A (very) short introduction to R

Paul Torfs & Claudia Brauer

Hydrology and Quantitative Water Management Group

Wageningen University, The Netherlands

3 March 2014

1 Introduction

R is a powerful language and environment for statistical computing and graphics. It is a public domain (a so called "GNU") project which is similar to the commercial S language and environment which was developed at Bell Laboratories (formerly AT&T, now Lucent Technologies) by John Chambers and colleagues. R can be considered as a different implementation of S, and is much used in as an educational language and research tool.

The main advantages of R are the fact that R is freeware and that there is a lot of help available online. It is quite similar to other programming packages such as MatLab (not freeware), but more user-friendly than programming languages such as C++ or Fortran. You can use R as it is, but for educational purposes we prefer to use R in combination with the RStudio interface (also freeware), which has an organized layout and several extra options.

This document contains explanations, examples and exercises, which can also be understood (hopefully) by people without any programming experience. Going through all text and exercises takes about 1 or 2 hours. Examples of frequently used commands and error messages are listed on the last two pages of this document and can be used as a reference while programming.

2 Getting started

2.1 Install R

To install R on your computer (legally for free!), go to the home website of R[*]:

```
http://www.r-project.org/
```

and do the following (assuming you work on a windows computer):
- click `download CRAN` in the left bar
- choose a download site
- choose `Windows` as target operation system
- click `base`
- choose `Download R 3.0.3 for Windows` [†] and choose default answers for all questions

It is also possible to run R and RStudio from a USB stick instead of installing them. This could be useful when you don't have administrator rights on your computer. See our separate note "How to use portable versions of R and RStudio" for help on this topic.

2.2 Install RStudio

After finishing this setup, you should see an "R" icon on you desktop. Clicking on this would start up the standard interface. We recommend, however, to use the RStudio interface. [‡] To install RStudio, go to:

```
http://www.rstudio.org/
```

and do the following (assuming you work on a windows computer):
- click `Download RStudio`
- click `Download RStudio Desktop`
- click `Recommended For Your System`
- download the `.exe` file and run it (choose default answers for all questions)

2.3 RStudio layout

The RStudio interface consists of several windows (see Figure 1).

- Bottom left: **console window** (also called **command window**). Here you can type simple commands after the ">" prompt and R will then execute your command. This is the most important window, because this is where R actually does stuff.

- Top left: **editor window** (also called **script window**). Collections of commands (scripts) can be edited and saved. When you don't get

---

[†]At the moment of writing 3.0.3 was the latest version. Choose the most recent one.

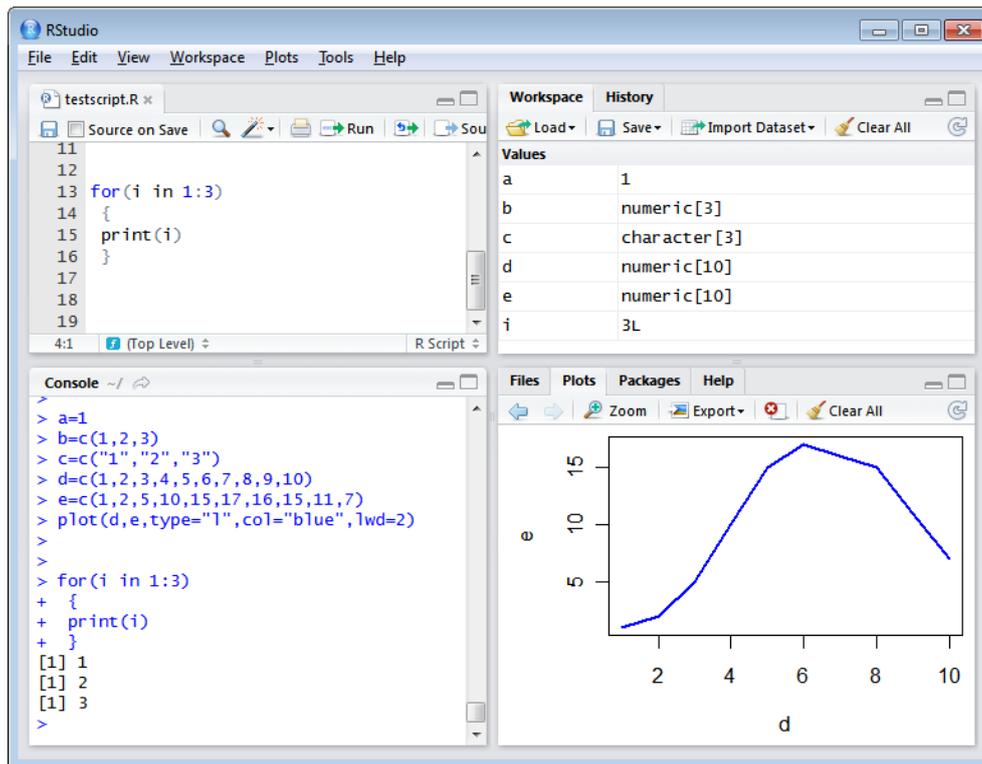[‡]There are many other (freeware) interfaces, such as Tinn-R.

