


Artificial Neural Networks (Gerstner). Exercises for week 9

Variations of SARSA and continuous space

Exercise 1. Monte-Carlo Batch versus Bootstrap Batch

example trials:

- 1: $\mathbf{s}, \mathbf{a2} \rightarrow r=0.2, \mathbf{s}', \mathbf{a4} \rightarrow r=0$
- 2: $\mathbf{s}', \mathbf{a3} \rightarrow r=1$
- 3: $\mathbf{s}', \mathbf{a4} \rightarrow r=0$
- 4: $\mathbf{s}', \mathbf{a3} \rightarrow r=1$
- 5: $\mathbf{s}, \mathbf{a1} \rightarrow r=0$
- 6: $\mathbf{s}', \mathbf{a4} \rightarrow r=0$
- 7: $\mathbf{s}', \mathbf{a4} \rightarrow r=0.5$
- 8: $\mathbf{s}', \mathbf{a3} \rightarrow r=1$
- 9: $\mathbf{s}, \mathbf{a2} \rightarrow r=0.2, \mathbf{s}', \mathbf{a4} \rightarrow r=0.5$
- 10: $\mathbf{s}, \mathbf{a1} \rightarrow r=0$



The aim is to estimate Q-values on the graph as shown in the figure on top. We run 10 exploration trials with the rewards as indicated in the figure. Some trials start in state s others in state s' .

- a. Use Monte-Carlo in batch version to calculate the cumulative expected reward $R(s, a)$ by averaging along all trials that start at s, a (discount factor $\gamma = 1$) for all state-action pairs.
- b. Use Q-learning in batch version (= dynamic programming) to evaluate the estimated Q-values $Q(s, a)$ by exploiting the Bellman equation.
Hint: Start at the terminal state and work up backwards.
- c. Which of the two algorithms performs best on this example?

Exercise 2. Q-values for continuous states (in class)

We approximate the state-action value function $Q(s, a)$ by a weighted sum of basis functions (BF):

$$Q(s, \tilde{a}) = \sum_j w_{\tilde{a}j} \Phi(s - s_j),$$

where $\Phi(x)$ is the BF “shape”, and the s_j ’s represent the centers of the BFs.

Calculate

$$\frac{dQ(s, a)}{dw_{ai}},$$

the gradient of $Q(s, a)$ along w_{ai} for a specific weight linking the basis function i to the action a .

Exercise 3. Eligibility traces

In week 7 we applied, in the exercises, the SARSA algorithm to the case of a linear track with actions 'up' and 'down'. We found that it takes a long time to propagate the reward information through state space. The eligibility trace is introduced as a solution to this problem.

Reconsider the exercise from week 7, but include an eligibility trace: for each state s and action a , a memory $e(s, a)$ is stored. At each time step, all the memories are reduced by a factor λ : $e(s, a) = \lambda e(s, a)$, except for the memory corresponding to the current state s^* and action a^* , which is incremented:

$$e(s^*, a^*) = \lambda e(s^*, a^*) + 1. \quad (1)$$

Now, unlike the case without eligibility trace, all Q-values are updated at each time step according to the rule

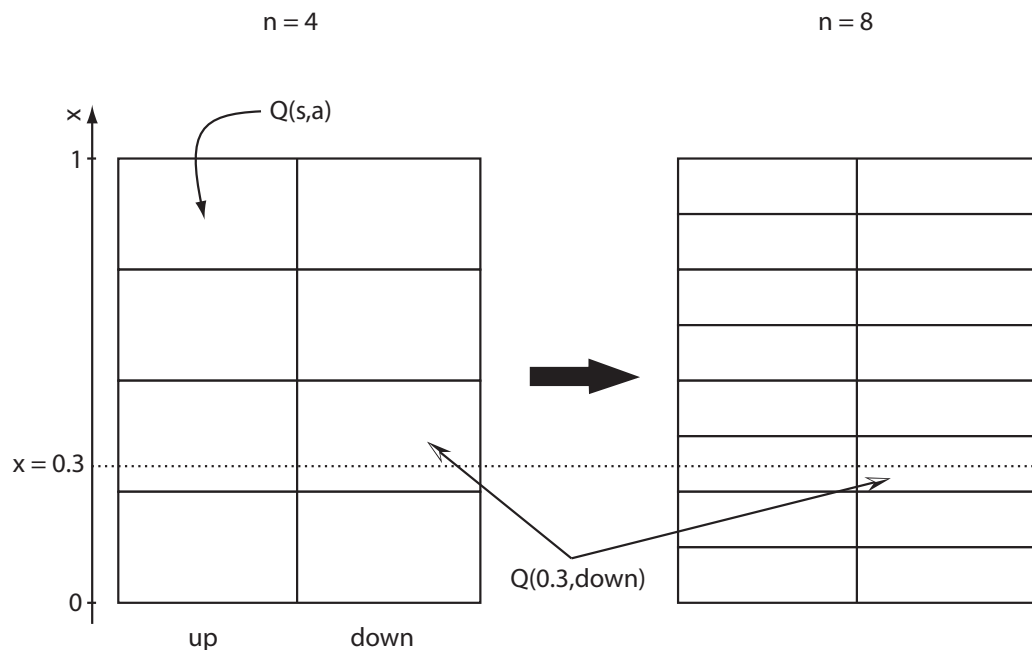
$$\forall(s, a) \quad \Delta Q(s, a) = \eta [r - (Q(s^*, a^*) - Q(s', a'))] e(s, a). \quad (2)$$

