

Reinforcement Learning

Olivier Buffet, Olivier Pietquin, and Paul Weng

Abstract Reinforcement learning (RL) is a general framework for adaptive control, which has proven to be efficient in many domains, e.g., board games, video games or autonomous vehicles. In such problems, an agent faces a sequential decision-making problem where, at every time step, it observes its state, performs an action, receives a reward and moves to a new state. An RL agent learns by trial and error a good policy (or controller) based on observations and numeric reward feedback on the previously performed action. In this chapter, we present the basic framework of RL and recall the two main families of approaches that have been developed to learn a good policy. The first one, which is value-based, consists in estimating the value of an optimal policy, value from which a policy can be recovered, while the other, called policy search, directly works in a policy space. Actor-critic methods can be seen as a policy search technique where the policy value that is learned guides the policy improvement. Besides, we give an overview of some extensions of the standard RL framework, notably when risk-averse behavior needs to be taken into account or when rewards are not available or not known.

Olivier Buffet
INRIA, Université de Lorraine, CNRS, UMR 7503 - LORIA, Nancy, France
e-mail: `olivier.buffet@loria.fr`

Olivier Pietquin
Univ. Lille, CNRS, Centrale Lille, Inria, UMR 9189 - CRIStAL, Lille; France, now with DeepMind, London, UK
e-mail: `olivier.pietquin@univ-lille1.fr`

Paul Weng
Shanghai Jiao Tong University, University of Michigan-Shanghai Jiao Tong University Joint Institute, Shanghai, China
e-mail: `paul.weng@sjtu.edu.cn`

1 Introduction

Reinforcement learning (RL) is a general framework for building autonomous agents (physical or virtual), which are systems that make decisions without human supervision in order to perform a given task. Examples of such systems abound: expert backgammon player [Tesauro, 1995], dialogue systems [Singh et al., 1999], acrobatic helicopter flight [Abbeel et al., 2010], human-level video game player [Mnih et al., 2015], go player [Silver et al., 2016] or autonomous driver [Bojarski et al., 2016]. See also Chapter 11 of Volume 2 and Chapters 10 and 12 of Volume 3.

In all those examples, an agent faces a sequential decision-making problem, which can be represented as an interaction loop between an agent and an environment. After observing its current situation, the agent selects an action to perform. As a result, the environment changes its state and provides a numeric reward feedback about the chosen action. In RL, the agent needs to learn how to choose good actions based on its observations and the reward feedback, without necessarily knowing the dynamics of the environment.

In this chapter, we focus on the basic setting of RL that assumes a single learning agent with full observability. Some work has investigated the partial observability case (see [Spaan, 2012] for an overview of both the model-based and model-free approaches). The basic setting has also been extended to situations where several agents interact and learn simultaneously (see [Busoniu et al., 2010] for a survey). RL has also been tackled with Bayesian inference techniques, which we do not mention here for space reasons (see [Ghavamzadeh et al., 2015] for a survey).

In Section 2, we recall the Markov decision process model on which RL is formulated and the RL framework, along with some of their classic solution algorithms. We present two families of approaches that can tackle large-sized problems for which function approximation is usually required. The first, which is value-based, is presented in Section 3. It consists in estimating the value function of an optimal policy. The second, called policy search, is presented in Section 4. It searches for an optimal policy directly in a policy space. In Section 5, we present some extensions of the standard RL setting, namely extensions to the case of unknown rewards and risk-sensitive RL approaches. Finally, we conclude in Section 6.

2 Background for RL

Before presenting the RL framework, we recall the Markov decision process (MDP) model, on which RL is based. See also Chapter 17 of this volume and Chapter 10 of Volume 2.

Markov decision process. MDPs and their multiple variants (e.g., Partially Observable MDP or POMDP) [Puterman, 1994] have been proposed to represent and solve sequential decision-making problems under uncertainty. An MDP is defined as a tuple $\mathcal{M} = \langle \mathbf{S}, \mathbf{A}, T, R, \gamma, H \rangle$ where \mathbf{S} is a set of states, \mathbf{A} is a set of actions, transition function $T(s, a, s')$ specifies the probability of reaching state s' after performing

action a in state s , reward function $R(s, a) \in \mathbb{R}$ yields the immediate reward after performing action a in state s , $\gamma \in [0, 1]$ is a discount factor and $H \in \mathbb{N} \cup \{\infty\}$ is the horizon of the problem, which is the number of decisions to be made. An immediate reward, which is a scalar number, measures the value of performing an action in a state. In some problems, it can be randomly generated. In that case, $R(s, a)$ is simply the expectation of the random rewards. In this MDP formulation, the environment is assumed to be stationary. Using such an MDP model, a system designer needs to define the tuple \mathcal{M} such that an optimal policy performs the task s/he wants.

Solving an MDP (i.e., *planning*) amounts to finding a controller, called a *policy*, which specifies which action to take in every state of the environment in order to maximize the expected discounted sum of rewards (standard decision criterion). A policy π can be deterministic (i.e., $\pi(s) \in A$) or randomized (i.e., $\pi(\cdot | s)$ is a probability distribution over \mathbf{A}). It can also be stationary or time-dependent, which is useful in finite-horizon or non-stationary problems.

A t -step history (also called trajectory, rollout or path) $h = (s_1, a_1, s_2, \dots, s_{t+1}) \in (\mathbf{S} \times \mathbf{A})^t \times \mathbf{S}$ is a sequence of past states and actions. In the standard case, it is valued by its return defined as $\sum_t \gamma^{t-1} R(s_t, a_t)$. As a policy induces a probability distribution over histories, the *value function* $v^\pi : \mathbf{S} \rightarrow \mathbb{R}$ of a policy π is defined by:

$$v_H^\pi(s) = \mathbb{E}_\pi \left[\sum_{t=1}^H \gamma^{t-1} R(S_t, A_t) \mid S_1 = s \right],$$

where \mathbb{E}_π is the expectation with respect to the distribution induced by π in the MDP, and S_t and A_t are random variables respectively representing a state and an action at a time step t . We will drop subscript H if there is no risk of confusion. The value function can be computed recursively. For deterministic policy π , we have:

$$\begin{aligned} v_0^\pi(s) &= 0, \\ v_t^\pi(s) &= R(s, \pi(s)) + \gamma \sum_{s' \in \mathbf{S}} T(s, \pi(s), s') v_{t-1}^\pi(s'). \end{aligned}$$

In a given state, policies can be compared via their value functions. Interestingly, in standard MDPs, there always exists an optimal deterministic policy whose value function is maximum in every state. Its value function is said to be optimal.

In the infinite horizon case, when $\gamma < 1$, v_t^π is guaranteed to converge to v^π , which is the solution of the *Bellman evaluation equations*:

$$v^\pi(s) = R(s, \pi(s)) + \gamma \sum_{s' \in \mathbf{S}} T(s, \pi(s), s') v^\pi(s'). \quad (1)$$

Given v^π , a better policy can be obtained with the following improvement step:

$$\pi'(s) = \operatorname{argmax}_{a \in \mathbf{A}} R(s, a) + \gamma \sum_{s' \in \mathbf{S}} T(s, a, s') v^\pi(s'). \quad (2)$$

The policy iteration algorithm consists in alternating between a policy evaluation step (1) and a policy improvement step (2), which converges to the optimal value function $v^* : \mathbf{S} \rightarrow \mathbb{R}$.

Alternatively, the optimal value function $v_H^* : \mathbf{S} \rightarrow \mathbb{R}$ can also be iteratively computed for any horizon H by:

$$\begin{aligned} v_0^*(s) &= 0 \\ v_t^*(s) &= \max_{a \in \mathbf{A}} R(s, a) + \gamma \sum_{s' \in \mathbf{S}} T(s, a, s') v_{t-1}^*(s'). \end{aligned} \quad (3)$$

In the infinite horizon case, when $\gamma < 1$, v_t^* is guaranteed to converge to v^* , which is the solution of the *Bellman optimality equations*:

$$v^*(s) = \max_{a \in \mathbf{A}} R(s, a) + \gamma \sum_{s' \in \mathbf{S}} T(s, a, s') v^*(s'). \quad (4)$$

In that case, (3) leads to the value iteration algorithm.

Two other related functions are useful when solving an RL problem: the action-value function $Q_t^\pi(s, a)$ (resp. the optimal action-value function $Q_t^*(s, a)$) specifies the value of choosing an action a in a state s at time step t and assuming policy π (resp. an optimal policy) is applied thereafter, i.e.,

$$Q_t^x(s, a) = R(s, a) + \gamma \sum_{s' \in \mathbf{S}} T(s, a, s') v_{t-1}^x(s') \quad \text{where } x \in \{\pi, *\}.$$

Reinforcement learning. In the MDP framework, a complete model of the environment is assumed to be known (via the transition function) and the task to be performed is completely described (via the reward function). The RL setting has been proposed to tackle situations when those assumptions do not hold. An RL agent searches for (i.e., during the *learning phase*) a best policy while interacting with the unknown environment by trial and error. In RL, the standard decision criterion used to compare policies is the same as in the MDP setting. Although the reward function is supposed to be unknown, the system designer has to specify it.

