EE520 Data Driven Analysis of Complex Systems

Erik Bollt

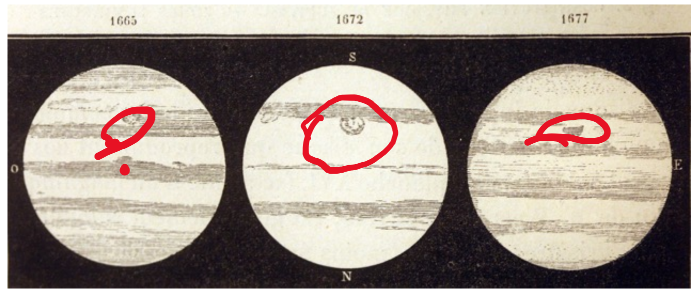

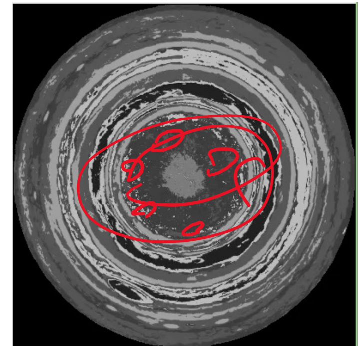

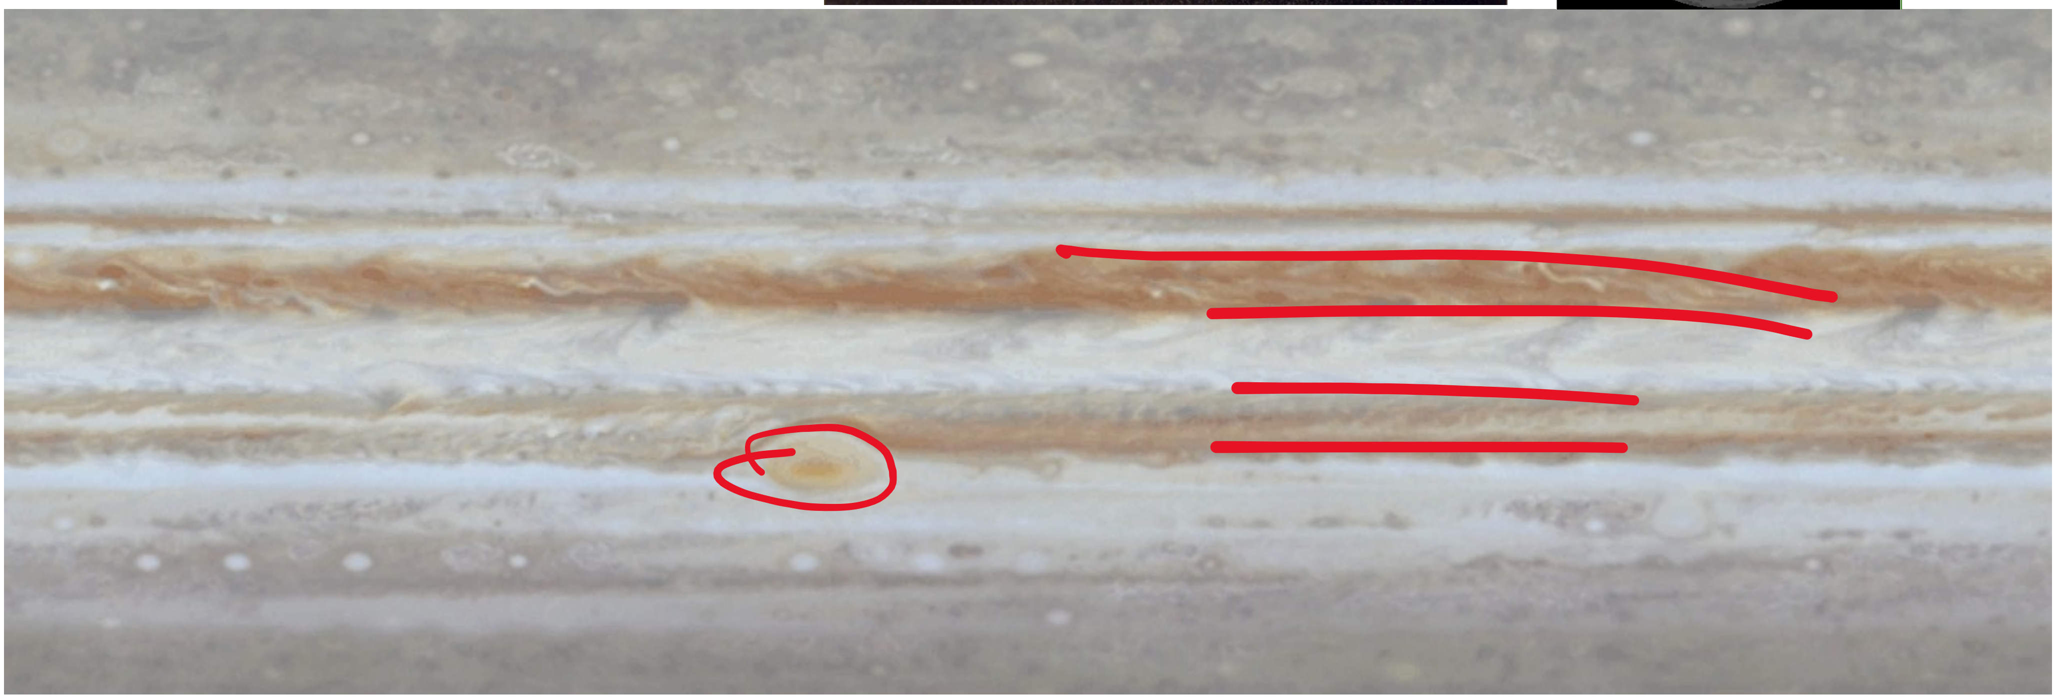

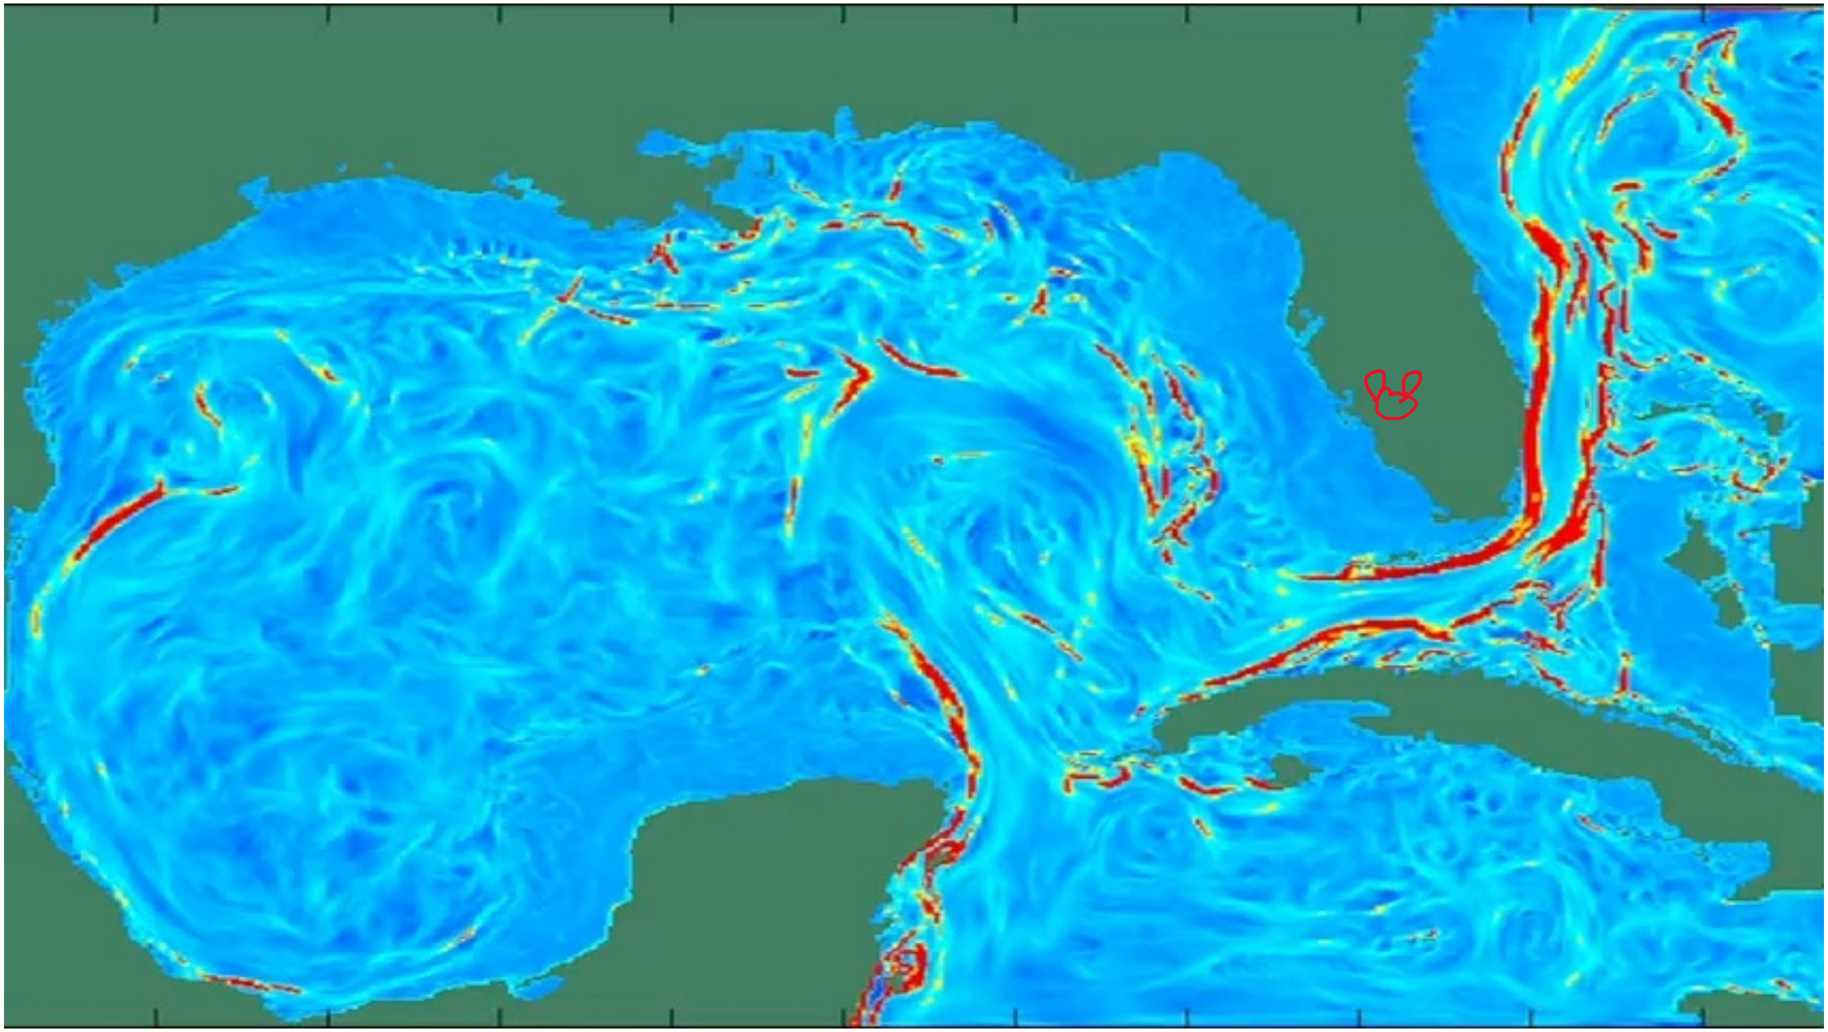

Data as an array

$\underline{X} =$

$\underline{X}_{p \times q} \in \mathbb{R}^{p \times q}$ matrix

$[\underline{X}]_{l,m}$ = one pixel

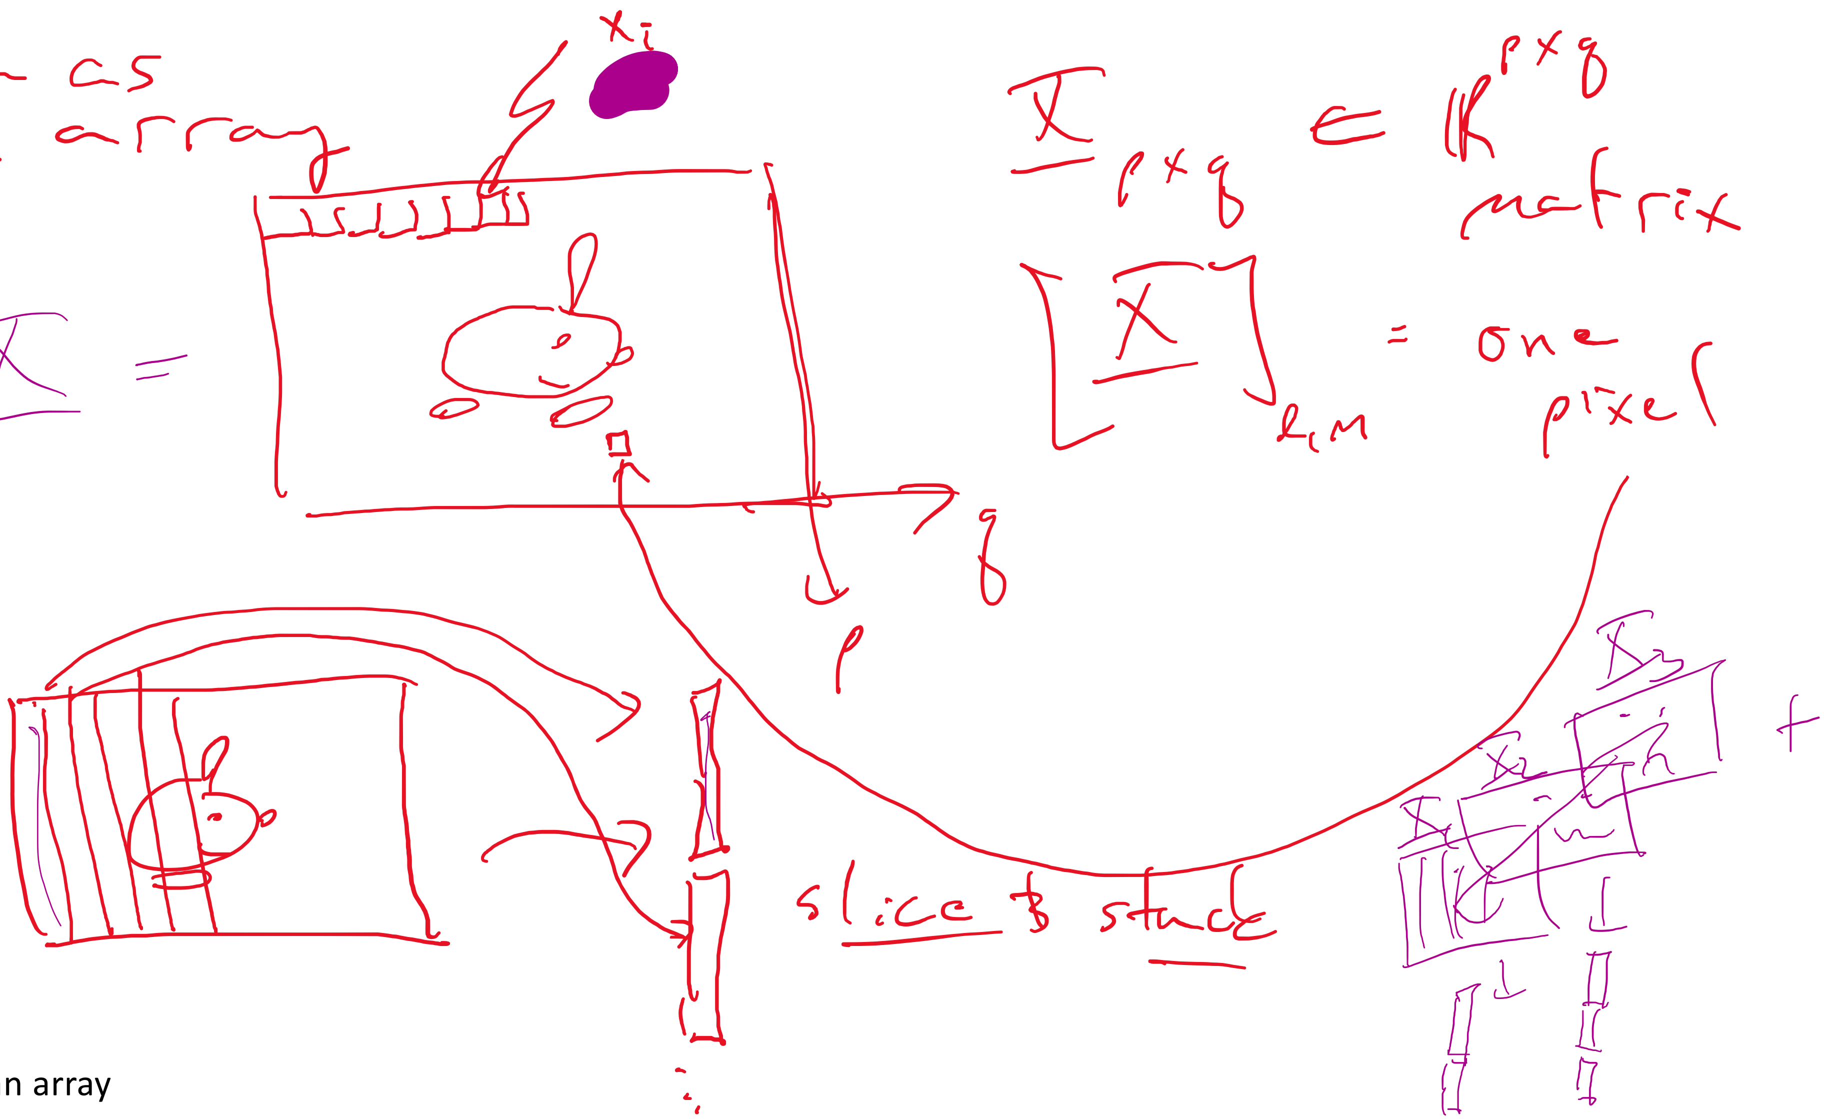

Data as an array

On Matrix Multiplication

$$A = \begin{pmatrix} a_{1,1} & a_{1,2} & \cdots & a_{1,n} \\ a_{2,1} & a_{2,2} & \cdots & a_{2,n} \\ \vdots & & \ddots & \vdots \\ a_{m,1} & a_{m,2} & \cdots & a_{m,n} \end{pmatrix}, \text{ and each } a_{i,j} \in \mathbb{C}$$

$$\begin{aligned} L(\mathbf{z}) : \mathbb{R}^n &\to \mathbb{R}^m \\ \mathbf{z} &\mapsto \mathbf{z}' = A\mathbf{z}, \end{aligned}$$

in terms of the usual matrix times vector multiplication,

$$[\mathbf{z}]'_i = \sum_{j=1}^{n} A_{i,j}[\mathbf{z}]_j, \text{ for each } i = 1, \cdots, m,$$

Geometric

