
Soft-Robust Actor-Critic Policy-Gradient

Esther Derman

Technion, Israel

estherderman@technion.ac.il

Timothy A. Mann

Google Deepmind, UK

timothymann@google.com

Daniel J. Mankowitz

Technion, Israel

danielm@campus.technion.ac.il

Shie Mannor

Technion, Israel

shie@ee.technion.ac.il

Abstract

Robust Reinforcement Learning aims to derive an optimal behavior that accounts for model uncertainty in dynamical systems. However, previous studies have shown that by considering the worst case scenario, robust policies can be overly conservative. Our *soft-robust* framework is an attempt to overcome this issue. In this paper, we present a novel Soft-Robust Actor-Critic algorithm (SR-AC). It learns an optimal policy with respect to a distribution over an uncertainty set and stays robust to model uncertainty but avoids the conservativeness of robust strategies. We show the convergence of SR-AC and test the efficiency of our approach on different domains by comparing it against regular learning methods and their robust formulations.

1 INTRODUCTION

Markov Decision Processes (MDPs) are commonly used to model sequential decision making in stochastic environments. A strategy that maximizes the accumulated expected reward is then considered as optimal and can be learned from sampling. However, besides the uncertainty that results from stochasticity of the environment, model parameters are often estimated from noisy data or can change during testing [Mannor et al., 2007; Roy et al., 2017]. This second type of uncertainty can significantly degrade the performance of the optimal strategy from the model’s prediction.

Robust MDPs were proposed to address this problem [Iyengar, 2005; Nilim and El Ghaoui, 2005; Tamar et al., 2014]. In this framework, a transition model is assumed to belong to a known uncertainty set and an optimal strategy is learned under the worst parameter realizations. Although the robust approach is computationally efficient

when the uncertainty set is state-wise independent, compact and convex, it can lead to overly conservative results [Mannor et al., 2012, 2016; Xu and Mannor, 2012; Yu and Xu, 2016].

For example, consider a business scenario where an agent’s goal is to make as much money as possible. It can either create a startup which may make a fortune but may also result in bankruptcy. Alternatively, it can choose to live off school teaching and have almost no risk but low reward. By choosing the teaching strategy, the agent may be overly conservative and not account for opportunities to invest in his own promising projects. Our claim is that one could relax this conservativeness and construct a softer behavior that interpolates between being aggressive and robust. Ideally, the *soft-robust* agent should stay agnostic to outside financing uncertainty but still be able to take advantage of the startup experience.

This type of dilemma can be found in various domains. In the financial market, investors seek a good trade-off between low risk and high returns regarding portfolio management [Mitchell and Smetters, 2013]. In strategic management, product firms must choose the amount of resources they put into innovation. A conservative strategy would then consist of innovating only under necessary conditions [Miller and Friesen, 1982].

In this paper, we focus on learning a *soft-robust policy* (defined below) by incorporating soft-robustness into an online actor-critic algorithm and show its convergence properties. Existing works mitigate conservativeness of robust MDP either by introducing coupled uncertainties [Mannor et al., 2012, 2016] or by assuming prior information on the uncertainty set [Xu and Mannor, 2012; Yu and Xu, 2016]. They use dynamic programming techniques to estimate a robust policy. However, these methods present some limiting restrictions such as non-scalability and offline estimation. Besides being computationally more efficient than batch learning [Wiering and van Otterlo, 2012], the use of an online algorithm is of significant interest in

robust MDPs because it can detect non-adversarial state-actions pairs along a trajectory and result in less conservative results, something which cannot be performed when solving the planning problem [Lim et al., 2016]. Other works have attempted to incorporate robustness into an online algorithm for policy optimization [Mankowitz et al., 2018; Tamar et al., 2015]. Although these approaches can deal with large domains, a sampling procedure is required for each critic estimate in Tamar et al. [2015], which differs from the strictly-speaking actor-critic. In Mankowitz et al. [2018], the authors introduce a robust version of actor-critic policy-gradient but its convergence results are only shown for the actor updates. Moreover, these works target the robust solution which may be too conservative. We review all existing methods in Section 7 and compare them to our approach.

To the best of our knowledge, our proposed work is the first attempt to incorporate a soft form of robustness into an online algorithm that has convergence guarantees besides being computationally scalable. We deal with the curse of dimensionality by using function approximation that parameterizes the expected value within a space of much smaller dimension than the state space. By fixing a distribution over the uncertainty set, the induced soft-robust actor-critic learns a locally optimal policy in an online manner. Under mild assumptions on the set of distributions and uncertainty set, we show that our novel Soft-Robust Actor-Critic (SR-AC) algorithm converges. We test the performance of soft-robustness on different domains, including a large state space with continuous actions. As far as we know, no other work has previously incorporated robustness into continuous action spaces.

Our specific contributions are: (1) A soft-robust derivation of the objective function for policy-gradient; (2) An SR-AC algorithm that uses stochastic approximation to learn a variant of distributionally robust policy in an online manner; (3) Convergence proofs of SR-AC; (4) An experiment of our framework to different domains that shows the efficiency of soft-robust behaviors in a continuous action space as well. All proofs can be found in the Appendix.

2 BACKGROUND

In this section, we introduce the background material related to our soft-robust approach.

Robust MDP A robust MDP is a tuple $\langle \mathcal{X}, \mathcal{A}, r, \mathcal{P} \rangle$ where \mathcal{X} is a finite state-space, \mathcal{A} is a finite set of actions, $r : \mathcal{X} \times \mathcal{A} \rightarrow \mathbb{R}$ is the immediate reward function which is deterministic and bounded and \mathcal{P} is a set of transition matrices. We assume that \mathcal{P} is structured as a cartesian product $\otimes_{x \in \mathcal{X}} \mathcal{P}_x$, which is known as the

rectangularity assumption [Nilim and El Ghaoui, 2005]. Given a state $x \in \mathcal{X}$, the uncertainty set \mathcal{P}_x is a family of transition models $p_x \in \mathcal{P}_x$ we represent as vectors in which the transition probabilities of each action are arranged in the same block. For $x, y \in \mathcal{X}$ and $a \in \mathcal{A}$, denote by $p(x, a, y)$ the probability of getting from state x to state y given action a .

