noise is centered at $(0, 0)$ with standard deviation 0.1, and clipped to be in $[-0.3, 0.3]$. The per-cell noise is illustrated by the black gradients in Figure 4. As Figure 5 shows, there is little difference between learning the metric for the 31-state MDP versus learning it for its continuous variant. Adding noise does not seem to hurt performance, and in fact seems to be helpful. We hypothesize that noise may be acting as a form of regularization, but this requires further investigation. In the supplemental material we include a figure exploring using the metric approximant for aggregating states in the continuous MDP setting with promising results.

Atari 2600

To evaluate the performance of our learning algorithm on an existing large problem, we take a set of reinforcement learning (RL) agents trained on three Atari 2600 games from the Arcade Learning Environment (Bellemare et al. 2013). The RL agents were obtained from the set of trained agents provided with the Dopamine library (Castro et al. 2018). Because our methods are designed for deterministic MDPs, we only used those trained without sticky actions⁴ (Machado et al. 2018) (evaluated in Section 4.3 in (Castro et al. 2018)); the trained checkpoints were provided for only three games: Space Invaders, Pong, and Asterix. We used the Rainbow

⁴Sticky actions add stochasticity to action outcomes.




Figure 5: Metric errors for the learned metric on the GridWorld as training progresses.

agent (Hessel et al. 2018) as it is the best performing of the provided Dopamine agents. We used the penultimate layer of the trained Rainbow agent as the representation ϕ .

To approximate the on-policy bisimulation metric d_{\sim}^{π} we loaded a trained agent and ran it in evaluation mode for each respective game, filling up the replay buffer while doing so. Once 10,000 transitions have been stored in the replay buffer, we begin sampling mini-batches and update our approximant ψ_{θ}^{π} using the target and loss defined previously. (note that we still continue populating our replay buffer). We ran our experiments using a network of two hidden layers of dimension 16, with $\gamma = 0.99$, $C = 500$, $b = 128$, and increased the β term from 0 to 1 by a factor of 0.99 every time the target network was updated. We used the Adam optimizer (Kingma and Ba 2015) with a learning rate of $7.5e^{-5}$ (except for Pong where we found 0.001 yielded better results). We trained the networks for around 600K steps, although in practice we found that about half that many steps were needed to reach a stable approximant. The configuration file specifying the full hyperparameter settings as well as the learning curves are provided in the supplemental.

After training we evaluated our approximant $\psi_{\theta}^{\pi}([\phi(s), \phi(t)])$ by fixing s to be the first state in the game and varying t throughout the episode; that is, we evaluate how similar the other states of the game are to the initial state w.r.t. our metric approximant. In Figure 6 we display the first 500 steps of one evaluation run on Space Invaders; as can be seen, the learned metric captures more meaningful differences between frames (start of episodes, enemy alien destroyed) that go beyond simple pixel differences. Interestingly, when sorting the frames by distance, the frames furthest away from s are typically those where the agent is about to be killed. It is worth noting that the way states are encoded in Dopamine is by stacking the last four frames; in our visualization we are only displaying the top frame in this stack. We observed similar results for Asterix and Pong; we include these and more extensive results, as well as videos for the three games, in the supplemental material.

Related work

There are a number of different notions of state similarity that have been studied over the years. Li, Walsh, and Littman (2006) provide a unified characterization and analysis of many of them. MDP-homomorphisms (Ravindran and Barto 2003) do not require behavioral equivalence under the same action labels, and this idea was extended to a metric by Taylor, Precup, and Panangaden (2009).

Ferns et al. (2006) introduced a sampling-based approximation to d_{\sim} which exchanges the computation of the Wasserstein with an instance of the assignment problem from network optimization; although Castro (2011) derived a PAC-bound for this approximant, the number of samples required is still prohibitive. Bacci et al. (2013a) exploit the underlying structure in \mathcal{S} to compute d_{\sim} .

Although there has been some work in concurrency theory to approximate large systems via ‘polynomially accurate’ simulations (Segala and Turrini 2007), they make no use of function approximators in the form of neural networks. We believe our use of neural networks may grant our approach greater generalizability.

