

# Exercises that Practice and Extend Skills with R

John Maindonald

April 15, 2009

**Note:** Asterisked exercises (or in the case of “IV: Examples that Extend or Challenge”, set of exercises) are intended for those who want to explore more widely or to be challenged. The subdirectory **scripts** at <http://www.math.anu.edu.au/~courses/r/exercises/scripts/> has the script files.

Also available are Sweave (**.Rnw**) files that can be processed through R to generate the  $\LaTeX$  files from which pdf’s for all or some subset of exercises can be generated. The  $\LaTeX$  files hold the R code that is included in the pdf’s, output from R, and graphics files.

There is extensive use of datasets from the *DAAG* and *DAAGextras* packages. Other required packages, aside from the packages supplied with all binaries, are:

**randomForest** (XII:rdiscrim-lda; XI:rdiscrim-ord; XIII: rdiscrim-trees; XVI:r-largish), **mlbench** (XIII:rdiscrim-ord), **e1071** (XIII:rdiscrim-ord; XV:rdiscrim-trees), **ggplot2** (XIII: rdiscrim-ord), **ape** (XIV: r-ordination), **mcclust** (XIV: r-ordination), **oz** (XIV: r-ordination).

## Contents

<b>I</b>	<b>R Basics</b>	<b>5</b>
1	Data Input	5
2	Missing Values	5
3	Useful Functions	6
4	Subsets of Dataframes	6
5	Scatterplots	7
6	Factors	8
7	Dotplots and Stripplots ( <i>lattice</i> )	8
8	Tabulation	9
9	Sorting	9
10	For Loops	10
11	The <code>paste()</code> Function	10
12	A Function	10
<b>II</b>	<b>Further Practice with R</b>	<b>11</b>
1	Information about the Columns of Data Frames	11
2	Tabulation Exercises	11
3	Data Exploration – Distributions of Data Values	12
4	The <code>paste()</code> Function	12
5	Random Samples	13
6	*Further Practice with Data Input	14

<b>III</b>	<b>Informal and Formal Data Exploration</b>	<b>15</b>
1	Rows with Missing Data – Are they Different	15
2	Comparisons Using Q-Q Plots	16
<b>IV</b>	<b>*Examples that Extend or Challenge</b>	<b>17</b>
1	Further Practice with Data Input	17
2	Graphs with logarithmic scales	17
3	Information on Workspace Objects	18
4	Different Ways to Do a Calculation – Timings	18
5	Functions – Making Sense of the Code	19
6	A Regression Estimate of the Age of the Universe	20
7	Use of <code>sapply()</code> to Give Multiple Graphs	21
8	The Internals of R – Functions are Pervasive	21
<b>V</b>	<b>Data Summary – Traps for the Unwary</b>	<b>23</b>
1	Multi-way Tables	23
2	Weighting Effects – Example with a Continuous Outcome	25
3	Extraction of <code>nassCDS</code>	26
<b>VI</b>	<b>Populations &amp; Samples – Theoretical &amp; Empirical Distributions</b>	<b>27</b>
1	Populations and Theoretical Distributions	27
2	Samples and Estimated Density Curves	28
3	*Normal Probability Plots	30
4	Boxplots – Simple Summary Information on a Distribution	31
<b>VII</b>	<b>Informal Uses of Resampling Methods</b>	<b>33</b>
1	Bootstrap Assessments of Sampling Variability	33
2	Use of the Permutation Distribution as a Standard	34
<b>VIII</b>	<b>Sampling Distributions, &amp; the Central Limit Theorem</b>	<b>35</b>
1	Sampling Distributions	35
2	The Central Limit Theorem	38
<b>IX</b>	<b>Simple Linear Regression Models</b>	<b>41</b>
1	Fitting Straight Lines to Data	41
2	Multiple Explanatory Variables	42
<b>X</b>	<b>Extending the Linear Model</b>	<b>43</b>
1	A One-way Classification – Eggs in the Cuckoo’s Nest	43
2	Regression Splines – one explanatory variable	45
3	Regression Splines – Two or More Explanatory Variables	46
4	Errors in Variables	47
<b>XI</b>	<b>Multi-level Models</b>	<b>49</b>
1	Description and Display of the Data	49

2	Multi-level Modeling	51
3	Multi-level Modeling – Attitudes to Science Data	53
4	*Additional Calculations	53
5	Notes – Other Forms of Complex Error Structure	54
<b>XII</b>	<b>Linear Discriminant Analysis vs Random Forests</b>	<b>55</b>
1	Accuracy for Classification Models – the Pima Data	55
2	Logistic regression – an alternative to lda	60
3	Data that are More Challenging – the crx Dataset	61
4	Use of Random Forest Results for Comparison	62
5	Note – The Handling of NAs	63
<b>XIII</b>	<b>Discriminant Methods &amp; Associated Ordinations</b>	<b>65</b>
1	Discrimination with Multiple Groups	65
<b>XIV</b>	<b>Ordination</b>	<b>71</b>
1	Australian road distances	71
2	If distances must first be calculated ...	73
3	Genetic Distances	73
4	*Distances between fly species	75
5	*Rock Art	76
<b>XV</b>	<b>Trees, SVM, and Random Forest Discriminants</b>	<b>77</b>
1	rpart Analyses – the Pima Dataset	77
2	rpart Analyses – Pima.tr and Pima.te	79
3	Analysis Using <i>svm</i>	81
4	Analysis Using <i>randomForest</i>	82
5	Class Weights	83
6	Plots that show the “distances” between points	83
7	Further Examples	84
<b>XVI</b>	<b>Data Exploration and Discrimination – Largish Dataset</b>	<b>85</b>
1	Data Input and Exploration	85
2	Tree-Based Classification	88
3	Use of <code>randomForest()</code>	89
4	Further comments	90
<b>A</b>	<b>Appendix – Use of the Sweave (.Rnw) Exercise Files</b>	<b>91</b>



## Part I

# R Basics

## 1 Data Input

### *Exercise 1*

The file **fuel.txt** is one of several files that the function `datafile()` (from *DAAG*), when called with a suitable argument, has been designed to place in the working directory. On the R command line, type `library(DAAG)`, then `datafile("fuel")`, thus:<sup>a</sup>

```
> library(DAAG)
> datafile(file="fuel") # NB datafile, not dataFile
```

Alternatively, copy **fuel.txt** from the directory **data** on the DVD to the working directory.

