

# CCREPE: Compositionality Corrected by PErmutation and REnormalization

Emma Schwager, George Weingart, Craig Bielski, Curtis Huttenhower

April 16, 2015

## Contents

---

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>ccrepe</b>	<b>2</b>
2.1	General functionality . . . . .	2
2.2	Arguments . . . . .	2
2.3	Output . . . . .	3
2.4	Usage . . . . .	3
2.5	Example 1 . . . . .	4
2.6	Example 2 . . . . .	6
2.7	Example 3 . . . . .	8
2.8	Example 4 . . . . .	10
2.9	Example 5 . . . . .	12
<b>3</b>	<b>nc.score</b>	<b>12</b>
3.1	General Functionality . . . . .	12
3.2	Arguments . . . . .	13
3.3	Output . . . . .	13
3.4	Usage . . . . .	13
3.5	Example 1 . . . . .	13
3.6	Example 2 . . . . .	15
3.7	Example 3 . . . . .	16
<b>4</b>	<b>References</b>	<b>17</b>

## 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](mailto: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.

`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`.

`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.

```
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 1.930303e-05 0.51197854 0.25647049
## Feature 2 1.930303e-05           NA 0.06259337 0.02157606
## Feature 3 5.119785e-01 6.259337e-02           NA 0.36023066
## Feature 4 2.564705e-01 2.157606e-02 0.36023066           NA
##
## $z.stat
##           Feature 1 Feature 2   Feature 3   Feature 4
## Feature 1           NA  4.272804 -0.6557600 -1.1347729
## Feature 2  4.272804           NA -1.8620690 -2.2977494
## Feature 3 -0.655760 -1.862069           NA  0.9149257
## Feature 4 -1.134773 -2.297749  0.9149257           NA
##
## $sim.score
##           Feature 1   Feature 2   Feature 3   Feature 4
## Feature 1           NA  0.9999647 -0.42079543 -0.89802315
## Feature 2  0.9999647           NA -0.42254492 -0.89716039
## Feature 3 -0.4207954 -0.4225449           NA -0.02120812
## Feature 4 -0.8980232 -0.8971604 -0.02120812           NA
##
## $q.values
##           Feature 1   Feature 2   Feature 3   Feature 4
## Feature 1           NA 0.0002743704 1.2128644 0.9113583
## Feature 2 0.0002743704           NA 0.2965643 0.1533394
## Feature 3 1.2128644269 0.2965642744           NA 1.0240530
```

```
## Feature 4 0.9113583198 0.1533394267 1.0240530 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(1,7)))

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.
```

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 0.05682613 0.3204305 0.3629247 0.12021341 0.2676471 0.9697994
## Feature 2 0.07099285 0.3243621 0.3994233 0.20464132 0.1636875 0.8253625
## Feature 3 0.00121750 0.7257399 0.8379971 0.02776478 0.2375613 0.3599061
## Feature 4 0.59666601 0.1750136 0.1780236 0.52871478 0.5470131 0.3070759
##           Feature 7
## Feature 1 0.12158216
## Feature 2 0.08851789
## Feature 3 0.42375403
## Feature 4 0.08090836
##
## $z.stat
##           Feature 1 Feature 2 Feature 3 Feature 4 Feature 5 Feature 6
## Feature 1 -1.904646  0.9935735 -0.9098063  1.5538784  1.108497 -0.03785992
## Feature 2 -1.805523  0.9855335 -0.8426517  1.2684387  1.392776 -0.22065312
## Feature 3  3.234748 -0.3507980  0.2044561 -2.2005939 -1.181104 -0.91554394
## Feature 4  0.529201 -1.3562689  1.3468653 -0.6299691 -0.602242  1.02137687
##           Feature 7
## Feature 1  1.5481667
## Feature 2  1.7032682
## Feature 3 -0.7999254
## Feature 4 -1.7454397
##
## $sim.score
##           Feature 1 Feature 2 Feature 3 Feature 4 Feature 5
```

