Foundations of Web Personalization and Recommender Systems

Jill Freyne & Shlomo Berkovsky

25 July 2015

Outline

Part 1: Introduction, Information Overload, User Modelling

Part 2: Personalization for Information Filtering, Information Access & Content Delivery

Part 3: Recommender Systems

Part 1:
Introduction,
Information Overload,
User Modelling

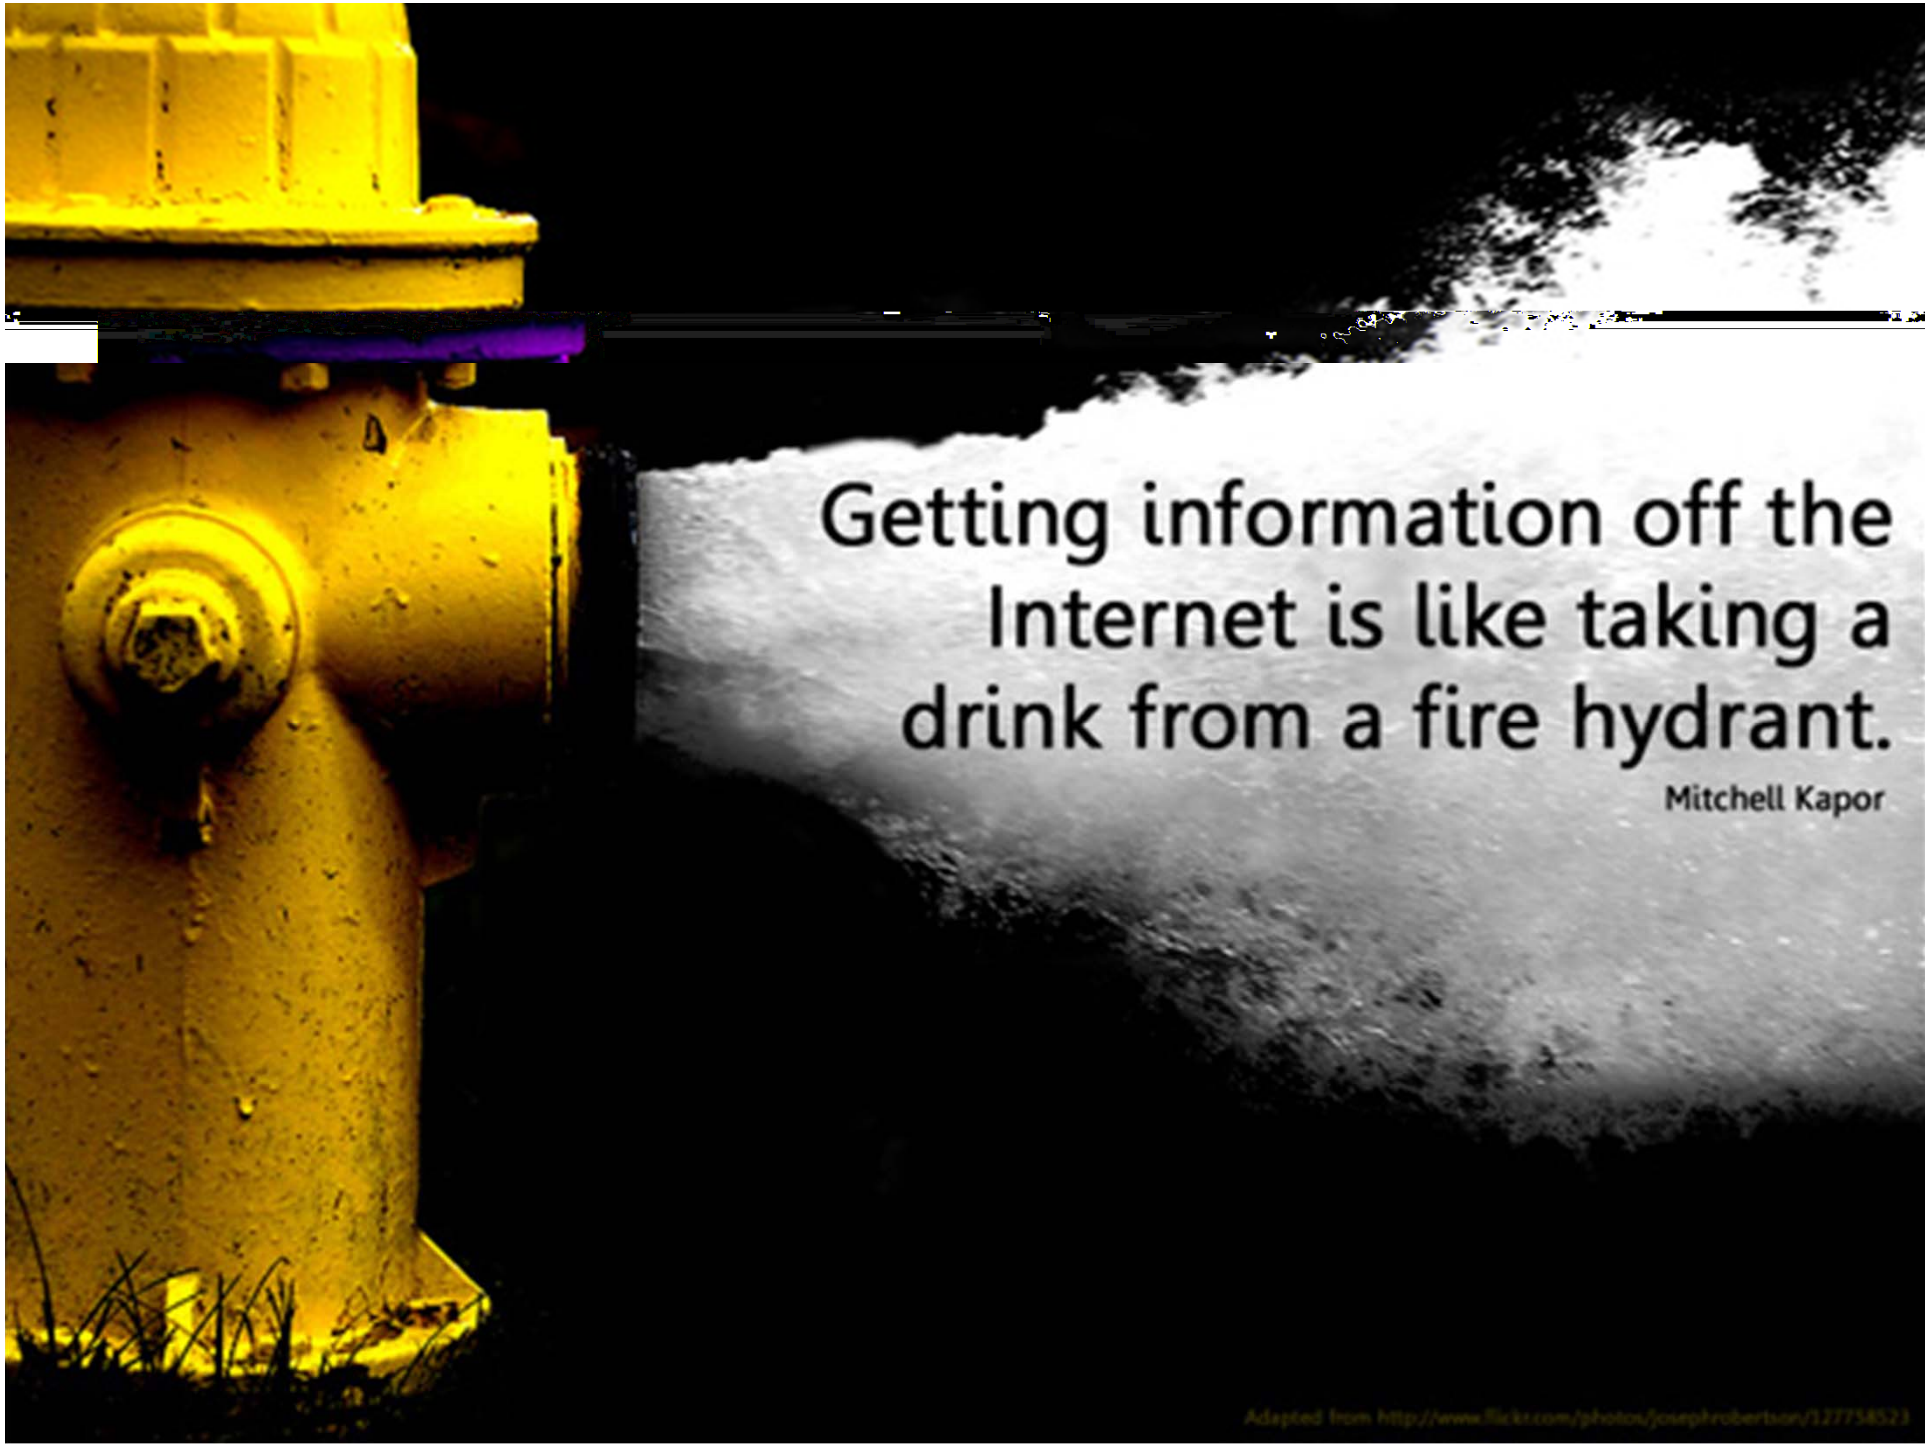
Getting information off the Internet is like taking a drink from a fire hydrant.
Mitchell Kapor

Information Overload

- Information presented at a rate too fast for a person to process

- The state of having too much information to make a decision or remain informed a topic

Online Information Overload

- Every time we go online, we are overwhelmed by the available options
  - Web Search….which search result is most relevant to my needs?
  - Entertainment….which movie should I download, which restaurant should I eat at?
  - E-commerce….which product is best for me? which holiday will I enjoy most?
  - News….which news stories are most interesting to me?
  - Health….which healthy meals will I enjoy? which types of exercise should I try? what doctor can I trust?

Google
News

U.S. edition

News

- Top Stories
  - Charleston
  - Golden State Warriors
  - U.S. Open (golf)
  - Pope Francis
  - Alex Rodriguez
  - Chicago Blackhawks
  - Donald Trump
  - Boston Red Sox
  - James Eagan Holmes
  - Mumbai
- Sydney, New South Wa...
- Suggested for you
- World
- U.S.
- Business
- Technology
- Entertainment
- Sports
- Science
- Health
- Spotlight

Top Stories

Reuters

See realtime coverage

NRA executive suggests slain Charleston pastor to blame for gun deaths

Reuters - 1 hour ago

DALLAS A National Rifle Association executive in Texas has come under fire for suggesting that a South Carolina lawmaker and pastor slain with eight members of his congregation bears some of the blame for his opposition to permitting concealed ...

Dylann Roof's friend: 'He never said anything racist' BBC News

Dylann Roof talked of 'hurting a bunch of people' before shootings, says friend The Guardian

Related
Charleston »
South Carolina »

Featured: Dylann Storm Roof's friend took gun away during 'crazy' bigoted rant 2 weeks ...
New York Daily News

In Depth: Raw emotion as victims' families address Charleston suspect Miami Herald

Wikipedia: Charleston church shooting

Pittsburgh ... Chicago Su... Daily Beast Quartz Huffington ... Foster's Da... Pittsburgh ...

Ynetnews

US report finds Iran threat undiminished as nuke

Ynetnews - 3 hours ago

Islamic Republic's support for terrorist proxies did not decrease last year, and even expanded in some ways, says US gov't.

NPR

Leaving Brooklyn, Bernie Sanders Found Home In Vermont

NPR - 42 minutes ago

This story is part of NPR's series Journey Home. We're going to the places that presidential candidates call home and finding out what those places tell us about how they see the world.

Charleston Church Shooting Renews Confederate Flag Debate

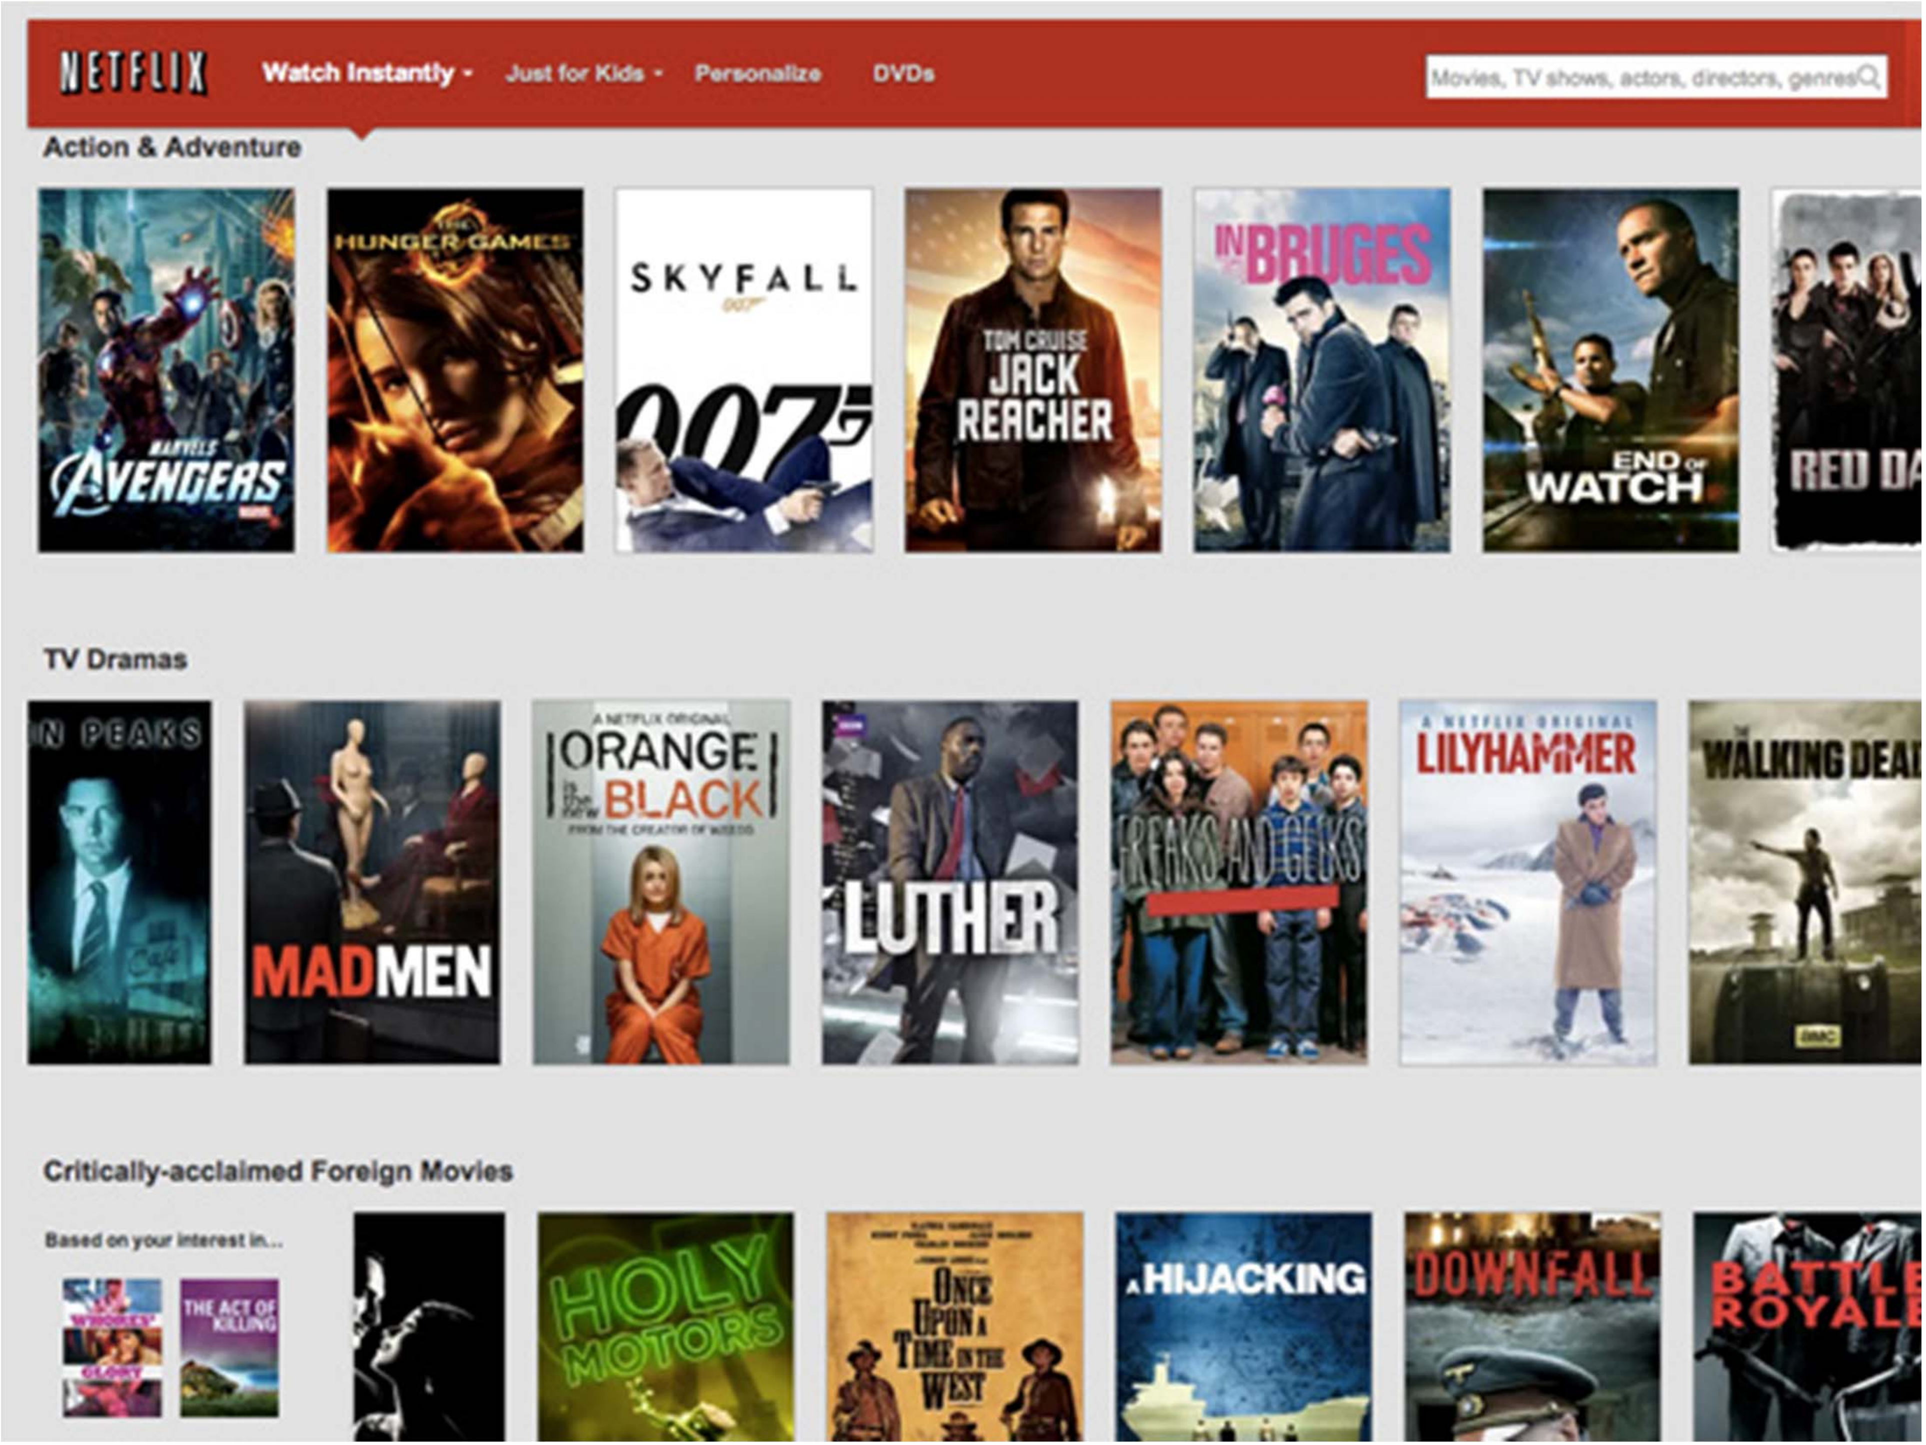
NETFLIX
Watch Instantly
Just for Kids
Personalize
DVDs
Movies, TV shows, actors, directors, genres
Action & Adventure
AVENGERS
THE HUNGER GAMES
SKYFALL
007
TOM CRUISE
JACK REACHER
IN BRUGES
END OF WATCH
TV Dramas
MAD MEN
ORANGE is the new BLACK
LUTHER
LILYHAMMER
THE WALKING DEAD
Critically-acclaimed Foreign Movies
Based on your interest in...
THE ACT OF KILLING
HOLY MOTORS
ONCE UPON A TIME IN THE WEST
A HIJACKING
DOWNFALL

