

---

# Structural MRI statistics using RMINC v0.5

---

*Author:*  
Jason LERCH

April 16, 2008

# Preface

This mini-book attempts to provide a general introduction to the statistics side of structural brain imaging, with a heavy emphasis on practical worked examples. It also introduces a particular toolkit, RMINC, designed to make running these types of statistical analyses easier. It is targeted at the general user, who may or may not have some statistical background, but does have some data they want analysed in a straightforward way. It is not meant to be a complete handbook on statistics, but hopefully will provide enough of a primer to get by, at least for a little while.

This mini-book exists for a number of reasons. I have over the years been asked multiple questions relating to structural brain imaging and statistics, and have had the chance to learn answers to those questions from countless people. This book thus exists as an attempt to put some of those answers down on paper. Secondly, writing this book is part and parcel of the development of RMINC; it is easier to write code useable by others if one documents it first, and then writes code to fit the documentation.

The book will likely be an incomplete work in progress for a long while yet. The  $\text{\LaTeX}$ source for this book are packaged along with RMINC itself, and contributors are most welcome!

# Contents

<b>1</b>	<b>Introduction</b>	<b>3</b>
1.1	Installing the tools . . . . .	3
1.1.1	R . . . . .	3
1.1.2	MINC . . . . .	3
1.1.3	RMINC . . . . .	4
1.2	Overview of the analysis process . . . . .	4
1.3	Data used throughout this book . . . . .	4
<b>2</b>	<b>Preparing the data</b>	<b>5</b>
2.1	Types of datasets . . . . .	5
2.1.1	Voxel based morphometry . . . . .	5
2.1.2	Deformation based morphometry . . . . .	5
2.1.3	Other . . . . .	5
2.2	Input data . . . . .	5
<b>3</b>	<b>Descriptive statistics</b>	<b>7</b>
3.1	Writing results to file . . . . .	8
3.2	Creating summaries by group . . . . .	9
<b>4</b>	<b>Linear Models</b>	<b>11</b>
4.1	First linear model . . . . .	11
4.2	Plotting voxels . . . . .	12
4.3	Creating images . . . . .	14
4.4	Using subsets . . . . .	14
4.5	Multiple Comparisons . . . . .	15
<b>5</b>	<b>Bits and Pieces</b>	<b>17</b>
5.1	Correlations . . . . .	17
5.2	non-parametric statistics . . . . .	17
5.3	mixed effects models . . . . .	17
5.4	Cleaning up . . . . .	17
<b>6</b>	<b>Advanced Topics</b>	<b>18</b>
6.1	Running arbitrary R functions . . . . .	18

# Chapter 1

## Introduction

The process of analysing brain imaging data is typically comprised of a series of stages. The study is designed with various choices made about the biology question that is to be addressed and the data necessary to answer the questions thus posed. Then the data is then acquired, and once that is completed, the images are processed in various automatic, semi-automatic, or manual ways and then analysed.

This book deals mainly with the final part, the data analysis, though there will be several side-tracks into the other topics. A single example will be used throughout: a mouse brain imaging study comparing two genotypes. The methods described herein should be easily transferable to any other structural imaging study which looks at brain shape, tissue classification, or signal intensities.

### 1.1 Installing the tools

All the analyses will be performed using RMINC, which is a library designed to handle MINC volumes inside the R statistical environment. All the tools needed are freely available, and should run on just about any computer/operating system. Installation and setup is described in some more detail below.

#### 1.1.1 R

Quick background about R.

Installing R.

Where to find further reading.

#### 1.1.2 MINC

Quick background about MINC.

In order for RMINC to work, the MINC libraries have to be compiled as shared libraries. When configuring, make sure the following flag is added on:

```
./configure --enable-shared=yes
```

Where to find more information.

### 1.1.3 RMINC

Quick background about RMINC.

Installing RMINC.

## 1.2 Overview of the analysis process

The data analysis process usually proceeds in the following way. First the input images are assessed for correctness; any obvious processing errors are removed from any subsequent analyses. The question of what constitutes an outlier is often a tricky one. In order to avoid the temptation to manipulate the data in a biased way it is best if the person who reviews the input data is blind about the categorization of each particular dataset.

