

Reinforcement Learning with Parameterized Actions

Warwick Masson and Pravesh Ranchod

School of Computer Science and Applied Mathematics
University of Witwatersrand
Johannesburg, South Africa
warwick.masson@students.wits.ac.za
pravesh.ranchod@wits.ac.za

George Konidaris

Department of Computer Science
Duke University
Durham, North Carolina 27708
gdk@cs.duke.edu

Abstract

We introduce a model-free algorithm for learning in Markov decision processes with parameterized actions—discrete actions with continuous parameters. At each step the agent must select both which action to use and which parameters to use with that action. We introduce the Q-PAMDP algorithm for learning in these domains, show that it converges to a local optimum, and compare it to direct policy search in the goal-scoring and Platform domains.

1 Introduction

Reinforcement learning agents typically have either a discrete or a continuous action space (Sutton and Barto 1998). With a discrete action space, the agent decides which distinct action to perform from a finite action set. With a continuous action space, actions are expressed as a single real-valued vector. If we use a continuous action space, we lose the ability to consider differences in kind: all actions must be expressible as a single vector. If we use only discrete actions, we lose the ability to finely tune action selection based on the current state.

A parameterized action is a discrete action parameterized by a real-valued vector. Modeling actions this way introduces structure into the action space by treating different kinds of continuous actions as distinct. At each step an agent must choose both which action to use and what parameters to execute it with. For example, consider a soccer playing robot which can kick, pass, or run. We can associate a continuous parameter vector to each of these actions: we can kick the ball to a given target position with a given force, pass to a specific position, and run with a given velocity. Each of these actions is parameterized in its own way. Parameterized action Markov decision processes (PAMDPs) model situations where we have distinct actions that require parameters to adjust the action to different situations, or where there are multiple mutually incompatible continuous actions.

We focus on how to learn an action-selection policy given pre-defined parameterized actions. We introduce the Q-PAMDP algorithm, which alternates learning action-selection and parameter-selection policies and compare it to

a direct policy search method. We show that with appropriate update rules Q-PAMDP converges to a local optimum. These methods are compared empirically in the goal and Platform domains. We found that Q-PAMDP out-performed direct policy search and fixed parameter SARSA.

2 Background

A Markov decision process (MDP) is a tuple $\langle S, A, P, R, \gamma \rangle$, where S is a set of states, A is a set of actions, $P(s, a, s')$ is the probability of transitioning to state s' from state s after taking action a , $R(s, a, r)$ is the probability of receiving reward r for taking action a in state s , and γ is a discount factor (Sutton and Barto 1998). We wish to find a policy, $\pi(a|s)$, which selects an action for each state so as to maximize the expected sum of discounted rewards (the return).

The value function $V^\pi(s)$ is defined as the expected discounted return achieved by policy π starting at state s

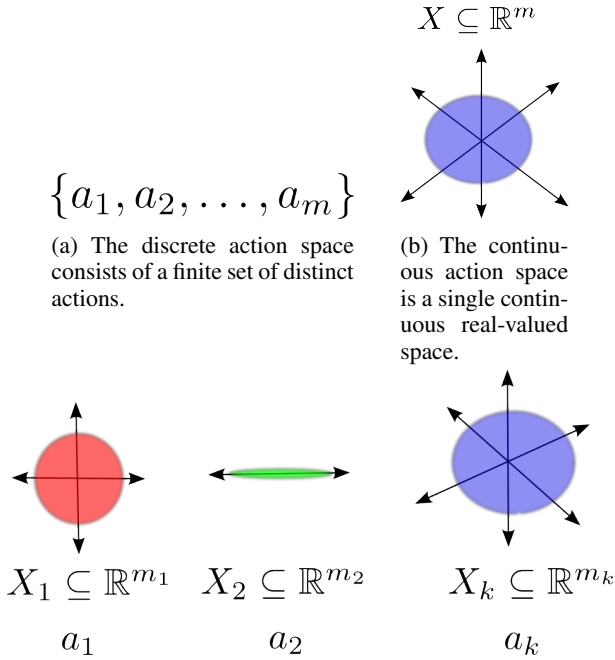
$$V^\pi(s) = \mathbb{E}_\pi \left[\sum_{t=0}^{\infty} \gamma^t r_t \right].$$

Similarly, the action-value function is given by

$$Q^\pi(s, a) = \mathbb{E}_\pi [r_0 + \gamma V^\pi(s')],$$

as the expected return obtained by taking action a in state s , and then following policy π thereafter. While using the value function in control requires a model, we would prefer to do so without needing such a model. We can approach this problem by learning Q , which allows us to directly select the action which maximizes $Q^\pi(s, a)$. We can learn Q for an optimal policy using a method such as Q-learning (Watkins and Dayan 1992). In domains with a continuous state space, we can represent $Q(s, a)$ using parametric function approximation with a set of parameters ω and learn this with algorithms such as gradient descent SARSA(λ) (Sutton and Barto 1998).

