

CCREPE: Compositionality Corrected by PERmutation and RENormalization

Emma Schwager, George Weingart, Craig Bielski, Curtis Huttenhower

April 10, 2025

Contents

1	Introduction	1
2	ccrepe	2
2.1	General functionality	2
2.2	Arguments	2
2.3	Output	4
2.4	Usage	4
2.5	Example 1	4
2.6	Example 2	6
2.7	Example 3	8
2.8	Example 4	9
2.9	Example 5	11
3	nc.score	12
3.1	General Functionality	12
3.2	Arguments	12
3.3	Output	13
3.4	Usage	13
3.5	Example 1	13
3.6	Example 2	14
3.7	Example 3	15
4	References	16

1 Introduction

ccrepe is a package for analysis of sparse compositional data. Specifically, it determines the significance of association between features in a composition, using any similarity measure (e.g. Pearson correlation, Spearman correlation, etc.) The CCREPE methodology stands

for Compositionality Corrected by Renormalization and Permutation, as detailed below. The package also provides a novel similarity measure, the N-dimensional checkerboard score (NC-score), particularly appropriate to compositions derived from microbial community sequencing data. This results in p-values and false discovery rate q-values corrected for the effects of compositionality. The package contains two functions `ccrepe` and `nc.score` and is maintained by the Huttenhower lab (ccrepe-users@googlegroups.com).

2 ccrepe

`ccrepe` is the main package function. It calculates compositionality-corrected p-values and q-values for a user-selected similarity measure, operating on either one or two input matrices. If given one matrix, all features (columns) in the matrix are compared to each other using the selected similarity measure. If given two matrices, each feature in the first are compared against all features in the second.

2.1 General functionality

Compositional data induces spurious correlations between features due to the nonindependence of values that must sum to a fixed total. CCREPE abrogates this when determining the significance of a similarity measure for each feature pair using two main steps, permutation/renormalization and bootstrapping. First, given two features to compare, CCREPE generates a null distribution of the similarity expected just due to compositionality by iteratively permuting one feature, renormalizing all samples in the composition to their previous sum, and computing the resulting similarity measures. Second, CCREPE bootstraps over sample subsets in order to assess confidence in the "true" similarity measure. Finally, the two resulting distributions are compared using a pooled-variance Z-test to give a compositionality-corrected p-value. False discovery rate q-values are additionally calculated using the Benjamin-Hochberg-Yekutieli procedure. For greater detail, see [Faust et al. \[2012\]](#) and [Schwager and Colleagues](#).

CCREPE employs several filtering steps before the data are processed. It removes any missing subjects using `na.omit`: in the two dataset case, any subjects missing in *either* dataset will be removed. Any subjects or features which are all zero are removed as well: an all-zero subject cannot be normalized (its sum is 0) and an all-zero feature has standard deviation 0 (in addition to being uninteresting biologically).

2.2 Arguments

- x** First *dataframe* or *matrix* containing relative abundances. Columns are features, rows are samples. Rows should therefore sum to a constant. Row names are used for identification if present.
- y** Second *dataframe* or *matrix* (optional) containing relative abundances. Columns are features, rows are samples. Rows should therefore sum to a constant. If both **x** and **y** are specified, they will be merged by row names. If no row names are specified for either or both datasets, the default is to merge by row number.

CCREPE: Compositionality Corrected by PERmutation and RENormalization

sim.score Similarity measure, such as `cor` or `nc.score`. This can be either an existing R function or user-defined. If the latter, certain properties should be satisfied as detailed below (also see examples). The default similarity measure is Spearman correlation.

