

Introduction

STAT598z: Introduction to Statistical Computing

Vinayak Rao

Purdue University

Logistics

- **Class** MWF 330PM-420PM, WALC 1121
- **Website** <https://purdue.brightspace.com/d2l/home/220693>
- **Campuswire** <https://campuswire.com/p/GF174ADA8> (Class code: 5023)
- **Instructor** Vinayak Rao
- **Office** Math212
- **Email** varao@purdue.edu (*use Campuswire unless it is confidential*)
- **Office Hours** TBA
- **TA** Haoyun Yin
- **Office Hours** TBA

Comp. statistics vs stat. computing

Computational statistics or statistical computing, is that the question?, Lauro, C., *Comp Stat and Data Analysis* 23 (1996).

Statistical Computing: Application of Comp. Sci. to Statistics

- Tools: programming, software, data structures and their manipulation, hardware (GPUs, parallel architectures)
- E.g. Releasing software/ sharing analysis

Computational statistics: Design of algorithms for implementing statistical methods on computers

- Statistical methodology
- E.g. writing a paper for a statistics journal

Comp. statistics vs stat. computing

Computational statistics or statistical computing, is that the question?, Lauro, C., *Comp Stat and Data Analysis* 23 (1996).

Statistical Computing: Application of Comp. Sci. to Statistics

- Tools: programming, software, data structures and their manipulation, hardware (GPUs, parallel architectures)
- E.g. Releasing software/ sharing analysis

Computational statistics: Design of algorithms for implementing statistical methods on computers

- Statistical methodology
- E.g. writing a paper for a statistics journal

This course: more of former

STAT545: more of the latter

Goals of the course

Broadly: to learn programming for Statistics/Data Science

- Only very basic programming background required.
- Perhaps not the best class for those already good at this

Our focus will be on

- the R programming language
- *statistical* rather than *general-purpose* computing
- R for reproducible research rather than ad hoc analysis

Topics covered (tentative)

- R fundamentals (data-structures, commands, flow control)
- R packages
- R plotting (ggplot2)
- Debugging with R
- Writing efficient R code
- R Markdown and dynamic documents
- Object-oriented programming
- Functional programming
- Optimization
- Regression/LASSO
- Monte Carlo Methods

Advanced topics (depending on how things progress):

- Interactive applications with R shiny
- Introduction to R internals
- Programming with Stan

Textbooks

- *“The Art of R Programming: A Tour of Statistical Software Design”*, Norman Matloff
- *“R for Data Science”*, Garrett Golemund and Hadley Wickham. ([Amazon](#) but also available [free](#))

Also useful:

- *“Software for Data Analysis”*, John M. Chambers
- *“An Introduction to R”* (The R [manual](#))
- *“Advanced R”*, Hadley Wickham

Grading

- **Homework:** 30%
- **Midterm I:** 25%
- **Midterm II:** 25%
- **Project:** 20%

Your final grade will be curved

- If you do well, you will get an A (you are not competing with your classmates)

Homework

(Approximately) biweekly assignments, involving reading, writing and programming

Are vital to doing well in the exams

Late policy:

Seven extension days can be applied anytime to homeworks/projects (no explanation needed)

- Use must be stated explicitly at the time of submission at the top of the HW
- Cannot be rearranged after they are applied to a submission
- **Other than this, no late submissions will be accepted**

Homework

(Approximately) biweekly assignments, involving reading, writing and programming

Are vital to doing well in the exams

Late policy:

Seven extension days can be applied anytime to homeworks/projects (no explanation needed)

- Use must be stated explicitly at the time of submission at the top of the HW
- Cannot be rearranged after they are applied to a submission
- **Other than this, no late submissions will be accepted**

You may discuss problems with other students, but must:

- write your own solution independently
- name students you had significant discussions with

[Purdue's guide on academic integrity](#)

Programming

Central to modern statistics/data analysis. We want:

- computers to do what we don't want to do ourselves
- computers to do what we actually want them to do

Programming

Central to modern statistics/data analysis. We want:

- computers to do what we don't want to do ourselves
- computers to do what we actually want them to do

Programming involves:

- **Correctness:** getting computers to do what we want
- **Efficiency:** low compute and (more imp.) human time
- **Clarity:** Donald Knuth: “treat a program as ... addressed to human beings rather than to a computer”
 - Especially important with messy data

The R programming language

A programming language and environment for statistics

A GNU project available as Free software.