Use `file.show()` to examine the file.<sup>b</sup> Check carefully whether there is a header line. Use the R Commander menu to input the data into R, with the name **fuel**. Then, as an alternative, use `read.table()` directly. (If necessary use the code generated by the R Commander as a crib.) In each case, display the data frame and check that data have been input correctly.

Note: If the file is elsewhere than in the working directory a fully specified file name, including the path, is necessary. For example, to input **travelbooks.txt** from the directory **data** on drive **D:**, type

```
> travelbooks <- read.table("D:/data/travelbooks.txt")
```

For input to R functions, forward slashes replace backslashes.

<sup>a</sup>This and other files used in these notes for practice in data input are also available from the web page <http://www.maths.anu.edu.au/~johnm/datasets/text/>.

<sup>b</sup>Alternatively, open the file in R's script editor (under Windows, go to File | Open script...), or in another editor.

### *Exercise 2*

The files **molclock1.txt** and **molclock2.txt** are in the **data** directory on the DVD.<sup>a</sup>

As in Exercise 1, use the R Commander to input each of these, then using `read.table()` directly to achieve the same result. Check, in each case, that data have been input correctly.

<sup>a</sup>Again, these are among the files that you can use the function `datafile()` to place in the working directory.

## 2 Missing Values

### *Exercise 3*

The following counts, for each species, the number of missing values for the column **root** of the data frame **rainforest** (*DAAG*):

```
> library(DAAG)
> with(rainforest, table(complete.cases(root), species))
```

For each species, how many rows are “complete”, i.e., have no values that are missing?

### *Exercise 4*

For each column of the data frame **Pima.tr2** (*MASS*), determine the number of missing values.

### 3 Useful Functions

#### Exercise 5

The function `dim()` returns the dimensions (a vector that has the number of rows, then number of columns) of data frames and matrices. Use this function to find the number of rows in the data frames `tinting`, `possum` and `possumsites` (all in the *DAAG* package).

#### Exercise 6

Use the functions `mean()` and `range()` to find the mean and range of:

- (a) the numbers 1, 2, ..., 21
- (b) the sample of 50 random normal values, that can be generated from a normal distribution with mean 0 and variance 1 using the assignment `y <- rnorm(50)`.
- (c) the columns `height` and `weight` in the data frame `women`.  
[The *datasets* package that has this data frame is by default attached when R is started.]

Repeat (b) several times, on each occasion generating a new set of 50 random numbers.

#### Exercise 7

Repeat exercise 6, now applying the functions `median()` and `sum()`.

### 4 Subsets of Dataframes

#### Exercise 8

Use `head()` to check the names of the columns, and the first few rows of data, in the data frame `rainforest` (*DAAG*). Use `table(rainforest$species)` to check the names and numbers of each species that are present in the data. The following extracts the rows for the species *Acmena smithii*

```
> library(DAAG)
> Acmena <- subset(rainforest, species=="Acmena smithii")
```

The following extracts the rows for the species *Acacia mabellae* and *Acmena smithii*

```
> AcSpecies <- subset(rainforest, species %in% c("Acacia mabellae",
+                                             "Acmena smithii"))
```

Now extract the rows for all species except *C. fraseri*.

#### Exercise 9

Extract the following subsets from the data frame `ais` (*DAAG*):

- (a) Extract the data for the rowers.
- (b) Extract the data for the rowers, the netballers and the tennis players.
- (c) Extract the data for the female basketballers and rowers.

## 5 Scatterplots

### Exercise 10

Using the `Acmena` data from the data frame `rainforest`, plot `wood` (wood biomass) vs `dbh` (diameter at breast height), trying both untransformed scales and logarithmic scales. Here is suitable code:

```
> Acmena <- subset(rainforest, species=="Acmena smithii")
> plot(wood ~ dbh, data=Acmena)
> plot(wood ~ dbh, data=Acmena, log="xy")
```

### Exercise 11\*

Use of the argument `log="xy"` to the function `plot()` gives logarithmic scales on both the  $x$  and  $y$  axes. For purposes of adding a line, or other additional features that use  $x$  and  $y$  coordinates, note that logarithms are to base 10.

```
> plot(wood~dbh, data=Acmena, log="xy")
> ## Use lm() to fit a line, and abline() to add it to the plot
> Acmena10.lm <- lm(log10(wood) ~ log10(dbh), data=Acmena)
> abline(Acmena10.lm)

> ## Now print the coefficients, for a log10 scale
> coef(Acmena10.lm)
> ## For comparison, print the coefficients for a natural log scale
> Acmena.lm <- lm(log(wood) ~ log(dbh), data=Acmena)
> coef(Acmena.lm)
```

Write down the equation that gives the fitted relationship between `wood` and `dbh`.

### Exercise 12

The `orings` data frame gives data on the damage that had occurred in US space shuttle launches prior to the disastrous Challenger launch of January 28, 1986. Only the observations in rows 1, 2, 4, 11, 13, and 18 were included in the pre-launch charts used in deciding whether to proceed with the launch. Add a new column to the data frame that identifies rows that were included in the pre-launch charts. Now make three plots of `Total` incidents against `Temperature`:

- Plot only the rows that were included in the pre-launch charts.
- Plot all rows.
- Plot all rows, using different symbols or colors to indicate whether or not points were included in the pre-launch charts.

Comment, for each of the first two graphs, whether and open or closed symbol is preferable. For the third graph, comment on the your reasons for choice of symbols.

Use the following to identify rows that hold the data that were presented in the pre-launch charts:

```
> included <- logical(23) # orings has 23 rows
> included[c(1,2,4,11,13,18)] <- TRUE
```

The construct `logical(23)` creates a vector of length 23 in which all values are `FALSE`. The following are two possibilities for the third plot; can you improve on these choices of symbols and/or colors?

```
> plot(Total ~ Temperature, data=orings, pch=included+1)
> plot(Total ~ Temperature, data=orings, col=included+1)
```

*Exercise 13*

Using the data frame `oddbooks`, use graphs to investigate the relationships between:

(a) weight and volume; (b) density and volume; (c) density and page area.

