Analytics, Inference and Computation in Cosmology: Exercises on Bayesian Inference

Roberto Trotta, Imperial College London

Sept 2018

Contents

1 Bayesian Reasoning

A cohort chemistry undergraduates are screened for a dangerous medical condition called *Bacillum Bayesianum* (BB). The incidence of the condition in the population (i.e., the probability that a randomly selected person has the disease) is estimated at about 1%. If the person has BB, the test returns positive 95% of the time. There is also a known 5% rate of false positives, i.e. the test returning positive even if the person is free from BB. One of your friends takes the test and it comes back positive. Here we examine whether your friend should be worried about her health.

(i) Translate the information above in suitably defined probabilities. The two relevant propositions here are whether the test returns positive (denote this with a + symbol) and whether the person is actually sick (denote this with the symbol $BB = 1$. Denote the case when the person is healthy as $BB = 0$).

(ii) Compute the conditional probability that your friend is sick, knowing that she has tested positive, i.e., find $P(BB = 1|+)$.

(iii) Imagine screening the general population for a very rare disease, whose incidence in the population is 10^{-6} (i.e., one person in a million has the disease on average, i.e. $P(BB = 1) = 10^{-6}$). What should the reliability of the test (i.e., $P(+|BB = 1)$) be if we want to make sure that the probability of actually having the disease after testing positive is at least 99%? Assume first that the false positive rate $P(+|BB = 0)$ (i.e, the probability of testing positive while healthy), is 5% as in part (a). What can you conclude about the feasibility of such a test?

2 Bayesian Parameter Estimation

2.1 Coin Tossing: Binomial Distribution

A coin is tossed N times and heads come up H times.

(i) What is the likelihood function? Identify clearly the parameter, θ, and the data.

(ii) What is a reasonable, non-informative prior on θ?

(iii) Compute the posterior probability for θ. Recall that θ is the probability that a single flip will give heads. This integral will prove useful:

$$\int_0^1 d\theta\theta^N(1-\theta)^M = \frac{\Gamma(N+1)\Gamma(M+1)}{\Gamma(N+M+2)}. \tag{1}$$

(iv) Determine the posterior mean and standard deviation of θ.

(v) Plot your results as a function of H for $N = 10, 100, 1000$.

2.2 The Gaussian Linear Model

This problem takes you through the steps to derive the posterior distribution for a quantity of interest θ, in the case of a Gaussian prior and Gaussian likelihood, for the 1-dimensional case.

2.2.1 Theory

Let us assume that we have made N independent measurements, $\hat{x} = \{\hat{x}_1, \hat{x}_2, \ldots, \hat{x}_N\}$ of a quantity of interest θ (this could be the temperature of an object, the distance of a galaxy, the mass of a planet, etc). We assume that each of the measurements in independently Gaussian distributed with known experimental standard deviation σ. Let us denote the sample mean by $\bar{x}$, i.e.

$$\bar{x} = \frac{1}{N}\sum_{i=1}^{N}\hat{x}_i. \tag{2}$$

Before we do the experiment, our state of knowledge about the quantity of interest θ is described by a Gaussian distribution on θ, centered around 0 (we can always choose the units in such a way that this is the case). Such a prior might come e.g. from a previous experiment we have performed. The new experiment is however much more precise, i.e. $\Sigma \gg \sigma$. Our prior state of knowledge be written in mathematical form as the following Gaussian pdf:

$$p(\theta) \sim \mathcal{N}(0, \Sigma^2). \tag{3}$$

(i) Write down the likelihood function for the measurements and show that it can be recast in the form:

$$\mathcal{L}(\theta) = L_0 \exp\left(-\frac{1}{2}\frac{(\theta - \bar{x})^2}{\sigma^2/N}\right), \tag{4}$$

where L_0 is a constant that does not depend on θ.

(ii) By using Bayes theorem, compute the posterior probability for θ after the data have been taken into account, i.e. compute $p(\theta|\hat{x})$. Show that it is given by a Gaussian of mean $\bar{x}\frac{\Sigma^2}{\Sigma^2+\sigma^2/N}$ and variance $\left[\frac{1}{\Sigma^2} + \frac{N}{\sigma^2}\right]^{-1}$.
Hint: you may drop the normalization constant from Bayes theorem, as it does not depend on θ

(iii) Show that as $N \to \infty$ the posterior distribution becomes independent of the prior.

(iv) Show that as $N \to \infty$ the mean of the posterior distribution converges to the MLE of the mean for θ. This means that for a large number of measurements, the Bayesian result matches the frequentist MLE result.

