

Region Enhanced Neural Q-learning for Solving Model-based POMDPs

Marco A. Wiering, *IEEE member*, and Thijs Kooi

Abstract—To get a robot to perform tasks autonomously, the robot has to plan its behavior and make decisions based on the input it receives. Unfortunately, contemporary robot sensors and actuators are subject to noise, rendering optimal decision making a stochastic process. To model this process, partially observable Markov decision processes (POMDPs) can be applied. In this paper we introduce the RENQ algorithm, a new POMDP algorithm that combines neural networks for estimating Q-values with the construction of a spatial pyramid over the state space. RENQ essentially uses region-based belief vectors together with state-based belief vectors, and these are used as inputs to the neural network trained with Q-learning. We compare RENQ to Qmdp and Perseus, two state-of-the-art algorithms for approximately solving model-based POMDPs. The results on three different maze navigation tasks indicate that RENQ outperforms Perseus on all problems and Qmdp if the problem becomes larger.

I. INTRODUCTION

In robotics, a major goal is getting a robot to learn to perform a task autonomously. This task can involve getting a robot from a start to a goal position. A possible approach to this problem is to use *reinforcement learning* (RL) [23]. Reinforcement learning originated from early work in cybernetics, statistics, psychology and neuroscience, but lately has received a lot of attention from the Artificial Intelligence (AI) and machine learning disciplines [7]. It can be seen as a form of machine learning, but is different from supervised learning methods in the sense that the agent does not learn from correct input-output examples, provided by an external supervisor, but has to learn from feedback given by the environment. The feedback the agent receives is typically represented as a numerical value, where a positive reward is given for the display of a desired behavior and a negative reward for an undesired action. The robot's task is to develop a model of what action to take in a given state, thereby maximizing its long term reward.

Unfortunately the robot's actuators do not always act according to the instructions they have been given. When it has to move right it sometimes moves left or bumps into a wall and stays in the same place. This uncertainty in transitions can be modeled using a *Markov decision process* (MDP). On top of the uncertainty in the robot's actuators, there is noise in its sensor readings as well. This partial observability of the world can be captured in a generalization of an MDP, called a *partially observable Markov decision process* (POMDP).

Marco Wiering and Thijs Kooi are with the Department of Artificial Intelligence, University of Groningen, the Netherlands. (email: {M.Wiering,T.Kooi}@ai.rug.nl).

Although robotic navigation is one application where a POMDP needs to be solved, the method is widely applicable to other problems. Givon and Grosfeld-Nir [5] used the POMDP model for computing optimal termination times of TV shows. An application in another field is provided by Hoey et al. [6], who use a POMDP model to handle the uncertainty in observations from a monitor, assisting people with dementia washing their hands.

In this paper we will present RENQ, a novel approach for solving model-based POMDPs. RENQ uses a neural network in combination with Q-learning [24], [25], where the belief-state is given as input to the neural network. RENQ enhances the state-based belief vector input of the network by constructing a spatial pyramid over the state space [9], a method derived from machine vision. At every level of the pyramid, the average belief of a subset of the state space is computed and the enhanced belief state is presented to the neural network. We compare the RENQ algorithm to Qmdp [11], a method known to be fast in handling large state spaces and Perseus [21], an efficient state-of-the-art point-based value iteration algorithm. We test the algorithms on three different maze navigation tasks and show how RENQ outperforms the other two methods.

Outline. This paper is divided into 6 sections. In section II, we will discuss the basic framework of Markov decision processes, followed by a brief description of value iteration and Q-learning, two techniques for solving MDPs and reinforcement learning problems. In section III the POMDP model will be presented along with two algorithms for handling POMDPs. Next we discuss RENQ, the new method based on Q-learning and neural networks, combined with a spatial pyramid approach. Section V will cover the experimental setup and results acquired with the three POMDP algorithms. A conclusion and discussion will be presented in section VI.

II. MARKOV DECISION PROCESS

In this section, we will start by giving a formal definition of the MDP model, followed by a description of value iteration. Then, we describe the Q-learning algorithm which is part of the RENQ algorithm.

A. Formal Description

An MDP is characterized by:

- a finite set of states $s \in \mathcal{S}$
- a finite set of actions $a \in \mathcal{A}$
- a transition function $T(s, a, s')$, specifying the probability of ending in state s' after taking action a in state s

- a reward function $R(s, a)$, providing the scalar reward the agent will receive for executing action a in state s

The MDP framework assumes a full map of the environment is known to the agent and treats time and sets \mathcal{S} and \mathcal{A} as discrete. For reinforcement learning algorithms, the MDP does not have to be known, but we assume a model of the POMDP is given in our experiments. Otherwise, the problem would become even harder.

The Markov property entails the fact that the state of the environment and the reward the agent receives at time $t + 1$ is stochastically determined by the state of the agent at time t and the action the agent takes. This is called a first order Markov process [19]:

$$P(s_t, r_t | s_0, a_0, \dots, s_{t-1}, a_{t-1}) = P(s_t, r_t | s_{t-1}, a_{t-1}) \quad (1)$$

The agent's task is to maximize its long term reward. Due to the stochasticity of the problem, a mapping is needed from states to actions. We call such a mapping a *policy* and denote it as $\pi(s)$. An optimal policy π^* maximizes the long term reward intake. In order to compute the optimal policy, the agent will assign a certain value for being in a state or performing some action in a state.

