The GroupLens Research Project:

Collaborative Filtering Recommender Systems

Joseph A. Konstan
University of Minnesota
konstan@cs.umn.edu
http://www.grouplens.org

About me …

Associate Professor of Computer Science & Engineering, Univ. of Minnesota

Ph.D. (1993) from U.C. Berkeley

- GUI toolkit architecture

Teaching Interests: HCI, GUI Tools

Research Interests: General HCI, and ...

- Collaborative Information Filtering
- Multimedia Authoring and Systems
- Web Automation
- Visualization and Information Management

The Problem: Information Overload

Too many

- research papers
- books
- movies
- web pages
- … even Usenet News articles!

Recommenders

Tools to help identify worthwhile stuff

- Filtering interfaces
  - ➔E-mail filters, clipping services
- Recommendation interfaces
  - ➔Suggestion lists, "top-n," offers and promotions
- Prediction interfaces
  - ➔Evaluate candidates, predicted ratings

History of Recommender Systems

The Early Years …

Why cave dwellers survived

How editors are like cave dwellers

Critics, critics, everywhere

Information Filtering

Information retrieval

- Dynamic information need
- Static content base

Information filtering

- Static information need
- Dynamic content base

Collaborative Filtering

Premise

- Information needs more complex than keywords or topics: quality and taste

Small Community: Manual

- Tapestry – database of content & comments
- Active CF – easy mechanisms for forwarding content to relevant readers

Automated CF

The GroupLens Project
(Resnick et al. CSCW '94)

- ACF for Usenet News
  - ➔users rate items
  - ➔users are correlated with other users
  - ➔personal predictions for unrated items
- Nearest-Neighbor Approach
  - ➔find people with history of agreement
  - ➔assume stable tastes

Usenet Trial

(Miller et al. Usenet '97;
Konstan et al. CACM Mar. '97)

Medium-scale Usenet trial

- seven weeks
- 250 users; 47,569 ratings; over 600,000 predictions
- variety of newsgroups
  - ➔moderated and unmoderated
  - ➔technical and recreational
- gathered reading activity as well as ratings

Does it Work?

Yes: The numbers don't lie!

- Usenet trial: rating/prediction correlation
  - ➔rec.humor: 0.62 (personalized) vs. 0.49 (avg.)
  - ➔comp.os.linux.system: 0.55 (pers.) vs. 0.41 (avg.)
  - ➔rec.food.recipes: 0.33 (pers.) vs. 0.05 (avg.)
- Significantly more accurate than predicting average or modal rating.
- Higher accuracy when partitioned by newsgroup

It Works Meaningfully Well!

Relationship with User Behavior

- Twice as likely to read 4/5 than 1/2/3

Users *Like* GroupLens

- Some users stayed 12 months after the trial!

ACF Blossomed

1995

- Ringo (later Firefly)
- Bellcore Video Recommender

1996 Recommender Systems Workshop

Early commercialization

- Agents Inc. (later Firefly)
- Net Perceptions

new issues of scale and performance!

Today

Broad research community

- live research systems
- substantial integration with machine learning, information filtering

