

Dimensionality Reduction

PCA

Aarti Singh

So far

- 1) MLE / MAP
- 2) Nonparametric density

Machine Learning 10-315
April 4, 2022

Slides Courtesy: Tom Mitchell, Eric Xing, Lawrence Saul



MACHINE LEARNING DEPARTMENT



High-Dimensional data

- High-Dimensions = Lot of Features

Document classification

Features per document =

thousands of words/unigrams

millions of bigrams, contextual
information



Surveys - Netflix

480189 users x 17770 movies



	movie 1	movie 2	movie 3	movie 4	movie 5	movie 6
Tom	5	?	?	1	3	?
George	?	?	3	1	2	5
Susan	4	3	1	?	5	1
Beth	4	3	?	2	4	2

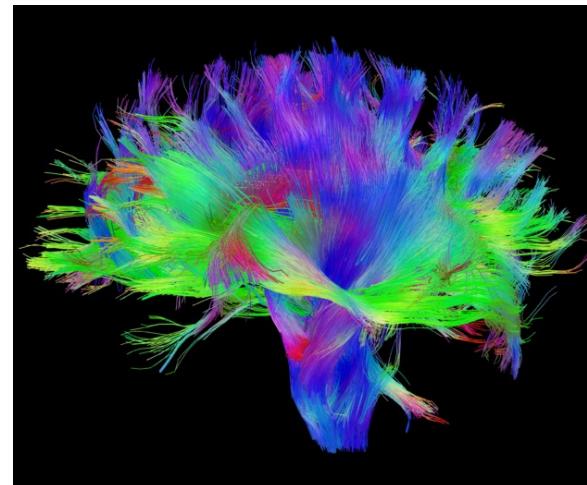
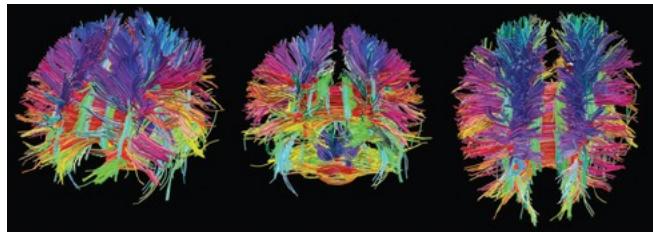
High-Dimensional data

- High-Dimensions = Lot of Features

High resolution images
millions of pixels



Diffusion scans of Brain
300,000 brain fibers



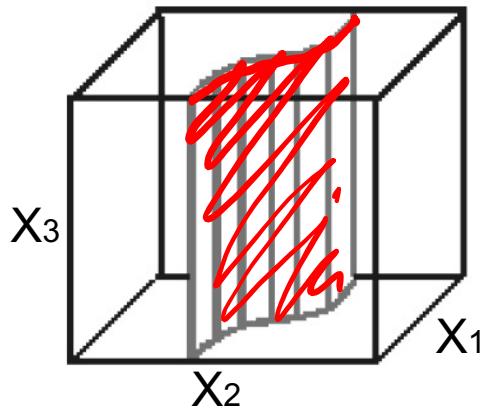
Curse of Dimensionality

- Why are more features bad?
 - Redundant features (not all words are useful to classify a document)
more noise added than signal
 - Hard to interpret and visualize
 - Hard to store and process data (computationally challenging)
 - Complexity of decision rule tends to grow with # features. Hard to learn complex rules as it needs more data (statistically challenging)

Dimensionality Reduction

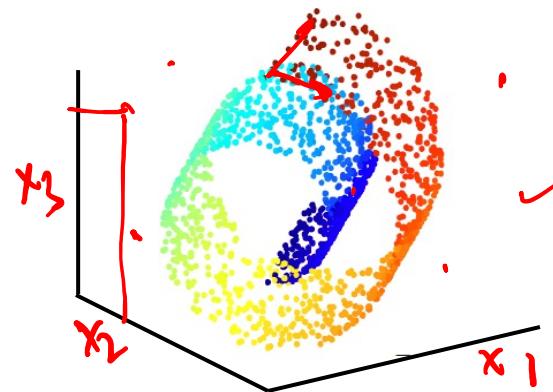
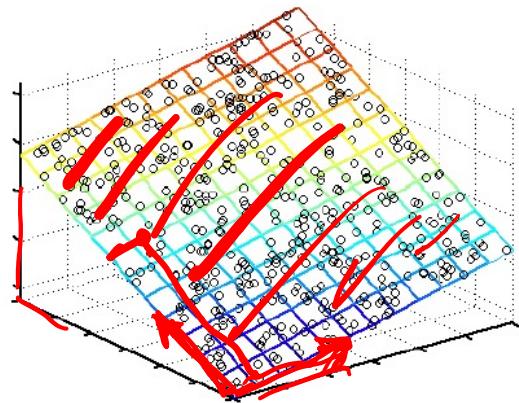
- **Feature Selection** – Only a few features are relevant to the learning task

—



X₃ - Irrelevant

- Latent features – Some linear/nonlinear combination of features provides a more efficient representation than observed features



Feature Selection

- One Approach: **Regularization (MAP)**

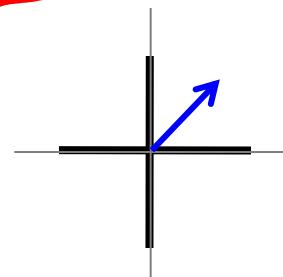
Integrate feature selection into learning objective by penalizing number of features with non-zero weights

$$\widehat{W} = \arg \min_W \sum_{i=1}^n -\log P(Y_i|X_i; W) + \lambda \|W\|$$

-ve log likelihood
penalty

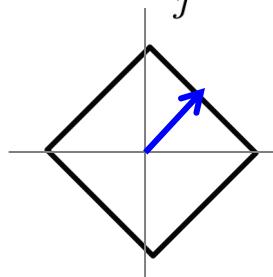
$\sum_{(i)} \underline{w_{(i)} x_{(i)}}$
 $\|W\|_1 = \sum_{(i)} |w_{(i)}|$
 $\|W\|_1$

