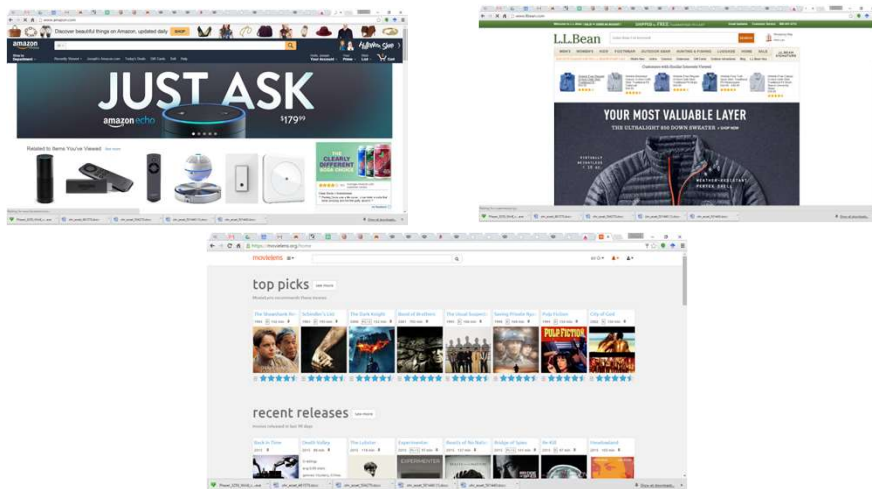


HCI and Recommender Systems

Joseph A. Konstan
konstan@umn.edu

What's wrong with these?



The screenshot shows the Amazon.com homepage. At the top, there is a navigation bar with the Amazon logo, a search bar, and links for 'Your Account', 'Prime', and 'List'. Below the navigation bar is a large banner for the Amazon Echo smart speaker, with the text 'JUST ASK' in large white letters and 'amazon echo \$179.99' below it. Underneath the banner, there is a section titled 'Related to Items You've Viewed' showing various smart home devices like smart speakers, smart light bulbs, and smart switches. To the right of this section is a product listing for 'THE CLEARLY DIFFERENT SODA CHOICE' featuring cans of Clear Zero and Fountains. The browser's address bar shows 'www.amazon.com' and the page title is 'Discover beautiful things on Amazon, updated daily'.