Figure 1 The editor, workspace, console and plots windows in RStudio.

this window, you can open it with `File →`
`New → R script`

Just typing a command in the editor window
is not enough, it has to get into the command
window before R executes the command. If
you want to run a line from the script window
(or the whole script), you can click `Run` or
press `CTRL+ENTER` to send it to the command
window.

- Top right: **workspace / history window**.
  In the workspace window you can see which
  data and values R has in its memory. You
  can view and edit the values by clicking on
  them. The history window shows what has
  been typed before.

- Bottom right: **files / plots / packages /
  help window**. Here you can open files, view
  plots (also previous plots), install and load
  packages or use the help function.

You can change the size of the windows by drag-
ging the grey bars between the windows.

2.4 Working directory

Your *working directory* is the folder on your com-
puter in which you are currently working. When
you ask R to open a certain file, it will look in the
working directory for this file, and when you tell
R to save a data file or figure, it will save it in the
working directory.

Before you start working, please set your work-
ing directory to where all your data and script files
are or should be stored.

Type in the command window:
`setwd("directoryname")`. For example:

```
> setwd("M:/Hydrology/R/")
```

Make sure that the slashes are forward slashes and
that you don't forget the apostrophes (for the rea-
son of the apostrophes, see section 10.1). R is case
sensitive, so make sure you write capitals where
necessary.

Within RStudio you can also go to `Tools / Set
working directory`.

2.5 Libraries

R can do many statistical and data analyses. They
are organized in so-called *packages* or *libraries*.
With the standard installation, most common
packages are installed.

To get a list of all installed packages, go to the
packages window or type `library()` in the console

window. If the box in front of the package name is ticked, the package is loaded (activated) and can be used.

There are many more packages available on the R website. If you want to install and use a package (for example, the package called "geometry") you should:

• Install the package: click `install packages` in the packages window and type `geometry` or type `install.packages("geometry")` in the command window.

• Load the package: check box in front of `geometry` or type `library("geometry")` in the command window.

3 Some first examples of R commands

3.1 Calculator

R can be used as a calculator. You can just type your equation in the command window after the ">":

```
> 10^2 + 36
```

and R will give the answer

```
[1] 136
```

ToDo

Compute the difference between 2014 and the year you started at this university and divide this by the difference between 2014 and the year you were born. Multiply this with 100 to get the percentage of your life you have spent at this university. Use brackets if you need them.

If you use brackets and forget to add the closing bracket, the ">" on the command line changes into a "+". The "+" can also mean that R is still busy with some heavy computation. If you want R to quit what it was doing and give back the ">", press `ESC` (see the reference list on the last page).

3.2 Workspace

You can also give numbers a name. By doing so, they become so-called variables which can be used later. For example, you can type in the command window:

```
> a = 4
```

You can see that `a` appears in the workspace window, which means that R now remembers what `a` is.[§] You can also ask R what `a` is (just type `a` `ENTER` in the command window):

```
> a
[1] 4
```

or do calculations with `a`:

```
> a * 5
[1] 20
```

If you specify `a` again, it will forget what value it had before. You can also assign a new value to `a` using the old one.

```
> a = a + 10
> a
[1] 14
```

To remove all variables from R's memory, type

```
> rm(list=ls())
```

or click "clear all" in the workspace window. You can see that RStudio then empties the workspace window. If you only want to remove the variable `a`, you can type `rm(a)`.

ToDo

Repeat the previous ToDo, but with several steps in between. You can give the variables any name you want, but the name has to start with a letter.

3.3 Scalars, vectors and matrices

Like in many other programs, R organizes numbers in *scalars* (a single number – 0-dimensional), *vectors* (a row of numbers, also called arrays – 1-dimensional) and *matrices* (like a table – 2-dimensional).

