Machine Learning
Exercise 8

Marc Toussaint
TAs: Janik Hager, Philipp Kratzer
Machine Learning & Robotics lab, U Stuttgart
Universitätsstraße 38, 70569 Stuttgart, Germany

June 6, 2019

(DS BSc students should nominally achieve 8 Pts on this sheet.)

1 PCA derived (6 Points)

For data $D = \{x_i\}_{i=1}^n$, $x_i \in \mathbb{R}^d$, we introduced PCA as a method that finds lower-dimensional representations $z_i \in \mathbb{R}^p$ of each data point such that $x_i \approx V z_i + \mu$. PCA chooses V, μ and z_i to minimize the reproduction error

$$\sum_{i=1}^n \|x_i - (V z_i + \mu)\|^2 \ .$$

We derive the solution here step by step.

a) Find the optimal latent representation z_i of a data point x_i as a function of V and μ. (1P)

b) Find *an* optimal offset μ. (1P)

(Hint: there is a whole subspace of solutions to this problem. Verify that your solution is consistent with (i.e. contains) $\mu = \frac{1}{n} \sum_i x_i$).

c) Find optimal projection vectors $\{v_i\}_{i=1}^p$, which make up the projection matrix

$$V = \begin{bmatrix} | & & | \\ v_1 & \dots & v_p \\ | & & | \end{bmatrix} \tag{1}$$

(2P)

Guide:

- Given a projection V, any vector can be split in orthogonal components which belong to the projected subspace and its complement (which we call W). $x = VV^\top x + WW^\top x$.
- For simplicity, let us work with the centered datapoints $\tilde{x}_i = x_i - \mu$.
- The optimal projection V is the one which minimizes the discarded components $WW^\top \tilde{x}_i$.

$$\widehat{V} = \underset{V}{\text{argmin}} \sum_{i=1}^n \|WW^\top \tilde{x}_i\|^2 = \sum_{i=1}^n \|\tilde{x}_i - VV^\top \tilde{x}_i\|^2 \tag{2}$$

- Don't try to solve computing gradients and setting them to zero. Instead, use the fact that $VV^\top = \sum_{i=1}^p v_i v_i^\top$, and the singular value decomposition of $\sum_{i=1}^n \tilde{x}_i \tilde{x}_i^\top = \tilde{X}^\top \tilde{X} = EDE^T$.

d) In the above, is the orthonormality of V an essential assumption? (1P)

e) Prove that you can compute the V also from the SVD of X (instead of $X^\top X$). (1P)

2 PCA and reconstruction on the Yale face database (5 Points)

On the webpage find and download the Yale face database http://ipvs.informatik.uni-stuttgart.de/mlr/marc/teaching/data/yalefaces.tgz. (Optionally use `yalefaces_cropBackground.tgz`, which is slightly cleaned version of the same dataset). The file contains gif images of 165 faces.

a) Write a routine to load all images into a big data matrix $X \in \mathbb{R}^{165\times 77760}$, where each row contains a gray image. In Octave, images can easily be read using `I=imread("subject01.gif");` and `imagesc(I);`. You can loop over files using `files=dir(".");` and `files(:).name`. Python tips:

```
import matplotlib.pyplot as plt
import scipy as sp
plt.imshow(plt.imread(file_name))
```

b) Compute the mean face $\mu = \frac{1}{n}\sum_i x_i$ and center the whole data matrix, $\tilde{X} = X - \mathbf{1}_n \mu^\top$.

c) Compute the singular value decomposition $\tilde{X} = UDV^\top$ for the centered data matrix.

In Octave/Matlab, use `[U, S, V] = svd(X, "econ")`, where the `"econ"` ensures you don't run out of memory. In python, use

```
import scipy.sparse.linalg as sla
u, s, vt = sla.svds(X, k=num_eigenvalues)
```

d) Find the p-dimensional representations $Z = \tilde{X} V_p$, where $V_p \in \mathbb{R}^{77760\times p}$ contains only the first p columns of V (Depending on which language / library you use, verify that the eigenvectors are returned in eigenvalue-descending order, otherwise you'll have to find the correct eigenvectors manually). Assume $p = 60$. The rows of Z represent each face as a p-dimensional vector, instead of a 77760-dimensional image.

e) Reconstruct all faces by computing $X' = \mathbf{1}_n \mu^\top + \boldsymbol{Z} V_p^\top$ and display them; Do they look ok? Report the reconstruction error $\sum_{i=1}^n \|x_i - x_i'\|^2$.

Repeat for various PCA-dimensions $p = 5, 10, 15, 20 \ldots$.