CS 540 Introduction to Artificial Intelligence
Linear Algebra and PCA

University of Wisconsin-Madison

Spring 2026 Sections 1 & 2

Announcements

- **HW 1 released:**
  - Due **Wednesday Feb 4 at 11:59PM**

- Class roadmap:

Linear Algebra & PCA
Logic
NLP
Machine Learning Introduction

Mostly Foundations

Outline

- Basics: vectors, matrices, operations
- Dimensionality reduction
- Principal Components Analysis (PCA)

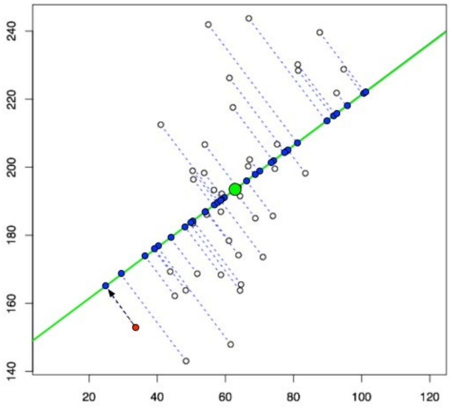

Lior Pachter

Linear Algebra

Review: Linear Algebra

- In AI/ML: building blocks for **all models**
  - e.g., linear regression; part of neural networks
- Vectors: list of values; point in space
- Matrices: Table of values; list of vectors

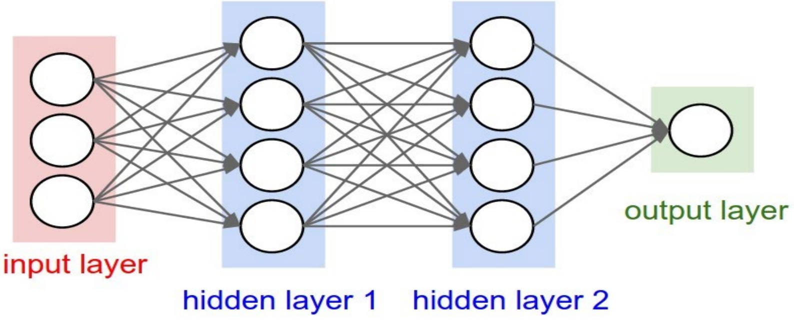

Stanford CS231n

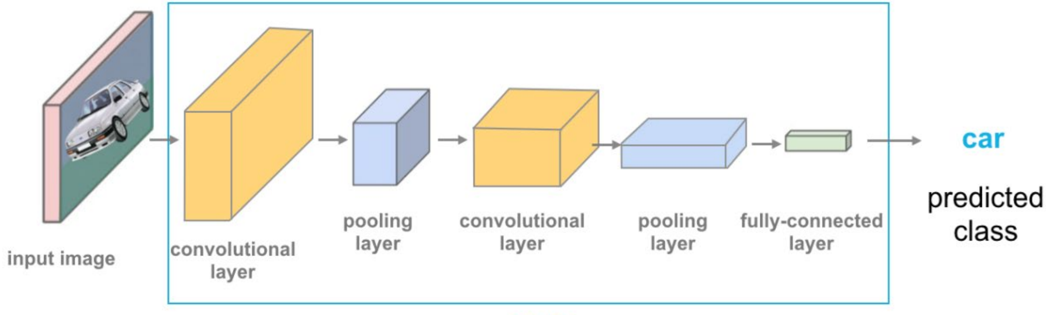

Cezanne Camacho

Matrix & Vector **Operations**

- **Matrix-Vector multiplication:**
  - **Linear transformation; plug in vector, get another vector**
  - Each entry in Ax is the inner product of a row of A with x

$$x \in \mathbb{R}^n, A \in \mathbb{R}^{m \times n}$$

$$Ax = \begin{bmatrix} \langle A_{1:}, x \rangle \\ \langle A_{2:}, x \rangle \\ \vdots \\ \langle A_{m:}, x \rangle \end{bmatrix} = \begin{bmatrix} A_{11}x_1 + A_{12}x_2 + \cdots + A_{1n}x_n \\ A_{21}x_1 + A_{22}x_2 + \cdots + A_{2n}x_n \\ \vdots \\ A_{m1}x_1 + A_{m2}x_2 + \cdots + A_{mn}x_n \end{bmatrix}$$

Matrix & Vector **Operations**

- Matrix multiplication
  - $A \in \mathbb{R}^{m\times n}, B \in \mathbb{R}^{n\times p}$, then $AB \in \mathbb{R}^{m\times p}$
  - "Composition" of linear transformations
  - Not commutative in general!

$$AB \neq BA$$

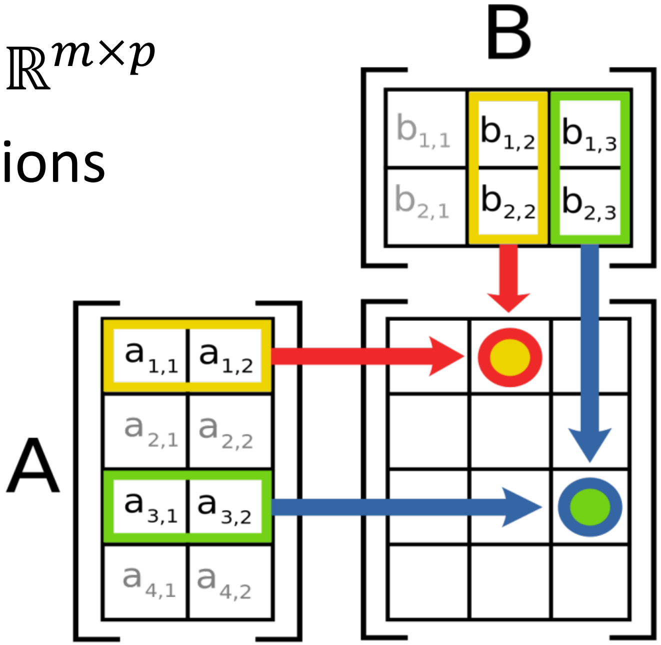

Wikipedia

Identity Matrix

- Like "1"
- Multiplying by it gets back the same matrix or vector

- Rows & columns are the "**standard basis vectors**" e_i

$$I = \begin{bmatrix} 1 & 0 & \dots & 0 \\ 0 & 1 & \dots & 0 \\ \vdots & \vdots & \dots & \vdots \\ 0 & 0 & \dots & 1 \end{bmatrix}$$

$e_1 \quad e_2 \quad e_n$

Break & Quiz

- **Q 3.2**: What is $\begin{bmatrix} 1 & 2 \\ 3 & 1 \\ 1 & 1 \end{bmatrix} \times \begin{bmatrix} 0 \\ 1 \end{bmatrix}$?
- A. $[-1\ 1\ 1]^T$
- B. $[2\ 1\ 1]^T$
- C. $[1\ 3\ 1]^T$
- D. $[1.5\ 2\ 1]^T$

Break & Quiz

- **Q 3.2**: What is $\begin{bmatrix} 1 & 2 \\ 3 & 1 \\ 1 & 1 \end{bmatrix} \times \begin{bmatrix} 0 \\ 1 \end{bmatrix}$?
- A. $[-1\ 1\ 1]^T$
- **B. $[2\ 1\ 1]^T$**
- C. $[1\ 3\ 1]^T$
- D. $[1.5\ 2\ 1]^T$