2.2.2 Two-dimensional Example

Now we specialize to the case $n = 2$, i.e. we have two parameters of interest, $\theta = \{\theta_1, \theta_2\}$ and the linear function we want to fit is given by

$$y = \theta_1 + \theta_2 x. \tag{5}$$

(In the formalism above, the basis vectors are $X^1 = 1, X^2 = x$).

Table 1 gives an array of $d = 10$ measurements $y = \{y_1, y_2, \ldots, y_{10}\}$, together with the values of the independent variable x_i. Assume that the uncertainty in the same for all measurements, i.e. $\tau_i = 0.1$ ($i = 1, \ldots, 10$). You may further assume that measurements are uncorrelated. The data set is shown in the left panel of Fig. 1

(i) Assume a Gaussian prior with Fisher matrix $P = \text{diag}\left(10^{-2}, 10^{-2}\right)$ for θ.

Find the posterior distribution for θ given the data, and plot it in 2 dimensions in the (θ_1, θ_2) plane (see right panel of Fig. 1).

Use the appropriate contour levels to demarcate 1, 2 and 3 sigma joint credible intervals of the posterior.

Table 1: Data sets for the Gaussian linear model exercise. You may assume that all data points are independently and identically distributed with standard deviation of the noise $\sigma = 0.1$.

x	y
0.8308	0.9160
0.5853	0.7958
0.5497	0.8219
0.9172	1.3757
0.2858	0.4191
0.7572	0.9759
0.7537	0.9455
0.3804	0.3871
0.5678	0.7239
0.0759	0.0964

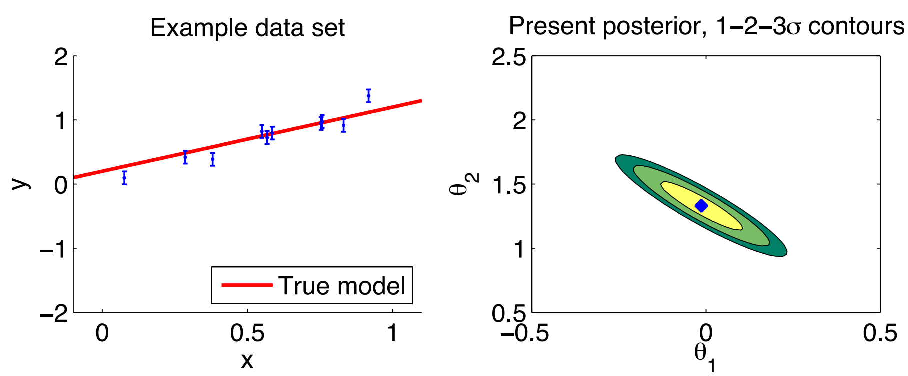

Figure 1: Left panel: data set for the Gaussian linear problem. The solid line shows the true value of the linear model from which the data have been generated, subject to Gaussian noise. Right panel: 2D credible intervals from the posterior distribution for the parameters. The the blue diamond is the Maximum Likelihood Estimator, whose value for this data set is $x = -0.0136, y = 1.3312$.

(ii) In a language of your choice, write an implementation of the Metropolis-Hastings Markov Chain Monte Carlo algorithm, and use it to obtain samples from the posterior distribution.

(iii) If you are already familiar with Metropolis-Hastings, write an implementation of Hamiltonian Monte Carlo instead.

(iv) Plot *equal weight* samples in the (θ_1, θ_2) space, as well as marginalized 1-dimensional posterior distributions for each parameter.

(v) Compare the credible intervals that you obtained from the MCMC with the analytical solution.

2.3 Poisson counts

2.3.1 Maximum Likelihood Approach

An astronomer measures the photon flux from a distant star using a very sensitive instrument that counts single photons. After one minute of observation, the instrument has collected $\hat{r}$ photons. One can assume that the photon counts, $\hat{r}$, are distributed according to the Poisson distribution. The astronomer wishes to determine λ, the emission rate of the source.

(i) What is the likelihood function for the measurement? Identify explicitly what is the unknown parameter and what are the data in the problem.

(ii) If the true rate is $\lambda = 10$ photons/minute, what is the probability of observing $\hat{r} = 15$ photons in one minute?

(iii) Find the Maximum Likelihood Estimate for the rate λ (i.e., the number of photons per minute). What is the maximum likelihood estimate if the observed number of photons is $\hat{r} = 10$?

