

CCREPE: Compositionality Corrected by Permutation and Renormalization

Emma Schwager, George Weingart, Craig Bielski, Curtis Huttenhower

October 13, 2014

Contents

1	Introduction	1
2	ccrepe	2
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	5
2.7	Example 3	7
2.8	Example 4	9
2.9	Example 5	10
3	nc.score	11
3.1	General Functionality	11
3.2	Arguments	12
3.3	Output	12
3.4	Usage	12
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](#).

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")
```

Note that this example corresponds to the default behavior.

`min.subj` Minimum number (count) of samples that must be nonzero 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` Subset of columns from `x` to work on. The default (NULL) uses all columns. Note that all features are used for normalization, but calculations are performed only with the requested subset.

`subset.cols.y` Subset of columns from `y` to work on. The default (NULL) uses all columns. Note that all features are used for normalization, but calculations are performed only with the requested subset.

`errthresh1` Maximum allowable probability of getting all zeros in a given bootstrapped column for the first dataset `x`. If a feature has a number of zeros that makes the probability of obtaining all zeros when sampling with replacement greater than this value, that feature will be excluded from the subsequent analysis. This is to ensure that the standard deviation of the bootstrap sample is non-zero. (default: 0.0001).

`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` or `subset.cols.y`. If TRUE but `subset.cols.y=NA`, returns all comparisons involving any features in `subset.cols.x`.

`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,
  errthresh1 = 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(10,0,1)
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")

```

```

test.output

## $p.values
##           Feature 1  Feature 2  Feature 3  Feature 4
## Feature 1           NA  0.05278275  0.4741754  0.1816230
## Feature 2  0.05278275           NA  0.3051871  0.4187276
## Feature 3  0.47417542  0.30518707           NA  0.3498009
## Feature 4  0.18162296  0.41872758  0.3498009           NA
##
## $z.stat
##           Feature 1  Feature 2  Feature 3  Feature 4
## Feature 1           NA  1.9366941 -0.7157019 -1.3357746
## Feature 2  1.9366941           NA -1.0253733 -0.8086307
## Feature 3 -0.7157019 -1.0253733           NA  0.9349755
## Feature 4 -1.3357746 -0.8086307  0.9349755           NA
##
## $q.values
##           Feature 1  Feature 2  Feature 3  Feature 4
## Feature 1           NA  0.7502461  1.123310  1.290781
## Feature 2  0.7502461           NA  1.445961  1.190346
## Feature 3  1.1233098  1.4459611           NA  1.243005
## Feature 4  1.2907808  1.1903463  1.243005           NA

```




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`.

```
##
## $sim.score
##      Feature 1  Feature 2  Feature 3  Feature 4
## Feature 1      NA  0.3575758 -0.7575758 -0.6848485
## Feature 2  0.3575758      NA -0.2363636 -0.1272727
## Feature 3 -0.7575758 -0.2363636      NA  0.2363636
## Feature 4 -0.6848485 -0.1272727  0.2363636      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(10,0,1)
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(10,0,1)
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)

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 0.11171782 0.20865723 0.9297895 0.92754410 0.3674114 0.7942380
## Feature 2 0.02011246 0.24007098 0.5698777 0.18512848 0.2212981 0.4208305
## Feature 3 0.11784868 0.01695918 0.6817767 0.53717583 0.0960038 0.1343974
## Feature 4 0.26705222 0.17639748 0.6100921 0.02767727 0.6824796 0.7128314
##           Feature 7
## Feature 1 0.9730405
## Feature 2 0.9374498
## Feature 3 0.6839566
## Feature 4 0.4213961