## 6 Factors

*Exercise 14*

Investigate the use of the function `unclass()` with a factor argument. Comment on its use in the following code:

```
> par(mfrow=c(1,2), pty="s")
> plot(weight ~ volume, pch=unclass(cover), data=allbacks)
> plot(weight ~ volume, col=unclass(cover), data=allbacks)
> par(mfrow=c(1,1))
```

[`mfrow=c(1,2)`: plot layout is 1 row  $\times$  2 columns; `pty="s"`: square plotting region.]

*Exercise 15*

Run the following code:

```
> gender <- factor(c(rep("female", 91), rep("male", 92)))
> table(gender)
> gender <- factor(gender, levels=c("male", "female"))
> table(gender)
> gender <- factor(gender, levels=c("Male", "female")) # Note the mistake
> # The level was "male", not "Male"
> table(gender)
> rm(gender) # Remove gender
```

The output from the final `table(gender)` is

```
gender
  Male female
    0     91
```

Explain the numbers that appear.

## 7 Dotplots and Stripplots (*lattice*)

*Exercise 16*

Look up the help for the lattice functions `dotplot()` and `stripplot()`. Compare the following:

```
> with(ant111b, stripchart(harvwt ~ site)) # Base graphics
> library(lattice)
> stripplot(site ~ harvwt, data=ant111b)
> stripplot(harvwt ~ site, data=ant111b)
> stripplot(harvwt ~ site, data=ant111b)
> stripplot(site ~ harvwt, data=ant111b)
```

*Exercise 17*

Check the class of each of the columns of the data frame `cabbages` (*MASS*). Do side by side plots of `HeadWt` against `Date`, for each of the levels of `Cult`.

```
> stripplot(Date ~ HeadWt | Cult, data=cabbages)
```

The lattice graphics function `stripplot()` seems generally preferable to the base graphics function `stripchart()`. It has functionality that `stripchart()` lacks, and a consistent syntax that it shares with other lattice functions.

*Exercise 18*

In the data frame `nsw74psid3`, use `stripplot()` to compare, between levels of `trt`, the continuous variables `age`, `educ`, `re74` and `re75`

It is possible to generate all the plots at once, side by side. A simplified version of the plot is:

```
> stripplot(trt ~ age + educ, data=nsw74psid1, outer=T, scale="free")
```

What are the effects of `scale = "free"`, and `outer = TRUE`? (Try leaving these at their defaults.)

## 8 Tabulation

*Exercise 19*

In the data set `nswpsdi1` (*DAAGxtras*), do the following for each of the two levels of `trt`:

- (a) Determine the numbers for each of the levels of `black`;
- (b) Determine the numbers for each of the levels of `hispanic`; item Determine the numbers for each of the levels of `marr` (married).

## 9 Sorting

*Exercise 20*

Sort the rows in the data frame `Acmena` in order of increasing values of `dbh`.

[Hint: Use the function `order()`, applied to `age` to determine the order of row numbers required to sort rows in increasing order of `age`. Reorder rows of `Acmena` to appear in this order.]

```
> Acmena <- subset(rainforest, species=="Acmena smithii")
> ord <- order(Acmena$dbh)
> acm <- Acmena[ord, ]
```

Sort the row names of `possumsites` (*DAAG*) into alphanumeric order. Reorder the rows of `possumsites` in order of the row names.

## 10 For Loops

### *Exercise 22*

- (a) Create a `for` loop that, given a numeric vector, prints out one number per line, with its square and cube alongside.
- (b) Look up `help(while)`. Show how to use a `while` loop to achieve the same result.
- (c) Show how to achieve the same result without the use of an explicit loop.

## 11 The `paste()` Function

### *Exercise 21*

Here are examples that illustrate the use of `paste()`:

```
> paste("Leo", "the", "lion")
> paste("a", "b")
> paste("a", "b", sep="")
> paste(1:5)
> paste(1:5, collapse="")
```

What are the respective effects of the parameters `sep` and `collapse`?

## 12 A Function

### *Exercise 23*

The following function calculates the mean and standard deviation of a numeric vector.

```
> meanANDsd <- function(x){
+   av <- mean(x)
+   sdev <- sd(x)
+   c(mean=av, sd = sdev) # The function returns this vector
+ }
```

Modify the function so that: (a) the default is to use `rnorm()` to generate 20 random normal numbers, and return the standard deviation; (b) if there are missing values, the mean and standard deviation are calculated for the remaining values.

## Part II

# Further Practice with R

### 1 Information about the Columns of Data Frames

#### *Exercise 1*

Try the following:

```
> class(2)
> class("a")
> class(cabbages$HeadWt)      # cabbages is in the datasets package
> class(cabbages$Cult)
```

Now do `sapply(cabbages, class)`, and note which columns hold numerical data. Extract those columns into a separate data frame, perhaps named `numtinting`.

[Hint: `cabbages[, c(2,3)]` is not the correct answer, but it is, after a manner of speaking, close!]

#### *Exercise 2*

Functions that may be used to get information about data frames include `str()`, `dim()`, `row.names()` and `names()`. Try each of these functions with the data frames `allbacks`, `ant111b` and `tinting` (all in *DAAG*).

For getting information about each column of a data frame, use `sapply()`. For example, the following applies the function `class()` to each column of the data frame `ant111b`.

```
> library(DAAG)
> sapply(ant111b, class)
```

For columns in the data frame `tinting` that are factors, use `table()` to tabulate the number of values for each level.