Increasing commercial application

- available commercial tools

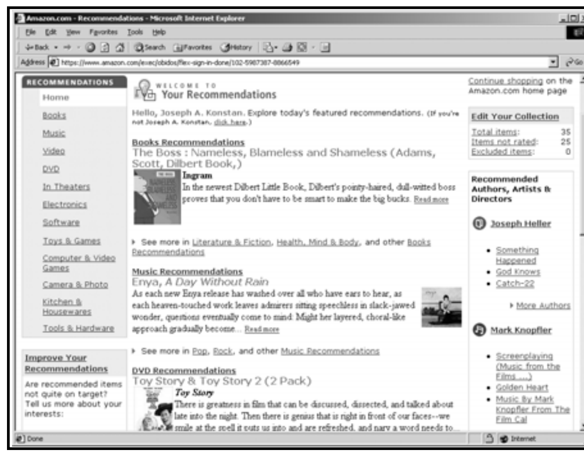

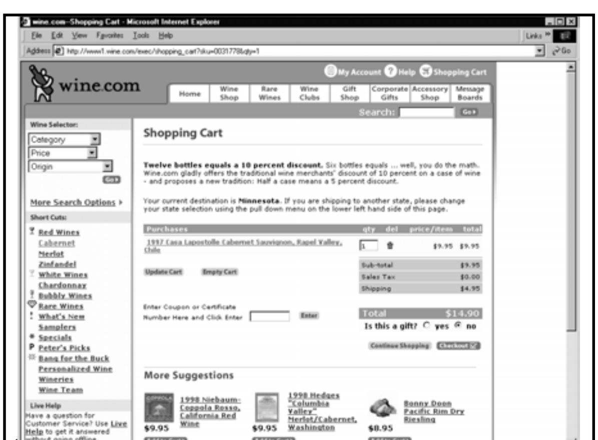

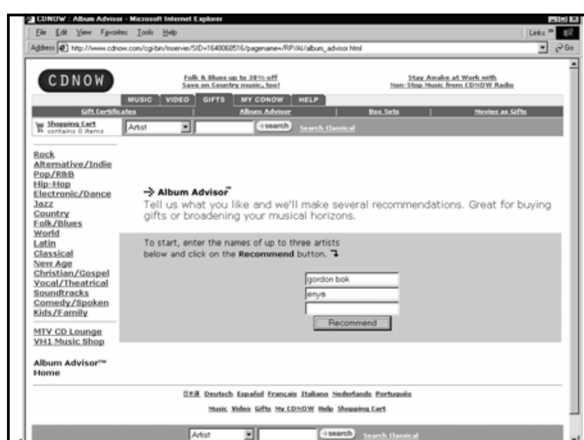


Ace Of Base
Cruel Summer
Fantasia 2000
Erasure
Abba-Esque
Vangelis
Opera Sauvage
See complete track list and more album info.


movielens
helping you find the right movies


movielens
helping you find the right movies


movielens
helping you find the right movies


movielens
helping you find the right movies


movielens
helping you find the right movies

movielens - Microsoft Internet Explorer

http://movielens.umn.edu/search.cgi

movielens

helping you find the right movies

Home · Your Ratings · To See List · Account Info · Groups · Reviews · Help · Privacy · Logoff

Search For Movie Title: Go | Select group: -- NO GROUP -- | Get Recommendations: Comedy, Any Year, Go

Rating more movies improves your predictions; you've rated **36** so far.
Currently displaying **Comedy** movies matching the search string **"Children"**.

[5] = Must See [4] = Will Enjoy It [3] = It's OK [2] = Fairly Bad [1] = Awful

PREDICTED RATING	YOUR RATING	GENRE	TITLE	REVIEWS	TO SEE LIST (3 movies)
★★★	? unseen	Drama	Children Are Watching us, The (Bambini ci guardano, I) (1942)		
★★★★½	? unseen	Drama	Children of Heaven, The (Bacheha-Ye Aseman) (1997)		
★★★★★	? unseen	Drama, Romance	Children of Paradise (Les enfants du paradis) (1945)		
★★★★	? unseen	Drama	Children of a Lesser God (1986)		
★★½	? unseen	Horror, Thriller	Children of the Corn (1984)		
★★	? unseen	Horror	Children of the Corn II: The Final Sacrifice (1993)		
★★	? unseen	Horror	Children of the Corn III (1994)		
★★½	? unseen	Horror	Children of the Corn IV: The Gathering (1996)		
★★★½	? unseen	Horror, Sci-Fi, Thriller	Children of the Damned (1963)		
★★★★	? unseen	Comedy	Children of the Revolution (1996)		
★★★★½	? unseen	Adventure, Sci-Fi	City of Lost Children, The (1995)		

Submit ratings

MovieLens | "Angela's Ashes" 2000 Review - Microsoft Internet Explorer

http://movielens.umn.edu/theater_reviews/3179.html

"ANGELA'S ASHES"
★★★★★

Drama
145 minutes | Rated: PG-13
Released by Paramount Pictures

Directed by Alan PArker
Written by Frank McCourt
Produced by David Brown

Starring Emily Watson & Robert Carlyle

All reviews are based on a 5 star rating.