In RL, value and action-value functions have to be estimated. For v^π of a given policy π , this can be done with the standard TD(0) evaluation algorithm, where the following update is performed after applying π in state s yielding reward r and moving to new state s' :

$$v_t^\pi(s) = v_{t-1}^\pi(s) - \alpha_t(s) (v_{t-1}^\pi(s) - (r + v_{t-1}^\pi(s'))), \quad (5)$$

where $\alpha_t(s) \in [0, 1]$ is a learning rate. For Q^π , the update is as follows, after the agent executed action a in state s , received r , moved to new state s' and executed action a' (chosen by π):

$$Q_t^\pi(s, a) = Q_{t-1}^\pi(s, a) - \alpha_t(s, a) (Q_{t-1}^\pi(s, a) - (r + \gamma Q_{t-1}^\pi(s', a'))), \quad (6)$$

where $\alpha_t(s, a) \in [0, 1]$ is a learning rate. This update leads to the SARSA algorithm (named after the variables s, a, r, s', a'). In the same way that the policy iteration algorithm alternates between an evaluation step and a policy improvement step, one can use the SARSA evaluation method and combine it with a policy improvement step. In practice, we do not wait for the SARSA evaluation update rule to converge to the actual value of the current policy to make a policy improvement step. We rather continuously behave according to the current estimate of the Q -function to generate a new transition. One common choice is to use the current estimate in a softmax (Boltzmann) function of temperature τ and behave according to a randomized policy:

$$\pi_t(a|s) = \frac{e^{Q_{\theta_t}(s,a)/\tau}}{\sum_b e^{Q_{\theta_t}(s,b)/\tau}}.$$

Notice that we chose to use the Bellman evaluation equations to estimate the targets. However we could also use the Bellman optimality equations in the case of the Q -function and replace $r + \gamma Q(s', a')$ by $r + \max_b Q(s', b)$. Yet this only holds if we compute the value Q^* of the optimal policy π^* . This gives rise to the Q -learning update rule, which directly computes the value of the optimal policy. It is called an *off-policy* algorithm (whereas SARSA is *on-policy*) because it computes the value function of another policy than the one that selects the actions and generates the transitions used for the update. The following update is performed after the agent executed action a (e.g., chosen according to the softmax rule) in state s , received r and moved to new state s' :

$$Q_t^*(s, a) = Q_{t-1}^*(s, a) - \alpha_t(s, a)(Q_{t-1}^*(s, a) - (r + \gamma \max_{a'} Q_{t-1}^*(s', a'))). \quad (7)$$

Updates (5), (6) and (7) can be proved to converge if the learning rates satisfy standard stochastic approximation conditions (i.e., $\sum_t \alpha_t = \infty$ and $\sum_t \alpha_t^2 < \infty$). Besides, for (6), temperature τ would also need to converge to 0 while ensuring sufficient exploration in order for SARSA to converge to the optimal Q -function. In practice, $\alpha_t(s, a)$ is often chosen constant, which would also account for the case where the environment is non-stationary.

Those two general framework (MDP and RL) have been successfully applied in many different domains. For instance, MDPs or their variants have been used in finance [Bäuerle and Rieder, 2011] or logistics [Zhao et al., 2010]. RL has been applied to soccer [Bai et al., 2013] or power systems [Yu and Zhang, 2013], to cite a few. To tackle real-life large-sized problems, MDP and RL have to be completed with other techniques, such as compact representations [Boutilier et al., 2000; Guestrin et al., 2004; van Otterlo, 2009] or function approximation [de Farias and Van Roy, 2003; Geist and Pietquin, 2011; Mnih et al., 2015].

3 Value-Based Methods with Function Approximation

In many cases, the state-action space is too large so as to be able to represent exactly the value functions v^π or the action-value function Q^π of a policy π . For this reason, function approximation for RL has been studied for a long time, starting with the seminal work of Bellman and Dreyfus [1959]. In this framework, the functions are parameterized by a vector of d parameters $\theta = [\theta_j]_{j=1}^d$, with $\theta \in \Theta \subset \mathbb{R}^d$ (we will always consider column vectors) and the algorithms will aim at learning the parameters from data provided in the shape of transitions $\{s_t, a_t, s'_t, r_t\}_{t=1}^N$ where s'_t is the successor state of s_t drawn from $T(s_t, a_t, \cdot)$. We will denote the parameterized versions of the functions as v_θ and Q_θ . Popular approximation schemes are linear function approximation and neural networks. The former gave birth to a large literature in the theoretical domain as it allows studying convergence rates and bounds (although it remains non-trivial). The latter, although already used in the 90's [Tesauro, 1995], has known a recent growth in interest following the Deep Learning successes in supervised learning.

The case of neural networks will be addressed in Section 3.4 but we will start with linear function approximation. In this particular case, a set of basis functions $\phi(\cdot) = [\phi_j(\cdot)]_{j=1}^d$ has to be defined by the practitioner (or maybe learned through unsupervised learning) so that the value functions can be approximated by:

$$v_\theta(s) = \sum_j \theta_j \phi_j(s) = \theta^\top \phi(s) \quad \text{or} \quad Q_\theta(s, a) = \sum_j \theta_j \phi_j(s, a) = \theta^\top \phi(s, a).$$

The vector space defined by the span of ϕ is denoted Φ .

Notice that the exact case in which the different values of the value functions can be stored in a table (tabular case) is a particular case of linear function approximation. Indeed, if we consider that the state space is finite and small ($s = \{s_k\}_{k=1}^{|\mathbf{S}|} \in \mathbf{S}$), then the value function can be represented in a table of $|\mathbf{S}|$ values $\{v_k \mid v_k = v(s_k)\}_{k=1}^{|\mathbf{S}|}$ where $|\mathbf{S}|$ is the number of states. This is equivalent to defining a vector of $|\mathbf{S}|$ parameters $\mathbf{v} = [v_k]_{k=1}^{|\mathbf{S}|}$ and a vector of $|\mathbf{S}|$ basis functions $\boldsymbol{\delta}(s) = [\delta_k(s)]_{k=1}^{|\mathbf{S}|}$ where $\delta_k(s) = 1$ if $s = s_k$ and 0 otherwise. The value function can thus be written $v(s) = \sum_k v_k \delta_k(s) = \mathbf{v}^\top \boldsymbol{\delta}(s)$.

3.1 Stochastic Gradient Descent Methods

3.1.1 Bootstrapped Methods

If one wanted to cast the Reinforcement Learning problem into a supervised learning problem (see Chapter 11 of this Volume and Chapter 12 of Volume 2), one could want to fit the parameters to the value function directly. For instance, to evaluate the value of a particular policy π , one would solve the following regression problem

(for some ℓ_p -norm and distribution μ over states):

$$\boldsymbol{\theta}^* = \underset{\boldsymbol{\theta}}{\operatorname{argmin}} \|v_{\boldsymbol{\theta}}^{\pi} - v^{\pi}\|_{p,\mu} = \underset{\boldsymbol{\theta}}{\operatorname{argmin}} \|v_{\boldsymbol{\theta}}^{\pi} - v^{\pi}\|_{p,\mu}^p$$

where $\|\cdot\|_{p,\mu}$ denotes the weighted ℓ_p -norm defined by $(\mathbb{E}_{\mu}|\cdot|^p)^{1/p}$, \mathbb{E}_{μ} is the expectation with respect to μ . Yet, as said before, we usually cannot compute these values everywhere and we usually only have access to some transition samples $\{s_t, a_t, s'_t, r_t\}_{t=1}^N$ generated according to distribution μ . So we could imagine casting the RL problem into the following minimization problem:

$$\boldsymbol{\theta}^* = \underset{\boldsymbol{\theta}}{\operatorname{argmin}} \frac{1}{N} \sum_{t=1}^N |v_{\boldsymbol{\theta}}^{\pi}(s_t) - v^{\pi}(s_t)|^p.$$

This cost function can be minimized by stochastic gradient descent (we will consider an ℓ_2 -norm):

$$\begin{aligned} \boldsymbol{\theta}_t &= \boldsymbol{\theta}_{t-1} - \frac{\alpha}{2} \nabla_{\boldsymbol{\theta}_{t-1}} (v_{\boldsymbol{\theta}_{t-1}}^{\pi}(s_t) - v^{\pi}(s_t))^2 \\ &= \boldsymbol{\theta}_{t-1} - \alpha \nabla_{\boldsymbol{\theta}_{t-1}} v_{\boldsymbol{\theta}_{t-1}}^{\pi}(s_t) (v_{\boldsymbol{\theta}_{t-1}}^{\pi}(s_t) - v^{\pi}(s_t)). \end{aligned}$$

Of course, it is not possible to apply this update rule as it is since we do not know the actual value $v^{\pi}(s_t)$ of the states we observe in the transitions. But, from the Bellman evaluation equations (1), we can obtain an estimate by replacing $v^{\pi}(s_t)$ by $r_t + \gamma v_{\boldsymbol{\theta}_{t-1}}^{\pi}(s_{t+1})$. Notice that this replacement uses bootstrapping as we use the current estimate of the target to compute the gradient. We finally obtain the following update rule for evaluating the current policy π :

$$\boldsymbol{\theta}_t = \boldsymbol{\theta}_{t-1} - \alpha \nabla_{\boldsymbol{\theta}_{t-1}} v_{\boldsymbol{\theta}_{t-1}}^{\pi}(s_t) (v_{\boldsymbol{\theta}_{t-1}}^{\pi}(s_t) - (r_t + \gamma v_{\boldsymbol{\theta}_{t-1}}^{\pi}(s'_t))).$$

In the case of linear function approximation, i.e., $v_{\boldsymbol{\theta}}^{\pi}(s) = \boldsymbol{\theta}^{\top} \boldsymbol{\phi}(s)$, we obtain:

$$\boldsymbol{\theta}_t = \boldsymbol{\theta}_{t-1} - \alpha \boldsymbol{\phi}(s_t) (\boldsymbol{\theta}_{t-1}^{\top} \boldsymbol{\phi}(s_t) - (r_t + \gamma \boldsymbol{\theta}_{t-1}^{\top} \boldsymbol{\phi}(s'_t))).$$

Everything can be written again in the case of the action-value function, which leads to the SARSA update rule with linear function approximation $Q_{\boldsymbol{\theta}}^{\pi}(s, a) = \boldsymbol{\theta}^{\top} \boldsymbol{\phi}(s, a)$:

$$\boldsymbol{\theta}_t = \boldsymbol{\theta}_{t-1} - \alpha \boldsymbol{\phi}(s_t, a_t) (\boldsymbol{\theta}_{t-1}^{\top} \boldsymbol{\phi}(s_t, a_t) - (r_t + \gamma \boldsymbol{\theta}_{t-1}^{\top} \boldsymbol{\phi}(s'_t, a'_t))).$$

Changing the target as in the Q-learning update, we obtain for $Q_{\boldsymbol{\theta}}^*(s, a) = \boldsymbol{\theta}^{\top} \boldsymbol{\phi}(s, a)$:

$$\boldsymbol{\theta}_t = \boldsymbol{\theta}_{t-1} - \alpha \boldsymbol{\phi}(s_t, a_t) \left(\boldsymbol{\theta}_{t-1}^{\top} \boldsymbol{\phi}(s_t, a_t) - \left(r_t + \gamma \max_b \boldsymbol{\theta}_{t-1}^{\top} \boldsymbol{\phi}(s'_t, b) \right) \right).$$

