Lecture 21: Bayesian estimation

- With squared error loss, data X, prior $\pi(\theta)$ the Bayes estimate of θ is

$$\mathrm{E}\left(\theta|X\right).$$

- For data $X = (X_1, \ldots, X_n)$ with the X_i iid Bernoulli(θ) and a Beta(α, β) prior I computed the joint density of X and θ, the marginal density of X and the posterior density of θ given X.
- Using $S = \sum_i X_i$:

$$f(x|\theta)\pi(\theta) = \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)}\theta^{\alpha+S-1}(1-\theta)^{\beta+n-S-1}$$

$$f(x) = \frac{\Gamma(\alpha+\beta)\Gamma(\alpha+S)\Gamma(\beta+n-S)}{\Gamma(\alpha)\Gamma(\beta)\Gamma(\alpha+\beta+n)}$$

$$\pi(\theta|X) = \frac{\Gamma(\alpha+\beta+n)}{\Gamma(\alpha+S)\Gamma(\beta+n-S)}\theta^{\alpha+S-1}(1-\theta)^{\beta+n-S-1}$$

- I did the $N(\mu, 1)$ with a $N(\nu, \tau^2)$ prior on μ.

Lecture 21

- In both cases the estimate is a weighted average of the usual estimate and the prior mean.
- The Bernoulli (Binomial) example illustrates the way we deal with mixed discrete and continuous variables.
- In the normal example I looked at completing the square.
- Many prediction, interpolation and extrapolation schemes are based on conditional expectations.
- I looked at a couple of improper priors leading, via Bayes' formula, to real posteriors.

Coverage in text and notes

- Chapters 11 and 12.
- Course slides on Bayes estimation.