

---

# Principal Component Analysis and Autoencoders

---

**Shuiwang Ji**  
Texas A&M University  
College Station, TX 77843  
sjj@tamu.edu

**Zhengyang Wang**  
Texas A&M University  
College Station, TX 77843  
zhengyang.wang@tamu.edu

**Ying Wei**  
Texas A&M University  
College Station, TX 77843  
ying.wei@tamu.edu

## 1 Introduction

This introduction of PCA aims at providing a complete, self-contained, and easy-to-understand introduction of PCA. The only required background would be college-level linear algebra. This document is based on lecture notes by Shuiwang Ji at Texas A&M University and can be used for undergraduate and graduate level classes.

## 2 Basic Concepts

### 2.1 Orthogonal Matrix

An orthogonal matrix is a square matrix whose columns and rows are orthogonal unit vectors, *i.e.*, orthonormal vectors. That is, if a matrix  $Q$  is an orthogonal matrix, we have

$$Q^T Q = Q Q^T = I.$$

It leads to  $Q^{-1} = Q^T$ , which is a very useful property as it provides an easy way to compute the inverse.

For an orthogonal  $n \times n$  matrix  $Q = [q_1, q_2, \dots, q_n]$ , where  $q_i \in \mathbb{R}^n$ ,  $i = 1, 2, \dots, n$ , it is easy to see that  $q_i^T q_j = 0$  when  $i \neq j$  and  $q_i^T q_i = 1$ . Furthermore, suppose  $Q_1 = [q_1, q_2, \dots, q_i]$  and  $Q_2 = [q_{i+1}, q_{i+2}, \dots, q_n]$ , we have  $Q_1^T Q_1 = I$ ,  $Q_2^T Q_2 = I$ , but  $Q_1 Q_1^T \neq I$ ,  $Q_2 Q_2^T \neq I$ .

### 2.2 Eigen-Decomposition

A square  $n \times n$  matrix  $S$  with  $n$  linearly independent eigenvectors can be factorized as

$$S = Q \Lambda Q^{-1},$$

where  $Q$  is the square  $n \times n$  matrix whose columns are eigenvectors of  $S$ , and  $\Lambda$  is the diagonal matrix whose diagonal elements are the corresponding eigenvalues. Note that only diagonalizable matrices can be factorized in this way. If  $S$  is a symmetric matrix, its eigenvectors are orthogonal. Thus  $Q$  is an orthogonal matrix and we have

$$S = Q \Lambda Q^T.$$

## 3 Singular Value Decomposition

### 3.1 Definition

The singular value decomposition (SVD) of an  $m \times n$  **real** matrix (without loss of generality, we assume  $m \geq n$ ) can be written as

$$R = U \Sigma V^T,$$

where  $\mathbf{U}$  is an orthogonal  $m \times m$  matrix,  $\mathbf{V}$  is an orthogonal  $n \times n$  matrix, and  $\tilde{\Sigma}$  is a diagonal  $m \times n$  matrix with non-negative real values on diagonal. To be specific, we have

$$\mathbf{U}^T \mathbf{U} = \mathbf{U} \mathbf{U}^T = \mathbf{I}_{m \times m},$$

$$\mathbf{V}^T \mathbf{V} = \mathbf{V} \mathbf{V}^T = \mathbf{I}_{n \times n},$$

$$\tilde{\Sigma} = \begin{bmatrix} \Sigma_{n \times n} \\ \mathbf{0} \end{bmatrix}_{m \times n}, \Sigma_{n \times n} = \begin{bmatrix} \sigma_1 & 0 & 0 & \dots & 0 \\ 0 & \sigma_2 & 0 & \dots & 0 \\ 0 & 0 & \sigma_3 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & \sigma_n \end{bmatrix}, \quad (1)$$

where  $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_n \geq 0$  are known as singular values. If  $\text{rank}(\mathbf{R}) = r$  ( $r \leq n$ ), we have  $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_r > 0$  and  $\sigma_{r+1} = \sigma_{r+2} = \dots = \sigma_n = 0$ .