ANGELA'S ASHES
CHRISTMAS

During this summer, we have been swept up as a nation by the likes of such reality shows as 'Survivor' and 'Big Brother'. Not many of us want to admit that we are turned on by the lives and perils of fellow human beings, but the truth is we are. Who knows why. Perhaps it is because it helps us as individuals, realize that we are not the only ones to have problems, or maybe it is because we are looking to identify with other people that may be like us in certain ways. I'm not going to pretend to know the reasons, but one thing that I do know is that I am affected in the same way as the countless others in this country by the phenomenon. Not only do I enjoy a good argument on 'Survivor' or a sexual discussion on 'Big Brother', but I also like a film that is based on a true story of a 'regular joe'. 'Angela's Ashes' was just that kind of film. It opened a window into the life of a young Irish boy named Frank McCourt.

Frank McCourt has become one of our nation's most read authors in the last few years, and 'Angela's Ashes' is the story of his life. It is the book that began his writing career and the film that brought that book to life. Frank, his parents (Emily Watson and Robert Carlyle) and siblings lived out a hard existence in a poor Irish community named Limerick, a place where the rain never seems to stop. We watch as Frank loses some of his siblings to disease and hunger, and his mother Angela deals with the situation of a husband who is more concerned about drinking booze than providing for his large family. Frank eventually grows up to be an intelligent and caring individual, and this story is his way of saying thank you to his mother and forgiving his father.

