Petter Mostad
Applied Mathematics and Statistics
Chalmers and GU

MVE550 Stochastic Processes and Bayesian Inference

Exam January 13, 2025, 8:30 - 12:30
Examiner: Petter Mostad
Isac Boström will be available by phone 0702470297,
and will visit the exam at 10.
Allowed aids: Chalmers-approved calculator
Total number of points: 30. At least 12 points are needed to pass.
See appendix for some information about some probability distributions.
All answers need to be explicitly computed or explicitly argued for.

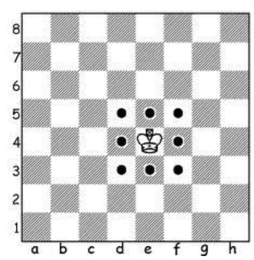

Figure 1: Illustration for question 1.

1. (3 points) Consider a chess board of 8×8 squares, and assume a single king moves around on the board. At every move, the king can move to a neighboring square, see Figure 1. The neighbors are the directly adjacent squares, even diagonally. So the number of neighbor squares is either 8, 5, or 3, depending on the current position of the king.

   (a) How many squares on the board have 8 neighbors? How many have 5? And how many have 3?

   (b) If a king starts in the upper left hand corner of the board, what is the expected number of moves until he returns to the same square?

2. (6 points) In a production process, there is a probability p that items manufactured are faulty. To investigate why they are faulty, Anton collects objects by checking each produced item and storing it if it is faulty, continuing until he has 10 faulty objects. If y denotes the number of non-faulty objects he needs to check in this process, y will have a Negative Binomial distribution.

   (a) Prove that the family of Beta distributions is a conjugate prior family for the probability parameter of the Negative Binomial distribution.

(b) Before collecting any items, Anton's beliefs about p is expressed in a Uniform distribution on the interval $(0, 1)$. During his first data collection, he ends up checking 297 non-faulty objects in order to obtain 10 faulty ones. Using this information, derive the formula for the probability distribution for y, the number of non-faulty objects checked the second time Anton collects objects to obtain 10 new faulty ones.

3. (8 points) Below, you may use the following: If $X_1 \sim \text{Gamma}(\alpha_1, \beta)$ and independently $X_2 \sim \text{Gamma}(\alpha_2, \beta)$ then $X_1 + X_2 \sim \text{Gamma}(\alpha_1 + \alpha_2, \beta)$. You may report your numerical answers either as numbers, or as R code (or pseudo-code if you cannot remember the exact syntax) computing such numbers: Such code should then only use elementary functions like $+, -, \cdot$ etc. and distribution functions like `dnorm, pnorm, qnorm` and similar.

   (a) Assume customers arrive as a Poisson process with the rate of 0.2 per minute. What is the probability of the event that there will be 10 customers during the next hour and that 4 of these will arrive during the last 10 minutes?

   (b) Assume customers arrive as a Poisson process with the rate of 0.2 per minute. If we know that exactly 10 customers arrive during an hour, what is the probability that 4 of these arrive during the last 10 minutes?

   (c) Let S_n be the arrival time of the n'th event for a Poisson process with parameter λ. Write down *and prove* the probability distribution for S_n.

   (d) Assume customers arrive as a Poisson process with the rate of 0.2 per minute. Given that the tenth customer arrives before 60 minutes, what is the expected arrival time of the tenth customer?

4. (6 points) Paul, Kim, and Anna are sharing a house. They have a common laundry room, but very little sense of organization. Whenever either of them wants to do laundry, they go to the laundry room to see if it is free; if it is not, they walk away and return at some other time. Each of them check the laundry room according to a Poisson process, with the rate of three times a week (assume for simplicity that they do not differentiate between night and day time in this household). The time Paul uses for his laundry is exponentially distributed with an expected value of 5 hours. The situation for Kim and Anna is similar, with expected laundry times 3 and 6 hours, respectively.

   (a) Describe the process above as a continuous-time discrete state space Markov chain: Explicitly list and describe the states and the generator matrix Q.

   (b) Is the process time reversible? Prove your answer.

   (c) In Kim's laundry there is a piece of clothing that will instantly wreck the washing-machine. Assuming nobody is currently doing laundry, what is the expected time until the machine is wrecked? You may represent your answer as an equation using matrix operations on matrices of numbers.

