

Example sheet 3
Frequentist inference
Data Science—DJW—2025/2026

Questions labelled * are more challenging.



Question 1. Sketch the cumulative distribution function, and calculate the density function, for this random variable:

```
def rx():  
    u1 = random.random()  
    u2 = random.random()  
    return min(u1,u2)
```

Question 2. We are given data $x = [x_1, \dots, x_m]$ which we believe is sampled from $\text{Exp}(\mu)$, and further data $y = [y_1, \dots, y_n]$ which we believe is sampled from $\text{Exp}(\nu)$.

- (a) We wish to test the hypothesis that $\mu = \nu$. Under this hypothesis, all the datapoints in x and y are sampled from a common distribution $\text{Exp}(\lambda)$, where λ is the common value of μ and ν . Find the maximum likelihood estimator $\hat{\lambda}$. Compute the p -value for this test, using the test statistic $\hat{\nu} - \hat{\mu}$, with parametric resampling.
- (b) Compute a 95% confidence interval for $\hat{\nu} - \hat{\mu}$ using parametric resampling. Explain carefully the difference between your sampling methods for (a) and (b).
- (c) Repeat parts (a) and (b) using non-parametric resampling.
- (d) We wish to test the hypothesis that $\mu \leq \nu$. Compute the p -value for this test, using the same test statistic as before.

[You can test your code using the online tester. There is a notebook with templates for answers and instructions for submission on the Course Materials webpage.]

Question 3 (Cardinality estimation).

- (a) Let X be the maximum of m independent $\text{Uniform}[0, 1]$ random variables. Find the likelihood $\text{Pr}_X(x; m)$.
- (b) Given an observed value x , what is the maximum likelihood estimator \hat{m} ?
- (c) Explain how to compute a 95% confidence interval for \hat{m} .

A common task in data processing is counting the number of distinct items in a collection. When the collection is large, we may wish to use approximate methods that only require $O(1)$ extra memory, such as the following: Given a collection of items a_1, \dots, a_n , compute the hash of each item $u_i = h(a_i)$, then compute $x = \max_i u_i$. If the hash function is well designed, then each u_i can be treated as if it were sampled from $\text{Uniform}[0, 1]$, and distinct items will yield independent samples. The more distinct items there are, the larger we expect x to be. Thus, we can use x to estimate the number of distinct items. Read more: <http://blog.notdot.net/2012/09/Dam-Cool-Algorithms-Cardinality-Estimation>

Question 4. We are given a dataset of (g_i, x_i, y_i) records, $i = 1, \dots, n$, where $g_i \in \{1, 2, 3\}$ is the group that record i belongs to, $x_i \in \mathbb{R}$ is a predictor variable, and $y_i \in \mathbb{R}$ is the response. We are interested in the model

$$Y_i \sim \alpha_{g_i} + \beta_{g_i} x_i + N(0, \sigma^2).$$

- (a) Explain how to fit this model to the dataset.
- (b) We wish to test the hypothesis that $\beta_1 = \beta_2 = \beta_3$. Suggest a test statistic, and describe how to conduct the test.

Question 5*. (a) I toss a coin n times and get x heads. My model is that the number of heads is $\text{Bin}(n, \theta)$ and I wish to test the null hypothesis that $\theta = 1/2$. Explain how to find the p -value for this test.

- (b) I make many attempts at a task, and I have no successes at all, just a string of failures. Modelling my attempts as independent random variables with success probability θ and failure probability $1 - \theta$, how many failures does it take for me to reject $\theta = 1/2$ at p -value 5%?

Question 6. We have a climate dataset of $(\mathbf{t}, \text{temp})$ pairs. Considered a model in which temperatures increase linearly,

$$\text{temp} \sim \alpha + \beta_1 \sin(2\pi\mathbf{t}) + \beta_2 \cos(2\pi\mathbf{t}) + \gamma(\mathbf{t} - 2000) + \text{Normal}(0, \sigma^2).$$

Let $\hat{\gamma}$ be the maximum likelihood estimator for the rate of temperature increase. Explain how to find a 95% confidence interval for $\hat{\gamma}$.

Question 7 (Confidence intervals for predictions). I have computed the maximum likelihood estimators for all the parameters in the model in question 6, and I have used them to define a temperature prediction function

```
def pred(t_new): return  $\hat{\alpha} + \hat{\beta}_1 \sin(2\pi t_{\text{new}}) + \hat{\beta}_2 \cos(2\pi t_{\text{new}}) + \hat{\gamma}(t_{\text{new}} - 2000)$ 
```

Modify this code so that in addition to predicting the temperature it also produces a 95% confidence interval for its prediction.

