



ELSEVIER

Contents lists available at ScienceDirect

Automatica

journal homepage: www.elsevier.com/locate/automatica

Natural actor–critic algorithms[☆]

Shalabh Bhatnagar^{a,*}, Richard S. Sutton^b, Mohammad Ghavamzadeh^c, Mark Lee^b

^a Department of Computer Science and Automation, Indian Institute of Science, Bangalore 560 012, India

^b The RLAI Laboratory, Department of Computing Science, University of Alberta, Edmonton, Alberta, Canada T6G 2E8

^c INRIA Lille - Nord Europe, Team Sequel, France

ARTICLE INFO

Article history:

Received 23 May 2007

Received in revised form

1 January 2009

Accepted 3 July 2009

Available online xxx

Keywords:

Actor–critic reinforcement learning algorithms

Policy-gradient methods

Approximate dynamic programming

Function approximation

Two-timescale stochastic approximation

Temporal difference learning

Natural gradient

ABSTRACT

We present four new reinforcement learning algorithms based on actor–critic, natural-gradient and function-approximation ideas, and we provide their convergence proofs. Actor–critic reinforcement learning methods are online approximations to policy iteration in which the value-function parameters are estimated using temporal difference learning and the policy parameters are updated by stochastic gradient descent. Methods based on policy gradients in this way are of special interest because of their compatibility with function-approximation methods, which are needed to handle large or infinite state spaces. The use of temporal difference learning in this way is of special interest because in many applications it dramatically reduces the variance of the gradient estimates. The use of the natural gradient is of interest because it can produce better conditioned parameterizations and has been shown to further reduce variance in some cases. Our results extend prior two-timescale convergence results for actor–critic methods by Konda and Tsitsiklis by using temporal difference learning in the actor and by incorporating natural gradients. Our results extend prior empirical studies of natural actor–critic methods by Peters, Vijayakumar and Schaal by providing the first convergence proofs and the first fully incremental algorithms.

© 2009 Elsevier Ltd. All rights reserved.

1. Introduction

Many problems of scientific and economic importance are optimal sequential decision problems and as such can be formulated as Markov decision processes (MDPs) (Bertsekas & Tsitsiklis, 1996; Rust, 1996; White, 1993). In some cases, MDPs can be solved analytically, and in many cases they can be solved iteratively by dynamic programming or linear programming. However, in other cases these methods cannot be applied either because the state space is too large, a system model is available only as a simulator, or no system model is available. It is in these cases that the techniques and algorithms of reinforcement learning (RL) may be helpful.

Reinforcement learning (Bertsekas & Tsitsiklis, 1996; Sutton & Barto, 1998) can be viewed as a broad class of sample-based methods for solving MDPs. In place of a model, these methods use sample trajectories of the system and the controller interacting, such as could be obtained from a simulation. It is not unusual in

practical applications for such a simulator to be available when an explicit transition-probability model of the sort suitable for use by dynamic or linear programming is not (Crites & Barto, 1998; Tesauro, 1995). Reinforcement learning methods can also be used with no model at all, by obtaining sample trajectories by direct interaction with the system (Kohl & Stone, 2004; Ng et al., 2004).

One of the biggest challenges to solve MDPs with conventional methods is handling large state (and action) spaces. This is sometimes known as the “curse of dimensionality” because of the tendency of the size of a state space to grow exponentially with the number of its dimensions. The computational effort required to solve an MDP thus increases exponentially with the dimension and cardinality of the state space. A natural and venerable way of addressing the curse is to approximate the value function and policy parametrically with a number of parameters much smaller than the size of the state space (Bellman & Dreyfus, 1959). However a straightforward application of such function-approximation methods to dynamic programming has not proved effective on large problems. Some work with RL and function approximation has also run into problems of convergence and instability (Baird, 1995; Boyan & Moore, 1995), but about a decade ago it was established that if trajectories were sampled according to their distribution under the target policy (the on-policy distribution) then convergence could be assured for linear feature-based function approximators (Sutton, 1996; Tadic, 2001; Tsitsiklis & Van Roy, 1997).

[☆] This paper was not presented at any IFAC meeting. This paper was recommended for publication in revised form by Associate Editor Giuseppe De Nicolao under the direction of Editor Ian R. Petersen.

* Corresponding author. Tel.: +91 80 2293 2987; fax: +91 80 2360 2911.

E-mail addresses: shalabh@csa.iisc.ernet.in (S. Bhatnagar), sutton@cs.ualberta.ca (R.S. Sutton), mohammad.ghavamzadeh@inria.fr (M. Ghavamzadeh), mlee@cs.ualberta.ca (M. Lee).

The rest of the paper is organized as follows. In Section 2 we present our RL framework and provide an overview of policy-gradient methods. In Section 3 we discuss policy-gradient methods with function approximation and present some preliminary results. We show here in particular that the minimum variance baseline for the action-value function corresponds to the state-value function and obtain a form of bias in gradient estimates that results from the use of function approximation. Our four actor-critic algorithms and their convergence analysis are presented in Sections 4 and 5, respectively. In Section 6 we discuss the relationship of our algorithms to the actor-critic algorithm of Konda and Tsitsiklis (2003) and to the natural actor-critic algorithm of Peters et al. (2003). Section 7 contains concluding remarks.

2. The policy-gradient framework

We consider the standard reinforcement learning framework (e.g., see Bertsekas & Tsitsiklis, 1996; Sutton & Barto, 1998) in which a learning agent interacts with a stochastic environment. The overall model we consider is that of a discrete time Markov decision process (MDP) with finite numbers of states and actions, and bounded rewards. We allow \mathcal{S} and \mathcal{A} to respectively denote the state and action spaces of this MDP. For simplicity, we assume that \mathcal{S} is the set $\mathcal{S} = \{1, \dots, n\}$. We denote by $s_t, a_t,$ and $r_t,$ the state, action, and reward at time $t,$ respectively. We assume that reward is stochastic, real-valued and uniformly bounded. For simplicity and ease of notation, we assume that all actions in \mathcal{A} are feasible in each state. The state transition probabilities for the environment will be characterized by $P(s, a, s') = \Pr(s_{t+1} = s' | s_t = s, a_t = a), \forall s, s' \in \mathcal{S}, a \in \mathcal{A}.$ Further, the single-stage expected reward when action a is taken in state s will be denoted $R(s, a) = \mathbf{E}[r_{t+1} | s_t = s, a_t = a].$

An admissible policy $\bar{\pi}$ is a decision rule that is described by a sequence of functions $\bar{\pi} = \{\mu_0, \mu_1, \dots\}$ such that each $\mu_t : \mathcal{S} \rightarrow \mathcal{A},$ with action $\mu_t(s)$ taken in state s at instant $t \geq 0.$ A stationary policy is a time invariant decision rule, i.e., one for which $\mu_t = \mu, \forall t \geq 0,$ for some $\mu : \mathcal{S} \rightarrow \mathcal{A}.$ Most often, one refers to the function μ itself as the stationary policy. A stationary randomized policy π that we refer to as simply a randomized policy is specified via a probability distribution $\pi(s, \cdot)$ over $\mathcal{A},$ for $s \in \mathcal{S}.$ Under the long-run average reward setting considered in this paper, it can be shown that a stationary optimal policy exists e.g., (see Puterman, 1994). Note that any stationary policy is trivially a randomized policy as well. We motivate the following discussion from the viewpoint of randomized policies as we consider a parameterized class of these in this paper. From now on, for simplicity, we shall refer to a randomized policy as a policy. For a given policy, the sequence of states produced by the MDP is a Markov chain. Throughout the paper we assume

(A1) Under any policy $\pi,$ the Markov chain resulting from the MDP is irreducible and aperiodic.

Let $d^\pi(s)$ denote the stationary probability of the Markov chain being in state $s \in \mathcal{S},$ and let $d^\pi = (d^\pi(s), s \in \mathcal{S}).$ Our aim is to find a policy π that maximizes the long-run average reward $J(\pi)$ given by

$$J(\pi) = \lim_{T \rightarrow \infty} \frac{1}{T} \mathbf{E} \left[\sum_{t=0}^{T-1} r_{t+1} \middle| \pi \right] \\ = \sum_{s \in \mathcal{S}} d^\pi(s) \sum_{a \in \mathcal{A}} \pi(s, a) R(s, a). \quad (2)$$

The limit in (2) is well defined by (A1). Let π^{opt} denote an optimal policy $\pi^{opt} = \arg \max_{\pi} J(\pi).$ Further, we shall denote by $Q^\pi(s, a),$

the expected differential reward associated with a state-action pair $(s, a),$ given policy $\pi,$ that is defined by $\forall s \in \mathcal{S}, a \in \mathcal{A},$

$$Q^\pi(s, a) = \sum_{t=1}^{\infty} \mathbf{E}[r_{t+1} - J(\pi) | s_0 = s, a_0 = a, \pi].$$

Likewise, we denote by $V^\pi(s),$ the expected differential reward associated with a state s when actions are selected according to policy $\pi.$ Here $V^\pi(s) = \sum_{a \in \mathcal{A}} \pi(s, a) Q^\pi(s, a).$ The Poisson equation under policy π is given by

$$J(\pi) + V^\pi(s) = \sum_{a \in \mathcal{A}} \pi(s, a) \left[R(s, a) + \sum_{s' \in \mathcal{S}} P(s, a, s') V^\pi(s') \right], \quad (3)$$

$s \in \mathcal{S}$ (Puterman, 1994). In policy-gradient methods, we define a class of parameterized randomized policies $\{\pi^\theta(s, \cdot), s \in \mathcal{S}, \theta \in \mathbb{R}^{d_1}\},$ estimate the gradient of the average reward with respect to the policy parameters θ from the observed states, actions, and rewards, and then improve the policy by adjusting its parameters in the direction of an estimate of the gradient of J with respect to $\theta.$ Because in this setting a policy π is represented by its parameters θ, J can be viewed as a function of θ and by abuse of notation, we let $J(\theta)$ denote the long-run average reward when the parameter is $\theta.$ In what follows, we shall interchangeably use $J(\pi)$ or $J(\theta)$ to denote the long-run average reward when the policy π or its associated parameter θ are to be emphasized. We also drop θ from $\pi^\theta,$ and simply denote this quantity as $\pi.$ The optimum parameter can now be obtained as $\theta^{opt} = \arg \max_{\theta} J(\theta).$ The following assumption is a standard requirement in policy-gradient methods.

