Question 1. (15 points)

Consider an algorithm learning in the mistake bound model. Prove that if the condition $\sum_{k=1}^{\infty} |r_k| \leq poly(n_X)$ ($n_X \in \mathbb{N}$) is satisfied, then from some time $K \in \mathbb{N}$ the agent will not make any mistakes, i.e.,

$$\exists K \in \mathbb{N}, \forall k \in \mathbb{N} : k > K \to r_k = 0$$

Question 2. (3 points)

Give two examples of a non-contingent conjunction, one tautologically true and one tautologically false. Do the same for disjunctions. Explain how an incomplete truth assignment to n propositional variables is represented by a conjunction and decide whether such a conjunction may be tautologically false. Explain why Winnow does not learn from incomplete truth assignments.

Question 3. (5 points)

Show where the assumption that the target hypothesis is a *monotone disjunction* is needed in the proof of Winnow mistake bound and explain how the proof would fail if assuming a general target concept on $X = \{ 0, 1 \}^n$.

Question 4. (2 points)

Give the lgg for all pairs from

$$\mathcal{H} = \{ \mathrm{p} \wedge \mathrm{q}, \mathrm{p} \wedge \neg \mathrm{q}, \mathrm{p} \vee \mathrm{q}, \mathrm{p} \vee \neg \mathrm{q} \}$$

whenever the lgg is defined for the pair.

Question 5. (3 points)

In concept classification, the generalization algorithm receives the sequence

$$x_1, x_2, \ldots, x_{10}$$

of observations (contingent conjunctions) where all x_k with odd indexes ($x_1, x_3, \ldots$) are positive examples, and all the others are negative. The agent's sequence of hypotheses is

$$h_1, h_2, \ldots, h_m$$

Determine

1. whether $h_2 = x_1$ or $h_2 = x_2$ or none of these options;
2. the sequence of hypotheses for the same agent that receives observations

$$x_1, x_1, x_3, x_3, \ldots, x_9, x_9$$