



# Personalization

Personalization is transforming our experience of the world

Youtube

Netflix 

Amazon

Spotify

Facebook

Many more...

Almost all have share a common trait where there are users that use the system and items that we want the user to look at.

A recommender system recommends items to a user based on what we think will be the most “useful” for the user.



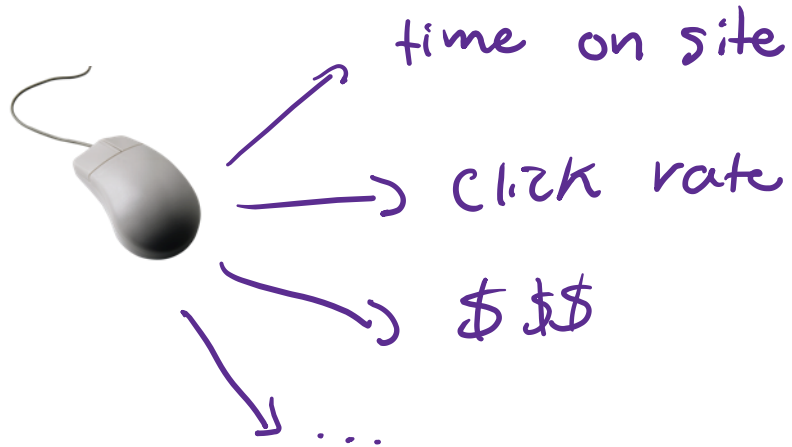
# Recommender System Challenges

# Types of Feedback

Explicit - User tells us what she likes



Implicit – We try to infer what she likes from usage data



# Top-k vs Diverse Outputs

Top-k recommendations might be very redundant

- Someone who likes Rocky I also will likely enjoy Rocky II, Rocky III, Rocky IV, Rocky V

Diverse Recommendations

- Users are multi-faceted & we want to hedge our bets
- Maybe recommend: Rocky II, Always Sunny in Philadelphia, Robin Hood



# Cold Start

When a new movie comes into our system, we don't know who likes it! This is called the **cold start** problem.

Generally, to solve we will need “side information”

- Genre, actors, if it's a sequel

Could also try to test users to see if they like it to learn quickly



# That's So Last Year

Interests change over time

- Is it 1967?
- Or 1977?
- Or 1998?
- Or 2011?

Models need flexibility to adapt to users

- Macro scale
- Micro scale (fads)



# Scalability

For  $N$  users and  $M$  movies, some approaches take  $\mathcal{O}(N^3 + M^3)$  time. This can be prohibitively slow for billions of users.

Big focus has been on:

- Efficient implementations
- Exact or faster approximate methods as needed
- *Societal impact*





# Popularity

*Solution 0*

# Popularity

Simplest Approach: Recommend whatever is popular

- Rank by global popularity (i.e. Avengers Endgame)

## Limitations

- No personalization
- Feedback loops



# Classification Model


*Solution 1*

logistic regression

=> Probability

## Learn a Classifier

Train a classifier to learn whether or not someone will like an item




### Pros

- Personalized
- Features can capture context (time of day, recent history, ...)
- Can even handle limited user history (age of user, location, ...)

# Learn a Classifier

Train a classifier to learn whether or not someone will like an item



## Cons

- Features might not be available or hard to work with
- Often doesn't perform well in practice when compared to more advanced techniques like collaborative filtering
- Can still lead to feedback loops.

# Roadmap

We will learn more advanced ideas for recommendation. Before that, we want to explore a slightly unrelated idea of dimensionality reduction.

- This will be useful in helping us think about some key ideas for what can make recommendation useful.
- A very general/powerful tool useful in other ML applications.

## Today:

- Recommender System Intro
- Dimensionality Reduction (PCA)

## Wednesday

- Matrix Factorization

# Dimensionality Reduction


# Embedding Pictures

$$16 \times 16 = 256 \Rightarrow \left[ \dots \right]_{256}$$


Example: Embed high dimensional data in low dimensions to visualize the data

- Goal: Similar images should be near each other.

