

## Recommender Systems

- Look at classic model and techniques
  - Items
  - Users
  - Recommend Items to Users
- Recommend new items based on:
  - similarity to items user liked in past: **individual history**  
“Content Filtering”
  - Liked by other users similar to this user: **collaborative history**  
“Collaborative Filtering”
  - Liked by other users: **crowd history**
    - easier case

1

## Recommender System attributes

- Need explicit or implicit ratings by user
  - Purchase is 0/1 rating
    - Movie tickets
    - Books
- Have focused category
  - examples: music, courses, restaurants
  - hard to cross categories with content-based
  - easier to cross categories with collaborative-based
    - users share tastes across categories?

2

## Content Filtering

- Items must have **characteristics**
- user values item
  - ⇒ values characteristics of item
- model each item as **vector** of weights of characteristics
  - much like vector-based IR
- user can give explicit preferences for certain characteristics

3

## Buy/no buy prediction method: similarity with centroid

- Average vectors of items user bought
  - **user's centroid**
- Find **similarity of new items** to user's centroid
- Decide **threshold** for “buy” recommendation

4

## Example

- user bought book 1 and book 2
- Average books bought = **(0, 1, 0.5, 0)**
- Score new books
  - dot product gives: score(A) = 0.5; score (B)= 1
- decide threshold for recommendation

	1 <sup>st</sup> person	romance	mystery	sci-fi
book 1	0	1	1	0
book 2	0	1	0	0
new book A	1	.5	0	0
new book B	0	1	0	.2

5

## Method issues

- Centroid best way to build a **preference vector**?
- What **metric** use for **similarity** between new items and preference vector?
  - Normalization?
- What if users give **ratings**?
  - Centroid per rating value?
- how include **explicit user preferences**
- How determine **threshold**?

6

## Example with explicit user preferences

How use scores of books bought?

Try: **preference vector**  $p$  where component  $k$  =

**user pref** for characteristic  $k$  if  $\neq 0$

**avg. comp.  $k$  of books bought** when user pref = 0

0 pref for user = "don't care"

$p = (0, 1, 0.5, -5)$

New scores?

$p \cdot A = 0.5$

$p \cdot B = 0$

	1 <sup>st</sup> per	rom	mys	sci-fi
<b>user pref</b>	<b>0</b>	<b>1</b>	<b>0</b>	<b>-5</b>
book 1	0	1	1	0
book 2	0	1	0	0
new A	1	.5	0	0
new B	0	1	0	.2

7

## Other methods: machine learning

- Major alternatives based on **classifiers**
  - Training set: items bought and not bought
  - Train classifier – many algorithms
  - Classify new item as buy/no buy
- Observations
  - Uses books not bought. Problems?
  - Multiple rating value
  - Can use multiple classes

8

## Limitations of Content Filtering

- Can only recommend items similar to those user rated highly
- New users
  - Insufficient number of rated items
- Only consider features explicitly associated with items
  - Do not include attributes of user

9

## Applying concepts to search

- Individual histories
  - Characterize individual by topic interest
    - Properties of objects interact with
  - Characterize query by related topics
    - Role of terms of query in topic
  - Modify query to bias to shared topics
  - Modify ranking to prefer shared topics

10

## Example study: Personalizing Web Search Using Long-term Browsing History (in *WSDM11*)

- Goal: **rerank**
  - top 50 results from Google query
- Strategy:
  - **score snippets** from search result against **user profile**
  - **rerank** based on snippet score
- Selection of info for user profile
  - list of visited URLs w/ number visits
  - list of past search queries and pages clicked
  - list of terms with **weights** for content of pages visited

11

## Personalizing Web Search Using Long-term Browsing History, cont

Studies **selection of methods** for

- user profile: what **sources of terms** use
- user profile: **weights for terms**
  - **tf-idf**
    - where get idf?
  - "**modified BM25**" - a "log odds measure"
- scoring
  - **language model** with adjustments for
    - **URLs previously visited**
    - **original rank of snippet in search**

worked best

performed best

12

## $W_{\text{modBM25}}$ weighting

$N$  = # documents on Web – estimated  
 $n_{t_i}$  = # docs on Web containing term  $t_i$  - estimated  
 $R$  = # documents in user browser history  
 $r_{t_i}$  = # docs in user browser history that contain term  $t_i$

$W_{\text{modBM25}}(t_i) =$

$$\log \left[ \frac{(r_{t_i} + 0.5)(N - n_{t_i} + 0.5)}{(n_{t_i} + 0.5)(R - r_{t_i} + 0.5)} \right]$$

13

## Scoring

$N_{s_i}$  = # unique words in snippet  $s_i$   
 $r_{s_i}$  = rank of snippet  $s_i$  in original search results  
 $n_i$  = # previous visits by user to web page with snippet  $s_i$   
 $w(t_k)$  = weight of term  $t_k$  in user profile  
 $w_{\text{total}}$  = sum of all term weights in user profile

$$\text{score}_{\text{lang. model}}(s_i) = \sum_{k=0}^{N_{s_i}} \log((w(t_k) + 1)/w_{\text{total}})$$

