

Lab 6

CSC 5930/9010 - Computer Vision

Description: In this lab you will practice KNN and SVM classification in Matlab. We will be using a subset of the CIFAR-10 dataset.

The CIFAR-10 dataset consists of 60000 32x32 colour images in 10 classes, with 6000 images per class. There are 50000 training images and 10000 test images.

The dataset is divided into five training batches and one test batch, each with 10000 images. The test batch contains exactly 1000 randomly-selected images from each class. The training batches contain the remaining images in random order, but some training batches may contain more images from one class than another. Between them, the training batches contain exactly 5000 images from each class.

Here are the classes in the dataset, as well as 10 random images from each:

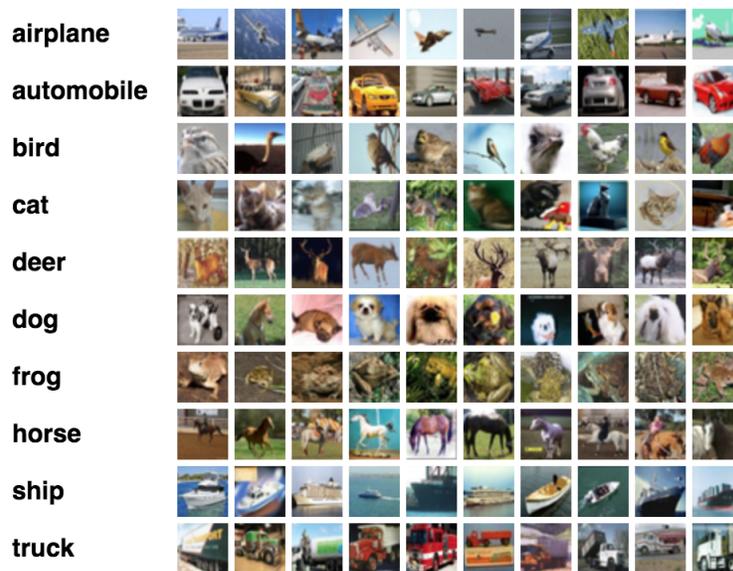


Figure 1: The classes are completely mutually exclusive. There is no overlap between automobiles and trucks. “Automobile” includes sedans, SUVs, things of that sort. “Truck” includes only big trucks. Neither includes pickup trucks. Each row of the array stores a 32x32 colour image. The first 1024 entries contain the red channel values, the next 1024 the green, and the final 1024 the blue. The image is stored in row-major order, so that the first 32 entries of the array are the red channel values of the first row of the image. labels – a list of 30000 numbers in the range 0-9. The number at index i indicates the label of the i th image in the array data.

Step 1 - Download the following MAT files from the website. “data30k.mat”, “labels30k.mat”, and “batches.meta.mat”. Explore the format of the data you have been given by double clicking the

mat file or by `load('filename.mat')`. Display some images and verify that they match the correct labels as shown in `label_names`.

You can display the first image by,

```
im1 = data(1,:);  
im1 = reshape(im1, [32 32 3]);  
imshow(imrotate(im1,-90));
```

Step 2 - We are going to use the Euclidean distance to perform a KNN match to a test image. Create a file called `knntest.m`. Take the first image and find the distance from that image to the first 100 images in the dataset (in the first 1000, 10000?). What is the closest match (K=1)? Recall that the euclidean distance is given by,

$$d(x, y) = \sqrt{(x_1 - y_1)^2 + (x_2 - y_2)^2 + \dots + (x_n - y_n)^2} \quad (1)$$

Step 3 - Once you have successfully implemented the KNN classifier we can build a multiclass SVM classifier using the machine learning tools built into matlab. First we define a support vector machine template,

```
t = templateSVM('Standardize',1)
```

Standardize indicates we will be centering the data (subtracting the mean as we did in PCA). Next, we will train the SVMs on the first 1000 data points.

```
Mdl = fitcecoc(double(data(1:1000,:)),labels(1:1000),'Learners',t,'Verbose',2)
```

This is taking the template `t` and the first 1000 data points with labels and training 45 multi-way SVMs. Why is it 45?

You can use the model, `Mdl` to make predictions about new images.

```
[label,score,~]=predict(Mdl,double(data(29012,:)));
```

This will take the 29,012 image and assign it a label and a score. How well is your SVM doing? If you would like a different model, I have trained a `Mdl.mat` on 29,000 images (1-2 hours of training time). You can experiment with this and see if it does better.