Lecture 10: Point Clouds, Eigenvectors, PCA

COMPSCI/MATH 290-04

Chris Tralie, Duke University

2/16/2016

Announcements

- Group Assignment 1 Part 1, Due This Sunday 11:55 PM
- Note on Assignment Difficulty
- Hackathon 2/26, 2/27, or 2/28 (upcoming poll)
- Choosing final projects before spring break
- Visitors Coming...

Table of Contents

What Is A Point Cloud?

Interactive Demo

What's In This Unit?

Point Cloud Registration (Alignment)

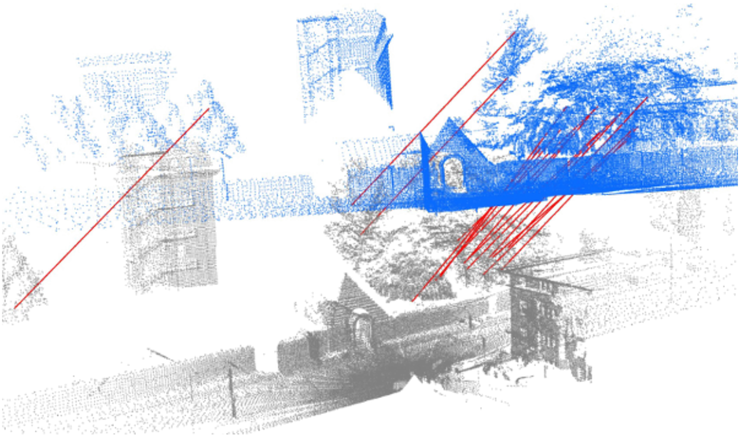

Courtesy of `http://www.pointclouds.org`

What's In This Unit?

"Shape Google" (Point Cloud Statistics)

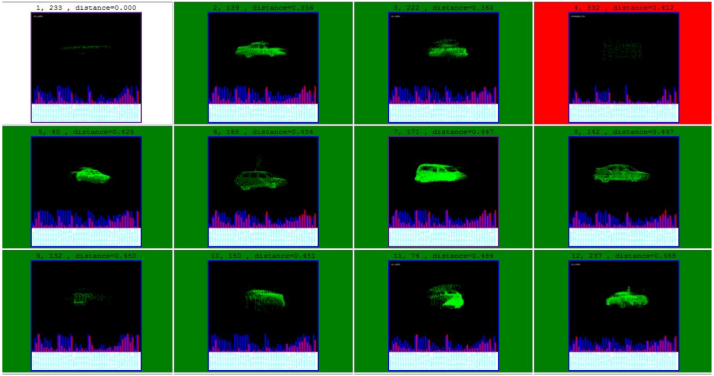

(my undergrad senior thesis)

What's In This Unit?

Symmetries

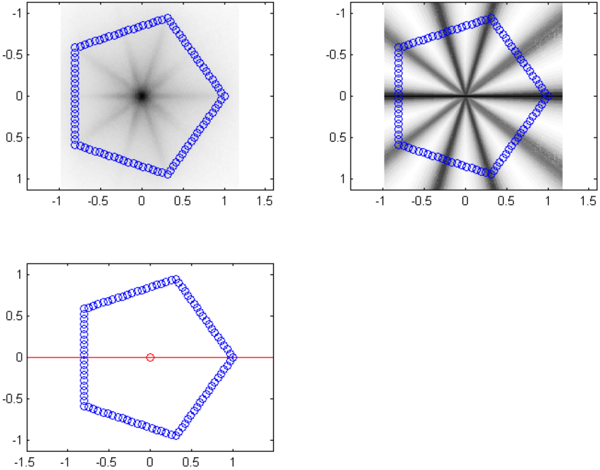

What's In This Unit?

Symmetries

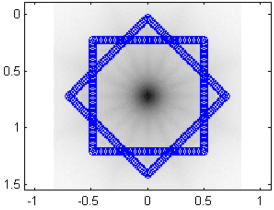

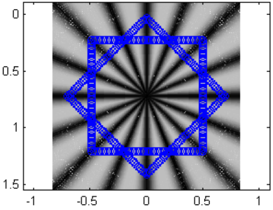

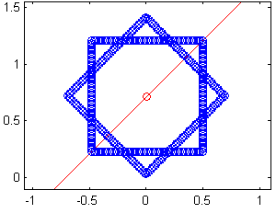

What's In This Unit?