Alan Parker, who last directed the epic

How It Works

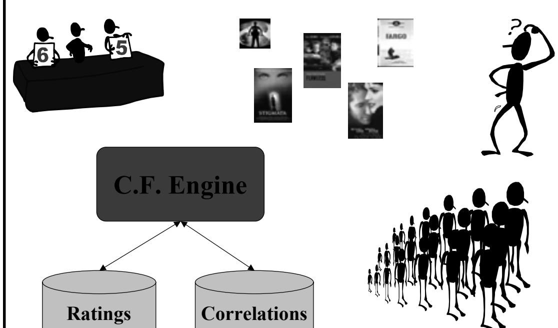

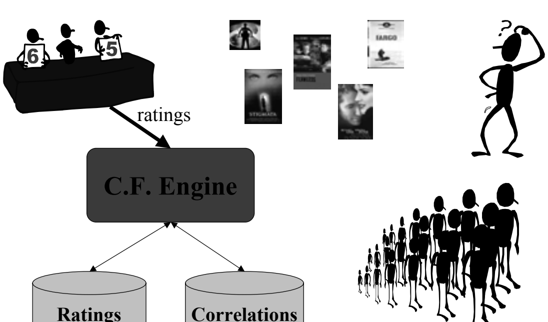

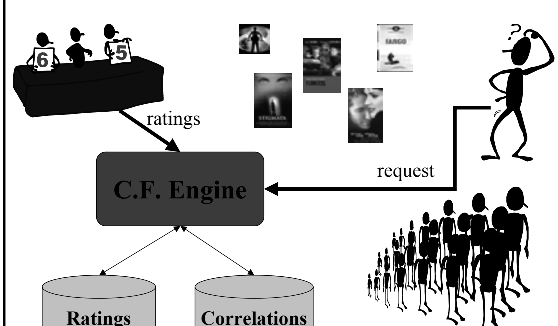

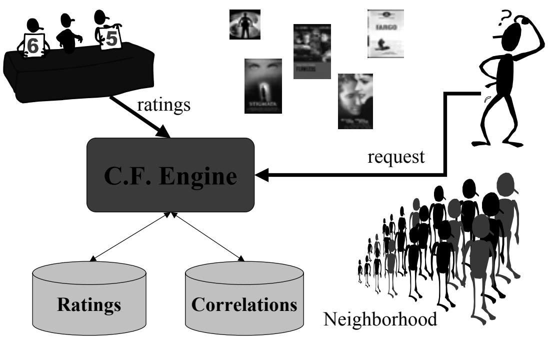

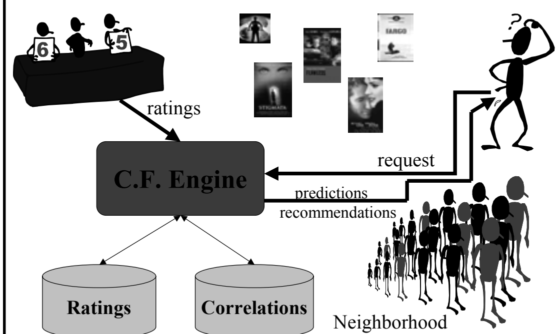

GroupLens Model of Information Filtering

- *Users* rate *Items.*
- Users are *correlated* with other users.
- *Predictions* made for an item's value to a particular user by combining ratings of highly correlated users who rated it.
- *Recommendations* for items for a particular user by identifying popular items among correlated users.

Understanding the Computation

	Hoop Dreams	Star Wars	Pretty Woman	Titanic	Blimp	Rocky XV
Joe	D	A	B	D	?	?
John	A	F	D		F	
Susan	A	A	A	A	A	A
Pat	D	A		C		
Jean	A	C	A	C		A
Ben	F	A				F
Nathan	D		A		A	

Understanding the Computation

	Hoop Dreams	Star Wars	Pretty Woman	Titanic	Blimp	Rocky XV
Joe	D	A	B	D	?	?
John	A	F	D		F	
Susan	A	A	A	A	A	A
Pat	D	A		C		
Jean	A	C	A	C		A
Ben	F	A				F
Nathan	D		A		A	

Understanding the Computation

	Hoop Dreams	Star Wars	Pretty Woman	Titanic	Blimp	Rocky XV
Joe	D	A	B	D	?	?
John	A	F	D		F	
Susan	A	A	A	A	A	A
Pat	D	A		C		
Jean	A	C	A	C		A
Ben	F	A				F
Nathan	D		A		A	

Understanding the Computation

	Hoop Dreams	Star Wars	Pretty Woman	Titanic	Blimp	Rocky XV
Joe	D	A	B	D	?	?
John	A	F	D		F	
Susan	A	A	A	A	A	A
Pat	D	A		C		
Jean	A	C	A	C		A
Ben	F	A				F
Nathan	D		A		A	

Understanding the Computation

	Hoop Dreams	Star Wars	Pretty Woman	Titanic	Blimp	Rocky XV
Joe	D	A	B	D	?	?
John	A	F	D		F	
Susan	A	A	A	A	A	A
Pat	D	A		C		
Jean	A	C	A	C		A
Ben	F	A				F
Nathan	D		A		A	

Understanding the Computation

	Hoop Dreams	Star Wars	Pretty Woman	Titanic	Blimp	Rocky XV
Joe	D	A	B	D	?	?
John	A	F	D		F	
Susan	A	A	A	A	A	A
Pat	D	A		C		
Jean	A	C	A	C		A
Ben	F	A				F
Nathan	D		A		A	

Understanding the Computation

	Hoop Dreams	Star Wars	Pretty Woman	Titanic	Blimp	Rocky XV
Joe	D	A	B	D	?	?
John	A	F	D		F	
Susan	A	A	A	A	A	A
Pat	D	A		C		
Jean	A	C	A	C		A
Ben	F	A				F
Nathan	D		A		A	

Recent and Current Research

Accuracy, Scale, and Sparsity

- Algorithm Performance and Metrics
- Filterbots
- Dimensionality Reduction Algorithms

Interfaces and User Experience

- Explaining Recommendations
- Ephemeral Recommendations
- PolyLens: Multi-User Recommendations
- MetaLens: Multi-Source Recommendations

Other Research (not covered in this talk)

- Distributed Recommenders (Sarwar)
- E-Commerce Recommender Applications (Schafer)
- User and Usage Studies

Accuracy, Scale, and Sparsity

Algorithm Performance and Metrics

(Herlocker et al., SIGIR '99; …)

Breese studied recommender algorithms
- k-nearest neighbors as good as any

We looked at relevant tuning parameters
- limiting neighborhood size important
- normalization of ratings very important
- most other parameters unimportant
  - ➔correlation measure, weightings

Which Metrics?

