


CS 523: Computer Graphics, Spring 2011

Shape Modeling


PCA Applications + SVD

Reminder: PCA

- Find principal components of data points
- Orthogonal directions that are dominant in the data (have variance extrema)



Scatter matrix $S = X X^T$



More applications of PCA

Morphable models of faces

- Data base of face scans: 3D geometry + texture (photo)




- 10,000 points in each scan
- x, y, z, R, G, B – 6 numbers for each point
- Thus, each scan is a $10,000 * 6 = \mathbf{60,000}$ -dimensional vector

See: V. Blanz and T. Vetter, A Morphable Model for the Synthesis of 3D Faces, SIGGRAPH 99

More applications of PCA

Morphable models of faces

- How to find interesting axes in this 60000-dimensional space?
 - axes that measure age, gender, etc...
 - There is hope: the faces are likely to be governed by a small set of parameters (much less than 60,000...)




FaceGen demo

Singular Value Decomposition



Geometric analysis of linear transformations

- We want to know what a linear transformation A does
- Need some simple and “comprehensible” representation of the matrix A
- Let’s look what A does to some vectors
 - Since $A(\alpha\mathbf{v}) = \alpha A(\mathbf{v})$, it’s enough to look at vectors \mathbf{v} of unit length





Geometric analysis of linear transformations

- A linear (non-singular) transform A always takes hyper-spheres to hyper-ellipses.




Geometric analysis of linear transformations

- Thus, one good way to understand what A does is to find which vectors are mapped to the “main axes” of the ellipsoid




Geometric analysis of linear transformations

- If A is symmetric: $A = V D V^T$, V orthogonal
- The eigenvectors of A are the axes of the ellipse




Symmetric matrix: eigendecomposition

- In this case A is just a scaling matrix. The **eigendecomposition** of A tells us which orthogonal axes it scales, and by how much




General linear transformations: Singular Value Decomposition

- In general A will also contain rotations, not just scales



$$A = [\mathbf{u}_1 \ \mathbf{u}_2 \ \dots \ \mathbf{u}_n] \begin{bmatrix} \sigma_1 & & & \\ & \sigma_2 & & \\ & & \ddots & \\ & & & \sigma_n \end{bmatrix} [\mathbf{v}_1 \ \mathbf{v}_2 \ \dots \ \mathbf{v}_n]^T$$

General linear transformations: Singular Value Decomposition



$$\mathbf{A} \mathbf{V} = \mathbf{U} \mathbf{\Sigma}$$

$$\begin{array}{c}
 \textit{orthonormal} \\
 \mathbf{A} [\mathbf{v}_1 \ \mathbf{v}_2 \ \dots \ \mathbf{v}_n] = [\mathbf{u}_1 \ \mathbf{u}_2 \ \dots \ \mathbf{u}_n] \begin{bmatrix} \sigma_1 & & & \\ & \sigma_2 & & \\ & & \ddots & \\ & & & \sigma_n \end{bmatrix}
 \end{array}$$

$$\mathbf{A} \mathbf{v}_i = \sigma_i \mathbf{u}_i, \quad \sigma_i \geq 0$$

Some history

- SVD was discovered by the following people:



E. Beltrami
(1835 – 1900)



M. Jordan
(1838 – 1922)



J. Sylvester
(1814 – 1897)



E. Schmidt
(1876-1959)




H. Weyl
(1885-1955)

SVD

- SVD exists for any matrix
- Formal definition:
 - For square matrices $A \in R^{n \times n}$, there exist orthogonal matrices $U, V \in R^{n \times n}$ and a diagonal matrix Σ , such that all the diagonal values σ_i of Σ are non-negative and


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



SVD


- The diagonal values of Σ are called the **singular values**. It is accustomed to sort them: $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_n$
- The columns of U ($\mathbf{u}_1, \dots, \mathbf{u}_n$) are called the **left singular vectors**. They are the axes of the ellipsoid.
- The columns of V ($\mathbf{v}_1, \dots, \mathbf{v}_n$) are called the **right singular vectors**. They are the preimages of the axes of the ellipsoid.

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




Reduced SVD

- For rectangular matrices, we have two forms of SVD. The reduced SVD looks like this:
 - The columns of U are orthonormal
 - Cheaper form for computation and storage



Full SVD

- We can complete U to a full orthogonal matrix and pad Σ by zeros accordingly



SVD

Applications

- There are stable numerical algorithms to compute SVD (albeit not cheap). Once you have it, you have many things:
 - Matrix inverse \rightarrow can solve square linear systems
 - Numerical rank of a matrix
 - Can solve linear least-squares systems
 - PCA
 - Many more...

