

Chapter 8

Data Clustering

8.1 Introduction

The purpose of clustering data is so that we can represent groups of data by an exemplar, or cluster center. This is one way to deal with data sets that would otherwise be too large to handle.

Clustering can be intuitively clear, but algorithmically difficult. For example, consider the data in the first graph in Figure 8.1. I think you would probably agree that there appears to be three natural clusterings.




Figure 8.1: We see data that have no class labels, but there seems to be a natural separation of the data into three clumps.

A clustering then is primarily building a function that can take a data point as its input, and then output which cluster (or class) the data point belongs to.

Data may already have labels- We would call a clustering of that type to be *supervised learning*, since there are examples of “correct” labels. In our example, and in general, we’ll be working with the unsupervised learning task, where the data points are given to us with no labels.

Definition: Unsupervised Clustering: Given a data set, X , whose elements are vectors $\mathbf{x}^{(i)} \in \mathbb{R}^n$, we want to construct a “membership function” which has its domain in \mathbb{R}^n and will output the cluster index (or label). Defining this function as m , we have, for the i^{th} data point, its associated class label g_i :

$$m(\mathbf{x}^{(i)}) = g_i$$

where g_i is an integer from 1 to k (k being the number of clusters desired). Later we’ll discuss how the algorithm might determine the value of k on its own, so for now, we’ll consider k to be given.

The biggest issue with unsupervised clustering is that the problem specification is very ill-posed, meaning that there are many, many ways that one might build such a function m . For two extreme examples, consider these two membership functions:

$$m_1(\mathbf{x}^{(i)}) = 1$$

where i is the index for the data points. Another function that is about as useful as m_1 is the following:

$$m_2(\mathbf{x}^{(i)}) = i$$


In this case, every data point is its own cluster. Neither situation is ideal, and we'll try to be somewhere in between these extremes.

In order to do determine some “best” clustering, that usually means that we'll be trying to minimize some error function, and the easiest way to do this is through *Voronoi Cells*:

Definition: Let $\{\mathbf{c}^{(i)}\}_{i=1}^k$ be points in \mathbb{R}^n . These points form k Voronoi Cells, where the j^{th} cell is defined as the set of points that are closer to cell j than any other cluster:

$$V_j = \left\{ \mathbf{x} \in \mathbb{R}^n \mid \|\mathbf{x} - \mathbf{c}^{(j)}\| \leq \|\mathbf{x} - \mathbf{c}^{(i)}\|, i = 1, 2, \dots, k \right\}$$

The points $\{\mathbf{c}^{(i)}\}_{i=1}^k$ are called *cluster centers*. In the uncommon occurrence that a point \mathbf{x} lies on the border between cells, it is customary to include it in the cell whose index is smaller (although one would fashion the decision on the problem at hand). The reader might note that a Voronoi cell has a piecewise linear border as shown.



Examples in “Nature”

1. In [6], Voronoi cells are used to define the “area potentially available around a tree”. That is, each tree in a stand represents the center of the cell.
2. Using a map of the campus with the emergency telephone boxes marked out and used as the centers, we could always tell where the closest phone is located.


We can draw a Voronoi diagram by hand: Between neighboring cells, draw a line (that will be erased at the end), then draw the perpendicular bisectors for each line drawn. This can get complicated fairly quickly, so we will be using Matlab or Python to produce the plots. The algorithms that do these plots are very interesting (see any text in Computation Geometry) but are beyond the scope of our text.

Matlab Example:

```
X=randn(10,2); %X is a tall matrix.
voronoi(X(:,1),X(:,2));
```

We can also have Matlab return the vertices of the Voronoi cells to plot them manually:

```
[vx,vy]=voronoi(X(:,1),X(:,2));
plot(vx,vy,'k-',X(:,1),X(:,2),'r*');
```



Python Example

In Python, we'll also build a Voronoi diagram using 10 random points in the plane. For this, we'll need `Voronoi` from `scipy`, and some other things for displaying the plot.

```
import matplotlib.pyplot as plt
from scipy.spatial import Voronoi, voronoi_plot_2d

A=np.random.rand(10,2) # Note that A must be a tall matrix.
vor=Voronoi(A)

# Plot the results
fig = voronoi_plot_2d(vor)
plt.show()
```