B. Value Functions

The return R_t of a state is defined as the cumulative reward the agent can expect to receive after reaching the given state at time step t . Mathematically, R_t is written as the sum over all rewards the agent receives at each time step, weighted by a discount factor γ , where $0 \leq \gamma < 1$:

$$R_t = r_{t+1} + \gamma r_{t+2} + \gamma^2 r_{t+3} + \dots = \sum_{k=0}^{\infty} \gamma^k r_{t+k+1} \quad (2)$$

Introducing a discount factor has two purposes: (1) it models the preference of the agent to immediate rewards as opposed to those received in the future, and (2) ensures the infinite sum is finite as long as $\gamma < 1$ and the rewards are bounded. When the discount factor is set close to 1, the agent will value future rewards greatly, whereas one close to 0 will make the agent focus on immediate rewards and value the future less.

The value of state s under policy π is defined as the expected discounted cumulative reward and is given by:

$$V^\pi(s) = E\left[\sum_{k=0}^{\infty} \gamma^k r_{t+k+1} | s_t = s\right] \quad (3)$$

In most situations it is desired to have knowledge of the value of an action in a certain state, we call this the Q-value, with $Q(s, a)$ providing the value of taking a in s , it is defined as:

$$Q^\pi(s, a) = E\left[\sum_{k=0}^{\infty} \gamma^k r_{t+k+1} | s_t = s, a_t = a\right] \quad (4)$$

Assuming the values of all successor states s' are known to the agent, Eq. (4) can be rewritten as the reward the agent receives plus the discounted value of s' , weighted by the probability of ending in s' , after taking action a in s :

$$Q^\pi(s, a) = \sum_{s'} T(s, a, s') [R(s, a) + \gamma V^\pi(s')] \quad (5)$$

This formula is a form of the Bellman equation named after Richard Bellman, who introduced it in 1957 [3]. With this function, we can iteratively update the value of all states, until it reaches a convergence criterion, resulting in an optimal state-value function $V^*(s)$, from which we can derive an optimal state action-value function $Q^*(s, a)$. Knowing the value of all states, the agent can select the action with the highest utility in every state, which will lead to an optimal policy. Value iteration is an algorithm that uses this concept.

C. Value Iteration

Value iteration is a (truncated) dynamic programming algorithm for computing optimal value functions and provides an exact solution for solving MDPs. The main idea behind this method is to compute the value of all $s \in \mathcal{S}$ iteratively, and to truncate the algorithm as soon as the difference in value of a state between two iterations: $\Delta = \max_{s \in \mathcal{S}} |V_i(s) - V_{i-1}(s)|$ drops below a threshold, where Δ is typically referred to as the Bellman residual. To approximate the value of a state, value iteration uses the Bellman equation in Eq. (5) as an update rule, see Algorithm 1.

Algorithm 1 Value iteration

```

Initialize  $V(s)$  and  $Q(s, a)$  arbitrarily
repeat
   $\Delta \leftarrow 0$ 
  for all  $s \in \mathcal{S}$  do
     $v \leftarrow V(s)$ 
    for all  $a \in \mathcal{A}$  do
       $Q(s, a) \leftarrow \sum_{s'} T(s, a, s') [R(s, a, s') + \gamma V(s')]$ 
    end for
     $V(s) \leftarrow \max_{a \in \mathcal{A}} Q(s, a)$ 
     $\Delta \leftarrow \max(\Delta, v - V(s))$ 
  end for
until  $\Delta < \theta$ 

```

Formally, the algorithm would need an infinite number of sweeps through the state space to converge to an optimal value function, but the optimal value can be approximated by aborting the algorithm if Δ is sufficiently small. A major drawback is that each iteration requires updating the value of every $s \in \mathcal{S}$, resulting in a computational complexity of $\mathcal{O}(|\mathcal{A}||\mathcal{S}|^2)$ per iteration. This is time consuming for problems with a large state space.

Once the algorithm is finished, the agent can use the values of state-action pairs to select the action with the best expected outcome:

$$\pi^*(s) = \arg \max_{a \in \mathcal{A}} Q^*(s, a) \quad (6)$$

D. Q-Learning

The introduction of Q-learning by Watkins in 1989 [24], [25], signified a great leap forward in the progress of the field of reinforcement learning. It is different from value iteration in the sense that it does not require an a-priori model of the environment and can therefore be used for more

applications. Furthermore, Q-learning can be combined with function approximators to solve large dimensional problems or problems involving continuous state spaces.

After each action taken, the agent evaluates the value of the action and uses this to update the current Q-value:

$$Q(s_t, a_t) \leftarrow Q(s_t, a_t) + \alpha \delta(s_t, a_t) \quad (7)$$

where α denotes the learning rate, to be decreased as the algorithm progresses, and $\delta(s_t, a_t)$ the TD-error, which is computed according to:

$$\delta(s_t, a_t) = r_{t+1} + \gamma \max_a Q(s_{t+1}, a) - Q(s_t, a_t) \quad (8)$$

It is known that for finite state and action spaces, Q-learning converges to the optimal $Q^*(s, a)$ as long as every state-action pair is visited infinitely often [25].