Question 8. If my p -value is 0.99, does that mean that H_0 is super-duper likely?

Question 9 (Choosing the null)*. In question 5, you explained how to test the null hypothesis $H_0 : \theta = 1/2$.

- (a) Explain how to test the null hypothesis $H_0 : \theta \geq 1/2$, and give code to compute the p -value.
- (b) Explain how to test the null hypothesis $H_0 : \theta \neq 1/2$, and give code to compute the p -value.
- (c) In both cases, investigate which values of x cause you to reject H_0 at threshold $p < 0.05$.

Hints and comments

Question 1. Work through exercise 5.3.4 in lecture notes, then apply the same strategy to this question.

Question 2. Check the definition of the Exponential distribution in lecture notes section 1.2, and watch out for difference in convention between maths (which refers to the rate parameter μ) and numpy (which refers to the scale parameter $1/\mu$). You should find the mles to be $\hat{\mu} = 1/\bar{x}$, $\hat{\nu} = 1/\bar{y}$, $\hat{\lambda} = 1/\bar{z}$ where z is the concatenation of x and y .

For part (a) look at exercise 9.3.2 in lecture notes. For (b) look at exercise 9.2.3. For (c) look at exercises 9.6.1 and 9.6.2. Think about whether your tests and confidence intervals should be one-sided or two-sided.

Part (d) asks you to consider a null hypothesis H_0 that imposes the constraint $\mu \leq \nu$ on the parameters. You need to maximize the likelihood $\Pr(\text{data}; \mu, \nu)$ under the H_0 constraint, as described in section 9.3 under “Tests on parameters”. (Here ‘data’ here refers to the entirety of the dataset, $[x_1, \dots, x_m, y_1, \dots, y_n]$.) Try fixing μ and solving for $\hat{\nu}(\mu)$ using maths, and then solving for $\hat{\mu}$ numerically. Or, you might spot how to solve it all with maths.

Question 3. Part (a) is similar to question 1. Part (b) is asking for the mle from a single datapoint, like exercise 1.3.1 from lecture notes. For part (c) you need to resample the dataset (i.e. the single value X) under fitted null hypothesis, and it’s easy to do this with parametric resampling. There’s a slick trick called Inversion Sampling, not taught in this course, that you may like to look up.

Question 4. In this question you’re given a parametric model, and asked to test a hypothesis that restricts the parameters. Look at section 9.3 of lecture notes under “Tests on parameters” for guidance, and at exercise 9.3.3 for inspiration,

How to choose the test statistic? This question tells us a general hypothesis H_1 , namely that $Y \sim \alpha_g + \beta_g x + N(0, \sigma^2)$; and it proposes a null hypothesis H_0 that is a restriction on the parameters of H_1 , namely that $\beta_0 = \beta_1 = \beta_2$. Can you think up a test statistic using the $\hat{\beta}_g$ parameters from H_1 ? Maybe also use $\hat{\beta}$ from H_0 ?

Next, ask yourself how your statistic would differ between the scenario where H_0 is true, and the scenario where H_0 isn’t true — this will tell you what “more extreme” means, in the definition of p -value, and hence whether to use a one-sided or two-sided test.

Question 5. For part (a), we’re given a general parametric model H_1 that says $X \sim \text{Bin}(n, \theta)$ and asked to test a hypothesis H_0 that restricts the parameters to $\theta = 1/2$. Look at section 9.3 of lecture notes under “Tests on parameters” for guidance. What test statistic shall we use? A good idea is to use the mle $\hat{\theta}$ from H_1 . You should find that $\hat{\theta} = x/n$, so in fact we might as well use x itself as the test statistic.

For this question, we can do much better than just giving pseudocode: we know the distribution that this test statistic will have under H_0 , so we can write out the p -value exactly in terms of the cdf of the Binomial distribution.

For part (b), just use your expression for the p -value from part (a), applied to data $x = 0$. Your expression will depend on n . Find the smallest n such that $p \leq 0.05$.

Question 6. Follow the general strategy from section 9.2 of lecture notes. In your answers for this question, it’s a good idea to use `sklearn` wherever reasonable—there’s no point going through lots of algebra, when there are fast easy routines that you can use. You can generate a synthetic dataset with `np.random.normal(loc=pred, scale= $\hat{\sigma}$)`, as in exercise 9.2.4 lines 14–15, and you can compute the predicted temperatures `pred` as in example 2.1.1 line 13.

Question 7. We want to generate a multiverse of synthetic datasets, and canvas the opinion of data scientists across this multiverse. If a parallel-universe data scientist sees dataset X^* , what value would they produce for `pred(tnew=2050)`? You just need to assemble a large collection of these predictions, then find a 95% confidence interval in the usual way.

