RECOMMENDER SYSTEMS

CSE435: *Intelligent Decision Support Systems*

Alexandra Coman

Sources

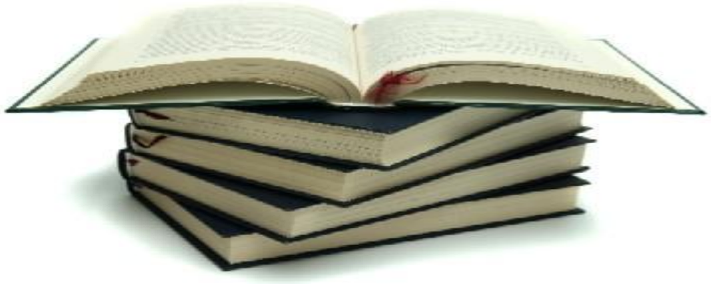

MAIN:

- Burke, R., 2002: “Interactive Critiquing for Catalog Navigation in e-Commerce”
- Burke, R., 2002: “Hybrid Recommender Systems: Survey and Experiments”

ADDITIONAL:

- Smyth, B. and McClave, P., 2001: “Similarity vs Diversity”
- McGinty, L. and Smyth, B., 2003: “On the Role of Diversity in Conversational Recommender Systems”
- McSherry, D. 2001: “Increasing Recommendation Diversity Without Loss of Similarity”
- McSherry, D. 2002: “Diversity-Conscious Retrieval”

A Definition

"any system that produces individualized recommendations as output or has the effect of guiding the user in a personalized way to interesting or useful objects in a large space of possible options."

[Burke, 2002]

Recommendation Problem

- *Problem:* user's need

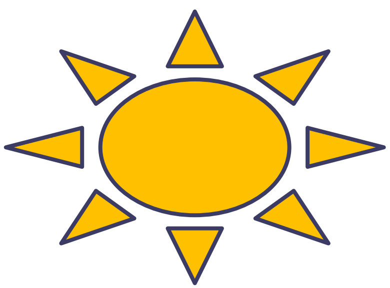

- *Solution:* match between user's need and an available product/service

Recommender Systems

emulate interaction with a salesperson, in view of choosing a product/service from the set of available ones

Related Systems

- Keyword-based search engines

- Other information retrieval systems

What sets recommender systems apart?

A Definition

"any system that produces ***individualized*** recommendations as output or has the effect of guiding the user in a ***personalized*** way to ***interesting*** or ***useful*** objects in a large space of possible options."
[Burke, 2002]

Recommender Systems

INDIVIDUALIZED

USEFUL

INTERESTING

When Necessary?

- Large, difficult to navigate search spaces
  - E.g. e-commerce inventories: "virtually unbounded", no shelf-space and paper-catalog cost limits
- Non-expert user, complex products

