

Chapter 14

Night 7: PCA and Eigenfaces

🔗 Learning Objectives

Concepts

- Understand the connection between PCA and eigenvalues and eigenvectors.
- Understand how to use PCA to carry out data compression.
- Understand the idea of using Eigenfaces to do facial recognition.

MATLAB skills

- Use “eig” to carry out a PCA.
- Implement facial recognition using PCA.
- Determine the accuracy of PCA for different numbers of principal components.

14.1 Principal Component Analysis Revisited

As we in class, PCA is an algorithm in which we express our original data along the eigenvectors corresponding to the largest eigenvalues of the covariance matrix. We examined the property of PCA that if we project our data onto these vectors, this will lead to maximizing the variance of the projected data. To refresh your memory further, here is the temperature plot for Boston, Sao Paolo, and Washington DC.




Figure 14.1: Temperatures in three cities and the eigenvectors of the covariance matrix.

We also briefly mentioned a second property of PCA, which is that it can be thought of as an optimal way to compress our data down to a smaller set of numbers. This idea, also known as dimensionality reduction, is going to be a view that we explore in this assignment. If you'd like, here are a few external resources on PCA:

- <http://www.cs.otago.ac.nz/>
- <http://dai.fmph.uniba.sk/courses/ml/sl/PCA.pdf>
- <https://deeplearning4j.org/eigenvector/linear>
- <http://www.cerebralmastication.com/2010/09/principal-component-analysis-pca-vs-ordinary-least-squares-ols-a-visual-explanation/>
- <http://stats.stackexchange.com/questions/2691/making-sense-of-principal-component-analysis-eigenvectors-eigenvalues/>

PCA in two dimensions

In general, PCA is conducted on data that is mean-centered (i.e., the data has had the mean of each variable subtracted out). To refresh your memory of PCA and scaffold the introduction of the view of PCA as compressing a dataset, let's think about a simple example data set \mathbf{D} .

$$\mathbf{D} = \begin{bmatrix} -1 & 3 \\ 1 & 4 \\ 3 & 4 \\ 7 & 5 \\ 10 & 9 \end{bmatrix} \quad (14.1)$$

Exercise 14.1

1. Create a plot of \mathbf{D} as a set of points in the xy -plane.
2. Define a matrix $\tilde{\mathbf{D}}$ which is the mean-centered version of \mathbf{D} and plot $\tilde{\mathbf{D}}$ as a set of points in the xy -plane
3. The principal components (\mathbf{p}_1 and \mathbf{p}_2) are the eigenvectors of the covariance matrix of the mean-centered $\tilde{\mathbf{D}}$. Compute \mathbf{p}_1 and \mathbf{p}_2 and plot them on top of the mean-centered data.
4. Compute the projection of your data onto the eigenvector which corresponds to the largest eigenvector, which in this case is \mathbf{p}_2 . This is the “reduced dimensionality” version of your data, called \mathbf{B} , which only include information about the projection along \mathbf{p}_2 . (We reduced the 2-dimensional data to 1-dimensional data.) Plot the original data \mathbf{D} and the reduced data \mathbf{B} .

Exercise 14.2

1. Can you recreate \mathbf{D} perfectly from \mathbf{B} ?
2. What would have happened if you had created \mathbf{B} using only information about the values along \mathbf{p}_1 instead of \mathbf{p}_2 ?
3. How might you quantify how well you can represent \mathbf{D} in this reduced dimensionality form?
4. If you received a new piece of data, how would you go about representing this as a linear combination of \mathbf{p}_1 and \mathbf{p}_2 ?

Data Compression via PCA

In this exercise you will perform a simple data compression exercise, similar to the one you did in a previous Night assignment. You will use temperature data from 3 cities over 10 years, as training data and use it to compress a year's worth of temperature data from 3 cities into a 2×365 matrix. In other words, you will represent 3×365 numbers (daily temperature data from 3 cities over 1 year), using 2×365 values. This compression is lossy, in that you will lose some information. However, by representing the data along the two most significant eigenvectors of the covariance matrix, you can reduce this data loss, because these two directions capture the bulk of the variation in the data set. Please note that while we have laid out the steps you need to take here quite explicitly, it is important for you to fully understand what each step does. You will be using very similar steps in your project.

*Exercises***Exercise 14.3**

1. Load the file `avg_temperatures_pt2.mat`. You will have 6 data vectors in your workspace. `b_tr`, `w_tr`, `s_tr` which represent 10 years of training data for the average daily temperatures in Boston, Washington DC, and Sao Paolo, respectively. The vectors `b_new`, `w_new`, `s_new` represent an additional year of data for the three cities – this is the data that you will compress using statistical knowledge of the previous 10 years of data. Create a covariance matrix \mathbf{R} using the 10 years worth of temperature data from Boston, Washington DC and Sao Paolo (in that order).
2. Perform an eigendecomposition of the matrix \mathbf{R} , and make a new matrix \mathbf{V}_p which has the 2 eigenvectors corresponding to the 2 largest eigenvalues of \mathbf{R} . You should use MATLAB's

`eig` function. Let these eigenvectors be \mathbf{v}_1 and \mathbf{v}_2 .

