Putting It All in Context

A Summary and Review of Beyond Personalization 2005

Sean M. McNee and Mark Van Setten

A Summary

- Fantastic turnout with 29 people attending!
- 14 full papers and 8 position statements
- This document contains notes from the whole workshop, including the Panel and the Breakout Sessions

Workshop Schedule

- Introductions & Minute Madness
- The Panel
- Breakout Sessions
- Results, Summary & Wrap up

Workshop Topics

- Understanding and trusting recommender systems
- User interfaces for recommender systems
- The future of recommendation algorithms and their metrics
- Social consequences and opportunities of recommender systems

"Reflections on the past, visions of the future"

A Panel on Recommender Systems

Panelists: Robin Burke, Alfred Kobsa, Hugo Liu, Ben Schafer, and special guest, Barry Smyth

Included are the slides from each panelists, some with a short comment (a panelist's reflection)

Robin Burke

- I was pleased to see such a wide variety of research approaches focused on recommender systems. In the 1997 introduction to the Communications of the ACM special issue on recommender systems, Resnick and Varian put forth a information management definition of a recommender system: "In a typical recommender system people provide recommendations as inputs, which the system then aggregates and directs to appropriate recipients.“ It is clear from the workshop that the field has now evolved to encompass a much wider set of system design possibilities. Recommender systems now are defined more by a certain kind of user interaction: namely that they are systems that produce individualized recommendations as output. Recommendations differ from, for example, search results in the complexity of their relationship to the user's stated preferences. We saw among the papers many different ideas on how recommendations can be generated and communicated to users. Given that this is the case, associating the workshop with Intelligent User Interfaces made a lot of sense -- we are all looking for intelligent ways to connect users to interesting and useful objects.