Many metrics used in published work
- Error metrics (MAE, MSE, RMSE)
- Decision-support metrics (ROC, errors)
- IR metrics (version of precision, recall)

We found that there are only two types
- Rank-sensitive, value-sensitive

All seem to work equally well and nearly identically, within type

Filterbots

The Inspiration:

Need selfless, consistent raters

Humans?

No: ratings robots.

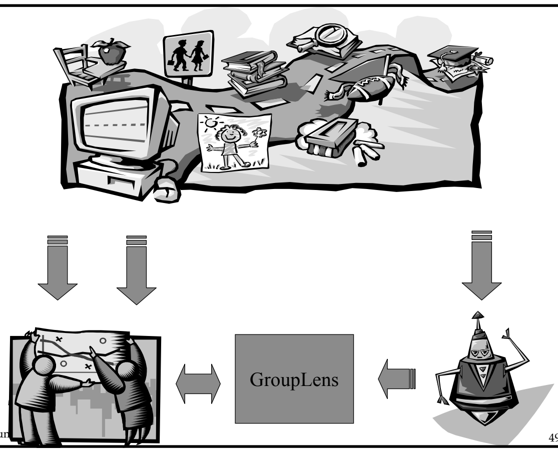

Filterbot Studies

Usenet Filterbots *(Sarwar et al. CSCW 98)*

- ◆Simple, non-personalized filterbots
  - ➔Spelling, Length, % new text
- ◆ One filterbot at a time

MovieLens Filterbots *(Good et al. AAAI '99)*

- ◆Personalized filterbots
  - ➔Learned from genre, cast, descriptions
- ◆Many filterbots per person

Lessons Learned

Even simple filterbots added value

C.F. best way to create a personal combination of filterbots

Filterbots better than a small community of users

Filterbots + users in CF better than either alone

Advantages of the FilterBot model

Combines best of agents *and* humans

- ◆ agents rate frequently, quickly, consistently
- ◆ humans add subjective taste and quality

Framework pulls out the best of each

- ◆ use only the bots that work; ignore the others
- ◆ use only the people who agree; ignore the others
- ◆ balance people and bots based on available ratings and agreement

Risks of Filterbots

- What if no humans read certain articles?
  - ◆"voluntary" censorship or quality control?
- What about rogue filterbots?
- What if people "prefer" filterbots to humans?

New Algorithms

(Sarwar et al., EC 00 & WebKDD 01)

Significant challenges

Scale

- ◆Number of users
- ◆Number of items

Sparsity

- ◆Small percentage of items experienced
- ◆Hard to find overlap with other users

Example Challenge

Synonymy

- Similar products treated differently
- Increases sparsity, loss of transitivity
- Results in poor quality

Example

- C_1 rates *recycled letter pads* High
- C_2 rates *recycled memo pads* High

➔Both of them like *Recycled office products*

Idea: Dimensionality Reduction

Latent Semantic Indexing

- Used by the IR community for document similarity
- Works well with similar vector space model
- Uses Singular Value Decomposition (SVD)

Main Idea

- Find (latent) "taste space"
- Represent users and items as points (vectors) in taste space
- Reduced space is dense and less-noisy

SVD: Mathematical Background

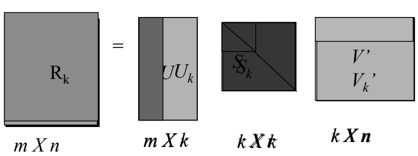

The reconstructed matrix $R_k = U_k.S_k.V_k'$ is the closest *rank-k* matrix to the original matrix *R*.