$Az = z'$

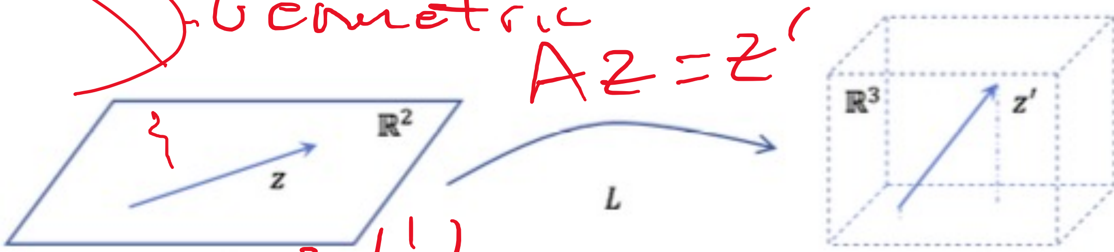

$n = 2$ $\quad z = \begin{pmatrix}1 \\ 2\end{pmatrix}$ $\quad m = 3$

- a vector has length & direction.

But as linear algebra

↳ i.e. Matrices × vectors

$$A_{2\times 2} = \begin{pmatrix} 1 & 2 \\ 3 & 4 \end{pmatrix}$$

$$A\begin{pmatrix}3 \\ 4\end{pmatrix} = \begin{pmatrix} 1 & 2 \\ 3 & 4 \end{pmatrix}\begin{pmatrix}3 \\ 4\end{pmatrix} = \begin{pmatrix} 1\cdot 3 + 2\cdot 4 \\ 3\cdot 3 + 4\cdot 4 \end{pmatrix} = \begin{pmatrix} 11 \\ 25 \end{pmatrix}$$

$z' = Az$

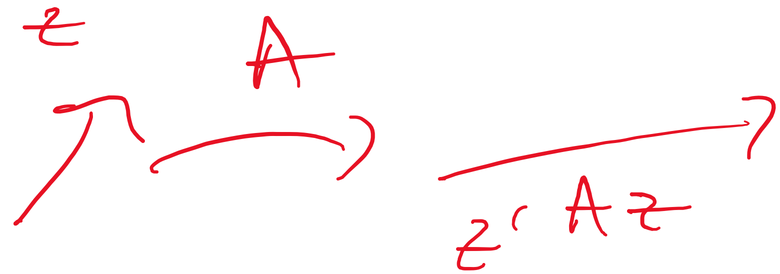

new direction, new length.

Eig for square.

- $\underset{2\times 2}{A} v = \lambda v$

Characterize matrices by knowing just these

- eig. special directions

Linear

- $Au = A(a_1 v_1 + a_2 v_2) = a_1 A v_1 + a_2 A v_2$

  $= a_1 \lambda_1 v_1 + a_2 \lambda_2 v_2$

$\det(A - \lambda I) = 0$

$(A - \lambda I)v = 0$

Matrix time circle =

all vectors of length 1.

? Matrix × circle ?! - But matrix times vector.

