

Chapter 7

Hierarchical cluster analysis

In Part 2 (Chapters 4 to 6) we defined several different ways of measuring distance (or dissimilarity as the case may be) between the rows or between the columns of the data matrix, depending on the measurement scale of the observations. As we remarked before, this process often generates tables of distances with even more numbers than the original data, but we will show now how this in fact simplifies our understanding of the data. Distances between objects can be visualized in many simple and evocative ways. In this chapter we shall consider a graphical representation of a matrix of distances which is perhaps the easiest to understand – a dendrogram, or tree – where the objects are joined together in a hierarchical fashion from the closest, that is most similar, to the furthest apart, that is the most different. The method of hierarchical cluster analysis is best explained by describing the algorithm, or set of instructions, which creates the dendrogram results. In this chapter we demonstrate hierarchical clustering on a small example and then list the different variants of the method that are possible.

Contents

The algorithm for hierarchical clustering

Cutting the tree

Maximum, minimum and average clustering

Validity of the clusters

Clustering correlations

Clustering a larger data set

The algorithm for hierarchical clustering

As an example we shall consider again the small data set in Exhibit 5.6: seven samples on which 10 species are indicated as being present or absent. In Chapter 5 we discussed two of the many dissimilarity coefficients that are possible to define between the samples: the first based on the matching coefficient and the second based on the Jaccard index. The latter index counts the number of ‘mismatches’ between two samples after eliminating the species that do not occur in either of the pair. Exhibit 7.1 shows the complete table of inter-sample dissimilarities based on the Jaccard index.

Exhibit 7.1 Dissimilarities, based on the Jaccard index, between all pairs of seven samples in Exhibit 5.6. For example, between the first two samples, A and B, there are 8 species that occur in on or the other, of which 4 are matched and 4 are mismatched – the proportion of mismatches is $4/8 = 0.5$. Both the lower and upper triangles of this symmetric dissimilarity matrix are shown here (the lower triangle is outlined as in previous tables of this type).

samples	A	B	C	D	E	F	G
A	0	0.5000	0.4286	1.0000	0.2500	0.6250	0.3750
B	0.5000	0	0.7143	0.8333	0.6667	0.2000	0.7778
C	0.4286	0.7143	0	1.0000	0.4286	0.6667	0.3333
D	1.0000	0.8333	1.0000	0	1.0000	0.8000	0.8571
E	0.2500	0.6667	0.4286	1.0000	0	0.7778	0.3750
F	0.6250	0.2000	0.6667	0.8000	0.7778	0	0.7500
G	0.3750	0.7778	0.3333	0.8571	0.3750	0.7500	0


The first step in the hierarchical clustering process is to look for the pair of samples that are the most similar, that is are the closest in the sense of having the lowest dissimilarity – this is the pair B and F, with dissimilarity equal to 0.2000. These two samples are then joined at a level of 0.2000 in the first step of the dendrogram, or clustering tree (see the first diagram of Exhibit 7.3, and the vertical scale of 0 to 1 which calibrates the level of clustering). The point at which they are joined is called a *node*.

We are basically going to keep repeating this step, but the only problem is how to calculate the dissimilarity between the merged pair (B,F) and the other samples. This decision determines what type of hierarchical clustering we intend to perform, and there are several choices. For the moment, we choose one of the most popular ones, called the *maximum*, or *complete linkage*, method: the dissimilarity between the merged pair and the others will be the maximum of the pair of dissimilarities in each case. For example, the dissimilarity between B and A is 0.5000, while the dissimilarity between F and A is 0.6250. hence we choose the maximum of the two, 0.6250, to quantify the dissimilarity between (B,F) and A. Continuing in this way we obtain a new dissimilarity matrix Exhibit 7.2.

