## Statistical Computing (36-350) Lecture 4: Writing and Calling Functions

Cosma Shalizi and Vincent Vu

12 September 2011

# Agenda

- Defining functions: Tying related commands into bundles
- Interfaces: Controlling what the function can see and do
- Example: Improving the lab's parameter estimation

Absolutely Essential Reading for Friday: Sec. 4.2 of the textbook

 ${\rm MERELY}\ {\rm USEFUL}\ {\rm READING};\ {\rm Chapter}\ 3$ 

CODE FROM THIS LECTURE: At class website, with comments

## Why Functions?

Data structures tie related values into one object Functions tie related commands into one object For example:

```
# "Robust" loss function, for outlier-resistant regression
# Inputs: vector of numbers (x)
# Outputs: vector with x<sup>2</sup> for small entries, |x| for large ones
psi.1 <- function(x) {
   psi <- ifelse(x<sup>2</sup> > 1, abs(x), x<sup>2</sup>)
   return(psi)
}
```

Our functions get used just like the built-in ones:

```
> z <- c(-0.5,-5,0.9,9)
> psi.1(z)
[1] 0.25 5.00 0.81 9.00
```

Go back to the declaration and look at the parts:

```
# "Robust" loss function, for outlier-resistant regression
# Inputs: vector of numbers (x)
# Outputs: vector with x<sup>2</sup> for small entries, |x| for large ones
psi.1 <- function(x) {
    psi <- ifelse(x<sup>2</sup> > 1, abs(x), x<sup>2</sup>)
    return(psi)
}
```

# **Interfaces**: the **inputs** or **arguments**; the **outputs** or **return value**

Calls other functions ifelse(), abs(), and operators ^ and >

could also call other functions we've written

return() says what the output is

alternately, return the last evaluation; I like explicit returns better

**Comments**: Not required by R, but a Very Good Idea One-line description of purpose; listing of arguments; listing of outputs

#### Named and default arguments

```
# "Robust" loss function, for outlier-resistant regression
# Inputs: vector of numbers (x), scale for crossover (c)
# Outputs: vector with x^2 for small entries, c|x| for large ones
psi.2 <- function(x,c=1) {
    psi <- ifelse(x^2 > c^2, c*abs(x), x^2)
    return(psi)
}
> identical(psi.1(z), psi.2(z,c=1))
[1] TRUE
```

#### Default values get used if names are missing:

```
> identical(psi.2(z,c=1), psi.2(z))
[1] TRUE
```

#### Named arguments can go in any order when explicitly tagged:

```
> identical(psi.2(x=z,c=2), psi.2(c=2,x=z))
[1] TRUE
```

(\* ) \* ) \* ) \* )

Image: A matrix

# **Checking Arguments**

 $\operatorname{PROBLEM}$ : Odd behavior when arguments aren't as we expect

```
> psi.2(x=z,c=c(1,1,1,10))
[1] 0.25 5.00 0.81 81.00
> psi.2(x=z,c=-1)
[1] 0.25 -5.00 0.81 -9.00
```

SOLUTION: Put little sanity checks into the code

```
# "Robust" loss function, for outlier-resistant regression
# Inputs: vector of numbers (x), scale for crossover (c)
# Outputs: vector with x<sup>2</sup> for small entries, c|x| for large ones
psi.3 <- function(x,c=1) {
    # Scale should be a single positive number
    stopifnot(length(c) == 1,c>0)
    psi <- ifelse(x<sup>2</sup> > c<sup>2</sup>, c*abs(x), x<sup>2</sup>)
    return(psi)
}
```

Arguments to stopifnot() are a series of expressions which should all evaluate to TRUE; execution halts, with error message, at *first* FALSE (try it!)

#### What the function can see and do

Argument names over-ride those in the larger environment, *inside* the function

Changes made inside the function don't propagate

```
> x <- 7
> y <- c("A","C","G","T","U")
> adder <- function(y) { x<- x+y; return(x) }
> adder(1)
[1] 8
> x
[1] 7
> y
[1] "A" "C" "G" "T" "U"
```

There *are* ways around this, but they are difficult and best avoided (see Chambers, ch. 5)

## Looking in the Environment

Any name not defined in the function will be looked for in the environment

What matters is the value when the function is called, not when defined

#### Respect the Interfaces!

Interfaces mark out a controlled inner environment for our code Interference with, or from, the rest of the system is only as allowed by interface

Good practice: explicitly give the function all the information it needs through the arguments; this minimizes the chances of confusion and error

Exception: true universals like  $\boldsymbol{\pi}$ 

Likewise, output should only be through the return value Will say more about breaking up tasks and about environments later

Further reading: Herbert Simon, The Sciences of the Artificial

## Example: Improving on Friday's Lab

We want to fit the statistical model

$$Y = y_0 N^a + \text{noise}$$

where Y is the per-capita "gross metropolitan product" of a city, N is its population, and  $y_0$  and a are parameters Approximate the derivative of error w.r.t a and move against it

$$MSE(a) \equiv \frac{1}{n} \sum_{i=1}^{n} (Y_i - y_0 N_i^a)^2$$
$$MSE'(a) \approx \frac{MSE(a+h) - MSE(a)}{h}$$
$$a_{t+1} - a_t \propto -MSE'(a)$$

The code given:

```
maximum.iterations <- 100
deriv.step <- 1/1000
step.scale <- 1e-12</pre>
stopping.deriv <- 1/100</pre>
iteration <- 0
deriv <- Inf
a <- 0.15
while ((iteration < maximum.iterations) && (deriv > stopping.deriv)) {
  iteration <- iteration + 1
  mse.1 <- mean((gmp$pcgmp - 6611*gmp$pop^a)^2)</pre>
  mse.2 <- mean((gmp$pcgmp - 6611*gmp$pop^(a+deriv.step))^2)</pre>
  deriv <- (mse.2 - mse.1)/deriv.step</pre>
  a <- a - step.scale*deriv
}
list(a=a,iterations=iteration,converged=(iteration < maximum.iterations))</pre>
```

# What's wrong with this?

- Not *encapsulated*: Re-run by cutting and pasting code but how much of it? Also, hard to make part of something larger
- *Inflexible*: To change initial guess at *a*, have to edit, cut, paste, and re-run
- *Error-prone*: To change the data set, have to edit, cut, paste, re-run, and hope that all the edits are consistent
- *Hard to fix*: should top when *absolute value* of derivative is small, but this stops when large and negative. Imagine having five copies of this and needing to fix same bug on each.

