Wintersemester 2010/2011, AST

# Cluster Analysis Übung

Assignment due date: 22.11.2010

Exercises on 8.11.2010 and 15.11.2010

# 1 Goals

For this Übung, you should:

- familiarise yourself with some of the common functions of the statistics package "R"
- employ and understand some of the clustering methods
- and submit a very brief report (questions in section 4)

# 2 Using the "R" package:

"R is a system for statistical computation and graphics. It consists of a language plus a runtime environment with graphics, a debugger, access to certain system functions, and the ability to run programs stored in script files." – R project page

First of all we create a directory to work in. The name of the directory is just a suggestion, feel free to choose any name you like. Type in the following commands from a shell prompt:

mkdir ast-cluster cd ast-cluster

We start R by running the following command from a shell prompt:

### /usr/local/zbh/bin/R

You should now see the R command prompt. Remember that you can save yourself a lot of typing by using the up and down arrows of the keyboard to cycle through previously entered commands and the tab key to automatically complete commands (just like you can at the shell prompt).

To store a value in a variable, we use the assignment operator "<-":

> answer <- 42

We can inspect the value of the variable **answer** by simply typing its name at the command prompt:

> answer [1] 42

The first part of the output ("[1]") indicates that this is the first row. This is of course irrelevant for simple numbers or vectors but will become useful later on when we deal with matrices. Use the function ls() to get a list of all variables in the current session. When you type in the name of a function without the parentheses and arguments, the function definition will be printed. Try this out by entering ls, but don't worry if you don't understand the code.

To create a vector, we will use the "c" function (an abbreviation of combine):

> foo <- c(1, 2, 3) > foo [1] 1 2 3

The contents of vectors can be accessed by using the vector name with the index in square brackets. IMPORTANT: in contrast to many programming languages, the index of the first element of a vector is "1". Accessing index "0" of any variable gives you its type.

> foo[0]
numeric(0)
> foo[1]
[1] 1

Matrices are constructed in a similar manner. The matrix constructor accepts a data vector and optional specifications, such as the number of rows or columns. The matrix is filled columnwise, unless the parameter **byrow** is set to **TRUE**. Please beware that if there are not enough entries in the data vector to fill the matrix, the entries of the data vector will be reused instead of raising an error message.

> bar <- matrix(c(1,2,3,4,5,6), nrow=3) > bar

	[,1]	[,2]
[1,]	1	4
[2,]	2	5
[3,]	3	6

As with vectors, you can access matrix elements using square brackets, but unlike C arrays, all indices are placed as a comma separated list in the first (and only) pair of square brackets. If an index is negative, the row or column corresponding to that (positive) index will be omitted.

> bar[3,2]
[1] 6
> bar[-3,2]
[1] 4 5

Simply leaving out a parameter selects the entire column or row:

> bar[,2] [1] 4 5 6

To concatenate matrices, you can use the **rbind** and **cbind** commands, which will try to append matrices and vectors row- and columnwise, respectively.

```
> rbind(bar, c(3.5,6.5))
```

		[,1]	[,2]	
	[1,]	1	4	
	[2,]	2	5	
	[3,]	3	6	
	[4,]	3.5	6.5	
>	<pre>&gt; cbind(bar,</pre>		c(7,8,	9))
		[,1]	[,2]	[,3]
	[1,]	1	4	7
	[2,]	2	5	8
	[3,]	3	6	9

As with the matrix constructor, if a vector does not contain enough elements, the first elements will be replicated at the missing positions. This does not hold true for matrices, though. Concatenating incompatible matrices will produce an error message.

### 2.1 Storing your R commands in a script file

Often it is more convenient to be able to store the R commands you want executed in a script file and run them from there. You can create the script file for R with your favourite text editor (if you don't have a favourite text editor you can use the program kate). The following command (to be executed from the shell prompt) will run the R commands stored in the file aufgabe1.r and store the output of R in the file out-aufgabe1:

### /usr/local/zbh/bin/Rscript aufgabe1.r > out-aufgabe1.txt

Please note that simply typing in the name of a variable at the interactive prompt will print its value, but to do the same in a script you will have to actually print the variable. For example, to print the variable x you would call print(x).

