751800 – Computational Statistics
Final Exam (open-book)

August 21, 2026

Exercise 1 is independent from the others. Exercise 2 must be solved before Exercises 3 and 4.

Submit your work as *single* R source file with name <your_name>.R. Insert comment lines between your code as follows:

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

Send your files to `tdenoeux@utc.fr` at the end of the exam. (No file sent after leaving the room will be accepted).

Exercise 1

The density function of the gamma distribution with shape $\alpha > 0$ and rate $\beta > 0$ is

$$f(x; \alpha, \beta) = \frac{\beta^\alpha}{\Gamma(\alpha)} x^{\alpha-1} e^{-\beta x} \mathbb{1}_{(0,+\infty)}(x),$$

where Γ is the gamma function. Generate an i.i.d. sample of size $n = 100$ from this distribution using the following code:

```
set.seed(20260821)
x <- rgamma(100, 2, 1)
```

We will use these data to estimate $\theta = (\alpha, \beta)$ and to compute confidence intervals on both parameters.

1. Compute the maximum likelihood estimate $\widehat{\theta}$ of θ and the Hessian matrix of the log-likelihood at $\widehat{\theta}$. (Use functions `dgamma` for the gamma density, and `optim` with `method="L-BFGS-B"` and `lower=c(1e-6,1e-6)` for the optimization).
2. Compute 95% normal-approximation confidence intervals on α and β.
3. Compute 95% confidence intervals on α and β using the bootstrap percentile method with the nonparametric bootstrap (use $B = 10^3$ bootstrap samples).

Exercise 2

The *Topp–Leone distribution* with shape parameter $\nu > 0$, denoted by $\mathsf{TL}(\nu)$, is the probability distribution on $[0,1]$ with density

$$f(x;\nu) = 2\nu\,(1-x)\left(2x - x^2\right)^{\nu-1}, \quad 0 \leq x \leq 1.$$

Its inverse c.d.f. is

$$F^{-1}(u;\nu) = 1 - \sqrt{1 - u^{1/\nu}}, \quad u \in]0,1[.$$

1. Write an R function to generate a sample of size n from this distribution using the probability integral transform.
2. Generate a sample of size $n = 1000$ with $\nu = 1.5$. Plot the histogram and superimpose the density.

Exercise 3

We consider the following realization of an i.i.d. random sample $X_1, \ldots, X_{10}$ from a geometric distribution with unknown parameter $0 < \theta < 1$:

$$1,\ 1,\ 0,\ 0,\ 0,\ 5,\ 5,\ 0,\ 0,\ 0$$

We recall the expression of the probability mass function (p.m.f.) of the geometric distribution:

$$p(x;\theta) = \theta(1-\theta)^x, \quad x = 0, 1, 2, \ldots$$

A Topp–Leone prior distribution for θ is assumed, with parameter $\nu = 1.5$. Denote the likelihood as $L(\theta;\boldsymbol{x})$ and the prior as $f(\theta)$.

1. Compute the maximum likelihood estimate $\widehat{\theta}$ of θ. (Use functions `dgeom` for the geometric p.m.f. and `optimize` for optimization).
2. Plot $q(\theta \mid \boldsymbol{x}) = f(\theta)L(\theta;\boldsymbol{x})$ and $e(\theta) = f(\theta)L(\widehat{\theta};\boldsymbol{x})$ as functions of θ.

3. Generate a sample of size $N = 1000$ from the posterior distribution $f(\theta \mid \boldsymbol{x})$ using the rejection sampling method. Draw a histogram of this sample.

4. Estimate the posterior expectation $\mathbb{E}(\theta \mid \boldsymbol{x})$ and its simulation standard error.

5. Compute a 95% credibility interval on θ.

Exercise 4

We consider the same data and the same model as in Exercise 3.

1. Construct an Metropolis-Hastings algorithm to sample from the posterior distribution $f(\theta \mid \boldsymbol{x})$ with an independence kernel, where the kernel is the prior distribution. Generate a Markov chain of size $N = 10^5$.

2. Plot the sample path, the histogram of simulated values and the autocorrelation function (use function `acf` and a burn-in period $D = 5000$).

3. Compute the posterior expectation of θ and its simulation standard error (use the batch-means method).