Once all the acceptable datasets are in place a series of descriptive statistics can be generated, usually consisting of means and standard deviations of all images in the study as well as of all the subgroupings. This is followed by generating statistical maps of the main variables of interest. These are then thresholded for significance while taking multiple comparisons into account. There is then often a series of steps in which new statistical models are analyzed and thresholded until the results become more understandable. This usually involves lots of plotting of individual datapoints.

## 1.3 Data used throughout this book

This book will consistently work with one dataset consisting of 5 male and 5 female C57Bl/6 mice, taken from a larger dataset published in a 2007 NeuroImage paper by Spring et al. The mice, all 12 weeks old, were scanned using an overnight T2-weighted FSE sequence, then all aligned into a common space using an automated image registration algorithm (i.e. deformation based morphometry). The final metric of interest was then the Jacobian determinant of the deformations needed to align each mouse to the final common atlas. These details are relatively unimportant for this book - the input might as well be voxel density maps from VBM - but at least it gives some background for those who care. If you want to follow along with the examples used in this book you can download the data at <http://launchpad.net/rminc>. Note that the data has been downsampled to 120 micron voxels (from the original 32 micron voxels) to keep the download within reasonable limits.

## Chapter 2

# Preparing the data

This chapter will briefly discuss how to generate structural imaging data useable for the statistical analyses described in the rest of the book.

### 2.1 Types of datasets

Describe what can be done.

#### 2.1.1 Voxel based morphometry

Some more detail on VBM.

#### 2.1.2 Deformation based morphometry

Some more detail on DBM.

#### 2.1.3 Other

Mention cortical thickness, manual segmentation, etc.

### 2.2 Input data

Once the files have been processed, the easiest way to proceed is by setting up a text file containing all the necessary information about each scan. This file should be comma or space separated, have one row per scan, with each column containing info about each scan. One of the columns should contain the filename pointing to the MINC volumes to be processed. The example from the five male and five female mice is the following:

```
Filename, Gender, coil, weight  
volumes/img_08nov05.0-fwhm1.0.mnc,Female,1,22.6
```

```

volumes/img_08nov05.1-fwhm1.0.mnc,Female,2,19.6
volumes/img_08nov05.2-fwhm1.0.mnc,Female,3,21.8
volumes/img_29sept05.0-fwhm1.0.mnc,Male,1,24.0
volumes/img_29sept05.2-fwhm1.0.mnc,Male,3,27.0
volumes/img_30sept05.0-fwhm1.0.mnc,Male,1,28.3
volumes/img_30sept05.1-fwhm1.0.mnc,Male,2,26.5
volumes/img_30sept05.2-fwhm1.0.mnc,Male,3,28.1
volumes/img_31oct05.0-fwhm1.0.mnc,Female,1,20.5
volumes/img_31oct05.2-fwhm1.0.mnc,Female,3,20.0

```

Notice how the first row contains a header. This is optional, but makes later access to the data easier and is therefore recommended.

The next step is to actually load this file into R. The steps are given below:

```

> library(RMINC)
> gf <- read.csv("control-file.csv")

```

The library commands load the RMINC library into R. The next line then reads the information describing this dataset from a comma-separated text file. The basic syntax of an R command is a variable name - which can be whatever you chose, within only a few limits - on the left hand side, the arrow (less than followed by a dash) indicating an assignment, and then the function call (in this case `read.csv` to read a comma separated value) with any arguments (in this case the filename) in parentheses. Strings - such as the filename in this case - are placed inside quotes.

```

> xtable(gf, caption = "GLIM File")

```

	Filename	Gender	coil	weight
1	volumes/img_08nov05.0-fwhm1.0.mnc	Female	1.00	22.60
2	volumes/img_08nov05.1-fwhm1.0.mnc	Female	2.00	19.60
3	volumes/img_08nov05.2-fwhm1.0.mnc	Female	3.00	21.80
4	volumes/img_29sept05.0-fwhm1.0.mnc	Male	1.00	24.00
5	volumes/img_29sept05.2-fwhm1.0.mnc	Male	3.00	27.00
6	volumes/img_30sept05.0-fwhm1.0.mnc	Male	1.00	28.30
7	volumes/img_30sept05.1-fwhm1.0.mnc	Male	2.00	26.50
8	volumes/img_30sept05.2-fwhm1.0.mnc	Male	3.00	28.10
9	volumes/img_31oct05.0-fwhm1.0.mnc	Female	1.00	20.50
10	volumes/img_31oct05.2-fwhm1.0.mnc	Female	3.00	20.00

Table 2.1: GLIM File

The little code fragment and table above just shows what the `gf` variable looks like after being read into R<sup>1</sup>

<sup>1</sup>`xtable` is only necessary for display purposes in this manual (which, by the way, is being written using Sweave, a tool for combining R with latex).

