## Practice Final 1

1. You observe a single sample $X$ of a Uniform$(-\Theta, \Theta)$ random variable with Bayesian Uniform$(0, 1)$ prior on the parameter $\Theta$.

   (a) If $X = 0.5$, what is the posterior PDF of $\Theta$?

   **Solution:** The posterior is proportional to

   $$f_{\Theta|X}(\theta|0.5) \propto \mathrm{P}(X = 0.5|\Theta = \theta) f_\Theta(\theta) = \begin{cases} 1/2\theta, & \text{if } 0.5 \le \theta \le 1 \\ 0, & \text{otherwise.} \end{cases}$$

   As $\int_{0.5}^{1} d\theta/2\theta = (\ln 2)/2$, the posterior equals $1/(\theta \ln 2)$ when $0.5 \le \theta \le 1$ and 0 otherwise.

   (b) What is the conditional expectation of $\Theta$ given $X = 0.5$?

   **Solution:** $\mathrm{E}[\Theta|X = 0.5] = \int_{-\infty}^{\infty} \theta \cdot f_{\Theta|X}(\theta|0.5) d\theta = \int_{0.5}^{1} d\theta/(\ln 2) = 1/(2\ln 2)$.

2. You observe one sample of a random variable that is either Uniform$(-2, 2)$ (null hypothesis) or Normal$(0, 1)$ (alternative hypothesis).

   (a) Describe the test with the smallest false negative error that achieves false positive error $1/4$.

   **Solution:** The likelihood ratio $f_{H_1}(x)/f_{H_0}(x)$ is proportional to $e^{-x^2/2}$ when $-2 \le x \le 2$ and infinite otherwise. As $e^{-x^2/2}$ is decreasing in $|x|$ the test should output $+$ if $|x| \ge 2$ or $|x| \le t$ for some threshold $t$ and $-$ otherwise. A false positive happens when $|\text{Uniform}(-2, 2)| \le t$, an event of probability $t/2$. To obtain false positive error $1/4$ we choose $t = 1/2$.

   (b) What is the false negative error of your test from part (a)?

   **Solution:** A false negative happens when $1/2 < |N| < 2$ for a Normal$(0, 1)$ random variable $N$. The probability of this event is

   $$\begin{aligned} \mathrm{P}(1/2 < |N| < 2) &= \mathrm{P}(|N| < 2) - \mathrm{P}(|N| \le 1/2) \\ &= (1 - 2\,\mathrm{P}(N \le -2)) - (1 - 2\,\mathrm{P}(N < -1/2)) \\ &= 2(\mathrm{P}(N < -1/2) - \mathrm{P}(N \le -2)) \\ &\approx 2(0.3085 - 0.0228) \\ &= 0.5716. \end{aligned}$$

3. The number of complaints that TV station X receives from its viewers in a given day is modeled by a Normal$(\mu, 10)$ random variable for some unknown $\mu$ that models viewer unhappiness. Station X received $M = 95$ complaints on Monday and $T = 130$ complaints on Tuesday.

   (a) Assuming $M$ and $T$ are independent, what kind of random variable is $T - M$?

   **Solution:** $T - M$ is a difference of independent normal random variables and is therefore normal with mean $\mu_T - \mu_M$ (the difference between viewer unhappiness) and standard deviation $\sqrt{10^2 + 10^2} = \sqrt{2} \cdot 10$.

(b) What is the p-value for the null hypothesis that unhappiness didn't increase?

**Solution:** Assuming $\mu_T = \mu_M$, $T - M$ is Normal$(0, \sqrt{2} \cdot 10)$. The test for the alternative hypothesis should accept if $T - M$ exceeds a threshold $t$ determined by the type I error and reject if not. The p-value of this test for the given data is

$$\mathrm{P}(\mathrm{Normal}(0, \sqrt{2} \cdot 10) \geq 130 - 95) = \mathrm{P}(\mathrm{Normal}(0, 1) \geq 2.475) \approx 0.0067.$$

4. You observe the following statistics in 100 tosses of a possibly unfair 3-sided die:

| outcome: | 1 | 2 | 3 |
|---|---|---|---|
| count: | 27 | 54 | 19 |

You want to test null hypothesis that outcomes 1 and 3 each have probability $\theta$ and outcome 2 has probability $1 - 2\theta$ for some unknown $\theta$.

