Statistical Science
2019, Vol. 34, No. 2, 202–205
https://doi.org/10.1214/19-STS704
Main article: https://doi.org/10.1214/18-STS674

Comment: Bayes, Oracle Bayes, and Empirical Bayes

Thomas A. Louis

1. INTRODUCTION

Brad Efron has done it again. He presents fascinating and insightful analyses that "open the box" on the properties of empirical Bayes methods. I especially like the exploratory data analysis theme, reminding us to look at the data, consider what information sources are relevant, and to conduct sensitivity analyses. These highlight the importance of computing diagnostics, and the dangers of black box modeling.

In what follows, I evaluate f-modeling (generate posterior summaries from the estimated marginal distribution of the data) and g-modeling (estimate the prior distribution and use Bayes' rule to obtain the posterior), consider Oracle Bayes, and address the choice between Bayes and empirical Bayes.

2. f- AND g-MODELING

Building on Efron (2014), Brad further compares f- and g-modeling as strategic approaches. While f-modeling is somewhat easier to implement, and the Robbins result for the Poisson (Robbins, 1983) is truly neat and showed how "empirical" can be wedded to "Bayes," g-estimation wins the day. Producing an effective g-model has its challenges, but the hard work pays off in that the (estimated) posterior distribution and generated summaries respect all constraints induced by prior to posterior mapping. There may be some models and goals for which f-modeling is competitive to g, but the situations are few and likely null when data aren't marginally i.i.d., in multivariate models, for goals such as histogram estimation and ranking (see below), or benchmarking (Bell, Datta and Ghosh, 2013). However, producing a good estimate of the X-marginal distribution is still very important; for example, it is central to assessing model fit (see Box, 1980).

Thomas A. Louis is Professor Emeritus, Department of Biostatistics, Johns Hopkins Bloomberg School of Public Health, 615 North Wolfe Street, Baltimore, Maryland 21205, USA (e-mail: tlouis@jhu.edu).

2.1 The Basic Poisson Model

In Section 5, Efron presents the Robbins (1983) f-modeling approach to estimating the posterior mean, $e_g(x)$, and variance, $v_g(x)$, of the Poisson rate parameter. That $v_g(x)$ is nonnegative implies that $e_g(x)$ is nondecreasing, and a nondecreasing $e_g(x)$ requires that,

$$\frac{f(x+2)f(x)}{f^2(x+1)} \geq \frac{x+1}{x+2}.$$

Directly estimating f does not ensure satisfaction of this or other conditions imposed by the representation, $e_g(x) = \frac{\int \theta^{x+1} g(\theta)\, d\theta}{\int \theta^{x} g(\theta)\, d\theta}$. For example, nonnegativity of the posterior fourth central moment also imposes restrictions on f. These and other restrictions are automatically satisfied in g-modeling, but require considerable machinations to be satisfied in f-modeling.

2.2 Corbet's Butterfly Data

Efron analyzes Corbet's butterfly data, comparing versions of f- and g-modeling for the Poisson rate parameter (θ) and its logarithm (λ). For comparison, I base g-modeling on the nonparametric, maximum likelihood estimate (NPML), implemented by the EM algorithm (Laird, 1982), starting the recursion with a sequence of 24 equiprobable mass points in the interval [0.1 to 36.0]. The recursion quickly converged the the three-point distribution in Table 1. It induces an X-marginal which is graphically close to the natural spline Poisson regression fit in Figure 4 of the article, but it gives less weight to small θ-values.

The g-NPML prior generates the posterior mean plots in Figure 1. In the left panel, the g-NPML line mimics the Robbins values displayed in Efron's Figure 5, but is monotone and respects other conditions imposed by the g-modeling approach. Zipf's/g-glm are plotted as a single line, even though in Efron's Figure 5 g-glm is slightly below Zipf's for large values of X.

Table 2 gives the (estimated) Bayes risk for the Robbins, g-NPML and g-glm priors with g-NPML less optimistic than Robbins, but more optimistic than g-glm.