Amazon Knows What You Like
An Introduction to Recommender Systems

Huizhi (Elly) Liang

1 June, 2012

THE AUSTRALIAN NATIONAL UNIVERSITY

Agenda

1. Introduction
2. Recommendation Approaches
3. Evaluation Approaches
4. Recommender Systems in Web 2.0
5. Conclusion

Introduction

Information Overload

- Catalog
  - Yahoo
- Search Engine
  - Google, Bing
- Recommender Systems
  - Amazon, Netflix
  - Google personalized news

Search Engine vs. Recommender System

- User will try search engine if
  - they have specific needs
  - they can use keywords to describe needs
- User will try recommender system if
  - they do not know what they want now
  - they can not use keywords to describe needs

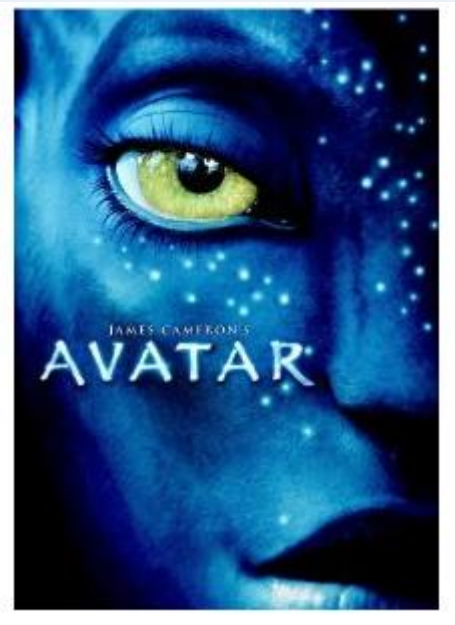

Customer Reviews

5 star	1,340	
4 star	289	
3 star	242	
2 star	120	
1 star	381	

2,372 reviews

3.9 out of 5 stars

Write a customer review

Customers Who Bought This Item Also Bought

Page 1 of 15

Tron: Legacy
Jeff Bridges
(550)
DVD
$13.61

Inception
Leonardo DiCaprio
(988)
DVD
$6.25

The Dark Knight (Single-Disc Widescreen ...
Christian Bale
(1,545)
DVD
$5.99

Star Trek (Single-Disc Edition)
Chris Pine
(1,267)
DVD
$7.30

300 (Single-Disc Widescreen Edition)
Gerard Butler
(1,262)
DVD
$6.55

Transformers Trilogy (Transformers ...
Shia Labeouf
(49)
DVD
$25.58

Iron Man 2 (Single-Disc Edition)
Robert Downey Jr.
(428)
DVD
$9.96

Are they effective ?

(Celma & Lamere, ISMIR 2007)

- Netflix
  - 2/3 rated movies are from recommendation
- Google News
  - 38% more click-through are due to recommendation
- Amazon
  - 35% sales are from recommendation

What can be recommended?

- Advertising messages
- Investment choices
- Restaurants
- Cafes
- Music tracks
- Movies
- TV programs
- Books
- Cloths
- Supermarket goods
- Tags
- News articles
- Online mates (Dating services)
- Future friends (Social network sites)
- Courses in e-learning
- Drug components
- Research papers
- Citations
- Code modules
- Programmers

Recommendation Approaches

Recommender systems

- Recommendation tasks
  - Top N Recommendation *(Precision, Recall, F1)*
  - Rating Prediction *(Mean Absolute Error, Root Mean Squared Error)*
- Recommendation approaches
  - Content based
    - *Term vector model*
    - *Latent Dirichlet Allocation (LDA)*
  - Collaborative Filtering (CF)
    - Memory based CF: User-KNN & Item-KNN
    - Model based CF: Matrix Factorization techniques
  - Hybrid

Content-based Method

- Text information
  - Web pages, blogs, micro-blogs, reviews, descriptions, comments, title, abstract, keywords
- Multi-media information
  - Audio, Video, Image
- Main Idea
  - Content matching between the topics of item and user
- Steps

