

The official seal of the University of Delaware, which is a circular emblem. It features a central shield with an open book. The book's pages are inscribed with the words 'GRAMM', 'PHILOSOPHY', 'RHETORIC', 'ETHICA', 'METAPHYSICS', 'LOGIC', 'MATHEMATICS', and 'PHYSICA'. The shield is surrounded by a circular border containing the Latin motto 'SOLUS MENSA' and the year '1743'.

FSAN/ELEG815: Statistical Learning

Gonzalo R. Arce

Department of Electrical and Computer Engineering
University of Delaware

2. Eigen Analysis, SVD, PCA, and Matrix Completion

Outline

Eigen Analysis

Eigen Properties

SVD

PCA

Matrix Completion

Introduction

Problem Formulation

Optimization Problem

Algorithms

Image Inpainting

Eigen Analysis

Objective: Utilize tools from linear algebra to characterize and analyze matrices, especially the correlation matrix

- ▶ The correlation matrix plays a large role in statistical characterization and processing.
- ▶ Previously result: \mathbf{R} is Hermitian.
- ▶ Further insight into the correlation matrix is achieved through eigen analysis
 - ▶ Eigenvalues and vectors
 - ▶ Matrix diagonalization
 - ▶ Application: Optimum filtering problems

Objective: For a Hermitian matrix \mathbf{R} , find a vector \mathbf{q} satisfying

$$\mathbf{R}\mathbf{q} = \lambda\mathbf{q}$$

- ▶ **Interpretation:** Linear transformation by \mathbf{R} changes the scale, but not the direction of \mathbf{q}
- ▶ **Fact:** A $M \times M$ matrix \mathbf{R} has M eigenvectors and eigenvalues

$$\mathbf{R}\mathbf{q}_i = \lambda_i\mathbf{q}_i \quad i = 1, 2, 3, \dots, M$$

To see this, note

$$(\mathbf{R} - \lambda\mathbf{I})\mathbf{q} = \mathbf{0}$$

For this to be true, the row/columns of $(\mathbf{R} - \lambda\mathbf{I})$ must be linearly dependent,

$$\Rightarrow \det(\mathbf{R} - \lambda\mathbf{I}) = 0$$

Note: $\det(\mathbf{R} - \lambda\mathbf{I})$ is a M th order polynomial in λ

- ▶ The roots of the polynomial are the eigenvalues $\lambda_1, \lambda_2, \dots, \lambda_M$

$$\mathbf{R}\mathbf{q}_i = \lambda_i\mathbf{q}_i$$

- ▶ Each eigenvector \mathbf{q}_i is associated with one eigenvalue λ_i
- ▶ The eigenvectors are not unique

$$\begin{aligned}\mathbf{R}\mathbf{q}_i &= \lambda_i\mathbf{q}_i \\ \Rightarrow \mathbf{R}(a\mathbf{q}_i) &= \lambda_i(a\mathbf{q}_i)\end{aligned}$$

Consequence: eigenvectors are generally normalized, e.g., $|\mathbf{q}_i| = 1$ for $i = 1, 2, \dots, M$

Example (General two dimensional case)

Let $M = 2$ and

$$\mathbf{R} = \begin{bmatrix} R_{1,1} & R_{1,2} \\ R_{2,1} & R_{2,2} \end{bmatrix}$$

Determine the eigenvalues and eigenvectors.

Thus

$$\begin{aligned} \det(\mathbf{R} - \lambda \mathbf{I}) &= 0 \\ \Rightarrow \begin{vmatrix} R_{1,1} - \lambda & R_{1,2} \\ R_{2,1} & R_{2,2} - \lambda \end{vmatrix} &= 0 \\ \Rightarrow \lambda^2 - \lambda(R_{1,1} + R_{2,2}) + (R_{1,1}R_{2,2} - R_{1,2}R_{2,1}) &= 0 \\ \Rightarrow \lambda_{1,2} &= \frac{1}{2} \left[(R_{1,1} + R_{2,2}) \pm \sqrt{4R_{1,2}R_{2,1} + (R_{1,1} - R_{2,2})^2} \right] \end{aligned}$$

Back substitution yields the eigenvectors:

$$\begin{bmatrix} R_{1,1} - \lambda & R_{1,2} \\ R_{2,1} & R_{2,2} - \lambda \end{bmatrix} \begin{bmatrix} q_1 \\ q_2 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$$

In general, this yields a set of linear equations. In the $M = 2$ case:

$$(R_{1,1} - \lambda)q_1 + R_{1,2}q_2 = 0$$

$$R_{2,1}q_1 + (R_{2,2} - \lambda)q_2 = 0$$

- ▶ Solving the set of linear equations for a specific eigenvalue λ_i yields the corresponding eigenvector, \mathbf{q}_i

Example (Two-dimensional white noise)

Let \mathbf{R} be the correlation matrix of a two-sample vector of zero mean white noise

$$\mathbf{R} = \begin{bmatrix} \sigma^2 & 0 \\ 0 & \sigma^2 \end{bmatrix}$$

Determine the eigenvalues and eigenvectors.

Carrying out the analysis yields eigenvalues

$$\begin{aligned} \lambda_{1,2} &= \frac{1}{2} \left[(R_{1,1} + R_{2,2}) \pm \sqrt{4R_{1,2}R_{2,1} + (R_{1,1} - R_{2,2})^2} \right] \\ &= \frac{1}{2} \left[(\sigma^2 + \sigma^2) \pm \sqrt{0 + (\sigma^2 - \sigma^2)^2} \right] = \sigma^2 \end{aligned}$$

and eigenvectors

$$\mathbf{q}_1 = \begin{bmatrix} 1 \\ 0 \end{bmatrix} \quad \text{and} \quad \mathbf{q}_2 = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$$

Note: The eigenvectors are unit length (and orthogonal).

Eigen Properties

Property (eigenvalues of \mathbf{R}^k)

If $\lambda_1, \lambda_2, \dots, \lambda_M$ are the eigenvalues of \mathbf{R} , then $\lambda_1^k, \lambda_2^k, \dots, \lambda_M^k$ are the eigenvalues of \mathbf{R}^k .

Proof: Note $\mathbf{R}\mathbf{q}_i = \lambda_i\mathbf{q}_i$. Multiplying both sides by \mathbf{R} $k - 1$ times,

$$\mathbf{R}^k \mathbf{q}_i = \lambda_i \mathbf{R}^{k-1} \mathbf{q}_i = \lambda_i^k \mathbf{q}_i$$

Property (linear independence of eigenvectors)

The eigenvectors $\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_M$, of \mathbf{R} are linearly independent, i.e.,

$$\sum_{i=1}^M a_i \mathbf{q}_i \neq \mathbf{0}$$

for all nonzero scalars a_1, a_2, \dots, a_M .

Property (Correlation matrix eigenvalues are real & nonnegative)

The eigenvalues of \mathbf{R} are real and nonnegative.

Proof:

$$\begin{aligned}\mathbf{R}\mathbf{q}_i &= \lambda_i\mathbf{q}_i \\ \Rightarrow \mathbf{q}_i^H\mathbf{R}\mathbf{q}_i &= \lambda_i\mathbf{q}_i^H\mathbf{q}_i \quad [\text{pre-multiply by } \mathbf{q}_i^H] \\ \Rightarrow \lambda_i &= \frac{\mathbf{q}_i^H\mathbf{R}\mathbf{q}_i}{\mathbf{q}_i^H\mathbf{q}_i} \geq 0\end{aligned}$$

Follows from the facts: \mathbf{R} is positive semi-definite and $\mathbf{q}_i^H\mathbf{q}_i = |\mathbf{q}_i|^2 > 0$

Note: In most cases, \mathbf{R} is positive definite and

$$\lambda_i > 0, \quad i = 1, 2, \dots, M$$

Property (Unique eigenvalues \Rightarrow orthogonal eigenvectors)

If $\lambda_1, \lambda_2, \dots, \lambda_M$ are unique eigenvalues of \mathbf{R} , then the corresponding eigenvectors, $\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_M$, are orthogonal.

Proof:

$$\begin{aligned} \mathbf{R}\mathbf{q}_i &= \lambda_i\mathbf{q}_i \\ \Rightarrow \mathbf{q}_j^H \mathbf{R}\mathbf{q}_i &= \lambda_i \mathbf{q}_j^H \mathbf{q}_i \end{aligned} \quad (*)$$

Also, since λ_j is real and \mathbf{R} is Hermitian

$$\begin{aligned} \mathbf{R}\mathbf{q}_j &= \lambda_j\mathbf{q}_j \\ \Rightarrow \mathbf{q}_j^H \mathbf{R} &= \lambda_j \mathbf{q}_j^H \\ \Rightarrow \mathbf{q}_j^H \mathbf{R}\mathbf{q}_i &= \lambda_j \mathbf{q}_j^H \mathbf{q}_i \end{aligned}$$

Substituting the LHS from (*)

$$\Rightarrow \lambda_i \mathbf{q}_j^H \mathbf{q}_i = \lambda_j \mathbf{q}_j^H \mathbf{q}_i$$

Thus

$$\begin{aligned}\lambda_i \mathbf{q}_j^H \mathbf{q}_i &= \lambda_j \mathbf{q}_j^H \mathbf{q}_i \\ \Rightarrow (\lambda_i - \lambda_j) \mathbf{q}_j^H \mathbf{q}_i &= 0\end{aligned}$$

Since $\lambda_1, \lambda_2, \dots, \lambda_M$ are unique

$$\mathbf{q}_j^H \mathbf{q}_i = 0 \quad i \neq j$$

$\Rightarrow \mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_M$ are orthogonal.

QED

Diagonalization of \mathbf{R}

Objective: Find a transformation that transforms the correlation matrix into a diagonal matrix.

