

Question 1

For each of the following examples of AI, categorize each as supervised, unsupervised, or reinforcement learning, and briefly explain your reasoning.

- a) A robot AI learns how to walk by attempting to move its joints randomly. Every time it either takes a step successfully or falls over, the AI learns what joint movements work and don't work.
- b) An email service's AI is given samples of spam emails and non-spam emails and learns to identify whether or not an email is spam.
- c) A biology lab's AI takes a collection of proteins and organizes them into groups based on their structure, so that similarly structured proteins end up in the same group.
- d) A social network's AI uses existing tagged photos of people to identify when those people appear in new photos.

Answers

- a) TODO
- b) TODO
- c) TODO
- d) TODO

Question 2

Imagine a regression AI that makes the following predictions for the following 5 data points:

- For data point 1, the true output is 2 and the AI predicted 4.
 - For data point 2, the true output is 4 and the AI predicted 5.
 - For data point 3, the true output is 4 and the AI predicted 3.
 - For data point 4, the true output is 5 and the AI predicted 2.
 - For data point 5, the true output is 6 and the AI predicted 5.
- a) What is the total L1 loss across all of these data points (i.e., the sum of all the individual L1 losses for each data point)?
- b) What is the total L2 loss across all of these data points (i.e., the sum of all the individual L2 losses for each data point)?

Answers

- a) TODO
- b) TODO

Question 3

- a) Of 0-1, L1, and L2 loss functions, which are generally most appropriate for classification problems and which are most appropriate for regression problems? Why? (You should give three answers, one for each of 0-1, L1, and L2).
- b) If Hypothesis 1 has a lower L1 loss and a lower L2 loss than Hypothesis 2 on a set of training data, why might Hypothesis 2 still be a preferable hypothesis?
- c) When computing the cost of a hypothesis, why is regularization useful?

Answers

- a) TODO
- b) TODO
- c) TODO

Question 4

- a) In the ϵ -greedy approach to action selection in reinforcement learning, why would setting $\epsilon = 0$ be a bad idea when training the AI?
- b) In the ϵ -greedy approach to action selection in reinforcement learning, why would setting $\epsilon = 1$ be a bad idea when training the AI?

Answers

- a) TODO
- b) TODO