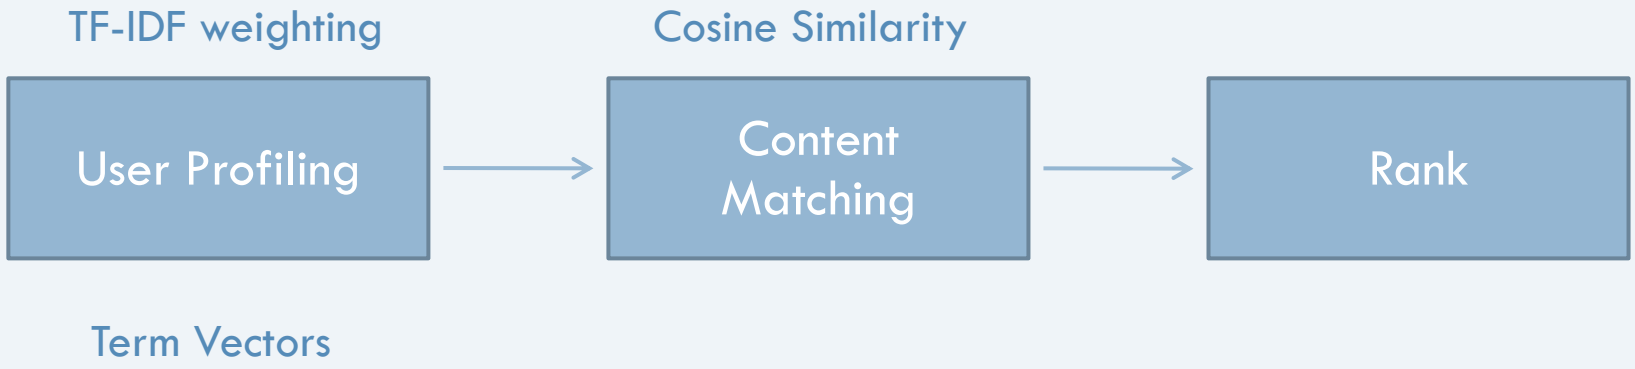

Example

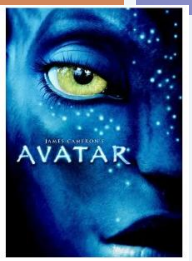

A reluctant hero. An epic journey. A Choice between the life he left behind and the incredible new world he's learned to call home. James Cameron's Avatar – the greatest adventure of all time.

Term Vectors

hero	…	time
1	…	1

dream	…	crime
2	…	1

Cosine Similarity

0

Recommendation

?

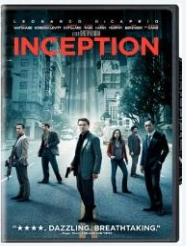

Acclaimed filmmaker Christopher Nolan directs an international cast in this sci-fictioner that travels around the globe and into the world of dreams. Dom Cobb (Leonardo DiCaprio) is the best there is at extraction: stealing valuable secrets inside the subconscious during the mind's vulnerable dream state. His skill has made him a coveted player in industrial espionage but also has made him a fugitive and cost him dearly. Now he may get a second chance if he can do the impossible: inception, planting an idea rather than stealing one. If they succeed, Cobb and his team could pull off the perfect crime. But no planning or expertise can prepare them for a dangerous enemy that seems to predict their every move. An enemy only Cobb could have seen coming.

Content-based Method

- Strength
  - Easy for text
  - Efficient & Scalable
  - Only rely on target user's history or information needs
- Weakness
  - Difficult for multimedia & Noise
  - Always recommend similar items &Poor Novelty
  - Attacks

Practice at home

- Recommender System based on tweets
  - Recommend tweets, users, news, topics
- Tips:
  - Twitter API
  - News feeds

"PM JG on SKY confirms the carbon tax will only apply to 500 companies in Australia". @Latika Bourke

✓ "The more we here about this #carbon tax the more it sounds like a camel - a horse designed by a committee. #AUSPOL "@demonspofforth

✓ "#Carbon price revealed: Finally, here's what it will cost #auspol Latest Business & Austr.. http://bit.ly/ngTS6V " @1southerncross

K-Nearest Neighborhood

- User-based Method (1994, GroupLens)
  - Rating information
  - Many people liked "Avatar"
  - Can you predict how much I like it?

