

Introduction to R

Network Analysis 2017

Sacha Epskamp

02-11-2017

Course information

- ▶ All information in syllabus on blackboard!
 - ▶ Tomorrow...
- ▶ Lectures every Monday and Thursday
 - ▶ Thursday will be more practical, but can still be a lecture
- ▶ Every week an assignment
 - ▶ Deadline **15:00** on Thursday!
 - ▶ Conceptual questions and practical questions
- ▶ Final Project (more info on Monday)
 - ▶ Paper and presentation in groups of 2. Topics may include:
 - ▶ Analyze your own dataset
 - ▶ Gather and analyze data
 - ▶ Review the literature (topics on blackboard)
 - ▶ Re-analyze an existing dataset

Individual Assignments

Each week, the assignment will be made available 15:00 on Thursday, and will be due 15:00 the next Thursday. Each assignment will contribute to 10% or 5% (first and last week) of your grade.

- ▶ Work on the assignments **alone**.
- ▶ Hand in a **PDF** file and an **.R** file.
- ▶ Make sure your PDF report is as standalone readable as possible. E.g., if you are asked to report a network, then include the plot in the PDF and not just say “look at .R file”.
- ▶ Assignments are due **before 15:00**. If you do not hand in an assignment before 15:00, you will get a 1.
- ▶ If you encounter any problems, or have any feedback, please let me know before the deadline, as then I can take it into account or help you.



What is R?

- ▶ R is a statistical programming language
 - ▶ Statistical analysis
 - ▶ Data visualization
 - ▶ Data mining
 - ▶ General programming
- ▶ It is *open-source*
 - ▶ Free as in “free beer” and “free speech”
 - ▶ Large active community around R
 - ▶ Many contributed packages

What do you need to know about R?

- ▶ In this course not much. We intend the methods we will describe to be usable by researchers with little experience in R
 - ▶ No programming
 - ▶ Cookbook method
- ▶ But it is strongly advised that you learn more about R if you want to do more in methodology or psychometrics.
- ▶ To install R and RStudio see links section on Blackboard

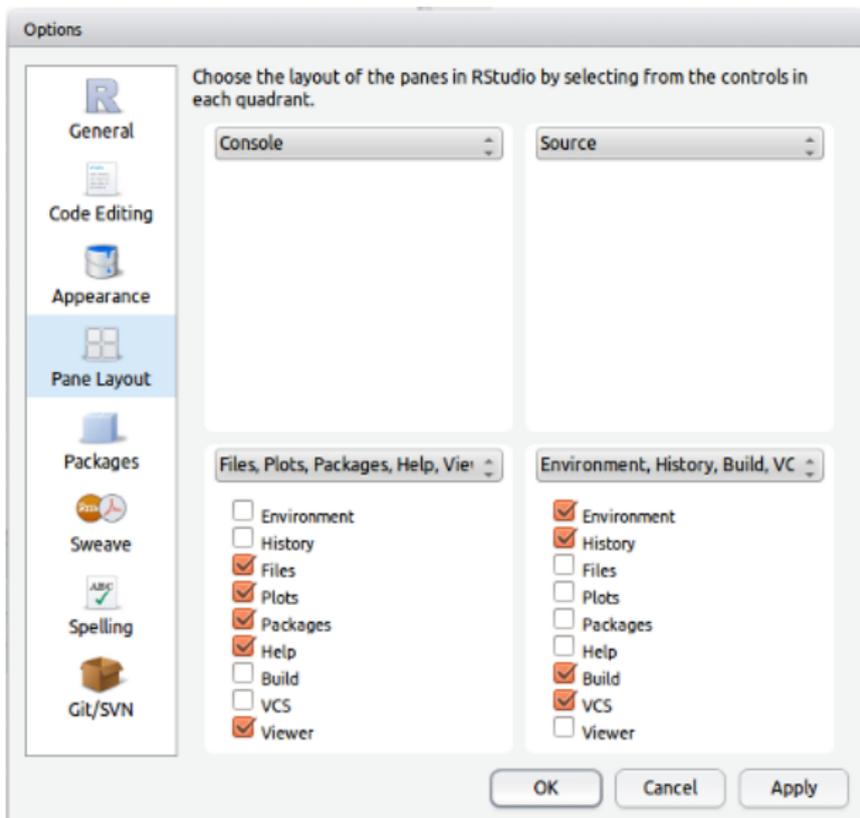
Installing R and RStudio

- ▶ R
 - ▶ This is the base programming language
 - ▶ Downloadable from <http://cran.r-project.org/>
 - ▶ It gives a program called R or RGui which we will **never** use
- ▶ RStudio
 - ▶ New and very good GUI for R
 - ▶ Downloadable from
<http://www.rstudio.com/products/rstudio/download/>
- ▶ Always use R from within RStudio!

First use of R

- ▶ We will use the environment RStudio for our work in R
- ▶ RStudio has 4 panes:
 - ▶ **Console** This is the actual R window, you can enter commands here and execute them by pressing enter
 - ▶ **Source** This is where we can edit *scripts*. It is where you should always be working. Control-enter sends selected codes to the console
 - ▶ **Plots/Help** This is where plots and help pages will be shown
 - ▶ **Workspace** Shows which objects you currently have
- ▶ Anything following a # symbol is treated as a comment!

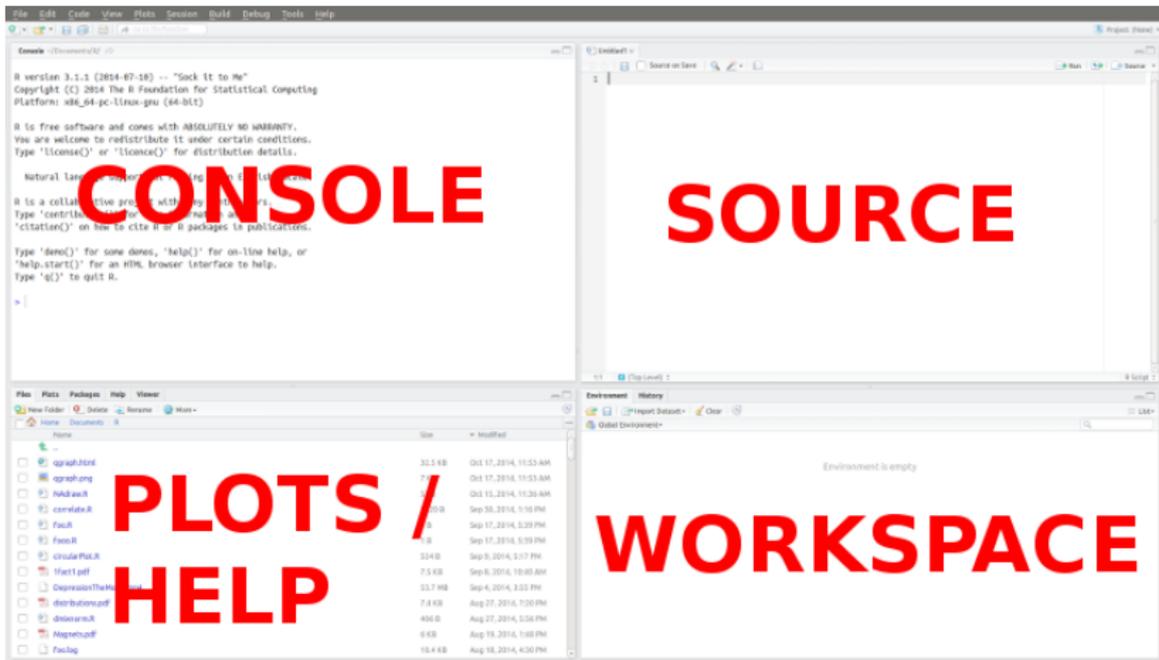
To change the layout of panes in RStudio, go to Tools -> Global Options -> Pane Layout. Make sure you set it like this:



The screenshot displays the RStudio IDE interface with the following components:

- Console:** Shows the R version 3.1.1 (2014-07-10) -- "Sock it to me" and copyright information. It also displays the R license and a list of help topics such as `license()`, `contributors()`, `citation()`, `demo()`, `help.start()`, and `q()`.
- Environment:** Shows a "Global Environment" which is currently empty.
- File Browser:** Lists files in the current directory, including R packages and PDF documents.

File Name	Size	Last Modified
igraph.html	32.5 KB	Oct 17, 2014, 11:53 AM
igraph.png	7 KB	Oct 17, 2014, 11:53 AM
igraph.aux	31 B	Oct 15, 2014, 11:36 AM
correlate.R	1020 B	Sep 30, 2014, 1:16 PM
Fault.R	1 B	Sep 17, 2014, 5:29 PM
Fault.R	1 B	Sep 17, 2014, 5:29 PM
circulaPlot.R	524 B	Sep 9, 2014, 5:17 PM
Ifact1.pdf	7.3 KB	Sep 8, 2014, 10:40 AM
DepressorTheMovie.mp4	53.7 MB	Sep 4, 2014, 3:25 PM
distributions.pdf	7.4 KB	Aug 27, 2014, 7:20 PM
dnlsurm.R	436 B	Aug 27, 2014, 5:54 PM
Magnets.pdf	6 KB	Aug 19, 2014, 1:48 PM
Faclog	10.4 KB	Aug 18, 2014, 4:20 PM



The R console

The image shows a screenshot of the RStudio application interface. The interface is divided into several panes:

- Console (top-left):** Displays the R version information (3.1.1 (2014-07-18)) and copyright notice. It also includes instructions on how to use the software, such as 'Type 'license()' or 'licence()' for distribution details.' and 'Type 'demo()' for some demos, 'help()' for on-line help, or 'help.start()' for an HTML browser interface to help. Type 'q()' to quit R.'
- Source (top-right):** A pane for editing R source code files.
- Plots / Help (bottom-left):** A pane for viewing plots and accessing help documentation. It shows a list of installed packages with columns for Name, Size, and Installed.
- Workspace (bottom-right):** A pane for managing the current R workspace, showing environment and history.

Four large red text labels are overlaid on the image:

- CONSOLE** (over the top-left pane)
- SOURCE** (over the top-right pane)
- PLOTS / HELP** (over the bottom-left pane)
- WORKSPACE** (over the bottom-right pane)

Name	Size	Installed
ape	302.7 KB	2014-11-20 16:10:00 AM
apeglm	7.1 MB	2014-11-20 16:10:00 AM
apeglm2	7.1 MB	2014-11-20 16:10:00 AM
apeglm3	7.1 MB	2014-11-20 16:10:00 AM
apeglm4	7.1 MB	2014-11-20 16:10:00 AM
apeglm5	7.1 MB	2014-11-20 16:10:00 AM
apeglm6	7.1 MB	2014-11-20 16:10:00 AM
apeglm7	7.1 MB	2014-11-20 16:10:00 AM
apeglm8	7.1 MB	2014-11-20 16:10:00 AM
apeglm9	7.1 MB	2014-11-20 16:10:00 AM
apeglm10	7.1 MB	2014-11-20 16:10:00 AM
apeglm11	7.1 MB	2014-11-20 16:10:00 AM
apeglm12	7.1 MB	2014-11-20 16:10:00 AM
apeglm13	7.1 MB	2014-11-20 16:10:00 AM
apeglm14	7.1 MB	2014-11-20 16:10:00 AM
apeglm15	7.1 MB	2014-11-20 16:10:00 AM
apeglm16	7.1 MB	2014-11-20 16:10:00 AM
apeglm17	7.1 MB	2014-11-20 16:10:00 AM
apeglm18	7.1 MB	2014-11-20 16:10:00 AM
apeglm19	7.1 MB	2014-11-20 16:10:00 AM
apeglm20	7.1 MB	2014-11-20 16:10:00 AM

The R console

- ▶ The console pane is where R operates
- ▶ In the console pane, we can type commands and press enter to have R evaluate them
 - ▶ To ask R the answer to $1 + 1$, we can type in `1 + 1` and press enter:

```
> 1 + 1
[1] 2
```

- ▶ The arrow-up key lets us go back to a previous command
- ▶ This is how R works, we write commands for R to evaluate
 - ▶ Commands can be as simple as `1 + 1` or much longer encoding an entire data analysis

In the sheets, we use gray boxes to demonstrate commands we send to R. The values that R returns are shown underneath the gray boxes followed by the ## symbols.

For example, the command:

```
> 1 + 1  
[1] 2
```

will be referred to in the sheets as:

```
1 + 1
```

```
## [1] 2
```

Comments in R

```
1 + 1 # Here I sum 1 and 1!
```

```
## [1] 2
```

- ▶ The # sign is used because R stops evaluating any command whenever it sees #
 - ▶ Everything after # is **not** an R command.
 - ▶ # can be used as comments to clarify your code
 - ▶ Do this as often as you can!

