CS786 Quiz 1

Q1. **Neuron electrochemistry (25 points)**

What is happening at the cell membrane of a neuron in each of the following five phases of the action potential illustrated below in the diagram?

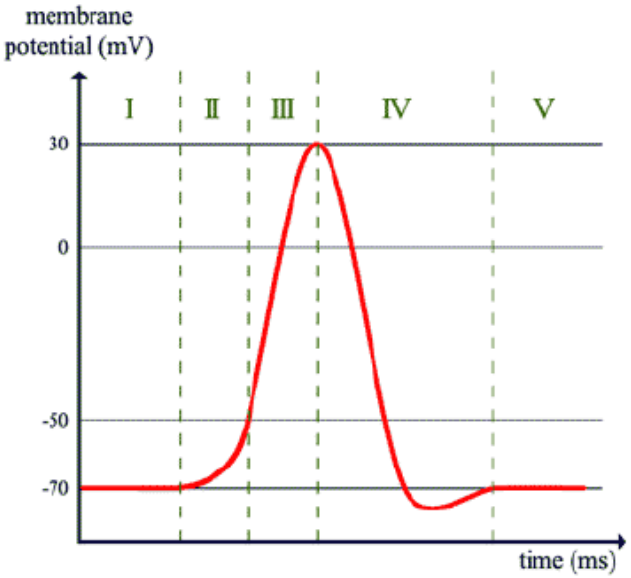

Phase I is the neuron at rest. There is a large concentration of Na+ ions outside the cell membrane (2 marks) and a large concentration of K+ ions inside the cell body (2 marks). The cell membrane contains voltage-gated ion channels, specific to the ion type (2 marks).

Phase II marks the onset of depolarization (1 mark). As a propagating action potential reaches this part of the membrane, it increases the instantaneous voltage at this location, opening both the sodium and potassium channels (2 marks). However, the Na channels open much more quickly than the K channels (2 marks), which means the Na outside the cell body flows into the cell body (2 marks), thus increasing the potential of the cell body (1 mark).

Phase III marks the completion of depolarization and the onset of repolarization (1 mark). At this point, the K channels have successfully opened, and K is flowing out from the cell body to the exterior of the cell (2 marks). This should cause the potential of the cell body to drop, since positive ions are leaving it (2 marks). But since sodium ions are still entering the cell body, the net voltage still increases, although more slowly now (1 mark).

Phase IV marks the completion of depolarization and the onset and completion of hyper-polarization (1 mark). In this phase, the Na concentrations are equalized, so there is no more Na transfer (1 mark), and the K transfer out of the cell body dominates the ion flow (2 marks). Hence, the potential drops, and in fact drops below the resting potential. Once the voltage becomes negative, the

voltage-gated K channels close (1 mark) and ATP-driven ion pumps take over to restore the original concentrations of ions on both sides of the membrane (2 marks).

In Phase V, the neuron has returned to its resting potential and is ready to fire again (1 mark)

Q2. What is David Redish's explanation for cocaine addiction? What facets of addiction does this explanation cover? How does this explanation differentiate between addiction and habituation? **(10 points)**

The explanation begins from the fact that cocaine inhibits the reuptake of dopamine at pre-synaptic autoreceptors in dopaminergic neurons. Since reuptake is what resets the baseline neurotransmitter concentration between two action potentials, inhibited reuptake means that there is excess free dopamine in the synaptic cleft between two neurons in a cocaine addict's brain (2 marks).

By the predictive coding hypothesis, dopamine encodes the temporal difference learning signal that determines how much to increase observers' preferences for specific actions. This TD signal has three components: additive influences from the experienced reward, and an expectation of what would happen in the future, and a subtractive influences from an expectation of what would happen in the present. The TD signal becomes negative when expected rewards do not materialize, which cause observers to reduce preference for actions bringing them to this state (2 marks).

However, since dopamine encodes the TD signal, excess dopamine concentration is decoded by post-synaptic neurons as a positive TD signal, which means that the preference for actions that produce the excess dopamine concentration (taking cocaine) cannot be unlearned, because the corresponding TD signal can never be negative (2 marks).