## Chapter 3

# Descriptive statistics

Before going into explicit hypothesis tests it is often useful to get a general feel for what the data looks like - this is where descriptive statistics come in. The most common functions include computing the mean and variance or standard deviation at every voxel. If the data is inherently divided into groups, such as patients and control, or, in our example dataset, males versus females, then the descriptive stats can also be grouped by those variables.

To start we can look at the mean Jacobian determinant at every voxel of all the data combined:

```
> library(RMINC)
> overall.mean <- mincMean(gf$Filename)

Method: mean
Number of volumes: 10
Volume sizes: 71 137 105
N GROUPS: 1.000000
In slice
 0  1  2  3  4  5  6  7  8  9 10 11 12 13 14 ...
Done

> overall.mean

Multidimensional MINC volume
Columns:
[1] "volumes/img_08nov05.0-fwhm1.0.mnc"
```

The `mincMean` function computes the mean at every voxel of a set of filenames specified as an argument. The output is in this case assigned to the `overall.mean` variable. Repeating the variable in the R session, as done above, causes a summary to be printed.

One thing to note about the R syntax above: the dollar symbol is used to access a specific column inside a data frame. What this means is that inside



the `gf` variable - which, remember, is the variable that was read in from the comma-separated values file which describes the dataset - each column has a name which can be accessed by that dollar variable. Here are some examples, first showing the entire contents of `gf` and then two separate columns alone:

```
> gf
      Filename Gender coil weight
1 volumes/img_08nov05.0-fwhm1.0.mnc Female    1  22.6
2 volumes/img_08nov05.1-fwhm1.0.mnc Female    2  19.6
3 volumes/img_08nov05.2-fwhm1.0.mnc Female    3  21.8
4 volumes/img_29sept05.0-fwhm1.0.mnc   Male    1  24.0
5 volumes/img_29sept05.2-fwhm1.0.mnc   Male    3  27.0
6 volumes/img_30sept05.0-fwhm1.0.mnc   Male    1  28.3
7 volumes/img_30sept05.1-fwhm1.0.mnc   Male    2  26.5
8 volumes/img_30sept05.2-fwhm1.0.mnc   Male    3  28.1
9 volumes/img_31oct05.0-fwhm1.0.mnc Female    1  20.5
10 volumes/img_31oct05.2-fwhm1.0.mnc Female    3  20.0

> gf$Filename
 [1] volumes/img_08nov05.0-fwhm1.0.mnc volumes/img_08nov05.1-fwhm1.0.mnc
 [3] volumes/img_08nov05.2-fwhm1.0.mnc volumes/img_29sept05.0-fwhm1.0.mnc
 [5] volumes/img_29sept05.2-fwhm1.0.mnc volumes/img_30sept05.0-fwhm1.0.mnc
 [7] volumes/img_30sept05.1-fwhm1.0.mnc volumes/img_30sept05.2-fwhm1.0.mnc
 [9] volumes/img_31oct05.0-fwhm1.0.mnc volumes/img_31oct05.2-fwhm1.0.mnc
10 Levels: volumes/img_08nov05.0-fwhm1.0.mnc ...

> gf$Gender
 [1] Female Female Female Male   Male   Male   Male   Male   Female Female
Levels: Female Male
```

So if, in the text file that describes the dataset, the column containing all the filenames was called “jacobians”, then the `mincMean` command would have been `mincMean(gf$jacobians)`.

If an incorrect column is specified - i.e. something which does not contain filenames - then you should receive an error.

### 3.1 Writing results to file

Once the means at every voxel have been computed, they can be written to file. This is done with command below:

```
> mincWriteVolume(overall.mean, "overall-mean.mnc")

Writing column 1 to file overall-mean.mnc
Sizes: 71 137 105
Range: 0.084176 -0.108390
```

The `mincWriteVolume` command takes two arguments in the above example - the variable containing the data, and a string giving the filename to which the data should be written to. This MINC file can then be read and viewed with the standard MINC tools such as `mincinfo`, `register`, `Display`, etc.