2.3.2 The On/Off Problem

Upon reflection, the astronomer realizes that the photon flux is the superposition of photons coming from the star plus "background" photons coming from other faint sources within the field of view of the instrument. The background rate is supposed to be known, and it is given by λ_b photons per minute. This can be estimated e.g. by pointing the telescope away from the source (the "off" measurement) and measuring the photon counts there, where the telescope is only picking up background photons. This estimate of the background comes with an uncertainty, of course, but we'll ignore this for now. She then points to the star again, measuring $\hat{r}_t$ photons in a time t_t (this is the "on" measurement).

(i) What is her maximum likelihood estimate of the rate λ_s from the star in this case? *Hint:* The total number of photons $\hat{r}_t$ is Poisson distributed with rate $\lambda = \lambda_s + \lambda_b$, where λ_s is the rate for the star.

(ii) What is the source rate (i.e., the rate for the star) if $\hat{r}_t = 30$, $t_t = 2$ mins, and $\lambda_b = 12$ photons per minute?

(iii) Is it possible that the measured average rate from the source (i.e., $\hat{r}_t/t_t$) is less than λ_b? Discuss what happens in this case and comment on the physicality of this result.

2.3.3 On/Off Problem: Bayesian version

We revisit the On/Off problem but this time from a Bayesian perspective, which fully and automatically accounts for uncertainty in the background rate estimate.

We consider first the "off" measurement, which collects n_{off} photons in a time t_{off}.

(i) Assuming a uniform prior on the background rate b, find the posterior distribution for b from the off measurement.

(ii) Now consider the "on" measurement, which collects a number n_{on} of photons during a time t_{on}. This is a measurement for the combined rate $s+b$ (where s denotes the source rate). Write down the likelihood function for this measurement.

(iii) Assume again a uniform prior on s, and a prior on b given by the posterior of the "off" measurement[1], find the (unnormalized) joint posterior distribution for s, b, and show that is is given by the expression:

$$p(s, b|n_{\text{on}}, t_{\text{on}}) \propto (s+b)^{n_{\text{on}}} b^{n_{\text{off}}} \exp(-s t_{\text{on}}) \exp(-b(t_{\text{on}} + t_{\text{off}})) \text{ for } s, b \geq 0. \tag{6}$$

(iv) Compute analytically the marginal posterior pdf for the signal, s, by integrating the joint posterior over b, i.e.

$$p(s|n_{\text{on}}, t_{\text{on}}) = \int_0^\infty p(s, b|n_{\text{on}}, t_{\text{on}}) db. \tag{7}$$

. Plot the resulting marginal distribution for the signal s for the following two cases, and compare the result with the MLE result:

(a) $n_{\text{on}} = 5, t_{\text{on}} = 1, n_{\text{off}} = 2, t_{\text{off}} = 1$

(b) $n_{\text{on}} = 2, t_{\text{on}} = 2, n_{\text{off}} = 3, t_{\text{off}} = 1$

(c) $n_{\text{on}} = 8, t_{\text{on}} = 1, n_{\text{off}} = 2, t_{\text{off}} = 4$

[1] The posterior for the "off" measurement can be used as prior on b for the "on" measurement. Alternatively, you can write down the joint posterior on s, b conditional on both measurements, with an ur-prior on b that is just the uniform prior (i.e., the prior that you used for the "off" measurement). Both procedures will give the same result, as they should (consistency of Bayesian reasoning is always in-built). Convince yourself that this is indeed the case!

Hint: use the binomial expansion: $(s+b)^{n_{\rm on}} = \sum_{k=0}^{n_{\rm on}} \binom{n_{\rm on}}{k} s^{n_{\rm on}-k} b^k$.

(v) Write a code to perform MCMC sampling of the joint posterior for s, b (in Python you may want to use the `PyMC` package). Plot equal-weight samples from the posterior in parameter space for $n_{\rm on} = 10, t_{\rm on} = 2$, $n_{\rm off} = 3, t_{\rm off} = 1$. Marginalize over b numerically and compare the resulting numerical estimate with the analytical result above.

3 Toy Cosmological Parameter Inference (Harder)

Supernovae type Ia can be used as standardizable candles to measure distances in the Universe. This series of problems explores the extraction of cosmological information from a simplified SNIa toy model.

The cosmological parameters we are interested in constraining are

$$\mathscr{C} = \{\Omega_m, \Omega_\Lambda, h\} \tag{8}$$

where Ω_m is the matter density (in units of the critical energy density) and Ω_Λ is the dark energy density, assumed here to be in the form of a cosmological constant, i.e. $w = -1$ at all redshifts. In the following, we will fix $h = 0.72$ for simplicity, where the Hubble constant today is given by $H_0 = 100h$km/s/Mpc.

