# Bayesian data analysis in practice: <br> Three simple examples 

Martin P. Tingley

## Introduction

These notes cover three examples I presented at Climatea on 25 October 2011. Matlab code is available by request to demonstrate the ideas in Sections 1 and 2, and to fit the model in Section 3 and perform various manipulations on the draws from the posterior.

Most of what follows is based on Bayes' rule:

$$
\begin{equation*}
P(A \mid B) \propto P(B \mid A) \cdot P(A) \tag{1}
\end{equation*}
$$

where the term to the left is called the posterior, and the terms to the right the likelihood and the prior, respectively.

Another helpful identity is,

$$
\begin{equation*}
P(A, B)=P(A \mid B) \cdot P(B) . \tag{2}
\end{equation*}
$$

In what follows, the short-hand notation $[A]$ denotes the distribution of the random variable $A$, while $[A \mid B]$ denotes the distribution of $A$ conditional on $B,[A \mid \cdot]$ the distribution of $A$ conditional on all other variables in the model, and $[A, B]$ the joint distribution of $A$ and $B$.

For additional information, see Gelman et al. [2003], Banerjee et al. [2004], Tingley and Huybers [2010a,b]

## 1 Bivariate normal with known covariance

Say we have $N$ independent and identically distributed (IID) draws from a bivariate normal distribution, and we seek inference on the mean vector. For the sake of simplicity, assume that the covariance matrix is known (put differently, all results in this section are conditional on the true value of the covariance matrix).

We'll use a normal prior, as it is conjugate (see below).
Data:

$$
\begin{equation*}
\mathbf{Y}_{i} \mid \boldsymbol{\mu} \sim \mathcal{N}(\boldsymbol{\mu}, \Sigma), \quad i=1 \ldots N . \tag{3}
\end{equation*}
$$

Prior:

$$
\begin{equation*}
\boldsymbol{\mu} \sim \mathcal{N}(\boldsymbol{\eta}, \Gamma) \tag{4}
\end{equation*}
$$

Apply Bayes' rule to solve for the posterior distribution of $\boldsymbol{\mu}$, using the notation $\mathbf{Y}_{\text {all }}$ to denote the collection of $N$ observations $\mathbf{Y}_{i}$ :

$$
\begin{equation*}
P\left(\boldsymbol{\mu} \mid \mathbf{Y}_{\text {all }}\right) \propto P\left(\mathbf{Y}_{\text {all }} \mid \boldsymbol{\mu}\right) \cdot P(\boldsymbol{\mu}) \tag{5}
\end{equation*}
$$

Now substitute in the likelihood and the prior, which are both normal, discarding all leading terms that do not contain $\boldsymbol{\mu}$ :

$$
\begin{equation*}
P\left(\boldsymbol{\mu} \mid \mathbf{Y}_{\text {all }}\right) \propto \prod_{1}^{N} \exp \left\{-\frac{1}{2}\left(\mathbf{Y}_{i}-\boldsymbol{\mu}\right)^{T} \Sigma^{-1}\left(\mathbf{Y}_{i}-\boldsymbol{\mu}\right)\right\} \cdot \exp \left\{-\frac{1}{2}(\boldsymbol{\mu}-\boldsymbol{\eta})^{T} \Gamma^{-1}(\boldsymbol{\mu}-\boldsymbol{\eta})\right\} \tag{6}
\end{equation*}
$$

Turn the product into a sum inside the exponential, use the notation $\overline{\mathbf{Y}}$ to denote the mean of the $\mathbf{Y}_{i}$, expand all terms, and discard any that do not contain $\boldsymbol{\mu}$ :
$P\left(\boldsymbol{\mu} \mid \mathbf{Y}_{\text {all }}\right) \propto \exp \left\{-\frac{1}{2}\left(-N \overline{\mathbf{Y}}^{T} \Sigma^{-1} \boldsymbol{\mu}-N \boldsymbol{\mu}^{T} \Sigma^{-1} \overline{\mathbf{Y}}+\boldsymbol{\mu}^{T} \Sigma^{-1} \boldsymbol{\mu}-\boldsymbol{\mu}^{T} \Gamma^{-1} \boldsymbol{\eta}-\boldsymbol{\eta} \boldsymbol{\Gamma}^{-1} \boldsymbol{\mu}^{T}+\boldsymbol{\mu}^{T} \Gamma^{-1} \boldsymbol{\mu}\right)\right\}$.