- Main idea
  - Pick 20-50 people that has similar taste with me
  - How much I like depends on how much they like
- Steps

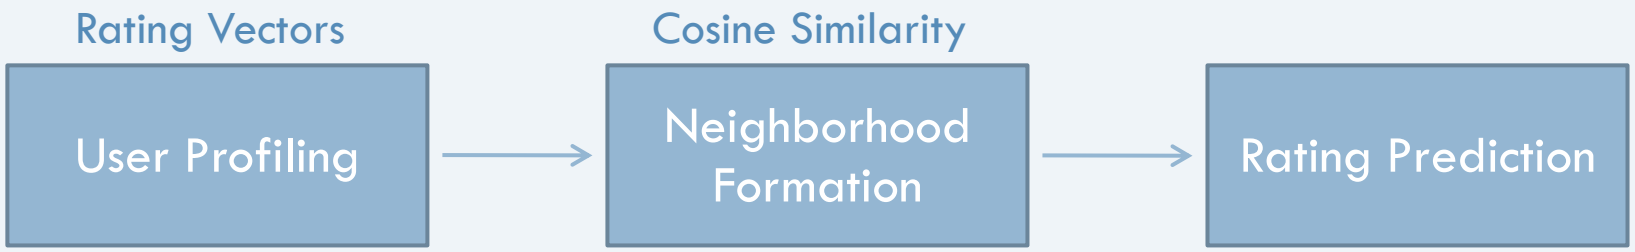

Example

- Step 1
  - Rating vector
- Step 2
  - Similarity of Elly and all other users
- Step 3
  - Select 2-nearest users: user 4&6
- Step 4
  - Average ratings of all neighbor users
  - **?**

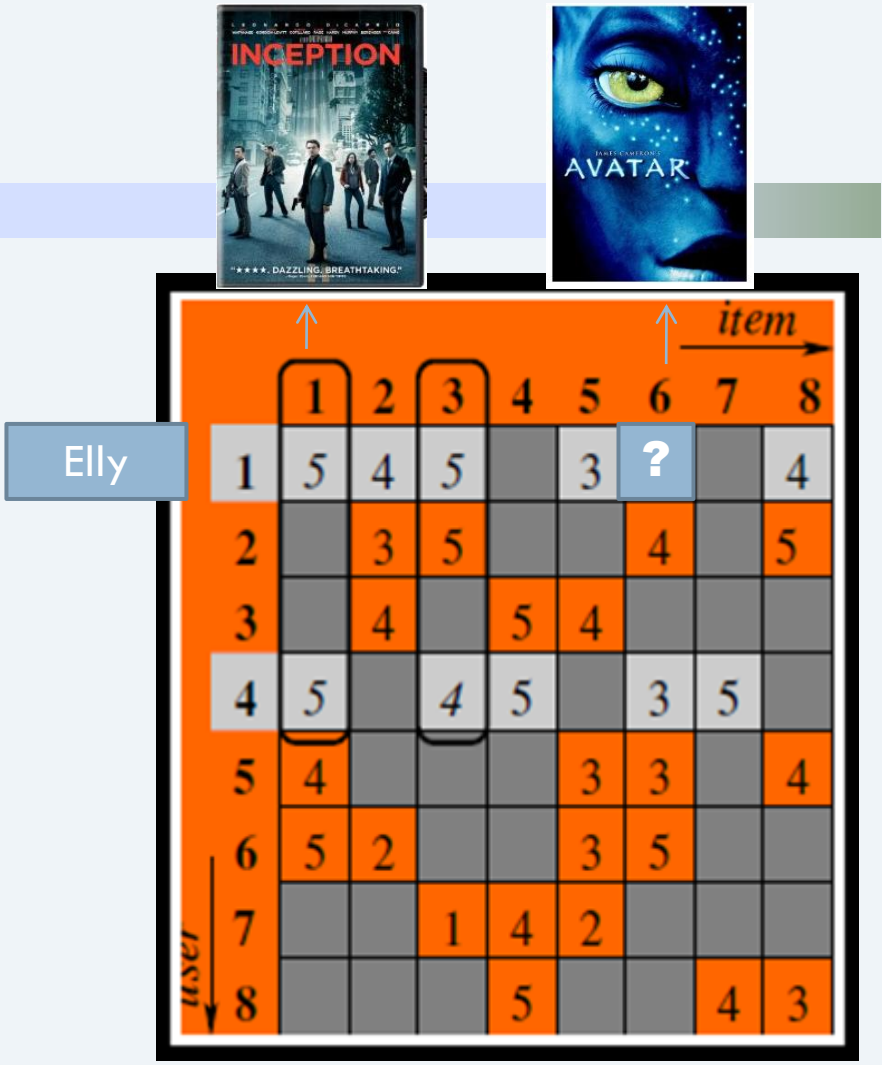

K-Nearest Neighborhood

- Item-based CF (2001, deployed at Amazon)
  - I have watched so many good & bad movies
  - Would you recommend me watching “Avatar”
- The idea
  - pick from my previous list 20-50 movies that share similar audience with “Avatar”
  - how much I will like depend on how much I liked those early movies
  - *People who have watched those movies also liked this movie*