Q-learning needs to visit all state-action pairs repeatedly to get reliable estimates for their values. Therefore it requires an exploration policy. In this paper we use the *softmax* action selection rule [23], also known as Boltzmann exploration. Softmax uses a Gibbs or Boltzmann distribution for acquiring the probability of an action:

$$\pi(s, a) = \frac{e^{Q(s, a)/\tau}}{\sum_{a' \in \mathcal{A}} e^{Q(s, a')/\tau}} \quad (9)$$

Where τ denotes the temperature parameter.

Consider again the problem of a robot trying to navigate itself from a start to a goal position. In all real world situations, there is noise in its sensor readings as well; the world is partially observable to the agent. On top of this some sensor readings might seem similar, due to a similar looking environment or due to the distortion caused by the noise in its sensors. This phenomenon is known as *perceptual aliasing*. The uncertainty in observations can be incorporated into the MDP model. The acquired result is called a partially observable Markov decision process, or POMDP.

III. PARTIALLY OBSERVABLE MARKOV DECISION PROCESS

A POMDP is a generalization of an MDP and models not only the stochasticity in transitions, but also in observations, rendering the state of the agent partially observable. The POMDP framework consists of the same set of states $s \in \mathcal{S}$, actions $a \in \mathcal{A}$, transition function $T(s, a, s')$ and reward function $R(s, a)$. On top of this a POMDP consists of a set of observations $z \in \mathcal{Z}$ and an observation function $O(s, a, z)$, providing the probability of observing z in state s , after executing action a . Similar to the MDP model all sets \mathcal{S} , \mathcal{A} and \mathcal{Z} are assumed to be discrete, although work in continuous spaces has been done (e.g., [15]).

In an MDP the agent acts according to what seems to be the best possible action for a given state, but since the agent is no longer certain of its location it has to estimate its position based on the input it receives and its actions taken. A common approach to do this is by tracking a *belief state* [2].

A. Belief states

A belief state \vec{b}_t is a probability distribution over \mathcal{S} , to model the belief of the agent at time t . The set of all possible belief states is referred to as the belief space \mathcal{B} . The belief of state s at time t is denoted as $b_t(s)$. Every time the agent takes an action, its belief state is updated. Given $O(s', a, z)$, the probability of observing z in successor state s' after action a and the transition probability $T(s, a, s')$, Bayes' theorem can be applied to update the belief of the agent:

$$b_a^z(s') = \eta O(s', a, z) \sum_{s \in \mathcal{S}} T(s, a, s') b(s) \quad (10)$$

Where η is a normalizing constant. The belief state effectively sums up all of the agent's past actions and observations and is therefore a Markovian signal and a sufficient statistic to base its actions on. Since the agent is no longer certain of its position, the expected reward for a belief state has to be weighted by the belief in all individual states:

$$R(\vec{b}, a) = \sum_{s \in \mathcal{S}} b(s) R(s, a) \quad (11)$$

The initial value function at $t = 0$ is given by:

$$V_0(\vec{b}) = \max_a \sum_{s \in \mathcal{S}} b(s) R(s, a) \quad (12)$$

The value of a belief state under a policy π is computed according to:

$$V^\pi(\vec{b}) = \sum_{s \in \mathcal{S}} b(s) V^\pi(s) \quad (13)$$

The key observation here is that this knowledge is sufficient to transform the POMDP to a continuous state MDP, where belief space \mathcal{B} represents the state space \mathcal{S} . Because the resulting MDP has a continuous state space, the problem is still very hard to solve optimally. In section IIIC we will further elaborate on this concept and show how value iteration can be applied in POMDPs.

Numerous algorithms have been developed for solving POMDPs [7], [11], [12], [13], most of these using some form of value iteration. Qmdp is one of these, applying value iteration in its most rudimentary form.

B. Qmdp

An easy method for solving model-based POMDPs is to make use of the Q-values of the underlying MDP, thereby ignoring the observation model [11]. By treating the belief space as if it were the state space in an MDP, the value of taking action a in belief state \vec{b} is given by:

$$Q(\vec{b}, a) = \sum_{s \in \mathcal{S}} b(s) Q_{MDP}(s, a) \quad (14)$$

Where Q_{MDP} denotes the Q-value of the underlying MDP. With these values in hand, Eq. (6) can be rewritten to select the action with the highest expected value:

$$\pi(\vec{b}) = \arg \max_a \left[\sum_{s \in \mathcal{S}} b(s) Q_{MDP}(s, a) \right] \quad (15)$$

The Qmdp algorithm is easily implementable and can be very fast in a problem with a large number of states. Furthermore, it has been applied with great success on particular maze navigation tasks [11]. A disadvantage however, is that an agent following this policy does not take information gathering actions. For a more exact solution we have to consider the observation model and adjust the value iteration algorithm to suit POMDPs. A short version of the methods involved will be provided in the next section.