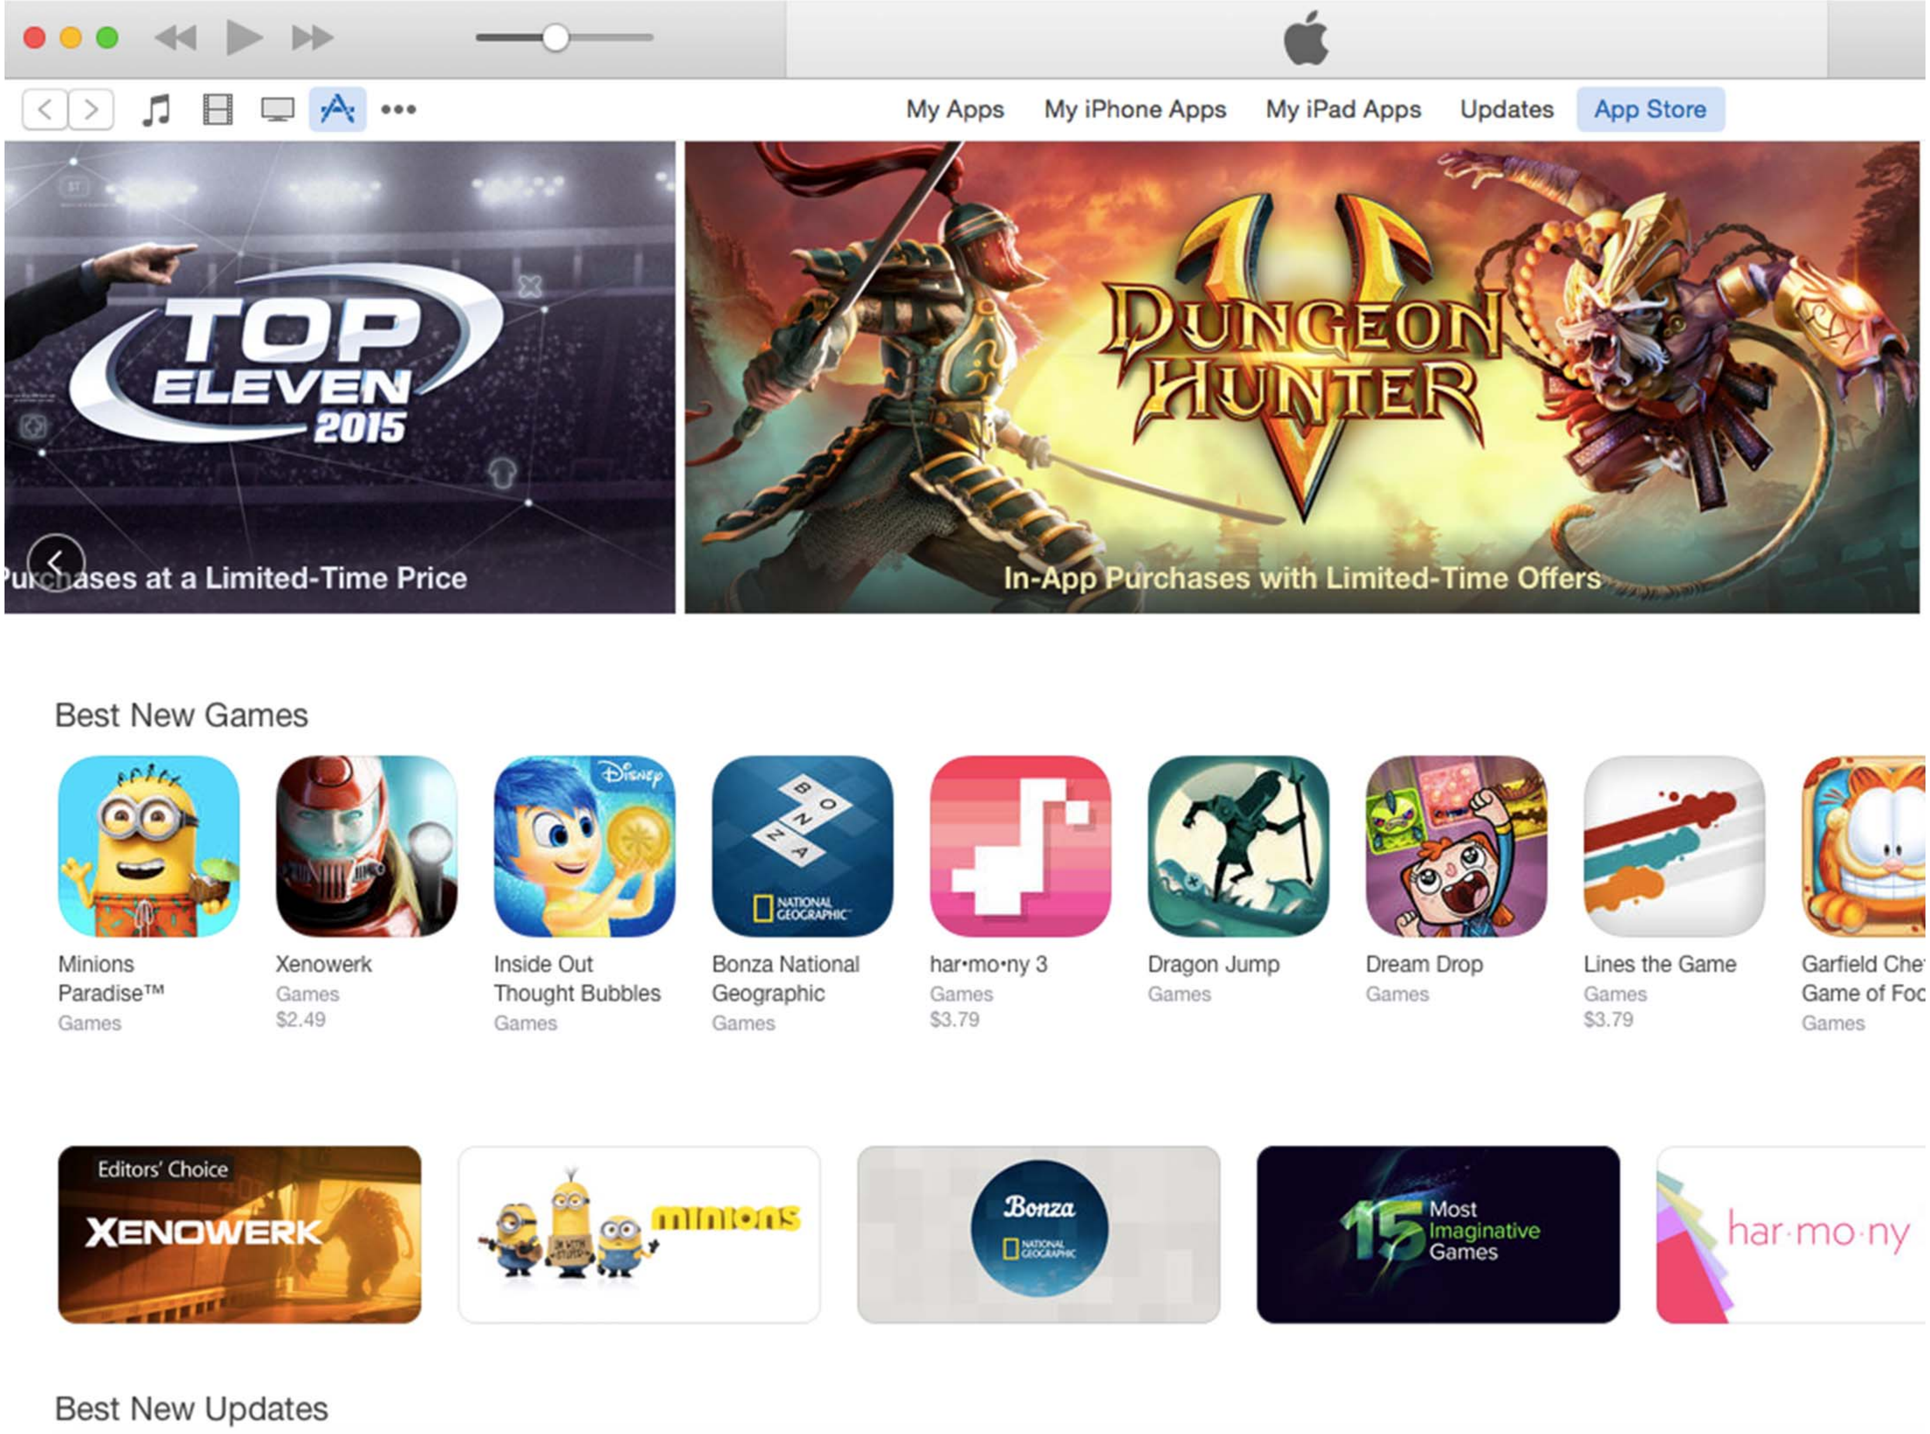

My Apps
My iPhone Apps
My iPad Apps
Updates
App Store
TOP ELEVEN 2015
Purchases at a Limited-Time Price
DUNGEON HUNTER V
In-App Purchases with Limited-Time Offers
Best New Games
Minions Paradise™
Games
Xenowerk
Games
$2.49
Inside Out Thought Bubbles
Games
Bonza National Geographic
Games
har•mo•ny 3
Games
$3.79
Dragon Jump
Games
Dream Drop
Games
Lines the Game
Games
$3.79
Garfield Che
Game of Foc
Games
Editors' Choice
XENOWERK
minions
Bonza
15 Most Imaginative Games
har·mo·ny
Best New Updates

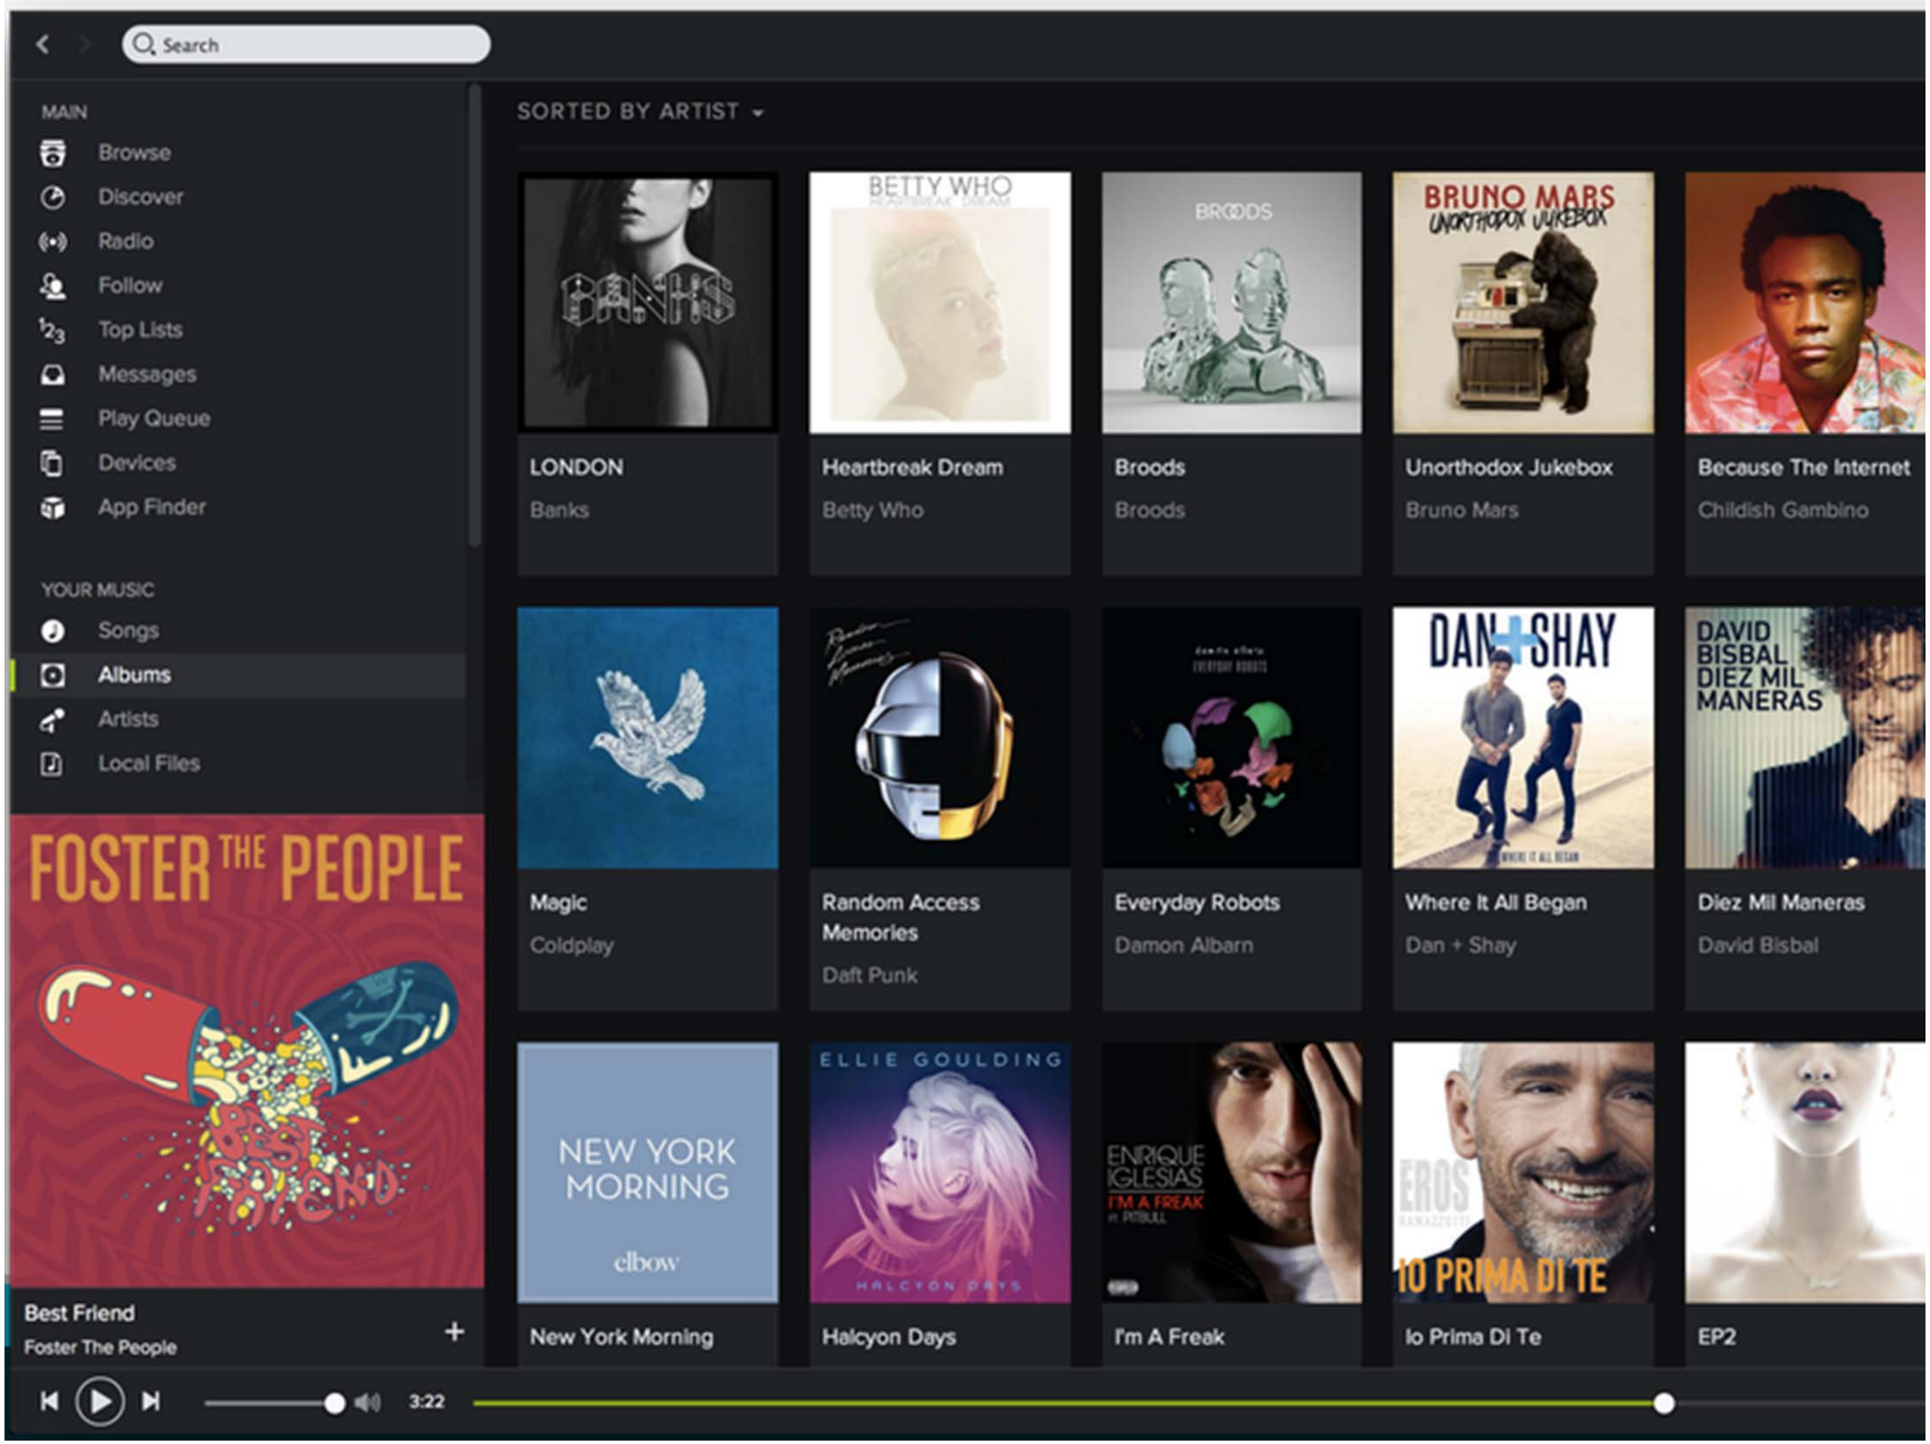
Search
MAIN
Browse
Discover
Radio
Follow
Top Lists
Messages
Play Queue
Devices
App Finder
YOUR MUSIC
Songs
Albums
Artists
Local Files
FOSTER THE PEOPLE
Best Friend
Foster The People
3:22
SORTED BY ARTIST
LONDON
Banks
Heartbreak Dream
Betty Who
Broods
Broods
Unorthodox Jukebox
Bruno Mars
Because The Internet
Childish Gambino
Magic
Coldplay
Random Access Memories
Daft Punk
Everyday Robots
Damon Albarn
Where It All Began
Dan + Shay
Diez Mil Maneras
David Bisbal
New York Morning
Halcyon Days
I'm A Freak
Io Prima Di Te
EP2

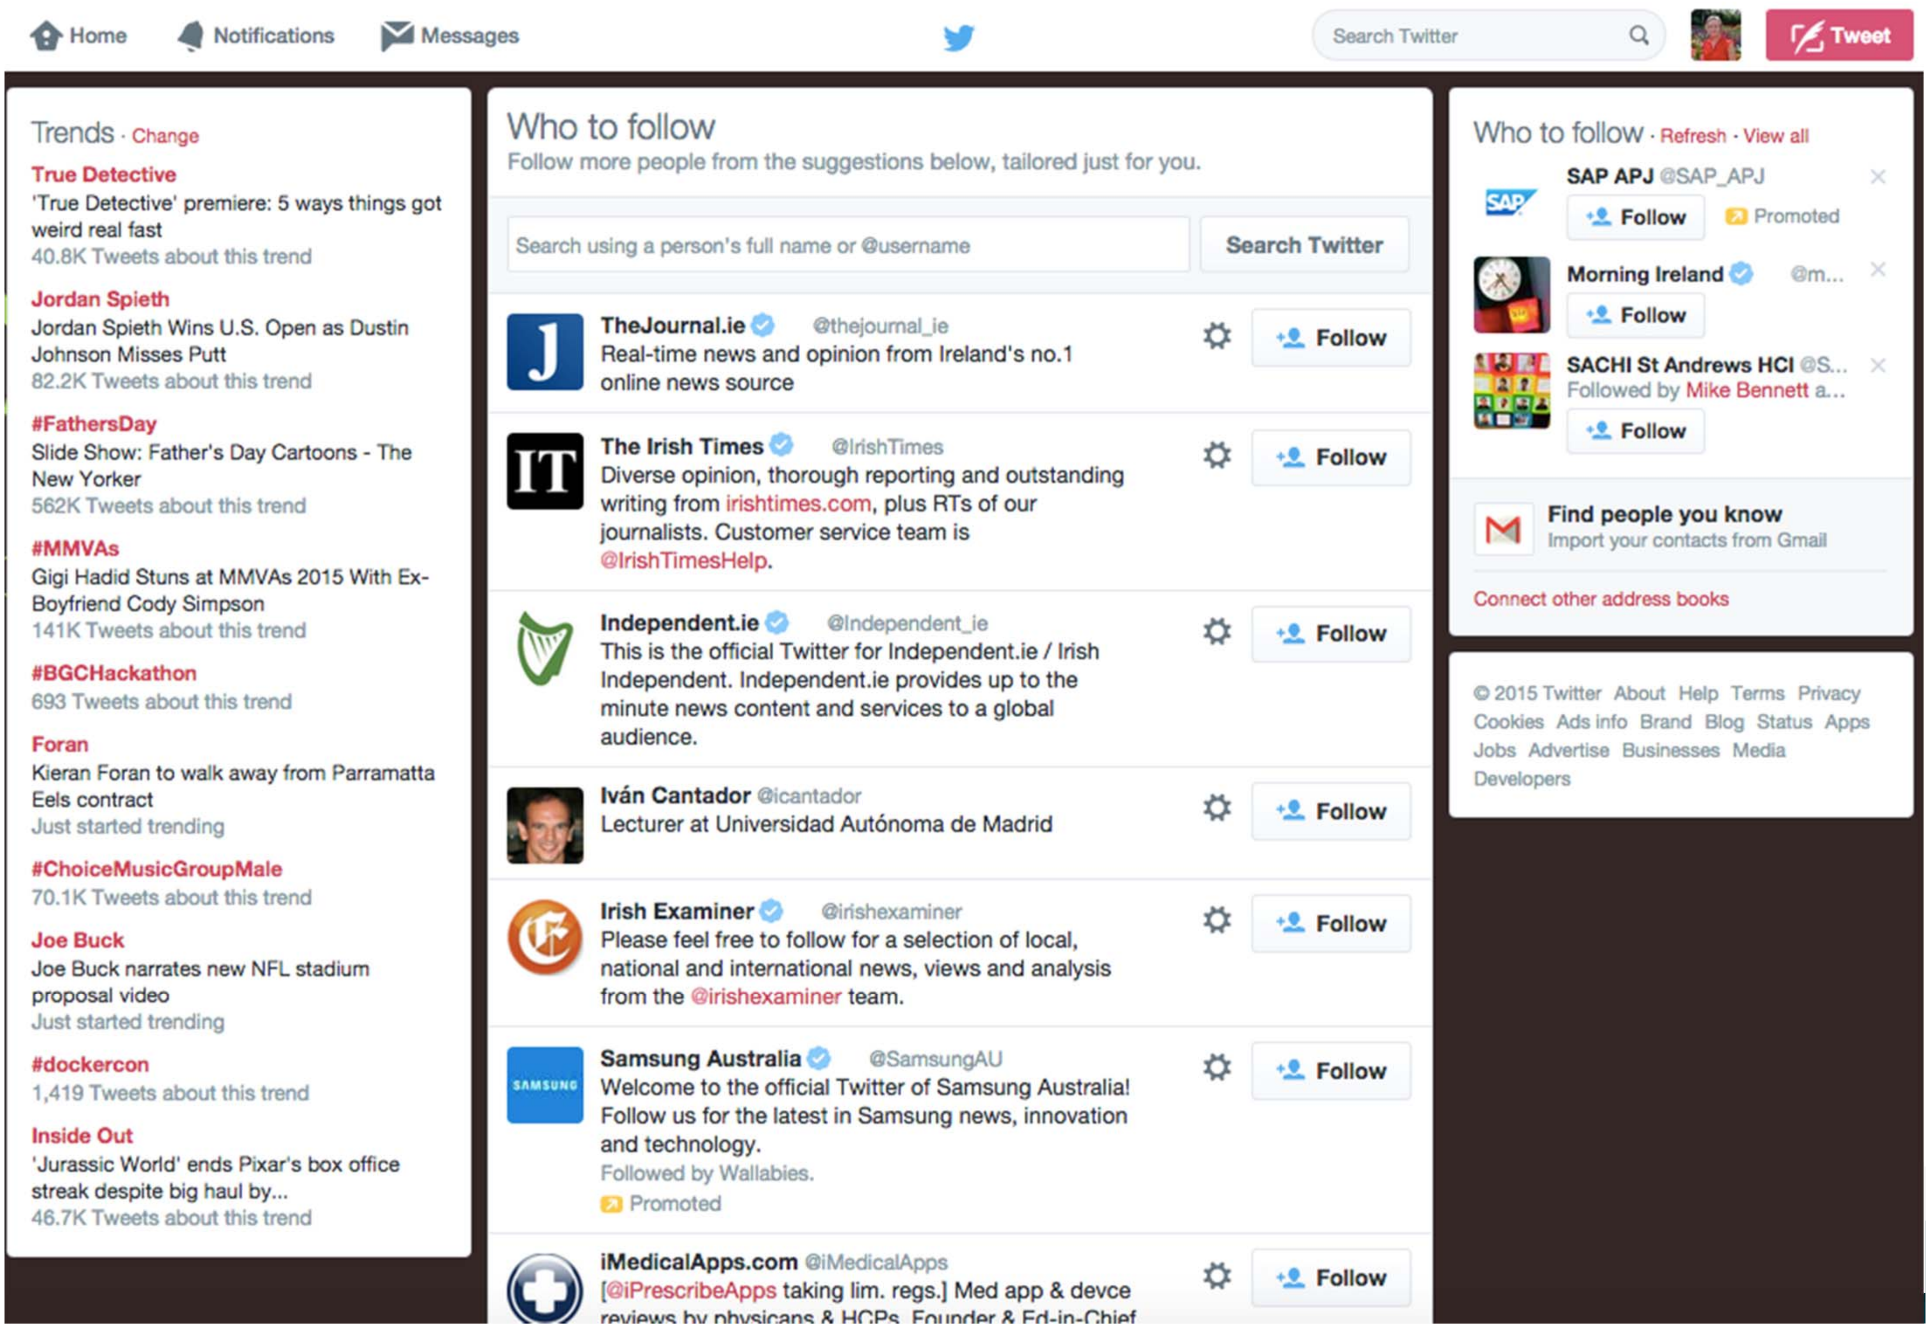
Home
Notifications
Messages
Search Twitter
Tweet
Trends · Change
True Detective
'True Detective' premiere: 5 ways things got weird real fast
40.8K Tweets about this trend
Jordan Spieth
Jordan Spieth Wins U.S. Open as Dustin Johnson Misses Putt
82.2K Tweets about this trend
#FathersDay
Slide Show: Father's Day Cartoons - The New Yorker
562K Tweets about this trend
#MMVAs
Gigi Hadid Stuns at MMVAs 2015 With Ex-Boyfriend Cody Simpson
141K Tweets about this trend
#BGCHackathon
693 Tweets about this trend
Foran
Kieran Foran to walk away from Parramatta Eels contract
Just started trending
#ChoiceMusicGroupMale
70.1K Tweets about this trend
Joe Buck
Joe Buck narrates new NFL stadium proposal video
Just started trending
#dockercon
1,419 Tweets about this trend
Inside Out
'Jurassic World' ends Pixar's box office streak despite big haul by...
46.7K Tweets about this trend
Who to follow
Follow more people from the suggestions below, tailored just for you.
Search using a person's full name or @username
Search Twitter
TheJournal.ie @thejournal_ie
Real-time news and opinion from Ireland's no.1 online news source
Follow
The Irish Times @IrishTimes
Diverse opinion, thorough reporting and outstanding writing from irishtimes.com, plus RTs of our journalists. Customer service team is @IrishTimesHelp.
Follow
Independent.ie @Independent_ie
This is the official Twitter for Independent.ie / Irish Independent. Independent.ie provides up to the minute news content and services to a global audience.
Follow
Iván Cantador @icantador
Lecturer at Universidad Autónoma de Madrid
Follow
Irish Examiner @irishexaminer
Please feel free to follow for a selection of local, national and international news, views and analysis from the @irishexaminer team.
Follow
Samsung Australia @SamsungAU
Welcome to the official Twitter of Samsung Australia! Follow us for the latest in Samsung news, innovation and technology.
Followed by Wallabies.
Promoted
Follow
iMedicalApps.com @iMedicalApps
[@iPrescribeApps taking lim. regs.] Med app & devce
Follow
Who to follow · Refresh · View all
SAP APJ @SAP_APJ
Follow
Promoted
Morning Ireland @m...
Follow
SACHI St Andrews HCI @S...
Followed by Mike Bennett a...
Follow
Find people you know
Import your contacts from Gmail
Connect other address books
© 2015 Twitter About Help Terms Privacy
Cookies Ads info Brand Blog Status Apps
Jobs Advertise Businesses Media
Developers

