# Representation Discovery for MDPs Using Bisimulation Metrics 

Sherry Shanshan Ruan, Gheorghe Comanici, Prakash Panangaden, Doina Precup<br>School of Computer Science,<br>McGill University, Montreal, QC, Canada<br>\{sherry, gcoman, prakash, dprecup\} @cs.mcgill.ca


#### Abstract

We provide a novel, flexible, iterative refinement algorithm to automatically construct an approximate statespace representation for Markov Decision Processes (MDPs). Our approach leverages bisimulation metrics, which have been used in prior work to generate features to represent the state space of MDPs. We address a drawback of this approach, which is the expensive computation of the bisimulation metrics. We propose an algorithm to generate an iteratively improving sequence of state space partitions. Partial metric computations guide the representation search and provide much lower space and computational complexity, while maintaining strong convergence properties. We provide theoretical results guaranteeing convergence as well as experimental illustrations of the accuracy and savings (in time and memory usage) of the new algorithm, compared to traditional bisimulation metric computation.


## Introduction

Solving large sequential decision problems modeled as Markov Decision Processes (MDPs) requires the use of approximations to represent the state space. Popular approximation methods include state aggregation, linear function approximation and kernel-based methods. In this paper we are mainly interested in state aggregation, in which the state space is partitioned into disjoint subsets and values are associated with each partition. The goal is to construct a partition incrementally, in such a way as to provide a good approximation to the true value function. One approach for this problem is to use bisimulation relations (Givan, Dean, and Greig 2003), also known as MDP homomorphisms (Ravindran and Barto 2002), or their relaxation as bisimulation metrics (Ferns, Panangaden, and Precup 2004). Bisimulation metrics in particular are attractive because they allow quantifying the approximation error for any state space partitioning, or more generally, any linear function approximator (Comanici and Precup 2012). However, bisimulation metric computation is very expensive (in the worst case, more expensive than performing dynamic programming in the original state space). Indeed, recent work (Ferns and Precup 2014) has shown that computing the metric amounts to

[^0]solving an MDP resulting from a coupling of the state space with itself; such a coupling has size quadratic in the number of states.
In this paper, we tackle this problem by proposing a significant improvement in how bisimulation metrics are computed. While we present our results in the context of MDPs, the algorithmic ideas can also benefit the verification community, which relies on bisimulation metrics in order to automatically verify the concordance of a model with a specification.

Our approach constructs an iteratively improving sequence of state space partitions, which converges in the limit to the bisimulation relation, just like previous algorithms. We prove that at each step the error of the value function computed over this partition (compared to the true optimal value function) is bounded. Since at each step, the value function approximation is computed over partitions rather than states, this approach can generate substantial space and computation time savings, as illustrated in our experiments described later.
The second contribution of the paper consists in an algorithm for asynchronous updates of the metric and the representation. We provide theoretical conditions which allow computational effort to be focused on parts of the state space where changes are happening rapidly, similar to successful asynchronous or distributed dynamic programming techniques such as Bertsekas and Tsitsiklis (1996), Bertsekas and Castanon (1989), Moore and Atkeson (1993) . Empirical results illustrate the use of heuristics that can substantially speed up the computation.

## Markov Decision Processes

An MDP is a tuple $(S, A, P, R)$ where $S$ is a finite state space $S, A$ is a finite action space, $P$ is a transition probability function $P: S \times A \times S \rightarrow[0,1]$, with $P_{s s^{\prime}}^{a}$ giving the probability that the system will end up in state $s^{\prime}$ when the action $a$ is performed in state $s$, and $R: S \times A \rightarrow \mathbb{R}$ is the immediate reward function, with $R_{s}^{a}$ giving the immediate reward for performing action $a$ in state $s$.

The main objective of MDP solvers is to compute value functions for different policies, i.e., strategies for choosing actions. A policy is denoted by $\pi: S \times A \rightarrow[0,1]$ where $\pi_{s}^{a}$ is the probability of choosing action $a$ in state $s$. We will define below the value function of $\pi$, which is dependent on a
discount factor $\gamma(0 \leq \gamma<1), R_{s}^{\pi}=\sum_{a} \pi_{s}^{a} R_{s}^{a}$, and $P_{s s^{\prime}}^{\pi}=$ $\sum_{a} \pi_{s}^{a} P_{s s^{\prime}}^{a}$. The value function is given by the following expectation over sample trajectories $X_{1}, X_{2}, X_{3}, \ldots, X_{n}, \ldots$ :

$$
V^{\pi}(s)=E\left[\sum_{i=1}^{\infty} \gamma^{i-1} R_{X_{i}}^{\pi} \mid X_{1}=s\right]=R_{s}^{\pi}+\gamma \sum_{s^{\prime}} P_{s s^{\prime}}^{\pi} V^{\pi}\left(s^{\prime}\right)
$$

Most algorithms for solving MDPs either use the model $(R, P)$ to find $V^{\pi}$ if it is available and allowed by the size of the problem, or estimate $V^{\pi}$ using samples $\left(s, a, r, s^{\prime}\right)$. In both cases, the computation of $V^{\pi}$ is performed in the space $\mathcal{F}_{S}$ of real valued functions over $S$, i.e., $\mathcal{F}_{S}:=\{f: S \rightarrow \mathbb{R}\}$.

The Bellman equation is a well-known result characterizing the value function $V^{\pi} \in \mathcal{F}_{S}$ as the fixed point of the following map $T^{\pi}: \mathcal{F}_{S} \rightarrow \mathcal{F}_{S}$ given by

$$
\begin{align*}
& T^{\pi}(f)=\mathcal{R}^{\pi}+\gamma \mathcal{P}^{\pi}(f) \\
& \text { with } \mathcal{R}^{\pi} \in \mathcal{F}_{S}, \mathcal{R}^{\pi}(s)=\sum_{a} \pi_{s}^{a} R_{s}^{a} \\
& \mathcal{P}^{\pi}: \mathcal{F}_{S} \rightarrow \mathcal{F}_{S}, \mathcal{P}^{\pi}(f)(s)=\sum_{a, s^{\prime}} \pi_{s}^{a} P_{s s^{\prime}}^{a} f\left(s^{\prime}\right) \tag{1}
\end{align*}
$$

The value $V^{*}$ associated with the best policy is the fixed point of the nonlinear Bellman optimality operator $T^{*}$ :

$$
\begin{equation*}
T^{*}(f)=\max _{a}\left(\mathcal{R}^{a}+\gamma \mathcal{P}^{a}(f)\right) \tag{2}
\end{equation*}
$$

where $\mathcal{R}^{a}$ and $\mathcal{P}^{a}$ are associated with the policy choosing $a$ deterministically.

