#### Recommender Systems and Adaptive Hypermedia

#### **Collaborative Filtering**





Costas Mourlas Associate Professor Univ. of Athens

#### Agenda

- Collaborative Filtering (CF)
  - Pure CF approaches
  - User-based nearest-neighbor
  - The Pearson Correlation similarity measure
  - Memory-based and model-based approaches
  - Item-based nearest-neighbor
  - The cosine similarity measure
  - Data sparsity problems
  - Recent methods (SVD, AssociationRuleMining,SlopeOne,RF-Rec, ...)
  - The Google News personalization engine
  - Discussion and summary
  - Literature

## Collaborative Filtering (CF)

The most prominent approach to generate recommendations

- used by large, commercial e-commerce sites
- well-understood, various algorithms and variations exist
- applicable in many domains (book, movies, DVDs, ..)
- Approach



- use the "wisdom of the crowd" to recommend items
- Basic assumption and idea
  - Users give ratings to catalog items (implicitly or explicitly)
  - Customers who had similar tastes in the past, will have similar tastes in the future

#### Pure CF Approaches

- 😼 Input
  - Only a matrix of given user-item ratings
- Output types
  - A (numerical) prediction indicating to what degree the current user will like or dislike a certain item
  - A top-N list of recommended items

## User-based nearest-neighbor collaborative filtering (1)

#### The basic technique

- Given an "active user" (Alice) and an item *i* not yet seen by Alice
  - find a set of users (peers/nearest neighbors) who liked the same items as Alice in the past **and** who have rated item *i*
  - $^{\scriptscriptstyle =}$  use, e.g. the average of their ratings to predict, if Alice will like item i
  - do this for all items Alice has not seen and recommend the best-rated

#### **Basic assumption and idea**

If users had similar tastes in the past they will have similar tastes in the future User preferences remain stable and consistent over time

## User-based nearest-neighbor collaborative filtering (2)

#### & Example

A database of ratings of the current user, Alice, and some other users is given:

	ltem1	ltem2	Item3	Item4	ltem5
Alice	5	3	4	4	?
User1	3	1	2	3	3
User2	4	3	4	3	5
User3	3	3	1	5	4
User4	1	5	5	2	1

Determine whether Alice will like or dislike *Item5*, which Alice has not yet rated or seen

## User-based nearest-neighbor collaborative filtering (3)

Some first questions



- How do we measure similarity?
  - How many neighbors should we consider?
- How do we generate a prediction from the neighbors' ratings?

	ltem1	ltem2	ltem3	ltem4	ltem5
Alice	5	3	4	4	?
User1	3	1	2	3	3
User2	4	3	4	3	5
User3	3	3	1	5	4
User4	1	5	5	2	1

#### Measuring user similarity (1)

- **A popular similarity measure in user-based CF: Pearson correlation** *a, b*: users
  - $r_{a,p}$  : rating of user a for item p

k × 1

- set of items, rated both by a and b
- Possible similarity values between -1 and 1

$$m(a,b) = \frac{\sum_{p \in P} (r_{a,p} - \bar{r}_a) (r_{b,p} - \bar{r}_b)}{\sqrt{\sum_{p \in P} (r_{a,p} - \bar{r}_a)^2} \sqrt{\sum_{p \in P} (r_{b,p} - \bar{r}_b)^2}}$$

#### Measuring user similarity (2)

- A popular similarity measure in user-based CF: Pearson correlation a, b : users
  - $r_{a,p}$ : rating of user a for item p

()

- ; set of items, rated both by a and b
- Possible similarity values between -1 and 1

	ltem1	ltem2	ltem3	ltem4	ltem5		
Alice	5	3	4	4	?		
User1	3	1	2	3	3		sim = 0,85
User2	4	3	4	3	5		sim = 0,00
User3	3	3	1	5	4	$\checkmark$	sim = 0,70
User4	1	5	5	2	1	$\checkmark$	sim = -0,79

#### Pearson correlation

## Takes differences in rating behavior into account



Works well in usual domains, compared with alternative measures
 – such as cosine similarity

