Petter Mostad
Applied Mathematics and Statistics
Chalmers and GU

MSA100 / MVE186 Computer Intensive Statistical Methods

Re-exam 2 January 2017, 14:00 - 18:00
Examiner: Petter Mostad, phone 031-772-3579
visits the exam at 15:00 annd 17:00
Allowed aids: None.
The appendix of this exam contains informaion about probability distributions.
Total number of points: 30. To pass, at least 12 points are needed

1. (1 points) What is a pseudo-random number, and how does it differ from a random number?

2. (3 points) For each of the following probability densities, indicate how you would simulate a random sample from it using each time only simulations from the Uniform$(0, 1)$ distribution as a starting point.

   (a) A Logistic$(-1, 1)$ distribution.

   (b) A Binomial$(12, 0.42)$ distribution.

   (c) A Beta$(4.5, 9.2)$ distribution.

3. (4 points) We would like to simulate a sample from a density that is proportional to the function $f(x, y)$ defined for all real x and y below:

$$f(x,y) = \frac{1}{(|x+y|/10+1)^4} \cdot (|x-y|+0.1)^2 \cdot e^{-7|x-y|}$$

   Level curves for the function are shown in the figure below.

   (a) Assume one implements Gibbs sampling in this situation. Outline the steps in the algorithm (do not compute the details). Give any arguments you see against using Gibbs sampling in this situation.

   (b) Propose another way to generate a sample from the density above. A hint is to stepwise transform the problem. Desribe in detail the steps in your algorithm. Argue why the alternative is better than the direct Gibbs sampling in (a).

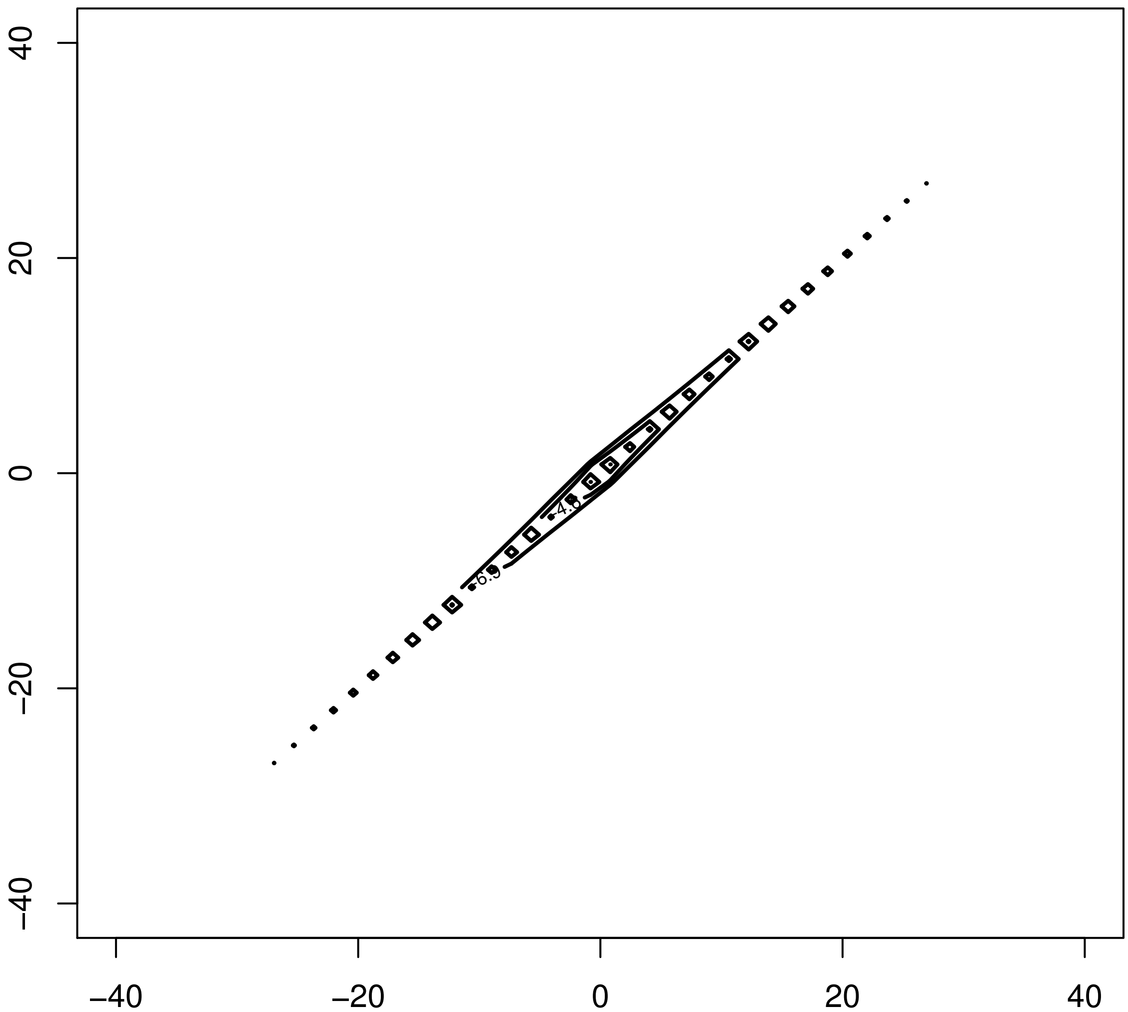

4. (2 points) Assume you have a sample $x_1, x_2, \ldots, x_m$ from a probability distribution with density function $\pi(x)$ on an interval $[a, b]$, and you want to estimate

$$I = \int_a^b g(x)\pi(x)\,dx$$

as well as produce an approximate 95% confidence interval for the estimate. Describe how you would compute the estimate and the approximate 95% confidence interval, as well as why the interval is an approximate 95% confidence interval. Make the assumptions you need to make along the way.