3.1.2 Residual Methods

Instead of using the Bellman equations to provide an estimate of the target after deriving the update rule, one could use it directly to define the loss function to be optimized. We would then obtain the following minimization problem:

$$\boldsymbol{\theta}^* = \operatorname{argmin}_{\boldsymbol{\theta}} \frac{1}{N} \sum_{t=1}^N \left(v_{\boldsymbol{\theta}}^{\pi}(s_t) - (r_t + \gamma v_{\boldsymbol{\theta}}^{\pi}(s'_t)) \right)^2.$$

This can also be seen as the minimization of the Bellman residual. Indeed the Bellman evaluation equations ($v^{\pi}(s) = \mathbb{E}_{\pi}[R(s, A) + \gamma v^{\pi}(S')]$) can be rewritten as $v^{\pi}(s) - \mathbb{E}_{\pi}[R(s, A) + \gamma v^{\pi}(S')] = 0$. So by minimizing the quantity $v^{\pi}(s) - \mathbb{E}_{\pi}[R(s, A) + \gamma v^{\pi}(S')]$, called the Bellman residual, we reach the objective of evaluating $v^{\pi}(s)$. Here, we take the observed quantity $r + \gamma v^{\pi}(s')$ as an unbiased estimate of its expectation. The Bellman residual can also be minimized by stochastic gradient descent as proposed by Baird et al. [1995] and the update rule becomes:

$$\boldsymbol{\theta}_t = \boldsymbol{\theta}_{t-1} - \alpha \nabla_{\boldsymbol{\theta}_{t-1}} \left(v_{\boldsymbol{\theta}_{t-1}}^{\pi}(s_t) - (r_t + \gamma v_{\boldsymbol{\theta}_{t-1}}^{\pi}(s'_t)) \right) \left(v_{\boldsymbol{\theta}_{t-1}}^{\pi}(s_t) - (r_t + \gamma v_{\boldsymbol{\theta}_{t-1}}^{\pi}(s'_t)) \right).$$

In the case of a linear approximation, we obtain:

$$\boldsymbol{\theta}_t = \boldsymbol{\theta}_{t-1} - \alpha \left(\boldsymbol{\phi}(s_t) - \gamma \boldsymbol{\phi}(s'_t) \right) \left(\boldsymbol{\theta}_{t-1}^T \boldsymbol{\phi}(s_t) - (r_t + \gamma \boldsymbol{\theta}_{t-1}^T \boldsymbol{\phi}(s'_t)) \right).$$

This approach, called R-SGD (for residual stochastic gradient descent), has a major flaw as it computes a biased estimate of the value-function. Indeed, $v_{\boldsymbol{\theta}}^{\pi}(s_t)$ and $v_{\boldsymbol{\theta}}^{\pi}(s'_t)$ are correlated as s'_t is the result of having taken action a_t chosen by $\pi(s_t)$ [Werbos, 1990]. To address this problem, Baird et al. [1995] suggest to draw two different next states s'_t and s''_t starting from the same state s_t and to update as follows:

$$\boldsymbol{\theta}_t = \boldsymbol{\theta}_{t-1} - \alpha \nabla_{\boldsymbol{\theta}_{t-1}} \left(v_{\boldsymbol{\theta}_{t-1}}^{\pi}(s_t) - (r_t + \gamma v_{\boldsymbol{\theta}_{t-1}}^{\pi}(s'_t)) \right) \left(v_{\boldsymbol{\theta}_{t-1}}^{\pi}(s_t) - (r_t + \gamma v_{\boldsymbol{\theta}_{t-1}}^{\pi}(s''_t)) \right).$$

Of course, this requires that a generative model or a simulator is available and that transitions can be generated on demand.

The same discussions as in previous section can apply to learning an action-value function. For instance, one could want to solve the following optimization problem to learn the optimal action-value function:

$$\boldsymbol{\theta}^* = \operatorname{argmin}_{\boldsymbol{\theta}} \frac{1}{N} \sum_{t=1}^N \left(Q_{\boldsymbol{\theta}}^*(s_t, a_t) - (r_t + \gamma \max_b Q_{\boldsymbol{\theta}}^*(s'_t, b)) \right)^2. \quad (8)$$

Yet this optimal residual cannot directly be minimized in the case of the Q -function as the max operator is not differentiable. Notice that a sub-gradient method can still be used.

3.2 Least-Squares Methods

Gradient descent was used to minimize the empirical norm of either the bootstrapping error or the Bellman residual in the previous section. As the empirical norm is generally using the ℓ_2 -norm and that linear function approximation is often assumed, another approach could be to find the least squares solution to these problems. Indeed, least squares is a powerful approach as it is a second-order type of method and offers a closed-form solution to the optimization problem. Although there is no method that explicitly applies least squares to the two aforementioned empirical errors, one can see the fixed-point Kalman Filter (FPKF) algorithm [Choi and Van Roy, 2006] as a recursive least squares method applied to the bootstrapping error minimization. Also, the Gaussian Process Temporal Difference (GPTD) [Engel et al., 2005] or the Kalman Temporal Difference (KTD) [Geist and Pietquin, 2010a] algorithms can be seen as recursive least squares methods applied to Bellman residual minimization. We invite the reader to refer to Geist and Pietquin [2013] for further discussion on this.

Yet, the most popular method inspired by least squares optimization does apply to a different cost function. The Least-Squares Temporal Difference (LSTD) algorithm [Bradtke and Barto, 1996] aims at minimizing:

$$\boldsymbol{\theta}^* = \operatorname{argmin}_{\boldsymbol{\theta}} \frac{1}{N} \sum_{i=1}^N (v_{\boldsymbol{\theta}}^{\pi}(s_i) - v_{\boldsymbol{\omega}^*}^{\pi}(s_i))^2,$$

where $\boldsymbol{\omega}^* = \operatorname{argmin}_{\boldsymbol{\omega}} \frac{1}{N} \sum_{i=1}^N (v_{\boldsymbol{\omega}}^{\pi}(s_i) - (r_i + \gamma v_{\boldsymbol{\omega}}^{\pi}(s'_i)))^2$ can be understood as a projection on the space Φ spanned by the family of functions ϕ_j 's used to approximate v^{π} . It can be seen as the composition of the Bellman operator and of a projection operator. This cost function is the so-called *projected Bellman residual*. When using linear function approximation, this optimization problem admits a closed-form solution:

$$\boldsymbol{\theta}^* = \left[\sum_{i=1}^N \boldsymbol{\phi}(s_i) [\boldsymbol{\phi}(s_i) - \gamma \boldsymbol{\phi}(s'_i)]^{\top} \right]^{-1} \sum_{i=1}^N \boldsymbol{\phi}(s_i) r_i.$$

Note that the projected Bellman residual can also be optimized with a stochastic gradient approach [Sutton et al., 2009].

Extensions to non-linear function approximation exist and rely on the kernel trick [Xu et al., 2007] or on statistical linearization [Geist and Pietquin, 2010b]. LSTD can be used to learn an approximate Q -function as well and can be combined with policy improvement steps into an iterative algorithm, similar to policy iteration, to learn an optimal policy from a dataset of sampled transitions. This gives rise to the so-called Least Squares Policy Iteration (LSPI) algorithm [Lagoudakis and Parr, 2003], which is one of the most popular batch-RL algorithm.

3.3 Iterative Projected Fixed-Point Methods

As we have seen earlier, dynamic programming offers a set of algorithms to compute value functions of a policy in the case the dynamics of the MDP is known. One of these algorithms, Value Iteration, relies on the fact that the Bellman equations define contraction operators when $\gamma < 1$. For instance, if we define the Bellman evaluation operator B^π such that $B^\pi Q(s, a) = R(s, a) + \gamma \mathbb{E}_\pi [Q(S', A') | S = s, A = a]$, one can show that iteratively applying B^π to a random initialization of Q converges to Q^π , because B^π defines a contraction for which the only fixed point is Q^π [Puterman, 1994]. The Bellman optimality operator B^* , defined as $B^* Q(s, a) = R(s, a) + \gamma \mathbb{E} [\max_b Q(S', b) | S = s, A = a]$, is also a contraction. The same holds for the sampled versions of the Bellman operators. For instance, let us define the sampled evaluation operator \hat{B}^* such that $\hat{B}^* Q(s, a) = r + \gamma \max_b Q(s', b)$, where the expectation has been removed (the sampled operator applies to a single transition). Unfortunately, there is no guarantee that this remains a contraction when the value functions are approximated. Indeed when applying a Bellman operator to an approximate Q_θ , the result might not lie in the space spanned by θ . One has thus to project back on the space Φ spanned by ϕ using a projection operator Π_Φ , i.e., $\Pi_\Phi f = \operatorname{argmin}_\theta \|\theta^\top \phi - f\|_2$. If the composition of Π_Φ and \hat{B}^π (or \hat{B}^*) is still a contraction, then recursively applying this composition to any initialization of θ still converges to a good approximate Q_θ^π (or Q_θ^*). Unfortunately, the exact projection is often impossible to get as it is a regression problem. For instance, one would need to use least squares methods or stochastic gradient descent to learn the best projection from samples. Therefore the projection operator itself is approximated and will result in some $\hat{\Pi}_\Phi$ operator. So the iterative projected fixed-point process is defined as:

$$Q_{\theta_i} = \hat{\Pi}_\Phi \hat{B}^\pi Q_{\theta_{i-1}} \quad \text{or} \quad Q_{\theta_i} = \hat{\Pi}_\Phi \hat{B}^* Q_{\theta_{i-1}}.$$

In practice, the algorithm consists in collecting transitions (e.g., $\{s_i, a_i, r_i, s'_i\}_{i=1}^N$), initialize θ_0 to some random value, compute a regression database by applying the chosen sampled Bellman operator (e.g., $\{\hat{B}^* Q_{\theta_0}(s_i, a_i) = r_i + \gamma \max_b Q_{\theta_0}(s_i, b)\}_{i=1}^N$), apply a regression algorithm to find the next value of parameters (e.g., $Q_{\theta_1} = \hat{\Pi}_\Phi \hat{B}^* Q_{\theta_0} = \operatorname{argmin}_\theta \frac{1}{N} \sum_{i=1}^N (Q_\theta(s_i, a_i) - \hat{B}^* Q_{\theta_0}(s_i, a_i))^2$) and iterate.