### 2 Tabulation Exercises

#### *Exercise 3*

In the data set `nswpsid1` (*DAAGxtras*) create a factor that categorizes subjects as: (i) black; (ii) hispanic; (iii) neither black nor hispanic. You can do this as follows:

```
> gps <- with(nswpsid1, 1 + black + hisp*2)
> table(gps)      # Check that there are no 3s, ie black and hispanic!
```

```
gps
 1   2   3
1816 862 109
```

```
> grouping <- c("other", "black", "hisp")[gps]
> table(grouping)
```

```
grouping
black  hisp other
 862   109 1816
```

*Exercise 4*

Tabulate the number of observations in each of the different districts in the data frame `rockArt` (*DAAGxtras*). Create a factor `groupDis` in which all Districts with less than 5 observations are grouped together into the category `other`.

```
> library(DAAGxtras)
> groupDis <- as.character(rockArt$District)
> tab <- table(rockArt$District)
> le4 <- rockArt$District %in% names(tab)[tab <= 4]
> groupDis[le4] <- "other"
> groupDis <- factor(groupDis)
```

### 3 Data Exploration – Distributions of Data Values

*Exercise 5*

The data frame `rainforest` (*DAAG* package) has data on four different rainforest species. Use `table(rainforest$species)` to check the names and numbers of the species present. In the sequel, attention will be limited to the species *Acmena smithii*. The following plots a histogram showing the distribution of the diameter at base height:

```
> library(DAAG)           # The data frame rainforest is from DAAG
> Acmena <- subset(rainforest, species=="Acmena smithii")
> hist(Acmena$dbh)
```

Above, frequencies were used to label the the vertical axis (this is the default). An alternative is to use a density scale (`prob=TRUE`). The histogram is interpreted as a crude density plot. The density, which estimates the number of values per unit interval, changes in discrete jumps at the breakpoints (= class boundaries). The histogram can then be directly overlaid with a density plot, thus:

```
> hist(Acmena$dbh, prob=TRUE, xlim=c(0,50)) # Use a density scale
> lines(density(Acmena$dbh, from=0))
```

Why use the argument `from=0`? What is the effect of omitting it?

[Density estimates, as given by R's function `density()`, change smoothly and do not depend on an arbitrary choice of breakpoints, making them generally preferable to histograms. They do sometimes require tuning to give a sensible result. Note especially the parameter `bw`, which determines how the bandwidth is chosen, and hence affects the smoothness of the density estimate.]

### 4 The `paste()` Function

*Exercise 6*

Here are examples that illustrate the use of `paste()`:

```
> paste("Leo", "the", "lion")
> paste("a", "b")
> paste("a", "b", sep="")
```

*Exercise 6, continued*

```

> paste(1:5)
> paste("a", 1:5)
> paste("a", 1:5, sep="")
> paste(1:5, collapse="")
> paste(letters[1:5], collapse="")
> ## possumsites is from the DAAG package
> with(possumsites, paste(row.names(possumsites), "(", altitude, ")", sep=""))

```

What are the respective effects of the parameters `sep` and `collapse`?

## 5 Random Samples

*Exercise 7*

By taking repeated random samples from the normal distribution, and plotting the distribution for each such sample, one can get an idea of the effect of sampling variation on the sample distribution. A random sample of 100 values from a normal distribution (with mean 0 and standard deviation 1) can be obtained, and a histogram and overlaid density plot shown, thus:

```

> y <- rnorm(100)
> hist(y, probability=TRUE) # probability=TRUE gives a y density scale
> lines(density(y))

```

Repeat several times

In place of the 100 sample values:

- Take 5 samples of size 25, then showing the plots.
- Take 5 samples of size 100, then showing the plots.
- Take 5 samples of size 500, then showing the plots.
- Take 5 samples of size 2000, then showing the plots.

(Hint: By preceding the plots with `par(mfrow=c(4,5))`, all 20 plots can be displayed on the one graphics page. To bunch the graphs up more closely, make the further settings `par(mar=c(3.1,3.1,0.6,0.6), mgp=c(2.25,0.5,0))`)

Comment on the usefulness of a sample histogram and/or density plot for judging whether the population distribution is likely to be close to normal.

Histograms and density plots are, for “small” samples, notoriously variable under repeated sampling. This is true even for sample sizes as large as 50 or 100.

*Exercise 8*

This explores the function `sample()`, used to take a sample of values that are stored or enumerated in a vector. Samples may be with or without replacement; specify `replace = FALSE` (the default) or `replace = TRUE`. The parameter `size` determines the size of the sample. By default the sample has the same size (length) as the vector from which samples are taken. Take several samples of size 5 from the vector `1:5`, with `replace=FALSE`. Then repeat the exercise, this time with `replace=TRUE`. Note how the two sets of samples differ.

*Exercise 9\**

If in Exercise 4 above a new random sample of trees could be taken, the histogram and density plot would change. How much might we expect them to change?

The bootstrap approach treats the one available sample as a microcosm of the population. Repeated with replacement samples are taken from the one available sample. This is equivalent to repeating each sample value an infinite number of times, then taking random samples from the population that is thus created. The expectation is that variation between those samples will be comparable to variation between samples from the original population.

- (a) Take repeated (5 or more) bootstrap samples from the Acmena dataset of Exercise 5, and show the density plots. [Use `sample(Acmena$dbh, replace=TRUE)`].
- (b) Repeat, now with the `cerealsugar` data from *DAAG*.

## 6 \*Further Practice with Data Input

One option is to experiment with using the R Commander GUI to input these data.

*Exercise 10\**

With a live internet connection, files can be read directly from a web page. Here is an example:

```
> webfolder <- "http://www.maths.anu.edu.au/~johnm/datasets/text/"
> webpage <- paste(webfolder, "molclock.txt", sep="")
> molclock <- read.table(url(webpage))
```

With a live internet connection available, use this approach to input the file `travelbooks.txt` that is available from this same web page.

## Part III

# Informal and Formal Data Exploration

Package: DAAGxtras

## 1 Rows with Missing Data – Are they Different

### *Exercise 1*

Look up the help page for the data frame `Pima.tr2` (*MASS* package), and note the columns in the data frame. The eventual interest is in using use variables in the first seven column to classify diabetes according to `type`. Here, we explore the individual columns of the data frame.

- (a) Several columns have missing values. Analysis methods inevitably ignore or handle in some special way rows that have one or more missing values. It is therefore desirable to check whether rows with missing values seem to differ systematically from other rows.

Determine the number of missing values in each column, broken down by `type`, thus:

```
> library(MASS)
> ## Create a function that counts NAs
> count.na <- function(x)sum(is.na(x))
> ## Check function
> count.na(c(1, NA, 5, 4, NA))
> ## For each level of type, count the number of NAs in each column
> for(lev in levels(Pima.tr2$type))
+   print(sapply(subset(Pima.tr2, type==lev), count.na))
```

The function `by()` can be used to break the calculation down by levels of a factor, avoiding the use of the `for` loop, thus:

```
> by(Pima.tr2, Pima.tr2$type, function(x)sapply(x, count.na))
```

- (b) Create a version of the data frame `Pima.tr2` that has `anymiss` as an additional column:

```
> missIND <- complete.cases(Pima.tr2)
> Pima.tr2$anymiss <- c("miss", "nomiss")[missIND+1]
```

For remaining columns, compare the means for the two levels of `anymiss`, separately for each level of `type`. Compare also, for each level of `type`, the number of missing values.