Recommender System

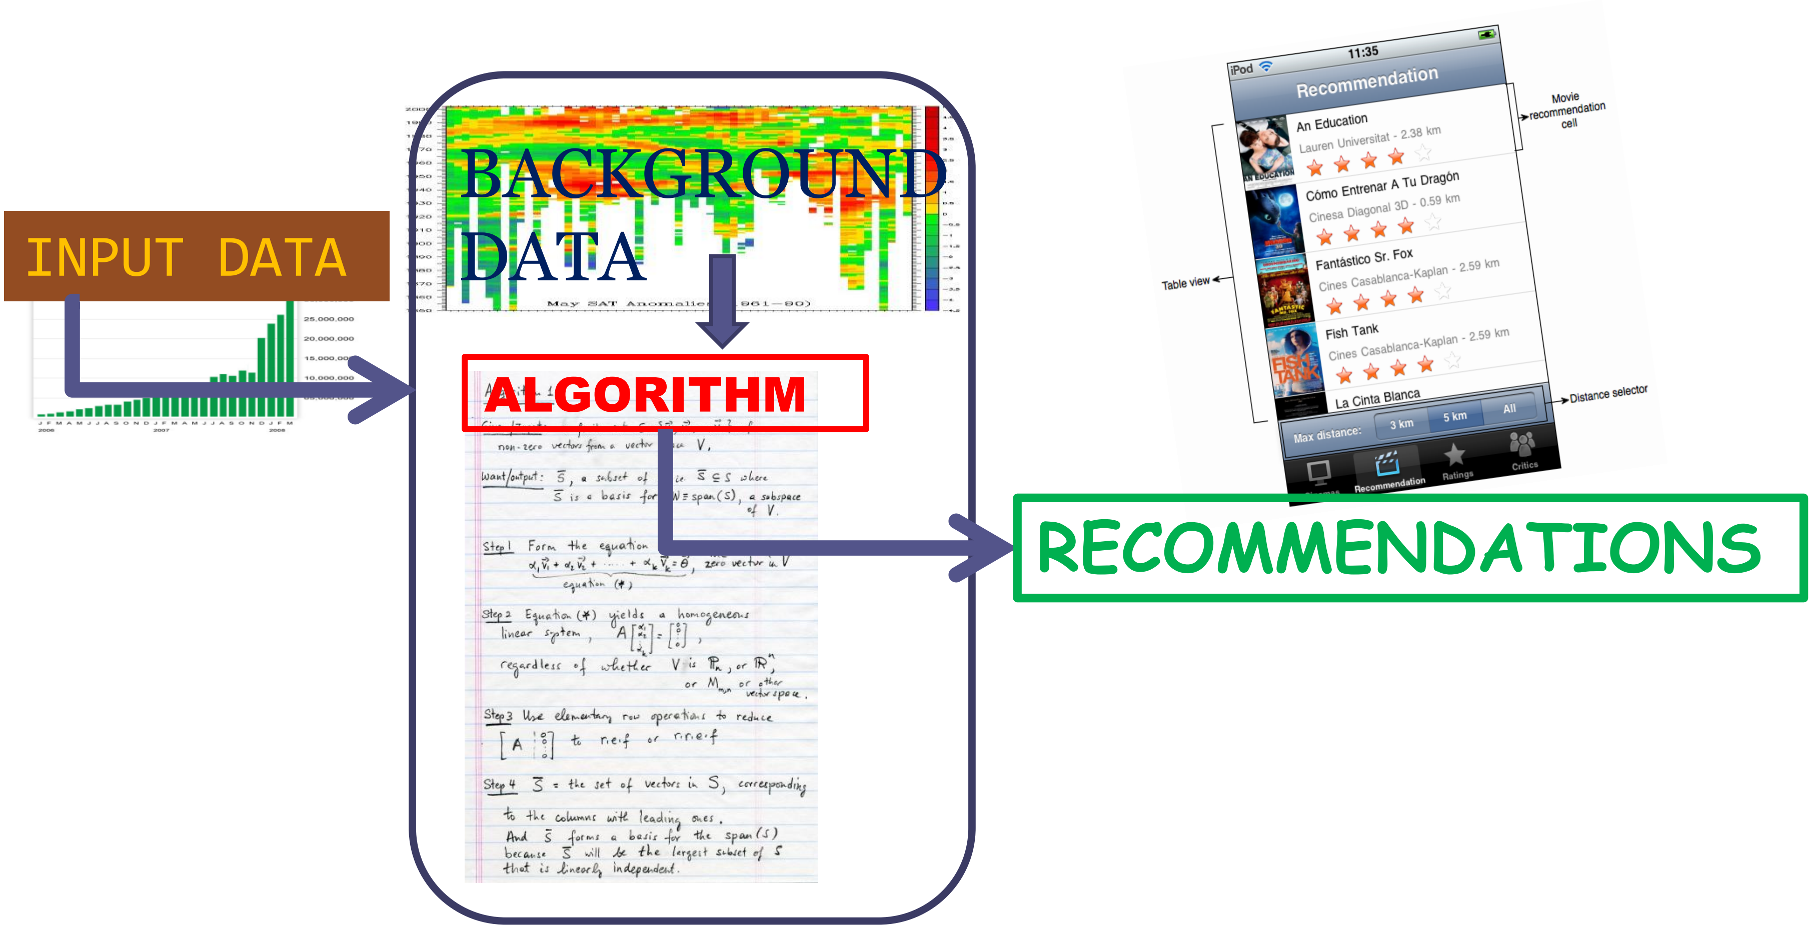

Data?

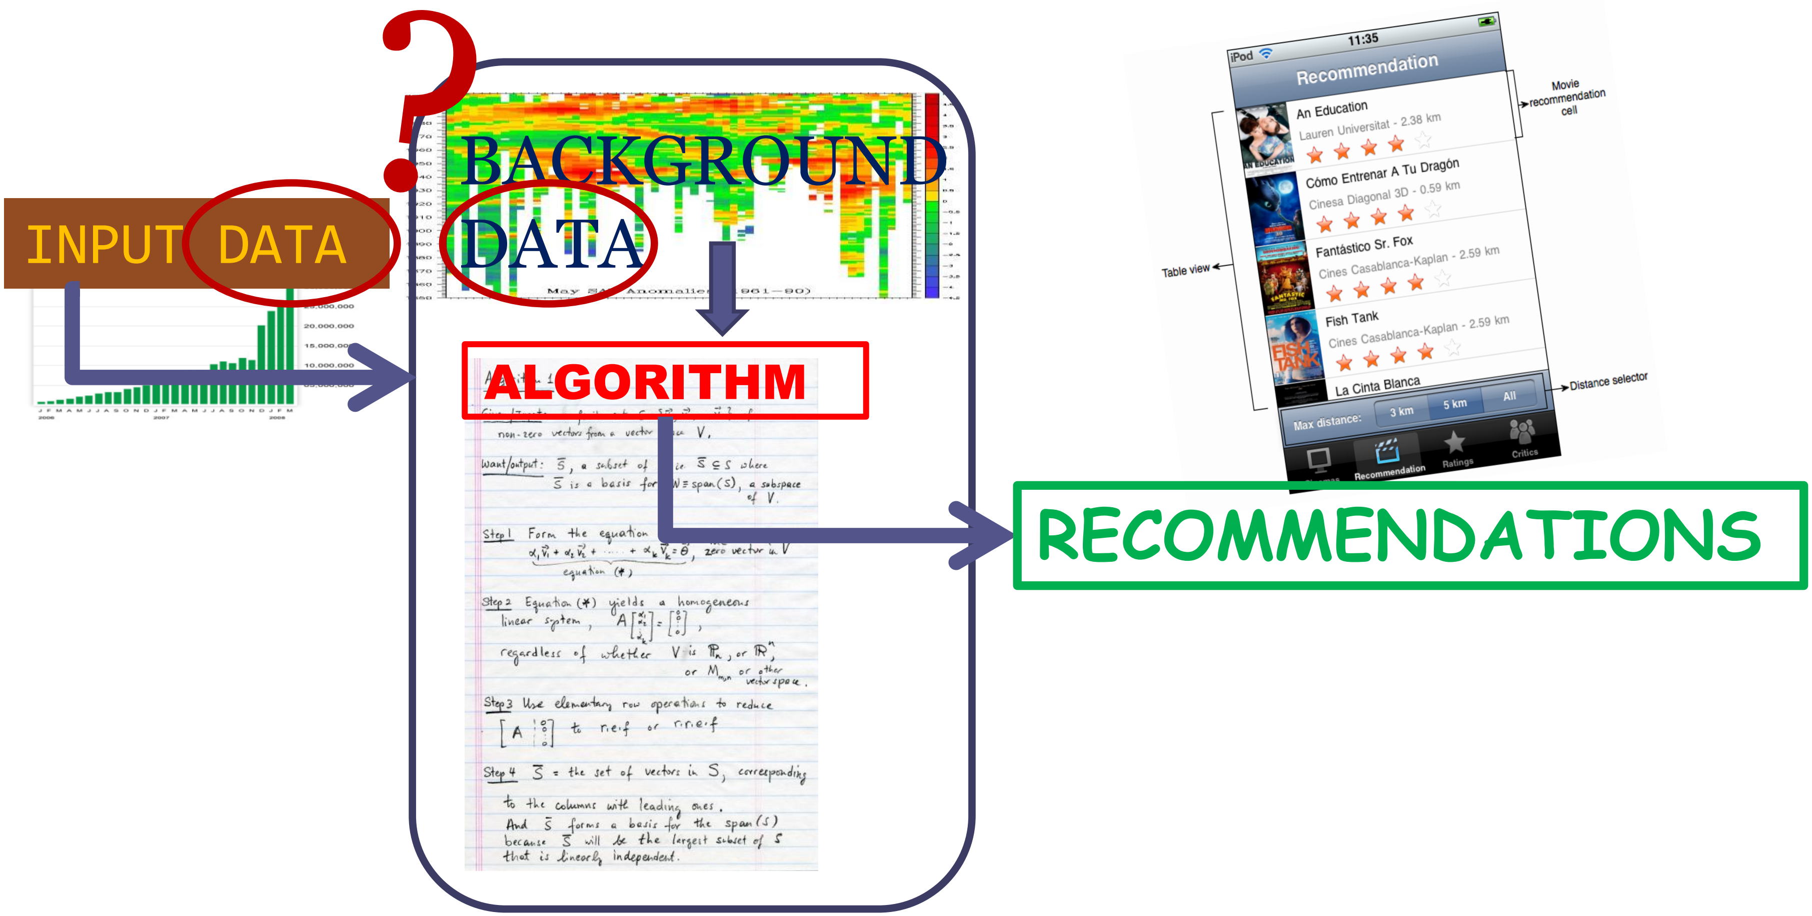
INPUT DATA
BACKGROUND DATA
ALGORITHM
RECOMMENDATIONS
Recommendation
An Education
Lauren Universitat - 2.38 km
Cómo Entrenar A Tu Dragón
Cinesa Diagonal 3D - 0.59 km
Fantástico Sr. Fox
Cines Casablanca-Kaplan - 2.59 km
Fish Tank
Cines Casablanca-Kaplan - 2.59 km
La Cinta Blanca
Max distance:
3 km
5 km
All
Movie recommendation cell
Table view
Distance selector