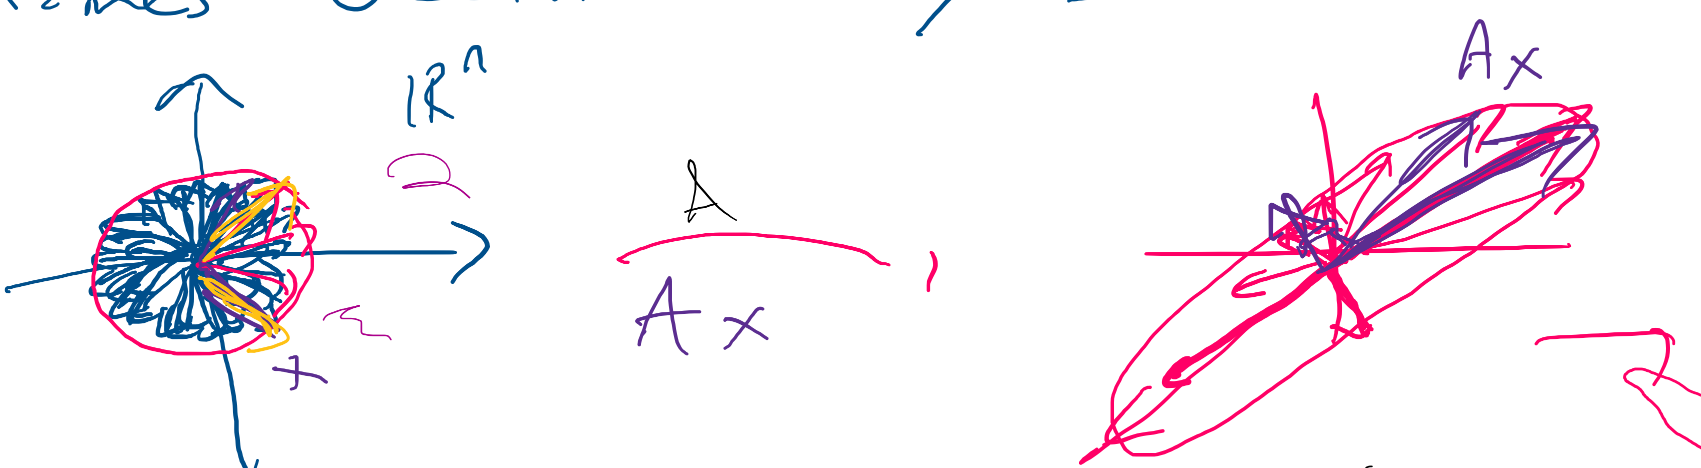

$S = \{ x \mid \|x\|_2 = 1,\ x \in E \cong \mathbb{R}^2 \}$; $A \cdot S = \{ y : y = Ax,\ x \in S \}$

Theorem 2.1.1 — Singular Value Decomposition. Let A be an $m \times n$ matrix whose entries come from the field $\mathcal{K}$, which is either the field of real numbers or the field of complex numbers. Then the singular value decomposition of A exists, and it takes the form of a product of matrices:

$$A_{m\times n} = U_{m\times m}\Sigma_{m\times n}V^*_{n\times n}, \tag{2.5}$$

where

- U is an $m \times m$ unitary matrix.
- Σ is a diagonal $m \times n$ matrix with non-negative real numbers on the diagonal.
- V is an $n \times n$ unitary matrix, and V^* is the conjugate transpose of V.

$$\Sigma := diag(\sigma_1, \sigma_2, \cdots, \sigma_p), p = min(m, n),$$
$$\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_p \geq 0.$$

$\Sigma = \begin{pmatrix} \sigma_1 & & 0 \\ & \sigma_2 & \\ 0 & & \ddots \end{pmatrix}$

$A = U \Sigma V^T$

$u_1^T u_1 = u_1 \cdot u_1 = 1$

$\begin{pmatrix} u_1^T u_1 = 1 & u_1^T u_2 = 0 & u_1^T u_3 = 0 & \cdots \\ u_2^T u_1 & u_2^T u_2 & \cdots & \end{pmatrix}$

$u_i^T u_j = \delta_{ij} = 1$ or 0

The singular values are the nonegative values: $\sigma_i \geq 0, i = 1, \cdots, n,$

The left singular vectors: u_i are the columns of $U = [u_1, u_2, ..., u_m]$.

The right singular vectors: v_i are the columns of $V = [v_1, v_2, ..., v_n]$.

Definition 2.1.1 — Singular values and singular vectors. The singular values of A are the scalar values, σ_i, and the columns of U and V have columns that are the corresponding i^{th} left and right singular vectors, u_i and v_i:

The singular values are the nonegative values: $\sigma_i \geq 0, i = 1, \cdots, n,$

The left singular vectors: u_i are the columns of $U = [u_1, u_2, ..., u_m]$.

The right singular vectors: v_i are the columns of $V = [v_1, v_2, ..., v_n]$.

$\|u\|^2 \cos\theta = \vec{u_1} \cdot \vec{u_1} = 1 \qquad u_2^T u_1 = 0$

- U unitary $\Leftrightarrow U^* U = U U^* = I$

$U^T U = \begin{pmatrix} | & | & & | \\ u_1 & u_2 & \cdots & u_m \\ | & | & & | \end{pmatrix}^T \begin{pmatrix} | & | & & | \\ u_1 & u_2 & \cdots & u_m \\ | & | & & | \end{pmatrix} =$

$= I = \begin{pmatrix} 1 & \\ & \ddots \end{pmatrix}$

Since V is orthogonal, then right multiplying Eq. (2.5) by V,

$$AV = U\Sigma V^* V = U\Sigma, \tag{2.8}$$

$$Ax = b$$

$$\begin{pmatrix} 1 & 2 & 3 \\ 2 & 4 & 6 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} = \begin{pmatrix} b_1 \\ b_2 \end{pmatrix}$$

if

A is square

$$AV = U\Sigma$$

$$[A]\,[v_1 v_2 \cdots v_n] = [u_1 u_2 \cdots u_n]\, diag(\sigma_1, \sigma_2, \cdots, \sigma_n).$$

■ **Example 2.1** Let $A = \begin{pmatrix} 1 & 2 & 3 \\ 2 & 4 & 6 \end{pmatrix}_{2\times 3}$. By SVD of the matrix A we have:

$$\begin{aligned} A &= U\Sigma V^T \\ &= \begin{pmatrix} \frac{1}{\sqrt{5}} & \frac{2}{\sqrt{5}} \\ \frac{-2}{\sqrt{5}} & \frac{1}{\sqrt{5}} \end{pmatrix} \begin{pmatrix} \sqrt{70} & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} \frac{1}{\sqrt{14}} & \sqrt{\frac{2}{7}} & \frac{3}{\sqrt{14}} \\ \frac{-3}{\sqrt{10}} & 0 & \frac{1}{\sqrt{10}} \\ \frac{-1}{\sqrt{35}} & \sqrt{\frac{5}{7}} & \frac{-3}{\sqrt{35}} \end{pmatrix}. \end{aligned} \tag{2.28}$$

We see that the second singular value, $\sigma_2 = 2$, meaning that number of non-zero singular values $r < \min\{m, n\}$. Such matrix is called rank deficient matrix. If we take the economy version (with $r = 1$) of the SVD we will have:

$$\begin{aligned} u_1\sigma_1 v_1^T &= \begin{pmatrix} \frac{1}{\sqrt{5}} \\ \frac{-2}{\sqrt{5}} \end{pmatrix} \left(\sqrt{70}\right) \begin{pmatrix} \frac{1}{\sqrt{14}} & \sqrt{\frac{2}{7}} & \frac{3}{\sqrt{14}} \end{pmatrix} \\ &\approx \begin{pmatrix} 1 & 2 & 3 \\ 2 & 4 & 6 \end{pmatrix} \end{aligned} \tag{2.29}$$

■

$$A^T A = V \Sigma^T \Sigma V^T \quad \cdot V$$

$$A^T A V = \Sigma^T \Sigma V$$

$$V = (v_1 \; v_2 \; \cdots \; v_n)$$

$$[A]\,[v_1 v_2 \cdots v_n] = [u_1 u_2 \cdots u_n]\, diag(\sigma_1, \sigma_2, \cdots, \sigma_n).$$

$$AA^T = U\Sigma \underbrace{V^T (V} \Sigma^T U^T)$$

$$= U \Sigma \Sigma^T U^T$$

$$(AA^T)\, U = U(\Sigma\Sigma^T) = (\Sigma\Sigma^T)\, U$$

$$U = \begin{pmatrix} | & | & & | \\ u_1 & u_2 & \cdots & u_m \\ | & | & & | \end{pmatrix}$$

Full ☆

The Economy SVD, and Reduced Rank SVD

The general SVD, Eq. (2.5) may be written in terms of submatrices.

Definition 2.1.3 — The Economy SVD. For any matrix $A \in \mathbb{R}^{m\times n}$, the general SVD Eq. (2.5) can be written in terms of smaller matrices,

$$A_{m\times n} = \hat{U}_{m\times n}\hat{\Sigma}_{n\times n}V^*_{n\times n}, \tag{2.21}$$

and $U = [\hat{U}_{m\times n}|\hat{U}_{(n-m)\times n}]$, written in terms of an orthogonal "buffer" matrix

$\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_r > \sigma_{r+1} = 0$

$= 0 \dots$

Definition 2.1.4 — Rank Deficient SVD. For a matrix $A \in \mathbb{R}^{m\times n}$ such that the SVD results in singular values

$$\sigma_r > \sigma_{r+1} = 0, \text{ for some } r < n. \tag{2.22}$$

then the SVD can be written in terms of an economy form as smaller matrices,

$$A_{m\times n} = \hat{U}_{m\times r}\hat{\Sigma}_{n\times n}V^*_{n\times r}, \tag{2.23}$$

and related to the general SVD Eq. (2.5) by $U = [\hat{U}_{m\times r}|\hat{U}_{(n-r)\times n}]$, but $r < n$.

Rank ill conditioned

Full,
Economy,
Truncated
SVD

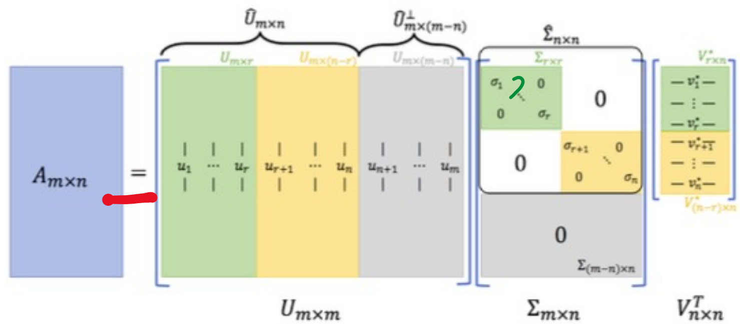

$m>n$

Figure 2.3: $m > n$ tall skinny

$\sigma_{r+1} = 0$

$\sigma_r = 10^{-6}$

$\sigma_{r+1} \sim 10^{-13}$

$< \varepsilon$

$m<n$

$A = [\quad]$

Recall that,

$$
\begin{aligned}
A_{m\times n} &= \hat{U}_{m\times n}\hat{\Sigma}_{n\times n}\hat{V}^T_{n\times n} \\
&= \begin{bmatrix} | & | & & | \\ u_1 & u_2 & \dots & u_n \\ | & | & & | \end{bmatrix} \begin{bmatrix} \sigma_1 & 0 & \dots & 0 \\ 0 & \sigma_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \sigma_n \end{bmatrix} \begin{bmatrix} - & v_1^T & - \\ - & v_2^T & - \\ - & \vdots & - \\ - & v_n^T & - \end{bmatrix}
\end{aligned} \tag{2.24}
$$

but $V^TV = I$, orthogonality allows:

$$A_{m\times n}\hat{V}_{n\times n} = \hat{U}_{m\times n}\hat{\Sigma}_{n\times n} \tag{2.25}$$

so,

$$A_{m\times n}\begin{bmatrix} | & | & & | \\ v_1 & v_2 & \dots & v_n \\ | & | & & | \end{bmatrix} = \begin{bmatrix} | & | & & | \\ u_1 & u_2 & \dots & u_n \\ | & | & & | \end{bmatrix} \begin{bmatrix} \sigma_1 & 0 & \dots & 0 \\ 0 & \sigma_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \sigma_n \end{bmatrix} \tag{2.26}$$

but this just states n-matrix times vector statements:

$$
\begin{aligned}
Av_1 &= \sigma_1 u_1 \\
Av_2 &= \sigma_2 u_2 \\
&\vdots \\
Av_n &= \sigma_n u_n
\end{aligned} \tag{2.27}
$$