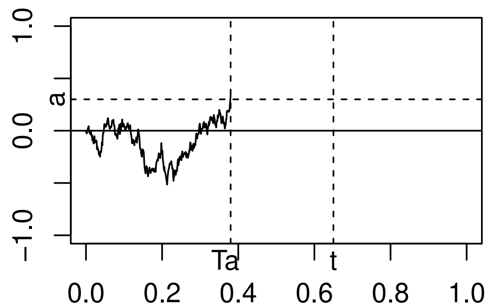

Figure 2: Illustration for question 5

5. (6 points) Let B_t denote Brownian motion and for $a > 0$ let T_a denote the first hitting time of a (the first time the process reaches a), i.e., $T_a = \min(t > 0 : B_t = a)$. See Figure 2. Prove each of the statements below:[1]

   (a) $\Pr(B_t > a \mid T_a < t) = \frac{1}{2}$

   (b) $2\Pr(B_t > a) = \Pr(T_a < t)$

   (c) T_a has the same probability distribution as $\left(\frac{1}{B_{1/a^2}}\right)^2$.

6. (1 point) What exactly is the Ising model, and how can you most efficiently generate a sample from it? Answer using at most three sentences.

[1]You may use each statement in the proof of the next one

Appendix: Some probability distributions

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 Gamma$(\alpha,\beta)$ distribution has expectation $\frac{\alpha}{\beta}$ and variance $\frac{\alpha}{\beta^2}$.

The Inverse Gamma distribution

If $x > 0$ has an Inverse 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 $x \mid \alpha,\beta \sim \text{Inverse-Gamma}(\alpha,\beta)$ and $\pi(x \mid \alpha,\beta) = \text{Inverse-Gamma}(x;\alpha,\beta)$. If $x \sim \text{Gamma}(\alpha,\beta)$ then $1/x \sim \text{Inverse-Gamma}(\alpha,\beta)$. The Inverse-Gamma$(\alpha,\beta)$ distribution has expectation $\frac{\beta}{\alpha-1}$ and variance $\frac{\beta^2}{(\alpha-1)^2(\alpha-2)}$.

The Geometric distribution

If $x \in \{1, 2, 3, \ldots\}$ 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 Negative Binomial distribution

A stochastic variable x taking on as possible values any nonnegative integer has a Negative Binomial distribution if its probability mass function is given by

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

where $r > 0$ and $p \in (0, 1)$ are parameters. We write $x \sim \text{Negative-Binomial}(r, p)$ 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 $\pi(x \mid \lambda) = \text{Poisson}(x; \lambda)$. The Poisson distribution has expectation λ and variance λ.

Petter Mostad
Applied Mathematics and Statistics
Chalmers

Suggested solutions for
MVE550 Stochastic Processes and Bayesian Inference
Exam January 13 2025

1. (a) 4 have 3 neighbors, $6 \cdot 4 = 24$ have 5 neighbors, and $6 \cdot 6 = 36$ have 8 neighbors.

   (b) The game represents a random walk on an undirected graph with 64 nodes, where 4 have 3 neighbors, 24 have 5 neighbors, and 36 have 8 neighbors. It is irreducible, and the number of edges, counting each edge twice, is $4 \cdot 3 + 24 \cdot 5 + 36 \cdot 8 = 420$. Thus, the limiting probability of being in the upper left hand corner is $3/420 = 1/140$, and the expected number of moves until revisiting the square is 140.

2. (a) If $y \mid p \sim$ Negative-Binomial$(y; r, p)$ and $p \sim$ Beta(α, β) then

$$\begin{aligned}
\pi(p \mid y) &\propto_p \pi(y \mid p)\pi(p) \\
&= \text{Negative-Binomial}(y; r, p)\,\text{Beta}(p; \alpha, \beta) \\
&\propto_p (1-p)^y p^r p^{\alpha-1}(1-p)^{\beta-1} \\
&= p^{\alpha+r-1}(1-p)^{\beta+y-1} \\
&\propto_p \text{Beta}(p; \alpha + r, \beta + y).
\end{aligned}$$

   This shows that if the prior is Beta-distributed, then the posterior is also Beta-distributed, and we have proved conjugacy.

   (b) To solve this, we need to find the predictive distribution for the Beta-Negative-Binomial conjugacy pair. We can compute

$$\begin{aligned}
\pi(y) &= \frac{\pi(y \mid p)\pi(p)}{\pi(p \mid y)} \\
&= \frac{\text{Negative-Binomial(y; r, p) Beta(p; }\alpha, \beta)}{\text{Beta}(p; \alpha + r, \beta + y)} \\
&= \frac{\binom{y+r-1}{y} \cdot (1-p)^y p^r \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} p^{\alpha-1}(1-p)^{\beta-1}}{\frac{\Gamma(\alpha+\beta+r+y)}{\Gamma(\alpha+r)\Gamma(\beta+y)} p^{\alpha+r-1}(1-p)^{\beta+y-1}} \\
&= \binom{y+r-1}{y} \cdot \frac{\Gamma(\alpha+r)}{\Gamma(\alpha)} \cdot \frac{\Gamma(\beta+y)}{\Gamma(\beta)} \cdot \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha+\beta+r+y)}.
\end{aligned}$$