In the algorithms that we work with, it will be convenient to have a function that will identify points within a given cluster- That is, will have the value 0 if the point is not in a given cluster, or 1 if it is. The “characteristic function” is typical for that purpose, and we'll define it as the following:

$$\chi_i(\mathbf{x}) = \begin{cases} 1 & \text{if } m(\mathbf{x}) = i \\ 0 & \text{otherwise} \end{cases}$$

One general way to measure the goodness of a clustering algorithm is to use a measure called the **distortion error** for a given cluster i , and the total distortion error¹. Given a clustering of N points, let N_i denote the number of points in cluster i . We first define the error for cluster i , then sum them all for the overall error. Here is the definition for the distortion error for the i^{th} cluster, where we notice that if the point $\mathbf{x}^{(k)}$ is not in cluster i , then we simply add 0:

$$E_i = \frac{1}{N_i} \sum_{k=1}^N \|\mathbf{x}^{(k)} - \mathbf{c}^{(i)}\|^2 \chi_i(\mathbf{x}^{(k)})$$

so that the overall distortion error is defined to be:

$$E_{\text{total}} = \sum_{k=1}^p E_k$$

Now that we have a way of measuring the goodness of a clustering, we want to have an algorithm that will minimize the function.

8.2 K-means clustering

A relatively fast method to perform the clustering is the k -means clustering. In some circles, this is called Lloyd's algorithm, and there was an update done by Linde, Buzo and Gray to call that the Linde-Buzo-Gray (LBG) algorithm [26].

The membership function for both k -means and LBG is defined so that we are in cluster i if we are in the Voronoi cell defined by the i^{th} cluster center. That is,

$$m(\mathbf{x}) = i \text{ iff } \|\mathbf{x} - \mathbf{c}^{(i)}\| < \|\mathbf{x} - \mathbf{c}^{(j)}\| \text{ } i \neq j, \text{ } j = 1 : p \tag{8.1}$$

In the rare circumstance that we have a tie, we'll choose to put the point with the center whose index is smallest (but this is ad-hoc).

¹Some authors do not square the norms, but it is more convenient for the theory, and does not change the end results.

The Euclidean Distance Matrix

You've seen that we'll need a function that will compute all of the inter-point distances for us. To be more precise, we'll need a function that will input two sets of points in \mathbb{R}^n , say there are M points in one set (organized in an $M \times n$ matrix), and N points in another (organized in an $N \times n$ matrix). The output will then be an $M \times N$ matrix where the (i, j) position is the distance between the point i in the first set and point j in the second set.

For Matlab/Octave, we'll write a short function to do this for us called `edm.m`. which will be available on the class website. For Python, you can compute the pairwise distances using the example below:

```
from sklearn.metrics.pairwise import euclidean_distances
```

```
X = [[0, 1], [1, 1]]
D=euclidean_distances(X, X) #D is 2 x 2 with distances.
```

The Algorithm

Keeping in mind that we'll be using Matlab for its implementation, we'll write the algorithm in matrix form.

Let X be a matrix of p data points in \mathbb{R}^n (so each data point is a row, and the matrix is $p \times n$), and let C denote a matrix of k centers in \mathbb{R}^n (each "center" is a row, and the matrix is $k \times n$).

At each pass of the algorithm, the membership function requires us to take $p \times k$ distance calculations, followed by p sorts (each point has to find its own cluster index), and this is where the algorithm takes most of the computational time.

K-Means Clustering

Given: p data points in \mathbb{R}^n and k cluster centers (as vectors in \mathbb{R}^n). Then:

- Sort the data into k sets by using the membership function in Equation 8.1.
- Re-set center c_i as the centroid (mean) of the data in the i^{th} set.
- Compute the distortion error.
- Repeat until the distortion error no longer decreases (either slows to nothing or starts increasing).

A Toy Example:

Here we use the k -means algorithm three times on a small data set, so you can have a test set for the code that you write.