Deterministic on-policy systems can be reduced to a graph. As such, our notion of π -bisimulation metrics bears a close relationship to graph similarity measures (Zager and Verghese 2008). However, graph similarity notions compare two full systems (graphs), as opposed to two nodes within a single graph, as we evaluate here. Nonetheless, the relationship warrants further investigation, which we leave for future work.

Perhaps most related to our sampling method is the on-the-fly methods introduced by Comanici, Panangaden, and Precup (2012). The authors replace the use of standard dynamic programming in their computation with something akin to asynchronous dynamic programming (Sutton and Barto 1998), where not all state-pairs are updated at each iteration, but rather \mathcal{S} is split into disjoint sets that are updated at different intervals. A few strategies for sampling state-pairs are discussed, of which the most similar to ours is the ‘uniform asynchronous update’.

Gelada et al. (2019) introduced DeepMDP latent models and established a close relationship with bisimulation metrics (specifically, Theorem 4 in their paper). Although closely related, there are some important differences. Their work deals

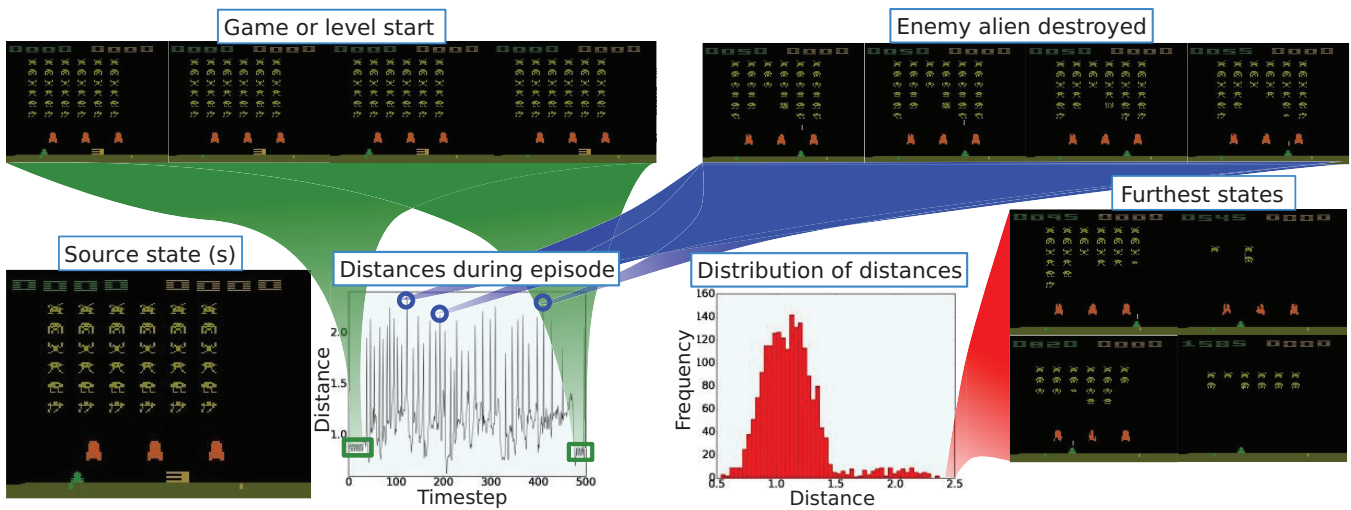


Figure 6: Evaluating the approximat to d^π with an eval run on Space Invaders. We plot the distance between the source state s (pictured in the bottom left) and every other state, highlighting the relatively low distances in game or level starts (green shading pointing to left and right side of the plot), as well as the peaks occuring when an enemy alien is destroyed (blue shading pointing to the distance peaks). On the bottom right we display the distribution of distances and the four furthest states for this run.

with state representations, where the distance between states is their distance in the representation space; by contrast, our proposed neural networks approximate the bisimulation metric between two states, independent of their representation. Further, the authors use the DeepMDP losses as an auxiliary task without a guarantee that their representations are consistent with their theoretical results. In our work we are able to show that our approximat is valid both quantitatively (GridWorld) and qualitatively (Atari 2600). Nonetheless, a natural extension of our work is to use the bisimulation losses we introduced as a means to learn better representations.