For problems with a continuous action space ($A \subseteq \mathbb{R}^m$), selecting the optimal action with respect to $Q(s, a)$ is non-trivial, as it requires finding a global maximum for a function in a continuous space. We can avoid this problem using a policy search algorithm, where a class of policies parameterized by a set of parameters θ is given, which transforms the problem into one of direct optimization over θ for an objective function $J(\theta)$. Several policy search approaches



(c) The parameterized action space has multiple discrete actions, each of which has a continuous parameter space.

Figure 1: Three types of action spaces: discrete, continuous, and parameterized.

exist, including policy gradient methods, entropy-based approaches, path integral approaches, and sample-based approaches (Deisenroth, Neumann, and Peters 2013).

Parameterized Tasks

A parameterized task is a problem defined by a task parameter vector τ given at the beginning of each episode. These parameters are fixed throughout an episode, and the goal is to learn a task dependent policy. Kober *et al.* (2012) developed algorithms to adjust motor primitives to different task parameters. They apply this to learn table-tennis and darts with different starting positions and targets. Da Silva *et al.* (2012) introduced the idea of a parameterized skill as a task dependent parameterized policy. They sample a set of tasks, learn their associated parameters, and determine a mapping from task to policy parameters. Deisenroth *et al.* (2014) applied a model-based method to learn a task dependent parameterized policy. This is used to learn task dependent policies for ball-hitting task, and for solving a block manipulation problem. Parameterized tasks can be used as parameterized actions. For example, if we learn a parameterized task for kicking a ball to position τ , this could be used as a parameterized action $\text{kick-to}(\tau)$.

3 Parameterized Action MDPs

We consider MDPs where the state space is continuous ($S \subseteq \mathbb{R}^n$) and the actions are parameterized: there is a finite set of discrete actions $A_d = \{a_1, a_2, \dots, a_k\}$, and each $a \in A_d$ has a set of continuous parameters $X_a \subseteq \mathbb{R}^{m_a}$. An action

is a tuple (a, x) where a is a discrete action and x are the parameters for that action. The action space is then given by

$$A = \bigcup_{a \in A_d} \{(a, x) \mid x \in X_a\},$$

which is the union of each discrete action with all possible parameters for that action. We refer to such MDPs as parameterized action MDPs (PAMDPs). Figure 1 depicts the different action spaces.

We apply a two-tiered approach for action selection: first selecting the parameterized action, then selecting the parameters for that action. The discrete-action policy is denoted $\pi^d(a|s)$. To select the parameters for the action, we define the action-parameter policy for each action a as $\pi^a(x|s)$. The policy is then given by

$$\pi(a, x|s) = \pi^d(a|s)\pi^a(x|s).$$

In other words, to select a complete action (a, x) , we sample a discrete action a from $\pi^d(a|s)$ and then sample a parameter x from $\pi^a(x|s)$. The action policy is defined by the parameters ω and is denoted by $\pi_\omega^d(a|s)$. The action-parameter policy for action a is determined by a set of parameters θ_a , and is denoted $\pi_{\theta_a}^a(x|s)$. The set of these parameters is given by $\theta = [\theta_{a_1}, \dots, \theta_{a_k}]$.

The first approach we consider is direct policy search. We use a direct policy search method to optimize the objective function.

$$J(\theta, \omega) = \mathbb{E}_{s_0 \sim D}[V^{\pi^\theta}(s_0)].$$

with respect to (θ, ω) , where s_0 is a state sampled according to the state distribution D . J is the expected return for a given policy starting at an initial state.

Our second approach is to alternate updating the parameter-policy and learning an action-value function for the discrete actions. For any PAMDP $M = \langle S, A, P, R, \gamma \rangle$ with a fixed parameter-policy π_θ^a , there exists a corresponding discrete action MDP, $M_\theta = \langle S, A_d, P_\theta, R_\theta, \gamma \rangle$, where

$$P_\theta(s'|s, a) = \int_{x \in X_a} \pi_\theta^a(x|s)P(s'|s, a, x)dx,$$

$$R_\theta(r|s, a) = \int_{x \in X_a} \pi_\theta^a(x|s)R(r|s, a, x)dx.$$

We represent the action-value function for M_θ using function approximation with parameters ω . For M_θ , there exists an optimal set of representation weights ω_θ^* which maximizes $J(\theta, \omega)$ with respect to ω . Let

$$W(\theta) = \arg \max_{\omega} J(\theta, \omega) = \omega_\theta^*.$$

We can learn $W(\theta)$ for a fixed θ using a Q-learning algorithm. Finally, we define for fixed ω ,

$$J_\omega(\theta) = J(\theta, \omega),$$

$$H(\theta) = J(\theta, W(\theta)).$$

$H(\theta)$ is the performance of the best discrete policy for a fixed θ .