$$\|W\|_0 = \#\{W_j > 0\}$$



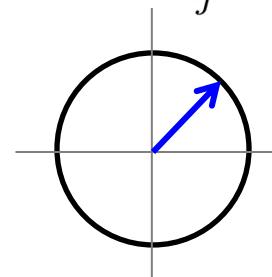
Minimizes # features chosen

$$\|W\|_1 = \sum_j |W_j|$$



Convex compromise

$$\|W\|_2 = \sum_j W_j^2$$



Small weights of features chosen

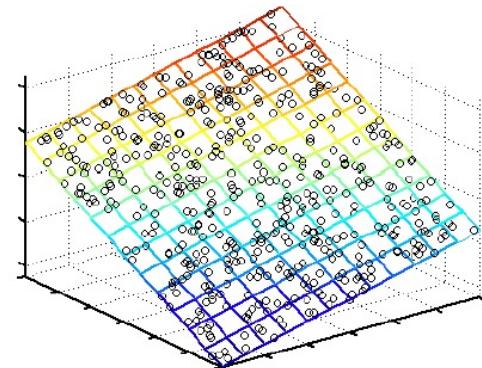
Latent Features

Combinations of observed features provide more efficient representation, and capture underlying relations that govern the data

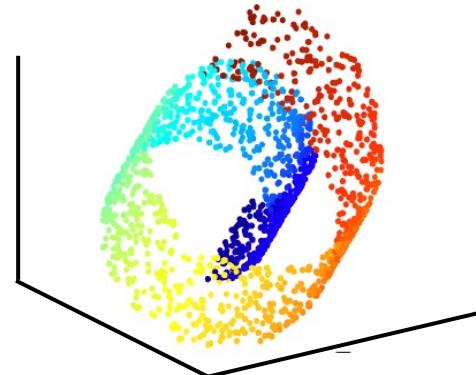
- E.g. Ego, personality and intelligence are hidden attributes that characterize human behavior instead of survey questions
- Topics (sports, science, news, etc.) instead of documents

Often may not have physical meaning

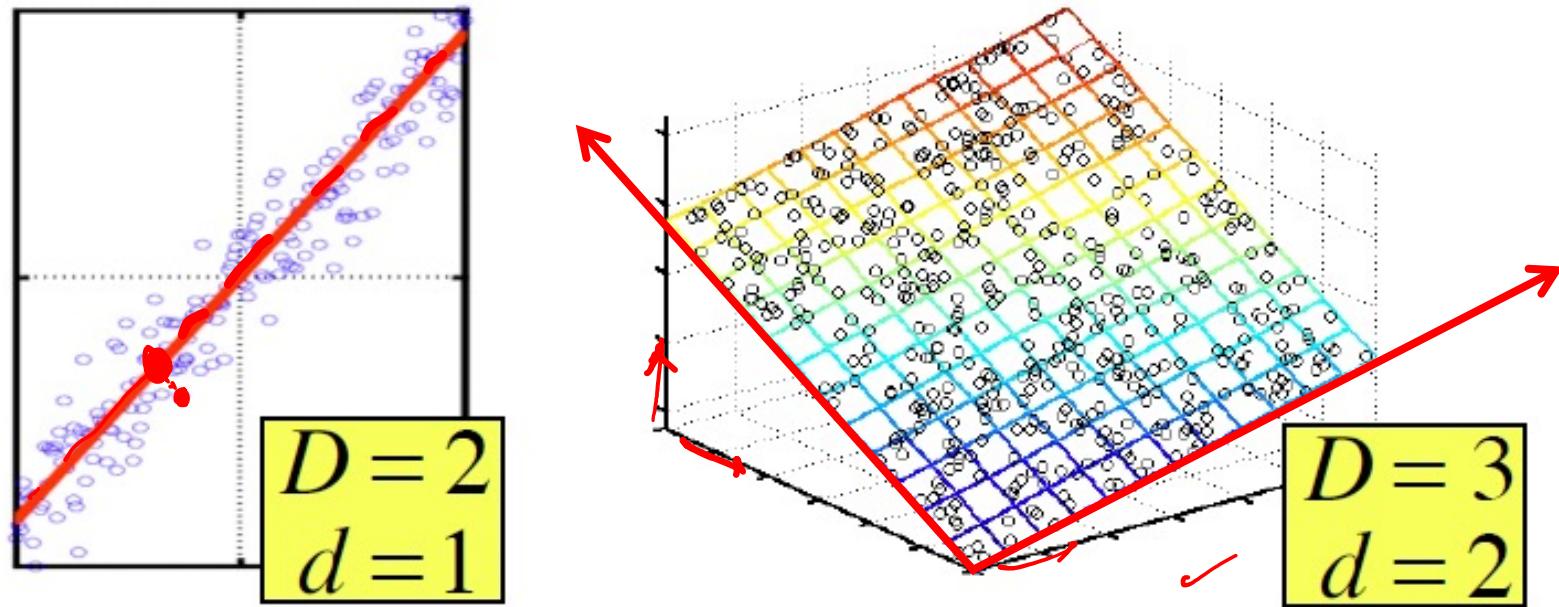
- Linear
 - Principal Component Analysis (PCA) ←
 - Factor Analysis
 - Independent Component Analysis (ICA)



- Nonlinear
 - Kernel PCA
 - Laplacian Eigenmaps, ISOMAP, LLE
 - Autoencoders



Principal Component Analysis (PCA)



When data lies on or near a low d -dimensional linear subspace, axes of this subspace are an effective representation of the data

Identifying the axes is known as Principal Components Analysis, and can be obtained by Eigen or Singular value decomposition

Data for PCA

Data $\underline{X} = [x_1, x_2, \dots, x_n]$ where each data point x_i is \underline{D} -dimensional vector

\underline{X} is $D \times n$ matrix

Assume data are centered i.e. sample mean

$$\frac{1}{n} \sum_{i=1}^n x_i = 0$$

What if data is not centered?