3. Create centered (i.e. subtract the mean), versions of the new temperature data vectors, and create a 3×365 matrix \mathbf{T} which has the centered temperatures of Boston, Washington DC and Sao Paolo as its rows (in that order). This matrix is a representation of the data you are now going to compress. Let the i -th column of \mathbf{T} be \mathbf{t}_i .
4. Take the dot product of each column of the matrix \mathbf{T} (which is a vector of the temperature of Boston, Washington DC and Sao Paolo for a given day) with the two eigenvectors in matrix \mathbf{V}_p , and save the values. Let this quantities be called α_{1i} and α_{2i} . In other words,

$$\alpha_{1i} = \mathbf{v}_1^T \mathbf{t}_i$$

$$\alpha_{2i} = \mathbf{v}_2^T \mathbf{t}_i$$

You can do this using matrix multiplications.

You should now have 365 different values for α_{1i} and α_{2i} , which are a compressed representation of 3×365 different temperature values. Moreover, these values are the components of the temperature data that lie in the directions of the two eigenvectors of the covariance matrix corresponding to the largest eigenvalues. From what we saw in the previous two classes, these vectors represent the two orthogonal directions in the data that have the most amount of variation, and hence the most "important" directions. Of course, there is a third direction (since the temperature vectors live in a 3-dimensional space), which we are discarding. But since this is the direction in which there is the least amount of variation in the data set, we do not lose too much information.

5. You can now check how well your compression worked, by using the values of α_{1i} and α_{2i} to reconstruct 365 different 3×1 vectors each representing the temperatures for the three cities over the 365 days. Let $\hat{\mathbf{t}}_i$ represent the reconstructed temperature vector on the i -th day. Using what you know about projections onto orthonormal vectors, reconstruct \mathbf{t}_i using α_{1i} , α_{2i} , \mathbf{v}_1 and \mathbf{v}_2 . Repeat this for all 365 days.
6. On the same axes, plot the original and reconstructed temperature for Boston. Repeat this for Washington DC and Sao Paolo. Observe how close the reconstructions are, for the different data sets.
7. How accurately do you think you can represent the data if you used 3 eigenvectors instead of 2?
8. If you feel inspired, repeat the above with temperature data for four different cities, and 2 or 3 different eigenvectors.

While this example can be thought of as a "toy" example where we are representing 3 dimensional data using 2 dimensions, there are many applications for which there may be many more dimensions in the data

for which accurate representations can be made using only a few dimensions. Additionally, you should note that such dimensionality reduction techniques are not just useful in compression, but they are also useful in speeding up computation. We can often get away with analyzing data over a small number of important dimensions, and this is an important technique when we deal with large amounts of data. Overall, these class of techniques is called Principal Component Analysis (PCA), since we are performing analysis along a few principal component directions of the data.

14.2 Face Data Compression via PCA

You are now ready to start applying PCA to face data. You have already seen this in a previous class assignment, except in that assignment you had the help of a genie. Load the MATLAB files `classdata_train.mat` and `classdata_test.mat`. These are the training and test datasets with photos of your classmates. The file contains some images and as well as the identity of the person in each image (coded as an integer from 1 to 89).

Remember that the principal eigenvectors of the covariance matrix tell you the directions of greatest variation in a data set and also the directions that optimally compress our data. Your job is to use the training images to build a model for your faces such that you can compress a many-pixelated test face image (pick one from the test image array) using a small set of image vectors (e.g., 10, 20, or 50).

Before you dig into this problem, think through how you would formalize the face data compression as a problem that you can solve with the linear algebra techniques that you've learned so far. There is no exercise to answer, but we want you to think through these steps before going further in the assignment.

- How you will choose your set of image vectors?
- Come up with the steps needed to do the above and write some pseudo code (e.g., load the data, vectorize the images, etc.).
- How could you tell whether your compression algorithm works (these could either be quantitative metrics, like root-mean squared error or qualitative metrics).

14.3 Eigenfaces for Face Recognition

It's time to bring it all together and finally plunge into the prosopagnosia (look it up) problem (or at least build some facial recognition software, but alliteration is fun). You will implement the eigenfaces algorithm to identify photos of your classmates. While it sounds fancy, you have almost all of the pieces needed to understand and implement Eigenfaces (the last necessary piece you will pick up momentarily). Here are the major steps in the Eigenfaces algorithm.

1. Use PCA to compute the k principal components of the training face images (the k eigenvectors with largest eigenvalues).
2. Project the training and test face images onto the k principal components. We'll call this the *facespace* representation of our original images.
3. For each of the test images, compute the closest match between the test image (represented as a k -dimensional vector in facespace) and the training images (again, in facespace). The notion of "closest match" here can be described in a few different ways, but the easiest thing to do is to use

the the Euclidean distance to define how far apart two points are. In this way, you would look for the training point that has the smallest Euclidean distance for a particular test point and predict the identity of the test point to be the same as the identity of this closest training point. This method of classification is known as **nearest neighbor classification** and it is the one new concept you need to implement Eigenfaces.

