Part 1: Recap on the Bayesian paradigm

Jean-Michel Marin

Université de Montpellier
Institut Montpelliérain Alexander Grothendieck (IMAG)

September, 22-27, Fréjus, France

The Bayesian paradigm

Bayes theorem = Inversion of probabilities

If A and B are events such that $\mathbb{P}(B) \neq 0$,

$$\mathbb{P}(A|B) = \frac{\mathbb{P}(A \cap B)}{\mathbb{P}(B)} = \frac{\mathbb{P}(B|A)\mathbb{P}(A)}{\mathbb{P}(B)} =$$

$$\frac{\mathbb{P}(B|A)\mathbb{P}(A)}{\mathbb{P}(A)\mathbb{P}(B|A) + \mathbb{P}(\bar{A})\mathbb{P}(B|\bar{A})}$$

The Bayesian paradigm

Subjectivism

Frank Plumpton Ramsey (1903-1930)

Bruno de Finetti (1906-1985)

Leonard Jimmie Savage (1921-1971)

The Bayesian paradigm

Given an iid sample $\mathscr{D}_n = (x_1, \ldots, x_n)$ from a density $f(x|\theta)$, depending upon an unknown parameter $\theta \in \Theta$, the associated likelihood function is

$$\ell(\theta|\mathscr{D}_n) = \prod_{i=1}^{n} f(x_i|\theta)$$

The Bayesian paradigm

When $\mathscr{D}_n$ is a normal $\mathscr{N}(\mu, \sigma^2)$ sample of size n and $\theta = (\mu, \sigma^2)$, we get

$$\begin{aligned}
\ell(\theta|\mathscr{D}_n) &= \prod_{i=1}^{n} \exp\{-(x_i-\mu)^2/2\sigma^2\}/\sqrt{2\pi}\sigma \\
&\propto \exp\left\{-\sum_{i=1}(x_i-\mu)^2/2\sigma^2\right\}/\sigma^n \\
&\propto \exp\left\{-\left(n\mu^2 - 2n\bar{x}\mu + \sum_{i=1} x_i^2\right)/2\sigma^2\right\}/\sigma^n \\
&\propto \exp\left\{-\left[n(\mu-\bar{x})^2 + s^2\right]/2\sigma^2\right\}/\sigma^n,
\end{aligned}$$

$\bar{x}$ denotes the empirical mean and $s^2 = \sum_{i=1}^{n}(x_i - \bar{x})^2$

The Bayesian paradigm

In the Bayesian approach θ is considered as a random variable

The Bayesian paradigm

In the Bayesian approach θ is considered as a random variable

In some sense, the likelihood function is transformed into a *posterior* distribution, which is a valid probability distribution on Θ

$$\pi(\theta|\mathscr{D}_n) = \frac{\ell(\theta|\mathscr{D}_n)\pi(\theta)}{\int \ell(\theta|\mathscr{D}_n)\pi(\theta)\,\mathrm{d}\theta}$$

The Bayesian paradigm

In the Bayesian approach θ is considered as a random variable

In some sense, the likelihood function is transformed into a *posterior* distribution, which is a valid probability distribution on Θ

$$\pi(\theta|\mathscr{D}_n) = \frac{\ell(\theta|\mathscr{D}_n)\pi(\theta)}{\int \ell(\theta|\mathscr{D}_n)\pi(\theta)\,\mathrm{d}\theta}$$

$\pi(\theta)$ is called the *prior* distribution and it has to be chosen to start the analysis

The Bayesian paradigm

The posterior density is a probability density on the parameter, which does not mean the parameter θ need be a genuine random variable

The Bayesian paradigm

The posterior density is a probability density on the parameter, which does not mean the parameter θ need be a genuine random variable

This density is used as an inferential tool, not as a truthful representation

The Bayesian paradigm

Two motivations:

- the prior distribution summarizes the *prior information* on θ. However, the choice of $\pi(\theta)$ is often decided on practical grounds rather than strong subjective beliefs
- the Bayesian approach provides a fully probabilistic framework for the inferential analysis, with respect to a reference measure $\pi(\theta)$

The Bayesian paradigm

Suppose $\mathscr{D}_n$ is a normal $\mathscr{N}(\mu, \sigma^2)$ sample of size n