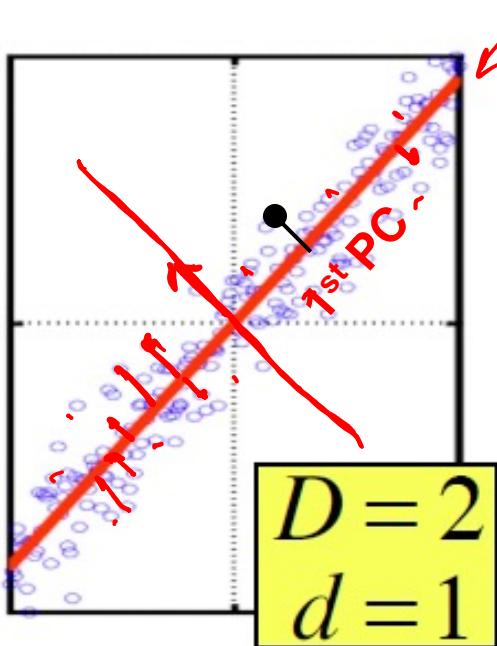
Subtract off sample mean from each data point

Since data matrix is centered, sample covariance matrix can be written as

$$S = \frac{1}{n} \underset{D \times D}{XX^\top}$$

$S_{ij} = \text{cov}_{i^{th} + j^{th} \text{ features}}$

Principal Component Analysis (PCA)



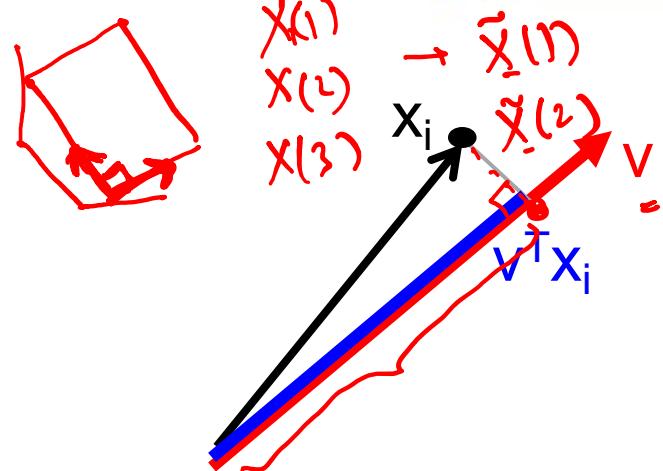
Principal Components (PC) are orthogonal directions that capture most of the variance in the data

1st PC – direction of greatest variability in data

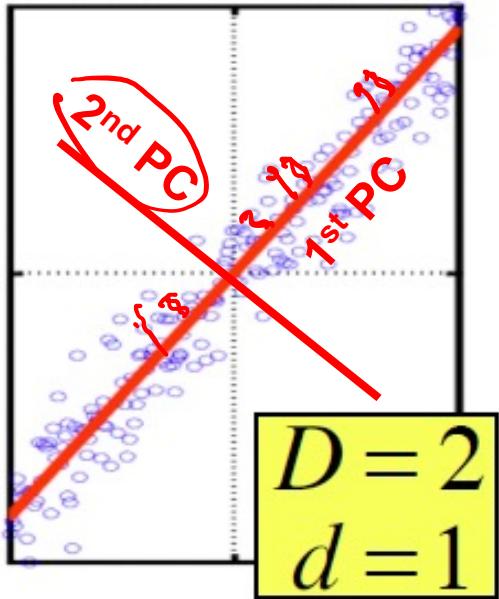
Projection of data points along 1st PC discriminate the data most along any one direction

Take a data point x_i (D -dimensional vector)

Projection of x_i onto the 1st PC v is $v^T x_i = \frac{v^T x_i}{\|v\|}$



Principal Component Analysis (PCA)

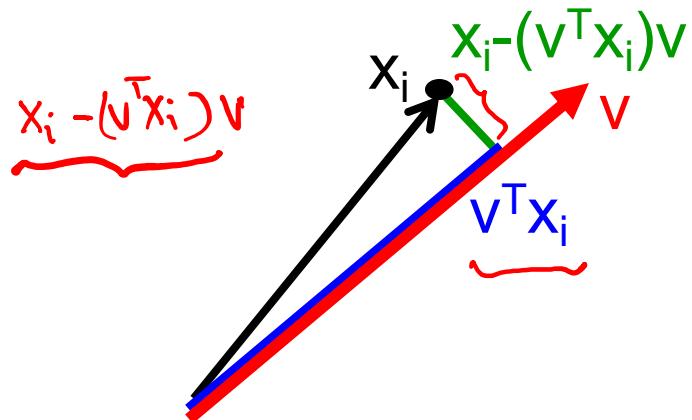


Principal Components (PC) are orthogonal unit norm directions that capture most of the variance in the data

1st PC – direction of greatest variability in data

2nd PC – Next orthogonal (uncorrelated) direction of greatest variability

(remove all variability in first direction, then find next direction of greatest variability)



And so on ...

$$D \rightarrow d < D$$

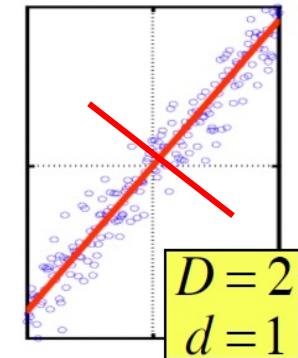
Principal Component Analysis (PCA)

Let v_1, v_2, \dots, v_d denote the principal components

Orthogonal and unit norm $v_i^T v_j = 0 \quad i \neq j$

$$\|v_i\|^2 = v_i^T v_i = 1$$

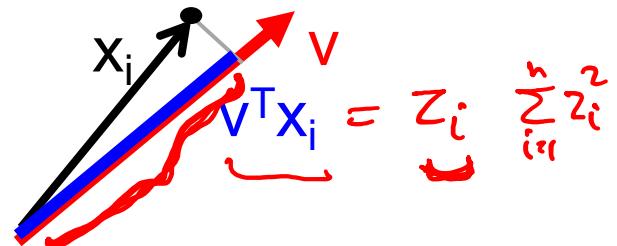
Find vector that maximizes sample variance of projection



$$\frac{1}{n} \sum_{i=1}^n (v^T x_i)^2 = \frac{v^T X X^T v}{n}$$

PCA

$$\max_v v^T X X^T v \quad \text{s.t.} \quad v^T v = 1$$



Pol:



- Is this a convex optimization problem?