Let $\lambda_1, \lambda_2, \dots, \lambda_M$ be unique eigenvalues of \mathbf{R} and take $\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_M$ to be the M orthonormal eigenvectors

$$\mathbf{q}_i^H \mathbf{q}_j = \begin{cases} 1 & i = j \\ 0 & i \neq j \end{cases}$$

Define $\mathbf{Q} = [\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_M]$ and $\mathbf{\Omega} = \text{diag}(\lambda_1, \lambda_2, \dots, \lambda_M)$. Then consider

$$\mathbf{Q}^H \mathbf{R} \mathbf{Q} = \begin{bmatrix} \mathbf{q}_1^H \\ \mathbf{q}_2^H \\ \vdots \\ \mathbf{q}_M^H \end{bmatrix} \mathbf{R} [\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_M]$$

$$\begin{aligned}
\mathbf{Q}^H \mathbf{R} \mathbf{Q} &= \begin{bmatrix} \mathbf{q}_1^H \\ \mathbf{q}_2^H \\ \vdots \\ \mathbf{q}_M^H \end{bmatrix} \mathbf{R} [\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_M] \\
&= \begin{bmatrix} \mathbf{q}_1^H \\ \mathbf{q}_2^H \\ \vdots \\ \mathbf{q}_M^H \end{bmatrix} [\lambda_1 \mathbf{q}_1, \lambda_2 \mathbf{q}_2, \dots, \lambda_N \mathbf{q}_M] \\
&= \begin{bmatrix} \lambda_1 & 0 & \dots & 0 \\ 0 & \lambda_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \lambda_M \end{bmatrix} \\
\Rightarrow \mathbf{Q}^H \mathbf{R} \mathbf{Q} &= \mathbf{\Omega} \quad (\text{eigenvector diagonalization of } \mathbf{R})
\end{aligned}$$

Property (Q is unitary)

Q is **unitary**, i.e., $\mathbf{Q}^{-1} = \mathbf{Q}^H$

Proof: Since the \mathbf{q}_i eigenvectors are **orthonormal**

$$\begin{aligned}\mathbf{Q}^H \mathbf{Q} &= \begin{bmatrix} \mathbf{q}_1^H \\ \mathbf{q}_2^H \\ \vdots \\ \mathbf{q}_M^H \end{bmatrix} [\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_M] = \mathbf{I} \\ \Rightarrow \mathbf{Q}^{-1} &= \mathbf{Q}^H\end{aligned}$$

Property (Eigen decomposition of R)

The correlation matrix can be expressed as

$$\mathbf{R} = \sum_{i=1}^M \lambda_i \mathbf{q}_i \mathbf{q}_i^H$$

Proof: The correlation diagonalization result states

$$\mathbf{Q}^H \mathbf{R} \mathbf{Q} = \mathbf{\Omega}$$

Isolating \mathbf{R} and expanding,

$$\begin{aligned} \mathbf{R} &= \mathbf{Q} \mathbf{\Omega} \mathbf{Q}^H = [\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_M] \mathbf{\Omega} \begin{bmatrix} \mathbf{q}_1^H \\ \mathbf{q}_2^H \\ \vdots \\ \mathbf{q}_M^H \end{bmatrix} \\ &= [\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_M] \begin{bmatrix} \lambda_1 \mathbf{q}_1^H \\ \lambda_2 \mathbf{q}_2^H \\ \vdots \\ \lambda_M \mathbf{q}_M^H \end{bmatrix} = \sum_{i=1}^M \lambda_i \mathbf{q}_i \mathbf{q}_i^H \end{aligned}$$

Aside (trace & determinant for matrix products)

Note $\text{trace}(\mathbf{A}) \triangleq \sum_i A_{i,i}$. Also,

$$\text{trace}(\mathbf{AB}) = \text{trace}(\mathbf{BA}) \quad \text{similarly} \quad \det(\mathbf{AB}) = \det(\mathbf{A})\det(\mathbf{B})$$

Property (Determinant–Eigenvalue Relation)

The determinant of the correlation matrix is related to the eigenvalues as follows:

$$\det(\mathbf{R}) = \prod_{i=1}^M \lambda_i$$

Proof: Using $\mathbf{R} = \mathbf{Q}\mathbf{\Omega}\mathbf{Q}^H$ and the above,

$$\begin{aligned} \det(\mathbf{R}) &= \det(\mathbf{Q}\mathbf{\Omega}\mathbf{Q}^H) \\ &= \det(\mathbf{Q})\det(\mathbf{Q}^H)\det(\mathbf{\Omega}) = \det(\mathbf{\Omega}) = \prod_{i=1}^M \lambda_i \end{aligned}$$

Property (Trace–Eigenvalue Relation)

The trace of the correlation matrix is related to the eigenvalues as follows:

$$\text{trace}(\mathbf{R}) = \sum_{i=1}^M \lambda_i$$

Proof: Note

$$\begin{aligned} \text{trace}(\mathbf{R}) &= \text{trace}(\mathbf{Q}\mathbf{\Omega}\mathbf{Q}^H) \\ &= \text{trace}(\mathbf{Q}^H\mathbf{Q}\mathbf{\Omega}) \\ &= \text{trace}(\mathbf{\Omega}) \\ &= \sum_{i=1}^M \lambda_i \end{aligned}$$

QED

Matrix-Vector Multiplication

Example in 2D:

$$\mathbf{A} = \begin{bmatrix} 2 & 1 \\ -1 & 1 \end{bmatrix} \quad \mathbf{x} = \begin{bmatrix} 1 \\ 3 \end{bmatrix}$$

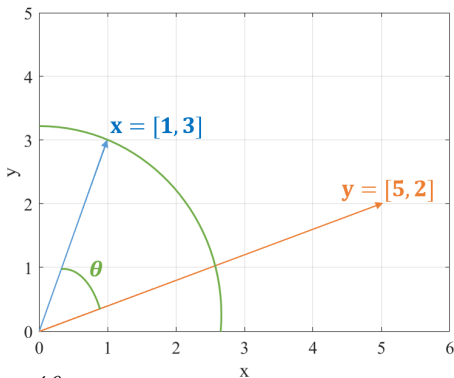
and,

$$\mathbf{y} = \mathbf{Ax} = \begin{bmatrix} 2 & 1 \\ -1 & 1 \end{bmatrix} \begin{bmatrix} 1 \\ 3 \end{bmatrix} = \begin{bmatrix} 5 \\ 2 \end{bmatrix}$$

What is the geometrical meaning of the matrix-vector multiplication?

Matrix-Vector Multiplication

$$\mathbf{y} = \mathbf{Ax} = \begin{bmatrix} 2 & 1 \\ -1 & 1 \end{bmatrix} \begin{bmatrix} 1 \\ 3 \end{bmatrix} = \begin{bmatrix} 5 \\ 2 \end{bmatrix}$$



- ▶ Rotates the vector $\angle \theta$
- ▶ Stretches the vector



Matrix-Vector Multiplication

To rotate \mathbf{x} by an angle θ , we pre-multiply by

$$\mathbf{A} = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix}$$


Stretch \mathbf{x} by factor α , pre-multiply by

$$\mathbf{A} = \begin{bmatrix} \alpha & 0 \\ 0 & \alpha \end{bmatrix}$$



Matrix-Vector Multiplication


Consider the vectors \mathbf{v}_1 and \mathbf{v}_2 depicting a circle. What happens to the circle under matrix multiplication?



2-D Circle

$$\mathbf{A}[\mathbf{v}_1 \ \mathbf{v}_2]$$


$$\mathbf{v}_i \in \mathbb{C}^2$$



3-D Sphere

$$\mathbf{A}[\mathbf{v}_1 \ \mathbf{v}_2 \ \mathbf{v}_3]$$

$$\mathbf{v}_i \in \mathbb{C}^3$$




n-D Hypersphere

$$\mathbf{A}[\mathbf{v}_1 \ \dots \ \mathbf{v}_n]$$

$$\mathbf{v}_i \in \mathbb{C}^n$$

Matrix-Vector Multiplication


What happens to the 2D circle under matrix multiplication?



Note: Orthogonality holds since they are all rotated by the same angle.

Matrix-Vector Multiplication

What happens to the n-D hyper-sphere under matrix multiplication?



n-dim Hyper-Sphere Mapping to n-dim Hyper-Ellipsoid

The mapping can be written as

$$\begin{aligned} \mathbf{A}\mathbf{v}_1 &= \sigma_1 \hat{\mathbf{u}}_1 \\ &\vdots \\ \mathbf{A}\mathbf{v}_n &= \sigma_n \hat{\mathbf{u}}_n \end{aligned}$$

Expressed in matrix form as

$$\underbrace{\begin{bmatrix} \mathbf{A} \end{bmatrix}}_{\mathbf{A} \in \mathbb{C}^{m \times n}} \underbrace{\begin{bmatrix} \mathbf{v}_1 & \mathbf{v}_2 & \dots & \mathbf{v}_n \end{bmatrix}}_{\mathbf{V} \in \mathbb{C}^{n \times n}} = \underbrace{\begin{bmatrix} \hat{\mathbf{u}}_1 & \hat{\mathbf{u}}_2 & \dots & \hat{\mathbf{u}}_n \end{bmatrix}}_{\hat{\mathbf{U}} \in \mathbb{C}^{m \times n}} \underbrace{\begin{bmatrix} \sigma_1 & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & \sigma_n \end{bmatrix}}_{\hat{\Sigma} \in \mathbb{C}^{n \times n}}$$

$$\mathbf{AV} = \hat{\mathbf{U}}\hat{\Sigma}$$

n-dim Hyper-Sphere Mapping to n-dim Hyper-Ellipsoid

Let $\mathbf{v}_1, \dots, \mathbf{v}_n$ be unitary orthonormal vectors, then $\mathbf{V} = [\mathbf{v}_1 \ \mathbf{v}_2 \ \dots \ \mathbf{v}_n]$ is a unitary transformation matrix, that is

$$\mathbf{V}^{-1} = \mathbf{V}^H.$$

Let $\hat{\mathbf{u}}_1, \dots, \hat{\mathbf{u}}_n$ be unitary orthonormal vectors, then $\hat{\mathbf{U}} = [\hat{\mathbf{u}}_1 \ \hat{\mathbf{u}}_2 \ \dots \ \hat{\mathbf{u}}_n]$ is a unitary transformation matrix, that is

$$\mathbf{U}^{-1} = \hat{\mathbf{U}}^H.$$

Reduced Singular Value Decomposition

The mapping is thus given by,

$$\mathbf{AV} = \hat{\mathbf{U}}\hat{\mathbf{\Sigma}}$$

Multiply both sides by \mathbf{V}^{-1} we obtain:

$$\mathbf{AVV}^{-1} = \hat{\mathbf{U}}\hat{\mathbf{\Sigma}}\mathbf{V}^{-1}$$

$$\mathbf{AVV}^H = \hat{\mathbf{U}}\hat{\mathbf{\Sigma}}\mathbf{V}^H$$

$$\mathbf{AI} = \hat{\mathbf{U}}\hat{\mathbf{\Sigma}}\mathbf{V}^H$$

$$\mathbf{A} = \hat{\mathbf{U}}\hat{\mathbf{\Sigma}}\mathbf{V}^H$$

where $\mathbf{\Sigma} = \text{diag}([\sigma_1, \sigma_2, \dots, \sigma_n])$, such that $\sigma_1 \geq \sigma_2 \geq \dots \sigma_p \geq 0$.

