


Computer Lab: K-means Clustering on Digits

The data we're looking at represents an extraction of handwritten digits, where the data was preprocessed to an 8×8 grid of pixels.

Here's a sample of 10 random points in the set, each visualized in the 8×8 grid:



When we load the `handdigits.mat` file (download it first), we'll have a matrix that is 1797×64 , and this represents the 1797 data points in \mathbb{R}^{64} . The “targets” vector is something we'll use later, but it tells us which digit the image is supposed to be.

Your goal is to cluster the points into 10 clusters using k-means and show the results. Be sure to run k-means several times to get the best clustering.

At the end, we'll try looking at the 10 cluster centers to see if they represent all the digits! This is the code towards the end that you shouldn't have to use as long as your cluster centers are in a matrix C that was output from the `kmeans` command (should be 10×64).

Here's a partial Matlab file to get you started:

```
load handdigits

%
% Fill in the missing command, but be sure your output for the cluster centers
% is the 10 x 64 matrix C for the image commands below to work.
```

```
%  
  
% Show the 10 clusters as 8 x 8 images (You don't need to edit this)  
for j=1:10  
    subplot(2,5,j)  
    temp=reshape(-C(j,:),8,8);  
    imagesc(temp');  
    colormap(gray); axis equal; axis off  
end
```