Problem Set 4 **(Graded)** —*Due Tuesday, November 5, before class starts*

For the Exercise Sessions on Oct 15 and Oct 29

Last name	First name	SCIPER Nr	Points

Problem 1: Lipschitz Bandits

Assume for the following that you have a bandit algorithm at your disposal that has an expected regret, call it R_n, bounded by $c\sqrt{Kn\log(n)}$, where K is the number of arms and n is the time horizon.

You have to design an algorithm for the following scenario. There are infinitely many bandits. More precisely the bandits are indexed by x, $x \in [0,1]$. Bandit x has mean $\mu(x)$ (which is unknown). But you do know that the various bandits are related in the sense that

$$|\mu(x) - \mu(y)| \leq L|x-y|, \tag{1}$$

where L is a known constant. This is known as the Lipschitz bandit problem due to the Lipschitz condition (1).

A natural approach to such a bandit problem is to discretize the space of bandits. I.e., assume that you pick K positions $0 \leq x_1 < x_2 < \cdots < x_K \leq 1$ and run your given bandit problem on these K bandits.

a) Bound the expected regret as a function of K, n, L and the placement of points.

b) For n and L fixed, minimize your expression with respect to K and the placement of points.

Hint: In order to simplify your computation, you might want to slightly loosen your bound.

Problem 2: MMSE Estimation

Consider the scenario where $p(x|d) = de^{-dx}$, for $x \geq 0$ (and zero otherwise), that is, the observed data x is distributed according to an exponential with mean $1/d$. Moreover, the desired variable d itself is also exponentially distributed, with parameter λ, that is, $p(d) = \lambda e^{-\lambda d}$.

(a) Find the MMSE estimator of d given x, and calculate the corresponding mean-squared error incurred by this estimator.

(b) Find the MAP estimator of d given x.

Problem 3: Conditional Independence and MMSE

For simplicity, throughout this problem, **all random variables are assumed to be zero-mean.**

(a) Show that if X and Y are conditionally independent given Z, then

$$\mathbb{E}[(X - \mathbb{E}[X|Z])(Y - \mathbb{E}[Y|Z])] = 0. \tag{2}$$

(b) Now let X and Y be jointly Gaussian (zero-mean). It is well known that if $\mathbb{E}[XY] = 0$, then X and Y are independent. Establish this fact starting from the observation that for (zero-mean) Gaussian random variables X and Y, we may always write $Y = \alpha X + W$, for some constant α, where W is zero-mean Gaussian *independent of* X. *Note: This prepares you for Part (c).*

(c) Let X, Y, Z be jointly Gaussian (and zero-mean, as throughout this problem). Prove that if

$$\mathbb{E}[(X - \mathbb{E}[X|Z])(Y - \mathbb{E}[Y|Z])] = 0, \tag{3}$$

then X and Y are conditionally independent given Z. *Hint:* Make sure to solve Part *(b)* first. Recall that for three jointly Gaussians X, Y, Z, we can always write $Y = \gamma X + \delta Z + V$, for some constants γ and δ, where V is Gaussian and independent of X and Z.

(d) Let X, Y, Z be jointly Gaussian (and zero-mean, as throughout this problem). Prove that X and Y are conditionally independent given Z if and only if

$$\mathbb{E}[XY]\mathbb{E}[Z^2] = \mathbb{E}[XZ]\mathbb{E}[YZ]. \tag{4}$$

(e) Continuing from Part *(d)*, let us simplify: $\mathbb{E}[X^2] = \mathbb{E}[Y^2] = \mathbb{E}[Z^2] = 1$, and use the notation $\rho = \mathbb{E}[XY]$. Define $a = \mathbb{E}[XZ]$ and $b = \mathbb{E}[YZ]$. Find

$$\arg\max_{a,b} \min_{f} \mathbb{E}[(Z - f(X,Y))^2], \tag{5}$$

where the inner minimum is over all measurable functions $f(x, y)$.

Problem 4: Fisher Information and Divergence

Suppose we are given a family of probability distributions $\{p(\,\cdot\,;\theta) : \theta \in \mathbb{R}\}$ on a set $\mathcal{X}$, parametrized by a real valued parameter θ. (Equivalently, a random variable X whose distribution depends on θ.) Assume that the parametrization is smooth, in the sense that

$$p'(x;\theta) := \frac{\partial}{\partial\theta} p(x;\theta) \quad \text{and} \quad p''(x;\theta) := \frac{\partial^2}{\partial\theta^2} p(x;\theta)$$

exist. (Note that the derivaties are with respect to the parameter θ, not with respect to x.) We will use the notation $\mathbb{E}_{\theta_0}[\cdot]$ to denote expectations when the parameter is equal to a particular value θ_0, i.e., $\mathbb{E}_\theta[g(X)] = \sum_x p(x;\theta) g(x)$.

Define the function $K(\theta, \theta') := D\big(p(\,\cdot\,;\theta) \| p(\,\cdot\,;\theta')\big)$.

(a) Show that for any θ_0, $\dfrac{\partial}{\partial\theta} K(\theta, \theta_0) = \sum_x p'(x;\theta) \log \dfrac{p(x;\theta)}{p(x;\theta_0)}$.

(b) Show that $\dfrac{\partial^2}{\partial\theta^2} K(\theta, \theta_0) = \sum_x p''(x;\theta) \log \dfrac{p(x;\theta)}{p(x;\theta_0)} + J(X;\theta)$ with

$$J(X;\theta) := \mathbb{E}_\theta\left[\big(p'(X;\theta)/p(X;\theta)\big)^2\right].$$

(c) Show that when θ is close to θ_0

$$K(\theta, \theta_0) = \tfrac{1}{2} J(X;\theta_0)(\theta - \theta_0)^2 + o((\theta - \theta_0)^2)$$

(d) Show that $J(X;\theta) = -\mathbb{E}_\theta\left[\frac{\partial^2}{\partial\theta^2} \log p(X;\theta)\right]$.