(A2) For any state-action pair $(s, a), \pi(s, a)$ is continuously differentiable in the parameter $\theta.$

Previous works (Baxter & Bartlett, 2001; Konda & Tsitsiklis, 2003; Marbach & Tsitsiklis, 2001; Sutton et al., 2000) have shown that the gradient of the average reward for parameterized policies that satisfy (A1) and (A2) is given by¹

$$\nabla J(\pi) = \sum_{s \in \mathcal{S}} d^\pi(s) \sum_{a \in \mathcal{A}} \nabla \pi(s, a) Q^\pi(s, a). \quad (4)$$

For the case of Markov processes with a parameterized infinitesimal generator, a similar expression was obtained by Cao and Chen (1997). Observe that if $b(s)$ is any given function of s (also called a *baseline*), then

$$\sum_{s \in \mathcal{S}} d^\pi(s) \sum_{a \in \mathcal{A}} \nabla \pi(s, a) b(s) \\ = \sum_{s \in \mathcal{S}} d^\pi(s) b(s) \nabla \left(\sum_{a \in \mathcal{A}} \pi(s, a) \right) \\ = \sum_{s \in \mathcal{S}} d^\pi(s) b(s) \nabla(1) = 0, \quad (5)$$

and thus, for any baseline $b(s),$ the gradient of the average reward can be written as

$$\nabla J(\pi) = \sum_{s \in \mathcal{S}} d^\pi(s) \sum_{a \in \mathcal{A}} \nabla \pi(s, a) [Q^\pi(s, a) \pm b(s)]. \quad (6)$$

The baseline $b(s)$ can be chosen in a way that the variance of the gradient estimates $\nabla J(\pi)$ is minimized (Greensmith, Bartlett, & Baxter, 2004).

The *natural* gradient, denoted $\tilde{\nabla} J(\pi),$ can be calculated by linearly transforming the *regular* gradient, $\nabla J(\pi),$ using the inverse Fisher information matrix of the policy: $\tilde{\nabla} J(\pi) = G(\theta)^{-1} \nabla J(\pi).$

¹ In the rest of the paper we use the notation ∇ to denote ∇_{θ} – the gradient with respect to the policy parameters.

The Fisher information matrix $G(\theta)$ can be seen (Bagnell & Schneider, 2003; Kakade, 2002; Peters et al., 2003) to be

$$G(\theta) = \mathbf{E}_{s \sim d^\pi, a \sim \pi} [\nabla \log \pi(s, a) \nabla \log \pi(s, a)^\top] \\ = \sum_{s \in \mathcal{S}} d^\pi(s) \sum_{a \in \mathcal{A}} \pi(s, a) \nabla \log \pi(s, a) \nabla \log \pi(s, a)^\top. \quad (7)$$

Matrix $G(\theta)$ plays an important role in the algorithms that use natural gradients (Kakade, 2002; Peters & Schaal, 2008). Here $\mathbf{E}_{s \sim d^\pi, a \sim \pi}[\cdot]$ denotes the expectation under the conditional joint distribution where states are first selected according to distribution d^π , and then given that a state s is selected, actions are selected according to distribution $\pi(s, \cdot)$. The Fisher information matrix is clearly positive definite (Kakade, 2002).

A well-studied example of parameterized randomized policies, which we use in our experiments, is the Gibbs (or Boltzmann) distribution having the form

$$\pi(s, a) = \frac{e^{\theta^\top \phi_{sa}}}{\sum_{a' \in \mathcal{A}} e^{\theta^\top \phi_{sa'}}}, \quad \forall s \in \mathcal{S}, \forall a \in \mathcal{A}, \quad (8)$$

where each ϕ_{sa} is a d_1 -dimensional feature vector for the state-action pair (s, a) .

3. Policy gradient with function approximation

Now consider the case in which the action-value function for a fixed policy π , Q^π , is approximated by a learned function approximator. If the approximation is sufficiently good, we might hope to use it in place of Q^π in Eqs. (4) and (6), and still point roughly in the direction of the true gradient. Sutton et al. (2000) showed that if the approximation \hat{Q}_w^π with parameter $w \in \mathbb{R}^{d_1}$ is compatible, i.e., $\nabla_w \hat{Q}_w^\pi(s, a) = \nabla \log \pi(s, a)$, and minimizes the mean-squared error

$$\mathcal{E}^\pi(w) = \sum_{s \in \mathcal{S}} d^\pi(s) \sum_{a \in \mathcal{A}} \pi(s, a) [Q^\pi(s, a) - \hat{Q}_w^\pi(s, a)]^2 \quad (9)$$

for parameter value w^* , then we can replace Q^π with $\hat{Q}_{w^*}^\pi$ in Eqs. (4) and (6). We work with linear approximation $\hat{Q}_w^\pi(s, a) = w^\top \psi_{sa}$ in which the ψ_{sa} 's are compatible features defined according to $\psi_{sa} = \nabla \log \pi(s, a)$. Convergence of a temporal difference critic under a linear approximation when trajectories are sampled according to their distribution under the target policy has been established earlier (Sutton, 1996; Tadic, 2001; Tsitsiklis & Van Roy, 1997). Note that compatible features are well defined under (A2). As an example, the compatible features for the Gibbs policy in Eq. (8) are $\psi_{sa} = \phi_{sa} - \sum_{a' \in \mathcal{A}} \pi(s, a') \phi_{sa'}$. The Fisher information matrix of Eq. (7) can be written using the compatible features as

$$G(\theta) = \mathbf{E}_{s \sim d^\pi, a \sim \pi} [\psi_{sa} \psi_{sa}^\top] = \sum_{s \in \mathcal{S}} d^\pi(s) \sum_{a \in \mathcal{A}} \pi(s, a) \psi_{sa} \psi_{sa}^\top. \quad (10)$$

Suppose $\mathcal{E}^\pi(w)$ denotes the mean-squared error

$$\mathcal{E}^\pi(w) = \sum_{s \in \mathcal{S}} d^\pi(s) \sum_{a \in \mathcal{A}} \pi(s, a) [Q^\pi(s, a) - w^\top \psi_{sa} - b(s)]^2 \quad (11)$$

of our compatible linear parameterized approximation, $w^\top \psi_{sa}$, and an arbitrary baseline $b(s)$. Let $w^* = \arg \min_w \mathcal{E}^\pi(w)$ denote the optimal parameter. Lemma 1 shows that the value of w^* does not depend on the given baseline $b(s)$; as a result the mean-squared error problems of Eqs. (9) and (11) have the same solutions. Next in Lemma 2, we show that if the parameter is set to be equal to w^* , then the resulting mean-squared error $\mathcal{E}^\pi(w^*)$ (now treated as a function of the baseline $b(s)$) is further minimized when $b(s) = V^\pi(s)$ (see also Chapter 11 of Meyn, 2007). In other words, the variance in the action-value-function estimator

is minimized if the baseline is chosen to be the value function itself.² The proofs of Lemmas 1 and 2 can be found in Bhatnagar et al. (2009, 2008).

Lemma 1. *The optimum weight parameter w^* for any given θ (policy π) satisfies³*

$$w^* = G(\theta)^{-1} \mathbf{E}_{s \sim d^\pi, a \sim \pi} [Q^\pi(s, a) \psi_{sa}].$$

Lemma 2. *For any given policy π , the minimum variance baseline $b^*(s)$ in the action-value-function estimator corresponds to the state-value function $V^\pi(s)$.*

From Lemma 1, $w^{*\top} \psi_{sa}$ is a least-squared optimal parametric representation for the action-value function $Q^\pi(s, a)$. On the other hand, from Lemma 2, the same is also a least-squared optimal parametric representation for the advantage function $A^\pi(s, a) = Q^\pi(s, a) - V^\pi(s)$. The mean-squared error (11) is seen to be minimized w.r.t. the baseline $b(s)$ for $b^*(s) = V^\pi(s)$, thereby making it more meaningful to consider $w^{*\top} \psi_{sa}$ to be the least-squared optimal parametric representation for the advantage function rather than the action-value function itself.

The temporal difference (TD) error δ_t is a random quantity that is defined according to

$$\delta_t = r_{t+1} - \hat{J}_{t+1} + \hat{V}_{s_{t+1}} - \hat{V}_{s_t}, \quad (12)$$

where \hat{V}_{s_i} is an unbiased estimate of the differential reward in states s_i , $i = t, t + 1$. Likewise, \hat{J}_{t+1} is an unbiased estimate of the average reward. Thus, in particular, these estimates satisfy $\mathbf{E}[\hat{V}_{s_t} | s_t, \pi] = V^\pi(s_t)$ and $\mathbf{E}[\hat{J}_{t+1} | s_t, \pi] = J(\pi)$, for any $t \geq 0$, respectively. We assume here that actions are chosen according to policy π . The next lemma is also a simple result that shows that δ_t is an unbiased estimate of the advantage function A^π , see Bhatnagar et al. (2009, 2008) for a proof. A proof of this lemma is also available in Peters et al. (2003) and Peters and Schaal (2008).

