

Practical Kernel-Based Reinforcement Learning

André M. S. Barreto

*Laboratório Nacional de Computação Científica
Petrópolis, Brazil*

AMSB@LNCC.BR

Doina Precup

Joelle Pineau

*School of Computer Science
McGill University
Montreal, Canada*

DPRECUP@CS.MCGILL.CA

JPINEAU@CS.MCGILL.CA

Editor: George Konidaris

Abstract

Kernel-based reinforcement learning (KBRL) stands out among approximate reinforcement learning algorithms for its strong theoretical guarantees. By casting the learning problem as a local kernel approximation, KBRL provides a way of computing a decision policy which converges to a unique solution and is statistically consistent. Unfortunately, the model constructed by KBRL grows with the number of sample transitions, resulting in a computational cost that precludes its application to large-scale or on-line domains. In this paper we introduce an algorithm that turns KBRL into a practical reinforcement learning tool. *Kernel-based stochastic factorization* (KBSF) builds on a simple idea: when a transition probability matrix is represented as the product of two stochastic matrices, one can swap the factors of the multiplication to obtain another transition matrix, potentially much smaller than the original, which retains some fundamental properties of its precursor. KBSF exploits such an insight to compress the information contained in KBRL's model into an approximator of fixed size. This makes it possible to build an approximation considering both the difficulty of the problem and the associated computational cost. KBSF's computational complexity is linear in the number of sample transitions, which is the best one can do without discarding data. Moreover, the algorithm's simple mechanics allow for a fully incremental implementation that makes the amount of memory used independent of the number of sample transitions. The result is a kernel-based reinforcement learning algorithm that can be applied to large-scale problems in both off-line and on-line regimes. We derive upper bounds for the distance between the value functions computed by KBRL and KBSF using the same data. We also prove that it is possible to control the magnitude of the variables appearing in our bounds, which means that, given enough computational resources, we can make KBSF's value function as close as desired to the value function that would be computed by KBRL using the same set of sample transitions. The potential of our algorithm is demonstrated in an extensive empirical study in which KBSF is applied to difficult tasks based on real-world data. Not only does KBSF solve problems that had never been solved before, but it also significantly outperforms other state-of-the-art reinforcement learning algorithms on the tasks studied.

Keywords: reinforcement learning, dynamic programming, Markov decision processes, kernel-based approximation, stochastic factorization

1. Introduction

Reinforcement learning provides a conceptual framework with the potential to materialize a long-sought goal in artificial intelligence: the construction of situated agents that learn how to behave from direct interaction with the environment (Sutton and Barto, 1998). But such an endeavor does not come without its challenges; among them, extrapolating the field’s basic machinery to large-scale domains has been a particularly persistent obstacle.

It has long been recognized that virtually any real-world application of reinforcement learning must involve some form of approximation. Given the mature stage of the supervised-learning theory, and considering the multitude of approximation techniques available today, this realization may not come across as a particularly worrisome issue at first glance. However, it is well known that the sequential nature of the reinforcement learning problem renders the incorporation of function approximators non-trivial (Bertsekas and Tsitsiklis, 1996).

Despite the difficulties, in the last two decades the collective effort of the reinforcement learning community has given rise to many reliable approximate algorithms (Szepesvári, 2010). Among them, Ormonet and Sen’s (2002) kernel-based reinforcement learning (KBRL) stands out for two reasons. First, unlike other approximation schemes, KBRL always converges to a unique solution. Second, KBRL is consistent in the statistical sense, meaning that adding more data improves the quality of the resulting policy and eventually leads to optimal performance.

Unfortunately, the good theoretical properties of KBRL come at a price: since the model constructed by this algorithm grows with the number of sample transitions, the cost of computing a decision policy quickly becomes prohibitive as more data become available. Such a computational burden severely limits the applicability of KBRL. This may help explain why, in spite of its nice theoretical guarantees, kernel-based learning has not been widely adopted as a practical reinforcement learning tool.

This paper presents an algorithm that can potentially change this situation. *Kernel-based stochastic factorization* (KBSF) builds on a simple idea: when a transition probability matrix is represented as the product of two stochastic matrices, one can swap the factors of the multiplication to obtain another transition matrix, potentially much smaller than the original, which retains some fundamental properties of its precursor (Barreto and Fragoso, 2011). KBSF exploits this insight to compress the information contained in KBRL’s model into an approximator of fixed size. Specifically, KBSF builds a model, whose size is independent of the number of sample transitions, which serves as an approximation of the model that would be constructed by KBRL. Since the size of the model becomes a parameter of the algorithm, KBSF essentially detaches the structure of KBRL’s approximator from its configuration. This extra flexibility makes it possible to build an approximation that takes into account *both* the difficulty of the problem *and* the computational cost of finding a policy using the constructed model.

KBSF’s computational complexity is linear in the number of sample transitions, which is the best one can do without throwing data away. Moreover, we show in the paper that the amount of memory used by our algorithm is independent of the number of sample transitions. Put together, these two properties make it possible to apply KBSF to large-scale problems in both off-line and on-line regimes. To illustrate this possibility in practice,

we present an extensive empirical study in which KBSF is applied to difficult control tasks based on real-world data, some of which had never been solved before. KBSF outperforms least-squares policy iteration and fitted Q -iteration on several off-line problems and SARSA on a difficult on-line task.

We also show that KBSF is a sound algorithm from a theoretical point of view. Specifically, we derive results bounding the distance between the value function computed by our algorithm and the one computed by KBRL using the same data. We also prove that it is possible to control the magnitude of the variables appearing in our bounds, which means that we can make the difference between KBSF’s and KBRL’s solutions arbitrarily small.

We start the paper presenting some background material in Section 2. Then, in Section 3, we introduce the stochastic-factorization trick, the insight underlying the development of our algorithm. KBSF itself is presented in Section 4. This section is divided in two parts, one theoretical and one practical. In Section 4.1 we present theoretical results showing that not only is the difference between KBSF’s and KBRL’s value functions bounded, but it can also be controlled. Section 4.2 brings experiments with KBSF on four reinforcement-learning problems: single and double pole-balancing, HIV drug schedule domain, and epilepsy suppression task. In Section 5 we introduce the incremental version of our algorithm, which can be applied to on-line problems. This section follows the same structure of Section 4, with theoretical results followed by experiments. Specifically, in Section 5.1 we extend the results of Section 4.1 to the on-line scenario, and in Section 5.2 we present experiments on the triple pole-balancing and helicopter tasks. In Section 6 we take a closer look at the approximation computed by KBSF and present a practical guide on how to configure our algorithm to solve a reinforcement learning problem. In Section 7 we summarize related works and situate KBSF in the context of kernel-based learning. Finally, in Section 8 we present the main conclusions regarding the current research and discuss some possibilities of future work.

The paper has three appendices. Appendix A has the proofs of our theoretical results. The details of the experiments that were omitted in the main body of the text are described in Appendix B. In Appendix C we provide a table with the main symbols used in the paper that can be used as a reference to facilitate reading.

Parts of the material presented in this article have appeared before in two papers published in the *Neural Information Processing Systems* conference (NIPS, Barreto et al., 2011, 2012). The current manuscript is a substantial extension of the aforementioned works.

2. Background

We consider the standard framework of reinforcement learning, in which an *agent* interacts with an *environment* and tries to maximize the amount of *reward* collected in the long run (Sutton and Barto, 1998). The interaction between agent and environment happens at discrete time steps: at each instant t the agent occupies a state $s_{(t)} \in S$ and must choose an action a from a finite set A . The sets S and A are called the *state* and *action spaces*, respectively. The execution of action a in state $s_{(t)}$ moves the agent to a new state $s_{(t+1)}$, where a new action must be selected, and so on. Each transition has a certain probability of occurrence and is associated with a reward $r \in \mathbb{R}$. The goal of the agent is to find a *policy* $\pi : S \mapsto A$, that is, a mapping from states to actions, that maximizes the expected

return. Here we define the *return* from time t as:

$$R_{(t)} = r_{(t+1)} + \gamma r_{(t+2)} + \gamma^2 r_{(t+3)} + \dots = \sum_{i=1}^T \gamma^{i-1} r_{(t+i)}, \quad (1)$$

where $r_{(t+1)}$ is the reward received at the transition from state $s_{(t)}$ to state $s_{(t+1)}$. The interaction of the agent with the environment may last forever ($T = \infty$) or until the agent reaches a terminal state ($T < \infty$); each sequence of interactions is usually referred to as an *episode*. The parameter $\gamma \in [0, 1)$ is the *discount factor*, which determines the relative importance of individual rewards depending on how far in the future they are received.

2.1 Markov Decision Processes

As usual, we assume that the interaction between agent and environment can be modeled as a *Markov decision process* (MDP, Puterman, 1994). An MDP is a tuple $M \equiv (S, A, P^a, R^a, \gamma)$, where P^a and R^a describe the dynamics of the task at hand. For each action $a \in A$, $P^a(\cdot|s)$ defines the next-state distribution upon taking action a in state s . The reward received at transition $s \xrightarrow{a} s'$ is given by $R^a(s, s')$, with $|R^a(s, s')| \leq R_{\max} < \infty$. Usually, one is interested in the expected reward resulting from the execution of action a in state s , that is, $r^a(s) = E_{s' \sim P^a(\cdot|s)}\{R^a(s, s')\}$.

Once the interaction between agent and environment has been modeled as an MDP, a natural way of searching for an optimal policy is to resort to *dynamic programming* (Bellman, 1957). Central to the theory of dynamic-programming is the concept of a *value function*. The value of state s under a policy π , denoted by $V^\pi(s)$, is the expected return the agent will receive from s when following π , that is, $V^\pi(s) = E_\pi\{R_{(t)}|s_{(t)} = s\}$ (here the expectation is over all possible sequences of rewards in (1) when the agent follows π). Similarly, the value of the state-action pair (s, a) under policy π is defined as $Q^\pi(s, a) = E_{s' \sim P^a(\cdot|s)}\{R^a(s, s') + \gamma V^\pi(s')\} = r^a(s) + \gamma E_{s' \sim P^a(\cdot|s)}\{V^\pi(s')\}$.

The notion of value function makes it possible to impose a partial ordering over decision policies. In particular, a policy π' is considered to be at least as good as another policy π if $V^{\pi'}(s) \geq V^\pi(s)$ for all $s \in S$. The goal of dynamic programming is to find an optimal policy π^* that performs no worse than any other. It is well known that there always exists at least one such policy for a given MDP (Puterman, 1994). When there is more than one optimal policy, they all share the same value function V^* .

When both the state and action spaces are finite, an MDP can be represented in matrix form: each function P^a becomes a matrix $\mathbf{P}^a \in \mathbb{R}^{|S| \times |S|}$, with $p_{ij}^a = P^a(s_j|s_i)$, and each function r^a becomes a vector $\mathbf{r}^a \in \mathbb{R}^{|S|}$, where $r_i^a = r^a(s_i)$. Similarly, V^π can be represented as a vector $\mathbf{v}^\pi \in \mathbb{R}^{|S|}$ and Q^π can be seen as a matrix $\mathbf{Q}^\pi \in \mathbb{R}^{|S| \times |A|}$.¹

When the MDP is finite, dynamic programming can be used to find an optimal decision-policy $\pi^* \in A^{|S|}$ in time polynomial in the number of states $|S|$ and actions $|A|$ (Ye, 2011).

1. Throughout the paper we will use the conventional and matrix notations interchangeably, depending on the context. When using the latter, vectors will be denoted by small boldface letters and matrices will be denoted by capital boldface letters. We will also use the same notation for all MDPs and associated components, distinguishing between them through the use of math accents. So, for example, if \bar{M} is an MDP, its transition functions and matrices will be referred to as \bar{P}^a and $\bar{\mathbf{P}}^a$, its expect-reward functions and vectors will be denoted by \bar{r}^a and $\bar{\mathbf{r}}^a$, its optimal decision policy will be $\bar{\pi}^*$, and so on (see Table 2).

Let $\mathbf{v} \in \mathbb{R}^{|S|}$ and let $\mathbf{Q} \in \mathbb{R}^{|S| \times |A|}$. Define the operator $\Gamma : \mathbb{R}^{|S| \times |A|} \mapsto \mathbb{R}^{|S|}$ such that $\Gamma \mathbf{Q} = \mathbf{v}$ if and only if $v_i = \max_j q_{ij}$ for all i . Also, given an MDP M , define $\Delta : \mathbb{R}^{|S|} \mapsto \mathbb{R}^{|S| \times |A|}$ such that $\Delta \mathbf{v} = \mathbf{Q}$ if and only if $q_{ia} = r_i^a + \gamma \sum_{j=1}^{|S|} p_{ij}^a v_j$ for all i and all a . The *Bellman operator* of the MDP M is given by $T \equiv \Gamma \Delta$. A fundamental result in dynamic programming states that, starting from $\mathbf{v}^{(0)} = \mathbf{0}$, the expression $\mathbf{v}^{(t)} = T \mathbf{v}^{(t-1)} = \Gamma \mathbf{Q}^{(t)}$ gives the optimal t -step value function, and as $t \rightarrow \infty$ the vector $\mathbf{v}^{(t)}$ approaches \mathbf{v}^* . At any point, the optimal t -step policy can be obtained by selecting $\pi_i^{(t)} \in \operatorname{argmax}_j q_{ij}^{(t)}$ (Puterman, 1994).

In contrast with dynamic programming, in reinforcement learning it is assumed that the MDP is unknown, and the agent must learn a policy based on transitions sampled from the environment. If the process of learning a decision policy is based on a fixed set of sample transitions, we call it *batch* reinforcement learning. On the other hand, in *on-line* reinforcement learning the computation of a decision policy takes place concomitantly with the collection of data (Sutton and Barto, 1998).

2.2 Kernel-Based Reinforcement Learning

Kernel-based reinforcement learning (KBRL) is a batch algorithm that uses a finite model approximation to solve a continuous MDP $M \equiv (\mathbb{S}, A, P^a, R^a, \gamma)$, where $\mathbb{S} \subseteq [0, 1]^{d_{\mathbb{S}}}$ and $d_{\mathbb{S}} \in \mathbb{N}_*^+$ is the dimension of the state space (Ormoneit and Sen, 2002). Let $S^a \equiv \{(s_k^a, r_k^a, \hat{s}_k^a) | k = 1, 2, \dots, n_a\}$ be sample transitions associated with action $a \in A$, where $s_k^a, \hat{s}_k^a \in \mathbb{S}$ and $r_k^a \in \mathbb{R}$. Let $\phi : \mathbb{R}^+ \mapsto \mathbb{R}^+$ be a Lipschitz continuous function satisfying $\int_0^1 \phi(x) dx = 1$. Let $k_{\tau}(s, s')$ be a kernel function defined as

$$k_{\tau}(s, s') = \phi\left(\frac{\|s - s'\|}{\tau}\right), \quad (2)$$

where $\tau \in \mathbb{R}$ and $\|\cdot\|$ is a norm in $\mathbb{R}^{d_{\mathbb{S}}}$ (for concreteness, the reader may think of $k_{\tau}(s, s')$ as the Gaussian kernel, although the definition also encompasses other functions). Finally, define the normalized kernel function associated with action a as

$$\kappa_{\tau}^a(s, s_i^a) = \frac{k_{\tau}(s, s_i^a)}{\sum_{j=1}^{n_a} k_{\tau}(s, s_j^a)}. \quad (3)$$

KBRL uses (3) to build a finite MDP whose state space \hat{S} is composed solely of the $n = \sum_a n_a$ states \hat{s}_i^a . We assume without loss of generality that the action space A is ordered and the sampled states are ordered lexicographically as $\hat{s}_i^a < \hat{s}_j^b \iff a < b$ or $(a = b \text{ and } i < j)$; if a given state $s \in \mathbb{S}$ occurs more than once in the set of sample transitions, each occurrence will be treated as a distinct state in the finite MDP. The transition functions of KBRL's model, $\hat{P}^a : \hat{S} \times \hat{S} \mapsto [0, 1]$, are given by:

$$\hat{P}^a(\hat{s}_i^b | s) = \begin{cases} \kappa_{\tau}^a(s, s_i^b), & \text{if } a = b, \\ 0, & \text{otherwise,} \end{cases} \quad (4)$$

where $a, b \in A$. Similarly, the reward functions of the MDP constructed by KBRL, $\hat{R}^a : \hat{S} \times \hat{S} \mapsto \mathbb{R}$, are

$$\hat{R}^a(s, \hat{s}_i^b) = \begin{cases} r_i^a, & \text{if } a = b, \\ 0, & \text{otherwise.} \end{cases} \quad (5)$$

Based on (4) and (5) we can define the transition matrices and expected-reward vectors of KBRL’s MDP. The matrices $\hat{\mathbf{P}}^a$ are derived directly from the definition of $\hat{P}^a(\hat{s}_i^b|s)$, replacing s with the sampled states \hat{s}_i^a (see Figure 2a). The vectors $\hat{\mathbf{r}}^a$ are computed as follows. Let $\mathbf{r} \equiv [(\mathbf{r}^1)^\top, (\mathbf{r}^2)^\top, \dots, (\mathbf{r}^{|A|})^\top]^\top \in \mathbb{R}^n$, where $\mathbf{r}^a \in \mathbb{R}^{n_a}$ are the vectors composed of the sample rewards, that is, the i^{th} element of \mathbf{r}^a is $r_i^a \in S^a$. Since $R^a(s, \hat{s}_i^b)$ does not depend on the start state s , we can write

$$\hat{\mathbf{r}}^a = \hat{\mathbf{P}}^a \mathbf{r}. \tag{6}$$

KBRL’s MDP is thus given by $\hat{M} \equiv (\hat{S}, A, \hat{\mathbf{P}}^a, \hat{\mathbf{r}}^a, \gamma)$.

Once \hat{M} has been defined, one can use dynamic programming to compute its optimal value function \hat{V}^* . Then, the value of any state-action pair of the continuous MDP can be determined as:

$$\hat{Q}(s, a) = \sum_{i=1}^{n_a} \kappa_\tau^a(s, s_i^a) \left[r_i^a + \gamma \hat{V}^*(\hat{s}_i^a) \right], \tag{7}$$

where $s \in \mathbb{S}$ and $a \in A$. Ormoneit and Sen (2002) have shown that, if $n_a \rightarrow \infty$ for all $a \in A$ and the kernel’s width τ shrink at an “admissible” rate, the probability of choosing a suboptimal action based on $\hat{Q}(s, a)$ converges to zero (see their Theorem 4).

As discussed, using dynamic programming one can compute the optimal value function of \hat{M} in time polynomial in the number of sample transitions n (which is also the number of states in \hat{M}). However, since each application of the Bellman operator \hat{T} is $O(n^2|A|)$, the computational cost of such a procedure can easily become prohibitive in practice. Thus, the use of KBRL leads to a dilemma: on the one hand one wants as much data as possible to describe the dynamics of the task, but on the other hand the number of transitions should be small enough to allow for the numerical solution of the resulting model. In the following sections we describe a practical approach to weigh the relative importance of these two conflicting objectives.

3. Stochastic Factorization

A *stochastic matrix* has only nonnegative elements and each of its rows sums to 1. That said, we can introduce the concept that will serve as a cornerstone for the rest of the paper:

Definition 1 *Given a stochastic matrix $\mathbf{P} \in \mathbb{R}^{n \times p}$, the relation $\mathbf{P} = \mathbf{D}\mathbf{K}$ is called a stochastic factorization of \mathbf{P} if $\mathbf{D} \in \mathbb{R}^{n \times m}$ and $\mathbf{K} \in \mathbb{R}^{m \times p}$ are also stochastic matrices. The integer $m > 0$ is the order of the factorization.*

This mathematical construct has been explored before. For example, Cohen and Rothblum (1991) briefly discuss it as a special case of nonnegative matrix factorization, while Cutler and Breiman (1994) study slightly modified versions of stochastic factorization for statistical data analysis. However, in this paper we will focus on a useful property of this type of factorization that has only recently been noted (Barreto and Fragoso, 2011).

3.1 Stochastic-Factorization Trick

Let $\mathbf{P} \in \mathbb{R}^{n \times n}$ be a transition matrix, that is, a square stochastic matrix, and let $\mathbf{P} = \mathbf{D}\mathbf{K}$ be an order m stochastic factorization. In this case, one can see the elements of




Figure 1: Reducing the dimension of a transition model from $n = 3$ states to $m = 2$ artificial states. Original states s_i are represented as big white circles; small black circles depict artificial states \bar{s}_h . The symbol ‘ \times ’ is used to represent nonzero elements. These figures have appeared before in the article by Barreto and Fragoso (2011).

\mathbf{D} and \mathbf{K} as probabilities of transitions between the states s_i and a set of m artificial states \bar{s}_h . Specifically, the elements in each row of \mathbf{D} can be interpreted as probabilities of transitions from the original states to the artificial states, while the rows of \mathbf{K} can be seen as probabilities of transitions in the opposite direction. Under this interpretation, each element $p_{ij} = \sum_{h=1}^m d_{ih}k_{hj}$ is the sum of the probabilities associated with m two-step transitions: from state s_i to each artificial state \bar{s}_h and from these back to state s_j . In other words, p_{ij} is the accumulated probability of all possible paths from s_i to s_j with a stopover in one of the artificial states \bar{s}_h . Following similar reasoning, it is not difficult to see that by *swapping* the factors of a stochastic factorization, that is, by switching from \mathbf{DK} to \mathbf{KD} , one obtains the transition probabilities between the artificial states \bar{s}_h , $\bar{\mathbf{P}} = \mathbf{KD}$. If $m < n$, $\bar{\mathbf{P}} \in \mathbb{R}^{m \times m}$ will be a compact version of \mathbf{P} . Figure 1 illustrates this idea for the case in which $n = 3$ and $m = 2$.

The stochasticity of $\bar{\mathbf{P}}$ follows immediately from the same property of \mathbf{D} and \mathbf{K} . What is perhaps more surprising is the fact that this matrix shares some fundamental characteristics with the original matrix \mathbf{P} . Specifically, it is possible to show that: (i) for each recurrent class in \mathbf{P} there is a corresponding class in $\bar{\mathbf{P}}$ with the same period and, given some simple assumptions about the factorization, (ii) \mathbf{P} is irreducible if and only if $\bar{\mathbf{P}}$ is irreducible and (iii) \mathbf{P} is regular if and only if $\bar{\mathbf{P}}$ is regular (for details, see the article by Barreto and Fragoso, 2011). We will refer to this insight as the “*stochastic-factorization trick*”:

Given a stochastic factorization of a transition matrix, $\mathbf{P} = \mathbf{DK}$, swapping the factors of the factorization yields another transition matrix $\bar{\mathbf{P}} = \mathbf{KD}$, potentially much smaller than the original, which retains the basic topology and properties of \mathbf{P} .

Given the strong connection between $\mathbf{P} \in \mathbb{R}^{n \times n}$ and $\bar{\mathbf{P}} \in \mathbb{R}^{m \times m}$, the idea of replacing the former by the latter comes almost inevitably. The motivation for this would be, of course, to save computational resources when $m < n$. For example, Barreto and Fragoso

(2011) have shown that it is possible to recover the stationary distribution of \mathbf{P} through a linear transformation of the corresponding distribution of $\bar{\mathbf{P}}$. In this paper we will use the stochastic-factorization trick to reduce the computational cost of KBRL. The strategy will be to summarize the information contained in KBRL’s MDP in a model of fixed size.

3.2 Reducing a Markov Decision Process

The idea of using stochastic factorization to reduce dynamic programming’s computational requirements is straightforward: given factorizations of the transition matrices \mathbf{P}^a , we can apply our trick to obtain a reduced MDP that will be solved in place of the original one. In the most general scenario, we would have one independent factorization $\mathbf{P}^a = \mathbf{D}^a \mathbf{K}^a$ for each action $a \in A$ and then use $\bar{\mathbf{P}}^a = \mathbf{K}^a \mathbf{D}^a$ instead of \mathbf{P}^a . However, in the current work we will focus on the particular case in which there is a single matrix \mathbf{D} , which will prove to be convenient both mathematically and computationally.

Obviously, in order to apply the stochastic-factorization trick to an MDP, we have to first *compute* the matrices involved in the factorization. Unfortunately, such a procedure can be computationally demanding, exceeding the number of operations necessary to calculate \mathbf{v}^* (Vavasis, 2009; Barreto et al., 2014). Thus, in practice we may have to replace the exact factorizations $\mathbf{P}^a = \mathbf{D} \mathbf{K}^a$ with approximations $\mathbf{P}^a \approx \mathbf{D} \mathbf{K}^a$. The following proposition bounds the error in the value-function approximation resulting from the application of our trick to approximate stochastic factorizations:

Proposition 2 *Let $M \equiv (S, A, \mathbf{P}^a, \mathbf{r}^a, \gamma)$ be a finite MDP with $|S| = n$ and $0 \leq \gamma < 1$. Let $\mathbf{D} \in \mathbb{R}^{n \times m}$ be a stochastic matrix and, for each $a \in A$, let $\mathbf{K}^a \in \mathbb{R}^{m \times n}$ be stochastic and let $\bar{\mathbf{r}}^a$ be a vector in \mathbb{R}^m . Define the MDP $\bar{M} \equiv (\bar{S}, A, \bar{\mathbf{P}}^a, \bar{\mathbf{r}}^a, \gamma)$, with $|\bar{S}| = m$ and $\bar{\mathbf{P}}^a = \mathbf{K}^a \mathbf{D}$. Then,*

$$\|\mathbf{v}^* - \Gamma \mathbf{D} \bar{\mathbf{Q}}^*\|_\infty \leq \xi_v \equiv \frac{1}{1-\gamma} \max_a \|\mathbf{r}^a - \mathbf{D} \bar{\mathbf{r}}^a\|_\infty + \frac{\bar{R}_{\text{dif}}}{(1-\gamma)^2} \left(\frac{\gamma}{2} \max_a \|\mathbf{P}^a - \mathbf{D} \mathbf{K}^a\|_\infty + \sigma(\mathbf{D}) \right), \quad (8)$$

where

$$\sigma(\mathbf{D}) = \max_i (1 - \max_j d_{ij}), \quad (9)$$

$\|\cdot\|_\infty$ is the maximum norm, and $\bar{R}_{\text{dif}} = \max_{a,i} \bar{r}_i^a - \min_{a,i} \bar{r}_i^a$.²

The proofs of most of our theoretical results are in Appendix A.1. We note that Proposition 2 is only valid for the maximum norm; in Appendix A.2 we derive another bound for the distance between \mathbf{v}^* and $\Gamma \mathbf{D} \bar{\mathbf{Q}}^*$ which is valid for any norm.

Our bound depends on two factors: the quality of the MDP’s factorization, given by $\max_a \|\mathbf{P}^a - \mathbf{D} \mathbf{K}^a\|_\infty$ and $\max_a \|\mathbf{r}^a - \mathbf{D} \bar{\mathbf{r}}^a\|_\infty$, and the “level of stochasticity” of \mathbf{D} , measured by $\sigma(\mathbf{D})$. When the MDP factorization is exact, we recover a computable version of Sorg and Singh’s (2009) bound for soft homomorphisms (see (32)). On the other hand, when \mathbf{D} is deterministic—that is, when all its nonzero elements are 1—expression (8) reduces to Whitt’s (1978) classical result regarding state aggregation in dynamic programming. Finally, if we have exact deterministic factorizations, the right-hand side of (8) reduces to

2. We recall that $\|\cdot\|_\infty$ induces the following norm over the space of matrices: $\|\mathbf{A}\|_\infty = \max_i \sum_j |a_{ij}|$.

zero. This also makes sense, since in this case the stochastic-factorization trick gives rise to an exact homomorphism (Ravindran, 2004).

Proposition 2 elucidates the basic mechanism through which one can use the stochastic-factorization trick to reduce the number of states in an MDP (and hence the computational cost of finding a policy using dynamic programming). One possible way to exploit this result is to see the computation of \mathbf{D} , \mathbf{K}^a , and $\bar{\mathbf{r}}^a$ as an optimization problem in which the objective is to minimize some function of $\max_a \|\mathbf{P}^a - \mathbf{D}\mathbf{K}^a\|_\infty$, $\max_a \|\mathbf{r}^a - \mathbf{D}\bar{\mathbf{r}}^a\|_\infty$, and possibly also $\sigma(\mathbf{D})$ (Barreto et al., 2014). Note though that addressing the factorization problem as an optimization may be computationally infeasible when the dimension of the matrices \mathbf{P}^a is large—which is exactly the case we are interested in here. To illustrate this point, we will draw a connection between the stochastic factorization and a popular problem known in the literature as *nonnegative matrix factorization* (Paatero and Tapper, 1994; Lee and Seung, 1999).

In a nonnegative matrix factorization the elements of \mathbf{D} and \mathbf{K}^a are greater or equal to zero, but in general no stochasticity constraint is imposed. Cohen and Rothblum (1991) have shown that it is always possible to derive a stochastic factorization from a nonnegative factorization of a stochastic matrix, which formally characterizes the former as a particular case of the latter. Unfortunately, nonnegative factorization is hard: Vavasis (2009) has shown that a particular version of the problem is in fact NP-hard.

Instead of solving the problem exactly, one can resort instead to an approximate nonnegative matrix factorization. However, since the number of states n in an MDP determines both the number of rows and the number of columns of the matrices \mathbf{P}^a , even the fast “linear” approximate methods run in $O(n^2)$ time, which is infeasible for large n (Barreto et al., 2014). One can circumvent this computational obstacle by exploiting structure in the optimization problem or by resorting to heuristics. In another article on the subject we explore both these alternatives at length (Barreto et al., 2014). However, in this paper we adopt a different approach. Since the model \hat{M} built by KBRL is itself an approximation, instead of insisting in finding a near-optimal factorization for \hat{M} we apply our trick to *avoid* the construction of \mathbf{P}^a and \mathbf{r}^a . As will be seen, this is done by applying KBRL’s own approximation scheme to \hat{M} .

4. Kernel-Based Stochastic Factorization

In Section 2 we presented KBRL, an approximation framework for reinforcement learning whose main drawback is its high computational complexity. In Section 3 we discussed how the stochastic-factorization trick can in principle be useful to reduce an MDP, as long as one circumvents the computational burden imposed by the calculation of the matrices involved in the process. We now show that by combining these two approaches we get an algorithm that overcomes the computational limitations of its components. We call it *kernel-based stochastic factorization*, or KBSF for short.

KBSF emerges from the application of the stochastic-factorization trick to KBRL’s MDP \hat{M} (Barreto et al., 2011). Similarly to Ormoneit and Sen (2002), we start by defining a “mother kernel” $\bar{\phi}(x) : \mathbb{R}^+ \mapsto \mathbb{R}^+$. In Section 4.1 we list our assumptions regarding $\bar{\phi}$. Here, it suffices to note that, since our assumptions and Ormoneit and Sen’s (2002) are not mutually exclusive, we can have $\phi \equiv \bar{\phi}$ (by using the Gaussian function in both

cases, for example). Let $\bar{S} \equiv \{\bar{s}_1, \bar{s}_2, \dots, \bar{s}_m\}$ be a set of *representative states*. Analogously to (2) and (3), we define the kernel $\bar{k}_\tau(s, s') = \bar{\phi}(\|s - s'\|/\bar{\tau})$ and its normalized version $\bar{\kappa}_\tau(s, \bar{s}_i) = \bar{k}_\tau(s, \bar{s}_i) / \sum_{j=1}^m \bar{k}_\tau(s, \bar{s}_j)$. We will use κ_τ^a to build matrices \mathbf{K}^a and $\bar{\kappa}_\tau$ to build matrix \mathbf{D} .

As shown in Figure 2a, KBRL’s matrices $\hat{\mathbf{P}}^a$ have a very specific structure, since only transitions ending in states $\hat{s}_i^a \in S^a$ have a nonzero probability of occurrence. Suppose that we want to apply the stochastic-factorization trick to KBRL’s MDP. Assuming that the matrices \mathbf{K}^a have the same structure as $\hat{\mathbf{P}}^a$, when computing $\hat{\mathbf{P}}^a = \mathbf{K}^a \mathbf{D}$ we only have to look at the sub-matrices of \mathbf{K}^a and \mathbf{D} corresponding to the n_a nonzero columns of \mathbf{K}^a . We call these matrices $\dot{\mathbf{K}}^a \in \mathbb{R}^{m \times n_a}$ and $\dot{\mathbf{D}}^a \in \mathbb{R}^{n_a \times m}$. The strategy of KBSF is to fill out matrices $\dot{\mathbf{K}}^a$ and $\dot{\mathbf{D}}^a$ with elements

$$\dot{k}_{ij}^a = \kappa_\tau^a(\bar{s}_i, s_j^a) \quad \text{and} \quad \dot{d}_{ij}^a = \bar{\kappa}_\tau(\hat{s}_i^a, \bar{s}_j). \quad (10)$$

Note that, based on $\dot{\mathbf{D}}^a$, one can easily recover \mathbf{D} as $\mathbf{D}^\top \equiv [(\dot{\mathbf{D}}^1)^\top, (\dot{\mathbf{D}}^2)^\top, \dots, (\dot{\mathbf{D}}^{|A|})^\top] \in \mathbb{R}^{n \times m}$. Similarly, if we let $\mathbf{K} \equiv [\dot{\mathbf{K}}^1, \dot{\mathbf{K}}^2, \dots, \dot{\mathbf{K}}^{|A|}] \in \mathbb{R}^{m \times n}$, then $\mathbf{K}^a \in \mathbb{R}^{m \times n}$ is matrix \mathbf{K} with all elements replaced by zeros except for those corresponding to matrix $\dot{\mathbf{K}}^a$ (see Figures 2b and 2c for an illustration). It should be thus obvious that $\hat{\mathbf{P}}^a = \mathbf{K}^a \mathbf{D} = \dot{\mathbf{K}}^a \dot{\mathbf{D}}^a$.

