Rotman

INTRO TO R

R Workshop - 1

March 20, 2022 Prepared by Jay Cao / TDMDAL

Website: https://tdmdal.github.io/r-tutorial-202122-winter/

Rotman School of Management
UNIVERSITY OF TORONTO

Plan for the 4 Sessions

- Overview (Today)

- Data Manipulation (tidyverse)
- Data Visualization (ggplot2)
- Modeling Workflow (tidymodels)

- Time Series & Finance Applications

Plan for Today

- Intro to Overview
  - What is R and what can R do?
  - Setup R
  - Motivation examples

- Overview of R programing and Data Science
  - Basics of R programming
  - Data science with R

- Learning Resources and Road Map

Goal for Today – Answer These Questions

- What's R?
- **What will I use R for?**

- How to Run R code?
- What are the R coding basics?
  - Data and programming structure
  - Data science workflow

- How to learn R (on my own)?
  - with get-started help from this mini-course

What's R?

- R = a language + an eco-system
  - A free and open-source programming language
  - An eco-system of many high-quality user-contributed libraries/packages
- In the past R is mostly known for its statistical analysis toolkits
- Nowadays R is capable of (and very good at) many other tasks
  - Tools that cover the whole data analysis workflow
  - Tools for web technology…

What can R do – Statistics & related

- Statistics & Econometrics
  - Regressions
  - Time series analysis
  - Bayesian inference
  - Survival analysis
  - …
- Numerical Mathematics
  - Optimization
  - Solver
  - Differential equations
  - …
- Finance
  - Portfolio management
  - Risk management
  - Option pricing
  - …
- …

See more R Empirical Finance Packages on R Task View - Finance

What can R do – Graphics (static ones)

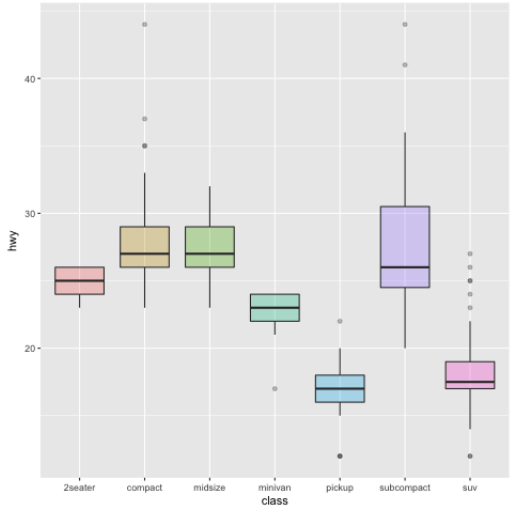

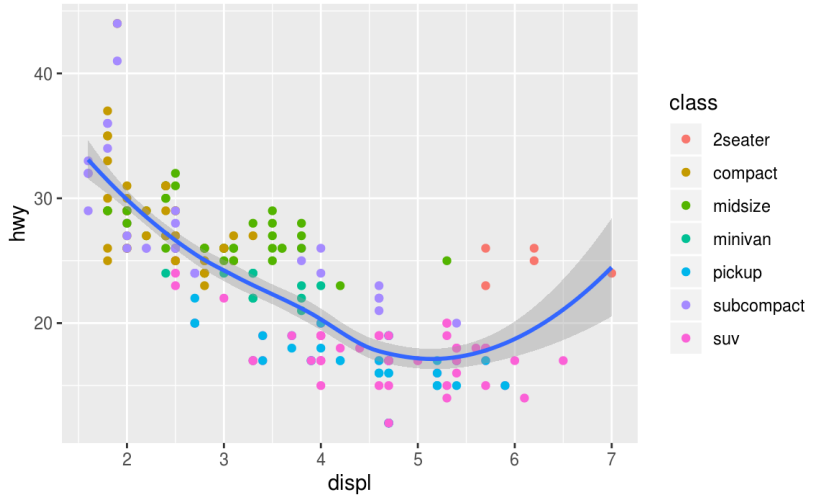

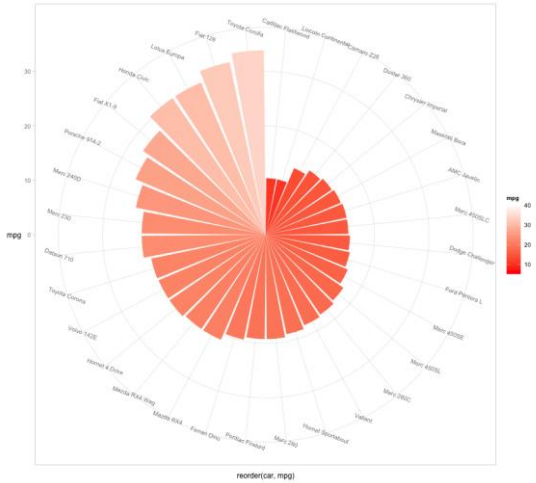

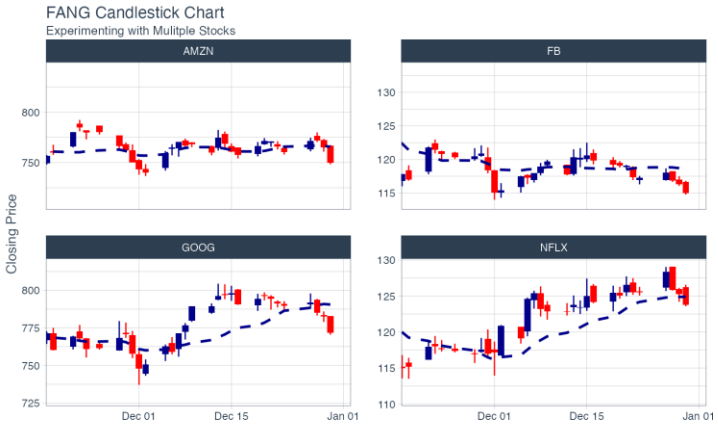

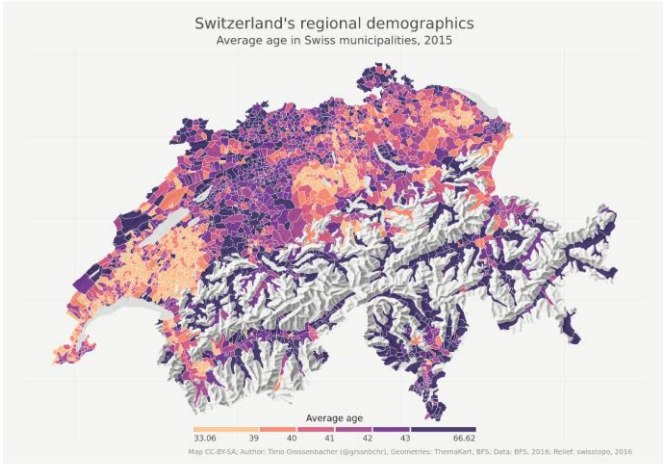

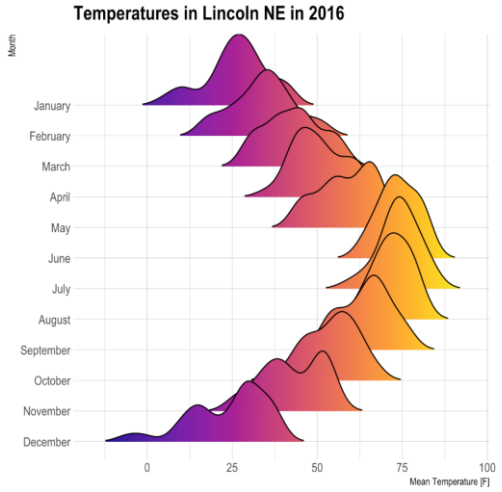

...

Ref. 1) https://www.r-graph-gallery.com/

2) https://timogrossenbacher.ch/2016/12/beautiful-thematic-maps-with-ggplot2-only/

What can R do – Graphics (dynamic ones)