Symmetries

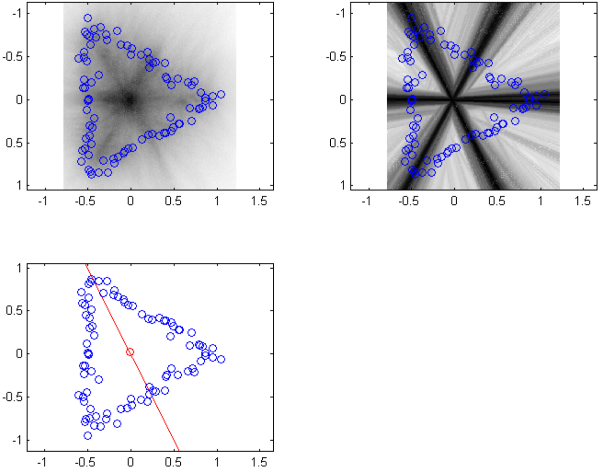

What's In This Unit?

Surface Reconstruction

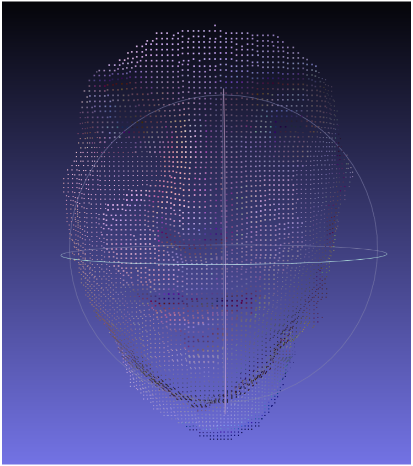

What's In This Unit?

Surface Reconstruction

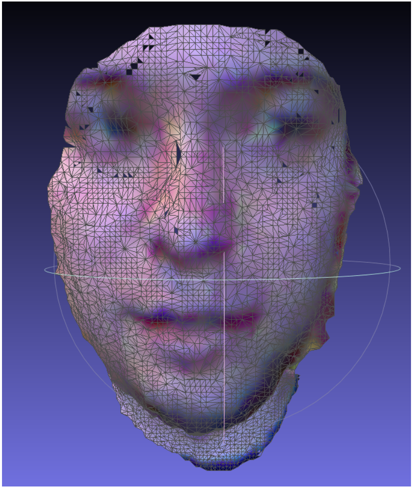

Table of Contents

Eigenvalues/Eigenvectors

$$Ax = \lambda x$$

A is $N \times N$ matrix
x is $N \times 1$ column matrix (vector)
λ is a scalar called the *eigenvalue* of x

Eigenvalues/Eigenvectors

$$Ax = \lambda x$$

A is $N \times N$ matrix
x is $N \times 1$ column matrix (vector)
λ is a scalar called the *eigenvalue* of x

- A does not change the direction of x

Eigenvalues/Eigenvectors Example: X Scale

$$\begin{bmatrix} 2 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 0 \\ 1 \end{bmatrix} =$$

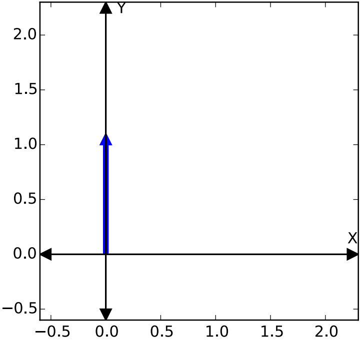

Before

Eigenvalues/Eigenvectors Example: X Scale

$$\begin{bmatrix} 2 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 0 \\ 1 \end{bmatrix} = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$$

Yes! $\lambda = 1$

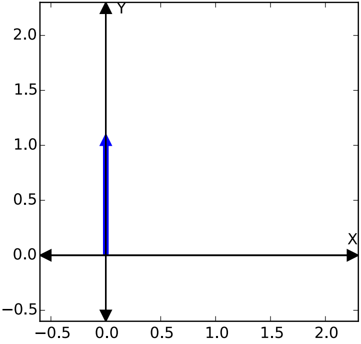

Before

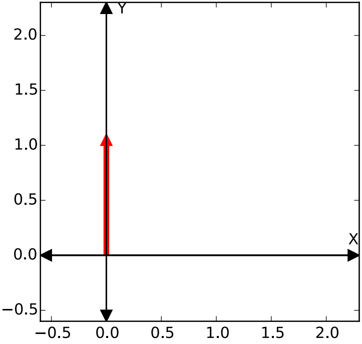

After

Eigenvalues/Eigenvectors Example: X Scale

$$\begin{bmatrix} 2 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 \\ 0 \end{bmatrix} =$$

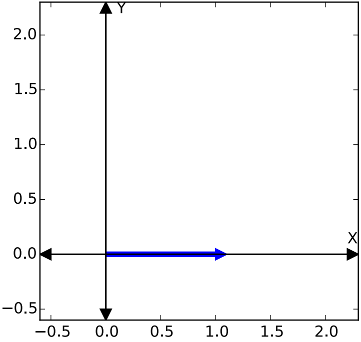

Before

Eigenvalues/Eigenvectors Example: X Scale

$$\begin{bmatrix} 2 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 \\ 0 \end{bmatrix} = \begin{bmatrix} 2 \\ 0 \end{bmatrix}$$

