PCA, THE SVD, AND THE GENERALIZED RAYLEIGH QUOTIENT

PCA

In the previous set of notes, we saw that PCA reduces to the following optimization problem:

$$\min_{P \in \mathcal{P}_k} \|X - PX\|_F \qquad \text{(PCA)}$$

where X is the $d \times n$ data matrix (column mean $= 0$), $\|\cdot\|_F$ is the Frobenius norm, and $\mathcal{P}_k$ is the set of rank k, $d \times d$ projection matrices. In this setting, the covariance matrix of the data is

$$S = \frac{1}{n}\sum_{i=1}^{n} x_i x_i^T = \frac{1}{n} X X^T$$

verify by comparing entries

We will derive the solution to PCA in two different ways, by connecting (PCA) to two other problems in matrix algebra.

$\boxed{\text{Connection to the SVD}}$

Every matrix X has a singular value decomposition (SVD)

$$X = U \Sigma V^T$$

where

columns of U: left singular vectors
columns of V: right singular vectors
$\sigma_1 \geq \sigma_2 \geq \dots \geq 0$: singular values

- $UU^T = U^TU = I_{d\times d}$
- $VV^T = V^TV = I_{n\times n}$
- $$\Sigma = \begin{cases} \begin{bmatrix} \sigma_1 & & & \\ & \ddots & & \\ & & \sigma_d & \end{bmatrix} & \text{if } n \geq d \\ \begin{bmatrix} \sigma_1 & & \\ & \ddots & \\ & & \sigma_n \\ & & \end{bmatrix} & \text{if } n < d \end{cases}$$

The SVD arises in the following theorem due to Eckart + Young.

Theorem Let X have rank $r \geq k$. The solution to

$$\min_{\substack{Z \in \mathbb{R}^{d\times n} \\ \operatorname{rank}(Z) = k}} \|X - Z\|_F^2 \qquad \text{(SVD)}$$

is $Z_k = U \Sigma_k V^T$, where Σ_k is Σ with

$\sigma_{k+1}, \sigma_{k+2}, \dots$ set to zero.

The proof of this result is given below.

To see the connection to (PCA), write

$$\Sigma_k = I_{k,d} \cdot \Sigma,$$

where

$$I_{k,d} = \begin{bmatrix} 1 & & & & & & \\ & 1 & & & & & \\ & & \ddots & & & & \\ & & & 1 & & & \\ & & & & 0 & & \\ & & & & & 0 & \\ & & & & & & \ddots \\ & & & & & & & 0 \end{bmatrix}. \qquad (d \times d)$$

Then

$$\begin{aligned} Z_k &= U \Sigma_k V^T \\ &= U (I_{k,d} \Sigma) V^T \\ &= U (I_{k,d} \cdot U^T X V) V^T \\ &= U_k U_k^T X \end{aligned}$$

where U_k contains the first k left singular vectors.

Clearly $U_k U_k^T \in \mathcal{P}_k$. Therefore $P = U_k U_k^T$ gives a solution to (PCA).

It remains to show that the left singular vectors are the eigenvectors of XX^T. But the eigenvalue decomposition of XX^T is

$$XX^T = U\Sigma V^T V \Sigma^T U^T$$
$$= U \Sigma \Sigma^T U^T$$
$$= U \Lambda U^T$$

where

$$\Lambda = \begin{cases} \begin{bmatrix} \sigma_1^2 & & \\ & \ddots & \\ & & \sigma_d^2 \end{bmatrix} & \text{if } n \geq d \\[2ex] \begin{bmatrix} \sigma_1^2 & & & & \\ & \ddots & & & \\ & & \sigma_n^2 & & \\ & & & 0 & \\ & & & & \ddots \\ & & & & & 0 \end{bmatrix} & \text{if } d > n. \end{cases}$$

Therefore, the SVD of X gives the PCA solution:

- principal eigenvectors = left singular vectors of X
- $\lambda_i = \frac{1}{n}$ (i^{th} singular value of X)

Proof of Eckart-Young Thm

Observe

$$\|X - Z\|_F = \|U\Sigma V^T - Z\|_F$$

$$= \|\Sigma - U^T Z V\|_F .$$

Denote $N = U^T Z V$, a $d \times n$ matrix of rank k.

A direct calculation gives

$$\|\Sigma - N\|_F^2 = \sum_{i,j} |\Sigma_{ij} - N_{ij}|^2$$

$$= \sum_{i=1}^{r} |\sigma_i - N_{ii}|^2 + \sum_{i>r} |N_{ii}|^2 + \sum_{i \neq j} |N_{ij}|^2 .$$

This is minimized (subject to N having rank k) when $N_{ii} = \sigma_i$ for $i = 1, \dots, k$, and all other $N_{ij} = 0$.

This implies $Z = Z_k$.

Generalized Rayleigh Quotient

We now present a second solution of (PCA).