The exponent is quadratic in $\boldsymbol{\mu}$, so immediately we recognize this as a normal distribution. All that remains is to complete the square to determine the mean and covariance. If you've never done this calculation, you should! Here's the answer:

$$
\begin{equation*}
\boldsymbol{\mu} \mid \mathbf{Y}_{\text {all }} \sim \mathcal{N}\left(\left(N \cdot \Sigma^{-1}+\Gamma^{-1}\right)^{-1}\left(N \cdot \Sigma^{-1} \overline{\mathbf{Y}}+\Gamma^{-1} \boldsymbol{\eta}\right),\left(N \cdot \Sigma^{-1}+\Gamma^{-1}\right)^{-1}\right) . \tag{8}
\end{equation*}
$$

Recall that the prior for $\boldsymbol{\mu}$ was also normal. The property of the prior and posterior having the same distributional form is known as conjugacy, and the use of conjugate or conditionally conjugate priors generally results in simpler calculations and computations. Priors are often chosen to be conjugate or conditionally conjugate for these practical (rather than scientific) reasons.

The posterior mean is an inverse variance weighted average of the mean of the data and the prior mean. To facilitate interpretation, make the simplifying assumption that $\Gamma=\frac{1}{c} \Sigma$, where $c$ is some constant. In other words, the prior covariance matrix is proportional to the covariance of the observations. In this case, the posterior simplifies to,

$$
\begin{equation*}
\boldsymbol{\mu} \left\lvert\, \mathbf{Y}_{\text {all }} \sim \mathcal{N}\left(\frac{N \overline{\mathbf{Y}}+c \boldsymbol{\eta}}{N+c}, \frac{1}{N+c} \Sigma\right)\right. \tag{9}
\end{equation*}
$$

The prior is thus equivalent to taking $c$ additional observations, with a mean of $\boldsymbol{\eta}$. If $c$ is small relative to $N$, the prior adds little, while if $c$ is large relative to $N$, the prior can dominate the posterior.

For general $\Gamma$, the prior can be interpreted as adding the same information as an additional observation, with value $\boldsymbol{\eta}$ and covariance matrix $\Gamma$.

Matlab Demo: Relative influence of prior and likelihood for bivariate normal.

## 2 Gibbs sampler for the bivariate normal

The end goal of (most) Bayesian data analysis is to infer the posterior distribution of a set of parameters. In the case presented in Section 1, the joint posterior was bivariate normal - a wellknown distributional form. In many cases (see Section 3 for an example), the joint distribution of the unknowns does not follow a standard distribution, and it becomes necessary to turn to various computational tools to draw samples from the posterior. The simplest such tool is the Gibbs sampler.

Say we have some nasty expression for the joint posterior of three random variables, $Y_{1}, Y_{2}$, and $Y_{3}$. The posterior does not follow a known distribution, so we can't just tell Matlab to produce samples for us. However, it may be the case that the distributions $\left[Y_{1} \mid Y_{2}, Y_{3}\right],\left[Y_{2} \mid Y_{1}, Y_{3}\right]$, and [ $\left.Y_{3} \mid Y_{1}, Y_{3}\right]$ (the full conditional posteriors; often abbreviated as $\left[Y_{1} \mid \cdot\right]$ and so on) are all well-known forms. In this case, the Gibbs sampler provides us with a recipe for producing (correlated) samples from the joint posterior of $\mathbf{Y}$ :

1. Produce/guess some initial value for each element of $\mathbf{Y}$, call them $\mathbf{Y}^{0}$.
2. For each sample $k=1 \ldots N$, draw $Y_{1}^{k}$ from $\left[Y_{1}^{k} \mid Y_{2}^{k-1}, Y_{3}^{k-1}\right]$, draw $Y_{2}^{k}$ from $\left[Y_{2}^{k} \mid Y_{1}^{k}, Y_{3}^{k-1}\right]$, and draw $Y_{3}^{k}$ from $\left[Y_{3}^{k} \mid Y_{1}^{k}, Y_{2}^{k}\right]$.