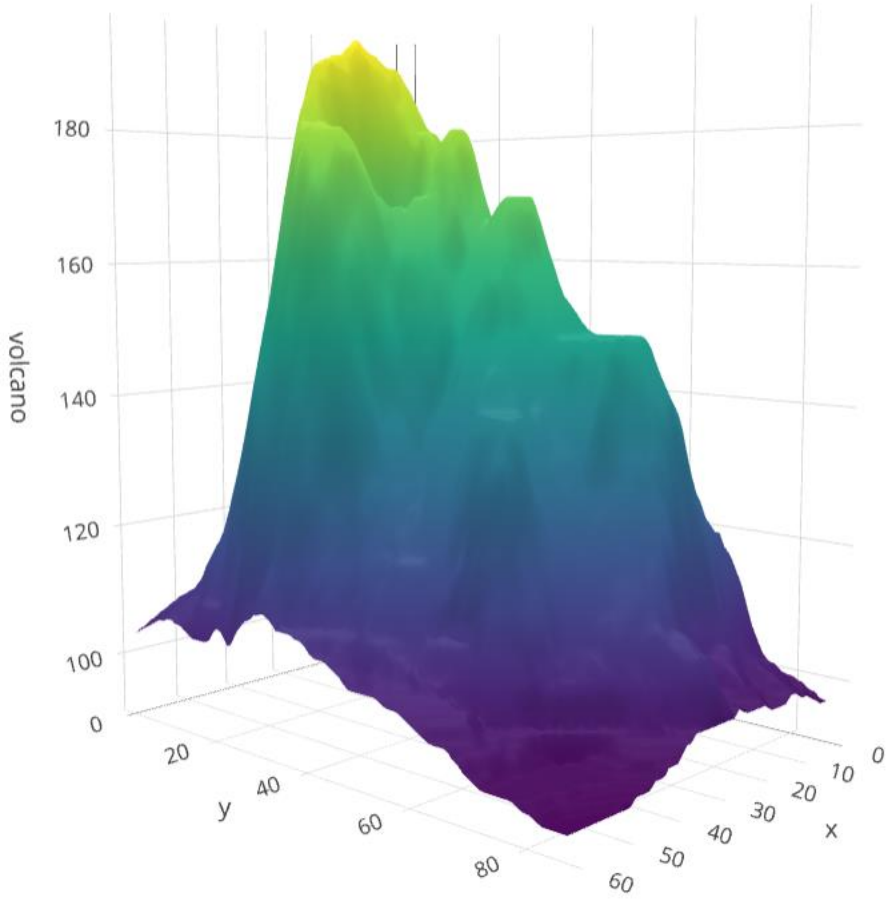

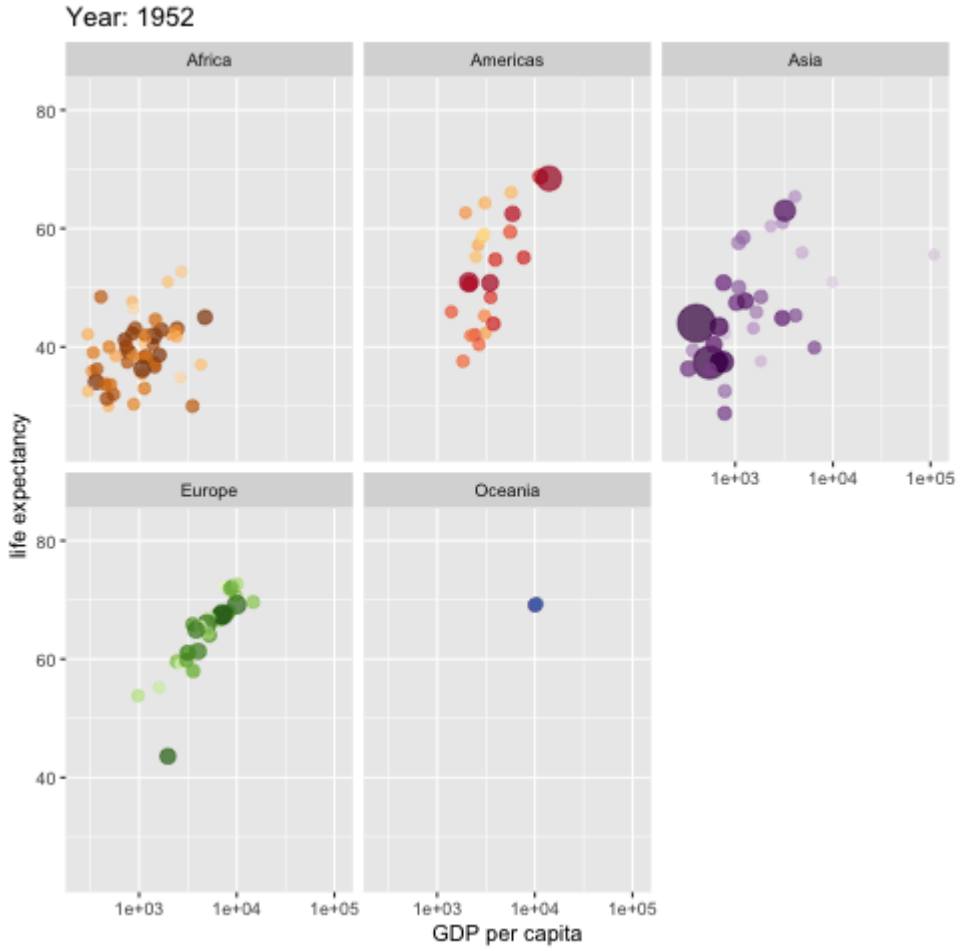

https://plot.ly/r/3d-surface-plots/;

https://github.com/thomasp85/gganimate;

What can R do – ML & NLP

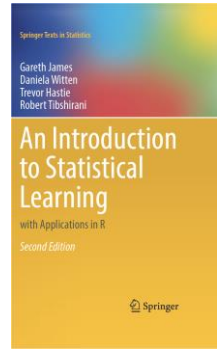

- Machine learning
  - Statistical learning (clustering, decision tree, etc.)
    - An Introduction to Statistical Learning (with Applications in R)
  - Deep learning (neural networks)
    - Interface to Keras and Tensorflow (via reticulate, an R to Python interface)
    - Torch for R (natively from R; similar as PyTorch in Python)

- Natural language processing (e.g., tidytext, topicmodels)

1. See more R Machine Learning Packages on R Task View - ML & Statistical Learning
2. See more R Natural Language Processing Packages on R Task View - NLP

What can R do – Web & Reporting

- Web technology
  - Web scraping (e.g., rvest)
  - API wrapper (e.g., Twitter: rtweet; bigquery: bigrquery; Quandl: Quandl)
  - Shiny web app (https://shiny.rstudio.com/)

- Reporting
  - R Markdown (write reports, slides, blogs, books, etc. See a gallery here.)

- … (see **R Task View** for more)

R vs Excel and BI Tools vs Python

- Excel & Business Intelligence (BI) Tools (e.g., Tableau, Power BI, etc.)
  - 2-D tables as basic data structure
  - Good UI (User Interface) and minimum programming
  - Limited modeling tools
  - Not easy to reproduce an analysis (because it's hard to store UI clicks)
  - Not flexible enough for complicated analysis problems, i.e., problems with
    - Many data cleaning steps/pipelines
    - Many different models to try

- Python
  - Python is more general purpose, R is more specialized in statistical analysis
  - R is much easier to learn (in my opinion)

Why learn R (What can R do **for You**)?

- Beyond Excel Data Analysis
  - I wish Excel could…
- Automate boring repeating tasks
  - e.g., daily data collection from different sources, weekly dashboard update
- Prototype ideas
  - e.g., a novel trading strategy, a new credit risk model
- Really, find anything that interests you and use R…

Plan for Today

- Intro to Overview
  - What is R and what can R do?
  - Setup R
  - Motivation examples
- Overview of R programing and Data Science
  - Basics of R programming
  - Data science with R
- Learning Resources and Road Map

Setup R

- R & RStudio on your computer (most of you should choose this one)
  - Install R (https://www.r-project.org/)
  - Install RStudio (https://rstudio.com/products/rstudio/download/)

- R & RStudio in the Cloud (run R without installation)
  - RStudio Cloud (https://rstudio.cloud/)
  - UofT JupyterHub RStudio (https://jupyter.utoronto.ca/hub/login)

- R & **Notebook** in the Cloud (run R without installation)
  - UofT JupyterHub Notebook (https://jupyter.utoronto.ca/hub/login)
  - Google Colab (https://colab.to/r)

What's RStudio?