Input: U, u, I and i

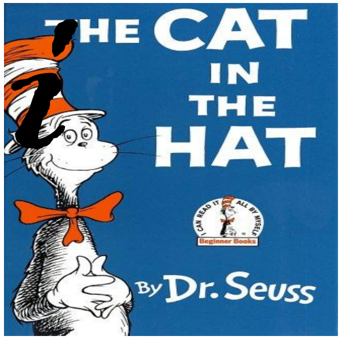

$\in$

$\in$

Input: U, u, I and i

- I : set of items over which recommendations might be made
- U : set of users whose preferences are known
- u : user for whom recommendations need to be generated
- i : item for which we would like to predict u's preference.

Output?

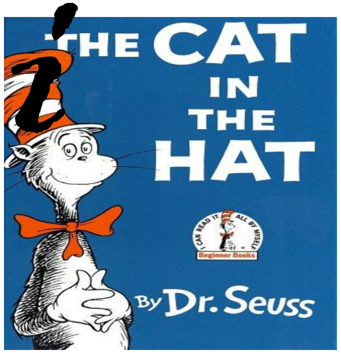

?

- Input: $\mathcal{U}$, u, $\mathcal{I}$, i
- Output: u's predicted preference for i
  - like/dislike (binary)
  - degree of preference (real)

Types of Recommender Systems

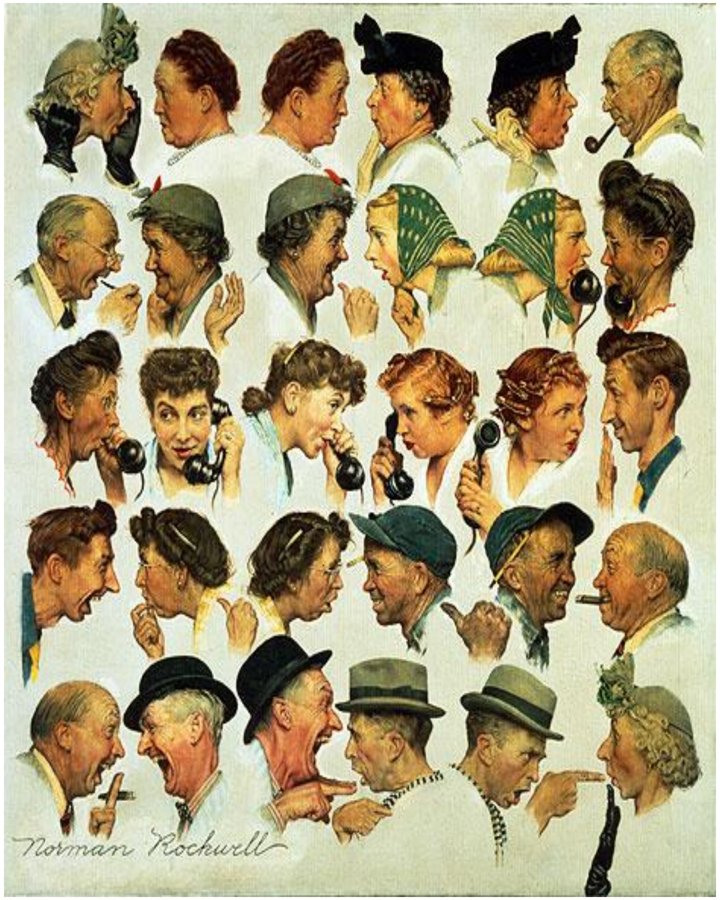

- Collaborative Systems
  - aggregate ratings or recommendations of objects
  - recognize commonalities between users on the basis of their ratings
  - generate new recommendations based on inter-user comparisons
  - possibly, use time-based discounting of ratings