In order to conclude the construction of KBSF’s MDP, we have to define the vectors of expected rewards $\bar{\mathbf{r}}^a$. As shown in expression (5), the reward functions of KBRL’s MDP, $\hat{R}^a(s, s')$, only depend on the ending state s' . Recalling the interpretation of the rows of \mathbf{K}^a as transition probabilities from the representative states to the original ones, illustrated in Figure 1, it is clear that

$$\bar{\mathbf{r}}^a = \dot{\mathbf{K}}^a \mathbf{r}^a = \mathbf{K}^a \mathbf{r}. \quad (11)$$

Therefore, the formal specification of KBSF’s MDP is given by $\bar{M} \equiv (\bar{S}, A, \dot{\mathbf{K}}^a \dot{\mathbf{D}}^a, \dot{\mathbf{K}}^a \mathbf{r}^a, \gamma) = (\bar{S}, A, \mathbf{K}^a \mathbf{D}^a, \mathbf{K}^a \mathbf{r}, \gamma) = (\bar{S}, A, \hat{\mathbf{P}}^a, \bar{\mathbf{r}}^a, \gamma)$.

As discussed in Section 2.2, KBRL’s approximation scheme can be interpreted as the derivation of a finite MDP. In this case, the sample transitions define both the finite state space \hat{S} and the model’s transition and reward functions. This means that the state space and the dynamics of KBRL’s model are inexorably linked: except maybe for degenerate cases, changing one also changes the other. By defining a set of representative states, KBSF decouples the MDP’s structure from its particular instantiation. To see why this is so, note that, if we fix the representative states, different sets of sample transitions will give rise to different models. Conversely, the same set of transitions can generate different MDPs, depending on how the representative states are defined.

A step by step description of KBSF is given in Algorithm 1. As one can see, KBSF is very simple to understand and to implement. It works as follows: first, the MDP \bar{M} is built as described above. Then, its action-value function $\bar{\mathbf{Q}}^*$ is determined through any dynamic programming algorithm. Finally, KBSF returns an approximation of $\hat{\mathbf{v}}^*$ —the optimal value function of KBRL’s MDP—computed as $\tilde{\mathbf{v}} = \Gamma \mathbf{D} \bar{\mathbf{Q}}^*$. Based on $\tilde{\mathbf{v}}$, one can compute an approximation of KBRL’s action-value function $\hat{Q}(s, a)$ by simply replacing \hat{V} for \tilde{V}^* in (7), that is,

$$\hat{Q}(s, a) = \sum_{i=1}^{n_a} \kappa_\tau^a(s, s_i^a) \left[r_i^a + \gamma \tilde{V}(\hat{s}_i^a) \right], \quad (12)$$

$$\hat{\mathbf{P}}^a = \begin{array}{c} \hat{s}_1^a \\ \hat{s}_2^a \\ \hat{s}_3^a \\ \hat{s}_1^b \\ \hat{s}_2^b \end{array} \begin{array}{c} \hat{s}_1^a \\ \hat{s}_2^a \\ \hat{s}_3^a \\ \hat{s}_1^a \\ \hat{s}_2^a \end{array} \begin{array}{c} \hat{s}_2^a \\ \hat{s}_2^a \\ \hat{s}_2^a \\ \hat{s}_2^a \\ \hat{s}_2^a \end{array} \begin{array}{c} \hat{s}_3^a \\ \hat{s}_3^a \\ \hat{s}_3^a \\ \hat{s}_3^a \\ \hat{s}_3^a \end{array} \begin{array}{c} \hat{s}_1^b \\ \hat{s}_1^b \\ \hat{s}_1^b \\ \hat{s}_1^b \\ \hat{s}_1^b \end{array} \begin{array}{c} \hat{s}_2^b \\ \hat{s}_2^b \\ \hat{s}_2^b \\ \hat{s}_2^b \\ \hat{s}_2^b \end{array} \left[\begin{array}{cc} \kappa_\tau^a(\hat{s}_1^a, s_1^a) & \kappa_\tau^a(\hat{s}_1^a, s_2^a) \\ \kappa_\tau^a(\hat{s}_2^a, s_1^a) & \kappa_\tau^a(\hat{s}_2^a, s_2^a) \\ \kappa_\tau^a(\hat{s}_3^a, s_1^a) & \kappa_\tau^a(\hat{s}_3^a, s_2^a) \\ \kappa_\tau^a(\hat{s}_1^b, s_1^a) & \kappa_\tau^a(\hat{s}_1^b, s_2^a) \\ \kappa_\tau^a(\hat{s}_2^b, s_1^a) & \kappa_\tau^a(\hat{s}_2^b, s_2^a) \end{array} \right],$$

$$\hat{\mathbf{P}}^b = \begin{array}{c} \hat{s}_1^a \\ \hat{s}_2^a \\ \hat{s}_3^a \\ \hat{s}_1^b \\ \hat{s}_2^b \end{array} \begin{array}{c} \hat{s}_1^a \\ \hat{s}_2^a \\ \hat{s}_3^a \\ \hat{s}_1^b \\ \hat{s}_2^b \end{array} \begin{array}{c} \hat{s}_1^b \\ \hat{s}_2^b \\ \hat{s}_3^b \\ \hat{s}_1^b \\ \hat{s}_2^b \end{array} \left[\begin{array}{cc} 0 & \kappa_\tau^a(\hat{s}_1^a, s_1^b) \\ 0 & \kappa_\tau^a(\hat{s}_2^a, s_1^b) \\ 0 & \kappa_\tau^a(\hat{s}_3^a, s_1^b) \\ 0 & \kappa_\tau^a(\hat{s}_1^b, s_1^b) \\ 0 & \kappa_\tau^a(\hat{s}_2^b, s_1^b) \end{array} \right]$$

(a) KBRL's matrices

$$\mathbf{D} = \begin{array}{c} \hat{s}_1^a \\ \hat{s}_2^a \\ \hat{s}_3^a \\ \hat{s}_1^b \\ \hat{s}_2^b \end{array} \begin{array}{c} \bar{s}_1 \\ \bar{s}_2 \\ \bar{s}_1 \\ \bar{s}_1 \\ \bar{s}_2 \end{array} \begin{array}{c} \bar{s}_2 \\ \bar{s}_2 \\ \bar{s}_2 \\ \bar{s}_2 \\ \bar{s}_2 \end{array} \left[\begin{array}{cc} \bar{\kappa}_\tau(\hat{s}_1^a, \bar{s}_1) & \bar{\kappa}_\tau(\hat{s}_1^a, \bar{s}_2) \\ \bar{\kappa}_\tau(\hat{s}_2^a, \bar{s}_1) & \bar{\kappa}_\tau(\hat{s}_2^a, \bar{s}_2) \\ \bar{\kappa}_\tau(\hat{s}_3^a, \bar{s}_1) & \bar{\kappa}_\tau(\hat{s}_3^a, \bar{s}_2) \\ \bar{\kappa}_\tau(\hat{s}_1^b, \bar{s}_1) & \bar{\kappa}_\tau(\hat{s}_1^b, \bar{s}_2) \\ \bar{\kappa}_\tau(\hat{s}_2^b, \bar{s}_1) & \bar{\kappa}_\tau(\hat{s}_2^b, \bar{s}_2) \end{array} \right],$$

$$\mathbf{K}^a = \begin{array}{c} \bar{s}_1 \\ \bar{s}_2 \end{array} \begin{array}{c} \hat{s}_1^a \\ \hat{s}_2^a \\ \hat{s}_3^a \end{array} \begin{array}{c} \hat{s}_1^a \\ \hat{s}_2^a \\ \hat{s}_3^a \end{array} \begin{array}{c} \hat{s}_1^b \\ \hat{s}_2^b \end{array} \begin{array}{c} \hat{s}_1^b \\ \hat{s}_2^b \end{array} \left[\begin{array}{cc} \kappa_\tau^a(\bar{s}_1, s_1^a) & \kappa_\tau^a(\bar{s}_1, s_2^a) \\ \kappa_\tau^a(\bar{s}_2, s_1^a) & \kappa_\tau^a(\bar{s}_2, s_2^a) \\ \kappa_\tau^a(\bar{s}_2, s_3^a) & \kappa_\tau^a(\bar{s}_2, s_3^a) \\ 0 & 0 \\ 0 & 0 \end{array} \right],$$

$$\mathbf{K}^b = \begin{array}{c} \bar{s}_1 \\ \bar{s}_2 \end{array} \begin{array}{c} \hat{s}_1^a \\ \hat{s}_2^a \\ \hat{s}_3^a \end{array} \begin{array}{c} \hat{s}_1^b \\ \hat{s}_2^b \end{array} \begin{array}{c} \hat{s}_1^b \\ \hat{s}_2^b \end{array} \left[\begin{array}{cc} 0 & \kappa_\tau^a(\bar{s}_1, s_1^b) \\ 0 & \kappa_\tau^a(\bar{s}_2, s_1^b) \\ 0 & \kappa_\tau^a(\bar{s}_2, s_2^b) \\ 0 & \kappa_\tau^a(\bar{s}_2, s_2^b) \end{array} \right].$$

(b) KBSF's sparse matrices

$$\dot{\mathbf{D}}^a = \begin{array}{c} \hat{s}_1^a \\ \hat{s}_2^a \\ \hat{s}_3^a \end{array} \begin{array}{c} \bar{s}_1 \\ \bar{s}_2 \\ \bar{s}_1 \end{array} \begin{array}{c} \bar{s}_2 \\ \bar{s}_2 \\ \bar{s}_2 \end{array} \left[\begin{array}{cc} \bar{\kappa}_\tau(\hat{s}_1^a, \bar{s}_1) & \bar{\kappa}_\tau(\hat{s}_1^a, \bar{s}_2) \\ \bar{\kappa}_\tau(\hat{s}_2^a, \bar{s}_1) & \bar{\kappa}_\tau(\hat{s}_2^a, \bar{s}_2) \\ \bar{\kappa}_\tau(\hat{s}_3^a, \bar{s}_1) & \bar{\kappa}_\tau(\hat{s}_3^a, \bar{s}_2) \end{array} \right],$$

$$\dot{\mathbf{K}}^a = \begin{array}{c} \bar{s}_1 \\ \bar{s}_2 \end{array} \begin{array}{c} \hat{s}_1^a \\ \hat{s}_2^a \\ \hat{s}_3^a \end{array} \begin{array}{c} \hat{s}_1^a \\ \hat{s}_2^a \\ \hat{s}_3^a \end{array} \left[\begin{array}{cc} \kappa_\tau^a(\bar{s}_1, s_1^a) & \kappa_\tau^a(\bar{s}_1, s_2^a) \\ \kappa_\tau^a(\bar{s}_2, s_1^a) & \kappa_\tau^a(\bar{s}_2, s_2^a) \\ \kappa_\tau^a(\bar{s}_2, s_3^a) & \kappa_\tau^a(\bar{s}_2, s_3^a) \end{array} \right],$$

$$\dot{\mathbf{D}}^b = \begin{array}{c} \hat{s}_1^b \\ \hat{s}_2^b \end{array} \begin{array}{c} \bar{s}_1 \\ \bar{s}_2 \end{array} \begin{array}{c} \bar{s}_2 \\ \bar{s}_2 \end{array} \left[\begin{array}{cc} \bar{\kappa}_\tau(\hat{s}_1^b, \bar{s}_1) & \bar{\kappa}_\tau(\hat{s}_1^b, \bar{s}_2) \\ \bar{\kappa}_\tau(\hat{s}_2^b, \bar{s}_1) & \bar{\kappa}_\tau(\hat{s}_2^b, \bar{s}_2) \end{array} \right],$$

$$\dot{\mathbf{K}}^b = \begin{array}{c} \bar{s}_1 \\ \bar{s}_2 \end{array} \begin{array}{c} \hat{s}_1^b \\ \hat{s}_2^b \end{array} \begin{array}{c} \hat{s}_1^b \\ \hat{s}_2^b \end{array} \left[\begin{array}{cc} \kappa_\tau^a(\bar{s}_1, s_1^b) & \kappa_\tau^a(\bar{s}_1, s_2^b) \\ \kappa_\tau^a(\bar{s}_2, s_1^b) & \kappa_\tau^a(\bar{s}_2, s_2^b) \end{array} \right].$$

(c) KBSF's dense matrices

 Figure 2: Matrices built by KBRL and KBSF for the case in which the original MDP has two actions, a and b , and $n_a = 3$, $n_b = 2$, and $m = 2$.

where $s \in \mathbb{S}$ and $a \in A$. Note that $\tilde{V}(\hat{s}_i^a)$ corresponds to one specific entry of vector $\tilde{\mathbf{v}}$, whose index is given by $\sum_{b=0}^{a-1} n_b + i$, where we assume that $n_0 = 0$.

Algorithm 1 Batch KBSF

Input: $S^a = \{(s_k^a, r_k^a, \hat{s}_k^a) | k = 1, 2, \dots, n_a\}$ for all $a \in A$ ▷ Sample transitions
 $\bar{S} = \{\bar{s}_1, \bar{s}_2, \dots, \bar{s}_m\}$ ▷ Set of representative states
Output: $\tilde{\mathbf{v}} \approx \hat{\mathbf{v}}^*$
for each $a \in A$ **do**
 Compute matrix $\dot{\mathbf{D}}^a$: $\dot{d}_{ij}^a = \bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, \bar{s}_j)$
 Compute matrix $\dot{\mathbf{K}}^a$: $\dot{k}_{ij}^a = \kappa_{\bar{\tau}}^a(\bar{s}_i, s_j^a)$
 Compute vector $\bar{\mathbf{r}}^a$: $\bar{r}_i^a = \sum_j \dot{k}_{ij}^a r_j^a$
 Compute matrix $\bar{\mathbf{P}}^a = \dot{\mathbf{K}}^a \dot{\mathbf{D}}^a$
Solve $\bar{M} \equiv (\bar{S}, A, \bar{\mathbf{P}}^a, \bar{\mathbf{r}}^a, \gamma)$ ▷ i.e., compute $\bar{\mathbf{Q}}^*$
Return $\tilde{\mathbf{v}} = \Gamma \mathbf{D} \bar{\mathbf{Q}}^*$, where $\mathbf{D}^\top = [(\dot{\mathbf{D}}^1)^\top, (\dot{\mathbf{D}}^2)^\top, \dots, (\dot{\mathbf{D}}^{|A|})^\top]$

As shown in Algorithm 1, the key point of KBSF’s mechanics is the fact that the matrices $\tilde{\mathbf{P}}^a = \mathbf{D}\mathbf{K}^a$ are never actually computed, but instead we directly solve the MDP \bar{M} containing m states only. This results in an efficient algorithm that requires only $O(nm|A|d_{\mathbb{S}} + \hat{n}m^2|A|)$ operations and $O(\hat{n}m)$ bits to build a reduced version of KBRL’s MDP, where $\hat{n} = \max_a n_a$. After the reduced model \bar{M} has been constructed, KBSF’s computational cost becomes a function of m only. In particular, the cost of solving \bar{M} through dynamic programming becomes polynomial in m instead of n : while one application of \hat{T} , the Bellman operator of \hat{M} , is $O(n\hat{n}|A|)$, the computation of \bar{T} is $O(m^2|A|)$. Therefore, KBSF’s time and memory complexities are only linear in n .

We note that, in practice, KBSF’s computational requirements can be reduced even further if one enforces the kernels $\kappa_{\bar{\tau}}^a$ and $\bar{\kappa}_{\bar{\tau}}$ to be sparse. In particular, given a fixed \bar{s}_i , instead of computing $\bar{\kappa}_{\bar{\tau}}(\bar{s}_i, s_j^a)$ for $j = 1, 2, \dots, n_a$, one can evaluate the kernel on a pre-specified neighborhood of \bar{s}_i only. Assuming that $\bar{\kappa}_{\bar{\tau}}(\bar{s}_i, s_j^a)$ is zero for all s_j^a outside this region, one can avoid not only computing the kernel but also storing the resulting values (the same reasoning applies to the computation of $\kappa_{\bar{\tau}}(\hat{s}_i^a, \bar{s}_j)$ for a fixed \hat{s}_i^a).

4.1 Theoretical results

Since KBSF comes down to the solution of a finite MDP, it always converges to the same approximation $\tilde{\mathbf{v}}$, whose distance to KBRL’s optimal value function $\hat{\mathbf{v}}^*$ is bounded by Proposition 2. Once $\tilde{\mathbf{v}}$ is available, the value of any state-action pair can be determined through (12). The following result generalizes Proposition 2 to the entire continuous state space \mathbb{S} :

Proposition 3 *Let \hat{Q} be the value function computed by KBRL through (7) and let \tilde{Q} be the value function computed by KBSF through (12). Then, for any $s \in \mathbb{S}$ and any $a \in A$, $|\hat{Q}(s, a) - \tilde{Q}(s, a)| \leq \gamma \xi_v$, with ξ_v defined in (8).*

Proof

$$\begin{aligned} |\hat{Q}(s, a) - \tilde{Q}(s, a)| &= \left| \sum_{i=1}^{n_a} \kappa_\tau^a(s, s_i^a) \left[r_i^a + \gamma \hat{V}^*(\hat{s}_i^a) \right] - \sum_{i=1}^{n_a} \kappa_\tau^a(s, s_i^a) \left[r_i^a + \gamma \tilde{V}(\hat{s}_i^a) \right] \right| \\ &\leq \gamma \sum_{i=1}^{n_a} \kappa_\tau^a(s, s_i^a) \left| \hat{V}^*(\hat{s}_i^a) - \tilde{V}(\hat{s}_i^a) \right| \leq \gamma \sum_{i=1}^{n_a} \kappa_\tau^a(s, s_i^a) \xi_v \leq \gamma \xi_v, \end{aligned}$$

where the second inequality results from the application of Proposition 2 and the third inequality is a consequence of the fact that $\sum_{i=1}^{n_a} \kappa_\tau^a(s, s_i^a)$ defines a convex combination. ■

Proposition 3 makes it clear that the approximation computed by KBSF depends crucially on ξ_v . In the remainder of this section we will show that, if the distances between sampled states and the respective nearest representative states are small enough, then we can make ξ_v as small as desired by setting $\bar{\tau}$ to a sufficiently small value.

4.1.1 GENERAL RESULTS

We assume that KBSF’s kernel $\bar{\phi}(x) : \mathbb{R}^+ \mapsto \mathbb{R}^+$ has the following properties:

- (i) $\bar{\phi}(x) \geq \bar{\phi}(y)$ if $x < y$,
- (ii) $\exists A_{\bar{\phi}} > 0, \lambda_{\bar{\phi}} \geq 1, B_{\bar{\phi}} \geq 0$ such that $A_{\bar{\phi}} \exp(-x) \leq \bar{\phi}(x) \leq \lambda_{\bar{\phi}} A_{\bar{\phi}} \exp(-x)$ if $x \geq B_{\bar{\phi}}$.

Assumption (ii) implies that the function $\bar{\phi}$ is positive and will eventually decay exponentially. Note that we assume that $\bar{\phi}$ is greater than zero everywhere in order to guarantee that $\bar{\kappa}_{\bar{\tau}}$ is well defined for any value of $\bar{\tau}$. It should be straightforward to generalize our results for the case in which $\bar{\phi}$ has finite support by ensuring that, given sets of sample transitions S^a and a set of representative states \bar{S} , $\bar{\tau}$ is such that, for any $\hat{s}_i^a \in S^a$, with $a \in A$, there is a $\bar{s}_j \in \bar{S}$ for which $\bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, \bar{s}_j) > 0$ (note that this assumption is naturally satisfied by the “sparse kernels” used in some of the experiments—see Appendix B).

Let $rs : \mathbb{S} \times \{1, 2, \dots, m\} \mapsto \bar{S}$ be a function that orders the representative states according to their distance to a given state s , that is,

$$rs(s, i) = \bar{s}_k \iff \bar{s}_k \text{ is the } i^{\text{th}} \text{ nearest representative state to } s. \quad (13)$$

Define

$$\begin{aligned} dist : \mathbb{S} \times \{1, 2, \dots, m\} &\mapsto \mathbb{R} \\ dist(s, i) &= \|s - rs(s, i)\|. \end{aligned} \quad (14)$$

We will show that, for any $\epsilon > 0$ and any $w \in \{1, 2, \dots, m\}$, there is a $\delta > 0$ such that, if $\max_{a,i} dist(\hat{s}_i^a, w) < \delta$, then we can set $\bar{\tau}$ in order to guarantee that $\xi_v < \epsilon$. To show that, we will need the following two lemmas:

Lemma 4 *For any $s_i^a \in S^a$ and any $\epsilon > 0$, there is a $\delta > 0$ such that $|\kappa_\tau^a(s, s_i^a) - \kappa_\tau^a(s', s_i^a)| < \epsilon$ if $\|s - s'\| < \delta$.*

Lemma 5 *Let $s \in \mathbb{S}$, let $m > 1$, and assume there is a $w \in \{1, 2, \dots, m - 1\}$ such that $\text{dist}(s, w) < \text{dist}(s, w + 1)$. Define*

$$W^w(s) \equiv \{k \mid \|s - \bar{s}_k\| \leq \text{dist}(s, w)\} \text{ and } \bar{W}^w(s) \equiv \{1, 2, \dots, m\} - W^w(s). \quad (15)$$

Then, for any $\alpha > 0$, $\sum_{k \in W^w(s)} \bar{\kappa}_{\bar{\tau}}(s, \bar{s}_k) < \alpha \sum_{k \in \bar{W}^w(s)} \bar{\kappa}_{\bar{\tau}}(s, \bar{s}_k)$ for $\bar{\tau}$ sufficiently small.

Lemma 4 is basically a continuity argument: it shows that, for any fixed s_i^a , $|\kappa_{\bar{\tau}}^a(s, s_i^a) - \kappa_{\bar{\tau}}^a(s', s_i^a)| \rightarrow 0$ as $\|s - s'\| \rightarrow 0$. Lemma 5 states that, if we order the representative states according to their distance to a fixed state s , and then partition them in two subsets, we can control the relative magnitude of the corresponding kernels's sums by adjusting the parameter $\bar{\tau}$ (we redirect the reader to Appendix A.1 for details on how to set $\bar{\tau}$). Based on these two lemmas, we present the main result of this section:

Proposition 6 *Let $w \in \{1, 2, \dots, m\}$. For any $\epsilon > 0$, there is a $\delta > 0$ such that, if $\max_{a,i} \text{dist}(\hat{s}_i^a, w) < \delta$, then we can guarantee that $\xi_v < \epsilon$ by making $\bar{\tau}$ sufficiently small.*

Proposition 6 tells us that, regardless of the specific reinforcement learning problem at hand, if the distances between sampled states \hat{s}_i^a and the respective w nearest representative states are small enough, then we can make KBSF's approximation of KBRL's value function as accurate as desired by setting $\bar{\tau}$ to a sufficiently small value (one can see how exactly to set $\bar{\tau}$ in the proof of the proposition). How small the maximum distance $\max_{a,i} \text{dist}(\hat{s}_i^a, w)$ should be depends on the particular choice of kernel $k_{\bar{\tau}}$ and on the sets of sample transitions S^a .

Note that a fixed number of representative states m imposes a minimum possible value for $\max_{a,i} \text{dist}(\hat{s}_i^a, w)$, and if this value is not small enough decreasing $\bar{\tau}$ may actually hurt the approximation (this is easier to see if we consider that $w = 1$). The optimal value for $\bar{\tau}$ in this case is again context-dependent. As a positive flip side of this statement, we note that, even if $\max_{a,i} \text{dist}(\hat{s}_i^a, w) > \delta$, it might be possible to make $\xi_v < \epsilon$ by setting $\bar{\tau}$ appropriately. Therefore, rather than as a practical guide on how to configure KBSF, Proposition 6 should be seen as a theoretical argument showing that KBSF is a sound algorithm, in the sense that in the limit it recovers KBRL's solution.

4.1.2 ERROR RATE

In the previous section we deliberately refrained from making assumptions on the kernel $\bar{\kappa}_{\bar{\tau}}$ used by KBSF in order to make Proposition 6 as general as possible. In what follows we show that, by restricting the class of kernels used by our algorithm, we can derive stronger results regarding its behavior. In particular, we derive an upper bound for ξ_v that shows how this approximation error depends on the variables of a given reinforcement learning problem. Our strategy will be to define an ‘‘admissible kernel’’ whose width is determined based on data.

We will need the following assumption:

- (iii) $|\phi(x) - \phi(y)| \leq C_{\phi}|x - y|$, with $C_{\phi} \geq 0$.

Assumption (iii) simply states that the function ϕ used to construct the kernel $\kappa_{\bar{\tau}}^a$ is Lipschitz continuous with constant C_{ϕ} . This is actually part of Ormoneit and Sen's (2002) Assumption 4, and is explicitly listed here for convenience only.

We will now define an auxiliary function which will be used in the definition of our admissible kernel. To simplify the notation, let

$$\varsigma_k^{a,i,j} \equiv |\kappa_\tau^a(\hat{s}_i^a, s_j^a) - \kappa_\tau^a(\bar{s}_k, s_j^a)|. \quad (16)$$

Based on (15) and (16), we define the following function:

$$\mathcal{F}(\bar{\tau}, w | S^a, \bar{S}, \kappa_\tau^a, \bar{k}_\tau) = \frac{R_{\max}}{(1-\gamma)^2} \left[\max_{a,i} \sum_{k \in \bar{W}^w(\hat{s}_i^a)} \bar{k}_\tau(\hat{s}_i^a, \bar{s}_k) \sum_{j=1}^{n_a} \varsigma_k^{a,i,j} + \frac{1}{2} \max_{a,i} (1 - \bar{k}_\tau(\hat{s}_i^a, rs(\hat{s}_i^a, 1))) \right], \quad (17)$$

where $\bar{\tau} > 0$, $w \in \{1, 2, \dots, m\}$, $R_{\max} = \max_{a,i} |r_i^a|$, and rs and \bar{W}^w are defined in (13) and (15), respectively. Note that the definition of \mathcal{F} makes it clear its dependency on the representative states, sets of sample transitions, and kernels adopted.

Lemma 5 implies that, as $\bar{\tau} \rightarrow 0$, $\bar{k}_\tau(\hat{s}_i^a, rs(\hat{s}_i^a, 1)) \rightarrow 1$ and $\sum_{k \in \bar{W}^w(\hat{s}_i^a)} \bar{k}_\tau(\hat{s}_i^a, \bar{s}_k) \rightarrow 0$ for all a, i , and w (see equation (35) in Appendix A.1 for a clear picture). Thus, for any w , $\mathcal{F}(\bar{\tau}, w) \rightarrow 0$ as $\bar{\tau} \rightarrow 0$. This leads to the following definition:

Definition 7 *Given $\epsilon > 0$ and $w \in \{1, 2, \dots, m\}$, an admissible kernel $\bar{k}_\tau^{\epsilon,w}$ is any kernel whose parameter $\bar{\tau}$ is such that $\mathcal{F}(\bar{\tau}, w) < \epsilon$.*

The definition above allows us to enunciate the following proposition:

Proposition 8 *Let $k_{\min} = \phi(\sqrt{d_{\mathbb{S}}}/\tau)$, where τ is the parameter used by KBRL's kernel κ_τ^a . Given $\epsilon > 0$ and $w \in \{1, 2, \dots, m\}$, suppose that KBSF is adopted with an admissible kernel $\bar{k}_\tau^{\epsilon,w}$. Then,*

$$\xi_v \leq \frac{2wC_\phi R_{\max}}{\tau k_{\min}(1-\gamma)^2} \max_{a,i} \text{dist}(\hat{s}_i^a, w) + \epsilon, \quad (18)$$

where C_ϕ is the Lipschitz constant of function ϕ appearing in Assumption (iii), $R_{\max} = \max_{a,i} |r_i^a|$, and $\gamma \in [0, 1)$ is the discount factor of the underlying MDP.

Proposition 8 shows that ξ_v is bounded by the maximum distance between a sampled state and the w^{th} closest representative state scaled by constants characterizing the MDP and the kernels used by KBSF.

Assuming that ϵ and w are fixed, among the quantities appearing in (18) we only have control over τ and $\max_{a,i} \text{dist}(\hat{s}_i^a, w)$ —the latter through the definition of the representative states. Regarding $\max_{a,i} \text{dist}(\hat{s}_i^a, w)$, the same observation made above applies here: given sample transitions S^a , a fixed value for $m < n$ imposes a lower bound on this term, and thus on the right-hand side of (18). As $m \rightarrow n$, this lower bound approaches some value ϵ' , with $0 \leq \epsilon' < \epsilon$. The fact that (18) is generally greater than zero even when $m = n$ reflects the fact that ξ_v depends on $\sigma(\mathbf{D})$, the level of stochasticity of \mathbf{D} (see (8) and (9)). Since Assumption (ii) implies that $\bar{k}_\tau(s, s') > 0$ for all $s, s' \in \mathbb{S}$, matrix \mathbf{D} will never become deterministic, no matter how small $\bar{\tau}$ is (see Figure 2b). If we replace \bar{k}_τ by a kernel with finite support, then we can set $\epsilon = 0$ in Definition 7, and therefore in the right-hand side of (18). Needless to say, this does not mean that using a kernel with finite support will lead to better performance in practice.

As for the definition of τ , we see that the right-hand side of (18) decreases as $\tau \rightarrow \infty$. This means that we can make the value function computed by KBSF arbitrarily close to the one computed by KBRL. Note though that increasing τ also changes the model constructed by KBRL, and an excessively large value for this parameter makes the resulting approximation meaningless (see (4) and (5)). Similarly, we might be tempted to always set w to 1 in order to minimize the right-hand side of (18). Note however that a kernel $\bar{\kappa}_{\bar{\tau}}^{\epsilon, w}$ that is admissible for $w > 1$ may not be so for $w = 1$, and in this case the bound would no longer be valid.

4.2 Empirical results

We now present a series of computational experiments designed to illustrate the behavior of KBSF in a variety of challenging domains. We start with a simple problem, the “puddle world”, to show that KBSF is indeed capable of compressing the information contained in KBRL’s model. We then move to more difficult tasks, and compare KBSF with other state-of-the-art reinforcement-learning algorithms. We start with two classical control tasks, single and double pole-balancing. Next we study two medically-related problems based on real data: HIV drug schedule and epilepsy-suppression domains.

All problems considered in this paper have a continuous state space and a finite number of actions, and were modeled as discounted tasks. The algorithms’s results correspond to the performance of the greedy decision policy derived from the final value function computed. In all cases, the decision policies were evaluated on challenging test states from which the tasks cannot be easily solved. The details of the experiments are given in Appendix B.

4.2.1 PUDDLE WORLD (PROOF OF CONCEPT)

In order to show that KBSF is indeed capable of summarizing the information contained in KBRL’s model, we use the puddle world task (Sutton, 1996). The puddle world is a simple two-dimensional problem in which the objective is to reach a goal region avoiding two “puddles” along the way. We implemented the task exactly as described by Sutton (1996), except that we used a discount factor of $\gamma = 0.99$ and evaluated the decision policies on a set of pre-defined test states surrounding the puddles (see Appendix B).

The experiment was carried out as follows: first, we collected a set of n sample transitions $(s_k^a, r_k^a, \hat{s}_k^a)$ using a random exploration policy (that is, a policy that selects actions uniformly at random). In the case of KBRL, this set of sample transitions defined the model used to approximate the value function. In order to define KBSF’s model, the states \hat{s}_k^a were grouped by the k -means algorithm into m clusters and a representative state \bar{s}_j was placed at the center of each resulting cluster (Kaufman and Rousseeuw, 1990). As for the kernels’s widths, we varied both τ and $\bar{\tau}$ in the set $\{0.01, 0.1, 1\}$ (see Table 1). The results reported represent the best performance of the algorithms over 50 runs; that is, for each n and each m we picked the combination of parameters that generated the maximum average return. We use the following convention to refer to specific instances of each method: the first number enclosed in parentheses after an algorithm’s name is n , the number of sample transitions used in the approximation, and the second one is m , the size of the model used to approximate the value function. Note that for KBRL n and m coincide.