### *Exercise 2*

- (a) Use strip plots to compare values of the various measures for the levels of `anymiss`, for each of the levels of `type`. Are there any columns where the distribution of differences seems shifted for the rows that have one or more missing values, relative to rows where there are no missing values?

Hint: The following indicates how this might be done efficiently:

```
> library(lattice)
> stripplot(anymiss ~ npreg + glu | type, data=Pima.tr2, outer=TRUE,
+           scales=list(relation="free"), xlab="Measure")
```

*Exercise 2, continued*

- (b) Density plots are in general better than strip plots for comparing the distributions. Try the following, first with the variable `npreg` as shown, and then with each of the other columns except `type`. Note that for `skin`, the comparison makes sense only for `type=="No"`. Why?

```
> library(lattice)
> ## npreg & glu side by side (add other variables, as convenient)
> densityplot( ~ npreg + glu | type, groups=anymiss, data=Pima.tr2,
+             auto.key=list(columns=2), scales=list(relation="free"))
```

## 2 Comparisons Using Q-Q Plots

*Exercise 3*

Better than either strip plots or density plots may be Q-Q plots. Using `qq()` from *lattice*, investigate their use. In this exercise, we use random samples from normal distributions to help develop an intuitive understanding of Q-Q plots, as they compare with density plots.

- (a) First consider comparison using (i) a density plot and (ii) a Q-Q plot when samples are from populations in which one of the means is shifted relative to the other. Repeat the following several times,

```
> y1 <- rnorm(100, mean=0)
> y2 <- rnorm(150, mean=0.5) # NB, the samples can be of different sizes
> df <- data.frame(gp=rep(c("first", "second"), c(100, 150)), y=c(y1, y2))
> densityplot(~y, groups=gp, data=df)
> qq(gp ~ y, data=df)
```

- (b) Now make the comparison, from populations that have different standard deviations. For this, try, e.g.

```
> y1 <- rnorm(100, sd=1)
> y2 <- rnorm(150, sd=1.5)
```

Again, make the comparisons using both density plots and Q-Q plots.

*Exercise 4*

Now consider the data set `Pima.tr2`, with the column `anymiss` added as above.

- (a) First make the comparison for `type="No"`.

```
> qq(anymiss ~ npreg, data=Pima.tr2, subset=type=="No")
```

Compare this with the equivalent density plot, and explain how one translates into the other. Comment on what these graphs seem to say.

- (b) The following places the comparisons for the two levels of `type` side by side:

```
> qq(anymiss ~ npreg | type, data=Pima.tr2)
```

Comment on what this graph seems to say.

NB: With `qq()`, use of “+” to get plots for the different columns all at once will not, in the current version of *lattice*, work.

## Part IV

# \*Examples that Extend or Challenge

## 1 Further Practice with Data Input

### *Exercise 1\**

For a challenging data input task, input the data from **bostonc.txt**.<sup>a</sup>

Examine the contents of the initial lines of the file carefully before trying to read it in. It will be necessary to change `sep`, `comment.char` and `skip` from their defaults. Note that `\t` denotes a tab character.

<sup>a</sup>Use `datafile("bostonc")` to place it in the working directory, or access the copy on the DVD.

### *Exercise 2\**

The function `read.csv()` is a variant of `read.table()` that is designed to read in comma delimited files such as may be obtained from Excel. Use this function to read in the file **crx.data** that is available from the web page <http://mlern.ics.uci.edu/databases/credit-screening/>.

Check the file **crx.names** to see which columns should be numeric, which categorical and which logical. Make sure that the numbers of missing values in each column are the number given in the file **crx.names**

With a live connection to the internet, the data can be input thus:

```
> crxpage <- "http://mlern.ics.uci.edu/databases/credit-screening/crx.data"
> crx <- read.csv(url(crxpage), header=TRUE)
```

## 2 Graphs with logarithmic scales

### *Exercise 3\**

Use of the argument `log="xy"` gives logarithmic scales on both the  $x$  and  $y$  axes. For purposes of adding a line, or other additional features that use  $x$  and  $y$  coordinates, note that logarithms are to base 10.

```
> plot(wood~dbh, data=Acmena, log="xy")
> ## Use lm() to fit a line, and abline() to add it to the plot
> Acmena10.lm <- lm(log10(wood) ~ log10(dbh), data=Acmena)
> abline(Acmena10.lm)

> ## Now print the coefficients, for a log10 scale
> coef(Acmena10.lm)
> ## For comparison, print the coefficients for a natural log scale
> Acmena.lm <- lm(log(wood) ~ log(dbh), data=Acmena)
> coef(Acmena.lm)
```

Write down the equation that gives the fitted relationship between wood and dbh.

### 3 Information on Workspace Objects

#### *Exercise 4\**

The function `ls()` lists, by default, the names of objects in the current environment. If used from the command line, it lists the objects in the workspace. If used in a function, it lists the names of the function's local variables

The following function lists the contents of the workspace:

```
> workls <- function()ls(name=".GlobalEnv")
> workls()
```

- If `ls(name=".GlobalEnv")` is replaced by `ls()`, the function lists the names of its local variables. Modify `workls()` so that you can use it to demonstrate this. [Hint: Consider adapting `if(is.null(name))ls()` for the purpose.]
- Write a function that calculates the sizes of all objects in the workspace, then listing the names and sizes of the largest ten objects.

### 4 Different Ways to Do a Calculation – Timings

#### *Exercise 5\**

This exercise will investigate the relative times for alternative ways to do a calculation. The function `system.time()` will provide timings. The numbers that are printed on the command line, if results are not assigned to an output object, are the user cpu time, the system cpu time, and the elapsed time.

First, create both matrix and data frame versions of a largish data set.

```
> xxMAT <- matrix(runif(480000), ncol=50)
> xxDF <- as.data.frame(xxMAT)
```

Repeat each of the calculations that follow several times, noting the extent of variation between repeats. If there is noticeable variation, make the setting `options(gcFirst=TRUE)`, and check whether this leads to more consistent timings.