A user-defined similarity measure should mimic the interface of `cor`:

1. Take either two *vector* inputs one *matrix* or *dataframe* input.
2. In the case of two inputs, return a single number.
3. In the case of one input, return a matrix in which the (i,j) th entry is the similarity score for column `i` and column `j` in the original matrix.
4. The resulting matrix (in the case of one input) must be symmetric.
5. The inputs must be named `x` and `y`.

sim.score.args An optional list of arguments for the measurement function. When given, they are passed to the `sim.score` function directly. For example, in the case of `cor`, the following would be acceptable:

```
sim.score.args = list(method="spearman", use="complete.obs")
```

min.subj Minimum number (count) of samples that must be non-missing in order to apply the similarity measure. This is to ensure that there are sufficient samples to perform a bootstrap (default: 20).

iterations The number of iterations for both bootstrap and permutation calculations (default: 1000).

subset.cols.x A vector of column indices from `x` to indicate which features to compare

subset.cols.y A vector of column indices from `y` to indicate which features to compare

errthresh If feature has number of zeros greater than $errthresh^{1/n}$, that feature is excluded

verbose If `TRUE`, print periodic progress of the algorithm through the dataset(s), as well as including more detailed debugging output. (default: `FALSE`).

iterations.gap If `verbose=TRUE`, the number of iterations between issuing status messages (default: 100).

distributions Optional output file for detailed log (if given) of all intermediate permutation and renormalization distributions.

compare.within.x A boolean value indicating whether to do comparisons given by taking all subsets of size 2 from `subset.cols.x` or to do comparisons given by taking all possible combinations of `subset.cols.x` and `subset.cols.y`. If `TRUE` but `subset.cols.y=NA`, returns all comparisons involving any features in `subset.cols.x`. This argument is only used when `y=NA`.

CCREPE: Compositionality Corrected by PERmutation and RENormalization

concurrent.output Optional output file to which each comparison will be written as it is calculated.

make.output.table A boolean value indicating whether to include table-formatted output.

2.3 Output

`ccrepe` returns a *list* containing both the calculation results and the parameters used:

sim.score *matrix* of similarity scores for all requested comparisons. The (i,j) th element corresponds to the similarity score of column i (or the i th column of `subset.cols.1`) and column j (or the j th column of `subset.cols.1`) in one dataset, or to the similarity score of column i (or the i th column of `subset.cols.1`) in dataset x and column j (or the j th column of `subset.cols.2`) in dataset y in the case of two datasets.

p.values *matrix* of the corrected p-values for all requested comparisons. The (i,j) th element corresponds to the p-value of the (i,j) th element of `sim.score`.

q.values *matrix* of the Benjamini-Hochberg-Yekutieli corrected p-values. The (i,j) th element corresponds to the p-value of the (i,j) th element of `sim.score`.

z.stat *matrix* of the z-statistics used in generating the p-values for all requested comparisons. The (i,j) th element corresponds to the z-statistic generating the (i,j) th element of `p.values`.

2.4 Usage

```
ccrepe(  
  x = NA,  
  y = NA,  
  sim.score = cor,  
  sim.score.args = list(),  
  min.subj = 20,  
  iterations = 1000,  
  subset.cols.x = NULL,  
  subset.cols.y = NULL,  
  errthresh = 1e-04,  
  verbose = FALSE,  
  iterations.gap = 100,  
  distributions = NA,  
  compare.within.x = TRUE,  
  concurrent.output = NA,  
  make.output.table = FALSE)
```

2.5 Example 1

An example of how to use `ccrepe` with one dataset.

CCREPE: Compositionality Corrected by PERmutation and REnormalization

```
data <- matrix(rlnorm(40,meanlog=0,sdlog=1),nrow=10,ncol=4)
data[,1] = 2*data[,2] + rnorm(10,0,0.01)
data.rowsum <- apply(data,1,sum)
data.norm <- data/data.rowsum
apply(data.norm,1,sum) # The rows sum to 1, so the data are normalized

## [1] 1 1 1 1 1 1 1 1 1 1

test.input <- data.norm

dimnames(test.input) <- list(c(
  "Sample 1", "Sample 2", "Sample 3", "Sample 4", "Sample 5",
  "Sample 6", "Sample 7", "Sample 8", "Sample 9", "Sample 10"),
  c("Feature 1", "Feature 2", "Feature 3", "Feature 4"))

test.output <- ccrepe(x=test.input, iterations=20, min.subj=10)
```

```
par(mfrow=c(1,2))
plot(data[,1],data[,2],xlab="Feature 1",ylab="Feature 2",main="Non-normalized")
plot(data.norm[,1],data.norm[,2],xlab="Feature 1",ylab="Feature 2",
      main="Normalized")
```