where s^*, a^* are the current state and action, and s', a' are the immediately following state and action.

We want to check whether the information about the reward propagates more rapidly. To find out, assume that the agent goes straight down in the first trial. In the second trial it uses a greedy policy. Calculate the Q-values after two complete trials and report the result.

Hint: Reset the eligibility trace to zero at the beginning of each trial.

Exercise 4. Eligibility traces in continuous space



The left part of the figure above shows a different representation of last week's "linear track" exercise: the vertical divisions represent different states, and the two columns correspond to the two possible actions available to the agent: go up or down. Each square represents a possible state-action combination, and thus a Q -value. (Note that the uppermost "up" action and the lowermost "down" action should be "greyed out": they are impossible. But this is not relevant to the rest of this exercise.)

Suppose now that the agent moves in a continuous 1-dimensional space $0 \leq x \leq 1$, with the target located at $x = 0$. Separate this state space into n equal bins of width $\Delta x = 1/n$. In each time step, the agent moves by one bin. Vary the discretization by varying n : $n = 4, 8, 16 \dots$

- a. Suppose that one action (such as move down) corresponds to one time step Δt in 'real time'. How should we rescale the parameter Δt , so that the speed $v = \Delta x / \Delta t$ remains constant when we change the discretization?
- b. We use an eligibility trace with decay parameter λ . How should we rescale λ , in order that the "speed of information propagation" in SARSA(λ) remains constant?

Hint: Consider for example, the Q-value at $x = 0.5$ after 2 complete learning trials.

Exercise 5. Gradient-based learning of Q-values

Assume again that the Q-values are expressed as a weighted sum of 400 basis functions: $Q(s, a) = \sum_{k=1}^{400} w_{ak} \Phi(s - s_k)$. For the moment the function Φ is arbitrary, but you may think of it as a Gaussian function. Note that s and s_k are usually vectors in \mathbb{R}^N in this case. There are 3 different actions so that the total number of weights is 1200. Now consider the error function $E_t = \frac{1}{2} \delta_t^2$, where

$$\delta_t = r_t + \gamma \cdot Q(s', a') - Q(s, a) \quad (3)$$

is the reward prediction error. Our aim is to optimize $Q(s, a)$ by changing the parameters w .

- a. Find a learning rule that minimizes the error function E_t by gradient descent. Consider the case where the actions a and a' are different.

Write down the learning rule. How many weights need to be updated in each time step?

- b. Find a learning rule that minimizes the error function E_t by gradient descent. Consider the case where the actions a and a' are the same.

Write down the learning rule.

Is there any difference to the case considered in a?

- c. Suppose that the input space is two-dimensional and you discretize the input in 400 small square 'boxes' (i.e., 20×20). The basis function $\Phi(s - s_k)$ is now the indicator function: it has a value equal to one if the current state s is in 'box' k and zero otherwise.

How do your results from (a) and (b) look like in this case?

- d. The learning rule in c is very similar to standard SARSA. What is the difference?

- e. Assume that $Q(s', a')$ in [Equation 3](#) does not depend on the weights. How does the learning rule look like for a, b and c. How is your result related to standard discrete SARSA?