Break & Quiz

- **Q 3.2**: What is $\begin{bmatrix} 1 & 2 \\ 3 & 1 \\ 1 & 1 \end{bmatrix} \times \begin{bmatrix} 0 \\ 1 \end{bmatrix}$?
- A. $[-1\ 1\ 1]^T$
- **B. $[2\ 1\ 1]^T$**
- C. $[1\ 3\ 1]^T$
- D. $[1.5\ 2\ 1]^T$

Check dimensions: answer must be 3 x 1 matrix (i.e., column vector).

$$\begin{bmatrix} 1 & 2 \\ 3 & 1 \\ 1 & 1 \end{bmatrix} \times \begin{bmatrix} 0 \\ 1 \end{bmatrix} = \begin{bmatrix} 0 * 1 + 1 * 2 \\ 0 * 3 + 1 * 1 \\ 0 * 1 + 1 * 1 \end{bmatrix} = \begin{bmatrix} 2 \\ 1 \\ 1 \end{bmatrix}$$

Break & Quiz

- **Q 3.3**: Given matrices $A \in \mathbb{R}^{m \times n}, B \in \mathbb{R}^{d \times m}, C \in \mathbb{R}^{p \times n}$ What are the dimensions of BAC^T

- A. *n x p*
- B. *d x p*
- C. *d x n*
- D. Undefined

Break & Quiz

- **Q 3.3**: Given matrices $A \in \mathbb{R}^{m \times n}, B \in \mathbb{R}^{d \times m}, C \in \mathbb{R}^{p \times n}$

What are the dimensions of BAC^T

- A. *n x p*
- **B. *d x p***
- C. *d x n*
- D. Undefined

Break & Quiz

- **Q 3.3**: Given matrices $A \in \mathbb{R}^{m \times n}, B \in \mathbb{R}^{d \times m}, C \in \mathbb{R}^{p \times n}$

What are the dimensions of BAC^T

- A. *n x p*
- **B. *d x p***
- C. *d x n*
- D. Undefined

To rule out (D), check that for each pair of adjacent matrices XY, the # of columns of X = # of rows of Y

Then, B has d rows so solution must have d rows. C^T has p columns so solution has p columns.

Break & Quiz

- **Q 3.4**: A and B are matrices, neither of which is the identity. Is $AB = BA$?

- A. Never
- B. Always
- C. Sometimes

Break & Quiz

- **Q 3.4**: A and B are matrices, neither of which is the identity. Is $AB = BA$?

- A. Never
- B. Always
- **C. Sometimes**

Break & Quiz

- **Q 3.4**: A and B are matrices, neither of which is the identity. Is $AB = BA$?

- A. Never
- B. Always
- **C. Sometimes**

Matrix multiplication is not necessarily commutative.

Matrix Inverses

- If for *A* there is a *B* such that $AB = BA = I$
  - Then *A* is invertible/nonsingular, B is its inverse
  - Some matrices are **not** invertible!

  - Usual notation: A^{-1}

$$\begin{bmatrix} 1 & 1 \\ 2 & 3 \end{bmatrix} \times \begin{bmatrix} 3 & -1 \\ -2 & 1 \end{bmatrix} = I$$

Eigenvalues & Eigenvectors

- For a square matrix *A*, solutions to $Av = \lambda v$
  - *v* (nonzero) is a vector: **eigenvector**
  - λ is a scalar: **eigenvalue**

  - Intuition: A is a linear transformation;
  - Can stretch/rotate vectors;
  - E-vectors: only stretched (by e-vals)

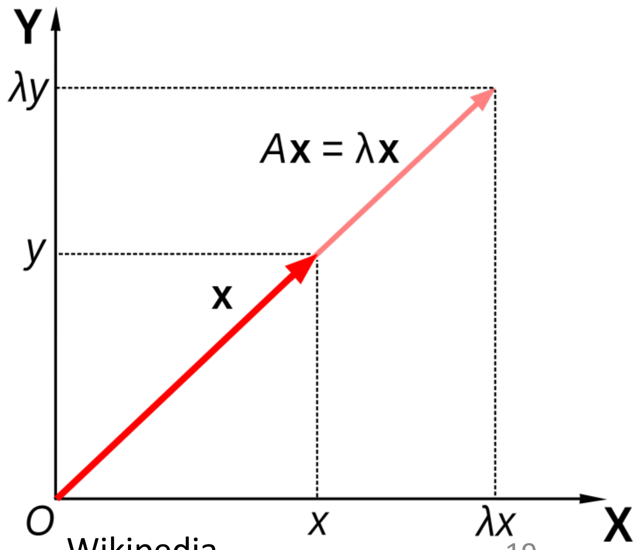

Wikipedia

Outline

- Basics: vectors, matrices, operations
- **Dimensionality reduction**
- Principal Components Analysis (PCA)

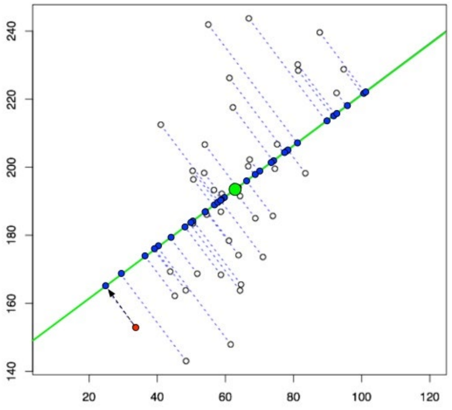

Lior Pachter

Dimensionality Reduction

- Vectors store features. Lots of features!
  - Document classification: thousands of words per doc
  - Netflix surveys: 480189 users x 17770 movies
  - MEG Brain Imaging: 120 locations x 500 time points x 20 objects

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

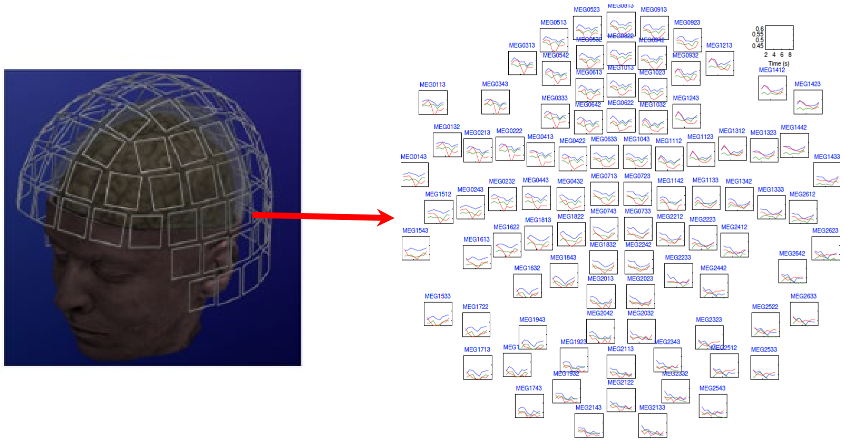

Dimensionality Reduction

Reduce dimensions

- Why?
  - Lots of features redundant
  - Storage & computation costs

CreativeBlog

- Goal: take $x \in \mathbb{R}^d \rightarrow x \in \mathbb{R}^r$ for $r << d$
  - But minimize information loss