Matrix inverse and solving linear systems

- Matrix inverse

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

$$A^{-1} = (U\Sigma V^T)^{-1} = (V^T)^{-1} \Sigma^{-1} U^{-1} =$$


$$= V \begin{pmatrix} \frac{1}{\sigma_1} & & & \\ & \ddots & & \\ & & \ddots & \\ & & & \frac{1}{\sigma_n} \end{pmatrix} U^T$$

- So, to solve $A\mathbf{x} = \mathbf{b}$

$$\mathbf{x} = V\Sigma^{-1}U^T\mathbf{b}$$

Matrix rank

- The rank of A is the number of non-zero singular values



Numerical rank

- If there are very small singular values, then A is close to being singular. We can set a threshold t , so that
$$\text{numeric_rank}(A) = \#\{\sigma_i \mid \sigma_i > t\}$$
- Using SVD is a numerically stable way! The determinant is not a good way to check singularity

PCA

- Construct the matrix X of the centered data points

$$X = \begin{pmatrix} | & | & \cdots & | \\ \mathbf{p}'_1 & \mathbf{p}'_2 & \cdots & \mathbf{p}'_n \\ | & | & & | \end{pmatrix}$$

- The principal axes are eigenvectors of $S = XX^T$

$$S = XX^T = U \begin{pmatrix} \lambda_1 & & & \\ & \ddots & & \\ & & & \lambda_d \end{pmatrix} U^T$$

PCA

- We can compute the principal components by SVD of X :

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


$$\begin{aligned} XX^T &= U\Sigma V^T (U\Sigma V^T)^T = \\ &= U\Sigma V^T V \Sigma U^T = U \underline{\Sigma^2} U^T \end{aligned}$$

- Thus, the **left singular vectors** of X are the principal components! We sort them by the size of the singular values of X .

Least-squares rotation with SVD


Shape matching

- We have two objects in correspondence
- Want to find the rigid transformation that aligns them



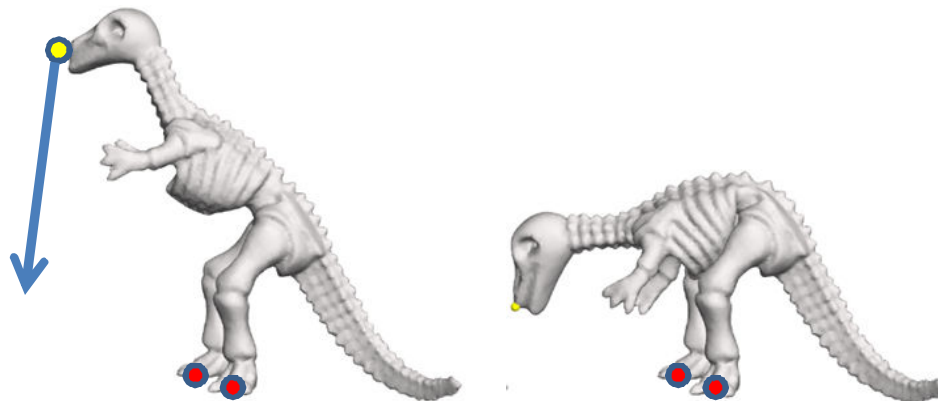

Shape matching

- When the objects are aligned, the lengths of the connecting lines are small



Optimal local rotation

- We will use this for mesh deformation



Shape matching – formalization

- Align two point sets

$$P = \{\mathbf{p}_1, \dots, \mathbf{p}_n\} \quad \text{and} \quad Q = \{\mathbf{q}_1, \dots, \mathbf{q}_n\}.$$

- Find a translation vector \mathbf{t} and rotation matrix \mathbf{R} so that

$$\sum_{i=1}^n \left\| (\mathbf{R}\mathbf{p}_i + \mathbf{t}) - \mathbf{q}_i \right\|^2 \quad \text{is minimized}$$

