

Part X

Ordination

In laboratory exercises 9, we performed a discriminant analysis and used the discriminant axes to provide a co-ordinate system in which to represent the data. If the groups in the data correspond to categorizations that will be of interest when the data are plotted, the low-dimensional representation that is obtained is likely to be superior to that obtained from ordination without regard to any such categorization.

The ordination method considered in this laboratory address two possibilities:

- Distances may be given, from which it is desired to recover a low-dimensional representation.
 - For example we may, as below, attempt to generate a map in which the distances on the map accurately reflect Australian road travel distances.
 - Or we may, based on genomic differences, derive genomic “distances” between, e.g., different insect species. The hope is that, with a judicious choice of the distance measure, the distances will be a monotone function of the time since the two species separated. We’d like to derive a 2 or 3-dimensional representation in which the distances accurately reflect the closeness of the evolutionary relationships.
- There may be no groupings of the data that are suitable for use in deriving a low-dimensional representation.

1 Australian road distances

The distance matrix that will be used is in the matrix `audists`, in the image file `audists.Rdata`.

Here is how the data can be read in from the text file:

```
audists <- read.table("audists.txt", sep="\t")
audists[is.na(audists)] <- 0
## Also, we will later use the data frame aulatlong
aulatlong <- read.table("aulatlong.txt", row.names=1)[,2:1]
aulatlong[,2] <- -aulatlong[,2]
colnames(aulatlong) <- c("latitude", "longitude")
```

Consider first the use of classical multi-dimensional scaling, as implemented in the function `cmdscale()`:

```
> load("audists.RData")
> aupoints <- cmdscale(audists)
> plot(aupoints)
> text(aupoints, labels = paste(rownames(aupoints)))
```

An alternative to `text(aupoints, labels=paste(rownames(aupoints)))`, allowing better placement of the labels, is `identify(aupoints, labels=rownames(aupoints))`. We can compare the distances in the 2-dimensional representation with the original road distances:

```
> audistfits <- as.matrix(dist(aupoints))
> misfit <- audistfits - audists
> for (j in 1:9) for (i in (j + 1):10) {
+   lines(aupoints[c(i, j), 1], aupoints[c(i, j), 2], col = "gray")
+   midx <- mean(aupoints[c(i, j), 1])
+   midy <- mean(aupoints[c(i, j), 2])
+   text(midx, midy, paste(round(misfit[i, j])))
+ }
> print(round(misfit))
```

	Adel.	Alice	Brisb	Brm	Cairns	Canb.	Darw.	Melb.	Perth	Sydn.
Adelaide	0	140	-792	-156	366	20	11	82	482	-273
Alice	140	0	-1085	-175	-41	76	-118	106	-26	-314
Brisbane	-792	-1085	0	198	319	-25	-233	-471	153	-56
Broome	-156	-175	198	0	527	-7	6	-65	990	70
Cairns	366	-41	319	527	0	277	-31	178	8	251
Canberra	20	76	-25	-7	277	0	-1	-241	372	-8
Darwin	11	-118	-233	6	-31	-1	0	-12	92	-58
Melbourne	82	106	-471	-65	178	-241	-12	0	301	-411
Perth	482	-26	153	990	8	372	92	301	0	271
Sydney	-273	-314	-56	70	251	-8	-58	-411	271	0

The graph is a tad crowded, and for detailed information it is necessary to examine the table.

It is interesting to overlay this “map” on a physical map of Australia.

```
> load("aulatlong.RData")
> library(oz)
> oz()
> points(aulatlong, col = "red", pch = 16, cex = 1.5)
> comparePhysical <- function(lat = aulatlong$latitude, long = aulatlong$longitude,
+   x1 = aupoints[, 1], x2 = aupoints[, 2]) {
+   fitlat <- predict(lm(lat ~ x1 + x2))
+   fitlong <- predict(lm(long ~ x1 + x2))
+   x <- as.vector(rbind(lat, fitlat, rep(NA, 10)))
+   y <- as.vector(rbind(long, fitlong, rep(NA, 10)))
+   lines(x, y, col = 3, lwd = 2)
+ }
> comparePhysical()
```