### 3.2 Relation to Eigen-Decomposition

The columns of  $\mathbf{U}$  (left-singular vectors) are orthonormal eigenvectors of  $\mathbf{R}\mathbf{R}^T$ , and the columns of  $\mathbf{V}$  (right-singular vectors) are orthonormal eigenvectors of  $\mathbf{R}^T\mathbf{R}$ . In other words, we have

$$\mathbf{R}\mathbf{R}^T = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^{-1},$$

$$\mathbf{R}^T\mathbf{R} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^{-1}.$$

It is easy to verify them as we have

$$\begin{aligned} \mathbf{R}^T\mathbf{R} &= (\mathbf{U} \begin{bmatrix} \Sigma \\ \mathbf{0} \end{bmatrix} \mathbf{V}^T)^T \mathbf{U} \begin{bmatrix} \Sigma \\ \mathbf{0} \end{bmatrix} \mathbf{V}^T = \mathbf{V} ([\Sigma \quad \mathbf{0}]) \begin{bmatrix} \Sigma \\ \mathbf{0} \end{bmatrix} \mathbf{V}^T = \mathbf{V} \Sigma^2 \mathbf{V}^T, \\ \mathbf{R}\mathbf{R}^T &= \mathbf{U} \begin{bmatrix} \Sigma \\ \mathbf{0} \end{bmatrix} \mathbf{V}^T (\mathbf{U} \begin{bmatrix} \Sigma \\ \mathbf{0} \end{bmatrix} \mathbf{V}^T)^T = \mathbf{U} \begin{bmatrix} \Sigma \\ \mathbf{0} \end{bmatrix} [\Sigma \quad \mathbf{0}] \mathbf{U}^T = \mathbf{U} \begin{bmatrix} \Sigma^2 & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \mathbf{U}^T, \end{aligned}$$

and  $\mathbf{V}^T = \mathbf{V}^{-1}$ ,  $\mathbf{U}^T = \mathbf{U}^{-1}$ .

Under what conditions are SVD and eigen-decomposition the same? First,  $\mathbf{R}$  is a symmetric matrix, *i.e.*,  $\mathbf{R} = \mathbf{R}^T$ . Second,  $\mathbf{R}$  is a positive semi-definite matrix, *i.e.*,  $\forall \mathbf{x} \in \mathbb{R}^n$ ,  $\mathbf{x}^T \mathbf{R} \mathbf{x} \geq 0$ .

The difference between  $\mathbf{\Lambda}$  in eigen-decomposition and  $\Sigma$  in SVD is that, the diagonal entries of  $\mathbf{\Lambda}$  can be negative, while the diagonal entries of  $\Sigma$  are non-negative. What are the fundamental reasons underlying this difference? Why the requirements on the singular values in SVD (non-negative and in sorted order) do not prevent the generality of SVD?

### 3.3 Compact SVD

If  $\text{rank}(\mathbf{R}) = r$  ( $r \leq n$ ), we have

$$\begin{aligned} \mathbf{R} &= \mathbf{U} \tilde{\Sigma} \mathbf{V}^T \\ &= [\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_r, \dots, \mathbf{u}_m] \begin{bmatrix} \sigma_1 & \dots & 0 & & & \\ \vdots & \ddots & \vdots & & \ddots & \\ 0 & \dots & \sigma_r & & & \\ & & & 0 & \dots & 0 \\ & & & \vdots & \ddots & \vdots \\ & & & 0 & \dots & 0 \end{bmatrix} \begin{bmatrix} \mathbf{v}_1^T \\ \mathbf{v}_2^T \\ \vdots \\ \mathbf{v}_r^T \\ \vdots \\ \mathbf{v}_n^T \end{bmatrix}. \end{aligned}$$