When σ^2 is known, if $\mu \sim \mathscr{N}\left(0, \sigma^2\right)$, then

$$\begin{aligned}
\pi(\mu|\mathscr{D}_n) &\propto \pi(\mu)\,\ell(\theta|\mathscr{D}_n) \\
&\propto \exp\{-\mu^2/2\sigma^2\}\ \exp\left\{-n(\bar{x}-\mu)^2/2\sigma^2\right\} \\
&\propto \exp\left\{-(n+1)\mu^2/2\sigma^2 + 2n\mu\bar{x}/2\sigma^2\right\} \\
&\propto \exp\left\{-(n+1)[\mu - n\bar{x}/(n+1)]^2/2\sigma^2\right\}
\end{aligned}$$

The Bayesian paradigm

Suppose $\mathscr{D}_n$ is a normal $\mathscr{N}(\mu, \sigma^2)$ sample of size n

When σ^2 is known, if $\mu \sim \mathscr{N}\left(0, \sigma^2\right)$, then

$$\begin{aligned}
\pi(\mu|\mathscr{D}_n) &\propto \pi(\mu)\,\ell(\theta|\mathscr{D}_n) \\
&\propto \exp\{-\mu^2/2\sigma^2\}\ \exp\left\{-n(\bar{x}-\mu)^2/2\sigma^2\right\} \\
&\propto \exp\left\{-(n+1)\mu^2/2\sigma^2 + 2n\mu\bar{x}/2\sigma^2\right\} \\
&\propto \exp\left\{-(n+1)[\mu - n\bar{x}/(n+1)]^2/2\sigma^2\right\}
\end{aligned}$$

$$\boxed{\mu|\mathscr{D}_n \sim \mathscr{N}\left(n\bar{x}/(n+1), \sigma^2/(n+1)\right)}$$

The Bayesian paradigm

When σ^2 is unknown, $\theta = (\mu, \sigma^2)$, if $\mu|\sigma^2 \sim \mathscr{N}\left(0, \sigma^2\right)$ and $\sigma^2 \sim \mathscr{I}\mathscr{G}(1, 1)$, then $\pi((\mu, \sigma^2)|\mathscr{D}_n) \propto \pi(\sigma^2) \times \pi(\mu|\sigma^2) \times f(\mathscr{D}_n|\mu, \sigma^2)$

$$\propto (\sigma^{-2})^{1/2+2} \exp\left\{-(\mu^2 + 2)/2\sigma^2\right\} \mathbf{1}_{\sigma^2>0}$$

$$(\sigma^{-2})^{n/2} \exp\left\{-\left(n(\mu - \overline{x})^2 + s^2\right)/2\sigma^2\right\}$$

The Bayesian paradigm

When σ^2 is unknown, $\theta = (\mu, \sigma^2)$, if $\mu|\sigma^2 \sim \mathscr{N}\left(0, \sigma^2\right)$ and $\sigma^2 \sim \mathscr{I}\mathscr{G}(1,1)$, then $\pi((\mu,\sigma^2)|\mathscr{D}_n) \propto \pi(\sigma^2) \times \pi(\mu|\sigma^2) \times f(\mathscr{D}_n|\mu,\sigma^2)$

$$\propto (\sigma^{-2})^{1/2+2} \exp\left\{-(\mu^2+2)/2\sigma^2\right\} \mathbf{1}_{\sigma^2>0}$$

$$(\sigma^{-2})^{n/2} \exp\left\{-\left(n(\mu-\overline{x})^2 + s^2\right)/2\sigma^2\right\}$$

$$\boxed{\mu|\mathscr{D}_n, \sigma^2 \sim \mathscr{N}\left(\frac{n\bar{x}}{n+1}, \frac{\sigma^2}{n+1}\right)}$$

$$\boxed{\sigma^2|\mathscr{D}_n \sim \mathscr{I}\mathscr{G}\left(\left\{1+\frac{n}{2}\right\}, \left\{1+\frac{s^2}{2}+\frac{n\bar{x}}{2(n+1)}\right\}\right)}$$

The Bayesian paradigm

Variability in σ^2 induces more variability in μ, the marginal posterior in μ being then a Student's t distribution

The Bayesian paradigm

Variability in σ^2 induces more variability in μ, the marginal posterior in μ being then a Student's t distribution

$$\boxed{\mu|\mathscr{D}_n \sim \mathscr{T}\left(n+2, \frac{n\bar{x}}{n+1}, \frac{2+s^2+(n\bar{x})/(n+1)}{(n+1)(n+2)}\right)}$$

Bayesian estimates

For a given loss function $\mathrm{L}\left(\theta, \hat{\theta}(\mathscr{D}_n)\right)$, we deduce a Bayesian estimate by minimizing the posterior expected loss:

$$\mathbb{E}^{\pi}_{\theta|\mathscr{D}_n}\left(\mathrm{L}\left(\theta, \hat{\theta}(\mathscr{D}_n)\right)\right)$$

