

PSM SPRING 2018, HOMEWORK 8

1. (*concerning Bayesian model estimation*) Like in the demo example in the LN, we consider a scenario where we are given a sample $X_1(\omega), \dots, X_N(\omega)$ of a real-valued RV X , and we assume that the true distribution P_X is a normal distribution with unit variance σ^2 but unknown mean μ , that is, we are faced with a 1-parametric family of models where the parameter is $\theta = \mu$. We furthermore assume a prior hyperdistribution for μ that is the normal distribution with zero mean and a standard deviation of 1. Furthermore, we consider a case where the sample size is just $N = 1$ and the (only) sample data point is $X_1(\omega) = -1$. Give a formula for the pdf of the posterior distribution and calculate the posterior mean estimate $\hat{\mu}$ numerical on your computer.
2. At our course homepage you find a link to a final exam (with solutions) from a bygone advanced course on ML. This exam consists in a step-by-step discussion of learning a so-called *trigram model* of English texts. Such trigram models are in widespread use in text data mining applications. Much of the themes and methods exhibited in that final exam are accessible to you and may be of interest for you - just “good to know” stuff. I recommend to go through this exam (skipping the tasks nr. 3 and 8 because they relate to methods that we did not cover in class) and try to understand the general set-up, the tasks specifications, and try to find answers yourself before looking at the solutions.