Singular Value Decomposition

▶ Reduced SVD

$$\begin{array}{c} \mathbf{A} \\ \mathbb{C}^{m \times n} \end{array} = \begin{array}{c} \hat{\mathbf{U}} \\ \mathbb{C}^{m \times n} \end{array} \begin{array}{c} \mathbf{\Sigma} \\ \mathbb{C}^{n \times n} \end{array} \begin{array}{c} \mathbf{V}^H \\ \mathbb{C}^{n \times n} \end{array}$$

▶ SVD

$$\begin{array}{c} \mathbf{A} \\ \mathbb{C}^{m \times n} \end{array} = \begin{array}{c} \mathbf{U} \\ \mathbb{C}^{m \times m} \end{array} \begin{array}{c} \mathbf{\Sigma} \\ \text{Zeros} \\ \mathbb{C}^{m \times n} \end{array} \begin{array}{c} \mathbf{V}^H \\ \mathbb{C}^{n \times n} \end{array}$$

Theorem 1

Every matrix $\mathbf{A} \in \mathbb{C}^{m \times n}$ has a singular value decomposition (SVD).

- ▶ Singular values σ_j are uniquely determined.
- ▶ If \mathbf{A} is square σ_j are distinct.
- ▶ \mathbf{u}_j and \mathbf{v}_j are also unique up to a complex sign. (unique if the complex sign is ignored)

SVD calculation

Start with $\mathbf{A}^T \mathbf{A}$:

$$\begin{aligned}\mathbf{A}^H \mathbf{A} &= (\mathbf{U} \Sigma \mathbf{V}^H)^H (\mathbf{U} \Sigma \mathbf{V}^H) \\ &= \mathbf{V} \Sigma \mathbf{U}^H \mathbf{U} \Sigma \mathbf{V}^H \\ \mathbf{A}^H \mathbf{A} \mathbf{V} &= \mathbf{V} \Sigma^2 \mathbf{V}^H \mathbf{V} \\ \mathbf{A}^H \mathbf{A} \mathbf{V} &= \mathbf{V} \Sigma^2\end{aligned}$$

Reduces to an eigenvalue decomposition problem of the form:

$$\underbrace{\mathbf{A}^T \mathbf{A}}_{\mathbf{B}} \mathbf{V} = \mathbf{V} \underbrace{\Sigma^2}_{\Lambda},$$

where Λ is a diagonal matrix with the eigenvalues of \mathbf{B} and \mathbf{V} corresponds to the eigenvectors of \mathbf{B} .

SVD calculation

How do we calculate \mathbf{U} :

$$\begin{aligned}
 \mathbf{A}\mathbf{A}^H &= (\mathbf{U}\mathbf{\Sigma}\mathbf{V}^H)(\mathbf{U}\mathbf{\Sigma}\mathbf{V}^H)^H \\
 &= \mathbf{U}\mathbf{\Sigma}\mathbf{V}^H\mathbf{V}\mathbf{\Sigma}\mathbf{U}^H \\
 \mathbf{A}\mathbf{A}^H\mathbf{U} &= \mathbf{U}\mathbf{\Sigma}^2\mathbf{U}^H\mathbf{U} \\
 \underbrace{\mathbf{A}\mathbf{A}^H}_{\mathbf{B}}\mathbf{U} &= \mathbf{U}\underbrace{\mathbf{\Sigma}^2}_{\mathbf{\Lambda}}
 \end{aligned}$$

Eigenvalue problem where $\mathbf{\Lambda}$ is a diagonal matrix with the eigenvalues of \mathbf{B} and \mathbf{U} corresponds to the eigenvectors of \mathbf{B} .

Netflix Movie Challenge

- ▶ Dataset: $n = 17,770$ movies (columns) and $m = 480,189$ customers (rows).
- ▶ Customers rated movies on a scale from 1 to 5. Matrix is very sparse with “only” 100 million of the ratings present in the training set.
- ▶ Goal: Predict the ratings for unrated movies.

Netflix Prize **COMPLETED**

Home Rules Leaderboard Update

Leaderboard

Showing Test Score. Click here to view past scores

Display top 20 leaders.

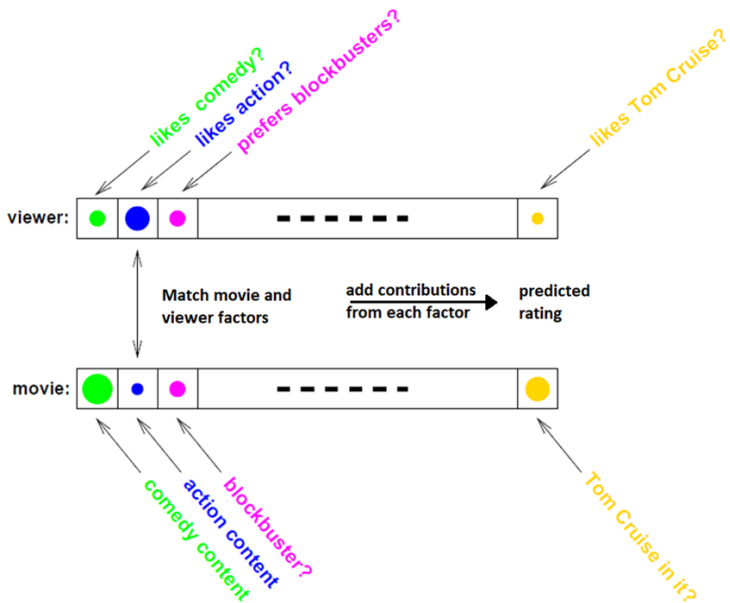
Rank	Team Name	Best Test Score	Improvement	Best Submit Time
Grand Prize - RMSE = 0.8817 - Winning Team: Bellkor's Pragmatic Chaos				
1	Bellkor's Pragmatic Chaos	0.8817	10.06	2009-07-26 18:19:28
2	The Ensemble	0.8927	10.06	2009-07-26 18:36:22
3	Claris (Pete Tass)	0.8982	9.96	2009-07-15 21:29:40
4	Claris Solutions and Velocity United	0.9088	9.84	2009-07-15 01:12:31
5	Velocity United	0.9191	9.81	2009-07-15 00:52:20
6	PragmaticTheory	0.9294	9.77	2009-06-24 12:06:58
7	Bellkor in BigChaos	0.9351	9.72	2009-05-13 08:14:09
8	Claris	0.9412	9.59	2009-07-04 17:18:43
9	Teedee	0.9522	9.48	2009-07-12 13:11:81
10	BigChaos	0.9623	9.47	2009-04-07 12:33:59
11	Claris Solutions	0.9653	9.47	2009-07-04 00:19:47
12	Bellkor	0.9824	9.46	2009-07-26 17:19:11
Progress Prize 2008 - RMSE = 0.8627 - Winning Team: Bellkor in BigChaos				
13	teedee	0.8642	9.27	2009-07-15 14:53:22
14	Claris	0.8643	9.26	2009-04-22 18:31:32
15	Claris	0.8651	9.18	2009-06-21 19:28:53
16	Pragmatic Theory	0.8653	9.18	2009-07-15 18:53:04
17	Just a Guy (J.A. Johnson)	0.8682	9.06	2009-05-24 10:02:04
18	J.Dennis.Su	0.8686	9.02	2009-03-07 17:16:17
19	Claris Connecticut	0.8688	9.02	2009-07-26 18:05:04
20	teedee	0.8688	9.02	2009-03-21 14:20:30
Progress Prize 2007 - RMSE = 0.8723 - Winning Team: Bellkor				
Cinematch 2006 - RMSE = 0.9525				

There are currently 81951 contestants on 41305 teams from 186 different countries. We have received 44574 valid submissions from 5169 different teams, 3 submissions in the last 24 hours. Questions about interpreting the leaderboard? Please read this.

- ▶ (2006) “Cinematch” algorithm used by Netflix RMSE=0.9525 over a large test set.
- ▶ Competition started in 2006, winner should improve this RMSE by at least 10%.
- ▶ 2009 “Bellkor’s Pragmatic Chaos,” uses a combination of many statistical techniques to win.

Movie Rating - A Solution

- Describe a movie as an array of factors, e.g. comedy, action...
- Describe each viewer using same factors, e.g. likes comedy, likes action, etc
- Rating based on match/mismatch
- More factors \rightarrow better prediction



Singular Value Decomposition Solution

Viewers rated movies on a scale from 1 to 5. 0 for movies that were not rated by the user.

- ▶ Each column j is a different movie
- ▶ Each row i is a different viewer
- ▶ Each element $a_{i,j}$ represents the rating of movie j by viewer i

	Movie 1	Movie 2	Movie 3	Movie 4	Movie 5
Viewer 1	0	1	0	0	5
Viewer 2	4	2	0	0	0
Viewer 3	0	0	3	3	0
Viewer 4	4	2	0	0	0
Viewer 5	0	0	0	0	5
Viewer 6	0	0	3	3	0
Viewer 7	1	0	0	0	4
Viewer 8	2	1	0	0	4
Viewer 9	1	0	0	0	4

$$\mathbf{A} = \begin{bmatrix} a_{1,1} & \cdots & \cdots & a_{1,n} \\ \vdots & \ddots & \ddots & \vdots \\ a_{m,1} & \cdots & \cdots & a_{m,n} \end{bmatrix}$$

Goal: Use SVD to predict unobserved data or the rating of a movie that hasn't come out yet.