   Now, the original prior corresponds to a Beta(1, 1) distribution, so the posterior for p after the first experiment becomes Beta$(1 + 10, 1 + 297) =$ Beta$(11, 298)$. Plugging $\alpha = 11$ and $\beta = 298$ into the predictive distribution above, together with $r = 10$, we get

$$\pi(y) = \binom{y+9}{y} \cdot \frac{\Gamma(21)}{\Gamma(11)} \cdot \frac{\Gamma(298+y)}{\Gamma(298)} \cdot \frac{\Gamma(309)}{\Gamma(319+y)}.$$

3. (a) There are unfortunately several different reasonable interpretations of the question. The intended interpretation was "What is the probability that there will be exactly 10 customers during th next hour and that exactly 4 of these will arrive during the last 10 minutes?"

The simplest solution is then to compute and multiply the independent probabilities that there will be 6 customers during the first 50 minutes and 4 customers during the last 10 minutes:

$$\text{Poisson}(6, 0.2 \cdot 50) \cdot \text{Poisson}(4, 0.2 \cdot 10) = e^{-10}\frac{10^6}{6!} \cdot e^{-2}\frac{2^4}{4!} = 0.005689086$$

R code could be

```
dpois(6,10)*dpois(4,2)
```

An alternative could be to first compute the probability that there are 10 customers all together and then the conditional probability that 4 of these show up in the last 10 minutes:

$$\text{Poisson}(10, 0.2 \cdot 60) \cdot \text{Binomial}(4, 10, 1/6) = e^{-12}\frac{12^{10}}{10!} \cdot \frac{10!}{4!6!}\left(\frac{1}{6}\right)^4\left(\frac{5}{6}\right)^6$$

simplifying to the same expression as above. R code could be

```
dpois(10,0.2*60)*dbinom(4,10,1/6)
```

Unfortunately, there are some other alternative reasonable interpretations:

- "What is the probability that there will be exactly 10 customers during the next hour and that 4 or more of these will arrive during the last 10 minutes?" An answer can for example be computed with

  ```
  dpois(10, 0.2*60)*sum(dbinom(4:10,10,1/6))
  ```

  yielding 0.007310076.
- "What is the probability that there will be 10 or more customers during the next hour and that exactly 4 of these will arrive during the last 10 minutes?" This is equivalent to 6 or more customers during the first 50 minutes and exactly 4 customers in the following 10 minutes. An answer can for example be computed with

  ```
  pgamma(50, 6, 0.2)*dpois(4,0.2*10)
  ```

  yielding 0.08417079.
- "What is the probability that there will be 10 or more customers during the next hour and that 4 or more of these will arrive during the last 10 minutes?" We may separate into the case where there are 10 or more customers in the last 10 minutes, and each of the cases where there are 4 through 9 customers in the last 10 minutes. The result can then be computed with

```
1 - ppois(9, 0.2*10) + sum(dpois(4:9, 0.2*10)*
     (1-ppois(5:0, 0.2*50)))
```

yielding 0.1356338.

(b) The conditional probability was in fact mentioned in the second solution above. (Note that, when the total number of customers is given, the arrival time of each one is uniformly distributed). We get

$$\text{Binomial}(4, 10, 1/6) = \frac{10!}{4!6!}\left(\frac{1}{6}\right)^4\left(\frac{5}{6}\right)^6 = 0.05426588$$

and in R

```
dbinom(4,10,1/6)
```

However, again, a possible interpretation is that we want the probability of 4 or more customers during the last 10 minutes. The solution then becomes

```
sum(dbinom(4:10,10,1/6))
```

yielding 0.06972784.

(c) The formula is

$$S_n \sim \text{Gamma}(n, \lambda).$$

To prove it note that

$$S_n = X_1 + X_2 + \cdots + X_n$$

where the X_i are independent and $X_i \sim \text{Exponential}(\lambda)$. Comparing the densities, we see that the Exponential(λ) distribution is equal to the Gamma($1, \lambda$) distribution. Using the hint in the question, the result follows.