Yes! $\lambda = 2$

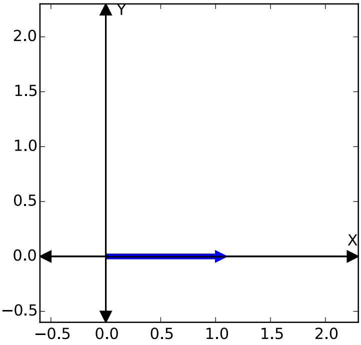

Before

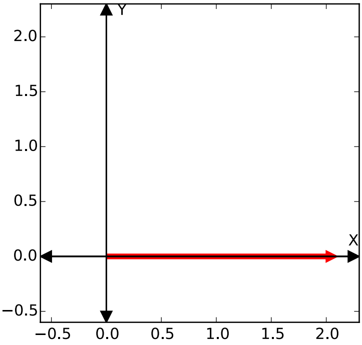

After

Eigenvalues/Eigenvectors Example: X Scale

$$\begin{bmatrix} 2 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 \\ 1 \end{bmatrix} =$$

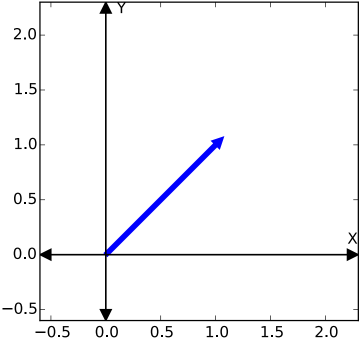

Before

Eigenvalues/Eigenvectors Example: X Scale

$$\begin{bmatrix} 2 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 \\ 1 \end{bmatrix} = \begin{bmatrix} 2 \\ 1 \end{bmatrix}$$

No, changes direction

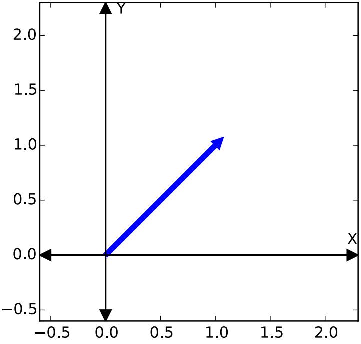

Before

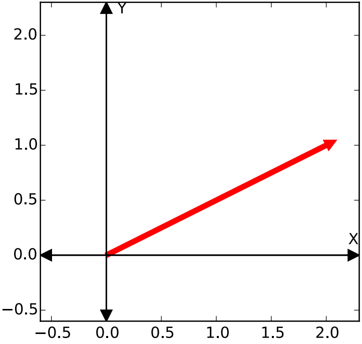

After

Eigenvalues/Eigenvectors Example: Shear

$$\begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 \\ 0 \end{bmatrix} =$$

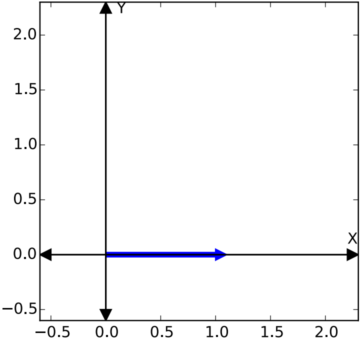

Before

Eigenvalues/Eigenvectors Example: Shear

$$\begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 \\ 0 \end{bmatrix} = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$$

Yes! $\lambda = 1$

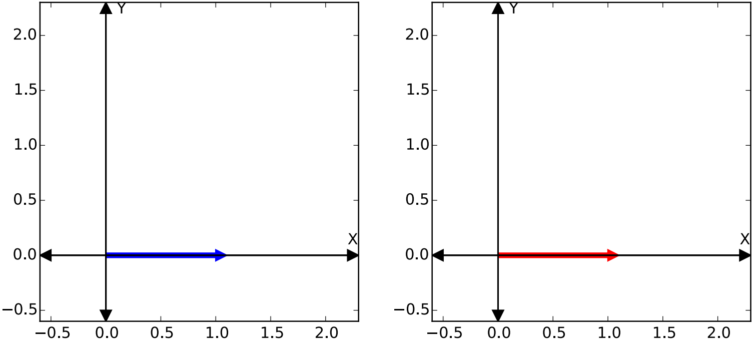

Before

After

Eigenvalues/Eigenvectors Example: Shear

$$\begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 0 \\ 1 \end{bmatrix} =$$

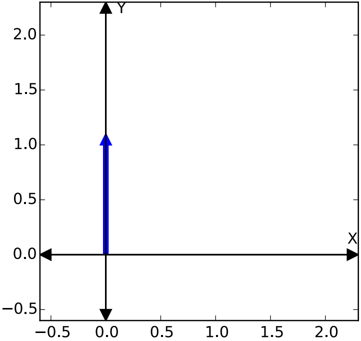

Before

Eigenvalues/Eigenvectors Example: Shear

$$\begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 0 \\ 1 \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \end{bmatrix}$$

