

Cluster Analysis: Tutorial with R

Jari Oksanen

February 2, 2012

Contents

1	Introduction	1
2	Hierarchical Clustering	1
2.1	Description of Classes	4
2.2	Numbers of Classes	5
2.3	Clustering and Ordination	6
2.4	Reordering a Dendrogram	7
2.5	Minimum Spanning Tree	9
2.6	Cophenetic Distance	9
3	Interpretation of Classes	10
3.1	Environmental Interpretation	10
3.2	Community Summaries	12
4	Optimized Clustering at a Given Level	15
4.1	Optimum Number of Classes	16
5	Fuzzy Clustering	17

1 Introduction

In this tutorial we inspect classification. Classification and ordination are alternative strategies of simplifying data. Ordination tries to simplify data into a map showing similarities among points. Classification simplifies data by putting similar points into same class. The task of describing a high number of points is simplified to an easier task of describing a low number of classes.

I did not have time to lecture on the classification methods. My slides are available on the Internet, and if you wish, I can talk about on classification. If you prefer, you can just go through this tutorial (and preferably my lecture slides), and you will probably get out of the class sooner. The choice is yours.

2 Hierarchic Clustering

The classification methods are available in standard R packages. The **vegan** package does not have many support functions for classification, but we still




Figure 1: Distance between two clusters A and B defined by single, complete and average linkage. Mark each of the linkage types in the connecting line. The fusion level in the cluster dendrogram would be the length of the corresponding connecting line of the linkage type.

load `vegan` to have access to its data sets and some of its support functions.¹

```
R> library(vegan)
R> data(dune)
```

Hierarchical clustering (function `hclust`) is in standard R and available without loading any specific libraries. Hierarchical clustering needs dissimilarities as its input. Standard R has function `dist` to calculate many dissimilarity functions, but for community data we may prefer **vegan** function `vegdist` with ecologically useful dissimilarity indices. The default index in `vegdist` is Bray–Curtis:

```
R> d <- vegdist(dune)
```

Ecologically useful indices in **vegan** have an upper limit of 1 for absolutely different sites (no shared species), and they are based on differences of abundances. In contrast, the standard Euclidean distance has no upper limit, but varies with the sum of total abundances of compared sites when there are no shared species, and uses squares of differences of abundances. There are many other ecologically useful indices in `vegdist`, but Bray–Curtis is usually not a bad choice.

There are several alternative clustering methods in the standard function `hclust`. We shall inspect three basic methods: single linkage, complete linkage and average linkage. All these start in the same way: they fuse two most similar points to a cluster. They differ in the way they combine clusters to each other, or new points to existing clusters (Fig. 1). In single linkage (a.k.a. nearest neighbour, or neighbour joining tree in genetics) the distance between two

¹If you do not have a package, but get an error message, you must install package using `install.packages("vegan")` or the installation menu.

clusters is the shortest possible distance among members of the clusters, or the best of the friends. In complete linkage (a.k.a. furthest neighbour) the distance between two clusters is the longest possible distance between the groups, or the worst among the friends. In average linkage, the distance between the clusters is the distance between cluster centroids. There are several alternative ways of defining the average and defining the closeness, and hence a huge number of average linkage methods. We only use one of these methods commonly known as UPGMA. The lecture slides discuss the methods in more detail.

In the following we will compare three different clustering strategies. If you want to plot three graphs side by side, you can divide the screen into three panels by

```
R> par(mfrow = c(1, 3))
```

This defines three panels side by side. You probably want to stretch the window if you are using this. Alternatively, you can have three panels above each other with

```
R> par(mfrow = c(3, 1))
```

You can get back to the single panel mode with

```
R> par(mfrow = c(1, 1))
```

You may also wish to use narrower empty margins for the panels:

```
R> par(mar = c(3, 4, 1, 1) + 0.1)
```

The single linkage clustering can be found with:

```
R> csin <- hclust(d, method = "single")
R> csin
```

```
Call:
hclust(d = d, method = "single")
```

```
Cluster method : single
Distance       : bray
Number of objects: 20
```

The dendrogram can be plotted with:

```
R> plot(csin)
```

The default is to plot an inverted tree with the root at the top, and branches hanging down. You can force the branches down to the base line giving the `hang` argument:

```
R> plot(csin, hang = -1)
```

The complete linkage and average linkage methods are found in the same way:

```
R> ccom <- hclust(d, method = "complete")
R> plot(ccom, hang = -1)
R> caver <- hclust(d, method = "aver")
R> plot(caver, hang = -1)
```



Figure 2: Vegemite is an Australian national delicacy made of yeast extract. The vegemite function was named because its output is just as dense as Vegemite.