Knowledge sources for recommendation

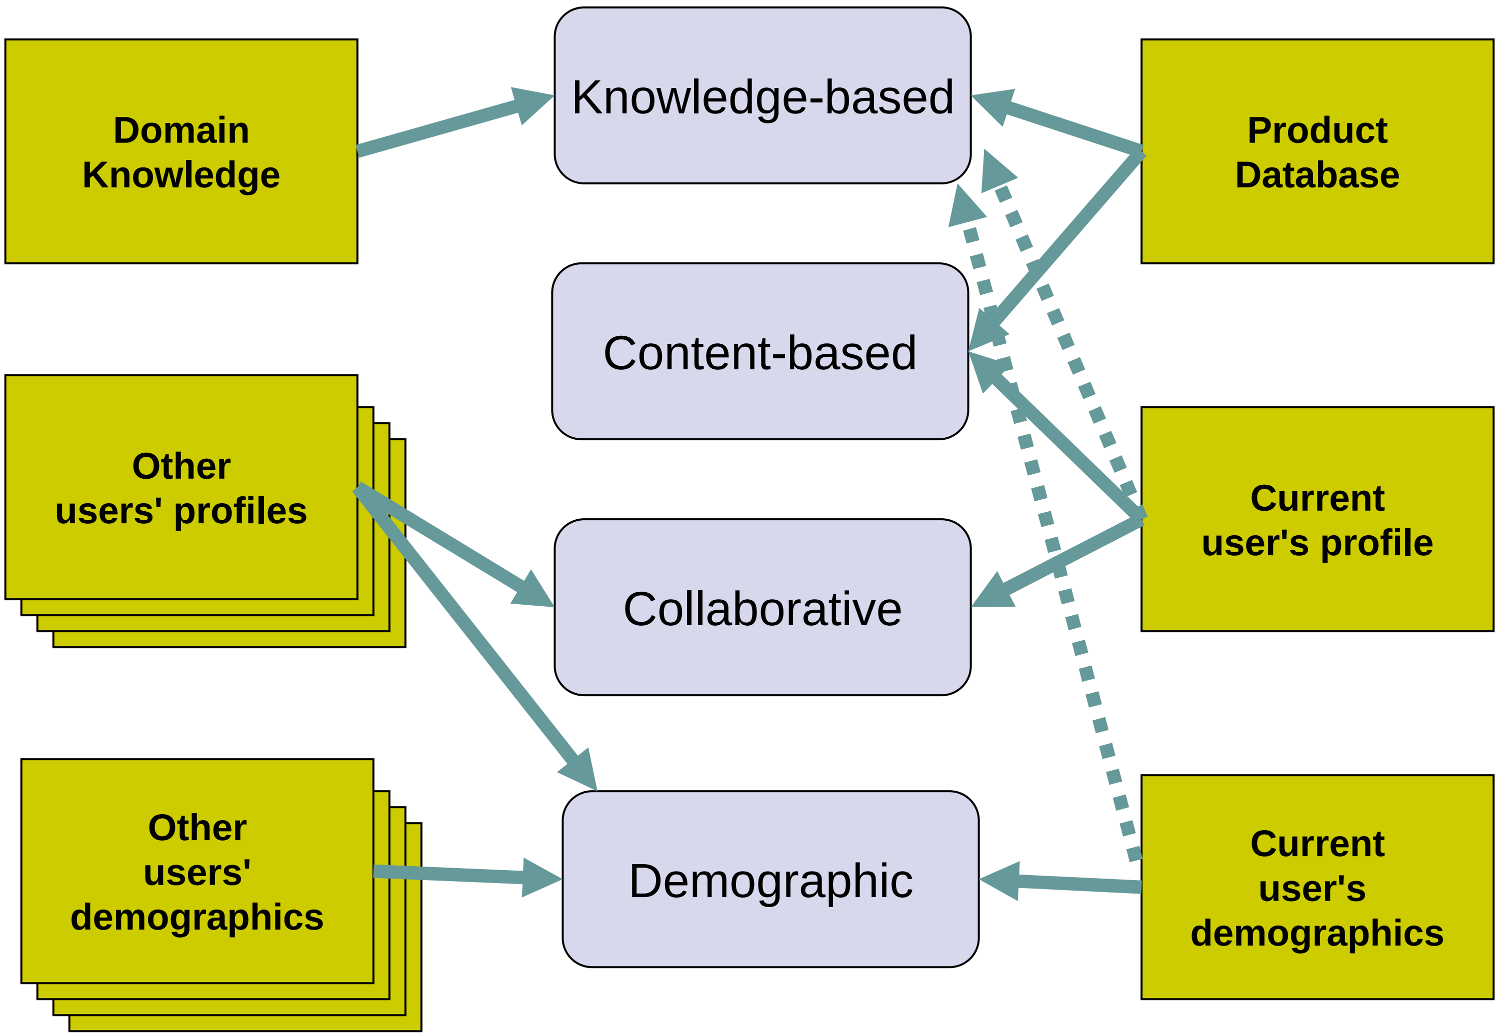

Alfred Kosba

Understanding and Trusting Recommender Systems

- Do users understand and trust the recommendations they receive from recommender systems?

  It seems so… They most likely don't recognize recommender systems though

- What kinds of information do recommenders need to provide to users to build trust

  Clear interface, brand recognition, consistent behavior, help, explanations, predictability, privacy assurances, third-party seals, testimonials,… (i.e., not very much "information")

- How difficult is it to regain trust in a recommender if it is lost?

  Very difficult if the recommender is central to the purpose of a website. Easy though if the recommender is peripheral.

Let’s not resurrect expert systems

Instead:

- Use recommender systems for low-involvement products only *
- Minimize “world knowledge” and “deep reasoning”
- Make recommender systems preferably peripheral

☛ All this will ascertain a low penalty for failure

**) Products that are bought fairly often and with a minimum of thought and effort because they are not of vital concern nor have any great impact on the consumer's life.*

Instead, blend recommender systems into the cacophony of "everyday recommenders"

(some of which people will start trusting more than others…)

"Wear a fancier shirt!"

"Don't burp!"

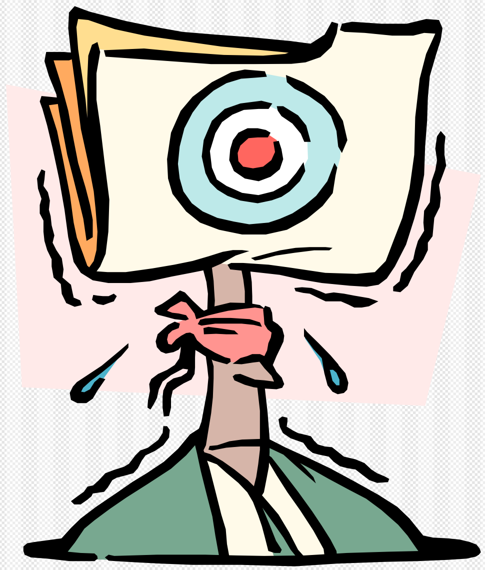

"Don't fly United!"

"Eat bran!"