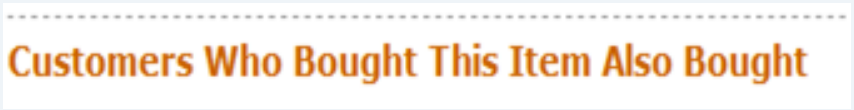

Example

- Step 1
  - Rating vector
- Step 2
  - Similarity of Avatar and all other items
- Step 3
  - Select 2-nearest items: item 1&5
- Step 4
  - Average ratings of all similar items
  - **?**

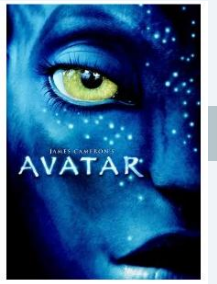

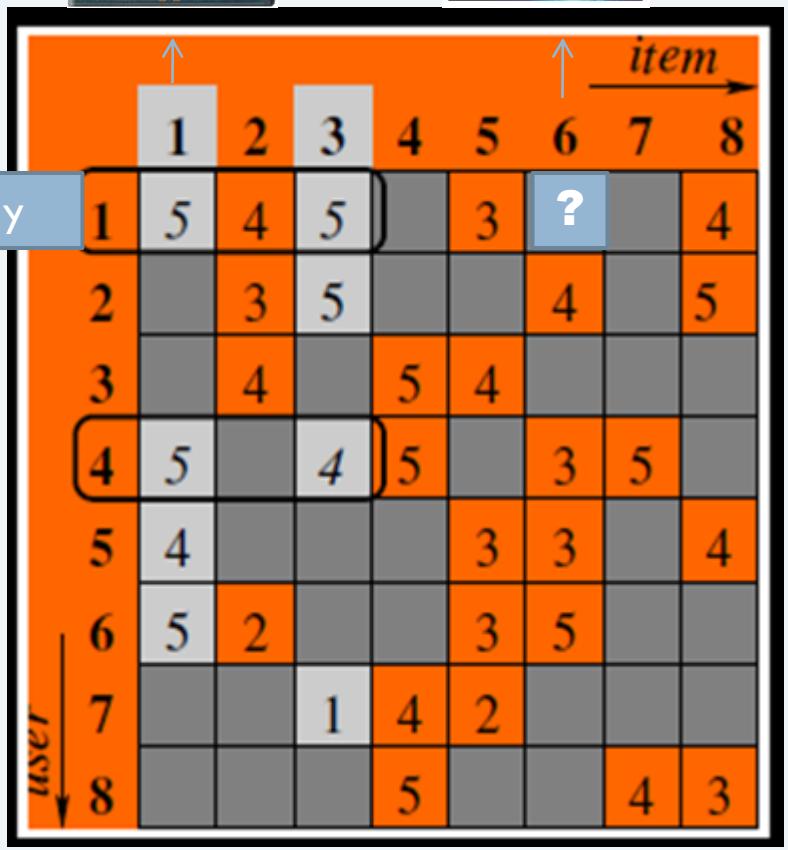

K-Nearest Neighborhood

- Strength
  - Content free
  - Good Novelty
  - Easy to understand
- Weakness
  - Cold start
    - New user
    - New item
  - Scalability

Matrix Factorization

Collaborative filtering (3)

~ [0.1 0.3 0.2 0.9 0.5 0.4 0.7 0.3 0.8 1.5]

- Matrix Factorization (2006, Netflix challgence)
  - You many have watched thousands of movies
  - But perhaps I can tell these movies belong to **10** groups, like Action, Sci-Fi, Animation, etc,…
  - So **10** numbers are enough to describe your taste
  - Likewise, "Titanic" has been watched by millions people, but perhaps …10 numbers are enough to describe its features
  - Magic: these hidden aspects can be discovered automatically by Matrix Factorization!

