

EE445 Mod3-Lec4: Principal Component Regression & Kernel PCA

Spectral Clustering

Outline

1. Principal Component Regression (PCR)
2. Kernel PCA
3. (time permitting) Spectral Clustering

Part 1: PCR

What is Principal Component Regression (PCR)?

- Combining PCA with linear regression leads to principal components regression (PCR).
- PCR = PCA + linear regression:
 - ▶ Choose how target number of principal components k
 - ▶ Use PCA to define a feature vector

$$\varphi(x^{(i)}) = (\langle x^{(i)}, u_1 \rangle, \dots, \langle x^{(i)}, u_k \rangle)$$

as in **Mod3-L3** containing the principal component scores $x^{(i)}$ —i.e., the projections onto the u_1, \dots, u_k principal components


- ▶ Use least-squares linear regression (as in **Module 2**) with this model:

$$y^{(i)} = \varphi(x^{(i)})^\top \theta + \varepsilon_i$$

Facts about PCR

- **Why use it?:** In regression analysis, when the independent variables appear multicollinearly (independent variables are correlated), the general effect of the classical regression method for least square estimates of regression coefficients will be poor, but principal component analysis can overcome this deficiency effectively.

Partial least squares
(PLS)
Supervised Method



PCR
Unsupervised
Method

- **PCR** works well when the directions in which the original predictors vary most are the directions that are predictive of the outcome

• **PCR** is very similar to **ridge regression** in a certain sense.

- Nice comparison between **PCR** and partial least squares

https://scikit-learn.org/stable/auto_examples/cross_decomposition/plot_pcr_vs_pls.html

$$\|K\alpha - y\|^2 + \frac{\lambda}{2} \| \alpha \|^2 \quad \text{"Ridge" Regression} \quad K(x_1, z) = x_1^\top z$$

$$K(x_1, z) = \exp(-\gamma \|x_1 - z\|^2)$$

- We saw Kernel regression with regularization in **Mod2 Lec4** where ℓ_2 -regularization was introduced to support numerical stability of the kernel matrix
- We can in general introduce a $\| \cdot \|_2$ regularization term to vanilla least squares and this is called *ridge regression*:

$$F(\theta) = \frac{1}{2} \|X\theta - y\|^2 + \frac{\lambda}{2} \| \theta \|^2 = \frac{1}{2} \theta^\top X^\top X \theta - y^\top X \theta + \frac{1}{2} y^\top y + \frac{\lambda}{2} \theta^\top \theta$$

First order: $\nabla F(\theta) = \underbrace{X^\top X \theta - X^\top y}_{=0} + \lambda \theta = 0 \Rightarrow \hat{\theta}_{\text{ridge}} = (X^\top X + \lambda I)^{-1} X^\top y$

- $\lambda \geq 0$: a tuning parameter that controls the strength of the penalty term
 - ▶ $\lambda = 0$: we get the linear regression estimate
 - ▶ $\lambda \rightarrow \infty$: we get $\hat{\theta}_{\text{ridge}} \rightarrow 0$
 - ▶ $\lambda \in (0, \infty)$: balancing fitting a linear model of y given data X , and **shrinking the coefficients** in θ

PCR and Ridge

$$\underline{V V^T = I}$$

- Write $X = \underline{USV^T}$

$$\begin{aligned} (X^T X + \lambda I)^{-1} &= (V S^T S V^T + \lambda I)^{-1} = (V S^T S V^T + \lambda \underline{V V^T})^{-1} \\ &= (V (S^T S + \lambda I) V^T)^{-1} = \underline{V (S^T S + \lambda I)^{-1} V^T} \end{aligned}$$

$$= \underline{V S^+ V^T}, \quad S^+ = \text{diag} \left(\frac{1}{\sigma_1^2 + \lambda}, \dots, \frac{1}{\sigma_n^2 + \lambda} \right)$$

$$\hat{\theta}_{\text{ridge}} = V S^+ \underline{V^T V S^T U^T} y = \underline{V S^+ S^T} U^T y = \sum_{i=1}^n \frac{\sigma_i}{\sigma_i^2 + \lambda} \underline{v_i u_i^T} y$$

$$= \underline{(X^T X + \lambda I)^{-1}} X^T y$$

- Hence, Ridge regression can be viewed as projecting the y vector onto the principal component directions and then shrinking the projection using λ

