



Newcastle University

MAS2905

Probability and Bayesian Methods with R

Part 1

Dr. Lee Fawcett

Semester 1, 2020/21

1 General Housekeeping

1.1 Lecturer information

The module leader for MAS2905, and lecturer for both semesters, is Dr. Lee Fawcett. If you have any questions about the module he can be contacted via email at lee.fawcett@newcastle.ac.uk. Lee is based in room 2.07 on the second floor of the Herschel Building. There will be several other people assisting with this module in semester 1, primarily in the computer practical sessions. They're all really friendly PhD students – Amy Green, Kieran Peel and Nicola Hewett – give them a shout in class if you're stuck with anything!



1.2 Classes

As you probably know, things will be very different in 2020/21. You will have a blended learning experience¹, meaning you will have some *Present-in-Person* activities on campus and some online activities. Your online activities might be *synchronous* or *non-synchronous*.

For semester 1 of MAS2905, this is how we will use these activities:

- **Present-in-Person (PiP):** A series of 90 minute supervised computer practical sessions in the Herschel PC cluster. These take place roughly every three weeks, will appear on your timetable and will be advertised on Canvas (and are detailed in Section 1.3). These sessions might be replaced with “nPiP” sessions – online versions of in-person computer sessions
- **Non-synchronous online (NSO):** Online activities posted to Canvas every week, covering what would have ordinarily been covered in standard face-to-face lectures. This will mainly consist of directed reading (e.g. lecture notes), completing exercises/practice questions, summaries/discussions of learning objectives, assessment feedback and other ‘fun’ activities (e.g. quizzes!). These activities must be completed in your own time, whenever you can fit them in, but it is *vital* that they are completed by their completion dates as advertised on Canvas (and as shown in Section 1.3)
- **Synchronous online (SO):** Timetabled sessions, broadcast live online via Zoom. No new material will be presented – instead, we will work through exercises based on the theory you have studied in the NSO activities. Some SO sessions will be dedicated ‘office hours’ held via Zoom, for those who want to drop-in (e.g. if you’re stuck on an assignment question). You should make every effort to attend these SO sessions when they appear on your timetable; the dates/times of these sessions are also advertised on Canvas (and shown in Section 1.3)

¹Your modules might move entirely online at short notice, as has happened for the first two weeks of term

Students can take part in other NSO activities, should they wish, including interacting with the Canvas discussion board; optional Zoom drop-in sessions; and one-to-one Zoom meetings. And of course, you will all be expected to prepare and complete homeworks in your own time and revise for tests!

Note that this arrangement is specific to semester 1 of this module. Our PiP sessions in semester 2 will be classroom-based. The specific arrangements for semester 2 will be discussed in the booklet for Part 2 of this module, distributed in January 2021.

1.3 Schedule

Week	w/b	Activities	Date	Time	Deadline?
1	19 Oct	SO Lecture NSO activities	Wed 21 Oct	12:00-13:00	— Completed by Mon 26 Oct
			After lecture, in your own time: Videos/exercises/reading		
2	26 Oct	nPiP Practical NSO activities	Tues 27 Oct	—90 minute intervals, check your timetable— Completed by Mon 2 Nov	
			After practical, in your own time: Videos/exercises/reading		
3	2 Nov	SO Lecture NSO activities	Wed 4 Nov	12:00-13:00	— Completed by Mon 9 Nov
		Test 1 released	Thu 5 Nov	16:00	16:00 Thu 12 Nov
4	9 Nov	SO Lecture NSO activities	Wed 11 Nov	12:00-13:00	—
		Test 1 due	Thu 12 Nov	16:00	16:00 Thu 12 Nov
—	16 Nov	—Buffer week, no new work scheduled—			
5	23 Nov	PiP Practical NSO activities	Tues 24 Nov	—90 minute intervals, check your timetable— Completed by Mon 30 Nov	
		Project released	Thurs 26 Nov	16:00	16:00 Thu 17 Dec
6	30 Nov	SO lecture NSO activities	2 Dec	12:00-13:00	— Completed by 7 Dec
			Some pre-lecture reading/videos; After lecture, in your own time: Exercises/project work		
7	7 Dec	SO lecture NSO activities	9 Dec	12:00-13:00	—
			Project drop-ins, times TBA Project work		
8	14 Dec	PiP Practical NSO activities Project due	Tues 15 Dec	—90 minute intervals, check your timetable— 16:00 Thu 17 Dec	
			Project work, practice test questions		
—		—Christmas holidays—			
9	11 Jan	SO Lecture NSO activities	Wed 13 Jan	12:00-13:00	—
		Test 2 released	Thu 14 Jan	10:00	16:00 Thu 14 Jan
		Test 2 due	Thu 14 Jan	16:00	16:00 Thu 14 Jan

1.4 Assessment

The full 20 credit MAS2905 module will be assessed via a formal written examination in summer 2021 (likely a '24 hour take-home paper') and some in-course assessment. The make-up is 50% summer exam, 50% in-course assessment. Further details:

- Summer exam, 100 marks in total
 - Semester 1 material: Section A (20 marks)
 - Semester 2 material: Section B (80 marks)
 - Answer *all* questions in *both* parts
- In-course assessment, 100 marks in total
 - Semester 1 material: Test 1 (20 marks), mini project (30 marks), test 2 (30 marks)
 - Semester 2 material: Two short homeworks (10 marks each)

Exam questions in section A, relating to the material from semester 1 of the module, will focus mainly on the theory and probability/statistics developed throughout the semester (rather than programming in R).

The semester 1 project will be R-based, but will require some pencil-and-paper calculations too. Test 1 (completed in Numbas) will be open for a 1-week period and will test your basic R skills, including dataframe querying, for loops, and the use of simple built-in functions to test your knowledge of probability calculations in R. Test 2 will cover later material from the course, including transformations of random variables and bivariate distributions. This test will require a mix of pencil-and-paper calculations and R work, and must be completed within a six hour window (the questions themselves are designed to take no longer than 90 minutes to complete).

Test-style questions, with solutions, will be made available in advance of both tests, to help you prepare.

More details about the semester 2 assessment will be discussed in the booklet for Part 2 of this module, distributed in January 2021.

1.5 Late work policy

In this module, deadline extensions can be requested for the mini project submitted in week 8 (by means of submitting a PEC form), and work submitted within 7 days of the deadline without good reason will be marked for reduced credit.

This module also contains a test worth more than 10% for which rescheduling can be requested (by means of submitting a PEC form).

For details of the policy (including procedures in the event of illness etc.) please consult the Mathematics, Statistics & Physics Community pages on Canvas, under: Assessment Information, Late Work and Missed Assessments.

1.6 Course materials

All course materials will be posted to the dedicated Canvas page for MAS2905. Here, you will find:

- A full copy of these lecture notes, with gaps
- A copy of the lecture notes with the gaps filled in, after the associated PiP/SO/NSO activities
- All slides used to support PiP/SO/NSO activities after these have taken place
- All video recordings
- Solutions to assessed exercises
- Any other materials, including additional handouts/worksheets, links to useful Numbas material, specimen exam papers and solutions, . . .

1.7 Books and other resources

At various points in the module you will be directed to additional reading, should you want extra examples or just another source of information relating to the material presented in the notes. This will primarily be text books (many of which are available to view online), but perhaps other sources (e.g. journal publications, magazine articles etc.).

Some useful books for semester 1 include:

- W.N. Venables and D.M. Smith: An Introduction To R. Free to download at:

<https://cran.r-project.org/doc/manuals/R-intro.pdf>

- D. Braun and D.J. Murdoch: A First Course In Statistical Programming With R
- S. Ross: A First Course in Probability
- N.A. Weiss, P.T. Holmes and H. Michael: A Course In Probability

More details about semester 2 reading material will be provided in the booklet for Part 2 of this module, distributed in January 2021.

1.8 General background for Part 1 of MAS2905

In Part 1 of the module we will learn about computational tools for performing probability calculations and statistical analysis. We will use the R programming language. R is very widely used in universities, businesses and other organizations due to its many advantages over other similar software packages. It is available free of charge, it is 'open source' and can be used for a very wide range of statistical and scientific applications. Libraries of R functions and data-structures have been created and are readily available on the web. These enable statisticians to perform a very wide range of tasks but with a minimum of additional programming. Learning to use R is important because:

- (i) It is the foremost tool in modern computational statistics
- (ii) It is a great way to learn general concepts in programming which can be used with other languages or in other contexts
- (iii) It can be used to illustrate ideas in probability and statistics

In MAS1802 last year we focused mainly on point (ii) above; in this module, we focus mainly on points (i) and (iii).

Given the COVID-19 situation in 2020 (and the industrial action that took place immediately before the closure of our campus), many of you will be worried about not having had the chance to learn R properly last year, in MAS1802. Thus, one of the first things we will do in MAS2905 is review some of the key concepts in the MAS1802 syllabus. This will roughly occupy the first two weeks of the module.

For the rest of semester 1, we will present a mixture of background theory in probability and statistics, together with elements of the R programming language. Thus, from around the third week of the course, activities will contain both some statistical theory and some R programming concepts. Most of the theory will be covered in your own time, during the non-synchronous online (NSO) activities that you are set in Canvas. In our synchronous online (SO) and Present-in-Person (PiP) sessions, we will put that theory into practice through worked examples, both written and on the computer in R.