(“Think free as in free speech, not free beer”: Richard Stallman)

The R programming language

A programming language and environment for statistics

A GNU project available as Free software.

(“Think free as in free speech, not free beer”: Richard Stallman)

You can (and should):

- Install R (available at <http://cran.r-project.org/>)
- Look at the R source code
- Modify the R source code (if you're feeling brave)

You will:

- Write clear, efficient and (hopefully) useful R code

A brief history of R

Based on Bell Labs' S language by John Chambers

Started by Ihaka and Gentleman at the Univ.of Auckland *R: A Language for Data Analysis and Graphics*, (1996)

A high-level interpreted language with convenient features for loading, manipulating and plotting data

A huge collection of user-contributed packages to perform a wide variety of tasks

Widely used in academia, and increasingly popular in industry

The R command prompt

Starting R begins a new session

R presents you with a command prompt or console

Can interact with R through the console:

- Enter command
- R processes command and prints output
- The command `q()` ends the session

```
In [ ]: 1 + 3
```

```
In [ ]: x <- rgamma(3,2,1); x # Generate Gamma(2,1) variables
```

```
In [ ]: x <- rnorm(1000)
        hist(x)
```

RStudio

RStudio provides a more convenient Integrated Development Environment (IDE) to interact with R

Layout includes

- an editor
- a console
- workspace/history tabs
- tabs for plots/packages/files etc

Convenient user interface: point-and-click, autocomplete, help etc.

You should install RStudio Desktop (available at rstudio.org)

[RStudio demo]

(run Alt-Shift-K for a list of all shortcuts)

R Scripting

While we often use R interactively, it is useful to do this through scripts

- Fewer errors
- Better reproducibility
- Can reuse useful sequences of operations
- Can build increasingly complicated sequence of operations

Ultimately, R is a full-fledged programming language for statistical computing: Treat it as such!

R scripting guidelines

Filenames should end with .R (e.g. denoise.R)

Scripts should have explanatory comments

Variables should have informative names

Scripts should be indented appropriately

See R style-guides from:

- [Google](#)
- [Hadley Wickham](#)

Learning R

We will look up a few useful R packages (e.g. ggplot, plyr)

The next part of the course aims to:

- Write clean, efficient and idiomatic R
- Understand why things done the way they are
- Be comfortable manipulating and presenting data

Dynamic documents and R Markdown

Take the idea of reproducible code to reproducible documents

Instead of working with R commands, work with an entire report

Report includes description of your problem, data and algorithm as well as embedded code and results

You can automatically “compile” the report, which will rerun your code, regenerate your results and form a new report

Allows collaborators to regenerate report on their computer

This is how we will be submitting homeworks

Jupyter notebook

Another nice system for dynamics notebooks is Jupyter notebook

Formerly called ipython notebooks, is still python based, but now supports more languages:

- Ju(lia)Py(thon)R

I made these slides using Jupyter

You can try installing it if you're prepared to deal with setting up python/python libraries

- Pros: you can play around with these slides
- Cons: I think RMarkdown and knitr is a bit more useful for some of the R programming we will do

To do

- Install R and RStudio
- Join class Campuswire webpage (look at *Content* on Brightspace)

Reading/Viewing:

- [A New York Times article](#)
- Sections 1 to 5 of: [A \(very\) short introduction to R, Torfs and Brauer](#)
- R : the good, the bad and the ugly, John Cook ([youtube](#))
- [Nature article on dynamic documents](#)