

COSC 4368: Fundamentals of Artificial Intelligence Spring 2019
 Problem Set2 (Individual Tasks) Version 5

Deadline: Su., March 31, 11a (3% bonus); Tuesday, April 2, 11p (the latest)

Remarks: About 50% of the available points for Problem Set2 are allocated to Task 10.

7. Game Theory Khadija (4pts)

There are 2 players. Player 1 has 10 dollars. She can choose to give (10 dollars to player 2 and 0 to player 1), share (5 dollars to each player), or keep (10 dollars to player 1 and 0 to player 2). After she makes her decision, which player 2 observes, player 2 can accept or reject. After accepting payoffs are as specified (let's assume utilities are in dollar amounts), and after rejecting everyone gets 0. i. Describe this game in its extensive form. ii. Describe this game in its normal form. iii. Find all (pure and mixed) Nash equilibria of this game.

8. Reinforcement Learning Romita (6 pts)

a) Assume the following world UVW is given:

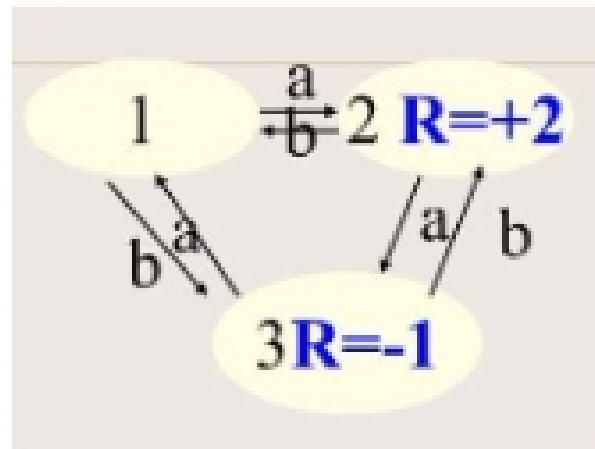



Fig. 1: UVW State Space.

Moreover, the current Q-table contains the following entries:

	Value
q(a,1)	1
q(b,1)	-0.5
q(a,2)	-0.5
q(b,2)	0
q(a,3)	0
q(b,3)	1



Assume the agent is currently in state 2 and her policy always applies action b in every state; how does the updated Q-Table look like after the agent has applied action b the second time¹ in state 2:

- i. when Q-Learning is used?
- ii. when SARSA is used?

Assume that the learning rate α is 0.5 and the discount rate γ is 1. Do not only report the updated value, but also give the formulas for the four Q-table updates.

- b) What **advantages** do you see in using SARSA over Q-Learning?
- c) What is the main difference between SARSA and Expected-SARSA? In which circumstances do you prefer Expected-SARSA over SARSA?
- d) How does reinforcement learning differ from supervised learning, e.g. learning neural networks?

9. NN and SVMs Romita (6 pts)

- a) What is the purpose of the Backpropagation algorithms; why is it needed?
- b) What factors determine the value of back-propagated errors in multi-layer NNs?
- c) Give a single² technical reason that explains the current popularity of Deep Neural Networks.
- d) The soft margin support vector machine solves the following optimization problem:

$$\operatorname{argmin} \frac{1}{2} \|\mathbf{w}\|^2 + C \sum_i \xi_i \quad \text{subject to } c_i(\mathbf{w} \cdot \mathbf{x}_i - b) \geq 1 - \xi_i \quad 1 \leq i \leq n.$$

What does the second term minimize? Depict all non-zero ξ_i in the figure below! What is the advantage of the soft margin approach over the linear SVM approach?

1 —after applying b-b-b-b

2 There might be multiple technical reasons!

Fig. 2: SVM Decision Boundaries for a Dataset containing Two Classes.

e) Referring to the figure above, explain how examples are classified correctly by SVMs! What is the relationship between γ_i and example i being classified correctly?