## 3.2 Creating summaries by group

Most often we are more interested in how the means break down by the grouping in this dataset. This can be done by adding another variable to the `mincMean` call:

```
> group.means <- mincMean(gf$Filename, gf$Gender)

Method: mean
Number of volumes: 10
Volume sizes: 71 137 105
N GROUPS: 2.000000
In slice
 0  1  2  3  4  5  6  7  8  9 10 11 12 13 14 ...
Done

> group.means

Multidimensional MINC volume
Columns:      Female Male
[1] "volumes/img_08nov05.0-fwhm1.0.mnc"
```

The *Gender* variable has two levels in it: *Male* and *Female*. So it will take the mean for all subjects in each group. These can then be written to file by specifying the column.

```
> mincWriteVolume(group.means, "male-mean.mnc", "Male")

Writing column Male to file male-mean.mnc
Sizes: 71 137 105
Range: 0.109804 -0.148817

> mincWriteVolume(group.means, "female-mean.mnc", "Female")

Writing column Female to file female-mean.mnc
Sizes: 71 137 105
Range: 0.193389 -0.174766
```

If the difference between the two columns is of interest, one can just subtract the two data columns:

```
> difference <- group.means[, "Male"] - group.means[, "Female"]
> mean(difference)
```

```
[1] -0.0002146645
```

```
> mincWriteVolume(difference, "diff.mnc", gf$Filename[1])
```

```
Sizes: 71 137 105
```

```
Range: 0.243403 -0.250164
```

Notice how *mincWriteVolume* now needs a third argument: the name of a minc-file which has the same dimensions as the data. By default commands such as *mincMean* will store that information; after the subtraction above, however, the result is just a series of numbers with all metadata removed, so it has to be specified when writing the data to file.

Of course means are not the only items of interest. Also computable are the standard-deviations, variances, and sums, as illustrated below. Just like *mincMean* a column of filenames is required and a grouping variable is optional.

```
> v <- mincVar(gf$Filename, gf$Gender)
```

```
Method: var
```

```
Number of volumes: 10
```

```
Volume sizes: 71 137 105
```

```
N GROUPS: 2.000000
```

```
In slice
```

```
0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 ...
```

```
Done
```

```
> s <- mincSd(gf$Filename)
```

```
Method: var
```

```
Number of volumes: 10
```

```
Volume sizes: 71 137 105
```

```
N GROUPS: 1.000000
```

```
In slice
```

```
0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 ...
```

```
Done
```

```
> s2 <- mincSum(gf$Filename, gf$Gender)
```

```
Method: sum
```

```
Number of volumes: 10
```

```
Volume sizes: 71 137 105
```

```
N GROUPS: 2.000000
```

```
In slice
```

```
0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 ...
```

```
Done
```

## Chapter 4

# Linear Models

Linear models represent the mainstay of structural brain imaging. Their essence is quite simple: the data at every voxel is modelled by a set of terms corresponding to extra information about each scan. One can then perform hypothesis tests on each linear model to calculate the significance of either the entire model or even the marginal significance of each term in the model. This can be used to ask the question of, for example, in which voxels the gender of the subject predicts the values at that voxel.

This approach is also known as *massively univariate statistics* - i.e. a separate linear model is calculated at every voxel, resulting in thousands or even millions of separate models for every statistical test applied to the images. The last step in analysing such data is thus often to account for these thousands of comparisons so that the results do not occur just by random chance.

### 4.1 First linear model

Let's start with a simple linear model - lets see where the Jacobian determinants contained in the files used in the male-female dataset are best modelled by the gender of the mouse.

```
> vs <- mincLm(Filename ~ Gender, gf)

Method: lm
Number of volumes: 10
Volume sizes: 71 137 105
N: 10 P: 2
In slice
 0  1  2  3  4  5  6  7  8  9 10 11 12 13 14 ...
Done

> vs
```

```

Multidimensional MINC volume
Columns:      F-statistic (Intercept) GenderMale
[1] "volumes/img_08nov05.0-fwhm1.0.mnc"

> mincWriteVolume(vs, "simple-lm.mnc", "GenderMale")

Writing column GenderMale to file simple-lm.mnc
Sizes: 71 137 105
Range: 9.581574 -8.354640

```

mincLm is the command to run linear models in RMINC. Its basic use is to provide a formula (same syntax as the R `lm` command) with the left side containing the filenames, the right side the variables to be regressed. The output of mincLm depends on the formula. There will always be a column of F-statistics, representing the significance of the entire model. Then there is one column for each of the terms in the model. The above linear model, relating the Jacobian determinant to gender, will thus have three columns:

**F-statistic** representing the significance of the entire model.

**(Intercept)** the intercept term - this term is rarely interesting, as it tests for whether the intercept is 0. There's no reason to believe it should be in most cases, so this value will be highly significant but meaningless.