For an extra challenge, write your code so that it accepts a vector-valued `tnew`.

Question 9. As in question 2(d), the null hypothesis expresses a constraint on the parameters, so we need to find the maximum likelihood under that constraint. For the test statistic, follow question 5, but think carefully about what ‘more extreme’ means. For the investigation, I suggest you do this numerically, and then (if you are bold and you remember the central limit theorem) try to answer mathematically.

Supplementary questions


These supplementary questions are not intended for supervision (unless your supervisor directs you otherwise). Some of them require careful maths, some are best answered with coding, some are philosophical.

Question 10. Sketch the cumulative distribution function, and calculate the density function, for this random variable:

```
def rx():  
    u = random.random()  
    return u * (1-u)
```

Question 11. A point lightsource at coordinates $(0, 1)$ sends out a ray of light at an angle Θ chosen uniformly in $(-\pi/2, \pi/2)$. Let X be the point where the ray intersects the horizontal line through the origin. What is the density of X ?

Note: This random variable is known as the Cauchy distribution. It is unusual in that it has no mean.



Question 12*. The dataset at https://www.cl.cam.ac.uk/teaching/current/DataSci/data/responsetime_ms.txt is a list of web server response times, measured in milliseconds.

- Plot the empirical cumulative distribution function (ecdf) of this sample.
- Plot the empirical tail distribution function (etdf = $1 - \text{ecdf}$), on a log-log plot.
- You should see that, for large enough response times, the etdf looks roughly like a (noisy) straight line on a log-log plot. Using this observation, estimate the 99.9%ile and 99.99%ile of response time.

Question 13. We are given a dataset x_1, \dots, x_n which we believe is drawn from $\text{Uniform}[0, \theta]$ where θ is unknown. Recall from Example Sheet 1 that the maximum likelihood estimator is $\hat{\theta} = \max_i x_i$. Find a 95% confidence interval for $\hat{\theta}$, both using parametric resampling and using non-parametric resampling.

Question 14. I implement the two resamplers from question 13. To test them, I generate 1000 values from $\text{Uniform}[0, \theta]$ with $\theta = 2$, and find a 95% confidence interval for $\hat{\theta}$. I repeat this 20 times. Not once does my confidence interval include the true value, $\theta = 2$, for either resampler. Explain.

Naive resampling (based on mle parameter estimates or on empirical distributions) is an heuristic, not a perfect procedure. It works well for 'central' statistics like averages or sums. It doesn't work well for certain types of extreme statistics (like the maximum of a dataset) nor for certain types of distribution (like the uniform).

The idea of resampling is that we want to simulate novel unseen versions of the dataset. The best way to do this is to use a model that we think is a good description for novel unseen data—in other words, to use a model that fits a holdout dataset well. (See section 9 of lecture notes for a longer discussion of generalization. That section of notes is non-examinable.) One ad hoc way to get better generalization in this case is to use an unbiased estimator for θ rather than a maximum likelihood estimator; though this is happenstance, not a general principle!

Question 15. Test the hypothesis that temperatures in Cambridge have not been changing, using a non-parametric test.

In lectures we looked at several examples of tests using parametric resampling. We also looked at one example of a test with non-parametric resampling, namely Fisher's permutation test. Example 8.6.2 in lecture notes gives another illustration of non-parametric sampling for hypothesis tests.

For this dataset, it's blindingly obvious that there is an annual cycle in temperatures, so your resampling strategy must respect this. If there were no global warming, and you wanted to simulate a January, how could you simulate it using the data in this dataset?

Second, the test statistic. You are at liberty to use any test statistic at all; it doesn't have to be linked to the resampling strategy. You might as well use $\hat{\gamma}$ from question 6.

Question 16. We have a dataset x_1, x_2, \dots, x_n , and we wish to model it as $\text{Normal}(\mu, \sigma^2)$ where μ and σ are unknown. How different are Bayesianist and frequentist confidence intervals for the mean? To be concrete, let's work with the first 10 values for `temp` in the climate dataset.