Shape matching – solution

- Solve translation and rotation separately
 - If (\mathbf{R}, \mathbf{t}) is the optimal transformation, then the point sets $\{\mathbf{R}\mathbf{p}_i + \mathbf{t}\}$ and $\{\mathbf{q}_i\}$ have the same centers of mass

$$\bar{\mathbf{p}} = \frac{1}{n} \sum_{i=1}^n \mathbf{p}_i \quad \bar{\mathbf{q}} = \frac{1}{n} \sum_{i=1}^n \mathbf{q}_i$$

$$\bar{\mathbf{q}} = \frac{1}{n} \sum_{i=1}^n (\mathbf{R}\mathbf{p}_i + \mathbf{t}) = \mathbf{R} \left(\frac{1}{n} \sum_{i=1}^n \mathbf{p}_i \right) + \mathbf{t} = \mathbf{R}\bar{\mathbf{p}} + \mathbf{t}$$

⇓

$$\mathbf{t} = \bar{\mathbf{q}} - \mathbf{R}\bar{\mathbf{p}}$$

Finding the rotation R

- To find the optimal R , we bring the centroids of both point sets to the origin

$$\mathbf{x}_i = \mathbf{p}_i - \bar{\mathbf{p}} \quad \mathbf{y}_i = \mathbf{q}_i - \bar{\mathbf{q}}$$

- We want to find R that minimizes

$$\sum_{i=1}^n \|\mathbf{R}\mathbf{x}_i - \mathbf{y}_i\|^2$$

Finding the rotation R

$$\begin{aligned} \sum_{i=1}^n \|\mathbf{R}\mathbf{x}_i - \mathbf{y}_i\|^2 &= \sum_{i=1}^n (\mathbf{R}\mathbf{x}_i - \mathbf{y}_i)^T (\mathbf{R}\mathbf{x}_i - \mathbf{y}_i) = \\ &= \sum_{i=1}^n \left(\underbrace{\mathbf{x}_i^T \mathbf{R}^T \mathbf{R} \mathbf{x}_i}_{\mathbf{I}} - \mathbf{y}_i^T \mathbf{R} \mathbf{x}_i - \mathbf{x}_i^T \mathbf{R}^T \mathbf{y}_i + \mathbf{y}_i^T \mathbf{y}_i \right) \end{aligned}$$

These terms do not depend on R ,
so we can ignore them in the minimization

Finding the rotation R

$$\min_{\mathbf{R}} \sum_{i=1}^n \left(-\mathbf{y}_i^T \mathbf{R} \mathbf{x}_i - \mathbf{x}_i^T \mathbf{R}^T \mathbf{y}_i \right) = \max_{\mathbf{R}} \sum_{i=1}^n \left(\mathbf{y}_i^T \mathbf{R} \mathbf{x}_i + \underbrace{\mathbf{x}_i^T \mathbf{R}^T \mathbf{y}_i}_{\text{this is a scalar}} \right)$$

$$\mathbf{x}_i^T \mathbf{R}^T \mathbf{y}_i = \left(\mathbf{x}_i^T \mathbf{R}^T \mathbf{y}_i \right)^T = \mathbf{y}_i^T \mathbf{R} \mathbf{x}_i$$

$$\Rightarrow \boxed{\operatorname{argmax}_{\mathbf{R}} \sum_{i=1}^n \mathbf{y}_i^T \mathbf{R} \mathbf{x}_i}$$

Finding the rotation R

$$\sum_{i=1}^n \mathbf{y}_i^T \mathbf{R} \mathbf{x}_i = \text{tr}(\mathbf{Y}^T \mathbf{R} \mathbf{X})$$

$$\text{tr}(\mathbf{A}) = \sum_{i=1}^n A_{ii}$$

$$\begin{array}{c}
 \begin{array}{|c|} \hline -\mathbf{y}_1^T- \\ \hline -\mathbf{y}_2^T- \\ \hline \vdots \\ \hline -\mathbf{y}_n^T- \\ \hline \end{array} \\
 \mathbf{Y}^T
 \end{array}
 \mathbf{R}
 \begin{array}{|c|c|c|c|} \hline | & | & \cdots & | \\ \hline \mathbf{x}_1 & \mathbf{x}_2 & & \mathbf{x}_n \\ \hline | & | & & | \\ \hline \end{array}
 =
 \begin{array}{c}
 \begin{array}{|c|} \hline -\mathbf{y}_1^T- \\ \hline -\mathbf{y}_2^T- \\ \hline \vdots \\ \hline -\mathbf{y}_n^T- \\ \hline \end{array} \\
 \mathbf{Y}^T
 \end{array}
 \begin{array}{|c|c|c|c|} \hline | & | & \cdots & | \\ \hline \mathbf{R}\mathbf{x}_1 & \mathbf{R}\mathbf{x}_2 & & \mathbf{R}\mathbf{x}_n \\ \hline | & | & & | \\ \hline \end{array}$$