An objection to `cmdscale()` is that it gives long distances the same weight as short distances. It is just as prepared to shift Canberra around relative to Melbourne and Sydney, as to move Perth. It makes more sense to give reduced weight to long distances, as is done by `sammon()` (*MASS*).

```
> aupoints.sam <- sammon(audists)

Initial stress      : 0.01573
stress after 10 iters: 0.00525, magic = 0.500
stress after 20 iters: 0.00525, magic = 0.500

> oz()
> points(aulatlong, col = "red", pch = 16, cex = 1.5)
> comparePhysical(x1 = aupoints.sam$points[, 1], x2 = aupoints.sam$points[,
+   2])
```

Notice how Brisbane, Sydney, Canberra and Melbourne now maintain their relative positions much better.

Now try full non-metric multi-dimensional scaling (MDS). This preserves only, as far as possible, the relative distances. A starting configuration of points is required. This might come from the configuration used by `cmdscale()`. Here, however, we use the physical distances.

```
> oz()
> points(aulatlong, col = "red", pch = 16, cex = 1.5)
> aupoints.mds <- isoMDS(audists, as.matrix(aulatlong))

initial value 11.875074
iter 5 value 5.677228
iter 10 value 4.010654
final value 3.902515
converged
```

```
> comparePhysical(x1 = aupoints.mds$points[, 1], x2 = aupoints.mds$points[,
+ 2])
```

Notice how the distance between Sydney and Canberra has been shrunk quite severely.

2 If distances must first be calculated ...

There are two functions that can be used to find distances – `dist()` that is in the base *statistics* package, and `daisy()` that is in the *cluster* package. The function `daisy()` is the more flexible. It has a parameter `stand` that can be used to ensure standardization when distances are calculated, and allows columns that are factor or ordinal. Unless measurements are comparable (e.g., relative growth, as measured perhaps on a logarithmic scale, for different body measurements), then it is usually desirable to standardize before using ordination methods to examine the data.

```
> library(cluster)
> load("diabetes.RData")
> data(diabetes)
> diadist <- daisy(diabetes[, -1], stand = TRUE)
> plot(density(diadist, from = 0))
```

3 Genetic Distances

Here, matching genetic DNA or RNA or protein or other sequences are available from each of the different species. Distances are based on probabilistic genetic models that describe how gene sequences change over time. The package *ape* implements a number of alternative measures. For details see `help(dist.ape)`.

3.1 Hasegawa's selected primate sequences

The sequences were selected to have as little variation in rate, along the sequence, as possible. The sequences are available from:

<http://evolution.genetics.washington.edu/book/primates.dna>. They can be read into R as:

```
> library(ape)
> webpage <- "http://evolution.genetics.washington.edu/book/primates.dna"
> primates.dna <- read.dna(url(webpage))
```

Now calculate distances, using Kimura's F84 model, thus

```
> primates.dist <- dist.dna(primates.dna, model = "F84")
```

We now try for a two-dimensional representation, using `cmdscale()`.

```
> primates.cmd <- cmdscale(primates.dist)
> eqscplot(primates.cmd)
> rtleft <- c(4, 2, 4, 2)[unclass(cut(primates.cmd[, 1],
+ breaks = 4))]
> text(primates.cmd[, 1], primates.cmd[, 2], row.names(primates.cmd),
+ pos = rtleft)
```

Now see how well the distances are reproduced:

```
> d <- dist(primates.cmd)
> sum((d - primates.dist)^2)/sum(primates.dist^2)
```

```
[1] 0.1977138
```

This is large enough (20%, which is a fraction of the total sum of squares) that it may be worth examining a 3-dimensional representation.

```
> primates.cmd <- cmdscale(primates.dist, k = 3)
> cloud(primates.cmd[, 3] ~ primates.cmd[, 1] * primates.cmd[,
+      2])
> d <- dist(primates.cmd)
> sum((d - primates.dist)^2)/sum(primates.dist^2)

[1] 0.1045168
```

