Computational Statistics

Course Design

- Interestingly, unlike other material discussed here, computational statistics can be taught *without the involvement of computers*.

“Statistical Computing Theory”

- Procession of algorithms
- Proofs of convergence
- Rates of convergence
- Algorithmic complexity
- etc.

Basic Goals

- Understand what goes on “under the hood”
- Program their own minimizers, samplers, etc.
  - Understand the difficulty in implementing these algorithms well
- Understand how computers work *and their limitations*
- Make better use of existing “black boxes” (optim, WinBUGS, etc.)

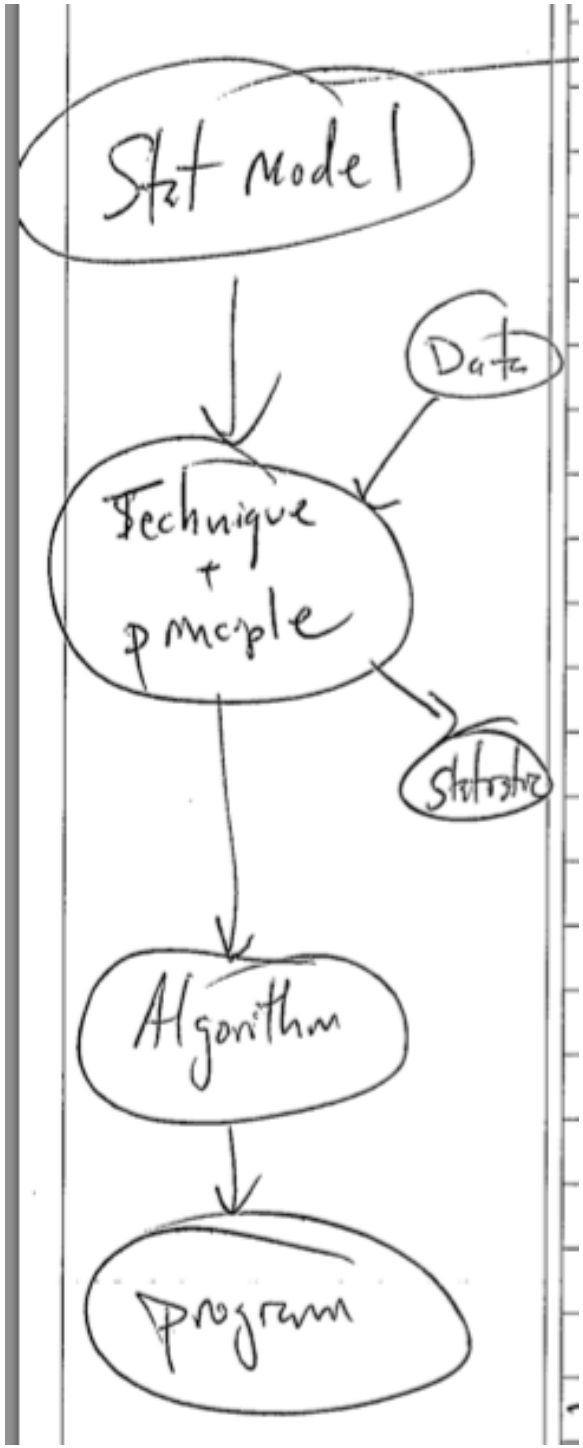
Stat model
Data
Technique
+
principle
Statistic
Algorithm
Program

Topics

- Solving nonlinear equations [f(x) = 0]
- General purpose minimization ("how to use optim")
  - Newton's method
  - Quasi-Newton (BFGS, FP)
  - Conjugate gradient
  - Simulated annealing (later)

Topics

- EM algorithm (Monte Carlo EM as an aside)
- Laplace approximation
- Quadrature
- Pseudo-random number generation
- Independent Monte Carlo (integration)
  - Rejection sampling
  - Importance sampling

Topics

- Markov Chain Monte Carlo
  - Metropolis-Hastings
  - Gibbs sampling
  - Variants
  - Diagnostic tools
- Parallel Computing
  - Basic embarassingly parallel computations
  - snow

If there's extra time

- Bootstrap
- Smoothing
- Perfect Sampling (as amusement)
- Additional Applications
  - Spatial statistics
  - Time series

Computational Aspects

- Visualize the geometry of optimization problems
  - Likelihoods, profile likelihoods
- Develop efficient implementations of MCMC algorithms
  - Profiling, C programming
  - Geometry of posterior sampling

Ex: Profile Likelihoods

Profile likelihood. Write a function called `makeProfLik` which takes a numeric vector `x` and returns a *function* of one argument named `sigma2`. The prototype for `makeProfLik` is

```
makeProfLik <- function(x) {
      ## function body
}
```

The function returned by `makeProfLik` should be the profile likelihood for the the variance parameter σ^2 of a normal distribution. The profile likelihood function, $\rho(\sigma^2)$, is simply

$$\rho(\sigma^2) = \max_{\mu} L(\mu, \sigma^2)$$

where L is the full likelihood function for a normal distribution. That is, for a given value of σ^2, we fix the value for σ^2 in L and maximize over μ. The function returned by `makeProfLik` should be a function of a *single argument only*, i.e. `sigma2`. The returned function should have no other arguments (hint: use lexical scoping).