\mathbf{X}

Finding the rotation R

$$\sum_{i=1}^n \mathbf{y}_i^T \mathbf{R} \mathbf{x}_i = \text{tr}(\mathbf{Y}^T \mathbf{R} \mathbf{X})$$

$$\text{tr}(\mathbf{A}) = \sum_{i=1}^n A_{ii}$$

The diagram illustrates the trace of a matrix product. On the left, a vertical column of vectors $\mathbf{y}_1^T, \mathbf{y}_2^T, \dots, \mathbf{y}_n^T$ is shown, with each vector enclosed in a box. To its right is a horizontal row of vectors $\mathbf{R}\mathbf{x}_1, \mathbf{R}\mathbf{x}_2, \dots, \mathbf{R}\mathbf{x}_n$, also enclosed in a box. An equals sign follows, leading to a large square box containing the resulting matrix of scalar products: $\mathbf{y}_1^T \mathbf{R} \mathbf{x}_1$ in the top-left corner, $\mathbf{y}_2^T \mathbf{R} \mathbf{x}_2$ in the second row, second column, and $\mathbf{y}_n^T \mathbf{R} \mathbf{x}_n$ in the bottom-right corner, with diagonal dots indicating the rest of the matrix.

Finding the rotation R

- Find R that maximizes

$$\text{tr}(Y^T R X) = \text{tr}(R X Y^T) \quad (\text{because } \text{tr}(AB) = \text{tr}(BA))$$

- Let's do SVD on $S = X Y^T$

$$S = X Y^T = U \Sigma V^T$$



$$\text{tr}(R X Y^T) = \text{tr}(\underbrace{R U}_{\text{orthogonal matrix}} \underbrace{\Sigma V^T}) = \text{tr}(\Sigma (\underbrace{V^T R U}_{\text{orthogonal matrix}}))$$

Finding the rotation R

- We want to maximize

$$\text{tr}\left(\Sigma\left(\underbrace{V^T R U}\right)\right)$$

orthogonal matrix
all entries ≤ 1

| | | |
|------------|------------|------------|
| σ_1 | | |
| | σ_2 | |
| | | σ_3 |

| | | |
|----------|----------|----------|
| m_{11} | \dots | |
| \vdots | m_{22} | \vdots |
| | \dots | m_{33} |

$$\text{tr}\left(\Sigma\left(V^T R U\right)\right) = \sum_{i=1}^3 \sigma_i m_{ii} \leq \sum_{i=1}^3 \sigma_i$$

Finding the rotation R

$$\text{tr}(\Sigma(V^T R U)) = \sum_{i=1}^3 \sigma_i m_{ii} \leq \sum_{i=1}^3 \sigma_i$$

- Our best shot is $m_{ii} = 1$, i.e. to make $V^T R U = I$

$$V^T R U = I$$

$$R U = V$$

$$R = V U^T$$

Summary of rigid alignment

- Translate the input points to the centroids

$$\mathbf{x}_i = \mathbf{p}_i - \bar{\mathbf{p}} \quad \mathbf{y}_i = \mathbf{q}_i - \bar{\mathbf{q}}$$

- Compute the “covariance matrix”

$$\mathbf{S} = \mathbf{X}\mathbf{Y}^T = \sum_{i=1}^n \mathbf{x}_i \mathbf{y}_i^T$$

- Compute the SVD of \mathbf{S}


$$\mathbf{S} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$$

- The optimal orthogonal \mathbf{R} is

$$\mathbf{R} = \mathbf{V}\mathbf{U}^T$$


Sign correction

- It is possible that $\det(VU^T) = -1$: sometimes reflection is the best orthogonal transform




Sign correction

- It is possible that $\det(VU^T) = -1$: sometimes reflection is the best orthogonal transform



Sign correction

- It is possible that $\det(VU^T) = -1$: sometimes reflection is the best orthogonal transform




- To restrict ourselves to rotations only:
take the last column of V (corresponding to the smallest singular value) and invert its sign.
- Why? See the PDF...