No, changes direction

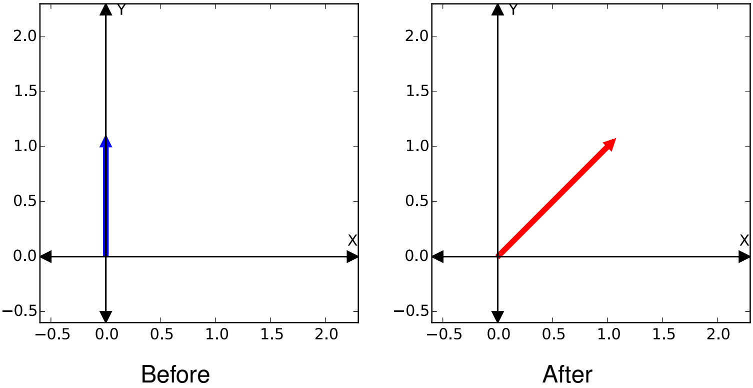

Eigenvalues/Eigenvectors Example: Flip X, Scale Y

$$\begin{bmatrix} -1 & 0 \\ 0 & 2 \end{bmatrix} \begin{bmatrix} 1 \\ 0 \end{bmatrix} =$$

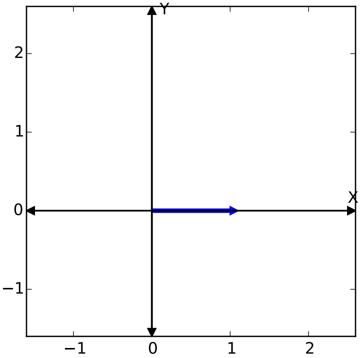

Before

Eigenvalues/Eigenvectors Example: Flip X, Scale Y

$$\begin{bmatrix} -1 & 0 \\ 0 & 2 \end{bmatrix} \begin{bmatrix} 1 \\ 0 \end{bmatrix} = \begin{bmatrix} -1 \\ 0 \end{bmatrix}$$

Yes! $\lambda = -1$ (tricky)

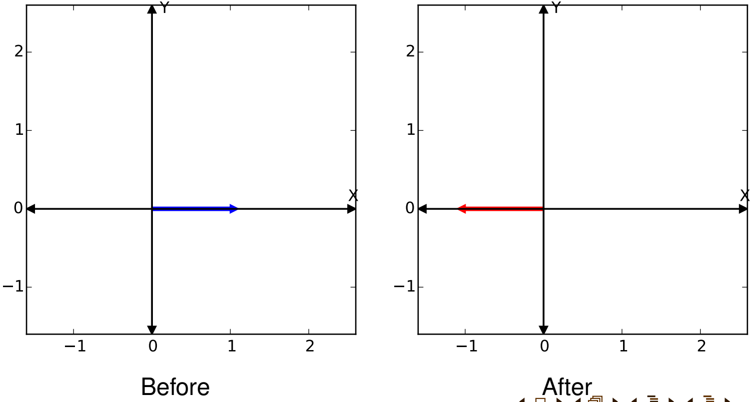

Eigenvalues/Eigenvectors Example: Flip X, Scale Y

$$\begin{bmatrix} -1 & 0 \\ 0 & 2 \end{bmatrix} \begin{bmatrix} 0 \\ 1 \end{bmatrix} =$$

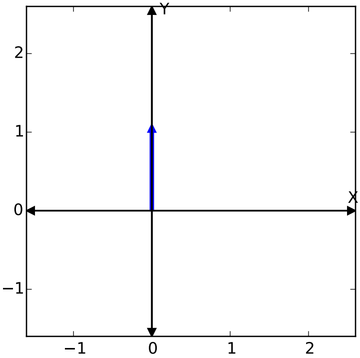

Before

Eigenvalues/Eigenvectors Example: Flip X, Scale Y

$$\begin{bmatrix} -1 & 0 \\ 0 & 2 \end{bmatrix} \begin{bmatrix} 0 \\ 1 \end{bmatrix} = \begin{bmatrix} 0 \\ 2 \end{bmatrix}$$

Yes! $\lambda = 2$

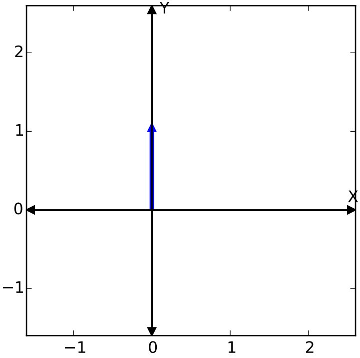

Before

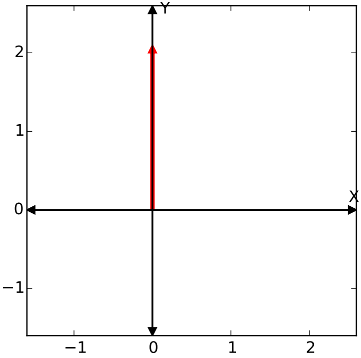

After

Eigenvalues/Eigenvectors Example: Flip X, Scale Y

$$\begin{bmatrix} -1 & 0 \\ 0 & 2 \end{bmatrix} \begin{bmatrix} 1 \\ 1 \end{bmatrix} =$$

