Computational Statistics
Final Exam

June 18, 2019

Write your code in a *single file* with name <your_name>.R. Insert comment lines between your code as follows:

```
\#---------------
\# Exercise 1
\# Q1
...
\# Q2
...
\#---------------
\# Exercise 2
\# Q1
...
```

Send your file to `tdenoeux@utc.fr`.

Exercise 1

The following data are assumed to be an i.i.d. sample from a Cauchy distribution:

0.31, 0.99, 0.54, 0.60, -1.08, -1.13, 0.92, 0.60, 1.13, 5.74, 1.14, 2.46, -11.28, 1.37, -17.58, 0.69, 0.20, 1.04, 1.10, 0.36

The density function of the Cauchy distribution with location parameter x_0 and scale parameter $\gamma > 0$ is

$$f(x) = \frac{1}{\pi\gamma\left[1 + \left(\frac{x-x_0}{\gamma}\right)^2\right]}.$$

1. Compute the maximum likelihood estimate of $\theta = (x_0, \gamma)$. (Use functions `dcauchy` for the Cauchy density and `optim` for the optimization).
2. Compute 95% confidence intervals on x_0 and γ using the bootstrap percentile method (use $B = 1000$ bootstrap samples).

Exercise 2

We consider the following realization from an i.i.d. random sample $X_1, \dots, X_{10}$ from an exponential distribution $\mathcal{E}(\theta)$ with rate θ:

0.345, 0.386, 0.279, 0.031, 0.177, 0.038, 0.450, 0.083, 0.217, 0.673

We recall that the density of $\mathcal{E}(\theta)$ is $f(x) = \theta \exp(-\theta x) I(x \geq 0)$. A lognormal prior distribution for θ is assumed: $\log \theta \sim \mathcal{N}(\log(6), 0.5^2)$. Denote the likelihood as $L(\theta; \mathbf{x})$ and the prior as $\pi(\theta)$. The MLE of θ is $\widehat{\theta} = 1/\overline{x}$.

1. Plot $q(\theta \mid \mathbf{x}) = \pi(\theta) L(\theta; \mathbf{x})$ and $e(\theta) = \pi(\theta) L(\widehat{\theta}; \mathbf{x})$ as a function of θ.
2. Generate a sample of size $N = 1000$ from the posterior distribution $f(\theta \mid \mathbf{x})$ using the rejection sampling method. Draw a histogram of this sample.
3. Compute a 95% confidence interval on the posterior expectation $\mathbb{E}(\theta|\mathbf{x})$.
4. Repeat the same operations as in the two previous questions, using now the SIR method.
5. Compare graphically the two samples using function `qqplot`.

Exercise 3

We consider the same data and the same model as in Exercise 2.

1. Construct an Metropolis-Hastings algorithm to sample from the posterior distribution $f(\theta \mid \mathbf{x})$ with an independence kernel, where the kernel is the prior distribution. Generate a Markov chain of size $n = 10,000$.
2. Plot the sample path, the histogram of simulated values and the autocorrelation function (use function `acf`).
3. Compute the posterior expectation of θ and its simulation standard error (use the batch-means method).