Let the data set $X^{5 \times 2}$ be the matrix whose rows are given by the data points in the plane:

$$\left\{ \begin{bmatrix} 1 \\ 2 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \end{bmatrix}, \begin{bmatrix} -1 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 1 \end{bmatrix}, \begin{bmatrix} -1 \\ -1 \end{bmatrix} \right\}$$

You might try drawing these on a plane. You may see a clustering of three points to the right, and a clustering of two points to the left.

We will use two cluster centers, (typically initialized randomly from the data) in this case:

$$\begin{bmatrix} 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 1 \\ 1 \end{bmatrix}$$

The EDM of distances will be $D^{5 \times 2}$ and is the distance between our 5 data points and 2 cluster centers.

The membership function then looks for the minimum in each row, shown as the vector M below:

$$D = \begin{bmatrix} \sqrt{2} & 1 \\ 0 & 1 \\ \sqrt{2} & \sqrt{5} \\ 1 & 0 \\ \sqrt{5} & \sqrt{8} \end{bmatrix}, \quad M = \begin{bmatrix} 2 \\ 1 \\ 1 \\ 2 \\ 1 \end{bmatrix}$$

Resorting into two clusters,

$$\left\{ \begin{bmatrix} 0 \\ 1 \end{bmatrix}, \begin{bmatrix} -1 \\ 0 \end{bmatrix}, \begin{bmatrix} -1 \\ -1 \end{bmatrix} \right\} \quad \left\{ \begin{bmatrix} 1 \\ 2 \end{bmatrix}, \begin{bmatrix} 1 \\ 1 \end{bmatrix} \right\}$$

The cluster centers are re-set as the centroids:

$$\begin{bmatrix} -2/3 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 3/2 \end{bmatrix}$$

After one more sort and centroid calculation, we get:

$$\begin{bmatrix} -1 \\ -1/2 \end{bmatrix}, \begin{bmatrix} 2/3 \\ 4/3 \end{bmatrix}$$

which does not change afterward (so the clusters also remain the same).

The basic algorithm may have a few shortcomings. For example,

- The centers may be (or become) empty- That is, there are no data points closest to a given center. This is a good chance of this happening if you simply choose random numerical values as the initial centers, that's why we recommend choosing random *data points* as initial centers- You'll have at least one point in each cluster that way.
- We had to pre-define the number of clusters. Is it possible to construct an algorithm that will grow or shrink the number of clusters based on some performance measure? Yes, and that change is the LBG algorithm.

Matlab Note

If the "Statistics Toolbox" is available for your copy of Matlab, then `kmeans` is a built-in command.

Matlab's clustering routine expects the data to be given row-wise, so if we have p points in \mathbb{R}^n , then Matlab expects the data matrix to be $p \times n$.

Here's the previous example using the stats toolbox:

```
X=[1,0,-1,1,-1;2,1,0,1,-1]';
C=[0,1;1,1]';
[M,C1]=kmeans(X,2,'start',C);
```

If we didn't have an initial clustering C , the command is a lot shorter. For example, with just the data and number of clusters, we would have:

```
X=[1,0,-1,1,-1;2,1,0,1,-1]';
[M,C]=kmeans(X,2);
```

Matlab gives us several ways of initializing the clusters- See the documentation.

Exercises

1. Fill in the missing details from the example in the text: Fill in the EDM and the vector M after the second sort, and continue the example one more iteration to show that the clusters do not change. Hint: To compute the distance matrix, use the points as they were originally ordered.
2. Given p scalars, x_1, x_2, \dots, x_p , show (using calculus) that the number μ that minimizes the function:

$$E(\mu) = \sum_{k=1}^p (x_k - \mu)^2$$

is the mean, $\mu = \bar{x}$. Hint: We think of x_1, \dots, x_p as being fixed numerical values, so E is a function of only one variable μ .