(d) The 10th arrival has distribution

$$S_{10} \sim \text{Gamma}(10, 0.2)$$

and if we know that the arrival happens before 60 minutes, the truncated distribution has density

$$\pi(s) = \frac{\text{Gamma}(s, 10, 0.2)}{\int_0^{60} \text{Gamma}(s, 10, 0.2)\, ds}.$$

The conditional expectation can then be computed with

$$
\begin{aligned}
\mathrm{E}(S_{10} \mid S_{10} < 60) &= \frac{\int_0^{60} s \,\mathrm{Gamma}(s; 10, 0.2)\, ds}{\int_0^{60} \mathrm{Gamma}(s; 10, 0.2)\, ds} \\
&= \frac{\int_0^{60} s \frac{0.2^{10}}{\Gamma(10)} s^9 \exp(-0.2s)\, ds}{\int_0^{60} \mathrm{Gamma}(s; 10, 0.2)\, ds} \\
&= \frac{\frac{\Gamma(11)}{\Gamma(10)} \cdot \frac{0.2^{10}}{0.2^{11}} \int_0^{60} \mathrm{Gamma}(s, 11, 0.2)\, ds}{\int_0^{60} \mathrm{Gamma}(s; 10, 0.2)\, ds} \\
&= \frac{10}{0.2} \cdot \frac{\int_0^{60} \mathrm{Gamma}(s, 11, 0.2)\, ds}{\int_0^{60} \mathrm{Gamma}(s; 10, 0.2)\, ds}
\end{aligned}
$$

This can be computed for example with the R code

```
50*pgamma(60, 11, 0.2)/pgamma(60, 10, 0.2)
```

yielding 43.08103.

4. (a) We can use 4 states:

   1. Nobody uses the laundry room.
   2. Paul uses the laundry room.
   3. Kim uses the laundry room.
   4. Anna uses the laundry room.

   The hourly rate at which either of them checks the laundry room is $3/(24 \cdot 7) = 1/56$, and the hourly rates at which they finish their laundry are $1/5, 1/3$, and $1/6$, respectively. (One may also choose to use daily rates or weekly rates, but all rates must have the same unit). With these states and rates the generator matrix becomes

$$
Q = \begin{bmatrix} \frac{-3}{56} & \frac{1}{56} & \frac{1}{56} & \frac{1}{56} \\ 1/5 & -1/5 & 0 & 0 \\ 1/3 & 0 & -1/3 & 0 \\ 1/6 & 0 & 0 & -1/6 \end{bmatrix}
$$

   (b) As the rates graph is a tree, with state 1 in the center and states 2, 3, and 4 representing branches, the process must be time reversible.

   (c) Making state 3 absorbing, the remaining part of the generator matrix becomes

$$
V = \begin{bmatrix} -3/56 & 1/56 & 1/56 \\ 1/5 & -1/5 & 0 \\ 1/6 & 0 & -1/6 \end{bmatrix}.
$$

The fundamental matrix is $F = (-V)^{-1}$ and the answer to the question is the sum of the top row of F. That can be computed with

$$\begin{bmatrix} 1 & 0 & 0 \end{bmatrix} \begin{bmatrix} 3/56 & -1/56 & -1/56 \\ -1/5 & 1/5 & 0 \\ -1/6 & 0 & 1/6 \end{bmatrix}^{-1} \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}.$$

5. (a) As T_a is a stopping time, $\{B_{T_a+s} - a\}_{s\geq 0}$ is Brownian motion. By symmetry of Brownian motion, $\Pr(B_{T_a+s} - a > 0) = \frac{1}{2}$ for any $s > 0$. If we assume $T_a < t$ we may choose $s = t - T_a$ to get $\Pr(B_t - a > 0 \mid T_a < t) = \frac{1}{2}$. In other words,

$$\Pr(B_t > a \mid T_a < t) = \frac{1}{2}$$

(b) Multiplying the statement from (a) with $2\Pr(T_a < t)$ on both sides yields

$$2\Pr(B_t > a \mid T_a < t)\Pr(T_a < t) = \Pr(T_a < t)$$

or

$$2\Pr(B_t > a, T_a < t) = \Pr(T_a < t).$$

But as $B_t > a$ implies $T_a < t$, the above yields

$$2\Pr(B_t > a) = \Pr(T_a < t).$$

(c) We may write, for any $t > 0$,

$$\Pr\left(\left(\frac{1}{B_{1/a^2}}\right)^2 < t\right) = \Pr\left(1 < tB^2_{1/a^2}\right) = \Pr\left(1 < \frac{1}{a^2}B_t^2\right) = \Pr(a^2 < B_t^2) = 2\Pr(a < B_t)$$

Combining with the statement from (b) we get, for any $t > 0$

$$\Pr(T_a < t) = \Pr\left(\left(\frac{1}{B_{1/a^2}}\right)^2 < t\right)$$

which show that the two have the same distribution.

6. The Ising model is a probability distribution on a set of states, where each state σ is a combination of states for nodes in a grid, and each node can have the state +1 or -1.

The probability of a state σ is proportional to $\exp(\beta \sum_{i\sim j} \sigma_i \sigma_j)$, where β is a parameter, σ_i and σ_j are the states at nodes i and j respectively, and $i \sim j$ means that i and j are neighbors in the grid.

To simulate a realization, one may use Gibbs sampling, but the most efficient method is to use perfect sampling, which is possible for this model.