(a) What is the maximum likelihood estimate of $\theta$ from your data?

**Solution:** The probability of observing this outcome is

$$f(\theta) = \mathrm{P}(N_1 = 27, N_2 = 54, N_3 = 19|\theta) \propto \theta^{27}(1 - 2\theta)^{54}\theta^{19} = \theta^{46}(1 - 2\theta)^{54} \propto (2\theta)^{46}(1 - 2\theta)^{54}.$$

Thus $f(\theta)$ is proportional to the PDF of a Beta$(47, 55)$ random variable evaluated at $2\theta$, which is maximized at $2\hat{\theta} = 46/(46 + 54) = 0.46$, so $\hat{\theta} = 0.23$ is the maximum likelihood estimate of $\theta$. (You could have reached the same conclusion with calculus.)

(b) Can you reject the null hypothesis at the 95% confidence level?

**Solution:** The expected frequencies of the three outcomes assuming $\theta = \hat{\theta}$ are 23, 54, and 23. The chi-square statistic is

$$X^2 = \frac{(27 - 23)^2}{23} + \frac{(54 - 54)^2}{54} + \frac{(19 - 23)^2}{23} \approx 1.391.$$

By Fisher's theorem for parametrized goodness-of-fit, under the null hypothesis $\mathrm{P}(X^2 \geq 1.391) \approx \mathrm{P}(\chi^2(1) \geq 1.391) \approx 0.238$ so it cannot be rejected at the 95% confidence level.

## Practice Final 2

1. You are observing a single sample of a random variable which is either Uniform$(0,2)$ $(\Theta = 0)$ or Uniform$(1,5)$ $(\Theta = 1)$. Your prior on $\Theta$ is $\mathrm{P}(\Theta = 0) = 1/5$ and $\mathrm{P}(\Theta = 1) = 4/5$.

   (a) What is the MAP estimator of $\Theta$?

   **Solution:** The ratio of posteriors is

$$\frac{f_{\Theta|X}(1|x)}{f_{\Theta|X}(0|x)} = \frac{\frac{4}{5} \cdot f_{X|\Theta}(x|1)}{\frac{1}{5} \cdot f_{X|\Theta}(x|0)} = \begin{cases} 0, & \text{if } 0 \leq x < 1 \\ (\frac{4}{5} \cdot \frac{1}{4})/(\frac{1}{5} \cdot \frac{1}{2}), & \text{if } 1 \leq x \leq 2 \\ \infty, & \text{if } 2 < x \leq 5. \end{cases}$$

   As $(\frac{4}{5} \cdot \frac{1}{4})/(\frac{1}{5} \cdot \frac{1}{2}) = 2 > 1$, the MAP estimate is 1 if $1 \leq x \leq 5$ and 0 if $0 \leq x < 1$.

   (b) What is the error probability $\mathrm{P}(MAP \neq \Theta)$?

   **Solution:** When $\Theta = 1$ the MAP estimate is always accurate and there is no error. An error only happens when $\Theta = 0$ and $X \geq 1$, giving an error probability of $\mathrm{P}(X \geq 1, \Theta = 0) = \mathrm{P}(X \geq 1|\Theta = 0)\,\mathrm{P}(\Theta = 0) = \frac{1}{2} \cdot \frac{1}{5} = \frac{1}{10}$.

2. Let $X_1, X_2$ be two independent samples of a Uniform$(0,1)$ random variable.

   (a) Calculate the PDF of the sample mean.