The source pane

The image shows a screenshot of the RStudio software interface. The interface is divided into several panes:

- Console (top-left):** Displays the R version (3.1.1) and copyright information. It contains the text: "It is free software and comes with ABSOLUTELY NO WARRANTY. We are welcome to redistribute it under certain conditions. Type 'license()' or 'licence()' for distribution details. Natural language processing with tm is a collaborative project with many contributors. Type 'contrib.authors()' or 'contrib.authors()' for a list of contributors. Type 'demo()' for some demos, 'help()' for on-line help, or 'help.start()' for an HTML browser interface to help. Type 'q()' to quit R." Below this text is a single line of code: `= |`.
- Source (top-right):** A large empty white area intended for editing R source code.
- Plots / Help (bottom-left):** A pane containing a list of installed R packages. The list includes: `ggplot2`, `ggplotify`, `ggplot2_2`, `ggplot2_3`, `ggplot2_4`, `ggplot2_5`, `ggplot2_6`, `ggplot2_7`, `ggplot2_8`, `ggplot2_9`, `ggplot2_10`, `ggplot2_11`, `ggplot2_12`, `ggplot2_13`, `ggplot2_14`, `ggplot2_15`, `ggplot2_16`, `ggplot2_17`, `ggplot2_18`, `ggplot2_19`, `ggplot2_20`, `ggplot2_21`, `ggplot2_22`, `ggplot2_23`, `ggplot2_24`, `ggplot2_25`, `ggplot2_26`, `ggplot2_27`, `ggplot2_28`, `ggplot2_29`, `ggplot2_30`, `ggplot2_31`, `ggplot2_32`, `ggplot2_33`, `ggplot2_34`, `ggplot2_35`, `ggplot2_36`, `ggplot2_37`, `ggplot2_38`, `ggplot2_39`, `ggplot2_40`, `ggplot2_41`, `ggplot2_42`, `ggplot2_43`, `ggplot2_44`, `ggplot2_45`, `ggplot2_46`, `ggplot2_47`, `ggplot2_48`, `ggplot2_49`, `ggplot2_50`, `ggplot2_51`, `ggplot2_52`, `ggplot2_53`, `ggplot2_54`, `ggplot2_55`, `ggplot2_56`, `ggplot2_57`, `ggplot2_58`, `ggplot2_59`, `ggplot2_60`, `ggplot2_61`, `ggplot2_62`, `ggplot2_63`, `ggplot2_64`, `ggplot2_65`, `ggplot2_66`, `ggplot2_67`, `ggplot2_68`, `ggplot2_69`, `ggplot2_70`, `ggplot2_71`, `ggplot2_72`, `ggplot2_73`, `ggplot2_74`, `ggplot2_75`, `ggplot2_76`, `ggplot2_77`, `ggplot2_78`, `ggplot2_79`, `ggplot2_80`, `ggplot2_81`, `ggplot2_82`, `ggplot2_83`, `ggplot2_84`, `ggplot2_85`, `ggplot2_86`, `ggplot2_87`, `ggplot2_88`, `ggplot2_89`, `ggplot2_90`, `ggplot2_91`, `ggplot2_92`, `ggplot2_93`, `ggplot2_94`, `ggplot2_95`, `ggplot2_96`, `ggplot2_97`, `ggplot2_98`, `ggplot2_99`, `ggplot2_100`.
- Workspace (bottom-right):** A pane showing the current workspace environment, which is currently empty.

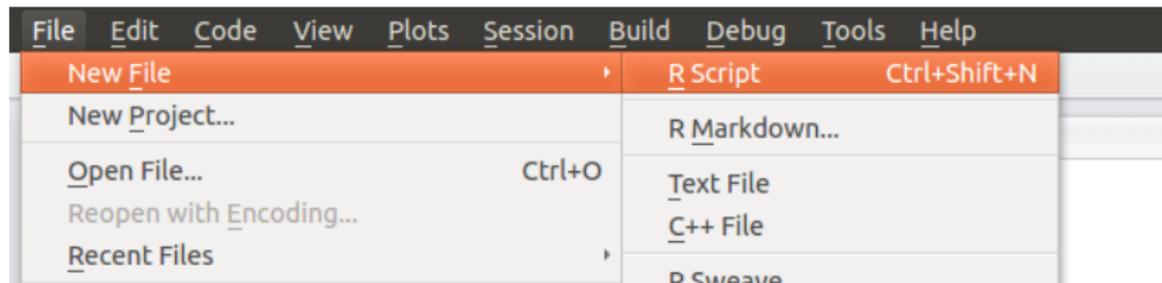
Large red text labels are overlaid on the image to identify the panes:

- CONSOLE** is overlaid on the top-left pane.
- SOURCE** is overlaid on the top-right pane.
- PLOTS / HELP** is overlaid on the bottom-left pane.
- WORKSPACE** is overlaid on the bottom-right pane.

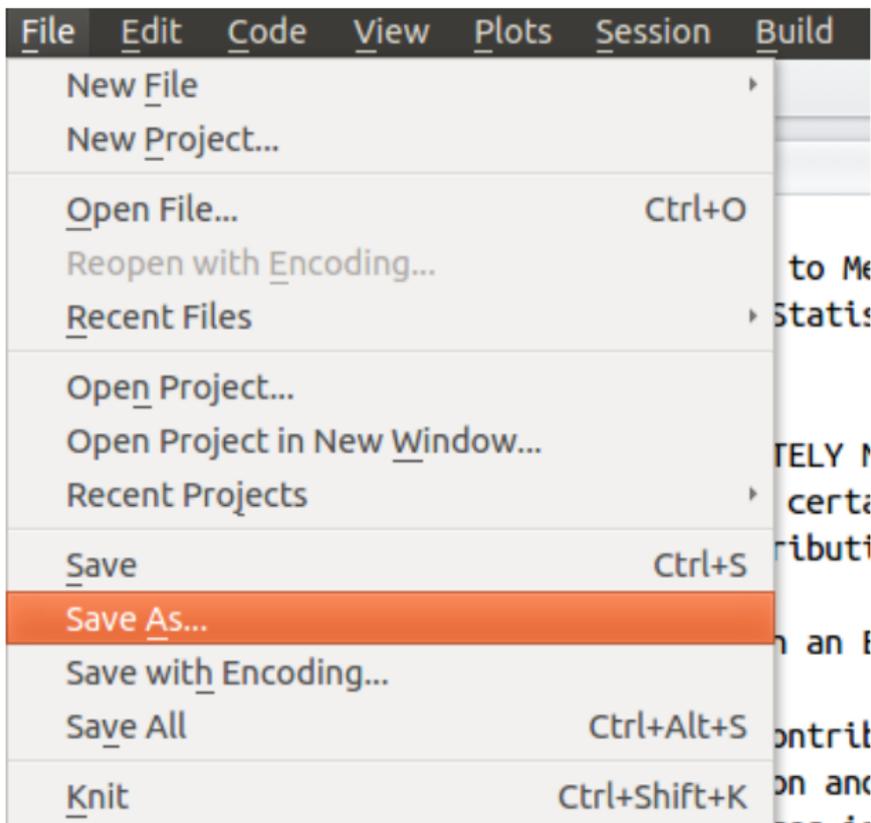
The source pane

- ▶ Often, you need to evaluate many commands sequentially
 - ▶ A command to read the data
 - ▶ Commands to transform the data
 - ▶ Commands to plot the data
 - ▶ Commands to analyze the data
 - ▶ ...
- ▶ You want to be able to save these commands so that you can later reuse them
- ▶ To do this, you need to write an *R script*
 - ▶ Plain-text file with the extension `.R` that contains all your codes
- ▶ R scripts are written and edited in the *source pane*