```

```
##
## $z.stat
##           Feature 1 Feature 2   Feature 3   Feature 4   Feature 5   Feature 6
## Feature 1 -1.590519  1.257266  0.08810966  0.09093518 -0.9013328  0.2608114
## Feature 2 -2.324243  1.174809 -0.56823168  1.32512867  1.2230837  0.8049813
## Feature 3  1.563868 -2.387592 -0.41003988  0.61708930  1.6645438 -1.4969839
## Feature 4  1.109877 -1.351931  0.50994201 -2.20183076 -0.4090818  0.3680562
##           Feature 7
## Feature 1  0.03379519
## Feature 2 -0.07847557
## Feature 3  0.40706993
## Feature 4  0.80400150
##
## $q.values
##           Feature 1 Feature 2 Feature 3 Feature 4 Feature 5 Feature 6
## Feature 1  2.445811  2.284041  3.914549  4.061299  2.872731  3.622512
## Feature 2  1.100793  2.189923  3.465609  2.251651  2.202193  3.071046
## Feature 3  2.150027  1.856416  3.731492  3.458899  2.627234  2.101664
## Feature 4  2.248657  2.413641  3.514894  1.009886  3.557466  3.392574
##           Feature 7
## Feature 1  3.804024
## Feature 2  3.800621
## Feature 3  3.403112
## Feature 4  2.882975
##
## $sim.score
##           Feature 1   Feature 2   Feature 3   Feature 4   Feature 5
## Feature 1 -0.4060606  0.3939394  0.04242424  0.260606061 -0.2969697
## Feature 2 -0.6969697  0.5515152 -0.16363636  0.418181818  0.2727273
## Feature 3  0.3696970 -0.5878788 -0.03030303  0.006060606  0.4787879
## Feature 4  0.2848485 -0.4909091  0.12727273 -0.624242424 -0.1030303
##           Feature 6   Feature 7
## Feature 1  0.2121212 -0.2242424
## Feature 2  0.1515152  0.2121212
## Feature 3 -0.6000000  0.3090909
## Feature 4  0.1393939  0.3696970
```

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(10,0,1)
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 0.002156567 0.7339854 0.4438846
## Feature 2 0.002156567           NA 0.2084638 0.1200695
## Feature 3 0.733985439 0.208463846           NA 0.8776635
## Feature 4 0.443884584 0.120069549 0.8776635           NA
##
## $z.stat
##           Feature 1 Feature 2  Feature 3  Feature 4
## Feature 1           NA  3.067778 -0.3398288  0.7656500
## Feature 2  3.0677780           NA -1.2578007 -1.5544817
## Feature 3 -0.3398288 -1.257801           NA  0.1539318
## Feature 4  0.7656500 -1.554482  0.1539318           NA
##
## $q.values
##           Feature 1  Feature 2 Feature 3 Feature 4
## Feature 1           NA 0.03065313 2.0865519 1.5773273
## Feature 2 0.03065313           NA 0.9876913 0.8533253
## Feature 3 2.08655191 0.98769133           NA 2.0791630
## Feature 4 1.57732731 0.85332532 2.0791630           NA
```

```
##
## $sim.score
##           Feature 1 Feature 2 Feature 3 Feature 4
## Feature 1         NA  0.7792208 -0.2857143 -0.8354978
## Feature 2  0.7792208         NA -0.3116883 -0.7489177
## Feature 3 -0.2857143 -0.3116883         NA -0.0952381
## Feature 4 -0.8354978 -0.7489177 -0.0952381         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(10,0,1)
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")

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
##
## $q.values
```




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.