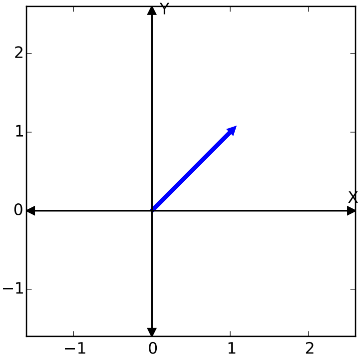

Before

Eigenvalues/Eigenvectors Example: Flip X, Scale Y

$$\begin{bmatrix} -1 & 0 \\ 0 & 2 \end{bmatrix} \begin{bmatrix} 1 \\ 1 \end{bmatrix} = \begin{bmatrix} -1 \\ 2 \end{bmatrix}$$

No, changes direction

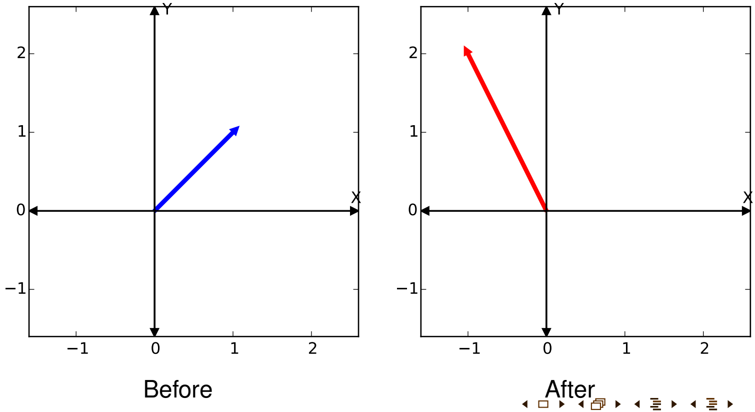

Rotation Matrix

What are the eigenvectors/eigenvalues of

$$A = \left[\begin{array}{cc} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{array} \right]$$

?

Linear Eigen-Decomposition

Eigenvectors/eigenvalues (λ_1, v_1), (λ_2, v_2)
Let

$$v = av_1 + bv_2$$

Then

$$Av = A(av_1 + bv_2) = aAv_1 + bAv_2$$

$$Av = \lambda_1 av_1 + \lambda_2 bv_2$$

Linear Eigen-Decomposition: Example

$$\begin{bmatrix} 4 & -1 \\ -1 & 4 \end{bmatrix}$$

$$v_1 = \begin{bmatrix} -1 \\ -1 \end{bmatrix}, \lambda_1 = 3$$

$$v_2 = \begin{bmatrix} -1 \\ 1 \end{bmatrix}, \lambda_2 = 5$$

Let $v = \begin{bmatrix} 0 \\ 2 \end{bmatrix}$

Spectral Theorem

If A is *symmetric*

that is $A = A^T$

then all of its eigenvectors are *orthogonal*

Spectral Theorem: Example

$$\begin{bmatrix} 1 & 2 \\ 2 & 1 \end{bmatrix}$$

$$v_1 = \begin{bmatrix} -1 \\ 1 \end{bmatrix}, \lambda_1 = -1$$

$$v_2 = \begin{bmatrix} 1 \\ 1 \end{bmatrix}, \lambda_2 = 3$$

Solving for Eigenvalues/Eigenvectors

Table of Contents

Point Cloud Centroid

Given a collection of points $X = \{\vec{x}_i\}|_{i=1}^{N}$
Centroid $\vec{c}$ is the "mean point"