Complexity

- Numerical SVD is an expensive operation $O(\min(mn^2, nm^2))$
- We always need to pay attention to the dimensions of the matrix we're applying SVD to.

SVD for animation compression



Chicken animation

See:

Representing Animations by Principal Components, M. Alexa and W. Muller, Eurographics 2000

Compression of Soft-Body Animation Sequences, Z. Karni and C. Gotsman, Computers&Graphics 28(1): 25-34, 2004


Key Point Subspace Acceleration and Soft Caching, M. Meyer and J. Anderson, SIGGRAPH 2007

Andrew Nealen, Rutgers, 2011

2/15/2011


3D animations

- Each frame is a 3D model (mesh)



3D animations

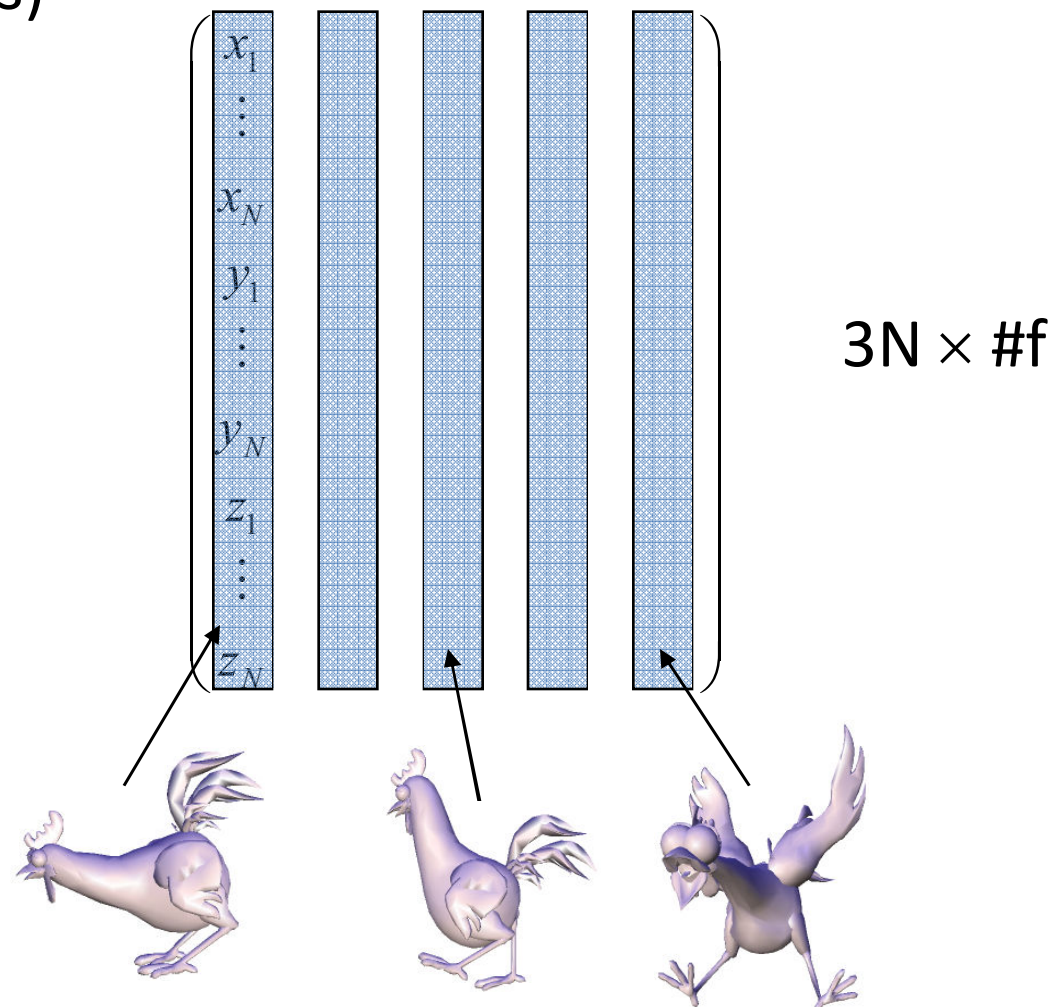
- Connectivity is usually constant (at least on large segments of the animation)
- The geometry changes in each frame → vast amount of data!