The trace of a square matrix is the sum of the

The trace of a square matrix is the sum of the diagonal entries. It satisfies the following properties:

- linearity: $\mathrm{tr}(C+D) = \mathrm{tr}(C) + \mathrm{tr}(D)$ for any two square matrices C and D
- invariance to cyclic permutations: $\mathrm{tr}(CD) = \mathrm{tr}(DC)$ as long as CD and DC are both well-defined.
- the trace of a matrix is the sum of its eigenvalues.
- for any matrix C, $\|C\|_F^2 = \mathrm{tr}(C^T C)$

These properties are easily verified. Now observe

$$\begin{aligned}
\|X - PX\|_F^2 &= \mathrm{tr}\left((X-PX)^T(X-PX)\right) \\
&= \mathrm{tr}(X^TX) - \mathrm{tr}(X^TPX) \\
&\quad - \mathrm{tr}(X^TP^TX) + \mathrm{tr}(X^TP^TPX) \\
&= \mathrm{tr}(X^TX) - \mathrm{tr}(X^TPX)
\end{aligned}$$

where we used $P = P^T$ and $P^2 = P$. Writing $P = AA^T$ where $A \in \mathcal{A}_k$, we need to maximize

$$\mathrm{tr}(X^T P X) = \mathrm{tr}(X^T A A^T X)$$

$$= \mathrm{tr}(A^T X X^T A)$$

The derivation of PCA is concluded by the following result, which I will refer to as the generalized Rayleigh quotient theorem (this terminology is not standard) because, in the special case $k=1$, it relates to the Rayleigh quotient.

Theorem Let C be a PSD matrix with eigenvalue decomposition $C = U \Lambda U^T$, where $U = [u_1 \cdots u_d]$. Then a solution of

$$\max_{A \in \mathcal{A}_k} \mathrm{tr}(A^T C A) \qquad \text{(GRQ)}$$

is $A = [u_1 \cdots u_k]$.

Proof of GRQ Theorem Introduce the change of variables

$$w_i = U^T a_i.$$

We know that $w_1, \ldots, w_k$ are orthonormal because

$$w_i^T w_j = a_i^T U U^T a_j = a_i^T a_j .$$

Then we need to maximize

$$tr(A^T C A) = \sum_{i=1}^{k} a_i^T C a_i$$

$$= \sum_{i=1}^{k} w_i^T \Lambda w_i$$

subject to $W = [w_1 \cdots w_k] \in \mathcal{A}_k$. Now

$$\sum_{i=1}^{k} w_i^T \Lambda w_i = \sum_{i=1}^{k} \sum_{j=1}^{d} (w_i^{(j)})^2 \lambda_j$$

$$= \sum_{j=1}^{d} \left[\sum_{i=1}^{k} (w_i^{(j)})^2 \right] \lambda_j$$

$$= \sum_{j=1}^{d} h_j \lambda_j$$

where $h_j = \sum_{i=1}^{k} (w_i^{(j)})^2$.

$\boxed{\text{Lemma}}$ $0 \le h_j \le 1 \quad \forall j$ and $\sum_{j=1}^{d} h_j = k$.

Proof: The second part is easy:

$$\sum_{j=1}^{d} h_j = \sum_{j=1}^{d} \left(\sum_{i=1}^{k} (w_i^{(j)})^2 \right)$$

$$= \sum^{k} \sum^{d} (w_i^{(j)})^2$$

$$= \sum_{i=1}^{k} \sum_{j=1}^{d} (w_i^{(j)})^2$$

$$= \sum_{i=1}^{k} (1)$$

$$= k.$$

$h_j \geq 0$ is also obvious. To show $h_j \leq 1$, let $w_{k+1}, \ldots, w_d$ extend $w_1, \ldots, w_k$ to an orthonormal basis. Consider the square matrix

$$M = \begin{bmatrix} w_1 & \cdots & w_d \end{bmatrix} \qquad (d \times d)$$

We know $M^T M = I$ by orthonormality. Therefore M^T is a left inverse of M, and so must also be a right inverse (a property of square matrices), meaning $MM^T = I$. This implies

$$h_j = \sum_{i=1}^{k} (w_i^{(j)})^2 \leq \sum_{i=1}^{d} (w_i^{(j)})^2 = 1. \qquad \blacksquare$$

We need to maximize

$$\sum_{j=1}^{d} h_j \lambda_j$$

with respect to the constraints imposed by the lemma. Since $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_d \geq 0$, this is accomplished by

$$h_j = \begin{cases} 1 & \text{if} \quad 1 \leq j \leq k \\ 0 & \text{otherwise} \end{cases}.$$

which in turn is achieved by

$$W = \left[\begin{array}{c} I_{k \times k} \\ \hdashline 0 \end{array} \right]$$

Therefore $A = UW = [u_1 \; \cdots \; u_k]$.

Note that the optimal A is not unique. Indeed if

$$W = \left[\begin{array}{c} \text{any set of} \\ \text{length } k \\ \text{orthonormal} \\ \text{vectors} \\ \hdashline 0 \end{array} \right]$$

then $A = UW$ also achieves the maximum. $\blacksquare$

An interesting question is when is $\langle A \rangle$ unique. You'll get to think about this on the homework.

The GRQ theorem can be used to derive PCA from the maximum variance approach described in the previous set of notes.

Furthermore, instead of a sequential definition of maximum variance, we can instead ask what $A \in \mathcal{A}_k$ maximizes the total variance

$$\sum_{i=1}^{k} a_i^T S a_i = tr(A^T S A).$$

The GRQ theorem again tells us that the principal eigenvectors of S provide the solution.

We will employ the GRQ theorem later in the course when we discuss spectral clustering.