Singular Value Decomposition Solution

We want to classify Movies and Viewers:

$$Movies = \begin{cases} \text{Category 1} \\ \text{Category 2} \\ \text{Category 3} \\ \vdots \end{cases}$$

Intuitively, if $Movie_1 \approx Movie_2$, these movies are similar (same category).

	<i>Movie 1</i>	<i>Movie 2</i>	<i>Movie 3</i>	<i>Movie 4</i>	<i>Movie 5</i>
Viewer 1	0	1	0	0	5
Viewer 2	4	2	0	0	0
Viewer 3	0	0	3	3	0
Viewer 4	4	2	0	0	0
Viewer 5	0	0	0	0	5
Viewer 6	0	0	3	3	0
Viewer 7	1	0	0	0	4
Viewer 8	2	1	0	0	4
Viewer 9	1	0	0	0	4

Categories are determined by matrix \mathbf{A} and SVD algorithm.

Singular Value Decomposition Solution

Now, consider that each movie belongs to more than one category e.g. half comedy and half action. This can be written as:

$$\begin{aligned} \text{Movie}_j &= v_1 \text{Cat1} + v_2 \text{Cat2} + \dots + v_n \text{Catn} \\ \text{s.t. } \|\mathbf{v}\|_2 &= 1 \end{aligned}$$

where the set of categories $\{\text{Cat}_j \in \mathbb{R}^{n \times 1}\}$ forms an orthonormal basis.

$$\text{Cat} = \begin{bmatrix} 1 & 0 & \dots & 0 \\ 0 & 1 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 1 \end{bmatrix}_{n \times n}$$

	Movie 1	Movie 2	Movie 3	Movie 4	Movie 5
Viewer 1	0	1	0	0	5
Viewer 2	4	2	0	0	0
Viewer 3	0	0	3	3	0
Viewer 4	4	2	0	0	0
Viewer 5	0	0	0	0	5
Viewer 6	0	0	3	3	0
Viewer 7	1	0	0	0	4
Viewer 8	2	1	0	0	4
Viewer 9	1	0	0	0	4

Singular Value Decomposition Solution

In the case of *Viewers*, we use the same *Movies*' categories:

$$\text{Movies} = \begin{Bmatrix} \text{Category 1} \\ \text{Category 2} \\ \text{Category 3} \\ \vdots \end{Bmatrix} = \text{Viewers}.$$

E.g. a viewer that loves comedy is represented with the same unit vector of the comedy category movies ($\text{Cat}_i \in \mathbb{R}^{1 \times n}$).

Each *Viewer* is represented as:

$$\text{Viewer}_i = u_1 \text{Cat}_1 + u_2 \text{Cat}_2 + \dots + u_n \text{Cat}_n$$

$$\text{s.t. } \|\mathbf{u}\|_2 = 1$$

	Movie 1	Movie 2	Movie 3	Movie 4	Movie 5
Viewer 1	0	1	0	0	5
Viewer 2	4	2	0	0	0
Viewer 3	0	0	3	3	0
Viewer 4	4	2	0	0	0
Viewer 5	0	0	0	0	5
Viewer 6	0	0	3	3	0
Viewer 7	1	0	0	0	4
Viewer 8	2	1	0	0	4
Viewer 9	1	0	0	0	4

Singular Value Decomposition Solution

If $m > n$ i.e # of Viewers $>$ # of Movies, each *Viewer* is represented as:

$$\begin{aligned} \text{Viewer}_i &= u_1 \text{Cat1} + u_2 \text{Cat2} + \cdots + u_n \text{Catn} + \cdots + u_m \text{Catm} \\ &\text{s.t. } \|\mathbf{u}\|_2 = 1 \end{aligned}$$

where $\text{Cat}_i \in \mathbb{R}^{1 \times m}$. Thus, useless categories vectors with zero rating value are added.

Singular Value Decomposition Solution


From Theorem 1:

There exist a unique decomposition into categories. Every matrix $\mathbf{A} \in \mathbb{C}^{m \times n}$ can be factorized as $\mathbf{A} = \hat{\mathbf{U}}\mathbf{\Sigma}\mathbf{V}^H$ where:

$$\begin{array}{c} \mathbf{A} \\ \mathbb{C}^{m \times n} \end{array} = \begin{array}{c} \hat{\mathbf{U}} \\ \mathbb{C}^{m \times n} \end{array} \begin{array}{c} \mathbf{\Sigma} \\ \mathbb{C}^{n \times n} \end{array} \begin{array}{c} \mathbf{V}^H \\ \mathbb{C}^{n \times n} \end{array}$$

Singular Value Decomposition Solution

We have more viewers than movies:



New categories are created. The new vectors are still unit vectors orthonormal to all the basis vectors but the ratings of these useless categories are zero.

Note: consider reduced SVD i.e. consider only useful categories.

Singular Value Decomposition Solution

$$\begin{array}{c}
 \mathbf{A} \\
 \mathbb{C}^{m \times n}
 \end{array}
 =
 \begin{array}{c}
 \hat{\mathbf{U}} \\
 \mathbb{C}^{m \times n}
 \end{array}
 \begin{array}{c}
 \mathbf{\Sigma} \\
 \mathbb{C}^{n \times n}
 \end{array}
 \begin{array}{c}
 \mathbf{V}^H \\
 \mathbb{C}^{n \times n}
 \end{array}$$

- ▶ Each row vector (\mathbf{u}_i) in $\hat{\mathbf{U}}$ represents the taste of a *Viewer_i* on the corresponding categories.

$$\hat{\mathbf{U}} = \begin{bmatrix} u_{1,1} & \cdots & \cdots & u_{1,n} \\ \vdots & \ddots & \ddots & \vdots \\ u_{m,1} & \cdots & \cdots & u_{m,n} \end{bmatrix} = \begin{bmatrix} \mathbf{u}_1 \\ \mathbf{u}_2 \\ \vdots \\ \mathbf{u}_m \end{bmatrix}$$

Singular Value Decomposition Solution

$$\begin{array}{c}
 \mathbf{A} \\
 \mathbb{C}^{m \times n}
 \end{array}
 =
 \begin{array}{c}
 \hat{\mathbf{U}} \\
 \mathbb{C}^{m \times n}
 \end{array}
 \begin{array}{c}
 \mathbf{\Sigma} \\
 \mathbb{C}^{n \times n}
 \end{array}
 \begin{array}{c}
 \mathbf{V}^H \\
 \mathbb{C}^{n \times n}
 \end{array}$$

- ▶ Each column (\mathbf{v}_j) in \mathbf{V}^H represents the content of a *Movie_j* on the corresponding categories.

$$\mathbf{V}^H = \begin{bmatrix} v_{1,1} & \cdots & \cdots & v_{1,n} \\ \vdots & \ddots & \ddots & \vdots \\ v_{n,1} & \cdots & \cdots & v_{n,n} \end{bmatrix} = \left[\mathbf{v}_1 \quad \mathbf{v}_2 \quad \cdots \quad \mathbf{v}_n \right]$$

Singular Value Decomposition Solution

$$\begin{array}{c}
 \mathbf{A} \\
 \mathbb{C}^{m \times n}
 \end{array}
 =
 \begin{array}{c}
 \hat{\mathbf{U}} \\
 \mathbb{C}^{m \times n}
 \end{array}
 \begin{array}{c}
 \mathbf{\Sigma} \\
 \mathbb{C}^{n \times n}
 \end{array}
 \begin{array}{c}
 \mathbf{V}^H \\
 \mathbb{C}^{n \times n}
 \end{array}$$

- ▶ Each singular value σ_{ii} in $\mathbf{\Sigma}$ computes how a viewer of category i rates a movie of the same category i .

$$\mathbf{\Sigma} = \begin{bmatrix} \sigma_{1,1} & 0 & \dots & 0 \\ \vdots & \sigma_{2,2} & \ddots & \vdots \\ 0 & 0 & \dots & \sigma_{n,n} \end{bmatrix}$$

Singular Value Decomposition Solution

The representation of each movie can be obtained by

$$\begin{aligned}
 \text{Movie}_j &= v_{1,j} \text{Cat1} + v_{2,j} \text{Cat2} + \cdots + v_{n,j} \text{Catn} && \text{s.t. } \|\mathbf{v}_j\|_2 = 1 \\
 &= v_{1,j} \begin{bmatrix} \sqrt{\sigma_{1,1}} \\ 0 \\ \vdots \\ 0 \end{bmatrix} + v_{2,j} \begin{bmatrix} 0 \\ \sqrt{\sigma_{2,2}} \\ \vdots \\ 0 \end{bmatrix} + \cdots + v_{n,j} \begin{bmatrix} 0 \\ 0 \\ \vdots \\ \sqrt{\sigma_{n,n}} \end{bmatrix} \\
 &= \sqrt{\Sigma} \mathbf{v}_j \in \mathbb{C}^{n \times 1}
 \end{aligned}$$

Singular Value Decomposition Solution


The representation of each viewer can be obtained by

$$\begin{aligned}
 \text{Viewer}_i &= u_{i,1}\text{Cat1} + u_{i,2}\text{Cat2} + \cdots + u_{i,n}\text{Catn} + \cdots + u_{i,m}\text{Catm} \\
 &\quad \text{s.t. } \|\mathbf{u}_i\|_2 = 1, \quad \text{Cat}_j = 0 \text{ for } j > n \rightarrow \text{useless categories} \\
 &= u_{i,1} \begin{bmatrix} \sqrt{\sigma_{1,1}} \\ 0 \\ \vdots \\ 0 \end{bmatrix}^H + u_{i,2} \begin{bmatrix} 0 \\ \sqrt{\sigma_{2,2}} \\ \vdots \\ 0 \end{bmatrix}^H + \cdots + u_{i,n} \begin{bmatrix} 0 \\ 0 \\ \vdots \\ \sqrt{\sigma_{n,n}} \end{bmatrix}^H \\
 &= \mathbf{u}_i \sqrt{\Sigma}^H \in \mathbb{C}^{1 \times n}
 \end{aligned}$$

Singular Value Decomposition Solution

Given the decomposition of a movie and a viewer, the rating is estimated by:

$$\begin{aligned}
 \text{Viewer}_i \text{Movie}_j &= u_{i,1}v_{1,j}\sigma_{1,1} + u_{i,2}v_{2,j}\sigma_{2,2} + \cdots + u_{i,n}v_{n,j}\sigma_{n,n} \\
 &= (\mathbf{u}_i \sqrt{\Sigma}^H)(\sqrt{\Sigma} \mathbf{v}_j) \\
 &= \mathbf{u}_i \Sigma \mathbf{v}_j
 \end{aligned}$$




Singular Value Decomposition - Example

Considering the rating from 60 viewers to 16 movies of 4 different genres (action, romance, sci-fi, comedy), we generate $\mathbf{A} \in \mathbb{R}^{60 \times 16}$


- ▶ Viewers rated movies on a scale from 1 to 5, 0 for movies that were not rated by the user.
- ▶ Observe the same 4 categories of viewers.