Algorithm 1 Q-PAMDP(k)

Input:Initial parameters θ_0, ω_0

Parameter update method P-UPDATE

Q-learning algorithm Q-LEARN

Algorithm: $\omega \leftarrow \text{Q-LEARN}^{(\infty)}(M_\theta, \omega_0)$ **repeat** $\theta \leftarrow \text{P-UPDATE}^{(k)}(J_\omega, \theta)$ $\omega \leftarrow \text{Q-LEARN}^{(\infty)}(M_\theta, \omega)$ **until** θ converges

Algorithm 1 describes a method for alternating updating θ and ω . The algorithm uses two input methods: P-UPDATE and Q-LEARN and a positive integer parameter k , which determines the number of updates to θ for each iteration. P-UPDATE(f, θ) should be a policy search method that optimizes θ with respect to objective function f . Q-LEARN can be any algorithm for Q-learning with function approximation. We consider two main cases of the Q-PAMDP algorithm: Q-PAMDP(1) and Q-PAMDP(∞).

Q-PAMDP(1) performs a single update of θ and then re-learns ω to convergence. If at each step we only update θ once, and then update ω until convergence, we can optimize θ with respect to H . In the next section we show that if we can find a local optimum θ with respect to H , then we have found a local optimum with respect to J .

4 Theoretical Results

We now show that Q-PAMDP(1) converges to a local or global optimum with mild assumptions. We assume that iterating P-UPDATE converges to some θ^* with respect to a given objective function f . As the P-UPDATE step is a design choice, it can be selected with the appropriate convergence property. Q-PAMDP(1) is equivalent to the sequence

$$\begin{aligned}\omega_{t+1} &= W(\theta_t) \\ \theta_{t+1} &= \text{P-UPDATE}(J_{\omega_{t+1}}, \theta_t),\end{aligned}$$

if Q-LEARN converges to $W(\theta)$ for each given θ .

Theorem 4.1 (Convergence to a Local Optimum). *For any θ_0 , if the sequence*

$$\theta_{t+1} = \text{P-UPDATE}(H, \theta_t), \quad (1)$$

converges to a local optimum with respect to H , then Q-PAMDP(1) converges to a local optimum with respect to J .

Proof. By definition of the sequence above $\omega_t = W(\theta_t)$, so it follows that

$$J_{\omega_t} = J(\theta, W(\theta)) = H(\theta).$$

In other words, the objective function J equals H if $\omega = W(\theta)$. Therefore, we can replace J with H in our update for θ , to obtain the update rule

$$\theta_{t+1} = \text{P-UPDATE}(H, \theta_t).$$

Therefore by equation 1 the sequence θ_t converges to a local optimum θ^* with respect to H . Let $\omega^* = W(\theta^*)$. As θ^* is a local optimum with respect to H , by definition there exists $\epsilon > 0$, s.t.

$$\|\theta^* - \theta\|_2 < \epsilon \implies H(\theta) \leq H(\theta^*).$$

Therefore for any ω ,

$$\begin{aligned}\left\| \begin{pmatrix} \theta^* \\ \omega^* \end{pmatrix} - \begin{pmatrix} \theta \\ \omega \end{pmatrix} \right\|_2 < \epsilon &\implies \|\theta^* - \theta\|_2 < \epsilon \\ &\implies H(\theta) \leq H(\theta^*) \\ &\implies J(\theta, \omega) \leq J(\theta^*, \omega^*).\end{aligned}$$

Therefore (θ^*, ω^*) is a local optimum with respect to J . \square

In summary, if we can locally optimize θ , and $\omega = W(\theta)$ at each step, then we will find a local optimum for $J(\theta, \omega)$. The conditions for the previous theorem can be met by assuming that P-UPDATE is a local optimization method such as a gradient based policy search. A similar argument shows that if the sequence θ_t converges to a global optimum with respect to H , then Q-PAMDP(1) converges to a global optimum (θ^*, ω^*) .