Lemma 3. *Under given policy π with actions chosen according to it, we have*

$$\mathbf{E}[\delta_t | s_t, a_t, \pi] = A^\pi(s_t, a_t).$$

By setting the baseline $b(s)$ equal to the value function $V^\pi(s)$, Eq. (6) can be written as

$$\nabla J(\pi) = \sum_{s \in \mathcal{S}} d^\pi(s) \sum_{a \in \mathcal{A}} \pi(s, a) \psi_{sa} A^\pi(s, a).$$

From Lemma 3, δ_t is an unbiased estimate of the advantage function $A^\pi(s, a)$. Thus, $\widehat{\nabla J}(\pi) = \delta_t \psi_{s_t a_t}$ is an unbiased estimate of $\nabla J(\pi)$. However, calculating δ_t requires having estimates, \hat{J} , \hat{V} , of the average reward and the value function. While an average reward estimate is simple enough to obtain given the single-stage reward function, the same is not necessarily true for the value function. We use function approximation for the value functions as well. Suppose f_s is a d_2 -dimensional feature vector for state s (for some $d_2 \geq 1$). We denote $f_s = (f_s(1), \dots, f_s(d_2))^\top$. One may then approximate $V^\pi(s)$ with $v^\top f_s$, where v is a d_2 -dimensional

² It is important to note that Lemma 2 is not about the minimum variance baseline for gradient estimation. It is about the minimum variance baseline of the action-value-function estimator.

³ This lemma is similar to Theorem 1 by Kakade (2002), except that we consider baseline $b(s)$ which again can be seen as additional basis functions in the sense of Peters et al. (2003) and Peters and Schaal (2008).

weight vector which can be tuned (for a fixed policy π) using a TD algorithm. In our algorithms, we then use

$$\delta_t = r_{t+1} - \hat{J}_{t+1} + v_t^\top f_{s_{t+1}} - v_t^\top f_{s_t} \quad (13)$$

as an estimate for the TD error, where v_t corresponds to the value function parameter at time t . From now on, unless explicitly mentioned, we shall consider δ_t to be defined according to (13). Let $\bar{V}^\pi(s)$ denote the quantity

$$\bar{V}^\pi(s) = \sum_{a \in \mathcal{A}} \pi(s, a) \left[R(s, a) - J(\pi) + \sum_{s' \in \mathcal{S}} P(s, a, s') v^{\pi^\top} f_{s'} \right], \quad (14)$$

where $v^{\pi^\top} f_{s'}$ is an estimate of the differential value function $V^\pi(s')$ that is obtained upon convergence of a TD recursion (above) viz., $\lim_{t \rightarrow \infty} v_t = v^\pi$ with probability one. Also, let δ_t^π denote the associated quantity

$$\delta_t^\pi = r_{t+1} - \hat{J}_{t+1} + v^{\pi^\top} f_{s_{t+1}} - v^{\pi^\top} f_{s_t}. \quad (15)$$

Here r_{t+1} and \hat{J}_{t+1} are the same as before. Then δ_t^π corresponds to a stationary estimate of the TD error (with function approximation) under policy π . We have the following analog of Theorem 1 of Sutton et al. (2000).

Lemma 4.

$$\mathbf{E}[\delta_t^\pi \psi_{s_t a_t} | \theta] = \nabla J(\pi) + \sum_{s \in \mathcal{S}} d^\pi(s) [\nabla \bar{V}^\pi(s) - \nabla v^{\pi^\top} f_s].$$

Proof. A simple calculation shows that

$$\begin{aligned} \mathbf{E}[\delta_t^\pi \psi_{s_t a_t} | \theta] &= \sum_{s \in \mathcal{S}} d^\pi(s) \sum_{a \in \mathcal{A}} \nabla \pi(s, a) \left[R(s, a) - J(\pi) \right. \\ &\quad \left. + \sum_{s' \in \mathcal{S}} P(s, a, s') v^{\pi^\top} f_{s'} - v^{\pi^\top} f_s \right]. \end{aligned} \quad (16)$$

Now from (14),

$$\begin{aligned} \nabla \bar{V}^\pi(s) &= \sum_{a \in \mathcal{A}} \nabla \pi(s, a) \left[R(s, a) - J(\pi) + \sum_{s' \in \mathcal{S}} P(s, a, s') v^{\pi^\top} f_{s'} \right] \\ &\quad + \sum_{a \in \mathcal{A}} \pi(s, a) \left[-\nabla J(\pi) + \sum_{s' \in \mathcal{S}} P(s, a, s') \nabla v^{\pi^\top} f_{s'} \right]. \end{aligned}$$

From (16) and the above, we get

$$\begin{aligned} \sum_{s \in \mathcal{S}} d^\pi(s) \nabla \bar{V}^\pi(s) &= \mathbf{E}[\delta_t^\pi \psi_{s_t a_t} | \theta] - \nabla J(\pi) \\ &\quad + \sum_{s \in \mathcal{S}} d^\pi(s) \sum_{a \in \mathcal{A}} \pi(s, a) \sum_{s' \in \mathcal{S}} P(s, a, s') \nabla v^{\pi^\top} f_{s'}. \end{aligned} \quad (17)$$

Now observe that $d^\pi(s)$ correspond to the stationary probabilities that satisfy, $\forall s \in \mathcal{S}$,

$$d^\pi(s) = \sum_{s'' \in \mathcal{S}} d^\pi(s'') p^\pi(s'', s), \quad \sum_{s'' \in \mathcal{S}} d^\pi(s'') = 1, \quad (18)$$

where $p^\pi(s'', s) = \sum_{a \in \mathcal{A}} \pi(s'', a) P(s'', a, s)$ are the transition probabilities of the resulting Markov chain under policy π . Hence,

$$\begin{aligned} \sum_{s \in \mathcal{S}} d^\pi(s) \sum_{a \in \mathcal{A}} \pi(s, a) \sum_{s' \in \mathcal{S}} P(s, a, s') \nabla v^{\pi^\top} f_{s'} \\ = \sum_{s' \in \mathcal{S}} \sum_{s \in \mathcal{S}} d^\pi(s) p^\pi(s, s') \nabla v^{\pi^\top} f_{s'} = \sum_{s' \in \mathcal{S}} d^\pi(s') \nabla v^{\pi^\top} f_{s'}. \end{aligned} \quad (19)$$

The claim now follows from (17). \square

Note that according to Theorem 1 of Sutton et al. (2000), $\mathbf{E}[\delta_t \psi_{s_t a_t} | \theta] = \nabla J(\pi)$, provided δ_t is defined according to (12). For the case with function approximation that we study, from Lemma 4, the quantity $\sum_{s \in \mathcal{S}} d^\pi(s) [\nabla \bar{V}^\pi(s) - \nabla v^{\pi^\top} f_s]$ may be viewed as the error or bias in the estimate of the gradient of average reward that results from the use of function approximation. It is interesting to observe that this does not depend on the differential reward $V^\pi(s)$ that is obtained as a solution to (3). We also have

Corollary 1. $\sum_{s \in \mathcal{S}} d^\pi(s) [\bar{V}^\pi(s) - v^{\pi^\top} f_s] = 0.$

Proof. This follows directly from the definition of $\bar{V}^\pi(s)$ in (14), the definition of $J(\pi)$ in (2), and an analogous equation as (19) with $v^{\pi^\top} f_{s'}$ in place of $\nabla v^{\pi^\top} f_{s'}$. \square

4. Actor–critic algorithms

We present four new actor–critic algorithms in this section. They update the policy parameters along the direction of the average reward gradient. While estimates of the regular gradient are used for this purpose in Algorithm 1, natural gradient estimates are used in Algorithms 2–4. Let $\hat{V}(s, v) = v^\top f_s$ denote the parameterized approximation to the differential value function in state s . One can also denote the same as $\hat{V}(v) = \Phi v$, where Φ is an $n \times d_2$ matrix whose k th column ($k = 1, \dots, d_2$) is $f(k) = (f_s(k), s \in \mathcal{S})^\top$. We make the following assumption as in Tsitsiklis and Van Roy (1999) (see also Tsitsiklis & Van Roy, 1997).

(A3) The basis functions $\{f(k), k = 1, \dots, d_2\}$ are linearly independent. In particular, $d_2 \leq n$ and Φ has full rank. Also, for every $v \in \mathbb{R}^{d_2}$, $\Phi v \neq e$, where e is the n -dimensional vector with all entries equal to one.

Let $\{\alpha_t\}$ and $\{\beta_t\}$ be two step-size schedules that satisfy (20). Further, let $\{\xi_t\}$ defined by $\xi_t = c\alpha_t$ for some $c > 0$ be the step-size schedule for the average reward recursions in our algorithms.

$$\sum_t \alpha_t = \sum_t \beta_t = \infty, \quad \sum_t (\alpha_t^2 + \beta_t^2) < \infty, \quad \alpha_t = o(\beta_t). \quad (20)$$

From the last condition in (20), $\beta_t \rightarrow 0$ faster than α_t . Thus the critic is a faster recursion than the actor.

We now present our actor–critic algorithms. For the actor updates in our algorithms, we use a projection operator $\Gamma : \mathbb{R}^{d_1} \rightarrow \mathbb{R}^{d_1}$ that projects any $x \in \mathbb{R}^{d_1}$ to a compact set $C = \{x \mid q_i(x) \leq 0, i = 1, \dots, s\} \subset \mathbb{R}^{d_1}$, where $q_i(\cdot), i = 1, \dots, s$ are real-valued, continuously differentiable functions on \mathbb{R}^{d_1} that represent the constraints specifying the (above) compact region. Here for each x on the boundary of C , the gradients of the active constraints are considered to be linearly independent. This is the setting considered for projection-based algorithms in Chapter 5 of Kushner and Clark (1978). For any $x \in \mathbb{R}^{d_1}$, $\Gamma(x) \in C$ and in particular for $x \in C$, $\Gamma(x) = x$ itself. As explained in Chapter 2 of Kushner and Clark (1978), any compact hyperrectangle in \mathbb{R}^{d_1} is a special case of C (above). The projection method is an often used technique to ensure boundedness of iterates in stochastic approximation algorithms, see for instance, Abounadi, Bertsekas, and Borkar (2001), where it has been used in the context of a stochastic shortest path Q-learning algorithm. Some discussion on this is also available in Tsitsiklis (1994). The other approach (that is also usually taken, which we do not follow) is to simply assume that the iterates (see below) (23), (29), (34) and (39) are bounded without the projection, and then show their convergence under this assumption. In our experiments, however (see Bhatnagar et al., 2009), we do not project the iterates to a constraint region as they are seen to remain bounded (without projection). In Remark 2 (that follows Theorem 1), we explain the difficulties in proving boundedness of iterates in the absence of the projection operator $\Gamma(\cdot)$.