That is

$$\vec{c}[k] = \frac{1}{N}\sum_{i=1}^{N} x_i[k]$$

Average each coordinate independently

Point Cloud Centroid

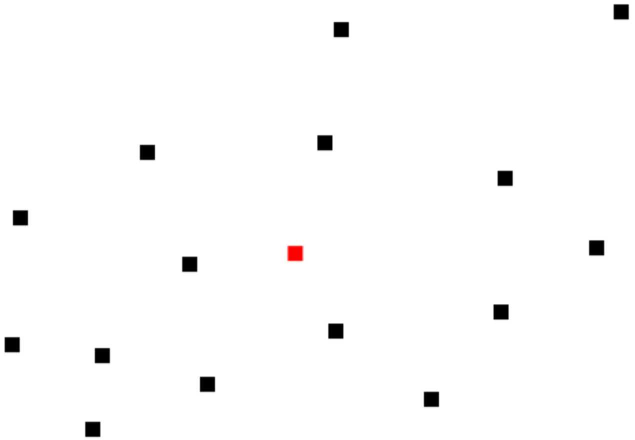

Point Cloud Centroid

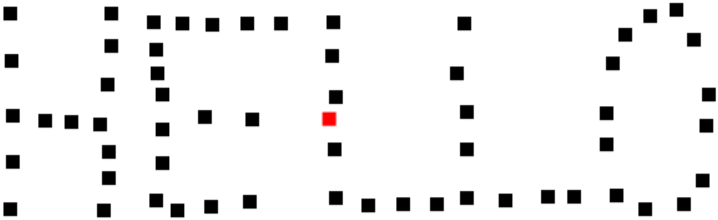

Directions of Variance

Organize point cloud into $N \times d$ matrix, each point along a row

$$X = \begin{bmatrix} - & \vec{x_1} & - \\ - & \vec{x_2} & - \\ - & \vec{x_3} & - \\ \ldots & \vdots & \ldots \\ - & \vec{x_N} & - \end{bmatrix}$$

