Reliability

It is ~~a~~ good practice to report a pt estimate (e.g. posterior mean, posterior mode), together w/ a measure of reliability, e.g. CI, HPD region, posterior variance, interval from Laplace approx.).

Confidence Interval

Frequentist CI is a probability statement about the interval.

$$p(\underbrace{\ell(Y)}_{\text{random}} < \theta < \underbrace{u(Y)}_{\text{random}} \mid \theta) = 1-\alpha$$

Bayesian CI is a probability statement about θ.

i.e. $p(\underbrace{\ell(y)}_{\text{obs. data}} < \theta < u(y) \mid Y=y) = 1-\alpha$

Ex: beta-binomial

likelihood : $p(y_1, \dots, y_n | \theta) = f(\theta) \propto \theta^{\Sigma y_i} (1-\theta)^{n - \Sigma y_i}$

prior : $p(\theta) \propto \theta^{a-1} (1-\theta)^{b-1}$

posterior $\propto$ likelihood $\times$ prior

$\sim \theta^{\boxed{a + \Sigma y_i} - 1} (1-\theta)^{\boxed{b + n - \Sigma y_i} - 1}$

$\theta | y_1, \dots, y_n \sim \text{beta}(\underbrace{a + \Sigma y_i}_{\alpha_n}, \underbrace{b + n - \Sigma y_i}_{\beta_n})$

qbeta(c(0.025, 0.975), α_n, β_n)

Ex: $\theta \sim \text{beta}(2, 2)$ prior

$n = 15$

$\Sigma y_i = 10$ data

$\theta | y_1, \dots, y_n \sim \text{beta}(\underbrace{2 + 10}_{12}, \underbrace{2 + 15 - 10}_{7})$

$E\, \theta | y_1, \dots, y_n = \frac{12}{12+7} = \frac{12}{19}$

$p(.41 < \theta < .82) = .95$

HPD: $\int_{\theta} p(\theta|y_1,\dots,y_n)\, d\theta \cdot \mathbb{1}_{p(\theta|\vec{y}) \geq c} = 1-\alpha$

Laplace Approximation

idea: fit a Gaussian to mode of posterior.

Method: Taylor expand $\log p(\theta|\vec{y})$ about $\hat{\theta}_{MAP}$

$\hat{\theta}_{MAP}$: $\frac{d}{d\theta} \log p(\theta|\vec{y}) = 0$ and locally concave; $\frac{d^2}{d\theta^2} \log p(\theta|\vec{y})\Big|_{\hat{\theta}=\varepsilon} < 0$

Define: $\log p(\theta|\vec{y}) \equiv L(\theta)$ for convenience

$$L(\theta) \approx L(\hat{\theta}) + \underbrace{L'(\hat{\theta})}_{0}(\theta-\hat{\theta}) + \frac{L''(\hat{\theta})}{2}(\theta-\hat{\theta})^2$$

$$p(\theta|\vec{y}) = \exp\{L(\theta)\} = e^{L(\hat{\theta})} \cdot \underbrace{e^{\frac{1}{2}L''(\hat{\theta})(\theta-\hat{\theta})^2}}_{\text{kernel of a normal}}$$

kernel of a normal
w/ mean $= \hat{\theta}$
& variance $= -1/L''(\hat{\theta})$

$\theta|\vec{y} \approx N\left(\hat{\theta}_{MAP}, -1/L''(\hat{\theta})\right)$

beta binomial

$$p(\theta \mid \vec{y}) = c\,\theta^{\alpha-1}(1-\theta)^{\beta-1}$$

$$L(\theta) = \log p(\theta \mid \vec{y}) = \log c + (\alpha-1)\log\theta + (\beta-1)\log(1-\theta)$$

$$(1-\theta)(\alpha-1) = \theta(\beta-1)$$

$$L'(\theta) = \frac{\alpha-1}{\theta} - \frac{\beta-1}{1-\theta}$$

$$L''(\theta) = \frac{-(\alpha-1)}{\theta^2} - \frac{\beta-1}{(1-\theta)^2}$$

Set $L'(\theta) = 0$ then $\hat{\theta}_{MAP} = \frac{\alpha-1}{\beta+\alpha-2}$

$$p(\theta \mid \vec{y}) \approx N\left(\hat{\theta}_{MAP}, \ {}^{-1}/L''(\hat{\theta}_{MAP})\right)$$