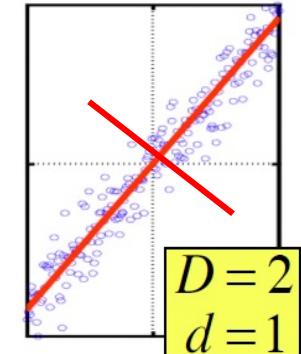
Principal Component Analysis (PCA)

Let v_1, v_2, \dots, v_d denote the principal components

Orthogonal and unit norm $v_i^T v_j = 0 \quad i \neq j$

$$v_i^T v_i = 1$$

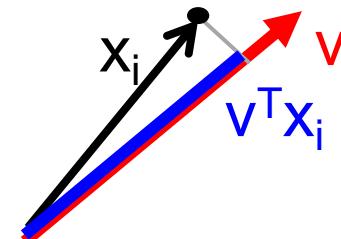
Find vector that maximizes sample variance of projection



$$\frac{1}{n} \sum_{i=1}^n (v^T x_i)^2 = \frac{v^T X X^T v}{n}$$

$$\max_v v^T X X^T v \quad \text{s.t.} \quad v^T v = 1 \quad \underline{\lambda}$$

Lagrangian: $\max_v v^T X X^T v - \lambda v^T v$



Wrap constraints into the objective function

$$\frac{\partial}{\partial v} = 0 \quad (X X^T - \lambda I) v = 0$$

$$\cancel{v^T v = 1} \quad \cancel{v^T X X^T v} = \lambda v^T v = \lambda$$

$$\Rightarrow \boxed{(X X^T) v = \lambda v}$$

Principal Component Analysis (PCA)

$$(XX^T)\mathbf{v} = \lambda \mathbf{v}$$

Therefore, \mathbf{v} is the eigenvector of sample covariance matrix XX^T

Sample variance of projection $\mathbf{v}^T \underbrace{XX^T}_{n} \mathbf{v} = \lambda \mathbf{v}^T \underbrace{\mathbf{v}}_{n} = \lambda / n$

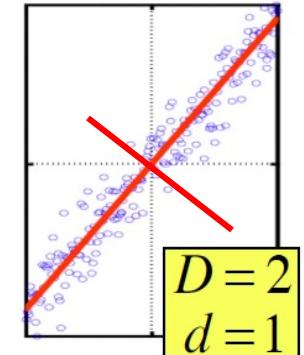
Thus, the eigenvalue λ denotes the amount of variability captured along that dimension (aka amount of energy along that dimension).

Eigenvalues $\lambda_1 > \lambda_2 > \lambda_3 > \dots$

The 1st Principal component \mathbf{v}_1 is the eigenvector of the sample covariance matrix XX^T associated with the largest eigenvalue λ_1

The 2nd Principal component \mathbf{v}_2 is the eigenvector of the sample covariance matrix XX^T associated with the second largest eigenvalue λ_2

And so on ...



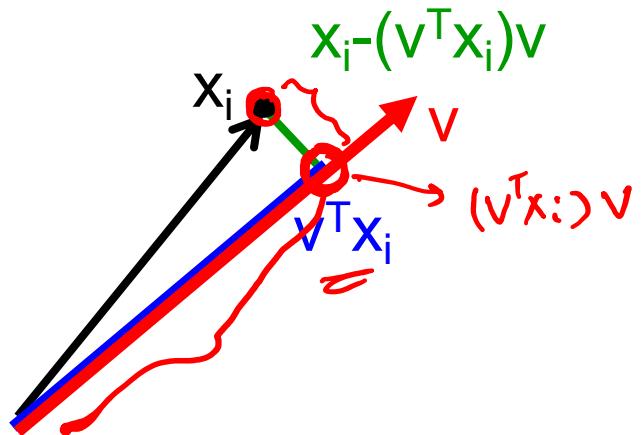
Another interpretation

Maximum Variance Subspace: PCA finds vectors v such that projections on to the vectors capture maximum variance in the data

$$\frac{1}{n} \sum_{i=1}^n (\mathbf{v}^T \mathbf{x}_i)^2 = \mathbf{v}^T \mathbf{X} \mathbf{X}^T \mathbf{v}$$

Minimum Reconstruction Error: PCA finds vectors v such that projection on to the vectors yields minimum MSE reconstruction

$$\frac{1}{n} \sum_{i=1}^n \|\mathbf{x}_i - (\mathbf{v}^T \mathbf{x}_i) \mathbf{v}\|^2$$



MSE in reconstructing the n points from projections onto v

$$\arg \min_v \frac{1}{n} \sum_{i=1}^n \|x_i - (v^T x_i)v\|^2$$

$$= \frac{1}{n} \sum_{i=1}^n \left[\underbrace{x_i^T x_i}_{\|x_i\|^2} + \underbrace{v^T (x_i^T v) (v^T x_i) v}_{(x_i^T v)^2} - \underbrace{2 x_i^T (v^T x_i) v}_{-2(x_i^T v)^2} \right]$$

$$\equiv \frac{1}{n} \sum_{i=1}^n - (x_i^T v)^2$$

$$\equiv \max_v \frac{1}{n} \sum_{i=1}^n (x_i^T v)^2$$
$$\frac{1}{n} v^T X^T X v$$

Dimensionality Reduction using PCA

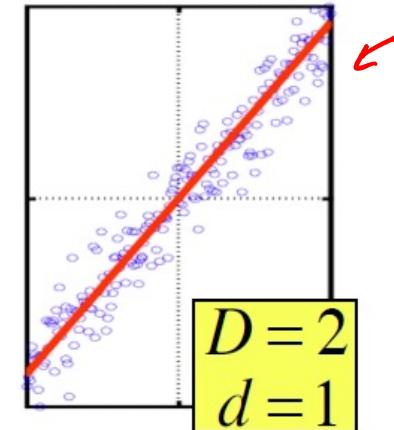
$$\mathbf{X} \mathbf{X}^T \quad D \times D$$

The eigenvalue λ denotes the amount of variability captured along that dimension.