$\mathbf{A} =$


60 Viewers




Singular Value Decomposition - Example



Singular Value Decomposition - Example



Singular Value Decomposition - Example



Singular Value Decomposition - Example

To estimate not rated movies (zero entries in \mathbf{A}), we use additional information: \mathbf{A} is known to be low-rank or approximately low-rank.

Thus, we are going to use the k -rank approximation of the matrix \mathbf{A} that is:


$$\hat{\mathbf{A}} = \mathbf{U}\hat{\Sigma}_k\mathbf{V}^H$$

where $\hat{\Sigma}_k$ has all but the first k singular values σ_{ii} set to zero.

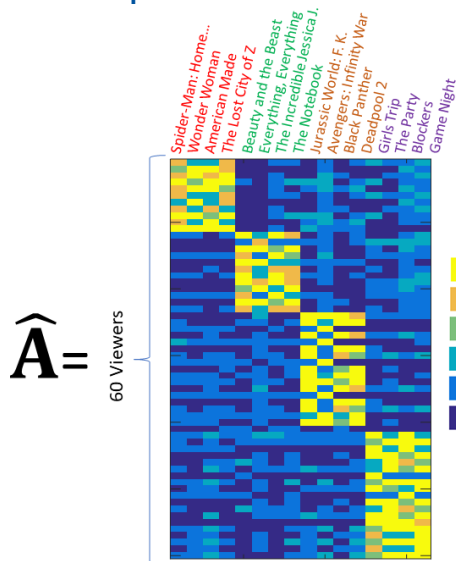

The ratings different from zero in \mathbf{A} are set to its original value.

Note: The ratings matrix \mathbf{A} is expected to be low-rank since user preferences can be described by a few categories (k), such as the movie genres.

Singular Value Decomposition - Example




Singular Value Decomposition - Example



Principal Component Analysis (PCA)

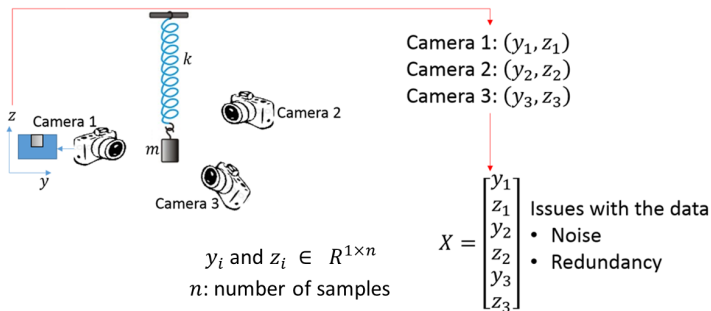
- ▶ Simple, method for extracting relevant information from confusing data sets.
- ▶ How to reduce a complex data set to a lower dimension?
- ▶ Consider a mass attached to a spring which oscillates as shown below.



What if we did not know that $F = ma$?


PCA - Motivation: Toy example

- ▶ Since we live in a 3D world \rightarrow use three cameras to capture data from the system.
- ▶ No information about the real x, y , and z axes \rightarrow camera positions are chosen arbitrarily.
- ▶ How do we get from this data set to a simple equation of z ?



PCA - Motivation: Toy example

- ▶ Three cameras give redundant information.
- ▶ Only one camera at a specific angle necessary to describe the system behavior.
- ▶ PCA is used to avoid redundancy.



Change of Basis


- ▶ PCA: Is there another basis, which is a linear combination of the original basis, that best represents the data set?
- ▶ Let \mathbf{X} be the original data set, where each column is a single measurements set.
- ▶ Let \mathbf{Y} be a linear transformation by \mathbf{P} , i.e. $\mathbf{Y} = \mathbf{P}\mathbf{X}$, where $\mathbf{X} = [\mathbf{x}_1 | \dots | \mathbf{x}_n]$ and $\mathbf{x}_i \in \mathbb{R}^{m \times 1}$ represents a sampled vector.

Implications:

- ▶ Geometrically \mathbf{P} is a rotation and a stretch which transforms \mathbf{X} into \mathbf{Y} .
- ▶ The rows of \mathbf{P} , $\{\mathbf{p}_1, \dots, \mathbf{p}_m\}$ are a set of new basis vectors for expressing the columns of \mathbf{X} .


What is the best way to re-express \mathbf{X} ?, what is a good choice for \mathbf{P} ?

Noise




- ▶ Signal and noise variances are depicted as σ_{signal}^2 and σ_{noise}^2 .
- ▶ The largest direction of variance is not along the natural basis but along the best-fit line.
- ▶ The directions with largest variances contain the dynamics of interest.
- ▶ Intuition: Find the direction indicated by σ_{signal} .

Redundancy



low redundancy



high redundancy

- ▶ Figures depict possible plots between two arbitrary measurement types r_1 and r_2 .
- ▶ Low redundancy \rightarrow uncorrelated recordings
- ▶ High redundancy \rightarrow correlated recordings, e.g. the sensors are too close or the measured variables are equivalent.
- ▶ If recordings are highly correlated it is not necessary to measure both of them.

PCA - Basic concepts

Let $\mathbf{a} = [a_1, a_2, \dots, a_n]$ and $\mathbf{b} = [b_1, b_2, \dots, b_n]$ be two sets of measurements.
Are they related?

If the mean of a and b is zero, then:



- Variance: How large the change is in each vector.

$$\sigma_a^2 = \frac{1}{n} \mathbf{a} \mathbf{a}^T = \frac{1}{n} \sum_i a_i^2$$

$$\sigma_b^2 = \frac{1}{n} \mathbf{b} \mathbf{b}^T = \frac{1}{n} \sum_i b_i^2$$

- Covariance: Statistical relationship between data in \mathbf{a} and \mathbf{b} .

$$\sigma_{ab}^2 = \frac{1}{n} \mathbf{a} \mathbf{b}^T = \frac{1}{n} \sum_i a_i b_i$$




Variance and Covariance

Let \mathbf{X} be defined as $\mathbf{X} = [\mathbf{x}_1^T | \dots | \mathbf{x}_m^T]$, where $\mathbf{x}_i \in \mathbb{R}^{n \times 1}$ is a column vector that corresponds to all measurements of a particular type. Then the covariance matrix is defined as:

$$\mathbf{C}_X = \frac{1}{n} \mathbf{X} \mathbf{X}^T$$

The covariance values reflect the noise and redundancy in the measurements.




Variance and Covariance

Recall \mathbf{C}_X is the covariance matrix of \mathbf{X} defined as

$$\mathbf{C}_X = \frac{1}{n} \mathbf{X} \mathbf{X}^T.$$

- ▶ Covariance matrix in the spring example is $\mathbf{C}_X \in \mathbb{R}^{6 \times 6}$.

$$\mathbf{C}_X = \begin{bmatrix} \sigma_{y_1 y_1}^2 & \sigma_{y_1 z_1}^2 & \sigma_{y_1 y_2}^2 & \sigma_{y_1 z_2}^2 & \sigma_{y_1 y_3}^2 & \sigma_{y_1 z_3}^2 \\ \sigma_{z_1 y_1}^2 & \sigma_{z_1 z_1}^2 & \sigma_{z_1 y_2}^2 & \sigma_{z_1 z_2}^2 & \sigma_{z_1 y_3}^2 & \sigma_{z_1 z_3}^2 \\ \sigma_{y_2 y_1}^2 & \sigma_{y_2 z_1}^2 & \sigma_{y_2 y_2}^2 & \sigma_{y_2 z_2}^2 & \sigma_{y_2 y_3}^2 & \sigma_{y_2 z_3}^2 \\ \sigma_{z_2 y_1}^2 & \sigma_{z_2 z_1}^2 & \sigma_{z_2 y_2}^2 & \sigma_{z_2 z_2}^2 & \sigma_{z_2 y_3}^2 & \sigma_{z_2 z_3}^2 \\ \sigma_{y_3 y_1}^2 & \sigma_{y_3 z_1}^2 & \sigma_{y_3 y_2}^2 & \sigma_{y_3 z_2}^2 & \sigma_{y_3 y_3}^2 & \sigma_{y_3 z_3}^2 \\ \sigma_{z_3 y_1}^2 & \sigma_{z_3 z_1}^2 & \sigma_{z_3 y_2}^2 & \sigma_{z_3 z_2}^2 & \sigma_{z_3 y_3}^2 & \sigma_{z_3 z_3}^2 \end{bmatrix}$$



- ▶ Diagonal: Variance measures; Off-diagonal: covariance between all pairs.
- ▶ \mathbf{C}_X is hermitian and symmetric, i.e. $\mathbf{C}_X = \mathbf{C}_X^T = \mathbf{C}_X^*$.