All materials for semester 1 of the module are in this booklet, including:

- Background theory that you will be expected to read yourself through NSO activities
- Pencil-and-paper exercises – some of which you will try out yourself in NSO activities, and some of which we will work through together in SO sessions
- Computer practical exercises – some of which you will try out yourself in NSO activities, and some of which we will work through together in our PiP sessions in the computer cluster

It is vital that, if possible, you attend your practical sessions (either PiP or nPiP). This is a great opportunity to ask us questions face-to-face, seek help from other students, and of course prepare

for the two tests. It is even more important that you follow instructions on Canvas and complete the other activities by the dates suggested – if you are told that some NSO activities should be completed before the next PiP computer practical, then make sure you do exactly that – the work set in the computer practical class will assume you have completed all work to-date!

2 R Refresher

2.1 Preliminaries

R is installed on all University machines. In this module we will be using RStudio, which is an R *Integrated Development Environment*, or IDE. This just makes R much easier to use and navigate, through the four windows that you see when you open RStudio (text editor, console window, file explorer and graphics display).

You might remember from MAS1802 that both R and RStudio are free, so you can install them on your own computer. Versions of R and RStudio are available for Windows, Apple Mac and Linux. See

<http://www.ncl.ac.uk/maths-physics/students/resources/>

for more details.

We will assume a limited working knowledge of R and RStudio. The first two weeks are designed to recap the basics from MAS1802. Specifically, we will review (i) data types and vectors in R; (ii) dataframe manipulation; (iii) data summaries; (iv) functions and control statements. Everything we do for the remainder of semester 1 will then assume you have a good grasp of these basics, so please work hard to make sure you're up-to-speed with everything as we progress through this Section of the notes! And of course, get in touch if you're not sure about anything – either by email or via the Canvas discussion board.

You might recall the `movies` dataset from MAS1802. This was obtained from the internet movie database (IMDB): <http://imdb.com/>, which is a website devoted to collecting movie data supplied by studios and fans. It claims to be the biggest movie database on the web and is run by *Amazon*. More information can be found online:

http://imdb.com/help/show_leaf?about

This includes information about the data collection process. This is a great dataset to use to demonstrate R's dataframe manipulation and querying capabilities – it is large, has many variables, and is interesting. Everyone loves movies! See Table 1 for an extract of the `movies` dataset.

Title	Year	Length	Budget	Voting statistics					Movie genre							
				Rating	Votes	r1	...	r10	mpaa	Action	Animation	Comedy	Drama	Documentary	Romance	Short
A.k.a. Cassius	1970	85	NA	5.7	43	4.5	...	14.5	PG	0	0	0	0	1	0	0
AKA	2002	123	NA	6.0	335	24.5	...	1.5	R	0	0	0	1	0	0	0
Alien Vs. Pred	2004	102	45000000	5.4	14651	4.5	...	4.5	PG-13	1	0	0	0	0	0	0
Abandon	2002	99	25000000	4.7	2364	4.5	...	4.5	PG-13	0	0	0	1	0	0	0
Abendland	1999	146	NA	5.0	46	14.5	...	24.5	R	0	0	0	0	0	0	0
Aberration	1997	93	NA	4.8	149	14.5	...	4.5	R	0	0	0	0	0	0	0
Abilene	1999	104	NA	4.9	42	0.0	...	24.5	PG	0	0	0	1	0	0	0
Ablaze	2001	97	NA	3.6	98	24.5	...	14.5	R	1	0	0	1	0	0	0
Abominable Dr	1971	94	NA	6.7	1547	4.5	...	14.5	PG-13	0	0	0	0	0	0	0
About Adam	2000	105	NA	6.4	1303	4.5	...	4.5	R	0	0	1	0	0	1	0

Table 1: A subset of ten rows of the movie data set. **Credit:** This data set was initially constructed by Hadley Wickham at <http://had.co.nz/>.

We have written an add-on R package for the module which will allow easy access to many interesting datasets (including the movies dataset). Installing the package is straightforward, and will make loading datasets used in the notes really easy:

```
1 install.packages("ncldata", repos = "http://R-Forge.R-project.org",
2 type = "source")
```

To load the package, type:

```
1 library(ncldata)
```

Then, to specifically load the movies dataset, type:

```
1 data(movies)
```

Alternatively, some datasets can be downloaded directly from the course Canvas page, and some from web sources; examples of how to do this will be given later in the notes.

2.2 Data types in R

R has a variety of data types. For example:

```
1 v = TRUE           # 'Logical' or 'boolean' - TRUE or FALSE
2 w = "fred"         # 'Character' or 'string'
3 x = 5.0             # 'Double' or 'numerical'
```

There are also some 'special' data types:

```
1 y = 5/0            # Returns Inf
2 z = y-y            # Returns NaN
```

Another important data type is NA, short for 'Not Available' and used for missing values. A summary of the data types in R is given in Table 2 below. Note the way R presents doubles in standard form (Example 4); notice also that anything declared inside inverted commas will be assumed by R to be of Character type.

Type	Example 1	Example 2	Example 3	Example 4
Doubles	2	3.1242	-45.6	4e-10
Logicals/Boolean	TRUE	FALSE		
Characters	"Fred"	"TRUE"	"Male"	"2"
Infinity	Inf	5/0		
Other	NaN	NA		

Table 2: A summary of some data types in R.

2.3 Vectors

Vectors are the most basic of all data structures, but are used in almost all R code. An R vector contains n values of the same type, where n can be zero. For example:

```

1 c(0, 1, 2, 3, 4, 5)
2 [1] 0 1 2 3 4 5
3 (first_vec = c(0, 1, 2, 3, 4, 5))
4 [1] 0 1 2 3 4 5
5 second_vec = c("Male", "Female", "Male")
6 [1] "Male" "Female" "Male"
7 x = 5.2
8 y = "Fred"

```

Some notes on the above code listing:

- In lines 1-2 we see a vector of length $n = 6$, containing the values $\{0, 1, \dots, 5\}$. Note that this vector has not been assigned for future use
- In lines 3-4, we see the same vector as in lines 1-2, but this vector has now been assigned to the object “first_vec”
- The outer round brackets used in line 3 are not necessary - they are used in these notes and slides to ensure the resulting vector is printed to the notes/screen, for your convenience so you can see what the code has produced
- “first_vec” is a vector of doubles
- “second_vec” is a vector of $n = 3$ characters
- Both x and y are vectors too – each of length $n = 1$ (x is a vector containing $n = 1$ double; y is a vector containing $n = 1$ character)

There are special functions in R that can be used to determine the variable type. For example:

```

1 is.double(x)
2 [1] TRUE
3 is.character(x)
4 [1] FALSE
5 is.character(y)
6 [1] TRUE
7 is.vector(first_vec)
8 [1] TRUE

```

And of course, we can use the length command to determine the length of a vector:

```

1 length(first_vec)
2 [1] 6
3 length(second_vec)
4 [1] 3

```

And you can use the `seq` command to generate sequence vectors very quickly:

```
1 (x1 = seq(1, 6))
2 [1] 1 2 3 4 5 6
3 (x2 = seq(-4, 4, by=2))
4 [1] -4 -2 0 2 4
5 (x3 = seq(3, -2, by=-0.5))
6 [1] 3.0 2.5 2.0 1.5 1.0 0.5 0.0 -0.5 -1.0 -1.5 -2.0
```

The following code listing demonstrates some other built-in R commands that can be very useful when applied to a vector – `rev`, `sort`, `sum`, `unique` and `which`. Make sure you understand what each of these is doing – especially `unique` and `which`!

```
1 rev(x1) # Reverses the order of the vector
2 [1] 6 5 4 3 2 1
3 sort(x3) # Sorts the vector into ascending order
4 [1] -2.0 -1.5 -1.0 -0.5 0.0 0.5 1.0 1.5 2.0 2.5 3.0
5 sort(second_vec) # Can be applied to characters for alphabetical order
6 [1] "Female" "Male" "Male"
7 sum(x2); sum(x3); sum(second_vec)
8 [1] 0
9 [1] 5.5
10 Error in sum(second_vec) : invalid 'type' (character) of argument
11 unique(second_vec) # Lists all unique elements of the vector
12 [1] "Male" "Female"
13 which(x2==2) # Gives index/indices of particular elements
14 [1] 4
15 which(x3<=1)
16 [1] 5 6 7 8 9 10 11
```

2.3.1 Vector operations

When our data are in a vector structure, we can apply standard operations to the entire vector. For example:

```
1 (x = seq(-4, 4))
2 [1] -4 -3 -2 -1 0 1 2 3 4
3 x*x
4 [1] 16 9 4 2 1 0 1 2 4 9 16
5 x-5
6 [1] -9 -8 -7 -6 -5 -4 -3 -2 -1
7 x + x
8 [1] -8 -6 -4 -2 0 2 4 6 8
```

2.3.2 Extracting elements from vectors

R has a number of useful methods that we can use to extract subsets of our data. For example, we could use square brackets to extract particular elements:

```

1 first_vec[2]
2 [1] 1
3 second_vec[2:3]
4 [1] "Female" "Male"
5 first_vec[4:2]
6 [1] 3 2 1
7 (z = 1:10)
8 [1] 1 2 3 4 5 6 7 8 9 10
9 z[z<3 | z>=7] # We only want values of z that are less than 3 or at least 7
10 [1] 1 2 7 8 9 10

```

We will come back to extraction from two-dimensional R objects (e.g. dataframes) shortly.