**GenderMale** the term testing whether the "Male" level of the Gender factor is significant. In this case this term is the most interesting and therefore the one written to file.

The output is placed into a variable that can be written to file in the same way as described in the descriptive statistics section.

## 4.2 Plotting voxels

```

> options(show.signif.stars = FALSE)
> voxel <- mincGetVoxel(gf$Filename, 44, 20, 52)
> summary(lm(voxel ~ Gender, gf))

```

Call:

```
lm(formula = voxel ~ Gender, data = gf)
```

Residuals:

Min	1Q	Median	3Q	Max
-0.13849	-0.04947	0.01240	0.04120	0.11775

Coefficients:

	Estimate	Std. Error	t value	Pr(> t )
(Intercept)	0.01900	0.03443	0.552	0.596

```
GenderMale -0.08765 0.04869 -1.800 0.110
```

```
Residual standard error: 0.07699 on 8 degrees of freedom  
Multiple R-Squared: 0.2883, Adjusted R-squared: 0.1993  
F-statistic: 3.241 on 1 and 8 DF, p-value: 0.1095
```

```
> vs[635093, ]
```

```
F-statistic (Intercept) GenderMale  
3.2408617 0.5517272 -1.8002394
```

The code above does the following: it gets the voxel from coordinates 44, 20, 52 for all subjects, then computes a linear model relating that voxel to Genotype using standard R functions. Lastly it prints the results from that same voxel as computed by `mincLm`<sup>1</sup>. This helps illustrate what the output of `mincLm` stores: the F-statistic is the same as can be found in the last line of the summary command, and the t-statistics for the Intercept and Genotype column can be found under "t-value" when using standard R functions.

`mincGetVoxel` needs three coordinates, given in voxel space in the same order as stored in the file. Just printing the voxel will show the corresponding world coordinates:

```
> voxel
```

```
[1] 0.13674168 0.04150733 0.07639956 -0.03271777 -0.12701763 -0.06638037  
[7] -0.09147905 -0.02570000 -0.04017891 -0.11949146
```

```
Voxel Coordinates: 44 20 52  
World Coordinates: -0.03 -5.79 1.08
```

If the coordinates are specified in world coordinates then `mincGetWorldVoxel` is what you want - it also takes three coordinates, this time in world space in `xspace,yspace,zspace` order:

```
> world.voxel <- mincGetWorldVoxel(gf$Filename, -6.27, -8.19, -4.2)  
> world.voxel
```

```
[1] 1.525891e-05 1.525891e-05 1.525891e-05 1.525891e-05 1.525891e-05  
[6] 1.525891e-05 1.525891e-05 1.525891e-05 1.525891e-05 1.525891e-05
```

```
Voxel Coordinates: 0 0 0  
World Coordinates: -6.27 -8.19 -4.2
```

<sup>1</sup>The actual number indexed here - 635093 - might appear odd. RMINC treats all MINC volumes as 1-dimensional arrays, so the actual index has to be computed by the following formula:  $(index_1 * size_2 + index_2) + size_3 + index_3$




Figure 4.1: Ray-traced slice




Figure 4.2: Box and whiskers plot from cursor location

### 4.3 Creating images

RMINC has the ability to call `ray_trace` to create images of individual slices corresponding to a specific voxel. For this to work `ray_trace`<sup>2</sup> has to be installed and present in the path, as does `MICe-minc-tools`<sup>3</sup>.

An example is given below - this will create an image of the slice corresponding to the voxel location along with a cross-hair over that voxel. The output can be seen in figure 4.1, along with a box-and-whiskers plot of the data at that voxel.

```
> mincRayTraceStats(voxel, "volumes/anatomy.mnc", vs, "GenderMale",
+   image.min = 350000, image.max = 1e+06, display = F)
```

```
Sizes: 71 137 105
Range: 9.581574 -8.354640
```

The `mincRayTraceStats` function needs the following arguments: a voxel, obtained by `mincGetVoxel` or `mincGetWorldVoxel`, the path towards a MINC image containing some background anatomy, the output of `mincLm` (`vs` in this case), and minimum and maximum values of the background anatomy.