NB: If your computer chokes on these calculations, reduce the dimensions of `xxMAT` and `xxDF`

- The following compares the times taken to increase each element by 1:

```
> system.time(invisible(xxMAT+1))[1:3]
> system.time(invisible(xxDF+1))[1:3]
```

- Now compare the following alternative ways to calculate the means of the 50 columns:

```
> ## Use apply() [matrix argument], or sapply() [data frame argument]
> system.time(av1 <- apply(xxMAT, 2, mean))[1:3]
> system.time(av1 <- sapply(xxDF, mean))[1:3]
> ## Use a loop that does the calculation for each column separately
> system.time({av2 <- numeric(50);
+             for(i in 1:50)av[i] <- mean(xxMAT[,i])
+             }[1:3]
> system.time({av2 <- numeric(50);
+             for(i in 1:50)av[i] <- mean(xxDF[,i])
+             }[1:3]
```

*Exercise 5\*, continued*

```
> ## Matrix multiplication
> system.time({colOFones <- rep(1, dim(xxMAT)[2])
+             av3 <- xxMAT %*% colOFones / dim(xxMAT)[2]
+             })[1:3]
```

Why is matrix multiplication so efficient, relative to equivalent calculations that use `apply()`, or that use for loops?

*Exercise 6\**

Pick one of the calculations in Exercise 5. Vary the number of rows in the matrix, keeping the number of columns constant, and plot each of user CPU time and system CPU time against number of rows of data.

## 5 Functions – Making Sense of the Code

*Exercise 7\**

Data in the data frame `fumig` (*DAAGxtras*) are from a series of trials in which produce was exposed to a fumigant over a 2-hour time period. Concentrations of fumigant were measured at times 5, 10, 30, 60, 90 and 120 minutes. Code given following this exercise calculates a concentration-time (c-t) product that measures exposure to the fumigant, leading to the measure `ctsum`.

Examine the code in the three alternative functions given below, and the data frame `fumig` (in the *DAAGxtras* package) that is given as the default argument for the parameter `df`. Do the following:

- Run all three functions, and check that they give the same result.
- Annotate the code for `calcCT1()` to explain what each line does.
- Are fumigant concentration measurements noticeably more variable at some times than at others?
- Which function is fastest? [In order to see much difference, it will be necessary to put the functions in loops that run perhaps 1000 or more times.]

**Code for 3 functions that do equivalent calculations**

```
> ## Function "calcCT1"
> "calcCT1" <-
+ function(df=fumig, times=c(5,10,30,60,90,120), ctcols=3:8){
+   multiplier <- c(7.5,12.5,25,30,30,15)
+   m <- dim(df)[1]
+   ctsum <- numeric(m)
+   for(i in 1:m){
+     y <- unlist(df[i, ctcols])
+     ctsum[i] <- sum(multiplier*y)/60
+   }
+   df <- cbind(ctsum=ctsum, df[, -ctcols])
+   df
+ }
```

```

> ##
> ## Function "calcCT2"
> "calcCT2" <-
+ function(df=fumig, times=c(5,10,30,60,90,120), ctcols=3:8){
+   multiplier <- c(7.5,12.5,25,30,30,15)
+   mat <- as.matrix(df[, ctcols])
+   ctsum <- mat%%multiplier/60
+   cbind(ctsum=ctsum, df[, -ctcols])
+ }
> ##
> ## Function "calcCT3"
> "calcCT3" <-
+ function(df=fumig, times=c(5,10,30,60,90,120), ctcols=3:8){
+   multiplier <- c(7.5,12.5,25,30,30,15)
+   mat <- as.matrix(df[, ctcols])
+   ctsum <- apply(mat, 1, function(x)sum(x*multiplier))/60
+   cbind(ctsum=ctsum, df[, -ctcols])
+ }

```

## 6 A Regression Estimate of the Age of the Universe

### *Exercise 8\**

Install the package *gamair* (from CRAN) and examine the help page for the data frame *hubble*. Type `data(hubble)` to bring the data into the workspace. (This is necessary because the *gamair* package, unlike most other packages, does not use the *lazy loading* mechanism for data.)

- Plot  $y$  (Velocity in  $\text{km sec}^{-1}$ ) versus  $x$  (Distance in Mega-parsec =  $3.09 \times 10^{-19}$  km).
- Fit a line, omitting the constant term; for this the `lm()` function call is

```

kmT0megaparsec <- 3.09*10^(-19)
lm(I(y*kmT0megaparsec) ~ -1 + x, data=hubble) # y & x both mega-parsecs

```

The inverse of the slope is then the age of the universe, in seconds. Divide this by  $60^2 \times 24 \times 365$  to get an estimate for the age of the earth in years.

[The answer should be around  $13 \times 10^9$  years.]

- Repeat the plot, now using logarithmic scales for both axes. Fit a line, now insisting that the coefficient of  $\log(x)$  is 1.0 (Why?) For this, specify

```

lm(log(y) ~ 1 + offset(log(x)), data=hubble)

```

Add this line to the plot. Again, obtain an estimate of the age of the universe. Does this give a substantially different estimate for the age of the universe?

- In each of the previous fits, on an untransformed scale and using logarithmic scales, do any of the points seem outliers? Investigate the effect of omitting any points that seem to be outliers?
- Does either plot seem to show evidence of curvature?  
[See further the note at the end of this set of exercises.]

Note: According to the relevant cosmological model, the velocity of recession of any galaxy from any other galaxy has been constant, independent of time. Those parts of the universe that started with the largest velocities of recession from our galaxy have moved furthest, with no change from the velocity just after after

time 0. Thus the time from the beginning should be  $s/v$ , where  $s$  is distance, and  $v$  is velocity. The slope of the least squares line gives a combined estimate, taken over all the galaxies included in the data frame `gamair`. More recent data suggests, in fact, that the velocity of recession is not strictly proportional to distance.

## 7 Use of `sapply()` to Give Multiple Graphs

### *Exercise 9\**

Here is code for the calculations that compare the relative population growth rates for the Australian states and territories, but avoiding the use of a loop:

```
> oldpar <- par(mfrow=c(2,4))
> invisible(
+ sapply(2:9, function(i, df)
+       plot(df[,1], log(df[, i]),
+           xlab="Year", ylab=names(df)[i], pch=16, ylim=c(0,10)),
+       df=austpop)
+ )
> par(oldpar)
```

Run the code, and check that it does indeed give the same result as an explicit loop.