By removing zero components, we obtain

$$\begin{aligned}
 \mathbf{R} &= \mathbf{U}_r \boldsymbol{\Sigma}_r \mathbf{V}_r^T \\
 &= [\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_r] \begin{bmatrix} \sigma_1 & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & \sigma_r \end{bmatrix} \begin{bmatrix} \mathbf{v}_1^T \\ \mathbf{v}_2^T \\ \vdots \\ \mathbf{v}_r^T \end{bmatrix} \\
 &= [\sigma_1 \mathbf{u}_1, \sigma_2 \mathbf{u}_2, \dots, \sigma_r \mathbf{u}_r] \begin{bmatrix} \mathbf{v}_1^T \\ \mathbf{v}_2^T \\ \vdots \\ \mathbf{v}_r^T \end{bmatrix} \\
 &= \sum_{i=1}^r \sigma_i \mathbf{u}_i \mathbf{v}_i^T,
 \end{aligned}$$

where  $\text{rank}(\sigma_i \mathbf{u}_i \mathbf{v}_i^T) = 1$ ,  $i = 1, 2, \dots, r$ .

### 3.4 Truncated SVD and Best Low-Rank Approximation

We can also approximate the matrix  $\mathbf{R}$  with the  $k$  largest singular values as

$$\mathbf{R}_k = \mathbf{U}_k \boldsymbol{\Sigma}_k \mathbf{V}_k^T = \sum_{i=1}^k \sigma_i \mathbf{u}_i \mathbf{v}_i^T.$$

Apparently,  $\mathbf{R} \neq \mathbf{R}_k$  unless  $\text{rank}(\mathbf{R}) = k$ . This approximation is the best in following sense:

$$\begin{aligned}
 \min_{\mathbf{B}: \text{rank}(\mathbf{B}) \leq k} \|\mathbf{R} - \mathbf{B}\|_F &= \|\mathbf{R} - \mathbf{R}_k\|_F = \sqrt{\sum_{i=k+1}^n \sigma_i^2}, \\
 \min_{\mathbf{B}: \text{rank}(\mathbf{B}) \leq k} \|\mathbf{R} - \mathbf{B}\|_2 &= \|\mathbf{R} - \mathbf{R}_k\|_2 = \sigma_{k+1},
 \end{aligned}$$

where  $\|\cdot\|_F$  denotes the Frobenius norm and  $\|\cdot\|_2$  denotes the spectral norm, defined as the largest singular value of the matrix. That is,  $\mathbf{R}_k$  is the best rank- $k$  approximation to  $\mathbf{R}$  in terms of both the Frobenius norm and spectral norm. Note the difference in terms of approximation errors when different matrix norms are used.

## 4 Principal Component Analysis

### 4.1 What is PCA?

Principal Component Analysis (PCA) is a statistical procedure that can be used to achieve feature (dimensionality) reduction. Note, feature reduction is different from feature selection. After feature reduction, we still use all the features, while feature selection selects a subset of features to use. The goal of PCA is to project the high-dimensional features to a lower-dimensional space with maximal variance and minimum reconstruction error simultaneously. We derive PCA based on maximizing variance, and then we show the solution also minimizes reconstruction error. In machine learning, PCA is an unsupervised learning technique, and therefore does not need labels.

To introduce PCA, we start from the simple case where PCA projects the features to a 1-dimensional space. Formally, suppose we have  $n$   $p$ -dimensional ( $p > 1$ ) features  $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n \in \mathbb{R}^p$ . Let  $\mathbf{a} \in \mathbb{R}^p$  represent a projection that  $\mathbf{a}^T \mathbf{x}_i = z_i$ ,  $i = 1, 2, \dots, n$  where  $z_1, z_2, \dots, z_n \in \mathbb{R}^1$ . PCA aims to solve

$$\mathbf{a}^* = \arg \max_{\|\mathbf{a}\|=1} \frac{1}{n} \sum_{i=1}^n (z_i - \bar{z})^2.$$

Note that the variance of the reduced data is

$$\frac{1}{n} \sum_{i=1}^n (z_i - \bar{z})^2,$$

which means that PCA tries to find the projection with the maximum variance in reduced data.