At timestep t , the agent is in state x_t and chooses an action a_t according to a stochastic policy $\pi : \mathcal{X} \rightarrow \mathcal{M}(\mathcal{A})$ that maps each state to a probability distribution over the action space, $\mathcal{M}(\mathcal{A})$ denoting the set of distributions over \mathcal{A} . It then gets a reward r_{t+1} and is brought to state x_{t+1} with probability $p(x_t, a_t, x_{t+1})$.

Policy-Gradient Policy-gradient methods are commonly used to learn an agent policy. A policy π is parametrized by θ and estimated by optimizing an objective function using stochastic gradient descent. A typical objective to be considered is the average reward function

$$\begin{aligned} J_p(\pi) &= \lim_{T \rightarrow +\infty} \mathbb{E}^p \left[\frac{1}{T} \sum_{t=0}^{T-1} r_{t+1} \mid \pi \right] \\ &= \sum_{x \in \mathcal{X}} d_p^\pi(x) \sum_{a \in \mathcal{A}} \pi(x, a) r(x, a) \end{aligned}$$

where r_t is the reward at time t , p an aperiodic and irreducible transition model under which the agent operates and d_p^π is the stationary distribution of the Markov process induced by p under policy π . The gradient objective has previously been shown to be

$$\nabla_\theta J_p(\pi) = \sum_{x \in \mathcal{X}} d_p^\pi(x) \sum_{a \in \mathcal{A}} \nabla_\theta \pi(x, a) Q_p^\pi(x, a)$$

where $Q_p^\pi(x, a)$ is the expected differential reward associated with state-action pair (x, a) . This gradient is then used to update the policy parameters according to: $\theta_{t+1} = \theta_t + \beta_t \nabla_\theta J_p(\pi)$, with β_t a positive step-size [Sutton et al., 2000].

Actor-Critic Algorithm Theoretical analysis and empirical experiments have shown that regular policy-gradient methods present a major issue namely high variance in the gradient estimates that results in slow convergence and inefficient sampling [Grondman et al., 2012]. First proposed by Barto et al. [1983], actor-critic methods attempt to reduce the variance by using a critic that estimates the value function. They borrow elements from both value function and policy-based methods. The value function estimate plays the role of a critic that helps evaluating the performance of the policy. As in policy-based methods, the actor then uses this signal to update policy parameters in the direction of a gradient estimate of a performance measure. Under appropriate conditions, the resulting algorithm is tractable and converges to a locally optimal policy [Bhatnagar et al., 2009; Konda and Tsitsiklis, 2000].

Deep Q-networks Deep Q-Networks (DQNs) have proven their capability of solving complex learning tasks such as Atari video games [Mnih et al., 2013]. The Q-learning of Watkins and Dayan [1992] typically learns a greedy or ϵ -greedy policy by updating the Q-function based on a TD-error. In Deep Q-learning [Mnih et al., 2013, 2015], a non-linear function such as a neural network is used as an approximator of the Q-function. It is referred to as a Q-network. The agent is then trained by optimizing the induced TD loss function thanks to stochastic gradient descent. Like actor-critic, DQN is an online algorithm that aims at finding an optimal policy. The main difference with actor-critic is that it is *off-policy*: it learns a greedy strategy while following an arbitrary behavior [Mnih et al., 2013].

Deep Deterministic Policy-Gradient Since DQN acts greedily at each iteration, it can only handle small action spaces. The Deep Deterministic Policy-Gradient (DDPG) is an *off-policy* algorithm that can learn behaviors in continuous action spaces [Lillicrap et al., 2016]. It is based on an actor-critic architecture that follows the same baseline as in DQN. The critic estimates the current Q-value of the actor using a TD-error while the actor is updated according to the critic. This update is based on the chain rule principle which establishes equivalence between the stochastic and the deterministic policy gradient [Silver et al., 2014].

3 SOFT-ROBUSTNESS

3.1 SOFT-ROBUST FRAMEWORK

Unlike robust MDPs that maximize the worst-case performance, we fix a prior on how transition models are distributed over the uncertainty set. A distribution over \mathcal{P} is denoted by ω and is structured as a cartesian product $\otimes_{x \in \mathcal{X}} \omega_x$. We find the same structure in Xu and Mannor [2012]; Yu and Xu [2016]. Intuitively, ω can be thought as the way the adversary distributes over different transition models. The product structure then means that this adversarial distribution only depends on the current state of the agent without taking into account its whole trajectory. This defines a probability distribution ω_x over \mathcal{P}_x independently for each state.

We further assume that ω is non-diffuse. This implies that the uncertainty set is non-trivial with respect to ω in a sense that the distribution does not affect zero mass to all of the models.

3.2 SOFT-ROBUST OBJECTIVE

Throughout this paper, we make the following assumption:

Assumption 3.1. *Under any policy π , the Markov chains resulting from any of the MDPs with transition laws $p \in \mathcal{P}$ are irreducible and aperiodic.*

Define d_p^π as the stationary distribution of the Markov chain that results from following policy π under transition model $p \in \mathcal{P}$.

Definition 3.1. *We call soft-robust objective or soft-robust average reward the function $\bar{J}(\pi) := \mathbb{E}_{p \sim \omega} [J_p(\pi)]$.*

The distribution ω introduces a softer form of robustness in the objective function because it averages over the uncertainty set instead of considering the worst-case scenario. It also gives flexibility over the level of robustness one would like to keep. A robust strategy would then consist of putting more mass on pessimistic transition models. Likewise, a distribution that puts all of its mass on one target model would lead to an aggressive behavior and result in model misspecification.

