The Normal-Gamma Model

PUBH 8442: Bayes Decision Theory and Data Analysis

Eric F. Lock
UMN Division of Biostatistics, SPH
elock@umn.edu

03/8/2021

Conjugate normal: σ^2 known

- Assume $\mathbf{y} = y_1, \ldots, y_n$ are iid Normal(μ, σ^2), and σ^2 known.
- If $p(\mu) = \text{Normal}(\mu_0, \tau^2)$, then

$$p(\mu \mid \mathbf{y}) = \text{Normal}\left(\frac{\sigma^2\mu_0 + n\tau^2\bar{y}}{\sigma^2 + n\tau^2}, \frac{\sigma^2\tau^2}{\sigma^2 + n\tau^2}\right)$$

- The *precision* is the reciprocal of the variance.
- Here, the posterior precision is $\frac{1}{\tau^2} + \frac{n}{\sigma^2}$:

$$\left(\frac{1}{\tau^2} + \frac{n}{\sigma^2}\right)^{-1} = \frac{\sigma^2\tau^2}{\sigma^2 + n\tau^2}$$

  - $\frac{1}{\tau^2}$ is the *prior precision*
  - $\frac{n}{\sigma^2}$ is the *data precision*

Conjugate normal: μ known

- Assume $\mathbf{y} = y_1, \ldots, y_n$ are iid Normal(μ, σ^2), with μ known and σ^2 unknown.
- Would like to make inference about σ^2
- Sampling model / likelihood:

$$p(\mathbf{y} \mid \sigma^2) \propto \left(\sigma^2\right)^{-n/2} \exp\left\{-\frac{n}{2\sigma^2}\left[\frac{1}{n}\sum_{i=1}^{n}(y_i - \mu)^2\right]\right\}$$

- The conjugate prior for σ^2 is the *inverse gamma* (IG) distribution
- If X $\sim$ Gamma(α, β), then $1/X \sim IG(\alpha, \beta)$.

Conjugate normal: μ known

- For prior $p(\sigma^2) = IG(\alpha, \beta)$, $\alpha > 0, \beta > 0$,

$$p(\sigma^2) = \frac{\beta^\alpha}{\Gamma(\alpha)} (\sigma^2)^{-(\alpha+1)} e^{-(\beta/\sigma^2)}$$

- Then, the posterior for σ^2 is

$$p(\sigma^2 \mid \mathbf{y}) = IG\left(\alpha + \frac{n}{2}, \beta + \frac{n}{2} W\right)$$

where $W = \frac{1}{n} \sum (y_i - \mu)^2$. **Conjugate!**

$$\propto p(y \mid \sigma^2) \cdot p(\sigma^2)$$

Conjugate normal: μ known

- For $\sigma^2 \sim IG(\alpha, \beta)$, the prior mean and variance are

$$E(\sigma^2) = \frac{\beta}{\alpha - 1} \quad \text{for } \alpha > 1$$

$$Var(\sigma^2) = \frac{\beta^2}{(\alpha - 1)^2(\alpha - 2)} \quad \text{for } \alpha > 2$$

- Solving for α, β:

$$\alpha = \frac{[E(\sigma^2)]^2}{var(\sigma^2)} + 2 \quad \text{and} \quad \beta = E(\sigma^2)\left\{\frac{[E(\sigma^2)]^2}{var(\sigma^2)} + 1\right\}$$

- Could determine α, β by prior guesses for $E(\sigma^2)$ and $var(\sigma^2)$.

Conjugate normal: μ and σ^2 unknown

- Assume $\mathbf{y} = y_1, \ldots, y_n$ are iid Normal(μ, σ^2), with μ unknown and σ^2 unknown.
- A conjugate prior for μ and σ^2 is

$$p(\sigma^2) \sim IG(\alpha, \beta)$$

$$p(\mu \mid \sigma^2) \sim \text{Normal}(\mu_0, \sigma^2/\lambda)$$

- Note: prior for μ depends on σ^2.
- λ reflects confidence/influence of prior (larger λ, more confidence in prior hypermean μ_0).

Conjugate normal: μ and σ^2 unknown

- Given the normal-normal model with σ^2 known,

$$p(\mu \mid \sigma^2, \mathbf{y}) = N\left(\frac{\lambda\mu_0 + n\bar{y}}{n+\lambda}, \frac{\sigma^2}{n+\lambda}\right)$$

- And, with μ known,

$$p(\sigma^2 \mid \mu, \mathbf{y}) = IG\left(\alpha + \frac{n}{2} + \frac{1}{2}, \beta + \frac{n}{2}W + \lambda\frac{(\mu - \mu_0)^2}{2}\right)$$

- One can show, integrating out μ, that

$$p(\sigma^2 \mid \mathbf{y}) = IG\left(\alpha + \frac{n}{2}, \beta + \frac{1}{2}\sum(y_i - \bar{y})^2 + \frac{\lambda n(\bar{y} - \mu_0)^2}{2(\lambda + n)}\right)$$

$$\bar{y} := \hat{y}$$

$$p(\bar{y}, \sigma^2, \mu) = p(\bar{y} \mid \mu, \sigma^2) \cdot p(\mu \mid \sigma^2) \cdot p(\sigma^2)$$

$$p(\sigma^2 \mid \mu, \bar{y}) \propto p(\bar{y}, \sigma^2, \mu)$$

$$p(\sigma^2 \mid \bar{y}) \propto \int p(\bar{y}, \sigma^2, \mu) \, d\mu$$