3. Generalize the last exercise so that the scalars (and μ) are now vectors in \mathbb{R}^n . That is, given a fixed set of p vectors in \mathbb{R}^n , show that the vector μ that minimizes:

$$E(\mu) = \sum_{k=1}^p \|\mathbf{x}^{(i)} - \mu\|^2$$

is the mean (in \mathbb{R}^n).

4. Write the following as a Matlab function. The abbreviation EDM is for Euclidean Distance Matrix. Good programming style: Include comments!

```
function z=edm(w,p)
% A=edm(w,p)
% Input: w, number of points by dimension
% Input: p is number of points by dimension
% Output: Matrix z, number points in w by number pts in p
%         which is the distance from one point to another

[S,R] = size(w);
[Q,R2] = size(p);
p=p';
if (R ~= R2), error('Inner matrix dimensions do not match. '),end

z = zeros(S,Q);
if (Q<S)
    p = p';
    copies = zeros(1,S);
    for q=1:Q
        z(:,q) = sum((w-p(q+copies,:)).^2,2);
    end
else
    w = w';
    copies = zeros(1,Q);
    for i=1:S
        z(i,:) = sum((w(:,i+copies)-p).^2,1);
    end
end
z = z.^0.5;
```

5. If you don't have the `kmeans` function built-in, try writing one using the EDM function. It is fairly straightforward using the EDM function and the `min` function.

6. Will the cluster placement at the end of the algorithm be independent of where the clusters start? Answer this question using all the different possible pairs of initial points from our toy data set, and report your findings- In particular, did some give a better distortion error than others?

LBG modifications

We can think about ways to prune and grow the number of clusters, rather than making it a predefined quantity. Here are some suggestions for modifications you can try. This actually changes the k -means algorithm into the LBG algorithm.

- Split that cluster with the highest distortion measure, and continue to split clusters until the overall distortion measure is below some preset value. The two new cluster centers can be initialized a number of ways- Here is one option:

$$\mathbf{c}^{(i_1, i_2)} = \mathbf{c}^{(i)} \pm \epsilon$$

However, this may again lead to empty clusters.

- We can prune away clusters that either have a small number of points, or whose distortion measure is smaller than some pre-set quantity. It is fairly easy to do this- Just delete the corresponding cluster center from the array.

8.3 Neural Gas

So far, the clustering algorithms we have considered use arithmetic means to place cluster centers. This works well when the data consists of “clouds” (coming from a random process, for example). This process does not work well when the data lies on some curved manifold²

If we imagine all of the data to be on the boundary of a circle, for example, taking an average will place the cluster centers off the circle.



Figure 8.2: If the data lies on a curved surface, then using the mean may pull the centers (stars in the graph) out of the data.

But there is something that may be exploited about the cluster centers to help us in visualizing how the data lies in its space.

It would be helpful if we knew what kinds of *structure* the data may have. For example, in Figure 8.2, it would be helpful to know that the data is actually (locally) one dimensional.

The method we will use to determine the structure of the data will be to define connections (or edges) between clusters. Consider the data in Figure 8.3. It is this kind of structure that the clustering should reflect- Which cluster centers are “neighboring” points in the data? Just as importantly, which are *not* neighbors in the data?

²A manifold is generally defined to be an object that looks locally like \mathbb{R}^k for some k . For example, a path in three dimensions is locally 1 dimensional.




Figure 8.3: The star shaped points are the cluster centers. It would be nice to know which cluster centers belong together in a single “blob” of data, versus which cluster centers are in the other set of data.

In mathematical terminology, we want to find what is called a *topology preserving* clustering. For us, the *topology* of the clusters is defined by how the cluster centers are connected to each other (that is, so we know which cluster centers are neighbors and which are not).

More generally, we want the clustering to be what is called “topology preserving” in the sense that clusters defined to be neighbors in the cluster center topology are actually neighbors in the data, and vice versa. Before continuing, some graphical examples should help us with this concept.

In Figure 8.4, we see three topologies mapped to the data, which is a uniform distribution of points in \mathbb{R}^2 . In the first picture, the topology of the cells is a one dimensional set. In this situation, the cluster mapping is not topology preserving, because neighboring cells in the topology are not adjacent in the plane. In the second situation, we have a three-dimensional topology mapping to the plane. In this case, neighboring data points in \mathbb{R}^2 are mapped to non-neighboring cells in the topology. Only in the third picture do we see that both parts of the topology preserving mapping are satisfied.