Dimensionality Reduction

Examples: 3D to 2D

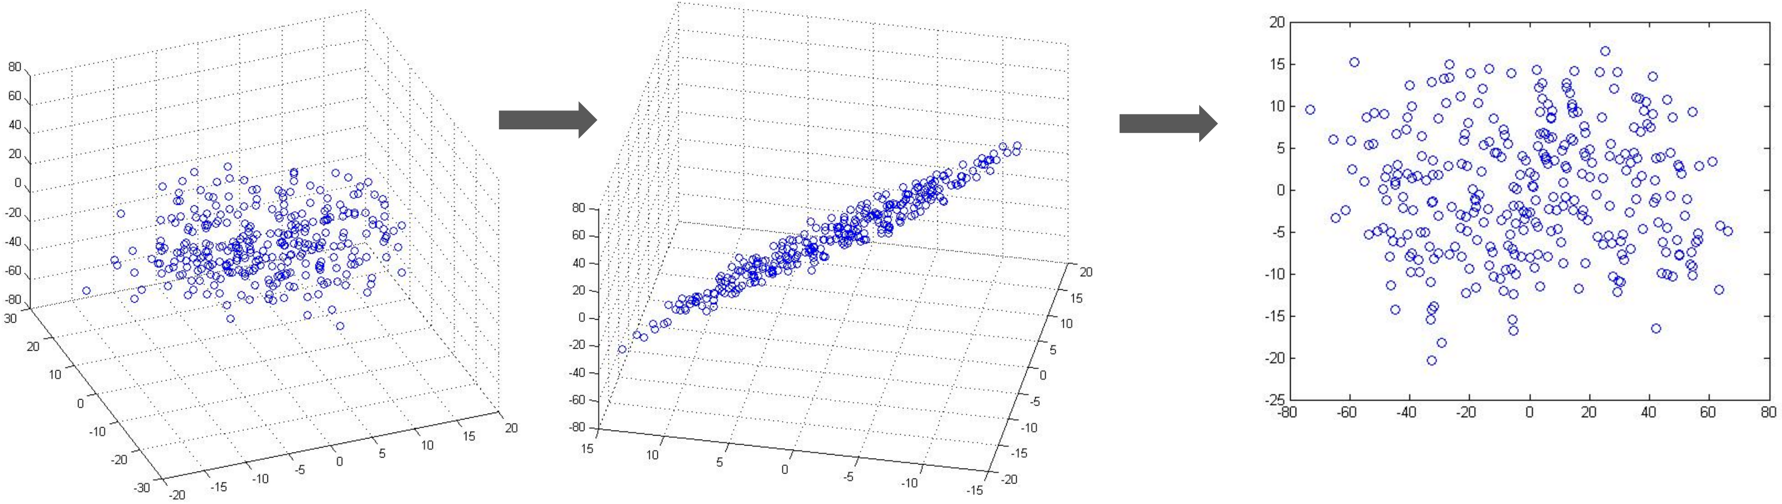

Andrew Ng

Break & Quiz

Q 2.1: What is the inverse of $A = \begin{bmatrix} 0 & 2 \\ 3 & 0 \end{bmatrix}$

A: $A^{-1} = \begin{bmatrix} -3 & 0 \\ 0 & -2 \end{bmatrix}$

B: $A^{-1} = \begin{bmatrix} 0 & \frac{1}{3} \\ \frac{1}{2} & 0 \end{bmatrix}$

C: Undefined / *A* is not invertible

Break & Quiz

Q 2.1: What is the inverse of $A = \begin{bmatrix} 0 & 2 \\ 3 & 0 \end{bmatrix}$

A: $A^{-1} = \begin{bmatrix} -3 & 0 \\ 0 & -2 \end{bmatrix}$

B: $A^{-1} = \begin{bmatrix} 0 & \frac{1}{3} \\ \frac{1}{2} & 0 \end{bmatrix}$

C: Undefined / *A* is not invertible

$$AA^{-1} = \begin{bmatrix} 0 & 2 \\ 3 & 0 \end{bmatrix} \begin{bmatrix} a & b \\ c & d \end{bmatrix} = \begin{bmatrix} 0*a+c*2 & 0*b+2*d \\ 3*a+c*0 & 3*b+0*d \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

$$\begin{aligned} 2c &= 1 \\ 3a &= 0 \\ 2d &= 0 \\ 3b &= 1 \end{aligned}$$

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix} = \begin{bmatrix} 0 & 1/3 \\ 1/2 & 0 \end{bmatrix}$$

Break & Quiz

Q 2.2: What are the eigenvalues of $A = \begin{bmatrix} 2 & 0 & 0 \\ 0 & 5 & 0 \\ 0 & 0 & 1 \end{bmatrix}$

A. -1, 2, 4
B. 0.5, 0.2, 1.0
C. 0, 2, 5
D. 2, 5, 1

Break & Quiz

Q 2.2: What are the eigenvalues of $A = \begin{bmatrix} 2 & 0 & 0 \\ 0 & 5 & 0 \\ 0 & 0 & 1 \end{bmatrix}$

A. -1, 2, 4
B. 0.5, 0.2, 1.0
C. 0, 2, 5
D. 2, 5, 1

Break & Quiz

Q 2.2: What are the eigenvalues of $A = \begin{bmatrix} 2 & 0 & 0 \\ 0 & 5 & 0 \\ 0 & 0 & 1 \end{bmatrix}$

A. -1, 2, 4
B. 0.5, 0.2, 1.0
C. 0, 2, 5
D. 2, 5, 1

Solution #1: You may recall from a linear algebra course that the eigenvalues of a diagonal matrix (in which only diagonal entries are non-zero) are just the entries along the diagonal. Hence D is the correct answer.

Break & Quiz

Q 2.2: What are the eigenvalues of $A = \begin{bmatrix} 2 & 0 & 0 \\ 0 & 5 & 0 \\ 0 & 0 & 1 \end{bmatrix}$

Solution #2: Use the definition of eigenvectors and values: $Av = \lambda v$

$$\begin{bmatrix} 2 & 0 & 0 \\ 0 & 5 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{matrix} v_1 \\ v_2 \\ v_3 \end{matrix} = \begin{bmatrix} 2v_1 + 0v_2 + 0v_3 \\ 0v_1 + 5v_2 + 0v_3 \\ 0v_1 + 0v_2 + 1v_3 \end{bmatrix} = \begin{bmatrix} 2v_1 \\ 5v_2 \\ v_3 \end{bmatrix} = \begin{bmatrix} \lambda v_1 \\ \lambda v_2 \\ \lambda v_3 \end{bmatrix}$$

A. -1, 2, 4
B. 0.5, 0.2, 1.0
C. 0, 2, 5
D. 2, 5, 1

Since A is a 3x3 matrix, A has 3 eigenvalues and so there are 3 combinations of values for λ and v that will satisfy the above equation. The simple form of the equations suggests starting by checking each of the standard basis vectors* as v and then solving for λ. Doing so gives D as the correct answer.

*Each standard basis vector $e_i \in \mathbb{R}^n$ is the vector in which all components are zero except component i is 1.

Break & Quiz

Q 2.3: Suppose we are given a dataset with *n*=10000 samples with 100-dimensional binary feature vectors. Our storage device has a capacity of 50000 bits. What's the lowest compression ratio we can use?

A. 20X
B. 100X
C. 5X
D. 1X

Break & Quiz

Q 2.3: Suppose we are given a dataset with n=10000 samples with 100-dimensional binary feature vectors. Our storage device has a capacity of 50000 bits. What's the lower compression ratio we can use?