Zero eigenvalues indicate no variability along those directions => data lies exactly on a linear subspace

Only keep data projections onto principal components with non-zero eigenvalues, say v_1, \dots, v_d where $d = \text{rank}(\mathbf{X} \mathbf{X}^T)$

$$d \leq D$$



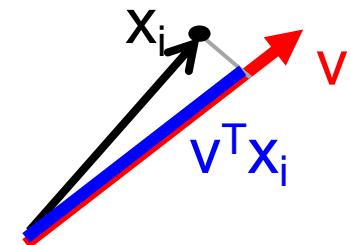
Original Representation
data point

$$\mathbf{x}_i = [x_i^1, x_i^2, \dots, x_i^D]^T \quad (D\text{-dimensional vector})$$

Transformed representation
projections

$$[v_1^T \mathbf{x}_i, v_2^T \mathbf{x}_i, \dots, v_d^T \mathbf{x}_i]^T \quad (d\text{-dimensional vector})$$

$$d < D$$

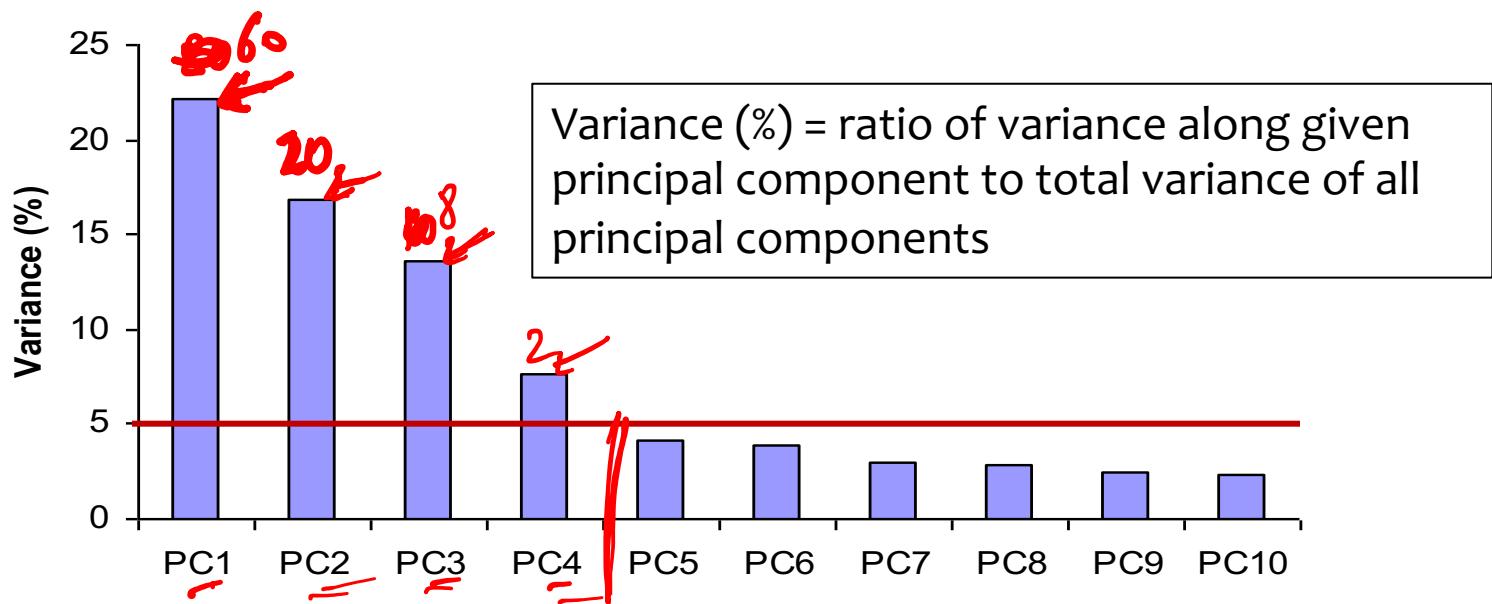


Dimensionality Reduction using PCA

In high-dimensional problem, data usually lies near a linear subspace, as noise introduces small variability

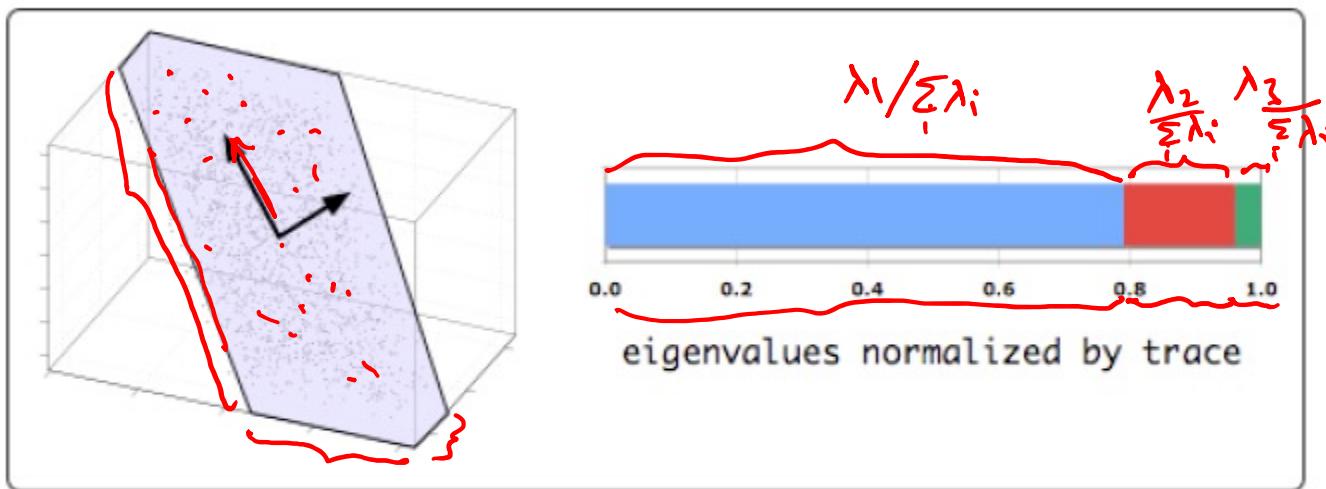
Only keep data projections onto principal components with **large** eigenvalues

Can *ignore* the components of lesser significance.