SVD for Collaborative Filtering

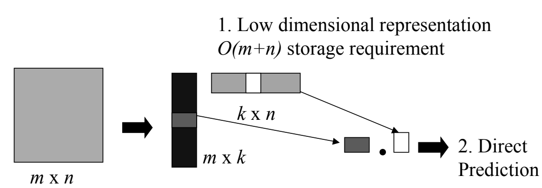

Experimental Setup

Data Sets

- MovieLens data (www.movielens.umn.edu)
  - ➔943 users, 1,682 items
  - ➔100,000 ratings on 1-5 Likert scale
  - ➔Used for *prediction* and *neighborhood* experiments
- E-commerce data
  - ➔6,502 users, 23,554 items
  - ➔97,045 purchases
  - ➔Used for *neighborhood* experiment

Results: Prediction

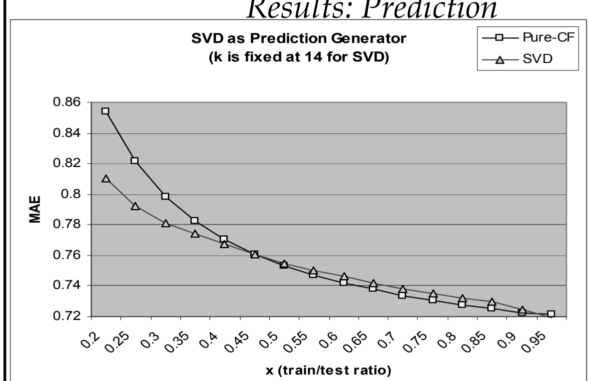

SVD Conclusions

Successful and promising approach

Several obstacles to overcome
- Incremental update
- Efficient "top-n" recommendations

Exploring SVD-based and other new algorithmic approaches

Interfaces and User Experience

Explaining Recommendations

(Herlocker et al. CSCW 2000)

Challenge: Belief
- *Why* should users believe the recommendations?
- *When* should users believe the recommendations?

Approach
- Explain recommendations
  - ➔Reveal data, process
  - ➔Corroborating data, track record

Two Studies

Pilot study of explanation feature
- Users liked explain
- Unclear whether they become more effective decision-makers

Comprehensive study of different explanation approaches
- Wide variation of effectiveness
- Some explanations hurt decision-making

Most Compelling Interfaces

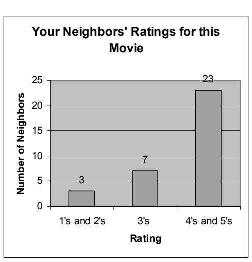

- *Simple* visual representations of neighbors ratings
- Statement of strong previous performance "MovieLens has predicted correctly 80% of the time for you"

Less Compelling Interfaces

- Anything with even minimal complexity
  - More than two dimensions
- Any use of statistical terminology
  - Correlation, variance, etc.

Addressing Ephemeral Needs
(Herlocker)

What is an ephemeral interest need?
- Immediate, temporary, dynamic

Current systems don't support this
- Assume interests will remain relatively constant
- Recommendations are relative to all your interests as a whole

One Simple Approach

User submits "theme" query
- Theme contains examples of items similar to those desired by the user

Set of potentially similar items identified
- Using item-to-item correlation in ratings space

Potentially similar items ranked based on traditional ACF predictions

Theme Creation

Theme Selection

Query Results

Results of Theme Query Study

Users were very positive about the theme query interface

Relevance of results were dependent on the "support threshold"
- Low support threshold => fewer relevant results

When results were relevant, users were positive overall

Even the users in the low support threshold groups indicated they would like to have the interface added to MovieLens

PolyLens: A Group Recommender

(O'Connor et al. Interact 2001)

Challenge: People watch movies together

Solution: A recommender for groups

Issues

- ◆Group formation, rules, composition
- ◆Recommender algorithm for groups
- ◆User interface

Goals

Explore group recommender design space

See if users would want and use a group recommender, at least for movies

Study behavior changes in group members

- ◆group vs. other users
- ◆new users via groups vs. other new users

Learn lessons about group recommenders

Design Issues

Characteristics of groups

- ◆public or private
- ◆many or few
- ◆permanent or ephemeral

Formation and evolution of groups

- ◆joining policy
- ◆administration and rights

Design Issues

What is a group recommendation?