This method finds its roots in early papers on dynamic programming [Samuel, 1959; Bellman et al., 1963] and convergence properties have been analyzed by Gordon [1995]. The most popular implementations use regression trees [Ernst et al., 2005] or neural networks [Riedmiller, 2005] as regression algorithms and have been applied to many concrete problems such as robotics [Antos et al., 2008].

3.4 Value-Based Deep Reinforcement Learning

Although the use of Artificial Neural Networks (ANN, see Chapter 12 of Volume 2) in RL is not new [Tesauro, 1995], there has been only a few successful attempts

to combine RL and ANN in the past. Most notably, before the recent advances in Deep Learning (DL) [LeCun et al., 2015], one can identify the work by Riedmiller [2005] as the biggest success of ANN as a function approximation framework for RL. There are many reasons for that, which are inherently due to the way ANN learns and assumptions that have to be made for both gradient descent and most value-based RL algorithms to converge. Especially, Deep ANNs (DNN) require a tremendous amount of data as they contain a lot of parameters to learn (typically hundreds of thousands to millions). To alleviate this issue, Tesauro [1995] trained his network to play backgammon through a self-play procedure. The model learned at iteration t plays again itself to generate data for training the model at iteration $t + 1$. It could thus reach super-human performance at the game of backgammon using RL. This very simple and powerful idea was reused in [Silver et al., 2016] to build the first artificial Go player that consistently defeated a human Go master. Yet, this method relies on the assumption that games can easily be generated on demand (backgammon and Go rules are simple enough even though the game is very complex). In more complex settings, the agent faces an environment for which it does not have access to the dynamics, maybe it cannot start in random states and has to follow trajectories, and it can only get transitions through actual interactions. This causes two major issues for learning with DNNs (in addition to intensive usage of data). First, gradient descent for training DNNs assume the data to be independent and identically distributed (i.i.d. assumption). Second, the distribution of the data should remain constant over time. Both these assumptions are normally violated by RL since transitions used to train the algorithms are part of trajectories (so next states are functions of previous states and actions, violating the i.i.d. assumption) and because trajectories are generated by a policy extracted from the current estimate of the value function (learning the value function influences the distribution of the data generated in the future). In addition, we also have seen in Section 3.1.2 that Bellman residual minimization suffers from the correlation between estimates of value functions of successive states. All these problems make RL unstable [Gordon, 1995].

To alleviate these issues, Mnih et al. [2015] used two tricks that allowed to reach super-human performances at playing Atari 2600 games from pixels. First, they made use of a biologically inspired mechanism, called experience replay [Lin, 1992], that consists in storing transitions in a Replay Buffer D before using them for learning. Instead of sequentially using these transitions, they are shuffled in the buffer and randomly sampled for training the network (which helps breaking correlation between successive samples). The buffer is filled on a first-in-first-out basis so that the distribution of the transitions is nearly stationary (transitions generated by old policies are discarded first). Second, the algorithm is based on asynchronous updates of the network used for generating the trajectories and a slow learning network. The slow learning network, called the target network, will be updated less often than the network that actually learns from the transitions stored in the replay buffer (the Q -network). This way, the update rule of the Q -network is built such that correlation between estimates of $Q(s, a)$ and $Q(s', a')$ is reduced. Indeed, the resulting algorithm (Deep Q-Network or DQN) is inspired by the gradient-descent up-

date on the optimal Bellman residual (8). But instead of using the double-sampling trick mentioned in Section 3.1.2, two different estimates of the Q -function are used. One according to the target network parameters (θ^-) and the other according to Q -network parameters (θ). The parameters of the Q -network are thus computed as:

$$\theta^* = \operatorname{argmin}_{\theta} \sum_{(s_t, a_t, s'_t, r_t) \in D} \left[\left(r_t + \gamma \max_b Q_{\theta^-}(s'_t, b) \right) - Q_{\theta}(s_t, a_t) \right]^2,$$

With this approach, the problem of non-differentiability of the max operator is also solved as the gradient is computed w.r.t. θ and not θ^- . Once in a while, the target network parameters are updated with the Q -network parameters ($\theta^- \leftarrow \theta^*$) and new trajectories are generated according to the policy extracted from Q_{θ^-} to fill again the replay buffer and train again the Q -network. The target network policy is actually a softmax policy based on Q_{θ^-} (see Section 3.1.1). Many improvements have been brought to that method since its publication, such as a prioritized replay mechanism [Schaul et al., 2016] that allows to sample more often from the replay buffer transitions for which the Bellman residual is larger, or the Double-DQN trick [Van Hasselt et al., 2016] used to provide more stable estimates of the max operator.

4 Policy-Search Approaches

Value-based approaches to RL rely on approximating the optimal value function V^* (typically using Bellman’s optimality principle), and then acting greedily with respect to this function. Policy Search algorithms directly optimize control policies, which typically depend on a parameter vector $\theta \in \Theta$ (and are thus noted π_{θ}), and whose general shape is user-defined.¹ Possible representations include linear policies, (deep) neural networks, radial basis function networks, and dynamic movement primitives (in robotics). Using such approaches avoids issues with discontinuous value functions, and makes it possible, in some cases, to deal with high-dimensional (possibly continuous) state and action spaces. They also allow providing expert knowledge through the shaping of the controller, or through example trajectories—to initialize the parameters.

In the following, we mainly distinguish between *model-free* and *model-based* algorithms—*i.e.*, depending on whether a model is being learned or not.

¹ This section is mainly inspired by [Deisenroth et al., 2011], although that survey focuses on a robotic framework.

4.1 Model-Free Policy Search

In model-free policy search, sampled trajectories are used directly to update the policy parameters. The discussion will follow the three main steps followed by the algorithms: (i) how they *explore* the space of policies, (ii) how they *evaluate* policies, and (iii) how policies are *updated*.

4.1.1 Policy Exploration

Exploring the space of policies implies either sampling the parameter vector the policy depends on, or perturbing the action choice of the policy. Often, the sampling of parameters takes place at the beginning of each episode (in episodic scenarios), and action perturbations are different at each time step, but other options are possible. Stochastic policies can be seen as naturally performing a step-based exploration in action space. Otherwise, the exploration strategy can be modeled as an *upper-level policy* $\pi_\omega(\theta)$ —sampling θ according to a probability distribution governed by parameter vector ω —, while the actual policy $\pi_\theta(a|s)$ is referred to as a *lower-level policy*. In this setting, the policy search aims at finding the parameter vector ω that maximizes the expected return given this vector. If $\pi_\omega(\theta)$ is a Gaussian distribution (common in robotics), then its covariance matrix can be diagonal—typically in step-based exploration—or not—which leads to more stability, but requires more samples—, meaning that the various parameters in θ can be treated in a correlated manner or not.

4.1.2 Policy Evaluation

Policy evaluation can also be step-based or episode-based. Step-based approaches evaluate each state-action pair. They have low variance and allow crediting several parameter vectors. They can rely on Q -value estimates, which can be biased and prone to approximation errors, or Monte-Carlo estimates, which can suffer from high variance. Episode-based approaches evaluate parameters using complete trajectories. They allow more performance criteria than step-based approaches—*e.g.*, minimizing the final distance to the target. They also allow for more sophisticated exploration strategies, but suffer all the more from noisy estimates and high variance that the dynamics are more stochastic.

4.1.3 Policy Update

Finally, the policy can be updated in rather different manners. We will discuss approaches relying on gradient ascents, inference-based optimization, information-theoretic ideas, stochastic optimization and path-integral optimal control.

Policy Gradient (PG) algorithms first require estimating the gradient. Some (episode-based) PG algorithms perform this estimate using a finite difference (FD) method by perturbing the parameter vector. Other algorithms instead exploit the *Likelihood ratio* trick, which allows estimating the gradient from a single trajectory, but requires a stochastic policy. These can be step-based as REINFORCE [Williams, 1992] or G(PO)MDP [Baxter and Bartlett, 2001; Baxter et al., 2001], or episode-based as PEPG [Sehnke et al., 2010].

Policy gradients also include natural gradient algorithms (NPG), i.e., algorithms that try to limit the distance between distributions $P_\theta(h)$ and $P_{\theta+\delta\theta}(h)$ using the KL divergence (estimated through the Fisher information matrix (FIM)). In step-based NPGs [Bagnell and Schneider, 2003; Peters and Schaal, 2008b], using appropriate (“*compatible*”) function approximation removes the need to estimate the FIM, but requires estimating the value function, which can be difficult. On the contrary, episodic Natural Actor-Critic (eNAC) [Peters and Schaal, 2008a] uses complete episodes, and thus only estimates $v(s_1)$. NAC [Peters and Schaal, 2008b] addresses infinite horizon problems, the lack of episodes leading to the use of Temporal Difference methods to estimate values.

Policy gradient usually applies to randomized policies. Recent work [Silver et al., 2014; Lillicrap et al., 2016] has adapted it to deterministic policies with a continuous action space, which can potentially facilitate the gradient estimation. Building on DQN, actor-critic methods have been extended to asynchronous updates with parallel actors and neural networks as approximators [Mnih et al., 2016].

Inference-based algorithms avoid the need to set learning rates. They consider that (i) the return R is an observed binary variable (1 meaning success),² (ii) the trajectory h is a latent variable, and (iii) one looks for the parameter vector that maximizes the probability of getting a return of 1. Then, an Expectation-Maximization algorithm can address this Bayesian inference problem. Variational inference can be used in the E-step of EM [Neumann, 2011], but most approaches rely on Monte-Carlo estimates instead, despite the fact that they perform maximum likelihood estimates over several modes of the reward function (and thus do not distinguish them). These can be episode-based algorithms as RWR [Peters and Schaal, 2007] (uses a linear upper-level policy) or CrKR [Kober et al., 2010] (a kernelized version of RWR, i.e., which does not need to specify feature vectors, but cannot model correlations). These can also be step-based algorithms as PoWER [Kober and Peters, 2010], which allows a more structured exploration strategy, and gives more influence to data points with less variance.