The `a` you defined before was a scalar. To define a vector with the numbers 3, 4 and 5, you need the function[¶] `c`, which is short for concatenate (paste together).

```
b=c(3,4,5)
```

Matrices and other 2-dimensional structures will be introduced in Section 6.

[§]Some people prefer te use `<-` instead of `=` (they do the same thing). `<-` consists of two characters, `<` and `-`, and represents an arrow pointing at the object receiving the value of the expression.

[¶]See next Section for the explanation of functions.

3.4 Functions

If you would like to compute the mean of all the elements in the vector `b` from the example above, you could type

```
> (3+4+5)/3
```

But when the vector is very long, this is very boring and time-consuming work. This is why things you do often are automated in so-called *functions*. Some functions are standard in R or in one of the packages. You can also program your own functions (Section 11.3). When you use a function to compute a mean, you'll type:

```
> mean(x=b)
```

Within the brackets you specify the *arguments*. Arguments give extra information to the function. In this case, the argument `x` says of which set of numbers (vector) the mean should computed (namely of `b`). Sometimes, the name of the argument is not necessary: `mean(b)` works as well.

ToDo

Compute the sum of 4, 5, 8 and 11 by first combining them into a vector and then using the function `sum`.

The function `rnorm`, as another example, is a standard R function which creates random samples from a normal distribution. Hit the `ENTER` key and you will see 10 random numbers as:

```
1 > rnorm(10)
2 [1] -0.949  1.342 -0.474 0.403
3 [5] -0.091 -0.379  1.015  0.740
4 [9] -0.639  0.950
```

• Line 1 contains the command: `rnorm` is the function and the 10 is an argument specifying how many random numbers you want — in this case 10 numbers (typing `n=10` instead of just `10` would also work).
• Lines 2-4 contain the results: 10 random numbers organised in a vector with length 10.

Entering the same command again produces 10 new random numbers. Instead of typing the same text again, you can also press the upward arrow key ($\uparrow$) to access previous commands. If you want 10 random numbers out of normal distribution with mean 1.2 and standard deviation 3.4 you can type

```
> rnorm(10, mean=1.2, sd=3.4)
```

showing that the same function (`rnorm`) may have different interfaces and that R has so called *named arguments* (in this case `mean` and `sd`). By the way, the spaces around the "," and "=" do not matter.

Comparing this example to the previous one also shows that for the function `rnorm` only the first argument (the number 10) is compulsory, and that R gives default values to the other so-called optional arguments.[‖]

RStudio has a nice feature: when you type `rnorm(` in the command window and press `TAB`, RStudio will show the possible arguments (Fig. 2).

3.5 Plots

R can make graphs. The following is a very simple [**] example:

```
1 > x = rnorm(100)
2 > plot(x)
```

• In the first line, 100 random numbers are assigned to the variable `x`, which becomes a vector by this operation.
• In the second line, all these values are plotted in the plots window.

ToDo

Plot 100 normal random numbers.

4 Help and documentation

There is a large amount of (free) documentation and help available. Some help is automatically installed. Typing in the console window the command

```
> help(rnorm)
```

gives help on the `rnorm` function. It gives a description of the function, possible arguments and the values that are used as default for optional arguments. Typing

```
> example(rnorm)
```

[‖]Use the help function (Sect. 4) to see which values are used as default.

[**]See Section 7 for slightly less trivial examples.

Figure 2 RStudio shows possible arguments when you press `TAB` after the function name and bracket.

gives some examples of how the function can be used.

An HTML-based global help can be called with:

```
> help.start()
```

or by going to the help window.

The following links can also be very useful:

• `http://cran.r-project.org/doc/manuals/R-intro.pdf` A full manual.

• `http://cran.r-project.org/doc/contrib/Short-refcard.pdf` A short reference card.

• `http://zoonek2.free.fr/UNIX/48_R/all.html`
A very rich source of examples.

• `http://rwiki.sciviews.org/doku.php`
A typical user wiki.

• `http://www.statmethods.net/`
Also called Quick-R. Gives very productive direct help. Also for users coming from other programming languages.

• `http://mathesaurus.sourceforge.net/`
Dictionary for programming languages (e.g. R for Matlab users).

• Just using Google (type e.g. "R rnorm" in the search field) can also be very productive.

ToDo

Find help for the `sqrt` function.

5 Scripts

R is an interpreter that uses a command line based environment. This means that you have to type commands, rather than use the mouse and menus. This has the advantage that you do not always have to retype all commands and are less likely to get complaints of arms, neck and shoulders.