Bayesian estimates

For a given loss function $\mathrm{L}\left(\theta, \hat{\theta}(\mathscr{D}_n)\right)$, we deduce a Bayesian estimate by minimizing the posterior expected loss:

$$\mathbb{E}^{\pi}_{\theta|\mathscr{D}_n}\left(\mathrm{L}\left(\theta, \hat{\theta}(\mathscr{D}_n)\right)\right)$$

To minimize the posterior expected loss is equivalent to minimize the Bayes risk, the frequentist risk integrated over the prior distribution

Bayesian estimates

For instance, for the L_2 loss function, the corresponding Bayes optimum is the expected value of θ under the posterior distribution,

$$\hat{\theta}(\mathscr{D}_n) = \int \theta\, \pi(\theta|\mathscr{D}_n)\, \mathrm{d}\theta = \frac{\int \theta\, \ell(\theta|\mathscr{D}_n)\, \pi(\theta)\, \mathrm{d}\theta}{\int \ell(\theta|\mathscr{D}_n)\, \pi(\theta)\, \mathrm{d}\theta}$$

Bayesian estimates

When no specific penalty criterion is available, the posterior expectation is often used as a default estimator, although alternatives are also available. For instance, the *maximum a posteriori estimator* (MAP) is defined as

$$\hat{\theta}(\mathscr{D}_n) \in \mathrm{argmax}_\theta \quad \pi(\theta|\mathscr{D}_n)$$

Bayesian estimates

When no specific penalty criterion is available, the posterior expectation is often used as a default estimator, although alternatives are also available. For instance, the *maximum a posteriori estimator* (MAP) is defined as

$$\hat{\theta}(\mathscr{D}_n) \in \text{argmax}_\theta \quad \pi(\theta|\mathscr{D}_n)$$

Similarity of with the maximum likelihood estimator: the influence of the prior distribution $\pi(\theta)$ on the estimate progressively disappears as the number of observations n increases

Conjugate prior

The selection of the prior distribution is an important issue in Bayesian statistics

Conjugate prior

The selection of the prior distribution is an important issue in Bayesian statistics

When prior information is available about the data or the model, it can be used in building the prior

Conjugate prior

The selection of the prior distribution is an important issue in Bayesian statistics

When prior information is available about the data or the model, it can be used in building the prior

In many situations, however, the selection of the prior distribution is quite delicate

Conjugate prior

The selection of the prior distribution is an important issue in Bayesian statistics

When prior information is available about the data or the model, it can be used in building the prior

In many situations, however, the selection of the prior distribution is quite delicate

Since the choice of the prior distribution has a considerable influence on the resulting inference, this inferential step must be conducted with the utmost care

Conjugate prior

Conjugate priors are such that the prior and posterior densities belong to the same parametric family

Conjugate prior

Conjugate priors are such that the prior and posterior densities belong to the same parametric family

An advantage when using a conjugate prior, is that one has to select only a few parameters to determine the prior distribution

Conjugate prior

Conjugate priors are such that the prior and posterior densities belong to the same parametric family

An advantage when using a conjugate prior, is that one has to select only a few parameters to determine the prior distribution

But the information known a priori may be either insufficient or incompatible with the structure imposed by conjugacy

Conjugate prior

Justifications

- Device of virtual past observations
- First approximations to adequate priors, backed up by robustness analysis
- But mostly... tractability and simplicity

Conjugate prior

$f(x\|\theta)$	$\pi(\theta)$	$\pi(\theta\|x)$
Normal $\mathcal{N}(\theta,\sigma^2)$	Normal $\mathcal{N}(\mu,\tau^2)$	$\mathcal{N}(\rho(\sigma^2\mu+\tau^2 x),\rho\sigma^2\tau^2)$ $\rho^{-1}=\sigma^2+\tau^2$
Poisson $\mathcal{P}(\theta)$	Gamma $\mathcal{G}(\alpha,\beta)$	$\mathcal{G}(\alpha+x,\beta+1)$
Gamma $\mathcal{G}(\nu,\theta)$	Gamma $\mathcal{G}(\alpha,\beta)$	$\mathcal{G}(\alpha+\nu,\beta+x)$
Binomial $\mathcal{B}(n,\theta)$	Beta $\mathcal{B}e(\alpha,\beta)$	$\mathcal{B}e(\alpha+x,\beta+n-x)$