Matrix Factorization

Singular Value Decomposition (SVD)

- Handy mathematical technique that has application to many problems
- Given any $m \times n$ matrix **R**, find matrices **U**, **I**, and **V** such that

  $\mathbf{I} = \mathbf{U}\mathbf{I}\mathbf{V}^T$

  U is $m \times r$ and orthonormal

  I is $r \times r$ and diagonal

  V is $n \times r$ and orthonormal
- Discard the smallest I values to get $R_{m,k}$ with k << r
  - *$R_{m,k}$ is an k-approximation of R based on k most important latent feature*
  - *Recommendation will be given based on R_k*

$$\begin{pmatrix} R \end{pmatrix} = \left(\begin{array}{c} \mathbf{U} \\ \hline \end{array}\right) \begin{pmatrix} I_1 & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & I_r \end{pmatrix} \begin{pmatrix} \mathbf{V} \end{pmatrix}^T$$

Example

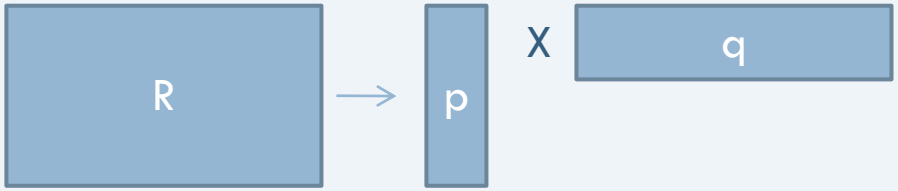

- Step 1
  - User's rating vector
- Step 2
  - Predictor

Mean Rating
User Bias
Item Bias

$$\hat{r}_{ui} = \mu + b_u + b_i + p_u^T q_i$$

Item factor vector

predicted Rating

User factor vector

- Step 3
  - Minimize cost function using gradient descent

Real Ratina

$$\sum_{(u,i)\in K^+ \cup K^-} (r_{ui} - \hat{r}_{ui}) + \lambda(||p_u||^2 + ||q_i||^2)$$

Positive Sample

Negative Sample

Hybrid approaches

- Any Recommendation approach has pros and cons
  - e.g. CF & CB both suffer from the cold start problem
    - but CF can recommend "outside the box" compared to Content-based approaches
- Hybrid recommender system combines two or more techniques to gain better performance with fewer drawbacks
- Hybrid methods:
  - Weighted: scores of several recommenders are combined together
  - Switching: switch between recommenders according to the current situation
  - Mixed: present recommendations that are coming from several recommenders
  - Cascade: One recommender refines the recommendations given by another

Practice at home

- Large rating dataset
  - MovieLens Dataset
  - Netflix Challenge
  - KDD Cup 2011, 2012
- Tips
  - Open source recommender systems
    - Mahout Taste
    - **MyMediaLite Recommender System Library**
- Question
  - How to make recommendation based on implicit ratings?

Evaluations Approaches

Experiment Metrics

- User Satisfaction
- Prediction Accuracy
- Coverage
- Diversity & Novelty
- Trust & Robust
- Real-time

Experiment Metrics

- User Satisfaction
  - Subjective metric
  - Measured by user survey or online experiments

Experiment Metrics

- Prediction Accuracy
  - Measured by offline experiments
    - Training (80%) & Test split (20%)
  - Top-N Recommendation
    - Precision: size (hit set)/size (recommendation list)
    - Recall: size(hit set)/size(test set)
  - Rating Prediction
    - Mean Absolute Error (MAE)
    - Root Mean Absolute Error (RMSE)

Experiment Metrics

- Diversity
  - Measure the ability of recommender system to cover users' different interests.