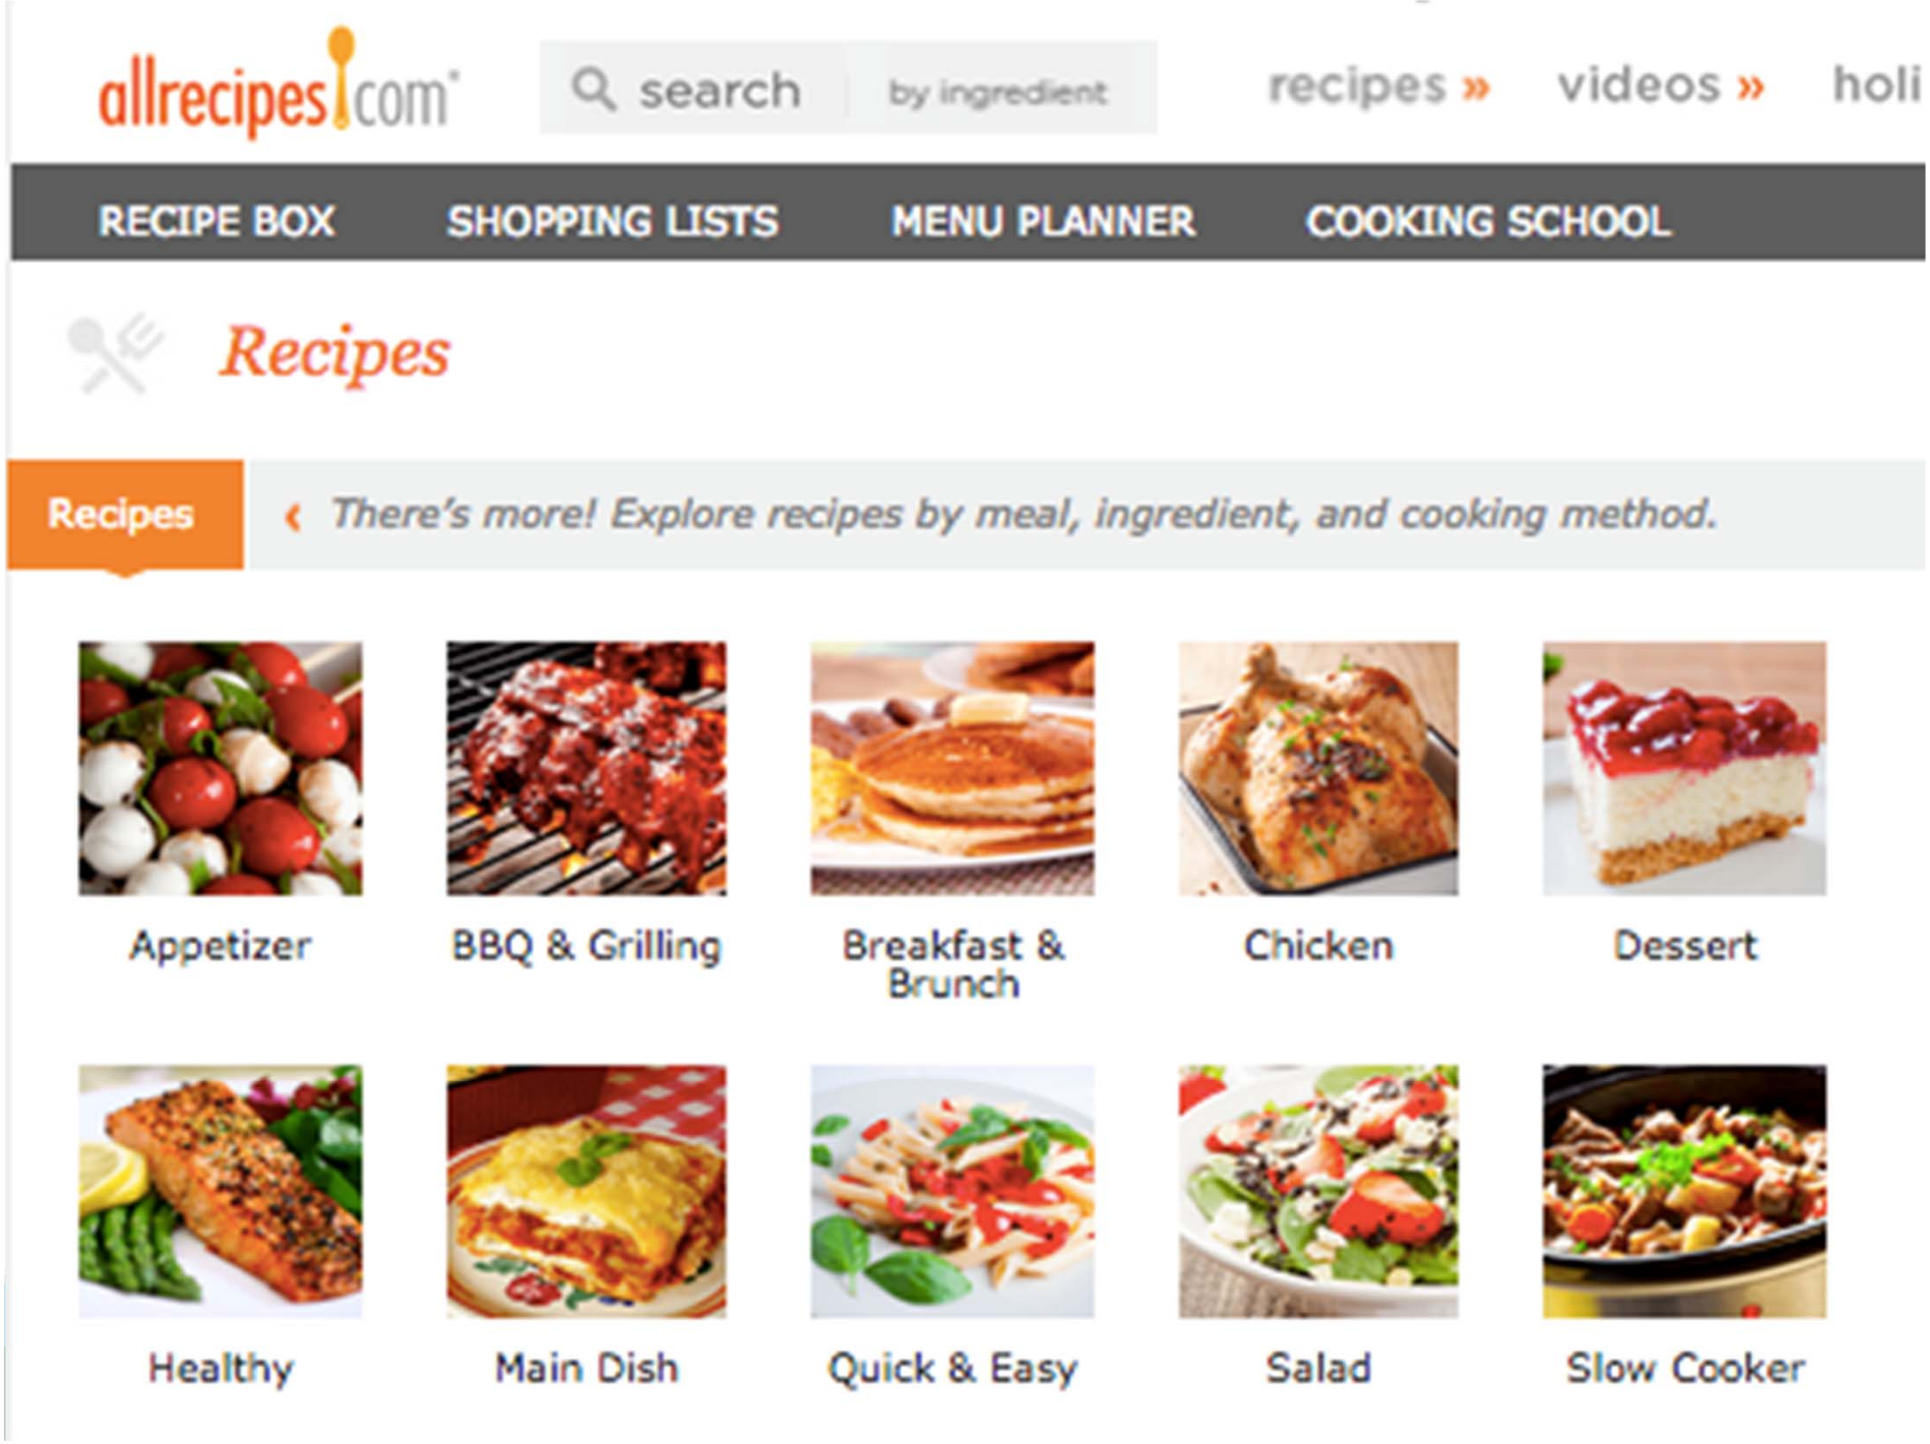
allrecipes.com
search
by ingredient
recipes »
videos »
holi
RECIPE BOX
SHOPPING LISTS
MENU PLANNER
COOKING SCHOOL
Recipes
Recipes
There's more! Explore recipes by meal, ingredient, and cooking method.
Appetizer
BBQ & Grilling
Breakfast & Brunch
Chicken
Dessert
Healthy
Main Dish
Quick & Easy
Salad
Slow Cooker

Information Overload Solutions

- Information Retrieval - assists users to locate online content
- Information Filtering – filters out irrelevant items from a user’s information stream
- Recommender Systems – highlight valuable items in a user's information stream

Personalization

Customization vs. Personalization

- The differentiator is the control over profile creation and presentation interface.
  - Customization = users control customization by specifying their preferences or requirements
  - Personalization = user profiles are created and service is personalized automatically by the system with minimal explicit control by the user

Personalization is…

- “… the ability to provide content and *services tailored to individuals* based on knowledge about their preferences and behavior” (tools and information)

- “… the capability to *customize customer communication* based on preferences and behaviors at the time of interaction [with the customer]” (communication)

- “… about *building customer loyalty and meaningful one-to-one relationship*; by understanding the needs of each individual and helping satisfy a goal that efficiently and knowledgeably addresses the individual’s need in a given context” (customer relationships)

Amazon and Personalization

- Jeff Bezos, Amazon CEO is credited with changing the way the world shops
- Deployed personalization on Amazon
  - *“If I have 3 million customers on the Web, I should have 3 million stores on the Web”*

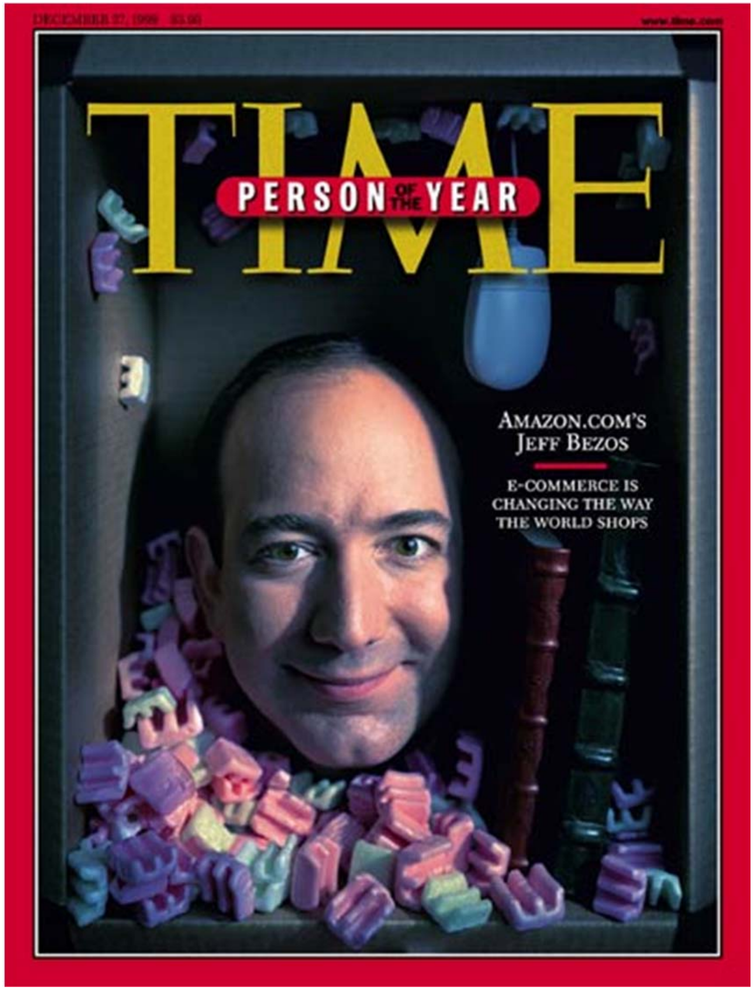

For example…..

- Amazon maintains profiles of all shoppers based on products
  - Purchased products, feedback, wish list, items browsed, …
- Rather than showing random or popular items, Amazon provides personalized recommendations for items to purchase

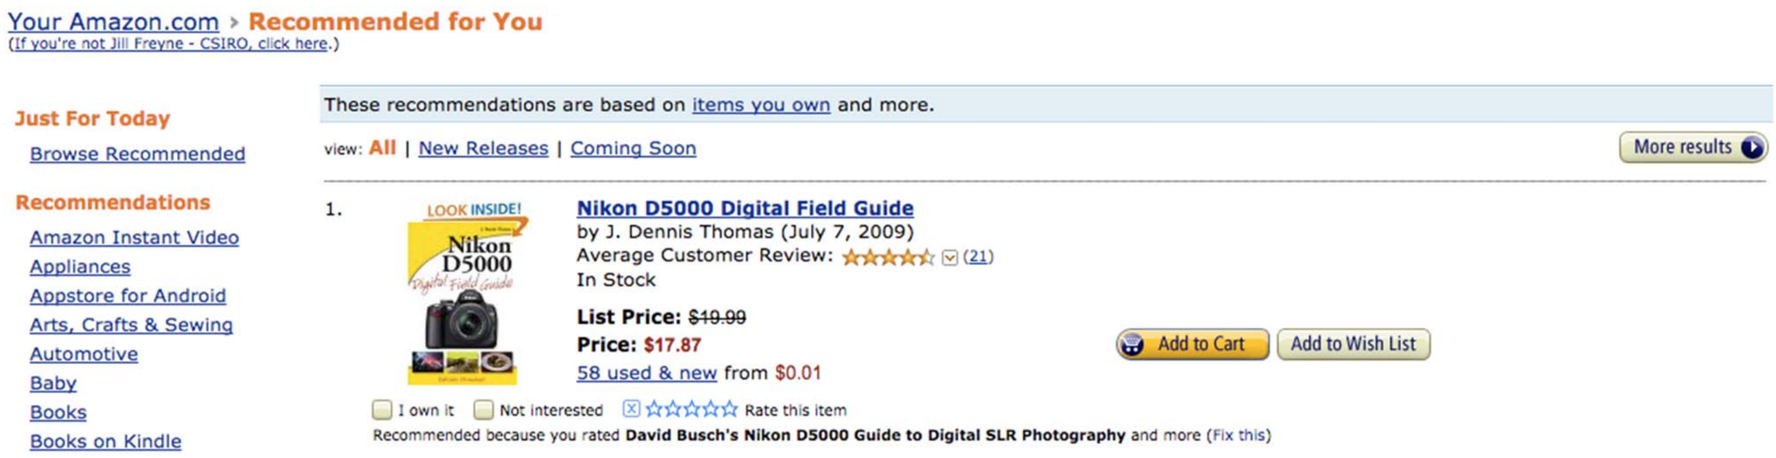

User Modeling and Personalization

- People leave traces on the internet...
  - What pages do they visit? How long do they visit for?
  - What search queries are they using?
  - What products do they buy?
  - What movies do they download?
  - Who are their online friends?

- User modelling is about making sense of this data
  - to gain an understanding of the characteristics, preferences, and needs of an individual user

- Personalization exploits user models (and context)
  - to filter information and provide personalized services that match the user's needs

Framework for Personalization

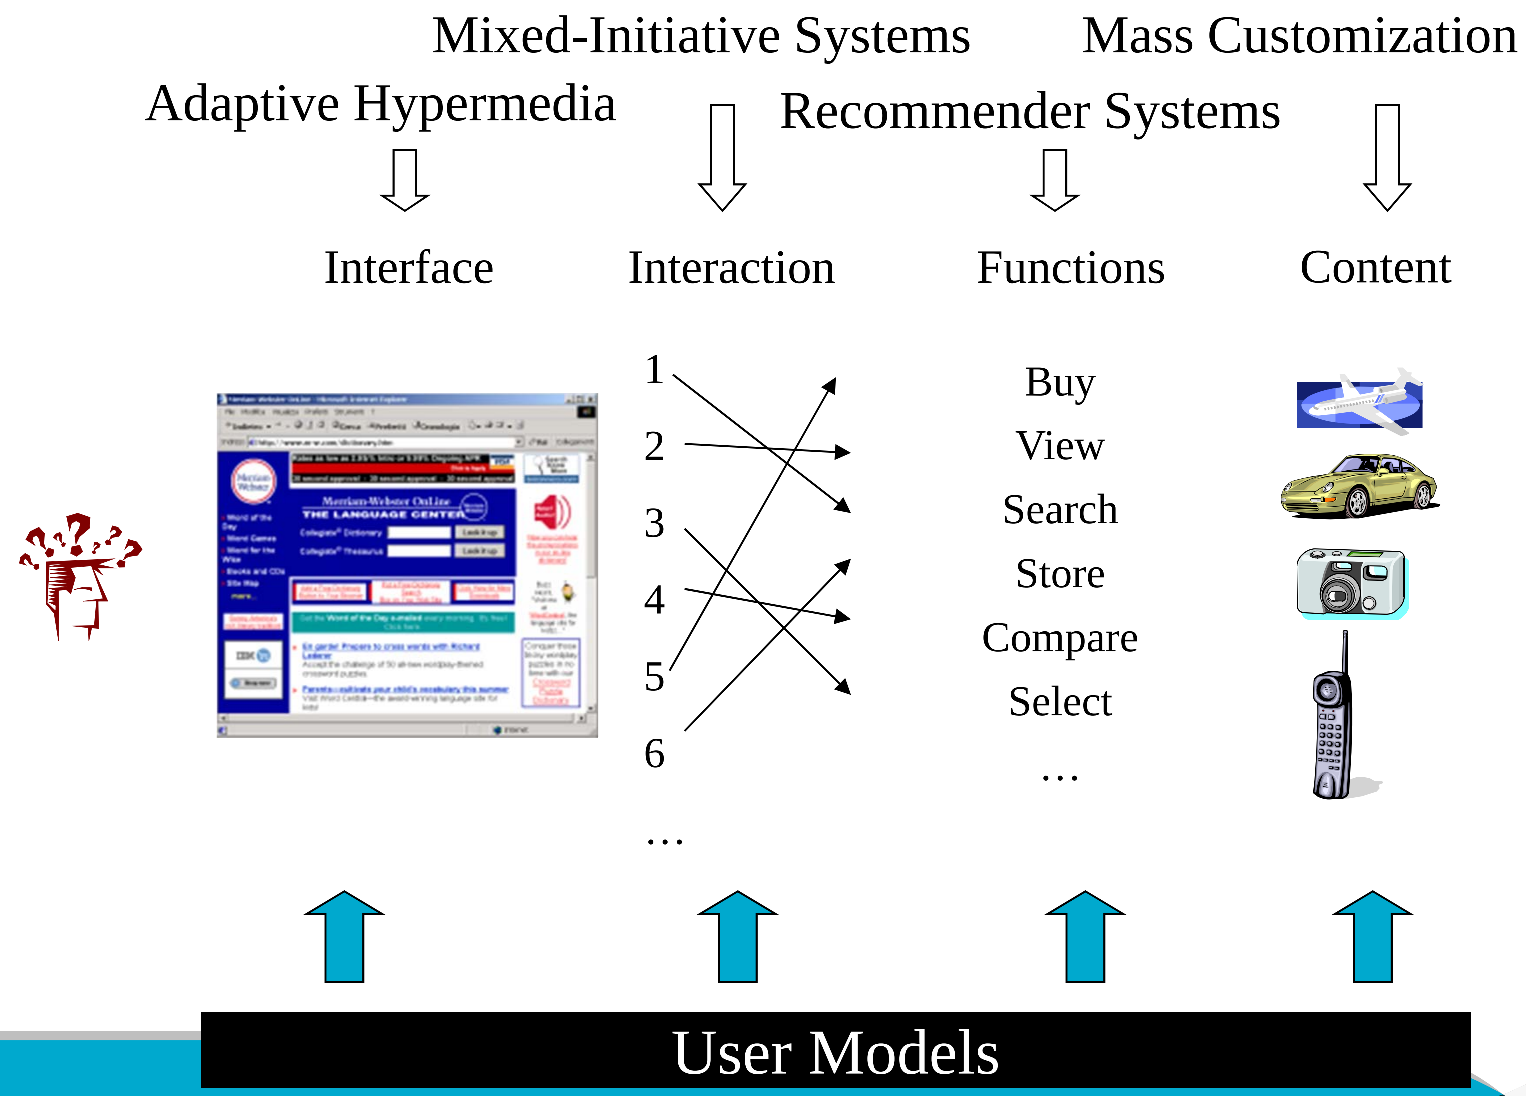

How is Personalization Achieved?