In Figure 3a and 3b we observe the effect of fixing the number of transitions n and varying the number of representative states m . As expected, KBSF’s results improve as $m \rightarrow n$. More surprising is the fact that KBSF has essentially the same performance as KBRL using models one order of magnitude smaller. This indicates that KBSF is summarizing well the information contained in the data. Depending on the values of n and m , such a compression may represent a significant reduction on the consumption of computational resources. For example, by replacing KBRL(8000) with KBSF(8000, 100), we obtain a decrease of approximately 99.58% on the number of operations performed to find a policy, as shown in Figure 3b (the cost of constructing KBSF’s MDP is included in all reported run times).

In Figures 3c and 3d we fix m and vary n . Observe in Figure 3c how KBRL and KBSF have similar performances, and both improve as n increases. However, since KBSF is using a model of fixed size, its computational cost depends only linearly on n , whereas KBRL’s cost grows with $n^2\hat{n}$, roughly. This explains the huge difference in the algorithms’s run times shown in Figure 3d.

4.2.2 SINGLE AND DOUBLE POLE-BALANCING (COMPARISON WITH LSPI)

We now evaluate how KBSF compares to other modern reinforcement learning algorithms on more difficult tasks. We first contrast our method with Lagoudakis and Parr’s (2003) least-squares policy iteration algorithm (LSPI). Besides its popularity, LSPI is a natural candidate for such a comparison for three reasons: it also builds an approximator of fixed size out of a batch of sample transitions, it has good theoretical guarantees, and it has been successfully applied to several reinforcement learning tasks.

We compare the performance of LSPI and KBSF on the pole balancing task. Pole balancing has a long history as a benchmark problem because it represents a rich class of unstable systems (Michie and Chambers, 1968; Anderson, 1986; Barto et al., 1983). The objective in this problem is to apply forces to a wheeled cart moving along a limited track in order to keep one or more poles hinged to the cart from falling over. There are several variations of the task with different levels of difficulty; among them, balancing two poles side by side is particularly hard (Wieland, 1991). In this paper we compare LSPI and KBSF on both the single- and two-poles versions of the problem. We implemented the tasks using a realistic simulator described by Gomez (2003). We refer the reader to Appendix B for details on the problems’s configuration.

The experiments were carried out as described in the previous section, with sample transitions collected by a random policy and then clustered by the k -means algorithm. In both versions of the pole-balancing task LSPI used the same data and approximation architectures as KBSF. To make the comparison with LSPI as fair as possible, we fixed the width of KBSF’s kernel κ_τ^α at $\tau = 1$ and varied $\bar{\tau}$ in $\{0.01, 0.1, 1\}$ for both algorithms. Also, policy iteration was used to find a decision policy for the MDPs constructed by KBSF, and this algorithm was run for a maximum of 30 iterations, the same limit used for LSPI.

Figure 4 shows the results of LSPI and KBSF on the single and double pole-balancing tasks. We call attention to the fact that the version of the problems used here is significantly harder than the more commonly-used variants in which the decision policies are evaluated on a single state close to the origin. This is probably the reason why LSPI achieves a

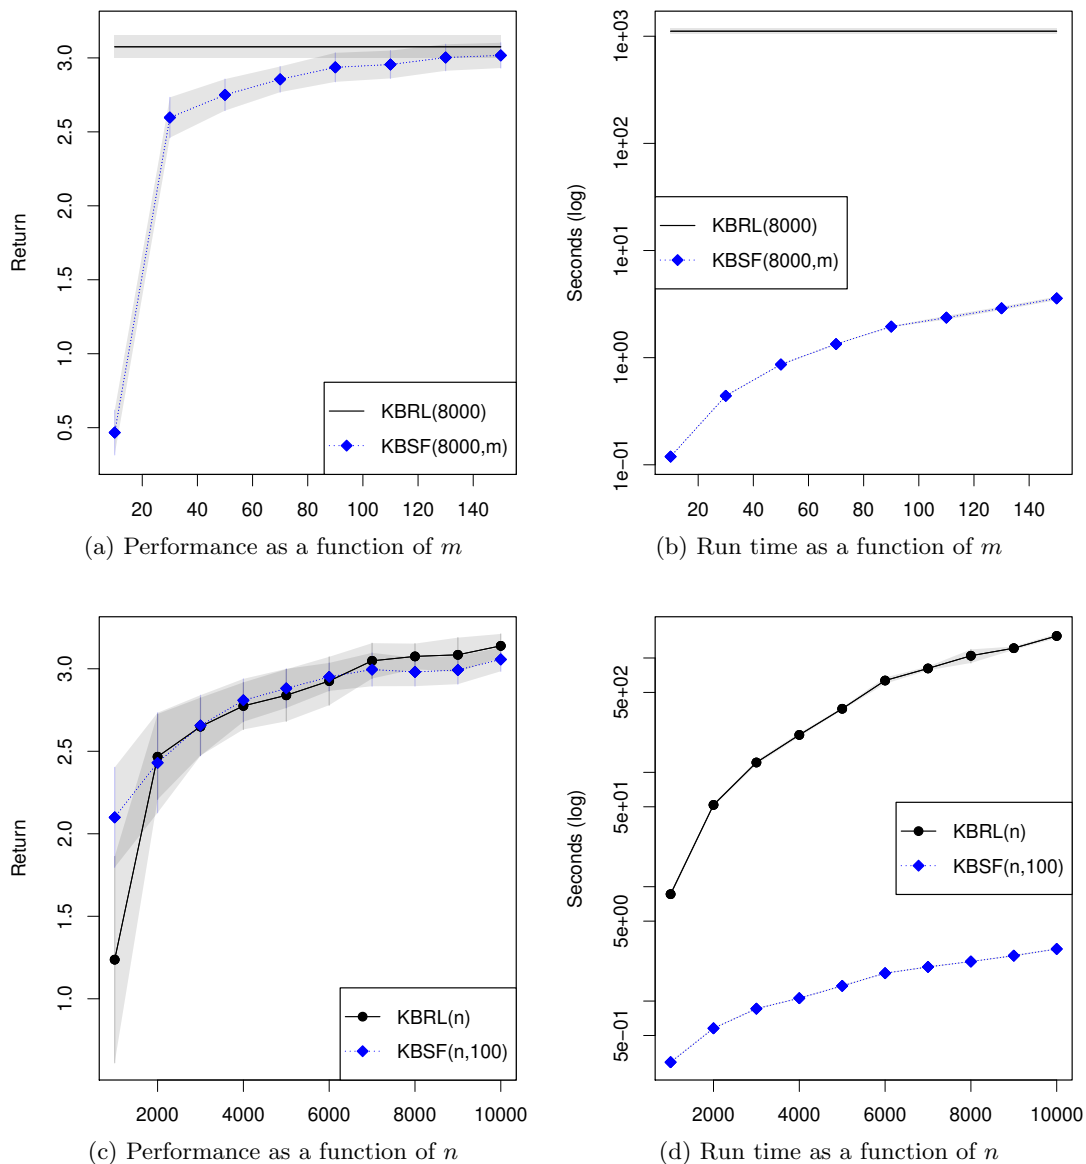


Figure 3: Results on the puddle-world task averaged over 50 runs. The algorithms were evaluated on a set of test states distributed over a region of the state space surrounding the “puddles” (details in Appendix B). The shadowed regions represent 99% confidence intervals.

success rate of no more than 60% on the single pole-balancing task, as shown in Figure 4a. In contrast, KBSF’s decision policies are able to balance the pole in 90% of the attempts, on average, using as few as $m = 30$ representative states.

The results of KBSF on the double pole-balancing task are still more impressive. As Wieland (1991) rightly points out, this version of the problem is considerably more difficult than its single pole variant, and previous attempts to apply reinforcement-learning techniques to this domain resulted in disappointing performance (Gomez et al., 2006). As shown in Figure 4c, KBSF(10^6 , 200) is able to achieve a success rate of more than 80%. To put this number in perspective, recall that some of the test states are quite challenging, with the two poles inclined and falling in opposite directions.

The good performance of KBSF comes at a relatively low computational cost. A conservative estimate reveals that, were KBRL(10^6) run on the same computer used for these experiments, we would have to wait for more than 6 *months* to see the results. KBSF(10^6 , 200) delivers a decision policy in less than 7 minutes. KBSF’s computational cost also compares well with that of LSPI, as shown in Figures 4b and 4d. LSPI’s policy evaluation step involves the update and solution of a linear system of equations, which take $O(nm^2)$ and $O(m^3|A|^3)$, respectively. In addition, the policy-update stage requires the definition of $\pi(\hat{s}_k^a)$ for all n states in the set of sample transitions. In contrast, at each iteration KBSF only performs $O(m^3)$ operations to evaluate a decision policy and $O(m^2|A|)$ operations to update it.

4.2.3 HIV DRUG SCHEDULE DOMAIN (COMPARISON WITH FITTED Q -ITERATION)

We now compare KBSF with the fitted Q -iteration algorithm (Ernst et al., 2005; Antos et al., 2007; Munos and Szepesvári, 2008). Fitted Q -iteration is a conceptually simple method that also builds its approximation based solely on sample transitions. Here we adopt this algorithm with an ensemble of trees generated by Geurts et al.’s (2006) extra-trees algorithm. This combination, which we call FQIT, generated the best results on the extensive empirical evaluation performed by Ernst et al. (2005).

We chose FQIT for our comparisons because it has shown excellent performance on both benchmark and real-world reinforcement-learning tasks (Ernst et al., 2005, 2006). In all experiments reported in this paper we used FQIT with ensembles of 30 trees. As detailed in Appendix B, besides the number of trees, FQIT has three main parameters. Among them, the minimum number of elements required to split a node in the construction of the trees, denoted here by η_{\min} , has a particularly strong effect on both the algorithm’s performance and computational cost. Thus, in our experiments we fixed FQIT’s parameters at reasonable values—selected based on preliminary experiments—and only varied η_{\min} . The respective instances of the tree-based approach are referred to as FQIT(η_{\min}).

We compare FQIT and KBSF on an important medical problem which we will refer to as the HIV drug schedule domain (Adams et al., 2004; Ernst et al., 2006). Typical HIV treatments use drug cocktails containing two types of medication: reverse transcriptase inhibitors (RTI) and protease inhibitors (PI). Despite the success of drug cocktails in maintaining low viral loads, there are several complications associated with their long-term use. This has attracted the interest of the scientific community to the problem of optimizing drug-scheduling strategies. One strategy that has been receiving a lot of attention recently

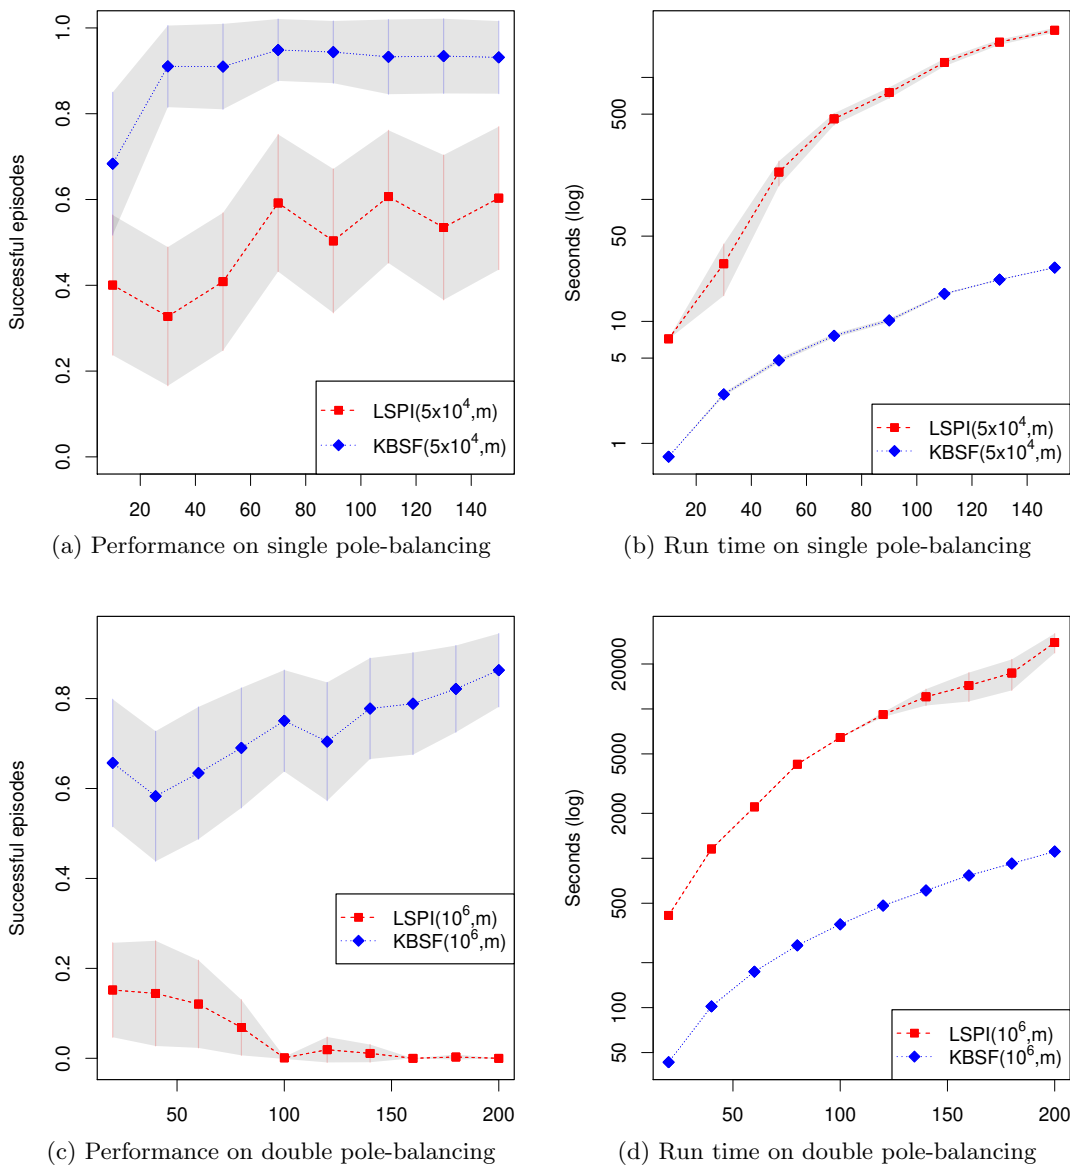


Figure 4: Results on the pole-balancing tasks, as a function of the number of representative states m , averaged over 50 runs. The values correspond to the fraction of episodes initiated from the test states in which the pole(s) could be balanced for 3000 steps (one minute of simulated time). The test sets were regular grids defined over the hypercube centered at the origin and covering 50% of the state-space axes in each dimension (see Appendix B). Shaded regions represent 99% confidence intervals.

is structured treatment interruption (STI), in which patients undergo alternate cycles with and without the drugs. Although many successful STI treatments have been reported in the literature, as of now there is no consensus regarding the exact protocol that should be followed (Bajaria et al., 2004).

The scheduling of STI treatments can be seen as a sequential decision problem in which the actions correspond to the types of cocktail that should be administered to a patient (Ernst et al., 2006). To simplify the problem’s formulation, it is assumed that RTI and PI drugs are administered at fixed amounts, reducing the actions to the four possible combinations of drugs: none, RTI only, PI only, or both. The goal is to minimize the viral load using as little drugs as possible. Following Ernst et al. (2006), we performed our experiments using a model that describes the interaction of the immune system with HIV. This model was developed by Adams et al. (2004) and has been identified and validated based on real clinical data. The resulting reinforcement learning task has a 6-dimensional continuous state space whose variables describe the overall patient’s condition.

We formulated the problem exactly as proposed by Ernst et al. (2006, see Appendix B for details). The strategy used to generate the data also followed the protocol proposed by these authors, which we now briefly explain. Starting from a batch of 6000 sample transitions generated by a random policy, each algorithm first computed an initial approximation of the problem’s optimal value function. Based on this approximation, a 0.15-greedy policy was used to collect a second batch of 6000 transitions, which was merged with the first.³ This process was repeated for 10 rounds, resulting in a total of 60000 sample transitions.

We varied FQIT’s parameter η_{\min} in the set $\{50, 100, 200\}$. For the experiments with KBSF, we fixed $\tau = \bar{\tau} = 1$ and varied m in $\{2000, 4000, \dots, 10000\}$ (in the rounds in which $m \geq n$ we simply used all states \hat{s}_i^a as representative states). As discussed in the beginning of this section, it is possible to reduce KBSF’s computational cost with the use of sparse kernels. In our experiments with the HIV drug schedule task, we only computed the $\mu = 2$ largest values of $k_{\tau}(\bar{s}_i, \cdot)$ and the $\bar{\mu} = 3$ largest values of $\bar{k}_{\bar{\tau}}(\hat{s}_i^a, \cdot)$ (see Appendix B.2). The representative states \bar{s}_i were selected at random from the set of sampled states \hat{s}_i^a (the reason for this will become clear shortly). Since in the current experiments the number of sample transitions n was fixed, we will refer to the particular instances of our algorithm simply as KBSF(m).

Before discussing the results, we point out that FQIT’s performance on the HIV drug schedule task is very good, comparable to that of Adams et al.’s (2004) approach, which uses a model of the task. FQIT’s results, along with KBSF’s, are shown in Figure 5. As shown in Figure 5a, the performance of FQIT improves when η_{\min} is decreased, as expected. In contrast, increasing the number of representative states m does not have a strong impact on the quality of KBSF’s solutions (in fact, in some cases the average return obtained by the resulting policies decreases slightly when m grows). Overall, the performance of KBSF on the HIV drug schedule task is not nearly as impressive as on the previous problems. For example, even when using $m = 10000$ representative states, which corresponds to one sixth of the sampled states, KBSF is unable to reproduce the performance of FQIT with $\eta_{\min} = 50$.

3. As explained by Sutton and Barto (1998), an ϵ -greedy policy selects the action with maximum value with probability $1 - \epsilon$, and with probability ϵ it picks an action uniformly at random.

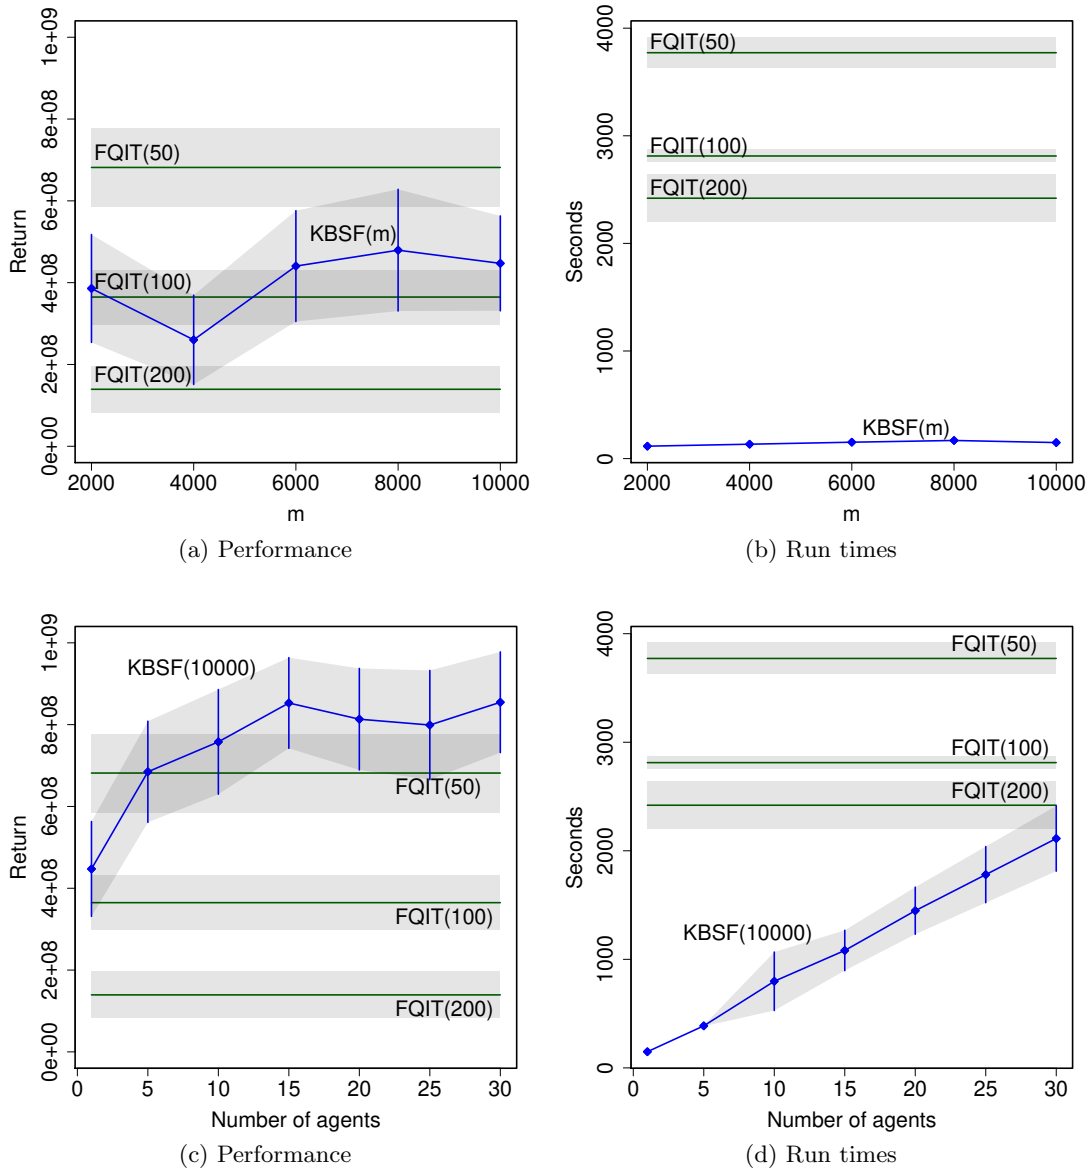


Figure 5: Results on the HIV drug schedule task averaged over 50 runs. The STI policies were evaluated for 5000 days starting from a state representing a patient’s unhealthy state (see Appendix B). The shadowed regions represent 99% confidence intervals.

On the other hand, when we look at Figure 5b, it is clear that the difference on the algorithms’s performance is counterbalanced by a substantial difference on the associated computational costs. As an illustration, note that KBSF(10000) is 15 times faster than FQTI(100) and 20 times faster than FQTI(50). This difference on the algorithms’s run times is expected, since each iteration of FQIT involves the construction (or update) of an ensemble of trees, each one requiring at least $O(n \log(n/\eta_{\min}))$ operations, and the improvement of the current decision policy, which is $O(n|A|)$ (Geurts et al., 2006). As discussed before, KBSF’s efficiency comes from the fact that its computational cost per iteration is independent of the number of sample transitions n .

Note that the fact that FQIT uses an ensemble of trees is both a blessing and a curse. If on the one hand this reduces the variance of the approximation, on the other hand it also increases the algorithm’s computational cost (Geurts et al., 2006). Given the big gap between FQIT’s and KBSF’s time complexities, one may wonder if the latter can also benefit from averaging over several models. In order to verify this hypothesis, we implemented a very simple model-averaging strategy with KBSF: we trained several agents independently, using Algorithm 1 on the same set of sample transitions, and then put them together on a single “committee”. In order to increase the variability within the committee of agents, instead of using k -means to determine the representative states \bar{s}_j we simply selected them uniformly at random from the set of sampled states \hat{s}_i^a (note that this has the extra benefit of reducing the method’s overall computational cost). The actions selected by the committee of agents were determined through “voting”—that is, we simply picked the action chosen by the majority of agents, with ties broken randomly.

We do not claim that the approach described above is the best model-averaging strategy to be used with KBSF. However, it seems to be sufficient to boost the algorithm’s performance considerably, as shown in Figure 5c. Note how KBSF already performs comparably to FQTI(50) when using only 5 agents in the committee. When this number is increased to 15, the expected return of KBSF’s agents is considerably larger than that of the best FQIT’s agent, with only a small overlap between the 99% confidence intervals associated with the algorithms’s results. The good performance of KBSF is still more impressive when we look at Figure 5d, which shows that even when using a committee of 30 agents this algorithm is faster than FQIT(200).

4.2.4 EPILEPSY-SUPPRESSION DOMAIN (COMPARISON WITH LSPI AND FITTED Q -ITERATION)

We conclude our empirical evaluation of KBSF by using it to learn a neuro-stimulation policy for the treatment of epilepsy. It has been shown that the electrical stimulation of specific structures in the neural system at fixed frequencies can effectively suppress the occurrence of seizures (Durand and Bikson, 2001). Unfortunately, *in vitro* neuro-stimulation experiments suggest that fixed-frequency pulses are not equally effective across epileptic systems. Moreover, the long term use of this treatment may potentially damage the patients’s neural tissues. Therefore, it is desirable to develop neuro-stimulation policies that replace the fixed-stimulation regime with an adaptive scheme.

The search for efficient neuro-stimulation strategies can be seen as a reinforcement learning problem. Here we study this problem using a generative model developed by Bush et al.

(2009) based on real data collected from epileptic rat hippocampus slices. Bush et al.’s model was shown to reproduce the seizure pattern of the original dynamical system and was later validated through the deployment of a learned treatment policy on a real brain slice (Bush and Pineau, 2009). The associated decision problem has a five-dimensional continuous state space and highly non-linear dynamics. At each time step the agent must choose whether or not to apply an electrical pulse. The goal is to suppress seizures as much as possible while minimizing the total amount of stimulation needed to do so.

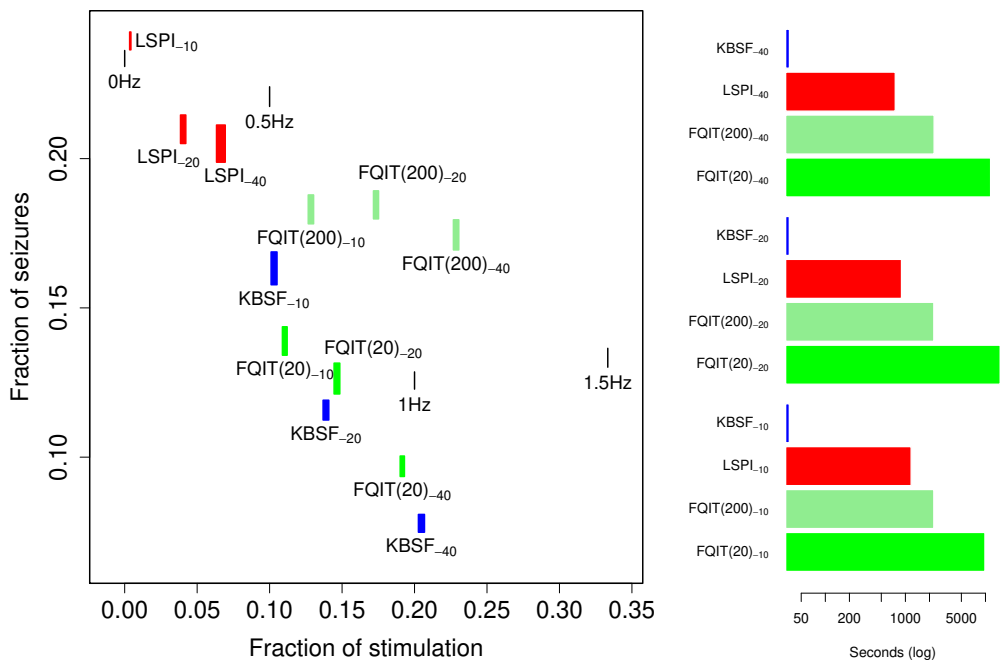
The experiments were performed as described in Section 4.2.1, with a single batch of sample transitions collected by a policy that selects actions uniformly at random. Specifically, the random policy was used to collect 50 trajectories of length 10000, resulting in a total of 500000 sample transitions. We use as a baseline for our comparisons the already mentioned fixed-frequency stimulation policies usually adopted in *in vitro* clinical studies (Bush and Pineau, 2009). In particular, we considered policies that apply electrical pulses at frequencies of 0 Hz, 0.5 Hz, 1 Hz, and 1.5 Hz.

We compare KBSF with LSPI and FQIT. For this task we ran both LSPI and KBSF with sparse kernels, that is, we only computed the kernels at the 6-nearest neighbors of a given state ($\mu = \bar{\mu} = 6$; see Appendix B.2 for details). This modification made it possible to use $m = 50000$ representative states with KBSF. Since for LSPI the reduction on the computational cost was not very significant, we fixed $m = 50$ to keep its run time within reasonable bounds. Again, KBSF and LSPI used the same approximation architectures, with representative states defined by the k -means algorithm. We fixed $\tau = 1$ and varied $\bar{\tau}$ in $\{0.01, 0.1, 1\}$. FQIT was configured as described in the previous section, with the parameter η_{\min} varying in $\{20, 30, \dots, 200\}$. In general, we observed that the performance of the tree-based method improved with smaller values for η_{\min} , with an expected increase in the computational cost. Thus, in order to give an overall characterization of FQIT’s performance, we only report the results obtained with the extreme values of η_{\min} .

Figure 6 shows the results on the epilepsy-suppression task. In order to obtain different compromises between the problem’s two conflicting objectives, we varied the relative magnitude of the penalties associated with the occurrence of seizures and with the application of an electrical pulse (Bush et al., 2009; Bush and Pineau, 2009). Specifically, we fixed the latter at -1 and varied the former with values in $\{-10, -20, -40\}$. This appears in the plots as subscripts next to the algorithms’s names. As shown in Figure 6a, LSPI’s policies seem to prioritize reduction of stimulation at the expense of higher seizure occurrence, which is clearly sub-optimal from a clinical point of view. FQIT(200) also performs poorly, with solutions representing no advance over the fixed-frequency stimulation strategies. In contrast, FQIT(20) and KBSF are both able to generate decision policies that are superior to the 1 Hz policy, which is the most efficient stimulation regime known to date in the clinical literature (Jerger and Schiff, 1995). However, as shown in Figure 6b, KBSF is able to do it at least 100 times faster than the tree-based method.

5. Incremental KBSF

As clear in the previous section, one characteristic of KBSF that sets it apart from other methods is its low demand in terms of computational resources. Specifically, both time and memory complexities of our algorithm are linear in the number of sample transitions n . In



(a) Performance. The length of the rectangles’s edges represent 99% confidence intervals.

(b) Run times (confidence intervals do not show up in logarithmic scale)

Figure 6: Results on the epilepsy-suppression problem averaged over 50 runs. The decision policies were evaluated on episodes of 10^5 transitions starting from a fixed set of 10 test states drawn uniformly at random.

terms of the number of operations performed by the algorithm, this is the best one can do without discarding transitions. However, in terms of memory usage, it is possible to do even better. In this section we show how to build KBSF's approximation incrementally, without ever having access to the entire set of sample transitions at once. Besides reducing the memory complexity of the algorithm, this modification has the additional advantage of making KBSF suitable for on-line reinforcement learning.

In the batch version of KBSF, described in Section 4, the matrices $\bar{\mathbf{P}}^a$ and vectors $\bar{\mathbf{r}}^a$ are determined using all the transitions in the corresponding sets S^a . This has two undesirable consequences. First, the construction of the MDP \bar{M} requires an amount of memory of $O(\hat{n}m)$, where $\hat{n} = \max_a n_a$. Although this is a significant improvement over KBRL's memory usage, which is lower bounded by $(\min_a n_a)^2 |A|$, in more challenging domains even a linear dependence on \hat{n} may be impractical. Second, in the batch version of KBSF the only way to incorporate new data into the model \bar{M} is to recompute the multiplication $\bar{\mathbf{P}}^a = \bar{\mathbf{K}}^a \bar{\mathbf{D}}^a$ for all actions a for which there are new sample transitions available. Even if we ignore the issue with memory usage, this is clearly inefficient in terms of computation. In what follows we present an incremental version of KBSF that circumvents these important limitations (Barreto et al., 2012).