The vertical axes of the cluster dendrogram shows the fusion level. The two most similar observations are combined first, and they are at the same level in all dendrograms. At the upper fusion levels, the scales diverge: they are the shortest dissimilarities among cluster members in single linkage, the longest possible dissimilarities in complete linkage, and the distances among cluster centroids in average linkage (Fig. 1).

2.1 Description of Classes

One problem with hierarchic clustering is that it gives a classification of observations (plots, sampling units), but it does not tell how these classes differ from each other. For community data, there is no information how the species composition differs between classes (we return to this subject in Chapter 3.2).

The **vegan** package has function `vegemite` (Fig. 2) that can produce compact community tables ordered by a dendrogram, ordination or environmental variables. With the help of these tables it is possible to see which species differ in classification:

```
R> vegemite(dune, caver)

      111211 11    11    1
      46507913243891856720
Potpal 2.2.....
Calcus 43.3.....
Elepal 4854.....4.....
Ranfla 2224...2...2.....
```

```

Airpra ....23.....
Empnig .....2.....
Agrsto 4745...548443.....
Junart .334.....44.....
Salrep ...5.3.....3.....
Hyprad ....25.....2.....
Chealb .....1.....
Alogen .4.....582753.....2.
Sagpro .....3.245.222.....
Cirarv .....2.....
Brarut .444.3..4222246262.2
Junbuf .....34...4....2..
Leoout 2.2226.2222325533353
Elyrep .....4..44.6..4..4.
Trirep 6.1..2.2312233225256
Poatri .2....2945645..64574
Antodo ....44.....432.4
Poapra ....1.42.45444323444
Lolper .....7..56427226656
Belper .....22...22..32
Rumace .....2...2..563..
Viclat .....21....1
Plalan ....2.....33555.3
Achmil ....2.1.....22234
Brohor .....3.....2.244
Tripra .....252..
  sites species
    20      30

```

The `vegemite` command will always use one-character columns. If the observed values do not fit one character, the `vegemite` refuses to work. With argument `scale` you can recode the values to one-character width.

2.2 Numbers of Classes

The hierarchic clustering methods produce all possible levels of classifications. The extremes are all observations in a single class, and each observation in its private class. The user normally wants to have a clustering into a certain number of classes. The fixed classification can be visually demonstrated with `rect.hclust` function:

```

R> plot(csin, hang = -1)
R> rect.hclust(csin, 3)
R> plot(ccom, hang = -1)
R> rect.hclust(ccom, 3)
R> plot(caver, hang = -1)
R> rect.hclust(caver, 3)

```

Single linkage has a tendency to chain observations: most common case is to fuse a single observation to an existing class: the single link is the nearest neighbour, and in a close neighbour is more probably in a large group than a

solitary point. Complete linkage has a tendency to produce compact bunches: complete link minimizes the spread within the cluster. The average linkage is between these two extremes.

We can extract classification in a certain level using function `cutree`:

```
R> cl <- cutree(ccom, 3)
R> cl

 2 13  4 16  6  1  8  5 17 15 10 11  9 18  3 20 14 19 12  7
 1  2  2  2  1  1  2  1  3  2  1  3  2  3  2  2  2  3  2  1
```

This gives a numeric classification vector of cluster identities. The clusters are numbered in the order the observations appear in the data: the first item will always belong to cluster 1, and the numbering does not match the dendrogram. We can tabulate the numbers of observations in each cluster:

```
R> table(cl)

cl
 1  2  3
 6 10  4
```

We can compare two clustering schemes by cross-tabulation which gives as a confusion matrix:

```
R> table(cl, cutree(csin, 3))

cl  1  2  3
 1  6  0  0
 2 10  0  0
 3  2  1  1
```

```
R> table(cl, cutree(caver, 3))

cl  1  2  3
 1  6  0  0
 2  6  4  0
 3  2  0  2
```

The confusion matrix tabulates the classifications against each other. The rows give the first classification, and the columns the second classification. If the classifications match and there is no “confusion”, each row and each column has only one non-zero entry, but if the classes are divided between several classes in the second classification, the row has several non-zero entries.