Create a script using file -> New File -> R script:

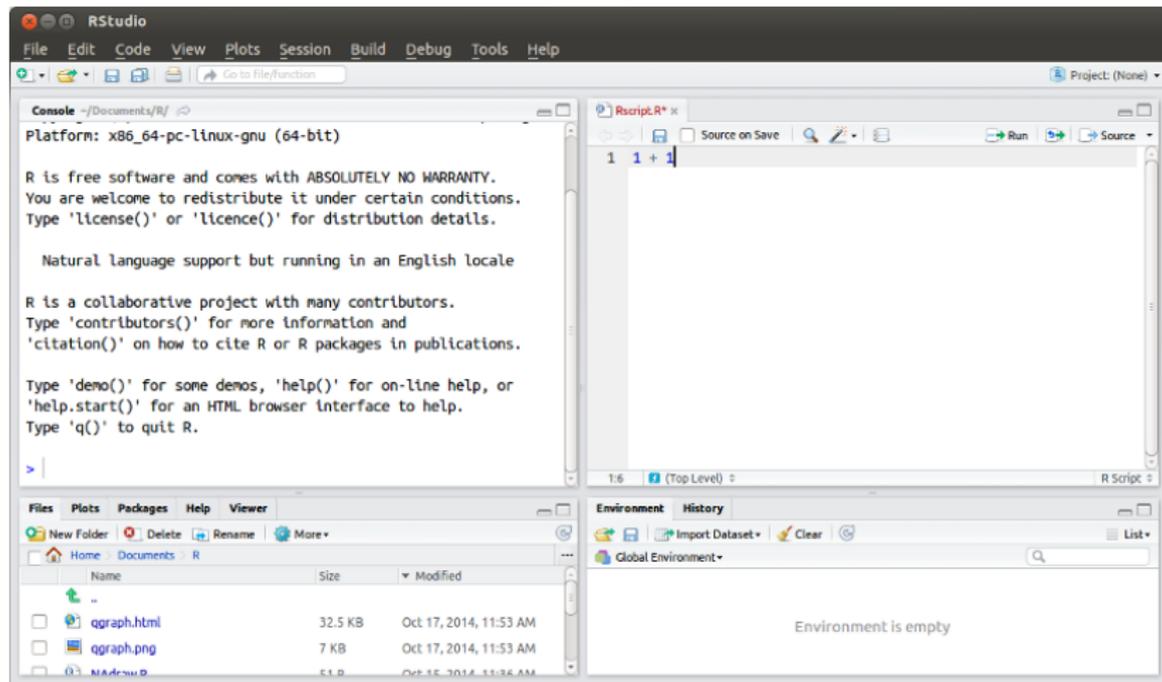


Save the script using file -> Save as:



- ▶ Save scripts with the extension .R.

Type **all** your commands into the script:



The screenshot shows the RStudio interface. The top menu bar includes File, Edit, Code, View, Plots, Session, Build, Debug, Tools, and Help. The toolbar contains icons for file operations and a 'Go to File/Function' search box. The 'Project: (None)' dropdown is visible in the top right.

The Console window (top left) displays the following text:

```
Platform: x86_64-pc-linux-gnu (64-bit)

R is free software and comes with ABSOLUTELY NO WARRANTY.
You are welcome to redistribute it under certain conditions.
Type 'license()' or 'licence()' for distribution details.

Natural language support but running in an English locale

R is a collaborative project with many contributors.
Type 'contributors()' for more information and
'citation()' on how to cite R or R packages in publications.

Type 'demo()' for some demos, 'help()' for on-line help, or
'help.start()' for an HTML browser interface to help.
Type 'q()' to quit R.

> |
```

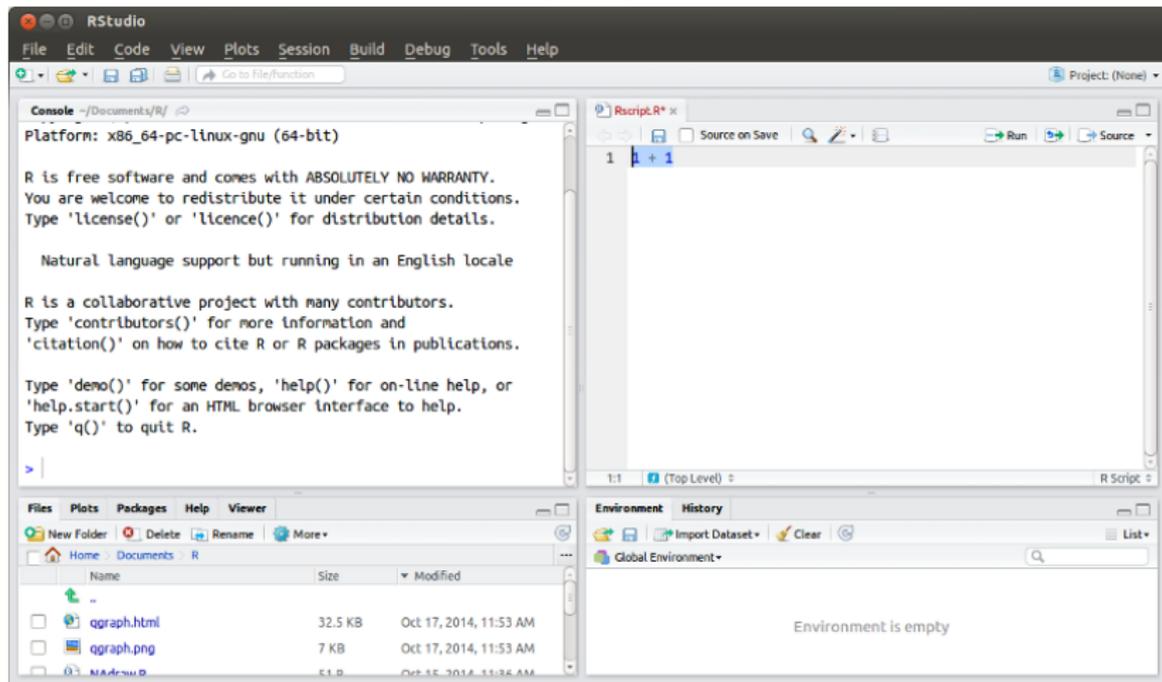
The Script Editor window (top right) shows a single line of code: `1 1 + 1`.