Covariance Matrix Interpretation

$$\mathbf{C}_x = \begin{bmatrix} \sigma_{y_1 y_1}^2 & \sigma_{y_1 z_1}^2 & \sigma_{y_1 y_2}^2 & \sigma_{y_1 z_2}^2 & \sigma_{y_1 y_3}^2 & \sigma_{y_1 z_3}^2 \\ \sigma_{z_1 y_1}^2 & \sigma_{z_1 z_1}^2 & \sigma_{z_1 y_2}^2 & \sigma_{z_1 z_2}^2 & \sigma_{z_1 y_3}^2 & \sigma_{z_1 z_3}^2 \\ \sigma_{y_2 y_1}^2 & \sigma_{y_2 z_1}^2 & \sigma_{y_2 y_2}^2 & \sigma_{y_2 z_2}^2 & \sigma_{y_2 y_3}^2 & \sigma_{y_2 z_3}^2 \\ \sigma_{z_2 y_1}^2 & \sigma_{z_2 z_1}^2 & \sigma_{z_2 y_2}^2 & \sigma_{z_2 z_2}^2 & \sigma_{z_2 y_3}^2 & \sigma_{z_2 z_3}^2 \\ \sigma_{y_3 y_1}^2 & \sigma_{y_3 z_1}^2 & \sigma_{y_3 y_2}^2 & \sigma_{y_3 z_2}^2 & \sigma_{y_3 y_3}^2 & \sigma_{y_3 z_3}^2 \\ \sigma_{z_3 y_1}^2 & \sigma_{z_3 z_1}^2 & \sigma_{z_3 y_2}^2 & \sigma_{z_3 z_2}^2 & \sigma_{z_3 y_3}^2 & \sigma_{z_3 z_3}^2 \end{bmatrix}$$

Off-diagonal terms

- ▶ If covariance is large then components are statistically dependent.
- ▶ If covariance is small then components are statistically independent.

Diagonal terms:

- ▶ If variance is large it contains a lot of information about the system.
- ▶ If variance is small it does not provide significant information about the system.

PCA

Goal: Change basis such that the covariance matrix of the data is diagonal.

- ▶ If off-diagonal terms ≈ 0 , the redundancies are eliminated.
- ▶ Diagonal terms represent the variance of each component.
- ▶ Components with large variance are the most representative.

$$\mathbf{C}_X = \begin{array}{c} \text{[A yellow square with a dashed white diagonal line from top-left to bottom-right, representing a diagonal matrix]} \end{array}$$

Looks like the SVD!

PCA and Eigenvalue Decomposition

How to solve the problem?

- ▶ Data set: $\mathbf{X} \in \mathbb{R}^{m \times n}$, where m is the number of measurement types and n is the number of samples.
- ▶ PCA : Find an orthonormal matrix \mathbf{P} in $\mathbf{Y} = \mathbf{P}\mathbf{X}$ such that $\mathbf{C}_Y = \frac{1}{n}\mathbf{Y}\mathbf{Y}^T$ is a diagonal matrix.
- ▶ The rows of \mathbf{P} are the principal components of \mathbf{X}

PCA and Eigenvalue Decomposition

We begin rewriting \mathbf{C}_Y in terms of the unknown variable.

$$\begin{aligned}\mathbf{C}_Y &= \frac{1}{n} \mathbf{Y} \mathbf{Y}^T \\ &= \frac{1}{n} (\mathbf{P} \mathbf{X}) (\mathbf{P} \mathbf{X})^T \\ &= \frac{1}{n} \mathbf{P} \mathbf{X} \mathbf{X}^T \mathbf{P}^T \\ &= \mathbf{P} \left(\frac{1}{n} \mathbf{X} \mathbf{X}^T \right) \mathbf{P}^T \\ &= \mathbf{P} \mathbf{C}_X \mathbf{P}^T\end{aligned}$$

PCA and Eigenvalue Decomposition

\mathbf{C}_X can be diagonalized by an orthogonal matrix of its eigenvectors since it is a symmetric matrix. Let $\mathbf{P} = \mathbf{Q}^T$, where \mathbf{Q} is a matrix with the eigenvectors of $\frac{1}{n}\mathbf{X}\mathbf{X}^T$, then:

$$\begin{aligned}\mathbf{C}_Y &= \mathbf{P}\mathbf{C}_X\mathbf{P}^T \\ &= \mathbf{P}(\mathbf{Q}\mathbf{\Omega}\mathbf{Q}^T)\mathbf{P}^T \\ &= \mathbf{P}(\mathbf{P}^T\mathbf{\Omega}\mathbf{P})\mathbf{P}^T \\ &= (\mathbf{P}\mathbf{P}^{-1})\mathbf{\Omega}(\mathbf{P}\mathbf{P}^{-1}) \\ &= \mathbf{\Omega}\end{aligned}$$

The transformation $\mathbf{Y} = \mathbf{P}\mathbf{X}$ diagonalizes the system. Covariance of \mathbf{Y} is a diagonal matrix with the eigenvalues of $\frac{1}{n}\mathbf{X}\mathbf{X}^T$.

PCA and SVD

The SVD of \mathbf{X} is given by $\mathbf{X} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$. Let $\mathbf{P} = \mathbf{U}^T$, then:

$$\mathbf{Y} = \mathbf{U}^T \mathbf{X},$$

The covariance matrix of \mathbf{Y} is given by:

$$\begin{aligned}\mathbf{C}_Y &= \frac{1}{n} \mathbf{Y}\mathbf{Y}^T \\ &= \frac{1}{n} \mathbf{U}^T \mathbf{X}\mathbf{X}^T \mathbf{U} \\ &= \frac{1}{n} \mathbf{U}^T \mathbf{U} \mathbf{\Sigma} \mathbf{V}^T \mathbf{V} \mathbf{\Sigma} \mathbf{U}^T \mathbf{U} \\ &= \frac{1}{n} \mathbf{\Sigma}^2\end{aligned}$$

PCA

- ▶ The transformation $\mathbf{Y} = \mathbf{U}^T \mathbf{X}$ diagonalized the system. Covariance of \mathbf{Y} is a diagonal matrix with the squared singular values of \mathbf{X} multiplied by a factor of $\frac{1}{n}$.
- ▶ It can be concluded that $\mathbf{\Sigma}^2 = \mathbf{\Omega}$, and $\sigma_i^2 = \lambda_i$.
- ▶ The principal components of the data matrix are given by \mathbf{U}^T .

Application: Face Recognition


- ▶ PCA in face recognition \triangleq Eigenfaces
- ▶ Intuition: Figure out the correlation between the rows/ columns of \mathbf{A} from the SVD.

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

- ▶ How important each direction is: $\mathbf{\Sigma}$
- ▶ Principal Directions: \mathbf{U}
- ▶ How each individual component (row/column) projects onto the principal components: \mathbf{V} .

Data in Face Recognition

The data matrix is constructed by vectorizing the face images as shown below, i.e. $\mathbf{A} = [\mathbf{A}_1^T | \mathbf{A}_2^T | \dots | \mathbf{A}_N^T]^T$. The matrix will be $N \times M$, where N is the number of images in the data base and M is the number of pixels of each image.



Example - Celebrity Images

Example, take 5 images of each celebrity: George Clooney, Bruce Willis, Margaret Thatcher and Matt Damon. In the example, $M = 240 * 160 = 38400$ and $N = 20$.



$$\mathbf{A} = \begin{bmatrix} \text{----- Image 1 -----} \\ \text{----- Image 2 -----} \\ \text{----- Image 3 -----} \\ \vdots \\ \text{----- Image 20 -----} \end{bmatrix}_{20 \times 38400}$$

Average Faces

How do the average of the faces of these celebrities look like?

$$\bar{\mathbf{a}}_i = \frac{1}{5} \sum_{j=1}^5 \mathbf{A}_j \quad \text{where} \quad \mathbf{A}_j \in \mathbb{R}^{1 \times M}$$



Average Faces

What defines George Clooney's face?



- ▶ Data matrix $\mathbf{A} \in \mathbb{R}^{N \times M}$ with the images of the example.
- ▶ Compute the correlation matrix of the features of the dataset, i.e. the pixels.
- ▶ The correlation matrix is $\mathbf{C} = \mathbf{A}^T \mathbf{A} \in \mathbb{R}^{M \times M}$, here $M = 38400$.
- ▶ High correlation values \rightarrow everybody has eyes, a nose and a mouth.
- ▶ Correlations between images of the same person will be higher.



Average Face

Eigendecomposition

- ▶ Obtain the eigenvalue decomposition of $\mathbf{C} = \mathbf{A}^T \mathbf{A}$. That is $\mathbf{C} = \mathbf{Q} \mathbf{\Omega} \mathbf{Q}^{-1}$.
- ▶ First eigenvectors $\mathbf{q}_i \in \mathbb{R}^{M \times 1}$ are called the principal components (eigenfaces).
- ▶ One can reconstruct each face as a weighted sum of the eigenvectors.



Representing Faces onto Basis


Each face $\mathbf{A}_i \in \mathbb{R}^{1 \times M}$ in the data set $\mathbf{A} = [\mathbf{A}_1^T | \mathbf{A}_2^T | \dots | \mathbf{A}_N^T]^T$, can be represented as a linear combination of the best K eigenvectors:

$$\mathbf{A}_i^T = \sum_{j=1}^K w_j \mathbf{q}_j, \text{ where } w_j = \mathbf{q}_j^T \mathbf{A}_i^T \quad (2)$$

 $K = 1$  $K = 5$  $K = 10$  $K = 15$  $K = 20$

Projection of the Average faces into the $K=20$ largest Eigenvectors

- ▶ \mathbf{Q} is $M \times M$, from now on let \mathbf{V} be the matrix formed by the first $K=20$ eigenvectors, i.e. $\mathbf{V} \in \mathbb{R}^{M \times K}$.
- ▶ Project the average faces $\bar{\mathbf{a}}_i \in \mathbb{R}^{1 \times M}$ onto the reduced eigenvector space, i.e. $\mathbf{p}_{\bar{\mathbf{a}}_i} = \bar{\mathbf{a}}_i \mathbf{V} \in \mathbb{R}^{1 \times K}$
- ▶ Projections for each face are characteristic of each average face and could be used for classification purposes.



