PROGRAMMING IN R









GENERAL LEARNING OBJECTIVE

Gain hands-on experience with, and a preliminary understanding of, the following elements of code in **both** R and Python:

- Variables
- **Data Structures**
- **Operators**
- Statements and Expressions

- Blocks (and Scope)
- **Functions**
- Logical (Control) Flow
- Libraries/Packages/Modules
- Inputs/Outputs
- Interpreters/Compilers







CONTENTS

- R Basics
- A Comparison of R and Python
- Resources for Programming
- **Exercises for Hands-On Learning**





PERSONAL LEARNING OBJECTIVES

Depending on where you are starting from, reaching the learning objectives for this session within the time frame of the session itself may be ambitious.

Pick a learning objective for yourself that you think is reasonable given your starting point.

You will have time to continue working on the exercises over the next weeks.







R BASICS









PROGRAMMING LANGUAGE REQUIREMENTS (DATA ANALYSIS)

Data-Centric:

- data analysts seek to uncover some information present in the data which is not obviously found during a first pass: focus on analysis, not on programming.
- The language should allow high-level concepts and questions related to data analysis to be expressed naturally.

Trustworthy:

Recipients of analyses are usually **far removed** from the details of implementation, which is why the language used for analysis must be trustworthy.







PERFORMANCE

Data analysis is centered on the analyst figuring out what questions to ask and what transformations or algorithms **may provide insights**:

the computer spends most of its time waiting for user input – it is more valuable to have a language that promotes human/analyst efficiency at this stage.

Once a **pipeline to answer the question** has been established, the algorithm can be run again and again every time the data set is updated:

data processing, not data analysis anymore – the non-creative process is **mechanical** and can be **automated**. This is when computer performance becomes paramount.









ABOUT R

R is an open source implementation of the (commercial) S language:

- designed for statistical and data analysis
- supports functional and object-oriented styles of programming
- data frame is a built-in type
- missing values are built-in (WARNING: read the documentation on NAs!!!)
- model formulas are first-class objects
- contains advanced statistical routines (external packages that must be loaded)
- state of the art graphics capabilities (ggplot2)
- has dialects (the tidyverse, in particular, is quite popular currently)









ABOUT R

The functional style is **not enforced**, although it is encouraged and natural in R.

The main unit of programming is the **function**.

Data frames are the **natural containers** for data (also, tibbles)

- "tables" whose rows are observations and columns are variables
- this is a **natural data model** (DBMS, spreadsheets) ...
- ... but there are restrictions when using tidy data (more on this later)









R REFERENCES & BOOKS

R Project: http://www.r-project.org (installation)

R Studio: http://www.rstudio.com (development environment)

R Cookbook: http://www.cookbook-r.com

The R Inferno: http://bit.ly/1mpZabc

Advanced R, by H. Wickham: https://adv-r.hadley.nz

R for Data Science, by G. Grolemund and H. Wickham: https://r4ds.had.co.nz







BASIC TYPES

Vector: logical, integer, double, double complex, string

- each has a missing value literal: NA
- NULL is not the same thing as NA
- scalars are vectors of length 1

Structured: list, factor, data frame, array, matrix, etc.







BASIC TYPES

Factor: represents categorical variable

- similar to vector of strings
- stored as an integer vector

Data Frame: represents observations of statistical variables

- behaves as a list of vectors
- also behaves as a matrix









BASIC OPERATORS

Usual arithmetic, comparison, logical operators are elementwise

For matrix multiplication use \%*\%

There are 5 assignment operators!!!

■ =, <- (best practice), ->, <<-, ->>

There are 3 indexing operators:

- [[]] for access to a single element by index or name
- [] for access to multiple elements by index or name
- \$ for access to single element by name









BASIC FLOW CONTROL

```
if(cond){expr}
if(cond){expr 1 else {expr2}}
for(var in seq) {expr}
while(cond) {expr}
repeat {expr}
break
next
```

Avoid loops if at all possible (they can be VERY SLOW...); vectorize instead!







FUNCTIONS

Functions are first-class objects

Function arguments passing can be complicated! (less so with the tidyverse)

- parameters are assigned by position or name
- use = to assign parameter by name
- lazy evaluation (symbol, function only evaluated when needed)

Hint: try using one of the *apply functions instead of a loop whenever possible









PRACTICE, PRACTICE, PRACTICE

In what follows we show how to do ... well, more data stuff using R.

The selection of problems is still not intended to be complete, but it provides a gloss of the myriad ways to approach data analysis with R.

Some of it overlaps with the other **Data Analysis Short Course** notebooks.

As always, do not hesitate to consult online resources for more examples (StackOverflow, R-bloggers, etc.).







We start by loading one of the standard pedagogical datasets used with R.

```
> data(cars)
```

It's a data frame with two variables, **speed** and **dist**.

```
> str(cars)
## 'data.frame': 50 obs. of 2 variables:
## $ speed: num 4 4 7 7 8 9 10 10 10 11 ...
   $ dist: num 2 10 4 22 16 10 18 26 34 17 ...
> summary(cars)
                      dist
##
       speed
  Min. : 4.0
                 Min. : 2.00
##
  1st Qu.:12.0
                 1st Qu.: 26.00
                 Median : 36.00
   Median :15.0
##
   Mean :15.4
                        : 42.98
                 Mean
   3rd Qu.:19.0
##
                 3rd Ou.: 56.00
   Max.
          :25.0
                  Max.
                        :120.00
```







We can compute the mean of each variable using the pre-built function colMeans().

```
> colMeans(cars)
## speed dist
## 15.40 42.98
```

But there is no such function to compute the minimum/maximum of each variable. instead, we use loops.

```
# min
> for(i in 1:2){
    print(min(cars[,i]))
## [1] 4
## [1] 2
```







```
# max
> for(i in 1:2){
    print(max(cars[,i]))
## [1] 25
## [1] 120
```

The use of loops was justified because the dataset is **small**, but it would be useful to know how to the computation using once of the *apply functions.

```
> sapply(cars, min)
## speed dist
##
> sapply(cars, max)
## speed dist
##
         120
```









```
> sapply(cars, mean)
## speed dist
## 15.40 42.98
```

Note that in each of the three cases, the result is a data frame.