One problem is that at each step we must re-learn $W(\theta)$ for the updated value of θ . We now show that if updates to θ are bounded and $W(\theta)$ is a continuous function, then the required updates to ω will also be bounded. Intuitively, we are supposing that a small update to θ results in a small change in the weights specifying which discrete action to choose. The assumption that $W(\theta)$ is continuous is strong, and may not be satisfied by all PAMDPs. It is not necessary for the operation of Q-PAMDP(1), but when it is satisfied we do not need to completely re-learn ω after each update to θ . We show that by selecting an appropriate α we can shrink the differences in ω as desired.

Theorem 4.2 (Bounded Updates to ω). *If W is continuous with respect to θ , and updates to θ are of the form*

$$\theta_{t+1} = \theta_t + \alpha_t \text{P-UPDATE}(\theta_t, \omega_t),$$

with the norm of each P-UPDATE bounded by

$$0 < \|\text{P-UPDATE}(\theta_t, \omega_t)\|_2 < \delta,$$

for some $\delta > 0$, then for any difference in $\omega \in \epsilon > 0$, there is an initial update rate $\alpha_0 > 0$ such that

$$\alpha_t < \alpha_0 \implies \|\omega_{t+1} - \omega_t\|_2 < \epsilon.$$

Proof. Let $\epsilon > 0$ and

$$\alpha_0 = \frac{\delta}{\|\text{P-UPDATE}(\theta_t, \omega_t)\|_2}.$$

As $\alpha_t < \alpha_0$, it follows that

$$\begin{aligned}\delta &> \alpha_t \|\text{P-UPDATE}(\theta_t, \omega_t)\|_2 \\ &= \|\alpha_t \text{P-UPDATE}(\theta_t, \omega_t)\|_2 \\ &= \|\theta_{t+1} - \theta_t\|_2.\end{aligned}$$

So we have

$$\|\theta_{t+1} - \theta_t\|_2 < \delta.$$

As W is continuous, this means that

$$\|W(\theta_{t+1}) - W(\theta_t)\|_2 = \|\omega_{t+1} - \omega_t\|_2 < \epsilon.$$

\square

In other words, if our update to θ is bounded and W is continuous, we can always adjust the learning rate α so that the difference between ω_t and ω_{t+1} is bounded.

With Q-PAMDP(1) we would P-UPDATE to optimize $H(\theta)$. One logical choice would be to use a gradient update. The next theorem shows that gradient of H is equal to the gradient of J if $\omega = W(\theta)$. This is useful as we can apply existing gradient-based policy search methods to compute the gradient of J with respect to θ . The proof follows from the fact that we are at a global optimum of J with respect to ω , and so the gradient $\nabla_{\omega} J$ is zero. This theorem requires that W is differentiable (and therefore also continuous).

Theorem 4.3 (Gradient of $H(\theta)$). *If $J(\theta, \omega)$ is differentiable with respect to θ and ω and $W(\theta)$ is differentiable with respect to θ , then the gradient of H is given by $\nabla_{\theta} H(\theta) = \nabla_{\theta} J(\theta, \omega^*)$, where $\omega^* = W(\theta)$.*

Proof. If $\theta \in \mathbb{R}^n$ and $\omega \in \mathbb{R}^m$, then we can compute the gradient of H by the chain rule:

$$\begin{aligned} \frac{\partial H(\theta)}{\partial \theta_i} &= \frac{\partial J(\theta, W(\theta))}{\partial \theta_i} \\ &= \sum_{j=1}^n \frac{\partial J(\theta, \omega^*)}{\partial \theta_j} \frac{\partial \theta_j}{\partial \theta_i} + \sum_{k=1}^m \frac{\partial J(\theta, \omega^*)}{\partial \omega_k^*} \frac{\partial \omega_k^*}{\partial \theta_i} \\ &= \frac{\partial J(\theta, \omega^*)}{\partial \theta_i} + \sum_{k=1}^m \frac{\partial J(\theta, \omega^*)}{\partial \omega_k^*} \frac{\partial \omega_k^*}{\partial \theta_i}, \end{aligned}$$

where $\omega^* = W(\theta)$. Note that as by definitions of W ,

$$\omega^* = W(\theta) = \arg \max_{\omega} J(\theta, \omega),$$

we have that the gradient of J with respect to ω is zero $\partial J(\theta, \omega^*) / \partial \omega_k^* = 0$, as ω is a global maximum with respect to J for fixed θ . Therefore, we have that

$$\nabla_{\theta} H(\theta) = \nabla_{\theta} J(\theta, \omega^*).$$

□

To summarize, if $W(\theta)$ is continuous and P-UPDATE converges to a global or local optimum, then Q-PAMDP(1) will converge to a global or local optimum, respectively, and the Q-LEARN step will be bounded if the update rate of the P-UPDATE step is bounded. As such, if P-UPDATE is a policy gradient update step then Q-PAMDP by Theorem 4.1 will converge to a local optimum and by Theorem 4.4 the Q-LEARN step will require a fixed number of updates. This policy gradient step can use the gradient of J with respect to θ .