## Linear value functions

In most cases, working with the original representation is unfeasible due to the large size of the state space. It is common to work instead with restrictions over the space of functions, $\Phi \subset \mathcal{F}_{S}$, for which searching for $V^{\pi}$ (or its approximation) becomes computationally feasible. It is also quite common to work with linear subspace $\Phi=$ $\operatorname{span}\left(\phi_{1}, \phi_{2}, \ldots, \phi_{k}\right)$, for some set of linearly independent finite subset $\Phi$ of $\mathcal{F}$ of dimension $k \ll S$.

Accurate representations: If one wants $V^{\pi}$ to be an element of $\Phi$, then it is sufficient to guarantee the accuracy of the reward and transition models: $\mathcal{R}^{\pi} \in \Phi$ and $\forall f \in \Phi, \mathcal{P}^{\pi}(f) \in \Phi$. Several methods have been developed to obtain accurate representations for a given policy $\pi$ (Parr et al. 2008b; 2008a). The methods described later in this paper are designed to provide accuracy for the model of any given policy.

Approximate representations: Given a $\Phi$, linear fixed point methods such as TD, LSTD, LSPE (Sutton 1988; Bradtke and Barto 1996; Yu and Bertsekas 2006) can be used to find the least squares fixed point approximation $V_{\Phi}^{\pi}$ of $V^{\pi}$, which is the fixed point of the alternate operator $T_{\Phi}^{\pi}$ :

$$
T_{\Phi}^{\pi} f:=\Pi_{\Phi}\left(\mathcal{R}^{\pi}+\gamma \mathcal{P}^{\pi} f\right)
$$

where $\Pi_{\Phi}$ is the orthogonal projection operator on $\Phi$. Using the fact that $\Pi_{\Phi}$ is itself linear, it is not hard to show that $T_{\pi}^{\Phi} f=\Pi_{\Phi} \mathcal{R}^{\pi}+\gamma \Pi_{\Phi} \mathcal{P}^{\pi} f$; therefore, $V_{\Phi}^{\pi}$ is the fixed point of the Bellman operator over the linear model $\left(\mathcal{R}_{\Phi}^{\pi}, \mathcal{P}_{\Phi}^{\pi}\right):=\left(\Pi_{\Phi} \mathcal{R}^{\pi}, \Pi_{\Phi} \mathcal{P}^{\pi}\right)$. For computational details using linear models for evaluation, see Parr et al. (2007; 2008) .

## Bisimulation relations and metrics

Probabilistic bisimulation is an equivalence relation between states of a process due to Larsen and Skou (1991). It was extended to MDPs with rewards by Givan et al. (2003); the metric analogue is due to Desharnais et al. $(1999,2004)$ and the extension of the metric to include rewards is from Ferns et al. (2004). Suppose we are given an equivalence relation $\sim$ on the state space $S$. We say that $\sim$ is a strong probabilistic bisimulation relation if the following two conditions are satisfied for any equivalent pair $s \sim s^{\prime}$ and any choice of action $a \in A: 1 . R_{s}^{a}=R_{s^{\prime}}^{a} ; 2 . P_{s}^{a}=P_{s^{\prime}}^{a}$, as probability measures over equivalence classes of $\sim$. The apparent circularity in this definition can be resolved by a fixed-point argument (Larsen and Skou 1991). The crucial point is that the immediate behaviour of bisimilar states is the same and it stays the same indefinitely.

Note that any equivalence relation $\sim$ determines a representation $B_{\sim}=\left\{\phi_{i}\right\}_{i=1}^{m}$, where $m$ is the number of equivalence classes of $\sim$ and each $\phi_{i}$ is the characteristic function over equivalence class $i$. For notational convenience, let $B_{\sim} \in\{0,1\}^{m \times|S|}$ be the matrix whose columns are the $\phi_{i}$ 's. It is not hard to check that $B^{T} e=e$ (where $e(s)=1, \forall s)$. Let $\delta_{s}$ be the characteristic function that is equal to 1 at state $s$ and 0 for all other states. Given any matrix $M, M \delta_{s}$ and $\delta_{s}^{T} M$ are the column and row respectively corresponding to $s$. Note that $s \sim s^{\prime}$ iff $B \delta_{s}=B \delta_{s^{\prime}}$ (i.e., $s$ and $s^{\prime}$ have the same representation). Another way to characterize a strong probabilistic bisimulation relation $\sim$ is through the following conditions: $\begin{cases}\mathcal{R}^{a} \in \operatorname{colspan}\left(B_{\sim}\right) & \forall a \\ \mathcal{P}^{a} \phi \in \operatorname{colspan}\left(B_{\sim}\right) & \forall a, \forall \phi \in B_{\sim}\end{cases}$
Note that checking whether $f \in \operatorname{colspan}\left(B_{\sim}\right)$ is equivalent to checking that $f(s)=f\left(s^{\prime}\right)$ for all pairs $s \sim s^{\prime}$, as $B_{\sim}$ consists of characteristic functions of equivalence classes.