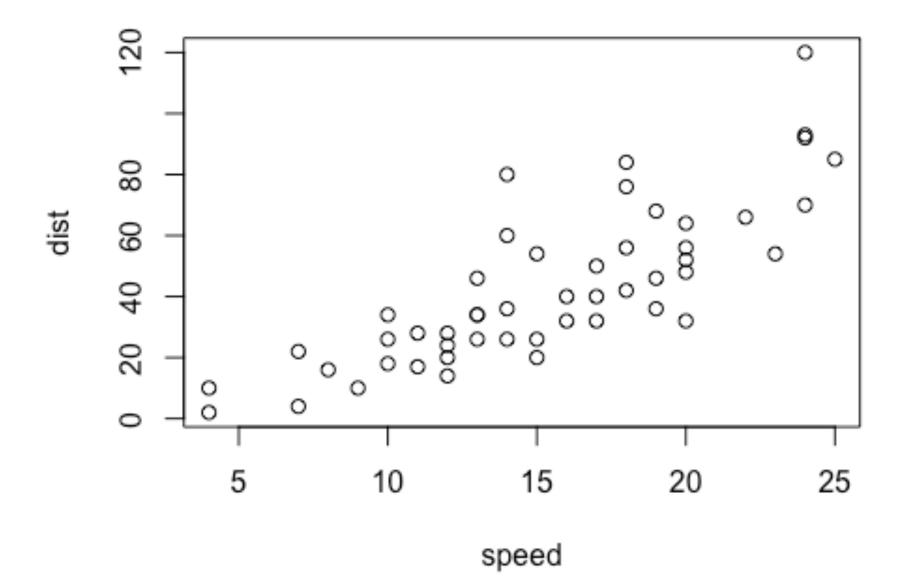
A dataset with 2 variables is easy to display.

```
> plot(cars)
```









We can see that the relationship between speed and distance has a strong linear component, which we can identify by running a linear regression (implemented in the function lm):

```
> req <- lm(dist ~ speed, data=cars)</pre>
```

In this assignment, reg is the **object** (or model) obtained when running **lm()** on the cars dataset.

We can summarize it and see what its attributes are.









```
> summary(reg)
## Call:
## lm(formula = dist ~ speed, data = cars)
##
## Residuals:
## Min 1Q Median 3Q Max
## -29.069 -9.525 -2.272 9.215 43.201
##
## Coefficients:
## Estimate Std. Error t value Pr(>|t|)
## (Intercept) -17.5791 6.7584 -2.601 0.0123 *
## speed 3.9324 0.4155 9.464 1.49e-12 ***
## ---
## Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 15.38 on 48 degrees of freedom
## Multiple R-squared: 0.6511, Adjusted R-squared: 0.6438
## F-statistic: 89.57 on 1 and 48 DF, p-value: 1.49e-12
```





```
> attributes(reg)
## $names
                                       "effects"
## [1] "coefficients" "residuals"
                                                       "rank"
## [5] "fitted.values" "assign"
                                       "qr"
                                                       "df.residual"
## [9] "xlevels"
                 "call"
                                       "terms"
                                                       "model"
## $class
## [1] "lm"
```

We can extract the coefficients of the linear model as follows:

```
> reg$coefficients
## (Intercept)
                  speed
## -17.579095 3.932409
```

and use them to plot the line of best fit over the display:

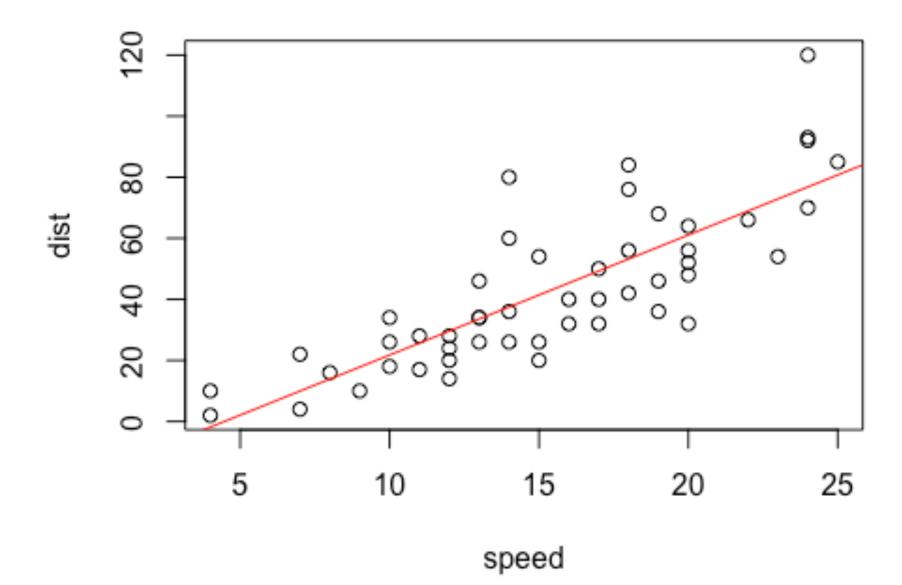
```
> plot(cars)
> abline(reg$coefficients, col="red")
```











R's matrix notation is not super intuitive, as we will see presently.

Let's generate 2 matrices that we will use in the following code blocks (the brackets around the code mean that the assignment is followed by an object display).

```
> (M <- matrix(1:12, nrow=3, ncol=4))</pre>
##
          [,1] [,2] [,3] [,4]
## [1,] 1 4 7 10
## [2,] 2 5 8 11
## [3,] 3 6 9 12
> (V <- matrix( runif(4), 4, 1))</pre>
##
                [,1]
## [1,] 0.02875853
## [2,] 0.07807830
## [3,] 0.31080912
## [4,] 0.54691573
```







When ncol and nrow are not specified, the default parameter ordering takes over.