With Q-PAMDP(∞) each step performs a full optimization on θ and then a full optimization of ω . The θ step would optimize $J(\theta, \omega)$, not $H(\theta)$, as we do update ω while we update θ . Q-PAMDP(∞) has the disadvantage of requiring global convergence properties for the P-UPDATE method.

Theorem 4.4 (Local Convergence of Q-PAMDP(∞)). *If at each step of Q-PAMDP(∞) for some bounded set Θ :*

$$\theta_{t+1} = \arg \max_{\theta \in \Theta} J(\theta, \omega_t),$$

$$\omega_{t+1} = W(\theta_{t+1}),$$

then Q-PAMDP(∞) converges to a local optimum.

Proof. By definition of W , $\omega_{t+1} = \arg \max_{\omega} J(\theta_{t+1}, \omega)$. Therefore this algorithm takes the form of direct alternating optimization. As such, it converges to a local optimum (Bezdek and Hathaway 2002). □

Q-PAMDP(∞) has weaker convergence properties than Q-PAMDP(1), as it requires a globally convergent P-UPDATE. However, it has the potential to bypass nearby local optima (Bezdek and Hathaway 2002).

5 Experiments

We first consider a simplified robot soccer problem (Kitano et al. 1997) where a single striker attempts to score a goal against a keeper. Each episode starts with the player at a random position along the bottom bound of the field. The player starts with the ball in possession, and the keeper is positioned between the ball and the goal. The game takes place in a 2D environment where the player and the keeper have a position, velocity and orientation and the ball has a position and velocity, resulting in 14 continuous state variables.

An episode ends when the keeper possesses the ball, the player scores a goal, or the ball leaves the field. The reward for an action is 0 for non-terminal state, 50 for a terminal goal state, and $-d$ for a terminal non-goal state, where d is the distance of the ball to the goal. The player has two parameterized actions: kick-to(x, y), which kicks to ball towards position (x, y); and shoot-goal(h), which shoots the ball towards a position h along the goal line. Noise is added to each action. If the player is not in possession of the ball, it moves towards it. The keeper has a fixed policy: it moves towards the ball, and if the player shoots at the goal, the keeper moves to intercept the ball.

To score a goal, the player must shoot around the keeper. This means that at some positions it must shoot left past the keeper, and at others to the right past the keeper. However at no point should it shoot at the keeper, so an optimal policy is discontinuous. We split the action into two parameterized actions: shoot-goal-left, shoot-goal-right. This allows us to use a simple action selection policy instead of complex continuous action policy. This policy would be difficult to represent in a purely continuous action space, but is simpler in a parameterized action setting.

We represent the action-value function for the discrete action a using linear function approximation with Fourier basis features (Konidaris, Osentoski, and Thomas 2011). As we have 14 state variables, we must be selective in which basis functions to use. We only use basis functions with two non-zero elements and exclude all velocity state variables. We use the soft-max discrete action policy (Sutton and Barto 1998). We represent the action-parameter policy π_{θ}^a as a normal distribution around a weighted sum of features $\pi_{\theta}^a(x|s) = \mathcal{N}(\theta_a^T \psi_a(s), \Sigma)$, where θ_a is a matrix of weights, and $\psi_a(s)$ gives the features for state s , and Σ is a fixed covariance matrix. We use specialized features for each action. For the shoot-goal actions we are using a simple linear basis $(1, g)$, where g is the projection of the keeper onto the goal line. For kick-to we use linear features $(1, bx, by, bx^2, by^2, (bx - kx) / \|b - k\|_2, (by -$

$ky)/\|b - k\|_2$, where (bx, by) is the position of the ball and (kx, ky) is the position of the keeper.

For the direct policy search approach, we use the episodic natural actor critic (eNAC) algorithm (Peters and Schaal 2008), computing the gradient of $J(\omega, \theta)$ with respect to (ω, θ) . For the Q-PAMDP approach we use the gradient-descent Sarsa(λ) algorithm for Q-learning, and the eNAC algorithm for policy search. At each step we perform one eNAC update based on 50 episodes and then refit Q_ω using 50 gradient descent Sarsa(λ) episodes.