The Files pane (bottom left) shows the directory structure: Home > Documents > R. A table lists files:

Name	Size	Modified
..		
qgraph.html	32.5 KB	Oct 17, 2014, 11:53 AM
qgraph.png	7 KB	Oct 17, 2014, 11:53 AM
Mdfswp	51 B	Oct 17, 2014, 11:56 AM

The Environment and History pane (bottom right) shows 'Global Environment' and the message 'Environment is empty'.

Select the commands you want to evaluate:



The screenshot shows the RStudio interface. The top menu bar includes File, Edit, Code, View, Plots, Session, Build, Debug, Tools, and Help. The toolbar contains icons for file operations and a 'Go to File/Function' search box. The 'Project: (None)' dropdown is visible in the top right.

The Console window (bottom left) displays the following text:

```
Platform: x86_64-pc-linux-gnu (64-bit)

R is free software and comes with ABSOLUTELY NO WARRANTY.
You are welcome to redistribute it under certain conditions.
Type 'license()' or 'licence()' for distribution details.

Natural language support but running in an English locale

R is a collaborative project with many contributors.
Type 'contributors()' for more information and
'citation()' on how to cite R or R packages in publications.

Type 'demo()' for some demos, 'help()' for on-line help, or
'help.start()' for an HTML browser interface to help.
Type 'q()' to quit R.

> |
```

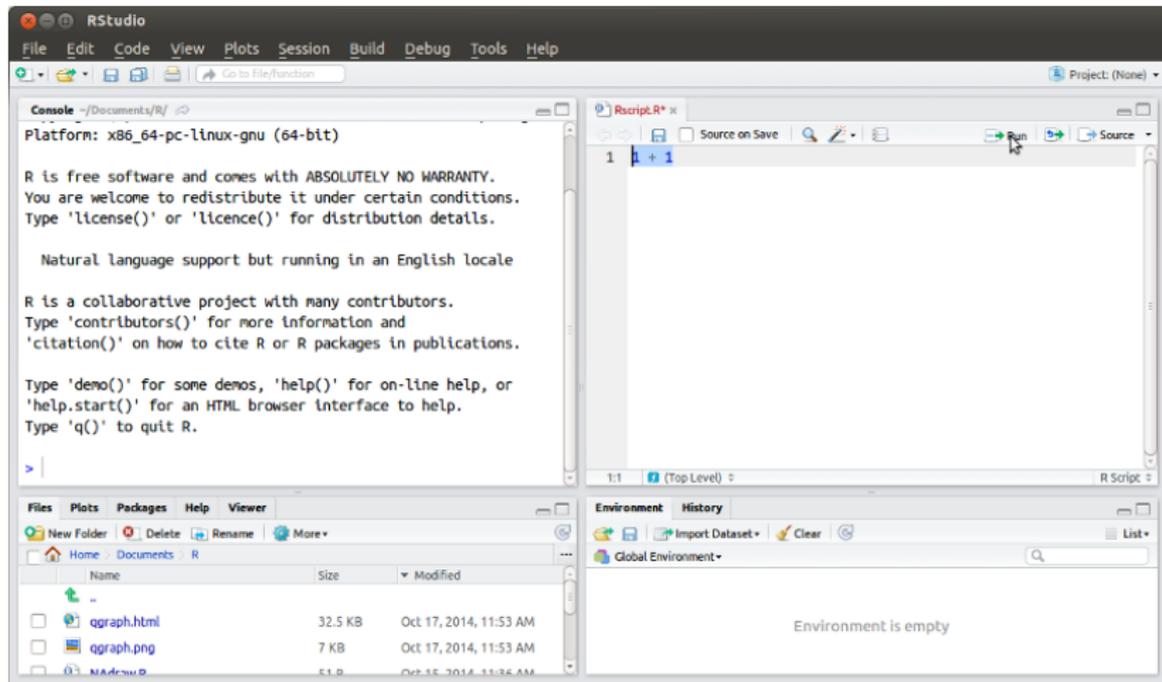
The Source Editor window (top right) shows a script named 'Rscript.R' with a single line of code: `1 + 1`. The status bar at the bottom of the editor indicates '1:1 (Top Level)' and 'R Script'.

The Files pane (bottom left) shows the current directory structure:

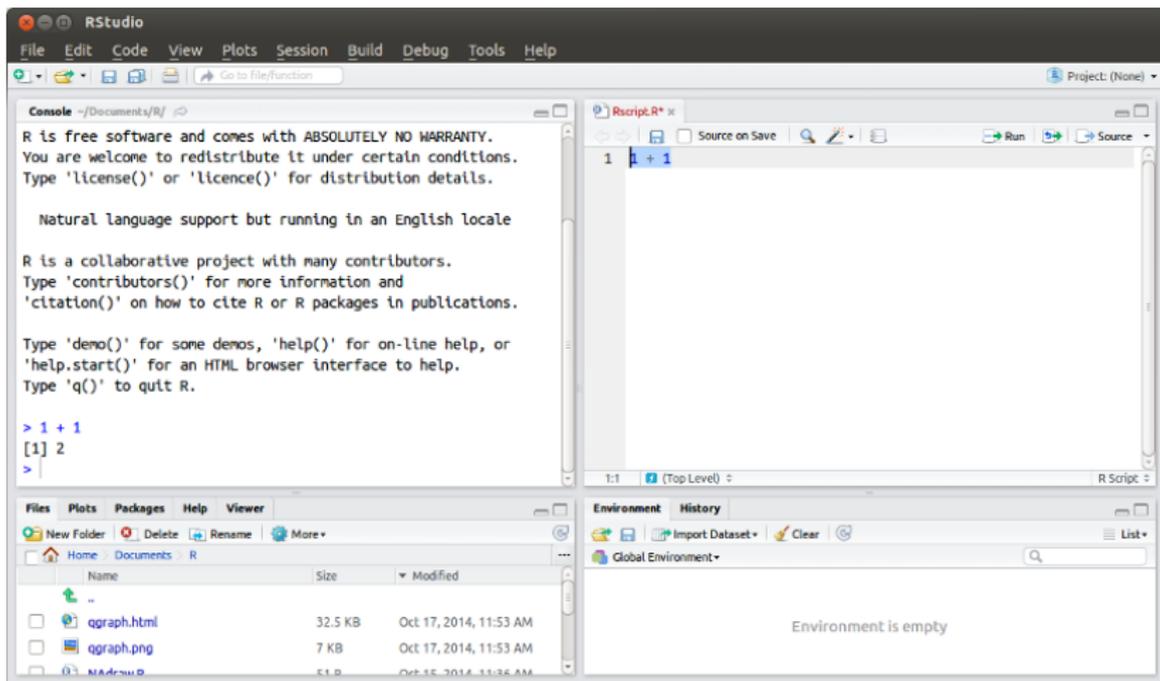
Name	Size	Modified
..		
qgraph.html	32.5 KB	Oct 17, 2014, 11:53 AM
qgraph.png	7 KB	Oct 17, 2014, 11:53 AM
Mdfcsw.D	51 B	Oct 15, 2014, 11:36 AM

The Environment and History pane (bottom right) shows 'Global Environment' and the message 'Environment is empty'.