PCR vs Ridge

$$\|\theta\|_1 = \sum_{i=1}^n |\theta_i|$$

- Ridge regression shrinks everything, but it never shrinks anything to zero.
- By contrast, **PCR** either does not shrink a component at all or shrinks it to zero.
- Yet another alternative is **LASSO**: assuming the features are orthonormal

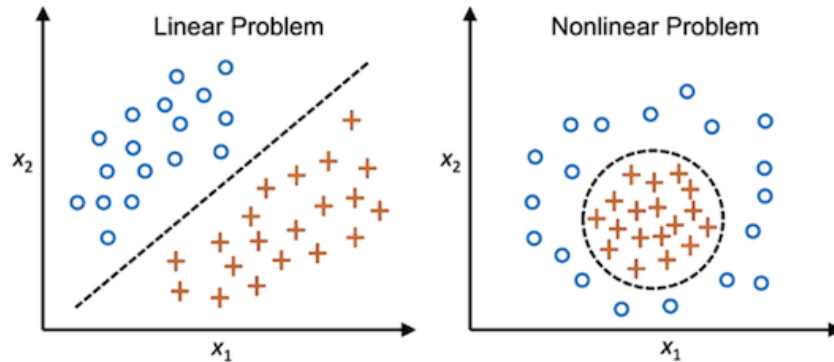
$$\min_{\theta} \frac{1}{2} \|X\theta - y\|^2 + \frac{\lambda}{2} \|\theta\|_1 \rightarrow \hat{\theta}_j = \hat{\theta}_j^{\text{ols}} \max \left\{ 0, 1 - \frac{m\lambda}{|\hat{\theta}_j^{\text{ols}}|} \right\}$$

which can set some coefficients to zero, and scales others

Part 2: Kernel PCA

Motivation

- when the data is non-linear, we may need a more complex polynomial function to separate the data



Recall the Kernel "trick"

Mod2-L4

$$\phi(x) = x^d$$

$$O(n^d) \rightarrow O(n)$$

- Essence of Kernel trick:

- ▶ If we can write down an algorithm only in terms of $\phi(x^{(i)})^\top \phi(x^{(j)})$, then we don't need to explicitly enumerate $\phi(x^{(i)})$ for all the $i \in \{1, \dots, m\}$
- ▶ Instead we can just compute $\underline{K(x^{(i)}, x^{(j)})} = \underline{\phi(x^{(i)})^\top \phi(x^{(j)})}$
- ▶ And we can even handle infinite dimensional feature maps $\phi(\cdot)$

- Examples:

- ▶ Linear kernel $K(x, z) = z^\top x$
- ▶ radial basis function $K(x, z) = \exp(-\gamma \|x - z\|^2)$ (Gaussian kernel)
- ▶ polynomial kernels $K(x, z) = (\underline{x^\top z + c})^p$

- Predictions only depend on training data through kernel function which is just a dot product.

$$\| \theta^\top x - y \|^2 \rightarrow \| k_d - y \|^2$$

Rewrite PCA using Kernel "trick" $\{\mathbb{R}^2\} \ni \theta = \sum \alpha x^{(i)}$

↓

- Let u_1, \dots, u_k be the k largest principal components, and define matrix $U = [u_1 \ \dots \ u_k]$ —i.e., the j -th column of U is the j -th eigenvector of the covariance matrix

$$\Sigma = \frac{1}{m} \sum_{i=1}^m \phi(x^{(i)}) \phi(x^{(i)})^\top$$

$$\Sigma = \frac{1}{m} \sum_{i=1}^m x^{(i)} (x^{(i)})^\top$$

- Claim:** Eigenvectors can be expressed as linear combination of features

Let (λ, v) be an eigenpair of Σ . $v = \sum_{i=1}^m \alpha_i \phi(x^{(i)}) \leftarrow$

$$\Sigma v = \frac{1}{m} \sum_{i=1}^m \phi(x^{(i)}) \phi(x^{(i)})^\top v = \lambda v \Rightarrow$$

$$v = \frac{1}{\lambda m} \sum_{i=1}^m \phi(x^{(i)}) \phi(x^{(i)})^\top v = \frac{1}{\lambda m} \sum_{i=1}^m \underbrace{(\phi(x^{(i)})^\top v)}_{\text{scalar}} \phi(x^{(i)})$$