Exhibit 7.2 Dissimilarities calculated after B and F are merged, using the ‘maximum’ method to recomputed the values in the row and column labelled (B,F).

samples	A	(B,F)	C	D	E	G
A	0	0.6250	0.4286	1.0000	0.2500	0.3750
(B,F)	0.6250	0	0.7143	0.8333	0.7778	0.7778
C	0.4286	0.7143	0	1.0000	0.4286	0.3333
D	1.0000	0.8333	1.0000	0	1.0000	0.8571
E	0.2500	0.7778	0.4286	1.0000	0	0.3750
G	0.3750	0.7778	0.3333	0.8571	0.3750	0

Exhibit 7.3 First two steps of hierarchical clustering of Exhibit 7.1, using the ‘maximum’ (or ‘complete linkage’) method.



The process is now repeated: find the smallest dissimilarity in Exhibit 7.2, which is 0.2500 for samples A and E, and then cluster these at a level of 0.25, as shown in the second figure of Exhibit 7.3. Then recomputed the dissimilarities between the merged pair (A,E) and the rest to obtain Exhibit 7.4. For example, the dissimilarity between (A,E) and (B,F) is the maximum of 0.6250 (A to (B,F)) and 0.7778 (E to (B,F)).

Exhibit 7.4 Dissimilarities calculated after A and E are merged, using the ‘maximum’ method to recomputed the values in the row and column labelled (A,E).

samples	(A,E)	(B,F)	C	D	G
(A,E)	0	0.7778	0.4286	1.0000	0.3750
(B,F)	0.7778	0	0.7143	0.8333	0.7778
C	0.4286	0.7143	0	1.0000	0.3333
D	1.0000	0.8333	1.0000	0	0.8571
G	0.3750	0.7778	0.3333	0.8571	0

In the next step the lowest dissimilarity in Exhibit 7.4 is 0.3333, for C and G – these are merged, as shown in the first diagram of Exhibit 7.6, to obtain Exhibit 7.5. Now the smallest dissimilarity is 0.4286, between the pair (A,E) and (B,G), and they are shown merged in the second diagram of Exhibit 7.6. Exhibit 7.7 shows the last two dissimilarity matrices in this process, and Exhibit 7.8 the final two steps of the construction of the dendrogram, also called a *binary* tree because at each step two objects (or clusters of objects) are merged. Because there are 7 objects to be clustered, there are 6 steps in the sequential process (i.e., one less) to arrive at the final tree where all objects are in a single cluster. For botanists that may be reading this: this is an upside-down tree, of course!

Exhibit 7.5 Dissimilarities calculated after C and G are merged, using the 'maximum' method to recompute the values in the row and column labelled (C,G).

samples	(A,E)	(B,F)	(C,G)	D
(A,E)	0	0.7778	0.4286	1.0000
(B,F)	0.7778	0	0.7778	0.8333
(C,G)	0.4286	0.7778	0	1.0000
D	1.0000	0.8333	1.0000	0

Exhibit 7.6 The third and fourth steps of hierarchical clustering of Exhibit 7.1, using the 'maximum' (or 'complete linkage') method. The point at which objects (or clusters of objects) are joined is called a node.


