Theory of Inference: Homework 3

1. Important but a bit tedious. Let $\boldsymbol{X} := \{X_1, \dots, X_m\} \in \boldsymbol{\mathcal{X}}$, as usual. Show that the PMF of the observables is

$$\mathrm{p}(\boldsymbol{y}) = \prod_{i=1}^{n} \mathrm{p}_X(y_i)$$

under the following conditions:

   (a) Simple observational model (SOM), i.e.

$$Y_i = X_i \qquad i = 1, \dots, n,$$

   (b) $X_1, \dots, X_m \overset{\text{iid}}{\sim} \mathrm{p}_X(\cdot)$, i.e.

$$\mathrm{p}(\boldsymbol{x}) = \prod_{i=1}^{m} \mathrm{p}_X(x_i).$$

2. Here is a exam-style revision question on decision theory and prediction. Each part is worth five marks.

   (a) Describe the statistical framework for analysing decision problems, including decision rules. Define what it means for an action or a decision rule to be optimal.

   (b) Consider the special case of a prediction problem: how does this differ from more general decision problems? Provide illustrations of prediction problems for the same random quantity X, but which are likely to differ in the prediction that is made.

   (c) Explain the motivation for producing 'generic' predictions, and justify the use of convex loss functions for such predictions.

   (d) Let $\boldsymbol{Y}$ be a set of observables with the statistical model $\boldsymbol{Y} \sim \mathrm{p}(\cdot\,; \theta)$ for some $\theta \in \Omega$. Define what is meant by an 'estimator', and what it means for an estimator to be 'admissible' (for simplicity, assume that Ω is finite).

(e) State, informally, Wald's theorem. Using a diagram, describe the classification of the space of all estimators, in terms of necessary and sufficient conditions for admissibility. Where does Stein's paradox locate the Maximum Likelihood (ML) estimator in this space?

Please hand in your answers for marking next Tue (17 Mar), at the lecture or by 5pm in the box outside my office door. I will return them in Thu's lecture.

Jonathan Rougier
Mar 2015