Information-theoretic approaches (see Chapter 2 of Volume 3) try to limit changes in trajectory distributions between two consecutive time steps, which could correspond to degradations rather than improvements in the policy. Natural PGs have the same objective, but need a user-defined learning rate. Instead, REPS [Peters et al., 2010] combines advantages from NPG (smooth learning) and EM-based algorithms (no learning-rate). Episode-based REPS [Daniel et al., 2012] learns a higher-level policy while bounding parameter changes by solving a constrained op-

² Transformations can bring us in this setting.

timization problem. Variants are able to adapt to multiple contexts or learn multiple solutions. Step-based REPS [Peters et al., 2010] solves an infinite horizon problem (rather than an episodic one), optimizing the average reward per time step. It requires enforcing the stationarity of state features, and thus solving another constrained optimization problem. A related recent method, TRPO [Schulman et al., 2015], which notably constrains the changes of $\pi(\cdot|s)$ instead of those of state-action distributions, proves to work well in practice.

Stochastic Optimization relies on black-box optimizers, and thus can easily be used for episode-based formulations, i.e., working with an upper-level policy $\pi_\omega(\theta)$. Typical examples are CEM [de Boer et al., 2005; Szita and Lörincz, 2006], CMA-ES [Hansen et al., 2003; Heidrich-Meisner and Igel, 2009], and NES [Wierstra et al., 2014], three evolutionary algorithms that maintain a parametric probability distribution (often Gaussian) $\pi_\omega(\theta)$ over the parameter vector. They sample a population of candidates, evaluate them, and use the best ones (weighted) to update the distribution. Many rollouts may be required for evaluation, as exemplified with the game of Tetris [Szita and Lörincz, 2006].

Path Integral (PI) approaches were introduced for optimal control, i.e., to handle non-linear continuous-time systems. They handle squared control costs and arbitrary state-dependent rewards. *Policy Improvement with PIs* (PI²) applies PI theory to optimize Dynamic Movement Primitives (DMPs), i.e., representations of movements with parameterized differential equations, using Monte-Carlo rollouts instead of dynamic programming.

4.2 Model-Based Policy Search

Typical model-based policy-search approaches repeatedly (i) sample real-world trajectories using a fixed policy; (ii) learn a forward model of the dynamics based on these samples (and previous ones); (iii) optimize this policy using the learned model (generally as a simulator). As can be noted, this process does not explicitly handle the exploration-exploitation trade-off as policies are not chosen so as to improve the model where this could be appropriate. We now discuss three important dimensions of these approaches: how to learn the model, how to make reliable long-term predictions, and how to perform the policy updates.

Model learning often uses probabilistic models. They first allow accounting for uncertainty due to sparse data (at least in some areas) or an inappropriate model class. In robotics, where action and state spaces are continuous, non-parametric probabilistic methods can be used such as Linearly Weighted Bayesian Regression (LWBR) or Gaussian Processes (GPs), which may suffer from increasing time and memory requirements. But probabilistic models can also be employed to represent stochastic dynamics. An example is that of propositional problems, which are often modeled as Factored MDPs [Boutilier et al., 1995], where the dynamics and rewards are DBNs whose structure is *a priori* unknown. A variety of approaches have been proposed, which rely on different representations (such as rule sets, decision trees,

Stochastic STRIPS, or PPDDL) [Degris et al., 2006; Pasula et al., 2007; Walsh et al., 2009; Lesner and Zanuttini, 2011]. See Chapter 10 of Volume 2.

Long-term predictions are usually required to optimize the policy given the current forward model. While the real world is its own best (unbiased) model, using a learned model has the benefit of allowing to control these predictions. A first approach, similar to paired statistical tests, is to always use the same random initial states and the same sequences of random numbers when evaluating different policies. It has been introduced for policy-search in the PEGASUS framework [Ng and Jordan, 2000] and drastically reduces the sampling variance. Another approach is, when feasible, to compute a probability distribution over trajectories using deterministic approximations such as linearization [Anderson and Moore, 2005], sigma-point methods (e.g., [Julier and Uhlmann, 2004]) or moment-matching.

Policy updates can rely on gradient-free optimization (e.g., Nelder-Mead method or hill-climbing) [Bagnell and Schneider, 2001], on sampling-based gradients (e.g., finite difference methods), as in model-free approaches, although they require many samples, or on analytical gradients [Deisenroth and Rasmussen, 2011], which require the model as well as the policy to be differentiable, scale favorably with the number of parameters, but are computationally involved.

5 Extensions: Unknown Rewards and Risk-sensitive Criteria

In the previous sections, we recalled different techniques for solving RL problems, with the assumption that policies are compared with the expected cumulated rewards as a decision criterion. However, rewards may not be scalar, known or numeric, and the standard criterion based on expectation may not always be suitable. For instance, multiobjective RL has been proposed to tackle situations where an action is evaluated over several dimensions (e.g., duration, length, power consumption for a navigation problem). The interested reader may refer to [Roijers et al., 2013] for a survey and refer to Chapter 16 of this volume for an introduction to multicriteria decision-making. For space reasons, we focus below only on three extensions: reward learning (Section 5.1), preference-based RL (Section 5.2) and risk sensitive RL (Section 5.3).

5.1 Reward Learning

From the system designer’s point of view, defining the reward function can be viewed as programming the desired behavior in an autonomous agent. A good choice of reward values may accelerate learning [Matignon et al., 2006] while an incorrect choice may lead to unexpected and unwanted behaviors [Randløv and Alstrøm, 1998]. Thus, designing this function is a hard task (e.g., robotics [Argall et al.,

2009], natural language parsers [Neu and Szepesvari, 2009] or dialogue systems [El Asri et al., 2012]).

When the reward signal is not known, a natural approach is to learn from demonstration. Indeed, in some domains (e.g., autonomous driving), it is much simpler for an expert to demonstrate how to perform a task rather than specify a reward function. Such an approach has been called apprenticeship learning [Abbeel and Ng, 2004], learning from demonstration [Argall et al., 2009], behavior cloning or imitation learning [Hussein et al., 2017]. Two families of techniques have been developed to solve such problems. The first group tries to directly learn a good policy from (near) optimal demonstrations [Argall et al., 2009; Pomerleau, 1989] while the second, called inverse RL (IRL) [Ng and Russell, 2000; Russell, 1998], tries to first recover a reward function that explains the demonstrations and then computes an optimal policy from it. The direct methods based on supervised learning usually suffer when the reward function is sparse and even more when dynamics is also perturbed [Piot et al., 2013].

As the reward function is generally considered to be a more compact, robust and transferable representation of a task than a policy [Abbeel and Ng, 2004; Russell, 1998], we only discuss reward learning approaches here.

As for many inverse problems, IRL is ill-posed: any constant function is a trivial solution that makes all policies equivalent and therefore optimal. Various solutions were proposed to tackle this degeneracy issue, differing on whether a probabilistic model is assumed or not on the generation of the observation. When the state and/or action spaces are large, the reward function is generally assumed to take a parametric form: $R(s, a) = f_{\theta}(s, a)$ for f_{θ} a parametric function of θ . One important case, called *linear features*, is when f is linear in θ , i.e., $R(s, a) = \sum_i \theta_i \phi_i(s, a)$ where ϕ_i are basis functions.

No generative model assumption. As underlined by Neu and Szepesvari [2009], many IRL methods can be viewed as finding the reward function R that minimizes a dissimilarity measure between the policy π_R^* optimal for R and the expert demonstrations. Most work assume a linear-feature reward function, with some exceptions that we mention below. Abbeel and Ng [2004] introduced the important idea of expected feature matching, which says that the expected features of π_R^* and those estimated from the demonstrations should be close. Thus, they notably proposed the projection method, which amounts to minimizing the Euclidean distance between those two expected features. Neu and Szepesvari [2007] proposed a natural gradient method for minimizing this objective function. Syed and Schapire [2008] reformulated the projection method problem as a zero-sum two-player game, with the nice property that the learned policy may perform better than the demonstrated one. Abbeel and Ng [2004]’s work was extended to the partially observable case [Choi and Kim, 2011].

Besides, Ratliff et al. [2006] proposed a max-margin approach enforcing that the found solution is better than any other one by at least a margin. Interestingly, the method can learn from multiple MDPs. It was later extended to the non-linear feature case [Ratliff et al., 2007].

Another technique [Klein et al., 2012; Piot et al., 2014] consists in learning a classifier based on a linearly parametrized score function to predict the best action for a state given the set of demonstrations. The learned score function can then be interpreted as a value function and can be used to recover a reward function.

Traditional IRL methods learn from (near) optimal demonstration. More recent approaches extend IRL to learn from other types of observations, e.g., a set of (non-necessarily optimal) demonstrations rated by an expert [El Asri et al., 2016; Burchfield et al., 2016], bad demonstrations [Sebag et al., 2016] or pairwise comparisons [da Silva et al., 2006; Wirth and Neumann, 2015]. In the latter case, the interactive setting is investigated with a reliable expert [Chernova and Veloso, 2009] or unreliable one [Weng et al., 2013].

Generative model assumption. Another way to tackle the degeneracy issue is to assume a probabilistic model on how observations are generated. Here, most work assumes that the expert policy is described by Boltzmann distributions, where higher-valued actions are more probable. Two notable exceptions are the work of Grollman and Billard [2011], which shows how to learn from failed demonstrations using Gaussian mixture models, and the Bayesian approach of Ramachandran and Amir [2007], with the assumption that state-action pairs in demonstrations follow such a Boltzmann distribution. This latter approach has been extended to Boltzmann distribution-based expert policy and for multi-task learning [Dimitrakakis and Rothkopf, 2011], and to account for multiple reward functions [Choi and Kim, 2012]. This Bayesian approach has been investigated to interactive settings where the agent can query for an optimal demonstration in a chosen state [Lopes et al., 2009] or for a pairwise comparison [Wilson et al., 2012; Akroun et al., 2013, 2014].