Projection of new images

- ▶ Test set: New image of Margaret Thatcher, Meryl Streep as Margaret Thatcher in “The Iron Lady”, Betty White.
- ▶ Project test images onto eigenvector space, $\mathbf{p} = \mathbf{x}\mathbf{V} \in \mathbb{R}^{1 \times K}$, where $\mathbf{x} \in \mathbb{R}^{1 \times M}$ is the new vectorized image and $\mathbf{V} \in \mathbb{R}^{M \times K}$ is the matrix with the first 20 eigenvectors of the database.
- ▶ Reconstruct images as $\hat{\mathbf{x}} = \mathbf{V}\mathbf{p}^T$.
- ▶ Error defined as the difference between the projection of the new image and the projection of the original Margaret Thatcher images $\mathbf{o}_j\mathbf{V}$ where $j = 1, \dots, 5$, that is

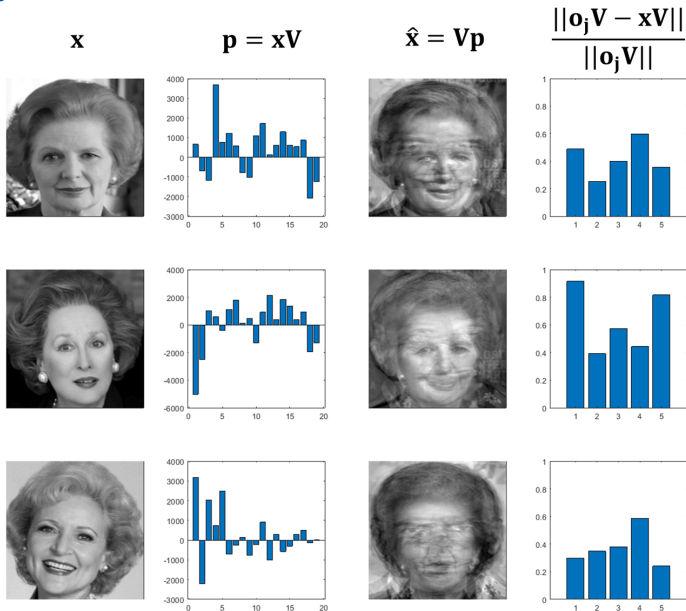
$$E_j = \frac{\|\mathbf{o}_j\mathbf{V} - \mathbf{x}\mathbf{V}\|}{\|\mathbf{o}_j\mathbf{V}\|},$$

where \mathbf{o}_j are the original images of the database, in this case the 5 images of Margareth Thatcher.



Projection of new images

Image depicts, from left to right


- ▶ Test images.
- ▶ Projection of the test images onto the eigenvector space $\mathbf{p} = \mathbf{xV}$.
- ▶ Reconstructed images using the first 20 eigenvectors of the database $\hat{\mathbf{x}} = \mathbf{Vp}^T$.
- ▶ Error of the projection with respect to each original Margareth Thatcher Image \mathbf{o}_j for $j = 1, \dots, 5$.




Projection of new images

 \mathbf{x}  $\mathbf{p} = \mathbf{xV}$  $\hat{\mathbf{x}} = \mathbf{Vp}$ 


Clooney



Willis



Thatcher



Damon

Matrix Approximations and Completion

Given an $m \times n$ matrix $\mathbf{Z} = \{z_{ij}\}$, find a matrix $\hat{\mathbf{Z}}$ that approximates \mathbf{Z} .

- ▶ $\hat{\mathbf{Z}}$ may have simpler structure.
- ▶ Missing entries in \mathbf{Z} , a problem known as *matrix completion*.

Approach based on optimization:

$$\hat{\mathbf{Z}} = \arg \min_{\mathbf{M} \in \mathbb{R}^{m \times n}} \|\mathbf{Z} - \mathbf{M}\|_F^2 \text{ subject to } \Phi(\mathbf{M}) \leq c \quad (3)$$


where $\|\mathbf{A}\|_F^2 = \sum \sum_{i,j} |a_{ij}|^2$ is the Frobenius Norm, and $\Phi(\cdot)$ is a constraint function that encourages $\hat{\mathbf{Z}}$ to be sparse in some sense.

Constraint $\Phi(\mathbf{Z})$	Resulting method
(a) $\ \hat{\mathbf{Z}}\ _{\ell_1} \leq c$	Sparse matrix approximation
(b) $\text{rank}(\hat{\mathbf{Z}}) \leq k$	Singular value decomposition
(c) $\ \hat{\mathbf{Z}}\ _* \leq c$	Convex matrix approximation


- ▶ (a) ℓ_1 -norm of all entries of $\hat{\mathbf{Z}}$. Leads to a soft-thresholding $\hat{z}_{ij} = \text{sign}(z_{ij})(|z_{ij}| - \gamma)_+$, where $\gamma > 0$ is such that $\sum_{i=1}^m \sum_{j=1}^n |\hat{z}_{ij}| = c$.
- ▶ (b) Bounds the rank of $\hat{\mathbf{Z}}$, or the number of nonzero singular values in $\hat{\mathbf{Z}}$. Approximation is non-convex, but solution found by computing the SVD and truncating it to its top k components.
- ▶ (c) Relaxes the rank constraint to a *nuclear norm* ($\|\mathbf{A}\|_* = \sum_{i=1}^{\min\{m,n\}} \sigma_i$). Solved by computing the SVD and soft-thresholding its singular values.

Motivation: Image Reconstruction from Incomplete Data

Reconstructed image



Incomplete image 50% of the pixels



Matrices with missing elements can be solved exactly using method (c), whereas methods based on (b) are more difficult to solve in general.

The Singular Value Decomposition

Given an $m \times n$ matrix \mathbf{Z} with $m \geq n$, its *singular value decomposition* takes the form

$$\mathbf{Z} = \mathbf{U}\mathbf{D}\mathbf{V}^T \quad (4)$$

- ▶ \mathbf{U} is an $m \times n$ orthogonal matrix ($\mathbf{U}^T \mathbf{U} = \mathbf{I}_n$) whose columns $\mathbf{u}_j \in \mathbb{R}^m$ are the *left singular vectors*.
- ▶ \mathbf{V} is an $n \times n$ orthogonal matrix ($\mathbf{V}^T \mathbf{V} = \mathbf{I}_n$) whose columns $\mathbf{v}_j \in \mathbb{R}^n$ are the *right singular vectors*.
- ▶ The $n \times n$ matrix \mathbf{D} is diagonal, with $d_1 \geq d_2 \geq \dots \geq d_n \geq 0$ known as the *singular values*.

The Singular Value Decomposition

- ▶ If columns of \mathbf{Z} are centered (zero mean), then the right singular vectors $\{\mathbf{v}_j\}_{j=1}^n$ define the *principal components* of \mathbf{Z} .
- ▶ The unit vector \mathbf{v}_1 yields the linear combination $\mathbf{s}_1 = \mathbf{Z}\mathbf{v}_1$ with highest sample variance among all possible choices of unit vectors.
- ▶ \mathbf{s}_1 is the *first principal component* of \mathbf{Z} , and \mathbf{v}_1 is the corresponding *direction* or *loading* vector.

The Singular Value Decomposition

Suppose $r \leq \text{rank}(\mathbf{Z}) = 800$, and let \mathbf{D}_r be a diagonal matrix with all but the first r diagonal entries of \mathbf{D} set to zero. The optimization problem

$$\hat{\mathbf{Z}}_r = \min_{\text{rank}(M)=r} \|\mathbf{Z} - \mathbf{M}\|_F \quad (5)$$


has a closed form solution $\hat{\mathbf{Z}}_r = \mathbf{U}\mathbf{D}_r\mathbf{V}^T \triangleq$ the rank- r SVD. $\hat{\mathbf{Z}}_r$ is sparse in the sense that all but r singular values are zero.




800 Singular Values



164 Singular Values



24 Singular Values



12 Singular Values

Matrix Completion

Problem Formulation: Recover an $m \times n$ matrix \mathbf{Z} when we only get to observe $p \ll mn$ of its entries.

- ▶ Impossible without additional information!
- ▶ Assumption: Matrix is known to be low-rank or approximately low-rank.
- ▶ Matrix Completion: Fill the missing entries.
- ▶ Used in: machine learning, computer vision...

Optimization Problem

- ▶ Observe the entries of the $m \times n$ matrix \mathbf{Z} indexed by the subset $\Omega \subset \{1, \dots, m\} \times \{1, \dots, n\}$.
- ▶ Seek the lowest rank approximating matrix $\hat{\mathbf{Z}}$ that interpolates the entries of \mathbf{Z}

$$\begin{aligned} & \text{minimize} && \text{rank}(\mathbf{M}) \\ & \text{subject to} && m_{ij} = z_{ij}, (i, j) \in \Omega, \end{aligned} \tag{6}$$

- ▶ Rank minimization problem is NP-hard.
- ▶ Forcing interpolation leads to overfitting.

Optimization Problem

- ▶ Better to allow \mathbf{M} to make some errors on the observed data:

$$\begin{aligned} & \text{minimize} && \text{rank}(\mathbf{M}) \\ & \text{subject to} && \sum_{(i,j) \in \Omega} (z_{ij} - m_{ij})^2 \leq \delta, \end{aligned} \tag{7}$$

or equivalently

$$\begin{aligned} & \text{minimize} && \sum_{(i,j) \in \Omega} (z_{ij} - m_{ij})^2, \\ & \text{rank}(\mathbf{M}) \leq r && \end{aligned} \tag{8}$$

- ▶ Both problems are non-convex, and exact solutions are generally not available.