Another useful way of extracting elements from vectors is to use logical vectors. Recall that the logical data type in R can take only the values TRUE and FALSE. Hence:

```

1 A = TRUE
2 B = FALSE
3 !A # Read as "Not A"
4 [1] FALSE
5 !B
6 [1] TRUE
7 A & B # Need "both A and B"
8 [1] FALSE
9 A | B # Need "either A or B"
10 [1] TRUE

```

Then we could declare a vector of such logicals to help us extract from another vector:

```

1 (logic1 = c(TRUE, FALSE, TRUE, FALSE))
2 [1] TRUE FALSE TRUE FALSE
3 (vec1 = seq(3, 6))
4 [1] 3 4 5 6
5 vec1[logic1] # Will only extract the elements in vec1 where logic1 is true
6 [1] 3 5

```

We can also sum a logical vector. For example:

```

1 sum(logic1)
2 [1] 2
3 sum(vec1[logic1])
4 [1] 8

```

Example 2.1: If you've not already got RStudio installed on the machine you're using, **you must do so now** by following the advice on our student resources webpage:

<http://www.ncl.ac.uk/maths-physics/students/resources/>

From this link, click on the link for our "Computing page", log-in using your University details, and then click on the link for "R". Here you will find some basic resources aimed at students who have never used R before; scroll down, and you'll find information about how to install R on your home PC.

In RStudio, commit the following lines of code. Both lines read in files stored on a webpage. **Do not copy-and-paste this code – you must type it in carefully!**

```
1 x = scan("http://www.mas.ncl.ac.uk/~nlf8/teaching/mas2905/x.txt")
2 y = as.logical(scan("http://www.mas.ncl.ac.uk/~nlf8/teaching/mas2905/y.txt",
  what="logical"))
```

Questions about x

- What type of data is in the vector x?
- How many values are in the vector x?
- What is the smallest value in the vector x?
- What is the value of x in position 3500?
- What is the 97th smallest value in x?
- Find $\sum_{i=5612}^{8664} x_i$.

Questions about y

- What type of data is in the vector y?
- What is the first value of y?
- What is the last value of y?
- How many times does TRUE occur?
- When is the 472nd occurrence of TRUE?
- How many occurrences of TRUE are there between positions 504 and 593 (inclusive)?

Your answers about x

-
-
-
-
-
-

Your answers about y

-
-
-
-
-
-

2.4 Dataframes

A dataframe is a special kind of object in R. We use dataframes for storing and managing datasets that have a rectangular structure. Typically the rows correspond to *cases* and the columns *variables*. The crucial difference between a dataframe and a matrix is that all values in a matrix must be of the same type, whereas in a dataframe we can have responses for cases on variables of different types. For example, we might have survey results on 10 individuals, who have each been asked 5 questions. This would give a dataframe with 10 rows and 5 columns (a so-called “10 by 5 dataframe”; generally, we refer to $r \times c$ dataframes). The variable types need not all be the same – the first column could be a vector of doubles recording the individuals’ height in metres; the second column could be a vector of logicals taking the value TRUE if they like Marmite; the third column could be a vector of characters taking the name of their favourite book, etc.

2.4.1 Creating and querying your own dataframe

You can construct your own dataframe using R’s `data.frame` command. For example:

```
1 height_m = c(1.81, 1.75, 1.58)
2 weight_kg = c(NA, NA, 70.3)
3 marmite = c(TRUE, FALSE, FALSE)
4 book = c("Midnight in Chernobyl", "To Kill a Mockingbird", "Life of Pi")
5
6 (small.survey = data.frame(height_m, weight_kg, marmite, book))
7   height_m weight_kg marmite      book
8 1      1.81      NA    TRUE Midnight in Chernobyl
9 2      1.75      NA   FALSE To Kill a Mockingbird
10 3      1.58     70.3   FALSE      Life of Pi
```

So, case 2 in our survey is 1.75 metres tall, does *not* like Marmite and their favourite book is “To Kill a Mockingbird”. We do not know the weights of cases 1 and 2 (missing values). Here we have 3 cases and 4 variables, so we have a 3×4 dataframe. If you had a large dataframe and wanted to quickly know its size, you could use the `dim` command (short for “dimension”). For example:

```
1 dim(small.survey)
2 [1] 3 4
```

We can extract elements from a dataframe using square brackets – just as we demonstrated in Section 2.3.2 for extraction from a vector. However, if the columns in our dataframe have labels, we can also use the `$` command to access columns directly. For example:

```
1 small.survey[2, 4]      # Will extract the element from the 2nd row, 4th column
2 [1] To Kill a Mockingbird
3
4 small.survey[, 4]      # Will extract the elements from all rows in the 4th column
5 [1] Midnight in Chernobyl To Kill a Mockingbird Life of Pi
6
7 small.survey$book      # Does exactly the same as line 4!
8 [1] Midnight in Chernobyl To Kill a Mockingbird Life of Pi
```

```

9 small.survey$book[2]      # Does exactly the same as line 1!
10 [1] To Kill a Mockingbird
11
12 small.survey[3, ]        # Will extract the elements from all columns in the 3rd row
13   height_m weight_kg marmite      book
14 3      1.58      70.3   FALSE Life of Pi

```

Though used less, the `subset` command can be applied to dataframes to extract information from selected columns only. For example:

```

1 (small.survey2 = subset(small.survey, select = 3:4))
2   marmite      book
3 1   TRUE Midnight in Chernobyl
4 2  FALSE To Kill a Mockingbird
5 3  FALSE      Life of Pi

```

And what if you wanted the height and favourite book of respondents who didn't like Marmite?

```

1 subset(small.survey, marmite == FALSE, select=c(1,4))
2   height_m      book
3 2      1.75 To Kill a Mockingbird
4 3      1.58      Life of Pi

```

2.4.2 The movies dataframe

The Movies dataset is available from within the `ncldata` package. Recall that you can access this in the following way:

```

1 install.packages("ncldata", repos="http://R-Forge.R-Project.org",
2 type="source")
3 library(ncldata)
4 data(movies)

```

- Line 1 installs the `ncldata` package to your machine. You will only need to do this once, and in PiP sessions in the computer cluster this has been done for you already. You will make use of this package for other datasets in Statistics modules, including semester 2 of this module
- Line 3 reads the package into your current working RStudio space – each time you need to access data from within the package you should use this line of code
- Line 4 then pulls the `movies` dataset from the `ncldata` package, ready for you to use

You can now explore the movies dataset:

```

1 dim(movies)      # Looks like there are 4988 movies and 23 variables
2 [1] 4988      23
3
4 head(movies)     # Dataframe too large to view; let's look at the first six rows
5
6      Title Year Rating  Votes Length  Budget BudgetCurrency
7 1      10,000 BC 2008   5.1 115607   109 1.05e+08          USD
8 2      101 Dalmatians (I) 1961   7.2 126498    79      NA          <NA>
9 3      101 Dalmatians (II) 1996   5.7  84527   103 7.50e+07          USD
10 4      102 Dalmatians 2000   4.9  29328   100 8.50e+07          USD
11 5      10 Cloverfield Lane 2016   7.2 226726   104 1.50e+07          USD
12 6 10 Things I Hate About You 1999   7.2 249949    97 3.00e+07          USD
13
14 mpaa Action Adventure Animation Biography Comedy Crime Documentary Drama
15 1 PG-13 TRUE TRUE FALSE FALSE FALSE FALSE FALSE FALSE TRUE
16 2 <NA> FALSE TRUE TRUE FALSE TRUE FALSE FALSE FALSE FALSE
17 3 <NA> FALSE TRUE FALSE FALSE TRUE TRUE FALSE FALSE FALSE
18 4 <NA> FALSE TRUE FALSE FALSE TRUE FALSE FALSE FALSE FALSE
19 5 PG-13 FALSE FALSE FALSE FALSE FALSE FALSE FALSE FALSE TRUE
20 6 PG-13 FALSE FALSE FALSE FALSE FALSE TRUE FALSE FALSE TRUE
21
22 Fantasy Horror Mystery Romance SciFi Short Thriller
23 1 FALSE FALSE FALSE FALSE FALSE FALSE FALSE
24 2 FALSE FALSE FALSE FALSE FALSE FALSE FALSE
25 3 FALSE FALSE FALSE FALSE FALSE FALSE FALSE
26 4 FALSE FALSE FALSE FALSE FALSE FALSE FALSE
27 5 FALSE TRUE TRUE FALSE FALSE FALSE FALSE
28 6 FALSE FALSE FALSE TRUE FALSE FALSE FALSE

```

From the above R code, we can see that the variables in the `movies` dataframe are:

- **Title** of the movie
- The **Year** the movie was released
- The movie's average **Rating** (out of 10), based on the number of **Votes** submitted
- The **Length** of the movie (in minutes)
- The movie **Budget** (this is in USD, as shown by **BudgetCurrency**, unless otherwise stated)
- The movie's rating as suggested by the Motion Picture Association of America (**mpaa**)
- Various logical variables indicating the movie's genre (e.g. **Animation**, **Horror** etc.)

For example, the movie "10 Cloverfield Lane" was released in 2016; it has an average rating of 7.2/10 from 226, 726 votes submitted to the Internet Movie Database (IMDB); the movie length is one hour 44 minutes; the budget was around US\$ 15M; it was rated PG-13 by the Motion Picture Association of America; and it has been classified as a drama/horror/mystery movie.

Examples of querying the movies dataframe

1. Note that there are some missing values – for example, the budget for the original version of 101 Dalmatians, released in 1961, is unknown. What if we wanted to know how many movies had known budgets?
2. What if we wanted to create a new dataset with only those movies made since the year 1998?
3. Or (harder) movies that cost more than US\$ 100,000 but are not R rated?

