

Indices and the apply commands in R

math612.0 A1: From numbers through algebra to calculus and linear algebra

Gunnar Stefansson (editor) with contributions from very many students

March 7, 2022

Giving names to elements

We can name elements of vectors and data frames in R using the "names" command.

Regular matrix indices and naming

A matrix is a table of numbers. Typical matrix indexing: `mat[i,j]`, `mat[1:2,]` etc

A matrix can have row and column names Indexing with row and column names: `mat["a","B"]`

Example: `dimnames(mat)<-list(c("a","b"),c("1","2","3"))`

The apply command

The apply command...

`apply(mat,2,sum)` – applies the sum function within each column

`apply(mat,1,mean)` – computes the mean within each row

The tapply command

Commonly one has a data vector and another vector of the same length giving categories for the measurements. In this case one often wants to compute the mean or variance (or median etc) within each category. To do this we use the tapply command in R.

Logical indexing

A logical vector consists of *TRUE* (1) or *FALSE* (0) values. These can be used to index vectors or matrices.

Lists, indexing lists

A list is a collection of objects. Thus, data frames are lists.

Copyright 2021, Gunnar Stefansson (editor) with contributions from very many students

This work is licensed under the Creative Commons Attribution-ShareAlike License. To view a copy of this license, visit

<http://creativecommons.org/licenses/by-sa/1.0/> or send a letter to Creative Commons, 559 Nathan Abbott Way, Stanford, California 94305, USA.