You can store your commands in files, the so-called *scripts*. These scripts have typically file names with the extension .R, e.g. `foo.R`. You can open an editor window to edit these files by clicking `File` and `New` or `Open file...` [††].

You can run (send to the console window) part of the code by selecting lines and pressing `CTRL+ENTER` or click `Run` in the editor window. If you do not select anything, R will run the line your cursor is on. You can always run the whole script with the console command `source`, so e.g. for the script in the file `foo.R` you type:

```
> source("foo.R")
```

You can also click `Run all` in the editor window or type `CTRL+SHIFT + S` to run the whole script at once.

ToDo

Make a file called `firstscript.R` containing R-code that generates 100 random numbers and plots them, and run this script several times.

6 Data structures

If you are unfamiliar with R, it makes sense to just retype the commands listed in this section. Maybe you will not need all these structures in the beginning, but it is always good to have at least a first glimpse of the terminology and possible applications.

6.1 Vectors

Vectors were already introduced, but they can do more:

[††]Where also the options `Save` and `Save as` are available.

```
1 > vec1 = c(1,4,6,8,10)
2 > vec1
3 [1]  1  4  6  8 10
4 > vec1[5]
5 [1] 10
6 > vec1[3] = 12
7 > vec1
8 [1]  1  4 12  8 10
9 > vec2 = seq(from=0, to=1, by=0.25)
10 > vec2
11 [1] 0.00 0.25 0.50 0.75 1.00
12 > sum(vec1)
13 [1] 35
14 > vec1 + vec2
15 [1]  1.00 4.25 12.50 8.75 11.00
```

- In line 1, a vector `vec1` is explicitly constructed by the concatenation function `c()`, which was introduced before. Elements in vectors can be addressed by standard `[i]` indexing, as shown in lines 4-5.
- In line 6, one of the elements is replaced with a new number. The result is shown in line 8.
- Line 9 demonstrates another useful way of constructing a vector: the `seq()` (sequence) function.
- Lines 10-15 show some typical vector oriented calculations. If you add up two vectors of the same length, the first elements of both vectors are summed, and the second elements, etc., leading to a new vector of length 5 (just like in regular vector calculus). Note that the function `sum` sums up the elements within a vector, leading to one number (a scalar).

6.2 Matrices

Matrices are nothing more than 2-dimensional vectors. To define a matrix, use the function `matrix`:

```
1 mat=matrix(data=c(9,2,3,4,5,6),ncol=3)
2 > mat
3      [,1] [,2] [,3]
4 [1,]    9    3    5
5 [2,]    2    4    6
```

The argument `data` specifies which numbers should be in the matrix. Use either `ncol` to specify the number of columns or `nrow` to specify the number of rows.

ToDo

Put the numbers 31 to 60 in a vector named `P` and in a matrix with 6 rows and 5 columns named `Q`. Tip: use the function `seq`. Look at the different ways scalars, vectors and matrices are denoted in the workspace window.

Matrix-operations are similar to vector operations:

```
1 > mat[1,2]
2 [1] 3
3 > mat[2,]
4 [1] 2 4 6
5 > mean(mat)
6 [1] 4.8333
```

- Elements of a matrix can be addressed in the usual way: `[row,column]` (line 1).
- Line 3: When you want to select a whole row, you leave the spot for the column number empty (the other way around for columns of course).
- Line 5 shows that many functions also work with matrices as argument.

6.3 Data frames

Time series are often ordered in *data frames*. A data frame is a matrix with names above the columns. This is nice, because you can call and use one of the columns without knowing in which position it is.

```
1 > t = data.frame(x = c(11,12,14),
2  y = c(19,20,21), z = c(10,9,7))
3 > t
4    x  y  z
5 1 11 19 10
6 2 12 20  9
7 3 14 21  7
8 > mean(t$z)
9 [1] 8.666667
10 > mean(t[["z"]])
11 [1] 8.666667
```

- In lines 1-2 a typical data frame called `t` is constructed. The columns have the names `x`, `y` and `z`.
- Line 8-11 show two ways of how you can select the column called `z` from the data frame called `t`.