Will turn this into a function and then improve it; comments omitted here, see online

#### First attempt, with logic fix:

```
estimate.scaling.exponent.1 <- function(a) {</pre>
  maximum.iterations <- 100
  deriv.step <- 1/1000
  step.scale <- 1e-12</pre>
  stopping.deriv <- 1/100</pre>
  iteration <- 0
  deriv <- Inf
  while ((iteration < maximum.iterations) && (abs(deriv) > stopping.deriv)) {
    iteration <- iteration + 1
    mse.1 <- mean((gmp$pcgmp - 6611*gmp$pop^a)^2)</pre>
    mse.2 <- mean((gmp$pcgmp - 6611*gmp$pop^(a+deriv.step))^2)</pre>
    deriv <- (mse.2 - mse.1)/deriv.step</pre>
    a <- a - step.scale*deriv
  }
  fit <- list(a=a,iterations=iteration,</pre>
    converged=(iteration < maximum.iterations))</pre>
  return(fit)
}
```

```
PROBLEM: All those magic constants! SOLUTION: Make them defaults
```

```
estimate.scaling.exponent.2 <- function(a, y0=6611, maximum.iterations=100,
    deriv.step = 1/100, step.scale = 1e-12, stopping.deriv = 1/100) {
  iteration <- 0
  deriv <- Inf
  while ((iteration < maximum.iterations) && (abs(deriv) > stopping.deriv)) {
    iteration <- iteration + 1
    mse.1 <- mean((gmp$pcgmp - y0*gmp$pop^a)^2)</pre>
    mse.2 <- mean((gmp$pcgmp - y0*gmp$pop^(a+deriv.step))^2)</pre>
    deriv <- (mse.2 - mse.1)/deriv.step</pre>
    a <- a - step.scale*deriv
  }
  fit <- list(a=a,iterations=iteration,</pre>
    converged=(iteration < maximum.iterations))</pre>
  return(fit)
}
```

EXERCISE: Experiment with different values of deriv.step

PROBLEM: Why type out the same calculation of the MSE twice? SOLUTION: Declare a function

```
estimate.scaling.exponent.3 <- function(a, y0=6611, maximum.iterations=100,
    deriv.step = 1/100, step.scale = 1e-12, stopping.deriv = 1/100) {
    iteration <- 0
    deriv <- Inf
    mse <- function(a) { mean((gmp$pcgmp - y0*gmp$pop^a)^2) }
    while ((iteration < maximum.iterations) && (abs(deriv) > stopping.deriv)) {
        iteration <- iteration + 1
        deriv <- (mse(a+deriv.step) - mse(a))/deriv.step
        a <- a - step.scale*deriv
    }
    fit <- list(a=a,iterations=iteration,
        converged=(iteration < maximum.iterations))
    return(fit)
}
```

mse() declared inside estimate.scaling.exponent.3(), so it won't be added to the global (terminal) environment, but it can see y0

PROBLEM: Locked in to using specific columns of gmp; shouldn't have to re-write just to compare two data sets SOLUTION: More arguments, with defaults

```
estimate.scaling.exponent.4 <- function(a, y0=6611, response=gmp$pcgmp,
    predictor = gmp$pop, maximum.iterations=100, deriv.step = 1/100,
    step.scale = 1e-12, stopping.deriv = 1/100) {
  iteration <- 0
  deriv <- Inf
  mse <- function(a) { mean((response - y0*predictor^a)^2) }</pre>
  while ((iteration < maximum.iterations) && (abs(deriv) > stopping.deriv)) {
    iteration <- iteration + 1
    deriv <- (mse(a+deriv.step) - mse(a))/deriv.step</pre>
    a <- a - step.scale*deriv
  3
  fit <- list(a=a,iterations=iteration,</pre>
    converged=(iteration < maximum.iterations))</pre>
  return(fit)
3
```

Respecting the interfaces: We could turn the while() loop into a for() loop, and nothing outside the function would care

```
estimate.scaling.exponent.5 <- function(a, y0=6611, response=gmp$pcgmp,
    predictor = gmp$pop, maximum.iterations=100, deriv.step = 1/100,
    step.scale = 1e-12, stopping.deriv = 1/100) {
    mse <- function(a) { mean((response - y0*predictor^a)^2) }
    for (iteration in 1:maximum.iterations) {
        deriv <- (mse(a+deriv.step) - mse(a))/deriv.step
        a <- a - step.scale*deriv
        if (abs(deriv) <= stopping.deriv) { break() }
    }
    fit <- list(a=a,iterations=iteration,
        converged=(iteration < maximum.iterations))
    return(fit)
}
```

# Summary

- Functions bundle related commands together into objects: easier to re-use, easier to modify, less risk of error, easier to think about
- Interfaces control what the function can see (arguments, environment) and change (its internals, its return value)
- Calling functions we define works just like calling built-in functions: named arguments, defaults

Next time: working with many functions