The *soft-robust differential reward* is given by $\bar{Q}^\pi(x, a) := \mathbb{E}_{p \sim \omega} [Q_p^\pi(x, a)]$ where

$$Q_p^\pi(x, a) := \mathbb{E}^p \left[\sum_{t=0}^{+\infty} r_{t+1} - J_p(\pi) \mid x_0 = x, a_0 = a, \pi \right].$$

Similarly, we introduce the quantity

$$\bar{V}^\pi(x) := \sum_{a \in \mathcal{A}} \pi(x, a) \bar{Q}^\pi(x, a) = \mathbb{E}_{p \sim \omega} [V_p^\pi(x)]$$

with $V_p^\pi(x) := \sum_{a \in \mathcal{A}} \pi(x, a) Q_p^\pi(x, a)$. We will interchangeably term it as *soft-robust expected differential reward* or *soft-robust value function*.

3.3 SOFT-ROBUST STATIONARY DISTRIBUTION

The above performance objective $\bar{J}(\pi)$ cannot as yet be written as an expectation of the reward over a stationary distribution because of the added measure ω on transition models. Define the average transition model as $\bar{p} := \mathbb{E}_{p \sim \omega} [p]$. It corresponds to the transition probability that results from distributing all transition models according to ω . In analogy to the transition probability that minimizes the reward for each given state and action in the robust transition function [Mankowitz et al., 2018], our average model rather selects the expected distribution over all the uncertainty set for each state and action. Under Assumption 3.1, we can show that the transition \bar{p} as defined is irreducible and aperiodic, which ensures the existence of a unique stationary law we will denote by \bar{d}^π .

Proposition 3.1 (Stationary distribution in the average transition model). *Under Assumption 3.1, the average transition matrix $\bar{p} := \mathbb{E}_{p \sim \omega}[p]$ is irreducible and aperiodic. In particular, it admits a unique stationary distribution.*

As in regular MDPs, the soft-robust average reward satisfies a Poisson equation, as it was first stated in the discounted reward case in Lemma 3.1 of Xu and Mannor [2012]. The following proposition reformulates this result for the average reward.

Proposition 3.2 (Soft-Robust Poisson equation).

$$\begin{aligned} \bar{J}(\pi) + \bar{V}^\pi(x) \\ = \sum_{a \in \mathcal{A}} \pi(x, a) \left(r(x, a) + \sum_{x' \in \mathcal{X}} \bar{p}(x, a, x') \bar{V}^\pi(x') \right) \end{aligned}$$

This Poisson equation enables us to establish an equivalence between the expectation of the stationary distributions over the uncertainty set and the stationary distribution of the average transition model, naming $\bar{d}^\pi(x) = \mathbb{E}_{p \sim \omega}[d_p^\pi(x)]$ with $x \in \mathcal{X}$. Indeed, we have the following:

Corollary 3.1. *Recall \bar{d}^π the stationary distribution for the average transition model \bar{p} . Then*

$$\bar{J}(\pi) = \sum_{x \in \mathcal{X}} \bar{d}^\pi(x) \sum_{a \in \mathcal{A}} \pi(x, a) r(x, a).$$

The goal is to learn a policy that maximizes the soft-robust average reward \bar{J} . We use a policy-gradient method for that purpose.

4 SOFT-ROBUST POLICY-GRADIENT

In policy-gradient methods, we consider a class of parametrized stochastic policies $\pi_\theta : \mathcal{X} \rightarrow \mathcal{M}(\mathcal{A})$ with $\theta \in \mathbb{R}^{d_\theta}$ and estimate the gradient of the objective function \bar{J} with respect to policy parameters in order to update the policy in the direction of the estimated gradient of \bar{J} . The optimal set of parameters thus obtained is denoted by

$$\theta^* := \arg \max_{\theta} \bar{J}(\pi_\theta).$$

When clear in the context, we will omit the subscript θ in π_θ for notation ease. We further make the following assumption, which is standard in policy-gradient literature:

Assumption 4.1. *For any $(x, a) \in \mathcal{X} \times \mathcal{A}$, the mapping $\theta \mapsto \pi_\theta(x, a)$ is continuously differentiable with respect to θ .*

Using the same method as in Sutton et al. [2000], we can derive the gradient of the soft-robust average reward thanks to the previous results.

Theorem 4.1 (Soft-Robust Policy-Gradient). *For any MDP satisfying previous assumptions, we have*

$$\nabla_{\theta} \bar{J}(\pi) = \sum_{x \in \mathcal{X}} \bar{d}^\pi(x) \sum_{a \in \mathcal{A}} \nabla_{\theta} \pi(x, a) \bar{Q}^\pi(x, a).$$

In order to manage with large state spaces, we also introduce a linear approximation of \bar{Q}^π we define as $f_w(x, a) := w^T \psi_{xa}$. Sutton et al. [2000] showed that if the features ψ_{xa} satisfy a compatibility condition and the approximation is locally optimal, then we can use it in place of \bar{Q}^π and still point roughly in the direction of the true gradient.

In the case of soft-robust average reward, this defines a soft-robust gradient update that possesses the ability to incorporate function approximation, as stated in the following result. The main difference with that of Sutton et al. [2000] is that we combine the dynamics of the system with distributed transitions over the uncertainty set.

Theorem 4.2 (Soft-Robust Policy-Gradient with Function Approximation). *Let $f_w : \mathcal{X} \times \mathcal{A} \rightarrow \mathbb{R}$ be a linear approximator of the soft-robust differential reward \bar{Q}^π . If f_w minimizes the mean squared error*

$$\mathcal{E}^\pi(w) := \sum_{x \in \mathcal{X}} \bar{d}^\pi(x) \sum_{a \in \mathcal{A}} \pi(x, a) \left[\bar{Q}^\pi(x, a) - f_w(x, a) \right]^2$$

and is compatible in a sense that $\nabla_w f_w(x, a) = \nabla_{\theta} \log \pi(x, a)$, then

$$\nabla_{\theta} \bar{J}(\pi) = \sum_{x \in \mathcal{X}} \bar{d}^\pi(x) \sum_{a \in \mathcal{A}} \nabla_{\theta} \pi(x, a) f_w(x, a)$$