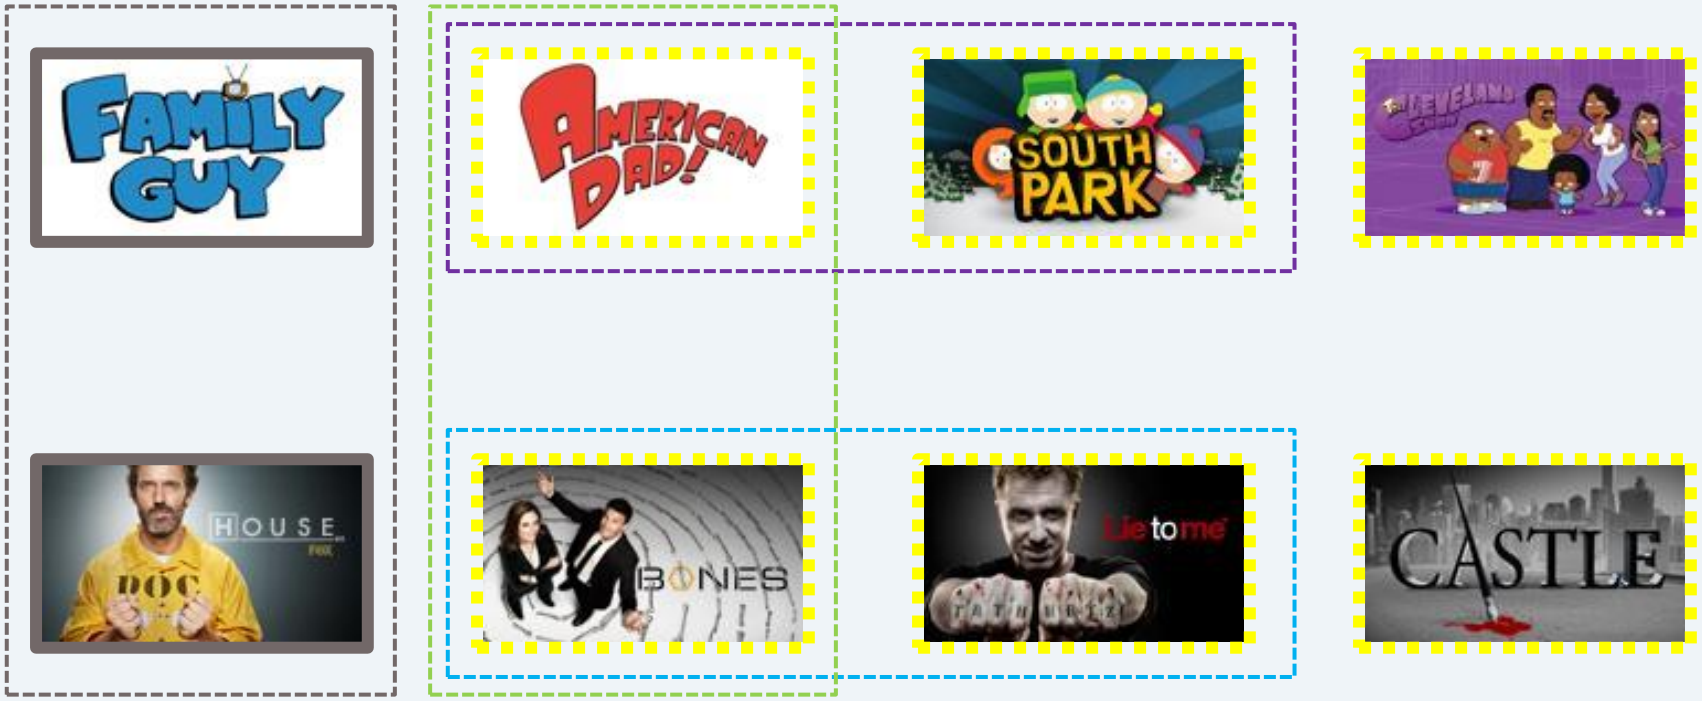

Experiment Metrics

- Coverage
  - Measure the ability of recommender system to recommend long-tail items.
    - Entropy, Gini Index
- Novelty
  - Measure the ability of recommender system to introduce long tail items to users.
  - Music Recommendation and Discovery in the Long Tail [Oscar Celma]
- Robust
  - The ability of recommender system to prevent attack.

Experiment Metrics

- Trust
  - If user trust recommender system, they will interact with it.
  - Ways to improve trust
    - Transparency
    - Social
    - Trust System (Epinion)