Conjugate prior

$f(x\|\theta)$	$\pi(\theta)$	$\pi(\theta\|x)$
Negative Binomial $\mathcal{N}eg(m,\theta)$	Beta $\mathcal{B}e(\alpha,\beta)$	$\mathcal{B}e(\alpha+m,\beta+x)$
Multinomial $\mathcal{M}_k(\theta_1,\ldots,\theta_k)$	Dirichlet $\mathcal{D}(\alpha_1,\ldots,\alpha_k)$	$\mathcal{D}(\alpha_1+x_1,\ldots,\alpha_k+x_k)$
Normal $\mathcal{N}(\mu,1/\theta)$	Gamma $\mathcal{G}a(\alpha,\beta)$	$\mathcal{G}(\alpha+0.5,\beta+(\mu-x)^2/2)$

Noninformative prior

Conjugate priors are nice to work with, but require hyperparameters's determination

Noninformative prior

Conjugate priors are nice to work with, but require hyperparameters's determination

One can opt for a completely different perspective and rely on so-called *noninformative* priors that aim at attenuating the impact of the prior on the resulting inference

Noninformative prior

Conjugate priors are nice to work with, but require hyperparameters's determination

One can opt for a completely different perspective and rely on so-called *noninformative* priors that aim at attenuating the impact of the prior on the resulting inference

These priors are fundamentally defined as coherent extensions of the uniform distribution

Noninformative prior

For unbounded parameter spaces, the densities of noninformative priors actually may fail to integrate to a finite number and they are defined instead as positive measures

Noninformative prior

For unbounded parameter spaces, the densities of noninformative priors actually may fail to integrate to a finite number and they are defined instead as positive measures

Generalized Bayesian estimators with improper prior distributions

Noninformative prior

Location models $x|\theta \sim f(x - \theta)$ are usually associated with flat priors $\pi(\theta) \propto 1$

Noninformative prior

Location models $x|\theta \sim f(x-\theta)$ are usually associated with flat priors $\pi(\theta) \propto 1$

Scale models $x|\theta \sim \frac{1}{\theta} f\left(\frac{x}{\theta}\right)$ are usually associated with the log-transform of a flat prior, that is, $\pi(\theta) \propto 1/\theta \times \mathbf{1}_{\theta>0}$

Jeffreys prior

In a more general setting, the noninformative prior favored by most Bayesians is the so-called **Jeffreys prior** which is related to the Fisher information matrix

$$I_x^F(\theta) = -\mathbb{E}\left(\frac{\partial^2 \log f(x|\theta)}{(\partial\theta)^2}\right)$$

by

$$\pi^J(\theta) \propto \sqrt{|I_x^F(\theta)|} \times \mathbf{1}_{\theta\in\Theta}\,,$$

where $|I|$ denotes the determinant of the matrix I

Jeffreys prior

Suppose $\mathscr{D}_n$ is a normal $\mathscr{N}(\mu, \sigma^2)$ sample of size n and $\theta = (\mu, \sigma^2)$

Jeffreys prior

Suppose $\mathscr{D}_n$ is a normal $\mathscr{N}(\mu, \sigma^2)$ sample of size n and $\theta = (\mu, \sigma^2)$

The Fisher information matrix leads to the Jeffreys prior

$$\pi^J(\mu, \sigma^2) \propto 1/\{(\sigma^2)\}^{3/2}\mathbf{1}_{\sigma^2>0}$$

Jeffreys prior

Suppose $\mathscr{D}_n$ is a normal $\mathscr{N}(\mu, \sigma^2)$ sample of size n and $\theta = (\mu, \sigma^2)$

The Fisher information matrix leads to the Jeffreys prior

$$\pi^J(\mu, \sigma^2) \propto 1/\{(\sigma^2)\}^{3/2}\mathbf{1}_{\sigma^2>0}$$

$$\boxed{\mu|\sigma^2, \mathscr{D}_n \sim \mathscr{N}\left(\bar{x}, \sigma^2/n\right)}$$

$$\boxed{\sigma^2|\mathscr{D}_n \sim \mathscr{I}\mathscr{G}\left(n/2, s^2/2\right)}$$

$$\boxed{\mu|\mathscr{D}_n \sim \mathscr{T}\left(n, \bar{x}, s^2/n^2)\right)}$$

Bayesian Credible Intervals

Since the Bayesian approach processes θ as a random variable, a natural definition of a confidence region on θ is to determine $C(\mathscr{D}_n)$ such that

$$\pi(\theta \in C(\mathscr{D}_n)|\mathscr{D}_n) = 1 - \alpha$$

where α is a predetermined level

Bayesian Credible Intervals

Since the Bayesian approach processes θ as a random variable, a natural definition of a confidence region on θ is to determine $C(\mathscr{D}_n)$ such that

$$\pi(\theta \in C(\mathscr{D}_n)|\mathscr{D}_n) = 1 - \alpha$$

where α is a predetermined level

The integration is done over the parameter space, rather than over the observation space