Choose a unit column vector direction $u \in \mathbb{R}^{d \times 1}$
Then

$$d = Xu$$

gives projections onto u

Directions of Variance

$$d = Xu$$

gives projections onto u

$$d^T d = (Xu)^T (Xu) = u^T (X^T X) u$$

Gives the sum of squared projections onto u

Directions of Variance: Example

1000 point example in 2D, centroid is origin

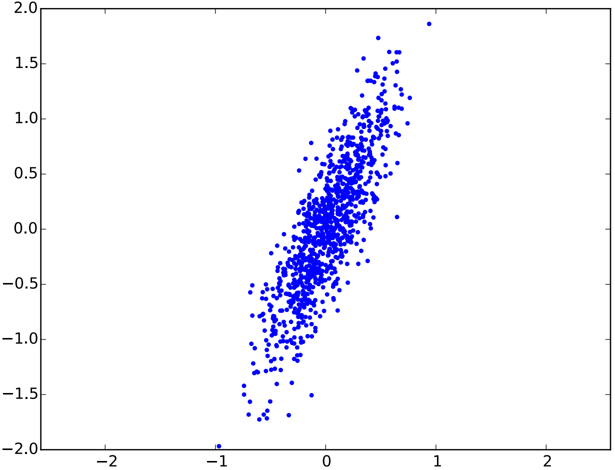

Directions of Variance: Example

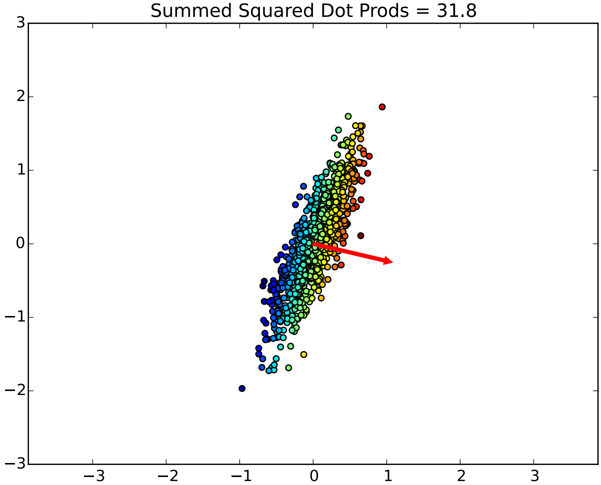

Directions of Variance: Example

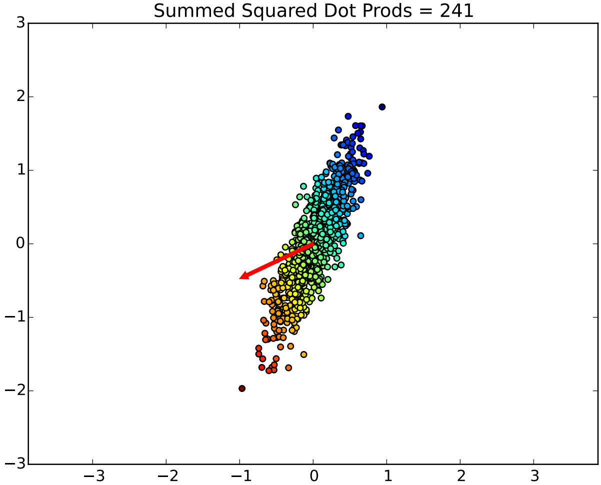

Directions of Variance: Example

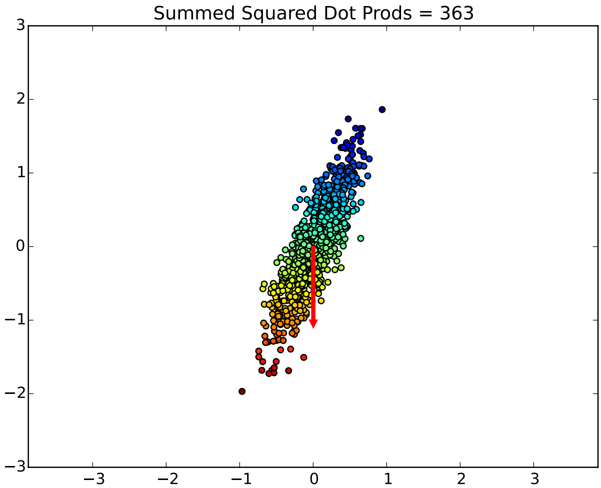

Directions of Variance: Example

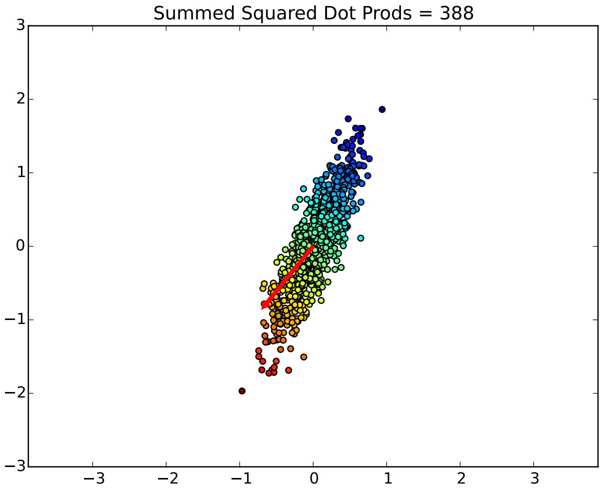

Directions of Variance: Example

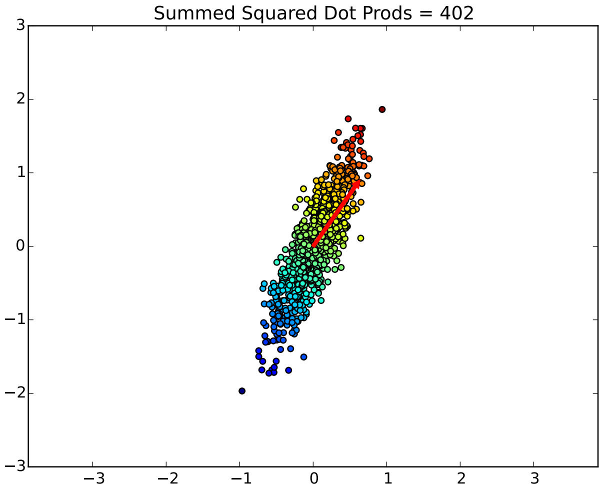

Directions of Variance: Some Math

$$d^T d = (Xu)^T (Xu) = u^T (X^T X) u = u^T A u$$

Directions of Variance: Some Math

$$d^T d = (Xu)^T (Xu) = u^T (X^T X) u = u^T A u$$

- Note that $A = X^T X$ is symmetric.

Directions of Variance: Some Math

$$d^T d = (Xu)^T (Xu) = u^T (X^T X) u = u^T A u$$

- Note that $A = X^T X$ is symmetric.
- It is also *positive semi-definite*. All eigenvalues are real and positive

Directions of Variance: Some Math

$$d^T d = (Xu)^T (Xu) = u^T (X^T X) u = u^T A u$$

- Note that $A = X^T X$ is symmetric.
- It is also *positive semi-definite*. All eigenvalues are real and positive
- Let (v_1, λ_1), (v_2, λ_2), ..., (v_d, λ_d) be the d eigenvalue/eigenvector pairs. Then a direction u can be written as

$$u = a_1 v_1 + a_2 v_2 + \ldots + a_d v_d$$

Directions of Variance: Some Math

$$u^T(X^TX)u = u^TAu$$

Let $u = a_1v_1 + a_2v_2 + \ldots + a_dv_d$
and all vs are unit-norm eigenvectors

Directions of Variance: Some Math

$$u^T(X^TX)u = u^TAu$$

Let $u = a_1v_1 + a_2v_2 + \ldots + a_dv_d$
and all vs are unit-norm eigenvectors
Then

$$u^TAu = (a_1v_1 + a_2v_2 + \ldots + a_dv_d)^TA(a_1v_1 + a_2v_2 + \ldots + a_dv_d)$$

Directions of Variance: Some Math

$$u^T(X^TX)u = u^TAu$$

Let $u = a_1v_1 + a_2v_2 + \ldots + a_dv_d$
and all vs are unit-norm eigenvectors
Then

$$u^TAu = (a_1v_1 + a_2v_2 + \ldots + a_dv_d)^TA(a_1v_1 + a_2v_2 + \ldots + a_dv_d)$$

$$u^TAu = (a_1v_1 + a_2v_2 + \ldots + a_dv_d)^T(a_1\lambda_1v_1 + a_2\lambda_2v_2 + \ldots + a_d\lambda_dv_d)$$

