

CFCS1: Assignment 1

January 21, 2010

Handed out:	22 January 2010
Due date:	1 February 2010

Introduction

For this assignment you will use Matlab to classify different sounds of human speech. From real recordings of speech, samples corresponding to three different vowel sounds have been extracted:

Number	Vowel name	Examples
1	ih	bit, hit
2	ae	bat, hat
3	uw	boot, fruit

Cognitive scientists have studied the way in which the human brain processes speech to distinguish between different sounds. For each sample, we have used this information to automatically extract **two numerical features** considered most relevant for distinguishing between vowel sounds. These features take the form of two-dimensional vectors. **For the purposes of this assignment, you do not need to understand how the vectors have been obtained.**

You are presented with 600 vectors, each corresponding to an example of a vowel sound. You will see that there is considerable variation amongst the vectors, even within a

single sound group: this is because the recordings have been taken from a range of speakers of different ages, different sexes, and with different accents. Your task will be to group the vectors into the three different sounds, using vector distance measures (Euclidean distance, Maximum-value distance and Cosine distance). The correct group for each vector has been given to you so that you can see how well each method works – but this information should not be used to “cheat” with the grouping process!

Getting started

Save the file `a1_data.mat` to your working directory. Open Matlab from your working directory, and load the data from this file using the command `load a1_data.mat`. This will create two variables:

- A matrix, `samples`. Each row `i`, obtained by `samples(i, :)`, is the 2-dimensional vector corresponding to the i th sample. For example, type `samples(2, :)` in the command line, you will get the 2nd sample: 0.5831 -0.1509.
- A column vector, `labels`. `labels(i)` gives the correct group of the i th sample, expressed as a number as in the table above. For example, type `labels(2)` in the command line, you will see the correct group label of the 2nd sample is ‘ 1 ’.

1 Displaying the vectors

1. In order to visualise the data you have been given, use Matlab to display the data in a scatter plot (using the `plot` command). Use the information in `labels` to display vectors from the three groups in different colours, on the same plot. Use red for group 1, green for group 2 and blue for group 3.

[5]

2. From your plot, which of the two numerical features looks to be the best for separating the data?

[2]

[*Hint*: the third argument to the `plot` command controls the type and colour of the plot. `'r*'` produces a red scatter plot, `'b*'` produces a blue scatter plot etc. For more information, type `help('plot')`. To check that your plot is correct: you should see that the data from group 1 forms a long, thin cluster and the data from group 2 forms a small cluster, whilst the data for group 3 is much more dispersed]

2 Grouping the vectors

Further work by cognitive scientists has determined typical vectors for each sound group. They are given by:

```
m1 = [0.496, 0.164]
m2 = [0.395, 0.201]
m3 = [0.306, 0.253]
```

3. Display these vectors as additional points on your scatter plot, using black squares.

[2]

You will group each sample vector as 1,2 or 3, according to which of the three typical vectors it is **closest to**, using various distance measures.

4. Write Matlab functions to group each sample vector, using:
 - (a) Standard (Euclidean) distance
 - (b) Maximum-value distance (the maximum-value distance between vector **a** and **b** is simply the magnitude of the largest element in vector **a-b**)
 - (c) Cosine distance

You will need to compare the groups chosen by your functions with the correct group supplied in `labels`. For each of the distance measures (a)-(c), produce a count of how many vectors were correctly grouped.

[12]

5. Give **two** possible reasons why these vector distance measures do not seem to work particularly well at grouping these samples.

[4]

To submit

For the assignment, you should submit:

- Your completed scatter plot.
- Written answers to the required questions.
- Printouts of all the Matlab functions you wrote.
- A (brief) transcription of your Matlab session, giving the commands you typed, together with the output that you used to answer the questions.

Please hand in a hard copies of the required material by 4:00pm on the due date to the Informatics Teaching Organisation, Level 4, Appleton Tower. If you have questions regarding the assignment, please contact Yansong Feng at yansong.feng@ed.ac.uk.