```
##           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
```

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

##           Feature 1 Feature 2 Feature 3 Feature 4
## Feature 1      NA      NA 0.1515152      NA
## Feature 2      NA      NA      NA      NA
## Feature 3 0.1515152      NA      NA      NA
## Feature 4      NA      NA      NA      NA

test.output.1$sim.score

##           Feature 1 Feature 2 Feature 3 Feature 4
## Feature 1      NA -0.3333333 0.1515152 0.01818182
## Feature 2 -0.3333333      NA      NA      NA
## Feature 3 0.1515151      NA      NA      NA
## Feature 4 0.01818182      NA      NA      NA

test.output.12.3$sim.score

##           Feature 1 Feature 2 Feature 3 Feature 4
## Feature 1      NA      NA 0.1515152      NA
## Feature 2      NA      NA -0.5757576      NA
## Feature 3 0.1515152 -0.5757576      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. The generalization of the checkerboard score is through defining general co-variation and co-exclusion patterns in ordinal data. With ordinal data, the checkerboard considers the 2×2 submatrices which are of the form:

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix}$$

In the cases that $a < b$ and $c < d$ (or $a > b$ and $c > d$), the submatrix is considered a co-variation pattern; conversely, if $a > b$ and $c < d$ (or $a < b$ and $c > d$) then the submatrix is considered a co-exclusion pattern. The number of co-exclusion patterns for n bins is adjusted based on the expected ratio of co-variation to co-exclusion patterns in any given sample. The adjustment factor is given by:

$$1.5 * \frac{n(n-1)}{n^2 - n + 1},$$

with the complete derivation being left to [Bielski \[2013\]](#).

The function as implemented here first performs basic quality control filtering of input relative abundance data. It then transforms relative abundances to ordinal values based on user-provided or default bin thresholds. It then computes a raw NC-score:

$$C_{ij} - 1.5 * \frac{n(n-1)}{n^2 - n + 1} * D_{ij}$$

where C_{ij} and D_{ij} are the total number of co-variation and co-exclusion patterns, respectively, between species i and j and where n is the number of bins.

The raw NC-score is then normalized to between -1 and 1 for n bins and s samples by dividing by the maximum possible NC-score, which is based on the maximum number of co-variation patterns possible. The normalized NC-score is then analogous to Pearson's ρ , where positive values indicate more co-variation than co-exclusion patterns and the magnitude of the score indicates the strength of the association between the species. The derivation of this normalization factor can be found in Bielski [2013]. The normalization factor is given by:

$$\binom{s}{2} - \left[s \bmod n * \binom{\lfloor \frac{s}{n} \rfloor + 1}{2} + (n - s \bmod n) * \binom{\lfloor \frac{s}{n} \rfloor}{2} \right].$$

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.

bins Either the number of bins to use or a *vector* specifying bin edges. If a single number is given, this is used as the number of bins with the `discretize` function of the package *infotheo*. If a *vector* is specified, the function `findInterval` is used to discretize the data. The default behavior is to use the defaults for the `discretize` function.

verbose Request verbose output.

min.abundance Minimum abundance threshold for quality control filtering. For a feature to be included, it must take a value of at least `min.abundance` in at least `min.samples` percent of samples.

min.samples Minimum sample threshold for quality control filtering. For a feature to be included, it must take a value of at least `min.abundance` in at least `min.samples` percent of samples.

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.

3.4 Usage

```
nc.score(
  x = NA,
  y = NA,
  bins = NA,
  verbose = FALSE,
  min.abundance = 1e-04,
  min.samples = 0.1)
```

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(10,0,1)
data.norm <- data/data.rowsum
apply(data.norm,1,sum) # The rows sum to 1, so the data are normalized

## [1] 0.3967281 0.7148575 1.2671215 1.1796668 1.1274304 0.9069392 1.5187038
## [8] 0.7483735 0.7064038 1.8589967

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         NA  0.7792208 -0.5324675 -0.5887446
## Feature 2  0.7792208         NA -0.5887446 -0.3116883
## Feature 3 -0.5324675 -0.5887446         NA -0.0952381
## Feature 4 -0.5887446 -0.3116883 -0.0952381         NA
```

```
test.output.num
## [1] 0.7792208
```

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(10,0,1)
data.norm <- data/data.rowsum
apply(data.norm,1,sum) # The rows sum to 1, so the data are normalized

## [1] 0.5298094 1.1699339 0.3420548 1.4865535 1.9076789 1.2918124 0.7006396
## [8] 2.0652879 2.0685193 0.5751604

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,bins=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 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      NA      0.6     -0.6      0.2
## Feature 2      0.6      NA     -0.6      0.2
## Feature 3     -0.6     -0.6      NA     -0.6
## Feature 4      0.2      0.2     -0.6      NA
```

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(10,0,1)
data.norm <- data/data.rowsum
apply(data.norm,1,sum) # The rows sum to 1, so the data are normalized

## [1] 0.9553070 2.5610701 2.2074180 1.0861213 1.0248568 0.4176919
## [7] 1.1582021 -0.2089967 2.2006628 0.3185923

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,bins = c(0.001,0.1,0.25,0.6))

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           NA  0.7750000 -0.08571429 -0.1785714
## Feature 2  0.77500000           NA -0.46361746 -0.2598753
## Feature 3 -0.08571429 -0.4636175           NA -0.1954262
## Feature 4 -0.17857143 -0.2598753 -0.19542620           NA
```

4 References

References

- Craig Bielski. Extending the Checkerboard Score to Ordinal Data: A Methodology for Detecting Species-Level Co-Variation and Co-Exclusion Patterns in the Human Microbiome. Master's thesis, Harvard School of Public Health, Boston, USA, 2013.
- 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.
- Emma Schwager and Colleagues. Detecting statistically significant associations between sparse and high dimensional compositional data. In Progress.