Types of Recommender Systems

- Demographic
  - categorize users based on personal attributes
  - make recommendations based on demographic classes

Types of Recommender Systems

- Content-based
  - objects defined by their associated features
  - learn profile of the user's interests based on the features present in objects the user has rated
  - long-term models, updated as more evidence about user preferences is observed

Types of Recommender Systems

- Utility-based
  - make suggestions based on a computation of the utility of each object for the user
  - employ constraint satisfaction techniques to locate the best match
  - no long-term generalizations about users

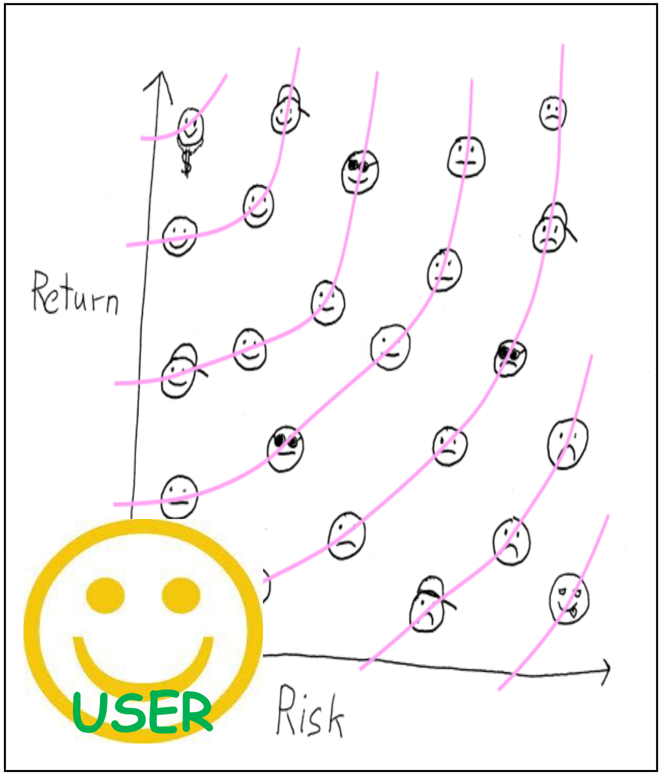

Types of Recommender Systems

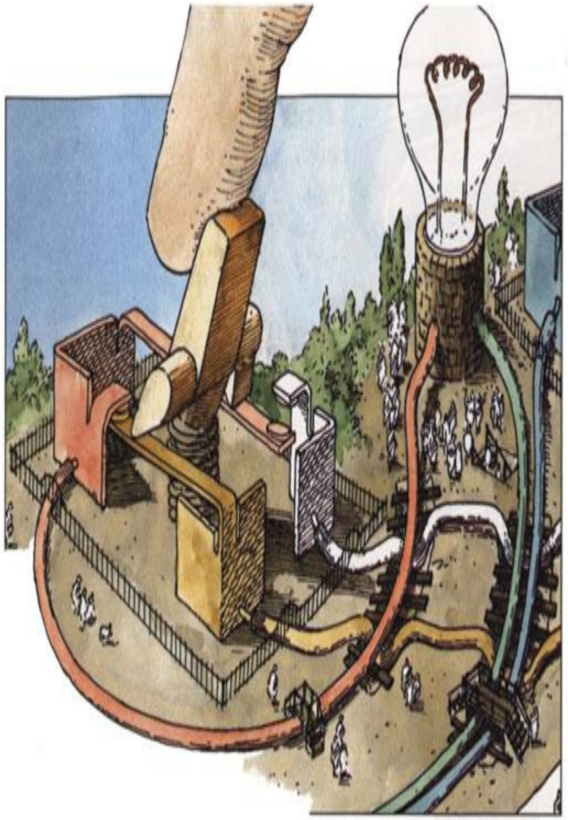

- Knowledge-based
  - functional knowledge: how a particular item meets a particular need
  - can reason about the relationship between a need and a possible recommendation
  - no long-term models

Most popular?

Technique	Background	Input	Process
Collaborative	Ratings from U of items in I.	Ratings from u of items in I.	Identify users in U similar to u, and extrapolate from their ratings of i.
Content-based	Features of items in I	u's ratings of items in I	Generate a classifier that fits u's rating behavior and use it on i.
Demographic	Demographic information about U and their ratings of items in I.	Demographic information about u.	Identify users that are demographically similar to u, and extrapolate from their ratings of i.
Utility-based	Features of items in I.	A utility function over items in I that describes u's preferences.	Apply the function to the items and determine i's rank.
Knowledge-based	Features of items in I. Knowledge of how these items meet a user's needs.	A description of u's needs or interests.	Infer a match between i and u's need.

Most popular

- Why?

Technique	Background	Input	Process
Collaborative	Ratings from U of items in I.	Ratings from u of items in I.	Identify users in U similar to u, and extrapolate from their ratings of i.
Content-based	Features of items in I	u's ratings of items in I	Generate a classifier that fits u's rating behavior and use it on i.
Demographic	Demographic information about U and their ratings of items in I.	Demographic information about u.	Identify users that are demographically similar to u, and extrapolate from their ratings of i.
Utility-based	Features of items in I.	A utility function over items in I that describes u's preferences.	Apply the function to the items and determine i's rank.
Knowledge-based	Features of items in I. Knowledge of how these items meet a user's needs.	A description of u's needs or interests.	Infer a match between i and u's need.

Strengths/Weaknesses?