Conclusion

We introduced new methods for computing and approximating bisimulation metrics in large deterministic MDPs. The first is an *exact* method that converges to the true metric asymptotically, and the second is a differentiable method for approximating the metric which we demonstrated can learn a good approximat even in continuous state spaces. Since their introduction, bisimulation metrics have been used for theoretical analysis or in MDPs of small-to-moderate size, but they have scarcely been used in larger problems. Our results open the doors for their application in systems with a large, and even continuous, state space.

One important avenue for research is to extend these results to stochastic systems. Computing the Wasserstein metric without access to a generative model is challenging for deep RL environments, as the next-state samples typically come from single trajectories in replay buffers. One possibility is to build a model of the transition dynamics from the transitions in the replay buffer and compute the Wasserstein metrics from this estimate.

Although the network architecture and hyperparameters used to train d^π are by no means optimal, the results we presented for the Atari 2600 domain are very promising and

suggest that bisimulation metrics can be used effectively for deep reinforcement learning. Some promising areas we are currently exploring are using bisimulation metrics as an auxiliary task for improved state representation, as a mechanism for compressing replay buffers, and as a tool for more efficient exploration.

Acknowledgements

The author would like to thank Marc G. Bellemare, Gheorghe Comanici, Marlos C. Machado, Doina Precup, Carles Gelada, as well as the rest of the Google Brain team in Montreal for helpful discussions. The author would also like to thank the anonymous reviewers for their useful feedback while reviewing this work.

References

- Bacci, G.; Bacci, G.; Larsen, K. G.; and Mardare, R. 2013a. Computing Behavioral Distances, Compositionally. In *Mathematical Foundations of Computer Science 2013*.
- Bacci, G.; Bacci, G.; Larsen, K. G.; and Mardare, R. 2013b. On-the-Fly Exact Computation of Bisimilarity Distances. In *Proceedings of the 19th Int. Conf. on Tools and Algorithms for the Construction and Analysis of Systems*.
- Baier, C.; Hermans, H.; Katoen, J.; and Wolf, V. 2006. Bisimulation and simulation relations for markov chains. In *Essays on Algebraic Process Calculi*, number 10 in Electronic Notes in Theoretical Computer Science, 73–78. Elsevier.
- Bellemare, M. G.; Naddaf, Y.; Veness, J.; and Bowling, M. 2013. The Arcade Learning Environment: An evaluation platform for general agents. *Jour. of AI Research* 47:253–279.
- Bellman, R. 1957. *Dynamic Programming*. Princeton, NJ, USA: University Press.