A. 20X
B. 100X
C. 5X
D. 1X

Break & Quiz

Q 2.3: Suppose we are given a dataset with n=10000 samples with 100-dimensional binary feature vectors. Our storage device has a capacity of 50000 bits. What's the lower compression ratio we can use?

A. **20X**
B. 100X
C. 5X
D. 1X

50,000 bits / 10,000 samples means compressed version must have 5 bits / sample.

Dataset has 100 bits / sample.

Must compress 20x smaller to fit on device.

Principal Components Analysis (**PCA**)

A type of dimensionality reduction approach

→For when data is **approximately lower dimensional**

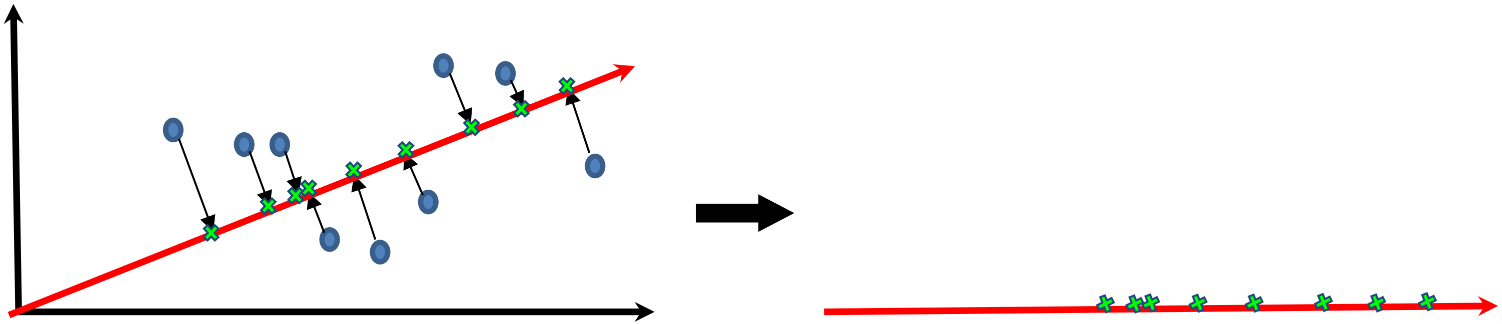

2D → 1D

Principal Components Analysis (**PCA**)

- Find axes $u_1, u_2, \ldots, u_n \in \mathbb{R}^d$ of a subspace
  - Will project to this subspace