Figure 1: Non-normalized and normalized associations between feature 1 and feature 2. In this case we would expect feature 1 and feature 2 to be associated. In the output we see this by the positive sim.score value in the [1,2] element of test.output\$sim.score and the small q-value in the [1,2] element of test.output\$q.values.

```
test.output

## $p.values
##           Feature 1   Feature 2   Feature 3   Feature 4
## Feature 1          NA 2.605617e-09 0.34603949 0.8247462
## Feature 2 2.605617e-09          NA 0.08277332 0.6054984
## Feature 3 3.460395e-01 8.277332e-02          NA 0.6335098
## Feature 4 8.247462e-01 6.054984e-01 0.63350979          NA
##
## $z.stat
##           Feature 1   Feature 2   Feature 3   Feature 4
```

CCREPE: Compositionality Corrected by PERmutation and REnormalization

```
## Feature 1      NA  5.9546923 -0.9422992 -0.2214446
## Feature 2  5.9546923      NA -1.7348164 -0.5165098
## Feature 3 -0.9422992 -1.7348164      NA  0.4767926
## Feature 4 -0.2214446 -0.5165098  0.4767926      NA
##
## $sim.score
##           Feature 1 Feature 2 Feature 3 Feature 4
## Feature 1      NA  0.9998807 -0.8109757 -0.3705813
## Feature 2  0.9998807      NA -0.8124773 -0.3681269
## Feature 3 -0.8109757 -0.8124773      NA -0.2428678
## Feature 4 -0.3705813 -0.3681269 -0.2428678      NA
##
## $q.values
##           Feature 1 Feature 2 Feature 3 Feature 4
## Feature 1      NA  3.703585e-08  1.6395179  1.953803
## Feature 2  3.703585e-08      NA  0.5882638  2.151616
## Feature 3  1.639518e+00  5.882638e-01      NA  1.800923
## Feature 4  1.953803e+00  2.151616e+00  1.8009227      NA
```

2.6 Example 2

An example of how to use `ccrepe` with two datasets.

```
data <- matrix(rlnorm(40,meanlog=0,sdlog=1),nrow=10,ncol=4)
data[,1] = 2*data[,2] + rnorm(10,0,0.01)
data.rowsum <- apply(data,1,sum)
data.norm <- data/data.rowsum
apply(data.norm,1,sum) # The rows sum to 1, so the data are normalized
## [1] 1 1 1 1 1 1 1 1 1 1

test.input <- data.norm

data2 <- matrix(rlnorm(105,meanlog=0,sdlog=1),nrow=15,ncol=7)
aligned.rows <- c(seq(1,4),seq(6,9),11,12) # The datasets dont need
                                           # to have subjects line up exactly
data2[aligned.rows,1] <- 2*data[,3] + rnorm(10,0,0.01)
data2.rowsum <- apply(data2,1,sum)
data2.norm <- data2/data2.rowsum
apply(data2.norm,1,sum) # The rows sum to 1, so the data are normalized
## [1] 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1

test.input.2 <- data2.norm

dimnames(test.input) <- list(paste("Sample",seq(1,10)),paste("Feature",seq(1,4)))
dimnames(test.input.2) <- list(paste("Sample",c(seq(1,4),11,seq(5,8),12,9,10,13,14,15)),paste("Feature",seq(1,4)))

test.output.two.datasets <- ccrepe(x=test.input, y=test.input.2, iterations=20, min.subj=10)

## Warning in preprocess_data(CA): Removing subjects Sample 11, Sample 12, Sample
13, Sample 14, Sample 15 from dataset y because they are not in dataset x.
```

CCREPE: Compositionality Corrected by PERmutation and REnormalization

Please note that we receive a warning because the subjects don't match - only paired observations.

```
par(mfrow=c(1,2))
plot(data2[aligned.rows,1],data[,3],xlab="dataset 2: Feature 1",ylab="dataset 1: Feature 3",main="Non-normalized")
plot(data2.norm[aligned.rows,1],data.norm[,3],xlab="dataset 2: Feature 1",ylab="dataset 1: Feature 3",
      main="Normalized")
```