2.3 Clustering and Ordination

We can use ordination to display the observed dissimilarities among points. A natural choice is to use metric scaling a.k.a. principal coordinates analysis (PCoA) that maps observed dissimilarities linearly onto low-dimensional graph. The metric scaling can be performed with standard R function `cmdscale`:

```
R> ord <- cmdscale(d)
```

We can display the results using **vegan** function `ordiplot` that can plot results of any **vegan** ordination function and many non-**vegan** ordination functions, such as `cmdscale`, `prcomp` and `princomp` (the latter for principal components analysis):

```
R> ordiplot(ord)
```

We got a warning because `ordiplot` tries to plot both species and sites in the same graph, and the `cmdscale` result has no species scores. We do not need to care about this warning.

There are many **vegan** functions to overlay classification onto ordination. For distinct, non-overlapping classes convex hulls are practical:

```
R> ordihull(ord, c1)
```

For overlapping classes we can use `ordispider`. If we are not interested in individual points, but on class centroids and average dispersions, we can use `ordiellipse`, both with the same arguments as `ordihull`.

The other clustering results can be seen with:

```
R> ordiplot(ord, dis = "si")
R> ordihull(ord, cutree(caver, 3))
R> ordiplot(ord, dis = "si")
R> ordicluster(ord, csin)
```

We set here explicitly the `display` argument to `display = "sites"` to avoid the annoying and useless warnings. The contrasting clustering strategies (nearest vs. furthest vs. average neighbour) are evident in the shapes of the clusters. Single linkage clusters are chained, complete linkage clusters are compact, and average linkage clusters between these two.

The **vegan** package has a special function to display the cluster fusions in ordination. The `ordicluster` function combines sites and cluster centroids similarly as the average linkage method:

```
R> ordiplot(ord, dis = "si")
R> ordicluster(ord, caver)
```

We can prune the top level fusions to highlight the clustering:

```
R> ordiplot(ord, dis = "si")
R> ordicluster(ord, caver, prune = 2)
```

2.4 Reordering a Dendrogram

The leaves of a dendrogram do not have a natural order. You can take a branch a turn it around its root, and the tree is the same (see Fig. 3).

R has two alternative dendrogram presentations: the `hclust` result object and a general `dendrogram` object. The cluster type can be changed with:

```
R> den <- as.dendrogram(caver)
```

The `dendrograms` are more general, and several methods are available for their manipulation and analysis. It is possible to re-order the leaves of a dendrogram so that they match as closely as possible an external variable.

In the following we rearrange the dendrogram so that the ordering of leaves corresponds as closely as possible with the first ordination axis:




Figure 3: A dendrogram is similar to a mobile: branches can turn around, but the mobile is the same. Alexander Calder: Red Mobile.

```
R> x <- scores(ord, display = "sites", choices = 1)
R> oden <- reorder(den, x)
```

We plot the dendrograms together:

```
R> par(mfrow = c(2, 1))
R> plot(den)
R> plot(oden)
R> par(mfrow = c(1, 1))
```

The reordered dendrogram may also give more regularly structured community table:

```
R> vegemite(dune, oden)

          1 11    11 111112
          76502183498321794650
Tripra 252.....
Rumace 365.....2..2.....
Plalan 5553.33.....2.....
Brohor 2.244...3.....
Achmil 22243.....12.....
Viclat ...1.21.....
Belper ..223.222.....
Lolper 66265726524..7.....
Poapra 432444354442.41.....
Antodo 2344.....44....
Poatri 54647..6554942...2..
Elyrep ..4.4..446...4.....
Trirep 2526532213223..26.1.
Cirarv .....2.....
Leoaut 3333555222322.262.22
Junbuf 2.....4.34.....
Brarut 2622.462222.4..3.444
Sagpro .....2..52224..3....
Alogen ...2..723558...4..
Chealb .....1.....
Hyprad .....2.....25....
Agrsto .....483454...4745
Junart .....44.....334
Salrep .....3.....3...5
```

```

Airpra .....23....
Empnig .....2....
Ranfla .....22....2224
Elepal .....4....4854
Potpal .....2.2.
Calcus .....43.3
  sites species
      20      30

```

2.5 Minimum Spanning Tree

In the graph theory, tree is a connected graph without loops, spanning tree is a tree connecting all points, and minimum spanning tree is the shortest of such trees. Minimum spanning tree (MST) is closely related to single linkage clustering, which also connects all points with minimum total connecting distance. However, MST really combines points, whereas R representations of single linkage clustering hierarchically connects clusters instead of single points.

