SANS2022-23-HW1

jorge.garcia.vidal

October 2022

1 Multivariate Gaussian random variables

Generate for $N = 100, 1000, 10000$ samples of a multivariate gaussian variable $X = (X_1, X_2)$ with $\mu = (1, 2)$, $\Sigma = diag(1, 4)$. Plot the scatterplots of the generated samples for the three values of N.

Plot the histograms for X_1 and X_2. Can you recognize distribution of the corresponding random variable? Justify your answer taking into account the theory on gaussian random variables.

Repeat for Σ given by:

$$\Sigma = \begin{bmatrix} \frac{\sqrt{2}}{2} & -\frac{\sqrt{2}}{2} \\ \frac{\sqrt{2}}{2} & \frac{\sqrt{2}}{2} \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & 4 \end{bmatrix} \begin{bmatrix} \frac{\sqrt{2}}{2} & \frac{\sqrt{2}}{2} \\ -\frac{\sqrt{2}}{2} & \frac{\sqrt{2}}{2} \end{bmatrix} \quad (1)$$

Explain what you see

2 Bayesian estimation of 2D Gaussian random variables

Note: For solving this, you need to find somewhere the adequate formulas. In the class we only explained the 1D case, and very briefly the multidimensional case. You are in a Master now.

We want to estimate the location of an static point located at point $(2, 3)$. We measure $N = 10$ samples of the location of the point. The measure error has a known variance-covariance matrix given by:

$$\Sigma = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \quad (2)$$

Generate $N = 10$ samples of our reading. Draw the scatter plot.

Estimate μ. Use a uniform prior for μ. Draw 3-D plot of the likelihood and posterior distributions.

Repeat, but assuming now a gaussian prior for μ with different values of the mean μ_0 and variance-covariance matrix Σ_0. Draw also 3-D plots for these priors.

3 Metropolis-Hastings Monte Carlo simulation of a multinomial distribution in a simplex

Note: you can use your own code or use the code I've wrote and that you can find in the raco. I'm not a great python programmer, sorry .

We want to estimate the posterior distribution $p(\theta|y)$ and predictive probability $p(y_{n+1} = 1|y)$ of a multinomial distribution. In this case there are closed expressions for these distributions, but we want to use MCMC.

For the auxiliary Markov-process that generates candidate next states, we start using a uniform distribution: $Q(\theta, \theta') = constant$ for all values of θ' in the simplex, and $Q(\theta, \theta') = 0$ for θ' outside the simplex. For generating samples fomr such uniform distribution we use $Dirichlet((1, 1, 1, .., 1))$.

Run for $T = 100, 1000, 10000$ the corresponding M-H simulation for different values of y, for instance:

$$y = [2/7, 1/7, 4/7],$$

$$y = [2/15, 1/15, 1/15, 1/15, 1/15, 2/15, 1/15, 0, 2/15, 1/15, 3/15].$$

Plot the obtained histograms for the averages of the marginals of the posterior $\mathbb{E}(\theta_i|y)$ and the predictive probability $p(y_{n+1} = 0|y)$. Include also in the plots the true values of the estimated parameter.

For computing η, the acceptance probability $\eta = min(1, \frac{p^*(\theta'|y)}{p^*(\theta|y)}$, in the code I use

$$\eta = min(1, exp(log(p^*(\theta'|y)) - log(p^*(\theta|y))),$$

why?.

This choice of Q has a clear problem: it generates too many candidates that are not taken (i.e. the rejection rate in the algorithm is too high). Explain why.

Think in another distribution for Q. In the code in the raco you can find a distribution I've tried. It seems to work, explain why. What is the meaning of the parameter h?. Try with $h = 1, 0.1, 0.001$ and explain the results.