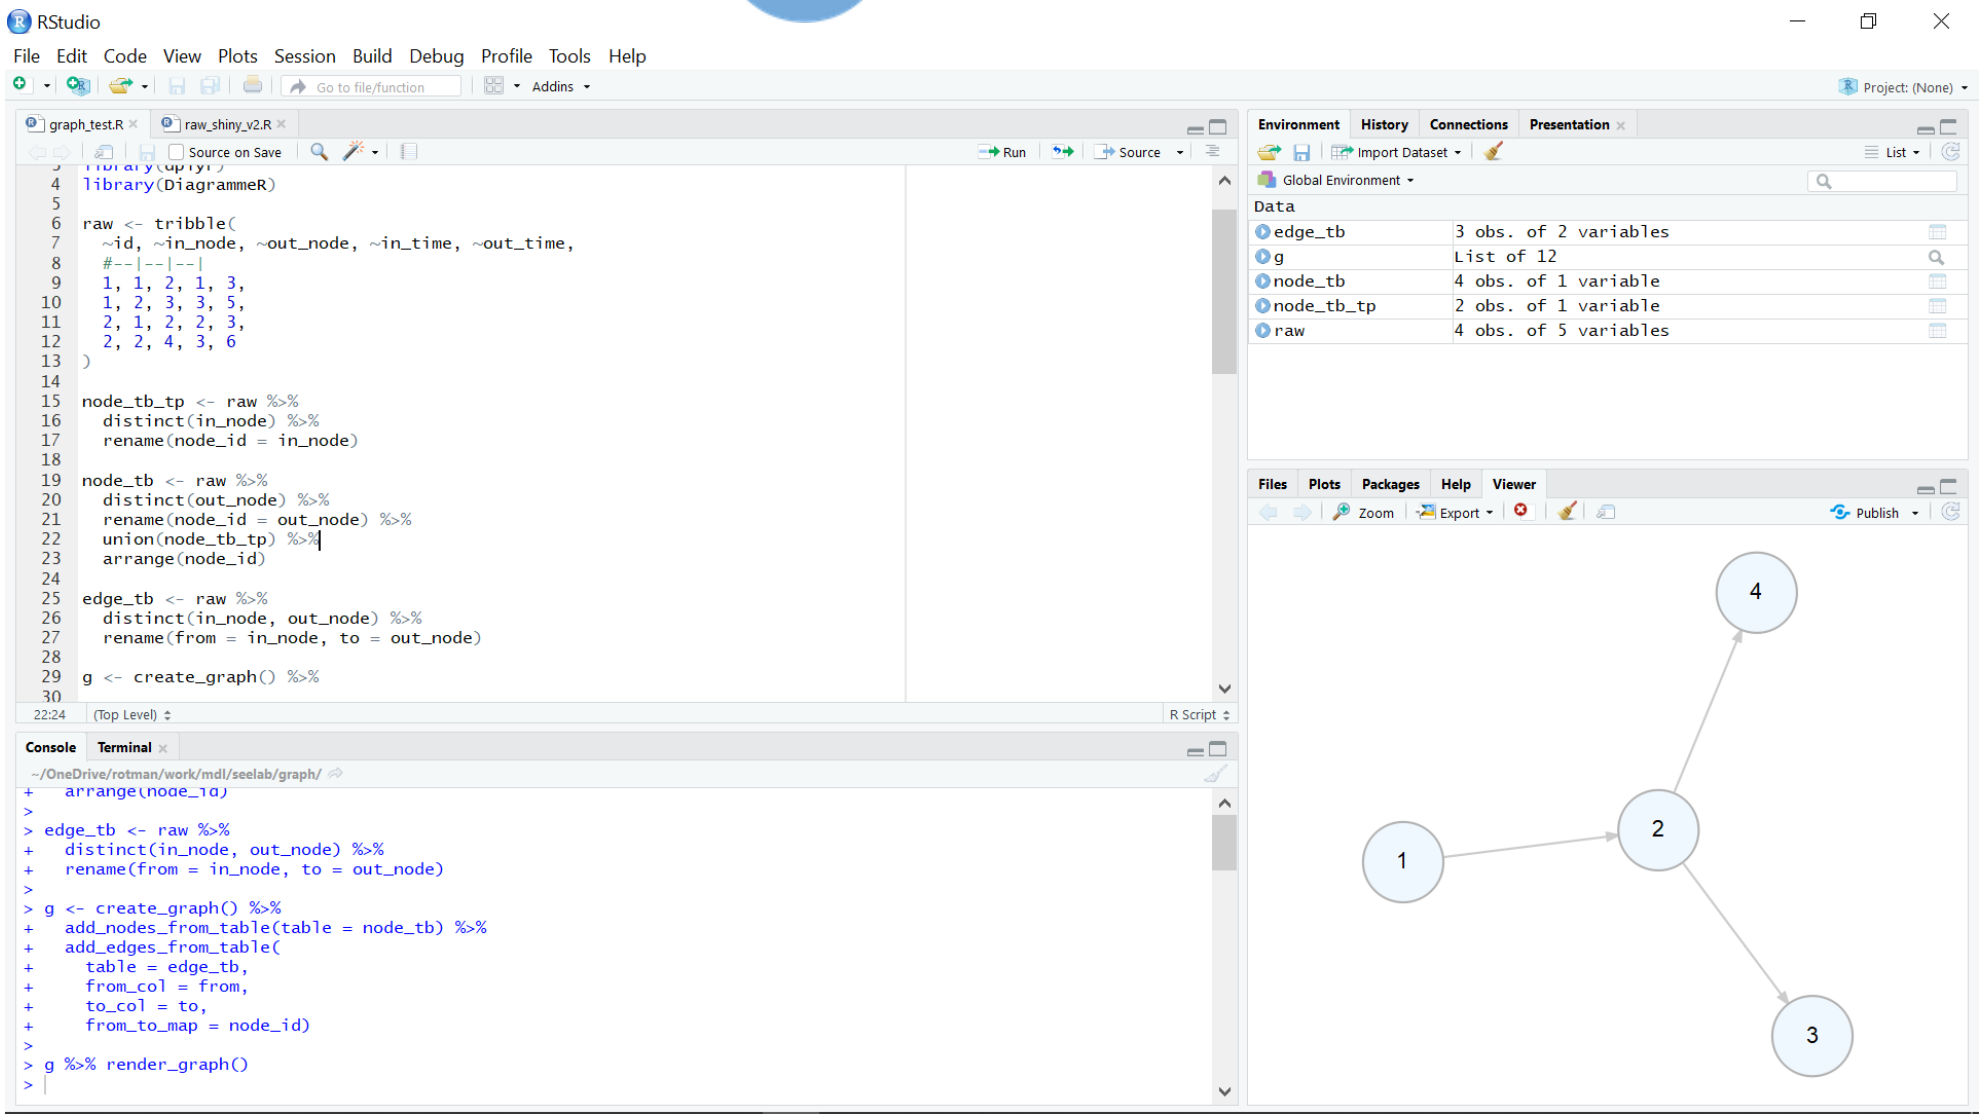

RStudio Cloud

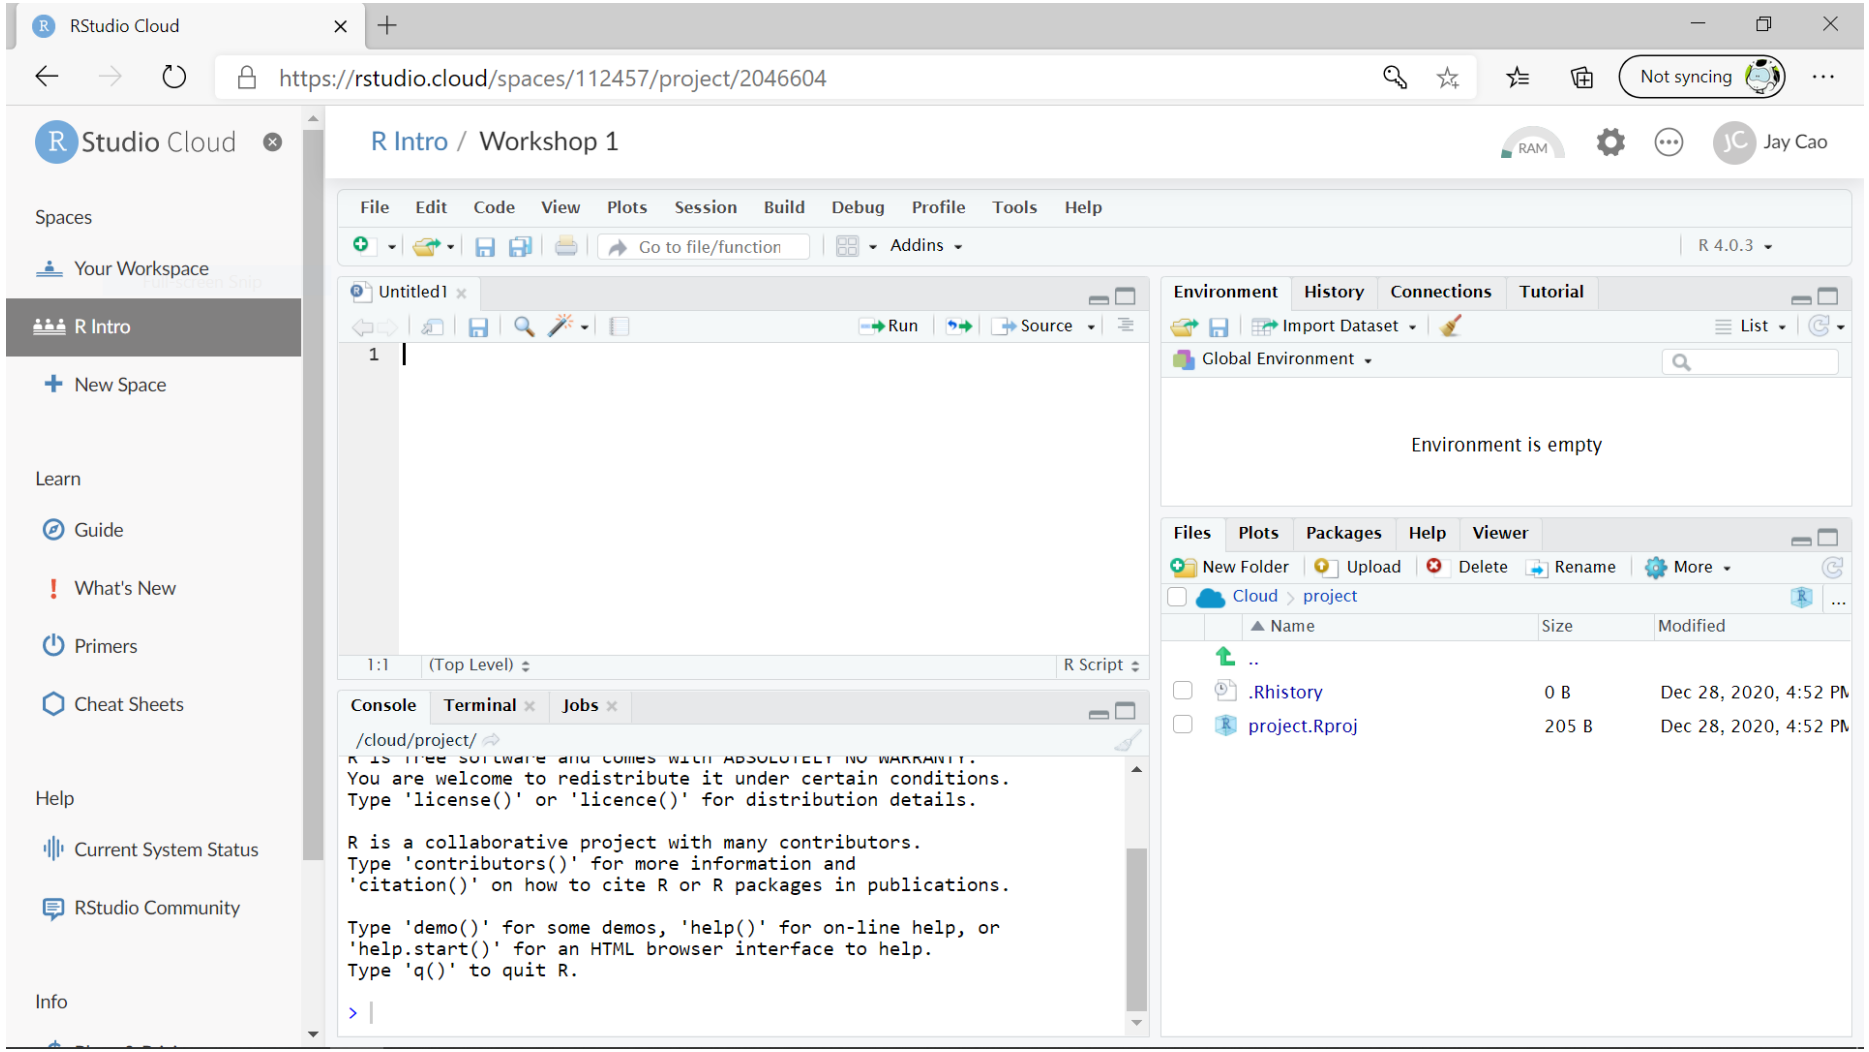

RStudio at UofT Jupyterhub

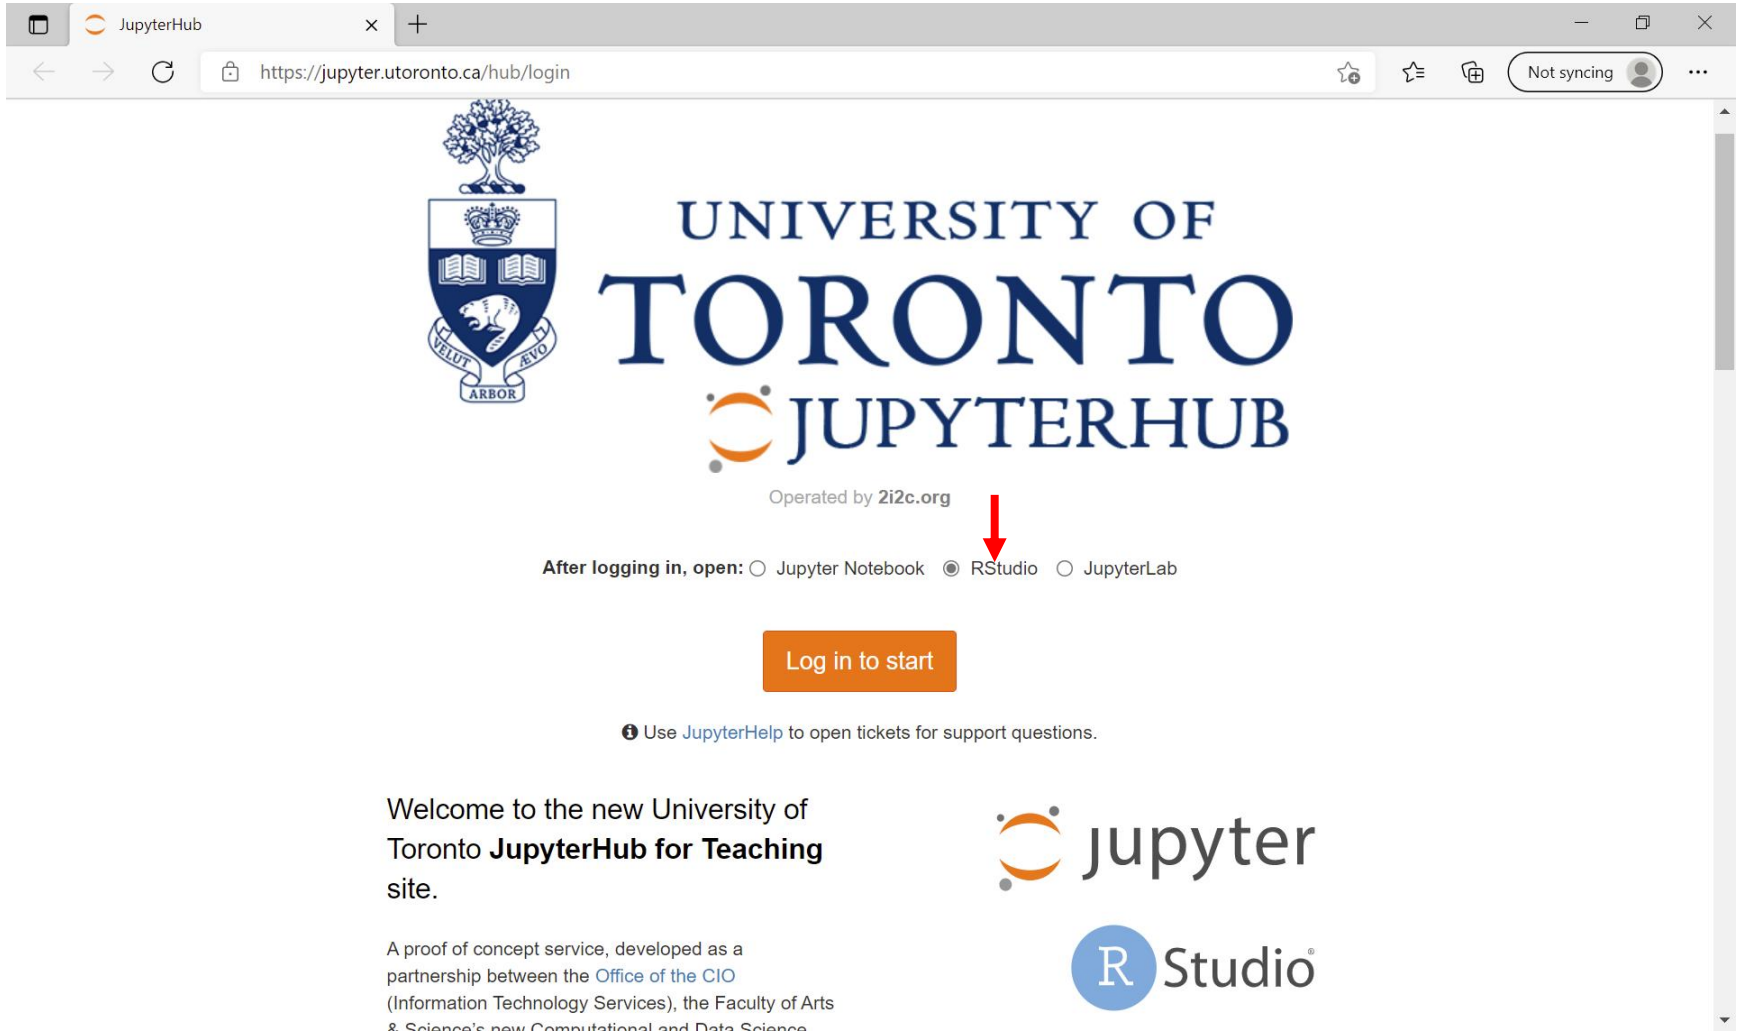

R Notebook in Google Colab

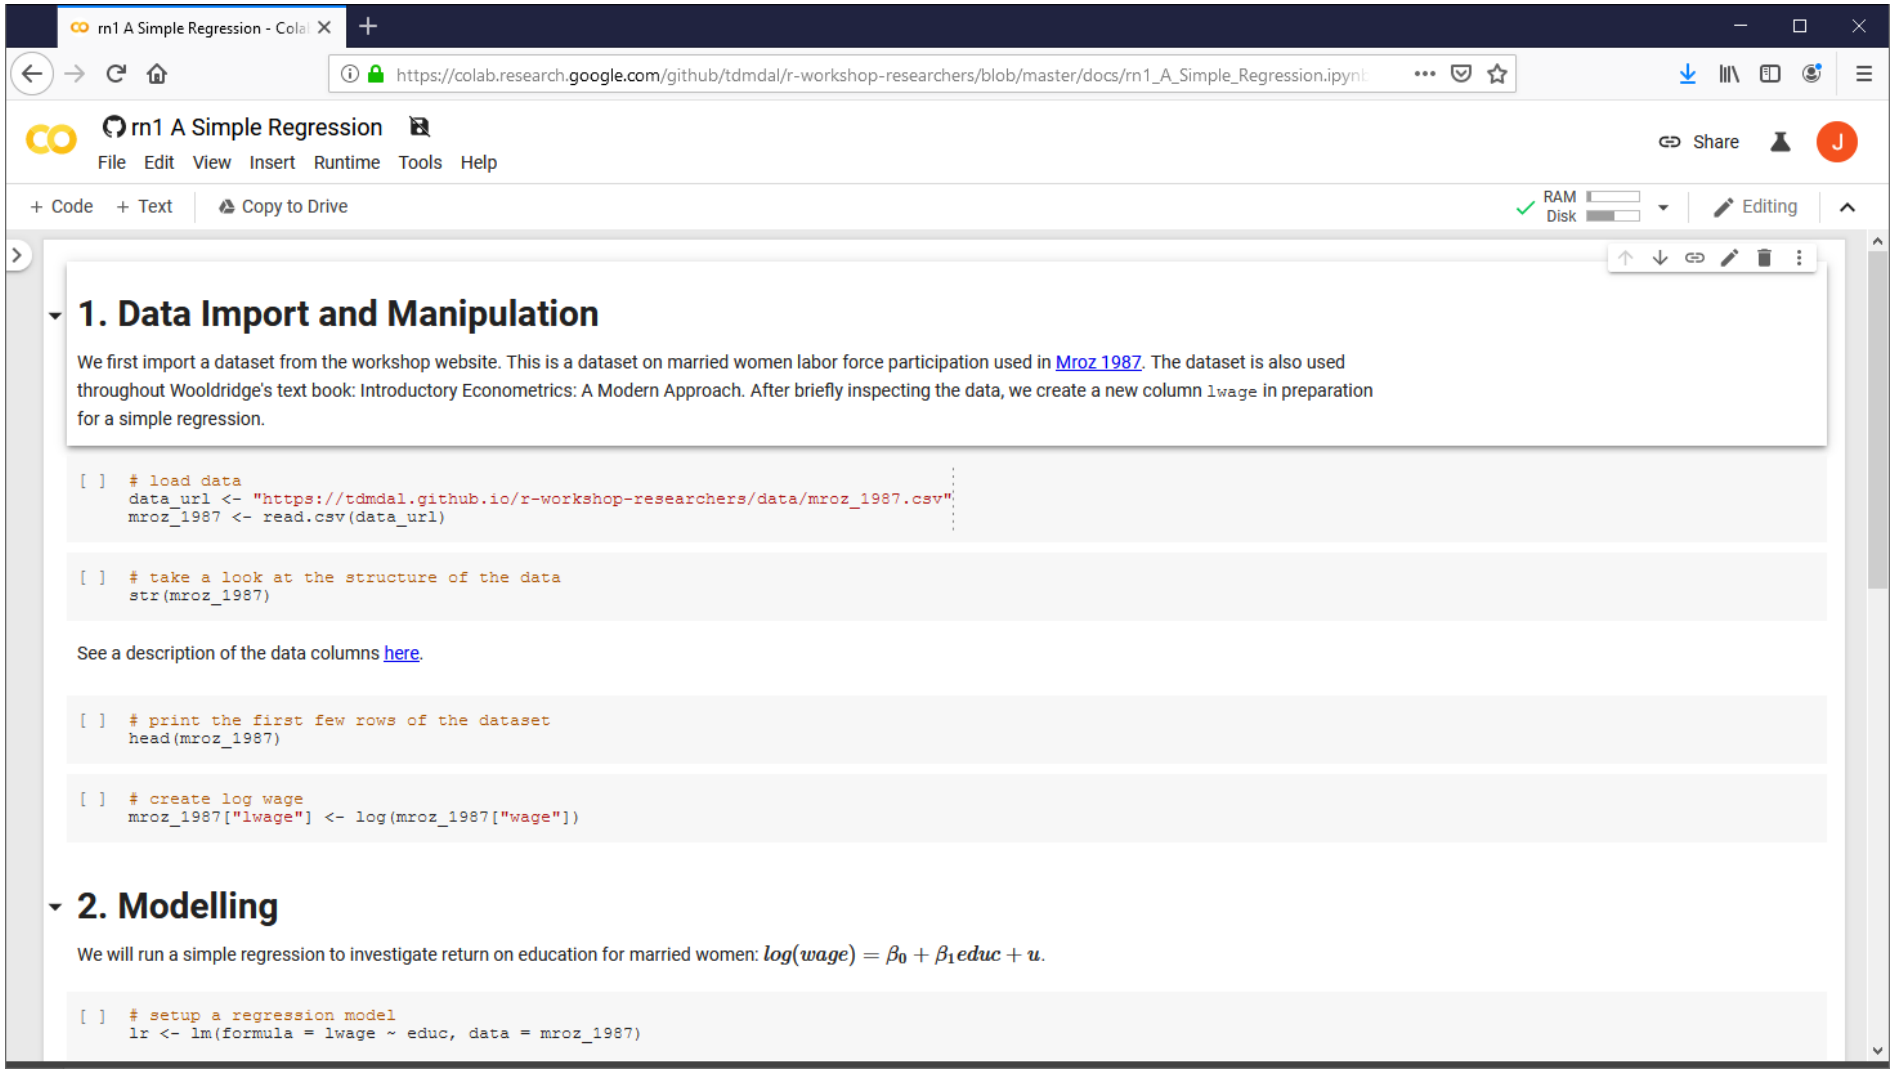

R Notebook at UofT Jupyterhub

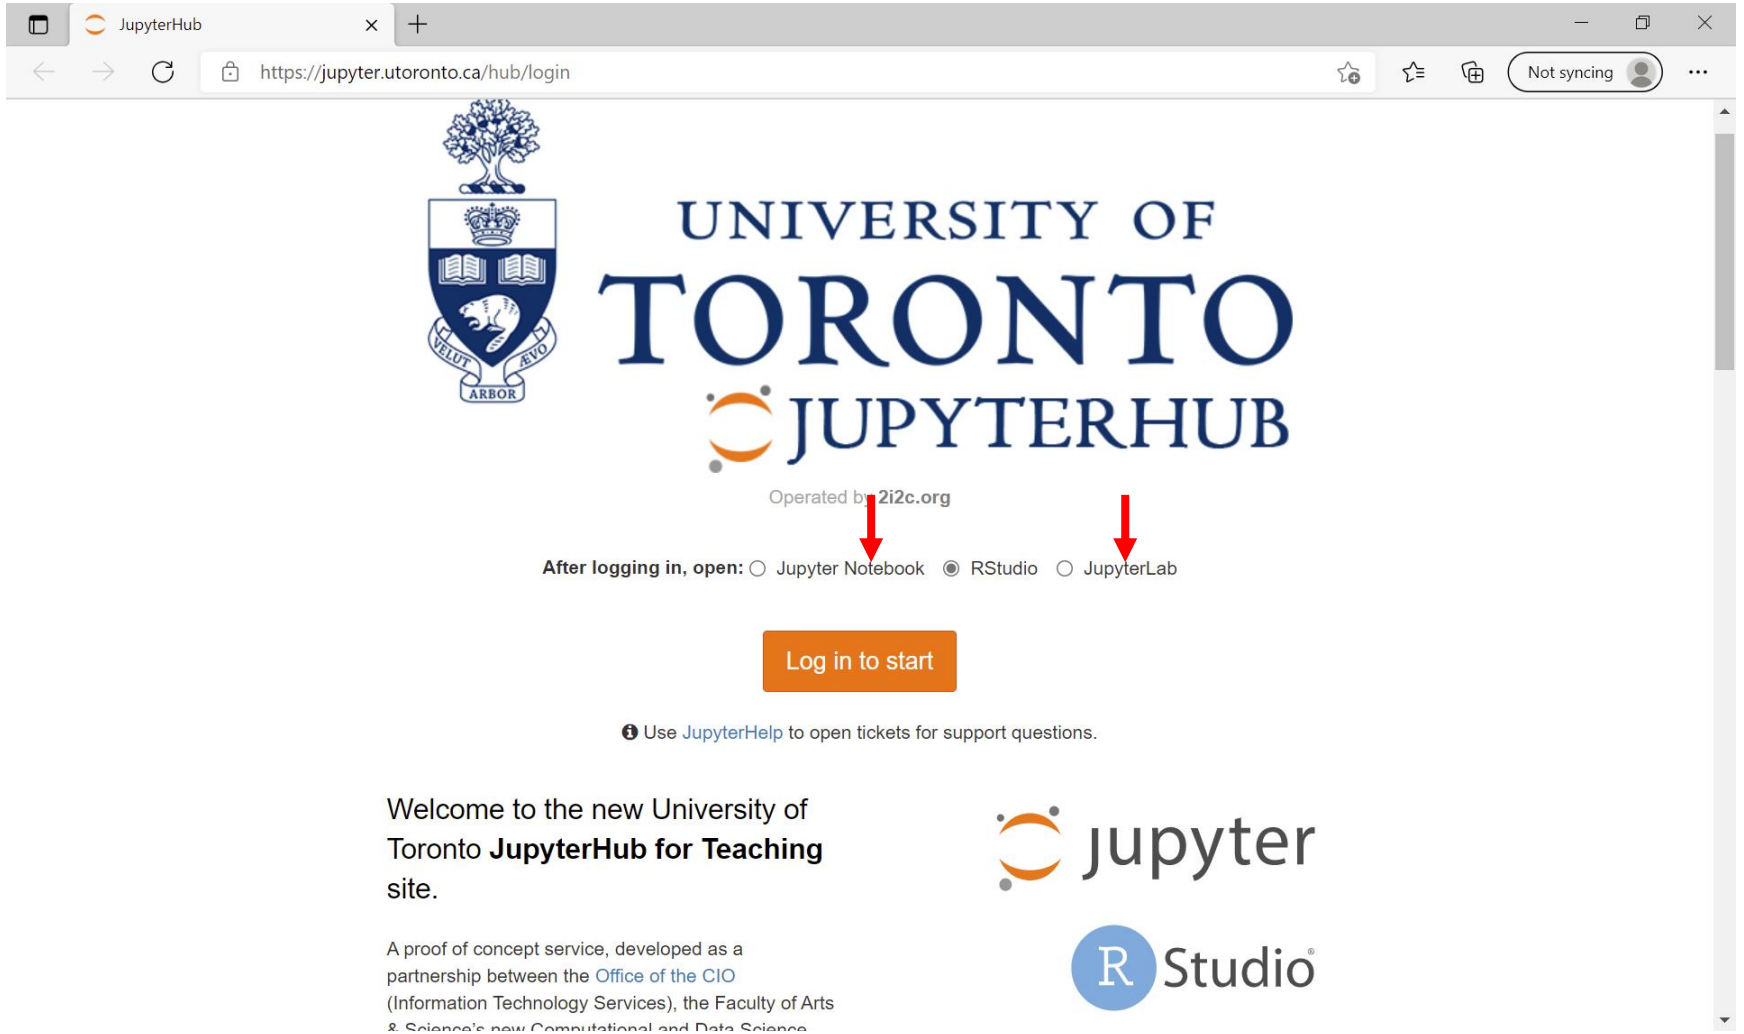

Plan for Today

- Intro to Intro
  - What is R and what can R do?
  - Setup R
  - Motivation examples

- Overview of R programing and Data Science
  - Basics of R programming
  - Data science with R

- Learning Resources and Road Map

A Few Examples

- Analyze portfolio performance

PerformanceAnalytics Package

- Look for trends in R community through Twitter

- Recognize handwritten digits - an example of deep learning

A Few Examples: What to Look For

- Focus on analysis workflow (by reading the code comments)
  - Import and manipulate data
  - Model data
  - Report and visualize results
- Don’t focus on R syntax
- Do notice everything is done in a sequential way
  - no conditional branching or looping

Plan for Today

- Intro to Intro

- Overview of R programing and Data Science
  - Basics of R programming
    - Expression & Assignment
    - Data Structure
    - Programming Structure (control flow & function)
    - Turn ideas into code
  - Data science with R

- Learning Road Map and Resources

Expression and Assignment

```