In an FRW cosmology defined by the parameters $\mathscr{C}$, the distance modulus μ (i.e., the difference between the apparent and absolute magnitudes, $\mu = m - M$) to a SN at redshift z is given by

$$\mu(z, \mathscr{C}) = 5 \log \left[\frac{D_L(z, \Omega_m, \Omega_\Lambda, h)}{\text{Mpc}} \right] + 25, \tag{9}$$

where D_L denotes the luminosity distance to the SN. Recalling that $D_L = c d_L / H_0$, We can rewrite this as

$$\mu(z, \mathscr{C}) = \eta + 5 \log d_L(z, \Omega_m, \Omega_\Lambda), \tag{10}$$

where

$$\eta = -5 \log \frac{100h}{c} + 25 \tag{11}$$

and c is the speed of light in km/s. We have defined the dimensionless luminosity distance

$$d_L(z, \Omega_m, \Omega_\Lambda) = \frac{(1+z)}{\sqrt{|\Omega_\kappa|}} \text{sinn}\{\sqrt{|\Omega_\kappa|} \int_0^z \mathrm{d}z' [(1+z')^3 \Omega_m + \Omega_\Lambda + (1+z')^2 \Omega_\kappa]^{-1/2}\}. \tag{12}$$

The curvature parameter is given by the constraint equation

$$\Omega_\kappa = 1 - \Omega_m - \Omega_\Lambda \tag{13}$$

and the function

$$\text{sinn}(x) = \begin{cases} x & \text{for a flat Universe } (\Omega_\kappa = 0); \\ \sin(x) & \text{for a closed Universe } (\Omega_\kappa < 0); \\ \sinh(x) & \text{for an open Universe } (\Omega_\kappa > 0). \end{cases} \tag{14}$$

We now assume that from each SNIa in our sample we get a measurement of the distance modulus with Gaussian noise[2], i.e., that the likelihood function for each SN i ($i = 1, \ldots, N$) is of the form

$$\mathcal{L}_i(z_i, \mathscr{C}, M) = \frac{1}{\sqrt{2\pi}\sigma_i} \exp\left(-\frac{1}{2} \frac{(\hat{\mu}_i - \mu(z_i, \mathscr{C}))^2}{\sigma_i^2} \right). \tag{15}$$

The observed distance modulus is given by $\hat{\mu}_i = \hat{m}_i - M$, where $\hat{m}_i$ is the observed apparent magnitude and M is the intrinsic magnitude of the SNIa. We assume that each SN observation is independent of all the others.

The provided data file[3] (`SNe_simulated.dat`) contains simulated observations from the above simplified model of $N = 300$ SNIa. The two columns give the redshift z_i and the observed apparent magnitude $\hat{m}_i$. The observational error is the same for all SNe, $\sigma_i = \sigma = 0.4$ mag for $i = 1, \ldots, N$.

A plot of the data set is shown in the left panel of Fig. 2. The characteristics of the simulated SNe are designed to mimic currently available datasets (see [6, 1, 5, 7, 3]).

(i) We assume that the intrinsic magnitude[4] is known and fix $M = M_0 = -19.3$ and that $h = 0.72$. We also assume that the observational error is known, given by the value above.

Using a language of your choice, write a code to carry out an MCMC sampling of the posterior probability for $(\Omega_m, \Omega_\Lambda)$ and plot the resulting 68% and 95% posterior regions, both in 2D and marginalized to 1D, using uniform priors on $(\Omega_m, \Omega_\Lambda)$ (be careful to define them explicitly).

You should obtain a result similar to the 2D plot shown in the right panel of Fig. 2.

(ii) † Add the quantity σ (the observational error) to the set of unknown parameters and estimate it from the data along with $\mathscr{C}$. Notice that since σ is a "scale parameter", the appropriate (improper) prior is $p(\sigma) \propto 1/\sigma$ (see [4] for a justification).

(iii) The location of the peaks in the CMB power spectrum gives a precise measurement of the angular diameter distance to the last scattering surface, divided by the sound horizon at decoupling. This approximately translates into an effective constraint (see [8], Fig. 20) on the following degenerate combination of Ω_m and Ω_Λ:

$$1.41\Omega_\Lambda + \Omega_m = 1.30 \pm 0.04. \tag{16}$$

---

[2] We neglect the important issue of applying the empirical corrections known as Phillip's relations to the observed light curve. This is of fundamental important in order to reduce the scatter of SNIa within useful limits for cosmological distance measurements, but it would introduce a technical complication here without adding to the fundamental scope of this exercise.

[3] Thanks to Marisa March for help with the simulation.