We can further improve our gradient estimate by reducing its variance. One direct method to do so is to subtract a baseline $b(x)$ from the previous gradient update. It is easy to show that this will not affect the gradient derivation. In particular, Bhatnagar et al. [2009] proved that the value function minimizes the variance. It is therefore a proper baseline to choose. We can thus write the following:

$$\begin{aligned} \nabla_{\theta} \bar{J}(\pi) &= \sum_{x \in \mathcal{X}} \bar{d}^\pi(x) \sum_{a \in \mathcal{A}} \nabla_{\theta} \pi(x, a) \left(\bar{Q}^\pi(x, a) - \bar{V}^\pi(x) \right) \\ &= \sum_{x \in \mathcal{X}} \bar{d}^\pi(x) \sum_{a \in \mathcal{A}} \pi(x, a) \psi_{xa} \bar{A}^\pi(x, a), \end{aligned} \tag{1}$$

where $\bar{A}^\pi(x, a)$ is the *soft-robust advantage function* defined by $\bar{A}^\pi(x, a) := \bar{Q}^\pi(x, a) - \bar{V}^\pi(x)$.

5 SOFT-ROBUST ACTOR-CRITIC ALGORITHM

In this section, we present our SR-AC algorithm which is defined as Algorithm 1. This novel approach incorporates a variation of distributional robustness into an online algorithm that effectively learns an optimal policy in a scalable manner. Under mild assumptions, the resulting two-timescale stochastic approximation algorithm converges to a locally optimal policy.

5.1 SR-AC ALGORITHM

An uncertainty set and a nominal model without uncertainty are provided as inputs. In practice, the nominal model and the uncertainty set can respectively be an estimate of the transition model resulting from data sampling and its corresponding confidence interval. A distribution ω over the uncertainty set is also provided. It corresponds to our prior information on the uncertainty set. The step-size sequences $(\alpha_t, \beta_t, \xi_t; t \geq 0)$ consist of small non-negative numbers properly chosen by the user (see Appendix for more details).

At each iteration, samples are generated using the nominal model and the current policy. These are utilized to update the soft-robust average reward (Line 5) and the critic (Line 7) based on an estimate of a soft-robust TD-error we detail further. In our setting, the soft-robust value function plays the role of the critic according to which the actor parameters are updated. We then exploit the critic to improve our policy by updating the policy parameters in the direction of a gradient estimate for the soft-robust objective (Line 8). This process is repeated until convergence.

5.2 CONVERGENCE ANALYSIS

We establish convergence of SR-AC to a local maximum of the soft-robust objective function by following an ODE approach [Kushner and Yin, 1997].

Consider \hat{V} and \hat{J} as unbiased estimates of \bar{V} and \bar{J} respectively. Calculating δ_t (Line 6 in Algorithm 1) requires an estimate of the soft-robust average-reward that can be obtained by averaging over samples given immediate reward r and distribution ω (Line 5). In order to get an estimate of the soft-robust differential value \hat{V} , we use linear function approximation. Considering φ as a d_2 -dimensional feature extractor over the state space \mathcal{X} , we may then approximate $\bar{V}^\pi(x)$ as $v^T \varphi_x$, where v is a d_2 -dimensional parameter vector that we tune using linear TD. This results in the following soft-robust TD-error:

$$\delta_t := r_{t+1} - \hat{J}_{t+1} + \sum_{x' \in \mathcal{X}} \bar{p}(x_t, a_t, x') v_t^T \varphi_{x'} - v_t^T \varphi_{x_t},$$

Algorithm 1 SR-AC

- 1: **Input:** \mathcal{P} - An uncertainty set; $\hat{p} \in \mathcal{P}$ - A nominal model; ω - A distribution over \mathcal{P} ; f_x - A feature extractor for the SR value function;
 - 2: **Initialize:** $\theta = \theta_0$ - An arbitrary policy parameter; $v = v_0$ - An arbitrary set of value function parameters; α_0, β_0, ξ_0 - Initial learning-rates; x_0 - Initial state
 - 3: **repeat**
 - 4: Act under $a_t \sim \pi_{\theta_t}(x_t, a_t)$
 Observe next state x_{t+1} and reward r_{t+1}
 - 5: **SR Average Reward Update:**
 $\hat{J}_{t+1} = (1 - \xi_t)\hat{J}_t + \xi_t r_{t+1}$
 - 6: **SR TD-Error:**
 $\delta_t = r_{t+1} - \hat{J}_{t+1} + \sum_{x' \in \mathcal{X}} \bar{p}(x_t, a_t, x') \hat{V}_{x'} - \hat{V}_{x_t}$
 - 7: **Critic Update:** $v_{t+1} = v_t + \alpha_t \delta_t \varphi_{x_t}$
 - 8: **Actor Update:** $\theta_{t+1} = \theta_t + \beta_t \delta_t \psi_{x_t a_t}$
 - 9: **until** convergence
 - 10: **Return:** SR policy parameters θ and SR value-function parameters v
-

where v_t corresponds to the current estimate of the soft-robust value function parameter.

As in regular MDPs, when doing linear TD learning, the function approximation of the value function introduces a bias in the gradient estimate [Bhatnagar et al., 2009]. Denoting it as e^π , we have $E[\widehat{\nabla_\theta J}(\pi) | \theta] = \nabla_\theta \bar{J}(\pi) + e^\pi$ (see Appendix). This bias term then needs to be small enough in order to ensure convergence.

Convergence of Algorithm 1 can be established by applying Theorem 2 from Bhatnagar et al. [2009] which exploits Borkar’s work on two-timescale algorithms [1997]. The convergence result is presented as Theorem 5.1.

Theorem 5.1. *Under all the previous assumptions, given $\epsilon > 0$, there exists $\delta > 0$ such that for a parameter vector $\theta_t, t \geq 0$ obtained using the algorithm, if $\sup_{\pi_t} \|e^{\pi_t}\| < \delta$, then the SR-AC algorithm converges almost surely to an ϵ -neighborhood of a local maximum of \bar{J} .*

6 NUMERICAL EXPERIMENTS

We demonstrate the performance of soft-robustness on various domains of finite as well as continuous state and action spaces. We used the existing structure of OpenAI Gym environments to run our experiments [Brockman et al., 2016].