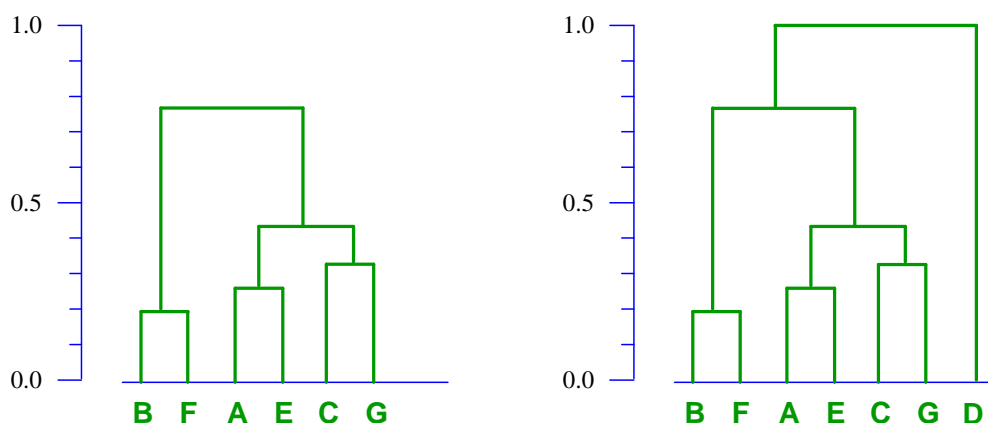


Exhibit 7.7 Dissimilarities calculated after C and G are merged, using the 'maximum' method to recompute the values in the row and column labelled (C,G).

samples	(A,E,C,G)	(B,F)	D
(A,E,C,G)	0	0.7778	1.0000
(B,F)	0.7778	0	0.8333
D	1.0000	0.8333	0

samples	(A,E,C,G,B,F)	D
(A,E,C,G,B,F)	0	1.0000
D	1.0000	0

Exhibit 7.8 The fifth and sixth steps of hierarchical clustering of Exhibit 7.1, using the ‘maximum’ (or ‘complete linkage’) method. The dendrogram on the right is the final result of the cluster analysis. In the clustering of n objects, there are $n-1$ nodes (i.e. 6 nodes in this case).



Cutting the tree

The final dendrogram on the right of Exhibit 7.8 is a compact visualization of the dissimilarity matrix in Exhibit 7.1, based on the presence-absence data of Exhibit 5.6. Interpretation of the structure of data is made much easier now – we can see that there are three pairs of samples that are fairly close, two of these pairs ((A,E) and (C,G)) are in turn close to each other, while the single sample D separates itself entirely from all the others. Because we used the ‘maximum’ method, all samples clustered below a particular level of dissimilarity will have inter-sample dissimilarities less than that level. For example, 0.5 is the point at which samples are exactly as similar to one another as they are dissimilar, so if we look at the clusters of samples below 0.5 – i.e., (B,F), (A,E,C,G) and (D) – then within each cluster the samples have more than 50% similarity, in other words more than 50% co-presences of species. The level of 0.5 also happens to coincide in the final dendrogram with a large jump in the clustering levels: the node where (A,E) and (C,G) are clustered is at level of 0.4286, while the next node where (B,F) is merged is at a level of 0.7778. This is thus a very convenient level to *cut* the tree. If the branches are cut at 0.5, we are left with the three clusters of samples (B,F), (A,E,C,G) and (D), which can be labelled types 1, 2 and 3 respectively. In other words, we have created a categorical variable, with three categories, and the samples are categorized as follows:

A	B	C	D	E	F	G
2	1	2	3	2	1	2

Checking back to Chapter 2, this is exactly the objective which we described in the lower right hand corner of the multivariate analysis scheme (Exhibit 2.2) – to reveal a categorical variable which underlies the structure of a data set.

Maximum, minimum and average clustering

The crucial choice when deciding on a cluster analysis algorithm is to decide how to quantify dissimilarities between two clusters. The algorithm described above was characterized by the fact that at each step, when updating the matrix of dissimilarities, the maximum of the between-cluster dissimilarities was chosen. This is also known as *complete linkage* cluster analysis, because a cluster is formed when all the dissimilarities ('links') between pairs of objects in the cluster are less than a particular level. There are several alternatives to complete linkage as a clustering criterion, and we only discuss two of these: minimum and average clustering.

The 'minimum' method goes to the other extreme and forms a cluster when only one pair of dissimilarities (not all) is less than a particular level – this is known as *single linkage* cluster analysis. So at every updating step we choose the minimum of the two distances and two clusters of objects can be merged when there is a single close link between them, irrespective of the other inter-object distances. In general, this is not a suitable choice for most applications, because it can lead to clusters that are quite heterogeneous internally, and the usual object of clustering is to obtain homogeneous clusters.

