Self test Questions for Slides Set #5a

Prof. K R Chowdhary

February 10, 2025

These questions should help review the key concepts presented in the slides!

Introduction to RL

1. How is human learning similar to reinforcement learning (RL)?
2. What are the four main components of RL?
3. Explain the concept of the "reward function" in RL.
4. In RL, how does a learner interpret feedback if there is no teacher?

State-Space and Task Design

5. Why is the state-space of real-world problems often infinitely large in RL?
6. Provide an example of a feature subset for an RL task and explain its relevance.
7. What is the challenge of task design in RL, specifically with feature selection?
8. For the task of traveling from point A to point B, what factors might influence the agent's policy?

Feature Set of the State-Space

9. List some features of the agent's state-space in a navigation task.
10. How might features like the weather or the type of shoes affect the agent's learning in RL?
11. How does considering a six-dimensional space differ from a two- or three-dimensional space in RL?

Functions in RL

12. What is the role of the transition function $T(p, a) \rightarrow q$ in RL?
13. How does the reward function $R(S \times A) \rightarrow \mathbb{R}$ contribute to an RL agent's learning?
14. What does the value function $V^{\pi}(p)$ represent in the context of RL?
15. How does the optimal policy relate to the value function in RL?

Q-Learning Example: 1D Grid World

16. In the 1D Grid World example, what is the goal of the agent?
17. What are the two possible actions that the agent can take in the 1D Grid World?
18. How does Q-learning help the agent make decisions in this environment?

Q-Learning Algorithm

19. Write down the Q-learning update rule. What does each component represent?
20. How does the agent learn using the Q-learning algorithm in this example?
21. What are the roles of the learning rate α and the discount factor γ in Q-learning?
22. Why is the exploration rate (ϵ) important in Q-learning?

Testing the Learned Behavior

23. What is the purpose of testing after training in the Q-learning example?
24. What is the significance of the reward structure in the 1D Grid World example?
25. How does the agent's training progress get tracked in the Q-learning example?