Now repeat the above with `sammon()` and `mds()`.

```
> primates.sam <- sammon(primates.dist, k = 3)

Initial stress      : 0.11291
stress after 10 iters: 0.04061, magic = 0.461
stress after 20 iters: 0.03429, magic = 0.500
stress after 30 iters: 0.03413, magic = 0.500
stress after 40 iters: 0.03409, magic = 0.500

> eqscplot(primates.sam$points)
> rtleft <- c(4, 2, 4, 2)[unclass(cut(primates.sam$points[,
+   1], breaks = 4))]
> text(primates.sam$points[, 1], primates.sam$points[, 2],
+      row.names(primates.sam$points), pos = rtleft)
```

There is no harm in asking for three dimensions, even if only two of them will be plotted.

```
> primates.mds <- isoMDS(primates.dist, primates.cmd, k = 3)

initial value 19.710924
iter 5 value 14.239565
iter 10 value 11.994621
iter 15 value 11.819528
iter 15 value 11.808785
iter 15 value 11.804569
final value 11.804569
converged

> eqscplot(primates.mds$points)
> rtleft <- c(4, 2, 4, 2)[unclass(cut(primates.mds$points[,
+   1], breaks = 4))]
> text(primates.mds$points[, 1], primates.mds$points[, 2],
+      row.names(primates.mds$points), pos = rtleft)
```

4 Distances between fly species

These data have been obtained from databases on the web. We extract them from an Excel `.csv` file and store the result in an object of class `"dist"`. Only below diagonal elements are stored, in column dominant order. For manipulating such an object as a matrix, should this be required, `as.matrix()` can be used to turn it into a square symmetric matrix. Specify, e.g., `dipdistM <- as.matrix(dipdist)`. For the clustering and ordination methods that are used here, storing it as a distance object is fine.

```
> diptera <- read.csv("Dipt.csv", comment = "", na.strings = c(NA,
+   ""))
> dim(diptera)
```

```
[1] 110 114
```

```
> species <- as.character(diptera[, 1])
> table(substring(species, 1, 1))

#
110

> species <- substring(species, 2)
> genus <- as.character(diptera[, 2])
> dipdist <- as.dist(diptera[1:110, 5:114])
> attributes(dipdist)$Labels <- species
> length(dipdist)
```

```
[1] 5995
```

Two of the functions that will be used – `sammon()` and `isoMDS()` – require all distances to be positive. For each zero distance, omit the second point. The function `zerodists()` can be used, in cases where there are one or more zero distances, to remove the second member of the pair from `dipdist`. This function is available from the usual web page.

```
> udist <- zerodists(omit = TRUE)
```

```
[1] 5565 5565
```

	Species1	zeros...1.	Species2
1	Bactrocera_philippinensis_(3)	30	Bactrocera_dorsalis_(3)
2	Lucilia_illustris_(4)	63	Lucilia_caesar_(4)
3	Sapromyza_indigena_(4)	86	Sapromyza_imitans_(3)
4	Sapromyza_laurisilvae_(4)	87	Sapromyza_inconspicua_(3)
zeros...2.			
1	23		
2	61		
3	84		
4	85		

Now use (i) classical metric scaling; (ii) `sammon` scaling; (iii) isometric multi-dimensional scaling, with i as the starting configuration; (iv) isometric multi-dimensional scaling, with ii as the starting configuration.

```
> dipt.cmd <- cmdscale(udist)
> genus <- sapply(strsplit(rownames(dipt.cmd), split = "_",
+   fixed = TRUE), function(x) x[1])
> nam <- names(sort(table(genus), decreasing = TRUE))
> genus <- factor(genus, levels = nam)
> plot(dipt.cmd, col = unclass(genus), pch = 16)
> dipt.sam <- sammon(udist)
> plot(dipt.sam$points, col = unclass(genus), pch = 16)
> dipt.mds <- isoMDS(udist, dipt.cmd)
> plot(dipt.mds$points, col = unclass(genus), pch = 16)
> dipt.mds2 <- isoMDS(udist, dipt.sam$points)
> plot(dipt.mds2$points, col = unclass(genus), pch = 16)
```