#### Making predictions

A common prediction function:

$$pred(a,p) = \overline{r_a} + \frac{\sum_{b \in N} sim(a,b) * (r_{b,p} - \overline{r_b})}{\sum_{b \in N} sim(a,b)}$$



Calculate, whether the neighbors' ratings for the unseen item *i* are higher or lower than their average

- $\checkmark$  Combine the rating differences use the similarity with a as a weight
- Add/subtract the neighbors' bias from the active user's average and use this as a prediction

## Improving the metrics / prediction function

- Not all neighbor ratings might be equally "valuable"
  - Agreement on commonly liked items is not so informative as agreement on controversial items
  - **Possible solution**: Give more weight to items that have a higher variance
- Value of number of co-rated items
  - Use "significance weighting", by e.g., linearly reducing the weight when the number of co-rated items is low
- Case amplification
  - Intuition: Give more weight to "very similar" neighbors, i.e., where the similarity value is close to 1.
- Neighborhood selection
  - Use similarity threshold or fixed number of neighbors

# Memory-based and model-based approaches

- User-based CF is said to be "memory-based"
  - the rating matrix is directly used to find neighbors / make predictions
  - does not scale for most real-world scenarios
    - large e-commerce sites have tens of millions of customers and millions of items
- Model-based approaches
  - based on an offline pre-processing or "model-learning" phase
  - at run-time, only the learned model is used to make predictions
  - models are updated / re-trained periodically
  - large variety of techniques used
  - model-building and updating can be computationally expensive
- *item*-based CF is an example for model-based approaches

#### Item-based collaborative filtering

#### Basic idea:

Use the similarity between items (and not users) to make predictions

🕹 Example:

Look for items that are similar to Item5

Take Alice's ratings for these items to predict the rating for

ltem5

ltem1	ltem2	ltem3	ltem4	ltem5
5	3	4	4	?
3	1	2	3	3
4	3	4	3	5
3	3	1	5	4
1	5	5	2	1
	Item1 5 3 4 3 1	Item1      Item2        5      3        3      1        4      3        3      3        1      5	Item1      Item2      Item3        5      3      4        3      1      2        4      3      4        3      1      2        4      3      4        1      5      5	Item1      Item2      Item3      Item4        5      3      4      4        3      1      2      3        4      3      4      3        3      1      2      3        1      5      5      2

#### The cosine similarity measure

#### Produces better results in item-to-item filtering

- Ratings are seen as vector in n-dimensional space
- Similarity is calculated based on the angle between the vectors  $sim(\vec{a}, \vec{b}) = \frac{\vec{a} \cdot \vec{b}}{|\vec{a}| * |\vec{b}|}$

#### Adjusted cosine similarity

take average user ratings into account, transform the original ratings U set of users who have rated both items a and b

$$sim(\vec{a},\vec{b}) = \frac{\sum_{u \in U} (r_{u,a} - \overline{r_u}) (r_{u,b} - \overline{r_u})}{\sqrt{\sum_{u \in U} (r_{u,a} - \overline{r_u})^2} \sqrt{\sum_{u \in U} (r_{u,b} - \overline{r_u})^2}}$$





## Making predictions

- A common prediction function:
- $pred(u,p) = \frac{\sum_{i \in ratedItem(u)} sim(i,p) * r_{u,i}}{\sum_{i \in ratedItem(u)} sim(i,p)}$ Weighborhood size is typically also limited to a specific size
- Not all neighbors are taken into account for the prediction
- An analysis of the MovieLens dataset indicates that "in most real-world situations, a neighborhood of 20 to 50 neighbors seems reasonable" (Herlocker et al. 2002)

## Pre-processing for item-based filtering

Item-based filtering does not solve the scalability problem itself

#### Pre-processing approach by Amazon.com (in 2003)

- Calculate all pair-wise item similarities in advance
- The neighborhood to be used at run-time is typically rather small, because only items are taken into account which the user has rated
- Item similarities are supposed to be more stable than user similarities

#### Memory requirements

- Up to  $N^2$  pair-wise similarities to be memorized (N = number of items) in theory
  - In practice, this is significantly lower (items with no co-ratings)
  - Further reductions possible
    - Minimum threshold for co-ratings
    - Limit the neighborhood size (might affect recommendation accuracy)