C. Value iteration in POMDPs

Value iteration can also be applied to compute solutions for POMDPs. Here, we will present a brief outline of the methods involved, as an introduction to the Perseus algorithm. For detailed descriptions we refer to Sondik [20] and Puterman [16].

Recall from section IIIA, that when acquiring the value function of a POMDP under a certain policy, the value of every state needs to be weighted by the agent's belief in the given state. For notational convenience Eq. (13) can be written as a dot product:

$$V^\pi(\vec{b}) = \vec{b} \cdot \alpha^\pi \quad (16)$$

Where $\alpha^\pi = \{V^\pi(s_1), V^\pi(s_2), \dots, V^\pi(s_n)\}$. In section IIB, a policy was described as a function specifying which action to take in a given state. Working towards an optimal policy, the agent needs to select the best action at every time step t :

$$V_t(\vec{b}) = \max_{\alpha \in \Gamma_t} \vec{b} \cdot \alpha \quad (17)$$

with $\Gamma_t = \{\alpha_1, \alpha_2, \dots, \alpha_k\}$. The state of the agent is a continuous function of all individual beliefs in a state. Assigning the belief of a state to every axis, plotting the belief state will result in an $|\mathcal{S}| - 1$ -dimensional hyperspace (probabilities sum to 1, thus the belief in $|\mathcal{S}| - 1$ states is sufficient to determine the entire belief state). All belief points are contained in a belief simplex Δ . With every region of the belief space, an optimal action is associated, this is represented by one of the α -vectors.

Again, applying the concept of weighting probabilities, we can combine the functions defined so far into a general formula for an optimal value function:

$$V^*(\vec{b}) = \max_{a \in \mathcal{A}} \left[\sum_{s \in \mathcal{S}} R(s, a) b(s) + \gamma \sum_{z \in \mathcal{Z}} V^*(b_a^z) \right] \quad (18)$$

where b_a^z is given by the belief function, defined in Eq. (10). For short, we can write: $V^* = HV^*$, with H defined as the Bellman backup operator.

Since for every region of the belief space, there is an α -vector optimizing the value, the optimal value function will be made out of a finite set of hyperplanes, building up the surface of the belief simplex. Sondik showed this function is *piecewise, linear and convex* (PWLC) for finite horizon POMDPs and is approximately PWLC for POMDPs with an infinite horizon [20].

Analogous to value iteration in MDPs, the value function is updated iteratively. In every iteration a set of α -vectors

will be added and a new value function, made up of the surface of all vectors will be computed. Every stage can be seen as a backup of the previous value function. At every backup stage, the vector parameterizing the surface of the value function can be computed according to:

$$\text{backup}(\vec{b}) = \alpha_{n+1}^b = \arg \max_{\alpha \in \Gamma_{n+1}} \vec{b} \cdot \alpha \quad (19)$$

Where $\Gamma_{n+1} = \{\alpha_1, \alpha_2, \dots, \alpha_l\}$ and $l = |HV_n|$, i.e., the number of vectors in the value function [21]. This notation will be useful for understanding the Perseus algorithm.

Value iteration is computationally expensive in POMDPs, because at each iteration the value of every point in the entire belief space is updated. Recently developed methods, known as point-based algorithms, have started working with restricting value iteration to a subset of the belief space [8], [12], [14], [27]. Perseus is one of such algorithms.

D. Perseus

Perseus is an approximate point-based value iteration algorithm for solving POMDPs and was introduced by Spaan and Vlassis in 2005 [21], [22]. The algorithm starts by performing a random walk through the environment, thereby sampling a set $B = \{b_1, b_2, \dots, b_m\}$ of reachable belief points. These points remain the same throughout the algorithm. This holds an advantage over other algorithms that work with the complete belief space in the sense that it only computes values for belief points that can actually be encountered by the agent.

The initial value function is set as a single vector, with all components set to $\frac{1}{1-\gamma} \min_{s,a} R(s, a)$, the minimal cumulative reward obtainable in the POMDP, guaranteed to be below V^* . Perseus introduces a new backup operator $\tilde{H}_{PERSEUS}$, and in every backup stage, tries to improve the value of all belief points, or at least makes sure that they do not decrease:

$$V_n(b) \leq V_{n+1}(b), \forall b \in B \quad (20)$$

It keeps track of the set of non-improved points \tilde{B} , and as long as \tilde{B} is not empty, samples uniformly at random a belief point b and computes $\alpha = \text{backup}(b)$. If the vector improves the value of b , it is added to the value function of V_{n+1} , otherwise a copy of V_n will be inserted. In an ideal situation, an increase in value of a belief point $b \in B$ will increase the value of many other points in B . Given the shape of the value function, such a method can be very effective in approximating solutions. The backup stage is given in Algorithm 2.