13 seconds, 3000 vertices/frame, 26 MB

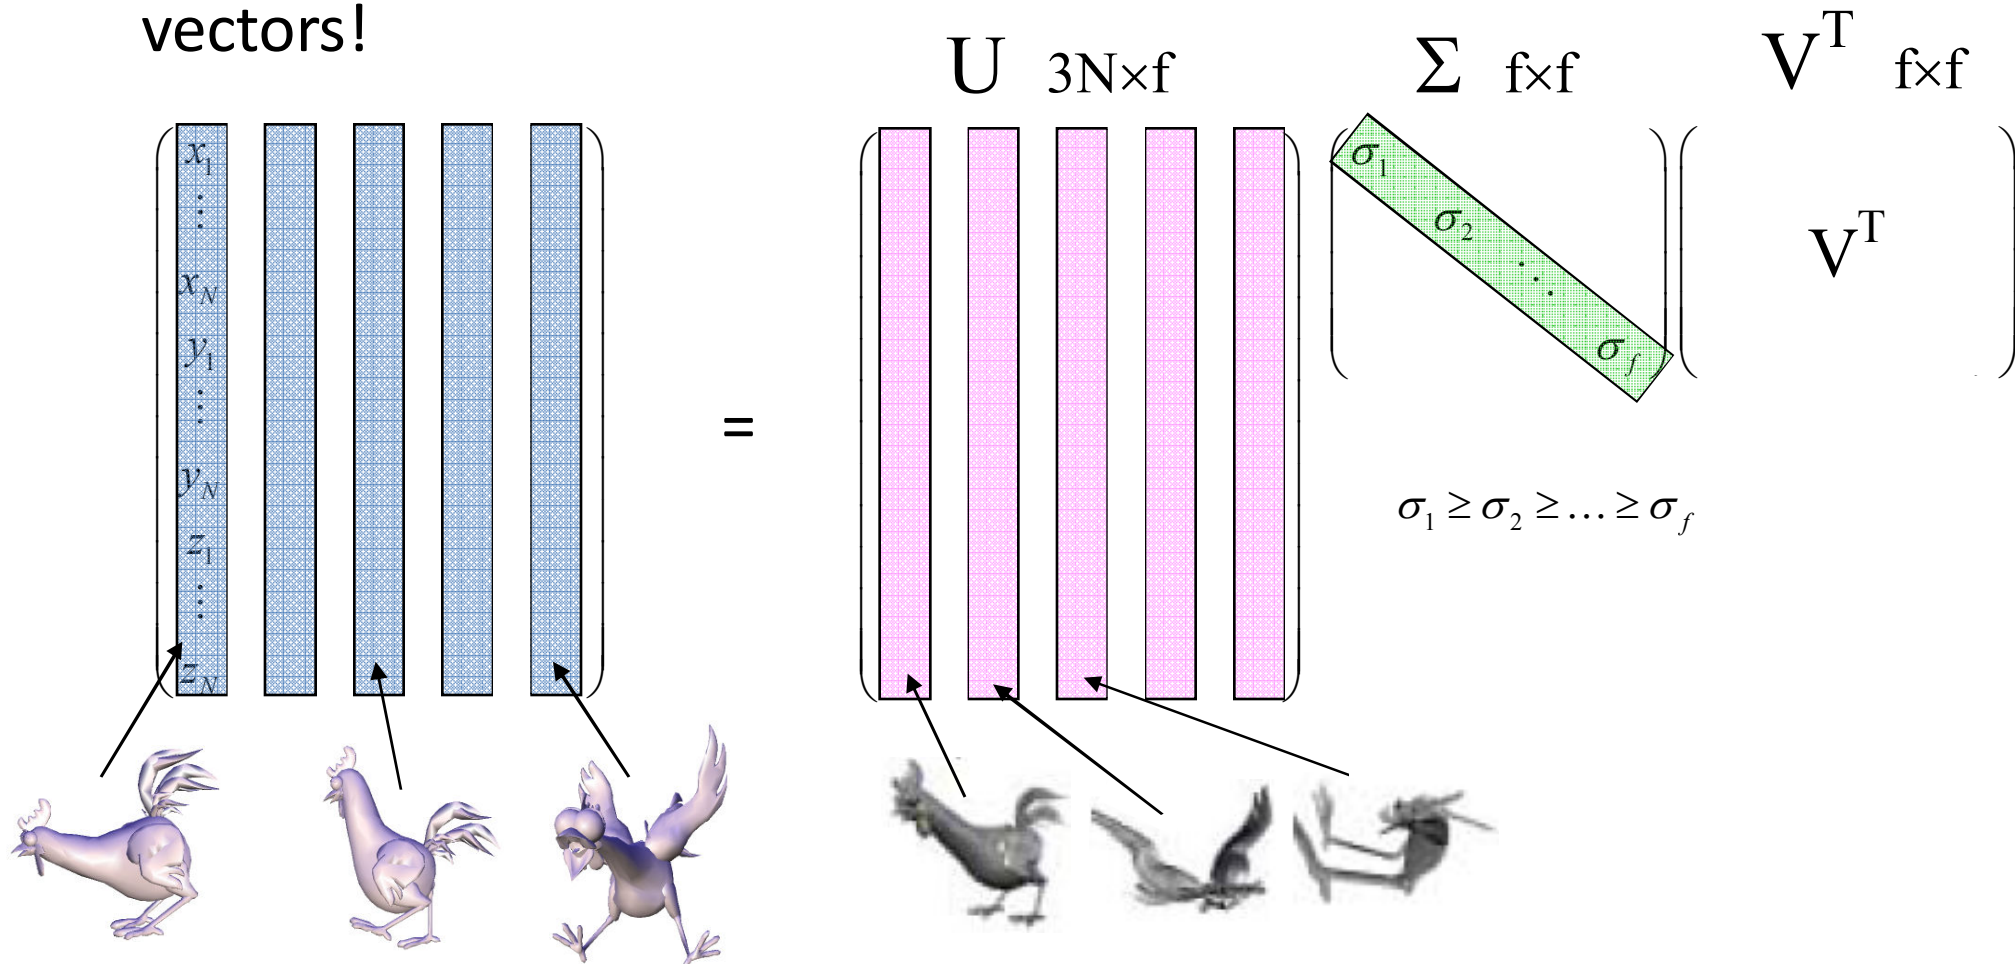
Animation compression by dimensionality reduction