### 4.4 Using subsets

It is quite common to want to run a linear model on only a subset of the data. This can be quite easily accomplished in `mincLm` using an extra subsetting specification:

<sup>2</sup><http://packages.bic.mni.mcgill.ca>

<sup>3</sup><http://wiki.phenogenomics.ca:8080/display/MICePub/MICe-minc-tools>

```

> vs <- mincLm(Filename ~ Gender, gf, coil == 1)

Method: lm
Number of volumes: 4
Volume sizes: 71 137 105
N: 4 P: 2
In slice
 0  1  2  3  4  5  6  7  8  9 10 11 12 13 14 ...
Done

> vs

Multidimensional MINC volume
Columns:      F-statistic (Intercept) GenderMale
[1] "volumes/img_08nov05.0-fwhm1.0.mnc"

```

This is the same linear model command as executed above, but this time using only mice scanned on RF coil number 1. The subset command works exactly the same way as for the standard `lm` command from R.

## 4.5 Multiple Comparisons

The example below illustrates the entire process involved in going running a linear model and correcting for multiple comparisons using the False Discovery Rate:

```

> vs <- mincLm(Filename ~ Gender, gf)

Method: lm
Number of volumes: 10
Volume sizes: 71 137 105
N: 10 P: 2
In slice
 0  1  2  3  4  5  6  7  8  9 10 11 12 13 14 ...
Done

> qvals <- mincFDR(vs, mask = "volumes/mask.mnc", method = "pFDR")

Computing FDR threshold for all columns
Sizes: 71 137 105
Start: 0 0 0
Count: 71 137 105
  Computing threshold for F-statistic
  Computing threshold for (Intercept)
  Computing threshold for GenderMale

> qvals

```



```

Multidimensional MINC volume
Columns:      F-statistic (Intercept) GenderMale
[1] "volumes/img_08nov05.0-fwhm1.0.mnc"
Degrees of Freedom: 1 8
FDR Thresholds:
      F-statistic (Intercept) GenderMale
0.01      NaN      NaN      NaN
0.05      NaN      NaN      NaN
0.1       NaN      NaN      NaN
0.15     8.775918      NaN  2.962418
0.2      5.352080      NaN  2.313456

> mincWriteVolume(qvals, "Gender-FDR.mnc", "GenderMale")

Writing column GenderMale to file Gender-FDR.mnc
Sizes: 71 137 105
Range: 1.000000 0.119666

```

The first command computes a linear model using `mincLm`. The results are then passed on to `mincFDR`, which computes the False Discovery Rate threshold separately for each of the terms in the linear model. Only results from within the mask specified as an optional argument to `mincFDR` are considered. The thresholds detected at different levels (0.01, 0.05, 0.10, 0.15, and 0.20) are then printed out. In this example seen above there is no data at a FDR level of 0.01, 0.05, or 0.10, but any t-statistic greater than 2.96 (or less than -2.96) would be significant at a 15% false positive level - i.e. 15% of the voxels above that threshold would be, on average, false positives. The "GenderMale" column is then written to file. Note that we use the positive false discovery rate above - the more standard false discovery rate is the default (i.e. if no method argument is specified for `mincFDR`)

## Chapter 5

# Bits and Pieces

5.1 Correlations

5.2 non-parametric statistics

5.3 mixed effects models

5.4 Cleaning up

## Chapter 6

# Advanced Topics

### 6.1 Running arbitrary R functions

RMINC has the capacity to run arbitrary R functions at every voxel of a set of files. This comes in quite handy when there are no easily wrapped functions that exist in RMINC but there is some existing R module you would like to try out. Some words of caution are in order, however:

- Running arbitrary functions involves writing your own small R function to wrap the code you want to use, which is a bit ugly.
- It is slow. Slower than molasses on a cold Georgia winter morning.

Here's an example of how it works:

The code above has two parts. The first is creating the function, which has the following properties:

- It takes no arguments.
- It works on the variable  $x$ . This function will be evaluated at every voxel, where  $x$  will be a vector containing the values at that voxel for all the files.
- It returns a vector.