Full,
Economy,
Truncated
SVD

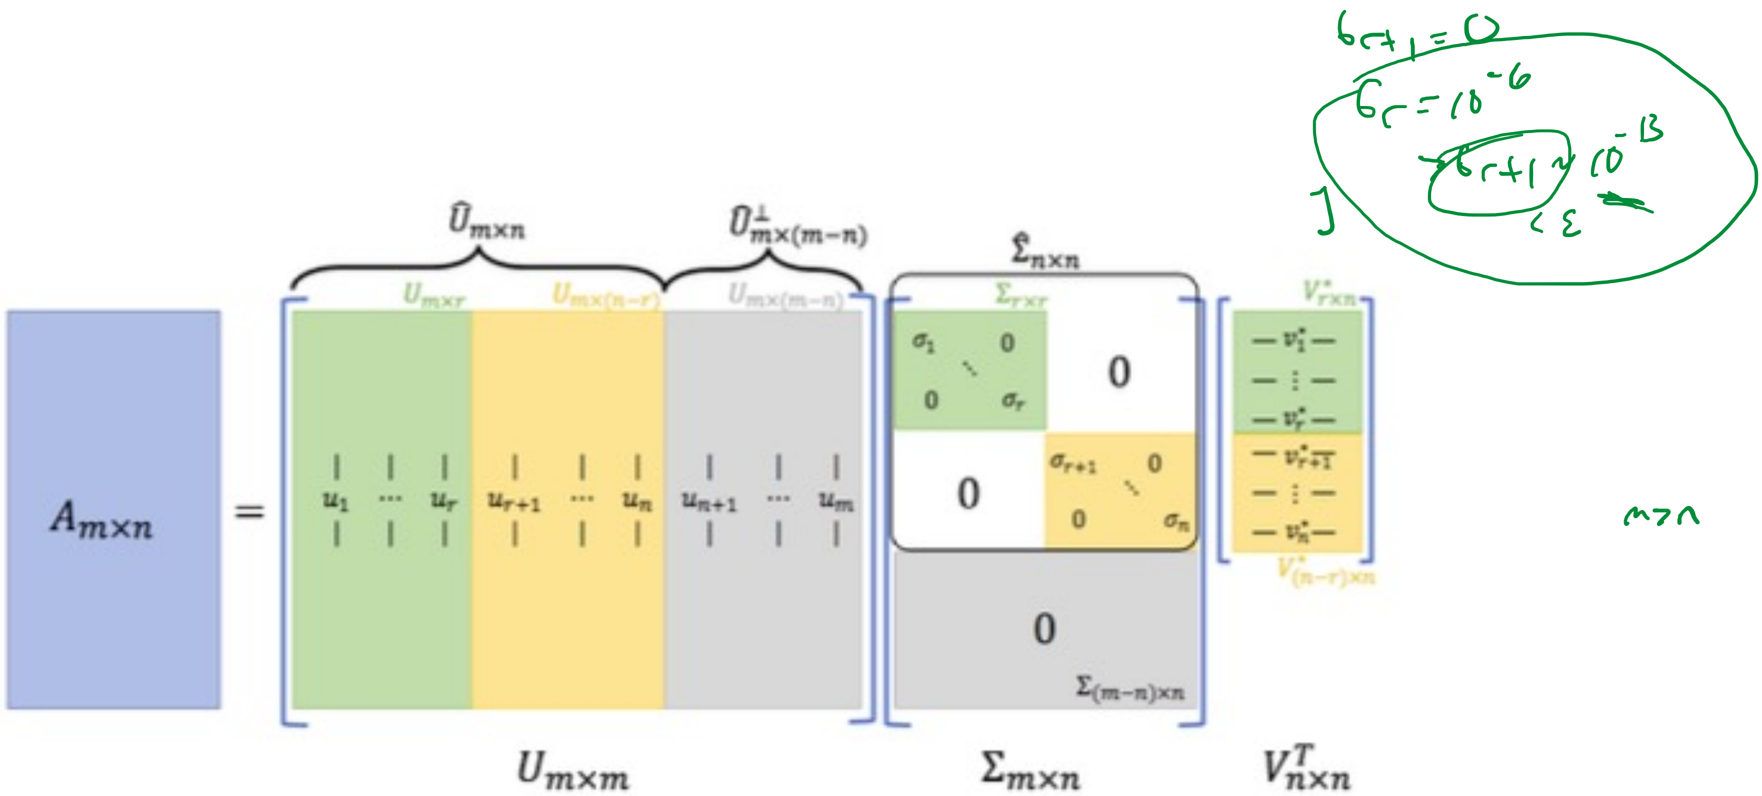

Figure 2.3: $m > n$ tall skinny

$A = U\Sigma V^T$

$AV = U\Sigma$

$Av_i = \sigma_i u_i$

Geometry:.

1. V^* rotates to a standard configuration.
2. Σ stretches each orthogonal axis to the major covariance axis of the corresponding ellipsoid, and
3. U rotates results back to the configuration that associates with A.

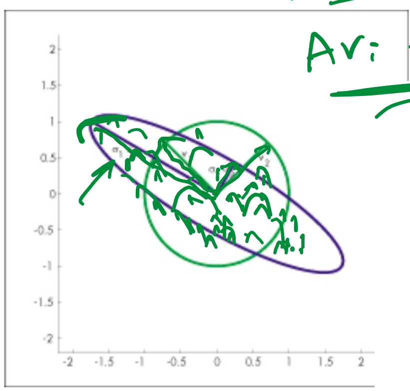

$\sigma_1 \sim 1.5$

$\sigma_2 \sim 0.5$

what if

$\sigma_1 = 2$

$\sigma_2 = 0.0003$

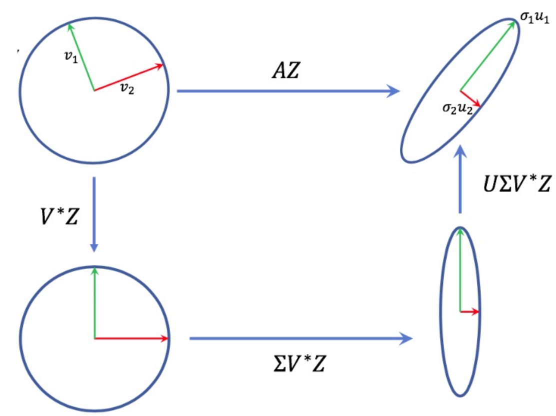

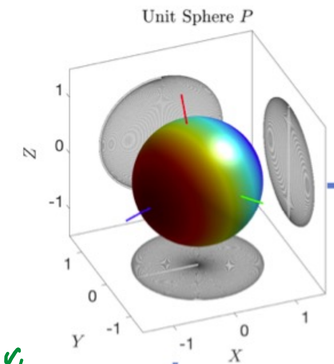

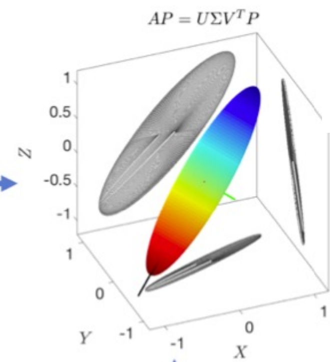


```
disp(A)

    0.8929    0.0417    0.1000
    0.3320    0.1077    0.0000
    0.8212    0.5951    0.0000

[U,S,V] = svd(A)

U = 3x3
   -0.6330    0.7469    0.2035
   -0.2590    0.0435   -0.9649
   -0.7296   -0.6635    0.1659

S = 3x3
    1.3438         0         0
         0    0.3912         0
         0         0    0.0208

V = 3x3
   -0.9304    0.3488   -0.1126
   -0.3635   -0.9176    0.1612
   -0.0471    0.1909    0.9805
```