Bisimulation metrics provide relaxations of bisimulation relations, which assign non-negative values to all pairs of states. Two states are bisimilar iff their distance is zero. For other pairs of states that are not bisimilar, the metric quantifies how different the states are from each other. Bounds relating bisimulation metrics to value functions have been derived in Ferns et al. (2004). The metric presented in their work is based on the Monge-Kantorovich metric $\mathcal{T}_{d}(\mu, \nu)$ for comparing two probability distributions $\mu$ and $\nu$ using a ground distance $d$ :

$$
\begin{aligned}
\Lambda(\mu, \nu) & =\left\{\lambda \geq 0 \text { s.t. } \lambda e=\mu \text { and } \lambda^{T} e=\nu\right\} \\
\mathcal{T}_{d}(\mu, \nu) & :=\max _{\lambda \in \Lambda(\mu, \nu)} \sum_{s, s^{\prime}}\left(\delta_{s}^{T} \lambda \delta_{s^{\prime}}\right)\left(\delta_{s} d \delta_{s^{\prime}}\right)
\end{aligned}
$$

Note that Ferns et al. (2004) does not use couplings; this more recent view is expressed in Ferns and Precup (2014). However, the definition using couplings is much more convenient for proving the results in the following sections. Note that $\mathcal{T}_{d}(\mu, \nu)$ is 0 if and only if $\mu$ and $\nu$ match as measures over classes of states at distance $d$ equal to 0 . One can use similar measures to compare the reward model, and we use the $L_{1}$ measure in the definition below (but any norm on
$\mathbb{R}$ could equally be used):

$$
\begin{aligned}
F(d)\left(s, s^{\prime}\right) & :=\max _{a}\left(\left|\mathcal{R}^{a}(s)-\mathcal{R}^{a}\left(s^{\prime}\right)\right|+\gamma \mathcal{T}_{d}\left(\delta_{s}^{T} \mathcal{P}^{a}, \delta_{s^{\prime}}^{T} \mathcal{P}^{a}\right)\right) \\
d^{*} & :=\sup _{n} F^{n}(0)
\end{aligned}
$$

Two theoretical results about $F$ allow one to compute $d^{*}$ iteratively: first, the collection of pseudo-metrics is equipped with an order, making it a complete lattice ${ }^{1}$; secondly, $F$ is shown to be a monotone map. Therefore, $F^{n}(0)$ increases to the solution $d^{*}$, and $d^{*}$ is a fixed point of $F$. Moreover, it is not hard to see that if we relate states $s$ and $s^{\prime}$ iff $d^{*}\left(s, s^{\prime}\right)=0$, then the corresponding relation is a bisimulation. Additionally, Ferns et al. (2004) state bounds which allow one to assess the quality of representations determining partitions over the state space ${ }^{2}$. Given such a representation $B$, if one computes the value function $V_{B}^{*}$ as the fixed point of $T_{B}^{\pi^{*}}$ instead of the fixed point of $T^{*}$, the approximation error is bounded as follows:

$$
\begin{equation*}
\left|V_{B}^{*}(s)-V^{*}(s)\right| \leq \frac{\delta_{s}^{T} \Pi_{B} d^{*} \delta_{s}}{1-\gamma}+\frac{\gamma \max _{s^{\prime}} \delta_{s^{\prime}}^{T} \Pi_{B} d^{*} \delta_{s^{\prime}}}{(1-\gamma)^{2}} \tag{3}
\end{equation*}
$$

Note that for representations that are based on characteristic functions, projection amounts to averaging. That is,

$$
\delta_{s}^{T} \Pi_{B} f=\sum_{i} \phi_{i}(s)\left(\phi^{T} f\right) /\left(\phi^{T} \phi\right)
$$

Hence, the value function approximation bounds depend on $\delta_{s}^{T} \Pi_{B} d^{*} \delta_{s}$, the average d* distance from a state s to all other states that have the same representation. It should be mentioned that these bounds are presented in terms of state aggregation maps in Ferns et al. (2004). See Parr et al. (2008) for a more detailed discussion on why solving a linear fixed-point solution (i.e., finding the fixed point of $T_{B}^{*}$ ) amounts to the same solution as solving a linear-model solution of an aggregate model (i.e., finding the fixed point of $T^{*}$ over the aggregate model). Note that these bounds are minimized by aggregating states which are "close" in terms of the bisimulation distance $d^{*}$, i.e., states that are close to being bisimilar. It should also be noted that $V_{B \sim}^{*}$ is exactly $V^{*}$ when $B_{\sim}$ represents a bisimulation relation. This is because $d^{*}$ is a pseudo-metric for which distance between states with the same representation is 0 , and as a consequence $\delta_{s}^{T} \Pi_{B_{\sim}} d^{*} \delta_{s}=0$.

Similar bounds for normalized features can be found in Comanici and Precup (2011). Moreover, their work presents ways of generating feature-based representations using spectral analysis methods with bisimulation metrics as ground metrics. We now proceed to analyze alternative characterizations of $d^{*}$ and describe a new algorithmic framework for computing $d^{*}$.

[^1]
## Iterative refinement algorithm

In this section we will characterize $d^{*}$ through a series of metrics over compact representations of the original MDPs. To achieve this, we will first define partitions and show how one can use these compact representations to compute iteratively improving approximations to $d^{*}$. A partition $\mathcal{B}$ is a basis $\left\{\phi_{i}\right\}_{i=1}^{m}$ such that $\phi_{i} \in\{0,1\}$ and $\sum_{i} \phi_{i}(s)=1, \forall s \in$ $S$. We will also use $B$ to denote the matrix whose columns are the elements in $\mathcal{B}$.

Let $\widehat{d}^{*}=\sup _{n} B_{n}^{T} d_{n} B_{n}$, where $d_{n}$ are metrics over a sequence of partitions $B_{n}$, defined inductively as follows:

- $B_{0}=\{e\} \quad(e(s)=1$ for every $s)$
- $B_{n}$ is a partition such that

$$
\begin{cases}\mathcal{R}^{a} \in \operatorname{colspan}\left(B_{n}\right) & \forall a \\ \mathcal{P}^{a} \phi, \phi \in \operatorname{colspan}\left(B_{n}\right) & \forall a, \forall \phi \in B_{n-1}\end{cases}
$$

- Define the metric $d_{n}$ based on the Monge-Kantorovich metric over partition $B_{n-1}$ with ground metric $d_{n-1}$. For $\phi, \phi^{\prime} \in B_{n}$, if $\phi(s)=1$ and $\phi^{\prime}\left(s^{\prime}\right)=1$, then

$$
\begin{aligned}
d_{n}\left(\phi, \phi^{\prime}\right)= & \max _{a}\left(\left|\mathcal{R}^{a}(s)-\mathcal{R}^{a}\left(s^{\prime}\right)\right|\right. \\
& \left.+\gamma \mathcal{T}_{d_{n-1}}\left(\delta_{s}^{T} \mathcal{P}^{a} B^{T}, \delta_{s^{\prime}}^{T} \mathcal{P}^{a} B^{T}\right)\right)
\end{aligned}
$$