For the first query:

```

1 x1 = is.na(movies$Budget)      # Will return TRUE if the budget is NA
2 head(x1)                      # Compare to dataset above – looks good
3 [1] FALSE TRUE FALSE FALSE FALSE FALSE
4
5 x2 = !is.na(movies$Budget)    # Flip things round: NA = FALSE
6 head(x2)
7 [1] TRUE FALSE TRUE TRUE TRUE TRUE
8
9 budget_known = movies$Budget[x2] # Apply logical vector to the movie budgets
10 head(budget_known)
11 [1] 1.05e+08 7.50e+07 8.50e+07 1.50e+07 3.00e+07 6.00e+06
12
13 length(budget_known)          # This will tell us how many movies have known budgets
14 [1] 3990
15
16 length(movies$Budget[!is.na(movies$Budget)]) # A bit more concise!
17 [1] 3990

```

Second query:

```

1 movies2 = subset(movies, Year >= 1998)
2
3 dim(movies2)
4 [1] 3478 23

```

And the last query:

```

1 movies3 = movies[movies$Budget > 100000 & movies$mpaa != "R", ]
2 dim(movies3)
3 [1] 3174 23

```

Example 2.2: Follow the steps at the start of Section 2.4.2 to load the `ncldata` package and to read in the `movies` dataframe. Then answer the following questions.

- (a) How many people submitted votes for the films “*10 Things I Hate About You*”, “*Midnight in the Garden of Good and Evil*” and “*Zoolander*”? **Answers:**_____
- (b) How many films are classified as “Action”? How many are classified as “Romance”? How many “Short” films are there in the database? **Answers:**_____
- (c) How many films were made between the Years 1970 and 2000 (inclusive)? **Answer:**_____
- (d) How many movies are longer than two hours in duration? **Answer:**_____
- (e) How many movies made since the year 2005 (not including 2005) are longer than three hours in duration? **Answer:**_____

Example 2.3: Load and explore the `music` dataset from the `ncldata` package by typing the following lines of code (Note: This will only work if you have correctly installed the `ncldata` package and read the package into your working RStudio space using the `library` command; see example 2.2 and Section 2.4.2):

```
1 data (music)
2 # What's this dataset all about? Let's find out
3 ?music
```

- (a) How many tracks do we have data for, and on how many variables? **Answers:**_____
- (b) How many tracks do we have from *The Beatles* in this dataframe? **Answer:**_____

What about *Lady GaGa*? **Answer:**_____
- (c) How many tracks have a tempo of at least 130 beats per minute? **Answer:**_____
- (d) Of the tracks with a tempo of at least 130 beats per minute, how many have been released since the year 2000 (including the year 2000)? **Answer:**_____
- (e) What is the most popular key for tracks in this database? **Answer:**_____

2.5 Graphical and numerical summaries

R has very powerful built-in functions for summarising datasets both graphically and numerically. Most of these are fairly intuitive, so in this section we provide only a brief review of material from stage 1.

For example, we might want to explore the distribution of budgets in the `movies` dataframe. As we know, there are some missing values in the dataset, but most of the built-in functions for data analysis in R have an argument we can invoke to ignore these:

```

1 # To find the mean of the movie budgets but ignoring the missing values
2 mean(movies$Budget, na.rm = TRUE)
3 [1] 41245023
4
5 # And the standard deviation
6 sd(movies$Budget, na.rm = TRUE)
7 [1] 68612118
8
9 # How about a range of numerical summaries with a single function?
10 summary(movies$Budget, na.rm = TRUE)
11      Min.    1st Qu.    Median      Mean   3rd Qu.      Max.      NA's
12 0.000e+00 9.000e+06 2.500e+07 4.125e+07 5.000e+07 1.900e+09      998

```

From the listing above we can make a few observations:

- The mean budget is larger than the median: ~US\$ 41M, compared to US\$ 25M. This indicates a positively skewed dataset, as the mean is being 'pulled up' by some very large budgets (of around USD\$1.9Bn?! Which movie is this?! Must be amazing...)
- The standard deviation seems extremely large; again, this could be due to a small number of extremely large budgets...
- ... or perhaps some unrealistically *small* budgets (of US\$ 0?!)

```

1 # Let's find out the position of the most expensive movie
2 which(movies$Budget == max(movies$Budget, na.rm=TRUE))
3 [1] 1211
4
5 # What's the name of this movie?
6 movies$Title[1211]
7 [1] "Enthiran"
8 # Never heard of it!
9
10 # Which movies were made for free?!
11 movies$Title[which(movies$Budget==0)]
12 [1] "Homefront"      "Just Friends"    "Knock Knock"     "Now You See Me"
13 [5] "Safe House"     "Snitch"

```

What about some graphical summaries? By far the most commonly-used graphics are histograms and boxplots. Of course, there are so many more graphical options available in R; you should consult Chapter 3 of your MAS1802 notes for a more comprehensive coverage.

```

1 # Partitions the plotting space: 1 row, 3 columns
2 par(mfrow = c(1,3))
3
4 # Now we'll try three different methods for producing histograms, based on class width/number of bars
5 # 1. Sturges' rule; the default, tries to find an optimal number of bars
6 hist(movies$Budget, col="wheat", breaks="sturges", xlab="Budget (USD)", main="
  Sturges' rule")
7 # 2. Freedman-Diaconis's rule; tries to optimise the class width
8 hist(movies$Budget, col="gold", breaks="fd", xlab="Budget (USD)", main="FD rule
  ")
9 # 3. Scott's rule; also tries to optimise the class width
10 hist(movies$Budget, breaks="scott", xlab="Budget (USD)", main="Scott's rule")

```

```

1 par(mfrow=c(1,2))
2 # Let's try a boxplot now
3 boxplot(movies$Budget, col="gold", main="Boxplot of movie budgets", horizontal
  =TRUE, xlab="Budget (USD)")
4
5 # Outliers are spoiling things a bit. Let's remove them
6 boxplot(movies$Budget, col="gold", main="Boxplot of movie budgets", horizontal
  =TRUE, xlab="Budget (USD)", outline=FALSE)

```

The histograms shown in Figure 1 reveal just how positively skewed the movie budget data are. They also show how care should be taken with the method chosen to produce the histogram. Of course, this is a subjective choice, and some experimentation might be necessary – *especially* with such a heavily skewed dataset. The boxplots also indicate just how skewed our dataset is; to get a handle on the distribution of most of our data, it might be preferable to produce the boxplot with the outliers removed.

Table 3 gives a summary of some R commands for commonly-used tables and graphs.

Command	Comment	Example
table	Contingency table	table(x)
barplot	Generate a bar chart	barplot(table(x))
hist	Histogram	hist(x)
plot	Scatterplot	plot(x, y)
points	Add points to a plot	points(x, y)
lines	Add lines to a plot	lines(x, y)
boxplot	Box and whisker plot	boxplot

Table 3: Summary of R commands for commonly-used tables/graphs/graphical features




Figure 1: Histograms for the movie budgets, constructed using three different methods (top row); boxplots (bottom row), with outliers removed (right).

Example 2.4: Load the `music` dataset into RStudio, just as you did in Example 2.3.

- (a) Find the mean and standard deviation for the duration of tracks in this dataset.
- (b) Produce histograms for track duration and danceability; put these histograms on the same panel of plots, use different colours for the histograms (you can choose whichever colours you like!) and make sure you provide suitable titles and labels on the axes.
- (c) You are asked to investigate whether or not the duration of tracks has changed over time.
 - (i) Use the `plot` command in R to produce a scatterplot of the year each track was released against track duration.
 - (ii) Try out the following code, and comment on what you see in the plots in part (c).

```
1 boxplot(music$duration ~ music$decade)
```

- (d) On average, how do energy scores for “Hip Hop” tracks compare to “Electronic” tracks?

Example 2.5: The code in the following listing will download COVID-19 data from the website of the *European Centre for Disease Prevention and Control*. Note: this dataset is dynamic – it is updated every day, so the results you get will differ to those shown in the listings below.

```
1 covid = read.csv("https://opendata.ecdc.europa.eu/covid19/casedistribution /  
2 csv", na.strings = "", fileEncoding = "UTF-8-BOM")
```

- (a) Use the `head` command to take a look at the first 6 rows of the dataframe, to get a feel for the information contained in the dataset. Use the `dim` command to find out the size of the dataframe.
- (b) We can use the `attach` command to directly access the columns in the dataframe by their headings. This might make it a bit easier to isolate data for certain countries. For example:

```
1 # Enable access to columns via their headings  
2 attach(covid)  
3  
4 # Find the total number of cases, and deaths, globally (Data as of 30/09/20)  
5 sum(cases); sum(deaths)  
6 [1] 33714595  
7 [1] 1008932  
8  
9 # What about Great Britain? (Data as of 30/09/20)  
10 sum(cases[countryterritoryCode=="GBR"], na.rm=TRUE)  
11 [1] 446156  
12 sum(deaths[countryterritoryCode=="GBR"], na.rm=TRUE)  
13 [1] 42072
```

```

1 # What about a plot of cases/deaths through time in Great Britain? See Figure 2
2 par(mfrow=c(1,2))
3 plot(rev(cases[countryterritoryCode=="GBR"]), type="l", xlab="Day", ylab=
4      "Confirmed cases", main="GBR COVID-19 cases")
5 plot(rev(deaths[countryterritoryCode=="GBR"]), type="l", xlab="Day", ylab=
6      "Confirmed deaths", main="GBR COVID-19 deaths")

```

- (i) Find the total number of confirmed cases, and deaths, due to COVID-19, for three other countries of your choosing.
- (ii) Why is it necessary to use the `rev` command in lines 3 and 4 in the above listing?
- (iii) For the three other countries you chose in part (i), produce plots showing confirmed cases and deaths through time.