   **Solution:** The PDF of the sum can be calculated by convolution:

$$f_{X_1+X_2}(x) = \int_{-\infty}^{\infty} f_{X_1}(t) f_{X_2}(x-t)dt = \int_0^1 f_{X_2}(x-t)dt.$$

   When $0 \leq x \leq 1$, $f_{X_2}(x-t)$ vanishes when $t > x$ so the integral evaluates to $\int_0^x dt = x$. When $1 < x \leq 2$, $f_{X_2}(x-t)$ vanishes when $t < x-1$ so the integral evaluates to $\int_{x-1}^1 dt = 2 - x$, so

$$f_{X_1+X_2}(x) = \begin{cases} x, & \text{if } 0 \leq x \leq 1, \\ 2 - x, & \text{if } 1 < x \leq 2. \end{cases}$$

   The sample mean $\overline{X}$ equals $(X_1 + X_2)/2$ so its PDF scales to

$$f_{\overline{X}}(x) = \begin{cases} 4x, & \text{if } 0 \leq x \leq 1/2, \\ 4(1-x), & \text{if } 1/2 < x \leq 1. \end{cases}$$

   (b) What is the probability that the actual mean is within $1/4$ of the sample mean?

   **Solution:** The actual mean is $1/2$, so the desired probability is $\mathrm{P}(|\overline{X} - 1/2| \leq 1/4)$. By symmetry of the PDF this is the same as

$$2\,\mathrm{P}(1/4 \leq \overline{X} \leq 1/2) = 2\int_{1/4}^{1/2} 4x dx = 4x^2\Big|_{1/4}^{1/2} = \frac{3}{4}.$$

3. Out of 100 customers polled (randomly with repetition) about the popularity of a new toothpaste, 60 were positive.

   (a) What is the sample mean and sample standard deviation of the popularity of the product?

   **Solution:** The sample consists of 60 ones (yes votes) and 40 zeros (no votes). The sample mean is $\overline{X} = 60/100 = 0.6$ and the sample standard deviation is $\sqrt{\overline{X}(1-\overline{X})} = \sqrt{0.6 \cdot 0.4} = \sqrt{0.24} \approx 0.49$.

   (b) Give a 90% confidence interval for the true popularity. Explain which formula you are using.

   **Solution:** We use the formula for the parameter (mean) of an indicator random variables from Lecture 6. As $\mathrm{P}(|\mathrm{Normal}(0,1)| \leq z) = 90\%$ for $z \approx 1.645$, the confidence interval is $(0.6 - z\sqrt{0.24/100}, 0.6 + z\sqrt{0.24/100}) \approx (0.519, 0.681)$.

4. A review of typists' work at a publishing company on a given day yielded the following data:

| | Alice | Bob | Charlie |
|---|---|---|---|
| morning shift typos | 18 | 16 | 10 |
| afternoon shift typos | 22 | 14 | 12 |

   You want to test the null hypothesis that the number of typos doesn't change between shifts. Assume that each data item comes from normal random variables with the same variance, and that the performance of different typists is independent.

   (a) Would you use the independent sample test or the paired t-test? Justify your answer.

   **Solution:** The paired t-test is more appropriate here because we cannot assume that the number of typos is independent of the typist. For example, if a given typist made many errors in the morning, we would expect them to make many errors in the afternoon as well.

   (b) Calculate the p-value for the test you chose in part (a).

   **Solution:** The difference between the number of afternoon and morning typos is $Z_A = 4$ for Alice, $Z_B = -2$ for Bob, and $Z_C = 2$ for Charlie. The sample mean and adjusted sample variance of this data is $\overline{Z} = 4/3$ and $S^2 = 9\frac{1}{3}$. Under the null hypothesis $T = \overline{Z}/(S/\sqrt{3})$ is a $t(2)$ random variable. As we are testing for a change in performance, the test should accept if $|T| \geq t$ for a threshold $t$ determined by the type I error. The p-value of this test is

$$\mathrm{P}(|T| \geq (4/3)/\sqrt{9\tfrac{1}{3}/3}) \approx \mathrm{P}(|t(2)| \geq 0.756) \approx 0.529.$$

## Practice Final 3

1. A hypothesis test T has a false positive (type I) error of 10%. Assume that the null hypothesis is true.

   (a) You run the test twice on independent data. What is the probability that the test produces inconsistent results?

   **Solution:** If we run the test $n$ times the number of outcomes that reject the null hypothesis (assuming it is true) is a Binomial$(n, 0.1)$ random variable $X$. Here $n$ equals two so the probability of inconsistent results is $\mathrm{P}(X = 1) = 2 \cdot 0.1 \cdot (1 - 0.1) = 0.18$.