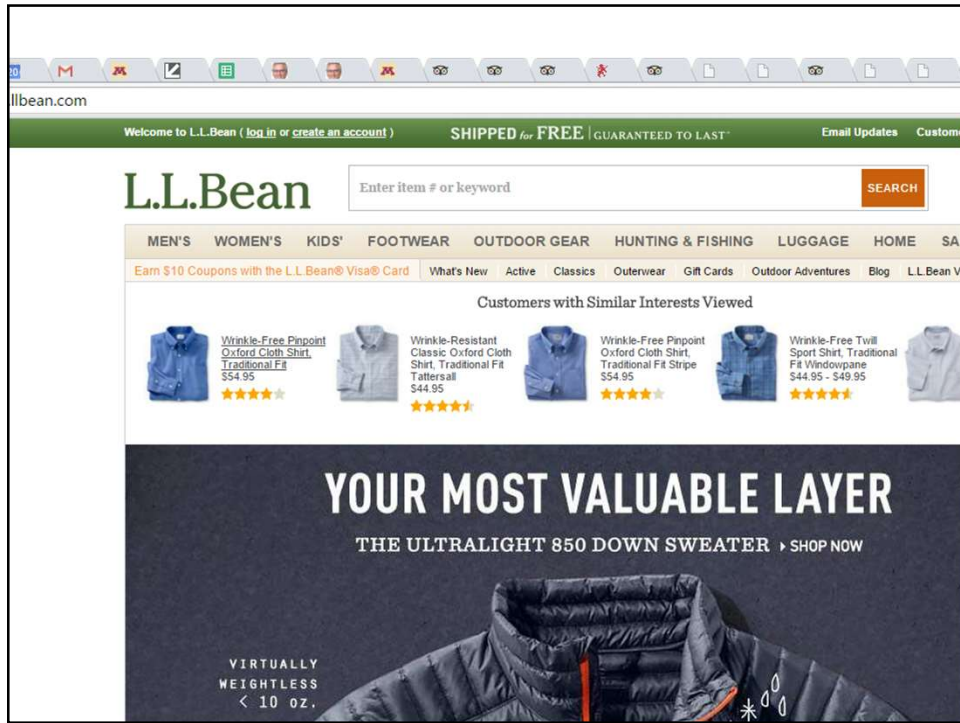
grøuplens | UNIVERSITY OF MINNESOTA

3

The screenshot shows the MovieLens website. The main heading is 'top picks' with a 'see more' link. Below this, there is a row of movie posters with their titles and release years: 'The Shawshank Redemption' (1994), 'Schindler's List' (1993), 'The Dark Knight' (2008), 'Band of Brothers' (2001), 'The Usual Suspects' (1995), 'Saving Private Ryan' (1998), 'Pulp Fiction' (1994), and 'City of God' (2002). Each poster has a star rating below it. Below the 'top picks' section is a section titled 'recent releases' with a 'see more' link. This section shows movies released in the last 90 days, including 'Back in Time' (2015), 'Death Valley' (2015), 'The Lobster' (2015), 'Experimenter' (2015), 'Beasts of No Nation' (2015), 'Bridge of Spies' (2015), 'Re-Kill' (2015), and 'Headland' (2015). The browser's address bar shows 'https://movielens.org/home' and the page title is 'movielens'.

grøuplens | UNIVERSITY OF MINNESOTA

4



What was wrong?

**ALL ARE SHOWING ME THE SAME
RECOMMENDATIONS AS LAST
TIME!!!**

What Metric Shows this Problem?

What Metric Shows this Problem?

» Our Challenge:

- Translate user experience into something quantitative that others can optimize for ...
- Two extremes (and lots of middle ground)
 - Theory-less experimentation
 - Optimize for sales in massive A/B tests
 - Theory-driven (and theory-building exploration)
 - Use, validate, and develop theories of user behavior

OK, Great, Now What?

» Examples

» Discussion

» Lessons

Example #1: How to Get Started?

- » And Moses came down from the mountain and said: “ten items shall ye rate before ye shall be holy and worthy of receiving recommendations.”
 - Well, more or less 😊
- » And the Industry responded: “Our Guests shall not be subject to such barriers of entry; we shall use implicit data, or simply not personalize.”
 - And the researchers were sad ☹

But Research Moved Forward

Al Mamunur Rashid, Istvan Albert, Dan Cosley, Shyong K. Lam, Sean M. McNee, Joseph A. Konstan, and John Riedl. 2002. Getting to know you: learning new user preferences in recommender systems. *Proc. IUI '02*

Sean M. McNee, Shyong K. Lam, Joseph A. Konstan, and John Riedl. 2003. Interfaces for eliciting new user preferences in recommender systems. *Proc. UM'03*.

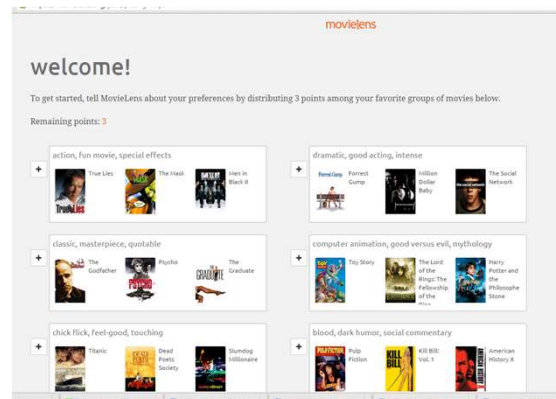
- » Which movies should we ask you about? Mix popularity and entropy.
 - Predicting what you've seen backfires!
- » User directed? Mixed results!
 - Slower, worse, perceived faster, liked.

Room for Innovation

Shuo Chang, F. Maxwell Harper, and Loren Terveen. 2015. Using Groups of Items for Preference Elicitation in Recommender Systems. *Proc. CSCW '15*.

» A new way to start:

- New model
- New algs
- Faster!
- Happier!



Example #2: Filter Bubble ...

Tien T. Nguyen, Pik-Mai Hui, F. Maxwell Harper, Loren Terveen, and Joseph A. Konstan. 2014. Exploring the filter bubble: the effect of using recommender systems on content diversity. *Proc. WWW '14*.