# expression
2 + sqrt(4) + log(exp(2)) + 2^2


# assignment
x <- 3
y <- (pi == 3.14)
```

R Data Structure - Overview

	Homogeneous	Heterogeneous
1-d	**Atomic vector**	**List**
2-d	Matrix	**Data frame**
n-d	Array	

http://adv-r.had.co.nz/Data-structures.html

R Data Structure - Overview

	Homogeneous	Heterogeneous
1-d	**Atomic vector**	**List**
2-d	Matrix	**Data frame**
n-d	Array	

http://adv-r.had.co.nz/Data-structures.html

Atomic Vectors

```
# create R vectors
vec_character <- c("Hello,", "World!")

vec_integer <- c(1L, 2L, 3L)

vec_double <- c(1.1, 2.2, 3.3)

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

Hello,	World!

1	2	3

1.1	2.2	3.3

TRUE	TRUE	FALSE

List

```
# create an R list
l1 <- list(
  1:3,
  "a",
  c(TRUE, FALSE, TRUE),
  c(2.3, 5.9)
)
```

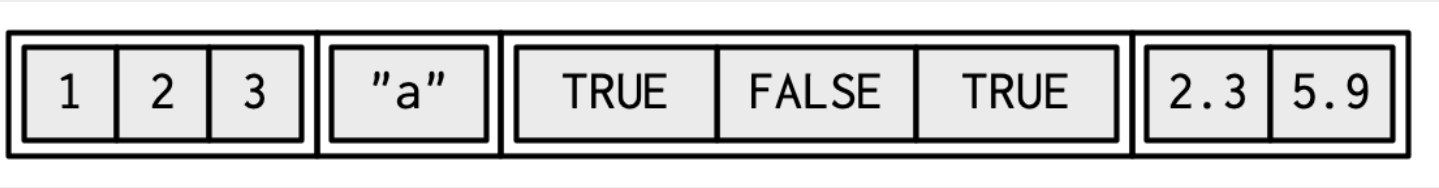

ref. https://adv-r.hadley.nz/vectors-chap.html#list-creating

Data Frame

```
# create a data frame
df1 <- data.frame(
  x = 1:3,
  y = letters[1:3],
  z = c(1.1, 2.2, 3.3)
)
```

x	y	z
1	"a"	1.1
2	"b"	2.2
3	"c"	3.3

Data Frame

```
# create a data frame
df1 <- data.frame(
  x = 1:3,
  y = letters[1:3],
  z = c(1.1, 2.2, 3.3)
)
```

x	y	z
1	"a"	1.1
2	"b"	2.2
3	"c"	3.3

Data Frame

```
# create a data frame
df1 <- data.frame(
  x = 1:3,
  y = letters[1:3],
  z = c(1.1, 2.2, 3.3)
)
```

x	y	z
1	"a"	1.1
2	"b"	2.2
3	"c"	3.3

A Cousin to Data Frame - Tibble