Without assuming a prior, Babes-Vroman et al. [2011] proposed to recover the expert reward function by maximum likelihood. The approach is able to handle the possibility of multiple intentions in the demonstrations. Furthermore, Nguyen et al. [2015] suggested an Expectation-Maximization approach to learn from demonstration induced by locally consistent reward functions.

To tackle the degeneracy issue, Ziebart et al. [2010] argued for the use of the maximum entropy principle, which states that among all solutions that fit the observations, the least informative one (i.e., maximum entropy) should be chosen, with the assumption that a reward function induces a Boltzmann probability distribution over trajectories. When the transition function is not known, Boularias et al. [2011] extended this approach by proposing to minimize the relative entropy between the probability distribution (over trajectories) induced by a policy and a baseline distribution under an expected feature matching constraint. Wulfmeier et al. [2015] extended this approach to the case where a deep neural network is used for the representation of the reward function, while Bogert et al. [2016] took into account non-observable variables.

5.2 Preference-Based Approaches

Another line of work redefines policy optimality directly based on pairwise comparisons of histories without assuming the existence of a scalar numeric reward function. This notably accounts for situations where reward values and probabilities are not commensurable. In this context, different decision criteria (e.g., quantile [Gilbert and Weng, 2016]) may be used. One popular decision model ([Yue et al., 2012; Fürnkranz et al., 2012]) is defined as follows: a policy π is preferred to another policy π' if

$$\mathbb{P}[h^\pi \succcurlyeq h^{\pi'}] \geq \mathbb{P}[h^{\pi'} \succcurlyeq h^\pi], \quad (9)$$

where \succcurlyeq is a preorder over histories, h^π is a random variable representing the history generated by policy π and therefore $\mathbb{P}[h^\pi \succcurlyeq h^{\pi'}]$ is the probability that a history generated by π is not less preferred than a history generated by π' . Based on (9), Fürnkranz et al. [2012] proposed a policy iteration algorithm. However, one crucial issue with (9) is that the concept of optimal solution is not well-defined as (9) can lead to preference cycles [Gilbert et al., 2015]. Busa-Fekete et al. [2014] circumvented this problem by refining this decision model with criteria from social choice theory. In [Gilbert et al., 2015], the issue was solved by considering mixed solutions: an optimal mixed solution is guaranteed to exist by interpreting it as a Nash equilibrium of a two-player zero-sum game. Gilbert et al. [2016] proposed a model-free RL algorithm based on a two-timescale technique to find such a mixed optimal solution.

5.3 Risk-Sensitive Criteria

Taking into account risk is important in decision-making under uncertainty (see Chapter 17 of this volume). The standard criterion based on expectation is risk-neutral. When it is known that a policy will only be used a few limited number of times, variability in the obtained rewards should be penalized. Besides, in some hazardous domains, good policies need to absolutely avoid bad or error states. In those two cases, preferences over policies need to be defined to be risk-sensitive.

In its simplest form, risk can directly be represented as a probability. For instance, Geibel and Wyszotzky [2005] adopted such an approach and consider MDP problems with two objectives where the first objective is the standard decision criterion and the second objective is to minimize the probability of reaching a set of bad states.

A more advanced approach is based on risk-sensitive decision criteria [Barbera et al., 1999]. Variants of Expected Utility [Machina, 1988], which is the standard risk-sensitive criterion, were investigated in two cases when the utility function is exponential [Borkar, 2010; Moldovan and Abbeel, 2012] and when it is quadratic [Tamar et al., 2012, 2013; Gosavi, 2014]. In the latter case, the criterion amounts to penalizing the standard criterion by the variance of the cumulated reward. While

the usual approach is to transform the cumulated reward, Mihatsch and Neuneier [2002] proposed to directly transform the temporal differences during learning.

Other approaches consider risk measures [Denuit et al., 2006] and in particular coherent risk measures [Artzner et al., 1999]. Value-at-risk, popular in finance, was considered in [Gilbert and Weng, 2016]. Policy gradient methods [Chow and Ghavamzadeh, 2014; Tamar et al., 2015b] were proposed to optimize Conditional Value-at-Risk (CVaR) and were extended to any coherent risk measure [Tamar et al., 2015a]. Jiang and Powell [2018] proposed dynamic quantile-based risk measures, which encompasses VaR and CVaR, and investigated an approximate dynamic programming scheme to optimize them.

In risk-constrained problems, the goal is to maximize the expectation of return while bounding a risk measure. For variance-constrained problems, Prashanth and Ghavamzadeh [2016] proposed an actor-critic algorithm. For CVaR-constrained problems, Borkar and Jain [2014] proposed a two-timescale stochastic approximation technique, while Chow et al. [2017] investigated policy gradient and actor-critic methods.

One important issue to consider when dealing with risk-sensitive criteria is that the Bellman optimality principle generally does not hold anymore: a sub-policy of an optimal risk-sensitive policy may not be optimal. However, in most cases, the Bellman optimality principle may be recovered by considering a state-augmented MDP, where the state includes the rewards cumulated so far [Liu and Koenig, 2006].

6 Conclusion

Recently, thanks to a number of success stories, reinforcement learning (RL) has become a very active research area. In this chapter, we recalled the basic setting of RL. Our focus was to present an overview of the main techniques, which can be divided into value-based and policy search methods, for solving large-sized RL problems with function approximation. We also presented some approaches for tackling the issue of unknown rewards that a system designer would encounter in practice and recalled some recent work in RL when risk-sensitivity needs to be taken into account in decision-making.

Currently RL still has too large sample and computational requirements for many practical domains (e.g., robotics). Research work is very active on improving RL algorithms along those two dimensions, notably by exploiting the structure of the problem [Kulkarni et al., 2016] or other a priori knowledge, expressed in temporal logic [Wen et al., 2017] for instance, or by reusing previous learning experience with transfer learning [Taylor and Stone, 2009], lifelong learning [Bou Ammar et al., 2015], multi-task learning [Wilson et al., 2007] or curriculum learning [Wu and Tian, 2017], to cite a few. Having more efficient RL algorithms is important as it will pave the way to more applications in more realistic domains.

References

- Abbeel, P., Coates, A., and Ng, A. Y. (2010). Autonomous helicopter aerobatics through apprenticeship learning. *International Journal of Robotics Research*, 29(13):1608–1639.
- Abbeel, P. and Ng, A. (2004). Apprenticeship learning via inverse reinforcement learning. In *International Conference Machine Learning*.
- Akrour, R., Schoenauer, M., and Sebag, M. (2013). Interactive robot education. In *ECML PKDD, Lecture Notes in Computer Science*.
- Akrour, R., Schoenauer, M., Souplet, J.-C., and Sebag, M. (2014). Programming by feedback. In *ICML*.
- Anderson, B. D. O. and Moore, J. B. (2005). *Optimal Filtering*. Dover Publications.
- Antos, A., Szepesvári, C., and Munos, R. (2008). Fitted q-iteration in continuous action-space mdps. In *Advances in neural information processing systems*, pages 9–16.
- Argall, B., Chernova, S., Veloso, M., and Browning, B. (2009). A survey of robot learning from demonstration. *Robotics and Autonomous Systems*, 57(5):469–483.
- Artzner, P., Delbaen, F., Eber, J., and Heath, D. (1999). Coherent measures of risk. *Mathematical Finance*, 9(3):203–228.
- Babes-Vroman, M., Marivate, V., Subramanian, K., and Littman, M. (2011). Apprenticeship learning about multiple intentions. In *ICML*.
- Bagnell, J. A. and Schneider, J. G. (2001). Autonomous helicopter control using reinforcement learning policy search methods. In *Proceedings of the International Conference on Robotics and Automation*, pages 1615–1620.
- Bagnell, J. A. and Schneider, J. G. (2003). Covariant policy search. In *Proceedings of the International Joint Conference on Artificial Intelligence*.
- Bai, A., Wu, F., and Chen, X. (2013). Towards a principled solution to simulated robot soccer. In *RoboCup-2012: Robot Soccer World Cup XVI, Lecture Notes in Artificial Intelligence*, volume 7500.
- Baird, L. et al. (1995). Residual algorithms: Reinforcement learning with function approximation. In *Proceedings of the twelfth international conference on machine learning*, pages 30–37.
- Barbera, S., Hammond, P., and Seidl, C. (1999). *Handbook of Utility Theory*. Springer.
- Bäuerle, N. and Rieder, U. (2011). *Markov Decision Processes with Applications to Finance*. Springer Science & Business Media.
- Baxter, J. and Bartlett, P. (2001). Infinite-horizon policy-gradient estimation. *Journal of Artificial Intelligence Research*, 15:319–350.
- Baxter, J., Bartlett, P., and Weaver, L. (2001). Experiments with infinite-horizon, policy-gradient estimation. *Journal of Artificial Intelligence Research*, 15:351–381.
- Bellman, R. and Dreyfus, S. (1959). Functional approximations and dynamic programming. *Mathematical Tables and Other Aids to Computation*, 13(68):247–251.