Since

$$\bar{z} = \frac{1}{n} \sum_{i=1}^n z_i = \frac{1}{n} \sum_{i=1}^n \mathbf{a}^T \mathbf{x}_i = \mathbf{a}^T \left( \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i \right) = \mathbf{a}^T \bar{\mathbf{x}},$$

the problem can be written as

$$\begin{aligned} \mathbf{a}^* &= \arg \max_{\|\mathbf{a}\|=1} \frac{1}{n} \sum_{i=1}^n (z_i - \bar{z})^2 \\ &= \arg \max_{\|\mathbf{a}\|=1} \frac{1}{n} \sum_{i=1}^n (\mathbf{a}^T \mathbf{x}_i - \bar{z})^2 \\ &= \arg \max_{\|\mathbf{a}\|=1} \frac{1}{n} \sum_{i=1}^n (\mathbf{a}^T \mathbf{x}_i - \mathbf{a}^T \bar{\mathbf{x}})^2 \\ &= \arg \max_{\|\mathbf{a}\|=1} \frac{1}{n} \sum_{i=1}^n \mathbf{a}^T (\mathbf{x}_i - \bar{\mathbf{x}}) (\mathbf{x}_i - \bar{\mathbf{x}})^T \mathbf{a} \\ &= \arg \max_{\|\mathbf{a}\|=1} \mathbf{a}^T \underbrace{\left( \frac{1}{n} \sum_{i=1}^n (\mathbf{x}_i - \bar{\mathbf{x}}) (\mathbf{x}_i - \bar{\mathbf{x}})^T \right)}_{p \times p \text{ covariance matrix}} \mathbf{a} \\ &= \arg \max_{\|\mathbf{a}\|=1} \mathbf{a}^T \mathbf{C} \mathbf{a}, \end{aligned}$$

where

$$\mathbf{C} = \frac{1}{n} \sum_{i=1}^n (\mathbf{x}_i - \bar{\mathbf{x}}) (\mathbf{x}_i - \bar{\mathbf{x}})^T,$$

denotes the covariance matrix.

What if we want to project the features to a  $k$ -dimensional space? Then the PCA problem becomes

$$\mathbf{A}^* = \arg \max_{\mathbf{A} \in \mathbb{R}^{p \times k}: \mathbf{A}^T \mathbf{A} = \mathbf{I}_k} \text{trace} \left( \mathbf{A}^T \mathbf{C} \mathbf{A} \right), \quad (2)$$

where  $\mathbf{A} = [\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_k] \in \mathbb{R}^{p \times k}$ . Note that when projecting onto  $k$ -dimensional space, PCA requires different projection vectors to be orthogonal. Also, the trace above is the sum of the variances after projecting the data to each of the  $k$  directions as

$$\text{trace} \left( \mathbf{A}^T \mathbf{C} \mathbf{A} \right) = \sum_{i=1}^k (\mathbf{a}_i^T \mathbf{C} \mathbf{a}_i).$$

Solving the problem in Eqn. (2) requires the follow theorem.

**Theorem. (Ky Fan)** Let  $\mathbf{H} \in \mathbb{R}^{n \times n}$  be a symmetric matrix with eigenvalues

$$\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n,$$

and the corresponding eigenvectors  $\mathbf{U} = [\mathbf{u}_1, \dots, \mathbf{u}_n]$ . Then

$$\lambda_1 + \dots + \lambda_k = \max_{\mathbf{A} \in \mathbb{R}^{n \times k}: \mathbf{A}^T \mathbf{A} = \mathbf{I}_k} \text{trace} \left( \mathbf{A}^T \mathbf{H} \mathbf{A} \right).$$

And the optimal  $\mathbf{A}^*$  is given by  $\mathbf{A}^* = [\mathbf{u}_1, \dots, \mathbf{u}_k] \mathbf{Q}$  with  $\mathbf{Q}$  an arbitrary orthogonal matrix. ■