6.1 DOMAINS

Single-step MDP We consider a simplified formulation of the startup vs teaching dilemma described in Section 1.

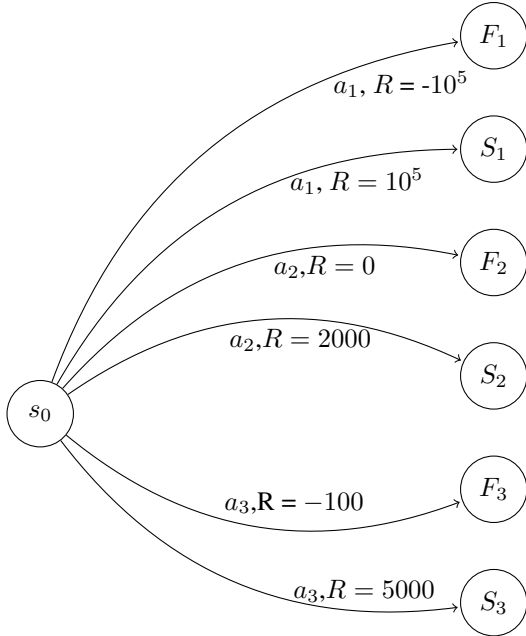


Figure 1: Illustration of the MDP with initial state s_0 . States F_1, F_2, F_3 correspond to failing scenarios for each action. The succeeding states are represented by states S_1, S_2, S_3 .

The problem is modeled as a 7-state MDP in which one action corresponds to one strategy. An illustration of this construction is given in Figure 1. At the starting state s_0 , the agent chooses one of three actions. Action a_1 [corresponds to the startup adventure] may lead it to a very high reward in case of success but can be catastrophic in case of failure. Action a_2 [corresponds to the teaching carrier] leads it to low positive reward in case of success with no possibility of negative reward. Action a_3 [corresponds to an intermediate strategy] can lead to an intermediate positive reward with a slight risk of negative reward. Depending on the action it chose and if it succeeded or not, the agent is brought to one of the six right-hand states and receives the corresponding reward. It is brought back to s_0 at the end of each episode. We assume the probability of success to be the same for all three actions.

Cart-Pole In the Cart-Pole system, the agent’s goal consists of balancing a pole atop a cart in a vertical position. It is modeled as a continuous MDP in which each state consists of a 4-tuple $\langle x, \dot{x}, \theta, \dot{\theta} \rangle$ which represents the cart position, the cart speed, the pole angle with respect to the vertical and its angular speed respectively. The agent can make two possible actions: apply a constant force either to the right or to the left of the pole. It gets a positive reward of 1 if the pole has not fallen down and if it stayed in the boundary sides of the screen. If it terminates, the agent receives a reward of 0. Since each episode lasts for

200 timesteps, the maximal reward an agent can get is 200 over one episode.

Pendulum In the inverted pendulum problem, a pendulum starts in a random position and the goal is to swing it up so that it stabilizes upright. The state domain consists in a 2-tuple $\langle \theta, \dot{\theta} \rangle$ which represents the pendulum angle with respect to the vertical and its angular velocity. At each timestep, the agent’s possible actions belong to a continuous interval $[-a, a]$ which represents the force level being applied. Since there is no specified termination, we establish a maximal number of 200 steps for each episode.

6.2 UNCERTAINTY SETS

For each experiment, we generate an uncertainty set \mathcal{P} before training. In the single-step MDP, we sample from 5 different probabilities of success using a uniform distribution over $[0, 1]$. In Cart-Pole, we sample 5 different lengths from a normal distribution centered at the nominal length of the pole which we fix at 0.3. We proceed similarly for Pendulum by generating 10 different masses of pendulum around a nominal mass of 2. Each corresponding model thus generates a different transition function. We then sample the average model by fixing ω as a realization of a Dirichlet distribution. A soft-robust update for the actor is applied by taking the optimal action according to this average transition function.

6.3 LEARNING ALGORITHMS

We trained the agent on the nominal model in each experiment. The soft-robust agent was learned using SR-AC in the single-step MDP. In Cart-Pole, we run a soft-robust version of a DQN algorithm. The soft-robust agent in Pendulum was trained using a soft-robust DDPG.

Soft-Robust AC We analyze the performance of SR-AC by training a soft-robust agent on the single-step MDP. We run a regular AC algorithm to derive an aggressive policy and learn a robust behavior by using a robust formulation of AC which consists in replacing the TD-error with a robust TD-error, as implemented in Mankowitz et al. [2018]. The derived soft-robust agent is then compared with the resulting aggressive and robust strategies respectively.

Soft-Robust DQN Robustness has already been incorporated in DQN [Di-Castro Shashua and Mannor, 2017]. The Q-network addressed there performs an online estimation of the Q-function by minimizing at each timestep

t the following robust TD-error:

$$\begin{aligned} \delta_{dq_n,t}^{rob} &:= r(x_t, a_t) - Q(x_t, a_t) \\ &+ \gamma \inf_{p \in \mathcal{P}} \sum_{x' \in \mathcal{X}} p(x_t, a_t, x') \max_{a' \in \mathcal{A}} Q(x', a'), \end{aligned}$$

where γ is a discount factor.

In our experiments, we incorporate a soft-robust TD-error inside a DQN that trains a soft-robust agent according to the induced loss function. The soft-robust TD-error for DQN is given by:

$$\begin{aligned} \delta_{dq_n,t}^{srob} &:= r(x_t, a_t) - Q(x_t, a_t) \\ &+ \gamma \sum_{x' \in \mathcal{X}} \bar{p}(x_t, a_t, x') \max_{a' \in \mathcal{A}} Q(x', a') \end{aligned}$$

We use the Cart-Pole domain to compare the resulting policy with the aggressive and robust strategies that were obtained from a regular and a robust DQN respectively.

