Petter Mostad
Applied Mathematics and Statistics
Chalmers and GU

MVE550 Stochastic Processes and Bayesian Inference

Re-exam August 23, 2021, 8:30 - 12:30
Examiner: Petter Mostad, phone 031-772-3579
Allowed aids: All aids are allowed.

For example you may access teaching material on any format and you may use R for computation. However, you are **not** allowed to communicate with any person other than the examiner and the exam guard. Total number of points: 30. To pass, at least 12 points are needed. You need to explain how you derive your answers, i.e., show the steps in computations, unless explicitly stated otherwise. There is an appendix containing information about some probability distributions.

1. (6 points) Assume p with $0 < p < 1$ is a random variable with a uniform distribution. Assume that the random variables $x_1, x_2, \dots,$ are independent given p, that they can have either 0 or 1 as their value, and that given p, the probability for 1 is p.
   (a) What is the marginal probability that $x_1 = 1$?
   (b) Assume that you have observed $x_1 = 1$, $x_2 = 0$, $x_3 = 0$. What is the probability distribution for p given this data?
   (c) Given that you have observed $x_1 = 1$, $x_2 = 0$, and $x_3 = 0$, what is the probability that $x_4 = 0$?

2. (8 points) Consider a branching process where the offspring are produced as follows: First, for $k = 0, 1, 2 \dots,$ we add k children with a probability $2/3^{k+1}$. Then, independently, an extra child is added with a probability a, where $0 < a < 1$.
   (a) Compute the probability generating function for the offspring distribution. Use a form in your answer that does not contain an infinite sum.
   (b) For what values of a is there a positive probability that the branching process will not go extinct?
   (c) Compute the probaility of extinction for any a.
   (d) Assume you have observed the offspring process above in N cases, and that in r of these cases, the number of offspring was zero. Assuming a uniform prior for a and given this information, write down a function of a that is proportional to the posterior density for a on the interval $[0, 1]$.

3. (8 points) Two types of particles arrive at a detector according to independent Poisson processes: Type alpha arrives at a rate of 2.1 per second and type beta at a rate of 4.9 per

second. You may answer the questions below with computed numbers, or for example with explicit R code for how to compute the result. Always explain how you reach your result.

(a) What is the probability that exactly 4 particles reach the detector in the interval between 4.0 and 4.5 seconds from the start of detections?

(b) Given that 5 particles reach the detector within the first second of detections, what is the probability that exactly 3 of these are alpha particles?

(c) Given that 7 beta particles arrive in the first 2 seconds, what is the probability that none of these beta particles arrive during the last half of this period?

(d) Given that 6 alpha particles arrive during the first two seconds, what is the probability that the 6th beta particle arrives in the interval $[1.7, 2]$ seconds?

(e) What is the probability that the first particle to arrive in the detector is an alpha particle?

4. (4 points) A celltype can exist in four states: A growth state, a resting state, a reproduction state, or it can be dead. It starts in the growth state. From the growth state it moves to the resting state with a rate of 0.3 and to the reproduction state with a rate 0.1, or it dies with a rate of 0.1. From the resting state it moves to the growth state with a rate of 0.2, or it dies with a rate of 0.05. From the reproduction state it moves to the resting state at a rate 0.5, or it dies at a rate of 0.3.

   (a) What is the expected life time of the cell? Compute the answer numerically, or describe in detail how such computations are done.

   (b) Write down the transition matrix for the corresponding *discrete-time* embedded Markov chain.

5. (4 points) The future price of a particular patent is modelled using a Brownian motion B_t with a drift: The price after t years is modelled as

$$V(t) = 8000 + 400t + 500B_t$$

   (a) What is the probability that the price is above 10000 after 2 years?

   (b) What is the expected time at which the price reaches 10000?

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 $\pi(x \mid \alpha, \beta) = \text{Beta}(x; \alpha, \beta)$.

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 Geometric distribution

If $x \in \{1, 2, 3, \dots\}$ has a Geometric distribution with parameter $p \in (0, 1)$, the probability mass function is

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

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

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 $x \in \{0, 1, 2, \dots\}$ 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 $\pi(x \mid \lambda) = \text{Poisson}(x; \lambda)$.

Petter Mostad
Applied Mathematics and Statistics
Chalmers

Suggested solutions for
MVE550 Stochastic Processes and Bayesian Inference
Re-exam August 23 2021

1. (a) One way to solve this is to argue using symmetry: We must have that the probability becomes 0.5. More computationally, we may compute the marginal with

$$P(x_1 = 1) = \int_0^1 P(x_1 = 1 \mid p)\pi(p)\,dp = \int_0^1 p\,dp = \left[\frac{1}{2}p^2\right]_0^1 = \frac{1}{2}.$$

But it is also possible to look at this as Bayesian statistics, where p has the prior Beta$(1, 1)$ corresponding to the uniform distribution, and where $x_i \sim$ Binomial$(1, p)$. Using the Beta-Binomial conjugacy, we may compute using the prior predictive, which according to the Compendium becomes

$$P(x_1 = 1) = \text{Beta-Binomial}(1; 1, 1, 1) = \binom{1}{1}\frac{B(1+1, 1-1+1)}{B(1,1)} = \frac{\Gamma(2)\Gamma(1)\Gamma(2)}{\Gamma(3)\Gamma(1)\Gamma(1)} = \frac{1}{2}$$

(b) Here it may be most easy to use the Beta-Binomial conjugacy. As the prior for p is Beta$(1, 1)$ and we update it with one observation of 1 and two of 0, the posterior becomes Beta$(2, 3)$.