Note that $B_{n}$ is purposely defined so that the value $d_{n}\left(\phi, \phi^{\prime}\right)$ is the same for any choice of states $s$ and $s^{\prime}$ with representations $\phi$ and, respectively $\phi^{\prime}$.
Theorem 1. Given the sequence $\left\{B_{n}, d_{n}\right\}_{n=1}^{\infty}$ defined above, it follows that $B_{n}^{T} d_{n} B_{n}=F^{n}(0)$. Then

$$
d^{*}=\sup _{n} B_{n}^{T} d_{n} B_{n}
$$

## This result is key in designing the new algorithm.

Proof. First, we will prove that for a partition $B$,

$$
\mathcal{T}_{d}(B \mu, B \nu)=\mathcal{T}_{B^{T} d B}(\mu, \nu)
$$

As mentioned before, the Monge-Kantorovich metric is an optimization function over the set of "couplings" of two measures. Remember that the set of couplings $\Lambda(\mu, \nu)$ is the set of maps $\lambda$ such that $\lambda e=\mu$ and $\lambda^{T} e=\nu$. Before we continue, it is helpful to point out a few properties of $B$ which will make the derivation much simpler: $B^{T} e=e$; if $\phi(s)=1, B \delta_{s}=\delta_{\phi}$. Now, given $\lambda \in \Lambda(\mu, \nu)$,

$$
\begin{aligned}
B \lambda B^{T} e & =B \lambda e=B \mu \\
\left(B \lambda B^{T}\right)^{T} e & =B \lambda^{T} e=B \nu
\end{aligned}
$$

Therefore $B \lambda B^{T} \in \Lambda(B \mu, B \nu)$. Now,

$$
\begin{aligned}
& \sum_{s, s^{\prime}} \delta_{s}^{T} \lambda \delta_{s^{\prime}} \delta_{s}^{T}\left(B^{T} d B\right) \delta_{s^{\prime}} \\
& =\sum_{\phi, \phi^{\prime}} \sum_{s, s^{\prime}} \phi(s) \phi^{\prime}\left(s^{\prime}\right) \delta_{s}^{T} \lambda \delta_{s^{\prime}}\left(B \delta_{s}\right)^{T} d\left(B \delta_{s^{\prime}}\right) \\
& =\sum_{\phi, \phi^{\prime}} \delta_{\phi}^{T} d \delta_{\phi^{\prime}} \sum_{s, s^{\prime}} \phi(s) \delta_{s}^{T} \lambda \delta_{s^{\prime}} \phi^{\prime}\left(s^{\prime}\right) \\
& =\sum_{\phi, \phi^{\prime}}\left(\delta_{\phi}^{T} d \delta_{\phi^{\prime}}\right)\left(\delta_{\phi}^{T} B \lambda B^{T} \delta_{\phi^{\prime}}\right)
\end{aligned}
$$

$$
\begin{aligned}
& \mathcal{T}_{B^{T} d B}(\mu, \nu)=\inf _{\lambda \in \Lambda(\mu, \nu)} \sum_{s, s^{\prime}} \delta_{s}^{T} \lambda \delta_{s^{\prime}} \delta_{s}^{T}\left(B^{T} d B\right) \delta_{s^{\prime}} \\
& =\inf _{\lambda \in \Lambda(\mu, \nu)} \sum_{\phi, \phi^{\prime}}\left(\delta_{\phi}^{T} d \delta_{\phi^{\prime}}\right)\left(\delta_{\phi}^{T} B \lambda B^{T} \delta_{\phi^{\prime}}\right) \\
& \geq \inf _{\bar{\lambda} \in \Lambda(B \mu, B \nu)} \sum_{\phi, \phi^{\prime}}\left(\delta_{\phi}^{T} d \delta_{\phi^{\prime}}\right)\left(\delta_{\phi}^{T} \bar{\lambda} \delta_{\phi^{\prime}}\right)=\mathcal{T}_{d}(B \mu, B \nu)
\end{aligned}
$$

```
Algorithm 1 partition_declust
    Given a partition \(B\) and one of its elements \(\phi \in B\), we
    want to find \(B_{\phi}\). We initialize it as \(B_{\phi} \leftarrow \emptyset\).
    for all \(s\) with \(\phi(s)=1\) do
        for all \(\phi^{\prime} \in B_{\phi}\) do
            choose \(s^{\prime}\) with \(\phi^{\prime}\left(s^{\prime}\right)=1\)
            if \(\forall a, \forall \phi^{\prime \prime} \in B,\left(\mathcal{P}^{a} \phi^{\prime \prime}\right)(s)=\left(\mathcal{P}^{a} \phi^{\prime \prime}\right)\left(s^{\prime}\right)\)
            then \(\phi^{\prime}(s) \leftarrow 1\) \{An equivalent condition can be
            used when declustering based on reward\}
        end for
        if \(\phi^{\prime}(s)=0, \forall \phi^{\prime} \in B_{\phi}\)
        then add a new element \(\widehat{\phi}\) to \(B_{\phi}\) and set \(\widehat{\phi}(s)=1\)
    end for
```

So $\mathcal{T}_{d}(B \mu . B \nu) \leq \mathcal{T}_{B^{T} d B}(\mu, \nu)$. Now we will prove the opposite. Define $\widehat{B}_{\mu}(s, \phi):=\phi(s) \mu(s) /(B \mu)(\phi)$. Note that $B \widehat{B}_{\mu}=I$ and $\widehat{B}_{\mu} B \mu=e$. Now let $\lambda \in \Lambda(B \mu, B \nu)$,

$$
\begin{aligned}
& \widehat{B}_{\mu} \lambda \widehat{B}_{\nu}^{T} e=\widehat{B}_{\mu} \lambda \widehat{B}_{\nu}^{T} B^{T} e=\widehat{B}_{\mu} \lambda e=\widehat{B}_{\mu} B \mu=\mu \\
& \left(\widehat{B}_{\mu} \lambda \widehat{B}_{\nu}^{T}\right)^{T} e=\widehat{B}_{\nu} \lambda^{T} \widehat{B}_{\mu}^{T} e=\nu
\end{aligned}
$$