Matrix Completion Using the Nuclear Norm

- ▶ Nuclear norm of $\mathbf{M}_{m \times n}$:

$$\|\mathbf{M}\|_* = \sum_{k=1}^n \sigma_k(\mathbf{M}) \quad (9)$$

- ▶ Convex relaxation of the rank minimization problem:

$$\begin{aligned} & \text{minimize} && \|\mathbf{M}\|_* \\ & \text{subject to} && m_{ij} = z_{ij}, (i, j) \in \Omega, \end{aligned} \quad (10)$$

- ▶ Whereas the rank counts the number of nonzero singular values, the nuclear norm sums their amplitude.
- ▶ Analogous to the ℓ_1 norm as a relaxation for the ℓ_0 norm as sparsity measure.

Notation

Given an observed subset Ω of matrix entries, define the projection operator as:

$$[P_{\Omega}(\mathbf{Z})]_{i,j} = \begin{cases} z_{ij} & \text{if } (i,j) \in \Omega \\ 0 & \text{otherwise} \end{cases}$$

P_{Ω} replaces the missing entries in \mathbf{Z} with zeros, and leaves the observed entries alone.

The optimization criterion is then :

$$\sum_{(i,j) \in \Omega} (z_{ij} - m_{ij})^2 = \|P_{\Omega}(\mathbf{Z}) - P_{\Omega}(\mathbf{M})\|_F^2 \quad (11)$$

where $\|\cdot\|_F$ is the Frobenius norm of a matrix defined as the element-wise sum of squares.

Singular Value Thresholding for Matrix Completion,⁺

- ▶ Solves the optimization problem:

$$\begin{aligned} & \text{minimize} && \|\mathbf{M}\|_* \\ & \text{subject to} && P_{\Omega}(\mathbf{M}) = P_{\Omega}(\mathbf{Z}), \end{aligned} \tag{12}$$

- ▶ The SVD of a matrix \mathbf{M} of rank r is:

$$\mathbf{M} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T, \quad \mathbf{\Sigma} = \text{diag}(\{\sigma_i\}_{1 \leq i \leq r}) \tag{13}$$


⁺Cai et al. (2010), SIAM Journal on Optimization, Vol. 20, No. 4

Singular Value Thresholding (SVT)

- For each $\tau \geq 0$, the soft-thresholding operator D_τ is defined as:

$$D_\tau(\mathbf{M}) = \mathbf{U}D_\tau(\mathbf{\Sigma})\mathbf{V}^T, \quad D_\tau(\mathbf{\Sigma}) = \text{diag}(\text{sgn}(\sigma_i) \{|\sigma_i| - \tau\}_+) \quad (14)$$

where $t, t_+ = \max(0, t)$. Operator applies soft-thresholding to the singular values of \mathbf{M} , effectively shrinking them towards zero.



SVT Algorithm - Shrinkage Iterations

Fix $\tau > 0$ and a sequence $\{\delta_k\}$ of positive step sizes. Starting with $\mathbf{Y}^0 = \mathbf{0}$, inductively define for $k = 1, 2, \dots$,

$$\begin{cases} \mathbf{M}^k = D_\tau(\mathbf{Y}^{k-1}) \\ \mathbf{Y}^k = \mathbf{Y}^{k-1} + \delta_k P_\Omega(\mathbf{Z} - \mathbf{M}^k) \end{cases}$$

$$\mathbf{M}^1 = D_\tau(\mathbf{Y}^0) = \mathbf{0}$$

$$\mathbf{M}^2 = D_\tau(\mathbf{Y}^1) = D_\tau(\delta_1 P_\Omega(\mathbf{Z}))$$

$$\begin{aligned} \mathbf{Y}^1 &= \mathbf{0} + \delta_1 P_\Omega(\mathbf{Z} - \mathbf{0}) \\ &= \delta_1 P_\Omega(\mathbf{Z}) \end{aligned}$$

$$\mathbf{Y}^2 = \delta_1 P_\Omega(\mathbf{Z}) + \delta_2 P_\Omega(\mathbf{Z} - D_\tau(\delta_1 P_\Omega(\mathbf{Z})))$$

until a stopping criterion is reached. At each step, we only need to compute an SVD and perform elementary matrix operations.


SVT Algorithm - Shrinkage Iterations



Image Inpainting - Convex Optimization Solver

With 70% of the Information.

Original Image



Noisy Image



Reconstructed






Image Inpainting - Convex Optimization Solver

With 50% of the Information. And multiple columns missing.

Original Image



Noisy Image



Reconstructed






Image Inpainting - Convex Optimization Solver

With 50% of the Information. PSNR=35.9 dB.

Original Image



Noisy Image



Reconstructed






Image Inpainting - SVT Algorithm⁺

With 50% of the Information. PSNR=38.1 dB.


Original Image



Noisy Image



Reconstructed



⁺Cai et al. (2010), SIAM Journal on Optimization, Vol. 20, No. 4

Text Removal - Convex Optimization Solver


Original Image



Noisy Image



Reconstructed



Netflix Movie Challenge - Revisited

- ▶ Dataset: $n = 17,770$ movies (columns) and $m = 480,189$ customers (rows).
- ▶ Customers rated movies on a scale from 1 to 5. Matrix is very sparse with “only” 100 million of the ratings present in the training set.
- ▶ Goal: Predict the ratings for unrated movies.

Netflix Prize **COMPLETED**

Home Rules Leaderboard Update

Leaderboard

Showing Test Score. Click here to view past scores

Display top 20 leaders.

Rank	Team Name	Best Test Score	Improvement	Best Submit Time
Grand Prize - RMSE = 0.8857 - Winning Team: Bellkor's Pragmatic Chaos				
1	Bellkor's Pragmatic Chaos	0.8857	10.06	2009-07-26 18:19:28
2	The Ensemble	0.8927	10.06	2009-07-26 18:36:22
3	Claris (Pete Tass)	0.8982	9.96	2009-07-15 21:29:40
4	Claris Solutions and Velocity Labs	0.9088	9.84	2009-07-15 01:12:31
5	Velocity Institute	0.9191	9.81	2009-07-15 00:52:20
6	PragmaticTheory	0.9294	9.77	2009-06-24 12:06:58
7	Bellkor in BigChaos	0.9351	9.72	2009-05-13 08:14:09
8	Claris	0.9512	9.59	2009-07-04 17:18:43
9	Feedz	0.9522	9.48	2009-07-12 13:11:81
10	BigChaos	0.9623	9.47	2009-04-07 12:33:59
11	Claris Solutions	0.9653	9.47	2009-07-04 00:19:47
12	Bellkor	0.9624	9.46	2009-07-26 17:19:11
Progress Prize 2008 - RMSE = 0.8627 - Winning Team: Bellkor in BigChaos				
13	EMPTOR	0.8642	9.27	2009-07-15 14:53:22
14	Claris	0.8643	9.26	2009-04-22 18:31:32
15	Claris	0.8651	9.16	2009-06-21 19:28:53
16	Pragmatic Theory	0.8653	9.16	2009-07-15 18:53:04
17	Just a Guy (J.A. Gibson)	0.8682	9.06	2009-05-24 10:02:04
18	J.Dennis.Su	0.8686	9.02	2009-03-07 17:16:17
19	Claris (Cristian)	0.8688	9.02	2009-07-26 18:05:14
20	EMPTOR	0.8688	9.02	2009-03-21 14:20:30
Progress Prize 2007 - RMSE = 0.8723 - Winning Team: Bellkor				
Cinematch 2006 - RMSE = 0.9525				

There are currently 81951 contestants on 41305 teams from 186 different countries. We have received 44574 valid submissions from 5169 different teams, 3 submissions in the last 24 hours. Questions about interpreting the leaderboard? Please read this.

- ▶ (2006) “Cinematch” algorithm used by Netflix RMSE=0.9525 over a large test set.
- ▶ Competition started in 2006, winner should improve this RMSE by at least 10%.
- ▶ 2009 “Bellkor’s Pragmatic Chaos,” uses a combination of many statistical techniques to win.

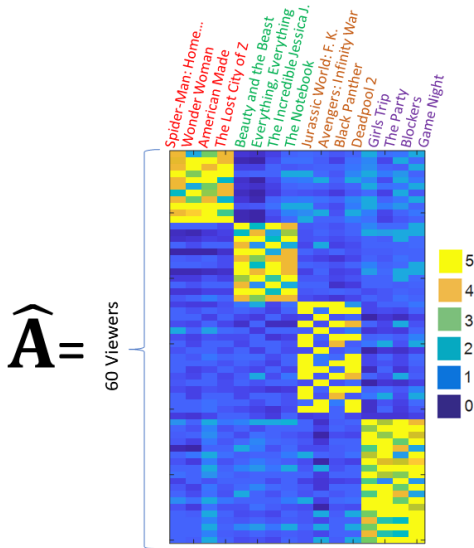
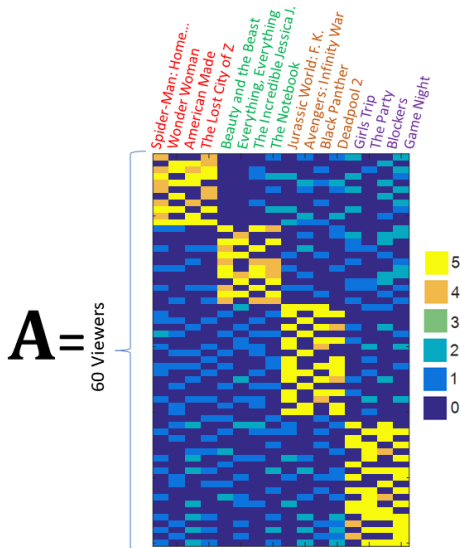
SVT Algorithm Solution

Use SVT Algorithm to estimate not rated movies (zero entries in \mathbf{A}), solving the optimization problem:

$$\begin{aligned} & \text{minimize} && \|\hat{\mathbf{A}}\|_* \\ & \text{subject to} && P_{\Omega}(\hat{\mathbf{A}}) = P_{\Omega}(\mathbf{A}), \end{aligned}$$

Note: The ratings matrix \mathbf{A} is expected to be low-rank since user preferences can be described by a few categories (k), such as the movie genres.

SVT Algorithm Solution



SVT Algorithm Solution