- The geometry of each frame is a vector in \mathbb{R}^{3N} space
($N = \text{\#vertices}$)




Animation compression by dimensionality reduction

- Find a few vectors of \mathbb{R}^{3N} that will best represent our frame vectors!



Animation compression by dimensionality reduction

- The first principal components are the important ones




Animation compression by dimensionality reduction

- Approximate each frame by linear combination of the first principal components
- The more components we use, the better the approximation
- Usually, the number of components needed is much smaller than f .


$$\begin{pmatrix} x_1 \\ \vdots \\ x_N \\ y_1 \\ \vdots \\ y_N \\ z_1 \\ \vdots \\ z_N \end{pmatrix} = \alpha_1 \mathbf{u}_1 + \alpha_2 \mathbf{u}_2 + \alpha_3 \mathbf{u}_3$$

Animation compression by dimensionality reduction

- Compressed representation:
 - The chosen principal component vectors
 - Coefficients α_i for each frame



Animation with only
2 principal components



Animation with
20 out of 400 principal
components

Eigenfaces

- Same principal components analysis can be applied to images




Eigenfaces

- Each image is a vector in $\mathbb{R}^{250 \cdot 300}$
- Want to find the principal axes – vectors that best represent the input database of images



Reconstruction with a few vectors

- Represent each image by the first few (n) principal components



$$\mathbf{v} = \alpha_1 \mathbf{u}_1 + \alpha_2 \mathbf{u}_2 + \dots + \alpha_n \mathbf{u}_n = (\alpha_1, \alpha_2, \dots, \alpha_n)$$

Face recognition

- Given a new image of a face, $\mathbf{w} \in \mathbb{R}^{250 \cdot 300}$
- Represent \mathbf{w} using the first n PCA vectors:


$$\mathbf{w} = \alpha_1 \mathbf{u}_1 + \alpha_2 \mathbf{u}_2 + \dots + \alpha_n \mathbf{u}_n = (\alpha_1, \alpha_2, \dots, \alpha_n)$$

