STAT J535: Chapter 6(b): Assessing Model Quality

David B. Hitchcock
E-Mail: hitchcock@stat.sc.edu

Spring 2012

Checking Model Adequacy

- Checking the adequacy of a Bayesian model involves:
  1. determining how sensitive the posterior is to the specification of the prior and the likelihood
  2. checking that the values we obtain in our sample fit those we would expect to see, given our posterior knowledge
  3. checking robustness to individual data values

Sensitivity Analysis

- Checking the sensitivity to the specification of the data model/likelihood should be done regularly, but rarely is.
- We might examine the effect on the posterior of choosing related data models (e.g., Poisson vs. negative binomial for count data).
- Far more often, we check the sensitivity of the posterior to the **prior** specification.
- Assume Poisson(θ_1) and Poisson(θ_2) models for the data.
- We might ask: What happens to the posterior when we:
  1. change the functional form of the prior?
  2. keep the same form, but change the parameter(s) of the prior?
- If the posterior is **robust** to such changes in the prior, we may be more comfortable with the posterior inferences we make.

Sensitivity Analysis

Example 1(a): Consider $X_1, \ldots, X_n \stackrel{\text{iid}}{\sim} N(\mu, \sigma^2)$ with σ^2 known.

- The conjugate prior for μ is $\mu \sim N(\delta, \tau^2)$.
- A noninformative prior for μ is $p(\mu) = 1$.
- Another choice of prior for μ might be a t-distribution centered at δ.
- How would the posterior change for these 3 prior choices?
- We could examine (1) plots of the posterior in each case, or (2) several posterior quantiles in each case.
- See `WinBUGS` example with Kenya lead data.

Local Sensitivity Analysis

- Unfortunately, it may be too difficult to examine a large class of prior specifications, especially when the target parameter $\boldsymbol{\theta}$ is multidimensional.
- **Local** sensitivity analysis simply focuses on how changes in the hyperparameter value(s) affect the posterior.
- **Example 1(a)**: $X_1, \ldots, X_n \overset{\text{iid}}{\sim} N(\mu, \sigma^2)$, σ^2 known.
- Conjugate prior for μ: $\mu \sim N(\delta, \tau^2)$
- Compare resulting posterior (the plot and/or quantiles) to the posterior from these priors:

$$\mu \sim N(\delta - \tau, \tau^2)$$
$$\mu \sim N(\delta + \tau, \tau^2)$$
$$\mu \sim N(\delta, 0.5\tau^2)$$
$$\mu \sim N(\delta, 2\tau^2)$$

See R example.

Local Sensitivity Analysis

- **Example 1(b)**: $X_1, \ldots, X_{200}$ are annual deaths from horse kicks for 10 Prussian cavalry corps for each of 20 years.
- Let $X_i \overset{\text{iid}}{\sim}$ Poisson(λ), and let $\lambda \sim$ Gamma(α, β) be the prior.
- Compare posteriors from these priors for λ:

$$\lambda \sim \text{Gamma}(2, 4)$$
$$\lambda \sim \text{Gamma}(4, 8)$$
$$\lambda \sim \text{Gamma}(1, 2)$$
$$\lambda \sim \text{Gamma}(0.1 \times 2, \sqrt{0.1} \times 4)$$
$$\lambda \sim \text{Gamma}(3 \times 2, \sqrt{3} \times 4)$$

See `R` example with Prussian horse kick data.
General recommendation when the posterior is highly sensitive to changes in prior specification:Choose a more "objective" prior (or be prepared to defend your prior knowledge!).