The 'average' method is an attractive compromise where dissimilarities are averaged at each step, hence the name *average linkage* cluster analysis. For example, in Exhibit 7.1 the first step of all types of cluster analysis would merge B and F. But then calculating the dissimilarity between A, for example, and (B,F) is where the methods distinguish themselves. The dissimilarity between A and B is 0.5000, and between A and F it is 0.6250. Complete linkage chooses the maximum: 0.6250; single linkage chooses the minimum: 0.5000; while average linkage chooses the average: $(0.5000+0.6250)/2 = 0.5625$.

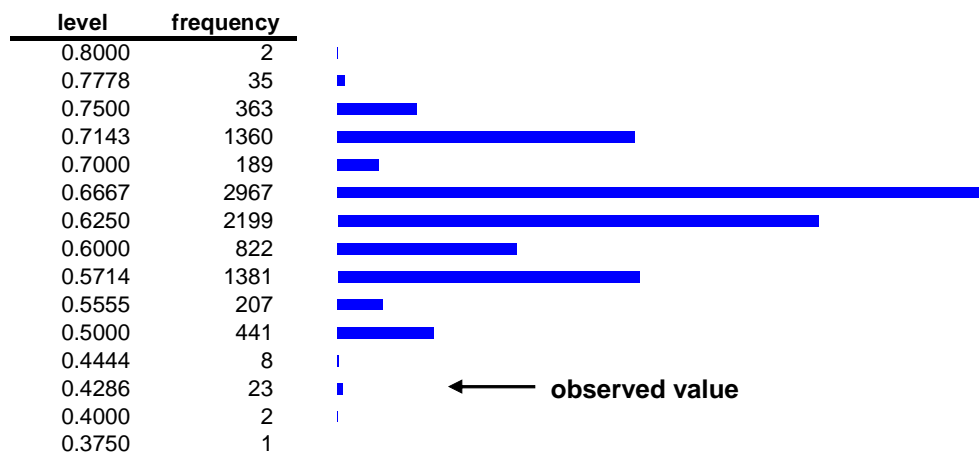
Validity of the clusters

If a cluster analysis is performed on a data matrix, a set of clusters can always be obtained, even if there is no actual grouping of the objects, in this case the samples. So how can we evaluate whether the three clusters in this example are not just any old three groups which we would have obtained on random data with no structure? There is a vast literature on validity of clusters (we give some references in the Bibliography, Appendix E) and here we shall explain one approach based on permutation testing. In our example, the three clusters were formed so that internally in each cluster formed by more than one sample the between-sample dissimilarities were all less than 0.5000. In fact, if we look at the result in the right hand picture of Exhibit 7.8, the cutpoint for three clusters can be brought down to the level of 0.4286, where (A,E) and (C,G) joined together. As in all statistical considerations of significance, we ask whether this is an unusual result or whether it could have arisen merely by chance. To answer this question we need an idea of what might have happened in chance results, so that we can judge our actual finding. This so-called "null distribution" can be generated through permuting the data in some reasonable way, evaluating the statistic of interest, and doing this many times (or for all permutations if this is feasible computationally) to obtain a distribution of the statistic. The statistic of interest could be that value at which the three clusters are formed, but we need to choose carefully

how we perform the permutations, and this depends on how the data were collected. We consider two possible assumptions, and show how different the results can be.

The first assumption is that the column totals of Table Exhibit 5.6 are fixed; that is, that the 10 species are present, respectively, 3 times in the 7 samples, 6 times, 4 times, 3 times and so on. Then the permutation involved would be to simply randomly shuffle the zeros and ones in each column to obtain a new presence-absence matrix with exactly the same column totals as before. Performing the complete linkage hierarchical clustering on this matrix leads to that value where the three cluster solution is achieved, and becomes one observation of the null permutation distribution. We did this 9999 times, and along with our actual observed value of 0.4286, the 10000 values are graphed in Exhibit 7.9 (we show it as a horizontal bar chart because there are only 15 different values observed of this value, shown here with their frequencies). The value we actually observed is one of the smallest – the number of permuted matrices that generates this value or a lower value is 26 out of 10000, so that in this sense our data are very unusual and the ‘significance’ of the three-cluster solution can be quantified with a p -value of 0.0026. The other 9974 random permutations all lead to generally higher inter-sample dissimilarities such that the level at which three-cluster solutions are obtained is 0.4444 or higher (0.4444 corresponds to 4 mismatches out of 9).