We assume the same scenario considered in Section 4: there is a set of sample transitions $S^a = \{(s_k^a, r_k^a, \hat{s}_k^a) | k = 1, 2, \dots, n_a\}$ associated with each action $a \in A$, where $s_k^a, \hat{s}_k^a \in \mathbb{S}$ and $r_k^a \in \mathbb{R}$, and a set of representative states $\bar{S} = \{\bar{s}_1, \bar{s}_2, \dots, \bar{s}_m\}$, with $\bar{s}_i \in \mathbb{S}$. Suppose now that we split the set of sample transitions S^a in two subsets S_1 and S_2 such that $S_1 \cap S_2 = \emptyset$ and $S_1 \cup S_2 = S^a$ (we drop the “ a ” superscript in the sets S_1 and S_2 to improve clarity). Without loss of generality, suppose that the sample transitions are indexed so that $S_1 \equiv \{(s_k^a, r_k^a, \hat{s}_k^a) | k = 1, 2, \dots, n_1\}$ and $S_2 \equiv \{(s_k^a, r_k^a, \hat{s}_k^a) | k = n_1 + 1, n_1 + 2, \dots, n_1 + n_2 = n_a\}$. Let $\bar{\mathbf{P}}^{S_1}$ and $\bar{\mathbf{r}}^{S_1}$ be matrix $\bar{\mathbf{P}}^a$ and vector $\bar{\mathbf{r}}^a$ computed by KBSF using only the n_1 transitions in S_1 (if $n_1 = 0$, we define $\bar{\mathbf{P}}^{S_1} = \mathbf{0} \in \mathbb{R}^{m \times m}$ and $\bar{\mathbf{r}}^{S_1} = \mathbf{0} \in \mathbb{R}^m$ for all $a \in A$). We want to compute $\bar{\mathbf{P}}^{S_1 \cup S_2}$ and $\bar{\mathbf{r}}^{S_1 \cup S_2}$ from $\bar{\mathbf{P}}^{S_1}$, $\bar{\mathbf{r}}^{S_1}$, and S_2 , without using the set of sample transitions S_1 .

We start with the transition matrices $\bar{\mathbf{P}}^a$. We know that

$$\begin{aligned} \bar{p}_{ij}^{S_1} &= \sum_{t=1}^{n_1} k_{it}^a d_{tj}^a = \sum_{t=1}^{n_1} \frac{k_\tau(\bar{s}_i, s_t^a)}{\sum_{l=1}^{n_1} k_\tau(\bar{s}_i, s_l^a)} \frac{\bar{k}_\tau(\hat{s}_t^a, \bar{s}_j)}{\sum_{l=1}^m k_\tau(\hat{s}_t^a, \bar{s}_l)} \\ &= \frac{1}{\sum_{l=1}^{n_1} k_\tau(\bar{s}_i, s_l^a)} \sum_{t=1}^{n_1} \frac{k_\tau(\bar{s}_i, s_t^a) \bar{k}_\tau(\hat{s}_t^a, \bar{s}_j)}{\sum_{l=1}^m k_\tau(\hat{s}_t^a, \bar{s}_l)}. \end{aligned}$$

To simplify the notation, define

$$z_i^{S_1} = \sum_{l=1}^{n_1} k_\tau(\bar{s}_i, s_l^a), \quad z_i^{S_2} = \sum_{l=n_1+1}^{n_1+n_2} k_\tau(\bar{s}_i, s_l^a), \quad \text{and} \quad b_{ij}^t = \frac{k_\tau(\bar{s}_i, s_t^a) \bar{k}_\tau(\hat{s}_t^a, \bar{s}_j)}{\sum_{l=1}^m k_\tau(\hat{s}_t^a, \bar{s}_l)},$$

with $t \in \{1, 2, \dots, n_1 + n_2\}$. Then, we can write

$$\bar{p}_{ij}^{S_1 \cup S_2} = \frac{1}{z_i^{S_1} + z_i^{S_2}} \left(\sum_{t=1}^{n_1} b_{ij}^t + \sum_{t=n_1+1}^{n_1+n_2} b_{ij}^t \right) = \frac{1}{z_i^{S_1} + z_i^{S_2}} \left(\bar{p}_{ij}^{S_1} z_i^{S_1} + \sum_{t=n_1+1}^{n_1+n_2} b_{ij}^t \right).$$

Now, defining $b_{ij}^{S_2} = \sum_{t=n_1+1}^{n_1+n_2} b_{ij}^t$, we have the simple update rule:

$$\boxed{\bar{p}_{ij}^{S_1 \cup S_2} = \frac{1}{z_i^{S_1} + z_i^{S_2}} \left(b_{ij}^{S_2} + \bar{p}_{ij}^{S_1} z_i^{S_1} \right)}. \quad (19)$$

We can apply similar reasoning to derive an update rule for the rewards \bar{r}_i^a . We know that

$$\bar{r}_i^{S_1} = \frac{1}{\sum_{l=1}^{n_1} k_\tau(\bar{s}_i, s_l^a)} \sum_{t=1}^{n_1} k_\tau(\bar{s}_i, s_t^a) r_t^a = \frac{1}{z_i^{S_1}} \sum_{t=1}^{n_1} k_\tau(\bar{s}_i, s_t^a) r_t^a.$$

Let $e_i^t = k_\tau(\bar{s}_i, s_t^a) r_t^a$, with $t \in \{1, 2, \dots, n_1 + n_2\}$. Then,

$$\bar{r}_i^{S_1 \cup S_2} = \frac{1}{z_i^{S_1} + z_i^{S_2}} \left(\sum_{t=1}^{n_1} e_i^t + \sum_{t=n_1+1}^{n_1+n_2} e_i^t \right) = \frac{1}{z_i^{S_1} + z_i^{S_2}} \left(z_i^{S_1} \bar{r}_i^{S_1} + \sum_{t=n_1+1}^{n_1+n_2} e_i^t \right).$$

Defining $e_i^{S_2} = \sum_{t=n_1+1}^{n_1+n_2} e_i^t$, we have the following update rule:

$$\boxed{\bar{r}_i^{S_1 \cup S_2} = \frac{1}{z_i^{S_1} + z_i^{S_2}} \left(e_i^{S_2} + \bar{r}_i^{S_1} z_i^{S_1} \right)}. \quad (20)$$

Since $b_{ij}^{S_2}$, $e_i^{S_2}$, and $z_i^{S_2}$ can be computed based on S_2 only, we can discard the sample transitions in S_1 after computing $\bar{\mathbf{P}}^{S_1}$ and $\bar{\mathbf{r}}^{S_1}$. In order to do that, we only have to keep the variables $z_i^{S_1}$. These variables can be stored in $|A|$ vectors $\mathbf{z}^a \in \mathbb{R}^m$, resulting in a modest memory overhead. Note that we can apply the ideas above recursively, further splitting the sets S_1 and S_2 in subsets of smaller size. Thus, we have a fully incremental way of computing KBSF's MDP which requires almost no extra memory.

Algorithm 2 shows a step-by-step description of how to update \bar{M} based on a set of sample transitions. Using this method to update its model, KBSF's space complexity drops from $O(\hat{n}m)$ to $O(m^2)$. Since the amount of memory used by KBSF is now independent of n , it can process an arbitrary number of sample transitions (or, more precisely, the limit on the amount of data it can process is dictated by time only, not space).

Instead of assuming that S_1 and S_2 are a partition of a fixed data set S^a , we can consider that S_2 was generated based on the policy learned by KBSF using the transitions in S_1 . Thus, Algorithm 2 provides a flexible framework for integrating learning and planning within KBSF. Specifically, our algorithm can cycle between learning a model of the problem based on sample transitions, using such a model to derive a policy, and resorting to this policy to collect more data. Algorithm 3 shows a possible implementation of this framework. In order to distinguish it from its batch counterpart, we will call the incremental version of our algorithm *i*KBSF. *i*KBSF updates the model \bar{M} and the value function $\bar{\mathbf{Q}}$ at fixed intervals t_m and t_v , respectively. When $t_m = t_v = n$, we recover the batch version of KBSF; when $t_m = t_v = 1$, we have a fully on-line method which stores no sample transitions.

Algorithm 3 allows for the inclusion of new representative states to the model \bar{M} . Using Algorithm 2 this is easy to do: given a new representative state \bar{s}_{m+1} , it suffices to set $z_{m+1}^a = 0$, $\bar{r}_{m+1}^a = 0$, and $\bar{p}_{m+1,j} = \bar{p}_{j,m+1} = 0$ for $j = 1, 2, \dots, m+1$ and all $a \in A$. Then, in the following applications of update rules (19) and (20), the dynamics of \bar{M} will naturally reflect the existence of state \bar{s}_{m+1} . Note that the inclusion of new representative states

Algorithm 2 Update KBSF's MDP

Input: $\bar{\mathbf{P}}^a, \bar{\mathbf{r}}^a, \mathbf{z}^a$ for all $a \in A$ ▷ Current model
 $S^a = \{(s_k^a, r_k^a, \hat{s}_k^a) | k = 1, 2, \dots, n_a\}$ for all $a \in A$ ▷ Sample transitions
Output: Updated \bar{M} and \mathbf{z}^a

for $a \in A$ **do**
 for $t = 1, \dots, n_a$ **do** $\bar{z}_t \leftarrow \sum_{l=1}^m \bar{k}_{\bar{\tau}}(\hat{s}_t^a, \bar{s}_l)$
 $n_a \leftarrow |S^a|$
 for $i = 1, 2, \dots, m$ **do**
 $z'_i \leftarrow \sum_{t=1}^{n_a} k_{\tau}(\bar{s}_i, s_t^a)$
 for $j = 1, 2, \dots, m$ **do**
 $b \leftarrow \sum_{t=1}^{n_a} k_{\tau}(\bar{s}_i, s_t^a) \bar{k}_{\bar{\tau}}(\hat{s}_t^a, \bar{s}_j) / \bar{z}_t$
 $\bar{p}_{ij} \leftarrow \frac{1}{z_i^a + z'_i} (b + \bar{p}_{ij} z_i^a)$ ▷ Update transition probabilities
 $e \leftarrow \sum_{t=1}^{n_a} k_{\tau}(\bar{s}_i, s_t^a) r_t^a$
 $\bar{r}_i \leftarrow \frac{1}{z_i^a + z'_i} (e + \bar{r}_i z_i^a)$ ▷ Update rewards
 $z_i^a \leftarrow z_i^a + z'_i$ ▷ Update normalization factor

Algorithm 3 Incremental KBSF (*i*KBSF)

$\bar{S} = \{\bar{s}_1, \bar{s}_2, \dots, \bar{s}_m\}$ ▷ Set of representative states
Input: t_m ▷ Interval to update model
 t_v ▷ Interval to update value function
Output: Approximate value function $\tilde{Q}(s, a)$

$\bar{\mathbf{P}}^a \leftarrow \mathbf{0} \in \mathbb{R}^{m \times m}, \bar{\mathbf{r}}^a \leftarrow \mathbf{0} \in \mathbb{R}^m, \mathbf{z}^a \leftarrow \mathbf{0} \in \mathbb{R}^m$, for all $a \in A$
 $\bar{\mathbf{Q}} \leftarrow$ arbitrary matrix in $\mathbb{R}^{m \times |A|}$
 $s \leftarrow$ initial state
 $a \leftarrow$ random action
for $t \leftarrow 1, 2, \dots$ **do**
 Execute a in s and observe r and \hat{s}
 $S^a \leftarrow S^a \cup \{(s, r, \hat{s})\}$
 if $(t \bmod t_m = 0)$ **then** ▷ Update model
 Add new representative states to \bar{M} using S^a ▷ This step is optional
 Update \bar{M} and \mathbf{z}^a using Algorithm 2 and S^a
 $S^a \leftarrow \emptyset$ for all $a \in A$ ▷ Discard transitions
 if $(t \bmod t_v = 0)$ update $\bar{\mathbf{Q}}$ ▷ Update value function
 $s \leftarrow \hat{s}$
 Select a based on $\tilde{Q}(s, a) = \sum_{i=1}^m \bar{k}_{\bar{\tau}}(s, \bar{s}_i) \bar{q}_{ia}$

does not destroy the information already in the model. This allows *i*KBSF to refine its approximation on the fly, as needed. One can think of several ways of detecting the need for new representative states. A simple strategy, based on Proposition 6, is to impose a maximum distance allowed between a sampled state \hat{s}_i^a and the nearest representative state, $\text{dist}(\hat{s}_i^a, 1)$. So, anytime the agent encounters a new state \hat{s}_i^a for which $\text{dist}(\hat{s}_i^a, 1)$ is above a given threshold, \hat{s}_i^a is added to the model as \bar{s}_{m+1} . In Section 5.2 we report experiments with *i*KBSF using this approach. Before that, though, we discuss the theoretical properties of the incremental version of our algorithm.

5.1 Theoretical results

As discussed, *i*KBSF does not need to store sample transitions to build its approximation. However, the computation of $\tilde{Q}(s, a)$ through (12) requires all the tuples $(s_i^a, r_i^a, \hat{s}_i^a)$ to be available. In some situations, it may be feasible to keep the transitions in order to compute $\tilde{Q}(s, a)$. However, if we want to use *i*KBSF to its full extend, we need a way of computing $\tilde{Q}(s, a)$ without using the sample transitions. This is why upon reaching state s at time step t *i*KBSF selects the action to be performed based on

$$\tilde{Q}_t(s, a) = \sum_{i=1}^m \bar{\kappa}_t(s, \bar{s}_i) \bar{Q}_t(\bar{s}_i, a), \quad (21)$$

where $\bar{Q}_t(\bar{s}_i, a)$ is the action-value function available to *i*KBSF at the t^{th} iteration (see Algorithm 3). Note that we do not assume that *i*KBSF has computed the optimal value function of its current model \bar{M}_t —that is, it may be the case that $\bar{Q}_t(\bar{s}_i, a) \neq \bar{Q}_t^*(\bar{s}_i, a)$.

Unfortunately, when we replace (12) with (21) Proposition 3 no longer applies. In this section we address this issue by deriving an upper bound for the difference between $\tilde{Q}_t(s, a)$ and $\hat{Q}_t(s, a)$, the action-value function that would be computed by KBRL using all the transitions processed by *i*KBSF up to time step t . In order to derive our bound, we assume that *i*KBSF uses a fixed set \bar{S} —meaning that no representative states are added to the model \bar{M} —and that it never stops refining its model, doing so at every iteration t (i.e., $t_m = 1$ in Algorithm 3). We start by showing the following lemma:

Lemma 9 *Let $M \equiv (S, A, \mathbf{P}^a, \mathbf{r}^a, \gamma)$ and $\tilde{M} \equiv (S, A, \tilde{\mathbf{P}}^a, \tilde{\mathbf{r}}^a, \gamma)$ be two finite MDPs. Then, for any $s \in S$ and any $a \in A$,*

$$|Q^*(s, a) - \tilde{Q}^*(s, a)| \leq \frac{1}{1 - \gamma} \max_a \|\mathbf{r}^a - \tilde{\mathbf{r}}^a\|_\infty + \frac{\gamma}{2(1 - \gamma)^2} R_{\text{dif}} \max_a \|\mathbf{P}^a - \tilde{\mathbf{P}}^a\|_\infty,$$

where $R_{\text{dif}} = \max_{a,i} r_i^a - \min_{a,i} r_i^a$.

Lemma 9 provides an upper bound for the difference in the action-value functions of any two MDPs having the same state space S , action space A , and discount factor γ .⁴ Our strategy will be to use this result to bound the error introduced by the application of the stochastic-factorization trick in the context of *i*KBSF.

4. Strehl and Littman’s (2008) Lemma 1 is similar to our result. Their bound is more general than ours, as it applies to any Q^π , but it is also slightly looser.

When $t_m = 1$, at any time step t *iKBSF* has a model \bar{M}_t built based on the t transitions observed thus far. As shown in the beginning of this section, \bar{M}_t exactly matches the model that would be computed by batch KBSF using the same data and the same set of representative states. Thus, we can think of matrices $\bar{\mathbf{P}}_t^a$ and vectors $\bar{\mathbf{r}}_t^a$ available at the t^{th} iteration of *iKBSF* as the result of the stochastic-factorization trick applied with matrices \mathbf{D}_t and \mathbf{K}_t^a . Although *iKBSF* does not explicitly compute such matrices, they serve as a theoretical foundation to build our result on.

Proposition 10 *Suppose *iKBSF* is executed with a fixed set of representative states \bar{S} using $t_m = 1$. Let \mathbf{D}_t , \mathbf{K}_t^a and $\bar{\mathbf{r}}_t^a$ be the matrices and the vector (implicitly) computed by this algorithm at iteration t . Then, if s is the state encountered by *iKBSF* at time step t ,*

$$|\hat{Q}_t(s, a) - \tilde{Q}_t(s, a)| \leq \frac{1}{1-\gamma} \max_a \|\hat{\mathbf{r}}_t^a - \mathbf{D}_t \bar{\mathbf{r}}_t^a\|_\infty + \frac{\bar{R}_{\text{dif},t}}{(1-\gamma)^2} \left(\frac{\gamma}{2} \max_a \|\hat{\mathbf{P}}_t^a - \mathbf{D}_t \mathbf{K}_t^a\|_\infty + \sigma(\mathbf{D}_t) \right) + \epsilon_{\bar{Q}_t}, \quad (22)$$

for any $a \in A$, where \tilde{Q}_t is the value function computed by *iKBSF* at time step t through (21), \hat{Q}_t is the value function computed by KBRL through (7) based on the same data, $\bar{R}_{\text{dif},t} = \max_{a,i} \bar{r}_{i,t}^a - \min_{a,i} \bar{r}_{i,t}^a$, $\sigma(\mathbf{D}_t) = \max_i (1 - \max_j d_{ij,t})$, and $\epsilon_{\bar{Q}_t} = \max_{i,a} |\bar{Q}_t^*(\bar{s}_i, a) - \bar{Q}_t(\bar{s}_i, a)|$.

Proposition 10 shows that, at any time step t , the error in the action-value function computed by *iKBSF* is bounded above by the quality and the level of stochasticity of the stochastic factorization implicitly computed by the algorithm. The term $\epsilon_{\bar{Q}_t}$ accounts for the possibility that *iKBSF* has not computed the optimal value function of its model at step t , either because $t_m \neq t_v$ or because the update of $\bar{\mathbf{Q}}$ in Algorithm 3 is not done to completion (for example, one can apply the Bellman operator \bar{T} a fixed number of times, stopping short of convergence). The restriction $t_m = 1$ is not strictly necessary if we are willing to compare $\hat{Q}_t(s, a)$ with $\hat{Q}_{t'}(s, a)$, where $t' = \lfloor (t + t_m)/t \rfloor$ (the next time step scheduled for a model update). However, such a result would be somewhat circular, since the sample transitions used to build $\hat{Q}_{t'}(s, a)$ may depend on $\hat{Q}_t(s, a)$. Finally, we note that, given the similarity between the right-hand sides of (8) and (22), it should be trivial to derive results analogous to Propositions 6 and 8 that also apply to *iKBSF*.

5.2 Empirical results

We now look at the empirical performance of the incremental version of KBSF. Following the structure of Section 4.2, we start with the puddle world task to show that *iKBSF* is indeed able to match the performance of batch KBSF without storing all sample transitions. Next we exploit the scalability of *iKBSF* to solve two difficult control tasks, triple pole-balancing and helicopter hovering. We also compare *iKBSF*'s performance with that of other reinforcement learning algorithms.

5.2.1 PUDDLE WORLD (PROOF OF CONCEPT)

We use the puddle world problem as a proof of concept (Sutton, 1996). In this first experiment we show that *iKBSF* is able to recover the model that would be computed by its batch counterpart. In order to do so, we applied Algorithm 3 to the puddle-world task using a random policy to select actions.

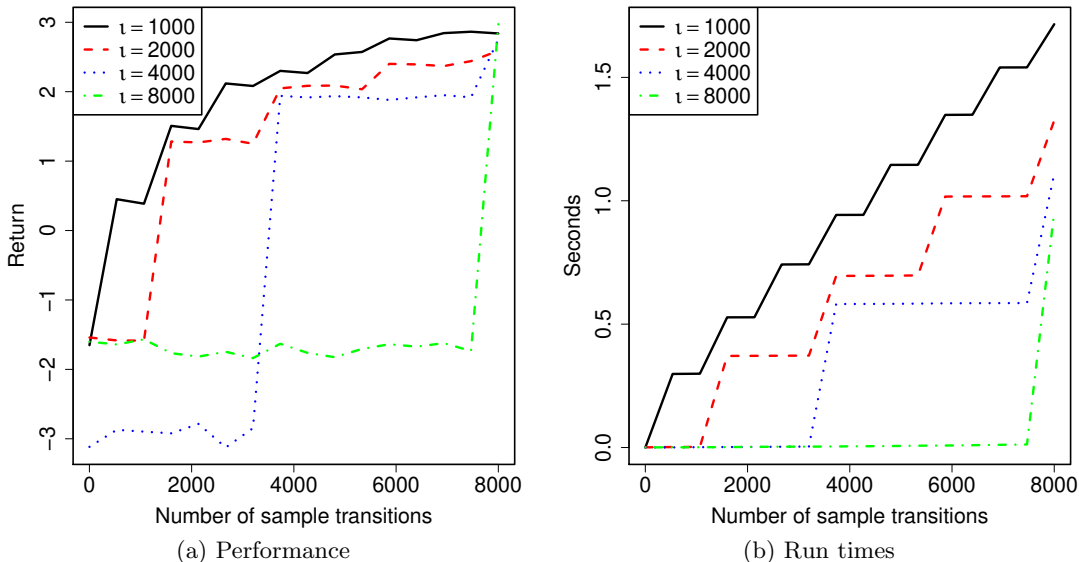


Figure 7: Results on the puddle-world task averaged over 50 runs. KBSF used 100 representative states evenly distributed over the state space and $t_m = t_v = \iota$ (see legends). Sample transitions were collected by a random policy. The agents were tested on two sets of states surrounding the “puddles” (see Appendix B).

Figure 7 shows the result of the experiment when we vary the parameters t_m and t_v . Note that the case in which $t_m = t_v = 8000$ corresponds to the batch version of KBSF, whose results on the puddle world are shown in Figure 3. As expected, the performance of KBSF policies improves gradually as the algorithm goes through more sample transitions, and in general the intensity of the improvement is proportional to the amount of data processed. More important, the performance of the decision policies after all sample transitions have been processed is essentially the same for all values of t_m and t_v , which confirms that *i*KBSF can be used as an instrument to circumvent KBSF’s memory demand. Thus, if one has a batch of sample transitions that does not fit in the available memory, it is possible to split the data in chunks of smaller sizes and still get the same value-function approximation that would be computed if the entire data set were processed at once. As shown in Figure 7b, there is only a small computational overhead associated with such a strategy (this results from unnormalizing and normalizing the elements of $\bar{\mathbf{P}}^a$ and $\bar{\mathbf{r}}^a$ multiple times through update rules (19) and (20)).

5.2.2 TRIPLE POLE-BALANCING (COMPARISON WITH FITTED Q -ITERATION)

As discussed in Section 4.2.2, the pole balancing task has been addressed in several different versions, and among them simultaneously balancing two poles is particularly challenging (Wieland, 1991). Figures 4c and 4d show that the batch version of KBSF was able to satisfactorily solve the double pole-balancing task. In order to show the scalability of the incremental version of our algorithm, in this section we raise the bar, adding a third pole

to the problem. We perform our simulations using the parameters usually adopted with the two-pole problem, with the extra pole having the same length and mass as the longer pole (Gomez, 2003, see Appendix B). This results in a difficult control problem with an 8-dimensional state space \mathbb{S} .

In our experiments with KBSF on the two-pole task we used 200 representative states and 10^6 sample transitions collected by a random policy. Here we start our experiment with triple pole-balancing using exactly the same configuration, and then we let i KBSF refine its model \bar{M} by incorporating more sample transitions through update rules (19) and (20). We also let i KBSF grow its model if necessary. Specifically, a new representative state is added to \bar{M} on-line every time the agent encounters a sample state \hat{s}_i^a for which $\bar{k}_\tau(\hat{s}_i^a, \bar{s}_j) < 0.01$ for all $j \in 1, 2, \dots, m$. This corresponds to setting a maximum allowed distance from a sampled state to the closest representative state, $\max_{a,i} \text{dist}(\hat{s}_i^a, 1)$.

Given the poor performance of LSPI on the double pole-balancing task, shown in Figures 4c and 4d, on the three-pole version of the problem we only compare KBSF with FQIT. We used FQIT with the same configuration adopted in Sections 4.2.3 and 4.2.4, with the parameter η_{\min} varying in the set $\{10000, 1000, 100\}$. As for KBSF, the widths of the kernels were fixed at $\tau = 100$ and $\bar{\tau} = 1$ and sparse kernels were used ($\mu = 50$ and $\bar{\mu} = 10$).

In order to show the benefits provided by the incremental version of our algorithm, we assumed that both KBSF and FQIT could store at most 10^6 sample transitions in memory. In the case of i KBSF, this is not a problem, since we can always split the data in subsets of smaller size and process them incrementally. Here, we used Algorithm 3 with a 0.3-greedy policy, $t_m = t_v = 10^6$, and $n = 10^7$. In the case of FQIT, we have two options to circumvent the limited amount of memory available. The first one is to use a single batch of 10^6 sample transitions. The other option is to use the initial batch of transitions to compute an approximation of the problem’s value function, then use an 0.3-greedy policy induced by this approximation to collect a second batch, and so on. Here we show the performance of FQIT using both strategies.

We first compare the performance of i KBSF with that of FQIT using a single batch of sample transitions. This is shown in Figure 8a and 8b. For reference, we also show the results of batch KBSF—that is, we show the performance of the policy that would be computed by our algorithm if we did not have a way of computing its approximation incrementally. As shown in Figure 8a, both FQIT and batch KBSF perform poorly in the triple pole-balancing task, with average success rates below 55%. These results suggest that the amount of data used by these algorithms is insufficient to describe the dynamics of the control task. Of course, we could give more sample transitions to FQIT and batch KBSF. Note however that, since they are batch learning methods, there is an inherent limit on the amount of data that these algorithms can use to construct their approximation. In contrast, the amount of memory required by i KBSF is independent of the number of sample transitions n . This fact together with the fact that KBSF’s computational complexity is only linear in n allow our algorithm to process a large amount of data in reasonable time. This can be clearly observed in Figure 8b, which shows that i KBSF can build an approximation using 10^7 sample transitions in under 20 minutes. As a reference for comparison, FQIT(1000) took an average of 1 hour and 18 minutes to process 10 times less data.

As shown in Figure 8a, i KBSF’s ability to process a large number of sample transitions allows our algorithm to achieve a success rate of approximately 80%. This is similar to the

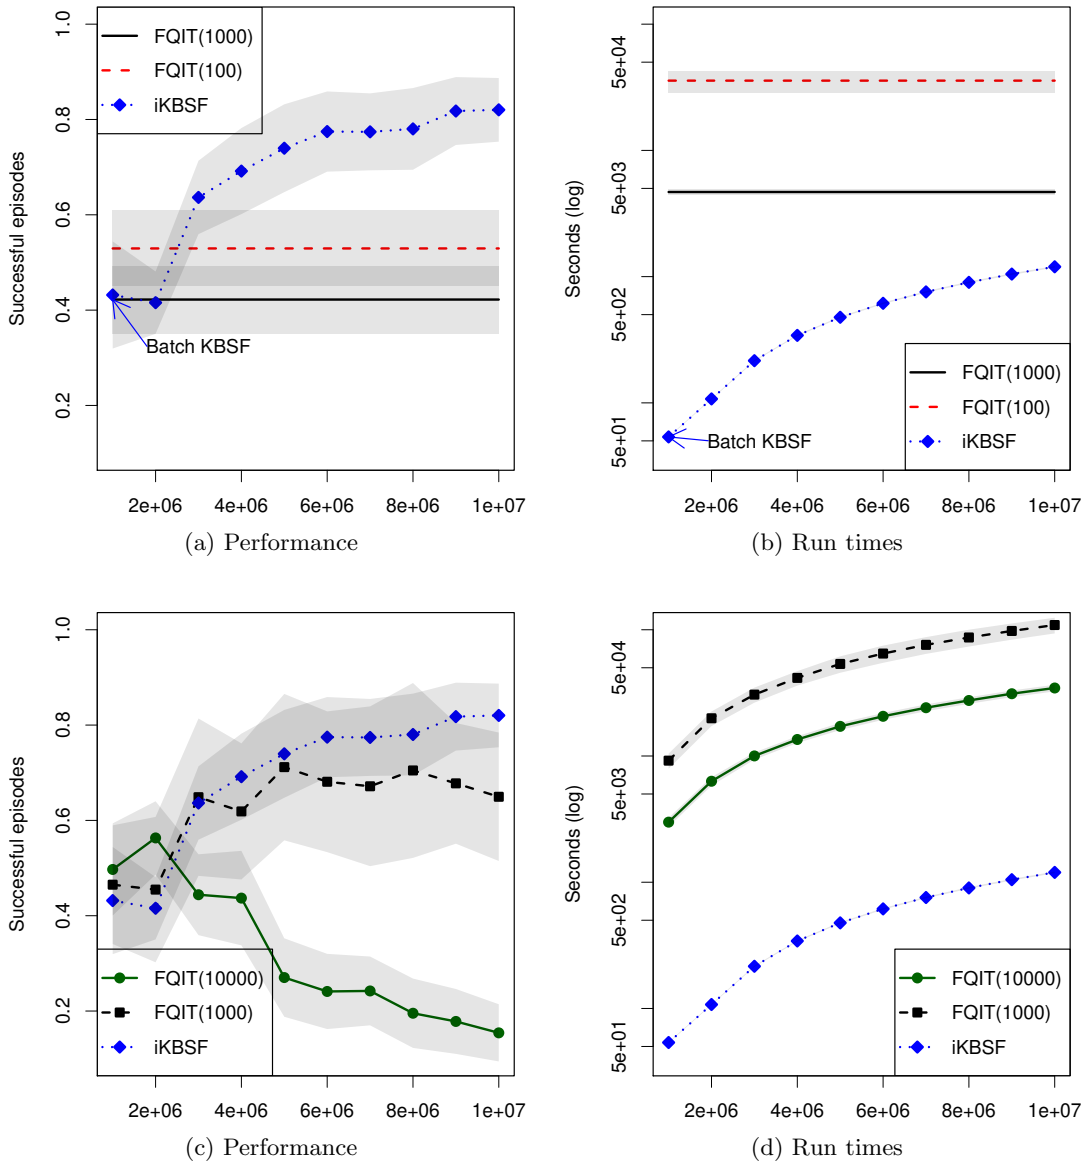


Figure 8: Results on the triple pole-balancing task, as a function of the number of sample transitions n , averaged over 50 runs. Figures 8a and 8b show results of FQIT when using a single batch of 10^6 transitions; in Figures 8c and 8d ten sets of 10^6 transitions were used in sequence by the algorithm (see text for details). The values correspond to the fraction of episodes initiated from the test states in which the 3 poles could be balanced for 3000 steps (one minute of simulated time). The test sets were regular grids of 256 cells defined over the hypercube centered at the origin and covering 50% of the state-space axes in each dimension (see Appendix B for details). Shaded regions represent 99% confidence intervals.




Figure 9: Number of representative states used by *i*KBSF on the triple pole-balancing task. Results were averaged over 50 runs (99% confidence intervals are almost imperceptible in the figure).

performance of batch KBSF on the two-pole version of the problem (*cf.* Figure 4). The good performance of *i*KBSF on the triple pole-balancing task is especially impressive when we recall that the decision policies were evaluated on a set of test states representing all possible directions of inclination of the three poles. In order to achieve the same level of performance with KBSF, approximately 2 Gb of memory would be necessary, even using sparse kernels, whereas *i*KBSF used less than 0.03 Gb of memory.

One may argue that the above comparison between FQIT and KBSF is not fair, since the latter used ten times the amount of data used by the former. Thus, in Figures 8c and 8d we show the results of FQIT using 10 batches of 10^6 transitions—exactly the same number of transitions processed by *i*KBSF. Here we cannot compare *i*KBSF with FQIT(100) because the computational cost of the tree-based approach is prohibitively large (it would take over 4 days only to train a single agent, not counting the test phase). When we look at the other instances of the algorithm, we see two opposite trends. Surprisingly, the extra sample transitions actually made the performance of FQIT(10000) *worse*. On the other hand, FQIT(1000) performs significantly better using more data, though still not as well as *i*KBSF (both in terms of performance and computing time).

To conclude, observe in Figure 9 how the number of representative states m grows as a function of the number of sample transitions processed by KBSF. As expected, in the beginning of the learning process m grows fast, reflecting the fact that some relevant regions of the state space have not been visited yet. As more and more data come in, the number of representative states starts to stabilize.

5.2.3 HELICOPTER HOVERING TASK (COMPARISON WITH SARSA)

In the previous two sections we showed how *i*KBSF can be used to circumvent the inherent memory limitations of batch learning. We now show how our algorithm performs in a fully

on-line regime. For that, we focus on a challenging reinforcement learning task in which the goal is to control an autonomous helicopter.