- Now find an image in the database whose representation in the PCA basis is the closest:

$$\mathbf{w}' = (\alpha'_1, \alpha'_2, \dots, \alpha'_n)$$

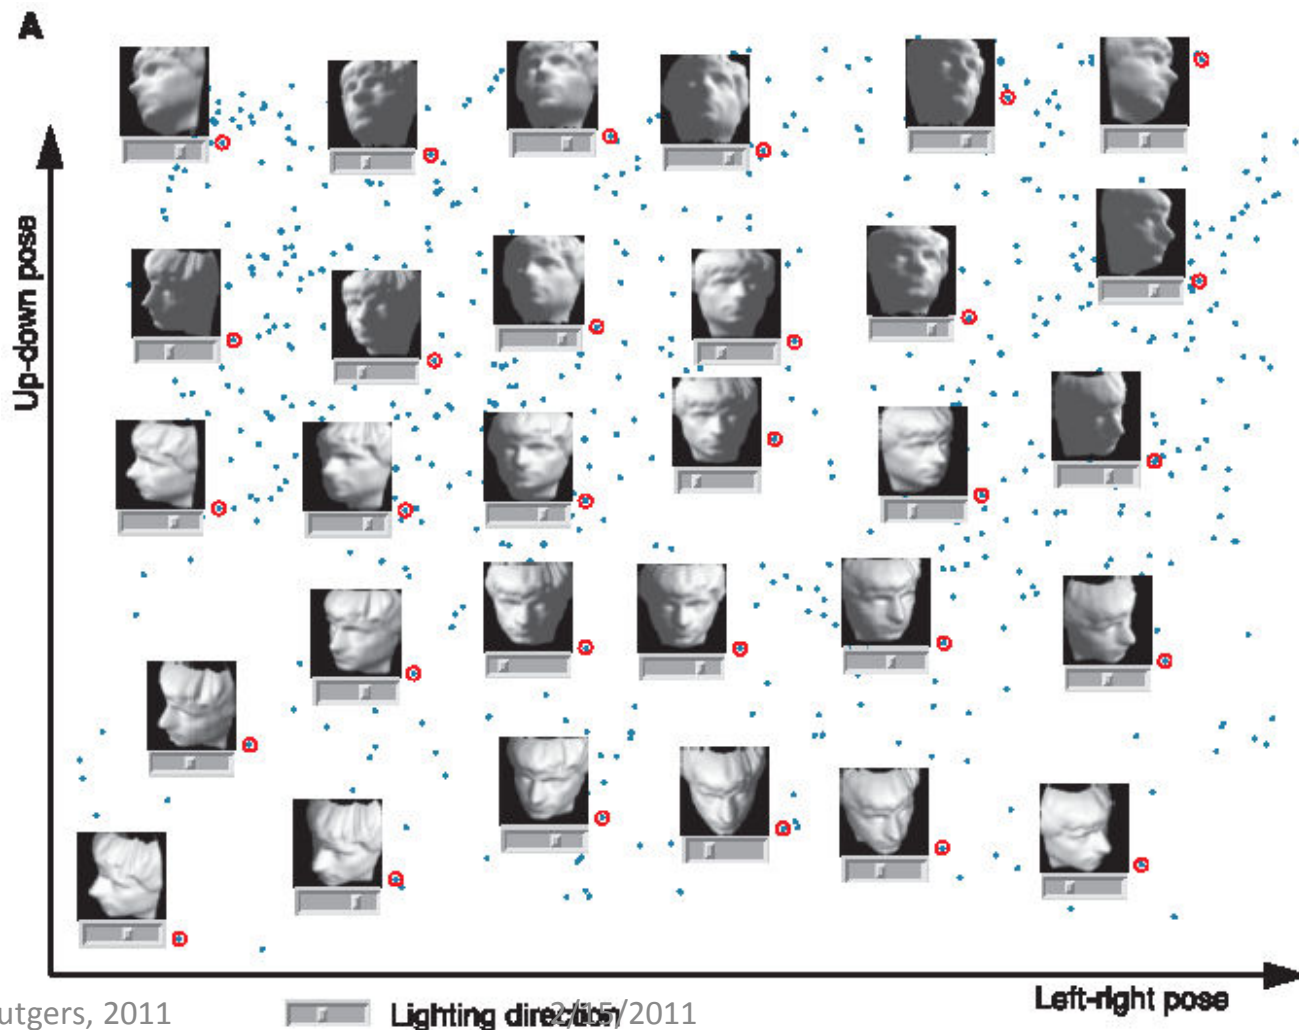
$\langle \mathbf{w}', \mathbf{w} \rangle$ is the largest

The angle between \mathbf{w} and \mathbf{w}' is the smallest



Non-linear dimensionality reduction

- More sophisticated methods can discover non-linear structures in the face datasets



Isomap,
Science, Dec. 2000