- Castro, P. S., and Precup, D. 2010. Using bisimulation for policy transfer in MDPs. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2010)*.
- Castro, P. S.; Moitra, S.; Gelada, C.; Kumar, S.; and Bellemare, M. G. 2018. Dopamine: A Research Framework for Deep Reinforcement Learning (arxiv.org/abs/1812.06110).
- Castro, P. S. 2011. *On planning, prediction and knowledge transfer in Fully and Partially Observable Markov Decision Processes*. Ph.D. Dissertation, McGill University.
- Chen, D.; van Breugel, F.; and Worrell, J. 2012. On the Complexity of Computing Probabilistic Bisimilarity. In *Found. of Software Science and Computational Structures*.
- Comanici, G.; Panangaden, P.; and Precup, D. 2012. On-the-Fly Algorithms for Bisimulation Metrics. In *Proc. of the 9th Int. Conf. on Quantitative Evaluation of Systems*.
- Desharnais, J.; Gupta, V.; Jagadeesan, R.; and Panangaden, P. 1999. Metrics for Labeled Markov Systems. In *CONCUR'99 Concurrency Theory*, 258–273.
- Ferns, N., and Precup, D. 2014. Bisimulation Metrics are Optimal Value Functions. In *Proceedings of the 30th Conference on Uncertainty in Artificial Intelligence*.
- Ferns, N.; Castro, P. S.; Precup, D.; and Panangaden, P. 2006. Methods for computing state similarity in Markov decision processes. In *Proceedings of the 22nd Conference on Uncertainty in Artificial Intelligence, UAI '06*.
- Ferns, N.; Panangaden, P.; and Precup, D. 2004. Metrics for Finite Markov Decision Processes. In *Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence*.
- Ferrer Fioriti, L. M.; Hashemi, V.; Hermanns, H.; and Turrini, A. 2016. Deciding probabilistic automata weak bisimulation: Theory and practice. *Form. Asp. Comput.* 28(1):109–143.
- Gelada, C.; Kumar, S.; Bucknam, J.; Nachum, O.; and Bellemare, M. G. 2019. DeepMDP: Learning Continuous Latent Space Models for Representation Learning. In *Proceedings of the 36th International Conference on Machine Learning*.
- Givan, R.; Dean, T.; and Greig, M. 2003. Equivalence notions and model minimization in Markov decision processes. *Artificial Intelligence* 147:163–223.
- Hessel, M.; Modayil, J.; van Hasselt, H.; Schaul, T.; Ostrovski, G.; Dabney, W.; Horgan, D.; Piot, B.; Azar, M.; and Silver, D. 2018. Rainbow: Combining Improvements in Deep Reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Jiang, N.; Kulesza, A.; Singh, S.; and Lewis, R. 2015. The Dependence of Effective Planning Horizon on Model Accuracy. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-15)*.
- Katoen, J.-P.; Kemna, T.; Zapreev, I.; and Jansen, D. N. 2007. Bisimulation minimisation mostly speeds up probabilistic model checking. In *Tools and Algorithms for the Construction and Analysis of Systems*, 87–101.
- Kingma, D. P., and Ba, J. 2015. Adam: A method for stochastic optimization. In *Proceedings of the International Conference on Learning Representations*.
- Li, L.; Walsh, T. J.; and Littman, M. L. 2006. Towards a unified theory of state abstraction for mdps. In *Proceedings of the Ninth International Symposium on Artificial Intelligence and Mathematics*, 531–539.
- Machado, M. C.; Bellemare, M. G.; Talvitie, E.; Veness, J.; Hausknecht, M.; and Bowling, M. 2018. Revisiting the Arcade Learning Environment: Evaluation protocols and open problems for general agents. *Journal of AI Research*.
- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Rusu, A. A.; Veness, J.; Bellemare, M. G.; Graves, A.; Riedmiller, M.; Fidjeland, A. K.; Ostrovski, G.; Petersen, S.; Beattie, C.; Sadik, A.; Antonoglou, I.; King, H.; Kumaran, D.; Wierstra, D.; Legg, S.; and Hassabis, D. 2015. Human-level control through deep reinforcement learning. *Nature*.
- Puterman, M. L. 1994. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. New York, NY, USA: John Wiley & Sons, Inc., 1st edition.
- Ravindran, B., and Barto, A. G. 2003. Relativized Options: Choosing the Right Transformation. In *Proceedings of the 20th International Conference on Machine Learning*.
- Ruan, S.; Comanici, G.; Panangaden, P.; and Precup, D. 2015. Representation discovery for mdps using bisimulation metrics. In *AAAI Conference on Artificial Intelligence*.
- Sangiorgi, D. 2011. *Introduction to Bisimulation and Coinduction*. Cambridge University Press.
- Santara, A.; Madan, R.; Ravindran, B.; and Mitra, P. 2019. Extra: Transfer-guided exploration. *CoRR* abs/1906.11785.
- Segala, R., and Turrini, A. 2007. Approximated computationally bounded simulation relations for probabilistic automata. In *Proceedings of the 20th IEEE Computer Security Foundations Symposium (CSF'07)*.
- Sutton, R. S., and Barto, A. G. 1998. *Introduction to Reinforcement Learning*. Cambridge, MA, USA: MIT Press.
- Taylor, J.; Precup, D.; and Panagaden, P. 2009. Bounding Performance Loss in Approximate MDP Homomorphisms. In *Advances in Neural Information Processing Systems 21*. 1649–1656.
- Villani, C. 2008. *Optimal Transport*. Springer-Verlag Berlin Heidelberg.
- Zager, L. A., and Verghese, G. C. 2008. Graph similarity scoring and matching. *Applied Mathematics Letters* 21(1):86–94.