Petter Mostad
Applied Mathematics and Statistics
Chalmers and GU

MVE550 Stochastic Processes and Bayesian Inference

Exam 19/8 2019, 14:00 - 18:00
Allowed aids: Chalmers-approved calculator.
Total number of points: 30. To pass, at least 12 points are needed.
There is an appendix containing information about some probability distributions.
Unless explicitly allowed, an answer is not complete without a supporting computation or argument.

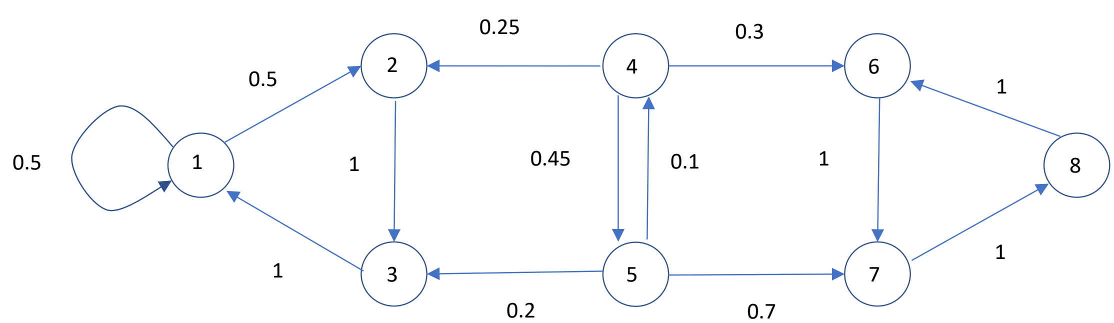

Figure 1: The transition graph for question 1.

1. (7 points) Consider the transition graph of Figure 1 for a discrete time Markov chain.

   (a) List the communication classes. Which of these classes are closed? (No explanation needed).

   (b) List the recurrent states. List the transient states. (No explanations are needed).

   (c) List all proper subsets of the eight nodes such that, if you consider the nodes in the subset and the transition probabilities between these nodes as indicated in the figure, you have an *ergodic* Markov chain. (Explain your conclusion.)

(d) Assume a chain starts at node 1. Compute the limit, as the number of steps goes to infinity, of the probability of being at node 2.

(e) Assume a chain starts at node 8. Compute $\lim_{n\to\infty} \frac{1}{n} \sum_{m=0}^{n-1} P^m_{8,7}$, where P is the transition matrix of the Markov chain.

2. (4 points) Assume Y has a negative Binomial distribution with parameters p and r (see the Appendix). Assume p has a uniform prior on the interval (0,1) and that r is fixed.

   (a) Find the name of and the parameter or parameters of the posterior density for p given an observation y.

   (b) Before any observations of y have been made, compute the expression for the marginal probability mass function for y, i.e., the probability mass function taking into account the uncertainty in p expressed in the prior.

3. (2 points) Assume a discrete-time Markov chain on the state space consisting of A, B, and C has been observed for 17 steps, with the following values:

$$B, A, C, C, B, B, C, A, A, C, B, B, A, B, C, B, A$$

   (a) Write down an estimate for the transition matrix P based on observed frequencies.

   (b) Assume we use a prior for the transition matrix consisting of a product of Dirichlet distributions, with all pseudo-counts equal to 1. What is the form of the posterior distribution given the data above? What is the expectation of this posterior?

4. (5 points) Assume a discrete probability distribution is specified with a probability vector $p = (p_1, p_2, \dots,)$. Let T be a transition matrix for this state space. The goal of the Metropolis-Hastings algorithm is to define a Markov chain $X_0, X_1, X_2, \dots$ with stationary distribution p.

   (a) Write down the Metropolis-Hastings algorithm where T is used for proposals.

   (b) Prove that the resulting chain $X_0, X_1, \dots$ is time-reversible.

   (c) Are there extra conditions needed to ensure that $X_0, X_1, \dots$ has p as a limiting distribution? If so, what is this condition or what are these conditions?

5. (5 points)

   (a) What is the definition of a Branching process?

   (b) How do you define a critical, supercritical, and a subcritical Branching process?

   (c) Assume the offspring distribution is Poisson with parameter λ. Find the probability generating function for the offspring distribution.

   (d) Assume $\lambda > 1$ and let s be the probability of extinction of the Branching process. Find and simplify an equation that s must satisfy, i.e., one which may be used to compute s for a given λ.

6. (3 points)

   (a) Explain briefly what Gibbs sampling is.

   (b) Explain briefly what Perfect sampling is.

   (c) For the matrix exponential, prove that $e^{(s+t)A} = e^{sA}e^{tA}$.

