

Data Analysis and Matrix Decomposition

Tristan Gaeta

University of Puget Sound

May 2, 2021

Definitions of matrix Multiplication

Suppose A is an $m \times n$ matrix with columns $\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_n$ and B is a $\mu \times n$ matrix with columns $\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_n$. We can define the matrix product AB^* as

$$AB^* = \sum_{k=1}^n \mathbf{a}_k \mathbf{b}_k^*$$

We can express any matrix product as summation of outer products.

The Singular Value Decomposition

Suppose X is an $m \times n$ matrix with singular values $\sigma_1, \sigma_2, \dots, \sigma_n$. Matrix U is unitary with columns $\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_m$ that span the column-space of X . Matrix V is unitary with columns $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$ that span the row-space of X . Matrix S is a diagonal matrix with singular values of X in decreasing order. There exists a decomposition such that

$$X = USV^*$$

Using our previous definition we can express the matrix X as

$$X = \sum_{k=1}^n \sigma_k \mathbf{u}_k \mathbf{v}_k^*$$

Because of the decreasing order of singular values, we can approximate matrix X using values.

Example of Data Compression

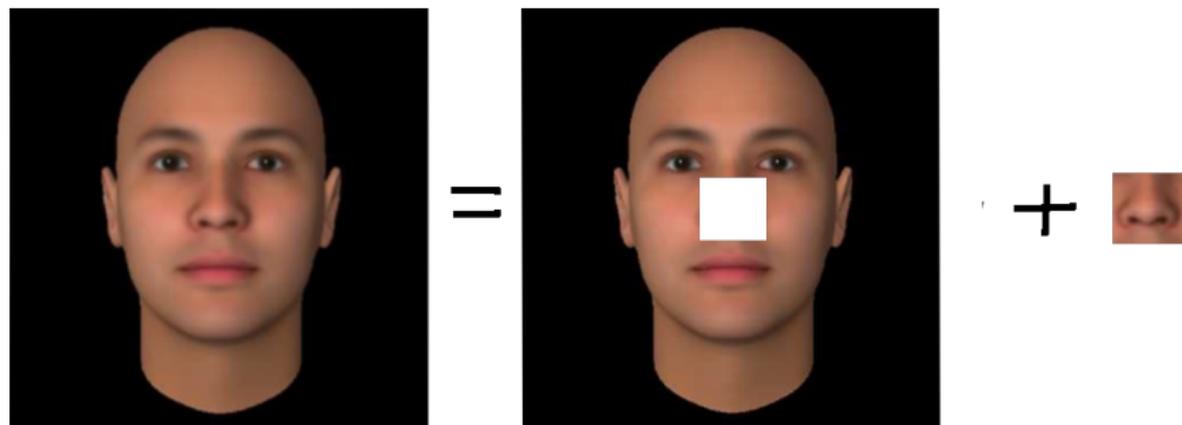
Suppose we have a set of several thousand high definition images of peoples' faces. We want to construct a new face as a linear combination of our other faces.

Instead of recording each individual pixel to represent a face, we can find clusters of pixels that are strongly correlated throughout the images, and use one value for their color. Doing this process, via the SVD, will reduce the dimensionality of the vector needed represent a face.

How, then, could we control the characteristics of the face we are designing?

Example of Data Compression

By breaking our images up into their different characteristics, we can create a span of different noses, mouths, eyes, etc. to choose from when designing our face.



Definitions of matrix Multiplication

Suppose A is an $n \times m$ matrix with columns $\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_m$ and B is a $n \times \mu$ matrix with columns $\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_\mu$. We can define any entry at row i and column j of the matrix product A^*B as

$$\begin{aligned}[A^*B]_{ij} &= \sum_{k=1}^n [A^*]_{ik} [B]_{kj} \\ &= \sum_{k=1}^n \overline{[\mathbf{a}_i]_k} [\mathbf{b}_j]_k \\ &= \langle \mathbf{a}_i, \mathbf{b}_j \rangle\end{aligned}$$

We can express any entry of a matrix product as an inner product of two vectors.

Key Aspects of the SVD

Suppose X is an $m \times n$ matrix with columns $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n$. We can define any entry at row i and column j of the correlation matrix X^*X as

$$[X^*X]_{ij} = \langle \mathbf{x}_i, \mathbf{x}_j \rangle$$

In other words each entry of the matrix, $[X^*X]_{ij}$, is the Hermitian inner product of the i th and the j th columns of the original matrix, X . Similarly the entries of the correlation matrix XX^* are inner products of two rows of matrix X .

The Netflix Prize

In 2006 the movie rental service Netflix, before creating their online streaming service, hosted a competition where, for the prize of one million dollars, contestants were asked to improve Netflix's current recommender system, Cinematch, by 10% or more. The participants were given a data set that contained the sparse ratings from 480,189 user accounts on 17,770 different movie titles. The set only contained 100,480,507 ratings total, each also include the date the rating was given.

The Netflix Prize

Lets construct the $m \times n$ data matrix, X , such that each column, \mathbf{x}_j , contains all of the ratings for one movie in the data set, in some order. The end goal is to fill in the unknown values of the data matrix X using correlations within users' movie preferences.

We can assume there exists a SVD for our data matrix, so if we would like to see what person i would rate item j we can see by the definition of matrix multiplication that

$$\begin{aligned} [X]_{ij} &= \sum_{k=1}^m [US]_{ik} [V^*]_{kj} \\ &= \sum_{k=1}^m [US]_{ik} \overline{[V]_{jk}} \\ &= \langle \hat{\mathbf{v}}_j, \hat{\mathbf{u}}_i \rangle \quad \hat{\mathbf{v}}_j \in \mathcal{R}(V), \hat{\mathbf{u}}_i \in \mathcal{R}(U) \end{aligned}$$

The Netflix Prize

Our goal is to approximate the vectors $\hat{\mathbf{v}}_j$ and $\hat{\mathbf{u}}_i$ for every unknown value, while keeping the known entries as close to their actual value as possible. We can achieve this by minimizing the system

$$\min_{\hat{\mathbf{X}}} \sum_{k_{ij} \in K} (k_{ij} - \langle \hat{\mathbf{v}}_j, \hat{\mathbf{u}}_i \rangle)^2 + \lambda (\|\hat{\mathbf{v}}_j\|^2 + \|\hat{\mathbf{u}}_i\|^2)$$

We can approximate a solution to this system by using Stochastic Gradient Descent or Alternating Least Squares. We can improve our model using biases and different data points provided in the Netflix data set.