Selected Topics in Combinatorics, tutorial 3

PAC-learning, VC-dimension, ε-nets

March 16, 2023

Problem 1. Consider the space $\mathbb{R}^2$ and the concept class $\mathcal{C}$ of all the rectangles in the plane with sides parallel to the coordinate axes. Prove that $\mathcal{C}$ is PAC-learnable with the following learner: given a sample S of points in $\mathbb{R}^2$ output the smallest rectangle from $\mathcal{C}$ that is consistent with S.

Problem 2. Prove that if the concept class $\mathcal{C}$ has infinite VC-dimension then it is not PAC-learnable.

Hint: Let $X = \{1, \ldots, 2m\}$ and $\mathcal{C} = 2^X$. Show that $\mathcal{C}$ is not PAC-learnable for $\varepsilon = \delta = 0.1$ with m samples.

Problem 3. Let $F(X_1, X_2, \ldots, X_k)$ be a fixed set-theoretic expression (using the operations of union, intersection, and difference) with variables $X_1, X_2, \ldots, X_k$ standing for sets; for instance,

$$F(X_1, X_2, X_3) = (X_1 \cup X_2 \cup X_3) \setminus (X_1 \cap X_2 \cap X_3).$$

Let $\mathcal{S}$ be a set system on a ground set X with $\mathrm{VC}(\mathcal{S}) = d < \infty$. Let

$$\mathcal{T} = \{F(S_1, S_2, \ldots, S_k) \colon S_1, S_2, \ldots, S_k \in \mathcal{S}\}.$$

Prove that $\mathrm{VC}(\mathcal{T}) = \mathcal{O}(kd \log k)$.

Hint: Show that for any set A we have $F(S_1, S_2, \ldots S_k) \cap A = F(S_1 \cap A, S_2 \cap A, \ldots, S_k \cap A)$.

Problem 4. Consider multilayered feedforward neural nets with a single output which use sgn(x) as an activation function (i.e. every neuron computes a linear threshold function of it inputs). Prove that any fixed architecture with n_0 inputs and ω parameters the class $\mathcal{C}$ of function that can be obtained by instantiating these parameters satisifies:

$$\mathrm{VC}(\mathcal{C}) \leq 2\omega \log_2(e\omega).$$

Problem 5. Let $(X, \mathcal{F})$ be a set system and let $Y \subseteq X$ be arbitrary. Consider the set system $(Y, \mathcal{F} \cap Y)$ where $\mathcal{F} \cap Y = \{F \cap Y \colon F \in \mathcal{F}\}$. Prove that $\mathrm{VC}(\mathcal{F} \cap Y) \leq \mathrm{VC}(\mathcal{F})$.

Definition 1. Let $(X, \mathcal{F})$ be a set system, let μ be a probability measure on X and let $\varepsilon > 0$ be given. We say that a set $A \subseteq X$ is an ε-*net* with respect to μ if for every $F \in \mathcal{F}$ with $\mu(F) \geq \varepsilon$ we have $F \cap A \neq \emptyset$.

Problem 6. Prove that for any $\varepsilon > 0$ and $d \in \mathbb{N}$ there is some $N = N(\varepsilon, d)$ with the following property: if $(X, \mathcal{F})$ is a set-system with $\mathrm{VC}(\mathcal{F}) \leq d$ and μ is a probability measure on X, then there is an ε-net $A \subseteq X$ with respect to μ and $\mathcal{F}$ of size $\leq N$.

Problem 7. Let $(X, \mathcal{F})$ be a set system such that $\mathrm{VC}(\mathcal{F})$ is infinite. Show that for every $\varepsilon > 0$ and for every $n \in \mathbb{N}$ we can find a finite set $Y \subseteq X$ of size n such that the smallest ε-net for $(Y, \mathcal{F} \cap Y)$ with respect to the uniform counting measure $\mu(S) = \frac{|S \cap Y|}{|Y|}$ is of size at least $(1 - \varepsilon)n$.