2) Learning

a) What is the main differences between reinforcement learning and inductive learning, also called learning from examples? [

b) We would like to predict the gender of a person based on two binary attributes: leg-cover (pants or skirts) and beard (beard or bare-faced). We assume we have a data set of 20000 individuals, 10000 of which are male and 10000 of which are female. 75% of the 10000 males are barefaced. Skirts are present on 50% of the females. All females are bare-faced and no male wears a skirt.

Compute the information gain of using the attribute leg-cover for predicting gender! [4]

$H(1/2,1/2)$

skirt: y (50) → $H(1,0)$; no (150) → $H(1/3,2/3)$

Gain(LegCover) = $H(1/2,1/2) - 3/4 * H(1/3,2/3)$

c) What is the key contribution of the backpropagation algorithm? What problems does it solve? Why is there no backpropagation algorithm for perceptrons?

... no intermediate layer where errors has to be determined.

3) Evolutionary Computing

a) What roll do crossover operators play in EC systems?

exploitation operators, explores all possible combinations of building blocks; does not introduce anything new.

b) What role do mutation operators play in EC systems?

exploration operator; might introduce a new feature that was extinct.

c) What advantages do you see in using probabilistic search algorithms—in contrast to deterministic search algorithms, such as simulated annealing or evolutionary computing?

a) create many good solution,

b) do not get stuck that easily

4) Reinforcement Learning [14]

a) Apply temporal difference learning[1] to the DEF World, depicted below, relying on the following assumptions: [4]

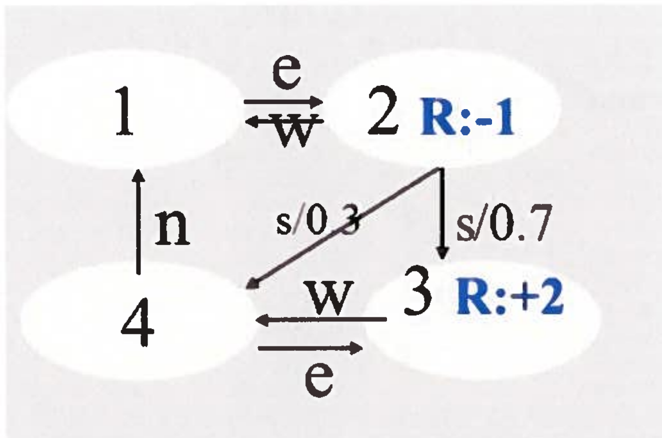

- If state 1 or 4 are visited a reward of 0 is obtained
- Utilities of the 4 states are initialized with 0
- The agent starts in state 2

The agent applies s(ending in state 3)-w-e. What are the utilities of states 2, 3, 4, and after those 3 operators have been applied? Do not only give the final result but also how you derived the final result including formulas used!

$U(2) = 0.5 \times 0 + 0.5 * (-1 + 0) = -0.5$

$U(3) = 0.5 \times 0 + 0.5 (+2 + 0) = +1$

$U(4) = 0.5 \times 0 + 0.5 (0 + 1 * 1) = 0.5$

c) Give the Bellman equation for state 2 of the DEF world! [2]

$U(2) = -1 + 1 * \max((0.3 * U(4) + 0.7 * U(3)), U(1))$

d) Assume you have a policy that always selects the action that leads to the state with the highest expected utility. Present arguments that this is usually not a good policy by describing scenarios in which this policy leads to suboptimal behavior of the agent [4]!

---

[1] and **not** Q-learning.

e) Assume you use temporal difference learning in conjunction with a random policy which choses actions randomly assuming a uniform distribution. Do you believe that the estimations obtained are a good measurement of the "goodness" of states, that tell an intelligent agent (*assume the agent is smart!!*) what states he/she should/should not visit? Give reasons for your answer! [3]

f) What role does the learning rate α play in temporal difference learning; how does running temporal difference learning with low values of α differ from running it with high values of α? [2]

g) Assume you run temporal difference/Q- learning with high values of γ—what are the implications of doing that? [2]

5) Naïve Bayes

Naïve Bayesian systems make the conditional independence assumption when for example computing P(D|S1,S2,S4). What assumptions are exactly made? What advantages do you see in the approach? What are the drawbacks of making the conditional independence assumption?

a) S1, S2, S4 are independent
S1|D, S2|D, S4|D are independent

b) Only probabilities concerning the association of single symptoms with diseases (and disease and symptom priors) have to be acquired ⇒ low knowledge acquisition cost

c) Asumption might be violated by the real data and errors are obtained in this case.