```
# load tibble library (part of tidyverse lib)
library(tibble)

# create a tibble
tb1 <- tibble(
  x = 1:3,
  y = letters[1:3],
  z = c(1.1, 2.2, 3.3)
)
```

x	y	z
1	"a"	1.1
2	"b"	2.2
3	"c"	3.3

https://r4ds.had.co.nz/tibbles.html#tibbles-vs.data.frame

Programming Structure: Control Flows

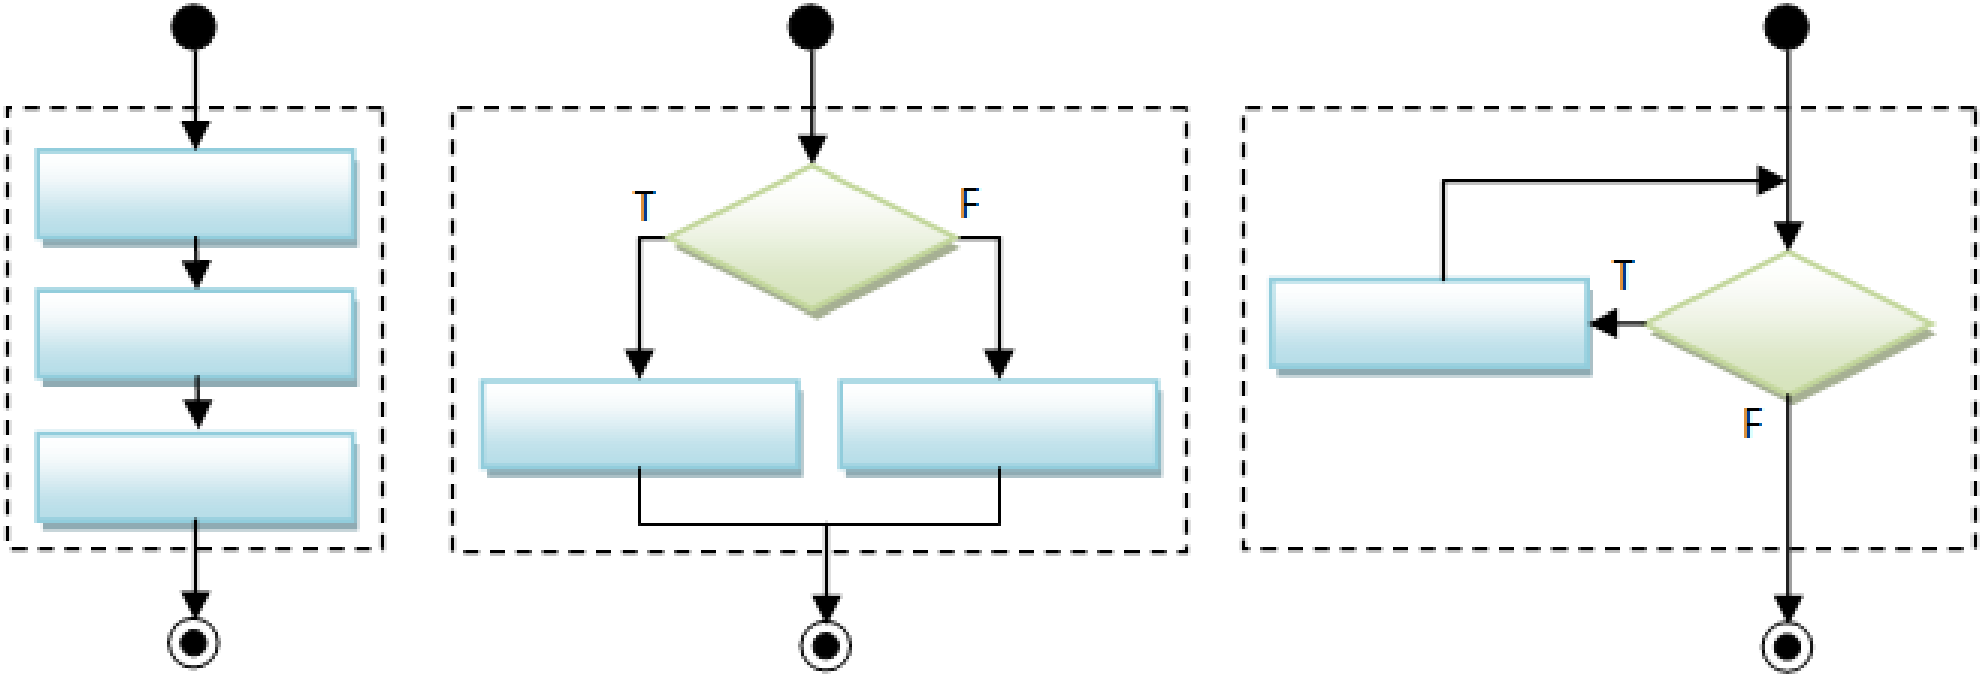

Today

Learn on your own (See Appendix)

Sequential

- Example: Sum of Squares

$$\sum_{t=1}^{3} t^2$$

```
# sum of squares
t <- 1:3
y <- sum(t^2)
print(y)
```

Sequential

- Example: Sum of Squares

$$\sum_{t=1}^{3} t^2$$

```
# sum of squares
t <- 1:3
y <- sum(t^2)
print(y)
```

t	1	2	3

Sequential

- Example: Sum of Squares

$$\sum_{t=1}^{3} t^2$$

```
# sum of squares
t <- 1:3
y <- sum(t^2)
print(y)
```

t	1	2	3
t^2	1	4	9
sum(t^2)	14		

Programming Structure: Functions

- What's a function
  - a logical block of code
  - input -> output
- Why write functions
  - Reusability
  - Abstraction
  - Maintainability
- Example: $\sum_{t=1}^{\boldsymbol{n}} t^2$

```
# sum of squares from 1 to n
ss <- function(n) {
  t <- 1:n
  sum(t^2)
}

# calling the ss() function
print(ss(2))
print(ss(3))
```

Programming Structure: Functions

- What's a function
  - a logical block of code
  - input -> output
- Why write functions
  - Reusability
  - Abstraction
  - Maintainability
- Example: $\sum_{t=1}^{n} t^2$