Strengths/Weaknesses

Technique	Pluses	Minuses
Collaborative filtering (CF)	A. Can identify cross-genre niches. B. Domain knowledge not needed. C. Adaptive: quality improves over time. D. Implicit feedback sufficient	I. New user ramp-up problem J. New item ramp-up problem K. "Gray sheep" problem L. Quality dependent on large historical data set. M. Stability vs. plasticity problem
Content-based (CN)	B, C, D	I, L, M
Demographic (DM)	A, B, C	I, K, L, M N. Must gather demographic information
Utility-based (UT)	E. No ramp-up required F. Sensitive to changes of preference G. Can include non-product features	O. User must input utility function P. Suggestion ability static (does not learn)
Knowledge-based (KB)	E, F, G H. Can map from user needs to products	P Q. Knowledge engineering required.

How do we get the most out of strengths and alleviate drawbacks?

Hybrid Recommender Systems

- Combine multiple methods in order to take advantage of strengths and alleviate drawbacks
- Weighted
  - scores/votes of several recommendation techniques combined together to produce a single recommendation
- Switching
  - system switches between recommendation techniques depending on the current situation
- Mixed
  - recommendations from several different recommenders presented at the same time

Hybrid Recommender Systems

- Feature combination
  - features from different recommendation data sources thrown together into a single recommendation algorithm
- Cascade
  - one recommender refines the recommendations given by another
- Feature augmentation
  - output from one technique is used as an input feature to another
- Meta-level
  - the model learned by one recommender is used as input to another

Which types are used by…

- Netflix?
- Amazon?
- Pandora?
- Last.fm?
- BookLamp?
- Others?