Note that in Eqn. (2), the covariance matrix  $\mathbf{C}$  is a symmetric matrix. Given the above theorem, we directly obtain

$$\begin{aligned} \lambda_1 + \dots + \lambda_k &= \arg \max_{\mathbf{A} \in \mathbb{R}^{n \times k}: \mathbf{A}^T \mathbf{A} = \mathbf{I}_k} \text{trace} \left( \mathbf{A}^T \mathbf{C} \mathbf{A} \right), \\ \mathbf{A}^* &= [\mathbf{u}_1, \dots, \mathbf{u}_k] \mathbf{Q}, \end{aligned}$$

where  $\lambda_1, \dots, \lambda_k$  are the  $k$  largest eigenvalues of the covariance matrix  $C$ , and the solution  $A^*$  is the matrix whose columns are corresponding eigenvectors.

It also follows from the above theorem that solutions to PCA are not unique, and they differ by an orthogonal matrix. We used the special case where  $Q = I$ , i.e.,  $A^* = [u_1, \dots, u_k]$ .

## 4.2 How to compute PCA?

Next, let's see how to compute the solutions above efficiently.

We define

$$\begin{aligned} X &= [x_1, x_2, \dots, x_n] \in \mathbb{R}^{p \times n}, \\ \bar{X} &= \frac{1}{n} X \mathbf{1}_n \in \mathbb{R}^{p \times 1}, \\ \tilde{X} &= (X - \bar{X} \mathbf{1}_n^T) \in \mathbb{R}^{p \times n}, \end{aligned}$$

where  $\mathbf{1}_n$  is the  $n$ -dimensional all-one vector. Here,  $\tilde{X}$  is the centered  $X$ , which is obtained by subtracting the mean from each column.

Then we have

$$C = \frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})(x_i - \bar{x})^T = \frac{1}{n} \tilde{X} \tilde{X}^T \in \mathbb{R}^{p \times p}.$$

If  $p$  is large, it is very costly to compute eigenvectors of  $C$  directly. Moreover, if  $p \gg n$ , there are known computational problems. However, in Section 3.2, we show that the SVD of  $\tilde{X}$  provides what we need. If the SVD of  $\tilde{X}$  is  $\tilde{X} = U \tilde{\Sigma} V^T$ , the columns of  $U$  (left-singular vectors) are orthonormal eigenvectors of  $\tilde{X} \tilde{X}^T$ . Note that we have  $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_n \geq 0$  in Eqn. (1) and thus, the first  $k$  columns of  $U$  correspond to the largest  $k$  eigenvalues of  $\tilde{X} \tilde{X}^T$ .

With the SVD of  $\tilde{X}$ , if we want to project features to a  $k$ -dimensional ( $k < p$ ) space, simply take the first  $k$  columns of  $U$  as the projection matrix. Generally, the process of computing PCA can be described as

$$\underbrace{X}_{p \times n} \rightarrow \underbrace{\tilde{X}}_{p \times n} \rightarrow \tilde{X} = U \tilde{\Sigma} V^T \rightarrow G = U_k \in \mathbb{R}^{p \times k}.$$

Then the projected features are

$$G^T \tilde{X} = Z \in \mathbb{R}^{k \times n}. \quad (3)$$

It is worth noting that the  $k$ -th row in  $Z$  is called the  $k$ -th principal components (PC). Therefore, PCA achieves feature reduction by keeping only the first  $k$  PCs.

Note that we project the centered data matrix  $\tilde{X}$  instead of the original data matrix  $X$  in Eqn (3). Maximal variance can be achieved in both cases, as the the covariance matrix  $C$  will not change. But the minimal reconstruction error can only be achieved when  $\tilde{X}$  is used, as shown in the next section. A common practice to avoid any confusion is to center the data before applying PCA, and use the centered data matrix in all computations.