Helicopters have unique control capabilities, such as low speed flight and in-place hovering, that make them indispensable instruments in many contexts. Such flexibility comes at a price, though: it is widely recognized that a helicopter is significantly harder to control than a fixed-wing aircraft (Ng et al., 2003; Abbeel et al., 2007). Part of this difficulty is due to the complex dynamics of the helicopter, which is not only non-linear, noisy, and asymmetric, but also counterintuitive in some aspects (Ng et al., 2003).

An additional complication of controlling an autonomous helicopter is the fact that a wrong action can easily lead to a crash, which is both dangerous and expensive. Thus, the usual practice is to first develop a model of the helicopter’s dynamics and then use the model to design a controller (Ng et al., 2003). Here we use the model constructed by Abbeel et al. (2005) based on data collected on actual flights of an XCell Tempest helicopter (see Appendix B). The resulting reinforcement learning problem has a 12-dimensional state space whose variables represent the aircraft’s position, orientation, and the corresponding velocities and angular velocities along each axis.

In the version of the task considered here the goal is to keep the helicopter hovering as close as possible to a fixed position. All episodes start at the target location, and at each time step the agent receives a negative reward proportional to the distance from the current state to the desired position. Because the tail rotor’s thrust exerts a sideways force on the helicopter, the aircraft cannot be held stationary in the zero-cost state even in the absence of wind. The episode ends when the helicopter leaves the hover regime, that is, when any of the state’s variables exceeds pre-specified thresholds.

The helicopter is controlled via a 4-dimensional continuous vector whose variables represent the longitudinal cyclic pitch, the latitudinal cyclic pitch, the tail rotor collective pitch, and the main rotor collective pitch. By adjusting the value of these variables the pilot can rotate the helicopter around its axes and control the thrust generated by the main rotor. Since KBSF was designed to deal with a finite number of actions, we discretized the set A using 4 values per dimension, resulting in 256 possible actions. The details of the discretization process are given below.

Here we compare i KBSF with the SARSA(λ) algorithm using tile coding for value function approximation (Rummery and Niranjan, 1994, Sutton, 1996—see Appendix B). We applied SARSA with $\lambda = 0.05$, a learning rate of 0.001, and 24 tilings containing 4^{12} tiles each. Except for λ , all the parameters were adjusted in a set of preliminary experiments in order to improve the performance of the SARSA agent. We also defined the action-space discretization based on SARSA’s performance. In particular, instead of partitioning each dimension in equally-sized intervals, we spread the break points unevenly along each axis in order to maximize the return obtained by the SARSA agent. The result of this process is described in Appendix B. The interaction of the SARSA agent with the helicopter hovering task was dictated by an ϵ -greedy policy. Initially we set $\epsilon = 1$, and at every 50000 transitions the value of ϵ was decreased in 30%.

The i KBSF agent collected sample transitions using the same exploration regime. Based on the first batch of 50000 transitions, $m = 500$ representative states were determined by the k -means algorithm. No representative states were added to i KBSF’s model after that.

Both the value function and the model were updated at fixed intervals of $t_v = t_m = 50000$ transitions. We fixed $\tau = \bar{\tau} = 1$ and $\mu = \bar{\mu} = 4$.


Figure 10 shows the results obtained by SARSA and KBSF on the helicopter hovering task. Note in Figure 10a how the average episode length increases abruptly at the points in which the value of ϵ is decreased. This is true for both SARSA and KBSF. Also, since the number of steps executed per episode increases over time, the interval in between such abrupt changes decreases in length, as expected. Finally, observe how the performance of both agents stabilizes after around 70000 episodes, probably because at this point there is almost no exploration taking place anymore.

When we compare KBSF and SARSA, it is clear that the former significantly outperforms the latter. Specifically, after the cut-point of 70000 episodes, the KBSF agent executes approximately 2.25 times the number of steps performed by the SARSA agent before crashing. Looking at Figures 10a and 10b, one may argue at first that there is nothing surprising here: being a model-based algorithm, KBSF is more sample efficient than SARSA, but it is also considerably slower (Atkeson and Santamaria, 1997). Notice though that the difference between the run times of SARSA and KBSF shown in Figure 10b is in part a consequence of the good performance of the latter: since KBSF is able to control the helicopter for a larger number of steps, the corresponding episodes will obviously take longer. A better measure of the algorithms’s computational cost can be seen in Figure 10c, which shows the average time taken by each method to perform one transition. Observe how KBSF’s computing time peaks at the points in which the model and the value function are updated. In the beginning KBSF’s MDP changes considerably, and as a result the value function updates take longer. As more data come in, the model starts to stabilize, accelerating the computation of \bar{Q}^* (we “warm start” policy iteration with the value function computed in the previous round). At this point, KBSF’s computational cost per step is only slightly higher than SARSA’s, even though the former computes a model of the environment while the latter directly updates the value function approximation.


To conclude, we note that our objective in this section was exclusively to show that KBSF can outperform a well-known on-line algorithm with compatible computational cost. Therefore, we focused on the comparison of the algorithms rather than on obtaining the best possible performance on the task. Also, it is important to mention that more difficult versions of the helicopter task have been addressed in the literature, usually using domain knowledge in the configuration of the algorithms or to guide the collection of data (Ng et al., 2003; Abbeel et al., 2007). Since our focus here was on evaluating the on-line performance of KBSF, we addressed the problem in its purest form, without using any prior information to help the algorithms solve the task (see Asbah et al.’s paper for experiments on the helicopter hovering task with a more specialized version of KBSF, 2013).

6. Discussion


In this section we deepen our discussion about KBSF. We start with an analysis of the approximation computed by our algorithm, in which we draw interesting connections with KBRL, and then we proceed to discuss some issues that come up when one uses KBSF in practice.



(a) Performance



(b) Run time



(c) Average time per step (time of an episode divided by the number of steps)

Figure 10: Results on the helicopter hovering task averaged over 50 runs. The learned controllers were tested from a fixed state (see text for details). The shadowed regions represent 99% confidence intervals.

6.1 A closer look at KBSF’s approximation

As outlined in Section 2, KBRL defines the probability of a transition from state \hat{s}_i^b to state \hat{s}_k^a as being $\kappa_\tau^a(\hat{s}_i^b, s_k^a)$, where $a, b \in A$ (see Figure 2a). Note that the kernel κ_τ^a is computed with the initial state s_k^a , and not \hat{s}_k^a itself. The intuition behind this is simple: since we know the transition $s_k^a \xrightarrow{a} \hat{s}_k^a$ has occurred before, the more “similar” \hat{s}_i^b is to s_k^a , the more likely the transition $\hat{s}_i^b \xrightarrow{a} \hat{s}_k^a$ becomes (Ormoneit and Sen, 2002).

From (10), it is clear that the computation of matrices \mathbf{K}^a performed by KBSF follows the same reasoning underlying the computation of KBRL’s matrices $\hat{\mathbf{P}}^a$; in particular, $\kappa_\tau^a(\bar{s}_j, s_k^a)$ gives the probability of a transition from \bar{s}_j to \hat{s}_k^a . However, when we look at matrix \mathbf{D} things are slightly different: here, the probability of a “transition” from \hat{s}_i^b to representative state \bar{s}_j is given by $\bar{\kappa}_\tau(\hat{s}_i^b, \bar{s}_j)$ —a computation that involves \bar{s}_j itself. If we were to strictly adhere to KBRL’s logic when computing the transition probabilities to the representative states \bar{s}_j , the probability of transitioning from \hat{s}_i^b to \bar{s}_j upon executing action a should be a function of \hat{s}_i^b and a state s' from which we knew a transition $s' \xrightarrow{a} \bar{s}_j$ had occurred. In this case we would end up with one matrix \mathbf{D}^a for each action $a \in A$. Note though that this formulation of the method is not practical, because the computation of the matrices \mathbf{D}^a would require a transition $(\cdot) \xrightarrow{a} \bar{s}_j$ for each $a \in A$ and each $\bar{s}_j \in \bar{S}$. Clearly, such a requirement is hard to fulfill even if we have a generative model available to generate sample transitions.

In this section we take a closer look at KBSF’s approximation, and provide two interpretations that support the way the matrices involved are built. We argue that KBSF can be seen as a kernel-based approximation of both the model and the value function computed by KBRL. Based on these interpretations we later discuss how KBSF can potentially be beneficial from a statistical point of view.

6.1.1 APPROXIMATING THE MODEL

We start by looking at how KBRL constructs its model. As shown in Figure 2a, for each action $a \in A$ the state \hat{s}_i^b has an associated stochastic vector $\hat{\mathbf{p}}_j^a \in \mathbb{R}^{1 \times n}$ whose nonzero entries correspond to the kernel $\kappa_\tau^a(\hat{s}_i^b, \cdot)$ evaluated at $s_k^a, k = 1, 2, \dots, n_a$. Since we are dealing with a continuous state space, it is possible to compute an analogous vector for any $s \in \mathbb{S}$ and any $a \in A$. Focusing on the nonzero entries of $\hat{\mathbf{p}}_j^a$, we define the function

$$\begin{aligned} \hat{\mathcal{P}}_{S^a} : \mathbb{S} &\mapsto \mathbb{R}^{1 \times n_a} \\ \hat{\mathcal{P}}_{S^a}(s) = \hat{\mathbf{p}}^a &\iff \hat{p}_i^a = \kappa_\tau^a(s, s_i^a) \text{ for } i = 1, 2, \dots, n_a. \end{aligned} \tag{23}$$

Clearly, full knowledge of the function $\hat{\mathcal{P}}_{S^a}$ allows for an exact computation of KBRL’s transition matrix $\hat{\mathbf{P}}^a$. Now suppose we do not know $\hat{\mathcal{P}}_{S^a}$ and we want to compute an approximation of this function in the points $\hat{s}_i^a \in S^a$, for all $a \in A$. Suppose further that we are only given a “training set” composed of m pairs $(\bar{s}_j, \hat{\mathcal{P}}_{S^a}(\bar{s}_j))$. One possible way of approaching this problem is to resort to kernel smoothing techniques. In this case, a particularly common choice is the so-called Nadaraya-Watson kernel-weighted estimator (Hastie et al., 2002, Chapter 6):

$$\bar{\mathcal{P}}_{S^a}(s) = \frac{\sum_{j=1}^m \bar{\kappa}_\tau(s, \bar{s}_j) \hat{\mathcal{P}}_{S^a}(\bar{s}_j)}{\sum_{j=1}^m \bar{\kappa}_\tau(s, \bar{s}_j)} = \sum_{j=1}^m \bar{\kappa}_\tau(s, \bar{s}_j) \hat{\mathcal{P}}_{S^a}(\bar{s}_j). \tag{24}$$

Contrasting the expression above with (10), we see that this is exactly how KBSF computes its approximation $\mathbf{D}\mathbf{K}^a \approx \hat{\mathbf{P}}^a$, with $\bar{\mathcal{P}}_{S^a}$ evaluated at the points $\hat{s}_i^b \in S^b$, $b = 1, 2, \dots, |A|$. In this case, $\bar{\kappa}_\tau(\hat{s}_i^b, \bar{s}_j)$ are the elements of matrix \mathbf{D} , and $\hat{\mathcal{P}}_{S^a}(\bar{s}_j)$ is the j^{th} row of matrix $\hat{\mathbf{K}}^a$. Thus, in some sense, KBSF uses KBRL’s own kernel approximation principle to compute a stochastic factorization of \hat{M} . One can easily extend the reasoning above to also include rewards by defining a function $\hat{\mathcal{M}}_{S^a} : \mathbb{S} \mapsto \mathbb{R}^{1 \times n_a + 1}$ such that $\hat{\mathcal{M}}_{S^a}(s) = [\hat{\mathcal{P}}_{S^a}(s), \hat{\mathcal{P}}_{S^a}(s)^\top \mathbf{r}^a]$.

The exposition above also makes it clear that it is possible to build the matrices $\hat{\mathbf{D}}^a$ using different sets of representative states \bar{S}^a (this corresponds to having $|A|$ training sets composed of pairs $(\bar{s}_j^a, \hat{\mathcal{P}}_{S^a}(\bar{s}_j^a))$). As long as all the sets \bar{S}^a have the same cardinality m , the stochastic-factorization trick can still be applied. However, in order for the approximation computed by KBSF to make sense, we would have to have one matrix \mathbf{D}^a for each action $a \in A$ (see in Figure 2b how matrix \mathbf{D} is computed). Unfortunately, when different matrices \mathbf{D}^a are used, our “core” theoretical result, Proposition 2, no longer applies. Although alternative error bounds are possible, as shown by Barreto (2014), they are in general significantly looser than (8). In any case, the extension of KBSF to use different sets of representative states \bar{S}^a may be an interesting topic for future research.

6.1.2 APPROXIMATING THE VALUE FUNCTION

We now turn our attention to the way KBRL computes its value function, and show an intuitive way to derive KBSF’s update equations. Based on (4) and (5), we see that value iteration’s update rule for KBRL’s MDP, $\hat{\Delta}\Gamma$, is

$$\hat{Q}_{t+1}(\hat{s}_i^c, a) = \sum_{j=1}^{n_a} \kappa_\tau^a(\hat{s}_i^c, s_j^a) \left[r_j^a + \gamma \max_b \hat{Q}_t(\hat{s}_j^a, b) \right], \quad (25)$$

where $a, b, c \in A$. Note that at any time step t the equation above can be used to compute an approximation of the t -step value function over the entire state space \mathbb{S} ; after convergence to \hat{Q}^* equation (25) gives rise to (7).

From a computational point of view, the potential problem with rule (25) is the fact that updating \hat{Q} takes $O(n\hat{n}|A|)$ operations (recall that $n = \sum_a n_a$ and $\hat{n} = \max_a n_a$). KBRL makes it possible to compute an approximate solution for an MDP with continuous state space \mathbb{S} by only updating the value of a finite subset $\hat{S} \subset \mathbb{S}$. It is reasonable to ask whether similar strategy can come to the rescue when \hat{S} is itself too large. Specifically, if the values of the states $\hat{s}_i^a \in \hat{S}$ can be approximated based on the values of a small set of states $\bar{S} \subset \mathbb{S}$, with $|\bar{S}| = m$, then only the latter must be updated during dynamic programming’s iterative process. Following this reasoning, we can replace $\hat{Q}_t(\hat{s}_j^a, b)$ with $\sum_{l=1}^m \bar{\kappa}_\tau(\hat{s}_j^a, \bar{s}_l) \hat{Q}_t(\bar{s}_l, b)$ and rewrite (25) as

$$\begin{aligned} \hat{Q}'_{t+1}(\bar{s}_i, a) &= \sum_{j=1}^{n_a} \kappa_\tau^a(\bar{s}_i, s_j^a) \left[r_j^a + \gamma \max_b \sum_{l=1}^m \bar{\kappa}_\tau(\hat{s}_j^a, \bar{s}_l) \hat{Q}'_t(\bar{s}_l, b) \right] \\ &= \sum_{j=1}^{n_a} \kappa_\tau^a(\bar{s}_i, s_j^a) r_j^a + \gamma \sum_{j=1}^{n_a} \kappa_\tau^a(\bar{s}_i, s_j^a) \max_b \sum_{l=1}^m \bar{\kappa}_\tau(\hat{s}_j^a, \bar{s}_l) \hat{Q}'_t(\bar{s}_l, b). \end{aligned} \quad (26)$$

Updating \hat{Q}' through (26) is $O(m\hat{n}|A|)$. This is already an improvement over the computational complexity of (25), but the dependence on \hat{n} still precludes the use of (26) in

scenarios involving many sample transitions, such as on-line problems, for example. Now, if we *swap* the positions of the ‘max’ operator and the kernel $\bar{\kappa}_{\bar{\tau}}$ in (26), we can get rid of the dependence on the number of transitions, in the following way:

$$\begin{aligned} \bar{Q}_{t+1}(\bar{s}_i, a) &= \sum_{j=1}^{n_a} \kappa_{\bar{\tau}}^a(\bar{s}_i, s_j^a) r_j^a + \gamma \sum_{j=1}^{n_a} \kappa_{\bar{\tau}}^a(\bar{s}_i, s_j^a) \sum_{l=1}^m \bar{\kappa}_{\bar{\tau}}(\hat{s}_j^a, \bar{s}_l) \max_b \bar{Q}_t(\bar{s}_l, b) \\ &= \sum_{j=1}^{n_a} \kappa_{\bar{\tau}}^a(\bar{s}_i, s_j^a) r_j^a + \gamma \sum_{l=1}^m \sum_{j=1}^{n_a} \kappa_{\bar{\tau}}^a(\bar{s}_i, s_j^a) \bar{\kappa}_{\bar{\tau}}(\hat{s}_j^a, \bar{s}_l) \max_b \bar{Q}_t(\bar{s}_l, b) \\ &= \bar{r}_i^a + \gamma \sum_{l=1}^m \bar{p}_{il}^a \max_b \bar{Q}_t(\bar{s}_l, b). \end{aligned} \tag{27}$$

Updating \bar{Q} through (27) takes only $O(m^2|A|)$ operations. It is not difficult to see that (27) is value iteration’s update rule for the MDP \bar{M} defined by KBSF, $\bar{\Delta}\Gamma$ (see Algorithm 1).

Summarizing, if we use a local kernel approximation in the update rule of KBRL’s MDP, potentially introducing some error, we get (26), which is faster to compute than (25). If in addition we change the order in which the operations are applied, we get KBSF’s update rule (27), which is in some sense a greater deviation from (25) than (26), but is even faster to compute. Therefore, KBSF can be interpreted as a computationally efficient way of applying kernel approximation, the basic principle behind KBRL, to KBRL itself.

6.1.3 CURSE OF DIMENSIONALITY

In this paper we emphasized the role of KBSF as a technique to reduce KBRL’s computational cost. However, it is equally important to ask whether our algorithm provides benefits from a statistical point of view. In particular, instead of trying to approximate KBRL’s solution, it may be possible to compute *better* solutions using the same amount of data.

Ormoneit and Sen (2002) showed that, in general, the number of sample transitions needed by KBRL to achieve a certain approximation accuracy grows exponentially with the dimension of the state space—a phenomenon usually referred to as the “curse of dimensionality” (Bellman, 1961). As with other methods, the only way to avoid such an exponential dependency is to exploit some sort of regularity in the problem’s structure—paraphrasing Ormoneit and Sen (2002), one can only “break” the curse of dimensionality by incorporating prior knowledge into the approximation. In this section we argue that KBSF can be seen as a strategy to do so. In particular, the definition of the representative states can be seen as a practical mechanism to incorporate additional assumptions about the continuous MDP besides the sample transitions.

The fact that KBRL’s sample complexity is exponential in d_S is not surprising. As noted by Ernst et al. (2005), KBRL can be seen as a particular case of the fitted Q -iteration algorithm, which solves a reinforcement learning problem by breaking it into a succession of supervised learning problems. If the MDP defined by KBRL is solved through value iteration, for example, each iteration of this algorithm can be seen as a non-parametric kernel-based regression (Barreto, 2014). It is well known that non-parametric kernel methods suffer from the curse of dimensionality (Stone, 1982; Györfi et al., 2002).

As mentioned above, one can circumvent the curse of dimensionality associated with non-parametric methods by exploiting regularities of the learning problem (Györfi et al.,

2002). For example, Kveton and Theodorou (2013) showed that, when the dynamics of the underlying MDP are factored, incorporating this knowledge into KBRL leads to a dramatic reduction on the number of sample transitions needed for learning a good decision policy. Another type of regularity arises when the “intrinsic dimension” of the state space is much smaller than $d_{\mathcal{S}}$ —that is, when the data lies close to a low-dimensional manifold. Farahmand et al. (2007) have shown that simple k -nearest neighborhood regression is manifold-adaptive, meaning that it naturally exploits this type of regularity in the approximation problem. Since KBRL’s approximation scheme is very similar to k -nearest neighborhood regression, it is likely that this algorithm also has such a property.

Another way to avoid the curse of dimensionality is to resort to parametric methods. Parametric methods circumvent the exponential growth of the number of samples by restricting the space of functions spanned by the approximator—which corresponds to making assumptions about the target function (Györfi et al., 2002). Since in KBRL the structure of the approximator is defined by the sampled states \hat{s}_i^a , one has little flexibility in controlling the induced function space (see (25)). Based on the discussions in Sections 6.1.1 and 6.1.2, one can see that KBSF can be cast as a kernel-based approximation in which the structure of the approximator is defined by the representative states—that is, the representative states play the role of basis functions or “features” (this architecture is similar to that of radial basis function networks, see Chapter 7 of Györfi et al.’s book, 2002). It should be clear, then, that when the representative states are fixed KBSF can be seen as a parametric model—and thus as a potential tool for helping KBRL circumvent the curse of dimensionality.

Parametric regression methods have a serious drawback, though: if the target function cannot be well approximated by any function in the space induced by the parametric model, the approximation error will be large regardless of the data. A popular strategy to overcome this issue without incurring in unacceptable sample complexity is to resort to adaptive methods that use the available data to adjust the complexity of the approximator to the problem at hand (Farahmand and Szepesvári, 2011). In the case of KBRL, the complexity of the induced function space can be controlled via the kernel’s width τ (Györfi et al., 2002). KBSF provides an alternative way of controlling the complexity of the approximator through the parameter m (this is akin to limiting the depth of a regression tree, for example—see Barreto’s paper for a detailed discussion, 2014). If we think in terms of the classical bias-variance analysis of statistical estimators, intuitively we are decreasing bias and increasing variance as $m \rightarrow n$ —assuming that the representative states are systematically defined by a reasonable method; see Section 6.2.2 (Hastie et al., 2002). Note that, unlike τ , m has a clear effect on the method’s computational cost.

Therefore, KBSF provides both a way of determining the structure of the approximator and extra flexibility in controlling the complexity of the resulting model. Of course, this does not automatically translate into reduced sample complexity. Even if the structural assumptions induced by KBSF arise in problems of interest, we must develop methods to configure the parameters of the algorithm appropriately (see Section 6.2.2). New theoretical results regarding the approximation properties of KBSF would probably be needed as well, since the current results were derived under the interpretation of our algorithm as a computational device to accelerate KBRL (and thus use its solution as a reference point). These might be interesting directions for future investigations.

6.2 Practical issues

During the execution of our experiments we observed several interesting facts about KBSF which are not immediate from its conceptual definition. In this section we share some of the lessons learned with the reader. We start by discussing the impact of deviating from the theoretical assumptions over the performance of our algorithm. We then present general guidelines on how to configure KBSF to solve reinforcement learning problems.

6.2.1 KBSF’S APPLICABILITY

The theoretical guarantees regarding KBRL’s solution assume that the initial states s_i^a in the transitions $(s_i^a, r_i^a, \hat{s}_i^a)$ are uniformly sampled from \mathbb{S} (Ormoneit and Sen, 2002, see Assumption 3). This is somewhat restrictive because it precludes the collection of data through direct interaction with the environment. Ormoneit and Sen conjectured that sampling the states s_i^a from an uniform distribution is not strictly necessary, and indeed later Ormoneit and Glynn (2002) relaxed this assumption for the case in which KBRL is applied to an average-reward MDP. In this case, it is only required that the exploration policy used to collect data chooses all actions with positive probability. As described in Sections 4.2 and 5.2, in our computational experiments we collected data through an ϵ -greedy policy (in many cases with $\epsilon = 1$). The good performance of KBSF corroborates Ormoneit and Sen’s conjecture and suggests that Ormoneit and Glynn’s results can be generalized to the discounted reward case, but more theoretical analysis is needed.

Ormoneit and Sen (2002) also make some assumptions regarding the smoothness of the reward function and the transition kernel of the continuous MDP (Assumptions 1 and 2). Unfortunately, such assumptions are usually not verifiable in practice. Empirically, we observed that KBSF indeed performs better in problems with “smooth dynamics”—loosely speaking, problems in which a small perturbation in s_i^a results in a small perturbation in \hat{s}_i^a , such as the pole balancing task. In problems with “rougher” dynamics, like the epilepsy-suppression task, it is still possible to get good results with KBSF, but in this case it is necessary to use more representative states and narrower kernels (that is, smaller values for $\bar{\tau}$). As a result, in problems of this type KBSF is less effective in reducing KBRL’s computational cost.

6.2.2 KBSF’S CONFIGURATION

KBSF depends on two basic definitions: the set of representative states and the widths of the kernels. As discussed in Section 6.1, the former can be seen as the definition of the “structure” of the approximator, while the latter are parameters of the approximation. In what follows we discuss the impact of each of them on KBSF’s performance.

Definition of Representative States In order to define the representative states we must determine their number, m , and their “position” in the state space. Both theory and practice indicate that KBSF’s performance gets closer to KBRL’s as m increases. Thus, a “rule of thumb” to define the number of representative states is to simply set m to the largest value allowed by the available computational resources (but see Section 6.1.3 for an alternative view).

As for the position of the representative states, looking at expression (24) we see that ideally \bar{s}_j would be such that the rows of the matrices \mathbf{K}^a would form a convex hull containing the rows of the corresponding $\hat{\mathbf{P}}^a$. However, it is easy to see that when $m < n$ such a set of states may not exist. Besides, even when it does exist, finding this set is not a trivial problem.

Instead of insisting on finding representative states that allow for an exact representation of the matrices $\hat{\mathbf{P}}^a$, it sounds more realistic to content oneself with an approximate solution for this problem. Assuming the dynamics of the underlying MDP are reasonably smooth, one strategy is to guarantee that every sampled state \hat{s}_i^a is close to at least one representative state \bar{s}_j . Since covering the entire state space is generally impractical, we want \bar{s}_j to reflect the distribution of the data, that is, we want the representative states to be in the regions of the state space where the data lies. This is why in our experiments we clustered the states \hat{s}_i^a and used the clusters’s centers as our representative states (this method is also used in the configuration of radial basis function networks, as discussed by Györfi et al., 2002). Despite its simplicity, this strategy usually results in good performance, as shown in Sections 4.2 and 5.2.

In our experiments we clustered the data using the popular k -means algorithm, which minimizes the average square distance from a given point to the center of the corresponding cluster. However, Proposition 6 states that KBSF’s approximation error can be controlled by $\max_{a,i} \text{dist}(\hat{s}_i^a, w)$, the maximum distance from a sampled state \hat{s}_i^a to the w^{th} nearest representative state. It is reasonable to ask whether KBSF’s performance would improve if we used a clustering method that directly minimizes this quantity.

In the k -center clustering problem one seeks a set of k clusters that minimize the maximum distance from a point to the corresponding cluster center (Gonzales, 1985). Thus, in this case we would be minimizing $\max_{a,i} \text{dist}(\hat{s}_i^a, 1)$. Although finding an exact solution for this problem is NP hard, there exist fast algorithms that compute good approximations in time linear in the number of points (Feder and Greene, 1988). In order to provide some intuition on how k -center clustering compares to k -means, we implemented an algorithm by Gonzales (1985) which is extremely simple: at each iteration one selects a new cluster center by simply picking the point whose distance to the closest cluster center is maximal. Although simple, this method provides a solution for the k -center problem whose cost is within a factor of two from the optimal solution, which is the best one can hope for in polynomial time (Feder and Greene, 1988).

Figure 11 shows the performance of KBSF on the puddle-world task using different strategies to define the representative states. The experiment was carried out exactly as the one described in Section 4.2.1. In addition to the results provided by k -means and k -centers, we also show KBSF’s performance when using representative states sampled uniformly at random and evenly distributed over the state space. As expected, the worst results correspond to the case in which representative states are picked at random. On the opposite extreme, representative states evenly distributed over \mathbb{S} result in the best performance. Unfortunately, in more realistic scenarios it is impractical to cover the entire state space with representative states, and this is precisely why one may have to resort to strategies like k -means and k -centers. The results shown in Figure 11 suggest that k -centers clustering may indeed be a better choice than k -means, although the difference in the resulting performance is not very significant. One advantage of k -center clustering is that





Figure 11: Comparison of different strategies to define the representative states. Results on the puddle-world task averaged over 50 runs. The errors around the mean correspond to the 99% confidence intervals. See Figure 3 for details.

it is very simple to implement and is also very fast. Besides, there are algorithms available that compute an approximation on-line, which allows their use with *i*KBSF (Charikar et al., 1997; Beygelzimer et al., 2006).

Clustering methods are a natural choice for defining the representative states because they summarize the distribution of the data, but of course they are not the only way to solve the problem. Ideally, we want a set of representative states that summarize the *dynamics* of the underlying dynamical system, and this may not reflect the spatial distribution of the data, regardless of the strategy used to collect transitions. As discussed in Section 6.1, the problem of defining representative states is in some sense akin to the problem of defining features in supervised learning (Guyon and Elisseeff, 2003). Such a connection emphasizes the difficulty of the problem and suggests that the best solution may be domain-dependent. On the other hand, as with the definition of features, the definition of representative states can be seen as an opportunity to incorporate prior knowledge about the domain of interest into the approximation model. For example, if one knows that some regions of the state space are more important than others, this information can be used to allocate more representative states to those regions. Similar reasoning applies to tasks in which the level of accuracy required from the decision policy varies across the state space.

Definition of Kernels’s Widths Given a well-defined strategy to select representative states, the use of KBSF requires the definition of two parameters, τ and $\bar{\tau}$. The kernels’s widths τ and $\bar{\tau}$ may have a strong effect on KBSF’s performance. To illustrate this point, we show in Figure 12 the results of this algorithm on the puddle world task when τ and $\bar{\tau}$ are varied in the set $\{0.01, 0.1, 1\}$ (these were the results used to generate Figure 3).



(a) Performance of KBSF(8000, ·)

 (b) Performance of KBRL(8000).
 Errors around the mean correspond to 99% confidence intervals.

Figure 12: The impact of the kernels’s widths on the performance of KBSF and KBRL. Results on the puddle-world task averaged over 50 runs. See Figure 3 for details.

Of course, the best combination of values for τ and $\bar{\tau}$ depends on the specific problem at hand and on the particular choice of kernels. Here we give some general advice as to how to set these parameters, based on both theory in practice. Since τ is the same parameter used by KBRL, it should decrease with the number of sample transitions n at an “admissible rate” (see Ormonoit and Sen’s Lemma 2, 2002). Analogously, Proposition 6 suggests that $\bar{\tau}$ should get smaller as $m \rightarrow n$ (details in the proof of the proposition in Appendix A). Empirically, we found out that a simple strategy that usually facilitates the configuration of KBSF is to rescale the data so that all the variables have approximately the same magnitude—which corresponds to using a weighted norm in the computation of the kernels. Using this strategy we were able to obtain good results with KBSF on all problems by performing a coarse search in the space of parameters in which we only varied the order of magnitude of τ and $\bar{\tau}$ (see Table 1).

Alternatively, one can fix τ and $\bar{\tau}$ and define the neighborhood used to compute $k_\tau(\bar{s}_j, \cdot)$ and $\bar{k}_{\bar{\tau}}(\hat{s}_i^a, \cdot)$. As explained in Appendix B.2, in some of our experiments we only computed $k_\tau(\bar{s}_j, \cdot)$ for the μ closest sampled states s_i^a from \bar{s}_j , and only computed $\bar{k}_{\bar{\tau}}(\hat{s}_i^a, \cdot)$ for the $\bar{\mu}$ closest representative states from \hat{s}_i^a . When using this approach, a possible way of configuring KBSF is to set τ and $\bar{\tau}$ to sufficiently large values (so as to guarantee a minimum level of overlap between the kernels) and then adjust μ and $\bar{\mu}$. The advantage is that adjusting μ and $\bar{\mu}$ may be more intuitive than directly configuring τ and $\bar{\tau}$ (*cf.* Table 1).

7. Previous work

In our experiments we compared KBSF with KBRL, LSPI, fitted Q -iteration, and SARSA, both in terms of computational cost and in terms of the quality of the resulting decision policies. In this section we situate our algorithm in the broader context of approximate reinforcement learning. Approximation in reinforcement learning is an important topic that has generated a huge body of literature. For a broad overview of the subject, we refer the reader to the books by Sutton and Barto (1998), Bertsekas and Tsitsiklis (1996), and Szepesvári (2010). Here we will narrow our attention to *kernel-based* approximation techniques.

We start by noting that the label “kernel based” is used with two different meanings in the literature. On one side we have kernel smoothing techniques like KBRL and KBSF, which use local kernels essentially as a device to implement smooth instance-based approximation (Hastie et al., 2002). On the other side we have methods that use reproducing kernels to implicitly represent an inner product in a high-dimensional state space (Schölkopf and Smola, 2002). Although these two frameworks can give rise to approximators with similar structures, they rest on different theoretical foundations. Since reproducing-kernels methods are less directly related to KBSF, we will only describe them briefly. We will then discuss the kernel smoothing approaches in more detail.