Figure 2: Daily cases and deaths due to COVID-19 in Great Britain.

2.6 Functions, for loops and if statements

The final part of this R Refresher will focus on what was probably the biggest part of the stage 1 R course: functions, for loops and if statements. In this section of the notes, we will focus on a single example that will pull all of these aspects of R programming together: How to fill a sticker album!

The 2018 *Panini* sticker album for the Fifa world cup football finals had a whopping 682 stickers to collect (including 40 awesome glitter stickers!). Each packet of stickers cost 80p and contained 5 stickers. The obvious question to ask is:

“How much does it cost to fill a Panini World Cup ‘18 sticker album?”

Now look at the code in the listing below.

```

1 # Create a vector for the sticker sequence - suppose for now that we only have 50 stickers
2 (stickers = seq(1, 50))      # Can actually be written more concisely as stickers = 1:50
3 [1]  1  2  3  4  5  6  7  8  9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24
4 [25] 25 26 27 28 29 30 31 32 33 34 35 36 37 38 39 40 41 42 43 44 45 46 47 48
5 [49] 49 50
6
7 # Now let's buy a packet!
8 (packet = sample(1:50, 5, replace=TRUE)) # Random sample of size 5 from {1, 2, ..., 50}
9 [1] 10 23 20 48 11                # We got stickers 10, 11, 20, 23 and 48!
10
11 # As a way of keeping track of the stickers we have collected, we could set the corresponding sticker
12 # sequence entry to zero
13 stickers[packet[1]] = 0      # Will set stickers[10] to 0
14 stickers[packet[2]] = 0      # Will set stickers[23] to 0
15 stickers[packet[3]] = 0      # Will set stickers[20] to 0
16 stickers[packet[4]] = 0      # Will set stickers[48] to 0
17 stickers[packet[5]] = 0      # Will set stickers[11] to 0
18
19 # That's a bit repetitive! We could put this inside a loop
20 for(j in 1:5){
21   stickers[packet[j]] = 0}
22
23 stickers
24 [1]  1  2  3  4  5  6  7  8  9  0  0 12 13 14 15 16 17 18 19  0 21 22  0 24
25 [25] 25 26 27 28 29 30 31 32 33 34 35 36 37 38 39 40 41 42 43 44 45 46 47  0
26 [49] 49 50

```

After running the code in the above listing a few times, you should start to get a feel for what's going on. The aim is to keep buying packets of stickers and setting the corresponding entry in the sticker sequence to zero as a way of keeping track of the stickers we have collected. The code above adjusts this sequence after buying just a single packet of stickers. We need to do this repeatedly, an unknown (and probably very large!) number of times. So, another for loop is required.

```

1 # Re-set the sticker sequence
2 (stickers = seq(1, 50))
3   [1]  1  2  3  4  5  6  7  8  9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24
4   [25] 25 26 27 28 29 30 31 32 33 34 35 36 37 38 39 40 41 42 43 44 45 46 47 48
5   [49] 49 50
6
7 # Now let's start a big for loop, in which we repeatedly buy packets of stickers
8 for(i in 1:10000){           # 10,000 is a big number - surely this will be enough packets?
9   packet = sample(1:50, 5, replace=TRUE)
10
11   for(j in 1:5){
12     stickers[packet[j]] = 0} # Close the j loop
13   }# Close the i loop
14 stickers
15   [1] 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
16   [37] 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0

```

Looking at the listing above, it looks like we've filled the sticker album after buying 10,000 packets of stickers – notice that all entries in the sticker sequence vector are zero. But, we don't know exactly when this happened. Did it happen after we bought the 10,000th packet? Or after we bought the 200th packet? Or the 1524th packet?! We need to 'break' the code as soon as all entries in the sticker sequence are zero.

```

1 (stickers = seq(1, 50))
2   [1]  1  2  3  4  5  6  7  8  9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24
3   [25] 25 26 27 28 29 30 31 32 33 34 35 36 37 38 39 40 41 42 43 44 45 46 47 48
4   [49] 49 50
5
6 for(i in 1:10000){
7   packet = sample(1:50, 5, replace=TRUE)
8
9   for(j in 1:5){
10    stickers[packet[j]] = 0}
11
12 # Now inside the for loop use an if statement to break the code when all entries in the sticker sequence are zero
13 if(sum(stickers)==0) {
14   break}
15 }
16
17 # When did the code break?
18 i
19 [1] 37

```

So this time it only took 37 packets of stickers to completely fill the album! Of course, in this toy example we only have 50 stickers to collect. For the *Panini* football sticker album, there are 682!

Really, what we'd like to do is run the code in the above listing repeatedly, to see what happens for different 'people' trying to fill their sticker albums. It might be a good idea to put all of this into a function now.

```

1 FillAlbum = function(N, k){           # N = something large, k = number of stickers to collect
2   stickers = seq(1, k)
3
4   for(i in 1:N){
5     packet = sample(1:k, 5, replace=TRUE)
6
7     for(j in 1:5){
8       stickers[packet[j]] = 0}
9
10    if(sum(stickers)==0) {
11      break}
12  }
13  return(i)}

```

Notice, from the code listing above, the general structure of a function in R:

- The word `function` (line 1), followed by round brackets and some “arguments” to be passed through the function (arguments that we must specify when we execute the function)
- An appropriate name for the function (`FillAlbum` here, see line 1)
- The use of braces surrounding the operations within the function (lines 1 and 13; all other braces relate to the for loops/if statement)
- A return statement (line 13), declaring what we would like the function to spit out

We can now execute our `FillAlbum` function to see how many packets of stickers different “people” might need to buy in order to fill their albums.

```

1 # Let's try for an album containing just 50 stickers again
2 FillAlbum(10000, 50)
3 [1] 52
4 FillAlbum(10000, 50)
5 [1] 32
6 FillAlbum(10000, 50)
7 [1] 40
8
9 # Now let's try for the full album of 682 stickers!
10 FillAlbum(10000, 682)
11 [1] 1290
12 FillAlbum(10000, 682)
13 [1] 1611
14 FillAlbum(10000, 682)
15 [1] 1252

```

Referring to the code in the listing above, we have simulated three people trying to fill the full *Panini* sticker album. Respectively, they needed to buy 1,290, 1,611 and 1,252 packets of stickers to do this. We could now execute this function many times, within another for loop.

```

1 # Let's simulate 1000 people trying to fill their albums
2 packets_needed = vector("numeric", 1000)
3
4 for(k in 1:1000){
5   packets_needed[k] = FillAlbum(10000, 682)}
6
7 # Each packet costs 80p
8 cost = packets_needed*0.8
9
10 mean(packets_needed); mean(cost)
11 [1] 974.609
12 [1] 779.6872
13
14 sd(packets_needed); sd(cost)
15 [1] 181.351
16 [1] 145.0808
17
18 par(mfrow=c(1,2))
19 hist(packets_needed, col="misty rose", xlab="Packets", main="Packets needed")
20 hist(cost, col="chocolate", xlab="Pounds", main="Cost")

```




Figure 3: Histograms showing the number of packets of stickers needed to fill the *Panini* football sticker album (left), and the associated cost (right).

This analysis is really insightful! We can see that, on average, it takes around 975 packets of stickers to fill the album (standard deviation 181), with an associated cost of around £780 (standard deviation £145). The histograms in Figure 3 reveal the full distribution for the number of packets needed and the associated cost. The larger the number of people simulated (we used 1,000 here), the more accurate this distribution will be!

3 Simulating Random Variables

3.1 Discrete distributions: Binomial, geometric and Poisson

The binomial distribution: If $X \sim \text{Bin}(n, p)$ then X has PMF (probability mass function) given by:


There is no closed formula for the CDF (cumulative distribution function; i.e. $\Pr(X \leq x)$).

The R commands `dbinom`, `pbinom`, and `rbinom` are used respectively to compute the PMF, compute the CDF, and simulate from a binomial distribution. Here are examples:

```
1 dbinom(5, 10, 0.7)      # Pr(X = 5), X ~ Bin(10, 0.7)
2 pbinom(4, 10, 0.7)      # Pr(X ≤ 4), X ~ Bin(10, 0.7)
3 rbinom(50, 10, 0.7)      # Sample a Bin(10, 0.7) distribution 50 times
```

Here's an example of how to create a bar plot of a sample from a binomial distribution:

```
1 x = rbinom(50, 10, 0.7)
2 plot(table(x), xlim = c(0, 10), xaxt = 'n', xlab = 'x', ylab = 'frequency')
3 axis(1, at = seq(0, 10))
```



The geometric distribution: If $Y \sim \text{Geom}(p)$ then Y has PMF and CDF given by:

Note that Y takes values in $1, 2, 3, \dots$. R uses a slightly different definition of the geometric distribution. We use the definition that Y is the number of Bernoulli trials with parameter p up to and including the first trial which is a success. The definition in R is the number of trials up to, but **not** including the first success, so in R geometric random variables take values $0, 1, 2, \dots$. Like the binomial distribution, the R commands `dgeom`, `pgeom`, and `rgeom` are used respectively to compute the PDF, compute the CDF, and simulate from the geometric distribution. However, they need to be adjusted to match our definition:

```
1 dgeom(4, 0.2)           # Pr(Y = 5), Y ~ Geom(0.2)
2 pgeom(2, 0.2)           # Pr(Y ≤ 3), Y ~ Geom(0.2)
3 1 + rgeom(100, 0.2)     # Sample a Geom(0.2) distribution 100 times
```

Here's a function to replace `dgeom` with our definition of the geometric distribution:

```
1 mydgeom = function(x, p) {
2   dgeom(x-1, p)}

```

The Poisson distribution: If $Z \sim \text{Po}(\lambda)$ then Z has PMF given by:

There is no closed formula for the CDF.