## 4.3 Have We Achieved the Minimal Reconstruction Error?

Now let's verify our solution achieves the other goal of PCA, which is to project the high-dimensional features to a lower-dimensional space with minimum reconstruction error.

First, we perform the verification in the case where we project and reconstruct the centered data matrix  $\tilde{X}$ .

Given the projection

$$G^T \tilde{X} = Z,$$

the reconstruction process is

$$\tilde{X} = GZ = GG^T \tilde{X},$$

and the reconstruction error is

$$\|\tilde{X} - \tilde{X}\|_F = \|\tilde{X} - GG^T \tilde{X}\|_F.$$

We will show that

$$GG^T \tilde{\mathbf{X}} = \mathbf{B}^* = \arg \min_{\mathbf{B}: \text{rank}(\mathbf{B}) \leq k} \|\tilde{\mathbf{X}} - \mathbf{B}\|_F, \quad (4)$$

which means  $\tilde{\mathbf{X}}$  gives the minimum reconstruction error.

In Section 3.4, we've already learned that  $\mathbf{B}^*$  is the truncated SVD of  $\tilde{\mathbf{X}}$ . So, we only need to show that  $GG^T \tilde{\mathbf{X}}$  is indeed the truncated SVD of  $\tilde{\mathbf{X}}$  as follows. Note that similar arguments can be made for the spectral norm.

Given  $\tilde{\mathbf{X}} = \mathbf{U}\tilde{\Sigma}\mathbf{V}^T$  and  $\mathbf{G} = \mathbf{U}_k$ , we have

$$\begin{aligned} GG^T \tilde{\mathbf{X}} &= \mathbf{U}_k \mathbf{U}_k^T \mathbf{U} \tilde{\Sigma} \mathbf{V}^T \\ &= \mathbf{U}_k \mathbf{U}_k^T [\mathbf{U}_k \quad \mathbf{U}_{-k}] \tilde{\Sigma} \mathbf{V}^T \\ &= \mathbf{U}_k [\mathbf{I} \quad \mathbf{0}] \tilde{\Sigma} \mathbf{V}^T \\ &= [\mathbf{U}_k \quad \mathbf{0}] \tilde{\Sigma} \mathbf{V}^T \\ &= \mathbf{U}_k \Sigma_k \mathbf{V}_k^T \\ &= \sum_{i=1}^k \sigma_i \mathbf{u}_i \mathbf{v}_i^T. \end{aligned}$$

As a result,  $\tilde{\mathbf{X}}$  gives the minimum reconstruction error, which is

$$\min_{\mathbf{B}: \text{rank}(\mathbf{B}) \leq k} \|\tilde{\mathbf{X}} - \mathbf{B}\|_F = \|\tilde{\mathbf{X}} - GG^T \tilde{\mathbf{X}}\|_F = \sqrt{\sum_{i=k+1}^n \sigma_i^2}.$$

Similarly, it can be shown that

$$\min_{\mathbf{B}: \text{rank}(\mathbf{B}) \leq k} \|\tilde{\mathbf{X}} - \mathbf{B}\|_2 = \|\tilde{\mathbf{X}} - GG^T \tilde{\mathbf{X}}\|_2 = \sigma_{k+1}.$$

Therefore, PCA projects the high-dimensional features to a lower-dimensional space with minimum reconstruction error in terms of both Frobenius and spectral norms.

Next, we repeat the verification above in the case where we project and reconstruct the original data matrix  $\mathbf{X}$ .

