HOMEWORK 1

Frequentism and Bayesianism: A Python-driven Primer

Jake VanderPlas*†

Abstract—This paper presents a brief, semi-technical comparison of the essential features of the frequentist and Bayesian approaches to statistical inference, with several illustrative examples implemented in Python. The differences between frequentism and Bayesianism fundamentally stem from differing definitions of probability, a philosophical divide which leads to distinct approaches to the solution of statistical problems as well as contrasting ways of asking and answering questions about unknown parameters. After an example-driven discussion of these differences, we briefly compare several leading Python statistical packages which implement frequentist inference using classical methods and Bayesian inference using Markov Chain Monte Carlo.[1]

Index Terms—statistics, frequentism, Bayesian inference

advanced Bayesian and frequentist diagnostic tests are left out in favor of illustrating the most fundamental aspects of the approaches. For a more complete treatment, see, e.g. [Wasserman2004] or [Gelman2004].

The Disagreement: The Definition of Probability

Fundamentally, the disagreement between frequentists and Bayesians concerns the definition of probability.

For frequentists, probability only has meaning in terms of **a limiting case of repeated measurements**. That is, if an astronomer measures the photon flux F from a given non-variable star, then measures it again, then again, and so on,

Jake VanderPlas

(Optional) Jupyter notebooks:

1. FreqBayes1
2. FreqBayes2
3. FreqBayes3
4. FreqBayes4
5. FreqBayes5

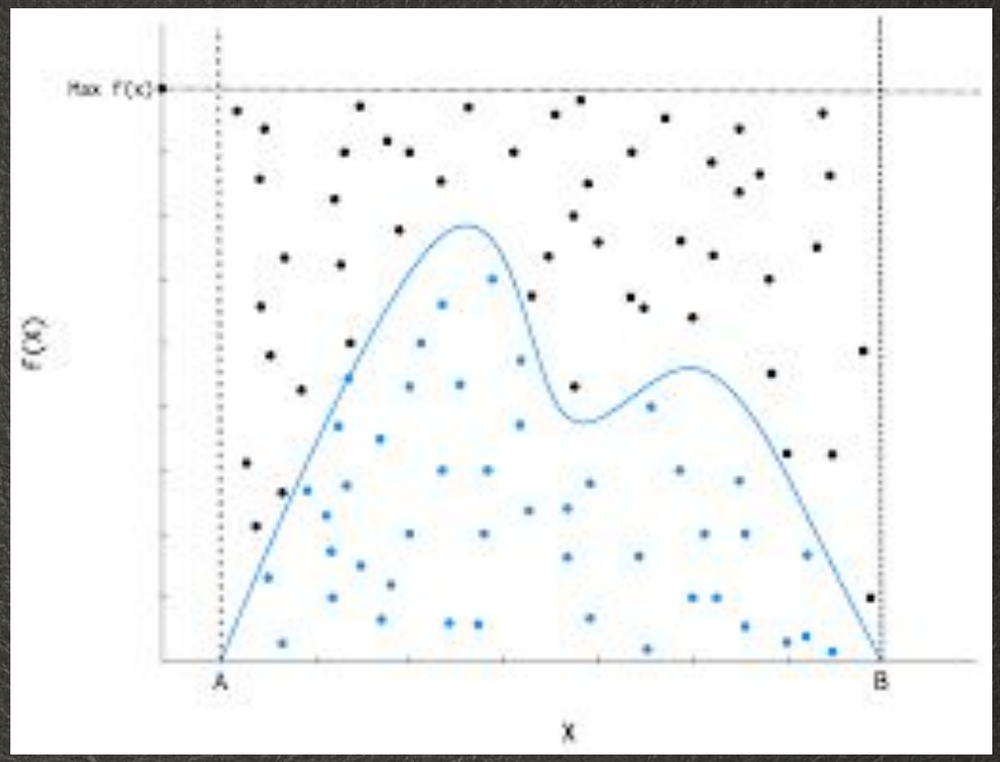

- Examples of Bayesian regression, hypothesis inference, and MCMC chains.