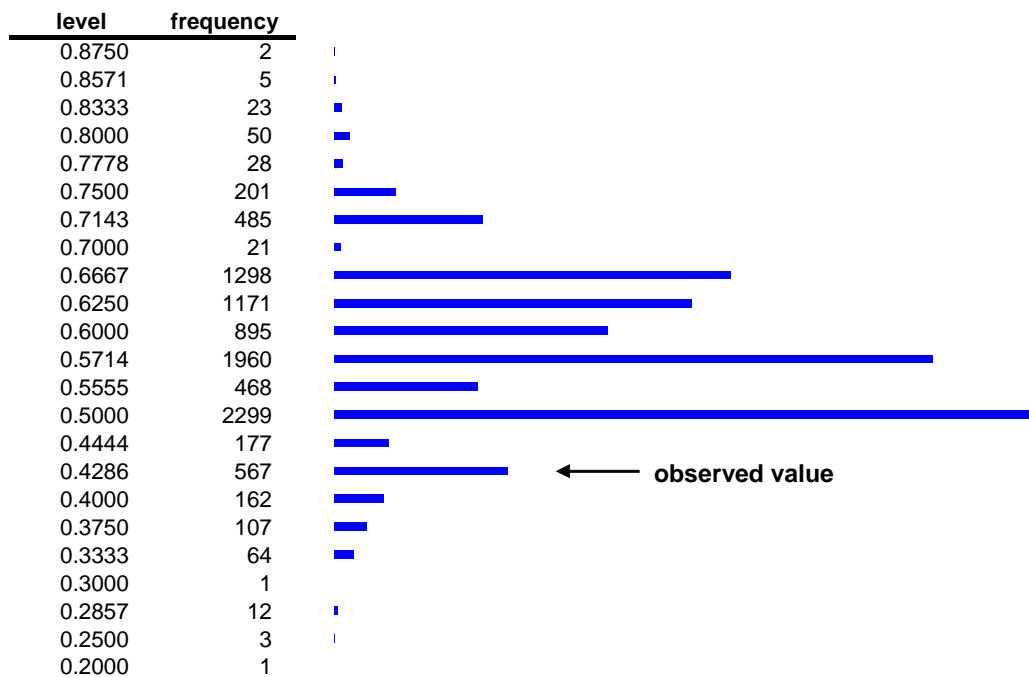
Exhibit 7.9 Bar chart of the 10000 values of the three-cluster solutions obtained by permuting the columns of the presence-absence data, including the value we observed in the original unpermuted data matrix.



The second and alternative possible assumption for the computation of the null distribution could be that the column margins are not fixed, but random; in other words, we relax the fact that there were exactly 3 samples that had species sp1, for example, and assume a binomial distribution for each column, using the observed proportion (3 out of 7 for species sp1) and the number of samples (7) as the binomial parameters. Thus there can be 0 up to 7 presences in each column, according to the binomial probabilities for each species. This gives a much wider range of possibilities for the null distribution, and leads us to a different conclusion about our three observed clusters. The permutation distribution

is now shown in Exhibit 7.10, and now our observed value of 0.4286 does not look so unusual, since 917 out of 10000 values in the distribution are less than or equal to it, giving an estimated P -value of 0.0917.

Exhibit 7.10 Bar chart of the 10000 values of the three-cluster solutions obtained by generating binomial data in each column of the presence-absence matrix, according to the probability of presence of each species.