Given the projection

$$\mathbf{G}^T \mathbf{X} = \mathbf{Z},$$

the reconstruction process is

$$\hat{\mathbf{X}} = \mathbf{G}\mathbf{Z} = \mathbf{G}\mathbf{G}^T \mathbf{X},$$

and the reconstruction error is

$$\|\mathbf{X} - \hat{\mathbf{X}}\|_F = \|\mathbf{X} - \mathbf{G}\mathbf{G}^T \mathbf{X}\|_F.$$

We note that

$$GG^T \mathbf{X} \neq \mathbf{B}^* = \arg \min_{\mathbf{B}: \text{rank}(\mathbf{B}) \leq k} \|\mathbf{X} - \mathbf{B}\|_F.$$

This is because  $GG^T \mathbf{X}$  is not the truncated SVD of  $\mathbf{X}$  (Note the difference here:  $\mathbf{G}$  is computed from the SVD of  $\tilde{\mathbf{X}}$ ).

Does it mean that we fail to achieve the minimum reconstruction error in this case? No. To see why, we derive

$$\begin{aligned} \|\mathbf{X} - \hat{\mathbf{X}}\|_F &= \|\mathbf{X} - \mathbf{G}\mathbf{G}^T \mathbf{X}\|_F \\ &= \|(\tilde{\mathbf{X}} + \bar{\mathbf{X}}) - \mathbf{G}\mathbf{G}^T (\tilde{\mathbf{X}} + \bar{\mathbf{X}})\|_F \\ &= \|(\tilde{\mathbf{X}} - \mathbf{G}\mathbf{G}^T \tilde{\mathbf{X}}) + (\bar{\mathbf{X}} - \mathbf{G}\mathbf{G}^T \bar{\mathbf{X}})\|_F. \end{aligned}$$

Note that the term  $(\bar{\mathbf{X}} - \mathbf{G}\mathbf{G}^T \bar{\mathbf{X}})$  represents the difference between the original mean matrix and the reconstructed mean matrix. Thus, if we project and reconstruct the original data matrix  $\mathbf{X}$ , the reconstruction error includes a term for reconstructing the mean matrix. However, the mean matrix is an offset term in the projection and should not be considered when computing the reconstruction error. To conclude, to show that PCA projects data with minimum reconstruction error, we should center the data first.

## 5 Connections with Autoencoders

Autoencoders [1] represent a family of fundamental machine learning methods, especially in unsupervised learning. In this section, we discuss the relationships between autoencoders and PCA. In the following discussion, we assume that  $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n] \in \mathbb{R}^{p \times n}$  is already centered, for the simplicity of notations.

### 5.1 What are Autoencoders?




Figure 1: A basic framework of autoencoders.

Autoencoders have become more and more popular with the development of deep learning. In particular, they play an important role in unsupervised learning tasks. Variants of autoencoders have been used to learn embeddings of data. For example, in natural language processing, word embedding models like *word2vec* can be seen as a variant of autoencoders. In addition, other variants have achieved great successes in data generation, like *variational autoencoders (VAE)* and *generative adversarial networks (GAN)*. Here, we discuss the original autoencoders.

A basic framework of autoencoders is illustrated in Figure 1. The outputs of this framework have the same dimension as inputs, and are supposed to reconstruct the inputs exactly. Without any constraint, the reconstruction task can be easily solved by directly copying. However, autoencoders come with a crucial restriction that the dimension of intermediate outputs should be smaller than inputs. Concretely, the framework of autoencoders can be divided into two parts: the *encoder*  $E(\cdot)$  and the *decoder*  $D(\cdot)$ . Given  $\mathbf{X} \in \mathbb{R}^{p \times n}$  as inputs, the encoder is supposed to output a reduced representation of the inputs, *i.e.*,  $E(\mathbf{X}) \in \mathbb{R}^{k \times n}$ , where  $k < p$ . And the decoder tries to reconstruct the inputs from this reduced representation. It is straightforward to see that the objective function of autoencoders is to minimize the reconstruction error:

$$L(\mathbf{X}, D(E(\mathbf{X}))),$$

where  $L(\cdot, \cdot)$  is a loss function that measures the reconstruction error.

### 5.2 PCA as Autoencoders

In Section 4.3, we've shown that PCA is able to achieve the minimum reconstruction error. Both PCA and autoencoders are unsupervised learning models with the goal of minimizing the reconstruction error. Are they connected? The answer is *yes* and we will show that PCA can fit into the framework of autoencoders.