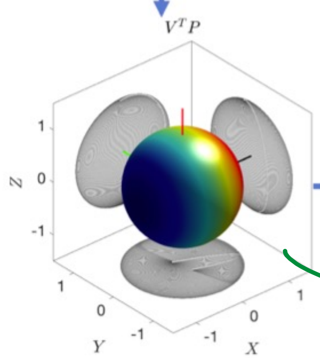

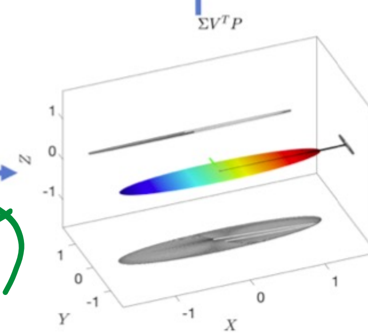

A

$Aw = U\Sigma V^T w$

$v_i^T w$

$= v_i \cdot w$

And ROM

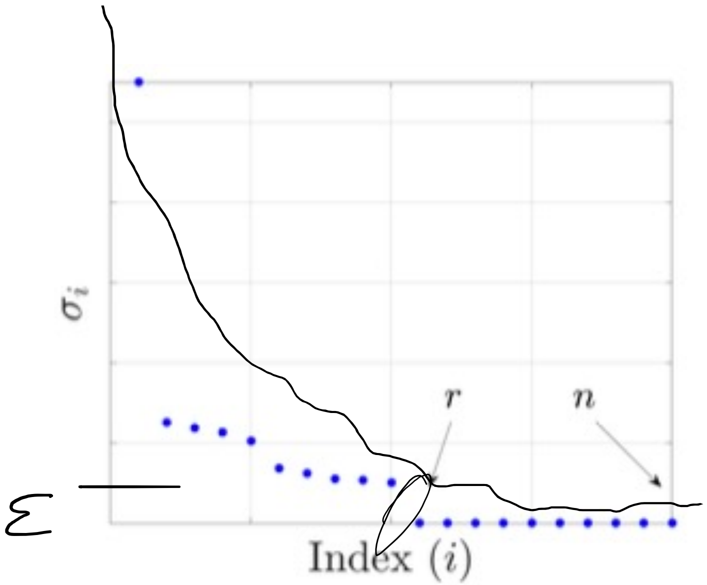

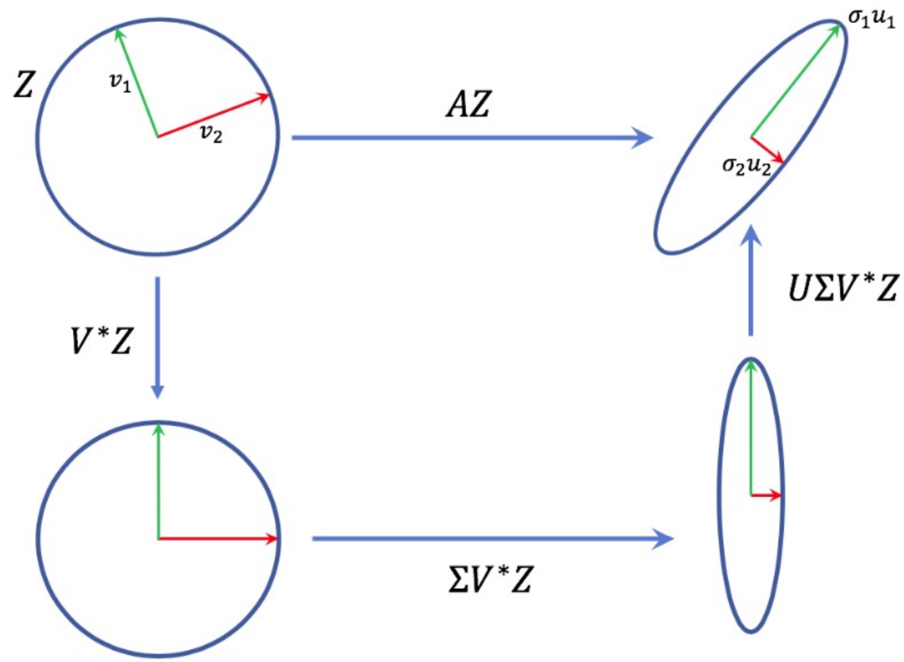

$\sigma_r > \varepsilon > \sigma_{r+1}$

~ 0

Bunny Compression

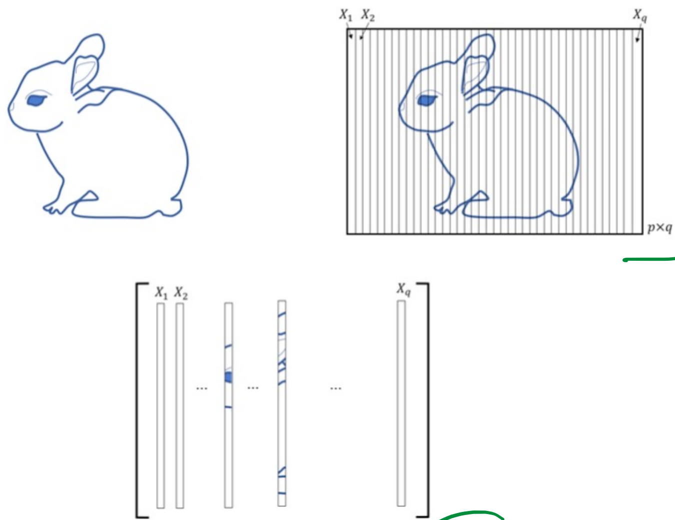

Figure 2.6: Caption

$\underline{X}$

$\underline{X}^T \underline{X}$

Covariance – notice the demean step

$$C_I = \frac{1}{n-1}\left(X - \tilde{X}\right)^T\left(X - \tilde{X}\right)$$

$A^T \quad A$

$\underline{X}$

$\underline{X}$

$\underline{X} - \tilde{\underline{X}}$

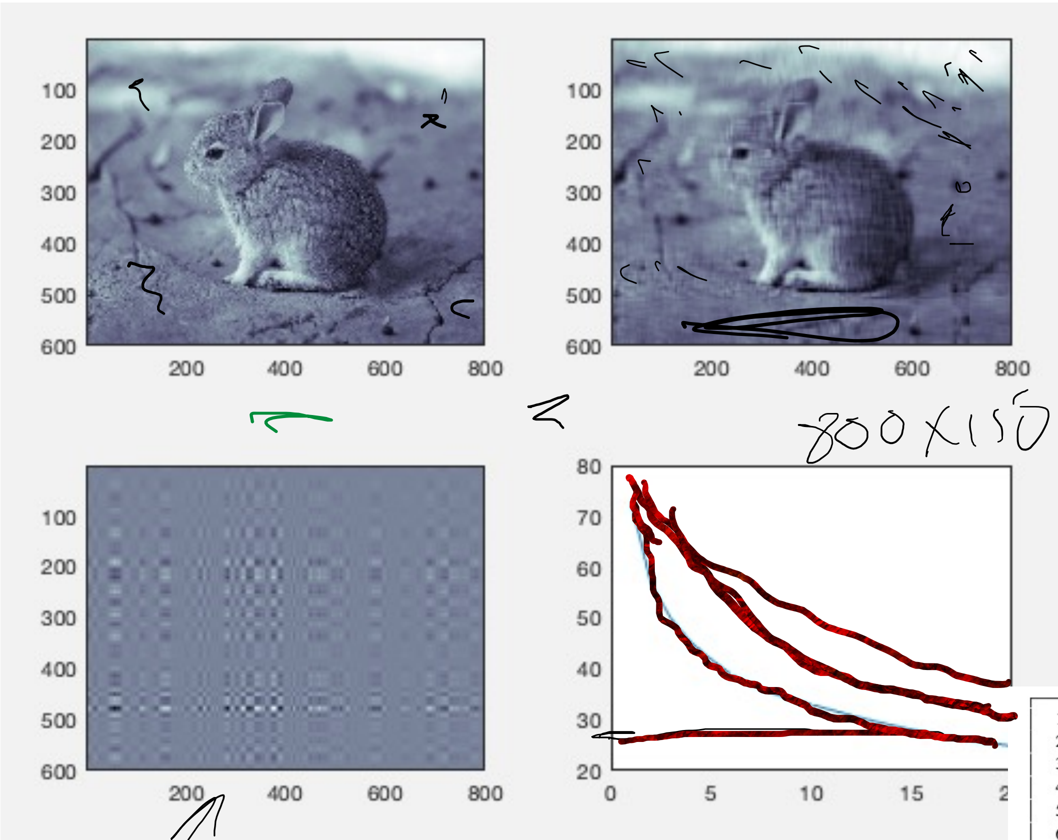

800x150

$$U(x,t) = \sum a_i(t)\phi_i(x)$$

$$|a_i(t)| \xrightarrow{i\to\infty} 0$$

as fast as possible.

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

800x150

$$\approx \hat{U} \hat{\Sigma} \hat{V}^T$$

800x10

```
I = imread('Bunny.jpg');

figure
subplot(1,2,1)
imshow(I)
xticks({}); yticks({});
pbaspect([1 1 1])
title('RGB Image')

I = rgb2gray(I);  %Convert the 3D RGB color to 1D grayscale
I = im2double(I); %Convert integer value to double (scaled ...
    from 0 to 1)

subplot(1,2,2)
imshow(I)
xticks({}); yticks({});
pbaspect([1 1 1])
title('Grayscale Image')
```

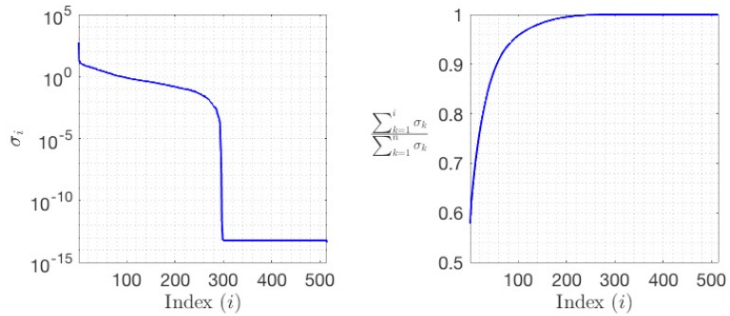

Figure 2.8: (Left) Singular Values. (Right) Energy

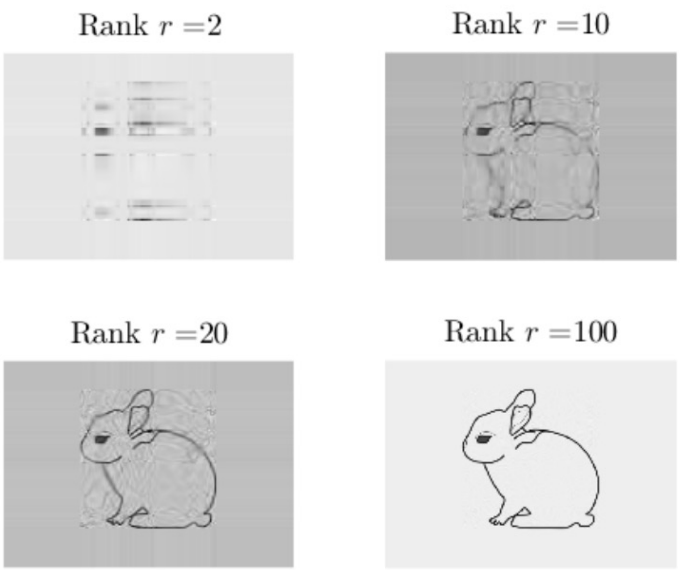

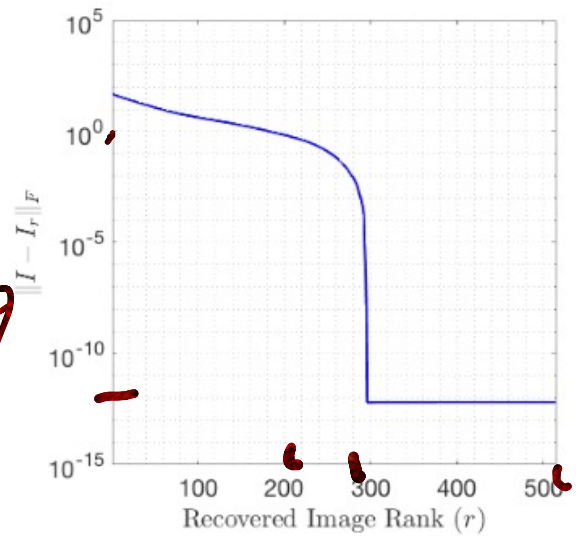

$$e_{rel} = \frac{\|I - I_i\|_F}{\|I\|_F}$$

: Distance $\|I - I_r\|_F$, where I_r is the recovered image using the reduced

Code 2.1: Read, convert, and display images.

```
I = imread('Bunny.jpg');

figure
subplot(1,2,1)
imshow(I)
xticks({}); yticks({});
pbaspect([1 1 1])
title('RGB Image')

I = rgb2gray(I);  %Convert the 3D RGB color to 1D grayscale
I = im2double(I); %Convert integer value to double (scaled ...
    from 0 to 1)

subplot(1,2,2)
imshow(I)
xticks({}); yticks({});
pbaspect([1 1 1])
title('Grayscale Image')
```

History

Gene Golub's license plate, photographed by Professor P. M. Kroonenberg of Leiden University.Gene Howard Golub (February 29, 1932 – November 16, 2007), Fletcher Jones Professor of Computer Science at Stanford University. His work made fundamental contributions that have made the singular value decomposition practical as one of the most powerful and widely used tools in modern matrix computation.

Lot's of Machine ~~lea~~ learning & Data Analysis is solving an ill-posed - optimize a cost function.

