


Algorithms and Data Structures

Introduction

Introduction




Learning goals

- Learn a thing
- Learn another thing

WHAT IS THIS LECTURE ABOUT

How to analyze and solve interesting computational problems in stats


We will mainly study techniques from applied maths (numerics, linear algebra, optimization) for that purpose, not so much from statistics itself.



STAT. COMPUTING VS. COMP. STATISTICS

Westlake (1994)

“Computational statistics relates to the advance of statistical theory and methods through the use of computational methods. This includes both the use of computation to explore the impact of theories and methods, and development of algorithms to make these ideas available to users.”



Statistical Computing

Computational tools for statistics using methods and techniques of computer science, numerics and optimization.


Computational Statistics

Applied statistics with the help of computational tools.

We will mainly study the former.

PRINCIPAL QUESTIONS OF COMPUTER SCIENCE

- What is a computer?
- What is an algorithm?
- What is computable?
- How to analyze algorithms? What is a “good” / “clever” algorithm?
- How can you describe the efficiency of an algorithm?




For only a few of these questions we have (appropriately much) time in the lecture, but we will at least discuss questions of efficiency.

NUMERICAL ANALYSIS

Subfield of mathematics that deals with the development and analysis of algorithms for continuous problems.


We are interested in these if:

- There is no analytic solution to a problem, or
- An exact solution to a problem is available, but the solution cannot be computed efficiently, or computational errors strongly impact the solution.




NUMERICAL COMPUTATIONS

- In this lecture we focus on how to solve statistical problems numerically with the help of a digital computer.
- Numbers are represented by fixed-length bit-strings, so there is only a finite set of numbers available.
- None of the common mathematical sets of numbers \mathbb{N} , \mathbb{Z} , \mathbb{Q} and \mathbb{R} is therefore entirely representable on the computer
- Reals will be replaced by approximate machine numbers.
- Even the basic arithmetic operations $+$, $-$, $*$, and $/$ can only be approximated.





OPTIMIZATION

- Huge and diverse subfield of applied maths
 - Linear optimization
 - Convex optimization
 - Discrete optimization
 - Black-box optimization
 - ...
- Many things in stats and nearly everything in machine learning is solved by optimizing a quality criterion and of the problems are often difficult enough to care about how this works.





EXAMPLE: SIMPLE COMPUTATION



How do we calculate $\cos(x)$ here?

Possible answer: approximated Taylor series!

EXAMPLE: EQUATION SOLVING / ROOT FINDING




EXAMPLE: EQUATION SOLVING / ROOT FINDING

/ 2

For a tolerance of 0.001, the numerical solution is:


```
f = function(x) { 2 + cos(x) + sin(2 * x) }
f0 = function(x) { f(x) - 2.5 }
uniroot(f0, c(0, 4), tol = 0.001)
## $root
## [1] 1.401727
##
## $f.root
## [1] -2.382814e-06
##
## $iter
## [1] 6
##
## $init.it
## [1] NA
##
```



EXAMPLE: EQUATION SOLVING / ROOT FINDING

/ 3

```
## $estim.prec  
## [1] 0.0005500599
```




EXAMPLE: SAMPLING

How does a computer generate randomness? How does R calculate:

```
rnorm(10)
## [1] 0.06758659 0.67908219 -0.26783828 0.77266612
## [5] -0.02887604 -0.46867480 -0.94277461 0.15943498
## [9] -0.73383393 -0.48259188
```

- How do we sample from a “non-standard” density?
- Wait for someone to create a CRAN package for it...?




EXAMPLE: LINEAR MODEL

How do we solve the normal equations?

$$\beta = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T y$$

- So invert matrix once (how does that work?), then 2 matrix multiplications? Nope! See chapter on matrix decompositions.
- How exact is that?
- How expensive is it for large n ?



EXAMPLE: LASSO

How do we optimize something like this?

$$\sum_{i=1}^n (\beta^T \mathbf{x}^{(i)} - y^{(i)})^2 + \lambda \sum_{j=1}^p |\beta_j|$$

- Analytically solvable? Answer: no!
- How does it work “numerically”?
- How expensive is it?