Algorithm 1 (Regular-Gradient Actor–Critic).

$$\hat{J}_{t+1} = (1 - \xi_t)\hat{J}_t + \xi_t r_{t+1}, \quad (21)$$

$$v_{t+1} = v_t + \alpha_t \delta_t f_{s_t}, \quad (22)$$

$$\theta_{t+1} = \Gamma(\theta_t + \beta_t \delta_t \psi_{s_t a_t}), \quad (23)$$

with δ_t as in (13). This is the only actor–critic algorithm presented in the paper that is based on the regular-gradient estimate. It stores two parameter vectors θ and v . Its per time-step computational cost is linear in the number of policy and value-function parameters.

The next algorithm is based on the natural-gradient estimate $\tilde{V}J(\theta_t) = G(\theta_t)^{-1} \delta_t \psi_{s_t a_t}$ in place of the regular-gradient estimate in Algorithm 1. We derive a procedure below for recursively estimating $G(\theta)^{-1}$ on a faster timescale. The above estimation is done on a faster scale so that convergence of the associated iterates is achieved prior to a θ -update. Suppose G_t^{-1} denote the t th estimate of $G(\theta)^{-1}$. Our procedure is obtained in a similar manner as the method described on pp. 147–152 of Widrow and Stearns (1985). The latter approach however considers the estimates as being obtained via a “fading memory” condition in which the most recent observation is given the highest weight. The weights themselves decrease geometrically over past observations. On the other hand, unlike Widrow & Stearns, 1985, we consider stationary averages that depend on parameter θ , that in turn gets updated along the “slower timescale”. This constitutes a natural setting for our algorithm. We show in Lemma 6 that $G_t^{-1} \rightarrow G(\theta)^{-1}$ as $t \rightarrow \infty$ with probability one. This is required for proving convergence of our algorithm. On the other hand, showing the same for the corresponding estimates in Widrow and Stearns (1985) does not seem possible because $G_t \not\rightarrow G(\theta)$ there.

We consider $G_t, t \geq 0$, defined as (the sample averages)

$$G_t = \frac{1}{t+1} \sum_{l=0}^t \psi_{s_l a_l} \psi_{s_l a_l}^\top.$$

Thus, one may obtain recursively

$$G_t = \left(1 - \frac{1}{t+1}\right) G_{t-1} + \frac{1}{t+1} \psi_{s_t a_t} \psi_{s_t a_t}^\top. \quad (24)$$

More generally, one may consider the recursion

$$G_t = (1 - \alpha_t) G_{t-1} + \alpha_t \psi_{s_t a_t} \psi_{s_t a_t}^\top, \quad (25)$$

where the step-size α_t is as before. This would correspond to a case of weighted averages (with the weights corresponding to the step-sizes α_t). However, through a stochastic approximation argument, one can see that (25) would asymptotically converge to $G(\theta)$, almost surely, if θ is held fixed. In fact, with an appropriate choice of $\{\alpha_t\}$, one can obtain faster convergence of iterates in (25) over those in (24). Using Sherman–Morrison matrix inversion lemma, one obtains

$$G_t^{-1} = \frac{1}{1 - \alpha_t} \left[G_{t-1}^{-1} - \alpha_t \frac{(G_{t-1}^{-1} \psi_{s_t a_t})(G_{t-1}^{-1} \psi_{s_t a_t})^\top}{1 - \alpha_t + \alpha_t \psi_{s_t a_t}^\top G_{t-1}^{-1} \psi_{s_t a_t}} \right]. \quad (26)$$

The following assumption is on the matrices G_t, G_t^{-1} .

(A4) We have $\sup_{t, \theta, s, a} \|G_t\|, \sup_{t, \theta, s, a} \|G_t^{-1}\| < \infty$.

This assumption will be used in proving the convergence of our Algorithms 2 and 4. It is similar to a corresponding requirement in the case of certain Hessian matrices in the Newton-based simulation optimization schemes in Bhatnagar (2005, 2007). A sufficient condition for both the requirements in (A4) is that (cf. pp. 35 of Bertsekas, 1999) for some scalars $c_1, c_2 > 0$,

$$c_1 \|z\|^2 \leq z^\top \psi_{sa} \psi_{sa}^\top z \leq c_2 \|z\|^2,$$

for all $s \in \mathcal{S}, a \in \mathcal{A}, z \in \mathbb{R}^{d_1}$ and θ . It is then easy to see that $\bar{c}_1 \|z\|^2 \leq z^\top G_t z \leq \bar{c}_2 \|z\|^2$, for all $t \geq 0$, and the eigenvalues of G_t lie between \bar{c}_1 and \bar{c}_2 . Here $\bar{c}_1 = \min(a, c_1)$ and $\bar{c}_2 = \max(a, c_2)$. Also, $\bar{c}_1, \bar{c}_2 > 0$. Hence, the procedure (below) does not get stuck at a nonstationary point. Under the above sufficient condition, (A4) follows from Propositions A.9 and A.15 of Bertsekas (1999).

Our second algorithm stores matrix G^{-1} and two parameter vectors θ and v . Its per time-step computational cost is linear in the number of value-function parameters and quadratic in the number of policy parameters.

Algorithm 2 (Natural-Gradient Actor–Critic with Fisher Information Matrix).

$$\hat{J}_{t+1} = (1 - \xi_t)\hat{J}_t + \xi_t r_{t+1}, \quad (27)$$

$$v_{t+1} = v_t + \alpha_t \delta_t f_{s_t}, \quad (28)$$

$$\theta_{t+1} = \Gamma(\theta_t + \beta_t G_t^{-1} \delta_t \psi_{s_t a_t}), \quad (29)$$

with δ_t as in (13). Also, the estimate of the inverse Fisher information matrix updated according to Eq. (26). Like Widrow and Stearns (1985), we let $G_0^{-1} = kI$, where I is a $d_1 \times d_1$ -dimensional identity matrix and $k > 0$. Thus G_0^{-1} and hence also G_0 are positive definite and symmetric matrices. From (25), $G_t, t \geq 1$ can be seen to be positive definite and symmetric because these are convex combinations of positive definite and symmetric matrices. Hence, $G_t^{-1}, t \geq 1$ are positive definite and symmetric matrices as well.

As we mentioned in Section 3, it is better to think of the compatible approximation $w^\top \psi_{sa}$ as an approximation of the advantage function rather than of the action-value function. In our next algorithm, we tune the weight parameters w in such a way as to minimize an estimate of the least-squared error $\mathcal{E}^\pi(w) = \mathbf{E}_{s \sim d_\pi, a \sim \pi} [(w^\top \psi_{sa} - A^\pi(s, a))^2]$. Note that the gradient of $\mathcal{E}^\pi(w)$ is

$$\nabla_w \mathcal{E}^\pi(w) = 2 \sum_{s \in \mathcal{S}} d^\pi(s) \sum_{a \in \mathcal{A}} \pi(s, a) [w^\top \psi_{sa} - A^\pi(s, a)] \psi_{sa}.$$

We use the following estimate of $\nabla_w \mathcal{E}^\pi(w)$.

$$\widehat{\nabla_w \mathcal{E}^\pi(w)} = 2(\psi_{s_t a_t} \psi_{s_t a_t}^\top w - \delta_t \psi_{s_t a_t}). \quad (30)$$

Hence, we update advantage parameters w along with value-function parameters v in the critic update of this algorithm as

$$w_{t+1} = w_t - \alpha_t \widehat{\nabla_w \mathcal{E}^\pi(w_t)} = w_t - \alpha_t (\psi_{s_t a_t} \psi_{s_t a_t}^\top w_t - \delta_t \psi_{s_t a_t}).$$

The factor 2 on the RHS of (30) does not play a role because of the diminishing step-size sequence $\alpha_t, t \geq 0$ and so has been dropped in the above recursion. We maximize the long-run average reward $J(\theta)$ along the slower timescale and use the natural-gradient estimate for this purpose. Like Peters et al. (2003) and Peters and Schaal (2008), the natural-gradient estimate that we use in the actor update of Algorithm 3 is $\tilde{V}J(\theta_t) = w_{t+1}$. This algorithm stores three parameter vectors, v, w , and θ . Its per time-step computational cost is linear in the number of value-function parameters and quadratic in the number of policy parameters.

Algorithm 3 (Natural-Gradient Actor–Critic with Advantage Parameters).

$$\hat{J}_{t+1} = (1 - \xi_t)\hat{J}_t + \xi_t r_{t+1}, \quad (31)$$

$$v_{t+1} = v_t + \alpha_t \delta_t f_{s_t}, \quad (32)$$

$$w_{t+1} = [I - \alpha_t \psi_{s_t a_t} \psi_{s_t a_t}^\top] w_t + \alpha_t \delta_t \psi_{s_t a_t}, \quad (33)$$

$$\theta_{t+1} = \Gamma(\theta_t + \beta_t w_{t+1}), \quad (34)$$

with δ_t as in (13). Although the estimates of $G(\theta)^{-1}$ are not explicitly computed and used in Algorithm 3, the convergence analysis of this algorithm in the next section shows that the overall scheme still moves in the direction of the natural gradient of average reward.