6.4 Lists

Another basic structure in R is a *list*. The main advantage of lists is that the "columns" (they're not really ordered in columns any more, but are more a collection of vectors) don't have to be of the same length, unlike matrices and data frames.

```
> L = list(one=1, two=c(1,2),
 five=seq(0, 1, length=5))
> L
$one
[1] 1
$two
[1] 1 2
$five
[1] 0.00 0.25 0.50 0.75 1.00
> names(L)
[1] "one"  "two"  "five"
> L$five + 10
[1] 10.00 10.25 10.50 10.75 11.00
```

- Lines 1-2 construct a list by giving names and values. The list also appears in the workspace window.
- Lines 3-9 show a typical printing (after pressing `L ENTER`).
- Line 10 illustrates how to find out what's in the list.
- Line 12 shows how to use the numbers.

7 Graphics

Plotting is an important statistical activity. So it should not come as a surprise that R has many plotting facilities. The following lines show a simple plot:

```
> plot(rnorm(100), type="l", col="gold")
```

Hundred random numbers are plotted by connecting the points by lines (the symbol between quotes after the `type=`, is the letter l, not the number 1) in a gold color.

Another very simple example is the classical statistical histogram plot, generated by the simple command

```
> hist(rnorm(100))
```
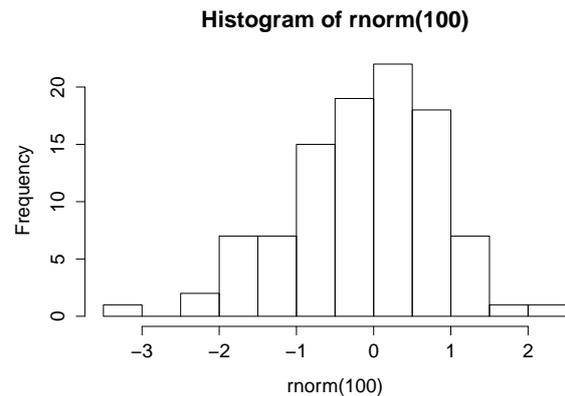
which generates the plot in Figure 3.

Histogram of rnorm(100)

Figure 3 A simple histogram plot.

The following few lines create a plot using the data frame `t` constructed in the previous ToDo:

```
plot(t$a, type="l", ylim=range(t),
 lwd=3, col=rgb(1,0,0,0.3))
lines(t$b, type="s", lwd=2,
 col=rgb(0.3,0.4,0.3,0.9))
points(t$c, pch=20, cex=4,
 col=rgb(0,0,1,0.3))
```

To learn more about formatting plots, search for `par` in the R help. Google "R color chart" for a pdf file with a wealth of color options.

To copy your plot to a document, go to the plots window, click the "Export" button, choose the nicest width and height and click `Copy` or `Save`.

8 Reading and writing data files

There are many ways to write data from within the R environment to files, and to read data from files. We will illustrate one way here. The following lines illustrate the essential:

```
1 > d = data.frame(a = c(3,4,5),
2  b = c(12,43,54))
3 > d
4   a  b
5 1 3 12
6 2 4 43
7 3 5 54
8 > write.table(d, file="tst0.txt",
9  row.names=FALSE)
10 > d2 = read.table(file="tst0.txt",
11  header=TRUE)
12 > d2
13   a  b
14 1 3 12
15 2 4 43
16 3 5 54
```

• In lines 1-2, a simple example data frame is constructed and stored in the variable `d`.

• Lines 3-7 show the content of this data frame: two columns (called `a` and `b`), each containing three numbers.

• Line 8 writes this data frame to a text file, called `tst0.txt` The argument `row.names=FALSE` prevents that row names are written to the file. Because nothing is specified about `col.names`, the default option `col.names=TRUE` is chosen and column names are written to the file. Figure 4 shows the resulting file (opened in an editor, such as Notepad), with the column names (`a` and `b`) in the first line.

• Lines 10-11 illustrate how to read a file into a data frame. Note that the column names are also read. The data frame also appears in the workspace window.

ToDo

Make a file called `tst1.txt` in Notepad from the example in Figure 4 and store it in your working directory. Write a script to read it, to multiply the column called `g` by 5 and to store it as `tst2.txt`.
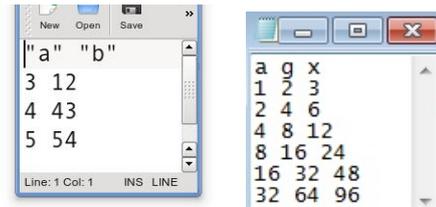

Figure 4 The files `tst0.txt` of section 8 (left) and `tst1.txt` from the ToDo below (right) opened in two text editors.

9 Not available data

ToDo

Compute the mean of the square root of a vector of 100 random numbers. What happens?