- These vectors are the **principal components**

- Want to preserve data
  - Minimize projection error

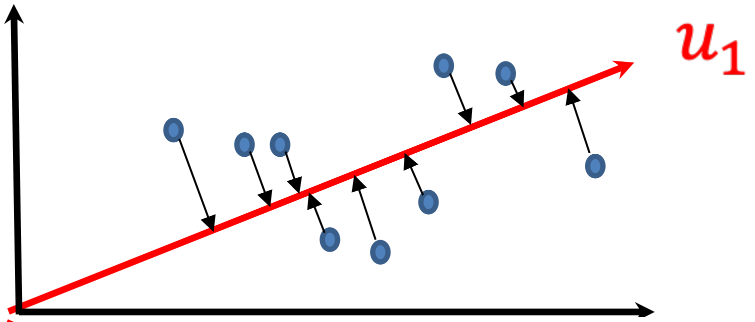

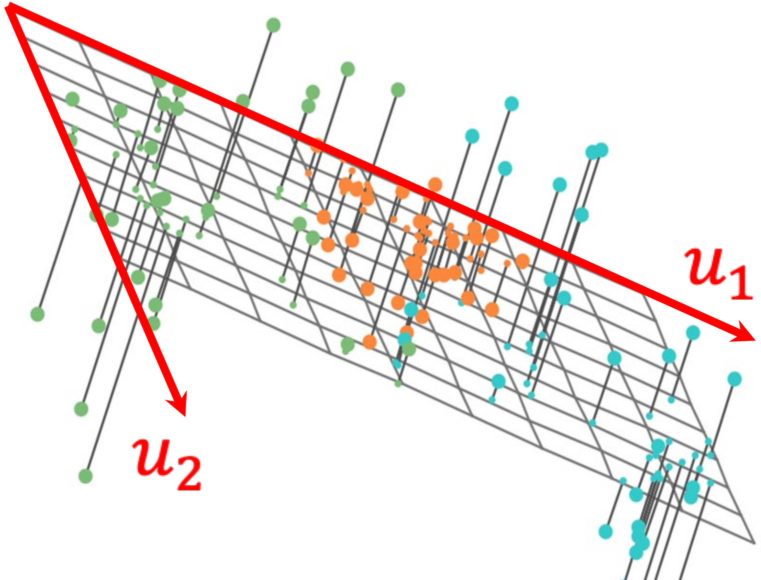

Projection: An Example

$x \in \mathbb{R}^2$

$d^2 = p^2 + e^2$

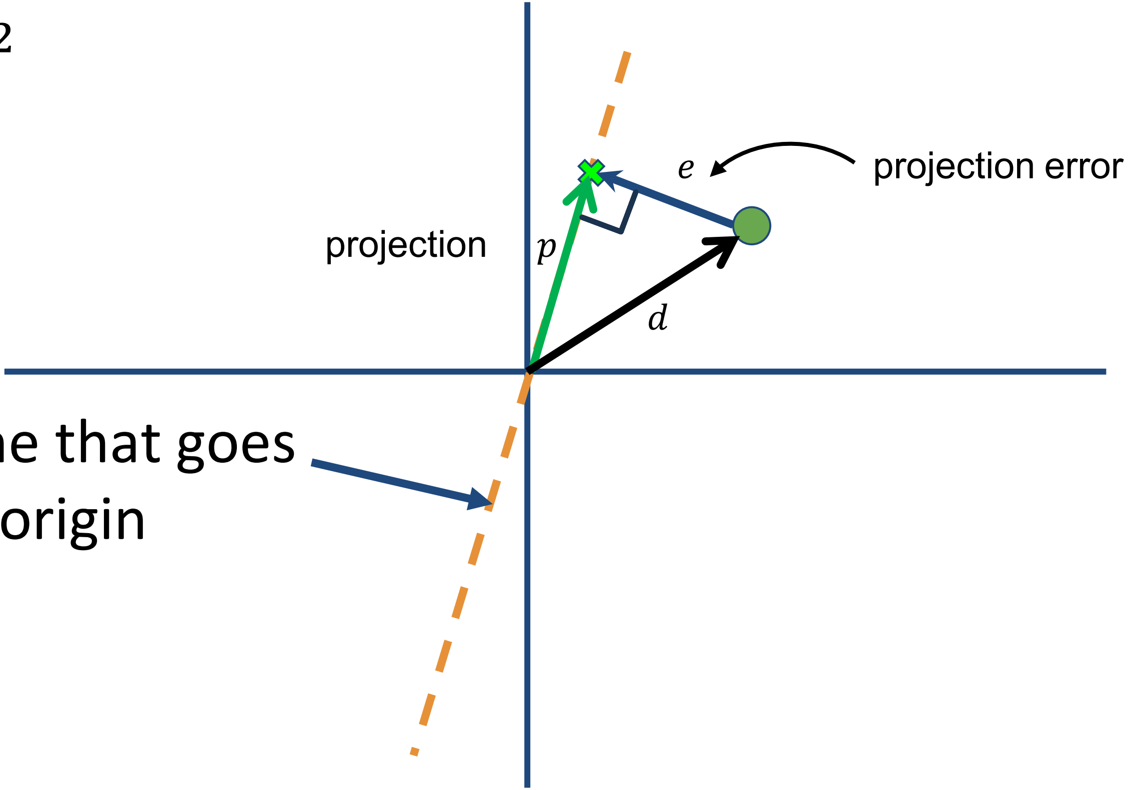

A random line that goes through the origin

Projection: An Example

$x \in \mathbb{R}^2$

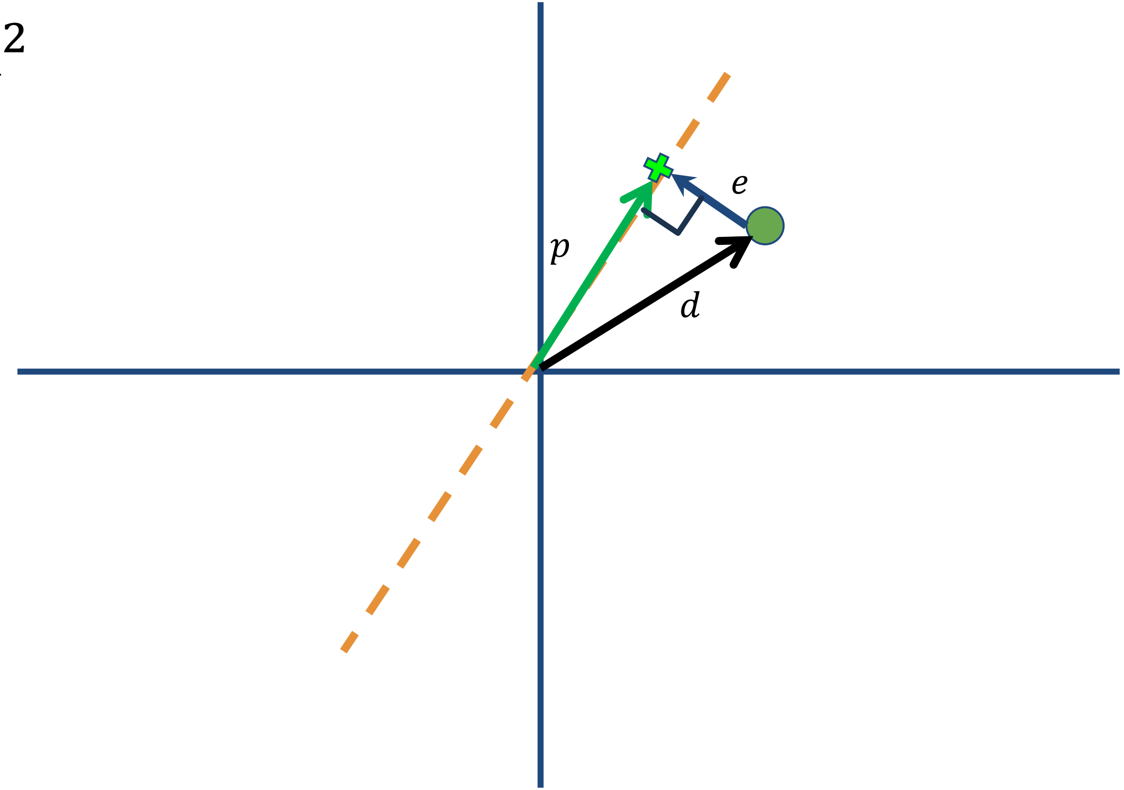

$d^2 = p^2 + e^2$

Projection: An Example

$x_1, x_2, \ldots, x_n \in \mathbb{R}^2$

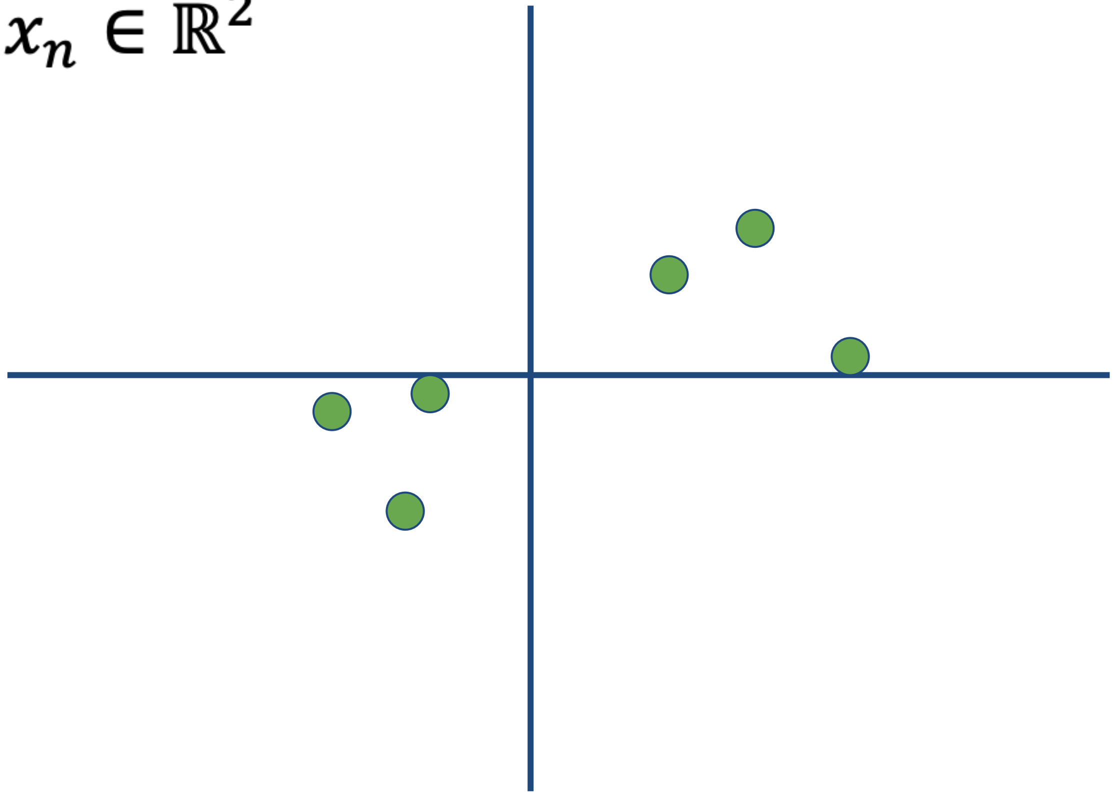

Projection: An Example