The basic idea of reproducing-kernel methods is to apply the “kernel trick” in the context of reinforcement learning (Schölkopf and Smola, 2002). Roughly speaking, the approximation problem is rewritten in terms of inner products only, which are then replaced by a properly-defined kernel. This modification corresponds to mapping the problem to a high-dimensional feature space, resulting in more expressiveness of the function approximator. Perhaps the most natural way of applying the kernel trick in the context of reinforcement learning is to “kernelize” some formulation of the value-function approximation problem (Xu et al., 2005; Engel et al., 2005; Farahmand, 2011). Another alternative is to approximate the *dynamics* of an MDP using a kernel-based regression method (Rasmussen and Kuss, 2004; Taylor and Parr, 2009). Following a slightly different line of work, Bhat et al. (2012) propose to kernelize the linear programming formulation of dynamic programming. However, this method is not directly applicable to reinforcement learning, since it is based on the assumption that one has full knowledge of the MDP. A weaker assumption is to suppose that only the reward function is known and focus on the approximation of the transition function. This is the approach taken by Grunewalder et al. (2012), who propose to embed the conditional distributions defining the transitions of an MDP into a Hilbert space induced by a reproducing kernel.

We now turn our attention to kernel-smoothing techniques, which are more closely related to KBRL and KBSF. Kroemer and Peters (2011) propose to apply kernel density estimation to the problem of policy evaluation. They call their method *non-parametric dynamic programming* (NPDP). If we use KBRL to compute the value function of a fixed policy, we see many similarities with NPDP, but also some important differences. Like KBRL, NPDP is statistically consistent. Unlike KBRL, which assumes a finite action space A and directly approximates the conditional density functions $P^a(s'|s)$, NPDP assumes that A is continuous and models the joint density $P(s, a, s')$. Kroemer and Peters (2011) showed that the value function of NPDP has a Nadaraya-Watson kernel regression form. Not

surprisingly, this is also the form of KBRL’s solution if we fix the policy being evaluated (*cf.* equation (7)). In both cases, the coefficients of the kernel-based approximation are derived from the value function of the approximate MDP. The key difference is the way the transition matrices are computed in each algorithm. As shown in (4), the transition probabilities of KBRL’s model are given by the kernel values themselves. In contrast, the computation of each element of NDPD’s transition matrix requires an integration over the continuous state space \mathbb{S} . In practice, this is done by numerical integration techniques that may be very computationally demanding (see for example the experiments performed by Grunewalder et al., 2012).

We directly compared NDPD with KBRL because both algorithms build a model whose number of states is dictated by the number of sample transitions n , and neither method explicitly attempts to keep n small. Since in this case each application of the Bellman operator is $O(n^2)$, these methods are not suitable for problems in which a large number of transitions is required, nor are they applicable to on-line reinforcement learning.⁵ There are however kernel-smoothing methods that try to avoid this computational issue by either keeping n small or by executing a number of operations that grows only linearly with n . These algorithms are directly comparable with KBSF.

One of the first attempts to adapt KBRL to the on-line scenario was that of Jong and Stone (2006). Instead of collecting a batch of sample transitions before the learning process starts, the authors propose to grow such a set incrementally, based on an exploration policy derived from KBRL’s current model. To avoid running a dynamic-programming algorithm to completion in between two transitions, which may not be computationally feasible, Jong and Stone (2006) resort to Moore and Atkeson’s (1993) “prioritized sweeping” method to propagate the changes in the value function every time the model is modified. The idea of exploiting the interpretation of KBRL as the derivation of a finite MDP in order to use tabular exploration methods is insightful. However, it is not clear whether smart exploration is sufficient to overcome the computational difficulties arising from the fact that the size of the underlying model is inexorably linked to the number of sample transitions. For example, even using sparse kernels in their experiments, Jong and Stone (2006) had to fix an upper limit for the size of KBRL’s model. In this case, once the number of sample transitions has reached the upper limit, all subsequent data must be ignored.

Following the same line of work, Jong and Stone (2009) later proposed to guide KBRL’s exploration of the state space using Brafman and Tennenholtz’s (2003) R-MAX algorithm. In this new paper the authors address the issue with KBRL’s scalability more aggressively. First, they show how to combine their approach with Dietterich’s (2000) MAX-Q algorithm, allowing the decomposition of KBRL’s MDP into a hierarchy of simpler models. While this can potentially reduce the computational burden of finding a policy, such a strategy transfer to the user the responsibility of identifying a useful decomposition of the task. A more practical approach is to combine KBRL with some stable form of value-function approximation. For that, Jong and Stone (2009) suggest the use of Gordon’s (1995) averagers. As shown in Appendix A.2, this setting corresponds to a particular case of KBSF in which representative states are selected among the set of sampled states \hat{s}_i^a . It should be noted that, even when using temporal abstraction and function approximation, Jong and Stone’s (2009) approach

5. We note that, incidentally, all the reproducing-kernel methods discussed in this section also have a computational complexity super-linear in n .

requires recomputing KBRL’s transition probabilities at each new sample, which can be infeasible in reasonably large problems.

Kveton and Theodorou (2012) propose a more practical algorithm to reduce KBRL’s computational cost. Their method closely resembles the batch version of KBSF. As with our algorithm, Kveton and Theodorou’s method defines a set of representative states \bar{s}_i that give rise to a reduced MDP. The main difference in the construction of the models is that, instead of computing a similarity measure between each sampled state \hat{s}_i^a and all representative states \bar{s}_j , their algorithm associates each \hat{s}_i^a with a single \bar{s}_j —which comes down to computing a hard aggregation of the state space \hat{S} . Such an aggregation corresponds to having a matrix \mathbf{D} with a single nonzero element per row. In fact, it is possible to rewrite Kveton and Theodorou’s (2012) algorithm using KBSF’s formalism. In this case, the elements of $\dot{\mathbf{D}}^a$ and $\dot{\mathbf{K}}^a$ would be defined as:

$$\dot{k}_{ij}^a = \kappa_{\bar{\tau}}^a(\bar{s}_i, rs(s_j^a, 1)), \quad \text{and} \quad \dot{d}_{ij}^a = \bar{\kappa}_0(rs(\hat{s}_i^a, 1), \bar{s}_j) \tag{28}$$

where $\bar{\kappa}_0$ is the normalized kernel induced by an infinitely “narrow kernel $\bar{k}_0(s, s')$ whose value is greater than zero if and only if $s = s'$ (recall from Section 4.1 that $rs(s, 1)$ gives the closest representative state from s). It is easy to see that we can make matrix \mathbf{D} computed by KBSF as close as desired to a hard aggregation by setting $\bar{\tau}$ to a sufficiently small value (see Lemma 5). More practically, we can simply plug (28) in place of (10) in Algorithm 1 to exactly recover Kveton and Theodorou’s method. Note though that, by replacing $\kappa_{\bar{\tau}}^a(\bar{s}_i, s_j^a)$ with $\kappa_{\bar{\tau}}^a(\bar{s}_i, rs(s_j^a, 1))$ in the computation of $\dot{\mathbf{K}}^a$, we would be in some sense deviating from KBRL’s framework. To see why this is so, observe that if the representative states \bar{s}_i are sampled from the set of states \hat{s}_i^a , the rows of matrix \mathbf{K}^a computed by KBSF would coincide with a subset of the rows of the corresponding KBRL’s matrix $\hat{\mathbf{P}}^a$ (cf. (23)). However, this property is lost if one uses (28) instead of (10).⁶

8. Conclusion

This paper presented KBSF, a reinforcement learning algorithm that results from the application of the stochastic-factorization trick to KBRL. KBSF summarizes the information contained in KBRL’s MDP in a model of fixed size. By doing so, our algorithm decouples the structure of the model from its configuration. This makes it possible to build an approximation which accounts for both the difficulty of the problem *and* the computational resources available.

One of the main strengths of KBSF is its simplicity. As shown in the paper, its uncomplicated mechanics can be unfolded into two update rules that allow for a fully incremental version of the algorithm. This makes the amount of memory used by KBSF independent of the number of sample transitions. Therefore, with a few lines of code one has a reinforcement-learning algorithm that can be applied to large-scale problems, in both off-line and on-line regimes.

KBSF is also a sound method from a theoretical point of view. As discussed, the distance between the value function computed by this algorithm and the one computed by KBRL

6. Note that this observation only means that KBSF is closer to KBRL in this strict sense; in particular, it does not imply that Kveton and Theodorou’s algorithm is not a principled method, nor does it mean that it will perform worse than KBSF.

is bounded by two factors: the quality and the level of stochasticity of the underlying stochastic factorization. We showed that both factors can be made arbitrarily small, which implies that, in theory, we can make KBSF’s solution as close to KBRL’s solution as desired.

Theoretical guarantees do not always translate into good performance in practice, though, either because they are built upon unrealistic assumptions or because they do not account for procedural difficulties that arise in practical situations. To ensure that this is not the case with our algorithm, we presented an extensive empirical study in which KBSF was successfully applied to different problems, some of them quite challenging. We also presented general guidelines on how to configure KBSF to solve a reinforcement learning problem.

For all the reasons listed above, we believe that KBSF has the potential of becoming a valuable resource for the solution of reinforcement learning problems. This is not to say that the subject has been exhausted. There are several possibilities for future research, some of which we now briefly discuss.

From an algorithmic point of view, perhaps the most pressing demand is for more principled methods to define the representative states. Incidentally, this also opens up the possibility of an automated procedure to set the kernel’s widths $\bar{\tau}$ based solely on data. Taking the idea one step further, one can think of having one distinct $\bar{\tau}_i$ associated with each kernel $\bar{\kappa}_{\bar{\tau}}(\cdot, \bar{s}_i)$ (which would make the similarity between KBSF’s approximation and radial basis function networks even stronger). Another important advance would be to endow *i*KBSF with more elaborate exploration strategies, maybe following the line of research initiated by Jong and Stone (2006, 2009).

From a theoretical perspective, it may be desirable to “detach” KBSF from KBRL and analyze the former on its own. In this paper we emphasized the view of our algorithm as a tool to apply KBRL in practice. This view is clearly reflected in our theoretical results, which show how KBSF’s solution deviates from KBRL’s and how we can control such a deviation. But KBRL is itself an approximation, so perhaps it makes sense to analyze KBSF independently from its precursor. This possibility is particularly appealing when we look at KBSF as an intermediate approach between fully nonparametric and parametric models, as discussed in Section 6.1.3.

Looking at KBSF against a broader context, a subject that deserves further investigation is the possibility of building an approximation based on multiple models. Model averaging is not inherently linked to KBSF, and in principle it can be used with virtually any reinforcement learning algorithm. However, KBSF’s low computational cost makes it particularly amenable to this technique. Since our algorithm is significantly faster than any method whose complexity per iteration is a function of the number of sample transitions, we can afford to compute several approximations and still have a solution in comparable time (see Section 4.2.3 and Barreto’s paper, 2014). Understanding how we can randomize the construction of the individual models and to what extent this can improve the quality of the resulting decision policy is a matter of interest.

We conclude by noting that KBSF represents one particular way in which the stochastic-factorization trick can be exploited in the context of reinforcement learning. In principle, any algorithm that builds a model based on sample transitions can resort to the same trick to leverage the use of the data. The basic idea remains the same: instead of estimating the transition probabilities between every pair of states, one focuses on a small set of representative states whose values are propagated throughout the state space based on some

notion of similarity. We believe that this general framework can potentially be materialized into a multitude of useful reinforcement learning algorithms.

Acknowledgments

The authors would like to thank Amir-massoud Farahmand and Yuri Grinberg for interesting discussions about approximation in reinforcement learning and related subjects. We also thank Branislav Kveton for insightful comments regarding the theory underlying kernel-based reinforcement learning, which were essential for the derivation of our Proposition 8. We are grateful to Keith Bush for making the epilepsy simulator available, and to Alicia Bendz and Ryan Primeau for helping in some of the computational experiments. Finally, we thank the anonymous reviewers for their valid comments and helpful suggestions to improve the paper. Funding for this research was provided by the NSERC Discovery Grant program.

Appendix A. Theoretical Results

A.1 Proofs

Proposition 2

Proof Let $\check{M} \equiv (S, A, \check{\mathbf{P}}^a, \check{\mathbf{r}}^a, \gamma)$, with $\check{\mathbf{P}}^a = \mathbf{D}\mathbf{K}^a$ and $\check{\mathbf{r}}^a = \mathbf{D}\bar{\mathbf{r}}^a$. From the triangle inequality, we know that

$$\|\mathbf{v}^* - \Gamma\mathbf{D}\bar{\mathbf{Q}}^*\|_\infty \leq \|\mathbf{v}^* - \check{\mathbf{v}}^*\|_\infty + \|\check{\mathbf{v}}^* - \Gamma\mathbf{D}\bar{\mathbf{Q}}^*\|_\infty, \quad (29)$$

where $\check{\mathbf{v}}^*$ is the optimal value function of \check{M} . Our strategy will be to bound $\|\mathbf{v}^* - \check{\mathbf{v}}^*\|_\infty$ and $\|\check{\mathbf{v}}^* - \Gamma\mathbf{D}\bar{\mathbf{Q}}^*\|_\infty$. In order to find an upper bound for $\|\mathbf{v}^* - \check{\mathbf{v}}^*\|_\infty$, we apply Whitt's (1978) Theorem 3.1 and Corollary (b) of his Theorem 6.1, with all mappings between M and \check{M} taken to be identities, to obtain

$$\|\mathbf{v}^* - \check{\mathbf{v}}^*\|_\infty \leq \frac{1}{1-\gamma} \left(\max_a \|\mathbf{r}^a - \mathbf{D}\bar{\mathbf{r}}^a\|_\infty + \frac{\gamma\bar{R}_{\text{dif}}}{2(1-\gamma)} \max_a \|\mathbf{P}^a - \mathbf{D}\mathbf{K}^a\|_\infty \right), \quad (30)$$

where we used the fact that $\max_{a,i} \check{r}_i^a - \min_{a,i} \check{r}_i^a \leq \bar{R}_{\text{dif}}$. It remains to bound $\|\check{\mathbf{v}}^* - \Gamma\mathbf{D}\bar{\mathbf{Q}}^*\|_\infty$. Since $\check{\mathbf{r}}^a = \mathbf{D}\bar{\mathbf{r}}^a$ and $\mathbf{D}\bar{\mathbf{P}}^a = \mathbf{D}\mathbf{K}^a\mathbf{D} = \check{\mathbf{P}}^a\mathbf{D}$ for all $a \in A$, the stochastic matrix \mathbf{D} satisfies Sorg and Singh's (2009) definition of a *soft homomorphism* between \check{M} and \bar{M} (see equations (25)–(28) in their paper). Applying Theorem 1 by the same authors, we know that

$$\|\Gamma(\check{\mathbf{Q}}^* - \mathbf{D}\bar{\mathbf{Q}}^*)\|_\infty \leq \frac{1}{1-\gamma} \sup_{i,t} (1 - \max_j d_{ij}) \bar{\delta}_i^{(t)}, \quad (31)$$

where $\bar{\delta}_i^{(t)} = \max_{j:d_{ij}>0,k} \bar{q}_{jk}^{(t)} - \min_{j:d_{ij}>0,k} \bar{q}_{jk}^{(t)}$ and $\bar{q}_{jk}^{(t)}$ are elements of $\bar{\mathbf{Q}}^{(t)}$, the optimal t -step action-value function of \bar{M} . Since $\|\Gamma\check{\mathbf{Q}}^* - \Gamma\mathbf{D}\bar{\mathbf{Q}}^*\|_\infty \leq \|\Gamma(\check{\mathbf{Q}}^* - \mathbf{D}\bar{\mathbf{Q}}^*)\|_\infty$ and, for all

$t > 0$, $\bar{\delta}_i^{(t)} \leq (1 - \gamma)^{-1}(\max_{a,k} \bar{r}_k^a - \min_{a,k} \bar{r}_k^a)$, we can write

$$\|\tilde{\mathbf{v}}^* - \Gamma \mathbf{D} \bar{\mathbf{Q}}^*\|_\infty \leq \frac{\bar{R}_{\text{dif}}}{(1 - \gamma)^2} \max_i (1 - \max_j d_{ij}) = \frac{\bar{R}_{\text{dif}}}{(1 - \gamma)^2} \sigma(\mathbf{D}). \quad (32)$$

Substituting (30) and (32) back into (29), we obtain (8). \blacksquare

Lemma 4

Proof Define the function

$$\psi_{\tau,s}^{a,i}(s') = \left| \frac{k_\tau(s, s_i^a)}{\sum_{j=1}^{n_a} k_\tau(s, s_j^a)} - \frac{k_\tau(s', s_i^a)}{\sum_{j=1}^{n_a} k_\tau(s', s_j^a)} \right| = \left| \frac{\phi(\|s - s_i^a\|/\tau)}{\sum_{j=1}^{n_a} \phi(\|s - s_j^a\|/\tau)} - \frac{\phi(\|s' - s_i^a\|/\tau)}{\sum_{j=1}^{n_a} \phi(\|s' - s_j^a\|/\tau)} \right|.$$

Since ϕ is continuous, it is obvious that $\psi_{\tau,s}^{a,i}(s')$ is also continuous in s' . The property follows from the fact that $\lim_{s' \rightarrow s} \psi_{\tau,s}^{a,i}(s') = 0$. \blacksquare

Lemma 5[†] *Let $s \in \mathbb{S}$, let $m > 1$, and assume there is a $w \in \{1, 2, \dots, m - 1\}$ such that $\text{dist}(s, w) < \text{dist}(s, w + 1)$. Define $W^w(s) \equiv \{k \mid \|s - \bar{s}_k\| \leq \text{dist}(s, w)\}$ and $\bar{W}^w(s) \equiv \{1, 2, \dots, m\} - W^w(s)$. Then, for any $\alpha > 0$, we can guarantee that*

$$\sum_{k \in \bar{W}^w(s)} \bar{k}_{\bar{\tau}}(s, \bar{s}_k) < \alpha \sum_{k \in W^w(s)} \bar{k}_{\bar{\tau}}(s, \bar{s}_k) \quad (33)$$

by making $\bar{\tau} < \varphi(s, w, m, \alpha)$, where

$$\varphi(s, w, m, \alpha) = \min(\varphi_1(s, w), \varphi_2(s, w, m, \alpha)) \quad (34)$$

and

$$\varphi_1(s, w) = \begin{cases} \frac{\text{dist}(s, w)}{B_{\bar{\phi}}}, & \text{if } B_{\bar{\phi}} > 0, \\ \infty, & \text{otherwise,} \end{cases} \quad \varphi_2(s, w, m, \alpha) = \begin{cases} \frac{\text{dist}(s, w) - \text{dist}(s, w + 1)}{\ln(\alpha w / (m - w) \lambda_{\bar{\phi}})}, & \text{if } \frac{\alpha w}{(m - w) \lambda_{\bar{\phi}}} < 1, \\ \infty, & \text{otherwise.} \end{cases}$$

Proof Expression (33) can be rewritten as

$$\frac{\sum_{k \in \bar{W}^w(s)} \bar{k}_{\bar{\tau}}(s, \bar{s}_k)}{\sum_{i=1}^m \bar{k}_{\bar{\tau}}(s, \bar{s}_i)} < \alpha \frac{\sum_{k \in W^w(s)} \bar{k}_{\bar{\tau}}(s, \bar{s}_k)}{\sum_{i=1}^m \bar{k}_{\bar{\tau}}(s, \bar{s}_i)} \iff \sum_{k \in \bar{W}^w(s)} \bar{k}_{\bar{\tau}}(s, \bar{s}_k) < \alpha \sum_{k \in W^w(s)} \bar{k}_{\bar{\tau}}(s, \bar{s}_k), \quad (35)$$

which is equivalent to

$$\sum_{k \in \bar{W}^w(s)} \bar{\phi} \left(\frac{\|s - \bar{s}_k\|}{\bar{\tau}} \right) < \alpha \sum_{k \in W^w(s)} \bar{\phi} \left(\frac{\|s - \bar{s}_k\|}{\bar{\tau}} \right). \quad (36)$$

[†]. We restate the lemma here showing explicitly how to define $\bar{\tau}$. This detail was omitted in the main body of the text to improve clarity.

Based on Assumption (i), we know that a sufficient condition for (36) to hold is

$$\bar{\phi}\left(\frac{\text{dist}(s, w+1)}{\bar{\tau}}\right) < \frac{\alpha w}{m-w} \bar{\phi}\left(\frac{\text{dist}(s, w)}{\bar{\tau}}\right). \quad (37)$$

Let $\beta = \alpha w/(m-w)$. If $\beta > 1$, then (37) is always true, regardless of the value of $\bar{\tau}$. We now show that, when $\beta \leq 1$, it is always possible to set $\bar{\tau}$ in order to guarantee that (37) holds. Let $z = \text{dist}(s, w)$ and let $\delta = \text{dist}(s, w+1) - z$. From Assumption (ii), we know that, if $B_{\bar{\phi}} = 0$ or $\bar{\tau} < z/B_{\bar{\phi}}$,

$$\frac{\bar{\phi}((z+\delta)/\bar{\tau})}{\bar{\phi}(z/\bar{\tau})} \leq \frac{\lambda_{\bar{\phi}} A_{\bar{\phi}} \exp(-(z+\delta)/\bar{\tau})}{A_{\bar{\phi}} \exp(-z/\bar{\tau})} = \frac{\lambda_{\bar{\phi}} \exp(-(z+\delta)/\bar{\tau})}{\exp(-z/\bar{\tau})}.$$

Thus, in order for the result to follow, it suffices to show that

$$\frac{\exp(-(z+\delta)/\bar{\tau})}{\exp(-z/\bar{\tau})} < \frac{\beta}{\lambda_{\bar{\phi}}}. \quad (38)$$

We know that, since $\delta > 0$, if $\beta/\lambda_{\bar{\phi}} = 1$ inequality (38) is true. Otherwise,

$$\begin{aligned} \frac{\exp(-(z+\delta)/\bar{\tau})}{\exp(-z/\bar{\tau})} < \frac{\beta}{\lambda_{\bar{\phi}}} &\iff \ln\left(\frac{\exp(-(z+\delta)/\bar{\tau})}{\exp(-z/\bar{\tau})}\right) < \ln\left(\frac{\beta}{\lambda_{\bar{\phi}}}\right) \\ &\iff -\frac{\delta}{\bar{\tau}} < \ln\left(\frac{\beta}{\lambda_{\bar{\phi}}}\right) \iff \bar{\tau} < -\frac{\delta}{\ln(\beta/\lambda_{\bar{\phi}})}. \end{aligned}$$

Thus, by taking $\bar{\tau} < -\delta/\ln(\beta/\lambda_{\bar{\phi}})$ if $B_{\bar{\phi}} > 0$, or $\bar{\tau} < \min(-\delta/\ln(\beta/\lambda_{\bar{\phi}}), z/B_{\bar{\phi}})$ otherwise, the result follows. \blacksquare

Proposition 6

Proof From (6) and (11), we know that

$$\|\hat{\mathbf{r}}^a - \mathbf{D}\bar{\mathbf{r}}^a\|_{\infty} = \|\hat{\mathbf{P}}^a \mathbf{r} - \mathbf{D}\mathbf{K}^a \mathbf{r}\|_{\infty} = \|(\hat{\mathbf{P}}^a - \mathbf{D}\mathbf{K}^a) \mathbf{r}\|_{\infty} \leq \|\hat{\mathbf{P}}^a - \mathbf{D}\mathbf{K}^a\|_{\infty} \|\mathbf{r}\|_{\infty}. \quad (39)$$

Thus, plugging (39) back into (8), it is clear that there is a $\nu > 0$ such that $\xi_v < \epsilon$ if

$$\max_a \|\hat{\mathbf{P}}^a - \mathbf{D}\mathbf{K}^a\|_{\infty} < \nu \quad (40)$$

and

$$\max_i (1 - \max_j d_{ij}) < \nu. \quad (41)$$

We start by showing that there is a $\delta > 0$ and a $\theta > 0$ such that expression (40) is true if $\max_{a,i} \text{dist}(\hat{s}_i^a, w) < \delta$ and $\bar{\tau} < \theta$. Let $\check{\mathbf{P}}^a = \mathbf{D}\mathbf{K}^a$ and let $\hat{\mathbf{p}}_i^a \in \mathbb{R}^{1 \times n}$ and $\check{\mathbf{p}}_i^a \in \mathbb{R}^{1 \times n}$ be the i^{th} rows of $\hat{\mathbf{P}}^a$ and $\check{\mathbf{P}}^a$, respectively. Then,

$$\begin{aligned} \|\hat{\mathbf{p}}_i^a - \check{\mathbf{p}}_i^a\|_{\infty} &= \sum_{j=1}^{n_a} |\hat{p}_{ij}^a - \sum_{k=1}^m d_{ik}^a \check{p}_{kj}^a| \\ &= \sum_{j=1}^{n_a} |\kappa_{\bar{\tau}}^a(\hat{s}_i^a, s_j^a) - \sum_{k=1}^m \bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, \bar{s}_k) \kappa_{\bar{\tau}}^a(\bar{s}_k, s_j^a)| \\ &= \sum_{j=1}^{n_a} \left| \sum_{k=1}^m \bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, \bar{s}_k) \kappa_{\bar{\tau}}^a(\hat{s}_i^a, s_j^a) - \sum_{k=1}^m \bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, \bar{s}_k) \kappa_{\bar{\tau}}^a(\bar{s}_k, s_j^a) \right| \\ &= \sum_{j=1}^{n_a} \left| \sum_{k=1}^m \bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, \bar{s}_k) [\kappa_{\bar{\tau}}^a(\hat{s}_i^a, s_j^a) - \kappa_{\bar{\tau}}^a(\bar{s}_k, s_j^a)] \right| \\ &\leq \sum_{j=1}^{n_a} \sum_{k=1}^m \bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, \bar{s}_k) \left| \kappa_{\bar{\tau}}^a(\hat{s}_i^a, s_j^a) - \kappa_{\bar{\tau}}^a(\bar{s}_k, s_j^a) \right|. \end{aligned} \quad (42)$$

Our strategy will be to show that, for any a, i , and j , there is a $\delta^{a,i,j} > 0$ and a $\theta^{a,i,j} > 0$ such that

$$\sum_{k=1}^m \bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, \bar{s}_k) |\kappa_{\bar{\tau}}^a(\hat{s}_i^a, s_j^a) - \kappa_{\bar{\tau}}^a(\bar{s}_k, s_j^a)| = \sum_{k=1}^m \bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, \bar{s}_k) \zeta_k^{a,i,j} < \frac{\nu}{n_a} \quad (43)$$

if $\text{dist}(\hat{s}_i^a, w) < \delta^{a,i,j}$ and $\bar{\tau} < \theta^{a,i,j}$ (recall that $\zeta_k^{a,i,j}$ had already been defined in (16)). To simplify the notation, we will use the superscript ‘ z ’ meaning ‘ a, i, j ’. From Lemma 4 we know that there is a $\delta^z > 0$ such that $\zeta_k^z < \nu/n_a$ if $\|\hat{s}_i^a - \bar{s}_k\| < \delta^z$. Let $W^z \equiv \{k \mid \|\hat{s}_i^a - \bar{s}_k\| < \delta^z\}$ and $\bar{W}^z \equiv \{1, 2, \dots, m\} - W^z$. Since we are assuming that $\text{dist}(\hat{s}_i^a, w) < \delta^z$, we know that $W^z \neq \emptyset$. In this case, we can write:

$$\sum_{k=1}^m \bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, \bar{s}_k) \zeta_k^z = \sum_{k \in W^z} \bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, \bar{s}_k) \zeta_k^z + \sum_{k \in \bar{W}^z} \bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, \bar{s}_k) \zeta_k^z.$$

Let

$$\zeta_{\min}^z = \begin{cases} \min_{k \in W^z} \{\zeta_k^z \mid \zeta_k^z > 0\} & \text{if } \max_{k \in W^z} \zeta_k^z > 0, \\ 0 & \text{otherwise} \end{cases} \quad \text{and} \quad \zeta_{\max}^z = \begin{cases} \max_{k \in \bar{W}^z} \zeta_k^z & \text{if } |W^z| < m, \\ 0 & \text{otherwise.} \end{cases}$$

If $\zeta_{\max}^z = 0$, inequality (43) is necessarily true, since $\sum_{k \in W^z} \bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, \bar{s}_k) \zeta_k^z \leq \max_{k \in W^z} \zeta_k^z < \nu/n_a$.

We now turn to the case in which $\zeta_{\max}^z > 0$. Suppose first that $\zeta_{\min}^z = 0$. In this case, we have to show that there is a $\bar{\tau}$ that yields

$$\sum_{k \in \bar{W}^z} \bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, \bar{s}_k) \zeta_k^z < \frac{\nu}{n_a}. \quad (44)$$

A sufficient condition for (44) to be true is

$$\sum_{k \in \bar{W}^z} \bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, \bar{s}_k) < \frac{\nu}{n_a \zeta_{\max}^z} \iff \frac{1}{\sum_{j=1}^m \bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, \bar{s}_j)} \sum_{k \in \bar{W}^z} \bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, \bar{s}_k) < \frac{\nu}{n_a \zeta_{\max}^z}. \quad (45)$$

Obviously, if $\zeta_{\max}^z \leq \nu/n_a$ inequality (45) is always true, regardless of the value of $\bar{\tau}$. Otherwise, we can rewrite (45) as

$$\sum_{k \in \bar{W}^z} \bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, \bar{s}_k) < \frac{\nu}{n_a \zeta_{\max}^z} \left(\sum_{j \in W^z} \bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, \bar{s}_j) + \sum_{k \in \bar{W}^z} \bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, \bar{s}_k) \right),$$

and, after a few algebraic manipulations, we obtain

$$\sum_{k \in \bar{W}^z} \bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, \bar{s}_k) < \frac{\nu}{n_a \zeta_{\max}^z - \nu} \sum_{k \in W^z} \bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, \bar{s}_k) \iff \sum_{k \in \bar{W}^z} \bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, \bar{s}_k) < \frac{\nu}{n_a \zeta_{\max}^z - \nu} \sum_{k \in W^z} \bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, \bar{s}_k). \quad (46)$$

We can guarantee that (46) is true by applying Lemma 5. Before doing so, though, let's analyze the case in which $\zeta_{\min}^z > 0$. Define

$$\beta^z = \frac{\nu}{n_a \sum_{k \in W^z} \bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, \bar{s}_k) \zeta_k^z} - 1 \quad (47)$$

(note that $\beta^z > 0$ because $\sum_{k \in W^z} \bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, \bar{s}_k) \zeta_k^z < \nu/n_a$). In order for (43) to hold, we must show that there is a $\bar{\tau}$ that guarantees that

$$\sum_{k \in \bar{W}^z} \bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, \bar{s}_k) \zeta_k^z - \beta^z \sum_{k \in W^z} \bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, \bar{s}_k) \zeta_k^z < 0. \quad (48)$$

A sufficient condition for (48) to hold is

$$\sum_{k \in \bar{W}^z} \bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, \bar{s}_k) < \frac{\beta^z \zeta_{\min}^z}{\bar{\zeta}_{\max}^z} \sum_{k \in W^z} \bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, \bar{s}_k). \quad (49)$$

Observe that expressions (46) and (49) only differ in the coefficient multiplying the right-hand side of the inequalities. Let

$$\alpha_1^z = \begin{cases} \nu/(\bar{\zeta}_{\max}^z n_a - \nu), & \text{if } \bar{\zeta}_{\max}^z > \nu/n_a \\ \infty, & \text{otherwise,} \end{cases}$$

and

$$\alpha_2^z = \begin{cases} \beta^z \zeta_{\min}^z / \bar{\zeta}_{\max}^z, & \text{if } \bar{\zeta}_{\max}^z > 0 \text{ and } \zeta_{\min}^z > 0, \\ \infty, & \text{otherwise.} \end{cases}$$

Let $\alpha^z < \min(\alpha_1^z, \alpha_2^z)$. Then, if we make $\theta^z = \varphi(\hat{s}_i^a, |W|, m, \alpha^z)$, with φ defined in (34), we can apply Lemma 5 to guarantee that (43) holds. Finally, if we let $\delta = \min_z \delta^z = \min_{a,i,j} \delta^{a,i,j}$ and $\theta = \min_z \theta^z = \min_{a,i,j} \theta^{a,i,j}$, we can guarantee that (43) is true for all a , i , and j , which implies that (40) is also true (see (42)).