Algorithm 2 The backup stage $V_{n+1} = \tilde{H}_{PERSEUS}$

$V_{n+1} \leftarrow \emptyset$
 $\tilde{B} \leftarrow B$
repeat
 Sample a belief point b uniformly at random from \tilde{B}
 and compute $\alpha = \text{backup}(b)$
 if $b \cdot \alpha \geq V_n(b)$ **then**
 add α to V_{n+1}
 else
 add $\alpha' = \arg \max_{\alpha \in \Gamma_n} b \cdot \alpha$ to V_{n+1}
 end if
 $\tilde{B} = \{\forall b \in B : V_{n+1}(b) < V_n(b)\}$
until $\tilde{B} = \emptyset$

This stage is performed iteratively, until some stopping criterion is met. This could be, analogous to regular value iteration, terminating the algorithm as soon as the maximum difference between two backup stages $\max_{b \in B} (V_{n+1}(b) - V_n(b))$, drops below a threshold.

The POMDP algorithms discussed so far, all make use of value iteration. We will now discuss some previous work on Q-learning in combination with reinforcement learning and neural networks, followed by the introduction of a new approach, combining several of these techniques.

IV. RENQ FOR SOLVING MODEL-BASED POMDPs

As described in section IIE, the Q-learning algorithm updates each state-action pair after executing an action. However, because the belief space is used as a state space, the number of possible states encountered is infinite. Therefore, to work with a lookup table for each belief state-action pair becomes impossible and there is need for a function approximator, which generalizes between these pairs and associated Q-values. Neural Networks (NN) provide such a method and are known to be a powerful formalism in function approximation, gaining success in a wide variety of applications [4], [10], [18], [26], [17].

To use neural nets to predict Q-values, there are two possible approaches. Either one network is used, with $|\mathcal{A}|$ output units, or a single net is assigned to every $a \in \mathcal{A}$, as used by Lin [10]. The advantage of the latter approach is that, when trying to obtain the Q-value of a given action, one can easily address the responsible network. Also this will reduce the untraining of weights, caused by changes in the state space. A possible disadvantage of this method is that it requires more weights and might therefore sometimes require more training data.

To each $a \in \mathcal{A}$, we assign an NN: Q_a^{NN} . The belief state at time t , \vec{b}_t , is fed to the networks as an $|\mathcal{S}|$ -dimensional input vector and a single output unit is used to predict the Q-value of an action in the given belief state. The number of hidden units is left as a parameter.

As opposed to normal neural network training, the networks in this case do not learn from correct examples, but from approximations. So as a target, the Q-value at the next time step was fed to the network and used to compute the

error. To obtain this target, we can rewrite Eq. (8) and update the network of the previously selected action in the following way:

$$Q(\vec{b}_t, a_t) = r_{t+1} + \gamma \max_{a \in \mathcal{A}} Q_a^{NN}(\vec{b}_{t+1}) \quad (21)$$

Instead of updating the Q-value as in Eq. (7), the backpropagation learning algorithm is used to update the weights of the networks with learning rate α . The method discussed so far, which we call BQNN, has been used in [26] and obtained very good results for some small model-based POMDPs. In the following section we will show that this method can be enhanced by using region-based belief vectors.

A. RENQ: Region Enhanced Neural Q-learning

The RENQ algorithm enhances the belief vector by adding the average beliefs that the agent is in a particular region of the state space. This information is easily obtainable and its use can be very profitable. By using this kind of state abstraction, RENQ can be seen as a new hierarchical RL algorithm.

Basically, RENQ is inspired by a technique for extracting information from images in object recognition, known as a *spatial pyramid* [9], [1]. In this method, the image is divided into regions and spatial features, e.g., a histogram, are computed for all regions. This approach is known to improve recognition performance greatly. RENQ uses this method in a novel way and applies the spatial pyramids to the state space of the POMDP. The approach works with several levels. At each level, the state space is divided into k regions of equal size. At level 1 the used region is equal to the original state space, where each state is a singleton region, thus $k = 1$. This is equal to the BQNN method used in [26]. Level 2 decomposes \mathcal{S} into 2×2 quadrants, making $k = 4$. Level 3 subsequently subdivides \mathcal{S} in 3×3 regions with $k = 9$, level 4 in 4×4 , etc.

For every k , $\zeta_k \subset \mathcal{S}$ the average belief value \bar{b} is computed:

$$\bar{b}(\zeta_k) = \frac{1}{|\zeta_k|} \sum_{s \in \zeta_k} b(s) \quad (22)$$

The enhanced belief vector b_L^\dagger is the combination of all belief sets B_k for every level k :

$$b_L^\dagger = \cup_{k=1}^L B_k \quad (23)$$

Subsequently, b_L^\dagger is fed as an input-vector to the networks and Q-values are estimated. The computation of this additional information might seem redundant, but as we will show, this can actually be very effective.

Example. Consider a 4×4 grid maze, with $|\mathcal{S}| = 12$. The grid can be divided into 4 square regions of 2×2 (level 2). For each region we will compute the average belief according to Eq. (22). A depiction of the general idea is provided by Figure 1.




Fig. 1. Level 2 spatial pyramid applied to the belief state of a POMDP

The belief function represents a belief of the agent, being in a certain state. The knowledge we add will also provide the agent with an estimate of its approximate position. If the goal of the agent is to get to a goal state somewhere in the upper right corner of the state space and the agent has a fair degree of faith its position is somewhere within the boundaries of the lower left corner, it is very likely it will steer itself either north or east.