Press the run button or control + enter on your keyboard
(command + enter for Mac users)



Now the command is sent to and evaluated in the R console!



► Never directly use the console! Always work like this!

The screenshot displays the RStudio environment with three main panels:

- Console:** Shows the R prompt and the following text:

```
'citation()' on how to cite R or R packages in publications.  
Type 'demo()' for some demos, 'help()' for on-line help, or  
'help.start()' for an HTML browser interface to help.  
Type 'q()' to quit R.  
  
> 1 + 1  
[1] 2  
> 1 + 1  
[1] 2  
>  
> # You can select and run multiple lines, including comments!  
> 2 + 2  
[1] 4  
> 3 + 3  
[1] 6  
> |
```
- Script Editor:** Shows a file named "Rscript.R" with the following code:

```
1 1 + 1  
2  
3 # You can select and run multiple lines, including comments!  
4 2 + 2  
5 3 + 3
```
- Files Panel:** Shows a file browser for the "Documents > R" directory with the following table:

Name	Size	Modified
..		
qgraph.html	32.5 KB	Oct 17, 2014, 11:53 AM
qgraph.png	7 KB	Oct 17, 2014, 11:53 AM
MAF-r.R	51 B	Oct 15, 2014, 11:38 AM
- Environment Panel:** Shows "Global Environment" with the message "Environment is empty".

Assign and use objects

- ▶ The `<-` operator can be used to store values into *objects*
 - ▶ An object can be given an arbitrary name
 - ▶ Alternatively `=` can be used to do the same thing
- ▶ an object can contain anything in R
- ▶ R expressions that are not stored in an object are *printed*

You can assign the number 1 to an object which I call a:

```
a <- 1
```

You can assign the number 1 to an object which I call a:

```
a <- 1
```

Now whenever I ask R What a is it gives me 1:

```
a
```

```
## [1] 1
```

You can assign the number 1 to an object which I call a:

```
a <- 1
```

Now whenever I ask R What a is it gives me 1:

```
a
```

```
## [1] 1
```

You can use this object in calculations:

```
a + 1
```

```
## [1] 2
```

All your objects are listed in the workspace pane:

The image shows a screenshot of the RStudio interface with four panes. The top-left pane is the Console, showing the R version 3.1.1 (64-bit) and copyright information. The top-right pane is the Source editor, which is currently empty. The bottom-left pane is the Plots/Help pane, displaying a file explorer view of the current project directory. The bottom-right pane is the Environment pane, which shows that the environment is empty.

CONSOLE

SOURCE

PLOTS / HELP

WORKSPACE

Name	Size	Modified
0	32.0 KB	2014-07-20 10:22 AM
1	7 B	2014-07-20 10:22 AM
2	7 B	2014-07-20 10:22 AM
3	270 B	2014-07-20 10:22 AM
4	4 B	2014-07-20 10:22 AM
5	18 B	2014-07-20 10:22 AM
6	300 B	2014-07-20 10:22 AM
7	7.0 KB	2014-07-20 10:22 AM
8	33.7 KB	2014-07-20 10:22 AM
9	7.0 KB	2014-07-20 10:22 AM
10	460 B	2014-07-20 10:22 AM
11	50 B	2014-07-20 10:22 AM
12	10.4 KB	2014-07-20 10:22 AM

Objects can have any name you want:

```
anyNameYouWant <- 10  
anyNameYouWant
```

```
## [1] 10
```

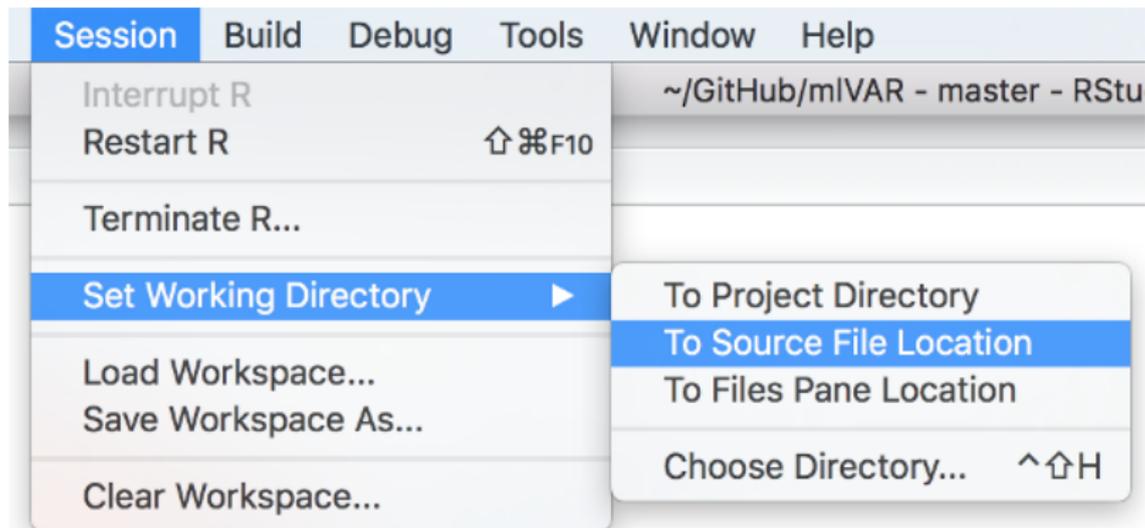
- ▶ Case sensitive! object, Object and OBJECT are **different** objects!
- ▶ It is recommended to use informative object names
 - ▶ sampleSize
 - ▶ rawData
 - ▶ nMales
- ▶ Use *camel case*
 - ▶ Start object name with lowercase letters, start new objects with lowercase and capitalize every word

Working Directory

The working directory is the folder on your computer in which R works:

```
getwd()  
setwd("~/Documents")  
getwd()
```

I recommend to never use `setwd()` or `getwd()` but to use the RStudio menu:



Write R codes assuming this is done (R script in same folder as data and/or output).