(c) Here we may use the posterior predictive of the Beta-Binomial: We get

$$P(x_4 = 0 \mid \text{data}) = \text{Beta-Binomial}(0; 1, 2, 3) = \binom{1}{0}\frac{B(0+2, 1-0+3)}{B(2,3)} = \frac{\Gamma(2)\Gamma(4)\Gamma(5)}{\Gamma(6)\Gamma(2)\Gamma(3)} = \frac{3}{5}.$$

An alternative is to use the find the marginal by integrating out the Beta$(2, 3)$ density derived in (b).

2. (a) If we write X_1 for the random variable representing the first set of children and X_2 for the possible extra child, we get

$$\begin{aligned} G(s) &= E(s^{X_1+X_2}) = E(s^{X_1}) \cdot E(s^{X_2}) = \left(\sum_{k=0}^{\infty} s^k \frac{2}{3^{k+1}}\right)(1 - a + sa) \\ &= \left(\frac{2}{3}\sum_{k=0}^{\infty}\left(\frac{s}{3}\right)^k\right)(1 - a + sa) = \frac{2}{3} \cdot \frac{1}{1 - s/3}(1 - a + sa) = 2\frac{1 - a + sa}{3 - s}. \end{aligned}$$

(b) One way to compute is to use that the expected value of the offspring distribution is $G'(1)$: As $G'(s) = 2(2a + 1)/(3 - s)^2$ we get $G'(1) = a + 1/2$. There is a positive probability that the branching process will not go extinct if and only if this expectation is larger than 1, i.e., if $a > 1/2$. It is also easy to compute the expectation $a + 1/2$ directly from the model (not using $G(s)$).

(c) We find the probability of extinction by finding the smallest positive root of $G(s) = s$, in other words

$$2\frac{1 + a + sa}{3 - s} = s$$

or

$$s^2 + (2a - 3)s + 2 - 2a = 0$$

Using that one root is necessarily $s = 1$ we can factor this to get that the other root is $s = 2 - 2a$. When $1/2 < a < 1$ this is a number in the interval $(0, 1)$, so it is the smallest positive root.

(d) Given a, the probability of no offspring is $\frac{2}{3} \cdot (1 - a)$ as it corresponds to no offspring in both ways of producing children. Thus we get

$$\begin{aligned}
\pi(a \mid \text{data}) &\propto \pi(\text{data} \mid a)\pi(a) \\
&\propto \pi(\text{no offspring} \mid a)^r \pi(\text{some offspring} \mid a)^{N-r} \\
&\propto \left(1 - \frac{2}{3}(1 - a)\right)^r \left(\frac{2}{3}(1 - a)\right)^{N-r} \\
&\propto (1 + 2a)^r (1 - a)^{N-r}
\end{aligned}$$

3. (a) Particles arrive as a Poisson process with a total rate of $2.1 + 4.9 = 7$ per second. Thus, in an interval of length 0.5 the number of particles is distributed according to $\text{Poisson}(7 \cdot 0.5) = \text{Poisson}(3.5)$. The probability for 4 particles can be computed with

```
dpois(4, 3.5)
```

giving 0.1888123.

(b) The number of alpha particles among these 5 particles is distributed as $\text{Binomial}(5, \frac{2.1}{2.1+4.9}) = \text{Binomial}(5, 0.3)$. The probability for 3 alpha particles can be computed with

```
dbinom(3, 5, 0.3)
```

giving 0.1323.

(c) The arrival time of each of the 7 particles is uniform on the interval. Thus the probability that it comes in the first half of the interval is $1/2$. The probability that all particles come in the first half is $(1/2)^7 = 0.0078125$.

(d) What happens with the alpha particles is irrelevant. The sixth beta particle to arrive has an arrival time that is distributed as $\text{Gamma}(6, 4.9)$. The probability that it arrives in the interval given can be computed with

```
pgamma(2, 6, 4.9)-pgamma(1.7, 6, 4.9)
```

giving 0.08779957.

(e) The probability is $\frac{2.1}{2.1+4.9} = 0.3$.

4. (a) Using the notation of Dobrow, we get

$$Q = \begin{bmatrix} -0.5 & 0.3 & 0.1 & 0.1 \\ 0.2 & -0.25 & 0 & 0.05 \\ 0 & 0.5 & -0.8 & 0.3 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

and thus

$$V = \begin{bmatrix} -0.5 & 0.3 & 0.1 \\ 0.2 & -0.25 & 0 \\ 0 & 0.5 & -0.8 \end{bmatrix}.$$

We have that he fundamental matrix is $F = -V^{-1}$ and that the answer is the sum of the top row of F. In R we can compute as follows:

```
V <- matrix(c(-0.5, 0.3, 0.1, 0.2, -0.25, 0 , 0, 0.5, -0.8), 3, 3, byrow=T)
print(sum(-solve(V)[1,]))
```

giving 12.2619.

(b) We get

$$\tilde{P} = \begin{bmatrix} 0 & 0.8 & 0.2 & 0.2 \\ 0.8 & 0 & 0 & 0.2 \\ 0 & 0.625 & 0 & 0.325 \\ 0 & 0 & 0 & 1 \end{bmatrix}.$$

5. (a) We have

$$P(8000 + 400 \cdot 2 + 500B_2 > 10000) = P(B_2 > 2.4)$$

As B_2 is normally distributed with expectation 0 and variance 2, this can be computed to be 0.0448, for example using the R command `1 - pnorm(2.4, 0, sqrt(2))`.

(b) We have that

$$8000 + 400 \cdot t + 500B_t > 10000$$

corresponds to $B_t > 4 - 0.8t$. If T is the smallest value such that $B_T = 4 - 0.8T$ then T is a stopping time. We may then compute (see Example 8.25 in Dobrow for details)

$$0 = E(B_0) = E(B_T) = 4 - 0.8E(T)$$

so solving for $E(T)$ gives $E(T) = 4/0.8 = 5$.