When you work with real data, you will encounter missing values because instrumentation failed or because you didn't want to measure in the weekend. When a data point is *not available*, you write `NA` instead of a number.

```
> j = c(1,2,NA)
```

Computing statistics of incomplete data sets is strictly speaking not possible. Maybe the largest value occurred during the weekend when you didn't measure. Therefore, R will say that it doesn't know what the largest value of `j` is:

```
> max(j)
[1] NA
```

If you don't mind about the missing data and want to compute the statistics anyway, you can add the argument `na.rm=TRUE` (Should I remove the NAs? Yes!).

```
> max(j, na.rm=TRUE)
[1] 2
```

10 Classes

The exercises you did before were nearly all with numbers. Sometimes you want to specify something which is not a number, for example the name of a measurement station or data file. In that case you want the variable to be a character string instead of a number.

An object in R can have several so-called *classes*. The most important three are *numeric*, *character* and *POSIX* (date-time combinations). You can ask R what class a certain variable is by typing `class(...)`.

10.1 Characters

To tell R that something is a character string, you should type the text between apostrophes, otherwise R will start looking for a defined variable with the same name:

```
> m = "apples"
> m
[1] "apples"
> n = pears
Error: object 'pears' not found
```

Of course, you cannot do computations with character strings:

```
> m + 2
Error in m + 2 : non-numeric argument to
binary operator
```

10.2 Dates

Dates and times are complicated. R has to know that 3 o'clock comes after 2:59 and that February has 29 days in some years. The easiest way to tell R that something is a date-time combination is with the function `strptime`:

```
1 > date1=strptime( c("20100225230000",
2 "20100226000000", "20100226010000"),
3 format="%Y%m%d%H%M%S")
4 > date1
5 [1] "2010-02-25 23:00:00"
6 [2] "2010-02-26 00:00:00"
7 [3] "2010-02-26 01:00:00"
```

• In lines 1-2 you create a vector with `c(...)`. The numbers in the vectors are between apostrophes because the function `strptime` needs character strings as input.
• In line 3 the argument `format` specifies how the character string should be read. In this case the year is denoted first (%Y), then the month (%m), day (%d), hour (%H), minute (%M) and second (%S). You don't have to specify all of them, as long as the format corresponds to the character string.

11 Programming tools

When you are building a larger program than in the examples above or if you're using someone else's scripts, you may encounter some programming statements. In this Section we describe a few tips and tricks.

11.1 If-statement

The *if-statement* is used when certain computations should *only* be done when a certain condition is met (and maybe something else should be done when the condition is not met). An example:

```
1 > w = 3
2 > if( w < 5 )
3    {
4    d=2
5    }else{
6    d=10
7    }
8 > d
9 2
```

• In line 2 a condition is specified: `w` should be less than 5.
• If the condition is met, R will execute what is between the first brackets in line 4.
• If the condition is *not* met, R will execute what is between the second brackets, after the `else` in line 6. You can leave the `else{...}`-part out if you don't need it.
• In this case, the condition is met and `d` has been assigned the value 2 (lines 8-9).

To get a subset of points in a vector for which a certain condition holds, you can use a shorter method:

```
1 > a = c(1,2,3,4)
2 > b = c(5,6,7,8)
3 > f = a[b==5 | b==8]
4 > f
5 [1] 1 4
```

● In line 1 and 2 two vectors are made.

● In line 3 you say that `f` is composed of those elements of vector `a` for which `b` equals 5 or `b` equals 8.

Note the double `=` in the condition. Other conditions (also called logical or Boolean operators) are `<`, `>`, `!=` ($\neq$), `<=` ($\leq$) and `>=` ($\geq$). To test more than one condition in one if-statement, use `&` if both conditions have to be met ("and") and `|` if one of the conditions has to be met ("or").

11.2 For-loop

If you want to model a time series, you usually do the computations for one time step and then for the next and the next, etc. Because nobody wants to type the same commands over and over again, these computations are automated in for-loops.

In a *for-loop* you specify what has to be done and how many times. To tell "how many times", you specify a so-called counter. An example:

```
> h = seq(from=1, to=8)
> s = c()
> for(i in 2:10)
    {
    s[i] = h[i] * 10
    }
> s
[1] NA 20 30 40 50 60 70 80 NA NA
```

● First the vector `h` is made.

● In line 2 an empty vector (`s`) is created. This is necessary because when you introduce a variable within the for-loop, R will not remember it when it has gotten out of the for-loop.

● In line 3 the for-loop starts. In this case, `i` is the counter and runs from 2 to 10.