```
# sum of squares from 1 to n
ss <- function(n) {
  t <- 1:n
  sum(t^2)
}

# calling the ss() function
print(ss(2))
print(ss(3))
```

Programming Structure: Functions

- What's a function
  - a logical block of code
  - input -> output
- Why write functions
  - Reusability
  - Abstraction
  - Maintainability
- Example: $\sum_{t=1}^{n} t^2$

```
# sum of squares from 1 to n
ss <- function(n) {
  t <- 1:n
  sum(t^2)  # return(sum(t^2))
}

# calling the ss() function
print(ss(2))
print(ss(3))
```

Turn Ideas into Code

- Solve problems using code: three main ingredients
  - Data Structure + Programming Structure + **Algorithm** (sorting, searching, optimization, etc.)
  - Examples
    - Sort a list of integers
    - Generate and solve Sudoku puzzles
    - Implement and backtest a trading strategy

- For us (data analysis tasks), in most cases, we solve problems by
  - Using algorithm implemented by other peoples (i.e., functions from R packages)
  - Combine algorithms (and data & programming structures) to achieve our goal (still not easy; need practices to write good code.)

Plan for Today

- Intro to Intro

- Overview of R programing and Data Science
  - Basics of R programming
  - Data science with R
    - A Typical data analysis workflow
    - Choice of R packages
    - An example: regression analysis

- Learning Road Map and Resources

Data Science/Analysis Workflow

- Use this workflow to organize your thoughts and code

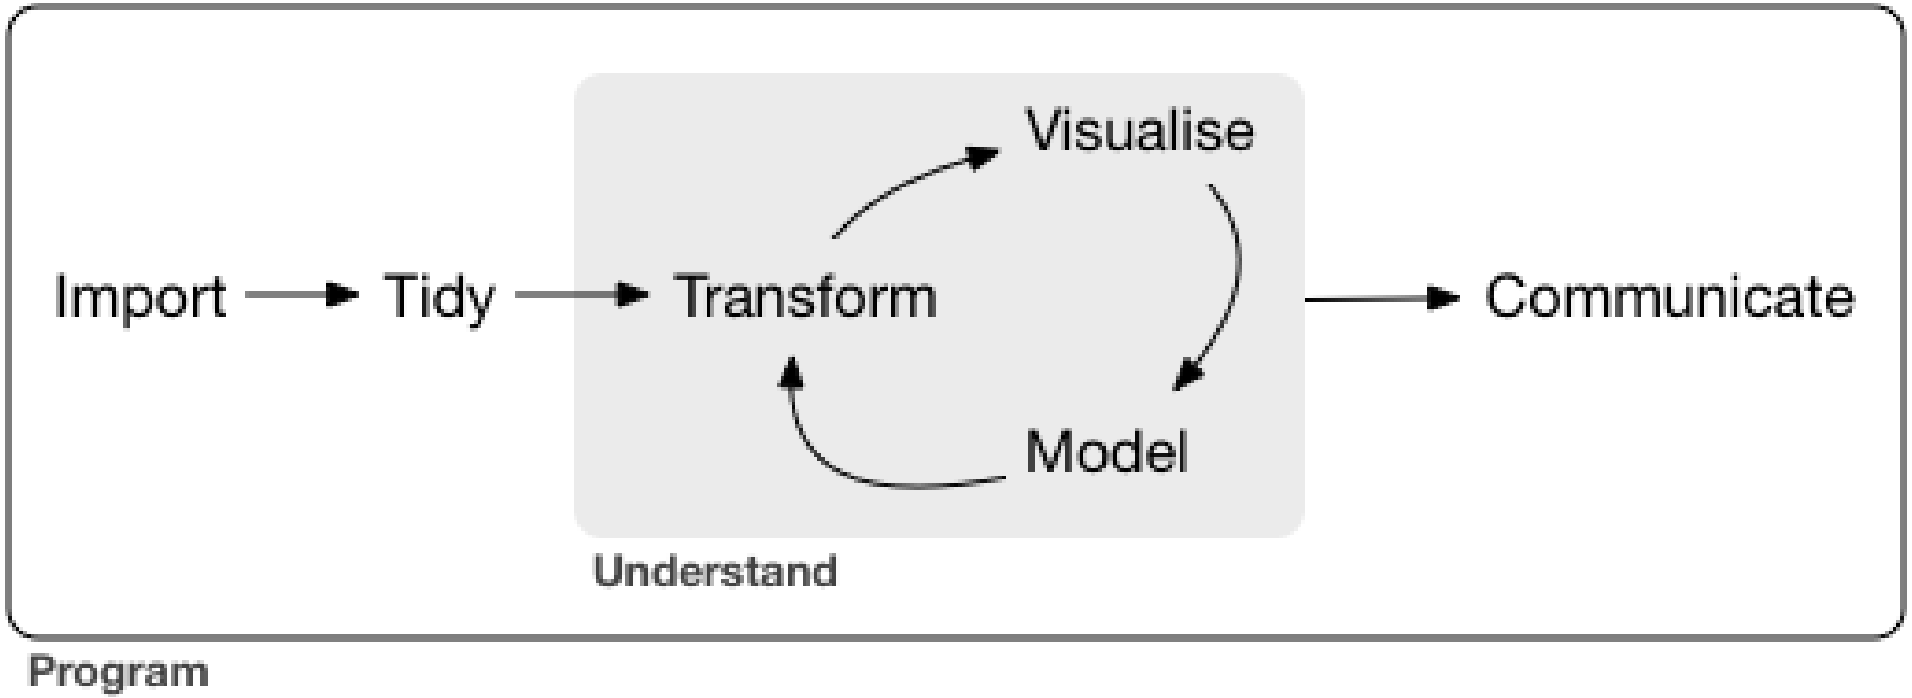

https://r4ds.had.co.nz/introduction.html

An Example: Housing Price & Clean Air

- Manipulate data
  - Load data
  - Create new columns
  - Filter columns and rows
- Build models
  - Multiple linear regressions
- Report and graph
  - Build a publication-ready table for regression results

```
Obs:   506

1. price      median housing price, $
2. crime      crimes committed per capita
3. nox        nitrous oxide, parts per 100 mill.
4. rooms      avg number of rooms per house
5. dist       weighted dist. to 5 employ centers
6. radial     accessibiliy index to radial hghwys
7. proptax    property tax per $1000
8. stratio    average student-teacher ratio
9. lowstat    % of people 'lower status'
```

Many Ways to Achieve the Same Goal

- The “pure” R way if possible
  - Mostly use functions/packages in R standard library (those shipped with R)
    - for structuring and manipulating data
    - for modeling if possible (e.g. regressions)
  - An example of a simple linear regression

- The “modern” way
  - Use specialized packages to manipulate data and assist modeling tasks
    - Data are stored in improved data structures (in most cases still compatible with base R data structure)
  - What we will focus on

R Packages: Many choices, which one to use

- Often, a task can be achieved using functions in different libraries
  - R is open and extensible!

- Example: load a csv file to a data frame (or its variations)
  - Use `read.csv()` function from the `utils` library in Base R
  - Use `read_csv()` function from the `readr` library
  - Use `fread()` function from the `data.table` library
  - Use `vroom()` from the `vroom` library

R Packages: Many choices, which one to use

- Start with the one most people use
- Choose one that is well maintained
  - check document, github, etc. for last update date
  - packages maintained by companies (e.g., RStudio Co.) or academic teams
- Choose one that suits your task

Great Choice for Data Science Work

- Tidyverse
  - "an opinionated collection of R packages designed for data science"
  - "All packages share an underlying design philosophy, grammar, and data structures."
  - Handle data manipulation, visualization, and much more
  - an eco-system: many package developers started to follow tidyverse principles too

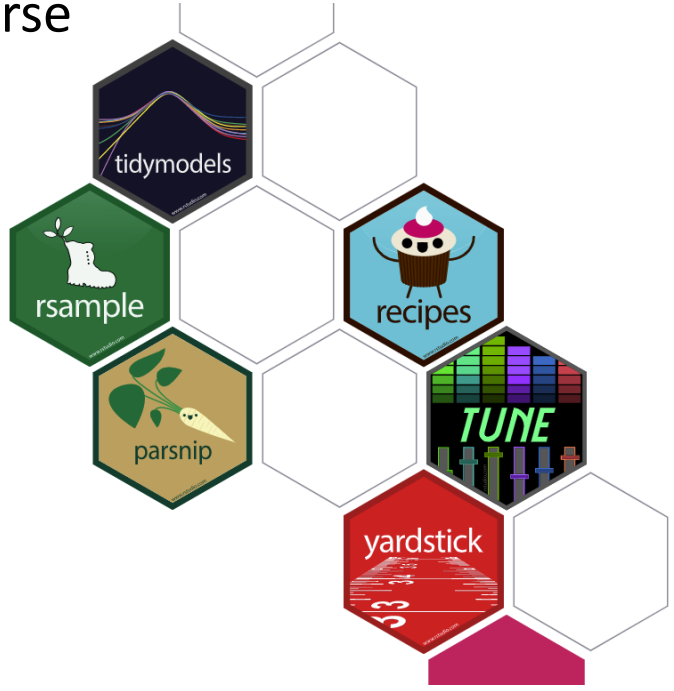

- Tidymodels
  - "a collection of packages for modeling and machine learning using tidyverse principles"
  - Manage modeling process but does not do modeling itself

Our Choice: the Regression Example

- Manipulate data (`tidyverse` eco-system)
  - Load data (`read_csv()` from the `readr`)
  - Create new columns (`mutate()` from `dplyr`)
  - Filter columns and rows (`select()` and `filter()` from `dplyr`)
- Build models
  - Multiple regression (`lm()` from `stats` library in R base)
- Report and graph
  - Build a publication-ready table (`huxreg()` from `huxtable` library)

Using R packages/libraries

- Install an R library (only need to install a library once)

```
install.packages("library_name")
```

- Load an R library (before you use a library)

```
library(library_name)
```

- CRAN (The Comprehensive R Archive Network)
  - CRAN Task Views

Load a CSV file

- `read_csv()` from the `readr`

```
read_csv(file)
```

```
e.g. hprice <- read_csv("hprice.csv")
```

- More about `read_csv()`
- More about `readr`

Load Data – Other file formats and sources

- readxl for Excel sheets
- haven for SPSS, Stata and SAS data
- jsonlite for JSON
- xml2 for XML
- httr for web APIs
- rvest for web scraping
- DBI for connecting to DataBase engine
- …

Load Data – Financial Dataset

- `tq_get()` from tidyquant library
  - collect financial and economic data from many online sources
    - Yahoo Finance, FRED, Quandl, Tiingo, Alpha Vantage, Bloomberg

- simfinapi library
  - download financial statements – balance sheet, cash flow and income statement – and adjusted daily price of stocks through the simfin project

- a few others (try to look for them yourselves…)

A Quick Detour - Load Data from Web API

- Modern Web API: An example using Yahoo Finance
  - Yahoo Finance API Specification
  - How does a call look like
    - Apple on Yahoo Finance from the web, https://finance.yahoo.com/quote/AAPL/history?p=AAPL
    - An API call for AAPL recommendation trend data, https://query2.finance.yahoo.com/v11/finance/quoteSummary/AAPL?formatted=true&modules=recommendationTrend&corsDomain=finance.yahoo.com
  - R packages to get data from Yahoo finance (API call wrappers), quantmod; tidyquant

- An R package to deal with raw Web API call, httr
  - You can build your own Yahoo Finance R package using it

A Quick Detour – Get Data From Bloomberg

- Bloomberg is working hard on its Web APIs
  - https://www.bloomberg.com/professional/blog/modernizing-capital-markets-tech-with-web-apis/

- Meanwhile, you can get data using its traditional server API
  - To connect, need authentication (i.e. hostname, userid, etc.)

- R packages
  - Rblpai (a wrapper of Bloomberg Server API)
  - Tidyquant's `tq_get(x, "rblpapi")` (a wrapper of Rblpai package)

Data Manipulation: `dplyr` basics

- Filter observations (rows): **`filter()`**

`filter(my_dataframe, condition1, …)`