- Bellman, R., Kalaba, R., and Kotkin, B. (1963). Polynomial approximation—a new computational technique in dynamic programming: Allocation processes. *Mathematics of Computation*, 17(82):155–161.
- Bogert, K., Lin, J. F.-S., Doshi, P., and Kulic, D. (2016). Expectation-maximization for inverse reinforcement learning with hidden data. In *AAMAS*.
- Bojarski, M., Testa, D. D., Dworakowski, D., Firner, B., Flepp, B., Goyal, P., Jackel, L. D., Monfort, M., Muller, U., Zhang, J., Zhang, X., and Zhao, J. (2016). End to end learning for self-driving cars. Technical report, NVIDIA.
- Borkar, V. and Jain, R. (2014). Risk-constrained Markov decision processes. *IEEE Transactions on Automatic Control*, 59(9):2574–2579.
- Borkar, V. S. (2010). Learning algorithms for risk-sensitive control. In *International Symposium on Mathematical Theory of Networks and Systems*.
- Bou Ammar, H., Tutunov, R., and Eaton, E. (2015). Safe policy search for lifelong reinforcement learning with sublinear regret. In *ICML*.
- Boularias, A., Kober, J., and Peters, J. (2011). Relative entropy inverse reinforcement learning. In *AISTATS*.
- Boutillier, C., Dearden, R., and Goldszmidt, M. (1995). Exploiting structure in policy construction. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*, pages 1104–1111.
- Boutillier, C., Dearden, R., and Goldszmidt, M. (2000). Stochastic dynamic programming with factored representations. *Artificial Intelligence*, 121(1-2):49–107.
- Bradtke, S. J. and Barto, A. G. (1996). Linear least-squares algorithms for temporal difference learning. *Machine Learning*, 22:33–57.
- Burchfield, B., Tomasi, C., and Parr, R. (2016). Distance minimization for reward learning from scored trajectories. In *AAAI*.
- Busa-Fekete, R., Szörenyi, B., Weng, P., Cheng, W., and Hüllermeier, E. (2014). Preference-based Reinforcement Learning: Evolutionary Direct Policy Search using a Preference-based Racing Algorithm. *Machine Learning*, 97(3):327–351.
- Busoniu, L., Babuska, R., and De Schutter, B. (2010). *Innovations in Multi-Agent Systems and Applications – 1*, volume 310, chapter "Multi-agent reinforcement learning: An overview," Chapter 7, pages 183–221. Springer.
- Chernova, S. and Veloso, M. (2009). Interactive policy learning through confidence-based autonomy. *Journal of Artificial Intelligence Research*, 34:1–25.
- Choi, D. and Van Roy, B. (2006). A generalized kalman filter for fixed point approximation and efficient temporal-difference learning. *Discrete Event Dynamic Systems*, 16(2):207–239.
- Choi, J. and Kim, K.-E. (2011). Inverse reinforcement learning in partially observable environments. *JMLR*, 12:691–730.
- Choi, J. and Kim, K.-E. (2012). Nonparametric bayesian inverse reinforcement learning for multiple reward functions. In *NIPS*.
- Chow, Y. and Ghavamzadeh, M. (2014). Algorithms for cvar optimization in MDPs. In *NIPS*.
- Chow, Y., Ghavamzadeh, M., Janson, L., and Pavone, M. (2017). Risk-constrained reinforcement learning with percentile risk criteria. *JMLR*, 18(1).

- da Silva, V. F., Costa, A. H. R., and Lima, P. (2006). Inverse reinforcement learning with evaluation. In *IEEE ICRA*.
- Daniel, C., Neumann, G., and Peters, J. (2012). Hierarchical relative entropy policy search. In *Proceedings of the International Conference of Artificial Intelligence and Statistics*, pages 273–281.
- de Boer, P., Kroese, D., Mannor, S., and Rubinstein, R. (2005). A tutorial on the cross-entropy method. *Annals of Operations Research*, 134(1):19–67.
- de Farias, D. and Van Roy, B. (2003). The linear programming approach to approximate dynamic programming. *Operations Research*, 51(6):850–865.
- Degrís, T., Sigaud, O., and Wuillemin, P.-H. (2006). Learning the structure of Factored Markov Decision Processes in reinforcement learning problems. In *Proceedings of the 23rd International Conference on Machine Learning*.
- Deisenroth, M. P., Neumann, G., and Peters, J. (2011). A survey on policy search for robotics. *Foundations and Trends in Robotics*, 2(1–2):1–142.
- Deisenroth, M. P. and Rasmussen, C. E. (2011). PILCO: A model-based and data-efficient approach to policy search. In *Proceedings of the International Conference on Machine Learning*, pages 465–472.
- Denuit, M., Dhaene, J., Goovaerts, M., Kaas, R., and Laeven, R. (2006). Risk measurement with equivalent utility principles. *Statistics and Decisions*, 24:1–25.
- Dimitrakakis, C. and Rothkopf, C. A. (2011). Bayesian multitask inverse reinforcement learning. In *EWRL*.
- El Asri, L., Laroche, R., and Pietquin, O. (2012). Reward function learning for dialogue management. In *STAIRS*.
- El Asri, L., Piot, B., Geist, M., Laroche, R., and Pietquin, O. (2016). Score-based inverse reinforcement learning. In *AAMAS*.
- Engel, Y., Mannor, S., and Meir, R. (2005). Reinforcement learning with gaussian processes. In *Proceedings of the 22nd international conference on Machine learning*, pages 201–208. ACM.
- Ernst, D., Geurts, P., and Wehenkel, L. (2005). Tree-based batch mode reinforcement learning. *Journal of Machine Learning Research*, 6(Apr):503–556.
- Fürnkranz, J., Hüllermeier, E., Cheng, W., and Park, S. (2012). Preference-based reinforcement learning: A formal framework and a policy iteration algorithm. *Machine Learning*, 89(1):123–156.
- Geibel, P. and Wysotzky, F. (2005). Risk-sensitive reinforcement learning applied to control under constraints. *JAIR*, 24:81–108.
- Geist, M. and Pietquin, O. (2010a). Kalman temporal differences. *Journal of artificial intelligence research*, 39:483–532.
- Geist, M. and Pietquin, O. (2010b). Statistically linearized least-squares temporal differences. In *Ultra Modern Telecommunications and Control Systems and Workshops (ICUMT), 2010 International Congress on*, pages 450–457. IEEE.
- Geist, M. and Pietquin, O. (2011). Parametric value function approximation: a unified view. In *ADPRL*.
- Geist, M. and Pietquin, O. (2013). Algorithmic survey of parametric value function approximation. *IEEE Transactions on Neural Networks and Learning Systems*, 24(6):845–867.

- Ghavamzadeh, M., Mannor, S., Pineau, J., and Tamar, A. (2015). Bayesian reinforcement learning: a survey. *Foundations and Trends in Machine Learning*, 8(5–6):359–492.
- Gilbert, H., Spanjaard, O., Viappiani, P., and Weng, P. (2015). Solving MDPs with skew symmetric bilinear utility functions. In *International Joint Conference in Artificial Intelligence (IJCAI)*, pages 1989–1995.
- Gilbert, H. and Weng, P. (2016). Quantile reinforcement learning. In *Asian Workshop on Reinforcement Learning*.
- Gilbert, H., Zanuttini, B., Viappiani, P., Weng, P., and Nicart, E. (2016). Model-free reinforcement learning with skew-symmetric bilinear utilities. In *International Conference on Uncertainty in Artificial Intelligence (UAI)*.
- Gordon, G. J. (1995). Stable function approximation in dynamic programming. In *Proceedings of the twelfth international conference on machine learning*, pages 261–268.
- Gosavi, A. A. (2014). Variance-penalized markov decision processes: Dynamic programming and reinforcement learning techniques. *International Journal of General Systems*, 43(6):649–669.
- Grollman, D. H. and Billard, A. (2011). Donut as i do: learning from failed demonstrations. In *IEEE ICRA*.
- Guestrin, C., Hauskrecht, M., and Kveton, B. (2004). Solving factored MDPs with continuous and discrete variables. In *AAAI*, pages 235–242.
- Hansen, N., Muller, S., and Koumoutsakos, P. (2003). Reducing the time complexity of the derandomized evolution strategy with covariance matrix adaptation (CMA-ES). *Evolutionary Computation*, 11(1):1–18.
- Heidrich-Meisner, V. and Igel, C. (2009). Neuroevolution strategies for episodic reinforcement learning. *Journal of Algorithms*, 64(4):152–168.
- Hussein, A., Gaber, M. M., Elyan, E., and Jayne, C. (2017). Imitation learning: a survey of learning methods. *ACM Computing Surveys*.
- Jiang, D. R. and Powell, W. B. (2018). Risk-averse approximate dynamic programming with quantile-based risk measures. *Mathematics of Operations Research*, 43(2):347–692.
- Julier, S. J. and Uhlmann, J. K. (2004). Unscented filtering and nonlinear estimation. *Proceedings of the IEEE*, 92(3):401–422.
- Klein, E., Geist, M., Piot, B., and Pietquin, O. (2012). Inverse reinforcement learning through structured classification. In *NIPS*.
- Kober, J., Oztop, E., and Peters, J. (2010). Reinforcement learning to adjust robot movements to new situations. In *Proceedings of the 2010 Robotics: Science and Systems Conference*.
- Kober, J. and Peters, J. (2010). Policy search for motor primitives in robotics. *Machine Learning*, pages 1–33.
- Kulkarni, T., Narasimhan, K. R., Saeedi, A., and Tenenbaum, J. (2016). Hierarchical deep reinforcement learning: Integrating temporal abstraction and intrinsic motivation. In *NIPS*.
- Lagoudakis, M. G. and Parr, R. (2003). Least-squares policy iteration. *Journal of machine learning research*, 4(Dec):1107–1149.

- LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. *Nature*, 521(7553):436–444.
- Lesner, B. and Zanuttini, B. (2011). Handling ambiguous effects in action learning. In *Proceedings of the 9th European Workshop on Reinforcement Learning*, page 12.
- Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., and Wierstra, D. (2016). Continuous control with deep reinforcement learning. In *ICLR*.
- Lin, L.-H. (1992). Self-improving reactive agents based on reinforcement learning, planning and teaching. *Machine learning*, 8(3/4):69–97.
- Liu, Y. and Koenig, S. (2006). Functional value iteration for decision-theoretic planning with general utility functions. In *AAAI*, pages 1186–1193. AAAI.
- Lopes, M., Melo, F., and Montesano, L. (2009). Active learning for reward estimation in inverse reinforcement learning. In *ECML/PKDD, Lecture Notes in Computer Science*, volume 5782, pages 31–46.
- Machina, M. (1988). Expected utility hypothesis. In Eatwell, J., Milgate, M., and Newman, P., editors, *The New Palgrave: A Dictionary of Economics*, pages 232–239. Macmillan.
- Matignon, L., Laurent, G. J., and Le Fort-Piat, N. (2006). Reward function and initial values: Better choices for accelerated goal-directed reinforcement learning. *Lecture notes in CS*, 1(4131):840–849.
- Mihatsch, O. and Neuneier, R. (2002). Risk-sensitive reinforcement learning. *Machine Learning*.
- Mnih, V., Badia, A. P., Mirza, M., Graves, A., Lillicrap, T. P., Harley, T., Silver, D., and Kavukcuoglu, K. (2016). Asynchronous methods for deep reinforcement learning. In *ICML*.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S., and Hassabis, D. (2015). Human-level control through deep reinforcement learning. *Nature*, 518:529–533.
- Moldovan, T. and Abbeel, P. (2012). Risk aversion Markov decision processes via near-optimal Chernoff bounds. In *NIPS*.
- Neu, G. and Szepesvari, C. (2007). Apprenticeship learning using inverse reinforcement learning and gradient methods. In *UAI*.
- Neu, G. and Szepesvari, C. (2009). Training parsers by inverse reinforcement learning. *Machine Learning*, 77:303–337.
- Neumann, G. (2011). Variational inference for policy search in changing situations. In *Proceedings of the International Conference on Machine Learning*, pages 817–824.
- Ng, A. and Russell, S. (2000). Algorithms for inverse reinforcement learning. In *ICML*. Morgan Kaufmann.
- Ng, A. Y. and Jordan, M. (2000). PEGASUS : A policy search method for large MDPs and POMDPs. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence*.

