Christoph F. Eick

Monday, April 8, 2019
Review for COSC 4368 Midtem2 Exam
Solution Sketches

1) Classification and Supervised Learning in General

a) What is overfitting? What are the characteristics of overfitting? What can be done in the context of decision trees to battle overfitting?

The model is too complex, the testing/generalization error is not optimal, although the training error is low.

b) What is the key contribution of the backpropagation algorithm? What problems does it solve?

It measures and associates an error with the nodes of the intermediate layers that then can be used to learn the weights of incoming connections of nodes of the intermediate layers.

c) What is the purpose of training, test and validation sets in Supervised Learning? What else can be said about their relationship?

Training set used for model learning
Validation set is used to learn the best parameters for the method that generates the model (e,g, C and kernel function parameters in the case of the SVM)
Test set is used to determine the accuracy of the learnt model; e.g. to determine accuracy, testing/generalization error
All three sets should be disjoint➔otherwise cheating

2) Reinforcement Learning

a) 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!

Not suitable for unknown worlds due its lack of exploration
Not suitable for changing worlds due to its lack of exploration

Other answers might deserve credit.

b) Assume the following world is given:

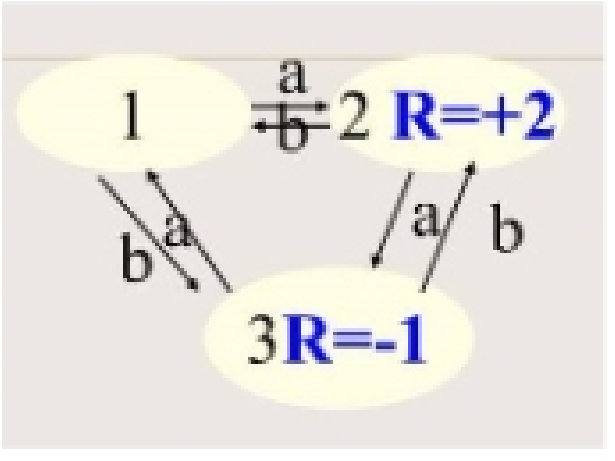

Moreover, the current Q-table contain the following entries:

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 fourth time assuming when Q-Learning is used? Assume that the learning rate α and the discount rate γ are both 0.5. Do not only report the updated value, but also give the formulas for the four Q-table updates. Assume that the q-table entries are initially set to 0.

$$Q(a,s) \leftarrow Q(a,s) + \alpha [R(s) + \gamma * Q(a',s') - Q(a,s)]$$

Q(b,2)= 0 + 0.5*(2 + 0 – 0)=1
Q(b,1)= 0+ 0.5*0=0
Q(b,3)=0+ 0.5*(-1+0.5*1-0)= □0.25
Q(b,2)=1+ 0.5*(2+1*0-1)=1.5

c) How does SARSA differ from Q-learning?
SARSA uses the chosen action in successor state in its q-table update when estimating the "utilities of the future", whereas Q-learning uses the optimal action in the successor state
Off policy vs. On policy Learning (add some points from the RL1 slide that discusses the difference)

3) SVMs [9]

a) What are the characteristics of hyperplanes that support vector machines learn from a training set? [3]

a. The margin of the learnt hyperplane is maximal
b. The hyperplane separates the examples of the two classes / minimizes the error in separating the examples of the two classes (in the case of the soft margin SVM).

b) The soft margin support vector machine solves the following optimization problem:

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

What does the first term minimize? Depict all non-zero ξ_i in the figure below! Depict all support vectors in the figure below---if example j is a support vector what is its value for ξ_j? What is the advantage of the soft margin approach over the linear SVM approach?

All other points have ξ_i values of 0!

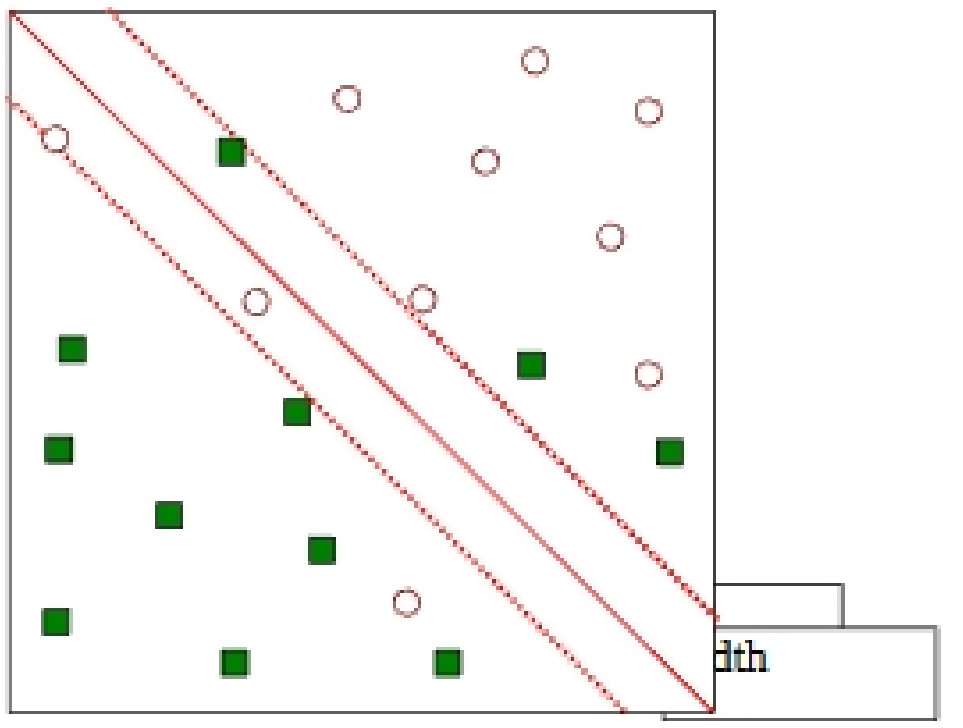

Minimize inverse margin; support vectors are example that reside on in the green class's and white class's hyperplane and carry the label of that hyperplane. Advantage: Soft margin SVM: can deal with datasets that contain examples that are not linearly separable.

What is a support vector?