In Algorithm 4, however, we explicitly estimate $G(\theta)^{-1}$ (as in Algorithm 2), and use it in the critic update for w . The overall scheme is again seen to follow the direction of the natural gradient of average reward. Here, we let

$$\tilde{\nabla}_w \mathcal{E}^\pi(w) = 2G_t^{-1}(\psi_{s_t a_t} \psi_{s_t a_t}^\top w - \delta_t \psi_{s_t a_t}) \quad (35)$$

be the estimate of the natural gradient of the least-squared error $\mathcal{E}^\pi(w)$. This also simplifies the critic update for w . Further, we remove the factor 2 from the natural-gradient estimate (35) because of diminishing α_t , $t \geq 0$, as before. Algorithm 4 stores a matrix G^{-1} and three parameter vectors, v , w , and θ . Its per time-step computational cost is linear in the number of value-function parameters and quadratic in the number of policy parameters.

Algorithm 4 (Natural-Gradient Actor–Critic with Advantage Parameters and Fisher Information Matrix).

$$\hat{J}_{t+1} = (1 - \xi_t)\hat{J}_t + \xi_t r_{t+1}, \quad (36)$$

$$v_{t+1} = v_t + \alpha_t \delta_t f_{s_t}, \quad (37)$$

$$w_{t+1} = (1 - \alpha_t)w_t + \alpha_t G_t^{-1} \delta_t \psi_{s_t a_t}, \quad (38)$$

$$\theta_{t+1} = \Gamma(\theta_t + \beta_t w_{t+1}), \quad (39)$$

with δ_t as in (13) and where the estimate of the inverse of the Fisher information matrix is updated according to (26). As with Algorithm 2, we let $G_0^{-1} = kI$ with $k > 0$.

5. Convergence analysis

We now present the convergence analysis of our algorithms. The analysis mainly follows the ordinary differential equation (ODE) approach (Benveniste, Metivier, & Priouret, 1990; Kushner & Clark, 1978; Kushner & Yin, 1997). Note that the problem we consider is a maximization and not a minimization problem. For the purpose of analysis, we consider an associated problem with costs defined as negative rewards and our aim is to minimize the associated long-run average cost. The negative of the minimum cost thus obtained then corresponds to the maximum reward in the original problem. This is useful in pushing through certain stability arguments and showing convergence of iterates. Our algorithms use function approximation and aim at finding the local maxima of the average rewards. All our convergence results are in the Euclidean norm. Further, for any matrix A , we define its norm as the induced matrix norm $\|A\| = \max_{\|x\|=1} \|Ax\|$.

5.1. Convergence analysis for Algorithm 1

We require Assumptions (A1)–(A3) here. As explained above, one may view $-r_{t+1}$ as the cost incurred at instant t in a transformed problem. Because of the above, a change occurs only in the actor recursion (23) due to this transformation, and it becomes

$$\theta_{t+1} = \Gamma(\theta_t - \beta_t \delta_t \psi_{s_t a_t}). \quad (40)$$

Recursions for the average reward (21), TD error (13), and critic (22), being fixed point recursions (see Tsitsikis & Van Roy, 1999), are left unchanged. For any given policy π (along the faster timescale), average reward (21), TD error (13), and critic (22) recursions correspond to the TD(λ) recursions in Tsitsikis and Van Roy (1999) with $\lambda = 0$.

Let D denote the diagonal matrix with elements $d^\pi(s_1), \dots, d^\pi(s_n)$ along its diagonal. Let P^π be the probability matrix $P^\pi = [p^\pi(s, s')]_{s, s' \in \mathcal{S}}$ for the Markov chain under policy π and R^π be the corresponding column vector of average rewards whose i th element is $\sum_{a \in \mathcal{A}} \pi(s_i, a)R(s_i, a)$. Also, let $T : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be the operator given by $T(J) = R^\pi - J(\pi)e + P^\pi J$. The proof of

convergence of TD(λ) by Tsitsikis and Van Roy (1999) is based on a result by Benveniste et al. (1990). We provide in Lemma 5 an alternative simpler proof of convergence under the same assumptions used by Tsitsikis and Van Roy (1999), using a recently developed result by Borkar and Meyn (2000). We consider $\lambda = 0$ to suit our algorithm. The proof however carries through quite easily for $\lambda > 0$ as well. We have

Lemma 5. For any given π and $\{\hat{J}_t\}$, $\{v_t\}$ as in (21) and (22), respectively, we have $\hat{J}_t \rightarrow J(\pi)$ and $v_t \rightarrow v^\pi$ with probability one, where

$$J(\pi) = \sum_{s \in \mathcal{S}} d^\pi(s) \sum_{a \in \mathcal{A}} \pi(s, a)R(s, a), \quad (41)$$

is the average reward under π and v^π is obtained as the unique solution to

$$\Phi^\top D \Phi v^\pi = \Phi^\top D T(\Phi v^\pi). \quad (42)$$

Proof. The proof is based on verifying the Assumptions (A1)–(A2) of Borkar and Meyn (2000). First consider the average reward recursion (21). The ODE describing the asymptotic behavior of this recursion corresponds to

$$\dot{\eta} = -\eta + \sum_{s \in \mathcal{S}} d^\pi(s) \sum_{a \in \mathcal{A}} \pi(s, a)R(s, a). \quad (43)$$

Let $f(\eta)$ denote the RHS of (43). Then $f(\eta)$ is Lipschitz continuous in η . Let $f_\infty(\eta) = \lim_{r \rightarrow \infty} \frac{f(r\eta)}{r}$. The function $f_\infty(\eta)$ exists and is simply $f_\infty(\eta) = -\eta$. The origin is clearly an asymptotically stable equilibrium for the ODE $\dot{\eta} = f_\infty(\eta)$.

Now consider recursions for TD error (13) and critic (22). Consider the following ODE (in vector-matrix notation) associated with them.

$$\dot{v} = \Phi^\top D(T(\Phi v) - \Phi v). \quad (44)$$

Let $g^1(v)$ denote the RHS of (44). Then $g^1(v)$ is also Lipschitz continuous in v . Further, for $g_\infty^1(v) \triangleq \lim_{r \rightarrow \infty} \frac{g^1(rv)}{r}$, it can be seen that $g_\infty^1(v)$ exists and equals $g_\infty^1(v) = \Phi^\top D(P^\pi - I)\Phi v$, where I is the identity matrix. Consider now the system

$$\dot{v} = g_\infty^1(v). \quad (45)$$

Note that the matrix P^π has a simple eigenvalue of one and its remaining eigenvalues have real parts that are less than one. Thus $(P^\pi - I)$ will have one eigenvalue of zero and other eigenvalues with negative real parts. Also, corresponding to the eigenvalue zero, the matrix $(P^\pi - I)$ has a left eigenvector d^{π^\top} and a right eigenvector $e = (1, \dots, 1)^\top$ (the n -dimensional unit vector), respectively. Thus, in principle, the set of asymptotically stable fixed points of (45) would correspond to the set $\{\alpha v \mid \Phi v = e \text{ and } \alpha \in \mathbb{R}, \alpha \neq 0\} \cup \{v = 0\}$. Now by the second part of Assumption (A3), $\Phi v \neq e$, for every $v \in \mathbb{R}^{d^2}$. Thus the only asymptotically stable equilibrium for (45) is the origin.

Next, define $N^1(t)$, $M^1(t)$, $t \geq 0$, according to $N^1(t) = r_{t+1} - \mathbf{E}[r_{t+1} \mid \mathcal{F}_1(t)]$, $M^1(t) = \delta_t f_{s_t} - \mathbf{E}[\delta_t f_{s_t} \mid \mathcal{F}_1(t)]$, respectively, where $\mathcal{F}_1(t) = \sigma(v_r, \hat{J}_r, M^1(r), N^1(r), r \leq t)$. It is easy to see that

$$\mathbf{E}[\|N^1(t+1)\|^2 \mid \mathcal{F}_1(t)] \leq C_1(1 + \|\hat{J}_t\|^2 + \|v_t\|^2),$$

$$\mathbf{E}[\|M^1(t+1)\|^2 \mid \mathcal{F}_1(t)] \leq C_2(1 + \|v_t\|^2 + \|\hat{J}_t\|^2),$$

$t \geq 0$, for some $C_1, C_2 < \infty$. In fact, quantities $N^1(t)$ can be directly seen to be uniformly bounded almost surely. Thus Assumptions (A1) and (A2) of Borkar and Meyn (2000) can be seen to be satisfied in the case of the average reward (21), TD error (13), and critic (22) recursions. From Theorem 2.1 of Borkar and Meyn (2000), average reward, TD error, and critic iterates are uniformly bounded with

$r \geq R$ and $t \geq T$, for some $R, T > 0$, the trajectories $\hat{\phi}(t)$ of the ODE $\dot{\theta}_t = \frac{h^1(r\theta_t)}{r}$ should lie within a ball of radius $1/2$ around the origin. This can be shown provided the origin is a unique asymptotically stable attractor for the above ODEs for all $r \geq R$. Again, it is not clear if this is the case here. Next, note that the methods described by Abounadi et al. (2001) and Tsitsiklis (1994) for stability of iterates are for different classes of algorithms, largely of the Q-learning type, and are not directly applicable in our setting.

