
Tutorial Exercises

Exercise 1 - 1

A number of k persons meet safely and socially distanced via zoom. Assume that the probability of a person to have his/her birthday is the same for every day of the year. Assume further that the number of days per year is always 365.

- What is the probability that at least q person's birthday is the first of January?
- What is the probability of at least two persons in the zoom meeting having their birthday on the same day?
- For which k is this probability larger than 50%?

Exercise 1 - 2

You take part in a gameshow. At one point in the show the host presents you with three doors, each hiding one prize. You get to choose one of the doors and get to keep whatever is behind it. Two of the doors are hiding a goat and one is hiding a sports car. After you have made your choice the host, who knows which door is hiding the car, opens one of the doors you have not chosen, making sure he is revealing a goat. Now he asks you if you want to stick to your original choice or if you would like to get what is behind the third door.

Should you change to the third door (assuming you prefer cars over goats)? Give a formal proof of your answer using Bayes' Theorem.

Hint: Assume the host to play fair and to always reveal a goat behind another door before one chooses.

Exercise 1 - 3

We want to prove some basic properties of the Poisson distribution $\mathcal{P}(x|\lambda) = \frac{\lambda^x e^{-\lambda}}{x!}$ with counts $x \in \mathbb{N}_0^+$ and rate $\lambda > 0$.

- Prove that this distribution is normalized.
- Prove that its mean is λ .
- Prove that its variance is λ .
- For a Poisson process with a rate of $\lambda = 1.5$, what is the probability to see 6 or more events? What is the probability of exactly 0 events?

Exercise 1 - 4

We want to prove some basic properties of the Uniform distribution on the interval $[a, b]$ with $a < b$:

$$\mathcal{P}(x|a, b) = \begin{cases} 1/(b-a) & a \leq x \leq b \\ 0 & \text{else.} \end{cases} \quad (1)$$

- Prove that the Uniform distribution is normalized.
- Prove that its mean is at the center of the interval.
- Prove that its variance is $\frac{(b-a)^2}{12}$.

Exercise 1 - 5

We want to empirically validate the Central Limit Theorem (CLT) for several distributions. It tells us that the average of random variables drawn from a distribution approaches normality. For one experiment we draw N samples from the distribution and average them. We repeat this M times. Calculate the theoretically expected parameters of the Gaussian with the N samples. Plot this distribution together with a normalized histogram of the $M = 100$ and $M = 100\,000$ experiments.

- a) Uniform distribution on the unit interval with $N = 10$.
- b) Bernoulli distribution with a probability $\mu = 0.3$, $N = 100$.
- c) Gaussian distribution with mean $m = 1$, variance $\sigma^2 = 1$, and $N = 2$.
- d) The CLT only holds if the original sampling distribution has a finite mean and variance. An example for which this is not the case is the Cauchy distribution:

$$\mathcal{P}(x|x_0, \gamma) = \frac{1}{\pi\gamma} \frac{\gamma^2}{(x - x_0)^2 + \gamma^2} \quad (2)$$

Perform the same experiments as above for $N = 100$ with $x_0 = 3$ and $\gamma = 3$ and discuss your findings.

Hint: You can generate a sample from the Cauchy distribution by first drawing a sample u from the uniform distribution over the unit interval and transform it according to:

$$x = \gamma \tan(\pi u - \pi/2) + x_0 \quad (3)$$

Exercise 1 - 6

We want to simulate the coin toss experiment and learn the probability of the coin to show either side.

- a) Set some ground truth $\mu_{\text{true}} \in (0, 1)$ and draw one sample from the Bernoulli distribution to simulate a coin-toss.
- b) A priori we assume a flat prior on the unit interval, i.e. a Beta distribution with parameters $\alpha = \beta = 1$. Plot this prior distribution and indicate the ground truth μ_{true} as vertical line.
- c) As the Beta distribution is the conjugate prior to the Bernoulli likelihood, the posterior will be again a Beta distribution with updated α and β parameters. Use the result from the experiment in the beginning and calculate the posterior. Add the curve of this distribution to the previous plot.
- d) We want to perform continual learning and use this updated distribution as prior for future experiments. Simulate three more coin tosses with the same rate, calculate the posterior distribution, and add it to the plot.
- e) Repeat this for another 100 tosses and discuss your findings.

Exercise 1 - 7

We want to implement the temperature measurement as discussed in the lecture. The prior distribution over temperatures is the Gaussian $P(x) = \mathcal{N}(x|x_0, T) = \mathcal{N}(x|x_0 = 295, T = 3^2)$ and the measurement noise is given by $\mathcal{P}(n) = \mathcal{N}(n|n_0 = 0, N = 1^2)$.

- a) Draw a true temperature from the prior distribution $x_{\text{true}} \sim \mathcal{P}(x)$, a noise realization $n \sim \mathcal{P}(n)$, and generate data by adding the true temperature and measurement noise $y = x_{\text{true}} + n$.

- b)** Plot the Prior distribution and indicate the data and ground truth as vertical lines.
- c)** Calculate posterior mean and variance and add the posterior distribution to the previous plot.
- d)** Repeat this procedure for various noise levels (but constant true temperature) and discuss your findings.