Bayesian Credible Intervals

Since the Bayesian approach processes θ as a random variable, a natural definition of a confidence region on θ is to determine $C(\mathscr{D}_n)$ such that

$$\pi(\theta \in C(\mathscr{D}_n)|\mathscr{D}_n) = 1 - \alpha$$

where α is a predetermined level

The integration is done over the parameter space, rather than over the observation space

The quantity $1 - \alpha$ thus corresponds to the probability that a random θ belongs to this set $C(\mathscr{D}_n)$, rather than to the probability that the random set contains the true value of θ

Bayesian Credible Intervals

Given this drift in the interpretation of a confidence set is called a *credible set* by Bayesians.

Bayesian Credible Intervals

Given this drift in the interpretation of a confidence set is called a *credible set* by Bayesians.

A standard credible set corresponds to the values of θ with the highest posterior values,

$$C(\mathscr{D}_n) = \{\theta;\, \pi(\theta|\mathscr{D}_n) \geqslant k_\alpha\}$$

where k_α is determined by the coverage constraint

Bayesian Credible Intervals

Given this drift in the interpretation of a confidence set is called a *credible set* by Bayesians.

A standard credible set corresponds to the values of θ with the highest posterior values,

$$C(\mathscr{D}_n) = \{\theta;\, \pi(\theta|\mathscr{D}_n) \geqslant k_\alpha\}$$

where k_α is determined by the coverage constraint

This region is called the **Highest Posterior Density** (HPD) region

Bayesian Credible Intervals

Once again, suppose $\mathscr{D}_n$ is a normal $\mathscr{N}(\mu, \sigma^2)$ sample of size n and $\theta = (\mu, \sigma^2)$

$$\boxed{\mu|\sigma^2, \mathscr{D}_n \sim \mathscr{N}\left(\bar{x}, \sigma^2/n\right)}$$

$$\boxed{\sigma^2|\mathscr{D}_n \sim \mathscr{I}\mathscr{G}\left(n/2, s^2/2\right)}$$

$$\boxed{\mu|\mathscr{D}_n \sim \mathscr{T}\left(n, \bar{x}, s^2/n^2)\right)}$$

Bayesian Credible Intervals

Once again, suppose $\mathscr{D}_n$ is a normal $\mathscr{N}(\mu, \sigma^2)$ sample of size n and $\theta = (\mu, \sigma^2)$

$$\boxed{\mu|\sigma^2, \mathscr{D}_n \sim \mathscr{N}\left(\bar{x}, \sigma^2/n\right)}$$

$$\boxed{\sigma^2|\mathscr{D}_n \sim \mathscr{I}\mathscr{G}\left(n/2, s^2/2\right)}$$

$$\boxed{\mu|\mathscr{D}_n \sim \mathscr{T}\left(n, \bar{x}, s^2/n^2)\right)}$$

Therefore, the credible interval of probability $1-\alpha$ on μ is

$$[\bar{x} - t_{1-\alpha/2,n}\sqrt{(n-1)s^2/n^2}, \bar{x} + t_{1-\alpha/2,n}\sqrt{(n-1)s^2/n^2}]$$

Bayesian model choice

When are comparing models with indices $k = 1, 2, \ldots, J$, we introduce a model indicator $\mathfrak{M}$ taking values in $\{1, 2, \ldots, J\}$ and representing the index of the "true" model

Bayesian model choice

When are comparing models with indices $k = 1, 2, \ldots, J$, we introduce a model indicator $\mathfrak{M}$ taking values in $\{1, 2, \ldots, J\}$ and representing the index of the “true” model

If $\mathfrak{M} = k$, the data $\mathscr{D}_n$ are generated from a statistical model $\mathfrak{M}_k$ with likelihood $\ell_k(\theta_k|\mathscr{D}_n)$ and parameter $\theta_k \in \Theta_k$

Bayesian model choice

Bayes procedures will depend on the posterior probabilities in the model space

$$\mathbb{P}^{\pi}(\mathfrak{M} = k|\mathscr{D}_n)$$

Bayesian model choice

The prior π is defined over the collection of model indices $\{1, 2, \ldots, J\}$, and, conditionally on the model index $\mathfrak{M}$, on the corresponding parameter space Θ_k