Figure 2: Non-normalized and normalized associations between feature 1 and feature 2. In this case we would expect feature 1 and feature 2 to be associated. In the output we see this by the positive `sim.score` value in the [1,2] element of `test.output$sim.score` and the small `q`-value in the [1,2] element of `test.output$q.values`.

```
test.output.two.datasets

## $p.values
##           Feature 1 Feature 2 Feature 3 Feature 4 Feature 5 Feature 6
## Feature 1 3.480021e-01 0.7411421 0.07955072 0.5126194 0.9737774 0.49018447
## Feature 2 5.915529e-01 0.9146306 0.05301933 0.6985538 0.9581907 0.39949764
## Feature 3 2.025718e-05 0.4515931 0.11159050 0.4296769 0.7575814 0.04317981
## Feature 4 7.601334e-02 0.5175823 0.69304432 0.2781050 0.8612398 0.01028782
##           Feature 7
## Feature 1 0.204441611
## Feature 2 0.503287814
## Feature 3 0.007088164
## Feature 4 0.340710881
##
## $z.stat
##           Feature 1 Feature 2 Feature 3 Feature 4 Feature 5 Feature 6
## Feature 1 -0.9384717 0.3303413 1.7532989 -0.6547645 0.03287103 -0.6900155
## Feature 2 -0.5365870 -0.1071997 1.9347635 -0.3872735 -0.05242414 -0.8425188
## Feature 3 4.2620371 0.7527618 -1.5910848 -0.7897447 0.30865838 -2.0219665
## Feature 4 -1.7743012 -0.6470767 -0.3947271 1.0845861 -0.17479620 2.5660020
##           Feature 7
## Feature 1 1.2689985
## Feature 2 0.6693256
## Feature 3 -2.6926735
## Feature 4 0.9527616
##
```

CCREPE: Compositionality Corrected by PERmutation and REnormalization

```
## $sim.score
##           Feature 1   Feature 2   Feature 3   Feature 4   Feature 5   Feature 6
## Feature 1 -0.1370560  0.07621881  0.6109183 -0.2811545 -0.01896809 -0.4805084
## Feature 2 -0.1289584  0.07659917  0.6059742 -0.2784191 -0.02017105 -0.4866583
## Feature 3  0.8563942  0.07722524 -0.3362744 -0.1945173  0.17851110 -0.5141231
## Feature 4 -0.3967216 -0.11509168 -0.3432120  0.3687399 -0.09043239  0.7431737
##           Feature 7
## Feature 1  0.32664961
## Feature 2  0.32908150
## Feature 3 -0.58173207
## Feature 4  0.05768553
##
## $q.values
##           Feature 1 Feature 2 Feature 3 Feature 4 Feature 5 Feature 6
## Feature 1 3.174467975  3.527313  1.243989  3.117403  3.806905  3.3535898
## Feature 2 3.237680426  3.850727  1.160739  3.475742  3.884710  3.3638858
## Feature 3 0.002217427  3.295538  1.526889  3.359575  3.455321  1.1816562
## Feature 4 1.386784461  2.981921  3.612535  3.044242  3.770982  0.3753813
##           Feature 7
## Feature 1 2.4865498
## Feature 2 3.2406929
## Feature 3 0.3879486
## Feature 4 3.3904996
```

2.7 Example 3

An example of how to use `ccrepe` with `nc.score` as the similarity score.

```
data <- matrix(rlnorm(40,meanlog=0,sdlog=1),nrow=10,ncol=4)
data[,1] = 2*data[,2] + rnorm(10,0,0.01)
data.rowsum <- apply(data,1,sum)
data.norm <- data/data.rowsum
apply(data.norm,1,sum) # The rows sum to 1, so the data are normalized
## [1] 1 1 1 1 1 1 1 1 1 1

test.input <- data.norm

dimnames(test.input) <- list(paste("Sample",seq(1,10)),paste("Feature",seq(1,4)))

test.output.nc.score <- ccrepe(x=test.input, sim.score=nc.score, iterations=20, min.subj=10)

par(mfrow=c(1,2))
plot(data[,1],data[,2],xlab="Feature 1",ylab="Feature 2",main="Non-normalized")
plot(data.norm[,1],data.norm[,2],xlab="Feature 1",ylab="Feature 2",
      main="Normalized")
```


