

STAT 713 sp 2023 Lec 05 slides

Bayes estimators

Karl B. Gregory

University of South Carolina

These slides are an instructional aid; their sole purpose is to display, during the lecture, definitions, plots, results, etc. which take too much time to write by hand on the blackboard. They are not intended to explain or expound on any material.

Suppose we observe data \mathbf{X} with distribution F depending on $\theta \in \Theta$.

Two paradigms:

- *Frequentist*: Treat θ as a fixed constant and write $\mathbf{X} \sim F(\mathbf{X}; \theta)$.
- *Bayesian*: Treat θ as a random variable and write $\mathbf{X}|\theta \sim F(\mathbf{X}|\theta)$.



The hierarchical model

Let $\mathbf{X} = (X_1, \dots, X_n)$ be our data and let θ be a rv in $\Theta \subset \mathbb{R}$. Assume

$$\begin{aligned}\mathbf{X}|\theta &\sim f(\mathbf{x}|\theta) \\ \theta &\sim \pi(\theta),\end{aligned}$$

where

- $f(\cdot|\theta)$ is the joint pmf/pdf of \mathbf{X} , conditional on θ (*data distribution*).
- $\pi(\cdot)$ is the marginal pmf/pdf of the parameter θ (*prior distribution*).

Posterior distribution

The *posterior distribution* of θ is the distribution of θ conditional on the data \mathbf{X} .

If $f(\mathbf{X}|\theta)$ is the joint pdf/pmf of $\mathbf{X}|\theta$, the conditional pdf/pmf of $\theta|\mathbf{X}$ is given by

$$\pi(\theta|\mathbf{X}) = \begin{cases} \frac{f(\mathbf{X}|\theta)\pi(\theta)}{\int_{\Theta} f(\mathbf{X}|\tilde{\theta})\pi(\tilde{\theta})d\tilde{\theta}} & \text{if } \theta \text{ is continuous} \\ \frac{f(\mathbf{X}|\theta)\pi(\theta)}{\sum_{\tilde{\theta} \in \Theta} f(\mathbf{X}|\tilde{\theta})\pi(\tilde{\theta})} & \text{if } \theta \text{ is discrete.} \end{cases}$$

Exercise: Derive the above.

Bayesian estimation with posterior mean

A typical Bayesian estimator of θ is the *posterior mean* of θ , which is

$$\hat{\theta}_{\text{Bayes}} = \mathbb{E}[\theta|\mathbf{X}] = \begin{cases} \int_{\Theta} \theta \cdot \pi(\theta|\mathbf{X}) d\theta & \text{if } \theta \text{ is continuous} \\ \sum_{\theta \in \Theta} \theta \cdot \pi(\theta|\mathbf{X}) & \text{if } \theta \text{ is discrete.} \end{cases}$$

Could also use the posterior mode or median.

Exercise: Show that the above estimator is the solution to


$$\hat{\theta}_{\text{Bayes}} = \underset{a}{\operatorname{argmin}} \mathbb{E}[(\theta - a)^2|\mathbf{X}].$$

Exercise: Suppose we have

$$Y|p \sim \text{Binomial}(n, p)$$
$$p \sim \text{Beta}(\alpha, \beta)$$

for some $\alpha > 0$ and $\beta > 0$.

- 1 Find the posterior distribution of $p|Y$.
- 2 Find an expression for $\hat{p}_{\text{Bayes}} = \mathbb{E}[p|Y]$.
- 3 Suppose $n = 10$ and $Y = 3$ is observed. Find the posterior mean of p under
 - ▶ $\alpha = 1, \beta = 1$
 - ▶ $\alpha = 4, \beta = 4$
 - ▶ $\alpha = 4, \beta = 10$



Shortcut to finding the posterior distribution

- 1 Find $g(\theta)$ such that $f(\mathbf{X}|\theta)\pi(\theta) \propto g(\theta)$.
- 2 Then $\pi(\theta|\mathbf{X}) = C \cdot g(\theta)$, where $C = [\int_{\Theta} g(\theta)d\theta]^{-1}$.

Can sometimes find C in Step 2 by identifying the distribution.

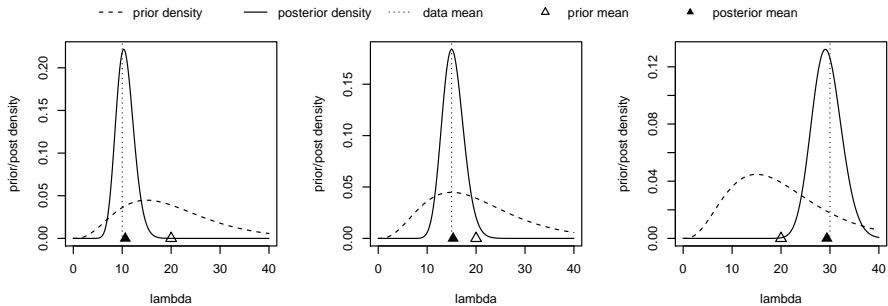
Exercise: Find the posterior distribution in the previous example with this strategy.