Hugo Liu

social consequences and opportunities of recommender systems

SOCIAL CAPTURE

- capturing and mirroring recommendations "in the wild"
  - e.g. friendster, chat forums, overheard conversations (?)
- ethnographics of the "wild" fold
  - publicity, intention, social politics

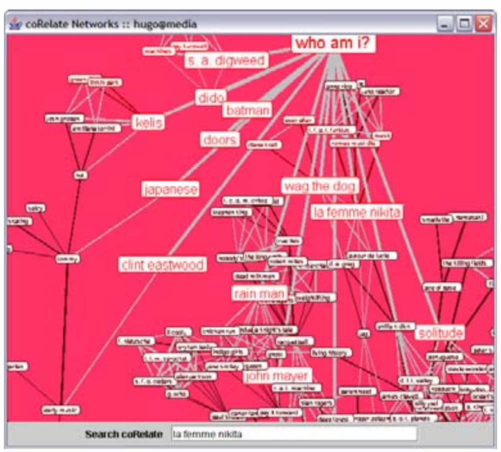

SENSE EVERYDAY CONTEXT

- tight integration ➔ consider aspects of a person's everyday context
  - mood, time, space, social context
  - ex: moviefone
- beyond aesthetic: what is *apropos*? (constraint satisfaction)
- prediction: sensing bottleneck

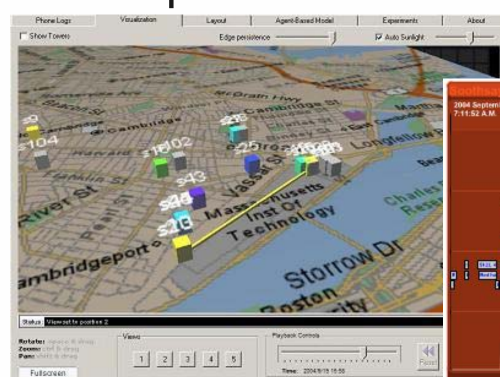

simplicity

at the mit media laboratory in cambridge, massachusetts

social consequences and opportunities of recommender systems

METAPHOR INVOCATION

- *either* Agent or Tool
- Agent's Social Contract
  - belief-desire-intention based
  - ability to explain/justify
  - expected to learn / intimate
- Tool's Design Contract
  - autoexplanatory: transparent
  - "fabric" – intrinsic semantics
  - community intention

PLENITUDE & IDEAL MARKETS

- recommending is *subjective search*
- social organisations
  - improve intraorg expert finding
  - overcome "knowledge discontinuity" problem
- more ideal consumer-product matchmaking
- actualisation of Plato's Plenitude: products directly mirror market desires, greater diversity

Figure 1 Nancy Caruso talks about protecting Boston's North End in our evolving documentary

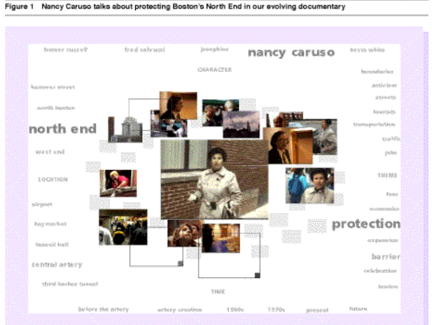

simplicity
at the mit media laboratory in cambridge, massachusetts

Ben Schafer

Interfaces

- Discussion Point : What can we do to create recommenders that
  - Feel more like “decision support” systems (one that can help me deliberate)
  - Don’t feel like recommenders
    - Dialogue based systems
    - Treat us as people rather than users
  - Are “personal” and portable

Barry Smyth

Barry Smyth, University College Dublin, Ireland

Recommend with Confidence

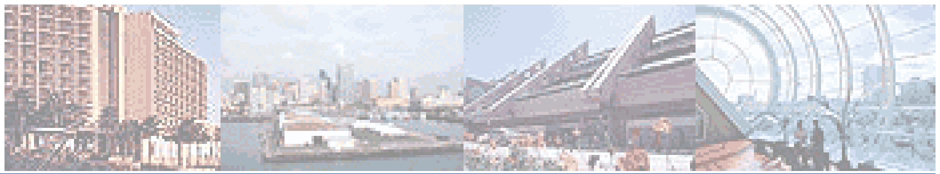

Explanation

- **Recommender Systems as Interactive Reasoners**
  - Success depends on the establishment of a solid relationship between recommender & user
  - Conflicts with traditional “black box” style recommenders
- **Explanations Build Trust**
  - Recommenders must attempt to explain and justify their suggestions if users are to trust these suggestions

Explanation

- **Examples of Explanations**
  - Content-Based Explanations (Keywords, feature comparisons, …)
  - Neighbourhood Explanations (Summary ratings/similarity information about neighbourhood profiles/cases)
  - Confidence values, …
- **See also:**
  - Explaining Recommendations: Satisfaction vs. Promotion (Bilgic & Mooney)
  - Explaining Collaborative Filtering Recommendations (Herlocker, Konstan, Riedl, CSCW’00)
  - ECCBR 2004 Workshop on Explanation in CBR (McSherry & Cunningham)