### 2.2 Getting help

Every function in R is thoroughly documented. You can access the documentation to a command with:

#### > help(plot)

This can be abbreviated to:

```
> ?plot
```

You can search inside the documentation with:

> ??plot

## 3 Exercises

Open the data file and assign the data to a variable (storing it as a matrix):

```
> cluster.datapath <- file.path('/home/matthies/uebung-clustering/sampledata')
> cluster.data <- as.matrix(read.table(cluster.datapath, sep=','))</pre>
```

Check the dimensionality of your data:

```
> dim(cluster.data)
[1] 100 2
```

This tells you that your data is a matrix with 100 rows and 2 columns. You can also plot your data with the plot command (see figure 1 for the result):



Figure 1: Plot of the test data set

### > plot(cluster.data)

To store the plot to the file myplot.png use the following commands:

```
> png('myplot.png')
> plot(cluster.data)
> dev.off()
```

The final dev.off() command finalises the plot and writes out the file. The next plot will again be shown on your graphical display.

You will be using two different clustering methods, k-means and hierarchical clustering. This can be done in R with the help of the kmeans and hclust functions. If you are unsure about the correct usage of these functions or simply curious, you can read their help documentation:

```
> help(kmeans)
> help(hclust)
```

### 3.1 K-means

Do several runs with k-means clustering, and use a different number of clusters as parameter for each run (replace <number of clusters> with the number of clusters in the following command):

```
> cluster.kmeans <- kmeans(cluster.data,<number of clusters>)
```

You can visualise them with:

```
> plot(cluster.data, col = cluster.kmeans$cluster)
> points(cluster.kmeans$centers, col = 1:6, pch=7, lwd=3)
```

The contents of the data structure can be inspected by entering cluster.kmeans. As you can see, the cluster subset is a vector of integers, each denoting the cluster membership. This cluster membership is used in the plot command to select the colour (the \$ symbol selects this substructure by name).

Repeat the clustering twice, but save both results in different variables, cluster.kmeans.1 and cluster.kmeans.2. To compare these two, create a table which denotes how often each element is matched in both vectors (also read help(table)):

#### > table(cluster.kmeans.1\$cluster, cluster.kmeans.2\$cluster)

The diagonal entries denote how often a data point was assigned to the same cluster in both runs, while the off-diagonal entries denote how often the cluster index didn't match. This table is therefore called a confusion matrix. Unfortunately, since the cluster centres are randomly initialised, even if the clustering is identical, the cluster numbering may be different. If we want to calculate the rate of mismatches for a given k-means pairing, we will have to first find a correct matching of clusters and then count the number of mismatches.

Note that the "#" symbol marks a comment that extends to the end of the line. This short function is adequate for well-behaved data sets and is the one we will use in the following. Play around with the different commands used in the function until you feel comfortable using them.

## **3.2** Hierarchical clustering

The hierarchical clustering algorithm in R needs a precomputed distance matrix, which contains the pairwise distances between data points. Fortunately, a function to generate a distance matrix is already implemented in R.

```
> cluster.dist <- dist(cluster.data, method='euclidean')</pre>
```

The resulting data structure is a lower triangular matrix, where the entry at (i,j) contains the distance between data points i and j. We can now run the hierarchical clustering algorithm:

### > cluster.hclust <- hclust(cluster.dist)</pre>

The output is a list of the order in which adjacent clusters are joined, with each data point being its own cluster at initialisation. This output can be plotted as a *dendrogram* by running **plot(cluster.hclust)**. To compare this to the k-means clustering, the output has to be reduced to a number of clusters. This can be achieved by using the **cutree** method, which produces the n last clusters to be joined.

```
> cluster.hcut <- cutree(cluster.hclust,<number of clusters>)
```

The cluster.hcut vector contains the same type of information as the previously used cluster.kmeans\$cluster, and can be compared with the same methods.

# 4 Assignment

Please submit a brief report for the following questions via email: ast\_uebung[at]zbh[dot]unihamburg[dot]de (please include your full name). Attach one file containing the answers to all assignments (only .txt or .pdf), and one file containing the source code for each assignment. Your source code should be runnable with **Rscript**. This means you should attach 4 files if you answer all three assignments.

1. Write a batch function to run and compare 500 pairs of k-means clusterings and calculate

the average mismatch rate. Where and why do they differ (if they do)? Is this consistent with your expectations? Elaborate.

- 2. Compare 500 pairs of k-means and hierarchical clusterings. Is there a systematic difference between the clustering results? Explain your observations.
- 3. You may add up to 5 data points to the data set (do not use outlandishly large values, stay within the general area of the existing data set). Can you add them in a way that breaks the hclust method, i.e. can you add the points in a way that will completely change the clustering? Where did you add them and why? Do they affect the k-means clustering? Start by finding out what kind of linkage method is used by the hclust function to join clusters.