

Assignment 2

IFT 6268, Fall 2013

Date posted: October 9, 2013

Due: October 16, 2013, at the beginning of class

1. (1/10) Show what happens to the covariance matrix, C , of a random variable \mathbf{x} , if we multiply \mathbf{x} by a matrix M . Use this to show that the covariance matrix of the ICA generative model with generative weights A is given by AA^T .
2. (2/10) The inverse of the ZCA transform is also called the *square-root of the data covariance matrix*. Explain why.
3. (3/10) Download the file

`www.iro.umontreal.ca/~memisevr/teaching/ift6268_2013/cifarmini_gray_images_train.txt`

and save it locally (this is the same file that you used in Assignment 1).

Crop patches (at least 1000), each of size 10×10 pixels, from random positions in the images. Make sure to crop patches from multiple different images (e.g. one randomly positioned patch per image).

DC center and contrast normalize each patch. Display some of the resulting patches as gray value images.

Compute the PCA whitening matrix and the ZCA whitening matrix of the (vectorized) patches. Display the 100 learned PCA and ZCA basis images as gray value images, sorted according to their corresponding eigenvalues.

Hints:

- As in the previous assignment, you may display a stack of filters using
`http://www.iro.umontreal.ca/~memisevr/teaching/ift6268_2013/dispims.py`
- Be sure to represent your data as floats, otherwise strange things will happen.
- You may get some very small eigenvalues. Make sure to avoid any numerical problems, for example, by adding a small constant when dividing.

What to hand in: Three images in total: One image showing example patches, one image showing the 100 PCA filters, one image showing the 100 ZCA filters. Do *not* hand in any program code.

4. (4/10) In this question, we will implement a very simple recognition pipeline using local features and your logistic regression code from the previous assignment. We will make use of an encoding scheme proposed by (Coates, Ng; ICML 2011). Download the files:

```
www.iro.umontreal.ca/~memisevr/teaching/ift6268_2013/cifarmini_gray_images_train.txt
www.iro.umontreal.ca/~memisevr/teaching/ift6268_2013/cifarmini_labels_train.txt
www.iro.umontreal.ca/~memisevr/teaching/ift6268_2013/cifarmini_gray_images_test.txt
www.iro.umontreal.ca/~memisevr/teaching/ift6268_2013/cifarmini_labels_test.txt
```

and save them locally. They are the same files as in Assignment 1.

To learn a recognition model from these images, follow these steps:

- (a) For each image (train and test), crop patches from all positions within the boundaries of the image. (You should get 484 patches per image, or $484 \times 4000 = 1936000$ patches in total).
- (b) DC center and contrast normalize each patch.
- (c) Compute the PCA whitening matrix from a subset of the *training* patches (you may want to use your code from the previous question). Do *not* use any test image patches for this!
- (d) Whiten all patches using your PCA matrix.
- (e) Designate 100 randomly chosen, whitened *training* patches as “prototypes”, which we call $\mathbf{f}_1, \dots, \mathbf{f}_{100}$ in the following. You may want to save them off in a separate matrix.
- (f) Transform each of your 4000 images into a feature vector as follows: For each image I , and for each patch \mathbf{x} within the image, compute a 100-dimensional feature vector, whose i^{th} component is given by:

$$\max(0, \mathbf{f}_i^T \mathbf{x})$$

Then represent each image, I , as the *mean* over its 484 feature vectors.

- (g) Each image is now represented by a 100-dimensional (mean) feature vector.
- (h) Train your *logistic regression* model using weightcost 0.0 on the *training* feature vectors, then test it on the test feature vectors.

Repeat the whole experiment using the ZCA whitening matrix instead of the PCA matrix. Then repeat the whole experiment using *no* whitening (simply replace your whitening matrix by the identity).

What to hand in:

- An image showing three training examples from each class.
- For each of the three runs (PCA, ZCA, no whitening), a plot of the *average* 100-dimensional image feature vector, averaged over the whole training set. (3 plots in total)
- Your classification error rate on the training and on the test data for (i) PCA whitening, (ii) ZCA whitening, (iii) no whitening (6 numbers in total).
- Do *not* hand in any program code.