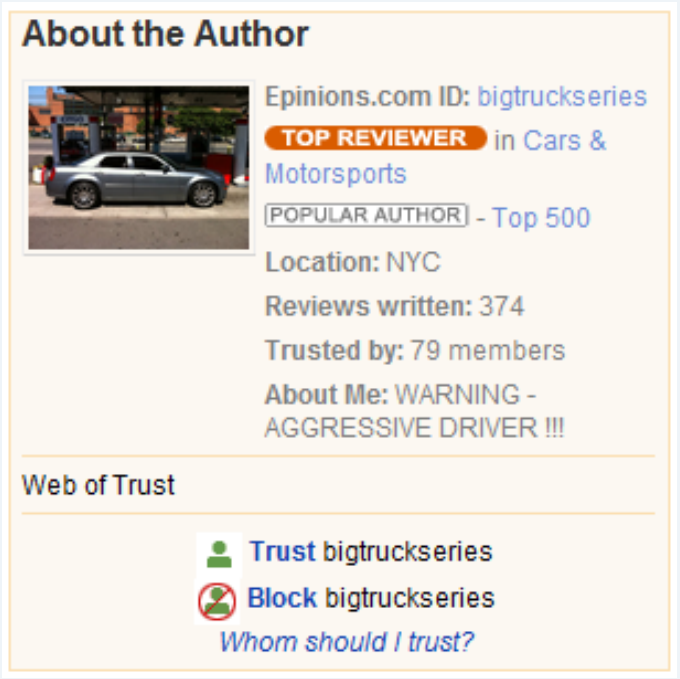

Experiment Metrics

- Real-time
  - Generate new recommendations when user have new behaviors immediately.

How to do trade-off

- Business goal
- Our belief
- Making new algorithms by 3 steps experiments:
  - Offline testing
  - User survey
  - Online testing

Recommender Systems in Web 2.0

Web 2.0

Web 2.0 and Social Media

- **Web 2.0 [Oreiley03]**
  - It's all about people
    - Joining online communities
    - Connecting to each other on social networks
    - Creating content as in wikis and blogs a
    - Annotating content with tags, comments, ratings

- **Social Media**
  - Refers to Web 2.0 sites that allow users to share and interact
  - Characterized by
    - User-centered design
    - User-generated content (e.g., tags)
    - Social networks and online communities

Amazon @ Web 2.0

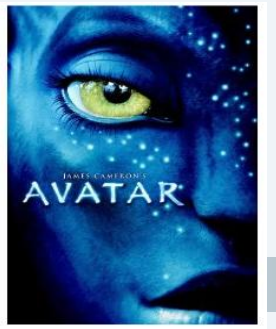

- Tags

Tags Customers Associate with This Product (What's this?)
Click on a tag to find related items, discussions, and people.

avatar (250)	environmentalism (82)	cameron (22)
james cameron (154)	3d (75)	gaia (12)
science fiction (135)	sam worthington (70)	See all 212 tags...
action thriller (84)	zoe saldana (57)	

- Social Networks

29,156,163 people like the **Avatar** movie page on Facebook.

Khoi-Nguyen Tran, Launa Lee and 31,439,318 others like this.

- Reviews

★★★★★ **Great Transfer to Blu-Ray** January 28, 2011

By Thomas Reed

I got this as soon as I found it available on the net. It will not be available commercially for some time and that, of course, means the price is WAY to high for most viewers. I was willing to be taken for a ride but if you do not just have to have it now I would recommend waiting until it is available everywhere.

The video quality is fantastic. I have a Samsung 40" 3D setup and the movie was just beautiful to watch. Not quite the same as IMAX but very close (size of screen being the only difference that I could see). The 3D is, to my eyes, exactly as good as the IMAX on-screen version. I am a huge fan of the movie but believe me I would tell you if the video quality was not great.

I would not hesitate to do the purchase again (even considering the huge rip-off in price at this time) but advise others to consider if you really have to have it right now or can wait awhile.

Great movie, almost unbelievable video transfer quality, and a price that is just not right!

Hope this helps.

21 Comments | **Was this review helpful to you?** Yes No

A question: Can we make use of these information sources?

- Blogs
- Social Media
- Online comments
- Online stores
- Review sites
- Locations
- Mobility