Definition 2.1.2 — Induced Norm. Suppose a vector norm $\|\cdot\|$ on $\mathcal{K}^m$ is given. Any matrix $A_{m\times n}$ induces a linear operator from $\mathcal{K}^n$ to $\mathcal{K}^m$ with respect to the standard basis, and one defines the corresponding induced norm or operator norm on the space $\mathcal{K}^{m\times n}$ of all $m \times n$ matrices as follows:

$$\|A\|_p = \sup_{x\neq 0} \frac{\|Ax\|_p}{\|x\|_p} \tag{2.14}$$

or, taking a vector x such that $\|x\|_p = 1$, then we have

$$\|A\|_p = \sup_{\|x\|_p=1} \|Ax\|_p \tag{2.15}$$

Some Special (Simple) Matrix Norms

The first 3 of these are induced norms, but the 4th is not.

- For $p = 1$:

$$\|A\|_1 = \max_{1\leq j\leq n} \sum_{i=1}^{m} |a_{ij}| \tag{2.16}$$

- For $p = \infty$:

$$\|A\|_\infty = \max_{1\leq i\leq m} \sum_{j=1}^{n} |a_{ij}| \tag{2.17}$$

- A special case is the spectral norm when $p = 2$, in which we have:

$$\|A\|_2 = \sqrt{\lambda_{max}(A^T A)} = \sigma_{max} \tag{2.18}$$

where σ_{max} is the maximum singular value of the matrix A.

- The Frobenius norm is given by:

$$\|A\|_F = \sqrt{\sum_{i=1}^{m}\sum_{j=1}^{n} |a_{ij}|^2} = \sqrt{\sum_{i=1}^{\min\{m,n\}} \sigma_i^2} \tag{2.19}$$

Theorem 2.1.2 For a matrix A, the product of the singular values of A, equals the absolute value of its determinant:

$$|det(A)| = \prod_{i=1}^{n} \sigma_i \tag{2.20}$$

$Ax = x'$

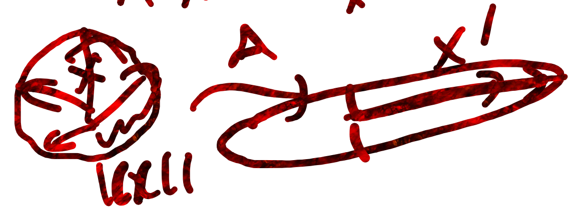

$p=1$: $\|(x_1, x_2)\|_1 = |x_1| + |x_2|$

$p=\infty$: $\|(x_1, x_2)\|_\infty = \max_i |x_i|$

$p=2$: $\|(x_1, x_2)\|_2 = \sqrt{x_1^2 + x_2^2}$

$\vec{x} = \begin{pmatrix} 1 \\ 3 \end{pmatrix}$; $\|x\|_1 = 1 + 3 = 4$

$\|x\|_\infty = 3$

$\|x\|_2 = \sqrt{1^2 + 3^2} = \sqrt{10}$

Definition 2.1.2 — Induced Norm. Suppose a vector norm $\|\cdot\|$ on $\mathcal{K}^m$ is given. Any matrix $A_{m\times n}$ induces a linear operator from $\mathcal{K}^n$ to $\mathcal{K}^m$ with respect to the standard basis, and one defines the corresponding induced norm or operator norm on the space $\mathcal{K}^{m\times n}$ of all $m \times n$ matrices as follows:

$$\|A\|_p = \sup_{x\neq 0} \frac{\|Ax\|_p}{\|x\|_p} \tag{2.14}$$

or, taking a vector x such that $\|x\|_p = 1$, then we have

$$\|A\|_p = \sup_{\|x\|_p=1} \|Ax\|_p \tag{2.15}$$

Some Special (Simple) Matrix Norms

The first 3 of these are induced norms, but the 4th is not.

- For $p = 1$:

$$\|A\|_1 = \max_{1\leq j\leq n} \sum_{i=1}^{m} |a_{ij}| \tag{2.16}$$

- For $p = \infty$:

$$\|A\|_\infty = \max_{1\leq i\leq m} \sum_{j=1}^{n} |a_{ij}| \tag{2.17}$$

- A special case is the spectral norm when $p = 2$, in which we have:

$$\|A\|_2 = \sqrt{\lambda_{max}(A^TA)} = \sigma_{max} \tag{2.18}$$

where σ_{max} is the maximum singular value of the matrix A.

- The Frobenius norm is given by:

$$\|A\|_F = \sqrt{\sum_{i=1}^{m}\sum_{j=1}^{n} |a_{ij}|^2} = \sqrt{\sum_{i=1}^{\min\{m,n\}} \sigma_i^2} \tag{2.19}$$

Theorem 2.1.2 For a matrix A, the product of the singular values of A, equals the absolute value of its determinant:

$$|det(A)| = \prod_{i=1}^{n} \sigma_i \tag{2.20}$$

Fun facts about matrix estimation (data estimation)

If A, $\sigma_1 \geq \dots \geq \sigma_r > \sigma_{r+1} = 0$

- $\mathrm{range}(A) = \mathrm{span}(u_1, u_2, \dots, u_r)$
- $\mathrm{null}(A) = \mathrm{span}(v_{r+1}, v_{r+2}, \dots, v_n)$
- $\|A\|_2 = \sigma_1$; $\|A\|_F = \sqrt{\sigma_1^2 + \sigma_2^2 + \dots + \sigma_r^2}$
- $A = \sum_{i=1}^{r} \sigma_i u_i v_i^* = \sigma_1 u_1 v_1^* + \sigma_2 u_2 v_2^* + \dots + \sigma_r u_r v_r^*$

rank-1 outer products.

$\Big\} \; 0 \xrightarrow{A}$

$\sigma_2 = 0$

$r = 1$

$A = U \Sigma V^T$

$$\begin{pmatrix} | & | & & | \\ u_1 & v_2 & \dots & v_n \\ | & | & & | \end{pmatrix}_{m \times n} \begin{pmatrix} \sigma_1 & \sigma_2 & \dots & \sigma_r \end{pmatrix} \begin{pmatrix} - & v_1^T & - \\ - & v_2^T & - \\ & \vdots & \end{pmatrix} \; 1 \times n$$

$w_1 \nearrow w_2$

$w_1^T w_2 = w_1 \cdot w_2$

$= \|w_1\| \|w_2\| \cos\theta$

$$\begin{pmatrix} 1 & 2 \\ 2 & 4 \end{pmatrix}$$

Matrix Estimation / Data Estimation. $A_{m \times n}$

A

- let $0 \le N \le r$ and $A_N = \sum_{i=1}^{N} \sigma_i u_i v_i^*$

(so we may be skipping some of them ... $\sum_{i=N+1}^{r}$

Then. $\|A - A_N\|_2 = \sigma_{N+1}$ (first one skipped)

(what if it zero?)

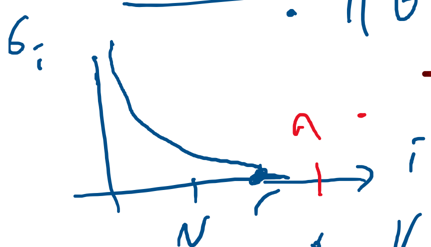

$e_N = \|A - A_N\|_F = \sqrt{\sigma_{N+1}^2 + \sigma_{N+2}^2 + \dots + \sigma_r^2}$

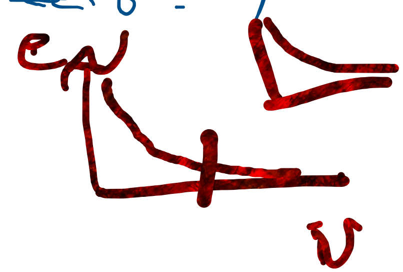

- EigenFace 1st Present

Eigenface

08/31/20

i^{th} picture

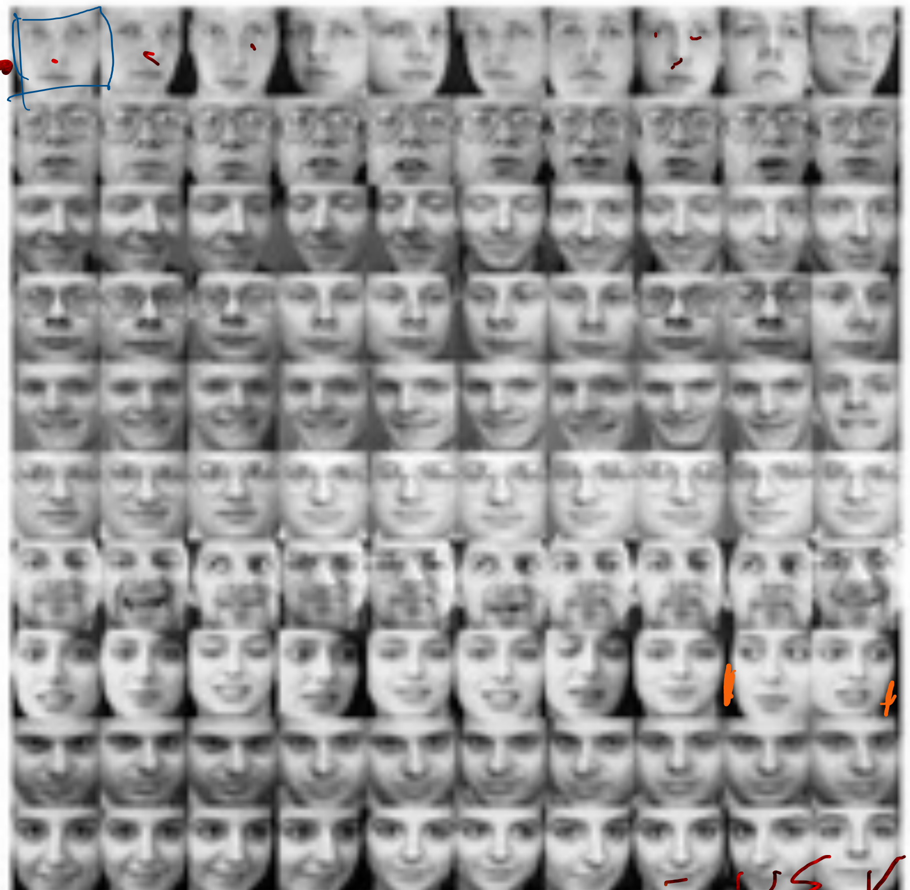

i — $p \times q$ array

reshape as vector

$$\begin{bmatrix} \\ x_i \\ \\ \end{bmatrix}_{M \times 1} \qquad e_1 = \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}$$

$M = pq$

$$\bar{X} = [\, x_1 \,|\, x_2 \,|\, \cdots \,|\, x_n \,]_{M \times n} = U \Sigma V^T$$

$p = 6000,\ q = 200$

- Remove variance

Register

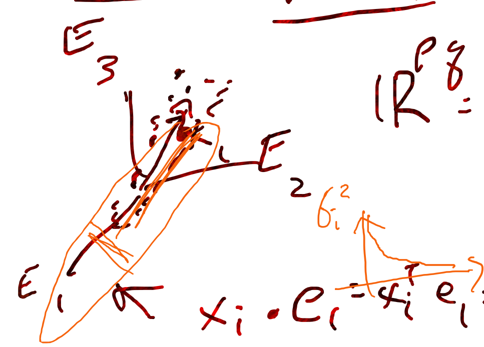

$\mathbb{R}^{pq} = \mathbb{R}^{r}$

$x_i \cdot e_1 = x_i^T e_1 = E$

Eigenfaces for Face Detection/Recognition

(M. Turk and A. Pentland, "Eigenfaces for Recognition", *Journal of Cognitive Neuroscience*, vol. 3, no. 1, pp. 71-86, 1991, hard copy)

• Face Recognition

- The simplest approach is to think of it as a template matching problem:

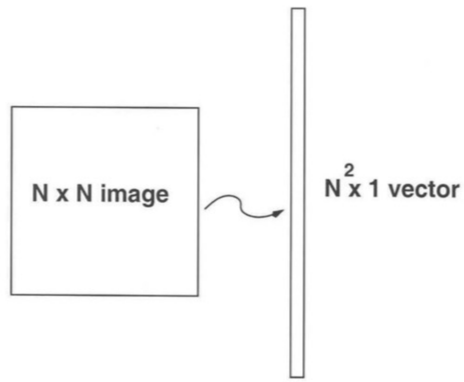

- Problems arise when performing recognition in a high-dimensional space.