Input data image




2-dimensional embedding





# Embedding Words



# Embedding Words

billmark mary  
bob jack stephen elizabeth  
tony jimmy richard henry alexander  
miss steve christopher richard charles  
joe tom harry john joseph francis maria  
mr. sam frank david james louis  
don arthur george thomas jean  
ray martin howard  
simon

Names

dr. ven al lee scott bush  
lewis  
r. a. wilson taylor johnson fox  
e. h. j. smith williams  
m. s. w. jones  
c. b. d. p. davis ford grant  
von bell

los angeles  
et da los san  
del el santa  
des hong  
core

months

june august  
february  
january september  
april october  
march

cape  
east  
south  
west  
southeast  
southwest  
northwest  
southern  
central  
northern  
western  
atlantic

latin norwegian

places

virginia missouri  
columbia indiana missouri  
colorado tennessee  
washington oregon illinois carolina  
california missouri  
houston philadelphia maryland georgia  
detroit detroit massachusetts virginia  
holliston ontario  
boston montreal  
sydney manchester cambridge  
london victoria  
berlin quebec  
moscow mexico scotland  
wales england  
canada ireland britain  
singapore australia sweden  
america norway france  
europe germany austria  
asia africa russia poland  
india japan rome  
korea china  
pakistan india  
egypt israel  
vietnam

usa philippine

families

canals

uss

# Dimensionality Reduction

Input data might have thousands or millions of dimensions!

**Dimensionality Reduction** is the the task of representing the data with a fewer number of dimensions, while keeping meaningful relations between data

- **Easier Learning:** fewer parameters, no curse of dimensionality
- **Visualization:** Hard to visualize more than 3D
- Discover “**intrinsic dimensionality**” of the data
  - High dimensional data is can sometimes be truly lower dimensional (i.e. redundant information)

*e.g., handwritten digits*



# PCA

One very popular dimensionality reduction algorithm is called **Principal Component Analysis (PCA)**.

Idea: Use a linear projection from  $d$ -dimensional data to  $k$ -dimensional data

- E.g. 1000 dimension word vectors to 3 dimensions

Choose the projection that minimizes **reconstruction error**

- Idea: The information lost if you were to "undo" the projection



# Linear Projections

$$x_i = (x_i[1], x_i[2], \dots, x_i[c])$$

Could do something like feature importance and find the subset of features with most meaningful information.

$$x_i \rightarrow z_i$$

A popular approach is to **create new features** that are **combinations of existing features**


$$\begin{aligned} u_i \quad z_i[1] &= \underline{2}x_i[1] + \underline{3}x_i[2] - \underline{7}x_i[3] + \dots \\ z_i[2] &= \underline{5}x_i[1] + \underline{0}x_i[2] + \underline{4}x_i[3] + \dots \\ &\dots \\ z_i[k] &= \dots \end{aligned}$$

Benefit: We can do this in the unsupervised setting when we only know  $x$  and not  $y$ !

# Linear Projection



$$K=1 \quad z_i[1] = u_1^T x_i = u_1[1]x_i[1] + u_1[2]x_i[2] + \dots$$

Project data into 1 dimension along a  $u_1$



# Reconstruction

Reconstruct original data only knowing the projection



Reconstruction Error



$$\hat{x}_i = z_i u_1$$

$$\|x_i - \hat{x}_i\|_2^2$$

# Projection w/ 2 Dimensions

projection of  $x_i$  on  $u_1$   $\rightarrow z_i[1] = u_1^T x_i$   
 projection of  $x_i$  on  $u_2$   $z_i[2] = u_2^T x_i$

What if projected onto  $d$  orthogonal vectors?



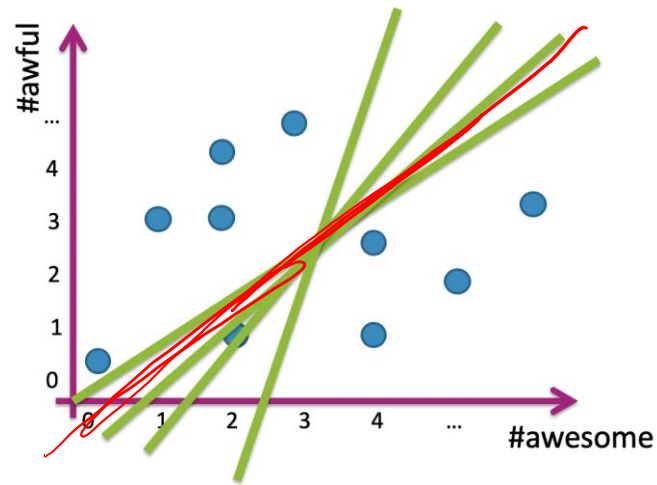
If both  $\Rightarrow$  perfect fit

$$\hat{x}_i = z_i[1] u_1 + z_i[2] u_2$$





# Which Line?

Given a dataset, how do we choose which line to project on?



