Question 1. (5 points)

Consider classification with $Y = \{ 0, 1 \}$, where $P_c(y^*|x)$ is the probability that y^* is the true class of x, and rewards are given as

$$r_k = \begin{cases} 0 \text{ if } y = y^* \\ -1 \text{ if } y = 0 \text{ and } y^* = 1 \\ -3 \text{ if } y = 1 \text{ and } y^* = 0 \end{cases}$$

Consider the policy

$$y(x) = \arg\max_y P_c(y|x)$$

is this policy necessarily optimal, i.e. does it always coincide with the policy

$$\bar{y}(x) = \arg\max_y \mathbb{E}\left(r|x, y\right)$$

? Justify your answer mathematically.

Question 2. (2 points)

Discuss how the *exploration-exploitation* dilemma manifests itself in the *concept learning* scenario. Specify the conditions on which the execution of random actions would (would not) be useful for a concept-learning agent.

Question 3. (2 points)

Consider an algorithm that learns *monotone* disjunctions (or monotone conjunctions) from n-tuples of Boolean attribute values corresponding to n propositional variables. How can you use that algorithm to learn *general* disjunctions (or conjunctions) without changing it? You may change the number of inputs. How will your solution change the mistake bound in the case of the Winnow algorithm? Consider the number s of literals in the target disjunction constant.

Question 4. (1 points)

Let h, h' be propositional conjunctions. Is $h' \models h$ equivalent to $h \subseteq h'$? Justify your answer.

Question 5. (4 points)

Let h, h' be contingent propositional conjunctions that prescribe policies by

$$y = h(x) = \begin{cases} 1 \text{ if } x \models h \\ 0 \text{ otherwise} \end{cases}$$

We say that h at least as general as h' if $h(x) = 1$ for any $x \in X$ such that $h'(x) = 1$. Is it true that $h' \models h$ if and only if h is at least as general as h'? Justify your answer.