In this example, the state space was divided into 4 regions. For larger problems, one might consider breaking up the problem into more and/or larger sections, i.e., higher levels, and add the average belief of these regions to the input vectors of the networks as well. This is left as a parameter. Naturally, variations on this scheme can be developed, for instance for problems that are not captured in a square state space POMDP, the dimensions of the spatial pyramid can be modified to suit the particular problem.

In the following section we will provide details about the benchmark problems used to test the algorithms, the optimal parameter settings found and show the results comparing the RENQ algorithm to BQNN, Qmdp, and Perseus.

V. EXPERIMENTS

To test the algorithms, we use 3 square maze navigation tasks of 4×4 , 10×10 and 22×22 , with $|\mathcal{S}| = 12$, $|\mathcal{S}| = 73$ and $|\mathcal{S}| = 344$, respectively. Every maze has one static goal position and every other unoccupied state can be the initial state. The objective of the agent is to reach the goal position as soon as possible and with every action the agent can only reach adjacent states. In every maze, we use the set of actions: $\mathcal{A} = \{\text{go left, go up, go right, go down}\}$, with a 20% chance that the selected action is changed by a random action from \mathcal{A} . In every of the 4 directions the agent can either observe a wall or an empty field, making the cardinality of $|\mathcal{Z}| = 2^4 = 16$. We added 10% noise to the observation in each separate direction, meaning that an observation is correct with probability $0.9^4 = 66\%$. The agent receives a reward of 100 for reaching the goal position and is penalized by -0.1 for every other action. In all mazes the only opportunity for the agent to get a reward is by reaching the goal state. We therefore chose to measure the number of steps to the goal position, since this does not depend on the size of the rewards.

We run the algorithms on an Intel Dual Core 2.33GHz, with 3.4 GB RAM. For Perseus, the Matlab implementation

available on Spaan's website is used and the algorithm is run on Matlab 2009b for Linux. The RENQ and Qmdp algorithms use self written C++ implementations.

A. Small Maze

The maze is depicted in Figure 2, with G as the goal position. The starting state can be any other unoccupied state. The entire maze is surrounded by a 1 block wall. For all algorithms, the discount factor is set to $\gamma = 0.7$.




Fig. 2. The small maze. G denotes the goal position. Note that we did not draw the walls around the maze.

RENQ. A simulation lasts for 100,000 learning steps. During an experiment, we perform 50 simulations. A run is finished if the goal is hit or if the agent performed 1000 actions during the run. The learning rate α of the neural network is set to 0.015. The neural network used 20 sigmoidal hidden units for each separate action network. We used Boltzmann exploration with $\tau = 1$. We also used the same parameters to test the performance of BQNN and Q-learning with neural networks on an MDP.

Perseus. Following Spaan, we ran Perseus 10 times, each with a different random seed. For the small maze, we sampled a set $|B| = 200$ belief points. With every simulation the algorithm perform 1000 episodes, starting from random positions. We let the algorithm run for 120 seconds, which proved to be enough for convergence. The average of the total 10.000 trajectories is computed along with the standard deviation.

Qmdp. We ran the algorithm 100 times, in each simulation letting the agent start at each different starting location. The average of all the $(12)(100) = 1200$ different episodes is computed, along with a success percentage, indicating how often the goal was found in an episode. As a stopping criterion for the value iteration part we use $\theta = 1E^{-6}$.

We also tested regular value iteration on an MDP. Here instead of using observations of the maze which creates the need for belief vectors, the single current state of the agent is fully observed and given to the agent. Of course solving this MDP is much easier, and we mainly did these experiments to compare value iteration for MDP (an optimal method), with Q-learning and neural networks to solve the MDP, to see how much worse neural network solutions are compared to the optimal solutions. The results are shown in Table I. For value iteration on the MDP (VI MDP) we did not compute standard deviations, since it always computes the same policies.

TABLE I
RESULTS ON THE SMALL MAZE.

Method	Final steps	Nr. Times Goal hit	% Success
VI MDP	3.86		100
RL+NN MDP	3.85 ± 0.04	25375 ± 533	100
BQNN	4.37 ± 0.09	22156 ± 967	100
RENQ LEVEL 2	4.37 ± 0.08	22172 ± 1001	100
RENQ LEVEL 3	4.36 ± 0.06	22307 ± 796	100
Qmdp	4.38 ± 0.68		100
Perseus	4.79 ± 0.03		100

Discussion. As can be seen in Table I, RENQ significantly outperforms Perseus at all levels, but performs the same as Qmdp. Furthermore, RENQ at Level 3 outperforms BQNN without region enhanced beliefs, since it finds the goal the most times during its learning process. Note that although all RENQ systems receive the same number of total steps, some methods may learn faster and hence hit the goal more often. We can also see that solving this MDP with Value Iteration (VI MDP) finds a solution of 3.86 steps on average, whereas using Q-learning and a neural network as function approximator (RL+NN MDP) learns the same optimal policy.