Directions of Variance: Some Math

$$u^T(X^TX)u = u^TAu$$

Let $u = a_1v_1 + a_2v_2 + \ldots + a_dv_d$
and all vs are unit-norm eigenvectors
Then

$$u^TAu = (a_1v_1 + a_2v_2 + \ldots + a_dv_d)^TA(a_1v_1 + a_2v_2 + \ldots + a_dv_d)$$

$$u^TAu = (a_1v_1 + a_2v_2 + \ldots + a_dv_d)^T(a_1\lambda_1v_1 + a_2\lambda_2v_2 + \ldots + a_d\lambda_dv_d)$$

Remembering that all vs are orthogonal

Directions of Variance: Some Math

$$u^T(X^TX)u = u^TAu$$

Let $u = a_1v_1 + a_2v_2 + \ldots + a_dv_d$
and all vs are unit-norm eigenvectors
Then

$$u^TAu = (a_1v_1 + a_2v_2 + \ldots + a_dv_d)^T A(a_1v_1 + a_2v_2 + \ldots + a_dv_d)$$

$$u^TAu = (a_1v_1 + a_2v_2 + \ldots + a_dv_d)^T(a_1\lambda_1v_1 + a_2\lambda_2v_2 + \ldots + a_d\lambda_dv_d)$$

Remembering that all vs are orthogonal

$$u^TAu = a_1^2\lambda_1 + a_2^2\lambda_2 + \ldots + a_d^2\lambda_d$$

Directions of Variance: Some Math

$$u^T A u = a_1^2 \lambda_1 + a_2^2 \lambda_2 + \ldots + a_d^2 \lambda_d$$

Assuming that $||u|| = 1 \implies (a_1^2 + a_2^2 + \ldots a_d^2) = 1$

Also assume that $\lambda_1 > \lambda_2 > \ldots > \lambda_d$

Raffle Point Question: What should the *a*s be to maximize the sum of squared dot products: $u^T A u$?

Principal Component Analysis

1. Stack points in rows.

2. Compute the “covariance matrix” $A = X^T X$

3. Compute eigenvalues/eigenvectors of A, sorted in decreasing order

Orthogonal directions of variance are in the eigenvectors

Sum of squared dot products with directions are associated eigenvalues

PCA Example

Largest direction of variance, $\lambda_1 = 422$

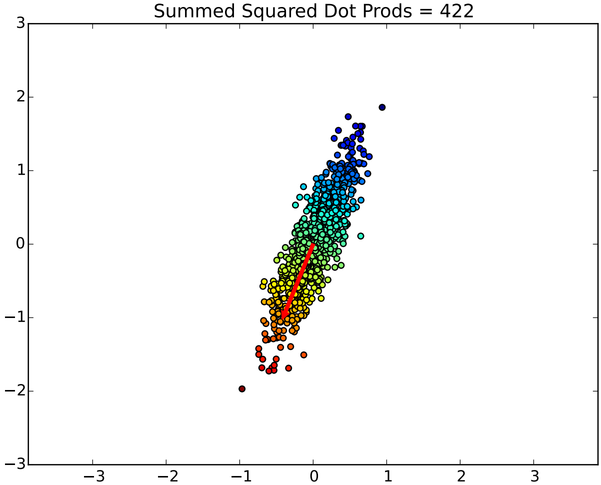

Smallest direction of variance, $\lambda_2 = 21.6$

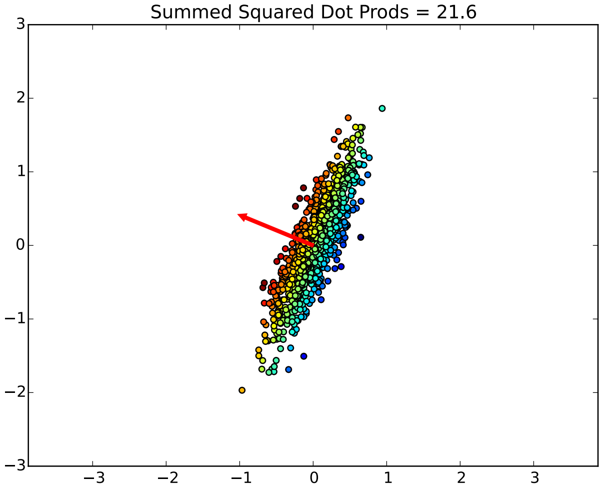

PCA: Interactive Examples 2D/3D

PCA Pitfalls

Outliers!
(Show Demo Again)

PCA Pitfalls

Close Eigenvectors
"Multiplicity issues" for "near-isotropic" point clouds