- Significant improvements can be achieved by first mapping the data into a *lower-dimensionality* space.

- How to find this lower-dimensional space?

• Main idea behind eigenfaces

- Suppose Γ is an N^2x1 vector, corresponding to an NxN face image I.

- The idea is to represent Γ (Φ=Γ - mean face) into a low-dimensional space:

$$\hat{\Phi} - mean = w_1 u_1 + w_2 u_2 + \cdots w_K u_K \ (K{<<}N^2)$$

Computation of the eigenfaces

Step 1: obtain face images $I_1, I_2, ..., I_M$ (training faces)

(**very important:** the face images must be *centered* and of the same *size*)

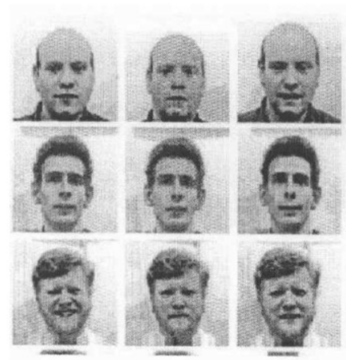

Step 2: represent every image I_i as a vector Γ_i

Step 3: compute the average face vector Ψ:

$$\Psi = \frac{1}{M} \sum_{i=1}^{M} \Gamma_i$$

Step 4: subtract the mean face:

$$\Phi_i = \Gamma_i - \Psi$$

Step 5: compute the covariance matrix C:

$$C = \frac{1}{M} \sum_{n=1}^{M} \Phi_n \Phi_n^T = AA^T \quad (N^2 \mathrm{x} N^2 \text{ matrix})$$

$C = X^T X$

where $A = [\Phi_1 \ \Phi_2 \cdots \Phi_M]$ $(N^2 \mathrm{x} M$ matrix$)$

Step 6: compute the eigenvectors u_i of AA^T

The matrix AA^T is very large --> not practical !!

Step 6.1: consider the matrix $A^T A$ (M x M matrix)

Step 6.2: compute the eigenvectors v_i of $A^T A$

$$A^T A v_i = \mu_i v_i$$

What is the relationship between us_i and v_i?

$A^T A v_i = \mu_i v_i$ => $AA^T A v_i = \mu_i A v_i$ =>

$CAv_i = \mu_i A v_i$ or $Cu_i = \mu_i u_i$ where $u_i = Av_i$

Thus, AA^T and $A^T A$ have the same eigenvalues and their eigenvectors are related as follows: $u_i = Av_i$!!

Note 1: AA^T can have up to N^2 eigenvalues and eigenvectors.

Note 2: $A^T A$ can have up to M eigenvalues and eigenvectors.

Note 3: The M eigenvalues of $A^T A$ (along with their corresponding eigenvectors) correspond to the M *largest* eigenvalues of AA^T (along with their corresponding eigenvectors).

Step 6.3: compute the M best eigenvectors of AA^T: $u_i = Av_i$

(**important:** normalize u_i such that $\|u_i\| = 1$)

Step 7: keep only K eigenvectors (corresponding to the K largest eigenvalues)

$X = U \Sigma V^T$

$U = [u_1 | u_2 | \ldots | u_s | u_{s+1} | \ldots | u_r | u_{r+1} \ldots u_n]$

- imagesc(g)
- g = reshape(U(:,i), p, q)

Representing faces onto this basis

- Each face (minus the mean) Φ_i in the training set can be represented as a linear combination of the best K eigenvectors:

$$\hat{\Phi}_i - mean = \sum_{j=1}^{K} w_j u_j, \quad (w_j = u_j^T \Phi_i)$$

(we call the u_j's *eigenfaces*)

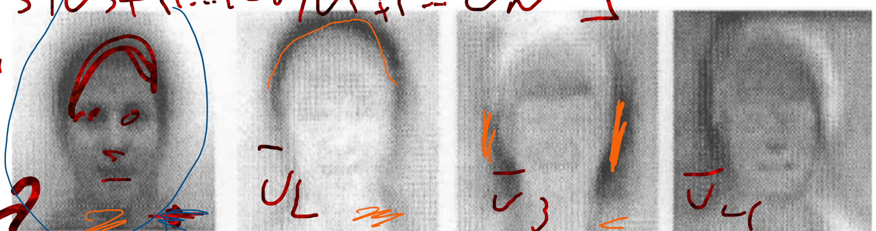

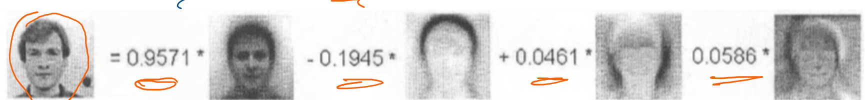

- Each normalized training face Φ_i is represented in this basis by a vector:

$$\Omega_i = \begin{bmatrix} w_1^i \\ w_2^i \\ \ldots \\ w_K^i \end{bmatrix}, \quad i = 1, 2, \ldots, M$$

$\|\vec{a}\|_2^2 = \sum a_i^2$

Parseval ineq.

$\|\vec{a}_i\|_2^2 \approx \|f\|_2^2 = \int |f(x)|^2 dx$

$= a_1 \cdot + a_2 + a_3 + \ldots$

Fourier

Taylor

Bunny

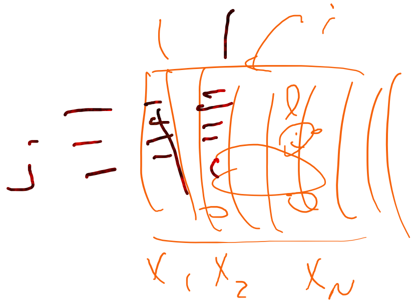

$\rightarrow X = [x_1 | x_2 | \dots | x_N]$ ← one column

$= U \Sigma V^T$

PCA of the bunny

$x_i = a_1^i u_1 + a_2^i u_2 + \dots + a_N^i u_N$

basis

full of as possible for the dataset

$a_1^i = x_i^T u_1 = \mathrm{comp}_{u_1} x_i = \dfrac{u_1 \cdot x_i}{\|u_1\|} = \|x_i\| \cos\theta$

- On PCA

- Optimal

On PCA Principal Component Analysis, Eigenface

-On Raleigh Ritz Quotient

-On Spectral Decomposition Theorem

-On Data Clouds

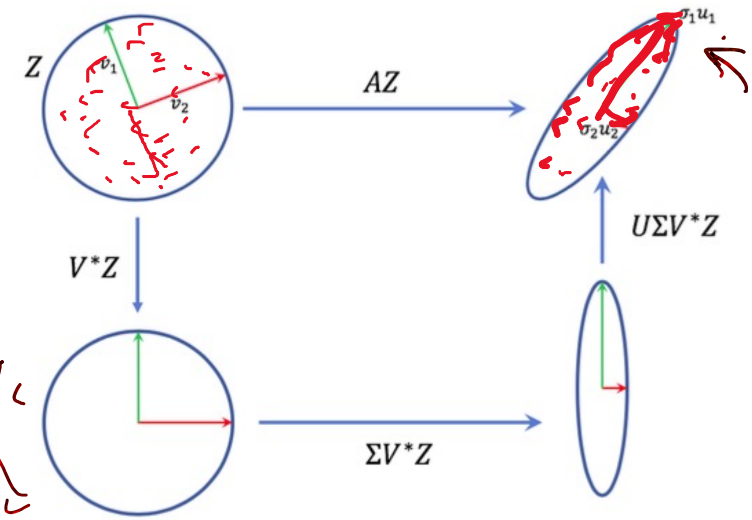

$$\underline{X} = [\,x_1 \mid x_2 \mid \dots \mid x_n\,]$$

$$\vec{x}_i = \underline{a^i_1}\,\vec{v}_1 + \underline{a^i_2}\,\vec{v}_2 + \dots + \underline{a^i_r}\,\vec{v}_r$$

v's as basis set.

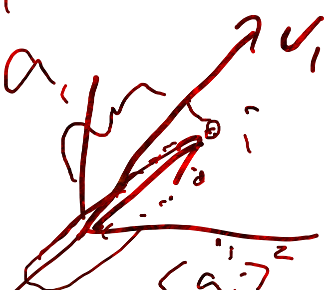

x_i = PCA gives basis set where $\langle a_j \rangle_i \downarrow$ as fast as possible vs any other basis set.

$$\langle a_1 \rangle_i = \frac{1}{n}\sum_{i=1}^{n} a_1^i$$

DNA Microarrays

Gene Expression

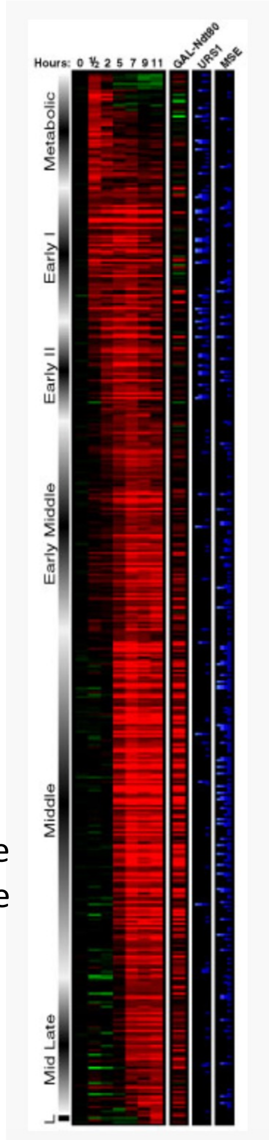

Microarray results that have been analyzed such that the colors were linked with expression and then similar gene profiles were grouped Together – budding yeast

cartoon illustrating an array of DNA snippets on a chip. The top portion depicts a possible nucleotide sequence for the DNA segment immobilized in the position indicated.