Soft-Robust DDPG Define μ_t as the estimated deterministic policy at step t . We incorporate robustness in DDPG by updating the critic network according to the following robust TD-error:

$$\begin{aligned} \delta_{ddpg,t}^{rob} &:= r(x_t, a_t) - Q(x_t, a_t) \\ &+ \gamma \inf_{p \in \mathcal{P}} \sum_{x' \in \mathcal{X}} p(x_t, a_t, x') Q(x', \mu(x_t)), \end{aligned}$$

Similarly, we incorporate soft-robustness in DDPG by using the soft-robust TD-error:

$$\begin{aligned} \delta_{ddpg,t}^{srob} &:= r(x_t, a_t) - Q(x_t, a_t) \\ &+ \gamma \sum_{x' \in \mathcal{X}} \bar{p}(x_t, a_t, x') Q(x', \mu(x_t)) \end{aligned}$$

We compare the resulting soft-robust DDPG with its regular and robust formulations in the Pendulum domain.

6.4 IMPLEMENTATION

For each experiment, we train the agent on the nominal model but incorporate soft-robustness during learning. A soft-robust policy is learned thanks to SR-AC in the single-step MDP. We use a linear function approximation with 5 features to estimate the value function. For Cart-Pole, we run a DQN using a neural network of 3 fully-connected hidden layers with 128 weights per layer and ReLU activations. In Pendulum, a DDPG algorithm learns a policy based on two target networks: the actor and the critic network. Both have 2 fully-connected hidden layers with 400 and 300 units respectively. We use a tanh activation for the actor and a ReLU activation for the critic output. We chose the ADAM optimizer to minimize all the induced loss functions. We used constant

learning rates which worked well in practice. Each agent was trained over 3000 episodes for the single-step MDP and Cartpole and tested over 600 episodes per parameter setting. For Pendulum, the agents were trained over 5000 episodes evaluated over 800 episodes per parameter setting. Other hyper-parameter values can be found in the Appendix.

6.5 RESULTS

Single-step MDP Figure 2 shows the evolution of the performance for all three agents during training. It becomes more stable along training time, which confirms convergence of SR-AC. We see that the aggressive agent performs best due to the highest reward it can reach on the nominal model. The soft-robust agent gets rewards in between the aggressive and the robust agent which performs the worst due to its pessimistic learning method.

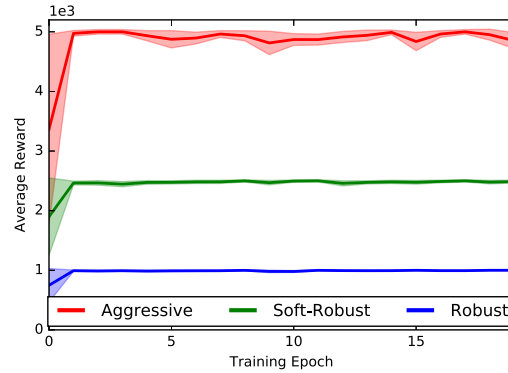


Figure 2: Comparison of robust, soft-robust and aggressive agents during training. One training epoch corresponds to 300 episodes.

The evaluation of each strategy is represented in Figure 3. As the probability of success gets low, the performance of the aggressive agent drops down below the robust and the soft-robust agents, although it performs best when the probability of success gets close to 1. The robust agent stays stable independently of the parameters but underperforms soft-robust agent which presents the best balance between high reward and low risk. We noticed that depending on the weighting distribution initially set, soft-robustness tends to being more or less aggressive (see Appendix). Incorporating a distribution over the uncertainty set thus gives significant flexibility on the level of aggressiveness to be assigned to the soft-robust agent.

Cart-Pole In Figure 4, we show the performance of all three strategies over different values of pole length during

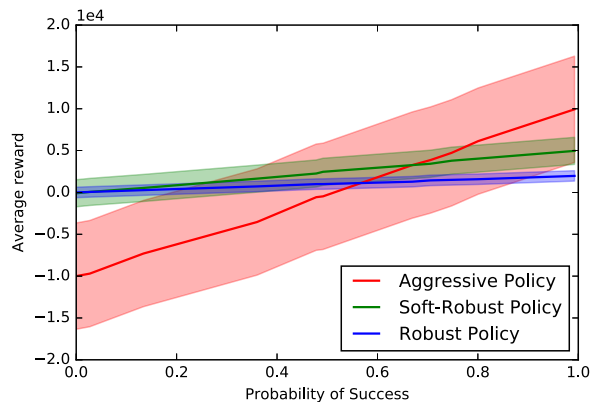


Figure 3: Average reward for AC, robust AC and SR-AC methods

testing. Similarly to our previous example, the non-robust agent performs well around the nominal model but its reward degrades on more extreme values of pole length. The robust agent keeps a stable reward under model uncertainty which is consistent with the results obtained in Di-Castro Shashua and Mannor [2017]; Mankowitz et al. [2018]. However, it is outperformed by the soft-robust agent around the nominal model. Furthermore, the soft-robust strategy shows an equilibrium between aggressiveness and robustness thus leading to better performance than the non-robust agent on larger pole lengths. We trained a soft-robust agent on other weighting distributions and noted that depending on its structure, soft-robustness interpolates between aggressive and robust behaviors (see Appendix).