1. Gathering information about the users
   - Explicitly – through direct user input
   - Implicitly – through monitoring user interactions

2. Exploiting this information to create a *user model* or *profile*
   - Dynamic vs. Static
   - Short term vs. Long term

3. Use the model to adapt some aspect of the system to reflect user needs, interests or preferences

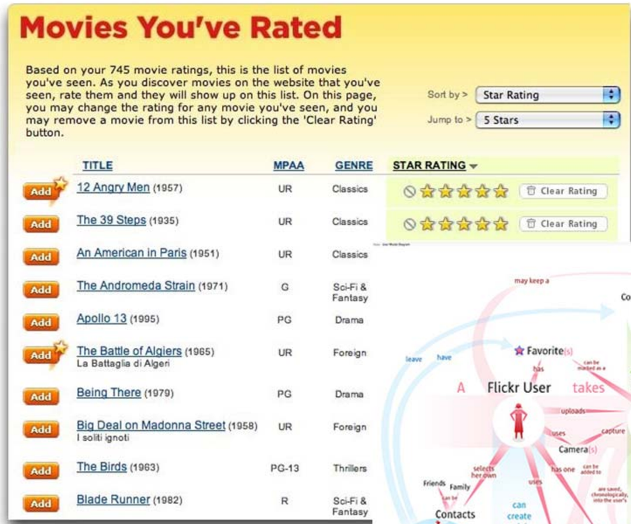

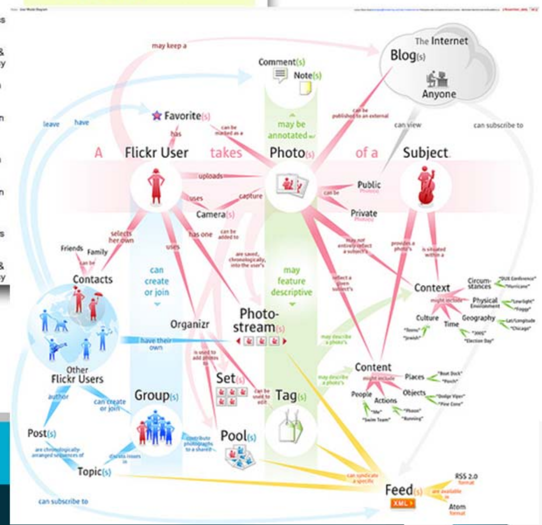

User Model Based Personalization

- 3 stages
  - User information collection
  - User profile construction
  - Exploitation of profile for personalization

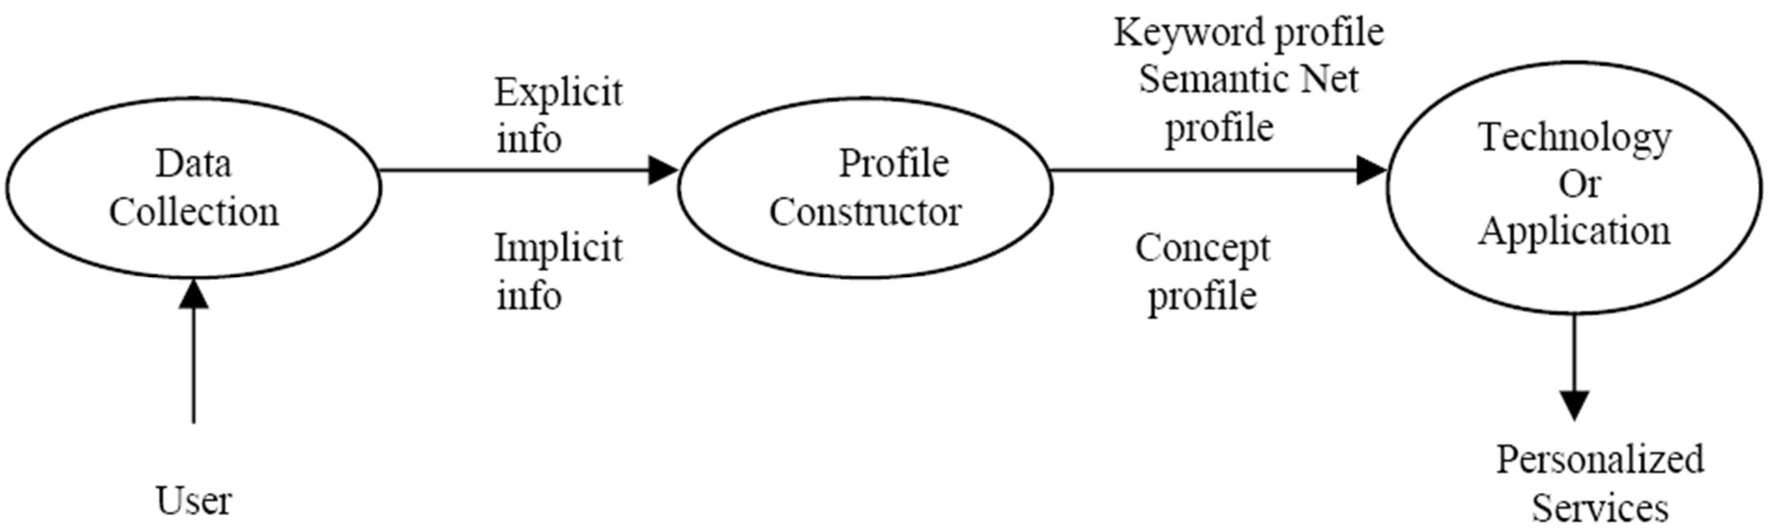

User Modelling for Personalization

- Different systems require different models
  - Sometimes you model the user in terms of their preferences and interests
    - Marketing a product to a user, returning search results, recommending tourist activities
  - Sometimes model user’s knowledge and goals
    - Adaptive educational systems, online tutorials, video lectures
  - Sometimes model fitness, health or medical conditions
- No generic user model structure

Explicit User Data Collection

- Relies on information provided by the user
  - usually through forms, drop down lists, and check boxes that allow users to select preferred options
- Often contains demographic information
  - birthday, location, interests, marital status, job …
- Typically accurate but requires time and effort

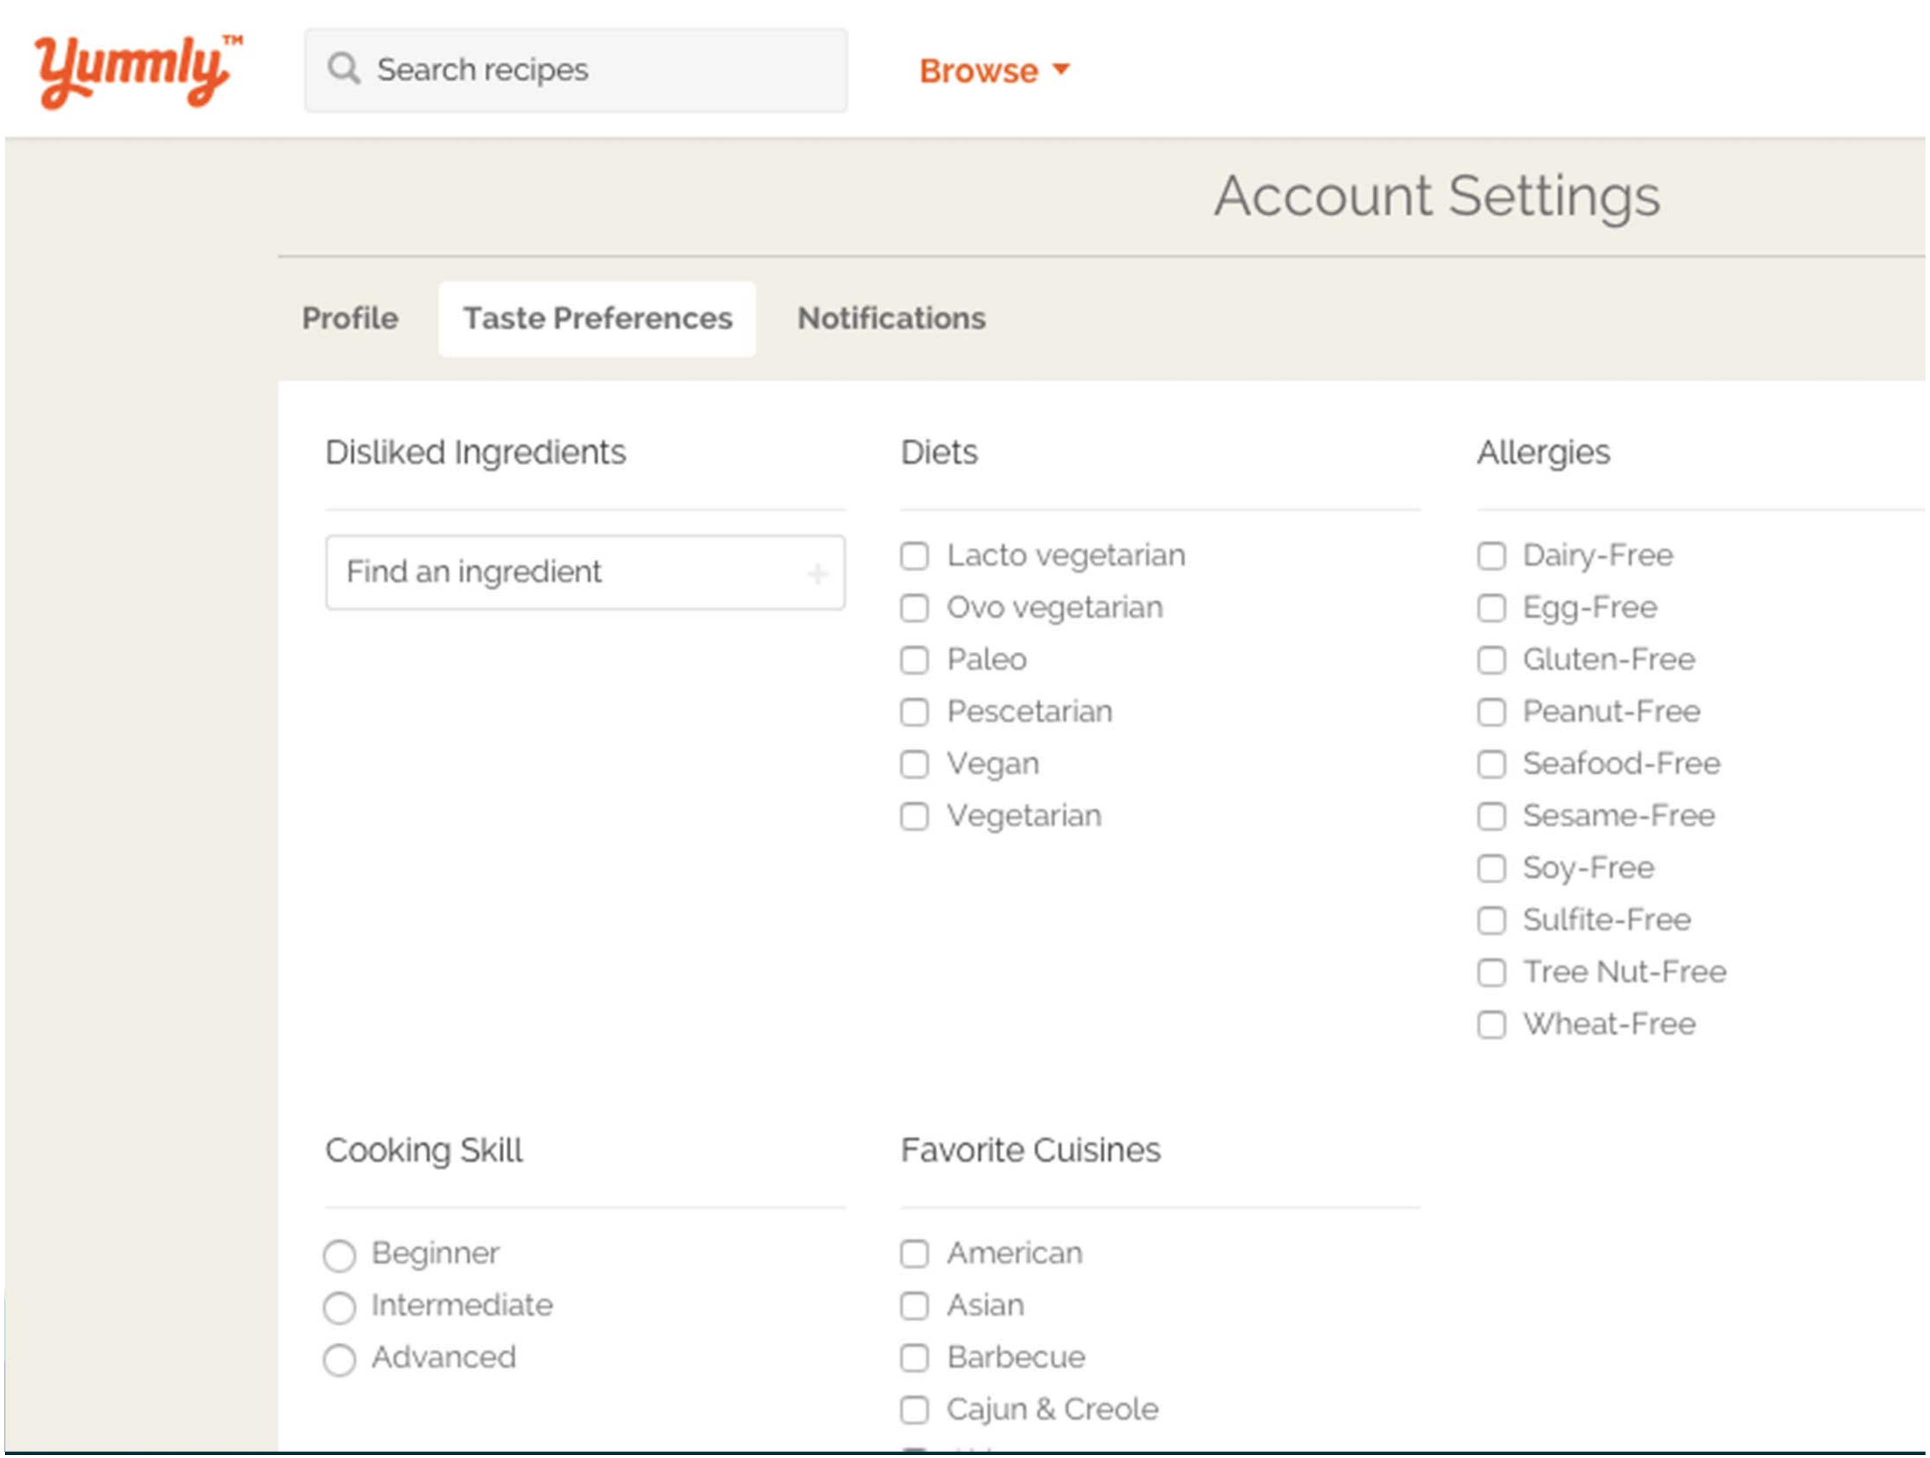
Yummly™
Search recipes
Browse
Account Settings
Profile
Taste Preferences
Notifications
Disliked Ingredients
Find an ingredient
Diets
Lacto vegetarian
Ovo vegetarian
Paleo
Pescetarian
Vegan
Vegetarian
Allergies
Dairy-Free
Egg-Free
Gluten-Free
Peanut-Free
Seafood-Free
Sesame-Free
Soy-Free
Sulfite-Free
Tree Nut-Free
Wheat-Free
Cooking Skill
Beginner
Intermediate
Advanced
Favorite Cuisines
American
Asian
Barbecue
Cajun & Creole

Explicit User Data Collection

- Commercial systems often look for explicit feedback, mostly ratings of symbolic scores

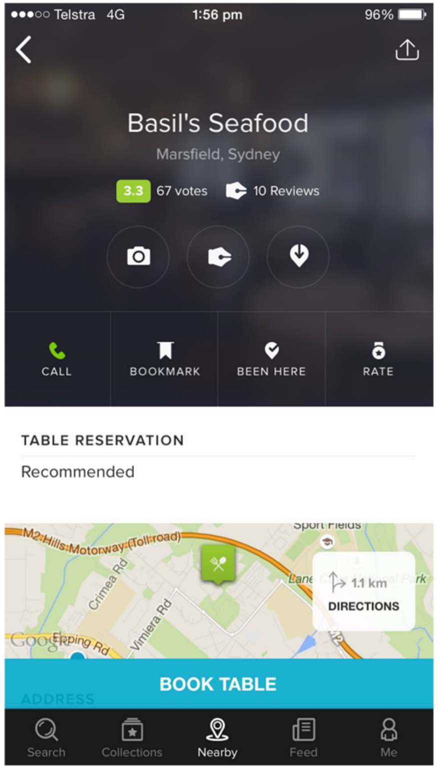

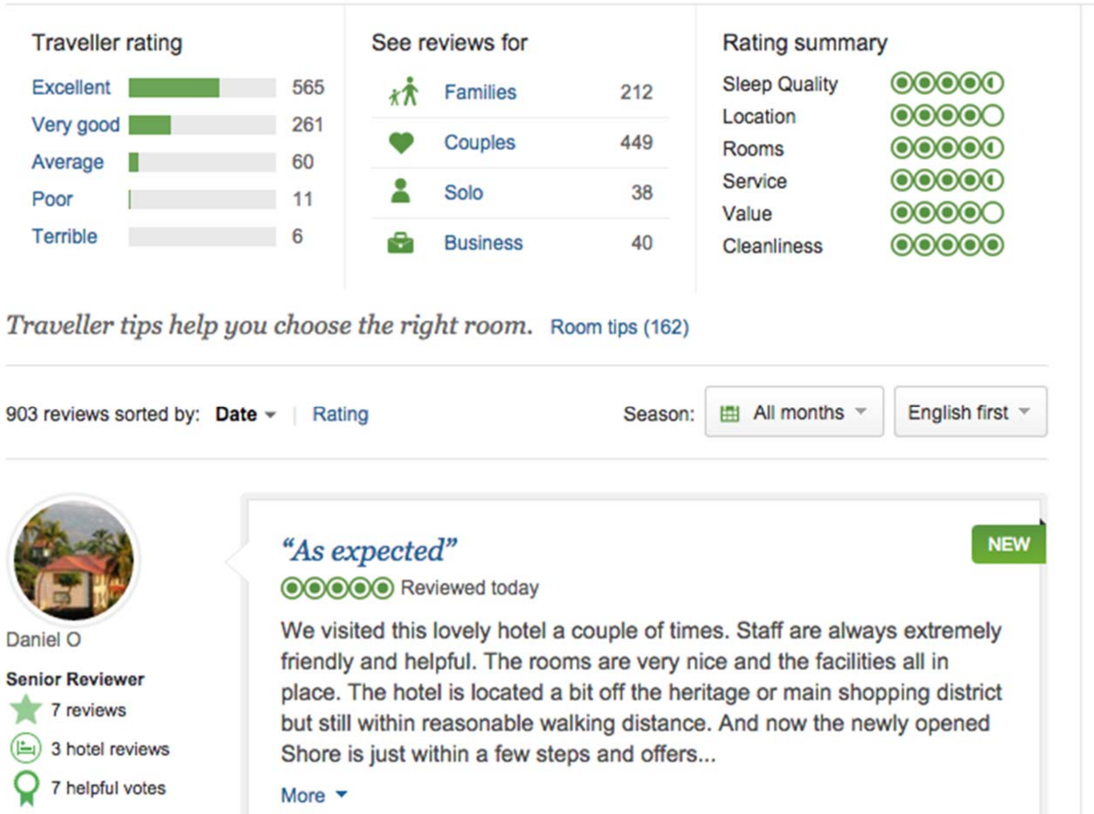

Implicit User Data Collection

- Derives user modeling data from observable user behavior
  - Monitor users interactions with the system and with other users
  - Learn/mine the required user data
- Examples
  - Browser cache, proxy servers, search logs, purchased items, examined products, bookmarked pages, links sent to friends, preferred brands, restaurants rated, followers/followees on social media, GPS data logged
- Typically less accurate than explicit data but does not require any extra-effort from users

Hybrid Data Collection

- Combines explicit and implicit methods
  - to leverage the benefits of both methods
- Typically achieves the highest accuracy
  - Many things are learned implicitly
  - User feedback is sought for uncertain/important data
- Used by many commercial systems

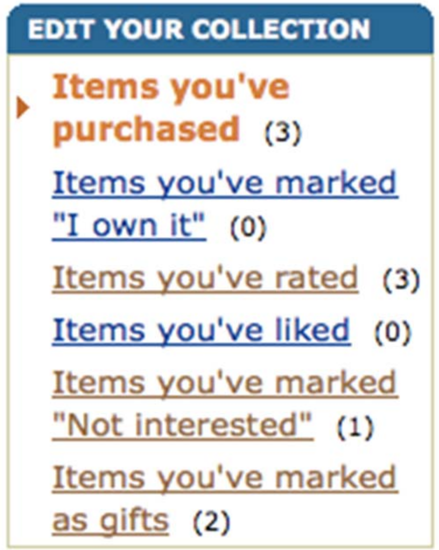

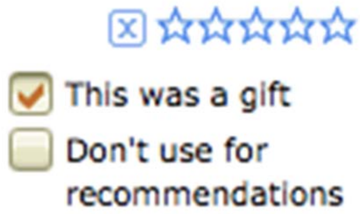

Emotion Based Modeling

- Relatively new direction in user modeling
- Experienced emotions reflect liked/disliked items
  - Explicit (sentiment analysis) and implicit (sensors)
  - Potentially very fine granularity

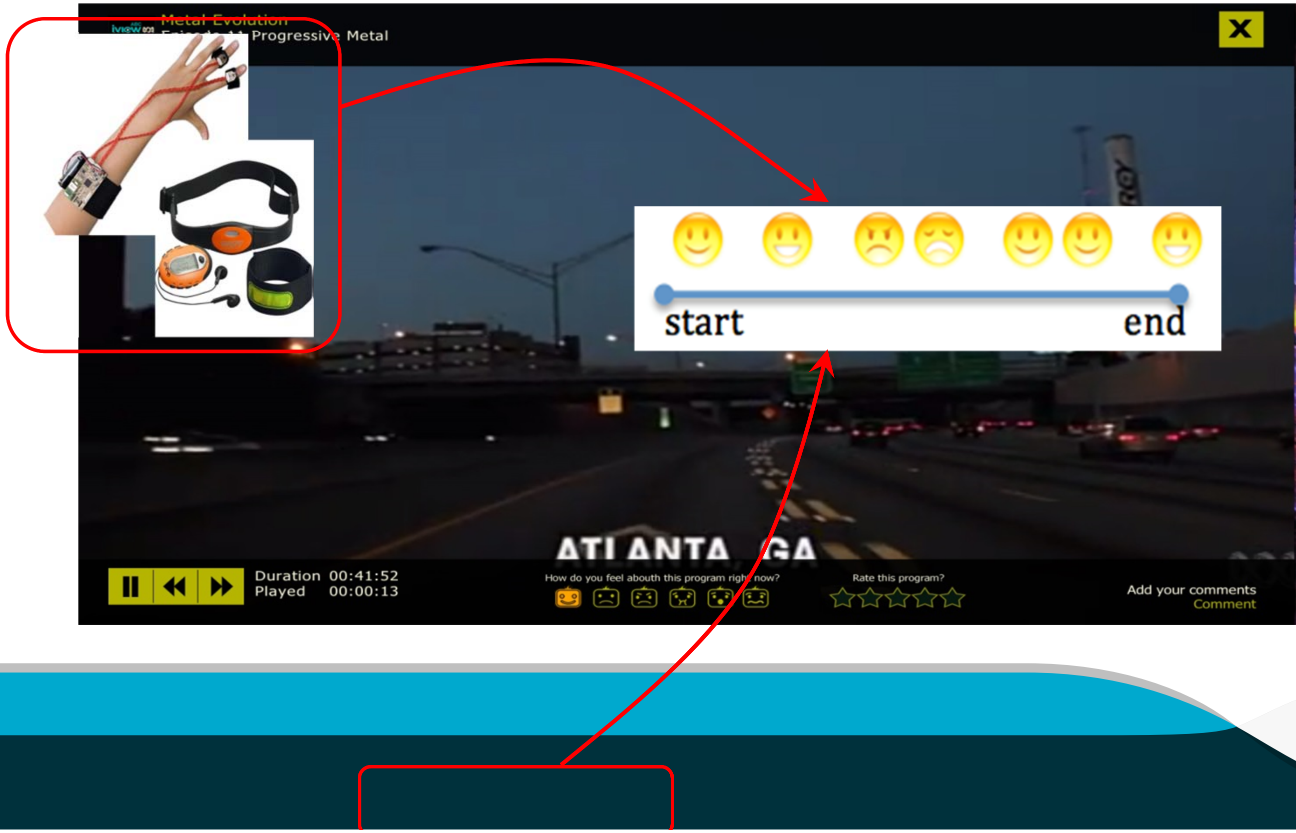

What can be modeled?