Confidence

- **Knowing when to say "I don't know"**
  - Most of us are able (if not always willing) to state our level of confidence in our views, opinions, and suggestions
  - Confident recommendations = Accurate recommendations !
  - By evaluating their confidence in a recommendation, a recommender system can choose when to decline from making a recommendation
- **See also:**
  - Measures of Solution Accuracy in Case-Based Reasoning Systems (Cheetham & Price, ECCBR 2004)

Breakout Sessions

- There were four breakout sessions, one for each topic of the workshop
- Notes are ‘raw’ and presented in order mentioned in the workshop

Understanding and Trusting (1)

- Explaining recommendations to users is important though hard, perhaps we can use bad recommendations to help
- There are differences in needs and context between users, we need to do more than just mess around with statistics, we really need to understand the needs of the users
- Confidence vs. risk: confidence is something that might be internally calculated by the recommender, risk however is what the user eventually experiences

Understanding and Trusting (2)

- Recommender security and robustness, it is possible to skew recommendations, we need to develop ways to detect and prevent compromises
- Can we trust (other) users? People can have agendas; not only other users but also the owners of the recommender system.
- There is a trust tradeoff of privacy versus the information that people provide to a recommender. There is a reward too.

Understanding and Trusting (3)

- Could provide a list of what is not recommended to a user instead of what is recommended. Could this be used to help explain the recommendation process to users? E.g. “X is not recommended because it misses y and z”.
- What are the needs of users? How can we use knowledge of these needs to determine what item features to use for the recommendations?

User Interfaces (1)

- One of the main user interface issues remains the lack of screen space compared to what is necessary to provide recommendations, scores and explanations (frustrating for a designer's point of view)
- Privacy vs. data gathering, how can we get users to share more information and while we protect their data and have the users feel good about it?
- What are the risk/benefit tradeoffs? Attention economics for recommenders.

User Interfaces (2)

- Guidelines for UIs in recommender systems are necessary; however recommender system researchers first need to study more basic UI design literature
- What is the user context for giving or receiving a recommendation? How is context domain dependent?
- How can we make users aware and sure about who is in control over their data in for example P2P environments?

User Interfaces (3)

- Determining the full context of a user is still an obstacle (e.g. how to determine the user’s emotional state?)
- Should explicit recommendations even be visible to users as being recommendations? (star ratings, etc.)

Metrics and Algorithms (1)

- Need real metrics, significant online testing; lack of theory on satisfaction, although this might be something very personal
- Need midlevel configurable components, to be combined together, or is it too individualized in each domain to do this?
- Need owner centric metrics, study if site meets owner objectives. How do owner objectives interfere with generated recommendations?

Metrics and Algorithms (2)

- Need metrics for security and vulnerability of recommender system
- Need algorithms for baskets of items, and in a time series
- Are the original ‘horrible’ algorithms are still the best? Is there a point in getting better? Why do users rate items? What is their context?

Metrics and Algorithms (3)

- Can recommenders scale? Multi- or cross-domain recommenders? Common sense recommenders?
- Can we use recommender systems to identify holes in a market?
- Results of dialogs with users?
- Can we create algorithms with formal properties? Can we combine these formal properties with user context?

Metrics and Algorithms (4)

- How can we learn more from machine learning, data mining and decision theory?
- Need a richer set of data with contextual information about ratings and user's needs
- Can recommender systems be used for educational purposes?

Social Implications (1)

- What are the best ways to combine recommenders and social networks?
- Recommenders can have personas: recommender as a person, as the system, etc. Can personas separate recommenders from the system they are imbedded in? Do we even want that to happen?

Social Implications (2)

- Understanding the user and context is important
- Use both positive and negative feedback in recommenders (especially in a social context), how can we acquire this? Should a recommender system give non optimal recommendations in order to get feedback on those items that are most informative?

Social Implications (3)

- Can recommender systems be social actors?
  Do people want to be social when online? Can we determine when do they want to be social and when not?
- People have different spheres of privacy, how do they translate across aspects of people’s lives? Can recommender respect them?

Summary and Conclusions

- Need to include user context when generating recommendations, gathering that information is difficult
- Trust and confidence are very important
- Strong concerns about recommender security and robustness

Finally…

- Papers will continue to be available at workshop website
  - http://www.grouplens.org/beyond2005/
- Join the collab@sims.berkeley.edu collaborative filtering mailing list
  - To join, send an e-mail to: majordomo@sims.berkeley.edu with ‘subscribe collab’ in the body
- Join the RecommenderSystems Yahoo! Group
  - To join, send an email to: RecommenderSystems-subscribe@yahoogroups.com