The Neural Gas Algorithm [27, 29, 28] is designed with a couple of big changes in mind (changes from k-means and LBG).

1. We don’t want to use arithmetic means if possible. It is more desirable to have cluster centers that are actually embedded in with the data (like in Figure 8.3) rather than every cluster outside of the data (like in Figure 8.2).
2. We want the algorithm to build connections between the cluster centers. Hopefully this will tell us if we have separated regions of data (like in Figure 8.3).

Embedding the Clusters

To solve the first problem, we are going to endow the data with the power to attract the cluster centers. That is, when a data point \mathbf{x} is chosen, we will determine which cluster center is closest (the “winner”), and we will move the cluster center towards that point. We don’t want to go very far, just a slight move in that direction.

Mathematically, for a given data point \mathbf{x} , find the winning center, \mathbf{c} and move in the direction $\mathbf{x} - \mathbf{c}$. The update rule would then look like:

$$\mathbf{c}_{\text{new}} = \mathbf{c}_{\text{old}} + h(t)(\mathbf{x} - \mathbf{c}_{\text{old}}) \tag{8.2}$$

where $h(t)$ will be a function determining how far to move. Normally, at the beginning of the clustering, h is relatively large. As the number of iterations grow larger, h gets very small (so that the movement at each iteration gets very slight).




Figure 8.4: Which mapping is topology preserving? Figure 1 shows a mapping that is not topology preserving, since neighboring cells are mapped to non-neighboring points in \mathbb{R}^2 . On the other hand, the middle clustering is not topology preserving, because neighboring points in \mathbb{R}^2 are mapped to non-neighboring cells in the topology. Only the third picture shows a topology preserving clustering.

This method would work very well, but it might result in a lot of empty cluster centers (centers that are never winners). This is where the concept of the “gas” comes into play- We think of the data as being in a viscous fluid, like jello.

How does this work in practice? As before, we first select a data point at random, and find the winning cluster center. We move the cluster center towards the data point. In a viscous fluid, that means that the cluster centers close to the winner will ALSO move toward the data point (just not as much). Therefore, instead of updating the single winning center using Equation 8.2, we update all centers.

The key point is that the winning center is attracted to \mathbf{x} with the strongest force, then the next closest center is attracted with the next strongest force, and so on. To do this, we’ll need to sort the centers by their distance to our data point \mathbf{x} .

You might have a negative reaction at this point- Given a lot of cluster centers, this is a lot of sorting. What is typical is that we will define a value k , and we will only update the k closest centers.

We define a way of setting this measure: We define s_i to be the number of centers closer to \mathbf{x} than $\mathbf{c}^{(i)}$. The easiest way to compute this is to sort the centers by their distance to \mathbf{x} , then put the first k indices in a vector \mathbf{v} . So, for example, the vector \mathbf{v} can be defined as:

$$\mathbf{v} = \{i_1, i_2, \dots, i_k\}$$

so the “winning” center has index i_i . Now we define the values s :

$$s_{i_1} = 0, \quad s_{i_2} = 1, \quad s_{i_3} = 2, \quad \dots, \quad s_{i_k} = k - 1$$

Example: Let $C = 0.1, 0.2, 0.4, 0.5$. If $x = 0.25$, then $i_1 = 2$, and $\mathbf{v} = \{2, 1, 3, 4\}$. The values of s : $s_2 = 0$, $s_1 = 1$, $s_3 = 2$, $s_4 = 3$.

Building Connections

Now that we have set up an algorithm that moves the centers into the data, we look at the problem of defining connections. The underlying idea is that we’ll connect a winning center to its neighbor, indexed by $\mathbf{v}(2)$. If that connection has not been selected for a long time, that probably means that over time, the centers have drifted apart, and we should remove the connection.

Therefore, to build connections, we will use two arrays (if the number of centers is k , they are both $k \times k$), M and T .