- Hence, finding the eigenvectors is equivalent to finding the coefficients α_i

Rewrite PCA using Kernel "trick" u_1, \dots, u_k

- That is, $u_j = \sum_{l=1}^m \alpha_{jl} \phi(x^{(l)})$ for each $j = 1, \dots, k$
- Let's find an expression for the α_j 's

$$\sum u_j = \lambda_j u_j$$

$$\frac{1}{m} \sum_{i=1}^m \phi(x^{(i)}) \phi(x^{(i)})^T \left(\sum_{l=1}^m \alpha_{jl} \phi(x^{(l)}) \right) = \lambda_j \sum_{l=1}^m \alpha_{jl} \phi(x^{(l)})$$

• Rewrite as

$$\frac{1}{m} \sum_{i=1}^m \phi(x^{(i)}) \left(\sum_{l=1}^m \alpha_{jl} K(x^{(i)}, x^{(l)}) \right) = \lambda_j \sum_{l=1}^m \alpha_{jl} \phi(x^{(l)})$$

$$\alpha_{jl} \underbrace{\phi(x^{(i)})^T \phi(x^{(l)})}_{=: K(x^{(i)}, x^{(l)})}$$

• Multiply by $\phi(x^{(s)})$

$$\frac{1}{m} \sum_{i=1}^m \underbrace{\phi(x^{(s)})^T \phi(x^{(i)})}_{\frac{1}{m} K^2 \alpha_j} \left(\sum_{l=1}^m \alpha_{jl} K(x^{(i)}, x^{(l)}) \right) = \lambda_j \sum_{l=1}^m \alpha_{jl} \underbrace{\phi(x^{(s)})^T \phi(x^{(l)})}$$

$$= \lambda_j K \alpha_j$$

$y^{(i)} = (\langle x^{(i)}, u_1 \rangle, \dots, \langle x^{(i)}, u_k \rangle)$ Rewrite PCA using Kernel "trick"

- By plugging in the kernel and rearranging we get:

$$\underline{K^2} \alpha_j = \underline{m \lambda_j} K \alpha_j \implies \boxed{K \alpha_j = m \lambda_j \alpha_j}$$

$$\alpha_j = \sum \alpha_{j\ell} \phi(x^{(\ell)})$$

We can remove a factor of K from both sides of the matrix (this only affect the eigenvectors with zero eigenvalue, which will not be a principal component anyway):

- Since $\|u_j\|^2 = 1$, we have that

$$\sum_{i=1}^m \sum_{\ell=1}^m \alpha_{j\ell} \alpha_{ji} \underbrace{\phi(x^{(i)})^T \phi(x^{(\ell)})}_{=K(x^{(i)}, x^{(\ell)})} = 1 \implies \alpha_j^T K \alpha_j = 1 \implies \alpha_j^T \alpha_j \cdot \lambda_j^{-m} = 1$$

$\|\alpha_j\|^2 = \frac{1}{\lambda_j^m}$

- Hence, for some feature vector x its projection on to the principal components is

$$\phi(x)^T u_j = \sum_{i=1}^m \alpha_{ji} \phi(x)^T \phi(x^{(i)}) = \sum_{i=1}^m \alpha_{ji} K(x, x^{(i)})$$

Summary of the Kernel Trick

- We showed that the eigenvectors (principal components) can be expressed as linear combination of the features with some coefficients α_i
- We showed that α_i 's are eigenvectors of the kernel matrix, and have some normalization property (norm equal to $1/(m\lambda_i)$)
- Showed the projection of a feature vector onto a principal component can be obtained via the kernel matrix

$$\phi(x)^\top u_j = \sum_{i=1}^m \alpha_{ji} K(x, x^{(i)})$$

Normalizing the Feature Space

$$K(x^{(i)}, x^{(j)}) = \phi(x^{(i)})^\top \phi(x^{(j)})$$

$$x^{(i)}, i=1, \dots, m$$

$$\tilde{x}^{(i)} = x^{(i)} - \mu$$

- In general $\phi(x^{(i)})$ may not be zero mean
- We can recenter the features just like before

$$\varphi(x^{(i)}) = \underbrace{\phi(x^{(i)})}_{\text{feature}} - \frac{1}{m} \sum_{j=1}^m \phi(x^{(j)})$$