● Everything between the curly brackets (line 5) is processed 9 times. The first time `i=2`, the second element of `h` is multiplied with 10 and placed in the second position of the vector `s`. The second time `i=3`, etc. In the last two runs, the 9th and 10th elements of `h` are requested, which do not exist. Note that these statements are evaluated without any explicit error messages.

ToDo

Make a vector from 1 to 100. Make a for-loop which runs through the whole vector. Multiply the elements which are smaller than 5 and larger than 90 with 10 and the other elements with 0.1.

11.3 Writing your own functions

Functions you program yourself work in the same way as pre-programmed R functions.

```
> fun1 = function(arg1, arg2 )
    {
    w = arg1 ^ 2
    return(arg2 + w)
    }
> fun1(arg1 = 3, arg2 = 5)
[1] 14

```

● In line 1 the function name (`fun1`) and its arguments (`arg1` and `arg2`) are defined.

● Lines 2-5 specify what the function should do if it is called. The return value (`arg2+w`) is shown on the screen.

● In line 6 the function is called with arguments 3 and 5.

ToDo

Write a function for the previous ToDo, so that you can feed it any vector you like (as argument). Use a for-loop in the function to do the computation with each element. Use the standard R function `length` in the specification of the counter. [a])

---

[a] Actually, people often use more for-loops than necessary. The ToDo above can be done more easily and quickly without a for-loop but with regular vector-computations.

12 Some useful references

12.1 Functions

This is a subset of the functions explained in the R reference card.

Data creation
- `read.table`: read a table from file. Arguments: `header=TRUE`: read first line as titles of the columns; `sep=","`: numbers are separated by commas; `skip=n`: don't read the first `n` lines.
- `write.table`: write a table to file
- `c`: paste numbers together to create a vector
- `array`: create a vector, Arguments: `dim`: length
- `matrix`: create a matrix, Arguments: `ncol` and/or `nrow`: number of rows/columns
- `data.frame`: create a data frame
- `list`: create a list
- `rbind` and `cbind`: combine vectors into a matrix by row or column

Extracting data
- `x[n]`: the n^{th} element of a vector
- `x[m:n]`: the m^{th} to n^{th} element
- `x[c(k,m,n)]`: specific elements
- `x[x>m & x<n]`: elements between `m` and `n`
- `x$n`: element of list or data frame named `n`
- `x[["n"]]`: idem
- `[i,j]`: element at i^{th} row and j^{th} column
- `[i,]`: row `i` in a matrix

Information on variables
- `length`: length of a vector
- `ncol` or `nrow`: number of columns or rows in a matrix
- `class`: class of a variable
- `names`: names of objects in a list
- `print`: show variable or character string on the screen (used in scripts or for-loops)
- `return`: show variable on the screen (used in functions)
- `is.na`: test if variable is `NA`
- `as.numeric` or `as.character`: change class to number or character string
- `strptime`: change class from character to date-time (POSIX)

Statistics
- `sum`: sum of a vector (or matrix)
- `mean`: mean of a vector
- `sd`: standard deviation of a vector

- `max` or `min`: largest or smallest element
- `rowSums` (or `rowMeans`, `colSums` and `colMeans`): sums (or means) of all numbers in each row (or column) of a matrix. The result is a vector.
- `quantile(x,c(0.1,0.5))`: sample the 0.1 and 0.5^{th} quantiles of vector `x`

Data processing
- `seq`: create a vector with equal steps between the numbers
- `rnorm`: create a vector with random numbers with normal distribution (other distributions are also available)
- `sort`: sort elements in increasing order
- `t`: transpose a matrix
- `aggregate(x,by=ls(y),FUN="mean")`: split data set `x` into subsets (defined by `y`) and computes means of the subsets. Result: a new list.
- `na.approx`: interpolate (in `zoo` package). Argument: vector with `NA`s. Result: vector without `NA`s.
- `cumsum`: cumulative sum. Result is a vector.
- `rollmean`: moving average (in the `zoo` package)
- `paste`: paste character strings together
- `substr`: extract part of a character string

Fitting
- `lm(v1~v2)`: linear fit (regression line) between vector `v1` on the y-axis and `v2` on the x-axis
- `nls(v1~a+b*v2, start=ls(a=1,b=0))`: non-linear fit. Should contain equation with variables (here `v1` and `v2` and parameters (here `a` and `b`) with starting values
- `coef`: returns coefficients from a fit
- `summary`: returns all results from a fit