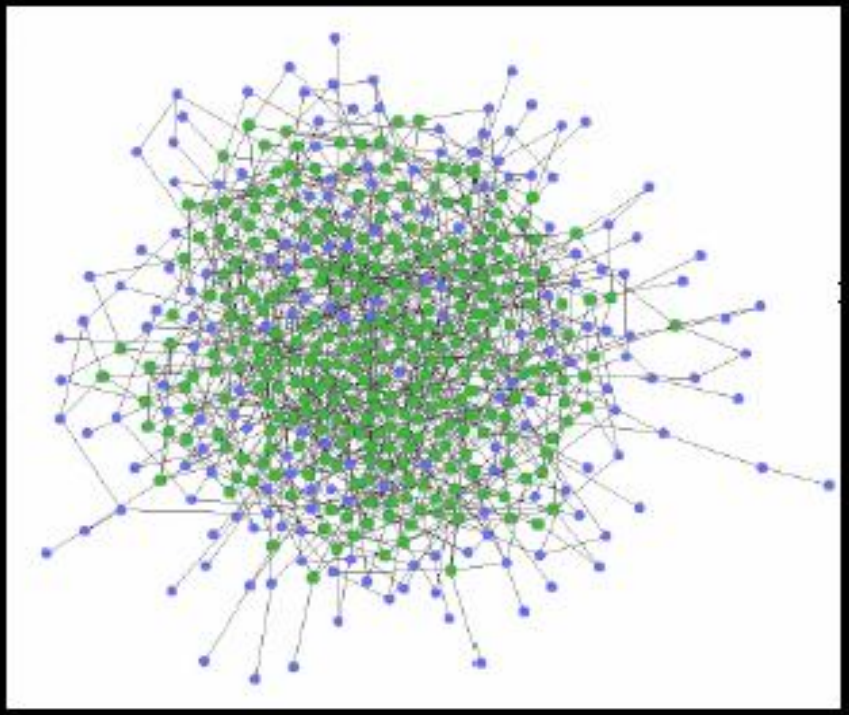

Recommender Systems & Web 2.0

Social Overload

- **Facebook** – largest social network site
  - 600,000,000 users, half log in every day
  - 35,000,000,000 online "friendships"
  - 900,000,000 objects people interact with
  - 30,000,000,000 shared content items / month

- **YouTube** – largest video sharing site
  - 2,000,000,000 views per day
  - 1,000,000 video hours uploaded per month

- **Twitter** – largest microblogging site
  - 200,000,000 users per month
  - 65,00,000 tweets per day (750 per second)
  - 8,000,000 followers of most popular user

Recommender Systems & Web 2.0

Social Recommender Systems

- Recommender Systems that target the social media domain
- Aim at coping with the challenge of social overload by presenting the most attractive and relevant content
- Also aim at increasing adoption and engagement
- Often apply personalization techniques

Ongoing Research Topic

Remain Open Questions

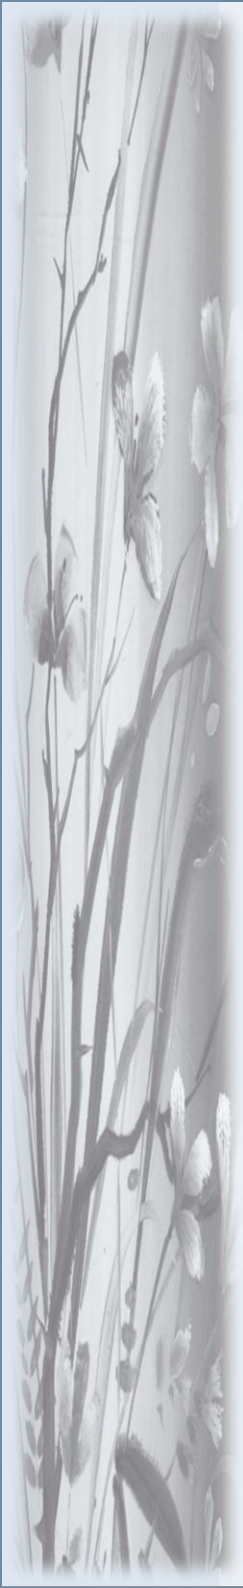

5 Conclusions

Conclusions

- Recommender Systems
  - Content Based
  - Collaborative Filtering
    - KNN based approach
    - Matrix factorization
- Evaluation
- Recommender Systems in Web 2.0

Further Exploration

- Will you implement and publish it as an app?
  - Undergraduate projects
- Will you participate in competitions?
  - KDD Cup 2013?
- Will you do research?
  - Master and PhD projects
  - www.ellyliang.com

Research opportunities with Dr Huizhi (Elly) LIANG

- Improving Recommendation Quality Through Entity Resolution
- Real-time Social Media Recommender Systems
- A Personalized News Recommender System based on Social Media
- A collaborative tagging system based on micro-blogs

Questions & Answers

huizhi.liang@anu.edu.au

www.ellyliang.com