- Nguyen, Q. P., Low, K. H., and Jaillet, P. (2015). Inverse reinforcement learning with locally consistent reward functions. In *NIPS*.
- Pasula, H. M., Zettlemoyer, L. S., and Kaelbling, L. P. (2007). Learning symbolic models of stochastic domains. *Journal of Artificial Intelligence Research*, 29:309–352.
- Peters, J., Mülling, K., and Altun, Y. (2010). Relative entropy policy search. In *Proceedings of the National Conference on Artificial Intelligence*.
- Peters, J. and Schaal, S. (2007). Applying the episodic natural actor-critic architecture to motor primitive learning. In *Proceedings of the European Symposium on Artificial Neural Networks*.
- Peters, J. and Schaal, S. (2008a). Natural actor-critic. *Neurocomputation*, 71(7–9):1180–1190.
- Peters, J. and Schaal, S. (2008b). Reinforcement learning of motor skills with policy gradients. *Neural Networks*, 4:682–697.
- Piot, B., Geist, M., and Pietquin, O. (2013). Learning from demonstrations: Is it worth estimating a reward function? In *ECML PKDD, Lecture Notes in Computer Science*.
- Piot, B., Geist, M., and Pietquin, O. (2014). Boosted and Reward-regularized Classification for Apprenticeship Learning. In *AAMAS*, pages 1249–1256, Paris, France.
- Pomerleau, D. (1989). Alvin: An autonomous land vehicle in a neural network. In *NIPS*.
- Prashanth, L. and Ghavamzadeh, M. (2016). Variance-constrained actor-critic algorithms for discounted and average reward mdps. *Machine Learning*.
- Puterman, M. (1994). *Markov decision processes: discrete stochastic dynamic programming*. Wiley.
- Ramachandran, D. and Amir, E. (2007). Bayesian inverse reinforcement learning. In *IJCAI*.
- Randløv, J. and Alstrøm, P. (1998). Learning to drive a bicycle using reinforcement learning and shaping. In *ICML, (1998)*.
- Ratliff, N., Bagnell, J., and Zinkevich, M. (2006). Maximum margin planning. In *ICML*.
- Ratliff, N., Bradley, D., Bagnell, J. A., and Chestnutt, J. (2007). Boosting structured prediction for imitation learning. In *NIPS*.
- Riedmiller, M. (2005). Neural fitted q iteration—first experiences with a data efficient neural reinforcement learning method. In *ECML*, volume 3720, pages 317–328. Springer.
- Roijers, D., Vamplew, P., Whiteson, S., and Dazeley, R. (2013). A survey of multi-objective sequential decision-making. *Journal of Artificial Intelligence Research*, 48:67–113.
- Russell, S. (1998). Learning agents for uncertain environments. In *Proceedings of the eleventh annual conference on Computational learning theory*, pages 101–103. ACM.
- Samuel, A. (1959). Some studies in machine learning using the game of checkers. *IBM Journal of Research and Development*, 3(3):210–229.

- Schaul, T., Quan, J., Antonoglou, I., and Silver, D. (2016). Prioritized experience replay. In *ICLR*.
- Schulman, J., Levine, S., Abbeel, P., Jordan, M., and Moritz, P. (2015). Trust region policy optimization. In *ICML*.
- Sebag, M., Akrou, R., Mayeur, B., and Schoenauer, M. (2016). Anti imitation-based policy learning. In *ECML PKDD, Lecture Notes in Computer Science*.
- Sehnke, F., Osendorfer, C., Rückstieß, T., Graves, A., Peters, J., and Schmidhuber, J. (2010). Parameter-exploring policy gradients. *Neural Networks*, 23(4):551–559.
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., van den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., Dieleman, S., Grewe, D., Nham, J., Kalchbrenner, N., Sutskever, I., Lillicrap, T., Leach, M., Kavukcuoglu, K., Graepel, T., and Hassabis, D. (2016). Mastering the game of Go with deep neural networks and tree search. *Nature*.
- Silver, D., Lever, G., Heess, N., Degris, T., Wierstra, D., and Riedmiller, M. (2014). Deterministic policy gradient algorithms. In *ICML*.
- Singh, S., Kearns, M., Litman, D., and Walker, M. (1999). Reinforcement learning for spoken dialogue systems. In *NIPS*.
- Spaan, M. T. (2012). *Reinforcement Learning*, chapter Partially Observable Markov Decision Processes. Springer.
- Sutton, R., Maei, H., Precup, D., Bhatnagar, S., Silver, D., Szepesvári, C., and Wiewiora, E. (2009). Fast gradient-descent methods for temporal-difference learning with linear function approximation. In *ICML*.
- Syed, U. and Schapire, R. E. (2008). A game-theoretic approach to apprenticeship learning. In *NIPS*.
- Szita, I. and Lörincz, A. (2006). Learning tetris using the noisy cross-entropy method. *Neural Computation*, 18:2936–2941.
- Tamar, A., Chow, Y., Ghavamzadeh, M., and Mannor, S. (2015a). Policy gradient for coherent risk measures. In *NIPS*.
- Tamar, A., Di Castro, D., and Mannor, S. (2012). Policy gradient with variance related risk criteria. In *ICML*.
- Tamar, A., Di Castro, D., and Mannor, S. (2013). Temporal difference methods for the variance of the reward to go. In *ICML*.
- Tamar, A., Glassner, Y., and Mannor, S. (2015b). Optimizing the CVaR via sampling. In *AAAI*.
- Taylor, M. E. and Stone, P. (2009). Transfer learning for reinforcement learning domains: A survey. *Journal of Machine Learning Research*, 10:1633–1685.
- Tesauro, G. (1995). Temporal difference learning and td-gammon. *Communications of the ACM*, 38(3):58–68.
- Van Hasselt, H., Guez, A., and Silver, D. (2016). Deep reinforcement learning with double q-learning. In *AAAI*, pages 2094–2100.
- van Otterlo, M. (2009). *The Logic of Adaptive Behavior*. IOS Press.
- Walsh, T., Szita, I., Diuk, C., and Littman, M. (2009). Exploring compact reinforcement-learning representations with linear regression. In *Proceedings of the 25th Conference on Uncertainty in Artificial Intelligence*.

- Wen, M., Papusha, I., and Topcu, U. (2017). Learning from demonstrations with high-level side information. In *IJCAI*.
- Weng, P., Busa-Fekete, R., and Hüllermeier, E. (2013). Interactive Q-learning with ordinal rewards and unreliable tutor. In *Workshop Reinforcement Learning with Generalized Feedback, ECML/PKDD*.
- Werbos, P. J. (1990). Consistency of hdp applied to a simple reinforcement learning problem. *Neural Networks*, 3:179–189.
- Wierstra, D., Schaul, T., Glasmachers, T., Sun, Y., Peters, J., and Schmidhuber, J. (2014). Natural evolution strategies. *JMLR*, 15:949–980.
- Williams, R. (1992). Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8(3):229–256.
- Wilson, A., Fern, A., Ray, S., and Tadepalli, P. (2007). Multi-task reinforcement learning: A hierarchical bayesian approach. In *ICML*.
- Wilson, A., Fern, A., and Tadepalli, P. (2012). A Bayesian approach for policy learning from trajectory preference queries. In *Advances in Neural Information Processing Systems*.
- Wirth, C. and Neumann, G. (2015). Model-free preference-based reinforcement learning. In *EWRL*.
- Wu, Y. and Tian, Y. (2017). Training agent for first-person shooter game with actor-critic curriculum learning. In *ICLR*.
- Wulfmeier, M., Ondruska, P., and Posner, I. (2015). Maximum entropy deep inverse reinforcement learning. In *NIPS, Deep Reinforcement Learning Workshop*.
- Xu, X., Hu, D., and Lu, X. (2007). Kernel-based least squares policy iteration for reinforcement learning. *IEEE Transactions on Neural Networks*, 18(4):973–992.
- Yu, T. and Zhang, Z. (2013). Optimal CPS control for interconnected power systems based on SARSA on-policy learning algorithm. *Power System Protection and Control*, pages 211–216.
- Yue, Y., Broder, J., Kleinberg, R., and Joachims, T. (2012). The k-armed dueling bandits problem. *Journal of Computer and System Sciences*, 78(5):1538–1556.
- Zhao, Q., Chen, S., Leung, S., and Lai, K. (2010). Integration of inventory and transportation decisions in a logistics system. *Transportation Research Part E: Logistics and Transportation Review*, 46(6):913–925.
- Ziebart, B., Maas, A., Bagnell, J., and Dey, A. (2010). Maximum entropy inverse reinforcement learning. In *AAAI*.

Contents

Reinforcement Learning	1
Olivier Buffet, Olivier Pietquin, and Paul Weng	
1 Introduction	2
2 Background for RL	2
3 Value-Based Methods with Function Approximation	6
3.1 Stochastic Gradient Descent Methods	6
3.2 Least-Squares Methods	9
3.3 Iterative Projected Fixed-Point Methods	10
3.4 Value-Based Deep Reinforcement Learning	10
4 Policy-Search Approaches	12
4.1 Model-Free Policy Search	13
4.2 Model-Based Policy Search	15
5 Extensions: Unknown Rewards and Risk-sensitive Criteria	16
5.1 Reward Learning	16
5.2 Preference-Based Approaches	19
5.3 Risk-Sensitive Criteria	19
6 Conclusion	20
References	21