- Plot the log likelihood function $\log \Pr(x_1, \dots, x_n | \mu, \sigma)$ as a function of μ and σ . (A code skeleton is provided in <https://github.com/damonjw/datasci/blob/master/ex/ex3.ipynb>.)
- Using frequentist resampling, generate 50 resampled datasets, find the maximum likelihood estimators $\hat{\mu}$ and $\hat{\sigma}$ for each, and show these 50 points on your plot.
- Using computational Bayesian methods, with priors $\mu \sim \text{Normal}(0, 10^2)$ and $\sigma \sim \Gamma(k = 2, \theta = 1)$ (where k and θ are as in the numpy documentation), sample 500 pairs from the prior distribution and show them on your plot. Then compute the posterior weights of these sampled pairs, and show the weighted pairs on your plot by setting the size of the plot marker in proportion to weight.
- Find the 95% confidence interval (for $\hat{\mu}$ in the frequentist case, and for $(\mu | \text{data})$ in the Bayesianist case), and show them on your plot.
- Repeat the exercise, using the first 100 values from the climate dataset.

You should see broadly similar outcomes, whether you're plotting frequentist samples of $(\hat{\mu}, \hat{\sigma})$ or whether you're plotting the Bayesianist samples that get non-negligible weight. When there are more datapoints, then the results are even more similar: there's a very narrow peak in the log likelihood plot, and the samples from both Bayesianist and frequentist approaches are heavily concentrated around this peak. (Though the naive computational Bayesian procedure we learnt in this course doesn't work very well when the log likelihood has such a sharp spike.)

Question 17. In hypothesis testing, what p -value would you expect if H_0 is true?

This is a mindbender! At first glance it's surprising that this question even has an answer that applies to any sort of hypothesis testing. And it's tricky to even work out what it's asking us to prove. Think of it this way ...

In frequentist inference, we decide on a sampling distribution X^* that tells us what the dataset might have been if H_0 were true. We then compute the p -value by an operation on $t(x)$ and on the histogram of $t(X^*)$.

Now, if H_0 were true, then the actual dataset x will look like a sample from X^* . If we perform the p -value operation not on the actual value $t(x)$ but on a typical value $t(X^*)$, what's the distribution we'll get for the p -value?

You can find the answer at https://en.wikipedia.org/wiki/Fisher's_method. The page also describes how the answer can be used to combine the results of several independent tests.

Question 18. We are given a dataset x_1, \dots, x_n . Our null hypothesis is that these values are drawn from $\text{Normal}(0, \sigma^2)$, where σ is an unknown parameter. Let

$$\hat{F}(x) = \frac{1}{n} \sum_{i=1}^n 1[x_i / \hat{\sigma} \leq x]$$

where $\hat{\sigma} = \sqrt{n^{-1} \sum_i x_i^2}$ is the maximum likelihood estimator for σ . If the null hypothesis is true, we'd expect $\hat{F}(x)$ to be reasonably close to $\Phi(x)$, the cumulative distribution function for $\text{Normal}(0, 1)$, for all x . Suggest how to test the hypothesis that the dataset is indeed drawn from $\text{Normal}(0, \sigma^2)$, using a test statistic based on \hat{F} and Φ .

This question is asking you to be creative in inventing a test statistic. If you don't feel creative, look up the Kolmogorov-Smirnov test on Wikipedia.

When we fit a linear model, there's an assumption that the residuals are normally distributed (as discussed in section 2.4). After fitting a linear model, it's always worth testing whether the residuals are indeed normally distributed, and this question gives you a way to do this test.

Question 19. A recent paper *Historical language records reveal a surge of cognitive distortions in recent decades* by Bollen et al., <https://www.pnas.org/content/118/30/e2102061118.full>, claims that depression-linked turns of phrase have become more prevalent in recent decades. This paper reports both confidence intervals and null hypotheses. Explain how it computes them, in particular (1) the readout statistic, (2) the sampling method.

Skim-read the whole paper, and read the Materials and Methods section closely. Note that the word 'bootstrapping' is another name for 'non-parametric resampling'. You can find a definition of z-score on Wikipedia, but it doesn't add anything to the explanation given in the paper.

In the notation used in this course, the dataset used in the paper is $(x_1, y_1), \dots, (x_k, y_k)$ where y_k is a vector

$$y_i = [y_{i,1855}, \dots, y_{i,2020}]$$

giving the prevalence of n -gram i in each year, and $x_i \in \{1, 2, 3, 4, 5\}$ is the number of words in that n -gram.

The readout statistic $t(x_1, \dots, x_k)$ is well hidden, and you will have to dig through the whole paper to find it.

Question 20. To allow for non-linear temperature increase, Example Sheet 1 suggested a model with a step function,

$$\text{temp} \sim \beta_1 \sin(2\pi t) + \beta_2 \cos(2\pi t) + \gamma_{\text{decade}} + \text{Normal}(0, \sigma^2).$$

Find a 95% confidence interval for $\hat{\gamma}_{2010s} - \hat{\gamma}_{1980s}$. Conduct a hypothesis test of whether $\gamma_{1980s} = \gamma_{2010s}$.