It remains to show that there is a $\omega > 0$ such that (41) is true if $\bar{\tau} < \omega$. Recalling that, for any i and any a ,

$$\max_j d_{ij}^a = \bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, rs(\hat{s}_i^a, 1)),$$

we want to show that

$$\bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, rs(\hat{s}_i^a, 1)) > (1 - \nu) \left[\bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, rs(\hat{s}_i^a, 1)) + \sum_{k=2}^m \bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, rs(\hat{s}_i^a, k)) \right],$$

which is equivalent to

$$(1 - \nu) \sum_{k=2}^m \bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, rs(\hat{s}_i^a, k)) < \nu \bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, rs(\hat{s}_i^a, 1)). \quad (50)$$

If $\nu \geq 1$, inequality (50) is true regardless of the particular choice of $\bar{\tau}$. Otherwise, we can rewrite (50) as

$$\sum_{k=2}^m \bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, rs(\hat{s}_i^a, k)) < \frac{\nu}{1 - \nu} \bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, rs(\hat{s}_i^a, 1)) \iff \sum_{k=2}^m \bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, rs(\hat{s}_i^a, k)) < \frac{\nu}{1 - \nu} \bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, rs(\hat{s}_i^a, 1)). \quad (51)$$

Let

$$\alpha = \begin{cases} \nu/(1 - \nu), & \text{if } \nu < 1, \\ \infty, & \text{otherwise.} \end{cases}$$

Then, if we make $\omega^{a,i} = \varphi(\hat{s}_i^a, 1, m, \alpha)$, with φ defined in (34), we can resort to Lemma 5 to

guarantee that (51) holds. As before, if we let $\omega = \min_{a,i} \omega^{a,i}$, we can guarantee that (41) is true. Finally, by making $\bar{\tau} = \min(\theta, \omega)$, the result follows. \blacksquare

Proposition 8

Proof

We start by plugging (39) into the definition of ξ_v , given in (8), to get:

$$\begin{aligned}
 \xi_v &\leq \frac{1}{1-\gamma} \max_a \|\hat{\mathbf{P}}^a - \mathbf{DK}^a\|_\infty \|\mathbf{r}\|_\infty + \frac{\bar{R}_{\text{dif}}}{(1-\gamma)^2} \left(\frac{\gamma}{2} \max_a \|\mathbf{P}^a - \mathbf{DK}^a\|_\infty + \sigma(\mathbf{D}) \right) \\
 &= \max_a \|\mathbf{P}^a - \mathbf{DK}^a\|_\infty \left(\frac{\|\mathbf{r}\|_\infty}{1-\gamma} + \frac{\gamma \bar{R}_{\text{dif}}}{2(1-\gamma)^2} \right) + \sigma(\mathbf{D}) \frac{\bar{R}_{\text{dif}}}{(1-\gamma)^2}. \\
 &\leq \max_a \|\mathbf{P}^a - \mathbf{DK}^a\|_\infty \left(\frac{R_{\text{max}}}{1-\gamma} + \frac{\gamma R_{\text{max}}}{(1-\gamma)^2} \right) + \sigma(\mathbf{D}) \frac{R_{\text{max}}}{2(1-\gamma)^2} \\
 &= \max_a \|\mathbf{P}^a - \mathbf{DK}^a\|_\infty \frac{R_{\text{max}}}{(1-\gamma)^2} + \sigma(\mathbf{D}) \frac{R_{\text{max}}}{2(1-\gamma)^2}, \tag{52}
 \end{aligned}$$

where we used the fact that $\bar{R}_{\text{dif}} \leq R_{\text{max}}/2$, which follows trivially from the observation that $\min_{a,i} \bar{r}_i^a \geq \min_{a,i} r_i^a$ and $\max_{a,i} \bar{r}_i^a \leq \max_{a,i} r_i^a$. To simplify the notation, let $\mathcal{R} = R_{\text{max}}/(1-\gamma)^2$. Then, from (42), the definition of $\zeta_k^{a,i,j}$ in (16), and the definition of $\sigma(\mathbf{D})$ in (9), we can rewrite (52) as

$$\begin{aligned}
 \xi_v &\leq \max_{a,i} \left(\sum_{j=1}^{n_a} \sum_{k=1}^m \bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, \bar{s}_k) \zeta_k^{a,i,j} \right) \mathcal{R} + \max_{a,i} [1 - \bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, rs(\hat{s}_i^a, 1))] \frac{\mathcal{R}}{2} \\
 &= \max_{a,i} \left(\sum_{k=1}^m \bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, \bar{s}_k) \sum_{j=1}^{n_a} \zeta_k^{a,i,j} \right) \mathcal{R} + \max_{a,i} [1 - \bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, rs(\hat{s}_i^a, 1))] \frac{\mathcal{R}}{2}
 \end{aligned}$$

Let $W^w(\hat{s}_i^a)$ and $\bar{W}^w(\hat{s}_i^a)$ be defined as in (15). Then,

$$\begin{aligned}
 \xi_v &\leq \max_{a,i} \left(\sum_{k \in W^w(\hat{s}_i^a)} \bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, \bar{s}_k) \sum_{j=1}^{n_a} \zeta_k^{a,i,j} + \sum_{k \in \bar{W}^w(\hat{s}_i^a)} \bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, \bar{s}_k) \sum_{j=1}^{n_a} \zeta_k^{a,i,j} \right) \mathcal{R} + \max_{a,i} (1 - \bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, rs(\hat{s}_i^a, 1))) \frac{\mathcal{R}}{2} \\
 &\leq \max_{a,i} \left(\sum_{k \in W^w(\hat{s}_i^a)} \bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, \bar{s}_k) \sum_{j=1}^{n_a} \zeta_k^{a,i,j} \right) \mathcal{R} \\
 &\quad + \underbrace{\max_{a,i} \left(\sum_{k \in \bar{W}^w(\hat{s}_i^a)} \bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, \bar{s}_k) \sum_{j=1}^{n_a} \zeta_k^{a,i,j} \right) \mathcal{R} + \max_{a,i} (1 - \bar{\kappa}_{\bar{\tau}}(\hat{s}_i^a, rs(\hat{s}_i^a, 1))) \frac{\mathcal{R}}{2}}_{\mathcal{F}(\bar{\tau}, w)}. \tag{53}
 \end{aligned}$$

We see that the two last terms of (53) coincide with the definition of $\mathcal{F}(\bar{\tau}, w)$, given in (17). Thus, if we make sure that $\bar{\kappa}_{\bar{\tau}}$ is an admissible kernel $\bar{\kappa}_{\bar{\tau}}^{\epsilon, w}$ (see Definition 7), we can

rewrite (53) as

$$\begin{aligned}
 \xi_v &\leq \max_{a,i} \left(\sum_{k \in W^w(\hat{s}_i^a)} \bar{\kappa}_{\bar{\tau}}^{\epsilon,w}(\hat{s}_i^a, \bar{s}_k) \sum_{j=1}^{n_a} \varsigma_k^{a,i,j} \right) \mathcal{R} + \epsilon \\
 &\leq \max_{a,i} \left(w \max_{k \in W^w(\hat{s}_i^a)} \sum_{j=1}^{n_a} \varsigma_k^{a,i,j} \right) \mathcal{R} + \epsilon \\
 &= w \max_{a,i} \left(\max_{k \in W^w(\hat{s}_i^a)} \sum_{j=1}^{n_a} |\kappa_{\tau}^a(\hat{s}_i^a, s_j^a) - \kappa_{\tau}^a(\bar{s}_k, s_j^a)| \right) \mathcal{R} + \epsilon. \tag{54}
 \end{aligned}$$

Let $z_i^a = \sum_l \kappa_{\tau}(\hat{s}_i^a, s_l^a)$ and $\bar{z}_k^a = \sum_l \kappa_{\tau}(\bar{s}_k, s_l^a)$. Then, for fixed a, i, j , and k :

$$\begin{aligned}
 |\kappa_{\tau}^a(\hat{s}_i^a, s_j^a) - \kappa_{\tau}^a(\bar{s}_k, s_j^a)| &= \left| \frac{\kappa_{\tau}(\hat{s}_i^a, s_j^a)}{\sum_l \kappa_{\tau}(\hat{s}_i^a, s_l^a)} - \frac{\kappa_{\tau}(\bar{s}_k, s_j^a)}{\sum_l \kappa_{\tau}(\bar{s}_k, s_l^a)} \right| \\
 &= \left| \frac{\kappa_{\tau}(\hat{s}_i^a, s_j^a)}{z_i^a} - \frac{\kappa_{\tau}(\bar{s}_k, s_j^a)}{\bar{z}_k^a} \right| \\
 &= \left| \frac{\bar{z}_k^a \kappa_{\tau}(\hat{s}_i^a, s_j^a) - z_i^a \kappa_{\tau}(\bar{s}_k, s_j^a)}{z_i^a \bar{z}_k^a} \right| \\
 &= \left| \frac{\bar{z}_k^a \kappa_{\tau}(\hat{s}_i^a, s_j^a) - z_i^a \kappa_{\tau}(\bar{s}_k, s_j^a) + z_i^a \kappa_{\tau}(\hat{s}_i^a, s_j^a) - z_i^a \kappa_{\tau}(\hat{s}_i^a, s_j^a)}{z_i^a \bar{z}_k^a} \right| \\
 &= \left| \frac{\kappa_{\tau}(\hat{s}_i^a, s_j^a) - \kappa_{\tau}(\bar{s}_k, s_j^a)}{\bar{z}_k^a} + \frac{\kappa_{\tau}(\hat{s}_i^a, s_j^a)(\bar{z}_k^a - z_i^a)}{z_i^a \bar{z}_k^a} \right| \\
 &\leq \frac{|\kappa_{\tau}(\hat{s}_i^a, s_j^a) - \kappa_{\tau}(\bar{s}_k, s_j^a)|}{\bar{z}_k^a} + \frac{\kappa_{\tau}(\hat{s}_i^a, s_j^a) |\bar{z}_k^a - z_i^a|}{z_i^a \bar{z}_k^a}.
 \end{aligned}$$

Thus,

$$\begin{aligned}
 \sum_{j=1}^{n_a} |\kappa_{\tau}^a(\hat{s}_i^a, s_j^a) - \kappa_{\tau}^a(\bar{s}_k, s_j^a)| &\leq \sum_{j=1}^{n_a} \frac{|\kappa_{\tau}(\hat{s}_i^a, s_j^a) - \kappa_{\tau}(\bar{s}_k, s_j^a)|}{\bar{z}_k^a} + \sum_{j=1}^{n_a} \frac{\kappa_{\tau}(\hat{s}_i^a, s_j^a) |\bar{z}_k^a - z_i^a|}{z_i^a \bar{z}_k^a} \\
 &= \sum_{j=1}^{n_a} \frac{|\kappa_{\tau}(\hat{s}_i^a, s_j^a) - \kappa_{\tau}(\bar{s}_k, s_j^a)|}{\bar{z}_k^a} + \sum_{j=1}^{n_a} \kappa_{\tau}^a(\hat{s}_i^a, s_j^a) \frac{|\bar{z}_k^a - z_i^a|}{\bar{z}_k^a} \\
 &= \sum_{j=1}^{n_a} \frac{|\kappa_{\tau}(\hat{s}_i^a, s_j^a) - \kappa_{\tau}(\bar{s}_k, s_j^a)|}{\bar{z}_k^a} + \frac{|\bar{z}_k^a - z_i^a|}{\bar{z}_k^a} \\
 &= \frac{1}{\sum_{l=1}^{n_a} \kappa_{\tau}(\bar{s}_k, s_l^a)} \left(\sum_{j=1}^{n_a} |\kappa_{\tau}(\hat{s}_i^a, s_j^a) - \kappa_{\tau}(\bar{s}_k, s_j^a)| + \left| \sum_{j=1}^{n_a} \kappa_{\tau}(\bar{s}_k, s_j^a) - \sum_{j=1}^{n_a} \kappa_{\tau}(\hat{s}_i^a, s_j^a) \right| \right) \\
 &\leq \frac{1}{\sum_{l=1}^{n_a} \kappa_{\tau}(\bar{s}_k, s_l^a)} \left(\sum_{j=1}^{n_a} |\kappa_{\tau}(\hat{s}_i^a, s_j^a) - \kappa_{\tau}(\bar{s}_k, s_j^a)| + \sum_{j=1}^{n_a} |\kappa_{\tau}(\bar{s}_k, s_j^a) - \kappa_{\tau}(\hat{s}_i^a, s_j^a)| \right) \\
 &= \frac{2}{\sum_{l=1}^{n_a} \kappa_{\tau}(\bar{s}_k, s_l^a)} \sum_{j=1}^{n_a} |\kappa_{\tau}(\hat{s}_i^a, s_j^a) - \kappa_{\tau}(\bar{s}_k, s_j^a)|. \tag{55}
 \end{aligned}$$

In order to derive our result, it suffices to replace (55) in (54). Before doing so, though, note that Assumption (iii) implies that

$$\begin{aligned}
 |k_\tau(s, s') - k_\tau(s, s'')| &= \left| \phi\left(\frac{\|s - s'\|}{\tau}\right) - \phi\left(\frac{\|s - s''\|}{\tau}\right) \right| \\
 &\leq \frac{C_\phi}{\tau} \left| \|s - s'\| - \|s - s''\| \right| \\
 &\leq \frac{C_\phi}{\tau} \left| \|s - s''\| + \|s'' - s'\| - \|s - s''\| \right| \\
 &= \frac{C_\phi}{\tau} \|s'' - s'\|.
 \end{aligned} \tag{56}$$

Proceeding with the substitution mentioned above, we have

$$\begin{aligned}
 \xi_v &\leq 2w \max_{a,i} \left(\max_{k \in W^w(\hat{s}_i^a)} \frac{1}{\sum_{l=1}^{n_a} k_\tau(\bar{s}_k, s_l^a)} \sum_{j=1}^{n_a} |k_\tau(\hat{s}_i^a, s_j^a) - k_\tau(\bar{s}_k, s_j^a)| \right) \mathcal{R} + \epsilon \\
 &\leq 2w \max_{a,i} \left(\max_{k \in W^w(\hat{s}_i^a)} \frac{1}{\sum_{l=1}^{n_a} k_\tau(\bar{s}_k, s_l^a)} \frac{C_\phi}{\tau} n_a \|\hat{s}_i^a - \bar{s}_k\| \right) \mathcal{R} + \epsilon \\
 &= \frac{2wC_\phi}{\tau} \max_{a,i} \left(\max_{k \in W^w(\hat{s}_i^a)} \frac{n_a \|\hat{s}_i^a - \bar{s}_k\|}{\sum_{l=1}^{n_a} k_\tau(\bar{s}_k, s_l^a)} \right) \mathcal{R} + \epsilon \\
 &\leq \frac{2wC_\phi}{\tau} \max_{a,i} \left(\max_{k \in W^w(\hat{s}_i^a)} \frac{n_a \|\hat{s}_i^a - \bar{s}_k\|}{n_a k_{\min}} \right) \mathcal{R} + \epsilon \\
 &= \frac{2wC_\phi}{\tau k_{\min}} \max_{a,i} \max_{k \in W^w(\hat{s}_i^a)} \|\hat{s}_i^a - \bar{s}_k\| \mathcal{R} + \epsilon.
 \end{aligned}$$

■

Lemma 9

Proof Let $\mathbf{q}_*^a, \tilde{\mathbf{q}}_*^a \in \mathbb{R}^{|S|}$ be the a^{th} columns of \mathbf{Q}^* and $\tilde{\mathbf{Q}}^*$, respectively. Then,

$$\begin{aligned}
 \|\mathbf{q}_*^a - \tilde{\mathbf{q}}_*^a\|_\infty &= \left\| \mathbf{r}^a + \gamma \mathbf{P}^a \mathbf{v}^* - \tilde{\mathbf{r}}^a - \gamma \tilde{\mathbf{P}}^a \tilde{\mathbf{v}}^* \right\|_\infty \\
 &\leq \|\mathbf{r}^a - \tilde{\mathbf{r}}^a\|_\infty + \gamma \left\| \mathbf{P}^a \mathbf{v}^* - \tilde{\mathbf{P}}^a \tilde{\mathbf{v}}^* \right\|_\infty \\
 &= \|\mathbf{r}^a - \tilde{\mathbf{r}}^a\|_\infty + \gamma \left\| \mathbf{P}^a \mathbf{v}^* - \tilde{\mathbf{P}}^a \mathbf{v}^* + \tilde{\mathbf{P}}^a \mathbf{v}^* - \tilde{\mathbf{P}}^a \tilde{\mathbf{v}}^* \right\|_\infty \\
 &\leq \|\mathbf{r}^a - \tilde{\mathbf{r}}^a\|_\infty + \gamma \left\| (\mathbf{P}^a - \tilde{\mathbf{P}}^a) \mathbf{v}^* \right\|_\infty + \gamma \left\| \tilde{\mathbf{P}}^a (\mathbf{v}^* - \tilde{\mathbf{v}}^*) \right\|_\infty \\
 &\leq \|\mathbf{r}^a - \tilde{\mathbf{r}}^a\|_\infty + \gamma \left\| (\mathbf{P}^a - \tilde{\mathbf{P}}^a) \mathbf{v}^* \right\|_\infty + \gamma \|\mathbf{v}^* - \tilde{\mathbf{v}}^*\|_\infty,
 \end{aligned} \tag{57}$$

where in the last step we used the fact that $\tilde{\mathbf{P}}^a$ is stochastic, and thus $\|\tilde{\mathbf{P}}^a \mathbf{v}\|_\infty \leq \|\mathbf{v}\|_\infty$ for any \mathbf{v} . We now provide a bound for $\|(\mathbf{P}^a - \tilde{\mathbf{P}}^a) \mathbf{v}^*\|_\infty$. Let $\mathbf{A} = \mathbf{P}^a - \tilde{\mathbf{P}}^a$. Then, for any i , $\sum_j a_{ij} = \sum_j (p_{ij}^a - \tilde{p}_{ij}^a) = \sum_j p_{ij}^a - \sum_j \tilde{p}_{ij}^a = 0$, that is, the elements in each row of \mathbf{A} sum to zero. Let a_i^+ be the sum of positive elements in the i^{th} row of \mathbf{A} and let $a_{\max}^+ = \max_i a_i^+$.

It should be clear that $\|\mathbf{A}\|_\infty = 2a_{\max}^+$. Then, for any i ,

$$\begin{aligned} \left| \sum_j a_{ij} v_j^* \right| &\leq \sum_{(j:a_{ij}>0)} a_{ij} v_{\max}^* + \sum_{(j:a_{ij}<0)} a_{ij} v_{\min}^* = a_i^+ v_{\max}^* - a_i^+ v_{\min}^* \leq a_{\max}^+ (v_{\max}^* - v_{\min}^*) \\ &\leq \frac{a_{\max}^+}{1-\gamma} (r_{\max}^a - r_{\min}^a) \leq \frac{a_{\max}^+ R_{\text{dif}}}{1-\gamma} = \frac{R_{\text{dif}}}{2(1-\gamma)} \|\mathbf{P}^a - \tilde{\mathbf{P}}^a\|_\infty, \end{aligned} \quad (58)$$

where we used the convention $v_{\max} = \max_i v_i$ (analogously for v_{\min}). As done in (30), we can resort to Whitt's (1978) Theorem 3.1 and Corollary (b) of his Theorem 6.1 to obtain a bound for $\|\mathbf{v}^* - \tilde{\mathbf{v}}^*\|_\infty$. Substituting such a bound and expression (58) in (57), we obtain

$$\begin{aligned} \|\mathbf{q}_*^a - \tilde{\mathbf{q}}_*^a\|_\infty &\leq \|\mathbf{r}^a - \tilde{\mathbf{r}}^a\|_\infty + \frac{\gamma R_{\text{dif}}}{2(1-\gamma)} \|\mathbf{P}^a - \tilde{\mathbf{P}}^a\|_\infty + \frac{\gamma}{1-\gamma} \left(\max_a \|\mathbf{r}^a - \tilde{\mathbf{r}}^a\|_\infty + \frac{\gamma R_{\text{dif}}}{2(1-\gamma)} \max_a \|\mathbf{P}^a - \tilde{\mathbf{P}}^a\|_\infty \right) \\ &\leq \max_a \|\mathbf{r}^a - \tilde{\mathbf{r}}^a\|_\infty + \frac{\gamma R_{\text{dif}}}{2(1-\gamma)} \max_a \|\mathbf{P}^a - \tilde{\mathbf{P}}^a\|_\infty + \frac{\gamma}{1-\gamma} \left(\max_a \|\mathbf{r}^a - \tilde{\mathbf{r}}^a\|_\infty + \frac{\gamma R_{\text{dif}}}{2(1-\gamma)} \max_a \|\mathbf{P}^a - \tilde{\mathbf{P}}^a\|_\infty \right). \end{aligned}$$

■

Proposition 10

Proof Let $\tilde{M}_t \equiv (\hat{S}_t, A, \tilde{\mathbf{P}}_t^a, \tilde{\mathbf{r}}_t^a, \gamma)$, with $\tilde{\mathbf{P}}_t^a = \mathbf{D}_t \mathbf{K}_t^a$ and $\tilde{\mathbf{r}}_t^a = \mathbf{D}_t \tilde{\mathbf{r}}_t^a$. From the triangle inequality, we know that

$$|\hat{Q}_t(s, a) - \tilde{Q}_t(s, a)| \leq |\hat{Q}_t(s, a) - \check{Q}_t^*(s, a)| + |\check{Q}_t^*(s, a) - \tilde{Q}_t^*(s, a)| + |\tilde{Q}_t^*(s, a) - \tilde{Q}_t(s, a)|, \quad (59)$$

where \hat{Q}_t and \tilde{Q}_t are defined in the proposition's statement, \check{Q}_t^* is the optimal action-value function of \tilde{M}_t , and $\tilde{Q}_t^*(s, a) = \sum_{i=1}^m \bar{\kappa}_{\bar{\tau}}(s, \bar{s}_i) \bar{Q}_t^*(\bar{s}_i, a)$ (the reader will forgive a slight abuse of notation here, since in general \check{Q}_t^* is not the optimal value function of any MDP). Our strategy will be to bound each term on the right-hand side of (59). Since \hat{M}_t is the model constructed by KBRL using all the data seen by i KBSF up to time step t , state s will correspond to one of the states \hat{s}_i^b in this MDP. Thus, from (7), we see that $\hat{Q}_t(s, a) = \hat{Q}_t^*(\hat{s}_i^b, a)$ for some i and some b . Therefore, applying Lemma 9 to \hat{M}_t and \tilde{M}_t , we can write

$$|\hat{Q}_t(s, a) - \check{Q}_t^*(s, a)| \leq \frac{1}{1-\gamma} \max_a \|\tilde{\mathbf{r}}_t^a - \mathbf{D}_t \tilde{\mathbf{r}}_t^a\|_\infty + \frac{\gamma}{2(1-\gamma)^2} \bar{R}_{\text{dif},t} \max_a \|\tilde{\mathbf{P}}_t^a - \mathbf{D}_t \mathbf{K}_t^a\|_\infty. \quad (60)$$

In order to bound $|\check{Q}_t^*(s, a) - \tilde{Q}_t^*(s, a)|$, we note that, since the information contained in the transition to state s has been incorporated to i KBSF's model \bar{M} at time t , $\tilde{Q}_t^*(s, a) = \sum_{i=1}^m d_{ti,t} \bar{Q}_t^*(\bar{s}_i, a)$, for any $a \in A$, where $d_{ti,t}$ is the element in the t^{th} row and i^{th} column of \mathbf{D}_t (see Figure 2b). In matrix form, we have $\tilde{\mathbf{Q}}_t^* = \mathbf{D}_t \bar{\mathbf{Q}}_t^*$. As \mathbf{D}_t is a soft homomorphism between \tilde{M}_t and \bar{M}_t , we can resort to Sorg and Singh's (2009) Theorem 1, as done in Proposition 2, to write:

$$|\check{Q}_t^*(s, a) - \tilde{Q}_t^*(s, a)| \leq \frac{\bar{R}_{\text{dif},t}}{(1-\gamma)^2} \sigma(\mathbf{D}_t) \quad (61)$$

(see (31) and (32)). Finally,

$$\begin{aligned} |\tilde{Q}_t^*(s, a) - \tilde{Q}_t(s, a)| &= \left| \sum_{i=1}^m \bar{\kappa}_{\bar{\tau}}(s, \bar{s}_i) \bar{Q}_t^*(\bar{s}_i, a) - \sum_{i=1}^m \bar{\kappa}_{\bar{\tau}}(s, \bar{s}_i) \bar{Q}_t(\bar{s}_i, a) \right| \\ &\leq \sum_{i=1}^m \bar{\kappa}_{\bar{\tau}}(s, \bar{s}_i) |\bar{Q}_t^*(\bar{s}_i, a) - \bar{Q}_t(\bar{s}_i, a)| \leq \epsilon_{\bar{Q}_t}, \end{aligned} \quad (62)$$

where the last step follows from the fact that $\sum_{i=1}^m \bar{\kappa}_{\bar{\tau}}(s, \bar{s}_i)$ is a convex combination. Substituting (60), (61), and (62) in (59), we obtain the desired bound. \blacksquare

A.2 Alternative error bound

In Section 3 we derived an upper bound for the approximation error introduced by the application of the stochastic-factorization trick. In this section we introduce another bound that has different properties. First, the bound is less applicable, because it depends on quantities that are usually unavailable in a practical situation (the fixed points of two contraction mappings). On the bright side, unlike the bound presented in Proposition 2, the new bound is valid for any norm. Also, it draws an interesting connection with an important class of approximators known as *averagers* (Gordon, 1995).

We start by deriving a theoretical result that only applies to stochastic factorizations of order n . We then generalize this result to the case in which the factorizations are of order $m < n$.

Lemma 11 *Let $M \equiv (S, A, \mathbf{P}^a, \mathbf{r}^a, \gamma)$ be a finite MDP with $|S| = n$ and $0 \leq \gamma < 1$. Let $\mathbf{E}\mathbf{L}^a = \mathbf{P}^a$ be $|A|$ stochastic factorizations of order n and let $\bar{\mathbf{r}}^a$ be vectors in \mathbb{R}^n such that $\mathbf{E}\bar{\mathbf{r}}^a = \mathbf{r}^a$ for all $a \in A$. Define the MDPs $\check{M} \equiv (S, A, \mathbf{L}^a, \bar{\mathbf{r}}^a, \gamma)$ and $\bar{\bar{M}} \equiv (S, A, \bar{\bar{\mathbf{P}}}^a, \bar{\mathbf{r}}^a, \gamma)$, with $\bar{\bar{\mathbf{P}}}^a = \mathbf{L}^a \mathbf{E}$. Then,*

$$\|\mathbf{v}^* - T\mathbf{E}\bar{\bar{\mathbf{v}}}^*\| \leq \xi'_v \equiv \frac{2\gamma}{1-\gamma} \|\mathbf{v}^* - \mathbf{u}\| + \frac{\gamma(1+\gamma)}{1-\gamma} \|\mathbf{v}^* - \check{\mathbf{v}}^*\|, \quad (63)$$

where $\|\cdot\|$ is a norm in \mathbb{R}^n and \mathbf{u} is a vector in \mathbb{R}^n such that $\mathbf{E}\mathbf{u} = \mathbf{u}$.

Proof The Bellman operators of M , \check{M} , and $\bar{\bar{M}}$ are given by $T = \Gamma\Delta$, $\check{T} = \Gamma\check{\Delta}$, and $\bar{\bar{T}} = \Gamma\bar{\bar{\Delta}}$. Note that $\mathbf{q}^a = \mathbf{r}^a + \gamma\mathbf{P}^a\mathbf{v} = \mathbf{E}\bar{\mathbf{r}}^a + \gamma\mathbf{E}\mathbf{L}^a\mathbf{v} = \mathbf{E}(\bar{\mathbf{r}}^a + \gamma\mathbf{L}^a\mathbf{v})$, where \mathbf{q}^a is the a^{th} column of \mathbf{Q} . Thus, $\Delta = \mathbf{E}\check{\Delta}$. Since \mathbf{E} is stochastic, we can think of it as one of Gordon's (1995) averagers given by $A(\mathbf{v}) = \mathbf{E}\mathbf{v}$, and then resort to Theorem 4.1 by the same author to conclude that $\bar{\bar{T}} = \mathbf{E}\check{T}$. Therefore,⁷

$$T\mathbf{v} = \Gamma\mathbf{E}\check{\Delta}\mathbf{v} \quad \text{and} \quad \bar{\bar{T}}\mathbf{v} = \mathbf{E}\Gamma\check{\Delta}\mathbf{v}. \quad (64)$$

7. Interestingly, the effect of swapping matrices \mathbf{E} and \mathbf{L}^a is to also swap the operators Γ and \mathbf{E} .

Using (64), it is easy to obtain the desired upper bound by resorting to the triangle inequality, the definition of a contraction map, and Denardo's (1967) Theorem 1:

$$\begin{aligned}
 \|\mathbf{v}^* - T\mathbf{E}\bar{\mathbf{v}}^*\| &\leq \gamma\|\mathbf{v}^* - \mathbf{E}\bar{\mathbf{v}}^*\| \leq \gamma(\|\mathbf{v}^* - \mathbf{u}\| + \|\mathbf{u} - \mathbf{E}\bar{\mathbf{v}}^*\|) \leq \gamma(\|\mathbf{v}^* - \mathbf{u}\| + \|\mathbf{u} - \bar{\mathbf{v}}^*\|) \\
 &\leq \gamma\left(\|\mathbf{v}^* - \mathbf{u}\| + \frac{1}{1-\gamma}\|\mathbf{u} - \mathbf{E}\Gamma\check{\Delta}\mathbf{u}\|\right) \leq \gamma\left(\|\mathbf{v}^* - \mathbf{u}\| + \frac{1}{1-\gamma}\|\mathbf{u} - \Gamma\check{\Delta}\mathbf{u}\|\right) \\
 &\leq \gamma\left[\|\mathbf{v}^* - \mathbf{u}\| + \frac{1}{1-\gamma}(\|\mathbf{u} - \check{\mathbf{v}}^*\| + \|\check{\mathbf{v}}^* - \Gamma\check{\Delta}\mathbf{u}\|)\right] \\
 &\leq \gamma\left[\|\mathbf{v}^* - \mathbf{u}\| + \frac{1}{1-\gamma}(\|\mathbf{u} - \check{\mathbf{v}}^*\| + \gamma\|\check{\mathbf{v}}^* - \mathbf{u}\|)\right] = \gamma\left[\|\mathbf{v}^* - \mathbf{u}\| + \frac{1+\gamma}{1-\gamma}\|\mathbf{u} - \check{\mathbf{v}}^*\|\right] \\
 &\leq \gamma\left[\|\mathbf{v}^* - \mathbf{u}\| + \frac{1+\gamma}{1-\gamma}(\|\mathbf{u} - \mathbf{v}^*\| + \|\mathbf{v}^* - \check{\mathbf{v}}^*\|)\right] \\
 &= \gamma\|\mathbf{v}^* - \mathbf{u}\| + \frac{\gamma(1+\gamma)}{1-\gamma}\|\mathbf{v}^* - \mathbf{u}\| + \frac{\gamma(1+\gamma)}{1-\gamma}\|\mathbf{v}^* - \check{\mathbf{v}}^*\| \\
 &= \frac{\gamma - \gamma^2 + \gamma + \gamma^2}{1-\gamma}\|\mathbf{v}^* - \mathbf{u}\| + \frac{\gamma(1+\gamma)}{1-\gamma}\|\mathbf{v}^* - \check{\mathbf{v}}^*\|.
 \end{aligned}$$

■

The derived upper bound depends on two fixed points: \mathbf{u} , a fixed point of \mathbf{E} , and $\check{\mathbf{v}}^*$, the unique fixed point of $\check{T} = \Gamma\check{\Delta}$. Since the latter is defined by $\bar{\mathbf{r}}^a$ and \mathbf{L}^a , the bound is essentially a function of the factorization terms, as expected. Notice that the bound is valid for any norm and any fixed point of \mathbf{E} (we may think of \mathbf{u} as the closest vector to \mathbf{v}^* in \mathbb{R}^n which satisfies this property). Notice also that the first term on the right-hand side of (63) is exactly the error bound derived in Gordon's (1995) Theorem 6.2. When $\mathbf{L}^a = \mathbf{P}^a$ and $\mathbf{r}^a = \bar{\mathbf{r}}^a$ for all $a \in A$, the operators T and \check{T} coincide, and hence the second term of (63) vanishes. This makes sense, since in this case $\bar{\mathbf{T}} = \mathbf{E}T$, that is, the stochastic-factorization trick reduces to the averager $A(\mathbf{v}) = \mathbf{E}\mathbf{v}$.