In other words, draw from each conditional posterior in turn, using the most recent draws of the other variables.

The first time I built a Gibbs sampler I thought it was magic. Let's build one to sample from a bivariate normal. (Matlab can, of course, draw directly from the joint distribution in this case). Assume that,

$$
\mathbf{Y} \sim \mathcal{N}\left(\binom{\mu_{1}}{\mu_{2}},\left(\begin{array}{cc}
\sigma_{1}^{2} & \rho \sigma_{1} \sigma_{2}  \tag{10}\\
\rho \sigma_{1} \sigma_{2} & \sigma_{2}^{2}
\end{array}\right)\right)
$$

The distribution of $Y_{1}$ conditional on a particular value $y_{2}$ of the random variable $Y_{2}$ is univariate normal, and likewise for $Y_{2} \mid y_{1}$ :

$$
\begin{align*}
& Y_{2} \left\lvert\, y_{1} \sim \mathcal{N}\left(\mu_{2}+\rho \frac{\sigma_{2}}{\sigma_{1}}\left(y_{1}-\mu_{1}\right), \sigma_{2}^{2}\left(1-\rho^{2}\right)\right)\right.  \tag{11}\\
& Y_{1} \left\lvert\, y_{2} \sim \mathcal{N}\left(\mu_{1}+\rho \frac{\sigma_{1}}{\sigma_{2}}\left(y_{2}-\mu_{2}\right), \sigma_{1}^{2}\left(1-\rho^{2}\right)\right) .\right. \tag{12}
\end{align*}
$$

The Gibbs sampler then proceeds by first initializing with some value $y_{1}^{0}$ (or $y_{2}^{0}$ ), and then iteratively drawing from the conditional distributions, always conditioning on the most recent draw of the other element of $Y$.

If some/all of the conditional posteriors do not follow well-known distributions, then the MetropolisHastings algorithm, or a variety of other tools, can be used to produce samples. See references for more details.

Matlab Demo: Gibbs sampler for the bivariate normal.

## 3 A simple hierarchical spatial model

The goal in this section is to infer a complete spatial field from noisy, incomplete, point referenced data. The Matlab example will use data from a number of weather stations to infer monthly average temperatures on a fine grid for the state of Colorado.

Let the vector $\mathbf{Y}$ refer to the spatial field at a set of locations $\mathbf{s}_{i}, i=1 \ldots N$, which includes both the set of points where there are observations and the set of points where inference will be made.

At the process level, assume that $\mathbf{Y}$ is Gaussian, with constant mean and exponential covariance:

$$
\begin{align*}
\mathbf{Y} \mid \mu, \phi, \sigma^{2} & \sim \mathcal{N}\left(\mu \mathbf{1}_{N}, \Sigma\right)  \tag{13}\\
\Sigma_{i j} & =\sigma^{2} \exp \left(-\phi| | \mathbf{s}_{i}-\mathbf{s}_{j} \|\right) \tag{14}
\end{align*}
$$

where $\left\|\mathbf{s}_{i}-\mathbf{s}_{j}\right\|$ is the distance between locations $\mathbf{s}_{i}$ and $\mathbf{s}_{j}$, and $\mathbf{1}_{N}$ is a vector of ones. Note that $\mu$ can be thought of as a regression coefficient for the covariate $\mathbf{1}$. Conceptually, we could replace $\mu \mathbf{1}$ with $\mu \mathbf{1}+\beta_{1} \mathbf{X}_{1}+\beta_{2} \mathbf{X}_{2}+\ldots$, where the $\mathbf{X}_{i}$ are covariates such as elevation, latitude, or whatever else may be relevant. The machinery developed below can easily be generalized to this situation.

For the purposes of this example, we'll assume $\phi$ is known and fixed. Inference on $\phi$ is certainly possible, and the references provide a number of examples. However, inference requires use of the Metropolis-Hastings algorithm, which has not been covered in these notes. Think of the MetropolisHastings algorithm as a tool to learn and use when one or more of the full conditional posteriors does not follow a known form.