Therefore $\widehat{B}_{\mu} \lambda \widehat{B}_{\nu}^{T} \in \Lambda(\mu, \nu)$. Now,

$$
\begin{aligned}
& \sum_{\phi, \phi^{\prime}} \delta_{\phi}^{T} \lambda \delta_{\phi^{\prime}} \delta_{\phi}^{T} d \delta_{\phi^{\prime}} \\
& =\sum_{\phi, \phi^{\prime}}\left(\delta_{\phi}^{T} \lambda \delta_{\phi^{\prime}} \delta_{\phi}^{T} d \delta_{\phi^{\prime}}\right)\left(\delta_{\phi}^{T} B \widehat{B}_{\mu} \delta_{\phi}\right)\left(\delta_{\phi^{\prime}}^{T} \widehat{B}_{\nu}^{T} B^{T} \delta_{\phi^{\prime}}\right) \\
& =\sum_{\phi, \phi^{\prime}}\left(\delta_{\phi}^{T} \lambda \delta_{\phi^{\prime}}^{T} \delta_{\phi}^{T} d \delta_{\phi^{\prime}}\right) \sum_{s, s^{\prime}} \phi(s) \phi\left(s^{\prime}\right)\left(\delta_{s}^{T} \widehat{B}_{\mu} \delta_{\phi}\right)\left(\delta_{\phi^{\prime}}^{T} \widehat{B}_{\nu}^{T} \delta_{s^{\prime}}\right) \\
& =\sum_{\phi, \phi^{\prime}} \delta_{\phi}^{T} d \delta_{\phi^{\prime}} \sum_{s, s^{\prime}} \phi(s) \phi\left(s^{\prime}\right)\left(\delta_{s}^{T} \widehat{B}_{\mu}\left(\delta_{\phi} \delta_{\phi}^{T} \lambda \delta_{\phi^{\prime}} \delta_{\phi^{\prime}}^{T} \widehat{B}_{\nu}^{T} \delta_{s^{\prime}}\right)\right. \\
& =\sum_{s, s^{\prime}}\left(\delta_{s}^{T} B^{T} d B \delta_{s^{\prime}}\right)\left(\delta_{s}^{T} \widehat{B}_{\mu} \lambda \widehat{B}_{\nu}^{T} \delta_{s^{\prime}}\right) \\
& \\
& \mathcal{T}_{d}(B \mu, B \nu)=\inf _{\lambda \in \Lambda(\mu, B)} \sum_{\phi, \phi^{\prime}} \delta_{\phi}^{T} \lambda \delta_{\phi^{\prime}} \delta_{\phi}^{T} d \delta_{\phi^{\prime}} \\
& \quad \inf _{\lambda \in \Lambda(B \mu, B \nu)} \sum_{s, s^{\prime}}\left(\delta_{s}^{T} B^{T} d B \delta_{s^{\prime}}\right)\left(\delta_{s}^{T} \widehat{B}_{\mu} \lambda \widehat{B}_{\nu}^{T} \delta_{s^{\prime}}\right) \\
& \sum_{s, s^{\prime}}\left(\delta_{s}^{T} B^{T} d B \delta_{s^{\prime}}\right)\left(\delta_{s}^{T} \bar{\lambda} \delta_{s^{\prime}}\right)=\mathcal{T}_{B^{T} d B}(\mu, \nu)
\end{aligned}
$$

Next, we can prove by induction that $B_{n}^{T} d_{n} B_{n}=F^{n}(0)$. We skip the base case which is trivial. Assuming the statement holds for $n$, let $\phi, \phi^{\prime} \in B_{n+1}$ such that $\phi(s)=1$ and $\phi^{\prime}\left(s^{\prime}\right)=1$.

$$
\begin{aligned}
& \left(B_{n+1}^{T} d_{n+1} B_{n+1}\right)\left(s, s^{\prime}\right)=d_{n+1}\left(\phi, \phi^{\prime}\right) \\
& =\max _{a \in A}\left(\left|\mathcal{R}^{a}(s)-\mathcal{R}^{a}\left(s^{\prime}\right)\right|+\gamma \mathcal{T}_{d_{n}}\left(\delta_{s}^{T} \mathcal{P}^{a} B^{T}, \delta_{s^{\prime}}^{T} \mathcal{P}^{a} B^{T}\right)\right) \\
& =\max _{a \in A}\left(\left|\mathcal{R}^{a}(s)-\mathcal{R}^{a}\left(s^{\prime}\right)\right|+\gamma \mathcal{T}_{B_{n}^{T} d_{n} B_{n}}\left(\delta_{s}^{T} \mathcal{P}^{a}, \delta_{s^{\prime}}^{T} \mathcal{P}^{a}\right)\right) \\
& =\max _{a \in A}\left(\left|\mathcal{R}^{a}(s)-\mathcal{R}^{a}\left(s^{\prime}\right)\right|+\gamma \mathcal{T}_{F^{n}(0)}\left(\delta_{s}^{T} \mathcal{P}^{a}, \delta_{s^{\prime}}^{T} \mathcal{P}^{a}\right)\right) \\
& =F^{n+1}(0)\left(s, s^{\prime}\right)
\end{aligned}
$$

Now we have the desired result, which is:

$$
d^{*}=\sup _{n} F^{n}(0)=\sup _{n} B_{n}^{T} d_{n} B_{n}
$$

```
Algorithm 2 Synchronous declustering
    \(B_{1} \leftarrow\) partition_declust of \(\{S\}\) based on reward
    while we seek a better approximation: \(i=1 \cdots \infty\) do
        \(B_{i+1} \leftarrow \emptyset\)
        for all \(\phi \in B_{i}\) do
            add partition_declust of \(\phi\) using transition
            \(P\) over \(B_{i}\) to \(B_{i+1}\)
        end for
    end while
```

Based on the result of the theorem above, one can compute $d^{*}$ using an iterative algorithm which generates partition refinements $B_{n}$ and computes a metric over these partitions. Refining $B_{n}$ to $B_{n+1}$ can be done efficiently using Algorithm 1 (note that this algorithm is presented for the transition map only, but it is easily modifiable to decluster based on the reward function). The complexity of computing $d_{n}$ given $d_{n-1}$ depends on the following terms: $\phi^{T} \phi$, the size of the block corresponding to $\phi ; B_{\phi}$, the set of blocks in $B_{n}$ containing states in $\phi ; B_{n}$, the size of the partition after $n$ iterations:
$O\left(\sum_{\phi \in B_{n-1}}\left(\phi^{T} \phi\right)\left|B_{\phi}\right||A|+\left|B_{n}\right|^{2}\left|B_{n-1}\right|^{2} \log \left|B_{n-1}\right||A|\right)$ The first term of the sum accounts for the construction of $B_{n}$ and the second part of the sum accounts for the computation of the metric $d_{n}$. As $n$ approaches $\infty$, the update algorithm runs in $O\left(|A|\left(|S|\left|B_{\sim}\right|+\left|B_{\sim}\right|^{4} \log \left|B_{\sim}\right|\right)\right)$, which is an upper bound for the update at any step. See Algorithm 3 for a simple asynchronous version of Algorithm 2 that attempts to maintain the computational cost away from the latter upper bound. Later in the paper we will discuss heuristics for choosing the update classes $B \subset B_{n}$ in a way that maintains desirable convergence properties, similar to the asynchronous approach proposed in Comanici et al. (2012).