- » Question: Do recommenders really lead to narrowing consumption?
- » Challenge: How to measure?
 - Identifying recommendation-takers
 - Measuring diversity of recent consumption
- » Surprising result: Consumption narrows, but less than without recommender!

Example #3: Towards Useful

- » Pause here for a long rant on the difference between data mining and recommendation!

Example #3: Towards Useful

- » Pause here for a long rant on the difference between data mining and recommendation!
 - Thanks! I feel better now
- » Looking at Diversity and Serendipity
 - Even the definitions are hard:
 - Diversity*: How different recommendations are from each other?
 - Serendipity*: How unexpected recommendations are?

Diversity and Serendipity

Cai-Nicolas Ziegler, Sean M. McNee, Joseph A. Konstan, and Georg Lausen. 2005. Improving recommendation lists through topic diversification. *Proc. WWW '05*.

Komal Kapoor, Vikas Kumar, Loren Terveen, Joseph A. Konstan, and Paul Schrater. 2015. "I like to explore sometimes": Adapting to Dynamic User Novelty Preferences. *Proc. RecSys '15*.

- » Early work: confirmed intuition that diversification can add value *even when decreasing accuracy*
- » Recent work by Kapoor/Kumar shows temporal changes in novelty-seeking among users

Giving Users Control ...

F. Maxwell Harper, Funing Xu, Harmanpreet Kaur, Kyle Condiff, Shuo Chang, and Loren Terveen. 2015. Putting Users in Control of their Recommendations. *Proc. RecSys '15*.

Michael D. Ekstrand, Daniel Kluver, F. Maxwell Harper, and Joseph A. Konstan. 2015. Letting Users Choose Recommender Algorithms: An Experimental Study. *Proc. RecSys '15*.

- » We've started giving users greater control over their recommendation algorithms

But Anchored in Understanding How User's See Recommendations

Michael D. Ekstrand, F. Maxwell Harper, Martijn C. Willemsen, and Joseph A. Konstan. 2014. User perception of differences in recommender algorithms. In *Proc. RecSys '14*.

- » Virtual lab experiment to explore user perception of recommendations, varying algorithms and comparing perceptions with analytic metrics
 - Found that users overall prefer less novelty but more diversity.

Next Steps: Psych + Temporal

- » Tien Nguyen working on studies showing links between recommendation preferences (e.g., diversity, serendipity, popularity) and Big-5 personality
- » Next looking at questions of temporal diversity; temporal changes in recommendation preferences; changes based on how people tune their algorithms.

Example #4: Recommending Work

Dan Cosley, Dan Frankowski, Loren Terveen, and John Riedl. 2007. SuggestBot: using intelligent task routing to help people find work in wikipedia. *Proc IUI '07*.

- » In 2006, Cosley started SuggestBot as a recommender that would suggest work to people within Wikipedia
- » Since then, we've continued to experiment with that system, but have tried to generalize that knowledge (along with many others).

Getting and Keeping Volunteers

Qian Zhao, Zihong Huang, F. Maxwell Harper, Loren Terveen, Joseph A. Konstan. Precision Crowdsourcing: Closing the Loop to Turn Information Consumers into Information Contributors. *Proc. CSCW 2016*.

Raghav Pavan Karumur, Tien Nguyen, Joseph A. Konstan. Early Activity Diversity: Assessing Newcomer Retention from First-Session Activity. *Proc CSCW 2016*.

- » Zhao and Huang carried out controlled experiments asking users to do work.
- » Karumur developed a new measure (early activity diversity) that predicts retention.

What else is going on ...

- » Back to algorithms
 - Putting a recommender into public library catalogs highlighted some serious problems with existing algorithms during new-user start-up
- » And new forms of input and output
 - Also giving us reasons to explore direct submission of “product-alike” data and ways to describe unknown items in terms of other, more well-known ones

Five Take-Away Messages

1. Ground research in real problems and user experiences; user challenges are research challenges
2. Mix engineering (solving the problem) with science (learning something general that can be applied elsewhere)
3. Mix methods: data and log analysis, surveys, lab and field experiments

Five Take-Away Messages

4. Collaborate widely (diversity of background leads to new insights).
5. Keep at it; some of these ideas have taken 10+ years to get from initial contributions to present form; some will take 5+ more years to get where we want to go!

Joseph A. Konstan
konstan@umn.edu

DISCUSSION