ECE531 Spring 2011 Midterm Examination

Instructions: This exam is worth a total of 300 points. You may consult one double-sized letter-sized sheet of notes (in your own handwriting) and you may use a calculator during the exam. Please show your work on each problem and box/circle your final answers. Points may be deducted for a disorderly presentation of your solution. This exam is closed-book.

1. 100 points total. During FDA certification, a drug test was tested on a large controlled population and found to have the following statistics:

   - For subjects that had *not* taken the drug, the drug test falsely indicated the presence of the drug 10% of the time, was inconclusive 20% of the time, and correctly detected the absence of the drug 70% of the time.
   - For subjects that had taken the drug, the drug test correctly detected the presence of the drug 80% of the time, was inconclusive 15% of the time, and failed to detect the drug 5% of the time.

   Given one outcome from this drug test from an uncontrolled test subject, we wish to design an optimal decision rule for deciding whether the test subject has taken the drug or not.

   (a) 30 points. Set up this hypothesis testing problem by explicitly defining the states, hypotheses, observations, and the conditional distributions on the observations for each state.

   (b) 40 points. Sketch the Pareto-optimal risk tradeoff surface as accurately as possible assuming the uniform cost assignment (UCA). Label all vertices. Hint: You do not need to find the whole set of achievable risk vectors here. The Pareto-optimal risk tradeoff surface can be found by only computing the conditional risks of "good" decision rules.

   (c) 30 points. Find the Neyman-Pearson decision rule for $\mathcal{H}_0$:absence versus $\mathcal{H}_1$:presence that gives a false positive probability of $\alpha = 0.1$. Simplify your decision rule as much as possible. What is the probability of detection?

2. 100 points total. Suppose you receive an observation Y drawn either from conditional distribution $p_0(y)$ or conditional distribution $p_1(y)$ shown in Figure 1. Hypothesis $\mathcal{H}_i$ is that the observation is drawn from $p_i(y)$ for $i \in \{0, 1\}$.

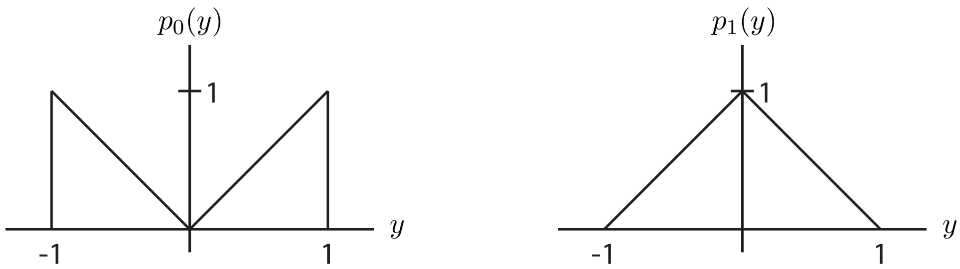

Figure 1: Conditional distributions for Problem 2.

(a) 40 points. Find the Neyman-Pearson detector for significance level $0 \le \alpha \le 1$.

(b) 20 points. Compute the probability of detection for your detector from part (a) as a function of the significance level $0 \le \alpha \le 1$. Hint: Check your result at $\alpha = 0$ and $\alpha = 1$.

(c) 40 points. Assume the uniform cost assignment and find a Bayes decision rule for a general prior.

In all cases, please simplify and describe your decision rules as explicitly as you can.

3. 100 points total. Suppose you observe one realization of a random variable $Y \in \mathbb{R}$ with $Y \sim \mathcal{N}(x, 1)$ and must decide between hypotheses

$$\begin{aligned} \mathcal{H}_0 &: x = 0 \\ \mathcal{H}_1 &: x \in \{1, 2\} \end{aligned}$$

(a) 50 points. Does a uniformly most powerful (UMP) decision rule of significance level α exist for this problem? To receive full credit, your reasoning must be correct. You do not need to find the UMP decision rule if it exists.

(b) 50 points. Suppose the prior probabilities of all three values of x are equal and the uniform cost assignment (UCA). Find the Bayes decision rule. Simplify and describe your Bayes decision rule as explicitly as you can.