Examples of R commands:

```
1 dpois(5, 3.5)           # Pr(Z = 5), Z ~ Po(3.5)
2 ppois(2, 3.5)           # Pr(Z ≤ 2), Z ~ Po(3.5)
3 rpois(100, 3.5)         # Sample a Po(3.5) distribution 100 times
```

Summary:

Distribution	Binomial	Poisson	Geometric 
PMF	<code>dbinom(...)</code>	<code>dpois(...)</code>	<code>dgeom(...)</code>
CDF	<code>pbinom(...)</code>	<code>ppois(...)</code>	<code>pgeom(...)</code>
sample	<code>rbinom(...)</code>	<code>rpois(...)</code>	<code>rgeom(...)</code>

3.2 Continuous random variables

R has functions for the uniform, exponential and normal distributions, where we now refer to the probability *density* function, rather than probability *mass* function:

Distribution	Uniform	Exponential	Normal
PDF	dunif(...)	dexp(...)	dnorm(...)
CDF	punif(...)	pexp(...)	pnorm(...)
Quantile	qunif(...)	qexp(...)	qnorm(...)
sample	runif(...)	rexp(...)	rnorm(...)

Examples of R commands:

```

1 dunif(3, 2, 5)      #  $f_X(3), X \sim U(2, 5)$ 
2 pexp(1.5, 5)        #  $F_Y(1.5), Y \sim \text{Exp}(5)$ 
3 rnorm(10, 0, 2)     # Sample a  $N(0, 2^2)$  distribution 10 times

```

For the standard uniform ($U(0, 1)$) and standard normal ($N(0, 1)$) distributions you don't need to provide the parameters $a = 0$, $b = 1$ and $\mu = 0$, $\sigma = 1$ respectively. For example:

```

1 runif(20)           # Samples a  $U(0, 1)$  distribution 20 times
2 pnorm(1.96)          #  $\Pr(Z < 1.96), Z \sim N(0, 1)$ 
3 [1] 0.9750021

```

Quantile functions such as qexp and qnorm solve equations of the form $F_X(\alpha) = p$, for α , given a probability p . For example:

```

1 qnorm(0.9750021)
2 [1] 1.96

```

Example 3.1: Suppose annual maximum wave heights observed off the coast at a flood-prone town are assumed Normally distributed, with mean 2 metres and standard deviation 0.5 metres.

- Write down the R command to find the probability that the largest wave height next year will exceed 3.25 metres. **Answer:** _____
- Write down the R command to estimate the height of a new sea wall such that we might expect the town to be flooded, on average, once per century. Why might our modelling assumption be invalid? **Answer:** _____

Example 3.2: Use R to generate a sample of size 10,000 from an exponential distribution with rate 2.5, and store this sample in the vector x.

- Use your vector x to estimate $\Pr(X \leq 0.5)$, where $X \sim \text{Exp}(2.5)$.
Answer: _____
- How does your estimate in part (a) compare to the true probability? Use R here – do not perform any calculations by hand. **Answer:** _____

Example 3.3: Write down the R code, and corresponding result, to answer each of the questions below. The first one has been done for you.

Model	Question	R code	Answer
$X \sim Po(3)$	Find $\Pr(X = 6)$.	<code>dpois(6,3)</code>	0.0504
	Find $\Pr(X < 5)$.		
	Generate 3 observations.		
$Y \sim Bin(10, 0.2)$	Find $\Pr(Y \leq 4)$.		
	Find $\Pr(Y \geq 4)$.		
	Find y such that $\Pr(Y \leq y) = 0.677$.		
$Z \sim N(0, 1)$	Find $\Pr(Z \leq 0.2)$.		
	Find $\Pr(Z > -2.5)$.		
	Generate 3 observations.		

3.3 Some more advanced probability models

The main reason to use computers for Mathematics and Statistics is to solve problems which cannot be done by hand or which require very lengthy calculations. This section contains a simple motivating example to illustrate this.

The number of arrivals per day at a university IT help-desk can be modelled using a Poisson distribution. However, it is thought that the mean of the Poisson distribution might vary from day-to-day. For example, the desk is often busier nearer the end of term due to more students having deadlines. The following model is proposed. Suppose the number of arrivals $X|\Lambda \sim Po(\Lambda)$ where the rate Λ is itself a random variable, with $\Lambda \sim Exp(c)$ for a constant $c = 0.05$. What can we say about the distribution of X ?

The random variable X does not have a distribution which we can immediately recognize. In fact, much can be done analytically, but we will use simulation to answer the following questions:


1. What are the expectation and variance of X ?
2. What is $\Pr(X > 30)$?

Here is code to simulate a number n of samples from the **distribution** of X :

```
1 arrivals = function(n, c = 0.05) {
2   x = vector(mode = 'numeric', length = n)
3   for (i in 1:n) {
4     lambda = rexp(1, rate = c)
5     x[i] = rpois(1, lambda)
6   }
7   return(x)
8 }
```

Here's how to create a bar plot of the sample and compute the **sample** mean and variance:

```
1 x = arrivals(500)
2 plot(table(x), xlim = c(0, 50), xaxt = 'n', xlab = 'x', ylab = 'frequency')
3 axis(1, at = seq(0, 50, 5))
```



```
1 mean(x); var(x)
2 [1] 19.618
3 [1] 382.1764
```

The mean and variance of the distribution of X will be close to the sample mean and variance that we see in the output above. To calculate $\Pr(X > 30)$ we calculate the proportion of times this occurs in the sample:

```
1 sum(x>30)/500
2 [1] 0.202
```

The bigger the value we use for n , the closer the sample distribution will be to the true distribution for X ; and the better our estimates of the mean, variance and probabilities will be.

3.4 Mixtures of normals

Suppose $\mu_1, \mu_2, \sigma_1 > 0, \sigma_2 > 0$ are fixed constants, and w_1, w_2 are positive constants with $w_1 + w_2 = 1$. Consider the following function:

$$f(x) = w_1 f_1(x) + w_2 f_2(x)$$

where f_1 and f_2 are the density functions for $Z_1 \sim N(\mu_1, \sigma_1^2)$ and $Z_2 \sim N(\mu_2, \sigma_2^2)$ respectively.


Check that $f(x)$ represents a valid probability density function:

If a random variable X has PDF corresponding to $f(x)$ we say it is a **mixture** of normal distributions $N(\mu_1, \sigma_1^2)$ and $N(\mu_2, \sigma_2^2)$ with weights w_1, w_2 . Note that X is **not** the sum of two normal random variables, i.e.

$$X \neq w_1 Z_1 + w_2 Z_2 \quad \text{where } Z_i \sim N(\mu_i, \sigma_i^2), \quad i = 1, 2.$$

The right hand side is actually another normal distribution (as seen in MAS1604), but X **does not have** the PDF of a normal random variable.

For example: suppose $\mu_1 = 3, \sigma_1 = 1$ and $\mu_2 = 6, \sigma_2 = 2$ with $w_1 = w_2 = 1/2$. Here is a plot of $f(x)$:



Consider how to sample X from this distribution:

1. Sample a random variable $J \in \{1, 2\}$ such that $\Pr(J = 1) = w_1$ and $\Pr(J = 2) = w_2$. The random variable J tells you which **component** of the mixture to sample X from.
2. If $J = 1$ then sample X from $N(\mu_1, \sigma_1^2)$, but if $J = 2$ then sample X from $N(\mu_2, \sigma_2^2)$.

Here is an R function to perform this sampling n times:

```

1  normal.mixture = function(n, mu1, sig1, w1, mu2, sig2, w2) {
2    p = c(w1, w2)
3    x = vector(mode = 'numeric', length = n)
4    for (i in 1:n) {
5      j = sample(c(1, 2), 1, prob = p)
6      if (j==1) {
7        x[i] = rnorm(1, mu1, sig1)
8      }
9      else {
10       x[i] = rnorm(1, mu2, sig2)
11     }
12   }
13   return(x)
14 }
```

Write down any notes about this function.

Normal mixture distributions can contain more than two components – the generalization is very straightforward. They offer a very flexible way to model data with different shaped distributions.

Example 3.4: Investigations into infants suffering from *severe idiopathic respiratory distress syndrome* have shown that whether the infant survives may be related to their weight at birth. Suppose that you are interested in developing a distribution for the birth weight X of such infants. You have a normal $N(2.3, 0.52^2)$ distribution for the birth weight (in kg) of infants who survive and a normal $N(1.7, 0.66^2)$ distribution for infants who die.

- (a) If you believe that the proportion of infants that survive is 0.6, write down a suitable Normal mixture distribution for X .
- (b) Produce some R code to sample from the normal mixture distribution you propose in part (a).
- (c) Use your code in part (b) to simulate a sample of size 1000 from your normal mixture distribution, and produce a histogram of your sample.

4 The Inverse CDF Method

Let $U \sim U(0, 1)$. Suppose $F(x)$ is a well-defined CDF which is invertible. Then the random variable $X = F^{-1}(U)$ has CDF $F(x)$.