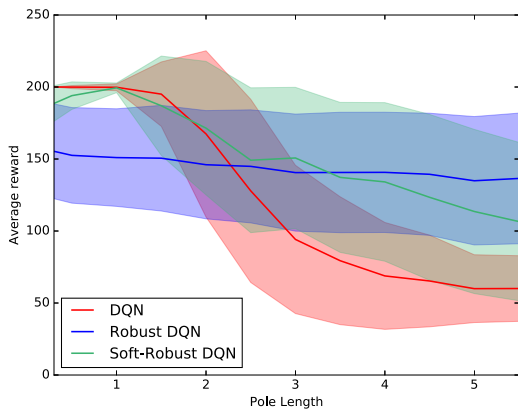


Figure 4: Average reward performance for DQN, robust DQN and soft-robust DQN

Pendulum Figure 5 shows the performance of all three

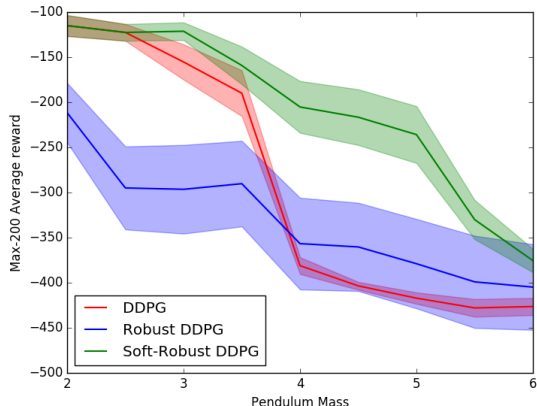


Figure 5: Max-200 episodes average performance for DDPG, robust DDPG and soft-robust DDPG

agents when evaluating them on different masses. Since the performance among different episodes is highly variable, we considered the best 200-episodes average reward as a performance measure. As seen in the figure, the robust strategy solves the task in a sub-optimal fashion, but is less affected by model misspecification due to its conservative strategy. The aggressive non-robust agent is more sensitive to model misspecification compared to the other methods as can be seen by its sudden dip in performance, below even that of the robust agent. The soft-robust solution strikes a nice balance between being less sensitive to model misspecification than the aggressive agent, and producing better performance compared to the robust solution.

7 RELATED WORK

This paper is related to several domains in RL such as robust and distributionally robust MDPs, actor-critic methods and online learning via stochastic approximation algorithms. Our work solves the problem of conservativeness encountered in robust MDPs by incorporating a variational form of distributional robustness. The SR-AC algorithm combines scalability to large scale state-spaces and online estimation of the optimal policy in an actor-critic algorithm. Table 1 compares our proposed algorithm with previous approaches.

Many solutions have been addressed to mitigate conservativeness of robust MDP. Mannor et al. [2012, 2016] relax the state-wise independence property of the uncertainty set and assume it to be coupled in a way such that the planning problem stays tractable. Another approach tends to assume *a priori* information on the parameter set. These

Table 1: Comparison of previous approaches with SR-AC

Reference	Scalable	Actor-Critic	Softly-Robust
SR-AC (this paper)	✓	✓	✓
Mankowitz et al. [2018]	✓	✗	✗
Lim et al. [2016]	✗	✗	✗
Yu and Xu [2016]	✗	✗	✓
Mannor et al. [2012, 2016]	✗	✗	✗
Tamar et al. [2015]	✓	✗	✗
Xu and Mannor [2012]	✗	✗	✓
Bhatnagar et al. [2009]	✓	✓	✗

methods include distributionally robust MDPs [Xu and Mannor, 2012; Yu and Xu, 2016] in which the optimal strategy maximizes the expected reward under the most adversarial distribution over the uncertainty set. For finite and known MDPs, under some structural assumptions on the considered set of distributions, this max-min problem reduces to classical robust MDPs and can be solved efficiently by dynamic programming [Puterman, 2009].

However, besides becoming untractable under large-sized MDPs, these methods use an offline learning approach which cannot adapt its level of protection against model uncertainty and may lead to overly conservative results. The work of Lim et al. [2016] solutions this issue and addresses an online algorithm that learns the transitions that are purely stochastic and those that are adversarial. Although it ensures less conservative results as well as low regret, this method sticks to the robust objective while strongly relying on the finite structure of the state-space. To alleviate the curse of dimensionality, we incorporate function approximation of the objective value and define it as a linear functional of features.

First introduced in Barto et al. [1983] and later addressed by Bhatnagar et al. [2009], actor-critic algorithms are online learning methods that aim at finding an optimal policy. We used the formulation of Bhatnagar et al. [2009] as a baseline for the algorithm we proposed. The key difference between their work and ours is that we incorporate soft-robustness. This relates in a sense to the Bayesian Actor-Critic setup in which the critic returns a complete posterior distribution of value functions using Bayes’ rule [Ghavamzadeh and Engel, 2007; Ghavamzadeh et al., 2015, 2016]. Our study keeps a frequentist approach, meaning that our algorithm updates return point estimates of the average value-function which prevents from tractability issues besides enabling the distribution to be more

flexible. Another major distinction is that the Bayesian approach incorporates a prior distribution on one model parameters whereas our method considers a prior on different transition models over an uncertainty set.

In Mankowitz et al. [2018]; Tamar et al. [2015], the authors incorporate robustness into policy-gradient methods. A sampling procedure is required for each critic estimate in Tamar et al. [2015], which differs from the strictly-speaking actor-critic. A robust version of actor-critic policy-gradient is introduced in Mankowitz et al. [2018] but its convergence guarantees are only shown for robust policy-gradient ascent. Both of these methods target the robust strategy whereas we seek a soft-robust policy that is less conservative while protecting itself against model uncertainty.

8 DISCUSSION

We have presented the SR-AC framework that is able to learn policies which keep a balance between aggressive and robust behaviors. SR-AC requires a stationary distribution under the average transition model and compatibility conditions for deriving a soft-robust policy-gradient. We have shown that this ensures convergence of SR-AC. This is the first work that has attempted to incorporate a soft form of robustness into an online actor-critic method. Our approach has been shown to be computationally scalable to large domains because of its low computational price. In our experiments, we have also shown that the soft-robust agent interpolates between aggressive and robust strategies without being overly conservative which leads it to outperform robust policies under model uncertainty even when the action space is continuous. Subsequent experiments should test the efficiency of soft-robustness on more complex domains.

The chosen weighting over the uncertainty set can be thought as the way the adversary distributes over different transition laws. In our current setting, this adversarial distribution stays constant without accounting for the rewards obtained by the agent. Future work should address the problem of learning the sequential game induced by an evolving adversarial distribution to derive an optimal soft-robust policy. Other extensions of our work may also consider non-linear objective functions such as higher order moments with respect to the adversarial distribution.