```
e.g., hprice_reg <- filter(hprice, price > 20000)
```

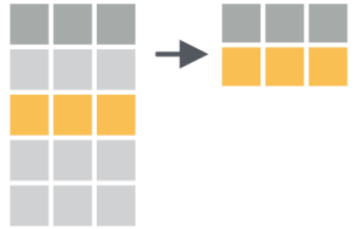

- Select variables (columns): **`select()`**

`select(my_dataframe, var1, …)`

```
e.g., hprice_reg <- select(hprice_reg, lprice, rooms)
```

- Create new variables: **`mutate()`**

`mutate(my_dataframe, new_var1 = expression1, …)`

```
e.g., hprice_reg <- mutate(hprice_reg, lprice = log(price))
```

Ref. dplyr data wrangling cheat sheet

Data Manipulation: Data Pipe (%>%)

```
hprice_reg <- filter(hprice, price > 20000)
hprice_reg <- mutate(hprice_reg, lprice = log(price))
hprice_reg <- select(hprice_reg, lprice, rooms)
```

```
hprice_reg <- hprice %>%
  filter(price > 20000) %>%
  mutate(lprice = log(price)) %>%
  select(lprice, rooms)
```

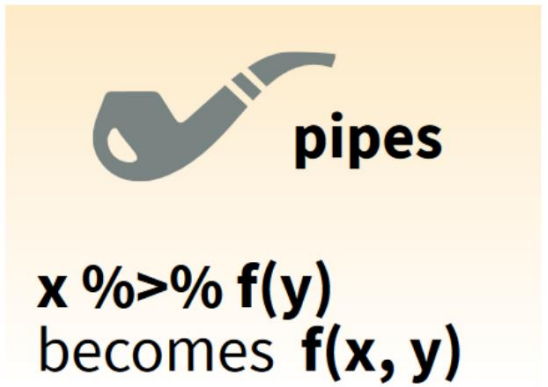

Ref. dplyr data wrangling cheat sheet

Regression

- Multiple regressions: `lm()` from `stats` library in base R

```
my_model <- lm(y ~ x1 + x2, data)
```

$$Y_i = \beta_0 + \beta_1 X_1 + \beta_2 X_2 + \epsilon_i$$

```
my_model <- lm(y ~ x1 + x2 + I(x1 * x2), data)
```

$$Y_i = \beta_0 + \beta_1 X_1 + \beta_2 X_2 + \beta_3 X_1 X_2 + \epsilon_i$$

- Regression result summary: `summary()`

Ref. https://faculty.chicagobooth.edu/richard.hahn/teaching/FormulaNotation.pdf

Report

- Summary table
  - Summary for lm(): `summary(my_model)`

- publication-ready table: `huxreg()` from `huxtable` library

```
huxtable(my_model1, my_model2, …)
```

Ref. https://hughjonesd.github.io/huxtable/huxreg.html

Report - Publish to The Web (1) / R Markdown

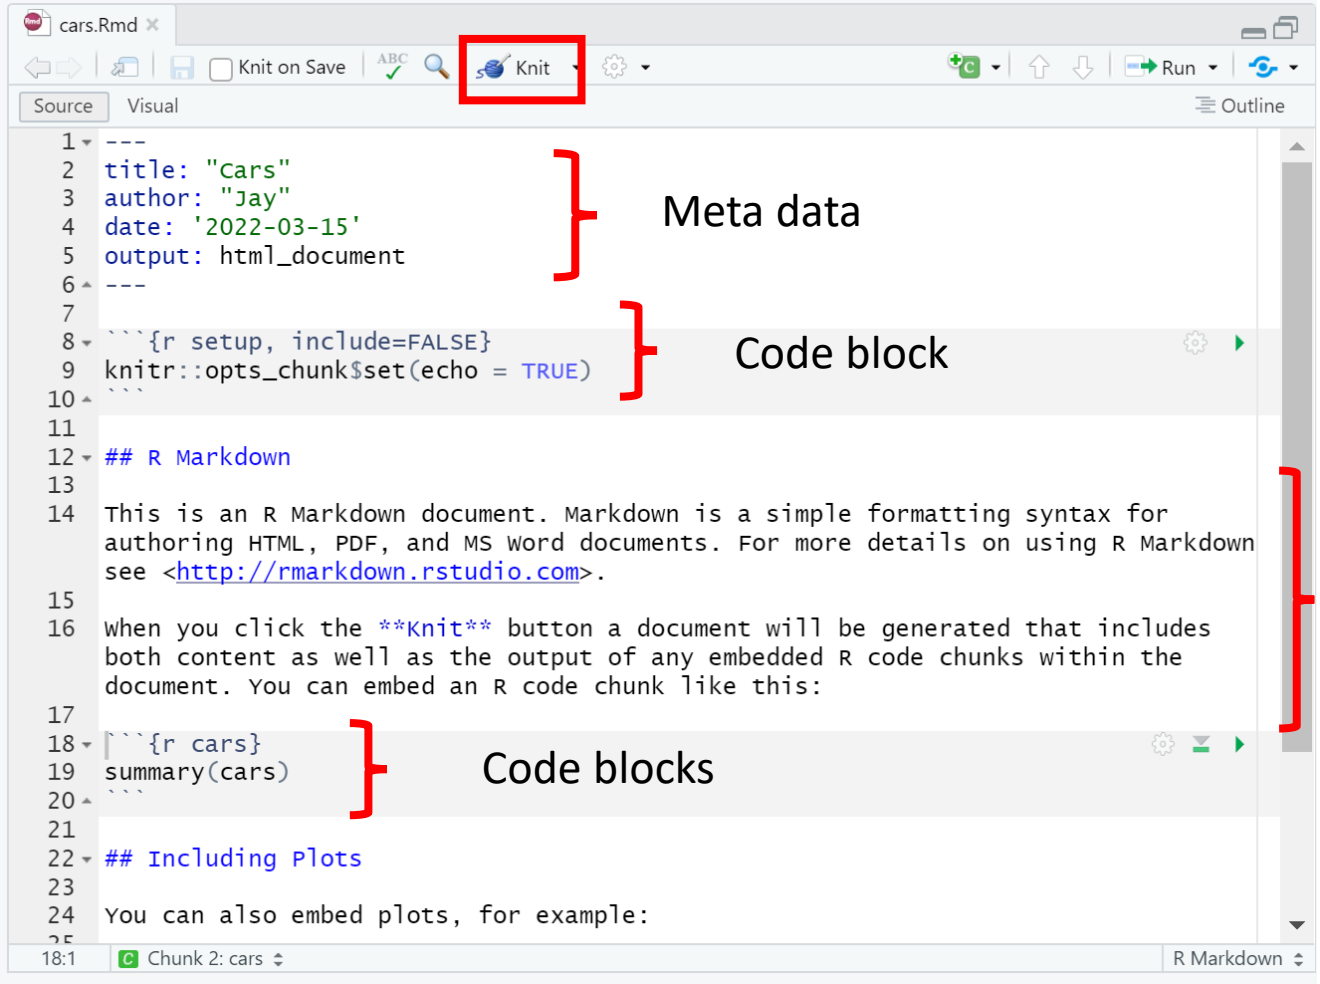


```
---
title: "Cars"
author: "Jay"
date: '2022-03-15'
output: html_document
---

```{r setup, include=FALSE}
knitr::opts_chunk$set(echo = TRUE)
```

R Markdown

This is an R Markdown document. Markdown is a simple formatting syntax for
authoring HTML, PDF, and MS Word documents. For more details on using R Markdown
see <http://rmarkdown.rstudio.com>.

When you click the **Knit** button a document will be generated that includes
both content as well as the output of any embedded R code chunks within the
document. You can embed an R code chunk like this:

```{r cars}
summary(cars)
```