Since M is 3×4 and V is 4×1 , the product MV exists (and is 3×1).

```
# use %*% to multiply matrices
> M %*% V
            [,1]
## [1,] 7.985893
## [2,] 8.950455
## [3,] 9.915016
```

The product VM does not exist, however, as the dimensions are not compatible.

```
# uncomment the following line to see the error message
> # V %*% M
```

The choice of %*% for matrix multiplication is not ... inspired; most languages use *.









What DOES * do in R?

```
> M * 2
        [,1] [,2] [,3] [,4]
## [1,] 2 8 14 20
## [2,] 4 10 16 22
## [3,] 6 12 18
                          24
```

Hah! Multiplication by a scalar. Good to know.

One of R's most nefarious habit is that it is not always compatible with resonable mathematical notation. What do you think the following line of code does?

$$> M * c(2, -2)$$









Apparently, it cycles through the arguments?!?

It's hard to imagine why this construction should yield a result without breaking down, and yet it does. Beware, then.

(It's probably a good idea to verify that code does what it's meant to along the way.)









Other familiar operations (like the transpose) are easy to compute:

is indeed the 4×3 transpose of M.







rbind and **cbind** are used to bind rows and columns, respectively.

```
> cbind(t(M), V)
## [,1] [,2] [,3] [,4]
## [1,] 1 2 3 0.02875853
## [2,] 4 5 6 0.07807830
## [3,] 7 8 9 0.31080912
## [4,] 10 11 12 0.54691573
> rbind(M,t(V))
##
              [,1] [,2] [,3] [,4]
## [1,] 1.00000000 4.0000000 7.0000000 10.0000000
## [2,] 2.00000000 5.0000000 8.0000000 11.0000000
## [3,] 3.00000000 6.0000000 9.0000000 12.0000000
## [4,] 0.02875853 0.0780783 0.3108091 0.5469157
# but this wont work because of dimension incompatibility; uncomment to test
> # cbind(M, V)
```







WORKING WITH STRINGS

What is usually called a **string** in other programming languages is a **character object** in R.

Character vectors are created by using double quotes (", preferred) or single quotes (', acceptable).

```
> "Come on, everybody!"
## [1] "Come on, everybody!"
```

In R, strings are scalar values, not vectors of characters.

```
> length("Come on, everybody!")
## [1] 1
```









WORKING WITH STRINGS

The combining function \mathbf{c} () creates a vector of strings.

```
> c("Come", "on", "," , "everybody", "!")
## [1] "Come" "on" "," "everybody"
                                                   "!"
```

We use **paste** or **paste0** to **concatenate** strings:

```
> paste("Come", "on", "," , "everybody", "!")
## [1] "Come on , everybody!"
> paste("Come", "on", "," , "everybody", "!", sep=" ")
## [1] "Come on , everybody!"
```

The function **strsplit()** does the opposite:

```
> strsplit("Come on, everybody!", ", ")
## [[1]]
## [1] "Come on" "everybody!"
```









R NOTEBOOKS & CHEATSHEETS

These examples, as well as numerous others, can be found in the R Notebooks

- R Basics
- More Data Stuff in R

There is no substitute for writing code, but these examples could be a good starting point for many analysts.

Take the time to run those examples at your console to get a sense for R in action.

RStudio publishes various cheatsheets, which can come in handy (see next slides), but remember that these cheatsheets may become obsolete with new R releases.











Getting Help

Accessing the help files

?mean

Get help of a particular function.

help.search('weighted mean')

Search the help files for a word or phrase.

help(package = 'dplyr')

Find help for a package.

More about an object

str(iris)

Get a summary of an object's structure.

class(iris)

Find the class an object belongs to.

Using Packages

install.packages('dplyr')

Download and install a package from CRAN.

library(dplyr)

Load the package into the session, making all its functions available to use.

dplyr::select

Use a particular function from a package.

data(iris)

Load a built-in dataset into the environment.

Working Directory

getwd()

Find the current working directory (where inputs are found and outputs are sent).

setwd('C://file/path')

Change the current working directory.

Use projects in RStudio to set the working directory to the folder you are working in.

Vectors Creating Vectors Join elements into 2 4 6 c(2, 4, 6) a vector An integer 2:6 2 3 4 5 6 sequence A complex 2.0 2.5 3.0 seq(2, 3, by=0.5)sequence 1 2 1 2 1 2 rep(1:2, times=3) Repeat a vector Repeat elements 111222 rep(1:2, each=3) of a vector

Vector Functions

sort(x)	rev(x)
Return x sorted.	Return x reversed.
table(x)	unique(x)
See counts of values.	See unique values.

Selecting Vector Elements

By Position

x[4]	The fourth elemen
	1110 1001111 010111011

```
x[-4] All but the fourth.
```

x l	12	4	ı	Elements	two	to	foui
^				LIGITICITIS	LVVO	ιO	IOUI

x[-(2:4)]	All elements except
X[-(2:4)]	two to four.

w[a/1	E \ 1	Elements one and
x[c(1,	5)]	five.

By Value

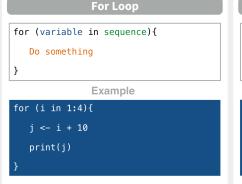
x[x == 10]	are equal to 10.
x[x < 0]	All elements less than zero.
x[x %in% c(1, 2, 5)]	Elements in the set

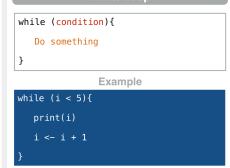
Elemente which

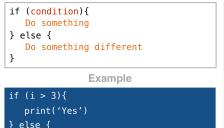
Named Vectors

x['apple']	Element with
	name 'apple'.

Programming While Loop

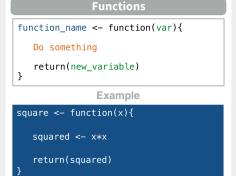






print('No')

If Statements



Reading and Writing Data

Also see the **readr** package.