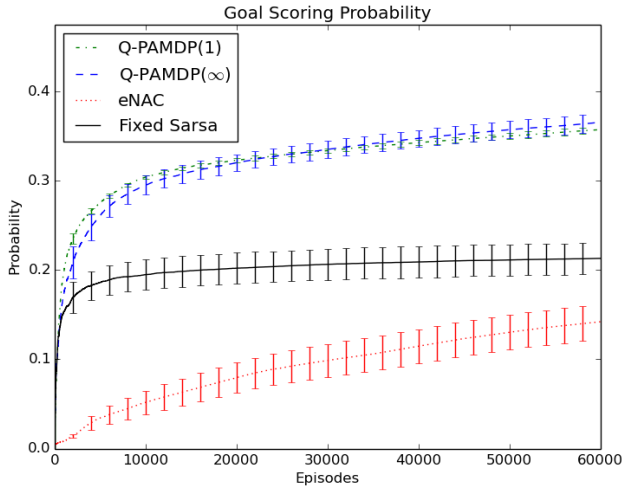


Figure 2: Average goal scoring probability, averaged over 20 runs for Q-PAMDP(1), Q-PAMDP(∞), fixed parameter Sarsa, and eNAC in the goal domain. Intervals show standard error.

Return is directly correlated with goal scoring probability, so their graphs are close to identical. As it is easier to interpret, we plot goal scoring probability in figure 2. We can see that direct eNAC is outperformed by Q-PAMDP(1) and Q-PAMDP(∞). This is likely due to the difficulty of optimizing the action selection parameters directly, rather than with Q-learning.

For both methods, the goal probability is greatly increased: while the initial policy rarely scores a goal, both Q-PAMDP(1) and Q-PAMDP(∞) increase the probability of a goal to roughly 35%. Direct eNAC converged to a local maxima of 15%. Finally, we include the performance of SARSA(λ) where the action parameters are fixed at the initial θ_0 . This achieves roughly 20% scoring probability. Both Q-PAMDP(1) and Q-PAMDP(∞) strongly outperform fixed parameter SARSA, but eNAC does not. Figure 3 depicts a single episode using a converged Q-PAMDP(1) policy—the player draws the keeper out and strikes when the goal is open.

Next we consider the Platform domain, where the agent starts on a platform and must reach a goal while avoiding enemies. If the agent reaches the goal platform, touches an enemy, or falls into a gap between platforms, the episode ends. This domain is depicted in figure 4. The reward for a step is the change in x value for that step, divided by the to-

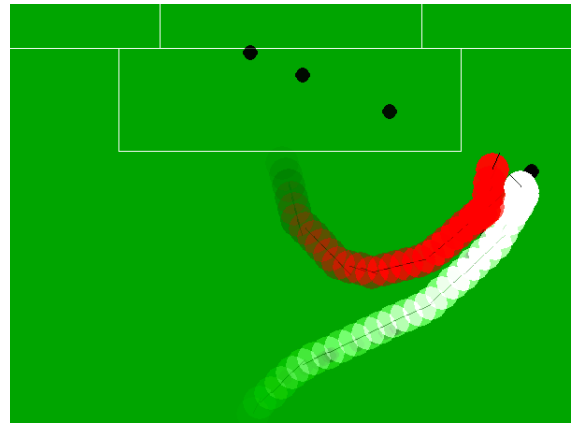


Figure 3: A robot soccer goal episode using a converged Q-PAMDP(1) policy. The player runs to one side, then shoots immediately upon overtaking the keeper.

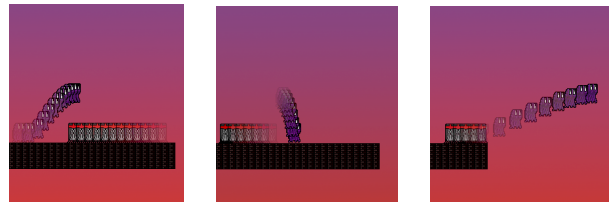


Figure 4: A screenshot from the Platform domain. The player hops over an enemy, and then leaps over a gap.

tal length of all the platforms and gaps. The agent has two primitive actions: run or jump, which continue for a fixed period or until the agent lands again respectively. There are two different kinds of jumps: a high jump to get over enemies, and a long jump to get over gaps between platforms. The domain therefore has three parameterized actions: run(dx), hop(dx), and leap(dx). The agent only takes actions while on the ground, and enemies only move when the agent is on their platform. The state space consists of four variables $(x, \dot{x}, ex, \dot{ex})$, representing the agent position, agent speed, enemy position, and enemy speed respectively. For learning Q_ω , as in the previous domain, we use linear function approximation with the Fourier basis. We apply a softmax discrete action policy based on Q_ω , and a Gaussian parameter policy based on scaled parameter features $\psi_a(s)$.