- modif. for URLs previously visited:  
 $\text{score}_{w/\text{URL}}(s_i) = \text{score}(s_i) * (1 + v * n_i)$  *parameter v*
- modif to acct. for orig. rank:  
 $\text{score}_{w/\text{orig}}(s_i) = \text{score}(s_i) * (1 / (1 + \log(r_{s_i})))$

14

## Personalizing Web Search Using Long-term Browsing History Evaluation

- “offline” evaluation:
  - relevance judgments by volunteers
  - used to select best of algorithmic variations
- online evaluation of best variations:
  - add-on to Browser by volunteers
  - interleave original results (no personalization) with results reranked by snippet score
  - record clicks by user – which list from

15

## Personalizing Web Search Using Long-term Browsing History Results

- Offline: normalized DCG, avg. of 72 queries
  - Google’s ranking w/out personalization: 0.502
  - best-performing of variations for reranking: 0.573
- Online
  - 8% queries: # clicks from original and reranked same
  - of rest: 60.5% queries: more clicks from reranked  
39.5% queries: more clicks from original

### Observation

- Reranking can be done **completely in browser** if enough space for data for user profile

16

## What we’ve just seen:

Recommender systems: Content Filtering  
Applying content filtering to search

## Now back to recommender systems:

**Collaborative Filtering**

17

## Collaborative Filtering

- Recommend new items liked by other **users similar to this user**
- need items already rated by user **and other users**
- don’t need characteristics of items
  - each rating by individual user becomes characteristic
- Can combine with item characteristics
  - hybrid content/collaborative

18

## Major method types

- **Nearest neighbor**
  - Use similarity function
  - Prediction based on previously rated items
- **Matrix Factorization**
  - “Latent factors”
  - Matrix decomposition
- Both use (user × item) matrix
  - vector similarity

19

## Example of nearest neighbor: Preliminaries

- Notation
  - $r(u,i)$  = rating of  $i^{\text{th}}$  item by user  $u$
  - $I_u$  = set of items rated by user  $u$
  - $I_{u,v}$  = set of items rated by both users  $u$  and  $v$
  - $U_{i,j}$  = set of users that rated items  $i$  and  $j$
- Adjust scales for user differences
  - Use average rating by user  $u$ :  

$$r_u^{\text{avg}} = (1/|I_u|) * \sum_{i \in I_u} r(u,i)$$
  - Adjusted ratings:  $r_{\text{adj}}(u,i) = r(u,i) - r_u^{\text{avg}}$

20

## One choice of similarity function: User Similarities

- similarity between users  $u$  and  $v$ 
  - Pearson correlation coefficient

$$\text{sim}(u,v) = \frac{\sum_{i \in I_{u,v}} (r_{\text{adj}}(u,i) * r_{\text{adj}}(v,i))}{(\sum_{i \in I_{u,v}} (r_{\text{adj}}(u,i))^2 * \sum_{i \in I_{u,v}} (r_{\text{adj}}(v,i))^2)^{1/2}}$$

21

## Predicting User's rating of new item: User-based

For item  $i$  not rated by user  $u$

$$r^{\text{pred}}(u,i) = r_u^{\text{avg}} + \frac{\sum_{v \in S} (\text{sim}(u,v) * r_{\text{adj}}(v,i))}{\sum_{v \in S} |\text{sim}(u,v)|}$$

$S$  can be all users who have rated  $i$  or just those users *most similar* to  $u$

22

## Collaborative filtering example

user ratings		book 1	book 2	book 3	book 4
user 1		5	1	2	0
user 2		x	5	2	5
user 3		3	1	x	2
user 4		4	0	2	?

adj. user ratings		book 1	book 2	book 3	book 4
user 1		3	-1	0	-2
user 2		x	1	-2	1
user 3		1	-1	x	0
user 4		2	-2	0	?

## Collaborative filtering example

- $\text{sim}(u1,u4) = (6+2)/(10*8)^{1/2} = .894$
- $\text{sim}(u2,u4) = (-2)/(5*4)^{1/2} = -.447$
- $\text{sim}(u3,u4) = (2+2)/(2*8)^{1/2} = 1$

- predict  $r(u4, \text{book}4) = 2 + \frac{(-2)*.894 + 1*(-.447) + 0*1}{.894 + .447 + 1}$   
 $= 2 - .955 \approx 1$