## Value function approximation

In this section we will analyze approximation errors when using the derived partitions to represent the value function (or the model) as opposed to using the original state space representations. Just as the optimal Bellman operator (Equation 2), it is not hard to prove (Ferns, Panangaden, and Precup 2004) that the bisimulation metric operator is also a contraction mapping and that

$$
\left\|d^{*}-F^{n}(0)\right\|_{\infty}<\gamma^{n} \max _{a}\left\|\mathcal{R}^{a}\right\|_{\infty}
$$

This result can be used to derive a bound for the approximation error induced when using intermediate partitions $B_{n}$. This bound is dependent on $M:=\max _{a}\left\|\mathcal{R}^{a}\right\|_{\infty}$,

$$
\left|V_{B_{n}}^{*}(s)-V^{*}(s)\right| \leq \frac{\gamma^{n} M}{1-\gamma}+\frac{\gamma^{n+1} M}{(1-\gamma)^{2}}=\frac{\gamma^{n} M}{(1-\gamma)^{2}}
$$

which is a special case of Equation 3. As far as nonoptimal policies are concerned, one can provide approximation bounds on a set of polices representable by $B_{n}$. That is, if $\pi^{a}(\cdot):=\pi(\cdot, a) \in \mathcal{F}_{S}$ has the property that $\operatorname{span}\left\{\pi^{a}: a \in A\right\} \subset \operatorname{colspan}\left(B_{n}\right)$, then one can easily verify in Equation 1 that $\Pi_{B_{n}} \mathcal{R}^{\pi}=\mathcal{R}^{\pi}$ and $\Pi_{B_{n}} \mathcal{P}^{\pi} f=$ $\mathcal{P}^{\pi} f, \forall f \in \operatorname{colspan}\left(B_{n-1}\right)$. Note that the set of all openloop policies satisfy this property. Using again the fact that
the Bellman operator is a contraction mapping,

$$
\left\|V_{B_{n}}^{\pi}-V^{\pi}\right\|_{\infty}=\left\|\left(T_{B_{n}}^{\pi}\right)^{n}(0)-V^{\pi}\right\|_{\infty} \leq \gamma^{n} \max _{a}\left\|\mathcal{R}^{a}\right\|_{\infty}
$$

## Asynchronous partitioning

In this section, we present an asynchronous partition algorithm. Just as in the previous section, we generate a sequence of partitions and metric over these, with the property that the corresponding metrics can be transformed to a sequence of metrics converging to the desired Kantorovich-based fixed point bisimulation metric.

Let $\widehat{d^{*}}=\sup _{n} B^{T} d_{n} B$, where $d_{n}$ are metrics over a sequence of partitions $B_{n}$, define inductively as follows:

- $B_{1}$ is a partition such that $\forall a, \mathcal{R}^{a} \in \operatorname{colspan}\left(B_{1}\right)$.
- $B_{n}$ is a partition that is built based on two (chosen) elements $\phi_{n, 1}$ and $\phi_{n, 2}$ in $B_{n-1}$. The partition $B_{n}$ should satisfy the property that $\phi(s)=\phi^{\prime}(s), \forall \phi \in B_{n}$ iff

$$
\begin{aligned}
& -\phi(s)=\phi\left(s^{\prime}\right) \forall \phi \in B_{n-1} \text { and } \\
& - \text { if } \phi_{i, 1}(s)+\phi_{i, 1}\left(s^{\prime}\right)+\phi_{i, 2}(s)+\phi_{i, 2}\left(s^{\prime}\right)=2 \\
& \text { then } \forall a \in A, \forall \phi \in B_{n-1}, \quad\left(\mathcal{P}^{a} \phi\right)(s)=\left(\mathcal{P}^{a} \phi\right)\left(s^{\prime}\right)
\end{aligned}
$$

- Define the metric $d_{n}$ based on the Kantorovich metric over partition $B_{n-1}$ with ground metric $d_{n-1}$. For $\phi, \phi^{\prime} \in$ $B_{n}$ and $s, s^{\prime}$ such that $\phi(s)=1, \phi^{\prime}\left(s^{\prime}\right)=1$, we define $d_{n}\left(\phi, \phi^{\prime}\right)$ as
if $\phi_{n, 1}(s)+\phi_{n, 2}(s)+\phi_{n, 1}\left(s^{\prime}\right)+\phi_{n, 2}\left(s^{\prime}\right)=2$, $\max _{a}\left(\left|\mathcal{R}^{a}(s)-\mathcal{R}^{a}\left(s^{\prime}\right)\right|+\gamma \mathcal{T}_{d_{n-1}}\left(\delta_{s}^{T} \mathcal{P}^{a} B^{T}, \delta_{s^{\prime}}^{T} \mathcal{P}^{a} B^{T}\right)\right)$ else $d_{n-1}\left(\widehat{\phi}, \widehat{\phi}^{\prime}\right)$ for $\widehat{\phi}, \widehat{\phi}^{\prime} \in B_{n-1}$ with $\widehat{\phi}(s)=1, \widehat{\phi}^{\prime}\left(s^{\prime}\right)=1$
Note that this metric is well defined: the choice of states in $s, s^{\prime}$ with $\phi(s)=1$ and $\phi\left(s^{\prime}\right)$ will not make any difference in either cases. This is guaranteed in the construction of the partition $B_{n}$, where we make sure that states for which the distance gets updated using the Kantorovich metric are mapped the same way in $B_{n-1}$ only if transitions to elements of $B_{n-1}$ are the same.