So, as in many situations in statistics, the result and decision depends on the initial assumptions. Could we have observed the presence of species s_1 less or more than 3 times in the 7 samples (and so on for the other species)? In other words, according to the binomial distribution with $n=7$, and $p=3/7$, the probabilities of observing k presences of species sp_1 ($k=0, 1, \dots, 7$) are:

0	1	2	3	4	5	6	7
0.020	0.104	0.235	0.294	0.220	0.099	0.025	0.003

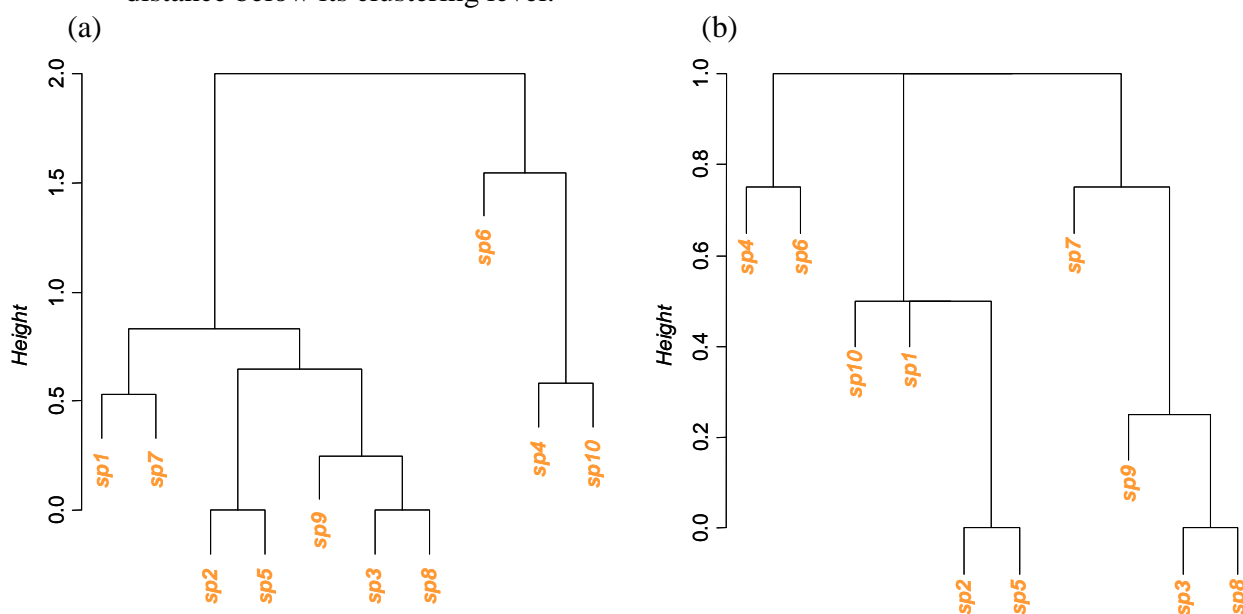
If this assumption (and similar ones for the other nine species) is realistic, then the cluster significance is 0.0917. However, if the first assumption is adopted (that is, the probability of observing 3 presences for species s_1 is 1 and 0 for other possibilities), then the significance is 0.0028. Our feeling is that perhaps the binomial assumption is more realistic, in which case our cluster solution could be observed in just over 9% of random cases – this gives us an idea of the validity of our results and whether we are dealing with real clusters or not. The value of 9% is a measure of ‘clusteredness’ of our samples in terms of the Jaccard index: the lower this measure, the more they are clustered, and the higher the measure, the more the samples lie in a continuum. Lack of evidence of

‘clusteredness’ does not mean that the clustering is not useful: we might want to divide up the space of the data into separate regions, even though the borderlines between them are ‘fuzzy’. And speaking of ‘fuzzy’, there is an alternative form of cluster analysis (fuzzy cluster analysis, not treated specifically in this book) where samples are classified fuzzily into clusters, rather than strictly into one group or another – this idea is similar to the fuzzy coding we described in Chapter 3.