Proof:

Example 4.1: Suppose we obtain two observations from a $U(0, 1)$ distribution: 0.1 and 0.85. Use these values to obtain two observations from $X \sim \text{Exp}(2)$.

Example 4.2: Write some R code to sample n observations from an $Exp(2)$ distribution. Do not use the built-in `rexp` function. Check your code by comparing a histogram of some simulated realisations against the PDF for an $Exp(2)$ distribution. Also check your code against results from the built-in `rexp` function.

Example 4.3: The *Pareto distribution*, named after the Italian civil engineer, economist, and sociologist Vilfredo Pareto, is a power-law probability distribution that is used in the description of social, scientific, geophysical, actuarial, and many other types of observable phenomena. It was originally used to describe the distribution of wealth in a society, fitting the trend that a large portion of wealth is held by a small fraction of the population.

If $X \sim \text{Pareto}(\sigma)$, then it has cumulative distribution function

$$F_X(x; \sigma) = 1 - (\sigma/x)^{\log_4 5}, \quad x \geq \sigma, \sigma > 0.$$

Write a function in R to produce realisations from this distribution, and test your function out for $\sigma = \{1, 2, 5\}$.

Example 4.4: The *Weibull distribution* is often used in survival data analysis to model lifetimes (e.g. the lifetimes of patients receiving a particular drug). If $X \sim \text{Weibull}(\lambda, \kappa)$, then it has cumulative distribution function

$$F_X(x; \lambda, \kappa) = 1 - e^{-(x/\lambda)^\kappa}, \quad x \geq 0, \lambda, \kappa > 0.$$

Write a function in R to produce realisations from this distribution, and test your function out (a recent study suggested that, for a particular cancer drug, $\lambda = 2$ and $\kappa = 1.5$ for survival times in months after starting this drug).

Example 4.5: Patients arriving at a doctor's surgery have appointment times 15 minutes apart. The random variable L measures how late a patient is for his/her appointment, with negative values denoting early arrivals. The PDF of L is

$$f_L(\ell) = \begin{cases} 0.0025(\ell + 20), & -20 \leq \ell \leq 0, \\ 0.05e^{-\ell/10}, & \ell > 0. \end{cases}$$

- (a) Sketch $f_L(\ell)$.
- (b) Show that the CDF of L is

$$F_L(\ell) = \begin{cases} 0, & \ell < -20, \\ 0.00125(\ell + 20)^2, & -20 \leq \ell \leq 0, \\ 1 - \frac{e^{-\ell/10}}{2}, & \ell > 0. \end{cases}$$

- (c) Suppose that 0.205 and 0.713 are two realized values of $U \sim U(0, 1)$. Show that the corresponding realized values of L are -7.19 and 5.55 under the inverse CDF simulation method.

5 Transformation of Random Variables


5.1 Transforming a normal random variable

The normal distribution is very widely used to model data. However, normal random variables take values on the entirety of $\mathbb{R}$ and they are symmetric around the mean. This is unrealistic in many situations – for example, if we want to model people's heights or weights. One way to obtain a positive random variable is to define

$$Y = e^X, \quad \text{where } X \sim N(\mu, \sigma^2).$$

Then Y is a random variable which is a **transformation** of X . What does the distribution of Y look like? We can use sampling to find out, at least approximately:

```
1 mu = 1.5
2 sigma = 0.3
3 n = 250
4 y = exp(rnorm(n, mean = mu, sd = sigma))
5 hist(y, xlab = 'y', ylab = 'frequency', main = '')
```




The random variable Y can only adopt positive values, and the distribution is positively skewed (long tail on the right). This kind of distribution is useful to model quantities like height – but what is the PDF of Y ? A first guess is $f_Y(y) = f_X(\log(y))$ since $\log$ maps Y back to X , but this is **not** the correct answer.

5.2 Computing PDFs for transformed random variables

Example 5.1: If $X \sim U(0, 1)$, what is the PDF of $Y = X^3$?

To confirm this result, we can simulate Y . Does the histogram look right?

```
1 y = (runif(500))^3  
2 hist(y, xlab = 'y', ylab = 'frequency', main = '')
```



Notes:

Example 5.2: If $X \sim N(0, 1)$, what is the PDF of $Y = X^2$?

Example 5.3: If $X \sim N(\mu, \sigma^2)$, what is the PDF of $Y = e^X$? This is the example used at the start of this section. We say Y has the **log-normal** distribution and write $Y \sim LN(\mu, \sigma^2)$.

5.3 The transformation method

These examples are all special cases of the following result.

If X has PDF $f_X(x)$ and $h(x)$ is a monotonically increasing or monotonically decreasing function then h is invertible and the PDF of $Y = h(X)$ is

$$f_Y(y) = f_X(h^{-1}(y)) \times \left| \frac{d(h^{-1}(y))}{dy} \right|$$

when y is in the range of $h(x)$ and zero otherwise.

This result allows direct calculation of $f_Y(y)$ and is called the **univariate transformation method**. We will see this result again in semester 2 of this module.

Proof:


Example 5.4: When $X \sim U(0, 1)$, use the transformation method to find the PDF of $Y = X^3$.

Example 5.5: When $X \sim U(-1, 1)$ and $Y = e^X$, use the transformation method to find the PDF of Y .

Example 5.6: Use the transformation method to find the PDF of the log-normal distribution.

Example 5.7: Suppose $X \sim U(-\frac{\pi}{2}, \frac{\pi}{2})$ and $Y = \tan(X)$. Use simulation to study the distribution of Y and explicitly compute the PDF.

```
1 y = tan(runif(100, -pi/2, pi/2))  
2 hist(y, xlab = 'y', ylab = 'frequency', main = '')
```



This distribution looks odd! Calculate the PDF:-

6 Bivariate distributions

In this section we consider general simulation from a bivariate distribution, as well as simulation from the bivariate Normal distribution. In this week's practical we will also look at *transformations* of bivariate random variables.

6.1 Sampling from bivariate distributions

Suppose we have a pair of continuous random variables (X, Y) with joint PDF $f_{X,Y}(x, y)$. Recall that

$$f_{X,Y}(x, y) = f_X(x)f_{Y|X}(y | X = x)$$

where $f_X(x)$ is the **marginal** PDF of X and $f_{Y|X}(y | X = x)$ is the **conditional** PDF of Y given $X = x$. One way to simulate realizations of the pair (X, Y) is:

1. Simulate a realization x of X from the marginal distribution of X .
2. Simulate a realization y of Y from the conditional distribution of Y given $X = x$.

In the **trivariate case** we can add a third step: simulate Z given $Y = y$ and $X = x$.

Example 6.1: Suppose $X \sim U(0.1, 0.5)$ and

$$f_{Y|X}(Y | X = x) = \begin{cases} xe^{-xy}, & \text{when } y \geq 0, \\ 0 & \text{otherwise.} \end{cases}$$

Write an R function to generate a $2 \times n$ **matrix** of samples from this distribution, and produce a scatter plot.

Note that, given that $X = x$, Y has an exponential distribution with parameter $\lambda = X$:

$$Y | X = x \sim \text{Exp}(x).$$


```

1 example4.1 = function(n) {
2   output = matrix(0, 2, n)
3   x = runif(n, 0.1, 0.5)
4   y = rexp(n, rate = x)
5   output[1, ] = x
6   output[2, ] = y
7   output
8 }
```

Make some notes about the R function.

To obtain a scatterplot:

```
1 test = example4.1(200)
2 plot(test[1,], test[2,], xlab = 'x', ylab = 'y')
```



Example 6.2: (Trivariate case.) Suppose that $X \sim U(-1, 1)$ and $Y \sim \text{Exp}(\lambda)$ independently of X and that

$$f_{Z|X,Y}(Z \mid X = x, Y = y) = \frac{1}{\sqrt{2\pi}y} e^{-\frac{1}{2}\left(\frac{z-x}{y}\right)^2}.$$

Write an R function to sample realizations (X, Y, Z) in the case $\lambda = 5$ and produce a scatterplot of (X, Z) . What would the plot look like if $\lambda = 1$?

Note that, given that $X = x, Y = y$ then Z has a normal distribution:


$$Z \mid X = x, Y = y \sim N(x, y^2).$$

```
1 example4.2 = function(n) {
2   output = matrix(0, 3, n)
3   x = runif(n, -1, 1)
4   y = rexp(n, rate = 5)
5   z = rnorm(n, mean = x, sd = y)
6   output[1, ] = x
7   output[2, ] = y
8   output[3, ] = z
9   output
10 }
```

Make some notes about the R function.

To obtain a scatterplot:

```
1 output = example4.2(2000)
2 plot(output[1,], output[3,], xlab = 'x', ylab = 'z')
```



6.2 The multivariate normal distribution

We met the bivariate normal distribution in MAS1604. The bivariate normal distribution is a generalization of the normal distribution to pairs of random variables, or equivalently, to a distribution on vectors in $\mathbb{R}^2$. The multivariate normal distribution is the analogous distribution on vectors in $\mathbb{R}^n$.

