General Principles of Human and Machine Learning:

Pop Quiz #3

Name: **Date:** **Grade:**

Questions:

1. What is one practical advantage of the policy gradient method over the classical Q-learning approach?

2. The temperature parameter in a softmax policy, and epsilon in the epsilon-greedy policy both regulate the exploration/exploitation trade-off. What then, is the difference between the two exploration strategies?

3. Use a real-world example to demonstrate how hierarchical RL (or the 'options' framework) enhances performance via transfer-learning.

4. Name one benefit and one disadvantage for each approach to concept learning (i.e., how would an agent benefit or be disadvantaged by having such a mechanism?):
    a. Rule-based approaches

    b. Similarity-based approaches

5. Define how prototypes and exemplars correspond to different theories of human similarity-based concept learning and explain how they might predict different patterns of generalization.

6. According to Shepard, why does generalization occur?

7. Which axioms do human similarity judgments often violate, which Tversky was able to account for with his contrast model? Briefly explain each axiom.

8. In Bayesian concept learning, what is the size principle?

9. What is the difference between discriminative and generative classifiers?

10. Name a ML method that captures characteristics of the following approaches to human concept learning
    a. Rule-based methods

    b. Prototype theory

Answer Key

1. A large action space is better handled by policy gradient methods.

2. Epsilon-greedy explores all non-optimal options with the same constant probability, while softmax explores them proportional to their values.

3. Any example involving the reuse (transfer) of representations or policies that prevents the need to re-learn a policy from scratch. E.g., ability to transfer learned sequence of sweetening tea to sequence of sweetening coffee. Or being faster at using a tennis racket after having learned to handle a squash racket.

4.

a. Rule based
    i. Benefits. Specificity makes them quick to generalize, easy to communicate, and compositionality allows us to productively combine rules to create new ones.
    ii. Disadvantages**.** Brittle, inflexible (even when accounting for exceptions), and fails to account for the strong influence of similar exemplars
b. Similarity-based
    i. Benefits. Can make on-the-fly generalizations by comparison to past stimuli, intuitive that stimuli with similar features are more likely to belong to the same category, simple method for evaluating class membership based on distance in feature space
    ii. Disadvantages: Fails to capture more structured concept representations, difficulties in choosing the right similarity metric, (when using metric similarity) fails to account for violations of symmetry and triangle inequality axioms

5.

a. Definition. Prototype theory condenses representations of categories into a single representation (based on a weighted contribution of features), whereas exemplar theory retains all previously encountered stimuli.
b. Different predictions.
    i. These theories predict different patterns of generalization based on sensitivity to outliers (i.e., objects that are atypical of their class), where prototype representations would largely ignore their influence,

whereas exemplar theory would weigh the influence of outliers more strongly.

ii. There are also important differences in terms of memory limitations and cognitive costs, where exemplar theory requires increasing memory capacity as a function of more data, with increased processing to compute each n x n similarity comparison. In contrast, prototype theory requires constant memory and constant processing time.

6. Generalization arises from uncertainty about the extent of consequential regions in psychological space. It is not a failure to distinguish stimuli, but because we expect stimuli belonging to the same consequential region to result in the same outcomes.

7. Symmetry and law of triangle inequality.

a. Symmetry states that the similarity (or distance) between two stimuli should the same either way you compare them d(x,x')=d(x',x)
b. Triangle inequality states that when comparing three different stimuli (A,B,C), the sum of any two comparisons (i..e, sides of the triangle) must be greater than or equal to the third comparison (i.e., third side of the triangle). [Any answer involving a basic understanding of this involving a 3-way comparison or referring to geometry/pythagoras should be accepted, since this is not a geometry course]

8. The size principle operates under the assumption of strong sampling, where smaller hypotheses that are consistent with the data are more likely. In effect, this is akin to a bias towards simpler and less permissive hypotheses.

9. Discriminative models learn a decision-boundary, while generative models learn the distribution of the data (for each class).

10.

a. Decision-trees
b. K-means clustering