At the data level, assume that the observations $Z_{j}, j=1 \ldots M$ can be expressed as the corresponding value of the field, plus IID white noise:

$$
\begin{equation*}
Z_{i} \mid Y_{i}, \tau^{2} \sim \mathcal{N}\left(Y_{i}, \tau^{2}\right) \tag{15}
\end{equation*}
$$

By defining an $M$ by $N$ selection matrix H of zeros and ones, we can write this in vector form:

$$
\begin{equation*}
\mathbf{Z} \mid \mathbf{Y}, \tau^{2} \sim \mathcal{N}\left(\mathbf{H Y}, \tau^{2} \mathbf{I}_{M}\right) \tag{16}
\end{equation*}
$$

where $\mathrm{I}_{M}$ is the $M$ by $M$ identity matrix.
At the prior level, assume conditionally conjugate priors for $\mu, \sigma^{2}$ and $\tau^{2}$ :

$$
\begin{align*}
\mu & \sim \mathcal{N}\left(\eta, \delta^{2}\right)  \tag{17}\\
\sigma^{2} & \sim I G\left(a_{\sigma}, b_{\sigma}\right)  \tag{18}\\
\tau^{2} & \sim I G\left(a_{\tau}, b_{\tau}\right) \tag{19}
\end{align*}
$$

$I G$ refers to the Inverse Gamma distribution, which is the conjugate prior for the variance parameters; see the Appendix for more information.

The joint posterior then follows from Bayes' rule:

$$
\begin{equation*}
\left[\mathbf{Y}, \mu, \sigma^{2}, \tau^{2} \mid \mathbf{Z}\right] \propto\left[\mathbf{Z} \mid \mathbf{Y}, \mu, \sigma^{2}, \tau^{2}\right]\left[\mathbf{Y} \mid \mu, \sigma^{2}, \tau^{2}\right][\mu]\left[\sigma^{2}\right]\left[\tau^{2}\right] . \tag{20}
\end{equation*}
$$

In order to sample from the joint posterior, we make use of a Gibbs sampler. To derive the conditional posteriors, plug in the distributional forms, and then, for each unknown, consider the resulting expression as a function of only that parameter. See Section 2 for how to do so in the case of a Normal mean. Denote the correlation matrix by $R$, i.e. $R=\sigma^{-2} \Sigma$. The resulting conditional posteriors are,

$$
\begin{align*}
\mathbf{Y} \mid \cdot & \sim \mathcal{N}\left(\left(\Sigma^{-1}+\tau^{-2} \mathbf{H}^{T} \mathbf{I}_{M} \mathbf{H}\right)^{-1} \cdot\left(\mu \Sigma^{-1} \mathbf{1}_{N}+\tau^{-2} \mathbf{H}^{\prime} \mathbf{Z}\right),\left(\Sigma^{-1}+\tau^{-2} \mathbf{H}^{T} \mathbf{I}_{M} \mathbf{H}\right)^{-1}\right)  \tag{21}\\
\mu \mid \cdot & \sim \mathcal{N}\left(\left(\mathbf{1}_{N}^{T} \Sigma^{-1} \mathbf{1}_{N}+\delta^{-2}\right)^{-1} \cdot\left(\mathbf{1}_{N}^{T} \Sigma^{-1} \mathbf{Y}+\eta / \delta^{2}\right),\left(\mathbf{1}_{N}^{T} \Sigma^{-1} \mathbf{1}_{N}+\delta^{-2}\right)^{-1}\right)  \tag{22}\\
\sigma^{2} \mid \cdot & \sim I G\left(N / 2+a_{1}, \frac{1}{2}\left(\mathbf{Y}-\mu \mathbf{1}_{N}\right)^{T} \mathbf{R}^{-1}\left(\mathbf{Y}-\mu \mathbf{1}_{N}\right)+b_{1}\right)  \tag{23}\\
\tau^{2} \mid \cdot & \sim I G\left(M / 2+a_{2}, \frac{1}{2}(\mathbf{Z}-\mathbf{H} \mathbf{Y})^{T}(\mathbf{Z}-\mathbf{H} \mathbf{Y})+b_{2}\right) . \tag{24}
\end{align*}
$$

Suggestion: go through these calculations, and make sure I haven't made a mistake!
To sample from the joint posterior, we start with some initial guess of the parameters $\mu, \sigma^{2}$ and $\tau^{2}$, and then successively draw from each of the conditional posteriors. After discarding an initial series of draws to allow the chain to 'burn in', we can then use the resulting ensemble of draws to perform just about any analysis we can think of.