6.2.1 The bivariate normal distribution

Suppose we have constants μ_x , μ_y , σ_x , σ_y and ρ such that $\sigma_x \geq 0$, $\sigma_y \geq 0$ and $-1 \leq \rho \leq 1$. Define the 2×2 matrix Σ by

$$\Sigma = \begin{pmatrix} \sigma_x^2 & \rho\sigma_x\sigma_y \\ \rho\sigma_x\sigma_y & \sigma_y^2 \end{pmatrix}.$$

Then define a joint probability density function by

$$f_{X,Y}(x, y) = \frac{1}{2\pi\sqrt{\det \Sigma}} \exp\left(-\frac{1}{2}Q(x, y)\right) \quad (1)$$

where

$$Q(x, y) = (\underline{x} - \underline{\mu})^T \Sigma^{-1} (\underline{x} - \underline{\mu}) \quad (2)$$

and the column vectors $\underline{x}$, $\underline{\mu}$ are defined by $\underline{x} = (x, y)^T$, $\underline{\mu} = (\mu_x, \mu_y)^T$. If random variables (X, Y) have joint probability density given by $f_{X,Y}$ above, then we say that (X, Y) have a **bivariate normal distribution** and write

$$(X, Y)^T \sim N_2(\underline{\mu}, \Sigma).$$

Some of the terms in the definition can be expanded:

$$\sqrt{\det \Sigma} = \sigma_x \sigma_y \sqrt{1 - \rho^2}$$

and

$$\begin{aligned}
 Q(x, y) &= \frac{1}{\sigma_x^2 \sigma_y^2 (1 - \rho^2)} (x - \mu_x \quad y - \mu_y) \begin{pmatrix} \sigma_y^2 & -\rho \sigma_x \sigma_y \\ -\rho \sigma_x \sigma_y & \sigma_x^2 \end{pmatrix} \begin{pmatrix} x - \mu_x \\ y - \mu_y \end{pmatrix} \\
 &= \frac{1}{1 - \rho^2} \left[\left(\frac{x - \mu_x}{\sigma_x} \right)^2 + \left(\frac{y - \mu_y}{\sigma_y} \right)^2 - \frac{2\rho(x - \mu_x)(y - \mu_y)}{\sigma_x \sigma_y} \right].
 \end{aligned}$$

However, equations (1) and (2) are **much** more important than the expansions just given.

It can be proved that the function $f_{X,Y}(x, y)$ in equation (1) integrates to 1 and therefore defines a valid joint PDF. The proof is omitted.

Remarks:

1. The vector $\underline{\mu} = (\mu_x, \mu_y)^T$ is called the **mean vector** and the matrix Σ is called the **covariance** matrix (or sometimes **variance-covariance** matrix).
2. Functions of the form $F(\underline{x}) = \underline{x}^T \Sigma^{-1} \underline{x}$ are called **quadratic forms**. Quadratic forms are functions $\mathbb{R}^n \rightarrow \mathbb{R}$ which satisfy certain properties. They crop up in several areas of mathematics and statistics.
3. The matrix Σ and its inverse Σ^{-1} are **positive definite**. A matrix A is positive definite if



$$\underline{x}^T A \underline{x} \geq 0$$

for all non-zero vectors $\underline{x}$.



4. It follows that when $\mu_x = \mu_y = 0$, $Q(x, y)$ is of positive definite quadratic form.

Over the next two pages we will visualize the bivariate Normal distribution using 3-D plots and contour plots.



$$\sigma_x = \sigma_y, \rho = 0$$





$$\sigma_x = 2\sigma_y, \rho = 0$$





$$2\sigma_x = \sigma_y, \rho = 0$$





$$\sigma_x = \sigma_y, \rho = 0.75$$



$$\sigma_x = \sigma_y, \rho = -0.75$$



$$2\sigma_x = \sigma_y, \rho = 0.75$$



The previous two pages show contour and surface plots for the density function on the xy -plane for different parameter values. Comments:

1. $Q(x, y) \geq 0$ with equality only when $\underline{x} = \underline{\mu}$. It follows that the density function has its mode at $\underline{x} = \underline{\mu}$. The above plots all have $\mu_x = \mu_y = 0$.
2. Changing the values of μ_x, μ_y does not change the shape of the plots, but corresponds to a translation of the xy -plane i.e. changing μ_x, μ_y just shifts the contours / surface to a new mode position.
3. The contours of equal density are **circular** when $\sigma_x = \sigma_y$ and $\rho = 0$ and **elliptical** when $\sigma_x \neq \sigma_y$ or $\rho \neq 0$.
4. σ_x and σ_y control the extent to which the distribution is **dispersed**.
5. The parameter ρ is the **correlation** of X, Y i.e. $\text{Cor}(X, Y) = \rho$. (The proof is omitted.) Thus for non-zero ρ , the contours are at an angle to the axes.

When (X, Y) has a bivariate normal distribution then given a general set $A \subseteq \mathbb{R}^2$, calculation of the probability $\Pr((X, Y) \in A)$ requires a computer. However, the marginal and conditional distributions of the bivariate normal distribution are particularly simple – one of the reasons the distribution is so useful in statistics.

Suppose $(X, Y)^T \sim N_2(\underline{\mu}, \Sigma)$. The following results hold.

1. The marginal distributions are normal:

$$X \sim N(\mu_x, \sigma_x^2) \quad \text{and} \quad Y \sim N(\mu_y, \sigma_y^2).$$

2. The conditional distributions are normal:

$$X|Y = y \sim N\left(\mu_x + \rho \frac{\sigma_x}{\sigma_y}(y - \mu_y), \sigma_x^2(1 - \rho^2)\right) \quad \text{and}$$

$$Y|X = x \sim N\left(\mu_y + \rho \frac{\sigma_y}{\sigma_x}(x - \mu_x), \sigma_y^2(1 - \rho^2)\right).$$

3. When $\rho = 0$, X and Y are independent.
4. Linear combinations of X and Y are also normally distributed:

$$aX + bY \sim N(a\mu_x + b\mu_y, a^2\sigma_x^2 + b^2\sigma_y^2 + 2ab\rho\sigma_x\sigma_y)$$

where a, b are constants. Thus a **linear combination** of normal random variables is also normally distributed itself.

Proofs are omitted.

6.2.2 Simulating bivariate normal random variables

R has built-in functions to simulate bivariate and multivariate normal random variables. However, we can simulate using the strategy introduced in section 6.1: first simulate from the marginal for X (or marginal for Y if preferred), and then simulate Y from the conditional $Y|X = x$.


Example 6.3: Suppose $(X, Y)^T \sim N_2(\underline{\mu}, \Sigma)$ where $\mu_x = 2$, $\mu_y = 3$, $\sigma_x = 1$, $\sigma_y = 1$ and $\rho = 0.5$. Simulate a sample of size 500 from this distribution and draw a scatter plot. Use simulation to find $\Pr(X^2 + Y^2 < 9)$.

The marginal distribution of X is $X \sim N(2, 1^2)$. Using the formula for the conditional

$$\begin{aligned} Y|X = x &\sim N\left(\mu_y + \rho \frac{\sigma_y}{\sigma_x}(x - \mu_x), \sigma_y^2(1 - \rho^2)\right) \\ &\sim N(3 + 0.5(x - 2), 0.75). \end{aligned}$$

To sample:

```
1 npts = 500
2 x = rnorm(npts, mean = 2, sd = 1)
3 y = rnorm(npts, mean = 3+0.5*(x-2), sd = sqrt(0.75))
```



To find $\Pr(X^2 + Y^2 < 9)$ approximately, count the number of points in the region:

```
1 npts = 10000
2 x = rnorm(npts, mean = 2, sd = 1)
3 y = rnorm(npts, mean = 3+0.5*(x-2), sd = sqrt(0.75))
4 f = x^2+y^2
5 sum(f<9)/npts
6 [1] 0.2877
```

6.2.3 The multivariate normal distribution

The multivariate normal distribution is a distribution on vectors in $\mathbb{R}^n$. Suppose that $\underline{X}$ is a random vector with n entries, i.e. $\underline{X} = (X_1, \dots, X_n)^T$. Then

$$\underline{X} \sim N_n(\underline{\mu}, \Sigma)$$

if $X_1, \dots, X_n$ have joint PDF given by

$$f_{\underline{X}}(\underline{x}) = \frac{1}{2\pi\sqrt{\det \Sigma}} \exp\left(-\frac{1}{2}Q(\underline{x})\right)$$

where

$$Q(\underline{x}) = (\underline{x} - \underline{\mu})^T \Sigma^{-1} (\underline{x} - \underline{\mu}).$$

This definition makes sense for any column vector $\underline{\mu} \in \mathbb{R}^n$ and any **positive definite** $n \times n$ matrix Σ . Some remarks:

1. The vector $\underline{\mu}$ is the **mean** of the distribution and Σ is called the **covariance** matrix.
2. All the marginal distributions of $\underline{X}$ are normal. (We do not specify their parameters here, however).
3. Similarly, all the conditional distributions of $\underline{X}$ are normal. (Again, we do not specify the parameters of these distributions here).

Example 6.4: Suppose $\underline{X} \sim N_3(0, \Sigma)$ where

$$\Sigma = \begin{bmatrix} 3 & 2 & 1 \\ 2 & 2 & 1 \\ 1 & 1 & 3 \end{bmatrix}.$$


Use R to sample from this distribution. Create a scatterplot of (X, Y) samples and compute the correlation.

In R, you use the `mvrnorm` function to sample:

```
1 library(MASS)
2 npts = 500
3 mu = c(0, 0, 0)
4 sigma = matrix(data = c(3,2,1,2,2,1,1,1,3), nrow = 3, byrow = TRUE)
5 x = mvrnorm(npts, mu, sigma)
```

Now `x` has dimension `npts`×`3`, so column 1 is X and column 2 is Y :

```
1 plot(x[,1], x[,2], xlab = 'x', ylab = 'y')
```



The sample correlation is:

```
1 cor(x[,1], x[,2])  
2 [1] 0.812629
```