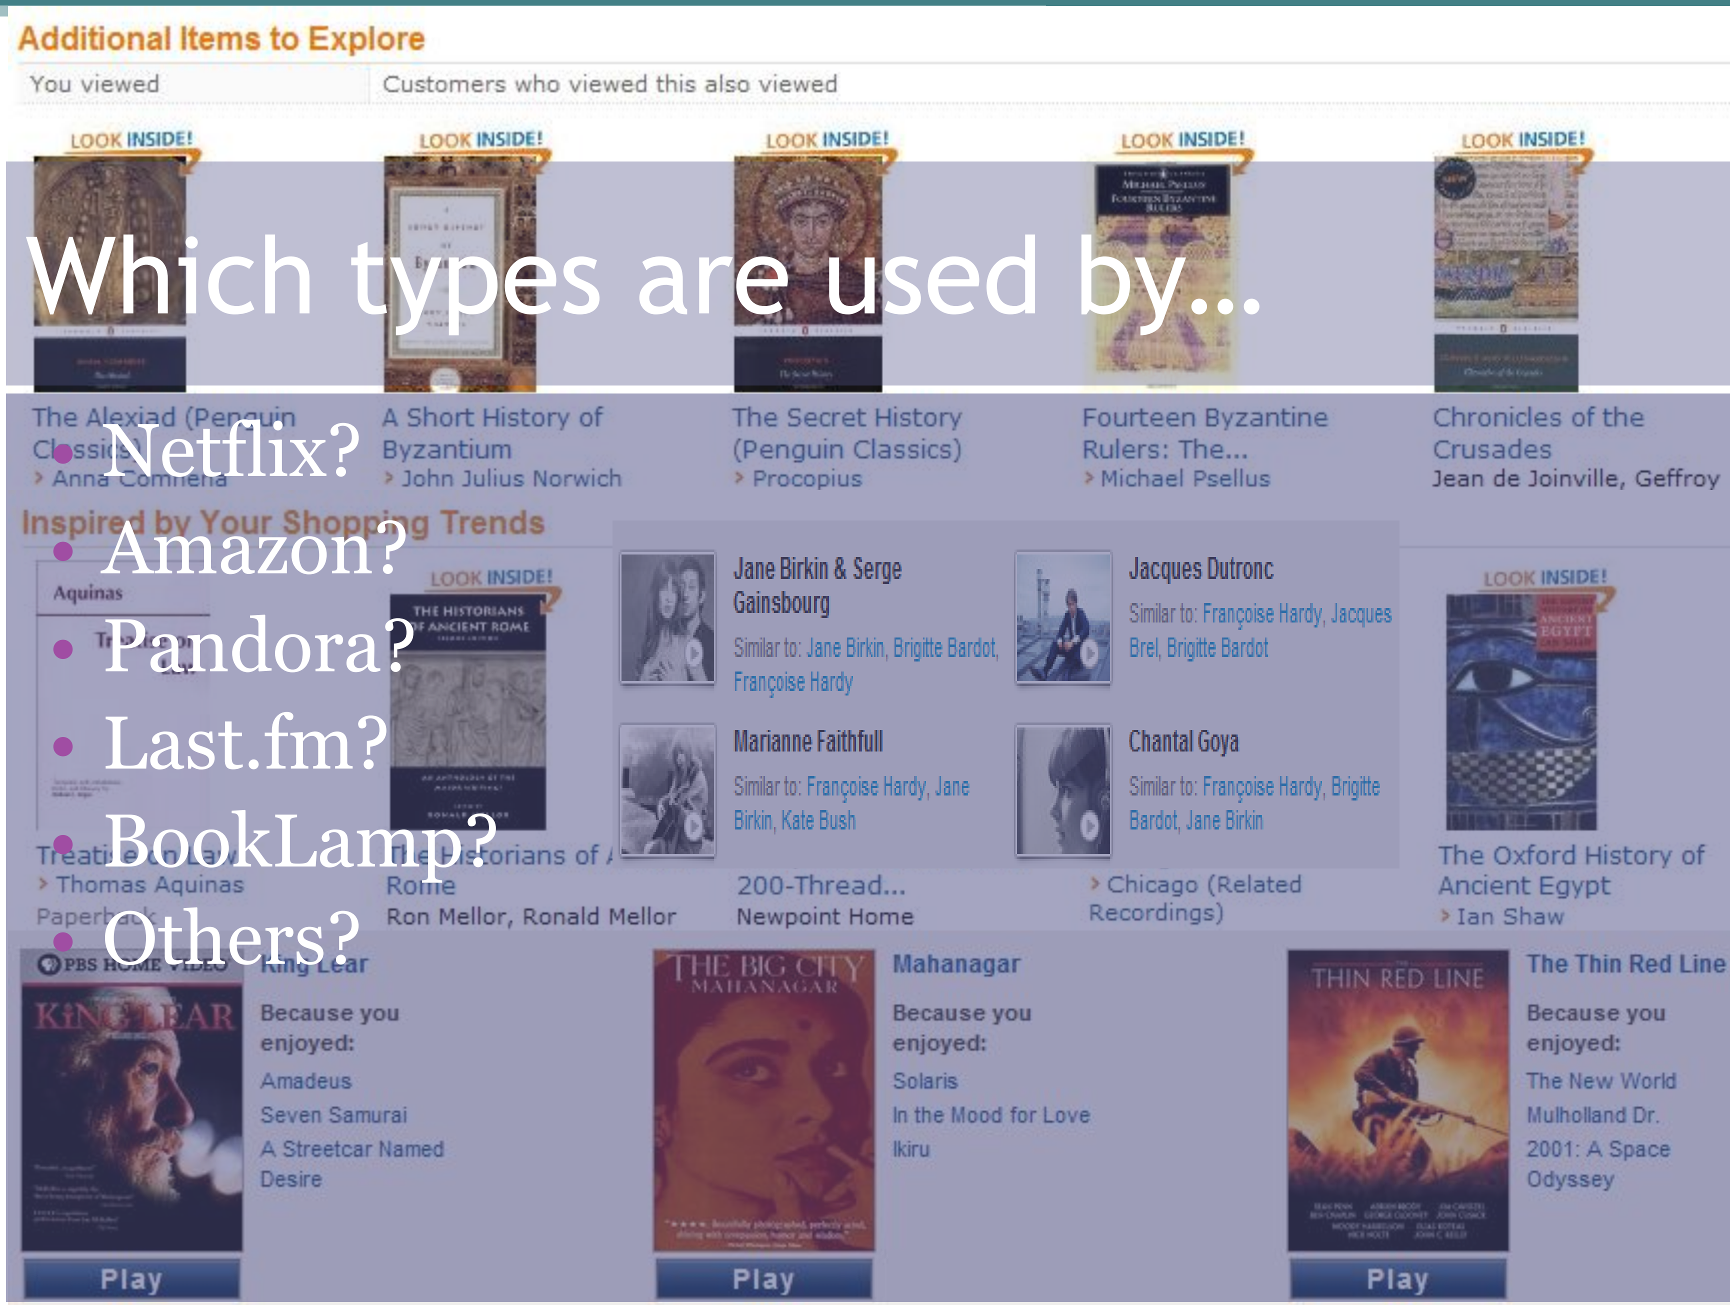

What’s Missing?

- So far, we have this:

What's Missing?

- So far, we have this:

What if not the right match?

What's Missing?

If no product of interest in recommended set, start over with reformulated search?

"NARROWING" PROBLEM

Solution?

- Ongoing interaction:

"The Women", George Cukor, 1939

Critiquing

- Navigation of solution space that solves the “narrowing” problem
- Interactive, incremental: does not require that the user have a completely specified need at the start
- Shoppers “learn” by exploring the product space
- Requirements not static, but constantly shifting

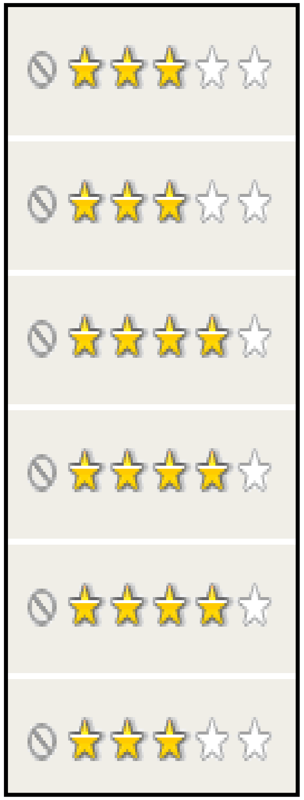

“FindMe” Critique-Based Retrieval

- CBR System
- Navigation system superior to repeated search: intuitive, natural, requires less effort
- Steps:
  - User chooses “source”/”entry point” from catalog (using search or direct name-based identification)
  - Perform case retrieval, find items most similar to source
  - Obtain “critique” of presented examples
  - Use critique to redirect the search, filtering the solution space along specified feature dimension

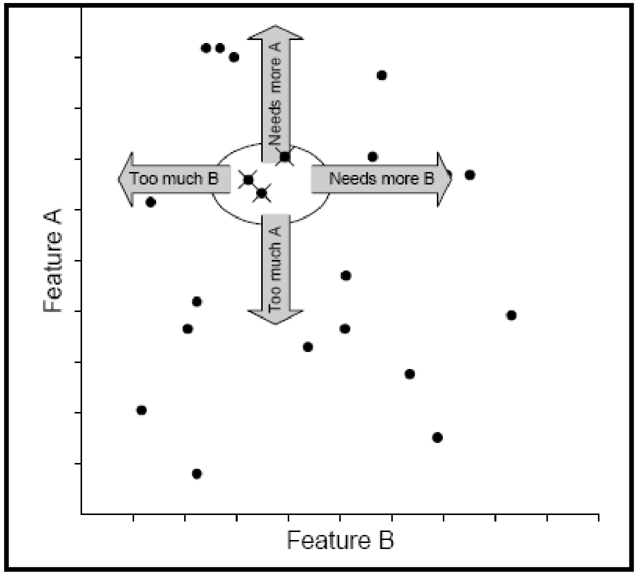

R.Burke: “Interactive Critiquing for Catalog Navigation in e-Commerce”