$x_1, x_2, \ldots, x_n \in \mathbb{R}^2$

A random line that goes through the origin

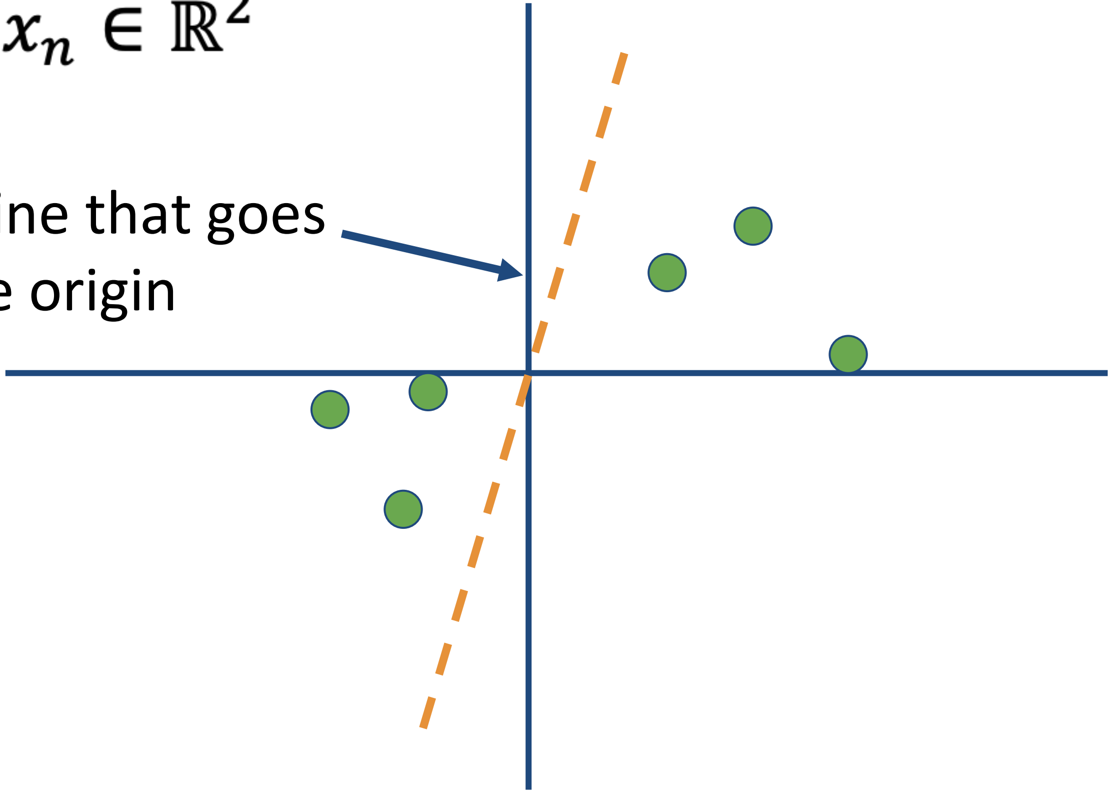

Projection: An Example

$x_1, x_2, \ldots, x_n \in \mathbb{R}^2$

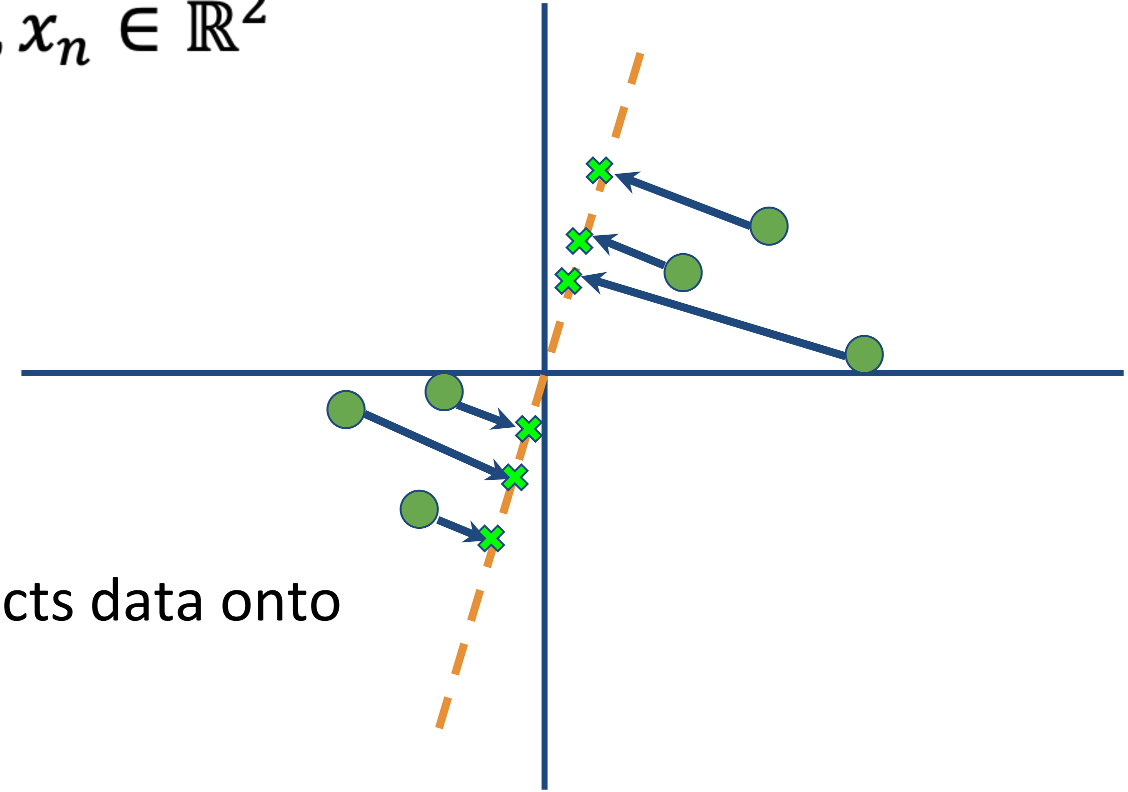

PCA projects data onto this line

Projection: An Example

$x_1, x_2, \ldots, x_n \in \mathbb{R}^2$

Goal: finding a line that minimizes the sum of squared distances/projection errors.

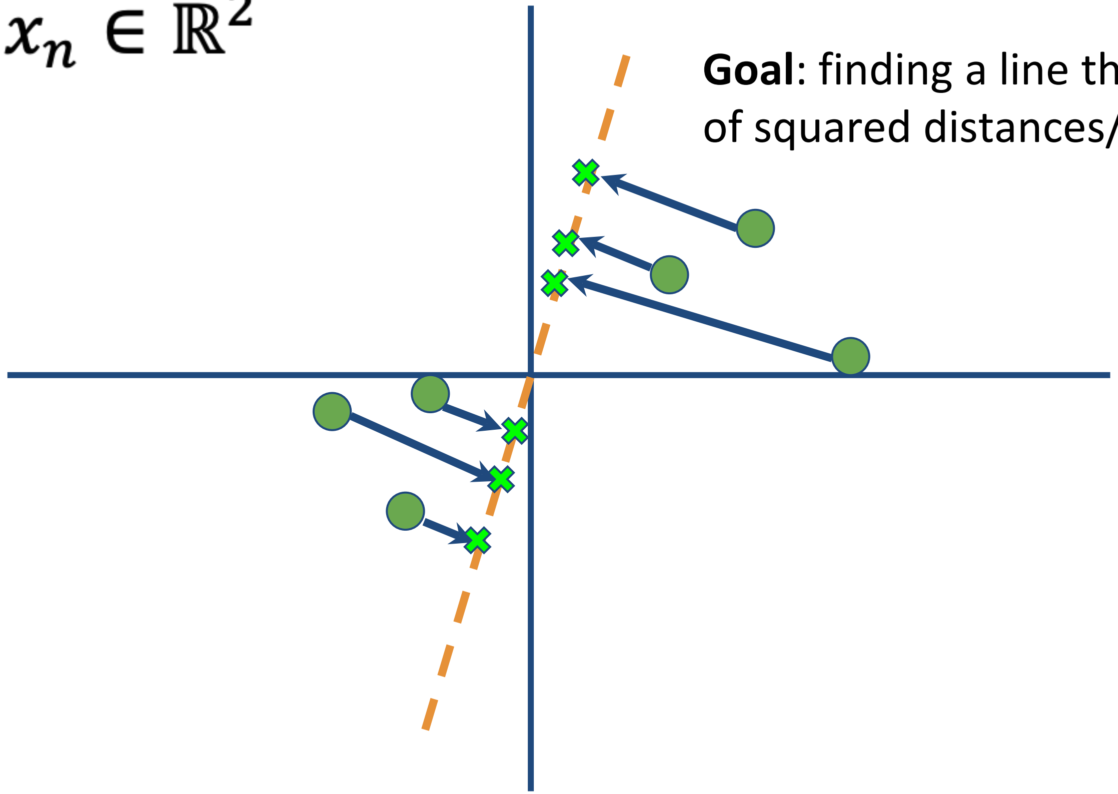

Projection: An Example

$x_1, x_2, \ldots, x_n \in \mathbb{R}^2$

Goal equivalent to **maximizing the variance of projected datapoints**.