Exercise: Assume for $\mathbf{X} = (X_1, \dots, X_n)$ the hierarchical model

$$\begin{aligned}\mathbf{X}|\lambda &\stackrel{\text{ind}}{\sim} \text{Poisson}(\lambda) \\ \lambda &\sim \text{Gamma}(\alpha, \beta)\end{aligned}$$

for some $\alpha > 0$ and $\beta > 0$.

- 1 Find the posterior distribution of $\lambda|\mathbf{X}$.
- 2 Find an expression for $\hat{\lambda}_{\text{Bayes}} = \mathbb{E}[\lambda|\mathbf{X}]$.
- 3 Under $\alpha = 4$ and $\beta = 5$ and a sample size of $n = 3$, compute the posterior mean of $\lambda|\mathbf{X}$ when $\bar{X}_n = 10, 15, 30$.



Theorem (Posterior distribution depends on data via a suff. stat.)

If, given θ , $T(\mathbf{X})$ is a sufficient statistic for θ , then the distribution of $\theta|\mathbf{X}$ is the same as the distribution of $\theta|T(\mathbf{X})$.

Exercise: Show why using the factorization theorem.

This means we have $\pi(\theta|\mathbf{X}) \propto f_T(T(\mathbf{X})|\theta)\pi(\theta)$, provided $T(\mathbf{X})$ is suff.

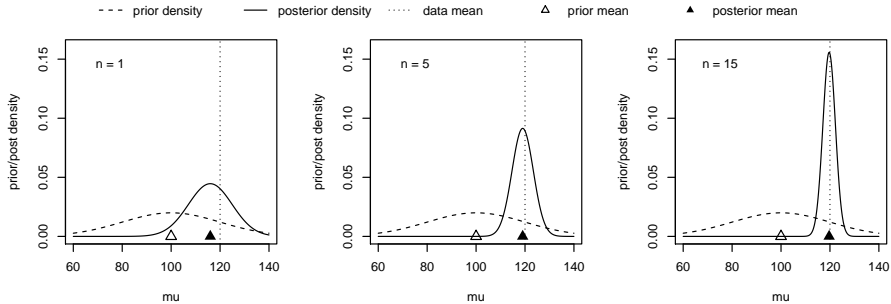
This can simplify calculations.

Exercise: Assume for $\mathbf{Y} = (Y_1, \dots, Y_n)$ the hierarchical model

$$\begin{aligned}\mathbf{Y}|\mu &\stackrel{\text{ind}}{\sim} \text{Normal}(\mu, \sigma^2) \\ \mu &\sim \text{Normal}(\mu_0, \tau^2)\end{aligned}$$

for some $\mu_0 \in \mathbb{R}$ and $\tau^2 > 0$, where σ^2 is a known constant.

- 1 Find the posterior distribution of $\mu|\mathbf{Y}$.
- 2 Find an expression for $\hat{\mu}_{\text{Bayes}} = \mathbb{E}[\mu|\mathbf{Y}]$.
- 3 Let $\sigma = 10$, $\tau = 20$, and $\mu_0 = 100$. Suppose that under sample sizes of $n = 1, 5, 15$, the sample mean $\bar{Y}_n = 120$ is observed. Compute the posterior mean of $\mu|\mathbf{Y}$ in each case.



Conjugacy

In the Bayesian setup, a *conjugate prior* is a prior distribution under which the posterior distribution belongs to the same family of distributions.

Examples we have seen:

- In Beta-Binomial model, posterior was Beta.
- In Gamma-Poisson model, posterior was Gamma.
- In Normal-Normal model, posterior was Normal.

We do not *need* to choose a conjugate prior; it just makes hand calculations easier.

Conjugacy mostly for textbook examples. In real life we use  power.

Exercise: Assume for $\mathbf{X} = (X_1, \dots, X_n)$ the hierarchical model

$$\begin{aligned}\mathbf{X}|\lambda &\stackrel{\text{ind}}{\sim} \text{Poisson}(\lambda) \\ \lambda &\sim \text{Uniform}(0, \lambda_{\max})\end{aligned}$$

for some $\lambda_{\max} > 0$.

- 1 Find the posterior density of $\lambda|\mathbf{X}$.
- 2 Is the prior a conjugate prior?
- 3 What happens to the posterior distribution of $\lambda|\mathbf{X}$ when $\lambda_{\max} \rightarrow \infty$?