Finally, we discuss the use of the stochastic Lyapunov function method (Kushner & Yin, 1997) for stability of iterates in (40). The prime requirement here is that there exists a real-valued nonnegative function $W(\cdot)$ that satisfies

$$\mathbf{E}[W(\theta_{t+1}) | \theta_t = \theta] - W(\theta) \leq -K(\theta)$$

for all $\theta \in Q_\lambda \triangleq \{\theta | W(\theta) \leq \lambda\}$, where $K(\theta) \geq 0$ is continuous on Q_λ . Then by Thm. 4.1, pp. 80–81 of Kushner and Yin (1997), the stability and convergence of iterates would follow. Hence consider the recursion (40). By Taylor’s expansion for “small” β_t assuming a smooth $W(\cdot)$, one gets

$$\mathbf{E}[W(\theta_{t+1}) | \theta_t] \approx W(\theta_t) - \beta_t \mathbf{E}[\delta_t^\pi \psi_{s_t a_t} | \theta_t]^\top \nabla W(\theta_t) + o(\beta_t). \quad (49)$$

It appears difficult to obtain such a $W(\cdot)$ here. On the other hand, if we use the look up table representation (viz., $d_2 = n$ in Assumption (A3) or that δ_t is as in (12)), then from Lemma 4 above, as also Theorem 1 of Sutton et al. (2000), one would get $\mathbf{E}[\delta_t \psi_{s_t a_t} | \theta] = \nabla_\theta J(\theta)$. Then $W(\theta) = J(\theta)$ would serve as a Lyapunov function and the iterates (40) (without the projection) will be bounded and almost surely convergent, in lieu of Theorem 4.1 of Kushner and Yin (1997). It is only because of the use of function approximation in the iterates that a Lyapunov function is hard to obtain. However, in our experiments, we do not use projection but still observe that the iterates remain bounded and convergence is achieved.

Note that if function approximation is not used, $J(\theta)$ also serves as a Lyapunov function for the ODE associated with (40) without the projection. When function approximation is used (as with our case), the above problem of finding a suitable Lyapunov function (now) for the associated ODE also carries over and it is difficult to suitably characterize the set of stable attractors.

Remark 1 and many of the arguments in the analysis of Alg. 1 are valid for the analysis of the other algorithms. We skip the details in such cases to avoid repetition.

5.2. Convergence analysis for Algorithm 2

The analysis in Lemma 5 of the recursions for the average reward (27), TD error (13), and critic (28) proceeds in the same manner as for Algorithm 1. We thus concentrate on showing convergence of the recursion for the inverse of the Fisher information matrix (26) and the actor recursion (29). We assume (A1)–(A4) for our analysis here. We have

Lemma 6. For any given parameter θ , G_t^{-1} , $t \geq 1$, in (26) satisfy $G_t^{-1} \rightarrow G(\theta)^{-1}$, as $t \rightarrow \infty$ with probability one.

Proof. It is easy to see from recursion (25) that $G_t \rightarrow G(\theta)$ as $t \rightarrow \infty$ with probability one, for any given θ held fixed. Now for fixed θ , we have,

$$\begin{aligned} \|G_t^{-1} - G(\theta)^{-1}\| &= \|G(\theta)^{-1}(G(\theta) - G_t)G_t^{-1}\| \\ &\leq \sup_\theta \|G(\theta)^{-1}\| \sup_{t,s,a} \|G_t^{-1}\| \cdot \|G(\theta) - G_t\| \\ &\rightarrow 0 \text{ as } t \rightarrow \infty, \end{aligned}$$

by (A4). In the above, I denotes the $d_2 \times d_2$ identity matrix. The inequality above follows from the property on induced matrix norms (see Proposition A.12 of Bertsekas & Tsitsiklis, 1989). The claim follows. \square

As with Algorithm 1, we consider again the transformed problem with rewards replaced with costs (see above). This transformation, however, only affects the actor recursion (29) that now becomes

$$\theta_{t+1} = \Gamma(\theta_t - \beta_t G_t^{-1} \delta_t \psi_{s_t a_t}). \quad (50)$$

We have

Theorem 2. Under Assumptions (A1)–(A4), given $\epsilon > 0$, $\exists \delta > 0$, such that for θ_t , $t \geq 0$ obtained using Algorithm 2, if $\sup_{\pi_t} \|e^{\pi_t}\| < \delta$, then $\theta_t \rightarrow Z^\epsilon$ as $t \rightarrow \infty$, with probability one.

Proof. As with the proof of Theorem 1, let $\mathcal{F}_3(t) = \sigma(\theta_r, r \leq t)$, $t \geq 0$. Note that

$$\theta_{t+1} = \Gamma(\theta_t - \beta_t \mathbf{E}[G(\theta_t)^{-1} \delta_t^\pi \psi_{s_t a_t} | \mathcal{F}_3(t)] - \beta_t \gamma_3(t) + \beta_t \xi_1(t)),$$

where $\gamma_3(t) = G(\theta_t)^{-1} \delta_t \psi_{s_t a_t} - \mathbf{E}[G(\theta_t)^{-1} \delta_t \psi_{s_t a_t} | \mathcal{F}_3(t)]$. In lieu of Lemmas 5 and 6, $\xi_1(t) = o(1)$. As before, the critic recursion (28) converges faster for a given policy π_t , corresponding to an actor update θ_t , and converges to v^{π_t} . For $t \geq 1$, let $M^3(t) = \sum_{r=0}^{t-1} \gamma_3(r)$. The quantities δ_t and $G(\theta_t)^{-1}$ are uniformly bounded from Lemmas 5 and 6, and from (A4) respectively. Now using (20), it can be seen (Bhatnagar & Kumar, 2004) that $\{M^3(t)\}$ is a convergent martingale sequence. Hence, $\sum_{r=n}^{\infty} \beta_r \gamma_3(r) \rightarrow 0$ a.s. as $n \rightarrow \infty$, with n_T as before (see proof of Theorem 1). As before, also note that

$$\begin{aligned} \mathbf{E}[G(\theta_t)^{-1} \delta_t^\pi \psi_{s_t a_t} | \theta_t] &= G(\theta_t)^{-1} \sum_{s \in \mathcal{S}} d^{\pi_t}(s) \\ &\times \sum_{a \in \mathcal{A}} \nabla \pi_t(s, a) \left[R(s, a) - J(\pi_t) + \sum_{s' \in \mathcal{S}} P(s, a, s') v^{\pi_t \top} f_{s'} \right]. \end{aligned}$$

Consider now the ODE associated with (50) (in lieu of Lemma 4)

$$\dot{\theta} = \hat{\Gamma}(-G(\theta)^{-1}(\nabla J(\pi) + e^\pi)). \quad (51)$$

Consider also the associated ODE

$$\dot{\theta} = \hat{\Gamma}(-G(\theta)^{-1} \nabla J(\pi)). \quad (52)$$

As with Theorem 1, from the Hirsch lemma, $\theta_t \rightarrow \mathcal{Y}$ as $t \rightarrow \infty$ w.p. 1. Now as $\sup_{\pi_t} \|e^{\pi_t}\| \rightarrow 0$, the trajectories of (51) converge to those of (52) uniformly on compact sets for the same initial condition. See the proof of Theorem 3 of Bhatnagar et al. (2009) for details. The claim follows. \square

5.3. Convergence analysis for Algorithm 3

As stated previously, the main idea in this algorithm is to minimize the least-squares error in estimating the advantage function via function approximation. The analysis of the average reward (31), TD error (13), and critic (32) recursions proceeds in the same manner as before (cf. Lemma 5). We thus concentrate on recursion (33) and the actor recursion (34). We require Assumptions (A1)–(A3) here. In the transformed problem (with costs in place of rewards), recursion (33) can be rewritten as

$$w_{t+1} = (I - \alpha_t \psi_{s_t a_t} \psi_{s_t a_t}^\top) w_t - \alpha_t \delta_t \psi_{s_t a_t}, \quad (53)$$

with the actor recursion (34) the same as before. Note that (53) moves on a faster timescale as compared to the actor recursion. Hence, on the timescale of the former recursion, one may consider the parameter θ_t to be fixed. We have the following result:

Lemma 7. Under a given parameter θ , the w_t , $t \geq 1$, in (53) satisfy $w_t \rightarrow -G(\theta)^{-1} \mathbf{E}[\delta_t^\pi \psi_{s_t a_t}]$ as $t \rightarrow \infty$ with probability one, where π is the policy corresponding to θ .

Proof. Consider the following ODE associated with (53) for given θ

$$\dot{w} = \mathbf{E}_{s_t \sim d^\pi, a_t \sim \pi} [-\psi_{s_t a_t} \psi_{s_t a_t}^T w - \delta_t^\pi \psi_{s_t a_t}] \triangleq g^2(w). \quad (54)$$

Note that $g^2(w)$ is Lipschitz continuous in w . Now let $g_\infty^2(w) = \lim_{r \rightarrow \infty} \frac{g^2(rw)}{r}$. The function $g_\infty^2(w)$ exists and can be seen to satisfy $g_\infty^2(w) = -G(\theta)w$. For the ODE $\dot{w} = -G(\theta)w$, the origin is an asymptotically stable equilibrium (because $G(\theta)$ is positive definite). Define now $\{M^4(t)\}$ as $M^4(t) = \hat{g}^2(w_t) - \mathbf{E}[\hat{g}^2(w_t) | \mathcal{F}_4(t)]$, where $\hat{g}^2(w_t) = -\psi_{s_t a_t} \psi_{s_t a_t}^T w - \delta_t^\pi \psi_{s_t a_t}$ and $\mathcal{F}_4(t) = \sigma(w_r, M^4(r), r \leq t)$. It is easy to see that there exists a constant $C_0 < \infty$ such that

$$\mathbf{E}[\|M^4(t+1)\|^2 | \mathcal{F}_4(t)] \leq C_0(1 + \|w_t\|^2),$$

for all $t \geq 0$. For the ODE (54), it can be easily verified that $w^\pi = -G(\theta)^{-1} \mathbf{E}[\delta_t^\pi \psi_{s_t a_t}]$ is an asymptotically stable equilibrium (see Lemma 8 of Bhatnagar et al., 2009). Now from Thm. 2.2 of Borkar and Meyn (2000), (53) converges with probability one to w^π . \square

We now consider the actor recursion (34), which is the slower recursion.