[Use of `invisible()` as a wrapper suppresses printed output that gives no useful information.]

Note that `lapply()` could be used in place of `sapply()`.

There are several subtleties here:

- (i) The first argument to `sapply()` can be either a list (which is, technically, a non-atomic vector) or a vector.<sup>1</sup> Here, we have supplied the vector `2:9`
- (ii) The second argument is a function. Here we have supplied an anonymous function that has two arguments. The argument `i` takes as its values, in turn, the successive elements in the first argument to `sapply`
- (iii) Where as here the anonymous function has further arguments, they are supplied as additional arguments to `sapply()`. Hence the parameter `df=austpop`.

## 8 The Internals of R – Functions are Pervasive

### *Exercise 10\**

The internals of the R parser's handling of arithmetic and related computations are close enough to the surface that users can experiment with them. This exercise will take a peek.

The binary arithmetic operators `+`, `-`, `*`, `/` and `^` are implemented as functions. (R is a functional language; albeit with features that compromise its purity as a member of this genre!) Try:

```
> "+"(2,5)
> "-"(10,3)
> "/"(2,5)
> "*"("+(5,2), "-"(3,7))
```

<sup>1</sup>By “vector” we usually mean an atomic vector, with “atoms” that are of one of the modes “logical”, “integer”, “numeric”, “complex”, “character” or “raw”. (Vectors of mode “raw” can for our purposes be ignored.)

*Exercise 10\*, continued*

There are two other binary arithmetic operators – `%%` and `%%/`. Look up the relevant help page, and explain, with examples, what they do. Try

```
> (0:25) %% 5
> (0:25) %%/ 5
```

Of course, these are also implemented as functions. Write code that demonstrates this.

Note also that `[]` is implemented as a function. Try

```
> z <- c(2, 6, -3, NA, 14, 19)
> "["(z, 5)
> heights <- c(Andreas=178, John=185, Jeff=183)
> "["(heights, c("Jeff", "John"))
```

Rewrite these using the usual syntax.

Use the function `"["()` to extract, from the data frame `possumsites` (*DAAG*), the altitudes for Byrangery and Conondale.

Note: Expressions in which arithmetic operators appear as explicit functions with binary arguments translate directly into postfix reverse Polish notation, introduced in 1920 by the Polish logician and mathematician Jan Lukasiewicz. Postfix notation is widely used in interpreters and compilers as a first step in the processing of arithmetic expressions. See the Wikipedia article “Reverse Polish Notation”.

## Part V

# Data Summary – Traps for the Unwary

Package: DAAGxtras

## 1 Multi-way Tables

	Small (<2cm)		Large (>=2cm)		Total	
	open	ultrasound	open	ultrasound	open	ultrasound
yes	81	234	192	55	273	289
no	6	36	71	25	77	61
Success rate	93%	87%	73%	69%	78%	83%

Table 1: Outcomes for two different types of surgery for kidney stones. The overall success rates (78% for open surgery as opposed to 83% for ultrasound) favor ultrasound. Comparison of the success rates for each size of stone separately favors, in each case, open surgery.

### Exercise 1

Table 1 illustrates the potential hazards of adding a multiway table over one of its margins. Data are from a study<sup>a</sup> that compared outcomes for two different types of surgery for kidney stones; A: **open**, which used open surgery, and B: **ultrasound**, which used a small incision, with the stone destroyed by ultrasound. The data can be entered into R, thus:

```
> stones <- array(c(81, 6, 234, 36, 192, 71, 55, 25), dim = c(2,
+ 2, 2), dimnames = list(Success = c("yes", "no"), Method = c("open",
+ "ultrasound"), Size = c("<2cm", ">=2cm")))
```

- Determine the success rate that is obtained from combining the data for the two different sizes of stone. Also determine the success rates for the two different stone sizes separately.
- Use the following code to give a visual representation of the information in the three-way table:

```
mosaicplot(stones, sort=3:1)
# Re-ordering the margins gives a more interpretable plot.
```

Annotate the graph to show the success rates?

- Observe that the overall rate is, for open surgery, biased toward the open surgery outcome for large stones, while for ultrasound it is biased toward the outcome for small stones. What are the implications for the interpretation of these data?

[Without additional information, the results are impossible to interpret. Different surgeons will have preferred different surgery types, and the prior condition of patients will have affected the choice of surgery type. The consequences of unsuccessful surgery may have been less serious than for ultrasound than for open surgery.]

<sup>a</sup>Charig, C. R., 1986. Comparison of treatment of renal calculi by operative surgery, percutaneous nephrolithotomy, and extracorporeal shock wave lithotripsy. *British Medical Journal*, 292:879–882

The relative success rates for the two different types of surgery, for the two stone sizes separately, can be calculated thus:

```
> stones[1, , ]/(stones[1, , ] + stones[2, , ])
```

To perform the same calculation after adding over the two stone sizes (the third dimension of the table), do

```
> stones2 <- stones[, , 1] + stones[, , 2]
> stones2[1, ]/(stones2[1, ] + stones2[2, ])
```

### 1.1 Which multi-way table? It can be important!

Each year the National Highway Traffic Safety Administration (NHTSA) in the USA collects, using a random sampling method, data from all police-reported crashes in which there is a harmful event (people or property), and from which at least one vehicle is towed. The data frame `nassCDS` (*DAAGxtras*) is derived from NHTSA data.<sup>2</sup>

The data are a sample. The use of a complex sampling scheme has the consequence that the sampling fraction differs between observations. Each point has to be multiplied by the relevant sampling fraction, in order to get a proper estimate of its contribution to the total number of accidents. The column `weight` (`national = national inflation factor` in the SAS dataset) gives the relevant multiplier.

Other variables than those included in `nassCDS` might be investigated – those extracted into `nassCDS` are enough for present purposes.