[4] In reality the SNe intrinsic magnitude is not fixed, but there is an "intrinsic dispersion" (even after Phillips' corrections) reflecting perhaps intrinsic variability in the explosion mechanism, or environmental parameters which are currently poorly understood.

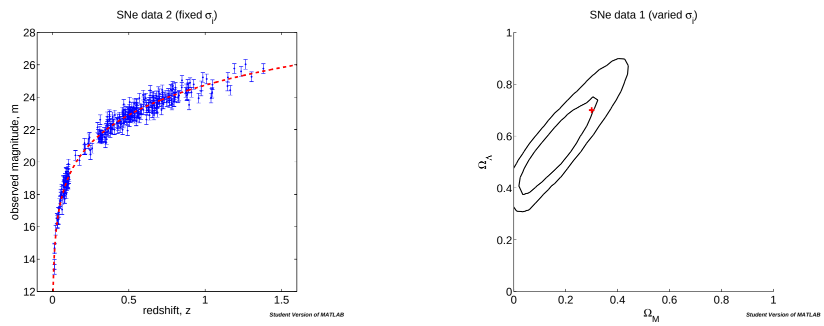

Figure 2: Left: Simulated SNIa dataset, `SNe_simulated.dat`. The solid line is the true underlying cosmology. Right: constraints on Ω_m, Ω_Λ from this dataset, with contours delimiting 2D joint 68% and 95% credible regions (uniform priors on the variables Ω_m, Ω_Λ, assuming $M = M_0$ fixed and $h = 0.72$). The red cross denotes the true value.

Add this constraint (assuming a Gaussian likelihood, with the above mean and standard deviation) to the SNIa likelihood and plot the ensuing combined 2D and 1D limits on $(\Omega_m, \Omega_\Lambda)$.

(iv) The measurement of the baryonic acoustic oscillation scale in the galaxy power spectrum at small redshift gives an effective constraint on the angular diameter distance D_A out to $z \sim 0.3$. This measurement can be summarized as [2]:

$$D_A(z = 0.57) = (1408 \pm 45) \text{ Mpc}. \tag{17}$$

Add this constraints (again assuming a Gaussian likelihood) to the above CMB+SNIa limits and plot the resulting combined 2D and 1D limits on $(\Omega_m, \Omega_\Lambda)$.
Hint: recall that $D_L(z) = (1 + z)^2 D_A(z)$.

References

[1] Amanullah, R., Lidman, C., Rubin, D., Aldering, G., Astier, P., et al.: Spectra and Light Curves of Six Type Ia Supernovae at 0.511 ¡ z ¡ 1.12 and the Union2 Compilation. Astrophys.J. **716**, 712–738 (2010). DOI 10.1088/0004-637X/716/1/712

[2] Anderson, L., et al.: The clustering of galaxies in the SDSS-III Baryon Oscillation Spectroscopic Survey: measuring D_A and H at z = 0.57 from the baryon acoustic peak in the Data Release 9 spectroscopic Galaxy sample. Mon. Not. Roy. Astron. Soc. **439**(1), 83–101 (2014). DOI 10.1093/mnras/stt2206

[3] Betoule, M., et al.: Improved cosmological constraints from a joint analysis of the SDSS-II and SNLS supernova samples. Astron. Astrophys. **568**, A22 (2014). DOI 10.1051/0004-6361/201423413

[4] Box, G.E.P., Tiao, G.C.: Bayesian Inference in Statistical Analysis. John Wiley & Sons, Chicester, UK (1992)

[5] Kessler, R., Becker, A., Cinabro, D., Vanderplas, J., Frieman, J.A., et al.: First-year Sloan Digital Sky Survey-II (SDSS-II) Supernova Results: Hubble Diagram and Cosmological Parameters. Astrophys.J.Suppl. **185**, 32–84 (2009). DOI 10.1088/0067-0049/185/1/32

[6] Kowalski, M., et al.: Improved Cosmological Constraints from New, Old and Combined Supernova Datasets. Astrophys.J. **686**, 749–778 (2008). DOI 10.1086/589937

[7] Rest, A., et al.: Cosmological Constraints from Measurements of Type Ia Supernovae discovered during the first 1.5 yr of the Pan-STARRS1 Survey. Astrophys. J. **795**(1), 44 (2014). DOI 10.1088/0004-637X/795/1/44

[8] Spergel, D.N., et al.: Wilkinson Microwave Anisotropy Probe (WMAP) three year results: implications for cosmology. Astrophys. J. Suppl. **170**, 377 (2007). DOI 10.1086/513700