You might *lose some information*, but if the eigenvalues are small, you don't lose much

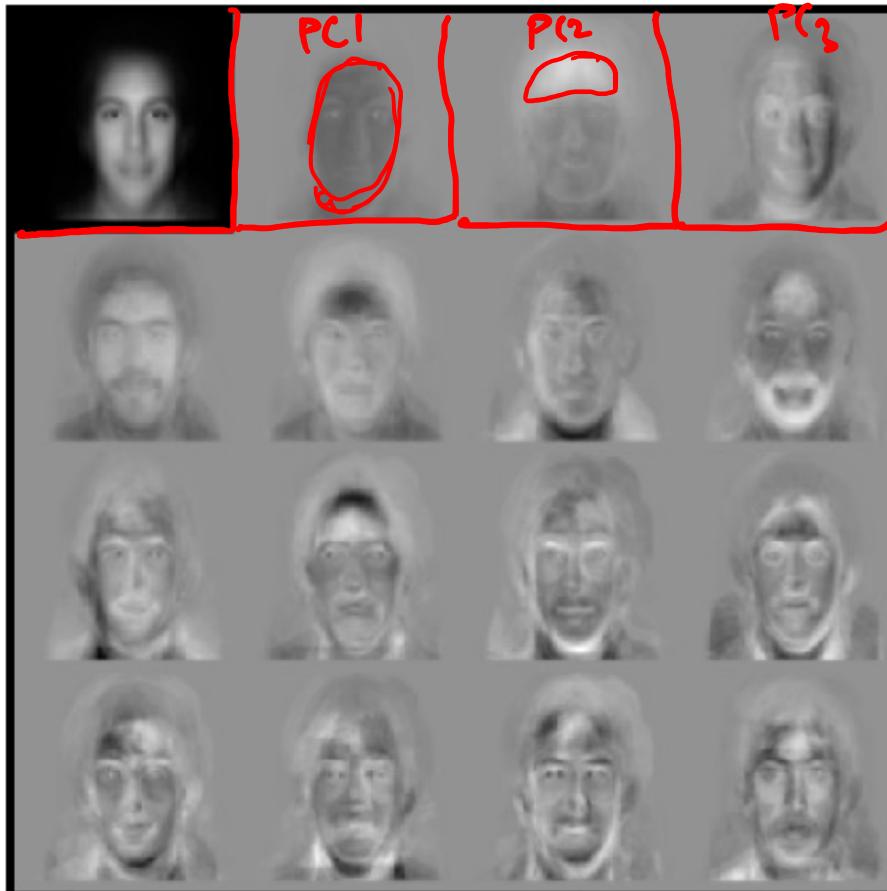
Example of PCA



Eigenvectors and eigenvalues of covariance matrix for $n=1600$ inputs in $d=3$ dimensions.

Example: faces

$$\begin{array}{l} \mathbf{X} \mathbf{X}^T \text{ } D \times D \\ \mathbf{v} \text{ } D \times 1 \end{array}$$



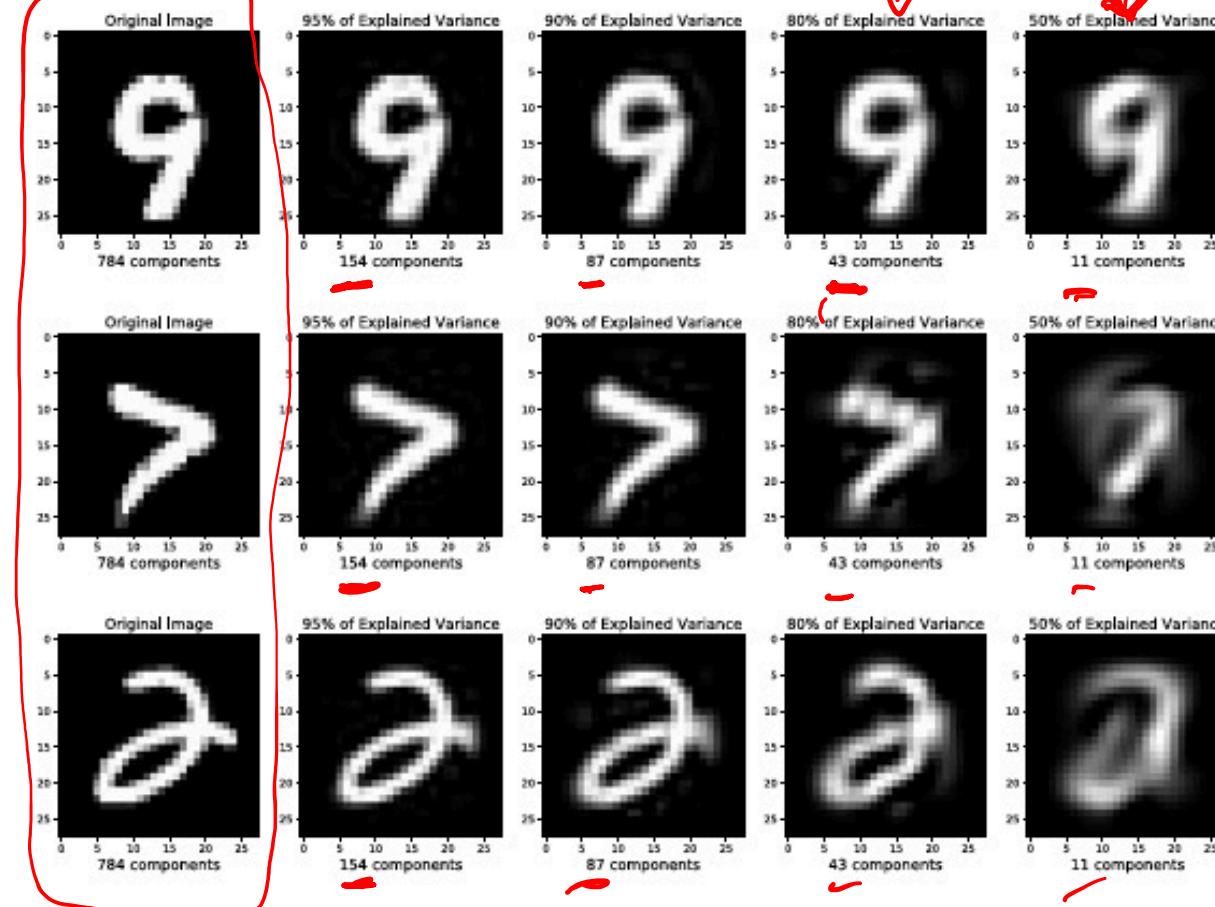
Eigenfaces
from 7562
images:

**top left image
is linear
combination
of rest.**

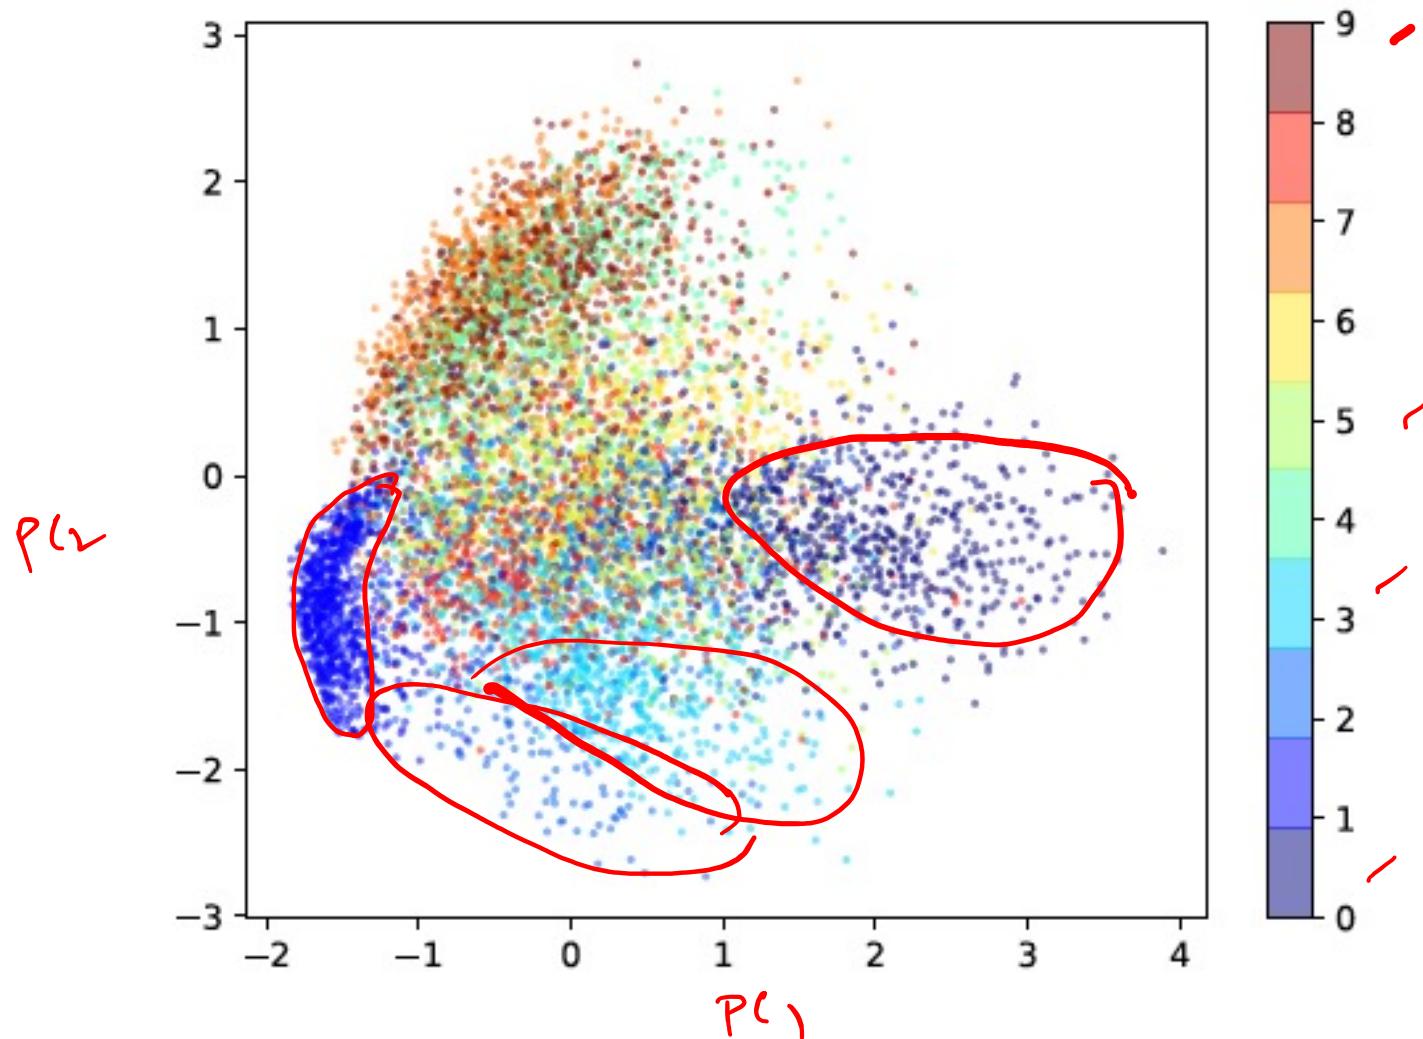
Sirovich & Kirby (1987)
Turk & Pentland (1991)