B. Middle-sized Maze

We use the 10×10 maze shown in Figure 3. For all algorithms, we set the discount factor to $\gamma = 0.95$.




Fig. 3. The Middle-sized maze. G denotes the goal position.

RENQ. A simulation lasts for 200,000 steps. The learning rate α of the neural network is again set to 0.015. The neural network used 60 sigmoidal hidden units for each separate action network. We used Boltzmann exploration with $\tau = 1$.

Perseus. Again we ran the algorithm 10 times, each with a different random seed and let the agent perform 1000 trajectories, each starting from a different random starting location. We sampled $|B| = 1000$ and let the value iteration stage run for 600 seconds. The average number of steps for the 1000 trajectories is computed along with a standard deviation.

Qmdp. We use 100 simulations, in each simulation consists of $|S|$ episodes and the average of all the $100|S| = 7300$ trajectories is computed. The results are shown in Table II.

TABLE II
RESULTS ON THE 10×10 MAZE.

Method	Final steps	Nr. Times Goal hit	% Success
VI MDP	11.3		100
RL+NN MDP	11.9 ± 0.3	14165 ± 436	100
BQNN	15.2 ± 0.4	10971 ± 208	100
RENQ LEVEL 2	15.1 ± 0.3	11526 ± 195	100
RENQ LEVEL 3	15.2 ± 0.3	11608 ± 200	100
RENQ LEVEL 4	15.2 ± 0.3	11606 ± 179	100
Qmdp	14.7 ± 0.6		100
Perseus	15.7 ± 0.1		100

Discussion. As can be seen in Table II, Qmdp performs the best for this maze, while RENQ significantly outperforms Perseus at all levels. RENQ at higher levels learns much faster than BQNN with only the state-based belief vector, since they hit the goal significantly more often. We can also see that solving this MDP with Value Iteration (VI MDP) finds a solution of 11.3 steps on average, whereas using Q-learning and a neural network as function approximator (RL+NN MDP) comes very close to this optimum.

C. Large Maze

For the large maze shown in Figure 4, we use a discount factor $\gamma = 0.99$.




Fig. 4. The Large maze. G denotes the goal position.

RENQ. A simulation lasts 2,000,000 steps. The learning rate α of the neural network is set to 0.01. The neural network used 60 sigmoid hidden units for each separate action network. We used Boltzmann exploration with $\tau = 1$.

Perseus. Continuing in the same fashion, we ran the algorithm 10 times, each with a different random seed. We sampled $|B| = 10000$ belief points and let the algorithm run for 2 hours. The average number of steps of 5 simulations is computed along with a standard deviation.

Qmdp. Again we use 100 simulations, in each simulation we let the algorithm start from all the $|S|$ different goal positions. The average number of steps of all the $100|S| = 34400$ episodes is computed, along with a success percentage. The results are shown in Table III.

TABLE III
RESULTS ON THE 22×22 MAZE.

Method	Final steps	Nr. Times Goal hit	% Success
VI MDP	21.4		100
RL+NN MDP	23.7 ± 2.4	51372 ± 5413	100
BQNN	397.8 ± 1438	27290 ± 7502	97
RENQ LEVEL 2	33.9 ± 0.5	37639 ± 1363	100
RENQ LEVEL 3	33.6 ± 0.5	38418 ± 1338	100
RENQ LEVEL 4	33.8 ± 0.4	38098 ± 1641	100
Qmdp	35.6 ± 1.0		99.3
Perseus	34.7 ± 0.3		100

Discussion. As can be seen in Table III, RENQ significantly outperforms Qmdp and Perseus at levels 2, 3, and 4. BQNN fails to learn a good policy in 2 of the 50 simulations. RENQ at Levels 3 and 4 performs the best of all POMDP methods. We can also see that solving this MDP with Value Iteration (VI MDP) finds a solution of 21.4 steps on average, whereas using Q-learning and a neural network as function approximator (RL+NN MDP) again comes close to this optimum. Note that for this largest problem, Qmdp may be a too simple algorithm and is outperformed by Perseus and RENQ. We did some preliminary experiments with even larger problems and noted that the difference between RENQ and Qmdp becomes even larger.

VI. CONCLUSION

The partially observable Markov decision process (POMDP) framework provides a model for decision making under uncertainty, caused by for instance, noise in a robot's actuators and sensor readings. In this paper we have presented RENQ, a novel approach combining techniques from machine vision with Q-learning and neural networks to approximate an optimal solution for POMDPs. We have shown that RENQ outperforms Qmdp, a simple POMDP algorithm, and Perseus, a state-of-the-art algorithm, when the maze problems become larger.

A problem that we kept as future work is to automatically create the optimal spatial layouts. We found that adding layers with more states never reduces performance for partially observable maze navigation problems. We also found that adding more hidden neurons usually improves performance a bit, although this is at the cost of more computational time. The system is also fairly robust to the learning rate, and we hardly experimented with it. The benchmark problems all consist of maze navigation tasks, where state transitions are only defined for adjacent states. It would be interesting to see how RENQ can be used for problems where this is not the case. Ultimately, the goal is to work towards a method providing effective learning behavior in a real world situation. We also want to study different hierarchical approaches to improve RENQ's learning speed in future work.