As mentioned above, one of the assumptions of Lemma 11 is that the factorizations $\mathbf{E}\mathbf{L}^a = \mathbf{P}^a$ are of order n . This is unfortunate, since the whole motivation behind the stochastic-factorization trick is to create an MDP with $m < n$ states. One way to obtain such a reduction is to suppose that matrix \mathbf{E} has $n - m$ columns with zeros only. Define $\mathcal{E} \subset \{1, 2, \dots, n\}$ as the set of columns of \mathbf{E} with at least one nonzero element and let \mathbf{H} be a matrix in $\mathbb{R}^{m \times n}$ such that $h_{ij} = 1$ if j is the i^{th} smallest element in \mathcal{E} and $h_{ij} = 0$ otherwise. The following proposition generalizes the previous result to a stochastic factorization of order m :

Proposition 12 *Suppose the assumptions of Lemma 11 hold. Let $\mathbf{D} = \mathbf{E}\mathbf{H}^\top$, $\mathbf{K}^a = \mathbf{H}\mathbf{L}^a$, and $\bar{\mathbf{r}}^a = \mathbf{H}\bar{\mathbf{r}}^a$, with \mathbf{H} defined as described above. Define the MDP $\bar{M} \equiv (\bar{S}, A, \bar{\mathbf{P}}^a, \bar{\mathbf{r}}^a, \gamma)$, with $|\bar{S}| = m$ and $\bar{\mathbf{P}}^a = \mathbf{K}^a\mathbf{D}$. Then, $\|\mathbf{v}^* - \Gamma\mathbf{D}\bar{\mathbf{Q}}^*\| \leq \xi'_v$, with ξ'_v defined in (63).*

Proof Let $\bar{\mathbf{q}}_*^a \in \mathbb{R}^m$ be the a^{th} column of $\bar{\mathbf{Q}}^*$. Then,

$$\begin{aligned}
 \mathbf{D}\bar{\mathbf{q}}_*^a &= \mathbf{D}(\bar{\mathbf{r}}^a + \gamma\bar{\mathbf{P}}^a\bar{\mathbf{v}}^*) = \mathbf{D}\bar{\mathbf{r}}^a + \gamma\mathbf{D}\mathbf{K}^a\mathbf{D}\bar{\mathbf{v}}^* = \mathbf{E}\mathbf{H}^\top\mathbf{H}\bar{\mathbf{r}}^a + \gamma\mathbf{E}\mathbf{H}^\top\mathbf{H}\mathbf{L}^a\mathbf{E}\mathbf{H}^\top\bar{\mathbf{v}}^* \\
 &= \mathbf{E}\bar{\mathbf{r}}^a + \gamma\mathbf{E}\bar{\mathbf{P}}^a\mathbf{H}^\top\bar{\mathbf{v}}^* = \mathbf{E}\bar{\mathbf{r}}^a + \gamma\mathbf{E}\bar{\mathbf{P}}^a\bar{\mathbf{v}}^* = \mathbf{E}(\bar{\mathbf{r}}^a + \gamma\bar{\mathbf{P}}^a\bar{\mathbf{v}}^*) = \mathbf{E}\bar{\mathbf{q}}_*^a,
 \end{aligned} \tag{65}$$

where the equality $\mathbf{E}\mathbf{H}^T\mathbf{H} = \mathbf{E}$ follows from the definition of \mathbf{H} . The identity $\bar{\mathbf{P}}^a\mathbf{H}^T\bar{\mathbf{v}}^* = \bar{\mathbf{P}}^a\bar{\mathbf{v}}^*$ is a consequence of the fact that the i^{th} column of $\bar{\mathbf{P}}^a$ only contains zeros if $i \notin \mathcal{E}$. Since the i^{th} state of \bar{M} is transient, the values of the recurrent states of \bar{M} , which effectively define the multiplication $\bar{\mathbf{P}}^a\bar{\mathbf{v}}^*$, will coincide with the values of the states of \bar{M} . Expression (65) then leads to $\mathbf{D}\bar{\mathbf{Q}}^* = \mathbf{E}\bar{\mathbf{Q}}^*$. Also, since $\mathbf{E}\bar{\mathbf{q}}_*^a = \mathbf{E}\bar{\mathbf{r}}^a + \gamma\mathbf{E}\mathbf{L}^a\mathbf{E}\bar{\mathbf{v}}^* = \mathbf{r}^a + \gamma\mathbf{P}^a\mathbf{E}\bar{\mathbf{v}}^*$, we know that $\mathbf{E}\bar{\mathbf{Q}}^* = \Delta\mathbf{E}\bar{\mathbf{v}}^*$. Putting these results together, we obtain $\|\mathbf{v}^* - \Gamma\mathbf{D}\bar{\mathbf{Q}}^*\| = \|\mathbf{v}^* - \Gamma\Delta\mathbf{E}\bar{\mathbf{v}}^*\| = \|\mathbf{v}^* - \mathbf{T}\mathbf{E}\bar{\mathbf{v}}^*\|$, and Lemma 11 applies. \blacksquare

The derived bound can be generalized to the case of approximate stochastic factorizations through the triangle inequality, as done in (29). However, if one resorts to Whitt’s (1978) results to bound the distance between \mathbf{v}^* and $\check{\mathbf{v}}^*$ —where $\check{\mathbf{v}}^*$ is the optimal value function of $\check{M} \equiv (S, A, \mathbf{D}\mathbf{K}^a, \mathbf{D}\bar{\mathbf{r}}^a, \gamma)$ —the compounded bound will no longer be valid for all norms, since (30) only holds for the infinity norm.

Appendix B. Details of the experiments

This appendix describes the details of the experiments omitted in the paper.

B.1 Tasks

Puddle World: The puddle-world task was implemented as described by Sutton (1996), but here the task was modeled as a discounted problem with $\gamma = 0.99$. All the transitions were associated with a zero reward, except those leading to the goal, which resulted in a reward of +5, and those ending inside one of the puddles, which lead to a penalty of -10 times the distance to the puddle’s nearest edge. If the agent did not reach the goal after 300 steps the episode was interrupted and considered as a failure. The algorithms were evaluated on two sets of states distributed over disjoint regions of the state space surrounding the puddles. The first set was a 3×3 grid defined over $[0.1, 0.3] \times [0.3, 0.5]$ and the second one was composed of four states: $\{0.1, 0.3\} \times \{0.9, 1.0\}$.

Pole Balancing: We implemented the simulator of the three versions of the pole-balancing task using the equations of motion and parameters given in the appendix of Gomez’s (2003) PhD thesis. For the integration we used the 4th order Runge-Kutta method with a time step of 0.01 seconds and actions chosen every 2 time steps. The problem was modeled as a discounted task with $\gamma = 0.99$. We considered the version of the task in which the angle between the pole and the vertical plane must be kept within $[-36^\circ, 36^\circ]$. In this formulation, an episode is interrupted and the agent gets a reward of -1 if the pole falls past a 36-degree angle or the cart reaches the boundaries of the track, located at 2.4m from its center. At all other steps the agent receives a reward of 0. In all versions of the problem an episode was considered a success if the pole(s) could be balanced for 3000 steps (one minute of simulated time). The test set was comprised of 81 states equally spaced in the region defined by $\pm[1.2\text{m}, 1.2/5\text{m}, 18^\circ, 75^\circ/\text{s}]$, for the single pole case, and by $\pm[1.2\text{m}, 1.2/5\text{m}, 18^\circ, 75^\circ/\text{s}, 18^\circ, 150^\circ/\text{s}]$ for the double-pole version of the problem. These values correspond to a hypercube centered at the origin and covering 50% of the state-space axes in each dimension (since the velocity of the cart and the angular velocity of the poles are theoretically not bounded, we defined the limits of these variables based on samples gen-

erated in simple preliminary experiments). For the triple pole-balancing task we performed our simulations using the parameters usually adopted with the two pole version of the problem, but we added a third pole with the same length and mass as the longer pole (Gomez, 2003). In this case the decision policies were evaluated on a test set containing 256 states equally distributed in the region $\pm[1.2\text{m}, 1.2/5\text{m}, 18^\circ, 75^\circ/s, 18^\circ, 150^\circ/s, 18^\circ, 75^\circ/s]$.

HIV drug schedule: The HIV drug schedule task was implemented using the system of ordinary differential equations (ODEs) given by Adams et al. (2004). Integration was carried out by the Euler method using a step size of 0.001 with actions selected at each 5000 steps (corresponding to 5 days of simulated time). As suggested by Ernst et al. (2006), the problem was modeled as a discounted task with $\gamma = 0.98$. All other parameters of the task, as well as the protocol used for the numerical simulations, also followed the suggestions of the same authors. In particular, we assumed the existence of 30 patients who were monitored for 1000 days. During the monitoring period, the content of the drug cocktail administered to each patient could be changed at fixed intervals of 5 days. Thus, in a sample transition $(s_i^a, r_i^a, \hat{s}_i^a)$: s_i^a is the initial patient condition, a is one of the four types of cocktails to be administered for the next 5 days, \hat{s}_i^a is the patient condition 5 days later, and r_i^a is a reward computed based on the amount of drug in the selected cocktail a and on the difference between the patient’s condition from s_i^a to \hat{s}_i^a (Ernst et al., 2006). The results reported in Section 4.2.3 correspond to the performance of the greedy policy induced by the value function computed by the algorithms using all available sample transitions. The decision policies (in this case STI treatments) were evaluated for 5000 days starting from an “unhealthy” state corresponding to a basin of attraction of the ODEs describing the problem’s dynamics (see the papers by Adams et al. and Ernst et al.).

Epilepsy suppression: We used a generative model developed by Bush et al. (2009) to perform our experiments with the epilepsy suppression task. The model was generated based on labeled field potential recordings of five rat brain slices electrically stimulated at frequencies of 0.0 Hz, 0.5 Hz, 1.0 Hz, and 2.0 Hz. The data was used to construct a manifold embedding which in turn gave rise to the problem’s state space. The objective is to minimize the occurrence of seizures using as little stimulation as possible, therefore there is a negative reward associated with both events (see Section 4.2.4). Bush et al.’s generative model is public available as an environment for the RL-Glue package (Tanner and White, 2009). In our experiments the problem was modeled as a discounted task with $\gamma = 0.99$. The decision policies were evaluated on episodes of 10^5 transitions starting from a fixed set of 10 test states drawn uniformly at random from the problem’s state space.

Helicopter hovering: In the experiments with the helicopter hovering task we used the simulator developed by Abbeel et al. (2005), which is available as an environment for the RL-Glue package (Tanner and White, 2009). The simulator was built based on data collected from two separate flights of an XCell Tempest helicopter. The data was used to adjust the parameters of an “acceleration prediction model”, which is more accurate than the linear model normally adopted by industry. The objective in the problem is to keep the helicopter hovering as close as possible to a specific location. Therefore, at each time step the agent gets a negative reward proportional to the distance from the target position. Since the problem’s original action space is $A \equiv [-1, 1]^4$, we discretized each dimension using 4 break points distributed unevenly over $[-1, 1]$. We tried several possible discretizations and picked the one which resulted in the best performance of the SARSA

agent (see Section 5.2.3). After this process, the problem’s action space was redefined as $A \equiv \{-0.25, -0.05, +0.05, +0.25\}^4$. The problem was modeled as a discounted task with $\gamma = 0.99$. The decision policies were evaluated in episodes starting from the target position and ending when the helicopter crashed.

B.2 Algorithms

In all experiments, we used

$$\phi(z) \equiv \bar{\phi}(z) \equiv \exp(-z) \tag{66}$$

to define the kernels used by KBRL, LSPI, and KBSF. In the experiments involving a large number of sample transitions we used sparse kernels, that is, we only computed the μ largest values of $k_\tau(\bar{s}_i, \cdot)$ and the $\bar{\mu}$ largest values of $\bar{k}_\tau(\hat{s}_i^a, \cdot)$. In order to implement this feature, we used a KD-tree to find the μ ($\bar{\mu}$) nearest neighbors of \bar{s}_i (\hat{s}_i^a) and only computed k_τ (\bar{k}_τ) in these states (Bentley, 1975). The value of k_τ and \bar{k}_τ outside this neighborhood was truncated to zero (we used specialized data structures to avoid storing those).

We now list a few details regarding the algorithms’s implementations which were not described in the paper:

- **KBRL** and **KBSF**: We used modified policy iteration to compute \hat{Q}^* (Puterman and Shin, 1978). The value function of a fixed policy π was approximated through value iteration using the stop criterion described by Puterman (1994, Proposition 6.6.5) with $\varepsilon = 10^{-6}$. Table 1 shows the parameters’s values used by KBSF across the experiments.
- **LSPI**: As explained above, LSPI used the kernel derived from (66) as its basis function. Following Lagoudakis and Parr (2003), we adopted one block of basis functions for each action $a \in A$. Singular value decomposition was used to avoid eventual numerical instabilities in the system of linear equations constructed at each iteration of LSPI (Golub and Loan, 1993).
- **Fitted Q-iteration and extra trees**: FQIT has four main parameters: the number of iterations, the number of trees composing the ensemble, the number of candidate cut-points evaluated during the generation of the trees, and the minimum number of elements required to split a node, denoted here η_{\min} . In general, increasing the first three improves performance, while η_{\min} has an inverse relation with the quality of the final value function approximation. Our experiments indicate that the following configuration of FQIT usually results in good performance on the tasks considered in this paper: 50 iterations (with the structure of the trees fixed after the 10th one), an ensemble of 30 trees, and $d_{\mathbb{S}}$ candidate cut points, where $d_{\mathbb{S}}$ is the dimension of the state space \mathbb{S} . The parameter η_{\min} has a particularly strong effect on FQIT’s performance and computational cost, and its correct value seems to be more problem-dependent. Therefore, in all of our experiments we fixed the parameters of FQIT as described above and only varied η_{\min} .
- **SARSA**: We adopted the implementation of SARSA(λ) available in the RL-Glue package (Tanner and White, 2009). The algorithm uses gradient descent temporal-difference learning to configure a tile coding function approximator.

| Problem | Section | \bar{s}_i | m | τ | $\bar{\tau}$ | μ | $\bar{\mu}$ |
|-------------|---------|-------------|--------------------------|------------------|------------------|----------|-------------|
| Puddle | 4.2.1 | k -means | {10, 30, ..., 150} | {0.01, 0.1, 0.1} | {0.01, 0.1, 0.1} | ∞ | ∞ |
| Puddle | 5.2.1 | evenly | 100 | {0.01, 0.1, 0.1} | {0.01, 0.1, 0.1} | ∞ | ∞ |
| Single Pole | 4.2.2 | k -means | {10, 30, ..., 150} | 1 | {0.01, 0.1, 0.1} | ∞ | ∞ |
| Double Pole | 4.2.2 | k -means | {20, 40, ..., 200} | 1 | {0.01, 0.1, 0.1} | ∞ | ∞ |
| Triple Pole | 5.2.2 | on-line | on-line | 100* | 1* | 50* | 10* |
| HIV | 4.2.3 | random | {2000, 4000, ..., 10000} | 1 | 1 | 2* | 3* |
| Epilepsy | 4.2.4 | k -means | 50000* | 1 | {0.01, 0.1, 0.1} | 6* | 6* |
| Helicopter | 5.2.3 | k -means | 500* | 1 | 1 | 4* | 4* |

Table 1: Parameters used by KBSF on the computational experiments. The values marked with an asterisk (*) were determined by trial and error on preliminary tests. The remaining parameters were kept fixed from the start or were defined based on a very coarse search.

Appendix C. Table of Symbols

Table 2 shows the main symbols used in the paper. Auxiliary symbols and functions whose use is restricted to a specific part of the text are not listed in the table.

References

Pieter Abbeel, Varun Ganapathi, and Andrew Y. Ng. Learning vehicular dynamics, with application to modeling helicopters. In *Advances in Neural Information Processing Systems (NIPS)*, 2005.

Pieter Abbeel, Adam Coates, Morgan Quigley, and Andrew Y. Ng. An application of reinforcement learning to aerobatic helicopter flight. In *Advances in Neural Information Processing Systems (NIPS)*, 2007.

Brian M. Adams, Harvey T. Banks, Hee-Dae Kwon, and Hien T. Tran. Dynamic multidrug therapies for HIV: optimal and STI control approaches. *Mathematical Biosciences and Engineering*, 1(2):223–41, 2004.

Charles W. Anderson. *Learning and Problem Solving with Multilayer Connectionist Systems*. PhD thesis, Computer and Information Science, University of Massachusetts, 1986.

Andras Antos, Remi Munos, and Csaba Szepesvari. Fitted Q-iteration in continuous action-space MDPs. In *Advances in Neural Information Processing Systems (NIPS)*, pages 9–16, 2007.

Anwarissa Asbah, Andre M. S. Barreto, Clement Gehring, Joelle Pineau, and Doina Precup. Reinforcement learning competition 2013: Controllability and kernel-based stochastic factorization. In *Proceedings of the ICML Workshop on the Reinforcement Learning Competition*, 2013.

Christopher G. Atkeson and Juan Carlos Santamaria. A comparison of direct and model-based reinforcement learning. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 3557–3564, 1997.

| Symbol | Meaning | Defined in Sec. |
|--------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| γ | Discount factor in $[0, 1]$ | 2 |
| A, a | Action space, generic action | 2 |
| S | Generic state space | 2 |
| $\mathbb{S}, d_{\mathbb{S}}$ | Continuous state space, dimension of \mathbb{S} | 2.2 |
| S^a | Sample transitions associated with action a , $S^a \equiv \{(s_i^a, r_i^a, \hat{s}_i^a) i = 1, 2, \dots, n_a\}$ | 2.2 |
| $s, s_i, s_{(t)}, s_i^a, \hat{s}_i^a$ | Generic state, i^{th} state of an MDP, state occupied at time t , i^{th} tuple in S^a | 2 |
| $r, r_{(t)}, r_i^a$ | Generic reward, reward received at time t , i^{th} sampled reward in S^a | 2 |
| n_a, n, \hat{n} | Number of sample transitions in S^a , $n = \sum_a n_a$, and $\hat{n} = \max_a n_a$ | 2.2 |
| $P^a, P^\pi, \mathbf{P}^a, \mathbf{P}^\pi$ | Transition function associated with a or π and their matrix counterparts | 2.1 |
| $R^a, R^\pi, \mathbf{r}^a, \mathbf{r}^\pi$ | Reward function associated with a or π and corresponding reward vectors | 2.1 |
| M | MDP, $M \equiv (S, A, P^a, R^a, \gamma)$ | 2.1 |
| π, π^* | Generic decision policy, optimal decision policy in M | 2 |
| $V^\pi, V^*, \mathbf{v}^\pi, \mathbf{v}^*$ | Value function of π , optimal value function of M , vector counterparts | 2.1 |
| $Q^\pi, Q^*, \mathbf{Q}^\pi, \mathbf{Q}^*$ | Action-value functions of π and M and matrix counterparts | 2.1 |
| Γ | Operator: $\Gamma \mathbf{Q} = \mathbf{v} \iff v_i = \max_j q_{ij}, \forall i$ | 2.1 |
| Δ | Operator associated with M : $\Delta \mathbf{v} = \mathbf{Q} \iff q_{ia} = r_i^a + \gamma \sum_{j=1}^{ S } p_{ij}^a v_j, \forall i, a$ | 2.1 |
| T | Bellman operator of M , given by $T \equiv \Gamma \Delta$ | 2.1 |
| $\phi, \bar{\phi}$ | Non-increasing functions in $\mathbb{R}^+ \mapsto \mathbb{R}^+$ used to construct the kernels | 2.2, 4 |
| k_τ, \bar{k}_τ | Kernel functions: $k_\tau(s, s') = \phi(\ s - s'\ /\tau)$ and $\bar{k}_\tau(s, s') = \bar{\phi}(\ s - s'\ /\bar{\tau})$ | 2.2, 4 |
| $\tau, \bar{\tau}$ | Scalars defining the “widths” of k_τ and \bar{k}_τ | 2.2, 4 |
| $\kappa_\tau^a, \bar{\kappa}_\tau$ | Normalized kernels: $\kappa_\tau^a(s, s_i^a) = \frac{k_\tau(s, s_i^a)}{\sum_{j=1}^{n_a} k_\tau(s, s_j^a)}$ and $\bar{\kappa}_\tau(s, \bar{s}_i) = \frac{\bar{k}_\tau(s, \bar{s}_i)}{\sum_{j=1}^m \bar{k}_\tau(s, \bar{s}_j)}$ | 2.2, 4 |
| \bar{S} | Set of representative states, $\bar{S} \equiv \{\bar{s}_1, \bar{s}_2, \dots, \bar{s}_m\}$ | 4 |
| \bar{s}_i | i^{th} representative state | 4 |
| m | Number of representative states | 4 |
| $\mathbf{D}, \dot{\mathbf{D}}^a$ | Stochastic matrix in $\mathbb{R}^{n \times m}$ and its $n_a \times m$ block associated with a | 4 |
| $\mathbf{K}^a, \dot{\mathbf{K}}^a$ | Stochastic matrix in $\mathbb{R}^{m \times n}$ and its dense version in $\mathbb{R}^{m \times n_a}$ | 4 |
| rs | Function: $rs(s, i) = \bar{s}_k \iff \bar{s}_k$ is the i^{th} closest representative state to s | 4.1 |
| $dist$ | Function: $dist(s, i) = \ s - rs(s, i)\ $ | 4.1 |

Table 2: List of main symbols used throughout the paper. We use the same notation to refer to all MDPs and the associated elements, and resort to math accents to distinguish between them. So, for example, if \bar{M} is an MDP, the associated elements are referred to as $\bar{\mathbf{P}}^a, \bar{\mathbf{r}}^a, \bar{\mathbf{v}}^*, \bar{\mathbf{Q}}^*, \bar{\pi}^*, \bar{\Delta}$, and \bar{T} .

- Seema H. Bajaria, Glenn Webb, and Denise E. Kirschner. Predicting differential responses to structured treatment interruptions during HAART. *Bulletin of Mathematical Biology*, 66(5):1093 – 1118, 2004.
- André M. S. Barreto. Tree-based on-line reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 2417–2423, 2014.
- André M. S. Barreto and Marcelo D. Fragoso. Computing the stationary distribution of a finite Markov chain through stochastic factorization. *SIAM Journal on Matrix Analysis and Applications*, 32:1513–1523, 2011.
- André M. S. Barreto, Doina Precup, and Joelle Pineau. Reinforcement learning using kernel-based stochastic factorization. In *Advances in Neural Information Processing Systems (NIPS)*, pages 720–728, 2011.
- André M. S. Barreto, Doina Precup, and Joelle Pineau. On-line reinforcement learning using incremental kernel-based stochastic factorization. In *Advances in Neural Information Processing Systems (NIPS)*, pages 1484–1492, 2012.
- André M. S. Barreto, Joelle Pineau, and Doina Precup. Policy iteration based on stochastic factorization. *Journal of Artificial Intelligence Research*, 50:763–803, 2014.
- Andrew G. Barto, Richard S. Sutton, and Charles W. Anderson. Neuronlike adaptive elements that can solve difficult learning control problems. *IEEE Transactions on Systems, Man, and Cybernetics*, 13:834–846, 1983.
- Richard E. Bellman. *Dynamic Programming*. Princeton University Press, 1957.
- Richard E. Bellman. *Adaptive Control Processes*. Princeton University Press, 1961.
- Jon Louis Bentley. Multidimensional binary search trees used for associative searching. *Communications of the ACM*, 18(9):509–517, 1975.
- Dimitri P. Bertsekas and John N. Tsitsiklis. *Neuro-Dynamic Programming*. Athena Scientific, 1996.
- Alina Beygelzimer, Sham Kakade, and John Langford. Cover trees for nearest neighbor. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 97–104, 2006.
- Nikhil Bhat, Ciamac Moallemi, and Vivek Farias. Non-parametric approximate dynamic programming via the kernel method. In *Advances in Neural Information Processing Systems (NIPS)*, 2012.
- Ronen I. Brafman and Moshe Tennenholtz. R-MAX - a general polynomial time algorithm for near-optimal reinforcement learning. *Journal of Machine Learning Research*, 3:213–231, 2003.
- Keith Bush, Joelle Pineau, and Massimo Avoli. Manifold embeddings for model-based reinforcement learning of neurostimulation policies. In *Proceedings of the ICML/UAI/COLT Workshop on Abstraction in Reinforcement Learning*, 2009.

- Keith Bush and Joelle Pineau. Manifold embeddings for model-based reinforcement learning under partial observability. In *Advances in Neural Information Processing Systems (NIPS)*, pages 189–197, 2009.
- Moses Charikar, Chandra Chekuri, Tomás Feder, and Rajeev Motwani. Incremental clustering and dynamic information retrieval. In *Proceedings of the Annual ACM Symposium on Theory of Computing*, pages 626–635, 1997.
- Joel E. Cohen and Uriel G. Rothblum. Nonnegative ranks, decompositions and factorizations of nonnegative matrices. *Linear Algebra and its Applications*, 190:149–168, 1991.
- Adele Cutler and Leo Breiman. Archetypal analysis. *Technometrics*, 36(4):338–347, 1994.
- Eric V. Denardo. Contraction mappings in the theory underlying dynamic programming. *SIAM Review*, 9(2):165–177, 1967.
- Thomas G. Dietterich. Hierarchical reinforcement learning with the MAXQ value function decomposition. *Journal of Artificial Intelligence Research*, 13:227–303, 2000.
- Dominique M. Durand and Marom Bikson. Suppression and control of epileptiform activity by electrical stimulation: a review. *Proceedings of the IEEE*, 89(7):1065–1082, 2001.
- Yaakov Engel, Shie Mannor, and Ron Meir. Reinforcement learning with Gaussian processes. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 201–208, 2005.
- Damien Ernst, Guy-Bart Stan, Jorge Gonçalves, and Louis Wehenkel. Clinical data based optimal STI strategies for HIV: a reinforcement learning approach. In *Proceedings of the IEEE Conference on Decision and Control (CDC)*, 2006.
- Damien Ernst, Pierre Geurts, and Louis Wehenkel. Tree-based batch mode reinforcement learning. *Journal of Machine Learning Research*, 6:503–556, 2005.
- Amir-massoud Farahmand. *Regularization in reinforcement learning*. PhD thesis, University of Alberta, 2011.
- Amir-massoud Farahmand and Csaba Szepesvári. Model selection in reinforcement learning. *Machine Learning Journal*, 85(3):299–332, 2011.
- Amir-massoud Farahmand, Csaba Szepesvári, and Jean-Yves Audibert. Toward manifold-adaptive learning. In *NIPS Workshop on Topology learning*, 2007.
- Tomás Feder and Daniel Greene. Optimal algorithms for approximate clustering. In *Proceedings of the Annual ACM Symposium on Theory of Computing*, pages 434–444, 1988.
- Pierre Geurts, Damien Ernst, and Louis Wehenkel. Extremely randomized trees. *Machine Learning*, 36(1):3–42, 2006.
- Gene H. Golub and Charles F. Van Loan. *Matrix Computations*. The Johns Hopkins University Press, 2nd edition, 1993.

- Faustino J. Gomez, Juergen Schmidhuber, and Risto Miikkulainen. Efficient non-linear control through neuroevolution. In *Proceedings of the European Conference on Machine Learning (ECML)*, pages 654–662, 2006.
- Faustino J. Gomez. *Robust non-linear control through neuroevolution*. PhD thesis, The University of Texas at Austin, 2003. Technical Report AI-TR-03-303.
- Teofilo Gonzales. Clustering to minimize the maximum intercluster distance. *Theoretical Computer Science*, 38:293–306, 1985.
- Geoffrey Gordon. Stable function approximation in dynamic programming. Technical Report CMU-CS-95-103, Computer Science Department, Carnegie Mellon University, 1995.
- Steffen Grunewalder, Guy Lever, Luca Baldassarre, Massi Pontil, and Arthur Gretton. Modelling transition dynamics in MDPs with RKHS embeddings. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 535–542, 2012.
- Isabelle Guyon and André Elisseeff. An introduction to variable and feature selection. *Journal of Machine Learning Research*, 3:1157–1182, 2003.
- László Györfi, Michael Kohler, Adam Krzyżak, and Harro Walk. *A Distribution-Free Theory of Nonparametric Regression*. Springer Verlag, 2002.
- Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer, 2002.
- Kristin Jerger and Steven Schiff. Periodic pacing an in vitro epileptic focus. *Journal of Neurophysiology*, (2):876–879, 1995.
- Nicholas Jong and Peter Stone. Kernel-based models for reinforcement learning in continuous state spaces. In *Proceedings of the International Conference on Machine Learning (ICML)—Workshop on Kernel Machines and Reinforcement Learning*, 2006.
- Nicholas Jong and Peter Stone. Compositional models for reinforcement learning. In *Proceedings of the European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 644–659, 2009.
- Leonard Kaufman and Peter J. Rousseeuw. *Finding Groups in Data: an Introduction to Cluster Analysis*. John Wiley and Sons, 1990.
- Oliver B. Kroemer and Jan R. Peters. A non-parametric approach to dynamic programming. In *Advances in Neural Information Processing Systems (NIPS)*, pages 1719–1727, 2011.
- Branislav Kveton and Georgios Theocharous. Kernel-based reinforcement learning on representative states. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 124–131, 2012.
- Branislav Kveton and Georgios Theocharous. Structured kernel-based reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2013.

- Michail G. Lagoudakis and Ronald Parr. Least-squares policy iteration. *Journal of Machine Learning Research*, 4:1107–1149, 2003.
- Daniel D. Lee and H. Sebastian Seung. Learning the parts of objects by non-negative matrix factorization. *Nature*, 401:788–791, 1999.
- Donald Michie and Roger Chambers. BOXES: An experiment on adaptive control. *Machine Intelligence 2*, pages 125–133, 1968.
- Andrew W. Moore and Christopher G. Atkeson. Prioritized sweeping: Reinforcement learning with less data and less time. *Machine Learning*, 13:103–130, 1993.
- Rémi Munos and Csaba Szepesvári. Finite-time bounds for fitted value iteration. *Journal of Machine Learning Research*, 9:815–857, 2008.
- Andrew Y. Ng, H. Jin Kim, Michael I. Jordan, and Shankar Sastry. Autonomous helicopter flight via reinforcement learning. In *Advances in Neural Information Processing Systems (NIPS)*, 2003.
- Dirk Ormoneit and Peter Glynn. Kernel-based reinforcement learning in average-cost problems. *IEEE Transactions on Automatic Control*, 47(10):1624–1636, 2002.
- Dirk Ormoneit and Šaunak Sen. Kernel-based reinforcement learning. *Machine Learning*, 49 (2–3):161–178, 2002.
- Pentti Paatero and Unto Tapper. Positive matrix factorization: A non-negative factor model with optimal utilization of error estimates of data values. *Environmetrics*, 5:111–126, 1994.
- Martin L. Puterman. *Markov Decision Processes—Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc., 1994.
- Martin L. Puterman and Moon C. Shin. Modified policy iteration algorithms for discounted Markov decision problems. *Management Science*, 24(11):1127–1137, 1978.
- Carl Edward Rasmussen and Malte Kuss. Gaussian processes in reinforcement learning. In *Advances in Neural Information Processing Systems (NIPS)*, pages 751–759, 2004.
- Balaraman Ravindran. *An Algebraic Approach to Abstraction in Reinforcement Learning*. PhD thesis, University of Massachusetts, 2004.
- Gavin Rummery and Mahesan Niranjana. On-line Q-learning using connectionist systems. Technical Report CUED/F-INFENG/TR 166, Engineering Department, Cambridge University, 1994.
- Bernhard Schölkopf and Alexander J. Smola. *Learning with Kernels*. MIT Press, 2002.
- Jonathan Sorg and Satinder Singh. Transfer via soft homomorphisms. In *Autonomous Agents & Multiagent Systems/Agent Theories, Architectures, and Languages*, pages 741–748, 2009.

- Charles J. Stone. Optimal global rates of convergence for nonparametric regression. *The Annals of Statistics*, 10:1040–1053, 1982.
- Alexander L. Strehl and Michael L. Littman. An analysis of model-based interval estimation for Markov decision processes. *Journal of Computer and System Sciences*, 74(8):1309–1331, 2008.
- Richard S. Sutton. Generalization in reinforcement learning: Successful examples using sparse coarse coding. In *Advances in Neural Information Processing Systems (NIPS)*, pages 1038–1044, 1996.
- Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, 1998.
- Csaba Szepesvári. *Algorithms for Reinforcement Learning*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers, 2010.
- Brian Tanner and Adam M. White. RL-Glue: Language-independent software for reinforcement-learning experiments. *Journal of Machine Learning Research*, 10:2133–2136, 2009.
- Gavin Taylor and Ronald Parr. Kernelized value function approximation for reinforcement learning. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 1017–1024, 2009. .
- Stephen A. Vavasis. On the complexity of nonnegative matrix factorization. *SIAM Journal on Optimization*, 20:1364–1377, 2009.
- Ward Whitt. Approximations of dynamic programs, I. *Mathematics of Operations Research*, 3(3):231–243, 1978.
- Alexis P. Wieland. Evolving neural network controllers for unstable systems. In *Proceedings of the International Joint Conference on Neural Networks*, pages 667–673, 1991.
- Xin Xu, Tao Xie, Dewen Hu, and Xicheng Lu. Kernel Least-Squares Temporal Difference Learning. *Information Technology*, pages 54–63, 2005.
- Yinyu Ye. The simplex and policy-iteration methods are strongly polynomial for the Markov decision problem with a fixed discount rate. *Mathematics of Operations Research*, 36(4): 593–603, 2011.