Theorem 3. Under Assumptions A1–(A3), given $\epsilon > 0$, $\exists \delta > 0$ such that for θ_t , $t \geq 0$, obtained using Algorithm 3, if $\sup_{\pi_t} \|e^{\pi_t}\| < \delta$, then $\theta_t \rightarrow \mathcal{Z}^\epsilon$ as $t \rightarrow \infty$, with probability one.

Proof. Note that the recursion (34) can be written as

$$\theta_{t+1} = \Gamma(\theta_t - \beta_t G(\theta_t)^{-1} \mathbf{E}[\delta_t^{\pi_t} \psi_{s_t a_t} | \theta_t] + \beta_t \xi_2(t)),$$

where $\xi_2(t) = o(1)$ by Lemma 7. The rest can be shown in a similar manner as Theorem 2. \square

5.4. Convergence analysis for Algorithm 4

As with Algorithm 2, we require Assumptions (A1)–(A4). The result in Lemma 6 continues to hold here and we get for fixed θ , $G_t^{-1} \rightarrow G(\theta)^{-1}$ as $t \rightarrow \infty$ with probability one. The recursions for average reward (36), TD error (13), and critic (37) are the same as before and have been analyzed earlier (cf. Lemma 5). We now concentrate on recursion (38) and the actor recursion (39). Under the transformed problem (with costs in place of rewards), recursion (38) can be rewritten as

$$w_{t+1} = (1 - \alpha_t)w_t - \alpha_t G_t^{-1} \delta_t \psi_{s_t a_t}, \quad (55)$$

with the actor recursion the same as before. An exactly similar result as Lemma 7 holds in this case as well (below); see Lemma 9 of Bhatnagar et al. (2009) for a proof.

Lemma 8. Under a given parameter θ , the w_t , $t \geq 1$, defined by (55) satisfy $w_t \rightarrow -G(\theta)^{-1} \mathbf{E}[\delta_t^\pi \psi_{s_t a_t} | \theta]$ as $t \rightarrow \infty$ with probability one, with π being the policy corresponding to θ . \square

We finally consider the actor recursion (39) and have the following result whose proof follows as in Theorems 2 and 3.

Theorem 4. Under Assumptions (A1)–(A4), given $\epsilon > 0$, $\exists \delta > 0$ such that for θ_t , $t \geq 0$, obtained using Algorithm 4, if $\sup_{\pi_t} \|e^{\pi_t}\| < \delta$, then $\theta_t \rightarrow \mathcal{Z}^\epsilon$ as $t \rightarrow \infty$, with probability one. \square

6. Relation to the previous algorithms

As we mentioned in Section 1, the actor–critic algorithms presented in this paper extend prior actor–critic methods, especially those of Konda and Tsitsiklis (2003) and of Peters et al. (2003). In this section, we discuss these relationships further.

Actor–Critic Algorithm of Konda and Tsitsiklis (2003): Contrary to Algorithms 2–4, this algorithm does not use estimates of the natural gradient in its actor’s update. It is somewhat similar to our Algorithm 1, but with some key differences. (1) Konda’s algorithm uses the Markov process of state–action pairs and thus its critic update is based on an action–value function. Algorithm 1 uses the state process and therefore its critic update is based on a value function. (2) Whereas Algorithm 1 uses TD error in both critic and actor recursions, Konda’s algorithm uses TD error only in its critic update. The actor recursion in Konda’s algorithm uses a Q -value estimate instead. Because the TD error is an unbiased estimate of the advantage function (Lemma 3), the actor recursion in Algorithm 1 uses estimates of advantages instead of Q -values, which may result in lower variances. (3) The convergence analysis of Konda’s algorithm is based on the martingale approach and aims at bounding error terms and directly showing convergence. Convergence to a local optimum is shown when TD(1) critic is used. For the case when $\lambda < 1$, they show that given $\epsilon > 0$, there exists λ close enough to one such that when a TD(λ) critic is used, one gets $\liminf_t |\nabla J(\theta_t)| < \epsilon$ with probability one. Unlike Konda and Tsitsiklis, we primarily use the ordinary differential equation (ODE) approach for our convergence analysis. Even though we also use martingale arguments in our analysis, these are restricted to showing that the noise terms asymptotically diminish and the resulting scheme can be viewed as a Euler–discretization of the associated ODE.

Natural Actor–Critic Algorithm of Peters et al. (2003): Algorithms 2–4 extend this algorithm by being fully incremental and by providing convergence proofs. Peters’s algorithm uses a least-squares TD method in its critic’s update, whereas our algorithms are all fully incremental. It is not entirely clear how to satisfactorily incorporate least-squares TD methods in a context in which the policy is changing. Our proof techniques do not immediately extend to this case. However, we use estimates of the advantage function in Algorithms 3 and 4 as in Peters’s algorithm.

7. Conclusions and future work

We have introduced and analyzed four actor–critic reinforcement learning algorithms utilizing linear function approximation. All the algorithms are based on existing ideas such as temporal difference learning, natural policy gradients, and two-timescale convergence analysis, but we combine them in new ways. The main contribution of this paper is the proof of convergence of the four algorithms to a local maximum in the space of policy and value–function parameters. Our four algorithms are the first actor–critic algorithms to be shown convergent that utilize both function approximation and bootstrapping, a combination which seems essential to large-scale applications of reinforcement learning.

Our Algorithms 2–4 are explorations of the use of natural gradients within an actor–critic policy–gradient architecture. The way we use natural gradients is distinctive in that it is totally incremental: the policy is changed on every time step yet we never reset the gradient computation as is done in the algorithm of Peters and Schaal (2008). Algorithm 3 is perhaps the most interesting of the three natural-gradient algorithms. It never explicitly stores an estimate of the inverse of the Fisher information matrix, and as a result, it requires less computation. In our empirical experiments (Bhatnagar et al., 2009), we found it easier to find

good parameter settings for Algorithm 3 than for the other natural-gradient algorithms, and perhaps because of this, it converged more rapidly than them and than Konda's and Tsitsiklis' algorithm. These empirical observations should be taken only as suggestive; more experiments to properly assess the relative performance of these algorithms must be carried out.

The most important potential extension of our results would be to characterize the quality of the converged solution. It may be possible to bound the performance loss due to bootstrapping and approximation error in a way similar to how it was bounded by Tsitsiklis and Van Roy (1997). Because of the use of function approximation, our convergence analysis would carry through for the case of continuously valued state–action spaces as well. It would be interesting to study empirical evaluations of our algorithms in such settings in order to evaluate their applicability in such scenarios. There are a number of other ways in which our results are limited and suggest future work. (1) There is the issue of rate of convergence. Ideally one would like analytic results but, short of that, it would be useful to conduct a thorough empirical study, varying parameters and schedules in a more extensive and sophisticated way than what we have done in Bhatnagar et al. (2009). (2) The algorithms could be extended to incorporate eligibility traces and least-squares methods. As discussed earlier, the former seems straightforward whereas the latter seems to require more fundamental extensions. (3) A thorough study of the sensitivity of our algorithms to the various system parameters and settings is needed. (4) A study of the choice of the basis functions for the critic to obtain a good estimate of the policy gradient needs to be done. (5) Application of these ideas and algorithms to a real-world problem is needed to assess their ultimate utility.