- User as an individual
  - Knowledge
  - Interests
  - Goals and motivation
  - Background
  - Personality and traits
  - Interactions with system
- Context-awareness of user models
  - Context and personalization

Knowledge

- .. of a subject or domain, changes over time
- Scalar models
  - Estimate user knowledge as a value
    - Either quantitative (e.g., 0-5) or qualitative (e.g., good, average, poor) scale
  - Often produced by user self evaluation
  - Allows the system to cluster/classify the user and adapt the service accordingly
- Structured (overlay) models
  - Represent user knowledge in fragments of the domain
    - Represents user knowledge as a fragment of the domain model that reflects the expert-level knowledge level.

Interests

- Important for Web information retrieval/filtering systems and for recommender systems

- Most popular approach weighted keywords
  - [(java, 6), (programming, 3), (tutorial, 1),…]

- More powerful: overlay model allows for different areas of interests to be modelled separately
  - News - interest in topics: world, local, sport, technology, …

- Semantics links can enrich the data and compensate for scarcity
  - Buenos Aires is in Argentina → if a user has interest in Buenos Aires news then the probability of interest in Argentinian news is high

Goals and motivation

- Goals represent the user’s immediate purposes
  - Motivation represents the reasons
- What does the user want to achieve?
  - Select product/service
  - Information need
  - Learning goal
  - Decision support
- Often, dynamic and highly context-dependent part of a user model

Background

- Background refers to the users experience outside the core domain of the system
- Usually explicit and stable
  - Does not change across sessions
- Possibilities
  - Profession, job responsibilities, experience in similar domain, language skills, …
- Often used for content customization
  - In encyclopaedias content can be adapted to varying languages and user education
- Basis for stereotypical modelling

Personality and traits

- Individual traits of user personality that define the user as an individual
  - Stable, determined using psychological and behavioral tests
- Examples
  - Personality traits – introvert/extrovert
  - Cognitive styles – holist/serialist
  - Cognitive factors – working memory capacity
  - Susceptibility to persuasion
  - Learning styles

Interaction with system

- Most widely used source of implicit models
  - Easier to obtain than explicit data
- User's feedback to personalized services
- Refine user models
- Close the feedback loop

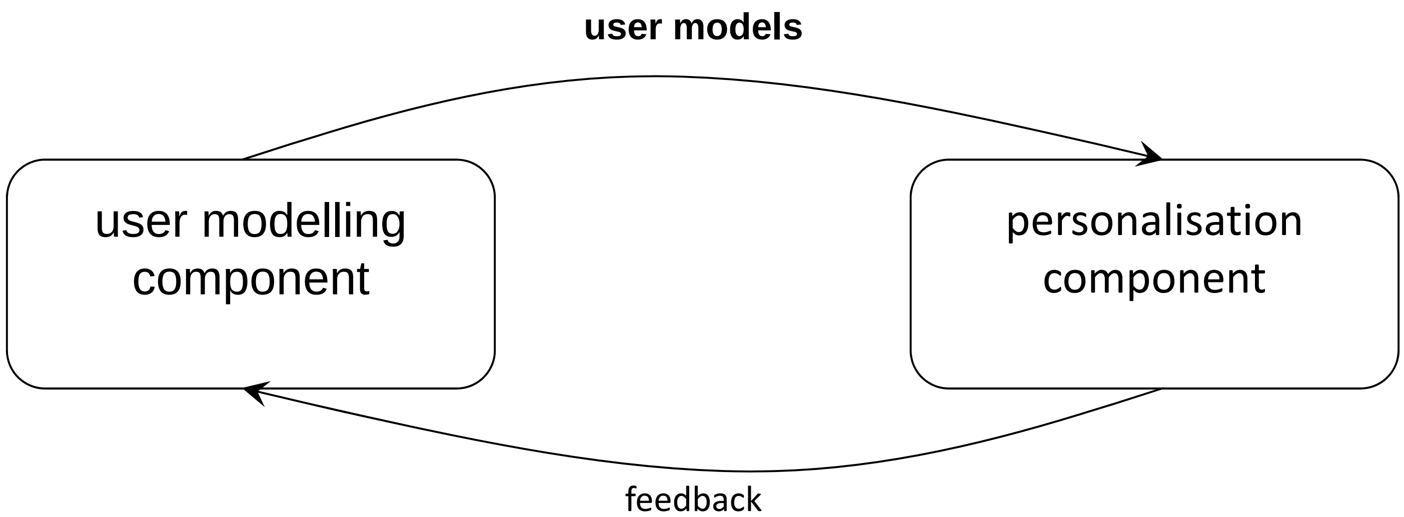

Context-Aware User Models

- Definition of context [Dey]
  - *"Any information that characterizes the situation of an entity. An entity can be a person, place, or object relevant to the interaction between a user and a system, including the user and the system."*
- What can be considered as context?
  - Location of the user, presence of other users, time of day, day of week, weather, temperature, mood, …
- Does context matter?
  - Cooking: alone vs. with kids
  - Music: happy vs. sad
  - Movie: home vs. theater
  - Vacation: summer vs. winter

Context-Aware User Models

- User preferences are not steady, but rather context-dependent
  - From User X Item/Service → Feedback
  - To User x Item/Service X Context → Feedback
- Context affects user feedback
- Only feedback-in-context is meaningful
  - Non-contextualized feedback is unreliable and may add noise
  - Most non-contextualized feedback assumes a default context
    - Default context = most likely context
    - Sometimes true, but often false

Part 2:

Personalization for Information Filtering, Information Access & Content Delivery

Framework for Personalization

Mixed-Initiative Systems

Mass Customization

Adaptive Hypermedia

Recommender Systems

Interface | Interaction | Functions | Content

1
2
3
4
5
6
…

Buy
View
Search
Store
Compare
Select
…

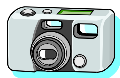

User Models

Personalized Information Access

- The Web has evolved into an ever growing public information source
  - Search engines help users to locate information
  - Information filtering systems hide irrevant information
  - Navigation tools direct users in the online space
- Social Networking sites
  - Extract the Wisdom of the Crowd BUT
  - Massively contribute to information overload

Personalized Search

- Most search engines are generic

JAVA™

Personalized Search

- Tailor the results to the individual user with the aim of better satisfying their needs
- May filter out irrelevant information or identify additional interesting information
- More expensive, needs to
  - Model the user and the document
- Tries to deal with the vocabulary gap
  - Query term do not match document terms
    - E.g. vehicle, car, automobile, …
- Personalization can occur at any of the 3 stages
  - Query entry
  - Search engine retrieval
  - Result ranking

Query Modification

- Terms in the query may not be the most suitable
  - too short, too general, ….
- Search engines often add or replace terms in queries
- If too few documents are returned, more can be found by adding terms to query
  - New keywords are found in the user profiles
    - past queries, words found in previously selected documents, etc

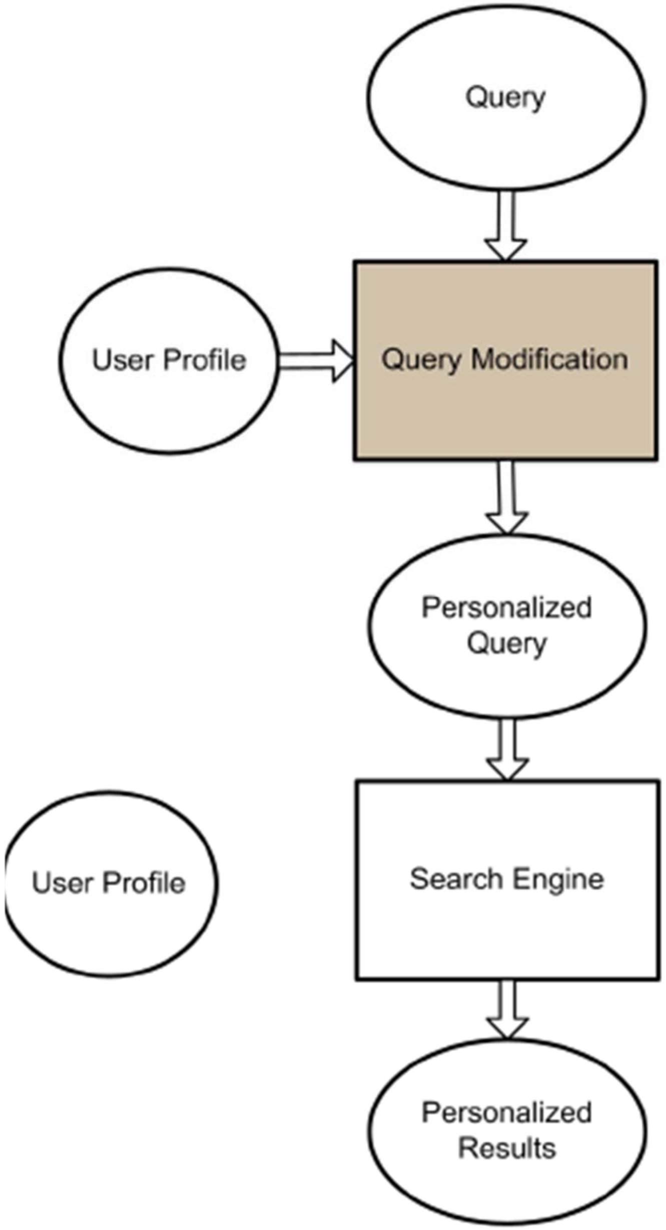

Retrieval Modification

- User profiles are used to score the documents identified as relevant
  - Similarity metric needed
    - VSM and dot product
- Documents are ranked according to the probability of the user to like them
  - Similarity to the user profile
  - Not only to the query
- Very expensive
  - Rarely done

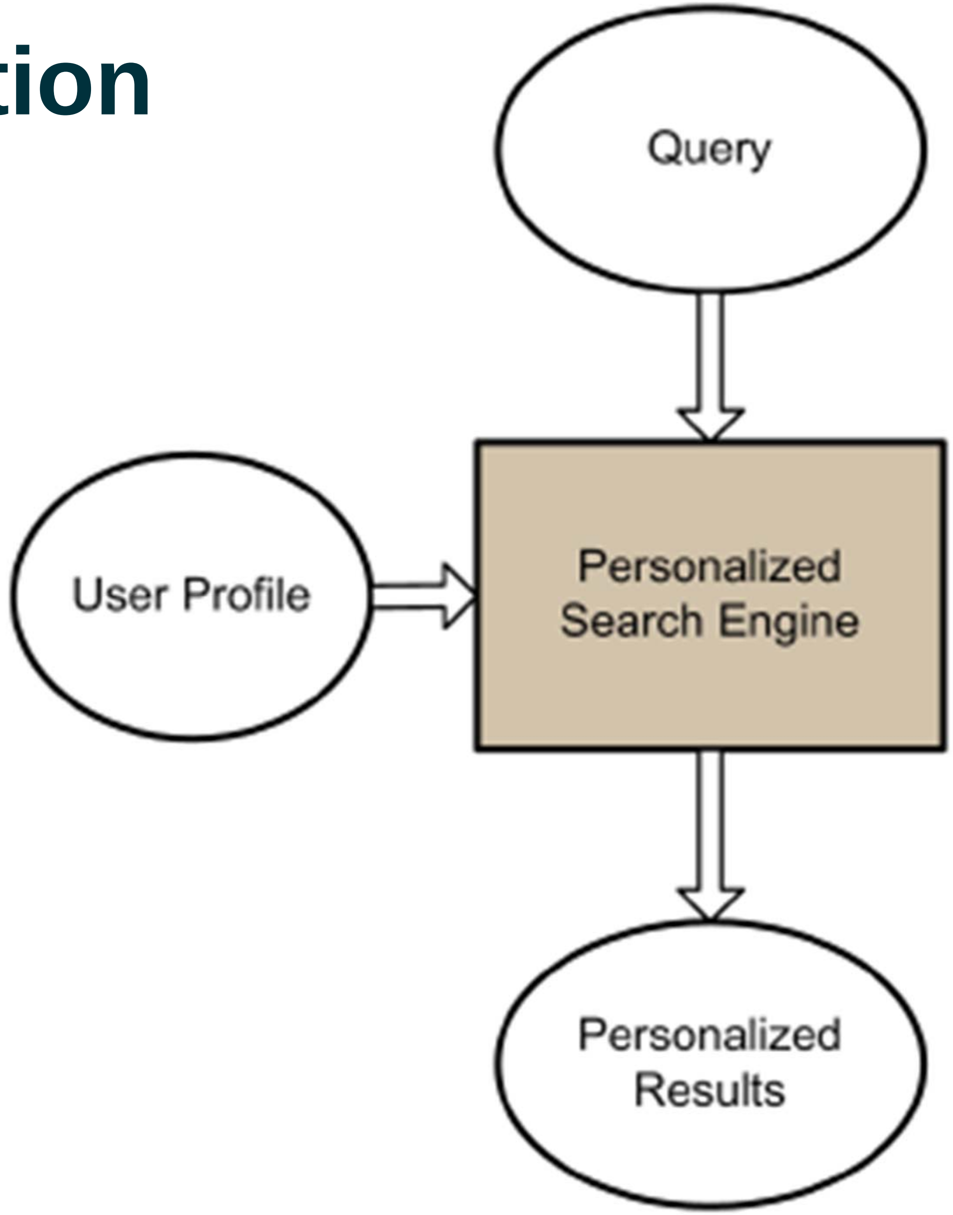

Re-Ranking of Results

- Identify relevant documents
- Compare them with the user profile
- Score them and rank according to the score
- Often performed on client side
  - user profiling component connects to search engine and re-ranks the results before displaying them to the user

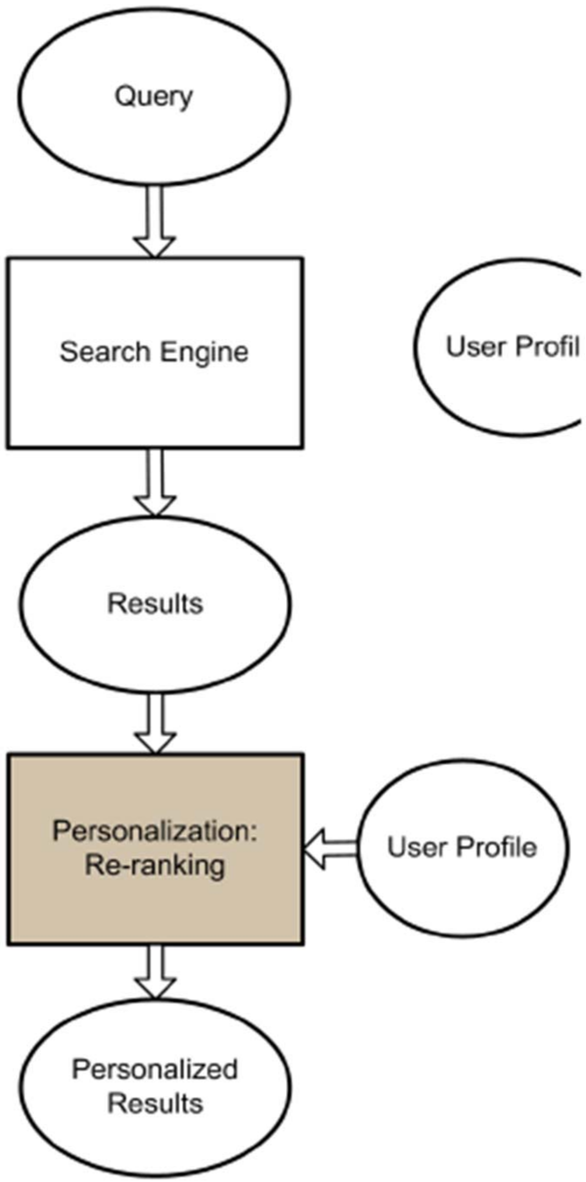

Search Histories

- Personalization determined by past searches
- Users are authenticated by accounts or cookies
  - No dedicated user modeling component
- If users enter short queries the profile could indicate the desired meaning
  - If a user has been entering queries about flights, accommodation, or vaccines, they are probably looking for a travel visa

VISA

Contextual Search

- Just In Time IR – JITIR
  - Find out what the user is doing and find results that help this activity
- Remembrance Agent
  - Monitors users while they use a text editor and retrieves documents related to the text they type
  - While you write your paper, it searches for relevant information even though you do not ask for this
- Watson
  - Tracks users across several applications: MS Office, web browsers etc
  - Starts a search for every open window
  - Results based on the windows or in combination

Location Based Search

- Results are tailored to user's geographical location
  - Even though this is not part of the query
- Done automatically through redirection across engines
  - Often switches the language
- Important for mobile search
- Results automatically invoke Maps

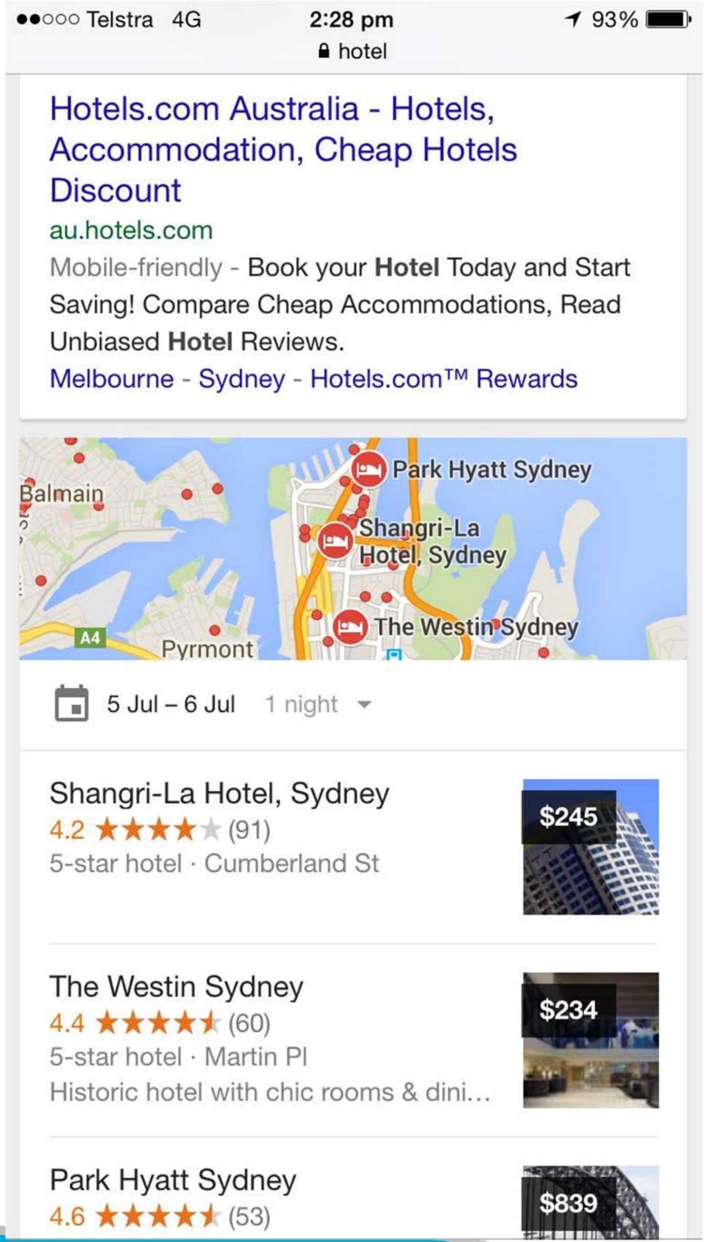

Collaborative Search Engines

- Search engines that personalize to groups and communities rather than to individuals
  - Shows group members the activity of the group as a whole so that members can learn from each other
  - *“Work done by others should leave traces that others can take advantage of when carrying out their work”*
- Often achieve success by
  - Retrieving results according to group profile and logs
  - Highlighting links to popular pages
  - Re-ranking with regards to community preferences and interaction history

HeyStaks Community Search

COMMUNITY IDENTIFICATION

People who share interests behave in the same way as they search and browse

HEYSTAKS

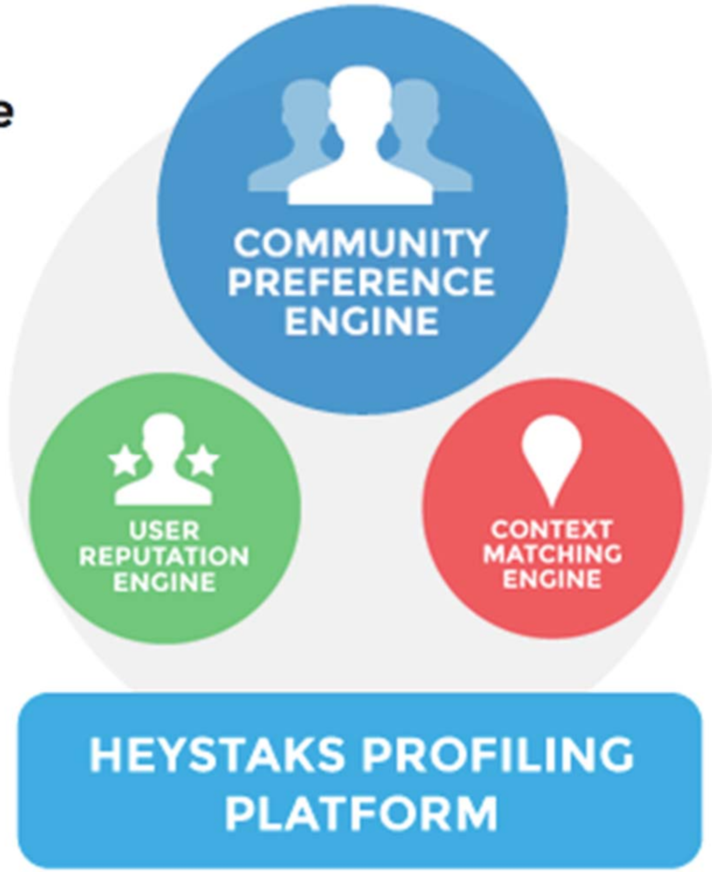

Personalized Navigation Support

- Showing users the way when they browse
- Helping users lost in the Web
  - Direct guidance
  - Sorting lists and links
  - Adding/changing/removing links
  - Adding textual annotations
  - Hiding/highlighting text
  - Increasing font size
  - Adapting images and maps
  - Many more…