“FindMe” Critique-Based Retrieval

- Steps:
  - User chooses “source”/”entry point” from catalog (using search or direct name-based identification)
  - Perform case retrieval, find items most similar to source
  - Obtain “critique” of presented examples
  - Use critique to redirect the search, filtering the solution space along specified feature dimension (e.g. “more B”)

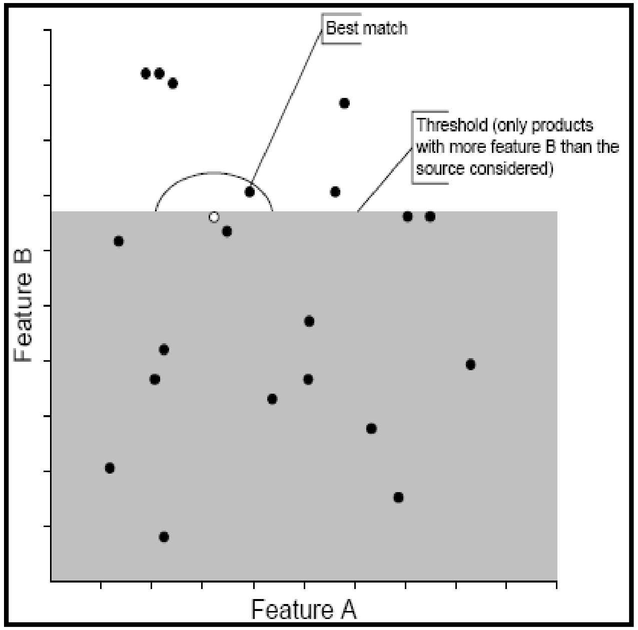

R.Burke: “Interactive Critiquing for Catalog Navigation in e-Commerce”

Recommendation and Critiquing

- ["Entry point"] A red polygon, please!
- Something with fewer sides?
- … but equilateral!
  (note: "red" feature lost)
- Something similar in green?
- … and with even fewer sides?

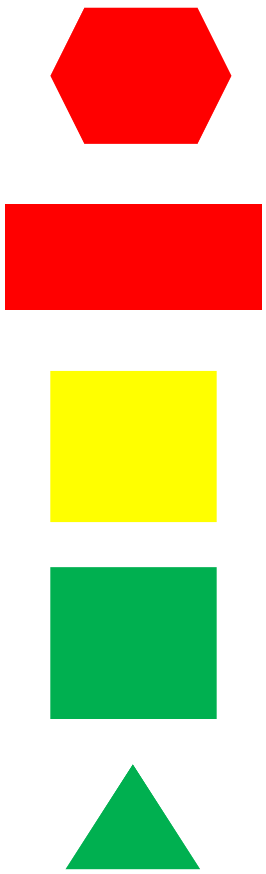

Similarity?

SIM(

, )=?

Similarity

- Not simple or uniform, situation-specific
  - E.g. SIM(car,boat)=1 if location near water
  - SIM(car, boat)=0 if location landlocked
- Must capture buyers' intuitive sense of what ought to be considered similar
- "Local Similarity metric": goal-oriented, multiple goals and their trade-offs considered
- Hierarchy of global similarity metrics : domain-specific

Entrée Chicago

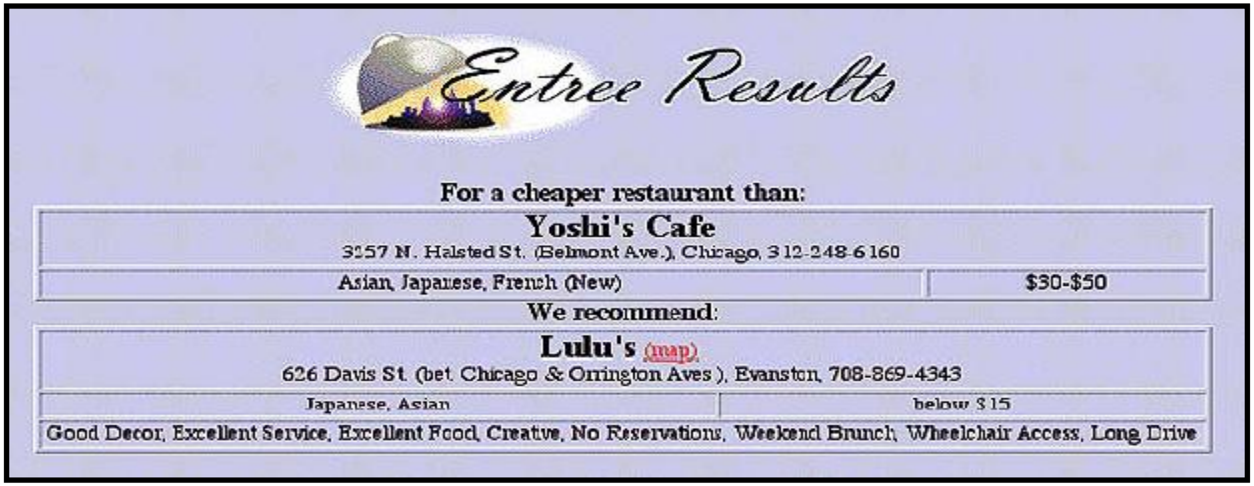

- Knowledge-based restaurant recommender
- “FindMe” System:
  - Similarity-based Recommendation
  - Critique-based Navigation

Entrée Chicago

- Goal hierarchy:
  - cuisine (I), price (II), quality (III), atmosphere (IV)
- Unfortunately, no longer online

A Demo Recommender System

- Movielens
  - http://www.movielens.org/login

What's the Potential Problem Here?