- ◆group user vs. combined individuals
- ◆social good functions

Privacy and interface issues

- ◆control over joining groups
- ◆withholding and recommendations
- ◆balancing between info overload and support

PolyLens

Design choices

- ◆ private, small, administered, invited groups
- ◆ combine individual recs with minimum misery
- ◆ high-information interface with opt-out

Group: Dantest Back To Individual Recommendations

TITLE	GENRE	REVIEWS	GROUP	YOUR	cosley@cs.umn.edu	cosley@quasar
Pixote (1981)	Drama		★★★★★	★★★★★	★★★★★	★★★★★
Wrong Trousers, The (1993)	Animation, Comedy		★★★★★	★★★★★	★★★★★	★★★★★
After Life (1998)	Drama		★★★★½	★★★★½	★★★★½	★★★★★
King of Masks, The (Bian Lian) (1996)	Drama		★★★★½	★★★★★	★★★★½	★★★★★

External invitations added by popular demand

Field Testing PolyLens

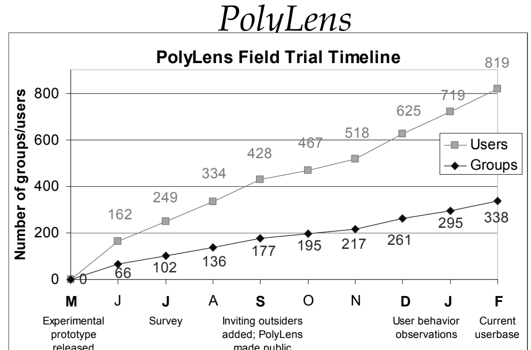

Survey and Usage Results

Satisfaction (95% like, 77% more useful)

Privacy not an issue (94% see, 93% share)
- individual recommendations “essential”

Groups reflect “real life” groups

New users via groups stayed 1.5x as often
- group vs. other users a wash

Many stillborn groups

Field Test Results and Lessons

Users like and use group recommenders
- groups have value for all members
- groups can help with outreach to new members

Users trade privacy for utility

Groups are both permanent and ephemeral

Users must be able to find each other

MetaLens: A Meta-Recommender
(Schafer)

Integrating multiple sources of information into a single recommendation list

What is the problem?

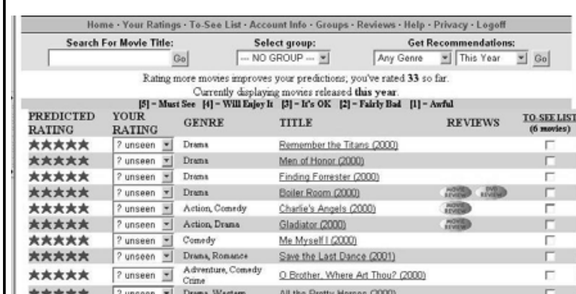

Home · Your Ratings · To-See List · Account Info · Groups · Reviews · Help · Privacy · Logoff

Search For Movie Title: [ ] Go | Select group: --- NO GROUP --- | Get Recommendations: Any Genre, This Year, Go

Rating more movies improves your predictions; you've rated **33** so far.
Currently displaying movies released **this year**.

[5] = Must See [4] = Will Enjoy It [3] = It's OK [2] = Fairly Bad [1] = Awful

PREDICTED RATING	YOUR RATING	GENRE	TITLE	REVIEWS	TO SEE LIST (6 movies)
★★★★★	? unseen	Drama	Remember the Titans (2000)		
★★★★★	? unseen	Drama	Men of Honor (2000)		
★★★★★	? unseen	Drama	Finding Forrester (2000)		
★★★★★	? unseen	Drama	Boiler Room (2000)		
★★★★★	? unseen	Action, Comedy	Charlie's Angels (2000)		
★★★★★	? unseen	Action, Drama	Gladiator (2000)		
★★★★★	? unseen	Comedy	Me Myself I (2000)		
★★★★★	? unseen	Drama, Romance	Save the Last Dance (2001)		
★★★★★	? unseen	Adventure, Comedy Crime	O Brother, Where Art Thou? (2000)		
★★★★★	? unseen	Drama, Western	All the Pretty Horses (2000)		