Exercise 14.4

Earlier in this assignment you wrote some pseudo code for face compression, which as you can see from the description of Eigenfaces above, gets you most of the way there. Before you actually implement Eigenfaces, we'd like you to extend your pseudocode to cover the whole Eigenfaces algorithm. In addition to the steps of Eigenfaces describe above, you should also think about the steps needed to calculate the accuracy of your system (i.e., how often does it get the person's identity correct).

Exercise 14.5

Implement the eigenfaces algorithm.

1. Your code, which can be a script, function, or livescript, should use the training and test sets of images provided: `classdata_train.mat` and `classdata_test.mat`, respectively.
2. You may want to start by identifying one face from the test set, but by the time you are done, your code should run through all of the test images and report the fraction it guesses correctly.
3. Test different numbers of eigenvectors. How many does it take to guess right most of the time?
4. If you'd like, time how long it takes your code to run. Can you do anything more efficiently to make it run faster? (use `tic` and `toc` in MATLAB for timing)
5. Visualize the first few eigenfaces. Can you interpret what they mean?
6. Generate a figure that depicts the success rate (accuracy at determining the identify of a person in an image) versus the number of eigenfaces used.
7. Generate a figure that depicts the success rate (accuracy at determining the identify of a person in an image) versus the number of eigenfaces used where the training data consists of only images of people not smiling and the test data consists of only images of people smiling (these are located in `classdata_non_smiles.mat` and `classdata_smiles.mat` respectively. Comment on the difference in performance on when using these files versus the data from the previous parts of this problem.

Guidelines:

- You should comment your code (use %) so others could read and understand it.
- Don't use the command `pca`, but instead build your algorithm using either the `eig` or `eigs` command (`eigs` computes just a few eigenvectors, which can be faster when you only care about the eigenvectors with large eigenvalues). We want you to think through all the steps involved in your facial recognition program, and that means doing the math "yourself".


(Optional) Extensions (These are not spelled out in much detail. We recommend you talk to a member of the teaching team before trying these (especially the second two)).

- Analyze the mistakes your algorithm makes (particularly when training on non-smiles and testing on smiles).
- Use Eigenfaces to do smile detection instead of identity recognition.
- Combine Eigenfaces with a classifier other than nearest neighbors (e.g., formulate an LSAE to create a series of one person versus everyone else detectors).
- Get your system working on live video.

Solution 14.1


1.

```
>> D = [-1 3; 1 4; 3 4; 7 5; 10 9]
>> plot(D(:,1),D(:,2), 'x')
>> axis ([-3 12 -3 12])
```



2.

```
>> mumatrix=[mean(D(:,1))*ones(5,1) mean(D(:,2))*ones(5,1)]
>> tildeD=D-mumatrix
>> plot(tildeD(:,1),tildeD(:,2), 'x')
>> axis ([-7 7 -7 7])
```




3.


```
>> [Vec,Diam]=eig(tildeD'*tildeD)
```



```
>> hold on
>> quiver(0,0,Vec(1,1),Vec(2,1))
>> quiver(0,0,Vec(1,2),Vec(2,2))
```



```
4. >> proj=tildeD*Vec(:,2)
>> B=proj*Vec(:,2) '
>> hold on
>> plot(tildeD(:,1),tildeD(:,2),'x')
>> plot(B(:,1),B(:,2),'o')
>> axis ([-6 8 -6 8])
```



Solution 14.2

1. No. If we write a data point $a\mathbf{p}_1 + b\mathbf{p}_2$, then the reduced dimension version is $b\mathbf{p}_2$. It's impossible to recover a , which is the information in the perpendicular direction.
2. We would get the information in the perpendicular direction, which we can interpret as the "error" in reducing the dimension from \mathbf{D} to \mathbf{B} .
3. You can use the error $\mathbf{B} - \mathbf{D}$.
4. We can write a new data point \mathbf{d} as

$$(\mathbf{d} \cdot \mathbf{p}_1)\mathbf{p}_1 + (\mathbf{d} \cdot \mathbf{p}_2)\mathbf{p}_2.$$

Solution 14.3

1.

```
>> A = (1/sqrt(7304))*[b_tr-mean(b_tr) w_tr-mean(w_tr) s_tr-mean(s_tr)];  
>> R=A'*A
```
2.

```
>> [V,D]=eig(R)  
>> Vp=[V(:,2) V(:,3)]
```
3.

```
>> T = [b_new-mean(b_new) w_new-mean(w_new) s_new-mean(s_new)]'
```
4.

```
>> alpha=Vp'*T;
```

Solution 14.5

A reference implementation of Eigenfaces is linked from the Canvas assignment page.