CCREPE: Compositionality Corrected by PERmutation and REnormalization




Figure 3: Non-normalized and normalized associations between feature 1 and feature 2. In this case we would expect feature 1 and feature 2 to be associated. In the output we see this by the positive `sim.score` value in the [1,2] element of `test.output$sim.score` and the small `q-value` in the [1,2] element of `test.output$q.values`. In this case, however, the `sim.score` represents the NC-Score between two features rather than the Spearman correlation.

```
test.output.nc.score

## $p.values
##           Feature 1  Feature 2 Feature 3 Feature 4
## Feature 1          NA 0.000480363 0.5587809 0.8405978
## Feature 2 0.000480363          NA 0.5368075 0.5102915
## Feature 3 0.558780889 0.536807521          NA 0.6486906
## Feature 4 0.840597763 0.510291518 0.6486906          NA
##
## $z.stat
##           Feature 1  Feature 2  Feature 3  Feature 4
## Feature 1          NA  3.4914741 -0.5846533 -0.2011289
## Feature 2  3.4914741          NA -0.6176478 -0.6583838
## Feature 3 -0.5846533 -0.6176478          NA  0.4555820
## Feature 4 -0.2011289 -0.6583838  0.4555820          NA
##
## $sim.score
##           Feature 1 Feature 2 Feature 3 Feature 4
## Feature 1          NA  1.00000 -0.59375 -0.34375
## Feature 2  1.00000          NA -0.59375 -0.34375
## Feature 3 -0.59375 -0.59375          NA  0.00000
## Feature 4 -0.34375 -0.34375  0.00000          NA
##
## $q.values
##           Feature 1  Feature 2 Feature 3 Feature 4
## Feature 1          NA 0.006827808 1.985607 1.991355
## Feature 2 0.006827808          NA 2.543367 3.626604
## Feature 3 1.985607049 2.543367337          NA 1.844078
## Feature 4 1.991355199 3.626603749 1.844078          NA
```

2.8 Example 4

An example of how to use `ccrepe` with a user-defined `sim.score` function.

CCREPE: Compositionality Corrected by PERmutation and REnormalization

```
data <- matrix(rlnorm(40,meanlog=0,sdlog=1),nrow=10,ncol=4)
data[,1] = 2*data[,2] + rnorm(10,0,0.01)
data.rowsum <- apply(data,1,sum)
data.norm <- data/data.rowsum
apply(data.norm,1,sum) # The rows sum to 1, so the data are normalized

## [1] 1 1 1 1 1 1 1 1 1 1

test.input <- data.norm

dimnames(test.input) <- list(paste("Sample",seq(1,10)),paste("Feature",seq(1,4)))

my.test.sim.score <- function(x,y=NA,constant=0.5){
  if(is.vector(x) && is.vector(y)) return(constant)
  if(is.matrix(x) && is.na(y)) return(matrix(rep(constant,ncol(x)^2),ncol=ncol(x)))
  if(is.data.frame(x) && is.na(y)) return(matrix(rep(constant,ncol(x)^2),ncol=ncol(x)))
  else stop('ERROR')
}

test.output.sim.score <- ccrepe(x=test.input, sim.score=my.test.sim.score, iterations=20, min.subj=10, sim
```

```
par(mfrow=c(1,2))
plot(data[,1],data[,2],xlab="Feature 1",ylab="Feature 2",main="Non-normalized")
plot(data.norm[,1],data.norm[,2],xlab="Feature 1",ylab="Feature 2",
      main="Normalized")
```




Figure 4: Non-normalized and normalized associations between feature 1 and feature 2. In this case we would expect feature 1 and feature 2 to be associated. Note that the values of sim.score are all 0.6 and none of the p-values are very small because of the arbitrary definition of the similarity score.

```
test.output.sim.score

## $p.values
##           Feature 1 Feature 2 Feature 3 Feature 4
## Feature 1         NA        NaN        NaN        NaN
## Feature 2         NaN         NA        NaN        NaN
## Feature 3         NaN         NaN         NA        NaN
## Feature 4         NaN         NaN         NaN         NA
```