MST can be found with **vegan** function `spantree`²

```
R> mst <- spantree(d)
```

MST can be overlaid onto an ordination with `lines` command:

```
R> ordiplot(ord, dis = "si")
R> lines(mst, ord)
```

Alternatively, the tree can be displayed with a `plot` command that tries to find a locally optimal configuration for points to reproduce the distances among points along the tree. Internally the function uses Sammon scaling (`sammon` function in the **MASS** package) to find the configuration of points. Sammon scaling is a variant of metric scaling trying to reproduce relative distances among points and it is optimal for showing the local (vs. global) structure of points.

```
R> plot(mst, type = "t")
```

```

Initial stress      : 0.03111
stress after 10 iters: 0.01302, magic = 0.500
stress after 20 iters: 0.01139, magic = 0.500
stress after 30 iters: 0.01118, magic = 0.500
stress after 40 iters: 0.01114, magic = 0.500

```

2.6 Cophenetic Distance

The estimated distance between two points is the level at which they are fused in the dendrogram, or the height of the root. A good clustering method correctly reproduces the actual dissimilarities. The distance estimated from a dendrogram is called cophenetic distance. The name echoes the origins of hierarchic clustering in old fashioned numeric taxonomy. Standard R function `cophenetic` estimates the distances among all points from a dendrogram.

²There are many other implementations MST in R, but the implementation in **vegan** probably is the fastest.

We can visually inspect the cophenetic distances against observed dissimilarities. In the following, `abline` adds a line with zero intercept and slope of one, or the equivalence line. We also set equal scaling on x and y axes (`asp = 1`):

```
R> plot(d, cophenetic(csin), asp = 1)
R> abline(0, 1)
R> plot(d, cophenetic(ccom), asp = 1)
R> abline(0, 1)
R> plot(d, cophenetic(caver), asp = 1)
R> abline(0, 1)
```

In single linkage clustering, the cophenetic distances are as long as or shorter than the observed distances: the distance between groups is the shortest possible distance between its members. In complete linkage clustering, the cophenetic distances are as long as or longer than observed distances: the distance between two groups is the longest possible distance between groups. In average linkage clustering, the cophenetic distance is the average of observed distances (cf. Fig. 1).

The correlation between observed and cophenetic distances is called the cophenetic correlation:

```
R> cor(d, cophenetic(csin))

[1] 0.6601692

R> cor(d, cophenetic(ccom))

[1] 0.6707479

R> cor(d, cophenetic(caver))

[1] 0.8168825
```

The ranking of these cophenetic correlations is not entirely random: it is guaranteed that average linkage (UPGMA) method maximizes the cophenetic correlation.

3 Interpretation of Classes

We commonly want to use classes for prediction of external variables — this is the idea om Finnish forest type system, EU Water Framework Directive and theory of gradient analysis. We make classes from community composition and the classes should describe community composition, and we assume the classes predict environmental conditions.

3.1 Environmental Interpretation

For interpreting the classes we take into use the the environmental data:

```
R> data(dune.env)
```

The `cutree` function produced a numerical classification vector, but we need a factor of classes:

```
R> c1 <- factor(c1)
```

We can visually display the differences in environmental conditions using `boxplot`. The only continuous variable we have is the thickness of the A1 horizon, but we can change the Moisture class to a numeric variable:

```
R> Moist <- with(dune.env, as.numeric(as.character(Moisture)))
```

Compare the result of this to a simpler looking alternative:

```
R> with(dune.env, as.numeric(Moisture))
```

```
[1] 1 4 2 4 1 1 4 1 2 4 2 1 3 1 2 4 4 4 3 1
```

The latter uses the internal representation (1,2,3,4) of the factor instead of its nominal values (1,2,4,5) that we can get with the added `as.character` command.

The boxplot is produced with:

```
R> boxplot(Moist ~ c1, notch = TRUE)
```

The boxplot shows the data extremes (whiskers and the points for extreme cases), lower and upper quartiles and the median. The notches show the approximate 95% confidence limits of the median: if the notches do not overlap, the medians probably are different at level $P < 0.05$. For small data sets like this, the confidence limits are wide, and you get a warning of waist going over the shoulders.