Definition: A Connection Matrix, M , is a matrix of 1’s and 0’s where

$$M_{ij} = \begin{cases} 1 & \text{If cell } i \text{ connected to } j \\ 0 & \text{Otherwise} \end{cases}$$

To build the connection matrix, once \mathbf{v} is determined, then M_{i_1, i_2} is set to 1, and the age of the connection T_{i_1, i_2} is set to 0. All the other ages have 1 added to them.

Lastly, we should check to see how many iterations have passed since each edge was constructed. If too much time has passed (pre-set by some parameter), then remove those connections.

Parameter List

The number of parameters we’ll use is a little large, but coding them is no problem. Here is a list of them, together with some standard values, where N is the number of data points.

- Number of data points: N
- Relative maximum distance to move, if we update the winning center: ϵ

$$\epsilon_{initial} = 0.3 \quad \epsilon_{final} = 0.05$$

- Parameter λ also gives some control over how much to move the winner and neighbors.

$$\lambda_{initial} = 0.2N \quad \lambda_{final} = 0.01$$

- The maximum amount of time before a connection is removed is:

$$T_{initial}^m = 0.1N \quad T_{final}^m = 2N$$

- Maximum number of iterations: `tmax`= 200*N*.

The parameters that have an “initial” and “final” time are updated using the power rule. If the initial value is $\alpha_{initial}$ and the final value is α_{final} , then on iteration i ,

$$\alpha_i = \alpha_{initial} \left(\frac{\alpha_{final}}{\alpha_{initial}} \right)^{i/tmax} \quad (8.3)$$

The Neural Gas Algorithm:

1. Select a data point \mathbf{x} at random, and find the winner, $c^{(i_1)}$.
2. Compute \mathbf{v} by finding the next $k - 1$ closest centers.
3. Update the k centers:

$$c^{(i_k)} = c^{(i_k)} + \epsilon \exp\left(\frac{-s_{i_k}}{\lambda}\right) (\mathbf{x} - c^{(i_k)})$$

4. Update the Connection and Time Matrices: Set $M_{i_1, i_2} = 1$, and $T_{i_1, i_2} = 0$.
Age all connections by 1, $T_{j, k} = T_{j, k} + 1$ for all j, k .
5. Remove all old connections. Set $M_{j, k} = 0$ if $T_{j, k} \geq T^m$
6. Repeat.

8.3.1 Matlab and Neural Gas

Matlab has not yet implemented a version of the Neural Gas algorithm. We will construct a suite of programs below.

The suite of programs we will use will be: `NeuralGas01.m`, `initng.m`, `paramUpdate.m`, and `plotng.m`.

The program `initng` is a program that will set all of the training parameters, and initialize the centers. This program should be edited at the beginning of a problem. The code is given below. Note the presence of the data structure, which allows us to wrap the parameters and the data together.

```
function C=initng(X)
[m,n]=size(X);
C.NC=50;           %Number of clusters (Depends on the data)
C.lr=[0.3 0.05];  %initial, final epsilon
C.nbr= [0.2*n 0.01]; %initial, final lambda
                  % SEE TRAINING NOTE 1

C.tcon=[0.1*n 2*n]; %initial, final time (for connection matrix)
                  % SEE TRAINING NOTE 2

C.tmax=200*n;     %max number of iterations
```

```

                % Shouldn't need to change tmax.

% The next two parameters should be left alone for now:
C.epochs=1;          %number of epochs (each epoch runs tmax iterations,
                    %and resets the training parameters after each.)
C.tflag=1           %Training flag: 1=Use Connection, 0=Do not use

% Now initialize centers randomly from the data:

Id=randperm(n);
C.cen= X(:,Id(1:C.NC)); %Initialize the centers randomly from the data
C.M=zeros(C.NC,C.NC);  %Initialize the connection matrix

```

The Neural Gas code itself is a bit lengthy, but the comments are included below so you know how to call the function.

```

function C=NeuralGas01(X,C)
%FUNCTION C=NeuralGas01(X,C)
% This is the Neural Gas clustering algorithm. There are
% two ways to call this program:
% [C,M]=NeuralGas01(X);
% With only one input argument (X is dimension by
% number of points), the program will call the function
% initng.m for the initialization.
%
% [C,M]=NeuralGas01(X,C)
% With two input parameters, we assume the centers
% have been initialized in the structure C. See the
% initng file for structure specifications.