CCREPE: Compositionality Corrected by PERmutation and RENormalization

```
##
## $z.stat
##           Feature 1 Feature 2 Feature 3 Feature 4
## Feature 1         NA      NaN      NaN      NaN
## Feature 2         NaN       NA      NaN      NaN
## Feature 3         NaN      NaN       NA      NaN
## Feature 4         NaN      NaN      NaN       NA
##
## $sim.score
##           Feature 1 Feature 2 Feature 3 Feature 4
## Feature 1         NA      0.6      0.6      0.6
## Feature 2         0.6      NA      0.6      0.6
## Feature 3         0.6      0.6      NA      0.6
## Feature 4         0.6      0.6      0.6      NA
##
## $q.values
##           Feature 1 Feature 2 Feature 3 Feature 4
## Feature 1         NA      NaN      NaN      NaN
## Feature 2         NaN       NA      NaN      NaN
## Feature 3         NaN      NaN       NA      NaN
## Feature 4         NaN      NaN      NaN      NA
```

2.9 Example 5

An example of how to use `ccrepe` when specifying column subsets.

```
data <- matrix(rlnorm(40,meanlog=0,sdlog=1),nrow=10,ncol=4)
data.rowsum <- apply(data,1,sum)
data.norm <- data/data.rowsum
apply(data.norm,1,sum) # The rows sum to 1, so the data are normalized
## [1] 1 1 1 1 1 1 1 1 1 1

test.input <- data.norm

dimnames(test.input) <- list(paste("Sample",seq(1,10)),paste("Feature",seq(1,4)))

test.output.1.3 <- ccrepe(x=test.input, iterations=20, min.subj=10, subset.cols.x=c(1,3))
test.output.1 <- ccrepe(x=test.input, iterations=20, min.subj=10, subset.cols.x=c(1), compare.within.x=
test.output.12.3 <- ccrepe(x=test.input, iterations=20, min.subj=10, subset.cols.x=c(1,2),subset.cols.y=c
test.output.1.3$sim.score

##           Feature 1 Feature 3
## Feature 1         NA -0.1410158
## Feature 3 -0.1410158         NA

test.output.1$sim.score

##           Feature 1 Feature 2 Feature 3 Feature 4
## Feature 1         NA 0.06541922 -0.1410158 -0.5522351
## Feature 2 0.06541922         NA      NA      NA
## Feature 3 -0.14101580         NA      NA      NA
## Feature 4 -0.55223514         NA      NA      NA
```

```
test.output.12.3$sim.score
##           Feature 1 Feature 2 Feature 3 Feature 4
## Feature 1          NA          NA -0.1410158      NA
## Feature 2          NA          NA  0.1230706      NA
## Feature 3 -0.1410158  0.1230706          NA      NA
## Feature 4          NA          NA          NA      NA
```

3 nc.score

The `nc.score` similarity measure is an N-dimensional extension of the checkerboard score particularly suited to similarity score calculations between compositions derived from ecological relative abundance measurements. In such cases, features typically represent species abundances, and the NC-score discretizes these continuous values into one of N bins before computing a normalized similarity of co-occurrence or co-exclusion. This can be used as a standalone function or with `ccrepe` as above to obtain compositionality-corrected p-values.

3.1 General Functionality

The NC-score is an extension to Diamond's checkerboard score (see [Cody and Diamond \[1975\]](#)) to ordinal data, and simplifies to a calculation of Kendall's τ on binned data instead of ranked data. Let two features in a dataset with n subjects be denoted by

$$\begin{bmatrix} x_1 & x_2 & \dots & x_n \\ y_1 & y_2 & \dots & y_n \end{bmatrix}.$$

The binning function maps from the original data to b numbered bins in $\{1, \dots, b\}$. Let the binning function be denoted by $B(\cdot)$. The co-variation and co-exclusion patterns are the same as concordant and discordant pairs in Kendall's τ . Considering a 2×2 submatrix of the form

$$\begin{bmatrix} B(x_i) & B(x_j) \\ B(y_i) & B(y_j) \end{bmatrix},$$

a co-variation pattern is counted when $(B(x_i) - B(x_j))(B(y_i) - B(y_j)) > 0$ and a co-exclusion pattern, conversely, when $(B(x_i) - B(x_j))(B(y_i) - B(y_j)) < 0$. The NC-score statistic for features x and y is then defined as

$$(\text{number of co-variation patterns}) - (\text{number of co-exclusion patterns}),$$

normalized by the Kendall's τ normalization factor accounting for ties described in [Kendall \[1970\]](#).

