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.,

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

Answer:

The assumption $\sum_{k=1}^{\infty} |r_k| \leq poly(n_X)$ ($n_X \in \mathbb{N}$) implies that $\sum_{k=1}^{\infty} |r_k|$ is convergent, i.e., (1) below. An elementary calculus lemma says that a series with non-negative terms is convergent iff it is upper-bounded. We only need to prove the 'only if' direction of it, i.e.,

$$\lim_{m\to\infty} \sum_{k=1}^{m} |r_k| < \infty \tag{1}$$

$$\Rightarrow$$

$$\exists b \in \mathbb{R}, \forall m \in \mathbb{N} : \sum_{k=1}^{m} |r_k| \leq b \tag{2}$$

Let $s_m = \sum_{k=1}^{m} |r_k|$ and $l = \lim_{m\to\infty} s_m$. From the definition of limit, $\forall \epsilon > 0, \exists N_\epsilon \in \mathbb{N}$ such that $\forall m > N_\epsilon : |s_m - l| < \epsilon$. Let $N = N_\epsilon$ for $\epsilon = 1$. Then

$$\forall m > N_\epsilon : |s_m - l| < 1$$

i.e.,

$$\forall m > N_\epsilon : l - 1 < s_m < l + 1$$

Since $l < \infty$ due to (1), the infinite set $\{ s_{N+1}, s_{N+2}, \dots \}$ is (lower and upper) bounded. Also the set $\{ s_1, s_2, \dots s_N \}$ is bounded because it is finite. So their union, i.e. $\{ s_1, s_2, \dots \}$ is bounded, thus $\forall m \in \mathbb{N} : s_m \leq b$ where $b \in \mathbb{R}$ is an upper bound. Therefore (2) holds. From (2) it follows that there is only a finite set of indexes k such that $|r_k| > 0$. Let K be the maximum of this finite set and we finally have $\forall k \in \mathbb{N} : k > K \rightarrow r_k = 0$.

Intuitively, one could admit that (1) is true even with an infinite number of mistakes $|r_k| \neq 0$ as long as the 'spaces' between the non-zero rewards in the k-indexed series are progressively larger, and the size of these spaces increases sufficiently quickly to guarantee convergence. The proof above shows that under the standard axiomatization of calculus, this is not the case and there indeed is a finite time K after which no mistakes are made.

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.

Answer:

	true	false
conjunction	$\emptyset$	$\mathrm{p} \wedge \neg \mathrm{p}$
disjunction	$\mathrm{p} \vee \neg \mathrm{p}$	$\emptyset$

Let $\mathrm{p}_1, \mathrm{p}_2, \dots, \mathrm{p}_n$ be propositional variables assigned the respective truth-values by observations $x = x^1, x^2, \dots, x^n$. A conjunction representing x includes literal p_i ($\neg\mathrm{p}_i$, respectively) iff $x^i = 1$ ($x^i = 0$). The two cases are mutually exclusive so a variable cannot be included both with and without negation, thus the conjunction cannot be tautologically false.

Winnow decision policy computes the dot product $h \cdot x$ for which the entire tuple x must be provided.

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 $X = \{0, 1\}^n$.

Answer:

The proof says:

> *Each promotion doubles h^i where i is one of the s indexes corresponding to the s atoms in the target disjunction.*

This would not be necessarily true if the target hypothesis $\bar{h}$ was not a monotone disjunction. Say $n = 2$, $x = (0, 1)$ and the target hypothesis $\bar{h} = \neg p_1$ is a *non-monotone* disjunction. Assume the agent's hypothesis is $h = (1, 1)$, so x is a false negative, thus h gets promoted. By the Winnow learning rule, only the component h^2 is doubled towards $(1, 2)$ but p_2 is not in the target disjunction.

The proof also deals with the marginal case

> *If on the other hand $s = 0$ then the target disjunction is tautologically false.*

This would also not be correct if the target hypothesis was e.g. a conjunction (the empty conjunction is tautologically true.)

Question 4. (2 points)

Give the lgg for all pairs from

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

whenever the lgg is defined for the pair.

Answer:

We defined lgg for a pair of conjunctions or a pair of disjunctions. For all $h \in \mathcal{H}$, the $\mathsf{lgg}(h, h) = h$. Then,

$$\mathsf{lgg}(p \wedge q, p \wedge \neg q) = \mathsf{lgg}(p \vee q, p \vee \neg q) = p$$

and since lgg is commutative, the remaining two defined pairs (flipping the arguments above) are also p.

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

$$h_1, h_2, \ldots, h_{10}$$

so if some hypothesis is unchanged for m time steps, then there is a subsequence of m identical hypotheses above.

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

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

Answer:

1. $h_2 = x_1$. By the description given in the lectures, the first mistake is made on the first positive example x_1, so $r_2 = -1$, at which time the hypothesis is updated so $h_2 = x_1$.

2. The sequence of hypotheses is the same as for the former agent because $h_{k+1} = h_k$ both if x_k is a negative example (as in every second observation in the former case) or if $h_k \subseteq x_k$ (as in every second observation in the latter case).