Extreme Example



# PCA Visualized

This process can work with any dimension data (go from  $d$  dimensions to  $k$  dimensions).

PCA Explained Visually:

<http://setosa.io/ev/principal-component-analysis/>




# Poll Everywhere

Think 

1 min

Assume we were trying to reduce the dimension of 3d data to 2d data, which of the following would be the result of PCA?  
Assume we are using the data from the visual demo and take pc1 and pc2 as the principal components.



[pollev.com/cs416](https://pollev.com/cs416)

**1:00**


# Poll Everywhere

Group 

2 min

[pollev.com/cs416](https://pollev.com/cs416)

Assume we were trying to reduce the dimension of 3d data to 2d data, which of the following would be the result of PCA?  
Assume we are using the data from the visual demo and take pc1 and pc2 as the principal components.



**2:00**

3:27



Brain Break



# PCA Algorithm

Input Data: An  $n \times d$  data matrix  $X$

- Each row is an example

$$X = n \begin{array}{|c|} \hline x_i \\ \hline \end{array} \quad d$$

1. Recenter Data: Subtract mean from each row

$$X_c \leftarrow X - \bar{X}[1:d]$$

$$\Sigma = \frac{1}{n} \sum_{i=1}^n x_{c,i} x_{c,i}^T$$

2. Compute spread/orientation: Compute covariance matrix  $\Sigma$

$$\Sigma[t,s] = \frac{1}{n} \sum_{i=1}^n x_{c,i}[t] x_{c,i}[s]$$

3. Find basis for orientation: Compute eigenvectors of  $\Sigma$


- Select  $k$  eigenvectors  $u_1, \dots, u_k$  with largest eigenvalues

4. Project Data: Project data onto principal


$$z_i[1] = u_1^T x_{c,i} = u_1[1] x_{c,i}[1] + \dots + u_1[d] x_{c,i}[d]$$

...

$$z_i[k] = u_k^T x_{c,i} = u_k[1] x_{c,i}[1] + \dots + u_k[d] x_{c,i}[d]$$



# Reconstructing Data



Using principal components and the projected data, you can reconstruct the data in the original domain.

$$\hat{x}_i[1:d] = \bar{X}[1:d] + \sum_{j=1}^k z_i[j] u_j$$

Tried to minimize  
 $\|\hat{x}_i[1:d] - x_i\|_2^2$

# Example: Eigenfaces

Apply PCA to face data

$$X = 64 \times 256$$

Input Data



$$x_i = \begin{array}{|c|} \hline \text{256} \\ \hline \end{array}$$

$$Z = 64 \times 25$$

$$u_1, \dots, u_{25}$$

Principal Components



$$z_i = \begin{array}{|c|} \hline \text{25} \\ \hline \end{array}$$



# Reconstructing Faces

Depending on context, it may make sense to look at either original data or projected data.

In this case, let's see how the original data looks after using more and more principal components for reconstruction.


- Each image shows additional 8 principal components



# Embedding Images

Other times, it does make sense to look at the data in the  
projected space! (Usually if  $k \leq 3$ )


tilt



expression

# Example: Genes

Dataset of genes of Europeans (3192 people; 500,568 loci) and their country of origin, ran PCA on the data and plotted 2 principal components.



# Poll Everywhere

Think 

1 min

Using the PCA algorithm described in class on this gene data is an example of a:

- Supervised Learning Problem
- ★  Unsupervised Learning Problem
- Neither
- I'm not sure

[pollev.com/cs416](http://pollev.com/cs416)

**1:00**

# Scaling Up

Covariance matrix  $\Sigma$  can be very large with high-dimensional data

- $\Sigma \in \mathbb{R}^{d \times d}$  which can be quite large for 10,000 features
- This means finding the principal components will be slow


In practice, you can use the Singular Value Decomposition (SVD)

- Can be used to find the  $k$  eigenvectors with largest eigenvalue
- Very fast implementations available
- Don't care if you know what SVD is, but you will likely use it in practice if you are doing anything with PCA.



# PCA Failure Modes

PCA assumes there is a lower dimensional **linear subspace** that represents the data well. Works some times, but can fail in practice.



May want to look into non-linear dimensionality reduction

- Manifold learning
- Popular: SDD Maps, Isomap, LLE, t-SNE

# PCA Recap

## Dimensionality Reduction

- Why and when it's important

## Principal Component Analysis (PCA)

- High level intuition for what the algorithm is doing
- Goal: Minimizing reconstruction error