Plotting
- `plot(x)`: plot `x` (y-axis) versus index number (x-axis) in a new window
- `plot(x,y)`: plot `y` (y-axis) versus `x` (x-axis) in a new window
- `image(x,y,z)`: plot `z` (color scale) versus `x` (x-axis) and `y` (y-axis) in a new window
- `lines` or `points`: add lines or points to a previous plot
- `hist`: plot histogram of the numbers in a vector
- `barplot`: bar plot of vector or data frame
- `contour(x,y,z)`: contour plot
- `abline`: draw line (segment). Arguments: `a,b` for intercept `a` and slope `b`; or `h=y` for horizontal line at `y`; or `v=x` for vertical line at `x`.
- `curve`: add function to plot. Needs to have an

x in the expression. Example: `curve(x^2)`
• `legend`: add legend with given symbols (`lty` or `pch` and `col`) and text (`legend`) at location (`x="topright"`)
• `axis`: add axis. Arguments: `side` – 1=bottom, 2=left, 3=top, 4=right
• `mtext`: add text on axis. Arguments: `text` (character string) and `side`
• `grid`: add grid
• `par`: plotting parameters to be specified before the plots. Arguments: e.g. `mfrow=c(1,3))`: number of figures per page (1 row, 3 columns); `new=TRUE`: draw plot over previous plot.

Plotting parameters
These can be added as arguments to `plot`, `lines`, `image`, etc. For help see `par`.
• `type`: `"l"`=lines, `"p"`=points, etc.
• `col`: color – `"blue"`, `"red"`, etc
• `lty`: line type – 1=solid, 2=dashed, etc.
• `pch`: point type – 1=circle, 2=triangle, etc.
• `main`: title - character string
• `xlab` and `ylab`: axis labels – character string
• `xlim` and `ylim`: range of axes – e.g. c(1,10)
• `log`: logarithmic axis – `"x"`, `"y"` or `"xy"`

Programming
• `function(arglist){expr}`: function definition: do `expr` with list of arguments `arglist`
• `if(cond){expr1}else{expr2}`: if-statement: if `cond` is true, then `expr1`, else `expr2`
• `for(var in vec) {expr}`: for-loop: the counter `var` runs through the vector `vec` and does `expr` each run
• `while(cond){expr}`: while-loop: while `cond` is true, do `expr` each run

12.2 Keyboard shortcuts

There are several useful keyboard shortcuts for RStudio (see `Help` → `Keyboard Shortcuts`):
• `CRL+ENTER`: send commands from script window to command window
• ↑ or ↓ in command window: previous or next command
• `CTRL+1`, `CTRL+2`, etc.: change between the windows

Not R-specific, but very useful keyboard shortcuts:
• `CTRL+C`, `CTRL+X` and `CTRL+V`: copy, cut and paste
• `ALT+TAB`: change to another program window
• ↑, ↓, ← or →: move cursor
• `HOME` or `END`: move cursor to begin or end of line
• `Page Up` or `Page Down`: move cursor one page up or down
• `SHIFT+↑/↓/←/→/HOME/END/PgUp/PgDn`: select

12.3 Error messages

• `No such file or directory` or `Cannot change working directory`
Make sure the working directory and file names are correct.
• `Object 'x' not found`
The variable `x` has not been defined yet. Define `x` or write apostrophes if `x` should be a character string.
• `Argument 'x' is missing without default`
You didn't specify the compulsory argument `x`.
• `+`
R is still busy with something or you forgot closing brackets. Wait, type `}` or `)` or press `ESC`.
• `Unexpected ')' in ")"` or `Unexpected '}' in "}"`
The opposite of the previous. You try to close something which hasn't been opened yet. Add opening brackets.
• `Unexpected 'else' in "else"`
Put the `else` of an if-statement on the same line as the last bracket of the "then"-part: `}else{`.
• `Missing value where TRUE/FALSE needed`
Something goes wrong in the condition-part (`if(x==1)`) of an if-statement. Is `x` NA?
• `The condition has length > 1 and only the first element will be used`
In the condition-part (`if(x==1)`) of an if-statement, a vector is compared with a scalar. Is `x` a vector? Did you mean `x[i]`?
• `Non-numeric argument to binary operator`
You are trying to do computations with something which is not a number. Use `class(...)` to find out what went wrong or use `as.numeric(...)` to transform the variable to a number.
• `Argument is of length zero` or `Replacement is of length zero`
The variable in question is `NULL`, which means that it is empty, for example created by `c()`. Check the definition of the variable.