Figure 5 shows the performance of eNAC, Q-PAMDP(1), Q-PAMDP(∞), and SARSA with fixed parameters. Both Q-PAMDP(1) and Q-PAMDP(∞) outperformed the fixed parameter SARSA method, reaching on average 50% and 65% of the total distance respectively. We suggest that Q-PAMDP(∞) outperforms Q-PAMDP(1) due to the nature of the Platform domain. Q-PAMDP(1) is best suited to domains with smooth changes in the action-value function with respect to changes in the parameter-policy. With the Platform domain, our initial policy is unable to make the first jump without modification. When the policy can reach the second platform, we need to drastically change the action-value function to account for this platform. Therefore, Q-

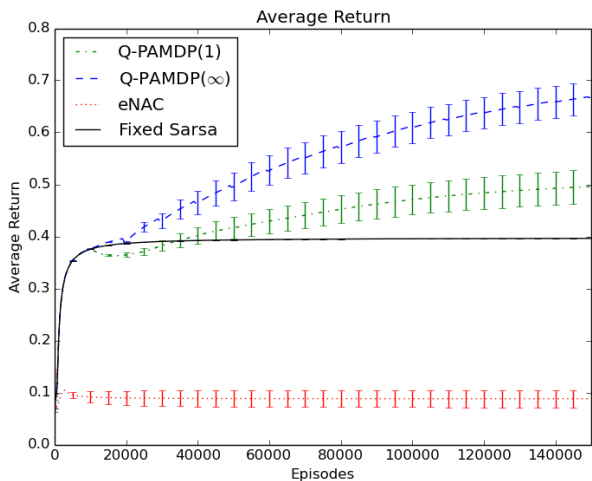


Figure 5: Average percentage distance covered, averaged over 20 runs for Q-PAMDP(1), Q-PAMDP(∞), and eNAC in the Platform domain. Intervals show standard error.

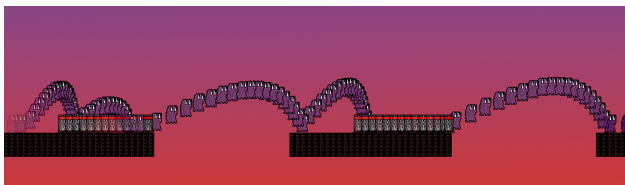


Figure 6: A successful episode of the Platform domain. The agent hops over the enemies, leaps over the gaps, and reaches the last platform.

PAMDP(1) may be poorly suited to this domain as the small change in parameters that occurs between failing to making the jump and actually making it results in a large change in the action-value function. This is better than the fixed SARSA baseline of 40%, and much better than direct optimization using eNAC which reached 10%. Figure 6 shows a successfully completed episode of the Platform domain.

6 Related Work

Hauskrecht *et al.* (2004) introduced an algorithm for solving factored MDPs with a hybrid discrete-continuous action space. However, their formalism has an action space with a mixed set of discrete and continuous components, whereas our domain has distinct actions with a different number of continuous components for each action. Furthermore, they assume the domain has a compact factored representation, and only consider planning.

Rachelson (2009) encountered parameterized actions in the form of an action to wait for a given period of time in his research on time dependent, continuous time MDPs (TMDPs). He developed XMDPs, which are TMDPs with a parameterized action space (Rachelson 2009). He developed a Bellman operator for this domain, and in a later paper mentions that the $TiMDP_{poly}$ algorithm can work with parameterized actions, although this specifically refers to the

parameterized wait action (Rachelson, Fabiani, and Garcia 2009). This research also takes a planning perspective, and only considers a time dependent domain. Additionally, the size of the parameter space for the parameterized actions is the same for all actions.

Hoey *et al.* (2013) considered mixed discrete-continuous actions in their work on Bayesian affect control theory. To approach this problem they use a form of POMCP, a Monte Carlo sampling algorithm, using domain specific adjustments to compute the continuous action components (Silver and Veness 2010). They note that the discrete and continuous components of the action space reflect different control aspects: the discrete control provides the “what”, while the continuous control describes the “how” (Hoey, Schroder, and Althothali 2013).

In their research on symbolic dynamic programming (SDP) algorithms, Zamani *et al.* (2012) considered domains with a set of discrete parameterized actions. Each of these actions has a different parameter space. Symbolic dynamic programming is a form of planning for relational or first-order MDPs, where the MDP has a set of logical relationships defining its dynamics and reward function. Their algorithms represent the value function as an extended algebraic decision diagram (XADD), and is limited to MDPs with pre-defined logical relations.