Input	Ouput	Description
<pre>df <- read.table('file.txt')</pre>	<pre>write.table(df, 'file.txt')</pre>	Read and write a delimited text file.
<pre>df <- read.csv('file.csv')</pre>	write.csv(df, 'file.csv')	Read and write a comma separated value file. This is a special case of read.table/ write.table.
<pre>load('file.RData')</pre>	<pre>save(df, file = 'file.Rdata')</pre>	Read and write an R data file, a file type special for R.

Types

Converting between common data types in R. Can always go from a higher value in the table to a lower value.

as.logical	TRUE, FALSE, TRUE	Boolean values (TRUE or FALSE)
as.numeric	1, 0, 1	Integers or floating point numbers.
as.character	'1', '0', '1'	Character strings. Generally preferred to factors.
as.factor	'1', '0', '1', levels: '1', '0'	Character strings with preset levels. Needed for some statistical models.

Maths Functions

Natural log.	sum(x)	Sum.
Exponential.	mean(x)	Mean.
Largest element.	median(x)	Median.
Smallest element.	quantile(x)	Percentage quantiles.
Round to n decimal places.	rank(x)	Rank of elements
Round to n significant figures.	var(x)	The variance.
Correlation.	sd(x)	The standard deviation.
	Exponential. Largest element. Smallest element. Round to n decimal places. Round to n significant figures.	Exponential. mean(x) Largest element. median(x) Smallest element. quantile(x) Round to n decimal places. Round to n significant figures.

Variable Assignment

> a <- 'apple'
> a
[1] 'apple'

The Environment

ls()	List all variables in the environment.
rm(x)	Remove x from the environment.
rm(list = ls())	Remove all variables from the environment.

You can use the environment panel in RStudio to browse variables in your environment.

Matrices

 $m \leftarrow matrix(x, nrow = 3, ncol = 3)$ Create a matrix from x.



Lists

 $l \leftarrow list(x = 1:5, y = c('a', 'b'))$

A list is a collection of elements which can be of different types.

1[[2]]
Second element

l[1] New list with only the first

element.

l\$x Element named

New list with only element named y.

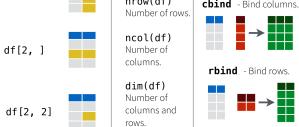
l['y']

Also see the **dplyr** package.

Data Frames

 $df \leftarrow data.frame(x = 1:3, y = c('a', 'b', 'c'))$ A special case of a list where all elements are the same length.

		List subsetting			
х	У		_		
1	a	df\$x	df[[2]]		
2	b	Understanding a data frame			
3	С	View(df)	See the full data frame.		
Matrix subsetting		head(df)	See the first 6 rows.		
df[, 2]		nrow(df) Number of rows.	cbind - Bind columns.		



Strings

Also see the **stringr** package.

paste(x, y, sep = ' ') Join multiple vectors together. paste(x, collapse = ' ') Join elements of a vector together. grep(pattern, x) Find regular expression matches in x. gsub(pattern, replace, x) Replace matches in x with a string. toupper(x) Convert to uppercase. tolower(x) Convert to lowercase.

Factors

factor(x)

nchar(x)

Turn a vector into a factor. Can set the levels of the factor and the order.

cut(x, breaks = 4)Turn a numeric vector into a factor by 'cutting' into sections.

Number of characters in a string.

Statistics

 $lm(y \sim x, data=df)$ Linear model.

 $qlm(y \sim x, data=df)$ Generalised linear model.

summary Get more detailed information out a model.

t.test(x, y) Perform a t-test for difference between means.

pairwise.t.test

Perform a t-test for paired data.

Test for a difference between proportions.

prop.test

aov Analysis of variance.

Distributions

	Random Variates	Density Function	Cumulative Distribution	Quantile
Normal	rnorm	dnorm	pnorm	qnorm
Poisson	rpois	dpois	ppois	qpois
Binomial	rbinom	dbinom	pbinom	qbinom
Uniform	runif	dunif	punif	qunif

Plotting

Also see the ggplot2 package.



plot(x, y) Values of x against y.

hist(x)Histogram of

Dates

See the **lubridate** package.

Advanced R Cheat Sheet

Created by: Arianne Colton and Sean Chen

Environment Basics

Environment – **Data structure** (with two components below) that powers lexical scoping

Create environment: env1<-new.env()

 Named list ("Bag of names") – each name points to an object stored elsewhere in memory.

If an object has no names pointing to it, it gets automatically deleted by the garbage collector.

- · Access with: Is('env1')
- Parent environment used to implement lexical scoping. If a name is not found in an environment, then R will look in its parent (and so on).
 - Access with: parent.env('env1')

Four special environments

- Empty environment ultimate ancestor of all environments
 - · Parent: none
 - · Access with: emptyenv()
- **2.** Base environment environment of the base package
 - Parent: empty environment
 - · Access with: baseenv()
- **3. Global environment** the interactive workspace that you normally work in
 - Parent: environment of last attached package
 - Access with: globalenv()
- Current environment environment that R is currently working in (may be any of the above and others)
 - Parent: empty environmentAccess with: environment()

Environments

Search Path

Search path – mechanism to look up objects, particularly functions.

- Access with: search() lists all parents of the global environment (see Figure 1)
- Access any environment on the search path: as.environment('package:base')

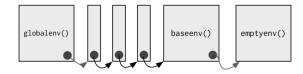


Figure 1 - The Search Path

- Mechanism: always start the search from global environment, then inside the latest attached package environment.
- New package loading with library()/require(): new package is attached right after global environment. (See Figure 2)
- Name conflict in two different package: functions with the same name, latest package function will get called.

search():

'.GlobalEnv' ... 'Autoloads' 'package:base'

library(reshape2); search()

'.GlobalEnv' 'package:reshape2' ... 'Autoloads' 'package:base'

NOTE: Autoloads: special environment used for saving memory by only loading package objects (like big datasets) when needed

Figure 2 - Package Attachment

Binding Names to Values

Assignment – act of binding (or rebinding) a name to a value in an environment.

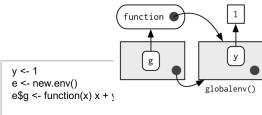
- <- (Regular assignment arrow) always creates a variable in the current environment
- 2. <<- (Deep assignment arrow) modifies an existing variable found by walking up the parent environments

Warning: If <<- doesn't find an existing variable, it will create one in the global environment.