Including Plots

You can also embed plots, for example:
```

Ref. 1) https://rmarkdown.rstudio.com/lesson-1.html; R Markdown Cheat Sheet

Report - Publish to The Web (2) / R Markdown

- Knit button for rendering
  - knitr package

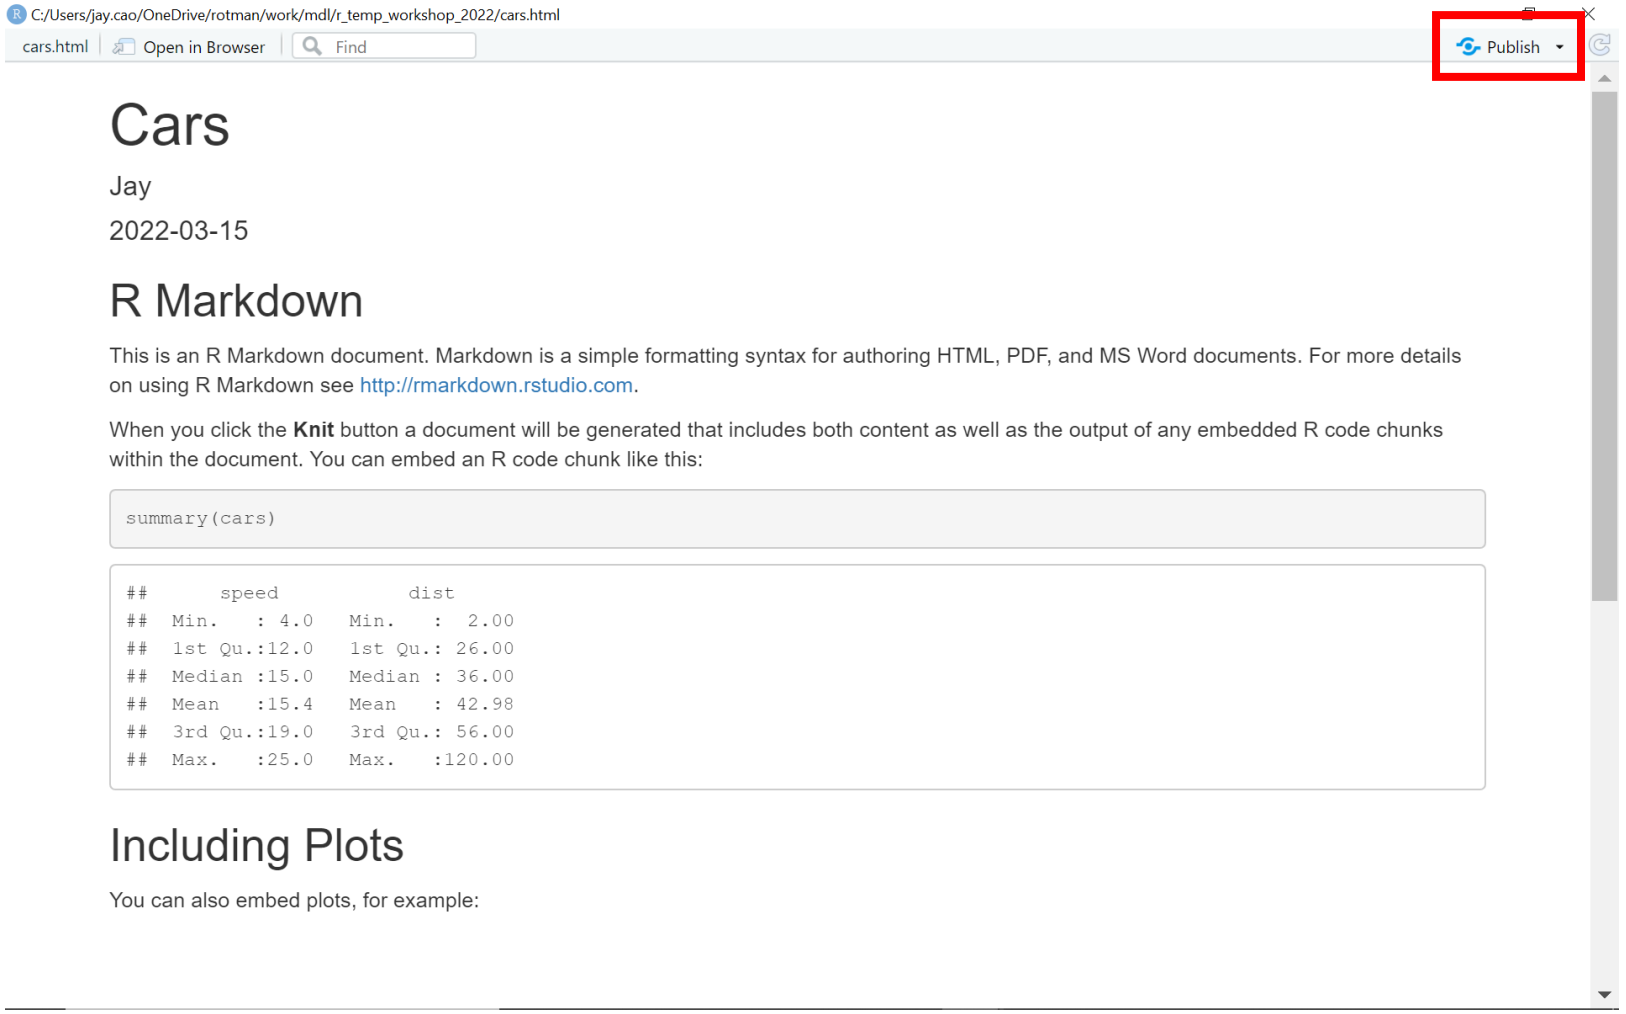

Report - Publish to The Web (3) / RPubs

- Publish button for publishing
- RPubs
  - Publish and share your RMarkdown output
- Demo

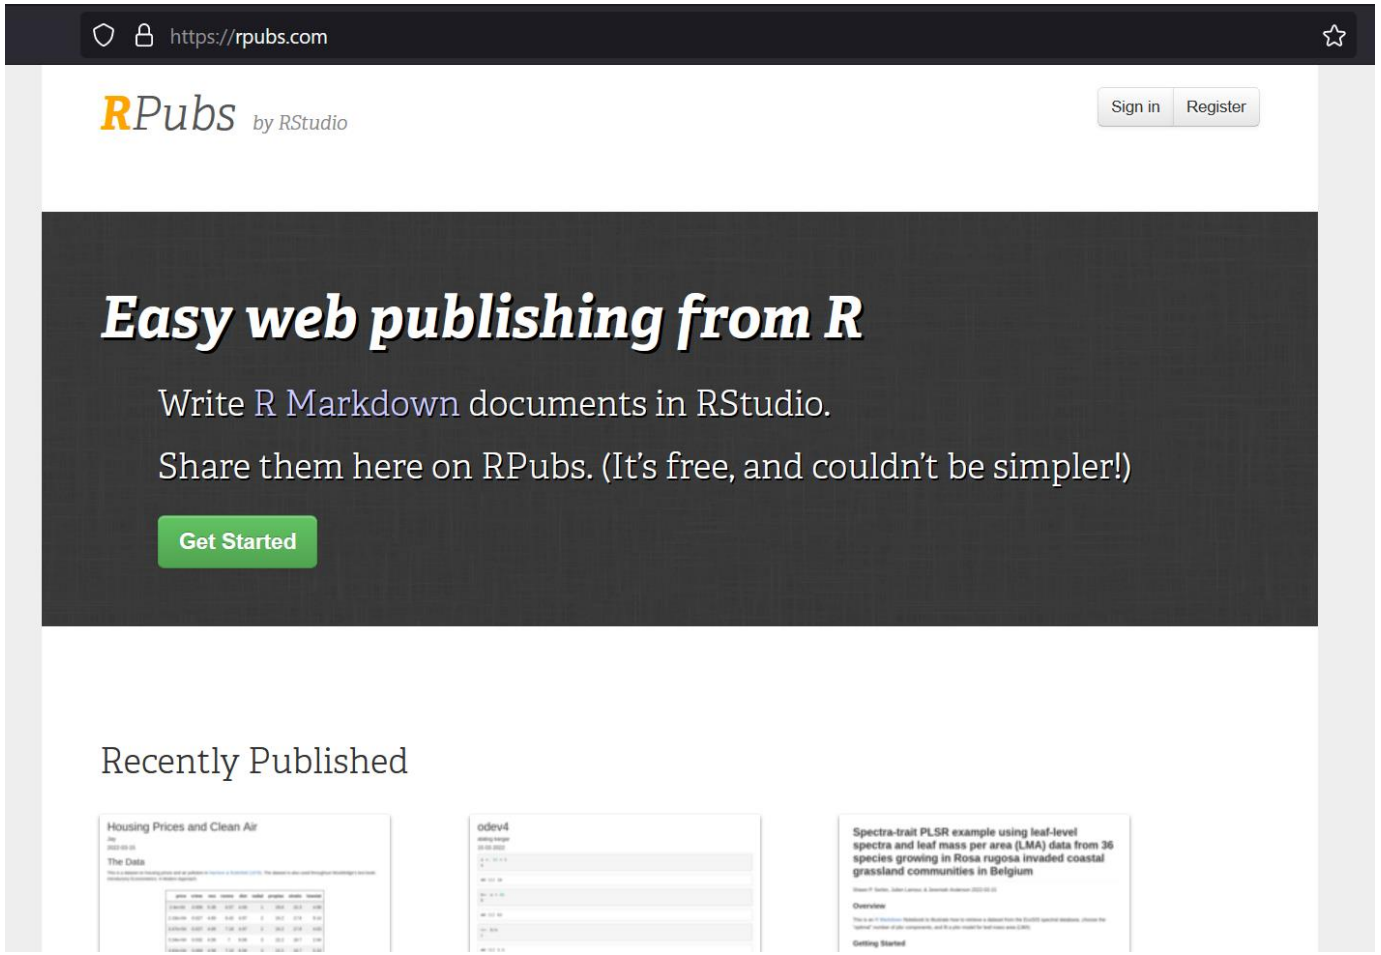

Plan for Today

- Intro to Intro

- Overview of R programing and Data Science
  - Basics of R programming
  - Data science with R

- Learning Resources and Road Map

Learning Road Map (Three Free Books)

- Step 1. Basic R programming skills (Never programmed before? Start here.)
  - Data and programming structure; how to turn an idea into code;
  - Book: Hands-On Programming with R

- Step 2. R Data Science skills
  - Data wrangling, modeling, and visualization/reporting; Best practice;
  - Book: R for Data Science

- Step 3. Take your R Skill to the next level
  - Book: Advanced R

Other free books check bookdown.org often

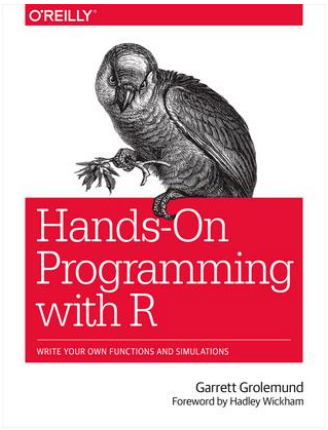

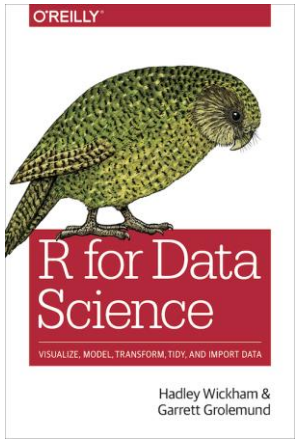

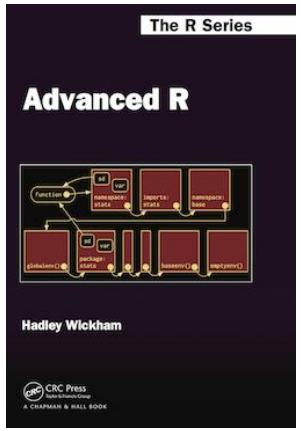

Learning Approach

- Learn the underlying principles
  - e.g., why organize data in a certain way
- Learn best practices
  - follow a consistent analysis workflow

Free Learning Resource

- RStudio Education
  - Choose Your Learning Paths
- RStudio Video Resources Site
- More free R books? Check bookdown.org often
- Coursera: Search R and learn
  - free for UofT students (mostly always free if you just audit the courses)
- Twitter (a few seeds: #rstat, @hadleywickham, @WeAreRLadies)

Appendix

- Programming Structure Continued
  - Conditional
  - Iteration

Conditional (if…else…)

```
if (cond) {
  # run here if cond is TRUE
} else {
  # run here if cond is FALSE
}
```

```
# y greater than 10?
if (y > 10) {
  print("greater than 10")
} else {
  print("less or equal to 10")
}
```

Conditional (if…else…)

```
if (cond) {
  # run here if cond is TRUE
} else {
  # run here if cond is FALSE
}
```

```
# y greater than 10?
if (y > 10) {
  print("greater than 10")
} else {
  print("less or equal to 10")
}
```

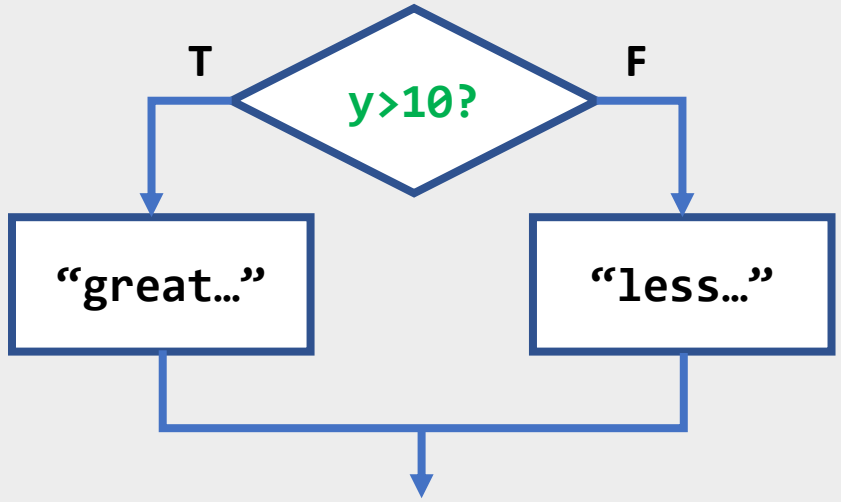

Conditional (if…else if…else…)

```
if (cond1) {
  # run here if cond1 is TRUE
} else if (cond2) {
  # run here if cond1 is FALSE but cond2 is TRUE
} else {
  # run here if neither cond1 nor cond2 is TRUE
}
```

Iteration

```
for (var in seq) {
  do something
}

while (cond) {
  do something if cond is TRUE
}
```

```
# sum of squares
t <- 1:3
y <- 0

for (x in t) {
  y <- y + x^2
}

print(y)
```