3.2 Arguments

- x** First numerical *vector*, or single *dataframe* or *matrix*, containing relative abundances. If the latter, columns are features, rows are samples. Rows should therefore sum to a constant.
- y** If provided, second numerical *vector* containing relative abundances. If given, **x** must be a *vector* as well.

CCREPE: Compositionality Corrected by PERmutation and RENormalization

nbins A non-negative integer of the number of bins to generate (cutoffs will be generated by the `discretize` function from the `infotheo` package).

bin.cutoffs A list of values demarcating the bin cutoffs. The binning is performed using the `findInterval` function.

use An optional character string giving a method for computing covariances in the presence of missing values. This must be (an abbreviation of) one of the strings "everything", "all.obs", "complete.obs", "na.or.complete", or "pairwise.complete.obs".

3.3 Output

`nc.score` returns either a single number (if called with two vectors) or a *matrix* of all pairwise scores (if called with a *matrix*) of normalized scores. This behaviour is precisely analogous to the `cor` function in R

3.4 Usage

```
nc.score(  
  x,  
  y = NULL,  
  use = "everything",  
  nbins = NULL,  
  bin.cutoffs=NULL)
```

3.5 Example 1

An example of using `nc.score` to get a single similarity score or a matrix.

```
data <- matrix(rlnorm(40,meanlog=0,sdlog=1),nrow=10,ncol=4)  
data.rowsum <- apply(data,1,sum)  
data[,1] = 2*data[,2] + rnorm(10,0,0.01)  
data.norm <- data/data.rowsum  
apply(data.norm,1,sum) # The rows sum to 1, so the data are normalized  
  
## [1] 0.8262706 1.8301154 2.3811958 1.1868199 1.3562002 0.9037219 1.5543819  
## [8] 1.3411014 0.7991590 0.8435433  
  
test.input <- data.norm  
  
dimnames(test.input) <- list(paste("Sample",seq(1,10)),paste("Feature",seq(1,4)))  
  
test.output.matrix <- nc.score(x=test.input)  
test.output.num <- nc.score(x=test.input[,1],y=test.input[,2])  
  
par(mfrow=c(1, 2))  
plot(data[,1],data[,2],xlab="Feature 1",ylab="Feature 2",main="Non-normalized")  
plot(data.norm[,1],data.norm[,2],xlab="Feature 1",ylab="Feature 2",  
      main="Normalized")
```




Figure 5: Non-normalized and normalized associations between feature 1 and feature 2 of the second example. Again, we expect to observe a positive association between feature 1 and feature 2. In terms of generalized checkerboard scores, we would expect to see more co-variation patterns than co-exclusion patterns. This is shown by the positive and relatively high value of the [1,2] element of `test.output.matrix` (which is identical to `test.output.num`)

```
test.output.matrix
##           Feature 1 Feature 2 Feature 3 Feature 4
## Feature 1      1.0000      1.0000     -0.4375      0.125
## Feature 2      1.0000      1.0000     -0.4375      0.125
## Feature 3     -0.4375     -0.4375      1.0000      0.125
## Feature 4      0.1250      0.1250      0.1250      1.000

test.output.num
## [1] 1
```