Bayesian model choice

The prior π is defined over the collection of model indices $\{1, 2, \ldots, J\}$, and, conditionally on the model index $\mathfrak{M}$, on the corresponding parameter space Θ_k

Choice of the prior model probabilities $\mathbb{P}^\pi(\mathfrak{M} = k)$

- in some cases, there is experimental or subjective evidence about those probabilities,
- typically, we are forced to settle for equal weights $\mathbb{P}^\pi(\mathfrak{M} = k) = 1/J$

Bayesian model choice

A key quantity, the integrated likelihood, also called the evidence

$$\mathbb{P}^\pi(\mathfrak{M} = k|\mathscr{D}_n) \propto \mathbb{P}^\pi(\mathfrak{M} = k) \int \ell_k(\theta_k|\mathscr{D}_n)\pi_k(\theta_k)\,\mathrm{d}\theta_k$$

Bayesian model choice

A key quantity, the integrated likelihood, also called the evidence

$$\mathbb{P}^{\pi}(\mathfrak{M} = k|\mathscr{D}_n) \propto \mathbb{P}^{\pi}(\mathfrak{M} = k) \int \ell_k(\theta_k|\mathscr{D}_n)\pi_k(\theta_k)\, d\theta_k$$

$\mathbb{P}^{\pi}(\mathfrak{M} = k|\mathscr{D}_n)$ **is the core object in Bayesian model choice, the default procedure is to select the model with the highest posterior probability**

Bayesian model choice

Why Bayesian inference embodies Occam's razor?

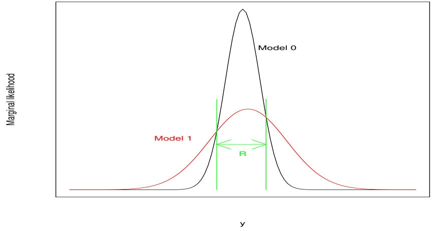

Bayesian model choice

Why Bayesian inference embodies Occam's razor?

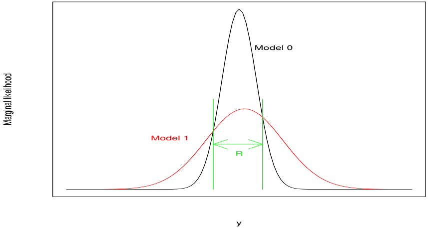

A simple model, like Model 0, makes only a limited range of predictions; a more powerful model, like Model 1, that has, for example, more free parameters, is able to predict a greater variety of data sets

Bayesian model choice

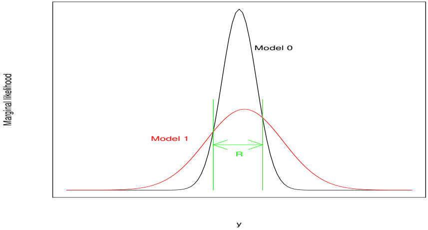

Suppose that equal prior probabilities have been assigned to the two models. Then, if the data set falls in region R, the less powerful model will be the more probable model

Bayesian model choice

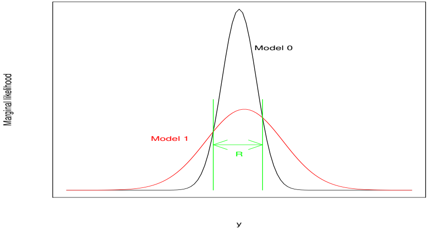

Suppose that equal prior probabilities have been assigned to the two models. Then, if the data set falls in region R, the less powerful model will be the more probable model

The marginal likelihood corresponds to a penalized likelihood!

Bayesian model choice

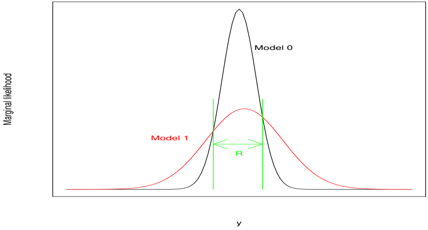

Suppose that equal prior probabilities have been assigned to the two models. Then, if the data set falls in region R, the less powerful model will be the more probable model

The marginal likelihood corresponds to a penalized likelihood!