$$Variance = \frac{1}{n-1}\sum_{i=1}^{n}\langle u, x_i\rangle^2$$

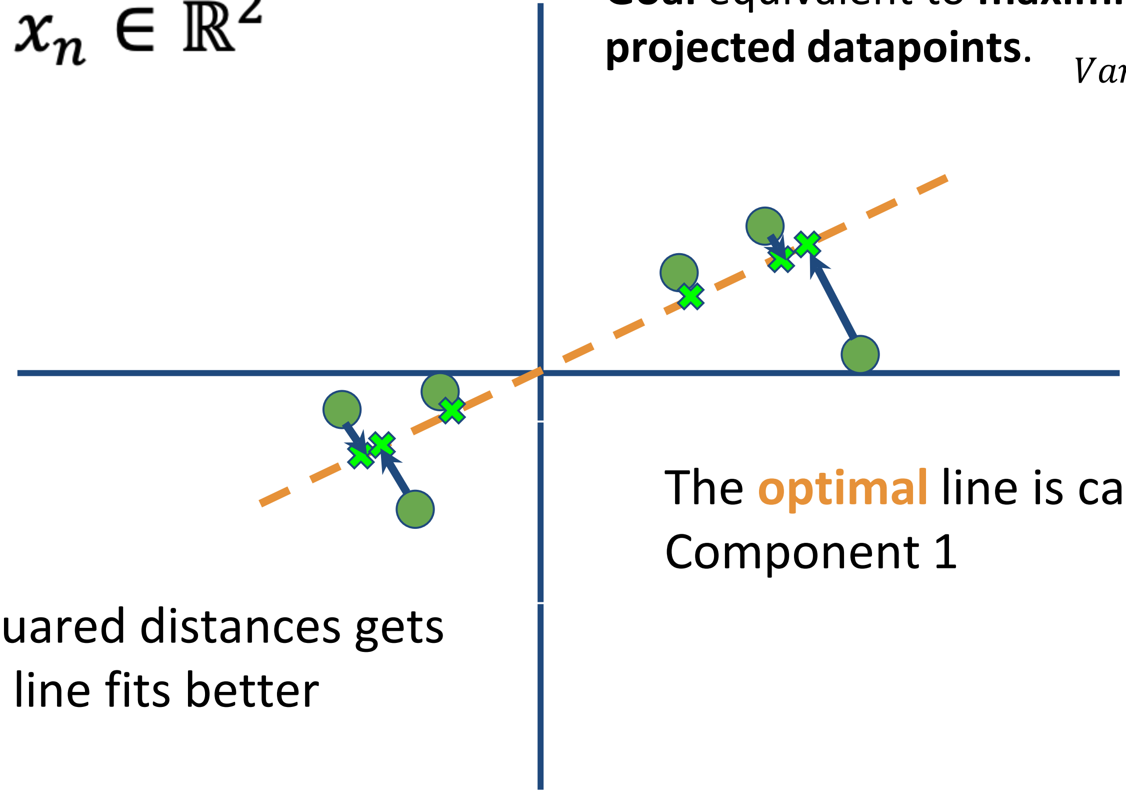

The **optimal** line is called Principal Component 1

The sum of squared distances gets smaller as the line fits better

PCA Procedure

Input: data $x_1, x_2, \dots \dots x_n \in \mathbb{R}^d$

Step 1: Center data at the origin so that $\hat{\mu} = 0$

Step 1.1: Calculate the mean $\hat{\mu} = \frac{1}{n}\sum_{i=1}^{n} x_i$

Step 1.2: Subtract the mean from the datapoints

$$x_1 = x_1 - \hat{\mu}$$

$$x_2 = x_2 - \hat{\mu}$$

....

$$x_n = x_n - \hat{\mu}$$

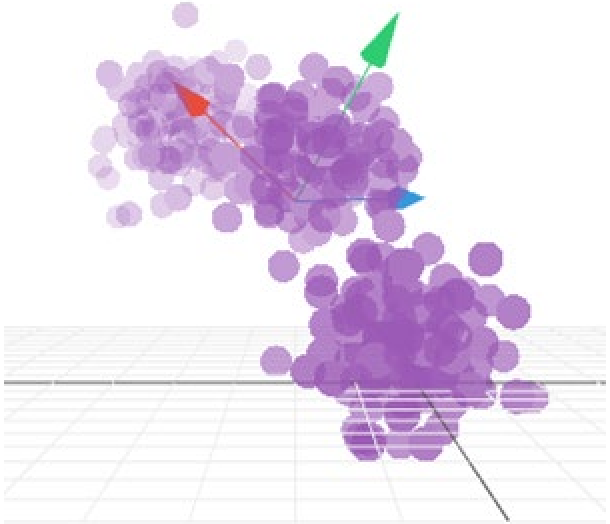

Victor Powell

PCA Procedure

Objective: $u = arg \max_{\|u\|=1} \frac{1}{n-1} \sum_{i=1}^{n} \langle u, x_i \rangle^2$

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

$$\underbrace{u^T \left(\frac{1}{n-1} \sum_{i=1}^{n} x_i x_i^T \right) u}_{\text{Covariance Matrix}}$$

PCA Procedure

Output: principal components $u_1, u_2, \dots \dots u_m \in \mathbb{R}^d$

The $\boldsymbol{m}$ **principal components** are:

- orthogonal
- the **top-$\boldsymbol{m}$ eigenvectors** of the covariance matrix $\left(S = \frac{1}{n-1} \sum_{i=1}^{n} x_i x_i^T\right)$

Step 2: Compute **covariance matrix**
Step 3: Compute **eigenvectors** and **eigenvalues** of covariance matrix
Step 4: **Sort eigenvalues** from high to low
Step 5: Select **top-$\boldsymbol{m}$ eigenvectors**