Clustering correlations on variables

Just like we clustered samples, so we can cluster variables in terms of their correlations, or distances based on their correlations as described in Chapter 6. The dissimilarity based on the Jaccard index can also be used to measure similarity between species – the index counts the number of samples that have both species of the pair, relative to the number of samples that have at least one of the pair, and the dissimilarity is 1 minus this index. Exhibit 7.11 shows the cluster analyses based on these two alternatives, for the columns of Exhibit 5.6, using the graphical output this time of the R function `hclust` for hierarchical clustering. The fact that these two trees are so different is no surprise: the first one is based on the correlation coefficient takes into account the co-absences, which strengthens the correlation, while the second does not. Both have the pairs (*sp2,sp5*) and (*sp3,sp8*) at zero dissimilarity because these are identically present and absent across the samples. Species *sp1* and *sp7* are close in terms of correlation, due to co-absences – *sp7* only occurs in one sample, which also has *sp1*, a species which is absent in four other samples. Notice in Exhibit 7.11(b) how species *sp10* and *sp1* both join the cluster (*sp2,sp5*) at the same level (0.5).


Exhibit 7.11 Complete linkage cluster analyses of (a) $1-r$ (1 minus the correlation coefficient between species); (b) Jaccard dissimilarity between species (1 minus the Jaccard similarity index). The R function `hclust` which calculates the dendrograms places the object (species) labels at a constant distance below its clustering level.



Clustering a larger data set


The more objects there are to cluster, the more complex becomes the result. In Exhibit 4.5 we showed part of the matrix of standardized Euclidean distances between the 30 sites of Exhibit 1.1, and Exhibit 7.12 shows the hierarchical clustering of this distance matrix, using complete linkage. There are two obvious places where we can cut the tree, at about level 3.4, which gives four clusters, or about 2.7, which gives six clusters. To get an idea

Exhibit 7.12 Complete linkage cluster analyses of the standardized Euclidean distances of Exhibit 4.5.



of the ‘clusteredness’ of these data, we performed a permutation test similar to the one described above, where the data are randomly permuted within their columns and the cluster analysis repeated each time to obtain 6 clusters. The permutation distribution of levels at which 6 clusters are formed is shown in Exhibit 7.13 – the observed value in Exhibit 7.12 (i.e., where (s2,s14) joins (s25,s23,s30,s12,s16,s27)) is 2.357, which is clearly not an unusual value. The estimated p -value according to the proportion of the distribution to the left of 2.357 in Exhibit 7.13 is $p = 0.3388$, so we conclude that these samples do not have a non-random cluster structure – they form more of a continuum, which will be the subject of Chapter 9.

Exhibit 7.13 Estimated permutation distribution for the level at which 6 clusters are formed in the cluster analysis of Exhibit 7.12, showing the value actually observed. Of the 10000 permutations, including the observed value, 3388 are less than or equal to the observed value, giving an estimated p -value for clusteredness of 0.3388.



SUMMARY: Hierarchical cluster analysis

1. Hierarchical cluster analysis of n objects is defined by a stepwise algorithm which merges two objects at each step, the two which have the least dissimilarity.
2. Dissimilarities between clusters of objects can be defined in several ways; for example, the maximum dissimilarity (complete linkage), minimum dissimilarity (single linkage) or average dissimilarity (average linkage).
3. Either rows or columns of a matrix can be clustered – in each case we choose the appropriate dissimilarity measure that we prefer.
4. The results of a cluster analysis is a binary tree, or dendrogram, with $n-1$ nodes. The branches of this tree are cut at a level where there is a lot of 'space' to cut them, that is where the jump in levels of two consecutive nodes is large.
5. A permutation test is possible to validate the chosen number of clusters, that is to see if there really is a non-random tendency for the objects to group together.