The BIC information criterium comes from an asymptotic Laplace approximation of the evidence

Bayesian model choice

Bayesian test and Bayesian model choice: the same problem

Bayesian model choice

Bayesian test and Bayesian model choice: the same problem

For instance, given a single observation $x \sim \mathscr{N}(\mu, \sigma^2)$ from a normal model where σ^2 is known

Bayesian model choice

Bayesian test and Bayesian model choice: the same problem

For instance, given a single observation $x \sim \mathcal{N}(\mu, \sigma^2)$ from a normal model where σ^2 is known

If $\mu \sim \mathcal{N}(\xi, \tau^2)$, the posterior distribution $\mu|x \sim \mathcal{N}(\xi(x), \omega^2)$ with

$$\xi(x) = \frac{\sigma^2\xi + \tau^2 x}{\sigma^2 + \tau^2} \quad \text{and} \quad \omega^2 = \frac{\sigma^2\tau^2}{\sigma^2 + \tau^2}$$

Bayesian model choice

If the question of interest is to decide whether μ is negative or positive, we can directly compute

$$\begin{aligned}\mathbb{P}^{\pi}(\mu < 0|x) &= \mathbb{P}^{\pi}\left(\frac{\mu - \xi(x)}{\omega} < \frac{-\xi(x)}{\omega}\right) \\ &= \Phi\left(-\xi(x)/\omega\right)\end{aligned}$$

where Φ is the normal cdf

Bayesian model choice

If the question of interest is to decide whether μ is negative or positive, we can directly compute

$$\begin{aligned}\mathbb{P}^{\pi}(\mu < 0|x) &= \mathbb{P}^{\pi}\left(\frac{\mu - \xi(x)}{\omega} < \frac{-\xi(x)}{\omega}\right) \\ &= \Phi\left(-\xi(x)/\omega\right)\end{aligned}$$

where Φ is the normal cdf

This computation does not seem to follow from the principles we just stated but it is only a matter of perspective

Bayesian model choice

We can derive the priors on both models from the original prior

Bayesian model choice

We can derive the priors on both models from the original prior

Deriving this posterior probability indeed means that, a priori, μ is negative with probability $\mathbb{P}^\pi(\mu < 0) = \Phi(-\xi/\tau)$ and that, in this model, the prior on μ is the truncated normal

$$\pi_1(\mu) = \frac{\exp\{-(\mu-\xi)^2/2\tau^2\}}{\sqrt{2\pi}\tau\Phi(-\xi/\tau)}\,\mathbb{I}_{\mu<0}$$

while μ is positive with probability $\Phi(\xi/\tau)$ and, in this second model, the prior on μ is the truncated normal

$$\pi_2(\mu) = \frac{\exp\{-(\mu-\xi)^2/2\tau^2\}}{\sqrt{2\pi}\tau\Phi(\xi/\tau)}\,\mathbb{I}_{\mu>0}$$

Bayesian model choice

The Bayes factor

$$B_{21}^{\pi}(\mathscr{D}_n) = \frac{\mathbb{P}^{\pi}(\mathfrak{M} = 2|\mathscr{D}_n)/\mathbb{P}^{\pi}(\mathfrak{M} = 1|\mathscr{D}_n)}{\mathbb{P}^{\pi}(\mathfrak{M} = 2)/\mathbb{P}^{\pi}(\mathfrak{M} = 1)}$$

Bayesian model choice

The Bayes factor

$$\mathrm{B}_{21}^{\pi}(\mathscr{D}_n) = \frac{\mathbb{P}^{\pi}(\mathfrak{M} = 2|\mathscr{D}_n)/\mathbb{P}^{\pi}(\mathfrak{M} = 1|\mathscr{D}_n)}{\mathbb{P}^{\pi}(\mathfrak{M} = 2)/\mathbb{P}^{\pi}(\mathfrak{M} = 1)}$$

While this quantity is a simple one-to-one transform of the posterior probability, it can be used for Bayesian model choice without first resorting to a determination of the prior weights of both models

Bayesian model choice

The Bayes factor

$$B_{21}^{\pi}(\mathscr{D}_n) = \frac{\mathbb{P}^{\pi}(\mathfrak{M}=2|\mathscr{D}_n)/\mathbb{P}^{\pi}(\mathfrak{M}=1|\mathscr{D}_n)}{\mathbb{P}^{\pi}(\mathfrak{M}=2)/\mathbb{P}^{\pi}(\mathfrak{M}=1)}$$

While this quantity is a simple one-to-one transform of the posterior probability, it can be used for Bayesian model choice without first resorting to a determination of the prior weights of both models

$$B_{21}^{\pi}(\mathscr{D}_n) = \frac{\int_{\Theta_2} \ell_2(\theta_2|\mathscr{D}_n)\pi_2(\theta_2)\,\mathrm{d}\theta_2}{\int_{\Theta_1} \ell_1(\theta_1|\mathscr{D}_n)\pi_1(\theta_1)\,\mathrm{d}\theta_1} = \frac{m_2(\mathscr{D}_n)}{m_1(\mathscr{D}_n)}$$