7. (4 points) A machine has three states: It works OK, it works in a stressed state, or it is broken. If it is OK, it wil stay OK for an exponentially distributed amount of time, with expectation 1000 hours. It will then go into the stressed state. If it is in a stressed state, it will break, according to a Poisson process with rate 0.1 per hour, or it will return to the OK state, according to an independent Poisson process with rate 0.5 per hour.

   (a) Write down the generator matrix Q.

   (b) If it is in the stressed state, what is the expected length of time it will stay in this state before it moves to another state[1]?

   (c) If it starts out OK, what is the expected time until it breaks?

Appendix: Some probability distributions

The Bernoulli distribution

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

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

We write $x \mid p \sim \text{Bernoulli}(p)$ and $\pi(x \mid p) = \text{Bernoulli}(x; p)$.

The Beta distribution

If $x \in [0, 1]$ has a Beta distribution with parameters 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 write $x \mid \alpha, \beta \sim \text{Beta}(\alpha, \beta)$ and $\pi(x \mid \alpha, \beta) = \text{Beta}(x; \alpha, \beta)$.

[1] In the original exam, the last part of the sentence ("before it moves to another state") was missing, making the question somewhat less clear.

The Beta-Binomial distribution

If $x \in \{0, 1, 2, \ldots, n\}$ has a Beta-Binomial distribution, with n a positive integer and parameters $\alpha > 0$ and $\beta > 0$, then the probability mass function is

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

We write $x \mid n, \alpha, \beta \sim \text{Beta-Binomial}(n, \alpha, \beta)$ and $\pi(x \mid n, \alpha, \beta) = \text{Beta-Binomial}(x; n, \alpha, \beta)$.

The Binomial distribution

If $x \in \{0, 1, 2, \ldots, n\}$ has a Binomial 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 write $x \mid n, p \sim \text{Binomial}(n, p)$ and $\pi(x \mid n, p) = \text{Binomial}(x; n, p)$.

The Dirichlet distribution

If $x = (x_1, x_2, \ldots, x_n)$ has a Dirichlet distribution, with $x_i \geq 0$ and $\sum_{i=1}^{n} x_i = 1$ and with parameters $\alpha = (\alpha_1, \ldots, \alpha_n)$ with $\alpha_1 > 0, \ldots, \alpha_n > 0$, then the density function is

$$\pi(x \mid \alpha) = \frac{\Gamma(\alpha_1 + \alpha_2 + \cdots + \alpha_n)}{\Gamma(\alpha_1)\Gamma(\alpha_2)\cdots\Gamma(\alpha_n)} p_1^{\alpha_1 - 1} p_2^{\alpha_2 - 1} \cdots p_n^{\alpha_n - 1}.$$

We write $x \mid \alpha \sim \text{Dirichlet}(\alpha)$ and $\pi(x \mid \alpha) = \text{Dirichlet}(x; \alpha)$.

The Exponential distribution

If $x \geq 0$ has an Exponential distribution with parameter $\lambda > 0$, then the density is

$$\pi(x \mid \lambda) = \lambda \exp(-\lambda x)$$

We write $x \mid \lambda \sim \text{Exponential}(\lambda)$ and $\pi(x \mid \lambda) = \text{Exponential}(x; \lambda)$. The expectation is $1/\lambda$ and the variance is $1/\lambda^2$.

The Gamma distribution

If $x > 0$ has a Gamma distribution with parameters $\alpha > 0$ and $\beta > 0$ then the density is

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

We write $x \mid \alpha, \beta \sim \text{Gamma}(\alpha, \beta)$ and $\pi(x \mid \alpha, \beta) = \text{Gamma}(x; \alpha, \beta)$.

The Negative Binomial distribution

If $x \in \{0, 1, 2, \ldots, n\}$ has a Negative Binomial distribution, with parameters r a positive integer and p satisflying $0 \leq p \leq 1$, then the probability mass function is

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

We write $x \mid r, p \sim \text{Negative-Binomial}(r, p)$ and $\pi(x \mid r, p) = \text{Negative-Binomial}(x; r, p)$.

The Normal distribution

If the real x has a Normal distribution with parameters μ and σ^2, its density is given by

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

We write $x \mid \mu, \sigma^2 \sim \text{Normal}(\mu, \sigma^2)$ and $\pi(x \mid \mu, \sigma^2) = \text{Normal}(x; \mu, \sigma^2)$.

The Poisson distribution

If $x \in \{0, 1, 2, \ldots\}$ has Poisson distribution with parameter $\lambda > 0$ then the probability mass function is

$$e^{-\lambda} \frac{\lambda^x}{x!}.$$

We write $x \mid \lambda \sim \text{Poisson}(\lambda)$ and $\pi(x \mid \lambda) = \text{Poisson}(x; \lambda)$.