Non Digital Objects

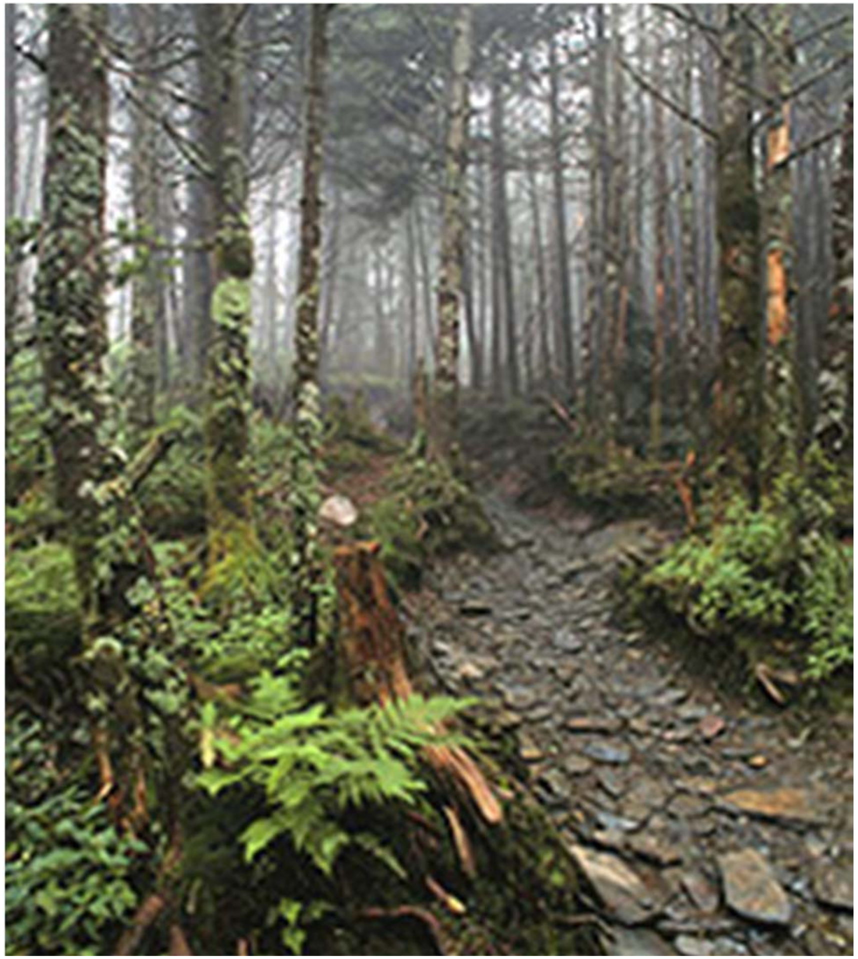

- Non digital objects are history rich through time and use
  - Information attaches itself in the form of wear
- Wear is a gradual and unavoidable change which occurs through interactions with an object
  - Easy to interpret as we are familiar with this form of clue
- Examples:
  - Books with highlighted text
  - Cookbooks with food stains
  - Roads with skid marks
  - Forests with worn paths

Digital Objects

- History poor!
- Users have no idea how a digital object was used in the past
- Collaborative systems show the history of digital objects by adding visual clues that reflect past interaction history.
- Examples:
  - Links followed
  - Pages read/visited
  - Text copied/printed
  - Pages bookmarked
  - Images saved
  - Form fields filled in
  - Tasks completed
  - Many more…

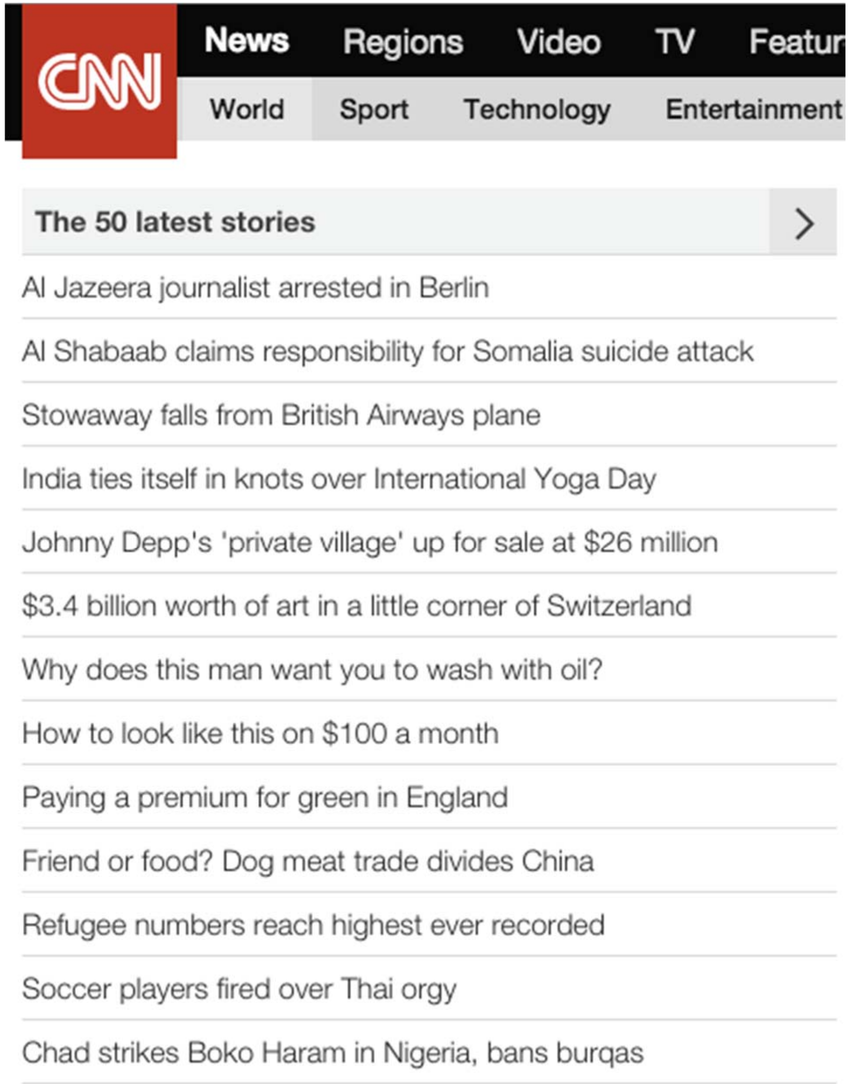

Direct Guidance

- Simplest form of navigation support
- Suggests the best “next step” or “next node” according to the user modeling data
- Two interface options
  - If link to next node is present – emphasise/highlight it for the user
  - If link is not there
    - create a link and add it to current page
    - provide instructions how to reach the next node
- Problem – does not support users who do not follow or ignore the guidance
  - mostly replaced by other techniques

Annotations and Signposts

- Annotations
  - Numbers appended to links to show how many times they have been followed
- Signposts – user feedback regarding past interaction history they've seen
  - Users may comment on pages or on paths in the social navigation display

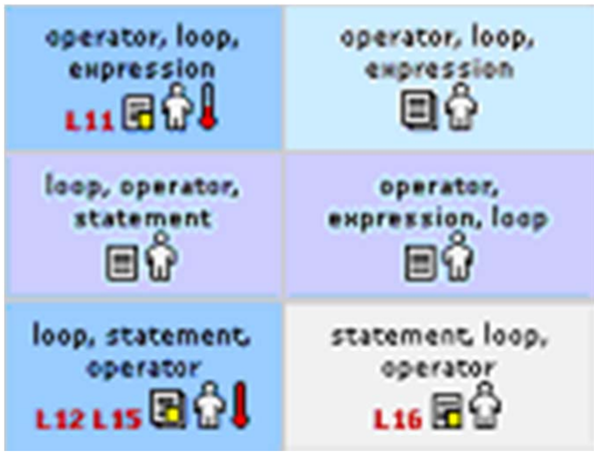

Link Annotation

- Augmenting links with visual cues
  - Give users an insight into the value of content/path behind the link
  - Enrich icons with mouseovers by providing elaborate textual explanation behind the annotation
- Annotation can be used to reflect the degree of relevance and irrelevance

Link Hiding

- Restrict navigation by removing links to content that the is not expected to be relevant/important
  - Consider the links, their visualization, their anchor text, and their functionality
    - Hide the link – remove all visualisation that tells the user that it is a link. The link is still there but hidden
      - “adaptive   is the …….”
      - “adaptive personalization is the …….”
    - Disable the link – remove the functionality of the link so that clicking will do does nothing
    - Remove even the anchor text
- Link hiding is unidirectional
  - Showing links that were hidden is OK
  - Hiding links that were shown frustrates users

Link Ordering

- Sorting or ordering links to prioritise links to relevant/interesting content
  - Reordering is based on user profile
  - User can manually reorder and this informs future decisions.
- Limited applicability
  - Menus and side bars – yes
  - Lists – yes
  - Text-embedded links – no
  - Structured content – no
- Each time the user visits a page it may be different
  - Poor usability

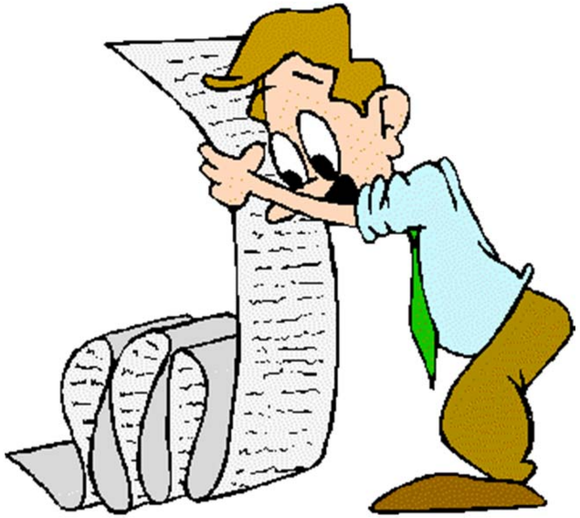

ALICE – Intelligent Tutoring System

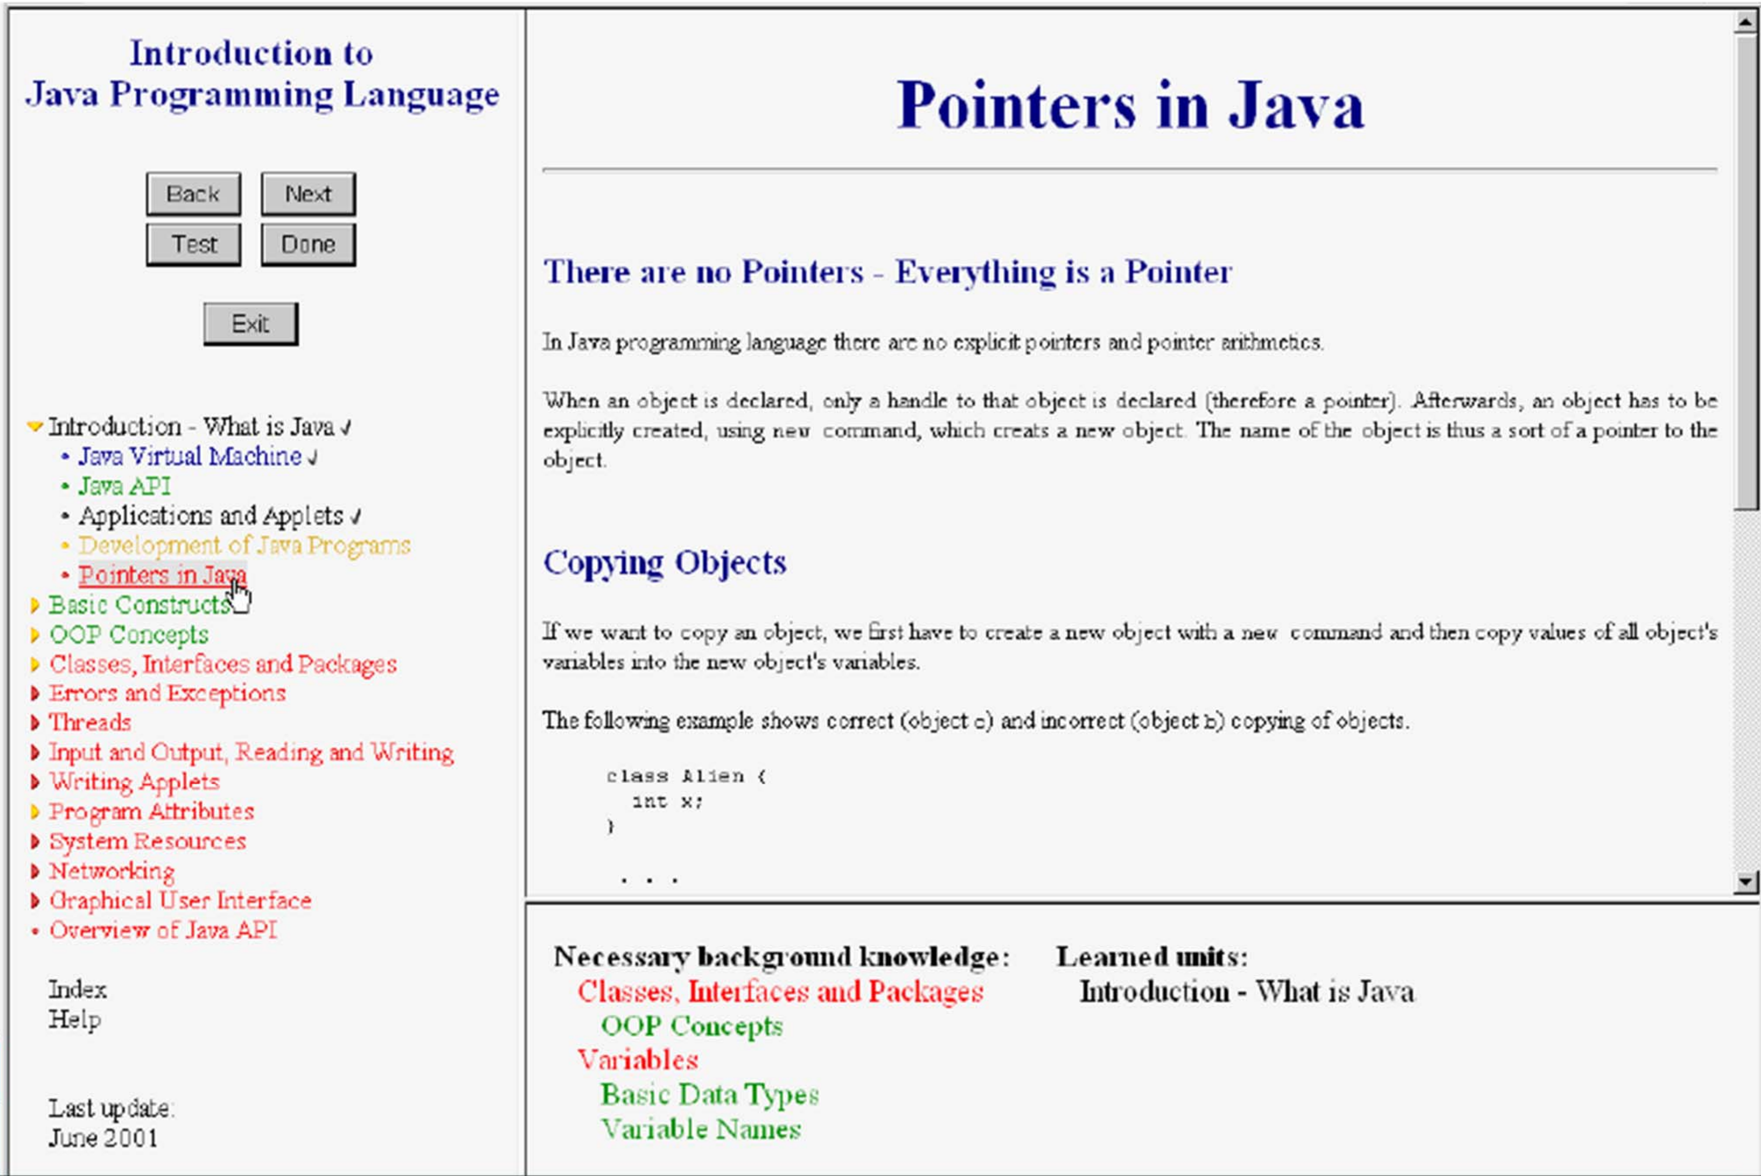

ALICE – Intelligent Tutoring System

- Adapts also to explicit users goal rather than only to user profile
- Uses color to highlight visited content
- Uses font size to show recommended content
  - Font size is determined by a cumulative score of the information behind a link
- Looks more than one step ahead
  - Evaluates outgoing links of the next node to other to see if they are useful
  - Conditional probabilistic model

Social Web Personalization

- Unprecedented volume of information
  - Huge contributor to the information overload
  - But non-negligible consumption medium as well
- Personalization use cases
  - News feed filtering and reordering
  - Preselection of tweets/posts
  - Recommendations of friends/followees
  - Recommendations of events/communities
  - Content ranking on behalf of users
  - Job/company suggestions
  - Many more…

Hurdles for Web Personalization

- Variability of user data
  - User goals change every day and every hour
    - User profiles must be updated to capture this
  - Personalization techniques must correctly identify user goals/need or they will push ill fitting information
- Privacy
  - Users do not want to be monitored and privacy is a critical issue
- The Social Web
- Mobile and ubiquitous use cases

Part 3:
Recommender Systems

Framework for Personalization

Mixed-Initiative Systems   Mass Customization

Adaptive Hypermedia   **Recommender Systems**

Interface	Interaction	Functions	Content

1
2
3
4
5
6
…

Buy
View
Search
Store
Compare
Select
…

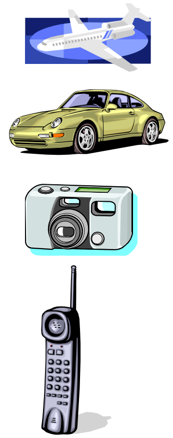

User Models

Recommender Systems

- Recommender Systems help to make choices without sufficient personal experience of the alternatives
  - suggest information items to the users
  - help to decide which product to purchase
- Original definition [1997]: in recommender systems people provide recommendations as inputs, which the system aggregates and directs to appropriate recipients
  - Aggregation recommendations and match with others searching for recommendations

Recommendation Main Steps

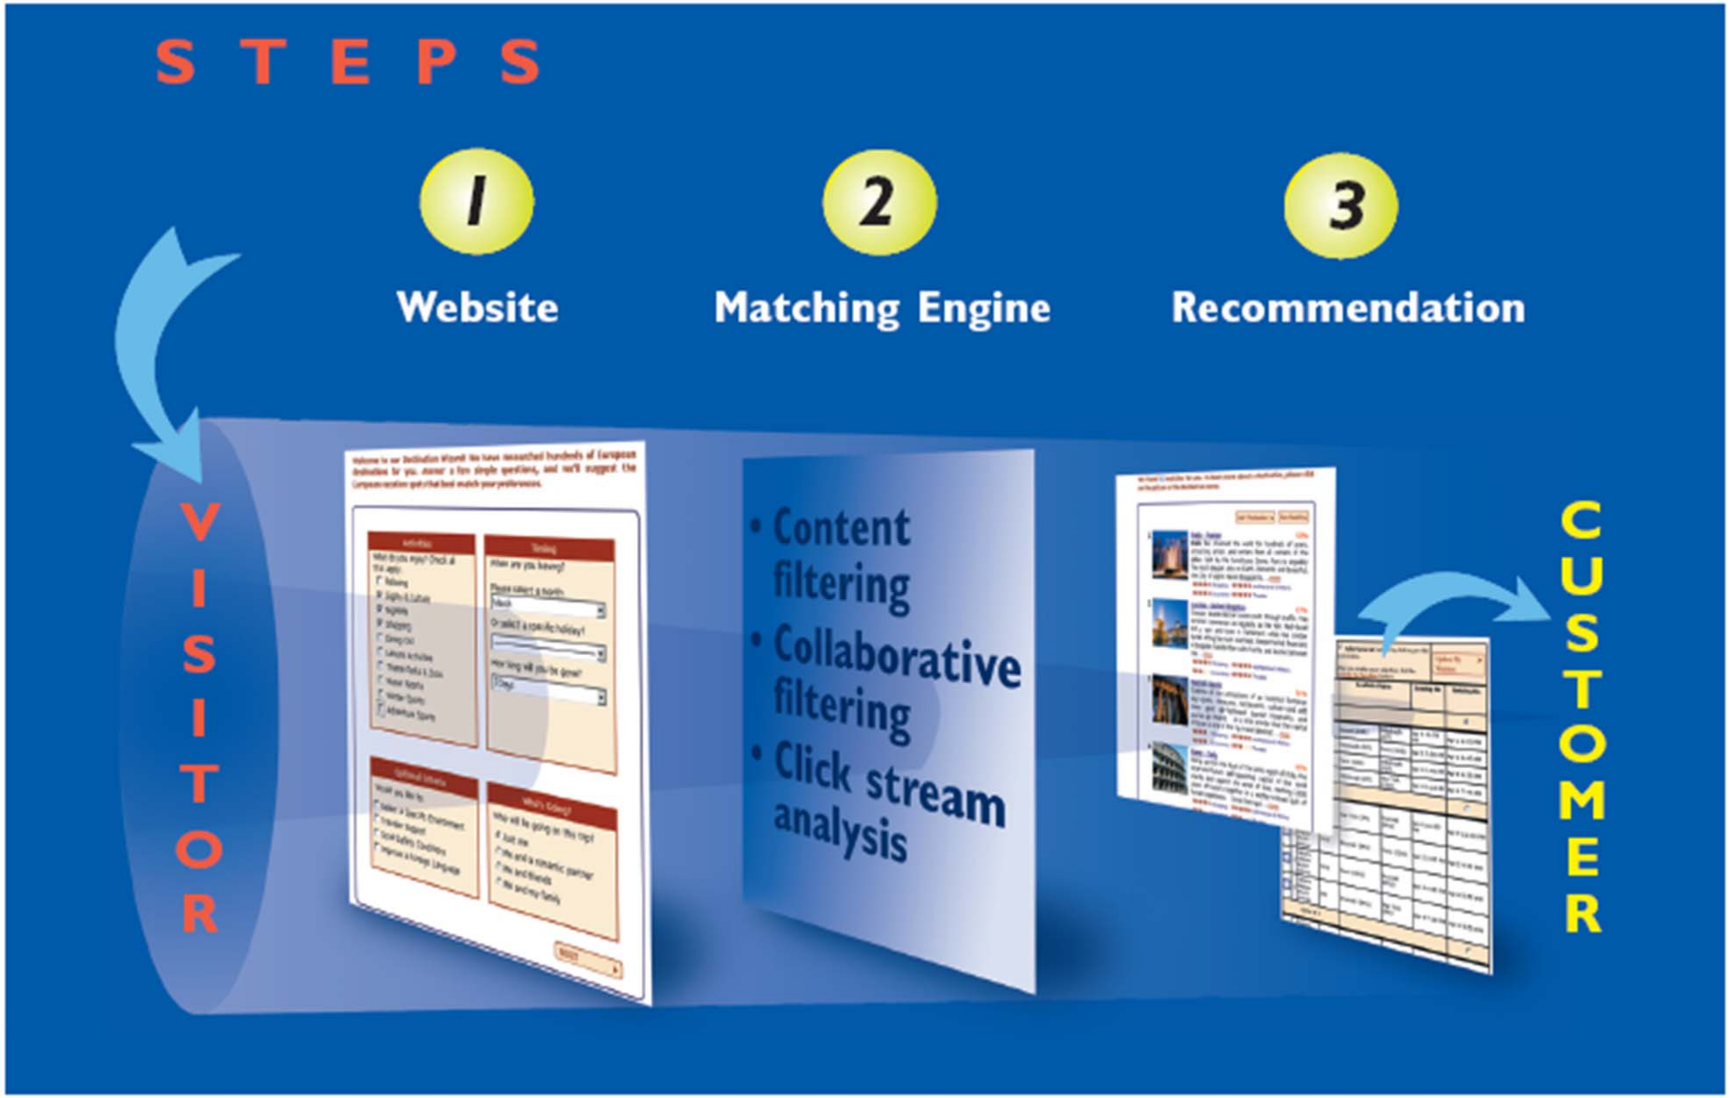

Examples

- Some examples found on the Web:
  1. **Amazon.com** – looks in the user past buying history and recommends product bought by users with similar buying behavior
  2. **Tripadvisor.com** - quotes past reviews given by a community of users
  3. **Activebuyersguide.com** – asks questions about features of the desired products to reduce the number of candidates
  4. **Trip.com** – asks questions about user constraints and preferences and shows options that satisfy these constraints
  5. **Smarterkids.com** – allows self-selection of a category of a user

“Core” Recommendation Techniques

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.

“Core” Recommendation Techniques

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.