Theorem 2. (Comanici, Panangaden, and Precup 2012) Given a sequence $\left\{K_{i}\right\}_{i=1}^{\infty}$ of subsets of $S \times S$ such that each pair from $S$ is represented infinitely often (i.e., $\bigcap_{n=1}^{\infty} \bigcup_{i=n}^{\infty} K_{i}=S \times S$ ), and $h_{n}$ is defined inductively as $h_{n}\left(s, s^{\prime}\right)=\left\{\begin{array}{ll}h_{n-1}\left(s, s^{\prime}\right) & \text { if }\left(s, s^{\prime}\right) \notin K_{n} \\ F\left(h_{n-1}\right)\left(s, s^{\prime}\right) & \text { if }\left(s, s^{\prime}\right) \in K_{n}\end{array}\right.$,
then $\sup _{n} h_{n}=d^{*}$.
Theorem 3. Let $\left\{B_{n}, \phi_{n, 1}, \phi_{n, 2}, d_{n}\right\}_{i=1}^{\infty}$ be a sequence of partitions and metrics generated under the strategy described above. Iffor every pair $s, s^{\prime} \in S$,
$\sum_{i=1}^{n} \max \left(0, \phi_{i, 1}(s)+\phi_{i, 2}(s)+\phi_{i, 1}\left(s^{\prime}\right)+\phi_{i, 2}\left(s^{\prime}\right)-1\right) \rightarrow \infty$
then $d^{*}=\sup _{n} B_{n}^{T} d^{n} B_{n}$.
Proof. We will use Theorem 2 in Comanici et al. (2012). For this, we will show that $h_{n}:=B_{n}^{T} d_{n} B_{n}$ can be computed using the algorithm in the paper cited. That is, at each time step, let $K_{n}$ be the subset of pairs for which


Figure 1: Asynchronous computation: A plot of the approximation error in the value function computation $\left(L_{\infty}\right.$ norm) as the size of the alternative representation is increased. This particular plot was generated on a Puddle World of size 4900.

```
Algorithm 3 Asynchronous declustering
    \(B_{1} \leftarrow\) partition_declust of \(\{S\}\) based on reward
    while we seek a better approximation: \(i=1 \cdots \infty\) do
        Use a heuristic to select subset \(B \in B_{i}\). Set \(B_{i+1} \leftarrow\)
        \(B_{i} \backslash B\)
        for all \(\phi \in B\) do
            add partition_declust of \(\phi\) using transition
            \(P\) over \(B_{i}\) to \(B_{i+1}\)
        end for
    end while
```

$\sum_{j=1,2} \phi_{i, j}(s)+\sum_{j=1,2}=2$.
If $\left(s, s^{\prime}\right) \in(S \times S) \backslash K$, by definition

$$
\begin{aligned}
h_{n+1}\left(s, s^{\prime}\right) & :=\left(B_{n+1}^{T} d_{n+1} B_{n+1}\right)\left(s, s^{\prime}\right) \\
& =\left(B_{n}^{T} d_{n} B_{n}\right)\left(s, s^{\prime}\right)=h_{n}\left(s, s^{\prime}\right)
\end{aligned}
$$

Now, for $\left(s, s^{\prime}\right) \in K$, by definition

$$
\begin{aligned}
& h_{n+1}\left(s, s^{\prime}\right):=\left(B_{n+1}^{T} d_{n+1} B_{n+1}\right)\left(s, s^{\prime}\right) \\
& =\max _{a}\left(\left|\mathcal{R}^{a}(s)-\mathcal{R}^{a}\left(s^{\prime}\right)\right|+\gamma \mathcal{T}_{d_{n}}\left(\delta_{s}^{T} \mathcal{P}^{a} B^{T}, \delta_{s^{\prime}}^{T} \mathcal{P}^{a} B^{T}\right)\right) \\
& =\max _{a}\left(\left|\mathcal{R}^{a}(s)-\mathcal{R}^{a}\left(s^{\prime}\right)\right|+\gamma \mathcal{T}_{B_{n}^{T}} d_{n} B_{n}\left(\delta_{s}^{T} \mathcal{P}^{a}, \delta_{s^{\prime}}^{T} \mathcal{P}^{a}\right)\right) \\
& =\max _{a}\left(\left|\mathcal{R}^{a}(s)-\mathcal{R}^{a}\left(s^{\prime}\right)\right|+\gamma \mathcal{T}_{h_{n}}\left(\delta_{s}^{T} \mathcal{P}^{a}, \delta_{s^{\prime}}^{T} \mathcal{P}^{a}\right)\right) \\
& =F\left(h_{n}\right)\left(s, s^{\prime}\right)
\end{aligned}
$$

Note also that

$$
\left(s, s^{\prime}\right) \in K \text { iff } \sum_{j=1,2} \phi_{i, j}(s)+\sum_{j=1,2} \phi_{i, j}\left(s^{\prime}\right)-1>0
$$

Therefore, we are guaranteed that every pair of states is selected infinitely often. The requirements of Theorem 2 are satisfied, so the algorithm in this section will compute the equivalent of an asynchronous bisimulation metric computation.


Figure 2: Puddle World - computing the metric. Left: A plot of the runtime as a function of state space size when computing the metric. We performed a comparison to previous work, where metrics are computed over the state space: the runtime jumps from 129 seconds on a 400 states environment, to 1375 seconds on 1600 states. Right: The number of features in the intermediate steps of the algorithm. We ran the algorithm for state spaces as large as 19600 states, but the number of features did not change substantially for sizes larger than 6400 states.

## Empirical illustration

One implication of the theoretical results presented in the previous sections is that the computational complexity of computing bisimulation based representations and corresponding metrics is mostly dependent on the intrinsic complexity of the reward function and transition models. To illustrate this, we computed bisimulation metrics using the procedure presented in Algorithm 2 over a series of MDPs that increases in size, but whose structure remains the same. That is, we vary the number of states in the original MDP, but not the main task, which is to find the path from any point to a given corner by avoiding the negative rewards obtained when navigating through a puddle in the middle of the environment. More states result in a finer discretization of the original continuous puddle world problem. The well known Puddle World problem has a state space consisting of a grid (of varying size for the purpose of illustration); the actions available are labeled by the 4 main compass directions, and they achieve movement in the corresponding direction with probability 0.85 , keep the state unchanged with probability 0.05 , and move in a random direction with probability 0.1 . Note that the grid has margins, which bounce back transitions that would take the agent outside the grid. The reward is dependent on the position of the state compared to a given puddle (Boyan and Moore 1995).