- The new kernel is $\tilde{K}(x^{(i)}, x^{(j)}) = \varphi(x^{(i)})^\top \varphi(x^{(j)})$ and without going through the details, we can derive an expression for the new kernel in matrix form

$$\tilde{K} \approx K - \frac{1}{m} \mathbb{1}_{m \times m} K - \frac{1}{m} K \mathbb{1}_{m \times m} + \frac{1}{m^2} \mathbb{1}_{m \times m} K \mathbb{1}_{m \times m}$$

$$\mathbb{1}_{m \times m} = \begin{bmatrix} 1 & \dots & 1 \\ \vdots & \ddots & \vdots \\ 1 & \dots & 1 \end{bmatrix}$$

Summary of Kernel PCA

- pick a kernel

- Construct $\tilde{K} = K - \frac{1}{n} \mathbb{1} K - \frac{1}{n} K \mathbb{1} + \frac{1}{n^2} \mathbb{1} K \mathbb{1}$

- Solve eigenvalue problem: $\tilde{K} \alpha_i = \lambda_i \alpha_i$

- for any x , $y^{(j)} = \sum_{l=1}^m \alpha_{jl} K(x, x^{(l)})$, $j = 1, \dots, k$

Example: De-noising images

Original data



Data corrupted with Gaussian noise



Result after linear PCA





Result after kernel PCA, Gaussian kernel




Part 3: Spectral Clustering

What is spectral clustering?

- Goal is to group points based on links in a graph



How to create the graph? $G = (E, V)$



- Given a distance metric, we compute the distance between each of our features
- It is common to use a Gaussian Kernel (RBF) to compute the similarity between features:

$$s_{ij} = S(i, j) = \exp\left(-\frac{\|x^{(i)} - x^{(j)}\|^2}{\sigma^2}\right)$$

- There are several popular constructions to transform a given set $x^{(1)}, \dots, x^{(m)}$ of data points with pairwise similarities s_{ij} or pairwise distances d_{ij} into a graph.
 - ▶ fully connected graph: Here we simply connect all points with positive similarity with each other, and we weight all edges $w_{ij} = s_{ij}$
 - ▶ k -nearest neighbor graphs: Here the goal is to connect vertex v_i with vertex v_j if v_j is among the k -nearest neighbors of v_i .
 - ▶ The ε -neighborhood graph: Here we connect all points whose pairwise distances are smaller than ε

Graph Review



- Consider a graph $G = (E, V)$ where $V = \{v_1, \dots, v_n\}$ is the vertex set and E is the set of edges

- G is weighted if there is some matrix W where w_{ij} is the non-negative weight on edge e_{ij}

$$W = \begin{bmatrix} & & & \\ & & & \\ & & & \\ & & & \end{bmatrix}$$

W is weighted adjacency matrix

• Graph Laplacian $L = D - W$

- $d_i = \sum_{j=1}^n w_{ij}$ degree of vertex i , $D = \text{diag}(d_1, \dots, d_n)$

Spectral Clustering Algorithm (un-normalized)

Input: Similarity matrix $S \in \mathbb{R}^{m \times m}$ (pairwise similarities between edges) and number of clusters k

Step 1: Build W based on graph

Step 2: Compute $L = D - W$

Step 3: Do PCA on graph Laplacian: u_1, \dots, u_k eigrs of L

Step 4: $V = [u_1 \dots u_k]$

Step 5: treat $y^{(i)}$ as i^{th} row of V



apply k -means to $y^{(i)}$, $i = 1, \dots, m \rightarrow C_1, \dots, C_k$

Output: $A_1, \dots, A_k = A_i = \{j \mid y^{(j)} \in C_i\}$

Example

go to `Mod3-Lec4.ipynb`