MovieLens

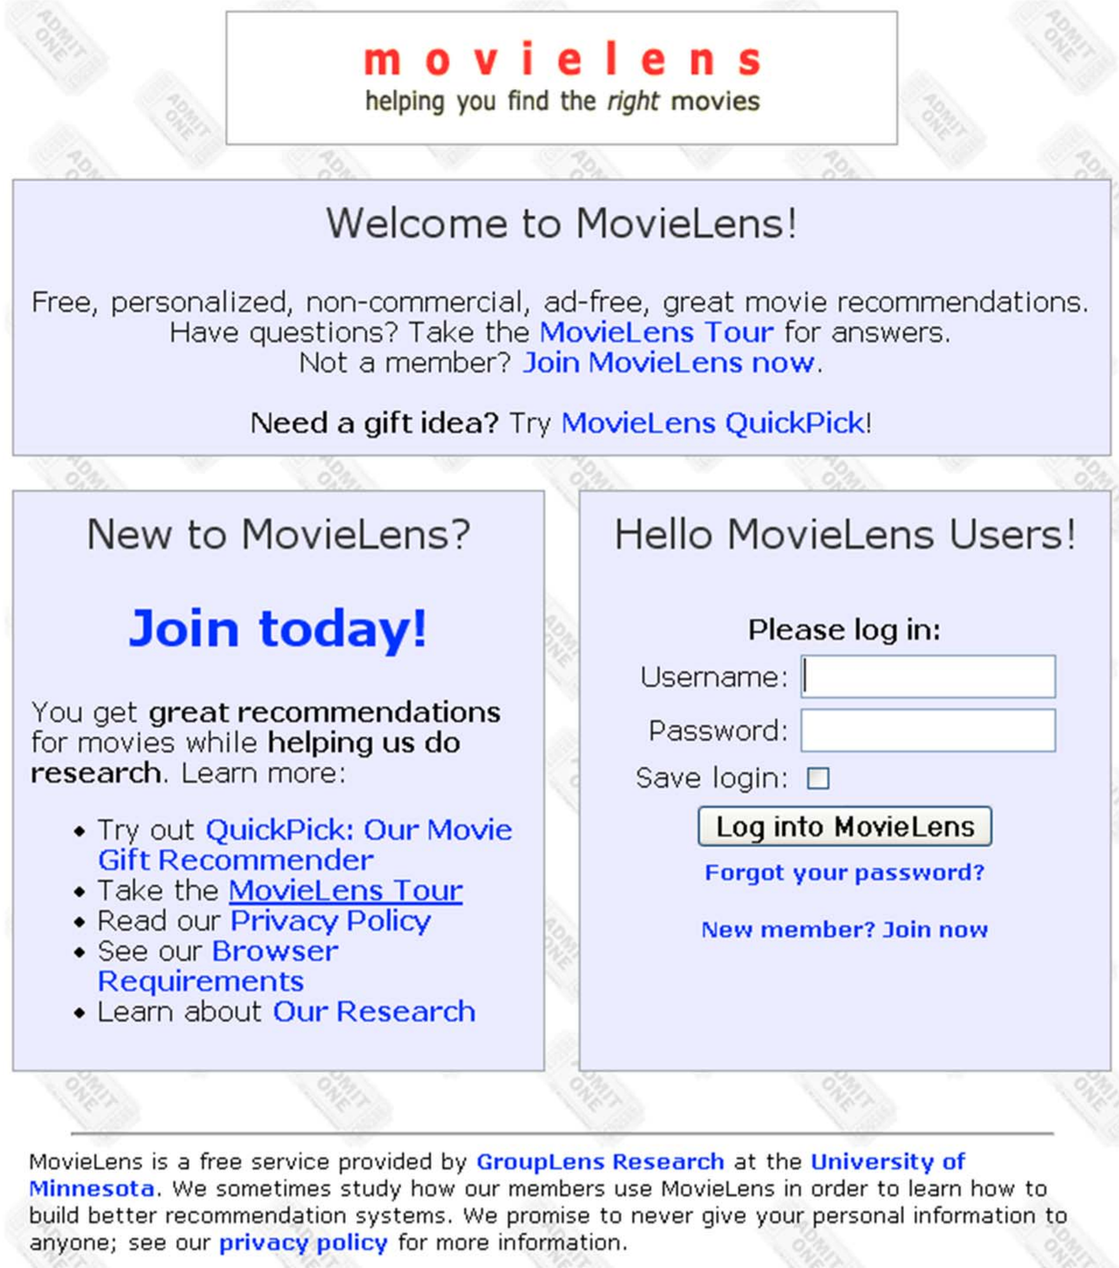
movielens
helping you find the right movies
Welcome to MovieLens!
Free, personalized, non-commercial, ad-free, great movie recommendations.
Have questions? Take the MovieLens Tour for answers.
Not a member? Join MovieLens now.
Need a gift idea? Try MovieLens QuickPick!
New to MovieLens?
Join today!
You get great recommendations for movies while helping us do research. Learn more:
Try out QuickPick: Our Movie Gift Recommender
Take the MovieLens Tour
Read our Privacy Policy
See our Browser Requirements
Learn about Our Research
Hello MovieLens Users!
Please log in:
Username:
Password:
Save login:
Log into MovieLens
Forgot your password?
New member? Join now
MovieLens is a free service provided by GroupLens Research at the University of Minnesota. We sometimes study how our members use MovieLens in order to learn how to build better recommendation systems. We promise to never give your personal information to anyone; see our privacy policy for more information.

Welcome to the new MovieLens!

Existing MovieLens users: We'd like to welcome you back to MovieLens, and let you know we have a new MovieLens FAQ you might want to read. We hope you like what you will see!

Take me to MovieLens!

New MovieLens users: Thank you for joining MovieLens! In order to generate personalized movie recommendations, we need to know a little about what movies you have already seen. MovieLens will now display several lists of movies. If you have seen any of the listed movies, please rate them using the rating scale shown below.

Ratings are on a scale of 1 to 5:

★★★★★ = Must See
★★★★☆ = Will Enjoy
★★★☆☆ = It's OK
★★☆☆☆ = Fairly Bad
★☆☆☆☆ = Awful

Remember: the more movies you rate, the more accurate MovieLens' predictions will be.

To rate a movie, just click on the pulldown next to the title of a movie you have seen. Blue stars will appear to indicate that your rating has been received.

★½☆☆☆	1.5 stars ▾	**Dude, Where's My Car? (2000)** DVD, VHS, info | imdb Comedy	☐

This image shows that the movie 'Dude, Where's My Car?' was rated 1.5 stars.

I'm ready to start rating!

So far you have rated **0** movies.
MovieLens needs at least **15** ratings from you to generate predictions for you.
Please rate as many movies as you can from the list below.

next >

	Your Rating	Movie Information
★★★	3.0 stars	**Austin Powers: International Man of Mystery (1997)** Action, Adventure, Comedy
★★★★	4.0 stars	**Contact (1997)** Drama, Sci-Fi
???	Not seen	**Crouching Tiger, Hidden Dragon (Wu Hu Zang Long) (2000)** Action, Adventure, Drama, Fantasy, Romance
???	Not seen	**Demolition Man (1993)** Action, Comedy, Sci-Fi
???	Not seen	**Eraser (1996)** Action, Drama, Thriller
???	Not seen	**Maverick (1994)** Action, Comedy, Western
★★★★½	4.5 stars	**Philadelphia (1993)** Drama
★★★½	3.5 stars	**Piano, The (1993)** Drama, Romance
???	Not seen	**Toy Story 2 (1999)** Adventure, Animation, Children, Comedy, Fantasy
★★★½	3.5 stars	**X-Men (2000)** Action, Adventure, Sci-Fi

next >

To get a new set of movies click the **next>** link.

Congratulations!

MovieLens can now generate personalized movie recommendations for you.

Start Using MovieLens

Remember, you can always keep rating movies you have seen. The more movies you rate, the better your predictions will be. We'd also like to tell you about some other features of MovieLens you might be interested in:

- **Getting recommendations.** MovieLens has shortcuts like Top Picks For You that provide you with quick access to common searches. You can use the Search tab to perform more advanced searches that filter by genre, date, and more, and save your favorite searches as personal shortcuts.

- **Your Wishlist.** Here you can keep track of movies you haven't yet seen. You can even print this list out and take it with you to your video store.

- **Movie buddies.** It can be a pain trying to decide what movie a group of people should see. Let MovieLens choose the right movie for you! You can add MovieLens users to be your buddies and be able to generate group movie recommendations

We will keep adding more great features as time goes on, so look for them!

Start Using MovieLens

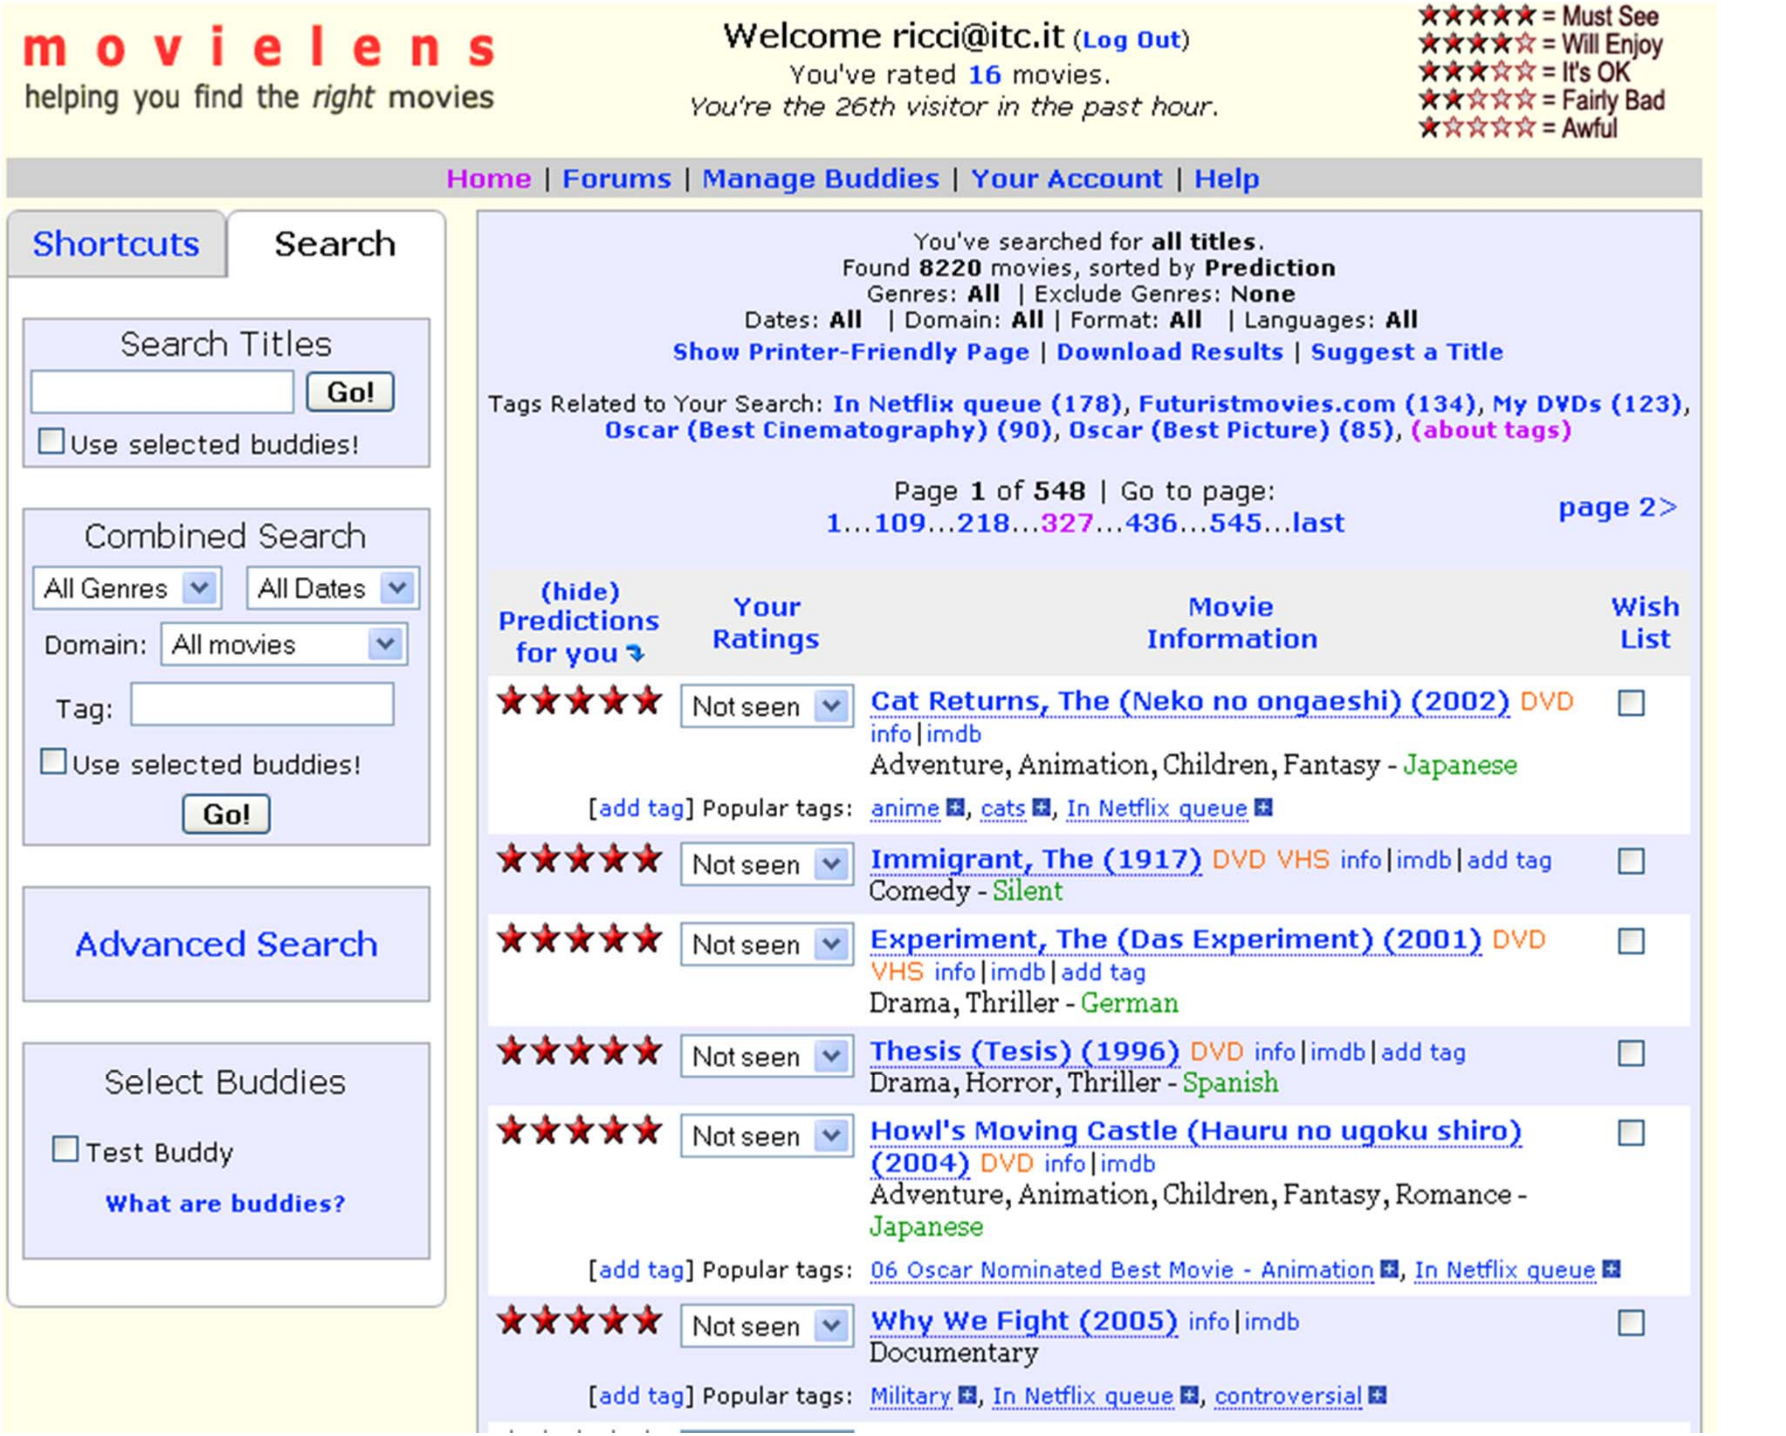
movielens
helping you find the right movies
Welcome ricci@itc.it (Log Out)
You've rated 16 movies.
You're the 26th visitor in the past hour.
★★★★★ = Must See
★★★★☆ = Will Enjoy
★★★☆☆ = It's OK
★★☆☆☆ = Fairly Bad
★☆☆☆☆ = Awful
Home | Forums | Manage Buddies | Your Account | Help
Shortcuts
Search
Search Titles
Go!
Use selected buddies!
Combined Search
All Genres
All Dates
Domain: All movies
Tag:
Use selected buddies!
Go!
Advanced Search
Select Buddies
Test Buddy
What are buddies?
You've searched for all titles.
Found 8220 movies, sorted by Prediction
Genres: All | Exclude Genres: None
Dates: All | Domain: All | Format: All | Languages: All
Show Printer-Friendly Page | Download Results | Suggest a Title
Tags Related to Your Search: In Netflix queue (178), Futuristmovies.com (134), My DVDs (123), Oscar (Best Cinematography) (90), Oscar (Best Picture) (85), (about tags)
Page 1 of 548 | Go to page:
1...109...218...327...436...545...last
page 2>
(hide) Predictions for you
Your Ratings
Movie Information
Wish List
Not seen
Cat Returns, The (Neko no ongaeshi) (2002) DVD info | imdb
Adventure, Animation, Children, Fantasy - Japanese
[add tag] Popular tags: anime, cats, In Netflix queue
Not seen
Immigrant, The (1917) DVD VHS info | imdb | add tag
Comedy - Silent
Not seen
Experiment, The (Das Experiment) (2001) DVD VHS info | imdb | add tag
Drama, Thriller - German
Not seen
Thesis (Tesis) (1996) DVD info | imdb | add tag
Drama, Horror, Thriller - Spanish
Not seen
Howl's Moving Castle (Hauru no ugoku shiro) (2004) DVD info | imdb
Adventure, Animation, Children, Fantasy, Romance - Japanese
[add tag] Popular tags: 06 Oscar Nominated Best Movie - Animation, In Netflix queue
Not seen
Why We Fight (2005) info | imdb
Documentary
[add tag] Popular tags: Military, In Netflix queue, controversial


m o v i e l e n s
helping you find the right movies
Welcome ricci@itc.it (Log Out)
You've rated 32 movies.
You're the 26th visitor in the past hour.
= Must See
= Will Enjoy
= It's OK
= Fairly Bad
= Awful
Home | Forums | Manage Buddies | Your Account | Help
Shortcuts
Search
- Top Picks For You
- Your Ratings
- Your Wishlist
- Newest Additions
- Rate Random Movies
- Most Often Rated
- Your Tags
- Suggest Title
- About Your Ratings
- New Drama
V for Vendetta (2006)
Inside Man (2006)
Match Point (2005)
- New DVDs
Capote (2005)
Walk the Line (2005)
Good Night, and Good...
- New Movies
V for Vendetta (2006)
Why We Fight (2005)
Inside Man (2006)
How to create shortcuts
Publish your shortcuts
V for Vendetta (2006)
Your Prediction:
Rate This Movie:
Not seen
Wish List:
Movie Information
Starring: Natalie Portman, Hugo Weaving, Stephen Rea, John Hurt
Directed by: James McTeigue
Genres: Action, Drama, Sci-Fi, Thriller
Language: English
Average rating: (4 stars)
Rated by: 128 users
Links: IMDb, Rotten Tomatoes
Movie Tags (more about tags)
Add and edit tags here
My Tags [edit]
- none
[add new tags]
Popular tags:
Click on this icon ( ) to add a tag to your list!
comic book (2)
revenge (1)
Alan Moore (1)
john hurt (1)
1984 (1)
guy fawkes (1)
Futuristmovies.com (1)
revolutionary (1)
disacknowledged (1)
Forum Posts
These posts mention V for Vendetta (2006)
Write about V for Vendetta (2006) in the MovieLens Forums!
Related Forum Posts
These posts mention movies similar to V for Vendetta (2006)
Topic
Author
Re: Fitting into movie groups
(shitdisturber)
Re: What's the last thing you watched an...
(PolarisDiB)
Re: Fitting into movie groups
(FarmerF)
Re: Fitting into movie groups
(Bec1029)
Re: Ask Dr. Vigilans
(Vigilans)
Ask Dr. Vigilans
(PolarisDiB)
Re: Fitting into movie groups
(Ryuukuro)
Re: What's the last thing you watched an...
(vargus)

User-based Collaborative Filtering

- Idea: users who agreed in the past are likely to agree in the future
- To predict a user's opinion for an item, use the opinions of like-minded users
  - Precisely, a (small) set of very similar users
- Similarity between users is decided by looking at their overlap in their past opinions
  - High overlap = strong evidence of similarity = high weight

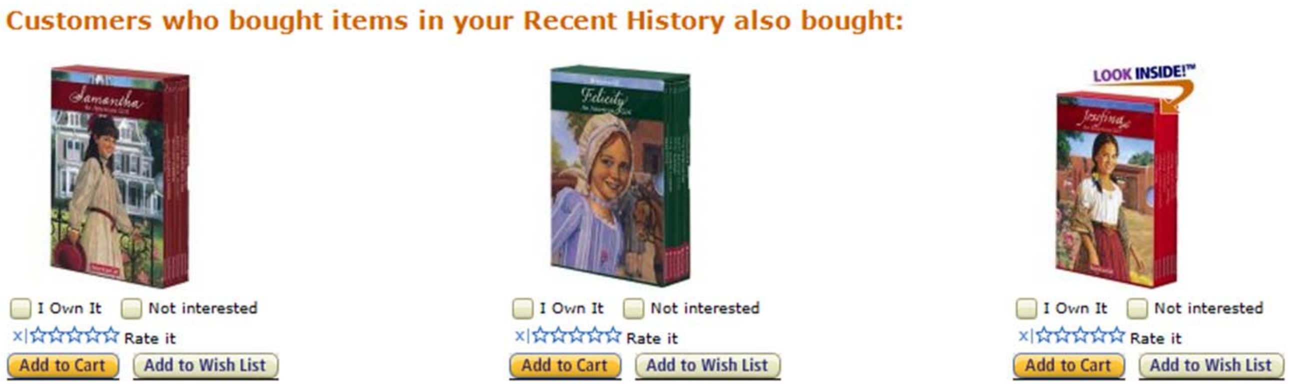

Collaborative Filtering

Collaborative Filtering consists of five steps:

1. For a target user (to whom a recommendation is produced) the set of his ratings is identified
2. The users similar to the target user (according to a similarity function) are identified
   - Cosine similarity, Pearson's correlation, Mean Squared Difference, or other similarity metrics
3. Items rated by similar users but not by the target user are identified
4. For each item a predicted rating is computed
   - Weighted according to users' similarity
5. Based on this predicted ratings a set of items is recommended