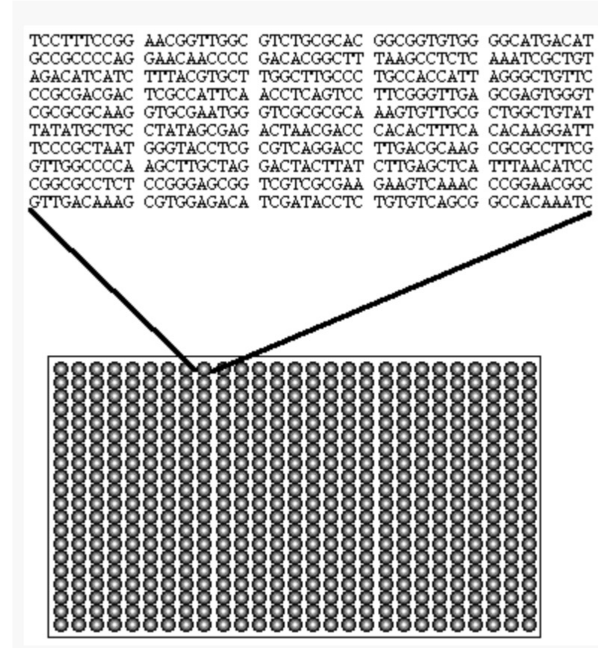

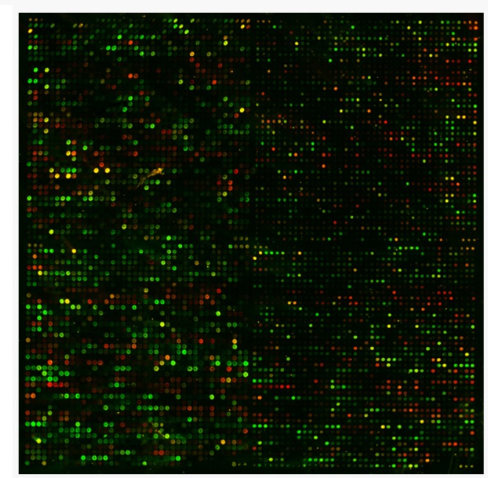

DNA Microarray chip containing the entire yeast genome

Morphological

- Height
- Weight
- Footsize
- Belt (waist) size
- Hand size
- Forearm size
- Head circumference
- Femur length

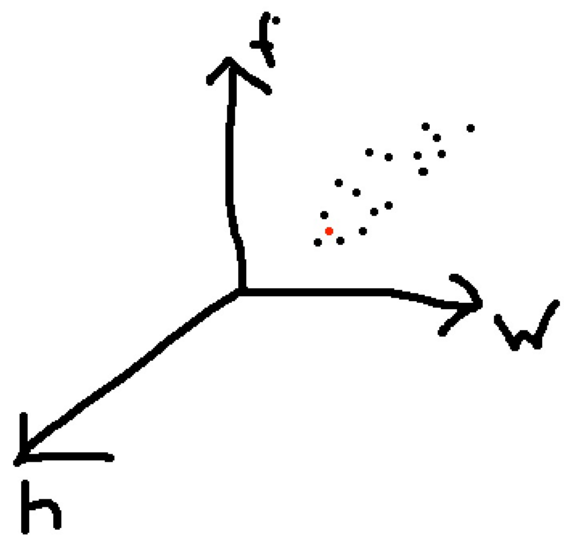

Interpreting as an ellipsoid in the high dimensional space is the simplest geometric interpretation of the data cloud and leads to simplification as major and minor axis, and even Reduced order model (ROM) (meaning a lower dimensional representation).

PCA, SVD, SDT – is optimal

Data for PCA - "Pretend data looks like an ellipsoid"

Ex: $\underline{X_i}$ ~ 4000×1 gene expression table for each i.

$i = 1 \ldots 216$ patients

$$X_i = \begin{pmatrix} x_i^1 \\ x_i^2 \\ \vdots \\ x_{4000} \end{pmatrix}$$

$y_i = 0$ or 1 "0" if not cancer "1" if cancer.

$$\underline{f}: \mathbb{R}^{4000} \to \mathbb{Z}_2 \qquad \mathbb{Z}_2 = \{0, 1\}.$$

$x \in \mathbb{R}^{4000}$

• Supervised vs. unsupervised.

unsupervised - just input - just structure / geometry.

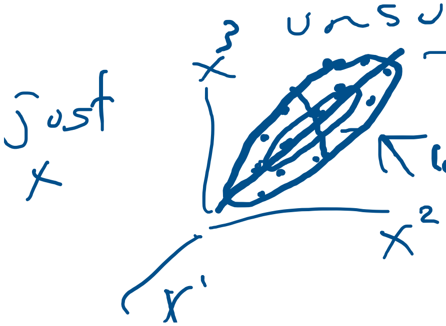

level sets of ρ

"Data Cloud" ~ distribution $\underline{\rho}(x) : \mathbb{R}^{4000} \to \mathbb{R}^+$ of R.V. $x \sim \underline{X}$

• supervised learning is descriptive $f : x \to y$ labels

Let A be a $n \times n$ symmetric matrix. From the spectral theorem, we know that there is an orthonormal basis $u_1, \cdots, u_n$ of $\mathbb{R}^n$ such that each u_j is an eigenvector of A. Let λ_j be the eigenvalue corresponding to u_j, that is,

$$Au_j = \lambda_j u_j.$$

Then

$$A = PDP^{-1} = PDP^T$$

where P is the orthogonal matrix $P = [u_1 \ \cdots \ u_n]$ and D is the diagonal matrix with diagonal entries $\lambda_1, \cdots, \lambda_n$. The equation $A = PDP^T$ can be rewritten as:

$$A = [u_1 \ \cdots \ u_n] \begin{bmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_n \end{bmatrix} \begin{bmatrix} u_1^T \\ \vdots \\ u_n^T \end{bmatrix}$$

$$= [\lambda_1 u_1 \ \cdots \ \lambda_n u_n] \begin{bmatrix} u_1^T \\ \vdots \\ u_n^T \end{bmatrix}$$

outer

$$= \lambda_1 u_1 u_1^T + \cdots + \lambda_n u_n u_n^T.$$

$n \times 1 \cdot 1 \times n$

$n \times n$

$\|u_i\|^2 = u_i \cdot u_i = u_i^T u_i$ scalar - inner product

The expression

$$A = \lambda_1 u_1 u_1^T + \cdots + \lambda_n u_n u_n^T.$$

is called the spectral decomposition of A. Note that each matrix $u_j u_j^T$ has rank 1 and is the matrix of projection onto the one dimensional subspace spanned by u_j. In other words, the linear map P defined by $P(x) = u_j u_j^T x$ is the orthogonal projection onto the subspace spanned by u_j.

- $A = B^T B$ is symmetric
$\Rightarrow$ spectral Decomp. theorem
i.e. also covariance metrices.

- A is pos. definite if
$\lambda_i > 0$ all i.

PCA as algorithm

(say $m = 4000$, $n = 216$)

- Data $\underline{\overline{X}} = \begin{pmatrix} | & | & & | \\ x_1 & x_2 & \dots & x_n \\ | & | & & | \end{pmatrix}_{m \times n}$ ith

- what if $x_i \sim \underline{\mathcal{N}(\bar{x}, \Sigma)}$ — covariance matrix.

$B = \underline{X} - \overline{B}$; $\overline{B} = \begin{pmatrix} | \\ | \\ | \end{pmatrix}_{n \times 1} \bar{x}^T_{1 \times n} = \underline{1}\,\bar{x}^T$. $\bar{x} = \begin{pmatrix} \bar{x}_1 \\ \bar{x}_2 \\ \vdots \\ \bar{x}_n \end{pmatrix}$

$\bar{x}_i = \frac{1}{n} \sum_{i=1}^{n} \underline{\overline{X}}_{ij}$

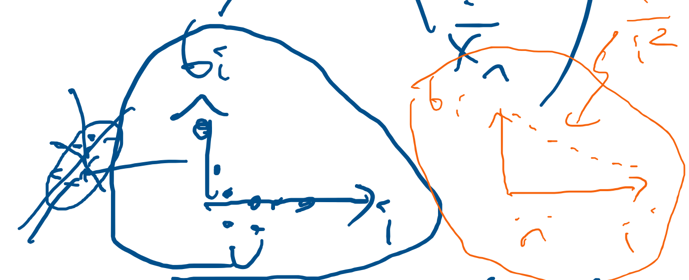

$B = U \Sigma V^T$; $U = [u_1\ u_2 \dots u_n]$

and u_1 is major axis — most energetic — feature that most explains data.

u_2 is first minor axis

Side Note:

$\frac{1}{i^2}$ is slowest converging to zero σ_i^2

i.e. $\frac{1}{i}$ is slowest converging σ_i

$\sigma_i \le \frac{1}{i}$

$$\sum_{i=1}^{\infty} \frac{1}{i^p} < \infty \text{ if } ~~p~~ \text{ (bad)}$$

$p \le 1$, $p = 1$ harmonic

σ_i^2 the power spectrum

$= 1 + \frac{1}{2^p} + \frac{1}{3^p} + \frac{1}{4^p} + \cdots \quad p = 1$

$B = \bar{X}^T$; let $C = \frac{1}{n-1} B^T B$ ~~everywhere~~

is covariance matrix.

$B = U \Sigma V^T$

$U = [u_1 | u_2 \ldots | u_n]$

$\sigma_1 u_1$

$\sigma_2 u_2$

$\sigma_3 \sim 0$

$\sigma_2 \gtrsim 0$

- $u_1 = \underset{\|u\|=1}{\text{argmax}}\ u^T B^T B u = \underset{u}{\text{argmax}} \frac{(u^T B^T)(B u)}{u^T u}$

$\|Bu\|_2^2$

Raleigh - Ritz quotient.

$Bu \cdot Bu = \|Bu\|_2^2$

- $u_2 = \underset{\substack{\|u\|=1 \\ u \perp u_1}}{\text{argmax}}\ u^T B^T B u$

- $CV = VD$ eigenvectors of C all stacked.

Ax optimize $x^T A x$, maximize. $\left(A = B^T B \quad x = 0 \right)$

$$x = \begin{pmatrix} x^1 \\ \vdots \\ x^n \end{pmatrix} \leftarrow j \qquad r(x) = \frac{x^T A x}{x^T x}$$

$$\frac{\partial r}{\partial x^j} = \frac{\partial}{\partial x^j}\left(r(x)\right) = \frac{\frac{\partial}{\partial x^j}(x^T A x) - x^T A x \frac{\partial}{\partial x^j}(x^T x)}{(x^T x)^2}$$

$$= \frac{2(Ax)_j}{x^T x} - \frac{(x^T A x)\, 2x_j}{(x^T x)^2} = \frac{2}{x^T x}\left(Ax - r(x)x\right)_j$$

$$\nabla r(x) = \frac{2}{x^T x}\left(Ax - r(x)x\right) = \frac{2}{x^T x}\left(A - r(x)\right)x = 0$$

$$A x = \underbrace{r(x)}_{\lambda}\, x \implies A x = \lambda x$$

Conclude

The x that optimizes $r(x) = \dfrac{x^T A x}{x^T x}$

is an eigenvector and $r(x)$ is its eigenvalue.

• S.D.T. for $A = B^T B = \sum_{i=1}^{n} \lambda_i u_i u_i^T$

are the v's

let $B = U \Sigma V^T$ svd.

$B^T B = V \Sigma^T \cancel{U^T U} \Sigma V^T = V \underbrace{\Sigma^T \Sigma}_{} V^T$

$$D = \begin{pmatrix} \sigma_1 & & & \\ & \ddots & & \\ & & \sigma_r & \\ & & & \ddots \; 0 \end{pmatrix} \left(\begin{array}{cccc|c} \sigma_1 & & & & \\ & \ddots & & & \\ & & \sigma_r & & 0 \\ & & & \ddots \; 0 & \end{array}\right) = \begin{pmatrix} \sigma_1^2 & & & & \\ & \sigma_2^2 & & & \\ & & \ddots & & \\ & & & \sigma_r^2 & \\ & & & & \ddots \; 0 \end{pmatrix}$$

$$= V D V^T = \sum_{i=1}^{r} \sigma_i^2 v_i v_i^T$$