You can use ANOVA for rigid numerical testing of differences among classes. In R, ANOVA is an analysis method for linear models (`lm`):

```
R> anova(lm(Moist ~ c1))
```

Analysis of Variance Table

Response: Moist

	Df	Sum Sq	Mean Sq	F value	Pr(>F)
c1	2	36.617	18.3083	12.359	0.0004854 ***
Residuals	17	25.183	1.4814		

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

This performs parametric ANOVA test. If we want to use non-parametric permutation tests with the same test statistic, we can use redundancy analysis:

```
R> anova(rda(Moist ~ c1))
```

Permutation test for rda under reduced model

Model: rda(formula = Moist ~ c1)

	Df	Var	F	N.Perm	Pr(>F)
Model	2	1.9272	12.359	199	0.005 **
Residual	17	1.3254			

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

This works nicely in this case because the *y*-variate (**Moist**) is in the current workspace: the command would not work for variates in data frames unless we attach the data frame or use the command as `with(varechem, . . .)`.

We can use confusion matrix to compare factor variables and classification:

```
R> with(dune.env, table(c1, Management))
```

```
      Management
c1   BF HF NM SF
  1   2  3  0  1
  2   0  2  3  5
  3   1  0  3  0
```

3.2 Community Summaries

We can use the **labdsv** package for analysing community classification.

```
R> library(labdsv)
```

The package has several functions for summarizing community composition in classes. The species occurrence frequencies in classes can be found with

```
R> const(dune, c1)
```

```
      1  2  3
Belper 0.50 0.2 0.25
Empnig 0.00 0.0 0.25
Junbuf 0.16 0.3 0.00
Junart 0.00 0.5 0.00
Airpra 0.00 0.0 0.50
Elepal 0.00 0.5 0.00
Rumace 0.50 0.2 0.00
Viclat 0.16 0.0 0.50
Brarut 0.66 0.8 0.75
Ranfla 0.00 0.6 0.00
Cirarv 0.00 0.1 0.00
Hyprad 0.00 0.0 0.75
Leoaut 0.83 0.9 1.00
Potpal 0.00 0.2 0.00
Poapra 1.00 0.5 0.75
Calcus 0.00 0.3 0.00
Tripra 0.50 0.0 0.00
Trirep 0.83 0.8 0.75
Antodo 0.66 0.0 0.50
Salrep 0.00 0.1 0.50
Achmil 1.00 0.0 0.25
Poatri 1.00 0.7 0.00
Chealb 0.00 0.1 0.00
Elyrep 0.50 0.3 0.00
Sagpro 0.00 0.5 0.50
Plalan 0.66 0.0 0.75
Agrsto 0.00 1.0 0.00
```

```
Lolper 1.00 0.4 0.50
Alogen 0.16 0.7 0.00
Brohor 0.66 0.1 0.00
```

and the mean abundances with

```
R> importance(dune, c1)
```

```

      1  2  3
Belper 1.16 0.4 0.50
Empnig 0.00 0.0 0.50
Junbuf 0.33 1.1 0.00
Junart 0.00 1.8 0.00
Airpra 0.00 0.0 1.25
Elepal 0.00 2.5 0.00
Rumace 2.33 0.4 0.00
Viclat 0.16 0.0 0.75
Brarut 2.00 2.4 3.25
Ranfla 0.00 1.4 0.00
Cirarv 0.00 0.2 0.00
Hyprad 0.00 0.0 2.25
Leoaut 2.83 1.9 4.50
Potpal 0.00 0.4 0.00
Poapra 3.50 1.9 2.00
Calcus 0.00 1.0 0.00
Tripra 1.50 0.0 0.00
Tlirep 3.33 2.0 1.75
Antodo 2.16 0.0 2.00
Salrep 0.00 0.5 1.50
Achmil 2.33 0.0 0.50
Poatri 4.66 3.5 0.00
Chealb 0.00 0.1 0.00
Elyrep 2.00 1.4 0.00
Sagpro 0.00 1.5 1.25
Plalan 3.00 0.0 2.00
Agrsto 0.00 4.8 0.00
Lolper 5.33 1.7 2.25
Alogen 0.33 3.4 0.00
Brohor 2.00 0.3 0.00
```