Acknowledgements

This work was partially funded by the Israel Science Foundation under contract 1380/16 and by the European Community’s Seventh Framework Programme (FP7/2007-2013) under grant agreement 306638 (SUPREL).

References

- Andrew G. Barto, Richard S. Sutton, and Charles W. Anderson. Neuronlike adaptive elements that can solve difficult learning control problems. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-13(5):834–846, 1983.
- Shalabh Bhatnagar, Richard Sutton, Mohammad Ghavamzadeh, and Mark Lee. *Natural Actor-Critic Algorithms*. Automatica, elsevier edition, 2009.
- Vivek S. Borkar. Stochastic Approximation with Two Timescales. *Systems and Control Letters*, 29:291–294, 1997.
- Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. OpenAI Gym. arXiv:1606.01540v1, 2016.
- Shirli Di-Castro Shashua and Shie Mannor. Deep Robust Kalman Filter. *arXiv preprint arXiv:1703.02310v1*, 2017.
- Mohammad Ghavamzadeh and Yaakov Engel. Bayesian Actor-Critic Algorithms. *Proceedings of the 24th international conference on Machine learning*, pages 297–304, 2007.
- Mohammad Ghavamzadeh, Shie Mannor, Joelle Pineau, and Aviv Tamar. Bayesian Reinforcement Learning: A Survey. *Foundations and Trends in Machine Learning*, 8(5-6):359–492, 2015.
- Mohammad Ghavamzadeh, Yaakov Engel, and Michal Valko. Bayesian Policy Gradient and Actor-Critic Algorithms. *Journal of Machine Learning Research*, 17(66):1–53, 2016.
- Ivo Grondman, Lucian Busoniu, Gabriel A.D. Lopes, and Robert Babuska. A Survey of Actor-Critic Reinforcement Learning: Standard and Natural Policy Gradients. *IEEE Transactions on Systems, Man, and Cybernetics—Part C: Applications and Reviews*, 42(1291-1307), 2012.
- Garud N. Iyengar. Robust Dynamic Programming. *Mathematics of Operations Research*, 30(2):257–280, 2005.
- Vijay R. Konda and John N. Tsitsiklis. Actor-Critic Algorithms. In *Advances in Neural Information Processing Systems*, volume 12, 2000.
- Harold J. Kushner and G. George Yin. *Stochastic Approximation Algorithms and Applications*. Springer Verlag, New York, 1997.
- Timothy P. Lillicrap, Jonathan J. Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous Control with Deep Reinforcement Learning. *arXiv:1509.02971*, US Patent App. 15/217,758, 2016.
- Shiau Hong Lim, Huan Xu, and Shie Mannor. Reinforcement Learning in Robust Markov Decision Processes. *Mathematics of Operations Research*, 41(4):1325–1353, 2016.
- Daniel J Mankowitz, Timothy A Mann, Shie Mannor, Doina Precup, and Pierre-Luc Bacon. Learning Robust Options. In *AAAI*, 2018.
- Shie Mannor, Duncan Simester, Peng Sun, and John N. Tsitsiklis. Bias and Variance Approximation in Value Function Estimates. *Management Science*, 53(2):308–322, 2007.
- Shie Mannor, Ofir Mebel, and Huan Xu. Lightning Does Not Strike Twice: Robust MDPs with Coupled Uncertainty. In *ICML*, 2012.
- Shie Mannor, Ofir Mebel, and Huan Xu. Robust MDPs with k-Rectangular Uncertainty. *Mathematics of Operations Research*, 41(4):1484–1509, 2016.
- Danny Miller and Peter H. Friesen. Innovation in conservative and entrepreneurial firms: Two models of strategic momentum. *Strategic Management Journal*, 1982.
- Olivia S. Mitchell and Kent Smetters. *The Market for Retirement Financial Advice*. Oxford University Press, first edition edition, 2013.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing Atari with Deep Reinforcement Learning: Technical Report. *DeepMind Technologies*, arXiv:1312.5602, 2013.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin Riedmiller, Andreas K. Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharshan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. Human-level control through deep reinforcement learning. *Nature*, 518:529–533, 2015.
- Arnab Nilim and Laurent El Ghaoui. Robust control of Markov decision processes with uncertain transition matrices. *Operations Research*, 53(5):783–798, 2005.
- Martin L. Puterman. *Markov decision processes: Discrete stochastic dynamic programming*, volume 414. Wiley Series in Probability and Statistics, 2009.
- Aurko Roy, Huan Xu, and Sebastian Pokutta. Reinforcement learning under Model Mismatch. *31st Conference on Neural Information Processing Systems*, 2017.
- David Silver, Guy Lever, Nicolas Heess, Thomas Degris, Daan Wierstra, and Martin Riedmiller. Deterministic Policy Gradient Algorithms. *ICML*, 2014.

- Richard S. Sutton, David McAllester, Satinder Singh, and Yishay Mansour. Policy Gradient Methods for Reinforcement Learning with Function Approximation. In *Advances in Neural Information Processing Systems*, volume 12, pages 1057–1063, 2000.
- Aviv Tamar, Shie Mannor, and Huan Xu. Scaling up robust mdps using function approximation. *ICML*, 32: 1401–1415, 2014.
- Aviv Tamar, Yinlam Chow, Mohammad Ghavamzadeh, and Shie Mannor. Policy gradient for coherent risk measures. In *Advances in Neural Information Processing Systems*, pages 1468–1476, 2015.
- Christopher J. C. H. Watkins and Peter Dayan. Q-learning. *Machine Learning*, 8:279–292, 1992.
- Marco Wiering and Martijn van Otterlo. *Reinforcement Learning: State-of-the-Art*. 12. Springer-Verlag Berlin Heidelberg, 2012.
- Huan Xu and Shie Mannor. Distributionally Robust Markov Decision Processes. *Mathematics of Operations Research*, 37(2):288–300, 2012.
- Pengqian Yu and Huan Xu. Distributionally Robust Counterpart in Markov Decision Processes. *IEEE Transactions on Automatic Control*, 61(9):2538 – 2543, 2016.