```

The function `paramUpdate` is simply a call to Equation (8.3). You won't need to change that.

The function `plotng01` is an example of how we can plot the edges of the graph that is produced using the algorithm. As presented, it only plots the two dimensional graph. You might look it over, but you won't need to change that either.

Training Notes: Neural Gas

To use the programs we've listed, first you need a set of data in X . Here's a really simple example of performing the Neural Gas algorithm in Matlab:

```

X=2*randn(2,300)-1; % NOTE: Data is dimension x num points!
C=NeuralGas01(X);

```

If you want to manually produce the final plot, you can do that with the following three commands:

```

plot(X(1,:),X(2,:), 'k. ');
plotng01(C);
plot(C.cen(1,:),C.cen(2,:), 'r*')
hold off

```

For practice, there are six sets of data online in the file `SixDatasets.mat`, plotted in Figure 8.5. The top left shows the first, and it involves concentric circles- A typical kind of "toy problem" that we use to test clustering routines. Data sets 1, 3, 5, 6 show distinct clusters on which we hopefully have connected centers.

Data set 2 has 300 data points to the left of zero, 100 data points to the right. A good clustering for that and data set 4 would be one in which the density of the data is somehow reflected in the density of the data.

To use these datasets, download the *mat* file from the class website, then use the `load` command. That loads in a “cell array” of data named `XX`. To get the data in a usable form, here’s how it would go in Matlab to load the smiley face data and try to cluster it:

```
load SixDatasets
X=XX{3};
C=NeuralGas01(X');
```

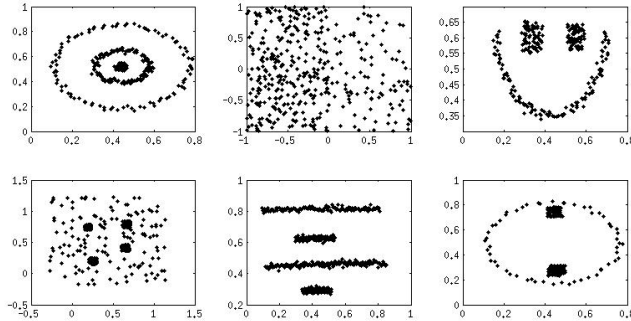


Figure 8.5: Six datasets with which to try using the Neural Gas clustering (going from left to right, top to bottom). Data sets 1, 3, 5, 6 show distinct clusters on which we hopefully have connected centers. Data set 2 has 300 data points to the left of zero, 100 data points to the right. A good clustering for that and data set 4 would be one in which the density of the data is somehow reflected in the density of the data.

Finally, a couple of training notes.

- Training Note 1: If ϵ and λ are too big, the cluster centers end up moving to the centroid of your data. If they are too small, the cluster centers won’t move into the data. Try changing these values to see if you can get the clusters to converge in the middle.
- Training Note 2: If the age at which we remove connections is too small, we end up with a lot of small, disconnected groups of centers. If the age is too big, you end up with a lot of redundant connections. What is just right? Try changing those values to see if you can see what we’re referring to.

8.3.2 Project: Neural Gas

This project explores one application of triangulating a data set: Obtaining a path preserving representation of a data set.

For example, suppose we have a set of points in \mathbb{R}^2 that represents a room. We have a robot that can traverse the room - the problem is, there are obstacles in its path. We would like a discretization of the free space so that the robot can plan the best way to move around the obstacles.

The data set is given in `obstacle1.mat`, and is plotted below. The areas where there are no points represent obstacles. **Assignment:** Use the Neural Gas clustering using 1,000 data points and 300 cluster centers to try to represent the obstacle-free area.




Figure 8.6: The data set that represents the "free" area. White regions indicate obstacles.