3.6 Example 2

An example of using `nc.score` with an arbitrary bin number.

```
data <- matrix(rlnorm(40,meanlog=0,sdlog=1),nrow=10,ncol=4)
data.rowsum <- apply(data,1,sum)
data[,1] = 2*data[,2] + rnorm(10,0,0.01)
data.norm <- data/data.rowsum
apply(data.norm,1,sum) # The rows sum to 1, so the data are normalized

## [1] 1.1557319 1.0101240 1.2338538 1.2184481 0.6827354 0.8007767 1.2987867
## [8] 1.2897660 1.1023362 1.0013838

test.input <- data.norm

dimnames(test.input) <- list(paste("Sample",seq(1,10)),paste("Feature",seq(1,4)))

test.output <- nc.score(x=test.input,nbins=4)

par(mfrow=c(1, 2))
plot(data[,1],data[,2],xlab="Feature 1",ylab="Feature 2",main="Non-normalized")
plot(data.norm[,1],data.norm[,2],xlab="Feature 1",ylab="Feature 2",
      main="Normalized")
```




Figure 6: Non-normalized and normalized associations between feature 1 and feature 2 of the second example. Again, we expect to observe a positive association between feature 1 and feature 2. In terms of generalized checkerboard scores, we would expect to see more co-variation patterns than co-exclusion patterns. This is shown by the positive and relatively high value in the [1,2] element of test.output. In this case, the smaller bin number yields a smaller NC-score because of the coarser partitioning of the data.

```
test.output

##           Feature 1  Feature 2  Feature 3  Feature 4
## Feature 1  1.0000000  1.0000000  0.2285714 -0.4285714
## Feature 2  1.0000000  1.0000000  0.2285714 -0.4285714
## Feature 3  0.2285714  0.2285714  1.0000000 -0.7428571
## Feature 4 -0.4285714 -0.4285714 -0.7428571  1.0000000
```

3.7 Example 3

An example of using `nc.score` with user-defined bin edges.

```
data <- matrix(rlnorm(40,meanlog=0,sdlog=1),nrow=10,ncol=4)
data.rowsum <- apply(data,1,sum)
data[,1] = 2*data[,2] + rnorm(10,0,0.01)
data.norm <- data/data.rowsum
apply(data.norm,1,sum) # The rows sum to 1, so the data are normalized

## [1] 0.8761373 0.9554427 1.0175996 1.3111933 1.8688742 1.3550883 0.7253182
## [8] 1.0288839 0.9343528 2.0943229

test.input <- data.norm

dimnames(test.input) <- list(paste("Sample",seq(1,10)),paste("Feature",seq(1,4)))

test.output <- nc.score(x=test.input,bin.cutoffs=c(0.1,0.2,0.3))
```

```
par(mfrow=c(1, 2))
plot(data[,1],data[,2],xlab="Feature 1",ylab="Feature 2",main="Non-normalized")
plot(data.norm[,1],data.norm[,2],xlab="Feature 1",ylab="Feature 2",
      main="Normalized")
```




Figure 7: Non-normalized and normalized associations between feature 1 and feature 2 of the second example. Again, we expect to observe a positive association between feature 1 and feature 2. In terms of generalized checkerboard scores, we would expect to see more co-variation patterns than co-exclusion patterns. This is shown by the positive and relatively high value in the [1,2] element of test.output. The bin edges specified here represent almost absent ([0,0.001)), low abundance ([0.001,0.1)), medium abundance ([0.1,0.25)), and high abundance ([0.6,1)).

```
test.output
```

```
##          Feature 1 Feature 2 Feature 3 Feature 4
## Feature 1  1.0000000  0.7198707 -0.5991713 -0.4934352
## Feature 2  0.7198707  1.0000000 -0.6391355 -0.4724045
## Feature 3 -0.5991713 -0.6391355  1.0000000  0.3714286
## Feature 4 -0.4934352 -0.4724045  0.3714286  1.0000000
```

4 References

References

- Martin Leonard Cody and Jared Mason Diamond. *Ecology and evolution of communities*. Harvard University Press, 1975.
- Karoline Faust, J Fah Sathirapongsasuti, Jacques Izard, Nicola Segata, Dirk Gevers, Jeroen Raes, and Curtis Huttenhower. Microbial co-occurrence relationships in the human microbiome. *PLoS computational biology*, 8(7):e1002606, 2012.
- M.G. Kendall. *Rank correlation methods*. Charles Griffin & Co., 1970.
- Emma Schwager and Colleagues. Detecting statistically significant associations between sparse and high dimensional compositional data. In Progress.