PCA Projection and Reconstruction

projection

For each x_i: $x_i^{pca} = \sum_{j=1}^{m} (u_j^T x_i) u_j$

reconstruction

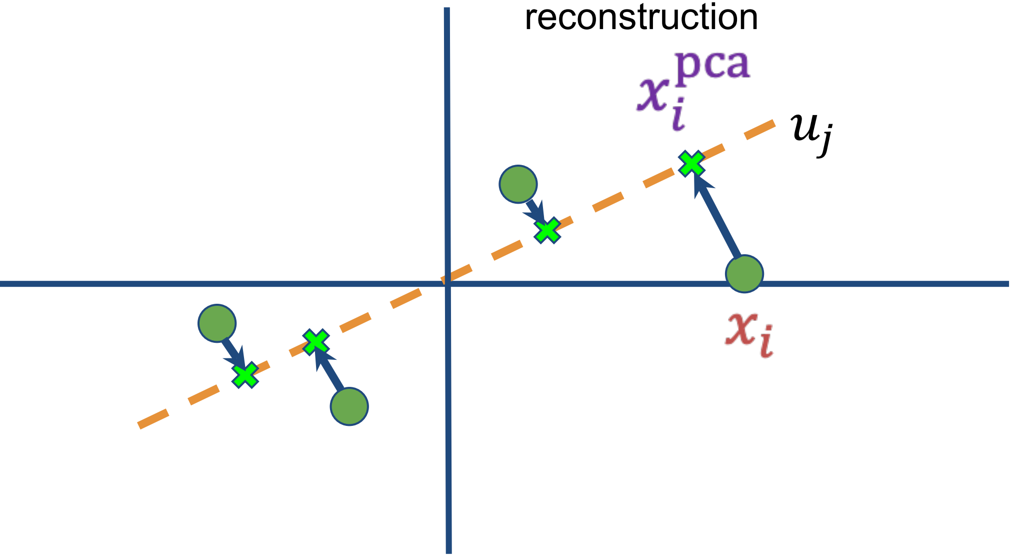

2D → 1D

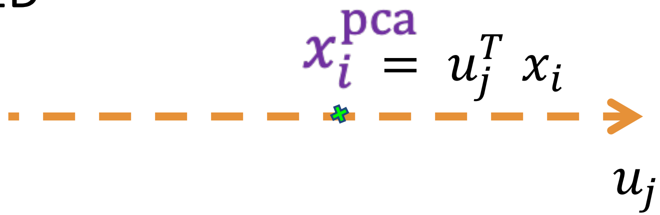

Many Variations

- PCA, Kernel PCA, ICA, CCA
  - Extract structure from high dimensional dataset
- Uses:
  - Visualization
  - Efficiency
  - Noise removal
  - Downstream machine learning use

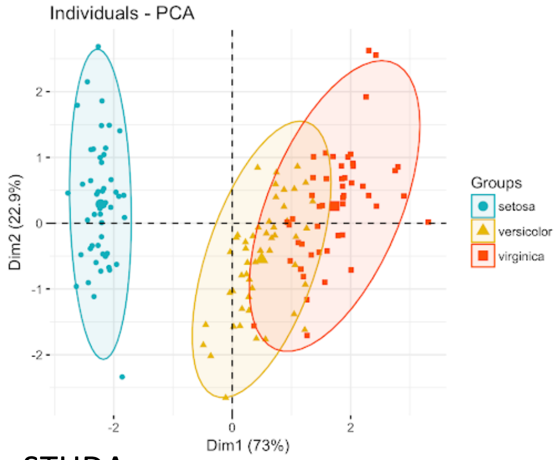

STHDA

Application: Image Compression

- Start with image; divide into 12x12 patches
  - That is, 144-D vector
  - **Original image:**

Application: Image Compression

- 6 principal components (as an image)

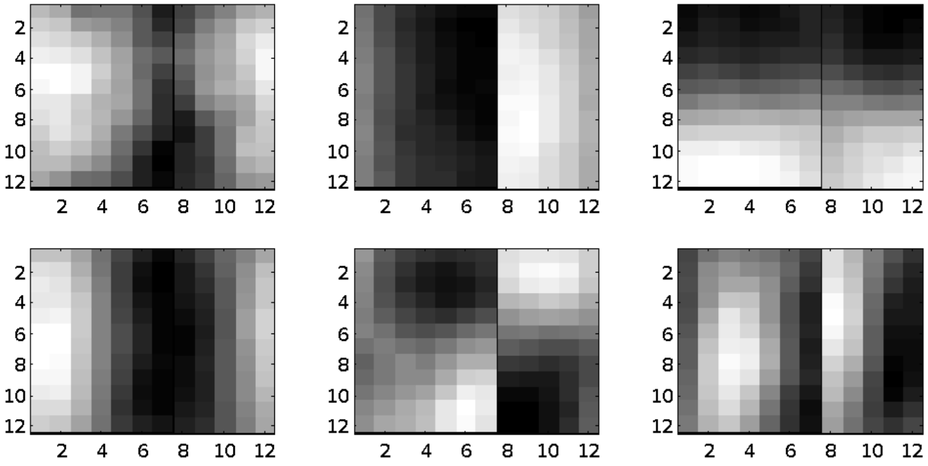

Application: Image Compression

- Project to 6D

Compressed

Original

Application: Exploratory Data Analysis

- [**Novembre et al. ’08**]: Take top two singular vectors of people x SNP matrix (POPRES)

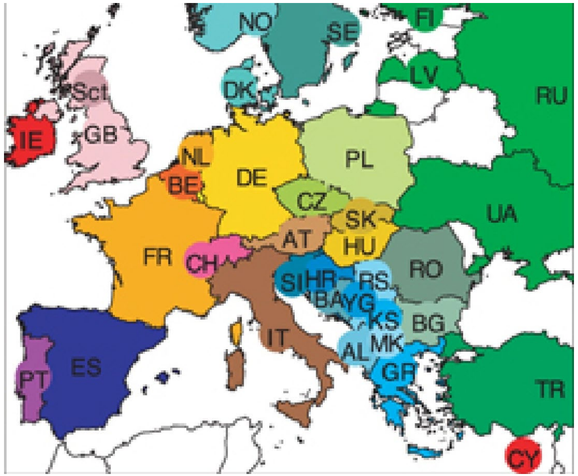

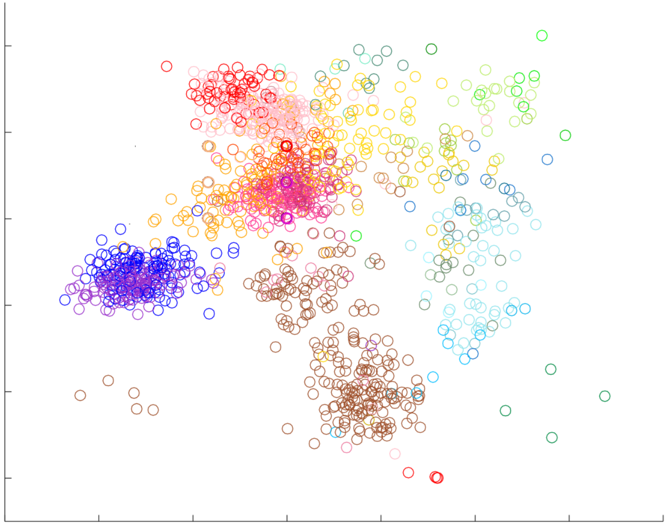

“Genes Mirror Geography in Europe”

Readings

- Vast literature on linear algebra.
- Local class: **Math 341**
- More on PCA (and other matrix methods in ML): **CS 532**

- **Suggested reading:**
  - Textbook: Artificial Intelligence: A Modern Approach (4th edition). Stuart Russell and Peter Norvig. Pearson, 2020. Appendix A
  - Lecture notes on PCA by Roughgarden and Valiant https://web.stanford.edu/class/cs168/l/l7.pdf