```
## Feature 1 -0.48838916  0.3901305 -0.35517480  0.47382793  0.2986995
## Feature 2 -0.49511106  0.3907386 -0.34867384  0.47957357  0.3036854
## Feature 3  0.79159785 -0.1772521  0.06736547 -0.74215805 -0.2946334
## Feature 4 -0.08375476 -0.4202849  0.48011511  0.05385856 -0.1531355
##           Feature 6  Feature 7
## Feature 1 -0.1549779  0.48314007
## Feature 2 -0.1598618  0.49223269
## Feature 3 -0.3971089 -0.09308043
## Feature 4  0.6724047 -0.66084358
##
## $q.values
##           Feature 1 Feature 2 Feature 3 Feature 4 Feature 5 Feature 6
## Feature 1 2.0734674  2.192221  2.090900  1.879859  2.092690  3.791353
## Feature 2 1.9427861  2.088582  2.186119  1.866734  1.990873  3.474896
## Feature 3 0.1332721  3.177689  3.397419  1.519619  2.000335  2.188705
## Feature 4 2.7213878  1.915765  1.771558  2.630687  2.603396  2.240912
##           Feature 7
## Feature 1  1.663605
## Feature 2  1.614917
## Feature 3  2.208843
## Feature 4  1.771307
```

## 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")
```






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 3.591133e-07 0.9125525 0.5890798
## Feature 2 3.591133e-07          NA 0.8290454 0.1505923
## Feature 3 9.125525e-01 8.290454e-01          NA 0.8574796
## Feature 4 5.890798e-01 1.505923e-01 0.8574796          NA
##
## $z.stat
##           Feature 1 Feature 2 Feature 3 Feature 4
## Feature 1          NA  5.089445 -0.1098196 -0.5401701
## Feature 2  5.0894448          NA -0.2159260 -1.4374423
## Feature 3 -0.1098196 -0.215926          NA  0.1795835
## Feature 4 -0.5401701 -1.437442  0.1795835          NA
##
## $sim.score
##           Feature 1 Feature 2 Feature 3 Feature 4
## Feature 1          NA  1.00000 -0.21875 -0.8125
## Feature 2  1.00000          NA -0.21875 -0.8125
## Feature 3 -0.21875 -0.21875          NA  0.0000
## Feature 4 -0.81250 -0.81250  0.00000          NA
##
## $q.values
##           Feature 1      Feature 2 Feature 3 Feature 4
## Feature 1          NA 5.104383e-06 2.161814 2.791031
## Feature 2 5.104383e-06          NA 2.945982 1.070249
```

```
## Feature 3 2.161814e+00 2.945982e+00      NA 2.437617
## Feature 4 2.791031e+00 1.070249e+00 2.437617      NA
```

## 2.8 Example 4

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

```
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,

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

##           Feature 1  Feature 3
## Feature 1          NA -0.3878881
## Feature 3 -0.3878881          NA

test.output.1$sim.score

##           Feature 1  Feature 2  Feature 3  Feature 4
## Feature 1          NA -0.2053181 -0.3878881 -0.1938784
## Feature 2 -0.2053181          NA          NA          NA
## Feature 3 -0.3878881          NA          NA          NA
## Feature 4 -0.1938784          NA          NA          NA

test.output.12.3$sim.score

##           Feature 1  Feature 2  Feature 3  Feature 4
## Feature 1          NA          NA -0.3878881          NA
## Feature 2          NA          NA -0.5915527          NA
## Feature 3 -0.3878881 -0.5915527          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  $\tau$  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  $\tau$ . Considering a  $2 \times 2$  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  $\tau$  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.
- `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] 1.7256506 0.4512462 0.9471490 1.2948010 1.2588558 1.0398838 1.4378445
## [8] 1.2644081 1.5075647 1.3011969

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   1.00000   1.00000   0.00000  -0.21875
## Feature 2   1.00000   1.00000   0.00000  -0.21875
## Feature 3   0.00000   0.00000   1.00000  -0.34375
## Feature 4  -0.21875  -0.21875  -0.34375   1.00000

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] 1.2648922 0.5592368 0.7803245 2.4944355 1.1214601 1.7917090 1.2315396
## [8] 0.5239878 0.9938638 1.0210325

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  1.00000000  1.00000000  0.1142857 -0.05714286
## Feature 2  1.00000000  1.00000000  0.1142857 -0.05714286
## Feature 3  0.11428571  0.11428571  1.0000000 -0.42857143
## Feature 4 -0.05714286 -0.05714286 -0.4285714  1.00000000
```

### 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] 0.7274986 0.5405594 1.2925826 0.7735517 1.0430980 1.7555114 1.5402231
## [8] 0.8766374 1.0716375 1.7556880

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 ( $[0.25,1)$ ).

```
test.output

##           Feature 1  Feature 2  Feature 3  Feature 4
## Feature 1  1.0000000  0.7533178 -0.2824942  0.1334075
## Feature 2  0.7533178  1.0000000 -0.1428571 -0.1517942
## Feature 3 -0.2824942 -0.1428571  1.0000000 -0.3946649
## Feature 4  0.1334075 -0.1517942 -0.3946649  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.