Function Environments

- Enclosing environment an environment where the function is created. It determines how function finds value.
 - Enclosing environment never changes, even if the function is moved to a different environment.
 - Access with: environment('func1')
- Binding environment all environments that the function has a binding to. It determines how we find the function.
 - · Access with: pryr::where('func1')

Example (for enclosing and binding environment):



- function g enclosing environment is the global environment,
- the binding environment is "e".
- Execution environment new created environments to host a function call execution.
 - Two parents :
 - I. Enclosing environment of the function
 - II. Calling environment of the function
 - Execution environment is thrown away once the function has completed.
- Calling environment environments where the function was called.
 - Access with: parent.frame('func1')
 - Dynamic scoping :
 - About : look up variables in the calling environment rather than in the enclosing environment
 - Usage : most useful for developing functions that aid interactive data analysis

Data Structures

	Homogeneous	Heterogeneous
1d	Atomic vector	List
2d	Matrix	Data frame
nd	Array	

Note: R has no 0-dimensional or scalar types. Individual numbers or strings, are actually vectors of length one, NOT scalars.

Human readable description of any R data structure :

str(variable)

Every Object has a mode and a class

- 1. **Mode**: represents how an object is stored in memory
 - 'type' of the object from R's point of view
 - · Access with: typeof()
- 2. Class: represents the object's abstract type
 - 'type' of the object from R's object-oriented programming point of view
 - · Access with: class()

	typeof()	class()
strings or vector of strings	character	character
numbers or vector of numbers	numeric	numeric
list	list	list
data.frame	list	data.frame

Factors

1. Factors are built on top of integer vectors using two attributes:

class(x) -> 'factor'
levels(x) # defines the set of allowed values

Useful when you know the possible values a variable may take, even if you don't see all values in a given dataset.

Warning on Factor Usage:

- Factors look and often behave like character vectors, they are actually integers. Be careful when treating them like strings.
- Most data loading functions automatically convert character vectors to factors. (Use argument stringAsFactors = FALSE to suppress this behavior)

Object Oriented (OO) Field Guide

Object Oriented Systems

R has three object oriented systems:

- S3 is a very casual system. It has no formal definition of classes. It implements generic function OO.
 - Generic-function OO a special type of function called a generic function decides which method to call.

Example:	drawRect(canvas, 'blue')
Language:	R

 Message-passing OO - messages (methods) are sent to objects and the object determines which function to call.

Example:	canvas.drawRect('blue')
Language:	Java, C++, and C#

- **2. S4** works similarly to S3, but is more formal. Two major differences to S3 :
 - Formal class definitions describe the representation and inheritance for each class, and has special helper functions for defining generics and methods.
 - Multiple dispatch generic functions can pick methods based on the class of any number of arguments, not just one.
- Reference classes are very different from S3 and S4:
 - Implements message-passing OO methods belong to classes, not functions.
 - Notation \$ is used to separate objects and methods, so method calls look like canvas\$drawRect('blue').

S

1. About S3:

- · R's first and simplest OO system
- Only OO system used in the base and stats package
- Methods belong to functions, not to objects or classes

2. Notation:

· generic.class()

	Date method for the generic - mean()
--	--------------------------------------

3. Useful 'Generic' Operations

- Get all methods that belong to the 'mean' generic:
 - Methods('mean')
- List all generics that have a method for the 'Date' class:
- methods(class = 'Date')
- **4. S3 objects** are usually built on top of lists, or atomic vectors with attributes.
 - · Factor and data frame are S3 class
 - · Useful operations:

Check if object is an S3 object	is.object(x) & !isS4(x) or pryr::otype()
Check if object inherits from a specific class	inherits(x, 'classname')
Determine class of any object	class(x)

Base Type (C Structure)

R base types - the internal C-level types that underlie the above OO systems.

- Includes: atomic vectors, list, functions, environments, etc.
- Useful operation: Determine if an object is a base type (Not S3, S4 or RC) is.object(x) returns FALSE
- Internal representation: C structure (or struct) that includes:
 - · Contents of the object
 - · Memory Management Information
 - Type
 - Access with: typeof()

Functions

Function Basics

Functions – objects in their own right All R functions have three parts:

body()	code inside the function
formals()	list of arguments which controls how you can call the function
environment()	"map" of the location of the function's variables (see "Enclosing Environment")

Every operation is a function call

- +, for, if, [, \$, { ...
- x + y is the same as `+`(x, y)

Note: the backtick (`), lets you refer to functions or variables that have otherwise reserved or illegal names.

Lexical Scoping

What is Lexical Scoping?

- Looks up value of a symbol. (see "Enclosing Environment")
- **findGlobals()** lists all the external dependencies of a function

```
f <- function() x + 1
codetools::findGlobals(f)
> '+' 'x'
environment(f) <- emptyenv()
```

error in f(): could not find function "+"

 R relies on lexical scoping to find everything, even the + operator.

Function Arguments

Arguments – passed by reference and copied on modify

- 1. Arguments are matched first by exact name (perfect matching), then by prefix matching, and finally by position.
- 2. Check if an argument was supplied: missing()

```
i <- function(a, b) {
  missing(a) -> # return true or false
}
```

Lazy evaluation – since x is not used stop("This is an error!")
never get evaluated.

```
f <- function(x) {
10
}
f(stop('This is an error!')) -> 10
```

4. Force evaluation

```
f <- function(x) {
force(x)
10
}
```

5. Default arguments evaluation

```
f <- function(x = ls()) {
    a <- 1
    x
}
```

f() -> 'a' 'x'	ls() evaluated inside f
f(ls())	ls() evaluated in global environment

Return Values

- Last expression evaluated or explicit return().
 Only use explicit return() when returning early.
- Return ONLY single object.
 Workaround is to return a list containing any number of objects.
- Invisible return object value not printed out by default when you call the function.

```
f1 <- function() invisible(1)
```

Primitive Functions

What are Primitive Functions?