REFERENCES

[1] A. Abdullah, R. C. Veltkamp, and M. A. Wiering. Spatial pyramids and two-layer stacking SVM classifiers for image categorization: a comparative study. In *IJCNN'09: Proceedings of the 2009 International Joint Conference on Neural Networks*, pages 1130–1137, 2009.

[2] K.J. Astrom. Optimal control of Markov decision processes with incomplete state estimation. *J. Math. Anal. Applic.*, 10:174–205, 1965.

[3] R. Bellman. *Dynamic Programming*. Princeton University Press, 1957.

[4] L.T. Dung, T. Komeda, and M. Takagi. Reinforcement learning for POMDPs using state classification. *Applied Artificial Intelligence*, 22(7&8):761–779, 2008.

[5] M. Givon and A. Grosfeld-Nir. Using partially observed Markov processes to select optimal termination time of tv shows. *Omega, International Journal of Management science*, 36:477–485, 2006.

[6] J. Hoey, P. Poupart, A. von Bertoldi, T. Craig, C. Boutilier, and A. Mihailidis. Automated handwashing assistance for persons with dementia using video and a partially observable Markov decision process. *Computer Vision and Image Understanding*, 2009.

[7] L.P. Kaelbling, M.L. Littman, and A.W. Moore. Reinforcement learning: A survey. *Journal of Artificial Intelligence Research*, 4:237–285, 1996.

[8] H. Kurniawati, D. Hsu, and W.S. Lee. SARSOP: Efficient point-based POMDP planning by approximating optimally reachable belief spaces. In *Proc. Robotics: Science and Systems*, 2008.

[9] S. Lazebnik, C. Schmid, and J. Ponce. Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories. In *CVPR*, pages 2169–2178, 2006.

[10] L.J. Lin. Self-improving reactive agents based on reinforcement learning, planning and teaching. *Machine Learning Journal*, 8(3/4), 1992. Special Issue on Reinforcement Learning.

[11] M.L. Littman, A.R. Cassandra, and L.P. Kaelbling. Learning policies for partially observable environments: scaling up. In *Proc. 12th International Conference on Machine Learning*, pages 362–370. Morgan Kaufmann, 1995.

[12] W.S. Lovejoy. A survey of algorithmic methods for partially observable Markov decision processes. *Annals of Operations Research*, 28(1):47–65, 1991.

[13] G.E. Monahan. A survey of partially observable Markov decision processes: Theory, models, and algorithms. *Management Science*, 28(1):1–16, jan 1982.

[14] J. Pineau, G.J. Gordon, and S. Thrun. Anytime point-based approximations for large POMDPs. *Journal of Artificial Intelligence Research (JAIR)*, 27:335–380, 2006.

[15] J.M. Porta, N. Vlassis, M.T.J. Spaan, and P. Poupart. Point-based value iteration for continuous POMDPs. *Journal of Machine Learning Research*, 7:2329–2367, 2006.

[16] M.L. Puterman. *Markov Decision Processes—Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc., New York, NY, 1994.

[17] M. Riedmiller. Neural fitted Q iteration - first experiences with a data efficient neural reinforcement learning method. In João Gama, Rui Camacho, Pavel Brazdil, Alípio Jorge, and Luís Torgo, editors, *Machine Learning: ECML 2005, 16th European Conference on Machine Learning, Porto, Portugal, October 3-7, 2005, Proceedings*, pages 317–328, 2005.

[18] G.A. Rummery and M. Niranjan. On-line Q-learning using connectionist systems. Technical report, October 04 1994.

[19] S.J. Russel and P. Norvig. *Artificial Intelligence: a modern approach*. Prentice Hall, Englewood Cliffs, NJ, 1994.

[20] E. J. Sondik. *The Optimal Control of Partially Observable Markov Decision Processes*. PhD thesis, Stanford University, 1971.

[21] M.T.J. Spaan. *Approximate planning under uncertainty in partially observable domains*. PhD thesis, University of Amsterdam, Amsterdam, The Netherlands, 2006.

[22] M.T.J. Spaan and N. Vlassis. Perseus: Randomized point-based value iteration for POMDPs. *Journal of Artificial Intelligence Research*, 24:195–220, 2005.

[23] R.S. Sutton and A.G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, MA, 1998.

[24] C. J. C. H. Watkins. *Learning from Delayed Rewards*. PhD thesis, Cambridge University, Cambridge, England, 1989.

[25] C. J. C. H. Watkins and P. Dayan. Q-learning. *Machine Learning Journal*, 8(3/4), May 1992. Special Issue on Reinforcement Learning.

[26] M.A. Wiering and H. van Hasselt. Ensemble algorithms in reinforcement learning. *IEEE Transactions on Systems, Man, and Cybernetics, Part B*, 38(4):930–936, 2008.

[27] W. Zhang and N.L. Zhang. Restricted value iteration: Theory and algorithms. *Journal of Artificial Intelligence Research (JAIR)*, 23:123–165, 2005.