24

### Another choice of similarity function: Item Similarities

- similarity between items  $i$  and  $j$ 
  - vector of ratings of users **in**  $U_{ij}$
  - cosine measure using adjusted ratings

$$\text{sim}(i,j) = \frac{\sum_{u \text{ in } U_{ij}} (r_{\text{adj}}(u,i) * r_{\text{adj}}(u,j))}{(\sum_{u \text{ in } U_{ij}} (r_{\text{adj}}(u,i))^2 \sum_{u \text{ in } U_{ij}} (r_{\text{adj}}(u,j))^2)^{1/2}}$$

25

### Predicting User's rating of new item: Item-based

For item  $i$  not rated by user  $u$

$$r^{\text{item-pred}}(u,i) = \frac{\sum_{j \text{ in } T} (\text{sim}(i,j) * r(u,j))}{\sum_{j \text{ in } T} |\text{sim}(i,j)|}$$

$T$  can be all items in  $I_u$  or just items *most similar* to  $i$

- Prediction uses only  $u$ 's ratings, but similarity uses other users' ratings

26

### Limitations

- May not have enough ratings for new users
- New items may not be rated by enough users
- Need "critical mass" of users
  - All similarities based on user ratings

But can take user "out of comfort zone"

27

### Applying nearest-neighbor collab. filtering concepts to search

- Collaborative histories
  - How determine user similarity?
    - Behavior on identical searches?
    - Overlap of general topic interests?
      - From overlapping behaviors
      - Hybrid content-based and behavior-based
    - Computational expense?
      - Argues for general topic-interest characterizations
  - How apply similarity?
    - Same search? Bias ranking?
    - Same topic of search? Bias topics of results?

28

### Example

from A Large-scale Evaluation and Analysis of Personalize Search Strategies (in WWW07)

- Goal: **rerank** search results
- Based on query log history – **clicks** as ratings
- Also uses 67 pre-defined **topic categories**
- Strategy:
  - get **similarity of users** based on **user history** of visited pages
  - find  $K$  most similar users to user doing search
    - $K$  nearest neighbor; use  $K=50$
  - calc. score for each result of search based on click history of  $K$  nearest neighbors
  - **rerank** results of search based on score

29

### Details

from A Large-scale Evaluation and Analysis of Personalize Search Strategies (in WWW07)

$P(u)$  = collection of Web pages visited by user  $u$  in the past

$$P(p|u) = \frac{\text{\# times } u \text{ clicked on page } p \text{ in past}}{\text{total \# times } u \text{ clicked on a page in past}}$$

$$w(p) = \log(\text{total \# users} / \text{\# users visited page } p)$$

"impact weight" - idf-like

$c(p)$  = "category vector" for page  $p$

do classification of page

vector gives confidence # for top 6 categories (other entries 0)

$$\text{User profile } c_i(u) = \sum_{p \text{ in } P(u)} P(p|u)w(p)c(p)$$

$$\text{User similarity } \text{sim}(u_1, u_2) = \frac{c_i(u_1) \cdot c_i(u_2)}{\|c_i(u_1)\| \|c_i(u_2)\|}$$

30

## Details

from A Large-scale Evaluation and Analysis of Personalized Search Strategies (in WWW07)

$S_k(u_a)$  denotes  $k$  nearest neighbors of user  $u_a$

click history:

$|clicks(q,p,u_s)|$  = # clicks on pg  $p$  by user  $u_s$  on past query  $q$

$|clicks(q,*,u_s)|$  = # clicks overall by user  $u_s$  on past query  $q$

the score of a page  $p$  for query  $q$  and user  $u$ :

$$S(q,p,u) = \frac{\sum_{u_s \in S_k(u)} \text{sim}(u_s, u) * |clicks(q,p,u_s)|}{\beta + \sum_{u_s \in S_k(u)} |clicks(q,*,u_s)|}$$

$\beta$  is a "smoothing factor"; taken to be 0.5

31

## Experiments

from A Large-scale Evaluation and Analysis of Personalized Search Strategies (in WWW07)

- Data set: **MSN query logs** 12 days August 2006  
sampled 10,000 distinct users  
used 11 days for training, last day for testing  
~ 4000 test queries
- Action, for each user and query
  - re-rank top 50 results using a "fusion" of original rank and order given by page scores  $S(q,p,u)$
- Evaluation: 2 metrics
  1. a DCG-like metric with clicking indicating relevance
  2. average rank of clicked items

32

## Results

from A Large-scale Evaluation and Analysis of Personalized Search Strategies (in WWW07)

- Good news:  
re-ranking improves over original ranking
- So-so news:  
improvement is 3.62% on queries where there is room for improvement
- Not so good news:  
non-collaborative personalization improves 3.68%

$$S(q,p,u) = \frac{|clicks(q,p,u)|}{\beta + |clicks(q,*,u)|}$$

33

## Where are we?

- Refinement/Personalization of results
- Study techniques of
  - Recommender systems
    - Content filtering
      - Applying content filtering to search
    - Collaborative filtering
      - Nearest neighbor methods
        - Applying nearest neighbor method to search
      - Matrix factorization methods

34