Consider a special autoencoder with the encoder and decoder defined as

$$\begin{aligned} E(\mathbf{X}) &= \mathbf{W}^T \mathbf{X} \in \mathbb{R}^{k \times n}, \\ D(E(\mathbf{X})) &= \mathbf{W}E(\mathbf{X}) = \mathbf{W}\mathbf{W}^T \mathbf{X} \in \mathbb{R}^{p \times n}, \end{aligned}$$

where  $\mathbf{W} \in \mathbb{R}^{p \times k}$ . Note that

$$\text{rank}(\mathbf{W}\mathbf{W}^T \mathbf{X}) \leq \min(\text{rank}(\mathbf{W}), \text{rank}(\mathbf{W}^T \mathbf{X})) \leq k$$

If the loss function  $L(\cdot, \cdot)$  is the mean square error, the objective of this autoencoder is equivalent to

$$\min_{\mathbf{W}: \text{rank}(\mathbf{W}\mathbf{W}^T \mathbf{X}) \leq k} \|\mathbf{X} - \mathbf{W}\mathbf{W}^T \mathbf{X}\|_F. \quad (5)$$

In an autoencoder,  $\mathbf{W}$  can be randomly initialized and optimized through gradient descent. However, comparing Eqn. (4) and Eqn. (5), it is easy to see that the optima can be achieved when  $\mathbf{W} = \mathbf{G}$ , where  $\mathbf{G}$  is given by PCA.

Therefore, a natural question is: Will gradient descent give the same results as PCA? The answer is two-fold: *yes* in terms of the minimal reconstruction error and *not necessary* in terms of optimal  $\mathbf{W}$ . Gradient descent converges to the global optima, resulting in the same minimal reconstruction error as PCA. However, as pointed out in Section 4.1, the solutions to PCA are not unique and can differ by an orthogonal matrix. Thus, the optimal  $\mathbf{W}$  can be different from a PCA solution  $\mathbf{G}$ . But one can be computed from another through an orthogonal matrix and the sets of possible solutions are the same.

To summarize, PCA is a special case of autoencoders, where the encoder and decoder are both one-layer linear transformations, which share the weights and do not have bias terms. A similar discussion can be found in [2].

## 6 Exercises

### 6.1 Eigen-decomposition and SVD

Given a symmetric matrix  $A \in \mathbb{R}^{3 \times 3}$ , suppose its eigen-decomposition can be written as

$$A = \begin{pmatrix} u_{11} & u_{12} & u_{13} \\ u_{21} & u_{22} & u_{23} \\ u_{31} & u_{32} & u_{33} \end{pmatrix} \begin{pmatrix} 3 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & -2 \end{pmatrix} \begin{pmatrix} u_{11} & u_{21} & u_{31} \\ u_{12} & u_{22} & u_{32} \\ u_{13} & u_{23} & u_{33} \end{pmatrix}.$$

What is the singular value decomposition of this matrix?

### 6.2 Ky Fan Theorem

Provide a complete proof of the Ky Fan Theorem.

## Acknowledgements

Our fundamental understanding of SVD and PCA is strongly inspired by Jieping Ye (University of Michigan), whose insights and passion on revealing the elegance between linear algebra and machine learning are unparalleled. We thank Lei Cai for proof-reading. This work was supported in part by National Science Foundation grants IIS-1908220, IIS-1908198, IIS-1908166, DBI-1147134, DBI-1922969, DBI-1661289, CHE-1738305, National Institutes of Health grant 1R21NS102828, and Defense Advanced Research Projects Agency grant N66001-17-2-4031.

## References

- [1] Hinton GE, Salakhutdinov RR (2006) Reducing the dimensionality of data with neural networks. *science* 313: 504–507.
- [2] Aggarwal CC (2018) *Neural Networks and Deep Learning: A Textbook*. Springer.