   (b) You run the test one more time. What is the probability that a majority (at least 2 out of 3) of test results reject the null hypothesis?

   **Solution:** Now $n$ equals 3 so the desired probability is $\mathrm{P}(X \geq 2) = 3 \cdot 0.1^2 \cdot (1-0.1) + 0.1^3 = 0.028$.

2. A random sample yields the following E.Coli bacteria counts in four one-litre samples of water from a river: 229, 354, 356, 344. Assume that the population counts are independent normal with the same mean and variance.

   (a) Find the sample mean and the adjusted sample variance.

   **Solution:** The sample mean is $\overline{X} = (229 + 354 + 356 + 344)/4 = 320.75$. The adjusted sample variance is $S^2 = ((229 - 320.75)^2 + \cdots + (344 - 320.75)^2)/3 \approx 3768.92$.

   (b) Construct a 95% confidence interval for the true mean of the population.

   **Solution:** The normalized sample mean $(\overline{X} - \mu)/(S/\sqrt{4})$ is a $t(3)$ type random variable. The confidence interval is of the form $(\overline{X} - zS/2, \overline{X} + zS/2)$ where $z$ is chosen so that $\mathrm{P}(|t(3)| > z) \leq 0.05$. This gives $z \approx 3.182$ and the confidence interval $(223, 418)$.

   (c) Give 95% upper confidence bound for the true variance of the population.

   **Solution:** The random variable $3S^2/\sigma^2$ is of type $\chi^2(3)$ so the confidence interval for the variance $\sigma^2$ has the form $(3S^2/z_+, 3S^2/z_-)$. For an upper confidence bound we set $z_+ = \infty$ and choose $z_-$ so that $\mathrm{P}(\chi^2(3) \geq z_-) = 0.95$, which gives $z_- \approx 0.352$ and an upper confidence bound of $3S^2/z_- \approx 32,121$.

3. A coin is tossed 252 times.

   (a) Let $h$ be the fraction of heads observed. Which values of $h$ would indicate that the coin is biased (either towards heads or tails) at the 90% confidence level?

   **Solution:** The number of observed heads is a Binomial$(252, 1/2)$ random variable. By the Central Limit Theorem its CDF is approximately that of a normal random variable $Z$ of mean $\mu = 252/2 = 126$ and standard deviation $\sigma = \sqrt{252 \cdot 1/2 \cdot 1/2} \approx 7.94$. To conclude bias at the 90% confidence level, $Z$ needs to be more than $z \approx 1.645$ standard deviations away from the mean, so it is either less than $\mu - z\sigma \approx 112.94$ or more than $\mu + z\sigma \approx 139.06$ so that $h < 112.94/252 \approx 0.448$ or $h > 139.06/252 \approx 0.552$.

(b) The coin comes up heads 133 times. Can you conclude that it is biased towards tails at the 90% confidence level?

**Solution:** The test for this alternative hypothesis should accept if the number of heads is less than some threshold $t$. In order to have a type I error of 50% or less $t$ cannot be greater than $\mu = 126$. As the number of observed heads is 133 a test at the 90% confidence level will reject the alternative hypothesis.

4. A normal random variable of mean zero has standard deviation either $\sigma = 1$ (null hypothesis) or $\sigma = 2$ (alternative hypothesis).

(a) Given one sample of this random variable, describe the test with false positive (type I) error of 20% that has the smallest false negative (type II) error.

**Solution:** The likelihood ratio is

$$\frac{f_1(x)}{f_0(x)} \propto \frac{\exp(-x^2/8)}{\exp(-x^2/2)} = \exp(x^2/2 - x^2/8) = \exp(3x^2/8).$$

This ratio increases with $|x|$ so by the Neyman-Pearson Lemma the test should accept the alternative hypothesis when $|x| \geq t$ for a suitable threshold $t$ and reject otherwise. The threshold is chosen so that $\mathrm{P}(|Z| \geq t) = 20\%$ for a standard normal $Z$, from where $t \approx 1.282$.

(b) What is the false negative (type II) error of your test?

**Solution:** The false negative error is $\mathrm{P}(|\mathrm{Normal}(0, 2)| < t) \approx P(|Z| < 0.641) \approx 0.478$.