The **labdsv** has function `indval` for indicator species analysis.³ The indicator species analysis combines frequency tables (`const`) and mean abundance tables (`importance`) and finds species that are significantly concentrated into specific classes. The model is fitted with:

```
R> mod <- indval(dune, c1)
```

The `indval` result object contains the following items:

```
R> names(mod)
```

³In old **labdsv** versions the function was called `duleg`.

```
[1] "relfrq" "relabu" "indval" "maxcls" "indcls" "pval"
```

The item `relfrq` is the same as the output of `const`, and `relabu` is based on the output of `importance`, but each row is normalized to unit sum so that the table show the distribution of total abundance in classes. These values are combined into indicator value (`indval`) of each species for each class, and `maxcls` gives the class to which species has the highest indicator value (`indcls`). Finally, vector `pval` gives the probability of finding higher indicator value in random permutations, and if this probability is low, the species is a significant indicator.

Each item can be inspected separately:

```
R> mod$maxcls
```

	Belper	Empnig	Junbuf	Junart	Airpra	Elepal	Rumace	Viclat	Brarut	Ranfla
Belper	1	3	2	2	3	2	1	3	3	2
Cirarv	2	3	3	2	1	2	1	1	1	3
Achmil	1	1	2	1	2	1	2	1	2	1

```
R> mod$pval
```

	Belper	Empnig	Junbuf	Junart	Airpra	Elepal	Rumace	Viclat	Brarut	Ranfla
Belper	0.334	0.186	0.320	0.051	0.041	0.045	0.072	0.074	0.762	0.013
Cirarv	1.000	0.007	0.031	0.356	0.130	0.209	0.025	0.454	0.174	0.089
Achmil	0.002	0.019	1.000	0.237	0.398	0.176	0.001	0.023	0.011	0.025

Often it is more practical to look at the summary. The "short" summary lists the significant indicator species with some extra information:

```
R> summary(mod)
```

	cluster	indicator_value	probability
Achmil	1	0.8235	0.002
Brohor	1	0.5797	0.025
Lolper	1	0.5745	0.023
Poatri	1	0.5714	0.019
Tripa	1	0.5000	0.025
Agrsto	2	1.0000	0.001
Alogen	2	0.6375	0.011
Ranfla	2	0.6000	0.013
Elepal	2	0.5000	0.045
Hyprad	3	0.7500	0.007
Airpra	3	0.5000	0.041
Leoaut	3	0.4874	0.031

```
Sum of probabilities = 6.265
```

```
Sum of Indicator Values = 13.19
```

Sum of Significant Indicator Values = 7.52

Number of Significant Indicators = 12

Significant Indicator Distribution

1 2 3
5 4 3

The "long" summary lists all species with their indicator values for classes, but replaces non-significant values with a dot:

```
R> summary(mod, type = "long")
```

	1	2	3
Belper	0.28	.	0.06
Empnig	.	.	0.25
Junbuf	.	0.23	.
Junart	.	0.50	.
Airpra	.	.	0.50
Elepal	.	0.50	.
Rumace	0.43	.	.
Viclat	.	.	0.41
Brarut	0.17	0.25	0.32
Ranfla	.	0.60	.
Cirarv	.	0.10	.
Hyprad	.	.	0.75
Leoaut	0.26	0.19	0.49
Potpal	.	0.20	.
Poapra	0.47	0.13	0.20
Calcus	.	0.30	.
Tripra	0.50	.	.
Trirep	0.39	0.23	0.19
Antodo	0.35	.	0.24
Salrep	.	.	0.38
Achmil	0.82	.	.
Poatri	0.57	0.30	.
Chealb	.	0.10	.
Elyrep	0.29	0.12	.
Sagpro	.	0.27	0.23
Plalan	0.40	.	0.30
Agrsto	.	1.00	.
Lolper	0.57	0.07	0.12
Alogen	.	0.64	.
Brohor	0.58	.	.

4 Optimized Clustering at a Given Level

Hierarchic clustering has a history: it starts from the bottom, and combines observations to clusters. At higher levels it can only fuse existing clusters to

others. Often we are most interested in classification at a low number of clusters, and these have a long historic burden. Once we have decided upon the number of classes, we can often improve the classification at a given number of clusters. A tool for this task is K -means clustering that produces an optimal solution for a given number of classes.