Matlab Demo: Implement this model for monthly Colorado temperature data. Note that the priors used in the demo for the parameters are strongly informative due to the paucity of data. It is worthwhile to change the priors (and also the value of $\phi$ ) and see how the results are affected.

This analysis can clearly be improved in a number of ways. We've left out some pretty important geographic co-variates, and have thrown away the temporal dimension of a space-time data set.

Assignment for the ambitious: Add in the relevant co-variates, include a temporal model with a seasonal cycle, and perform inference on $\phi$.

## Appendix: The Inverse-Gamma distribution

The pdf of the Inverse Gamma is of the form,

$$
\begin{equation*}
P\left(\sigma^{2} \mid a_{1}, b_{1}\right) \propto\left(\sigma^{2}\right)^{-a_{1}-1} \exp \left(-b_{1} / \sigma^{2}\right) . \tag{25}
\end{equation*}
$$

Note that this functional form is equivalent to treating the Normal pdf as a function of the variance parameter.

Consider estimating the variance from $M$ independent draws from a $\mathcal{N}\left(\mu, \sigma^{2}\right)$ distribution, where we assume that $\mu$ is known, and specify an $\operatorname{IG}(a, b)$ prior for $\sigma^{2}$. Think of this as one step
in a Gibbs sampler. Applying Bayes' rule, we have:

$$
\begin{equation*}
P\left(\sigma^{2} \mid, \mu, y_{i}, \ldots, y_{M}\right) \propto \prod_{i=1}^{M} P\left(y_{i} \mid \mu, \sigma^{2}\right) \cdot P\left(\sigma^{2} \mid a_{1}, b_{1}\right) \tag{26}
\end{equation*}
$$

Substituting in the distributional forms, we have,

$$
\begin{equation*}
P\left(\sigma^{2} \mid, \mu, y_{i}, \ldots, y_{M}\right) \propto\left(\sigma^{2}\right)^{-M / 2} \exp \left(-\frac{1}{\sigma^{2}} \frac{\sum_{i=1}^{M}\left(y_{i}-\mu\right)^{2}}{2}\right) \cdot\left(\sigma^{2}\right)^{-a_{1}-1} \exp \left(-b_{1} / \sigma^{2}\right) . \tag{27}
\end{equation*}
$$

Collecting terms yields,

$$
\begin{equation*}
P\left(\sigma^{2} \mid, \mu, y_{i}, \ldots, y_{m}\right) \propto\left(\sigma^{2}\right)^{-M / 2-a_{1}-1} \exp \left(-\frac{1}{\sigma^{2}}\left(\frac{\sum_{i=1}^{M}\left(y_{i}-\mu\right)^{2}}{2}+b_{1}\right)\right) \tag{28}
\end{equation*}
$$

The posterior is thus Inverse Gamma:

$$
\begin{equation*}
\sigma^{2} \mid, \mu, y_{i}, \ldots, y_{M} \sim I G\left(M / 2+a_{1}, \frac{1}{2} \sum_{i=1}^{M}\left(y_{i}-\mu\right)^{2}+b_{1}\right) . \tag{29}
\end{equation*}
$$

The Inverse Gamma prior for $\sigma^{2}$ can be interpreted as $2 a_{1}$ prior observations with an average squared deviation of $b_{1} / a_{1}$.

## References

S. Banerjee, B.P. Carlin, and A.E. Gelfand. Hierarchical modeling and analysis for spatial data. Chapman \& Hall, 2004.

Andrew Gelman, John B. Carlin, Hal S. Stern, and Donald B. Rubin. Bayesian Data Analysis. Chapman \& Hall/CRC, Boca Raton, 2 edition, 2003.
M.P. Tingley and P. Huybers. A Bayesian Algorithm for Reconstructing Climate Anomalies in Space and Time. Part 1: Development and applications to paleoclimate reconstruction problems. Journal of Climate, 23(10):2759-2781, 2010a.
M.P. Tingley and P. Huybers. A Bayesian Algorithm for Reconstructing Climate Anomalies in Space and Time. Part 2: Comparison with the Regularized Expectation-Maximization Algorithm. Journal of Climate, 23(10), 2010b.