Ex: Profile Likelihoods

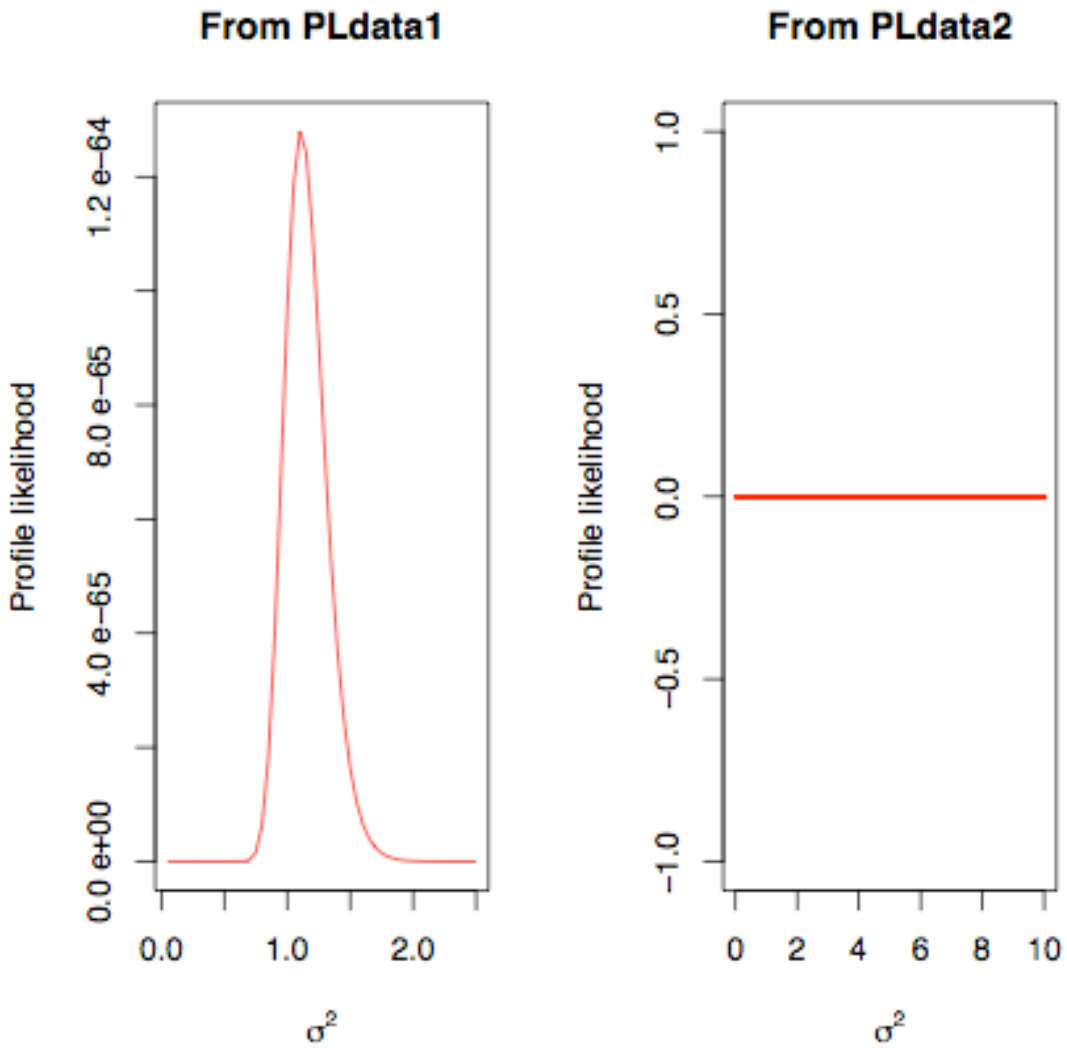

Textbooks

- *Numerical Analysis for Statisticians*, Ken Lange
- *Markov Chain Monte Carlo in Practice*, Gilks, Richardson, Spiegelhalter
- *Structured Computer Organization*, Tanenbaum