- ***Query: green, triangular***

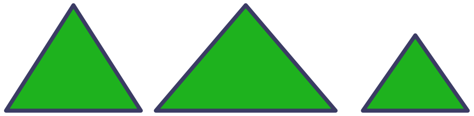

- ***Query: early twentieth-century novel by British author***

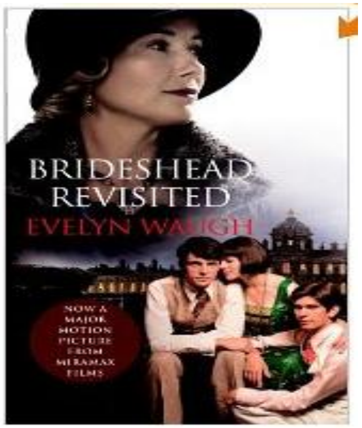

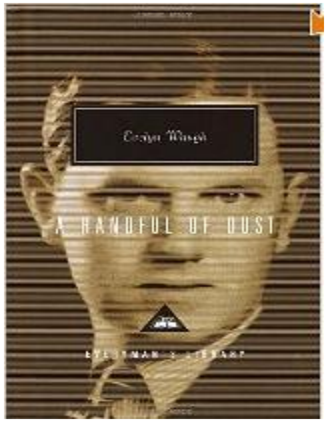

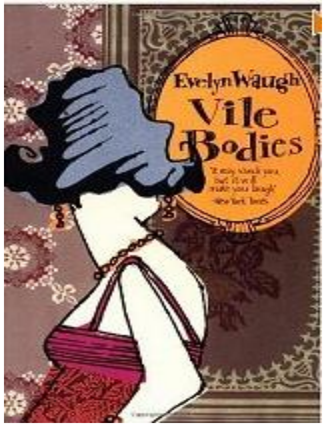

Diversity!!!

- *Query: green, triangular*

- *Query: early twentieth-century novel by British author*

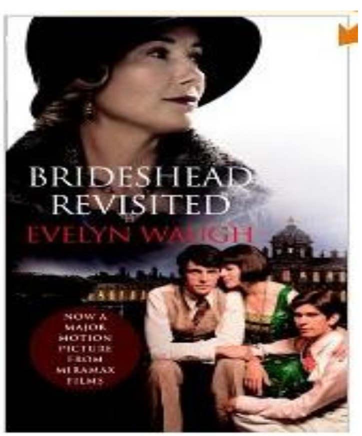

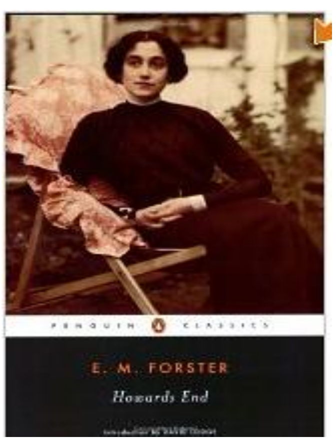

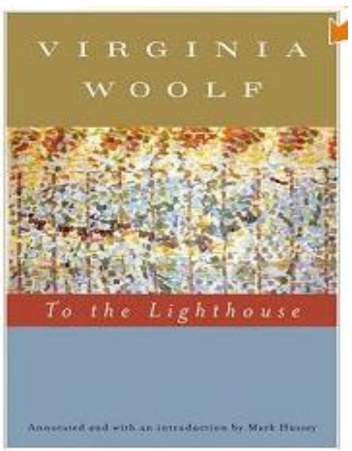

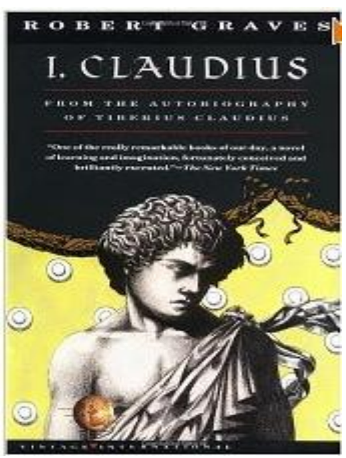

Diversity

- Variation of retrieved results
- Better sample of solution space
- Genuine alternatives for the user
- "Surprising" options that may otherwise not be considered

MUST be balanced with similarity!

Diversity-Aware Retrieval Algorithms

- No algorithm for optimal balance between Similarity and Diversity
- “Bounded Greedy” [Smyth & McClave]: a heuristic attempt, so far the most popular
- “Similarity Layers” [McSherry]
- “Adaptive Selection” [McGinty and Smyth]
  - doses similarity/diversity intelligently, based on user’s choices on each recommendation cycle

Current and Future Research Directions

- Mobile Recommender Systems (*MobyRek*)
- Recommender systems leveraging and enhancing Social Networks?

F. Ricci and Q. N. Nguyen, Critique-Based Mobile Recommender Systems, OEGAI Journal, 24(4):2005.

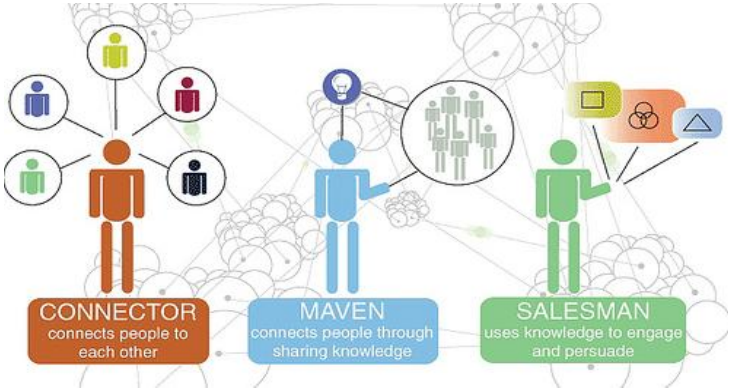

Summary

- Recommender systems help navigate vast product spaces, helping locate items that are interesting and useful to individual users
- Mostly used in business-to-person e-commerce contexts
- Multiple types of recommender systems, based on type and source of input/background data (can be combined into "hybrids")
- "Critiquing" makes navigation intuitive and efficient, emulating interaction with a salesperson; it helps reduce the "narrowing" effect
- Diversity a relevant, currently-explored issue, must be balanced with Similarity