A hierarchical MDP is an MDP where each action has subtasks. A subtask is itself an MDP with its own states and actions which may have their own subtasks. Hierarchical MDPs are well-suited for representing parameterized actions as we could consider selecting the parameters for a discrete action as a subtask. MAXQ is a method for value function decomposition of hierarchical MDPs (Dietterich 2000). One possibility is to use MAXQ for learning the action-values in a parameterized action problem.

7 Conclusion

The PAMDP formalism models reinforcement learning domains with parameterized actions. Parameterized actions give us the adaptability of continuous domains and to use distinct kinds of actions. They also allow for simple representation of discontinuous policies without complex parameterizations. We have presented three approaches for model-free learning in PAMDPs: direct optimization and two variants of the Q-PAMDP algorithm. We have shown that Q-PAMDP(1), with an appropriate P-UPDATE method, converges to a local or global optimum. Q-PAMDP(∞) with a global optimization step converges to a local optimum.

We have examined the performance of these approaches in the goal scoring domain and the Platformer domain. The robot soccer goal domain models the situation where a striker must out-manuever a keeper to score a goal. Of these, Q-PAMDP(1) and Q-PAMDP(∞) outperformed eNAC and fixed parameter SARSA. Q-PAMDP(1) and Q-PAMDP(∞) performed similarly well in terms of goal scoring, learning policies that score goals roughly 35% of the time. In the Platform domain we found that both Q-PAMDP(1) and Q-PAMDP(∞) outperformed eNAC and fixed SARSA.

References

- Bezdek, J., and Hathaway, R. 2002. Some notes on alternating optimization. In *Advances in Soft Computing*. Springer. 288–300.
- da Silva, B.; Konidaris, G.; and Barto, A. 2012. Learning parameterized skills. In *Proceedings of the Twenty-Ninth International Conference on Machine Learning*, 1679–1686.
- Deisenroth, M.; Englert, P.; Peters, J.; and Fox, D. 2014. Multi-task policy search for robotics. In *Proceedings of the Fourth International Conference on Robotics and Automation*, 3876–3881.
- Deisenroth, M.; Neumann, G.; and Peters, J. 2013. *A Survey on Policy Search for Robotics*. Number 12. Now Publishers.
- Dietterich, T. 2000. Hierarchical reinforcement learning with the MAXQ value function decomposition. *Journal of Artificial Intelligence Research* 13:227–303.
- Guestrin, C.; Hauskrecht, M.; and Kveton, B. 2004. Solving factored MDPs with continuous and discrete variables. In *Proceedings of the Twentieth Conference on Uncertainty in Artificial Intelligence*, 235–242.
- Hoey, J.; Schroder, T.; and Alhothali, A. 2013. Bayesian affect control theory. In *Proceedings of the Fifth International Conference on Affective Computing and Intelligent Interaction*, 166–172. IEEE.
- Kitano, H.; Asada, M.; Kuniyoshi, Y.; Noda, I.; Osawa, E.; and Matsubara, H. 1997. Robocup: A challenge problem for AI. *AI Magazine* 18(1):73.
- Kober, J.; Wilhelm, A.; Oztop, E.; and Peters, J. 2012. Reinforcement learning to adjust parametrized motor primitives to new situations. *Autonomous Robots* 33(4):361–379.
- Konidaris, G.; Osentoski, S.; and Thomas, P. 2011. Value function approximation in reinforcement learning using the Fourier basis. In *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence*, 380–385.
- Peters, J., and Schaal, S. 2008. Natural actor-critic. *Neurocomputing* 71(7):1180–1190.
- Rachelson, E.; Fabiani, P.; and Garcia, F. 2009. TiMDP-poly: an improved method for solving time-dependent MDPs. In *Proceedings of the Twenty-First International Conference on Tools with Artificial Intelligence*, 796–799. IEEE.
- Rachelson, E. 2009. *Temporal Markov Decision Problems: Formalization and Resolution*. Ph.D. Dissertation, University of Toulouse, France.
- Silver, D., and Veness, J. 2010. Monte-Carlo planning in large POMDPs. In *Advances in Neural Information Processing Systems*, volume 23, 2164–2172.
- Sutton, R., and Barto, A. 1998. *Introduction to Reinforcement Learning*. Cambridge, MA, USA: MIT Press.
- Watkins, C., and Dayan, P. 1992. Q-learning. *Machine learning* 8(3-4):279–292.
- Zamani, Z.; Sanner, S.; and Fang, C. 2012. Symbolic dynamic programming for continuous state and action MDPs. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence*.