Vectors

- ▶ A vector is an object that stores multiple values of the same mode
- ▶ Use `c(...)` to manually create a vector
 - ▶ For example, `c(1,2,3)` creates the vector `[1 2 3]`
- ▶ A colon can be used to create an integer sequence

Vectors

Create a vector using `c(...)`:

```
numericVector <- c(1, 5, 10)
numericVector
```

```
## [1] 1 5 10
```

Functions

- ▶ A *function* is a small program, it takes input, does something and gives you output
- ▶ It is always of the form `name(argument, argument, argument, ...)`
- ▶ A documentation can be found for every function using `?name`
 - ▶ The documentation tells you what you need to put into a function and what comes out of it
- ▶ You need to know the exact name of the function you want to use
 - ▶ Reference card
 - ▶ Use Google!

Functions

```
a <- c(5, 2, 10, 1)
# Compute mean of v1:
mean(a)
```

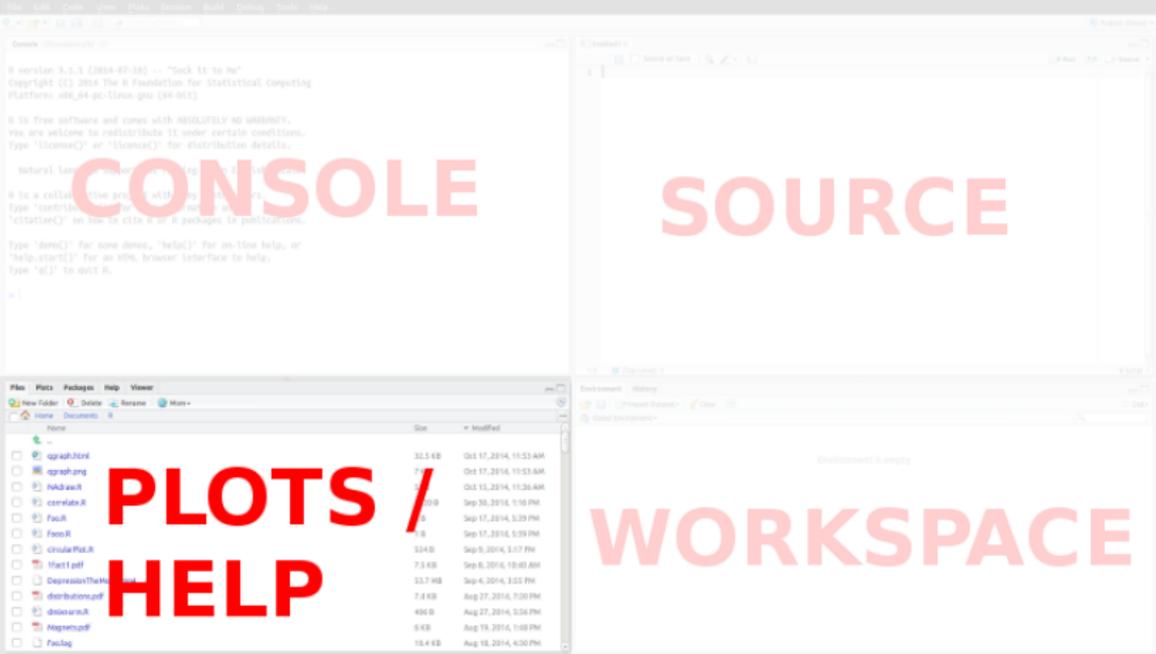
```
## [1] 4.5
```

```
# Compute sum of v1:
sum(a)
```

```
## [1] 18
```

```
# Help pages of these functions:
?mean
?sum
```

Help pages appear in the plots/help pane:



Object Modes

- ▶ There are different types of objects in R:
 - ▶ `numeric`
 - ▶ `character`
 - ▶ `logical`
- ▶ There are called *modes*

The numeric mode stores numbers:

```
numericObject <- 115  
numericObject + 1
```

```
## [1] 116
```

Strings

Anything between single or double quotation marks is stored as a character, which can be used to encode strings:

```
characterObject <- 'this is a string'  
characterObject
```

```
## [1] "this is a string"
```

```
characterObject2 <- '1'  
characterObject2 + 1
```

```
## Error in characterObject2 + 1: non-numeric argument to +
```

Logicals

The `logical` mode indicates a boolean object which can only be true (`TRUE`) or false (`FALSE`):

```
logicalObject <- TRUE  
logicalObject
```

```
## [1] TRUE
```

- ▶ Shortcuts for `TRUE` and `FALSE` are `T` and `F`
 - ▶ Avoid using these shortcuts!

Packages

- ▶ Packages are extensions contributed to R containing extra functions
 - ▶ Stored on the Comprehensive R Archive Network (CRAN)
 - ▶ Over 8000 contributed packages! All containing functions
- ▶ Packages can be installed using `install.packages()`
 - ▶ Needs to be done only once
- ▶ Afterwards they can be loaded using `library()`
- ▶ More information on packages on the CRAN site for that package:
 - ▶ `cran.r-project.org/package=PACKAGENAME`
 - ▶ Or search for CRAN PACKAGENAME

Packages

```
# Install package "qgraph"  
install.packages("qgraph")
```

```
# Load package "qgraph":  
library("qgraph")
```

Matrices

- ▶ A very important function for this course is `matrix()`. It creates a matrix, which is basically a two dimensional table.
- ▶ Technically, a matrix is a vector with two dimension attributes
- ▶ Rows indicate horizontal lines of cells
- ▶ Columns indicate vertical lines of cells
- ▶ The first argument of `matrix` is a vector to fill the matrix with, the second argument the number of rows and the third argument the number of columns (see also `?matrix`)
- ▶ Again they can be indexed with square brackets, but now need both row and column information, separated by a comma

Matrices

```
# A matrix:  
m <- matrix(1:9, nrow=3, ncol=3)  
m
```

```
##      [,1] [,2] [,3]  
## [1,]    1    4    7  
## [2,]    2    5    8  
## [3,]    3    6    9
```

Indexing matrices

```
# Index first row second column:
```

```
m[1,2]
```

```
## [1] 4
```

```
# Overwrite an element:
```

```
m[3,3] <- 0
```

```
m
```

```
##      [,1] [,2] [,3]  
## [1,]    1    4    7  
## [2,]    2    5    8  
## [3,]    3    6    0
```

Indexing matrices

```
# First row and all columns:
```

```
m[1,]
```

```
## [1] 1 4 7
```

```
# First 2x2 block:
```

```
m[1:2,1:2]
```

```
##      [,1] [,2]
```

```
## [1,]    1    4
```

```
## [2,]    2    5
```