1. Call C code directly with .Primitive() and contain no R code

```
print(sum) :
> function (..., na.rm = FALSE) .Primitive('sum')
```

- 2. formals(), body(), and environment() are all NULL
- 3. Only found in base package
- 4. More efficient since they operate at a low level

Influx Functions

What are Influx Functions?

- 1. Function name comes in between its arguments, like + or -
- 2. All user-created infix functions must start and end with %.

```
`%+%` <- function(a, b) paste0(a, b)
'new' %+% 'string'
```

3. Useful way of providing a default value in case the output of another function is NULL:

```
`%||%` <- function(a, b) if (!is.null(a)) a else b function_that_might_return_null() %||% default value
```

Replacement Functions

What are Replacement Functions?

- Act like they modify their arguments in place, and have the special name xxx <-
- Actually create a modified copy. Can use pryr::address() to find the memory address of the underlying object

```
`second<-` <- function(x, value) {
    x[2] <- value
    x
}
x <- 1:10
second(x) <- 5L
```

Subsetting returns a copy of the original data, NOT copy-on modified

Simplifying vs. Preserving Subsetting

1. Simplifying subsetting

 Returns the **simplest** possible data structure that can represent the output

2. Preserving subsetting

- Keeps the structure of the output the **same** as the input.
- When you use drop = FALSE, it's preserving

	Simplifying*	Preserving
Vector	x[[1]]	x[1]
List	x[[1]]	x[1]
Factor	x[1:4, drop = T]	x[1:4]
Array	x[1,] or x[, 1]	x[1, , drop = F] or x[, 1, drop = F]
Data frame	x[, 1] or x[[1]]	x[, 1, drop = F] or x[1]

<u>Simplifying behavior</u> varies slightly between different data types:

1. Atomic Vector

x[[1]] is the same as x[1]

2. List

- [] always returns a list
- Use [[]] to get list contents, this returns a single value piece out of a list

3. Factor

Drops any unused levels but it remains a factor class

4. Matrix or Array

 If any of the dimensions has length 1, that dimension is dropped

5. Data Frame

 If output is a single column, it returns a vector instead of a data frame

Data Frame Subsetting

Data Frame – possesses the characteristics of both lists and matrices. If you subset with a single vector, they behave like lists; if you subset with two vectors, they behave like matrices

1. Subset with a single vector: Behave like lists

```
df1[c('col1', 'col2')]
```

2. Subset with two vectors: Behave like matrices

```
df1[, c('col1', 'col2')]
```

The results are the same in the above examples, however, results are different if subsetting with only one column. (see below)

1. Behave like matrices

```
str(df1[, 'col1']) -> int [1:3]
```

· Result: the result is a vector

2. Behave like lists

```
str(df1['col1']) -> 'data.frame'
```

· Result: the result remains a data frame of 1 column

\$ Subsetting Operator

1. About Subsetting Operator

Useful shorthand for [[combined with character subsetting

```
x$y is equivalent to x[['y', exact = FALSE]]
```

2. Difference vs. [[

• \$ does partial matching, [[does not

```
x <- list(abc = 1)
x$a -> 1  # since "exact = FALSE"
x[['a']] ->  # would be an error
```

3. Common mistake with \$

• Using it when you have the name of a column stored in a variable

```
var <- 'cyl'
x$var

# doesn't work, translated to x[['var']]
# Instead use x[[var]]
```

Example

1. Lookup tables (character subsetting)

```
x <- c('m', 'f', 'u', 'f', 'f', 'm', 'm')
lookup <- c(m = 'Male', f = 'Female', u = NA)
lookup[x]
> m f u f f m m
> 'Male' 'Female' NA 'Female' 'Female' 'Male' 'Male'
unname(lookup[x])
> 'Male' 'Female' NA 'Female' 'Female' 'Male' 'Male'
```

Matching and merging by hand (integer subsetting) Lookup table which has multiple columns of information:

```
grades <- c(1, 2, 2, 3, 1)
info <- data.frame(
   grade = 3:1,
   desc = c('Excellent', 'Good', 'Poor'),
   fail = c(F, F, T)
)
```

First Method

```
id <- match(grades, info$grade)
info[id, ]
```

Second Method

```
rownames(info) <- info$grade
info[as.character(grades),]
```

- 3. Expanding aggregated counts (integer subsetting)
 - Problem: a data frame where identical rows have been collapsed into one and a count column has been added
 - **Solution**: rep() and integer subsetting make it easy to uncollapse the data by subsetting with a repeated row index: rep(x, y) rep replicates the values in x, y times.

```
df1$countCol is c(3, 5, 1)
rep(1:nrow(df1), df1$countCol)
> 1 1 1 2 2 2 2 2 3
```

4. Removing columns from data frames (character subsetting)

There are two ways to remove columns from a data frame:

•		
	Set individual columns to NULL	df1\$col3 <- NULL
	Subset to return only columns you want	df1[c('col1', 'col2')]

- 5. Selecting rows based on a condition (logical subsetting)
 - This is the most commonly used technique for extracting rows out of a data frame.

```
df1[df1$col1 == 5 & df1$col2 == 4, ]
```

Subsetting continued

Boolean Algebra vs. Sets (Logical and Integer Subsetting)

- Using integer subsetting is more effective when:
 - · You want to find the first (or last) TRUE.
 - You have very few TRUEs and very many FALSEs; a set representation may be faster and require less storage.
- **2. which()** conversion from boolean representation to integer representation

- Integer representation length: is always
 boolean representation length
- · Common mistakes:
- I. Use x[which(y)] instead of x[y]
- II. x[-which(y)] is not equivalent to
 x[!y]

Recommendation:

Avoid switching from logical to integer subsetting unless you want, for example, the first or last TRUE value

Subsetting with Assignment

 All subsetting operators can be combined with assignment to modify selected values of the input vector.