#### More on ratings – Explicit ratings

- Probably the most precise ratings
- Most commonly used (1 to 5, 1 to 7 Likert response scales)
- Research topics
  - Optimal granularity of scale; indication that 10-point scale is better accepted in movie dom.
  - An even more fine-grained scale was chosen in the joke recommender discussed by Goldberg et al. (2001), where a continuous scale (from –10 to +10) and a graphical input bar were used
    - No precision loss from the discretization
    - User preferences can be captured at a finer granularity
    - Users actually "like" the graphical interaction method
    - Multidimensional ratings (multiple ratings per movie such as ratings for actors and sound)

Main problems

- Users not always willing to rate many items
  - number of available ratings could be too small  $\rightarrow$  sparse rating matrices  $\rightarrow$  poor recommendation (quality
- How to stimulate users to rate more items?

#### More on ratings – Implicit ratings

- Typically collected by the web shop or application in which the recommender system is embedded
- When a customer buys an item, for instance, many recommender systems interpret this behavior as a positive rating
- Clicks, page views, time spent on some page, demo downloads ...
- Implicit ratings can be collected constantly and do not require additional efforts from the side of the user
- 🦌 Main problem
  - One cannot be sure whether the user behavior is correctly interpreted
  - For example, a user might not like all the books he or she has bought; the user also might have bought a book for someone else
- Implicit ratings can be used in addition to explicit ones; question of correctness of interpretation

#### Data sparsity problems

- Cold start problem
  - How to recommend new items? What to recommend to new users?
- Straightforward approaches
  - Ask/force users to rate a set of items
  - Use another method (e.g., content-based, demographic or simply non-personalized) in the initial phase
  - Default voting: assign default values to items that only one of the two users to be compared has rated (Breese et al. 1998)
- Alternatives
  - Use better algorithms (beyond nearest-neighbor approaches)
  - Example:
    - In nearest-neighbor approaches, the set of sufficiently similar neighbors might be too small to make good predictions
      - Assume "transitivity" of neighborhoods

## Example algorithms for sparse datasets

- Recursive CF (Zhang and Pu 2007)
  - Assume there is a very close neighbor n of u who however has not rated the target item i yet.
  - Idea:
    - Apply CF-method recursively and predict a rating for item *i* for the neighbor
      Use this predicted rating instead of the rating of a more distant direct
      neighbor

	ltem1	ltem2	ltem3	ltem4	Item5	
Alice	5	3	4	4	? 🗖	
User1	3	1	2	3	?	sim = 0.85
User2	4	3	4	3	5	Predict
User3	3	3	1	5	4	rating for
User4	1	5	5	2	1	User1

#### More model-based approaches

- Plethora of different techniques proposed in the last years,
  - Matrix factorization techniques, statistics
    - singular value decomposition, principal component analysis
  - Association rule mining
    - compare: shopping basket analysis
  - Probabilistic models

e.g.,

- clustering models, Bayesian networks, probabilistic Latent Semantic Analysis
- Various other machine learning approaches
- Costs of pre-processing
  - Usually not discussed
  - Incremental updates possible?

#### **Collaborative Filtering Issues**

well-understood, works well in some domains, no knowledge engineering required

🕹 Cons:''

🖕 Pros: 🗅

requires user community, sparsity problems, no integration of other knowledge sources, no explanation of results

#### What is the best CF method?

In which situation and which domain? Inconsistent findings; always the same domains and data sets; differences between methods are often very small (1/100)

#### How to evaluate the prediction quality?

MAE / RMSE: What does an MAE of 0.7 actually mean? Serendipity (novelty and surprising effect of recommendations) Not yet fully understood

#### What about multi-dimensional ratings?

## Google News personalization engine



Search News Search the Web Search and browse 4,500 news sources updated continuously.

