



Problem Sheet 4

Instructions: The problem sheets are designed to increase your understanding of the material taught in the lectures, as well as to prepare you for the final exam. You should attempt to solve the problems on your own after reading the lecture notes and other posted material, where applicable. Once you have given sufficient thought to a problem, if you are stuck, you are encouraged to discuss with others in the course and with the lecturer during office hours. Please avoid posting on Piazza until after the submission deadline. You are *not permitted* to search for solutions online. Questions marked with an asterisk (*) are optional.

1 Learning MONOTONE-DNF is equivalent to learning DNF

The class $\text{MONOTONE-DNF}_{n,s}$ over $\{0,1\}^n$ contains boolean functions that can be represented as DNF formulae with at most s terms over n variables, and where each term only contains positive literals. Then define,

$$\text{MONOTONE-DNF} = \bigcup_{n \geq 1} \bigcup_{s \geq 1} \text{MONOTONE-DNF}_{n,s}.$$

The class DNF is defined analogously, except that the literals may be positive or negative. An efficient learning algorithm is allowed time polynomial in n , s , $\frac{1}{\epsilon}$ and $\frac{1}{\delta}$. Show that if the class MONOTONE-DNF is efficiently PAC-learnable, then so is DNF.

Remark: We have shown in the lectures that MONOTONE-DNF is exactly learnable using membership and equivalence queries, and hence also PAC-learnable using membership queries (see Problem 2). On the other hand, there is no known algorithm for PAC-learning DNF even when membership queries are allowed. In fact, under a suitable cryptographic assumption, it has been shown that PAC-learning DNF with or without membership queries is equivalent (Angluin and Kharitonov, 1991).

2 From Exact Learning to PAC Learning with Membership Queries

Let C be a concept class that is exactly efficiently learnable using membership and equivalence queries. We will consider the learnability of C in the standard PAC framework. Prove that if in addition to access to the example oracle, $\text{EX}(c, D)$, the learning algorithm is allowed to make membership queries, then C is PAC-learnable. Formally, show that there exists a learning algorithm that for all $n \geq 1$, $c \in C_n$, D over X_n , $0 < \epsilon < 1/2$ and $0 < \delta < 1/2$, that with access to the oracle $\text{EX}(c, D)$ and the membership oracle for c and with inputs ϵ , δ and $\text{size}(c)$, outputs h that with probability at least $1 - \delta$ satisfies $\text{err}(h) \leq \epsilon$. The running time of L should be polynomial in n , $\text{size}(c)$, $\frac{1}{\epsilon}$ and $\frac{1}{\delta}$ and the h should be from a hypothesis class H that is polynomially evaluatable.



3 Learning Rectangles using Statistical Queries

We will consider an extension of the statistical query model, where in addition to making queries of the form (χ, τ) to the oracle $\text{STAT}(c, D)$, the learning algorithm is allowed access to *unlabelled* examples from D , *i.e.*, it may get points $x \in X$ drawn according to D , but not the labels $c(x)$.

1. Briefly argue why any concept that is learnable with access to $\text{STAT}(c, D)$ and unlabelled examples, is also learnable with access to the noisy example oracle, $\text{EX}^\eta(c, D)$.
2. Give an algorithm for learning axis-aligned rectangles in the plane using $\text{STAT}(c, D)$ and unlabelled examples.

References

Dana Angluin and Michael Kharitonov. When won't membership queries help? In *Proceedings of the twenty-third annual ACM symposium on Theory of computing*, pages 444–454. ACM, 1991.