- 2. Subsetting with nothing in conjunction with assignment :
 - Why: Preserve original object class and structure

df1[] <- lapply(df1, as.integer)

Debugging, Condition Handling and Defensive Programming

Debugging Methods

1. traceback() or RStudio's error inspector

 Lists the sequence of calls that lead to the error

2. browser() or RStudio's breakpoints tool

- Opens an interactive debug session at an arbitrary location in the code
- options(error = browser) or RStudio's "Rerun with Debug" tool
 - Opens an interactive debug session where the error occurred
 - Error Options:

options(error = recover)

 Difference vs. 'browser': can enter environment of any of the calls in the stack

options(error = dump and quit)

- Equivalent to 'recover' for noninteractive mode
- Creates last.dump.rda in the current working directory

In batch R process :

```
dump_and_quit <- function() {
    # Save debugging info to file
    last.dump.rda
    dump.frames(to.file = TRUE)
    # Quit R with error status
    q(status = 1)
}
options(error = dump_and_quit)</pre>
```

In a later interactive session:

load("last.dump.rda") debugger()

Condition Handling of Expected Errors

1. Communicating potential problems to users:

I. stop()

- · Action : raise fatal error and force all execution to terminate
- Example usage: when there is no way for a function to continue

II. warning()

- Action: generate warnings to display potential problems
- Example usage: when some of elements of a vectorized input are invalid

III. message()

- · Action : generate messages to give informative output
- Example usage: when you would like to print the steps of a program execution

2. Handling conditions programmatically:

try()

Action: gives you the ability to continue execution even when an error occurs

II. tryCatch()

 Action: lets you specify handler functions that control what happens when a condition is signaled

```
result = tryCatch(code,
error = function(c) "error",
warning = function(c) "warning",
message = function(c) "message"
)
```

Use conditionMessage(c) or c\$message to extract the message associated with the original error.

Defensive Programming

Basic principle: "fail fast", to raise an error as soon as something goes wrong

- 1. stopifnot() or use 'assertthat' package check inputs are correct
- 2. Avoid subset(), transform() and with() these are non-standard evaluation, when they fail, often fail with uninformative error messages.
- 3. Avoid [and sapply() functions that can return different types of output.
 - Recommendation: Whenever subsetting a data frame in a function, you should always use drop = FALSE

A COMPARISON OF R AND PYTHON









A BIT OF HISTORY

R:

- a successor to S
- developed by statisticians as a 'statistical programming language'
- built-in data structures and functionality intended to make working with data easier
- gained prominence as a free and open source alternative to expensive statistical software

Python:

- created in the early 90's but popularized in the oo's
- intended to be easy to read, easy to understand and easy to learn, relative to other OOLs
- has a massive base of open-source modules









COMPARISON

R:

- technically object oriented, but this tends to be a bit hidden in practice
- lends itself to quick interactive scripting, data exploration
- has special built-in notation for statistical models
- has a special data type the data frame for handling datasets

Python:

- object oriented
- lends itself to writing structured, predesigned computer code.
- intended to be a general programming language
- designed to create code that is easy to read





A NOTE: VECTORIZATION IN INTERPRETED LANGUAGES

High-level interpreted languages are slower than low-level/compiled languages.

To get around this, these languages will sometimes hand off (behind the scenes) certain types of operations to functions written in lower-level languages (like C).

In order to take advantage of this, the R and Python, communities emphasize a certain programming strategy when using lists/vectors/arrays.

In particular, they avoid cycling through each item of a list, and instead often use special functions that **map** a chosen function or operation to every item in the list.

This can run counter to habits gained when learning other languages.









SO MANY PACKAGES/MODULES!

The strength of both R and Python lies in their many technical packages and modules.

These allow a programmer to implement very sophisticated functionality simply by making a few function calls.

Let's open the RPackagesDemo and PythonPackagesDemo notebooks to see some of this in action.

Available CRAN Packages By Name

A B C D E F G H I J K L M N O P Q R S T U V W X Y Z

<u>A3</u>	Accurate, Adaptable, and Accessible
-----------	-------------------------------------

Error Metrics for Predictive Models

Access to Abbyy Optical Character abbyyR

Recognition (OCR) API

Tools for Approximate Bayesian abc

Computation (ABC)

Data Only: Tools for Approximate abc.data

Bayesian Computation (ABC)

Array Based CpG Region Analysis ABC.RAP

Pipeline

ABCanalysis Computed ABC Analysis

abcdeFBA ABCDE FBA: A-Biologist-Can-Do-

Everything of Flux Balance Analysis

with this package

ABCoptim Implementation of Artificial Bee Colony

(ABC) Optimization

Approximate Bayesian Computational ABCp2

Model for Estimating P2

abcrf Approximate Bayesian Computation via

Random Forests





RESOURCES FOR PROGRAMMING

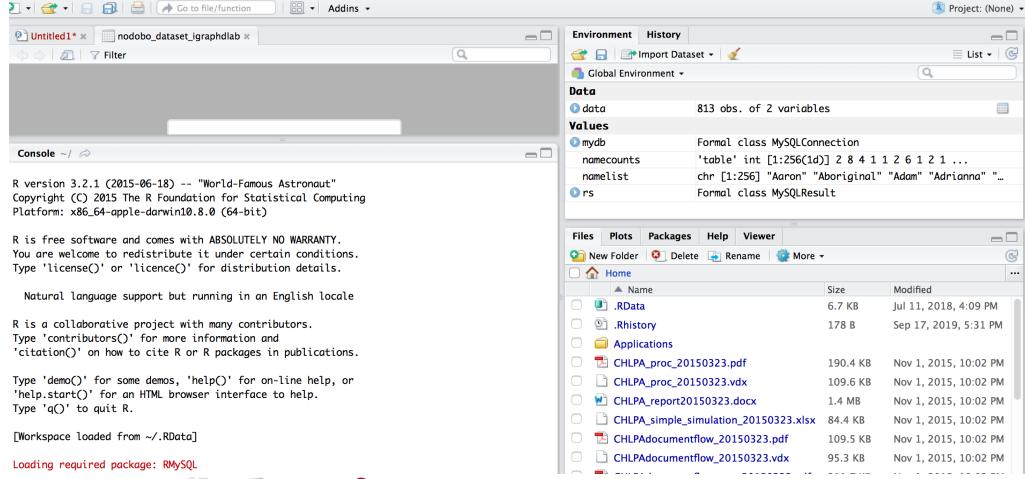








R STUDIO









JUPYTER NOTEBOOKS

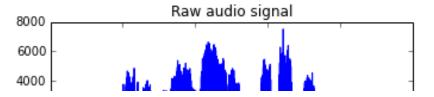
$$X_k = \sum_{n=0} x_n e^{-rac{2\pi i}{N}kn} \qquad k=0,\ldots,N-1$$