Example: MNIST digits

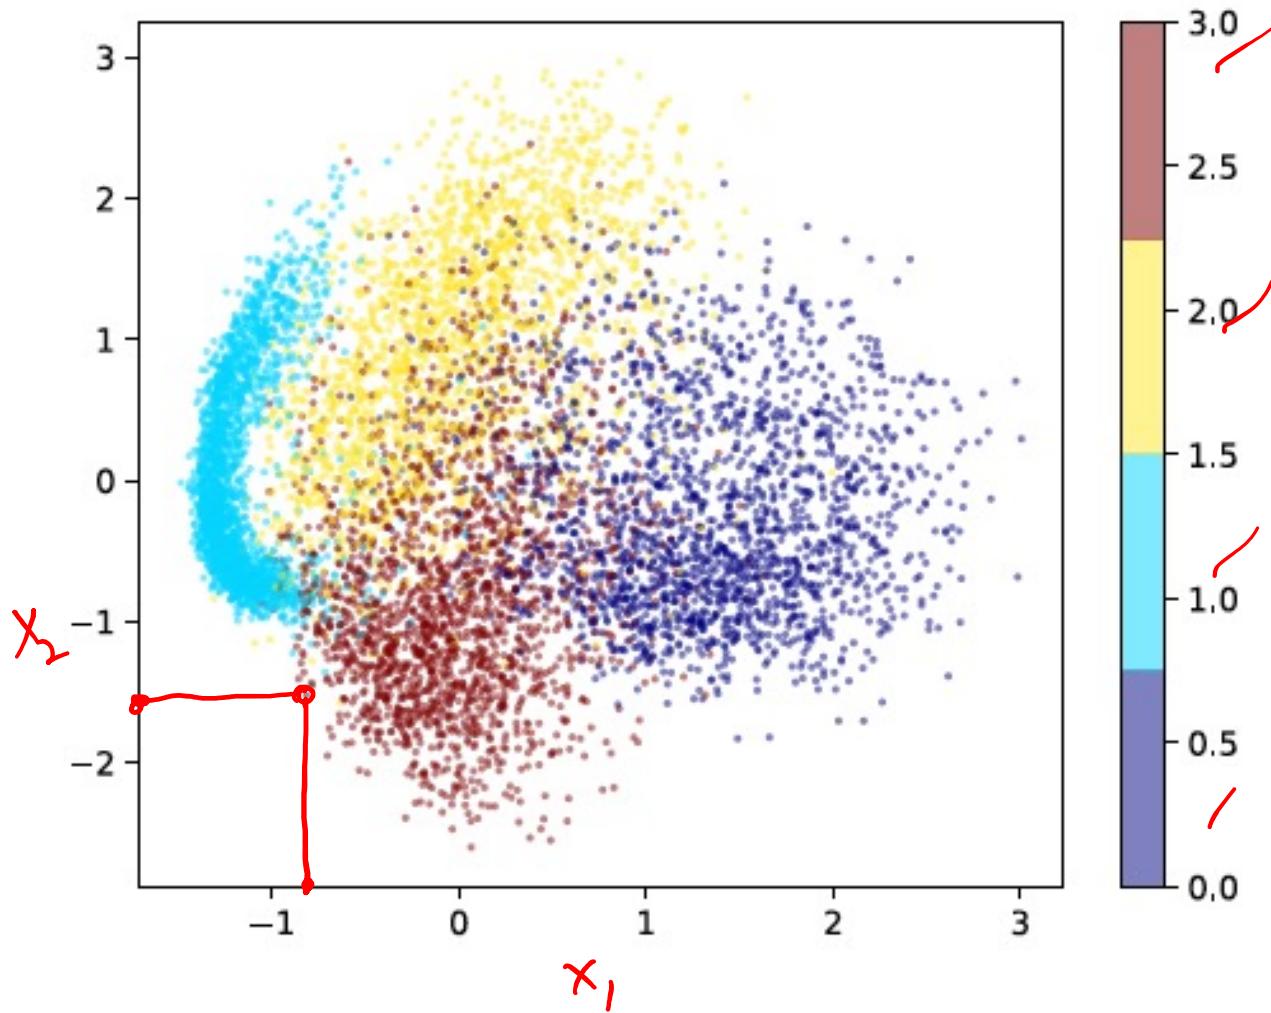
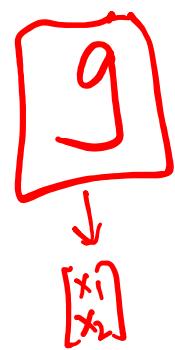
- 28×28 images = 784 PCA vectors
- Project to K dimensional space and then project back up



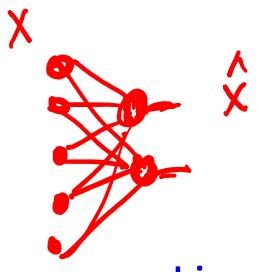
Projecting MNIST digits



Projecting MNIST digits



Unsupervised Dimensionality Reduction

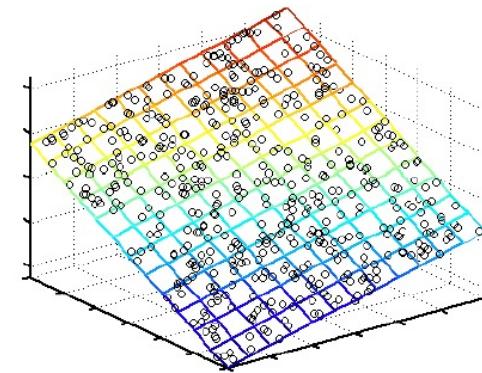


- Linear

→ **Principal Component Analysis (PCA)**

Factor Analysis

Independent Component Analysis (ICA)



- Nonlinear

Kernel PCA

Laplacian Eigenmaps, ISOMAP, LLE

→ Autoencoders