One solution

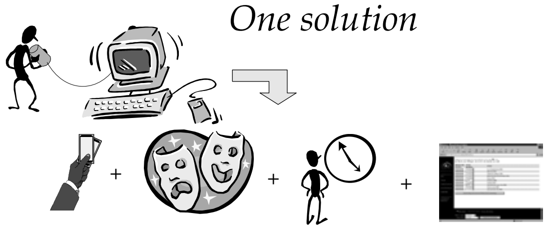

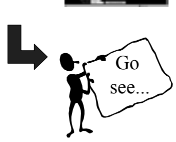

Meta-recommendation system
- MetaLens

Sources of Data

- Genre
- MPAA ratings
- Film length
- Objectionable Content
- Distributor
- Release Date
- Start/End Time
- Critical Reviews
- Average User Rating
- User’s personalized MovieLens prediction
- Distance to the Theater
- Special Accomodations
- Discounted Shows

Meta-Lens Score	Movie	Theater	Start Time	End Time	Movie- Lens Prediction	Genre	MPAA Rating	Run-time	Avg. User Rating
58.2	Best in Show (2000)	Riverview Theater	Unavailable	Unavailable	★★★★½	Comedy.	PG-13	90	4.0
53.6	Crouching Tiger, Hidden Dragon (2000)	Landmark Lagoon Cinema	7:10	9:10	★★★★★	Romance.	PG-13	120	4.5
53.6	You Can Count on Me (2000)	University Film Society	Unavailable	Unavailable	★★★★½	Drama.	R	109	4.4
51.5	The Family Man (2000)	Riverview Theater	Unavailable	Unavailable	★★★★½	Comedy and Romance.	PG-13	125	3.9
51.2	Enemy at the Gates (2001)	St. Anthony Main	7:20	9:46	★★★	Drama and Romance.	R	146	3.8
48.1	Requiem for a Dream (2000)	University Film Society	7:15	8:57	★★★★½	Drama.	NR	102	4.3
47.7	Billy Elliot (2000)	Historic Suburban World Cinema Grill Theatre	7:00	8:30	★★★★½	Drama and Musical/Performing Arts.	R	90	4.3
47.2	Pollock (2000)	Landmark Lagoon Cinema	7:15	9:12	★★★★	Drama.	R	117	3.7
46.9	The House of Mirth (2000)	Landmark Lagoon Cinema	7:00	9:20	★★★★	Drama and Romance.	PG	140	3.6
46.3	15 Minutes (2001)	St. Anthony Main	3:30	5:29	★★★★★	Thriller and Comedy.	R	119	3.4

What Have We Learned?

- Meta-recommenders can be built.
- Anecdotally, users like them.
- Some users make heavy use of them, and heavy users are most likely to make some use of them.

Conclusions and Future Work

Conclusions

Collaborative filtering works!

Lots of important issues:

- ◆ Algorithms
- ◆ Interfaces and User Experience
- ◆ Privacy
- ◆ Applications

Future Work

Better integration of collaborative and content filtering

Better support for community

Better understanding of user rewards, social role of recommenders

CF Under Diminishing Returns

Original goal of CF was to help people sift through the junk to find the good stuff.

Today, there may be so much good stuff that you need to sift even more.

Certain types of content yield diminishing returns, even with high quality

Portfolios of Content

What if my recommender knows which articles I've read, and can identify articles by topic?

What if it sees that I experience marginal returns from reading similar articles on a topic?

Could we downgrade some articles based on "lack of new content?" Could we discover which articles using collaborative filtering?

Temporal Collaborative Filtering

Today's CF systems may expire or degrade ratings, but do little to detect or predict changes in preference.

Ripe area with lots of commercial applications …

Wine for the Time

Seasonal taste – can we detect that a particular customer shifts wine tastes during hot and cold weather? Can we learn either the content, or separate profiles, reflecting these different tastes?

Evolving taste – can we help a wine newcomer build her palate? Could we identify wines that take her a step or two beyond her current ones? Can we do so by augmenting regular collaborative filtering with temporal models?

Acknowledgements

- This work is being supported by grants from the National Science Foundation, and by grants from Net Perceptions, Inc.
- Many people have contributed ideas, time, and energy to this project.

The GroupLens Research Project:

Collaborative Filtering Recommender Systems

Joseph A. Konstan
University of Minnesota
konstan@cs.umn.edu
http://www.grouplens.org