5. (2 points) Give a short explanation of the method of importance sampling.

6. (6 points) Assume you flip a coin n times, and let x denote the number of tails obtained. Let p be the probability of obtaining tails with this coin, and use Uniform(0, 1) as the prior for p.

   (a) Compute and name the posterior distribution for p given the observations, and find the posterior expectation of p. Also, write down the maximum likelihood estimate for p given the observations.

   (b) Assume you plan to flip the coin m additional times, and to record the number of tails in this new sequence as y. Compute an expression for the probability mass function of the predictive distribution for $y \mid x$.

   (c) Compute the expectation and variance of $y \mid x$.

7. (8 points) Define the following set of random variables:

$$
\begin{aligned}
Z &\sim \text{Bernoulli}(p) \\
X_0 &\sim \text{Normal}(\mu_0, 1) \\
X_1 &\sim \text{Normal}(\mu_1, 1) \\
X_2 \mid Z &\sim \begin{cases} \text{Normal}(\mu_0, 1) & \text{when } Z = 0 \\ \text{Normal}(\mu_1, 1) & \text{when } Z = 1 \end{cases}
\end{aligned}
$$

   Assume we have observed X_0, X_1, X_2 and want to find the Maximum Likelihood estimate for the parameters $\theta = (\mu_0, \mu_1, p)$ using the EM algorithm.

   (a) Write down the log-likelihood for the parameters $\theta = (\mu_0, \mu_1, p)$ given the full data[1] X_0, X_1, X_2, Z.

   (b) Assuming parameters $\theta' = (\mu_0', \mu_1', p')$ are fixed, compute the posterior probability that $Z = 1$ given the value of X_2.

   (c) Compute the expression for the $Q(\theta, \theta')$ function of the EM algorithm, i.e., do the E step of the EM algorithm.

   (d) Find the equations expressing $\hat{\theta}_{m+1}$ in terms of $\hat{\theta}_m$, i.e., the equations describing the iterative step of the EM algorithm.

8. (4 points)

   (a) Give a short explanation of the idea behind bootstrapping.

   (b) Assume you have observations $(x_1, y_1), \ldots, (x_n, y_n)$ of a dependent pair (X, Y) of random variables, and let $\hat{\rho}$ be the sample correlation coefficient computed from the data. Describe precisely how you would compute a Bootstrap estimate for the standard error of the sample correlation coefficient.

---

[1] In the original exam, the variables were erroneously listed as X_1, X_2, X_3, Z. Also, the specification of the model has been clarified from the original exam.

Appendix: Some probability distributions

The Bernoulli distribution

If $x \in \{0, 1\}$ has a Bernoulli(p) distribution, with $0 \leq p \leq 1$, then the probability mass function is

$$\pi(x) = p^x(1-p)^{1-x}.$$

The Beta distribution

If $x \geq 0$ has a Beta(α,β) distribution, with $\alpha > 0$ and $\beta > 0$, then the density is

$$\pi(x \mid \alpha,\beta) = \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} x^{\alpha-1}(1-x)^{\beta-1}.$$

We have $\mathbb{E}[x] = \frac{\alpha}{\alpha+\beta}$ and $\mathrm{Var}[x] = \frac{\alpha\beta}{(\alpha+\beta)^2(\alpha+\beta+1)}$.
The mode is $\frac{\alpha-1}{\alpha+\beta-2}$.

The Binomial distribution

If $x \in \{0, 1, 2, \dots, n\}$ has a Binomial(n, p) distribution, with n a positive integer and $0 \leq p \leq 1$, then the probability mass function is

$$\pi(x \mid n, p) = \binom{n}{x} p^x(1-p)^{n-x}.$$

We have $\mathbb{E}[x] = np$ and $\mathrm{Var}[x] = np(1-p)$.

The Gamma distribution

If x has a Gamma(α,β) distribution, with $\alpha > 0$ and $\beta > 0$, then the density is

$$\pi(x \mid \alpha\beta) = \frac{\beta^\alpha}{\Gamma(\alpha)} x^{\alpha-1} \exp(-\beta x)$$

The Logistic distribution

If x has a Logistic(μ, s) distribution, with $s > 0$, then the density is

$$\pi(x \mid \mu, s) = \frac{\exp\left(-\frac{x-\mu}{s}\right)}{s\left(1+\exp\left(-\frac{x-\mu}{s}\right)\right)^2}$$

and the cumulative density functino is given by

$$F(x) = \frac{1}{1+\exp\left(-\frac{x-\mu}{s}\right)}.$$