The left panel of Figure 2 shows the runtime of computing bisimulation metrics over partitions, for up to 14 iterations of Algorithm 2. Computation time stays roughly the same once the state space exceeds 4900 states. This is because the feature representation stays relatively unchanged in size and the metric computation is performed over the set of features and not over the entire state space. The right panel of Figure 2 shows the number of features obtained when computing bisimulation metrics using the procedure described in this paper. The key finding is that after a point, the number of features found is roughly constant even as the state space increases. This is because the complexity of the reward function and the transition system remains unchanged, and in particular, this domain is fairly simple.

To illustrate the importance of the asynchronous partition/metric update, we fixed the size of the Puddle World and we compared the value function approximation error as a function of the size of intermediate representations. In this case, for each partition we performed dynamic programming to compute the value function. In the asynchronous algorithm, we used a heuristic which selects first the largest block in order to update the partition/metric. This comes from the intuition that it would be advantageous to seek a representation that is as uniform as possible, instead of having a mix of large and small partitions. As can be seen in Figure 1, the asynchronous algorithm obtains representations of better quality in much earlier stages of the iterative framework. Note that we did not list the running time, as the time spent on obtaining such representations is relatively the same. As already discussed, runtime complexity is mostly dependent on the desired representation size.

## Conclusion and future work

We presented two new ways of describing bisimulation metrics from a theoretical perspective, and we used these to design novel iterative refinement algorithms. These algorithms provide substantial improvement in terms of time and memory usage, and more flexibility in terms of guiding the search for alternative state space representations for MDPs. As illustrated, the methods we propose are not nearly as sensitive to the size of the state space of the problem.

The approach presented in this paper opens the door to more specialized strategies to finding bisimulation-based MDP representations. We illustrated the advantage of using heuristic based search strategies, but the strategy we used (which attempts to keep the size of state partitions roughly the same) is very simple, and it is likely that more sophisticated approaches would work better. For example, one could try strategies similar to prioritized sweeping, which focus on areas of the state space where the metric is changing drastically. Investigating more sophisticated heuristics and applying them to larger problems is a worthwhile direction for future work.

## References

Bertsekas, D. P., and Castanon, D. A. 1989. Adaptive aggregation methods for infinite horizon dynamic programming. IEEE Transactions on Automatic Control 34.
Bertsekas, D. P., and Tsitsiklis, J. N. 1996. Neuro-Dynamic Programming. Athena Scientific, Bellman, MA.
Boyan, J. A., and Moore, A. W. 1995. Generalization in reinforcement learning: Safely approximating the value function. In Advances in Neural Information Processing Systems (NIPS) 7, 369-376. MIT Press.
Bradtke, S. J., and Barto, A. G. 1996. Linear LeastSquares algorithms for temporal difference learning. Machine Learning 22(1-3):33-57.
Comanici, G., and Precup, D. 2011. Basis function discovery using spectral clustering and bisimulation metrics. In the 25th AAAI Conference on Artificial Intelligence.
Comanici, G., and Precup, D. 2012. Basis function discovery using spectral clustering and bisimulation metrics. In Lecture Notes in Computer Science, Volume 7113.
Comanici, G.; Panangaden, P.; and Precup, D. 2012. On-the-fly algorithms for bisimulation metrics. In the 9th International Conference on Quantitative Evaluation of SysTems (QEST).
Desharnais, J.; Gupta, V.; Jagadeesan, R.; and Panangaden, P. 1999. Metrics for labeled Markov systems. In Proceedings of CONCUR99, number 1664 in Lecture Notes in Computer Science. Springer-Verlag.
Desharnais, J.; Gupta, V.; Jagadeesan, R.; and Panangaden, P. 2004. A metric for labelled Markov processes. Theoretical Computer Science 318(3):323-354.
Ferns, N., and Precup, D. 2014. Bisimulation metrics are optimal value functions. In the 30th Conference on Uncertainty in Artificial Intelligence.
Ferns, N.; Panangaden, P.; and Precup, D. 2004. Metrics for finite Markov decision precesses. In the 20th Conference on Uncertainty in Artificial Intelligence, 162-169.
Givan, R.; Dean, T.; and Greig, M. 2003. Equivalence notions and model minimization in Markov decision processes. Artificial Intelligence 147(1-2):163-223.
Larsen, K. G., and Skou, A. 1991. Bisimulation through probabilistic testing. Information and Computation 94:128.

Moore, A., and Atkeson, C. 1993. Prioritized sweeping: Reinforcement learning with less data and less real time. Machine Learning 13:103-130.
Parr, R.; Li, L.; Taylor, G.; Painter-Wakefield, C.; and Littman, M. L. 2008a. An analysis of linear models, linear value-function approximation, and feature selection for reinforcement learning. In International Conference on Machine Learning (ICML), 752-759.
Parr, R.; Painter-Wakefiled, H.; Li, L.; and Littman, M. L. 2008b. Analyzing feature generation for value function approximation. In International Conference on Machine Learning (ICML), 737-744.

Ravindran, B., and Barto, A. G. 2002. Model minimization in hierarchical reinforcement learning. In the 5th International Symposium on Abstraction, Reformulation and Approximation (SARA), 196-211.
Sutton, R. S. 1988. Learning to Predict by the Methods of Temporal Differences. Machine Learning 3(1):9-44.
Yu, H., and Bertsekas, D. 2006. Convergence results for some temporal difference methods based on least squares. Technical report, Laboratory for Information and Decision Systems, Massachusetts Institute of Technology.


[^0]:    Copyright (C) 2015, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

[^1]:    ${ }^{1} \mathrm{~A}$ complete lattice is a partially ordered set for which every subset has a greatest lower bound and a least upper bound. This is a crucial property in proving convergence to a fixed point of a monotone operator over such a space.
    ${ }^{2}$ A representation $\phi(s) \in\{0,1\}$ determines a partition if $\sum_{i} \phi_{i}(s)=1, \forall s$. The classes of the partition are determined by the sets $\{\{s \mid \phi(s)=1\} \mid \phi \in B\}$