K -means clustering has one huge limitation for community data: it works in Euclidean metric which rarely is meaningful for community data. In hierarchical clustering we could use non-Euclidean Bray-Curtis dissimilarity, but we do not have this choice here. However, we can work around some problems with appropriate standardization. Hellinger transformation has been suggested to be very useful for community data: square root of data divided by row (site) totals. The Hellinger transformation can be performed with **vegan** function `decostand` which has many other useful standardizations and transformations:

```
R> ckm <- kmeans(decostand(dune, "hell"), 3)
R> ckm$cluster

 2 13  4 16  6  1  8  5 17 15 10 11  9 18  3 20 14 19 12  7
 3  1  1  2  3  3  1  3  3  2  3  3  1  3  1  2  2  3  1  3
```

We can display the results in the usual way:

```
R> ordiplot(ord, dis = "si")
R> ordihull(ord, ckm$cluster, col = "red")
```

4.1 Optimum Number of Classes

The K -means clustering optimizes for the given number of classes, but what is the correct number of classes? Arguably there is no correct number, because the World is not made of distinct classes, but all classifications are patriarchal artifacts. However, some classifications may be less useful (and worse) than others.

Function `cascadeKM` in **vegan** finds some of these criteria for a sequence of K -means clusterings. The default method is the Calinski–Harabasz criterion which is the F ratio of ANOVA for the predictor variables (species in community data). The test uses Euclidean metric and we use again the Hellinger transformation (if you try without transformation or different transformations, you can see that transformations with `decostand` have capricious effects: obviously the World is not classy in this case).

```
R> ccas <- cascadeKM(decostand(dune, "hell"), 2, 15)
R> plot(ccas, sortq = TRUE)
```

We gave the smallest (2) and largest (15) number of inspected classes in the call, and sorted the plot by classification. Each class is shown by a different colour, and bends in the vertical colour bars show non-hierarchical clustering where the upper levels are no simple fusions of lower levels. The highest value of the criterion shows the optimal number of classes. This kind of graph is called Mondrian plot after an excellent Dutch-born artist Piet Mondriaan of *De Stijl* school (Fig. 4).



Figure 4: Piet Mondrian: Composition in red, blue and yellow.

5 Fuzzy Clustering

We have so far worked with classification methods which implicitly assume that there are distinct classes. The World is not made like that. If there are classes, they are vague and have intermediate and untypical cases. With one word, they are fuzzy.

Fuzzy classification means that each observation has a certain probability of belonging to a class. In the crisp case, it has probability 1 of belonging to one class, and probability 0 of belonging to any other class. In a fuzzy case, it has probability < 1 for the best class, and probabilities > 0 for several other classes (and the sum of these probabilities is 1).

Fuzzy classification is similar to K -means clustering in finding the optimal classification for a given number of classes, but the produced classification is fuzzy: the result is a probability profile of class membership.

The fuzzy clustering is provided by function `fanny` in package `cluster`⁴. Fuzzy clustering defaults to Euclidean metric, but current versions can accept any dissimilarities. In the following, we also need to use lower “membership exponent” (`memb.exp`): the default 2 gives complete fuzziness and fails here, and lower values give crisper classifications.⁵

```
R> library(cluster)
R> cfuz <- fanny(d, 3, memb.exp = 1.7)
```

The function returns an object with the following items:

```
R> names(cfuz)

[1] "membership" "coeff"      "memb.exp"   "clustering"
[5] "k.crisp"     "objective"  "convergence" "diss"
[9] "call"       "silinfo"
```

⁴All functions in the `cluster` package have stupid names.

⁵Old version of `fanny` may not have all these options and can fail; you should upgrade R to get a new version of the `cluster` package.

Here `membership` is the probability profile of belonging to a certain class, and `clustering` is the most likely crisp classification.

It is difficult to show the fuzzy results graphically, but here is one idea:

```
R> ordiplot(ord, dis = "si")
R> ordiplot(ord, dis = "si", type = "n")
R> stars(cfuz$membership, locatio = ord, draw.segm = TRUE,
+       add = TRUE, scale = FALSE, len = 0.1)
R> ordihull(ord, cfuz$clustering, col = "blue")
```

This uses `stars` function (with many optional parameters) to show the probability profile, and draws a convex hull of the crisp classification. The size of the sector shows the probability of the class membership, and in clear cases, one of the segments is dominant.