Bayesian model choice

The Ban on Improper Priors

Looking at the expression of the Bayes factor,

$$B_{21}^{\pi}(\mathscr{D}_n) = \frac{\int_{\Theta_2} \ell_2(\theta_2|\mathscr{D}_n)\pi_2(\theta_2)\,d\theta_2}{\int_{\Theta_1} \ell_1(\theta_1|\mathscr{D}_n)\pi_1(\theta_1)\,d\theta_1}$$

it is clear that, when either π_1 or π_2 are improper, it is impossible to normalise the improper measures in a unique manner

Bayesian model choice

The Ban on Improper Priors

Looking at the expression of the Bayes factor,

$$B_{21}^{\pi}(\mathscr{D}_n) = \frac{\int_{\Theta_2} \ell_2(\theta_2|\mathscr{D}_n)\pi_2(\theta_2)\,\mathrm{d}\theta_2}{\int_{\Theta_1} \ell_1(\theta_1|\mathscr{D}_n)\pi_1(\theta_1)\,\mathrm{d}\theta_1}$$

it is clear that, when either π_1 or π_2 are improper, it is impossible to normalise the improper measures in a unique manner

Therefore, the Bayes factor becomes completely arbitrary since it can be multiplied by one or two arbitrary constants

Bayesian model choice

Since improper priors are an essential part of the Bayesian approach, there are many proposals found in the literature to overcome this ban

Bayesian model choice

Since improper priors are an essential part of the Bayesian approach, there are many proposals found in the literature to overcome this ban

Most of those proposals rely on a device that transforms the improper prior into a proper probability distribution by exploiting a fraction of the data $\mathscr{D}_n$

Bayesian model choice

Since improper priors are an essential part of the Bayesian approach, there are many proposals found in the literature to overcome this ban

Most of those proposals rely on a device that transforms the improper prior into a proper probability distribution by exploiting a fraction of the data $\mathscr{D}_n$

The variety of available solutions is due to the many possibilities of removing the dependence on the choice of the portion of the data used in the first step

Bayesian model choice

Since improper priors are an essential part of the Bayesian approach, there are many proposals found in the literature to overcome this ban

Most of those proposals rely on a device that transforms the improper prior into a proper probability distribution by exploiting a fraction of the data $\mathscr{D}_n$

The variety of available solutions is due to the many possibilities of removing the dependence on the choice of the portion of the data used in the first step

The resulting procedures are called *pseudo-Bayes factors*, although some may actually correspond to true Bayes factors

Bayesian model choice

There is a major exception to this ban on improper priors that we can exploit

Bayesian model choice

There is a major exception to this ban on improper priors that we can exploit

If both models under comparison have parameters that have similar enough meanings to share the same prior distribution, as for instance a measurement error σ^2, then the normalisation issue vanishes

Bayesian model choice

There is a major exception to this ban on improper priors that we can exploit

If both models under comparison have parameters that have similar enough meanings to share the same prior distribution, as for instance a measurement error σ^2, then the normalisation issue vanishes

Note that we are not assuming that parameters are *common* to both models and thus that we do not contradict the earlier warning about different parameters to different models

Bayesian Model Averaging

The posterior probabilities in the model space can be used to average over the decisions coming from different models

Bayesian Model Averaging

The posterior probabilities in the model space can be used to average over the decisions coming from different models

Suppose that we are interested in the prediction of z and that, for model k, the predictive distribution of z is $\mathrm{g}_{\mathrm{k}}(z|\mathscr{D}_{\mathrm{n}})$

Bayesian Model Averaging

The posterior probabilities in the model space can be used to average over the decisions coming from different models

Suppose that we are interested in the prediction of z and that, for model k, the predictive distribution of z is $g_k(z|\mathscr{D}_n)$

The average predictive of z is

$$\sum_{k=1}^{J} \mathbb{P}^\pi(\mathfrak{M} = k|\mathscr{D}_n) g_k(z|\mathscr{D}_n)$$

Difficulties with the Bayesian paradigm

Prior difficulties:

- When we have prior informations, how to choose the prior distributions on the parameters of each model in a compatible way? What about the prior distribution in the models's space?
- When we do not have any prior information, **we can not use improper prior distribution.** Indeed, in that case, the models's posterior probabilities are only defined up to some arbitrary constants. How to choose the various prior distributions?

Difficulties with the Bayesian paradigm

Computational difficulties:

- How to approximate the various posterior probabilities?
- How to approximate the evidences?
- When the number of models in consideration is huge, how to explore the models's space?