References

- Abdulla, M. S., & Bhatnagar, S. (2007). Reinforcement learning based algorithms for average cost Markov decision processes. *Discrete Event Dynamic Systems: Theory and Applications*, 17(1), 23–52.
- Abounadi, J., Bertsekas, D., & Borkar, V. S. (2001). Learning algorithms for Markov decision processes. *SIAM Journal on Control and Optimization*, 40, 681–698.
- Aleksandrov, V., Sysoyev, V., & Shemeneva, V. (1968). Stochastic optimization. *Engineering Cybernetics*, 5, 11–16.
- Amari, S. (1998). Natural gradient works efficiently in learning. *Neural Computation*, 10(2), 251–276.
- Baird, L. (1993). Advantage updating. *Technical Report WL-TR-93-1146*, Wright Laboratory, OH.
- Baird, L. (1995). Residual algorithms: Reinforcement learning with function approximation. In *Proceedings of the twelfth international conference on machine learning* (pp. 30–37).
- Bagnell, J., & Schneider, J. (2003). Covariant policy search. In *Proceedings of international joint conference on artificial intelligence* (pp. 1019–1024).
- Barto, A., Sutton, R. S., & Anderson, C. (1983). Neuron-like elements that can solve difficult learning control problems. *IEEE Transactions on Systems, Man and Cybernetics*, 13, 835–846.
- Baxter, J., & Bartlett, P. (2001). Infinite-horizon policy-gradient estimation. *Journal of Artificial Intelligence Research*, 15, 319–350.
- Bellman, R. E., & Dreyfus, S. E. (1959). Functional approximations and dynamic programming. *Mathematical Tables and Other Aids to Computation*, 13, 247–251.
- Benveniste, A., Metivier, M., & Priouret, P. (1990). *Adaptive algorithms and stochastic approximations*. Berlin: Springer.
- Bertsekas, D. (1999). *Nonlinear programming*. Belmont, MA: Athena Scientific.
- Bertsekas, D., & Tsitsiklis, J. (1989). *Parallel and distributed computation*. New Jersey: Prentice Hall.
- Bertsekas, D., & Tsitsiklis, J. (1996). *Neuro-dynamic programming*. Belmont, MA: Athena Scientific.
- Bhatnagar, S., & Kumar, S. (2004). A simultaneous perturbation stochastic approximation based actor–critic algorithm for Markov decision processes. *IEEE Transactions on Automatic Control*, 49(4), 592–598.
- Bhatnagar, S. (2005). Adaptive multivariate three-timescale stochastic approximation algorithms for simulation based optimization. *ACM Transactions on Modeling and Computer Simulation*, 15(1), 74–107.
- Bhatnagar, S. (2007). Adaptive Newton-based multivariate smoothed functional algorithms for simulation optimization. *ACM Transactions on Modeling and Computer Simulation*, 18(1), 2:1–2:35.
- Bhatnagar, S., Sutton, R.S., Ghavamzadeh, M., & Lee, M. (2009). Natural actor-critic algorithms. Technical Report, *Department of Computing Science, University of Alberta, Canada* [<http://www.cs.ualberta.ca/research/techreports/2009/TR09-10.php>].
- Bhatnagar, S., Sutton, R. S., Ghavamzadeh, M., & Lee, M. (2008). Incremental natural actor-critic algorithms. *Advances in Neural Information Processing Systems*, 20, 105–112.
- Borkar, V. S. (1997). Stochastic approximation with two timescales. *Systems and Control Letters*, 29, 291–294.
- Borkar, V.S. (2008). Reinforcement learning – a bridge between numerical methods and Monte-Carlo. *Preprint*.
- Borkar, V. S., & Meyn, S. (2000). The O.D.E. method for convergence of stochastic approximation and reinforcement learning. *SIAM Journal on Control and Optimization*, 38(2), 447–469.
- Boyan, J. (1999). Least-squares temporal difference learning. In *Proceedings of the sixteenth international conference on machine learning* (pp. 49–56).
- Boyan, J., & Moore, A. (1995). Generalization in reinforcement learning: Safely approximating the value function. *Advances in Neural Information Processing Systems*, 7, 369–376.
- Brandiere, O. (1998). Some pathological traps for stochastic approximation. *SIAM Journal on Control and Optimization*, 36, 1293–1314.
- Bradtke, S., & Barto, A. (1996). Linear least-squares algorithms for temporal difference learning. *Machine Learning*, 22, 33–57.
- Cao, X., & Chen, H. (1997). Perturbation realization, potentials and sensitivity analysis of Markov processes. *IEEE Transactions on Automatic Control*, 42, 1382–1393.
- Crites, R., & Barto, A. (1998). Elevator group control using multiple reinforcement learning agents. *Machine Learning*, 33, 235–262.
- Farahmand, A. M., Ghavamzadeh, M., Szepesvári, Cs, & Mannor, S. (2009). Regularized policy iteration. *Advances in Neural Information Processing Systems*, 21, 441–448.
- Ghavamzadeh, M., & Mahadevan, S. (2003). Hierarchical policy gradient algorithms. In *Proceedings of the twentieth international conference on machine learning* (pp. 226–233).
- Ghavamzadeh, M., & Engel, Y. (2007a). Bayesian policy gradient algorithms. *Advances in Neural Information Processing Systems*, 19, 457–464.
- Ghavamzadeh, M., & Engel, Y. (2007b). Bayesian actor-critic algorithms. In *Proceedings of the twenty-fourth international conference on machine learning* (pp. 297–304).
- Glynn, P. (1990). Likelihood ratio gradient estimation for stochastic systems. *Communications of the ACM*, 33, 75–84.
- Gordon, G. (1995). Stable function approximation in dynamic programming. In *Proceedings of the twelfth international conference on machine learning* (pp. 261–268).
- Greensmith, E., Bartlett, P., & Baxter, J. (2004). Variance reduction techniques for gradient estimates in reinforcement learning. *Journal of Machine Learning Research*, 5, 1471–1530.
- Hirsch, M. (1989). Convergent activation dynamics in continuous time networks. *Neural Networks*, 2, 331–349.
- Kakade, S. (2002). A natural policy gradient. *Advances in Neural Information Processing Systems*, 14.
- Kohl, N., & Stone, P. (2004). Policy gradient reinforcement learning for fast quadrupedal locomotion. In *Proceedings of the IEEE international conference on robotics and automation* (pp. 2619–2624).
- Konda, V., & Borkar, V. S. (1999). Actor–critic like learning algorithms for Markov decision processes. *SIAM Journal on Control and Optimization*, 38(1), 94–123.
- Konda, V., & Tsitsiklis, J. (2003). On actor–critic algorithms. *SIAM Journal on Control and Optimization*, 42(4), 1143–1166.
- Kushner, H., & Clark, D. (1978). *Stochastic approximation methods for constrained and unconstrained systems*. New York: Springer Verlag.
- Kushner, H., & Yin, G. (1997). *Stochastic approximation algorithms and applications*. New York: Springer Verlag.
- Lagoudakis, M., & Parr, R. (2003). Least-squares policy iteration. *Journal of Machine Learning Research*, 4, 1107–1149.
- Marbach, P., & Tsitsiklis, J. (2001). Simulation-based optimization of Markov reward processes. *IEEE Transactions on Automatic Control*, 46, 191–209.
- Meyn, S. (2007). *Control techniques for complex networks*. Cambridge, UK: Cambridge Univ. Press.
- Ng, A., Coates, A., Diel, M., Ganapathi, V., Schulte, J., & Tse, B. et al. (2004). Inverted autonomous helicopter flight via reinforcement learning. In *International symposium on experimental robotics*.
- Pemantle, R. (1990). Nonconvergence to unstable points in urn models and stochastic approximations. *Annals of Probability*, 18, 698–712.
- Peters, J., Vijayakumar, S., & Schaal, S. (2003). Reinforcement learning for humanoid robotics. In *Proceedings of the third IEEE-RAS international conference on humanoid robots*.
- Peters, J., & Schaal, S. (2008). Natural actor-critic. *Neurocomputing*, 71(7–9), 1180–1190.
- Puterman, M. (1994). *Markov decision processes: Discrete stochastic dynamic programming*. New York: John Wiley.
- Richter, S., Aberdeen, D., & Yu, J. (2007). Natural actor-critic for road traffic optimization. *Advances in Neural Information Processing Systems*, 19, 1169–1176.
- Rust, J. (1996). Numerical dynamic programming in economics. In *Handbook of computational economics* (pp. 614–722). Amsterdam: Elsevier.
- Sutton, R.S. (1984). Temporal credit assignment in reinforcement learning. *Doctoral dissertation*, Amherst: University of Massachusetts.

- Sutton, R. S. (1988). Learning to predict by the method of temporal differences. *Machine Learning*, 3, 9–44.
- Sutton, R. S. (1996). Generalization in reinforcement learning: Successful examples using sparse coarse coding. *Advances in Neural Information Processing Systems*, 8, 1038–1044.
- Sutton, R. S., McAllester, D., Singh, S., & Mansour, Y. (2000). Policy gradient methods for reinforcement learning with function approximation. *Advances in Neural Information Processing Systems*, 12, 1057–1063.
- Sutton, R. S., & Barto, A. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.
- Tadic, V. (2001). On the convergence of temporal difference learning with linear function approximation. *Machine Learning*, 42(3), 241–267.
- Tesauro, G. (1995). Temporal difference learning and TD-Gammon. *Communications of the ACM*, 38, 58–68.
- Tsitsiklis, J. (1994). Asynchronous stochastic approximation and Q-learning. *Machine Learning*, 16, 185–202.
- Tsitsiklis, J., & Van Roy, B. (1997). An analysis of temporal-difference learning with function approximation. *IEEE Transactions on Automatic Control*, 42(5), 674–690.
- Tsitsiklis, J., & Van Roy, B. (1999). Average cost temporal-difference learning. *Automatica*, 35, 1799–1808.
- White, D. (1993). A survey of applications of Markov decision processes. *Journal of the Operational Research Society*, 44, 1073–1096.
- Widrow, B., & Stearns, S. (1985). *Adaptive signal processing*. Englewood Cliffs, NJ: Prentice-Hall.
- Williams, R. (1992). Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8, 229–256.



Shalabh Bhatnagar received his Ph.D. in Electrical Engineering from the Indian Institute of Science, Bangalore, in 1997. He was a postdoctoral research associate from 1997 to 2000 at the Institute for Systems Research, University of Maryland, College Park and from 2000 to 2001 at the Free University, Amsterdam, Netherlands. He is currently an Associate Professor in the Department of Computer Science and Automation, Indian Institute of Science, Bangalore. He has held visiting positions at the Indian Institute of Technology, Delhi and the RLAI research laboratory, University of Alberta, Canada.

His research interests are in the areas of stochastic control and simulation-based stochastic optimization with applications specifically in communication and wireless networks.



Richard S. Sutton is a professor and iCORE chair in the department of computing science at the University of Alberta. He is a fellow of the Association for the Advancement of Artificial Intelligence and co-author of the textbook *Reinforcement Learning: An Introduction*. Before joining the University of Alberta in 2003, he worked in industry at AT&T and GTE Labs, and in academia at the University of Massachusetts. He received a Ph.D. in computer science from the University of Massachusetts in 1984 and a BA in psychology from Stanford University in 1978. Rich's research interests center on the learning problems facing a decision-maker interacting with its environment, which he sees as central to artificial intelligence. He is also interested in animal learning psychology, in connectionist networks, and generally in systems that continually improve their representations and models of the world.



Mohammad Ghavamzadeh received a Ph.D. degree in Computer Science from the University of Massachusetts Amherst in 2005. From 2005 to 2008, he was a postdoctoral fellow at the Department of Computing Science at the University of Alberta, Canada. He has been a researcher at INRIA Lille - Nord Europe, Team SequeL in France since November 2008. The main objective of his research is to investigate the principles of scalable decision-making under uncertainty. In the last four years, Mohammad's research has been mostly focused on using recent advances in statistical machine learning, especially Bayesian reasoning and kernel methods, to develop more scalable reinforcement learning algorithms.



Mark Lee is a professional web developer and programmer, and the co-author of the book "C++ Programming for the Absolute Beginner". He received a B.Sc. degree in computer science from the University of Alberta in 2005.