This explanation covers why people addicted to cocaine are unable to change their preference for cocaine even in the face of other rewards and punishments (2 marks).

It differentiates addiction from habituation, in the sense that addiction arises from a physiological distortion of the neurotransmitter-based coding of preference-change signals in the brain, whereas habituation is simply a strongly instated preference (2 marks).

Q3. **Association(25 points)**

Kamin blocking is a classic conditioning experiment wherein a dog (say) is conditioned with food in the presence of a light first, then with food in the presence of both the light and a sound. On testing, it turns out that the dog fails to respond when tested with just the sound. Can the Rescorla-Wagner model predict this effect? Assume a maximum association strength of 100, a attention and salience parameters of 0.8 each, and sketch out a quantitative explanation for why the model does (or doesn't) explain the effect.

5 marks for getting to the right answer - Kamin blocking is explained by the RW model.

5 marks for an accurate statement, explaining all notation, of the RW model

$$\Delta V_X = \alpha_X \beta (\lambda - V_{total})$$

where V_X is the learned association of the stimulus X to the unconditioned stimulus US, α and β are salience and attention parameters respectively, and λ is the maximum possible associative strength for the observer (2 marks for just the equation with no notation).

7 marks for showing quantitatively that prior conditioning with one stimulus drives V(X) high and hence V(total) high also.

6 marks for showing quantitatively that when conditioned with a compound cue including a previously trained cue, the association learned is low, because λ - V(total) is low.

2 marks for showing that, when tested with the cue from the compound that was not previously trained, the association strength is low.

Q4. **Reinforcement (25 points)** See the diagram of wumpus world overleaf and answer the corresponding questions.

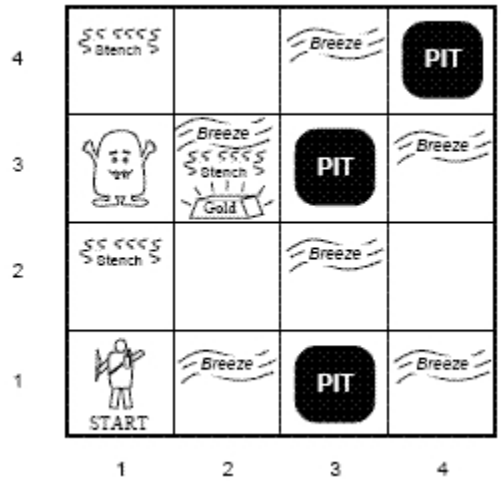

Wumpus world is a classic toy world from Russell & Norvig's AI book. In this question, we want to examine how (a) a Q-learner and (b) a SARSA-based agent, would fare if set loose in Wumpus world. The scoring rules are:

1. falling into a pit gives -100 points
2. encountering the wumpus gives -500 points
3. finding the gold gives +1000 points
4. encountering the stench of the wumpus gives -50 points
5. encountering a breeze gives -10 points

Let's say we always start from the cell marked START, and each episode terminates either once you find the gold, or have made 5 moves, can you sketch out a representative episode each for the Q-learner and the SARSA-based agent? Assume that your current knowledge of Q is represented by the diagram below.

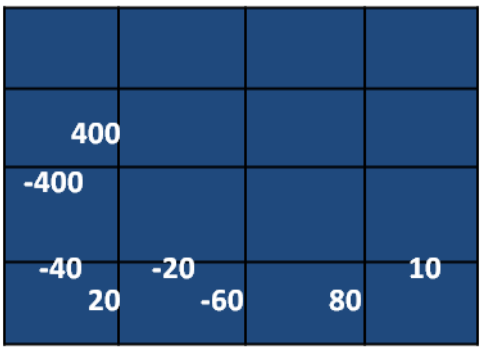

*Notation: numbers are placed close to one of the four box boundaries of each box. The number denotes Q(**s**,**a**), where **s** is the box where the number is located, and **a** is the action that leads from **s** to the box that is adjacent to **s** along the boundary of **s** that is closest to the number. No number present means this particular state-action pair has not been observed yet by the agent.*