News archive search | Advanced news search | Blog search

Auto-generated 13 minutes ag	JO
------------------------------	----





#### Edit this personalized page

Fed cuts key interest rate Los Angeles Times - all 510 news articles »

Obama on race Los Angeles Times - all 200 news articles »

US, Russia Politely Dug In Over Missile Defense Washington Post - all 1,096 news articles »

Sci-fi guru Sir Arthur C. Clarke dies Vancouver Sun - all 976 news articles »

Facebook Beefs Up Privacy Options, Readies Online Chat Washington Post - all 297 news articles »

Mills' Money Can't Buy Her Love E! Online - all 3,490 news articles »

Boeing confident of winning back tanker deal Reuters - all 200 news articles »

#### In The News

Dalai Lama Windows Vista Barack Obama Halle Berry



MSN UK News

## Google News portal (1)

- Aggregates news articles from several thousand sources
- Displays them to signed-in users in a personalized way
- Collaborative recommendation approach based on
  - the click history of the active user and
  - the history of the larger community
- Main challenges
  - Vast number of articles and users
    - Generate recommendation list in real time (at most one second)
    - Constant stream of new items
    - Immediately react to user interaction

Significant efforts with respect to algorithms, engineering, and parallelization are required

### Google News portal (2)

Pure memory-based approaches are not directly applicable and for model-based approaches, the problem of continuous model updates must be solved

A combination of model- and memory-based techniques is used

Model-based part: Two clustering techniques are used

Probabilistic Latent Semantic Indexing (PLSI) as proposed by (Hofmann 2004)

MinHash as a hashing method

 Memory-based part: Analyze story *co-visits* for dealing with new users
 Google's MapReduce technique is used for parallelization in order to make computation scalable

#### Literature (1)

- [Adomavicius and Tuzhilin 2005] Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions, IEEE Transactions on Knowledge and Data Engineering 17 (2005), no. 6, 734–749
- IBreese et al. 1998] Empirical analysis of predictive algorithms for collaborative filtering, Proceedings of the 14th Conference on Uncertainty in Artificial Intelligence (Madison, WI) (Gregory F. Cooper and Seraf´in Moral, eds.), Morgan Kaufmann, 1998, pp. 43–52
- [Gedikli et al. 2011] RF-Rec: Fast and accurate computation of recommendations based on rating frequencies, Proceedings of the 13th IEEE Conference on Commerce and Enterprise Computing - CEC 2011, Luxembourg, 2011, forthcoming
- [Goldberg et al. 2001] Eigentaste: A constant time collaborative filtering algorithm, Information Retrieval 4 (2001), no. 2, 133–151
- Golub and Kahan 1965] Calculating the singular values and pseudo-inverse of a matrix, Journal of the Society for Industrial and Applied Mathematics, Series B: Numerical Analysis 2 (1965), no. 2, 205–224
- [Herlocker et al. 2002] An empirical analysis of design choices in neighborhood-based collaborative filtering algorithms, Information Retrieval 5 (2002), no. 4, 287–310
- [Herlocker et al. 2004] Evaluating collaborative filtering recommender systems, ACM Transactions on Information Systems (TOIS) 22 (2004), no. 1, 5–53

#### Literature (2)

- [Hofmann 2004] Latent semantic models for collaborative filtering, ACM Transactions on Information Systems 22 (2004), no. 1, 89–115
- [Huang et al. 2004] Applying associative retrieval techniques to alleviate the sparsity problem in collaborative filtering, ACM Transactions on Information Systems 22 (2004), no. 1, 116–142
- [Koren et al. 2009] *Matrix factorization techniques for recommender systems*, Computer 42 (2009), no. 8, 30–37
- [Lemire and Maclachlan 2005] Slope one predictors for online rating-based collaborative filtering, Proceedings of the 5th SIAM International Conference on Data Mining (SDM '05) (Newport Beach, CA), 2005, pp. 471–480
- [Sarwar et al. 2000a] Application of dimensionality reduction in recommender systems a case study, Proceedings of the ACM WebKDD Workshop (Boston), 2000
- [Zhang and Pu 2007] A recursive prediction algorithm for collaborative filtering recommender systems, Proceedings of the 2007 ACM Conference on Recommender Systems (RecSys '07) (Minneapolis, MN), ACM, 2007, pp. 57–64