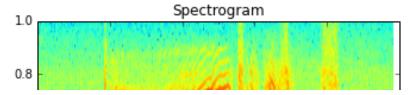
We begin by loading a datafile using SciPy's audio file support:

```
In [1]: from scipy.io import wavfile
        rate, x = wavfile.read('test mono.wav')
```

And we can easily view its spectral structure using matplotlib's builtin specgram routine:

```
In [2]: %matplotlib inline
        from matplotlib import pyplot as plt
        fig, (ax1, ax2) = plt.subplots(1, 2, figsize=(12, 4))
        ax1.plot(x); ax1.set title('Raw audio signal')
        ax2.specgram(x); ax2.set title('Spectrogram');
```











R NOTEBOOKS

You can use the provided R notebooks to:

- get a sense of what can be done
- gain exposure to many examples of the language syntax
- help you write your own code
- learn why the code works the way it does, and some theory behind the code

HCLUST()

Let's start by clustering the entire mtcars dataset, using the Euclidean distance metric, and plot the result. Hierarchical clustering is implemented in the cluster function hclust().

```
(hclustcars <- hclust(dist(mtcars)))</pre>
plot(hclustcars)
Call:
hclust(d = dist(mtcars))
```

Cluster method : complete Distance : euclidean

Number of objects: 32

The output of hclust gives us some information about the parameters being used to create the hierarchy. In this case the distance is Euclidean (as expected) and the cluster formation strategy (the linkage) is complete (these are the default settings).







ON-LINE RESOURCES

Stack Exchange/Stack Overflow/Cross Validated

Blogs (e.g. R Bloggers)

Official Sites:

- Python Software Foundation: https://www.python.org
- Comprehensive R Archive Network (CRAN): https://cran.r-project.org



CRAN Mirrors What's new? Task Views Search

The Comprehensive R Archive Network

Download and Install R

Precompiled binary distributions of the base system and contributed packages, Windows and Mac users most likely want one of these versions of R:

- Download R for Linux
- Download R for (Mac) OS X
- Download R for Windows

R is part of many Linux distributions, you should check with your Linux package management system in addition to the link above.









EXERCISES FOR HANDS-ON LEARNING









GETTING INTO PROGRAMMING

Develop/assess R skills by carrying out the following exercises (order unimportant):

You may choose to carry out each of the exercises separately, or to write a single program that carries out all of the individual exercises.

You will find much of the base code you need in the course notebooks and R markdown files, but you will need to tweak and add to this code to carry out the exercises. You will also find a lot of helpful information and code on the internet!







EXPRESSIONS, VARIABLES, DATA STRUCTURES, OPERATORS (1)

Create three variables and assign numerical values to each of these variables.

Then write one or more statements that carry out the following types of operations using these variables: addition, subtraction, multiplication, division, raising to a power.







EXPRESSIONS, VARIABLES, DATA STRUCTURES, OPERATORS (2)

Create three variables and assign string values to each of these variables.

Write a statement that joins the three strings into a single string. Write some code that prints the string.

Write some code that tests to see if a substring of your choosing is contained within the larger string.







EXPRESSIONS, VARIABLES, DATA STRUCTURES, OPERATORS (3)

Create three variables and assign lists to each of these variables. Join the three lists into a new list containing three distinct sub-lists (a list of three lists).

Create a list without sub-lists (all original list elements are part of a single larger list).

Create a fourth list by splitting this resulting list in half and assigning the second half of the list to a new variable.

Extract the last item of this list (it can either stay in the original list or be removed from it) and assign this element to a variable.







STATEMENTS, BLOCKS, CONTROL FLOW, LOGICAL OPERATORS

Write a statement that contains at least three nested blocks.

Use at least three of the following control flow options: if, if ... else, while, for, break, next, switch.









FUNCTIONS

Write a function that takes three arguments as input and returns one value.

Call the function with arguments of your choosing.









LIBRARIES/PACKAGES/MODULES

Execute the relevant command that shows a list of the packages that are currently installed in your R environment.

hint: use the internet, example notebooks and handouts to help you find the relevant command.

Use available documentation to determine what some of these do.

- hint: take a look at the R markdown files that are available you may notice some relevant information there.
- choose a module/package from this list and load the relevant package/module if necessary.

Write some code that uses functions and objects supplied by this package.









INPUTS/OUTPUTS (1)

Print to the standard output of R three sentences of your choosing, on three separate lines, using a single statement of code.







INPUTS/OUTPUTS (2)

Locate a comma separated values (CSV) file stored on your computer

(Hint - there should be a folder called Data in the main course directory).

Read this file into the notebook and store the results in one or more variables.







INPUTS/OUTPUTS (3)

Create a new file and write four lines in CSV format to this file.

In a separate statement, write four more lines to this existing file, without overwriting the original file.





INTERPRETERS/COMPILERS

Write enough code to generate at least five different error messages from R.

Copy these error messages into a markdown cell, and write a short note under each explaining the meaning of the error message, and how the code was fixed.





OPTIONAL EXERCISES

- Using a language of your choice, write a function that, when passed a dataset, reports 5 interesting pieces of information about the dataset. Load a dataset and run the function on this dataset.
- 2. Using a language of your choice, write two functions. The output of the first function should work as the input to the second function. The first function should read in a dataset and generate a subset of the dataset based on some chosen criteria. The second function should read in a dataset and provide summary data of some type for each column in the dataset. Load a dataset and run both functions on the dataset.