Number of rows:

```
nrow(m)
```

```
## [1] 3
```

Number of columns:

```
ncol(m)
```

```
## [1] 3
```

Lists

```
# A vector:
```

```
v1 <- c(5,10,1,3)
```

```
v1
```

```
## [1] 5 10 1 3
```

```
# A matrix:
```

```
m1 <- matrix(c(5, 2, 10, 1),2,2)
```

```
m1
```

```
##      [,1] [,2]
```

```
## [1,]    5   10
```

```
## [2,]    2    1
```

```
# Put them in a list:
```

```
l1 <- list(v1 = v1, m1 = m1)
```

Lists

```
str(l1)
```

```
## List of 2
```

```
## $ v1: num [1:4] 5 10 1 3
```

```
## $ m1: num [1:2, 1:2] 5 2 10 1
```

Indexing lists

```
# Index the vector v1:
```

```
l1$v1
```

```
## [1] 5 10 1 3
```

```
l1[['v1']]
```

```
## [1] 5 10 1 3
```

```
# Change an element in the matrix m1:
```

```
l1$m1[2,2] <- 0
```

```
l1$m1
```

```
##      [,1] [,2]
```

```
## [1,]    5   10
```

```
## [2,]    2    0
```

Data frames

- ▶ The `data.frame()` function can be used to create a *data frames*
- ▶ A data frame is a combination of a matrix and a list
 - ▶ Looks like a matrix and can be indexed as one
 - ▶ Contains variables as each row, which can be indexed as in lists
- ▶ This is the structure you will use when working with data
- ▶ Very similar to how data is stored in SPSS!

Data frames

A character vector:

```
sex <- c("male","female","male","female")
```

A logical vector:

```
exp <- c(TRUE,TRUE,FALSE,FALSE)
```

2 numeric vectors:

```
A <- c(5,10,1,3)
```

```
B <- 1:4
```

Put them in a data frame:

```
df1 <- data.frame(sex=sex,exp=exp,A=A,B=B)
```

Data frames

```
df1
```

```
##      sex  exp  A B
## 1  male  TRUE  5 1
## 2 female  TRUE 10 2
## 3  male FALSE  1 3
## 4 female FALSE  3 4
```

Indexing data frames

```
# Index the vector sex:
```

```
df1$sex
```

```
## [1] male   female male   female
```

```
## Levels: female male
```

```
df1[['sex']]
```

```
## [1] male   female male   female
```

```
## Levels: female male
```

```
# Subset of the data containing only A and B:
```

```
df1[,c('A', 'B')]
```

```
##      A B
```

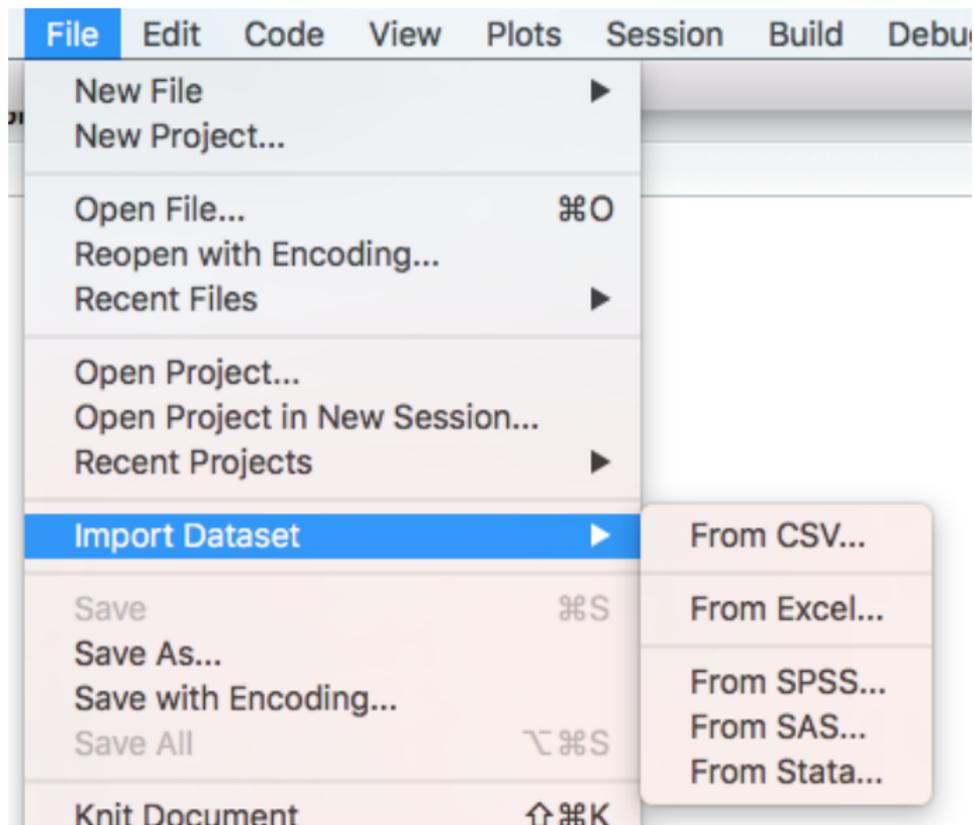
```
## 1    5 1
```

```
## 2   10 2
```

```
## 3    1 3
```

Importing data

Is now very easy!



Correlation matrix

Some functions take a data frame as input. Very important in this course is the function `cor`, which takes a data frame as input and computes the correlation matrix. The function `cov` computes the covariance (unstandardized correlations).

```
# Covariance matrix:  
cov(df1[,c('A', 'B')])
```

```
##           A           B  
## A 14.91667 -2.500000  
## B -2.50000  1.666667
```

```
# Correlation matrix:  
cor(df1[,c('A', 'B')])
```

```
##           A           B  
## A 1.0000000 -0.5013947  
## B -0.5013947  1.0000000
```

Missing data

Missing data is encoded in R with NA

```
a <- NA  
NA
```

```
## [1] NA
```

Missing data needs to be handled often using arguments in functions

Missing data

Check the help files! `?cor`

```
a <- c(1,2,NA,4)
b <- c(1,3,2,8)
cor(a,b)
```

```
## [1] NA
```

```
cor(a,b,use = "pairwise.complete.obs")
```

```
## [1] 0.9986254
```

Missing data

Or use `na.omit`

```
Data <- cbind(a,b)
cor(na.omit(Data),use = "pairwise.complete.obs")
```

```
##           a           b
## a 1.0000000 0.9986254
## b 0.9986254 1.0000000
```

Please work through the (very) short introduction to R (on blackboard and available at <https://cran.r-project.org/doc/contrib/Torfs+Brauer-Short-R-Intro.pdf>) *before* starting this assignment. You can skip section 11 on programming.