The following uses `xtabs()` to estimate numbers of front seat passengers alive and dead, classified by airbag use:

```
library(DAAGxtras)
> abtab <- xtabs(weight ~ dead + airbag, data=nassCDS)
> abtab
      airbag
dead   none   airbag
  alive 5445245.90 6622690.98
  dead  39676.02  25919.11
```

The function `prop.table()` can then be used to obtain the proportions in margin 1, i.e., the proportions dead, according to airbag use:

```
> round(prop.table(abtab, margin=2)["dead", ], 4)
      none airbag
0.0072 0.0039
## Alternatively, the following gives proportions alive & dead
## round(prop.table(abtab, margin=2), 4)
```

The above might suggest that the deployment of an airbag substantially reduces the risk of mortality. Consider however:

```
> abSBtab <- xtabs(weight ~ dead + seatbelt + airbag, data=nassCDS)
> ## Take proportions, retain margins 2 & 3, i.e. airbag & seatbelt
> round(prop.table(abSBtab, margin=2:3)["dead", , ], 4)
      seatbelt
airbag   none belted
  none  0.0176 0.0038
  airbag 0.0155 0.0021
```

The results are now much less favorable to airbags. The clue comes from examination of:

<sup>2</sup>They hold a subset of the columns from a corrected version of the data analyzed in the Meyer (2005) paper that is referenced on the help page for `nassCDS`. More complete data are available from one of the web pages <http://www.stat.uga.edu/~mmeyer/airbags.htm> (SAS transport file) or <http://www.maths.anu.edu.au/~johnm/datasets/airbags/> (R image file).

```
> margin.table(ASdtab, margin=2:3) # Add over margin 1
      airbag
seatbelt  none  airbag
  none 1366088.6 885635.3
  belted 4118833.4 5762974.8
```

In the overall table, the results without airbags are mildly skewed ( 4.12:1.37) to the results for `beltded`, while with airbags they are highly skewed ( 57.6:8.86) to the results for `beltded`.

#### *Exercise 2*

Do an analysis that accounts, additionally, for estimated force of impact (`dvcat`):

```
ASdvtab <- xtabs(weight ~ dead + seatbelt + airbag + dvcat,
                data=nassCDS)
round(prop.table(ASdvtab, margin=2:4) ["dead", , , ], 6)
## Alternative: compact, flattened version of the table
round(ftable(prop.table(ASdvtab, margin=2:4) ["dead", , , ]), 6)
```

It will be apparent that differences between `none` and `airbag` are now below any reasonable threshold of statistical detectability.

#### *Exercise 3*

The package *DAAGxtras* includes the function `excessRisk()`. Run it with the default arguments, i.e. type

```
> excessRisk()
```

Compare the output with that obtained in Exercise 2 when the classification was a/c seatbelt (and airbag), and check that the output agrees.

Now do the following calculations, in turn:

- Classify according to `dvcat` as well as `seatbelt`. All you need do is add `dvcat` to the first argument to `excessRisk()`. What is now the total number of excess deaths?  
[The categories are 0-9 kph, 10-24 kph, 25-39 kph, 40-54 kph, and 55+ kph]
- Classify according to `dvcat`, `seatbelt` and `frontal`, and repeat the calculations. What is now the total number of excess deaths?

Explain the dependence of the estimates of numbers of excess deaths on the choice of factors for the classification.

**Note:** ? argues that these data, tabulated as above, have too many uncertainties and potential sources of bias to give reliable results. He presents a different analysis, based on the use of front seat passenger mortality as a standard against which to compare driver mortality, and limited to cars without passenger airbags. In the absence of any effect from airbags, the ratio of driver mortality to passenger mortality should be the same, irrespective of whether or not there was a driver airbag. In fact the ratio of driver fatalities to passenger fatalities was 11% lower in the cars with driver airbags.

## 2 Weighting Effects – Example with a Continuous Outcome

#### *Exercise 4*

Table 2, shows data from the data frame `gaba` (*DAAGxtras*). For background, see the Gordon (1995) paper that is referenced on the help page for `gaba`. [Image files that hold the functions `plotGaba()` and `compareGaba()` are in the subdirectory <http://www.maths.anu.edu.au/~johnm/r/functions/> ]

	min	mbac	mpl	fbac	fpl
2	10	1.76	1.76	2.18	2.55
3	30	1.31	1.65	3.48	4.15
4	50	0.05	0.67	3.13	3.66
5	70	-0.57	-0.25	3.03	2.05
6	90	-1.26	-0.50	2.08	0.61
7	110	-2.15	-2.22	1.60	0.34
8	130	-1.65	-2.18	1.38	0.67
9	150	-1.68	-2.86	1.76	0.76
10	170	-1.68	-3.23	1.06	0.39

Table 2: Data (average VAS pain scores) are from a trial that investigated the effect of pentazocine on post-operative pain, with (mbac and fbac) and without (mpl and fpl) preoperatively administered baclofen. Data are in the data frame `gaba` (*DAAGxtras* package). Numbers of males and females on the two treatments were:

	baclofen	placebo
females	15	7
males	3	16

*Exercise 4, continued*

- What do you notice about the relative numbers on the two treatments?
- For each treatment, obtain overall weighted averages at each time point, using the numbers in Table 2 as weights. (These are the numbers you would get if you divided the total over all patients on that treatment by the total number of patients.) This will give columns `avbac` and `avplac` that can be added to the data frame.
- Plot `avbac` and `avplac` against time, on the same graph. On separate graphs, repeat the comparisons (a) for females alone and (b) for males alone. Which of these graphs make a correct and relevant comparison between baclofen and placebo (albeit both in the presence of pentazocine)?

### 3 Extraction of nassCDS

Here are details of the code used to extract these data.

```
nassCDS <- nass9702cor[,c("dvcat", "national", "dead", "airbag", "seatbelt",
                        "frontal", "male", "age.of.o", "yearacc")]
nassCDS$dead <- 2-nass_cds$dead # Ensures 0 = survive; 1 = dead
## Now make dead a factor
nassCDS$dead <- factor(c("alive", "dead")[nassCDS$dead+1])
names(nassCDS)[8] <- "ageOfocc"
names(nassCDS)[2] <- "weight"
table(nassCDS$seatbelt) # Check the values of seatbelt, & their order
## Now make seatbelt a factor. The first value, here 0, becomes "none"
## The second value, here 1, becomes "belted"
nassCDS$seatbelt <- factor(nassCDS$seatbelt, labels=c("none","belted"))
# NB labels (unlike levels) matches only the order, not the values
nassCDS$airbag <- factor(nassCDS$airbag, labels=c("none","airbag"))
```