Collaborative Filtering

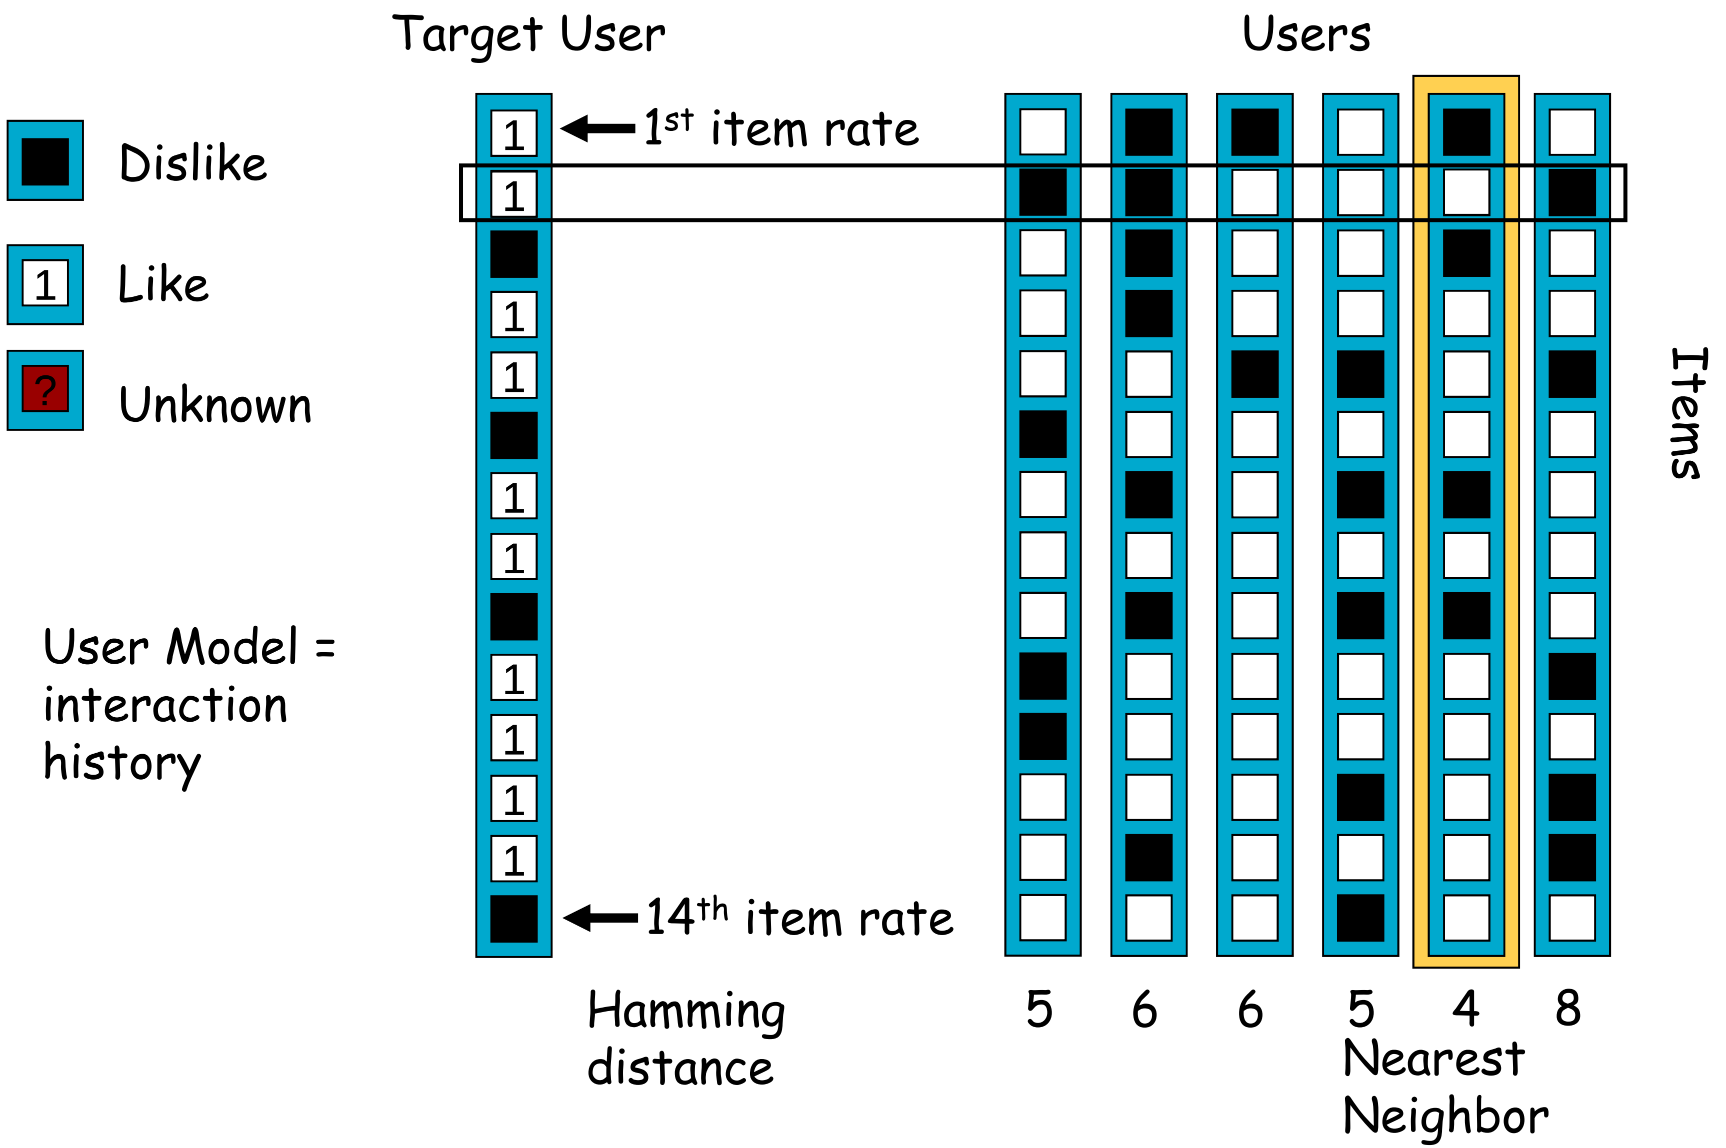

Collaborative Filtering

- **Pros:** requires minimal knowledge engineering efforts
  - Users and items have no structure or characteristics
- **Cons:**
  - Requires many explicit ratings to bootstrap
    - New user and new item problem
  - Does not explains recommendations
  - Does not support sequential decision making
  - Does not support bundle recommendation
  - Scalability
    - Quadratic computational time
    - Web-based recommender will struggle to provide real-time recommendations

Problems of CF: Sparsity

- Sparsity – large product sets and user ratings for a small percentage of them
  - Sparsity of real-life datasets: 98.69% and 99.94%
  - Amazon: millions of books and a user may have read hundreds
- Drift – popular items are recommended and there are no serendipitous recommendations
  - The usefulness of recommending popular items is questionable
    - Recommending top items is obvious for users
  - Recommending unpopular items
    - Is risky, but could be valuable for users

“Core” Recommendation Techniques

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.

Syskill & Webert User Interface

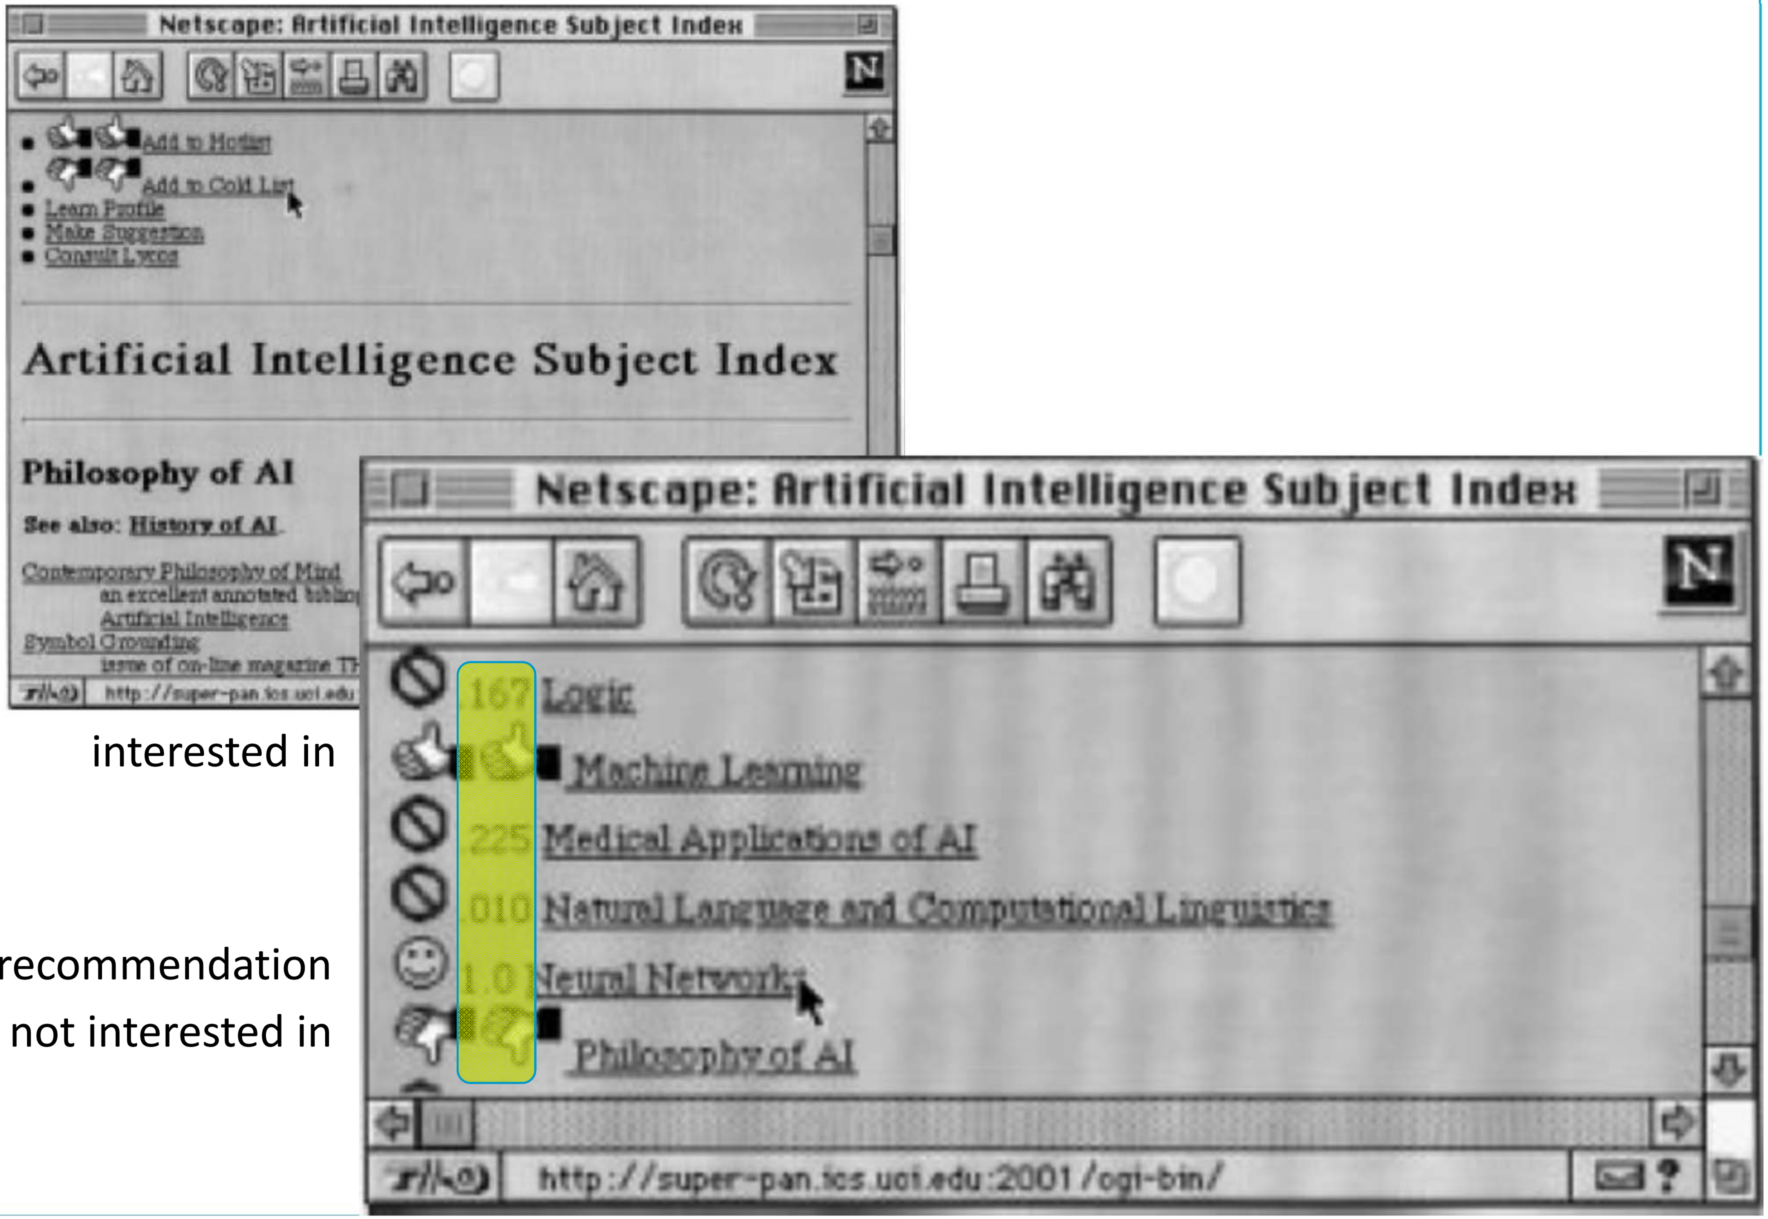

interested in

recommendation

not interested in

Content-Based Recommendations

- The system recommends items similar to those the user liked
  - Similarity is based on the content of items which that the user has evaluated
    - Very different from collaborative filtering
- Originated in Information Retrieval
  - Was used to retrieve similar textual documents
    - Documents are described by textual content
    - The user profile is structured in a similar way
    - Documents can be retrieved based on a comparison between their content and a user model
- Recommender implemented as a classifier
  - e.g., Neural Networks, Naive Bayes, C4.5, …

Content-Based Recommendations

- Assist users in finding items that satisfy their long-term recurring information needs
  - User profile describes long-term preferences
- Long-and short-term preferences can be combined
  - aggregate the level of interest as represented in the long-term and short-term profiles
- Long- and short-term recommendations can be combined
  - items satisfying short-term preferences can be sorted according to long-term preferences

Example Book Recommendation

Short Term

- Taking two weeks off
- Novel
- Interested in a Polish writer
- Travel book
- Would like to read about the meaning of the life

Long Term

- Dostoyevsky
- Stendhal
- Chekhov
- Pessoa
- Sedaris
- Auster
- Mann

User

Recommendation
Joseph Conrad, Heart of darkness

Problems of Content-Based Recommenders

- Only a shallow content analysis is performed
  - Images, video, music, …
- Certain textual features cannot be extracted
  - Quality, writing style, agreement, sentiments, …
    - If a page is rated positively, it could not necessarily be related to the presence of certain words
- Requires considerable domain knowledge
- Even less serendipity – recommends only similar items
  - Trustful but not very useful recommendations

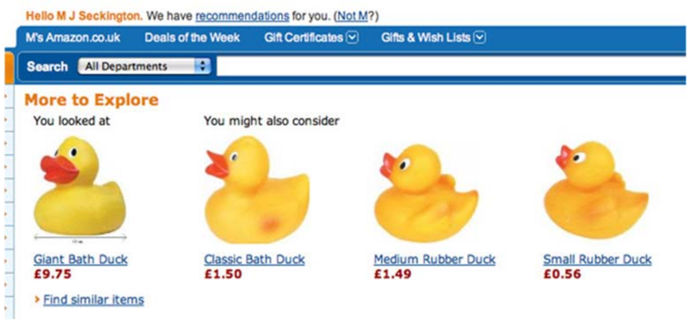

"Core" Recommendation Techniques

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.

Demographic recommendations

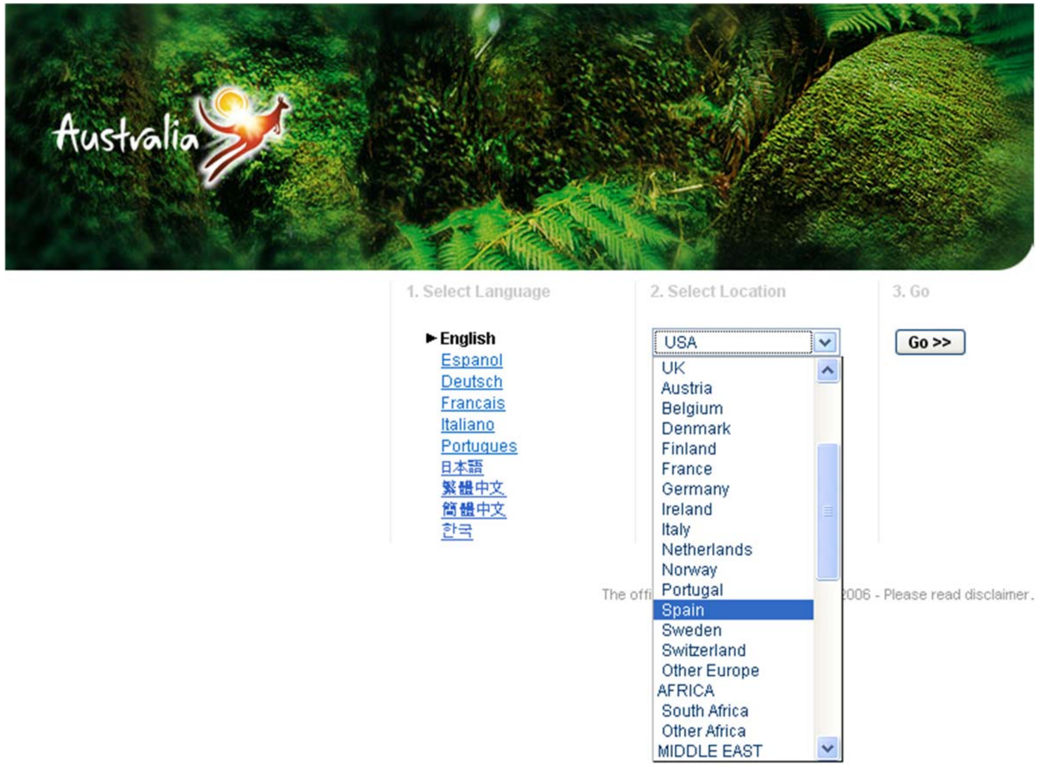

Demographic recommendations

- Collects demographic information about users
- Aggregates users into clusters
  - Using a similarity measure and data correlation
- Classifies each user to a cluster that contains the most similar users
- Generates cluster-based recommendation
  - Similar to CF but exploits demographic similarity

	gender	age	area code	education	employed	Dolce
Karen	F	15	714	HS	F	+
Lynn	F	17	714	HS	F	–
Chris	M	35	714	C	T	+
Mike	F	40	714	C	T	–
Jill	F	10	714	E	F	?

Problems of Demographic Recommenders

- Require domain engineering by human experts
- Involves expensive collection of demographic data
  - Severe privacy hazards – deferred by many users
- Efficient but
  - Does not track the changes in the population
- Demographic similarity does not necessarily imply preference similarity

“Core” Recommendation Techniques

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.

Utility related information

Utility methods

- Items are described using features $f_1, \ldots f_m$
  - E.g., price, size, various technical properties, …
- User is modeled using the same features
  - weights $u_1, \ldots, u_m$ – importance of each feature
  - scores $f_1, \ldots f_m$ – value of each feature
- Utility function combines the scores and weights into the overall degree of matching

$$U(u_1,\ldots,u_m,p_1,\ldots,p_m)=\sum_{j=1}^{m}u_j p_j$$

- Problems
  - How to ac
    – Do users know what they want?

"Core" Recommendation Techniques

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.

Knowledge-based recommenders

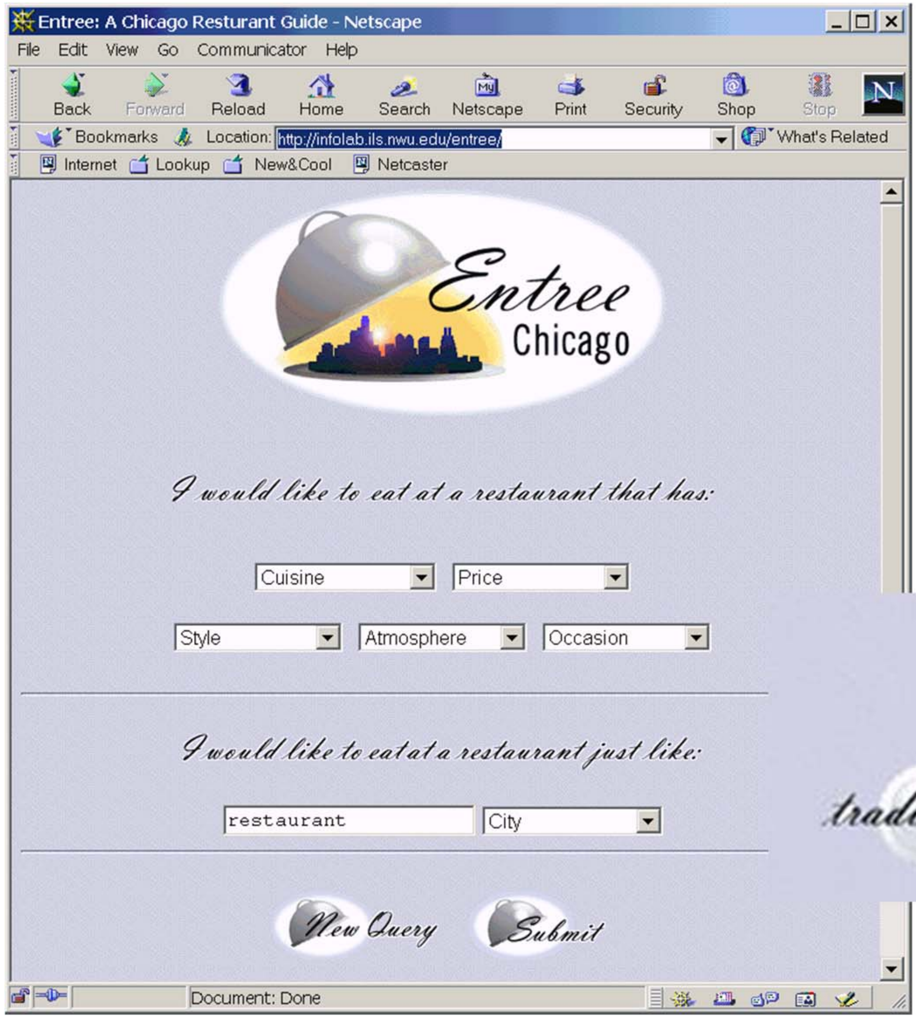

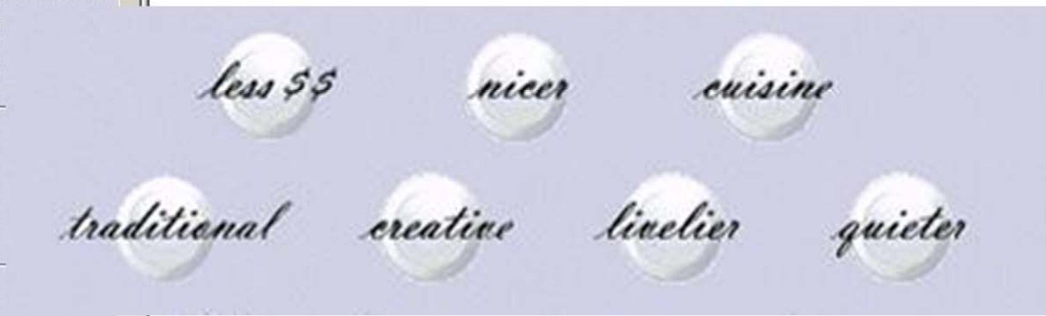

Knowledge-based recommenders

- Uses domain knowledge to identify items that meet user requirements
- A cycle of critique starts
  - If the user is not satisfied, he/she can criticize them
    - modify certain features
    - if the price is too high, ask for a cheaper restaurant
  - New recommendation cycle and the criticized features are considered the most important
- Problems of knowledge-based recommenders
  - Require heavy domain and item modeling
  - User model is barely used
    - Are the recommendation personalized?

Matrix Factorization

- On the map since the Netflix Prize Competition
  - Training data
    - 6 years of data: 2000-2005
    - 100M ratings of 480K users for 18K movies
  - Test data
    - Evaluation criterion: root mean squared error (RMSE)
    - Netflix Cinematch RMSE baseline 0.9514
  - Competition
    - 2700+ teams
    - $1M grand prize for 10% improvement on Cinematch
    - $50,000 annual progress prize for best improvement
  - Won by the Bellkor-Gravity team
    - Ensemble of more than 100 recommenders
    - Many of them based on Matrix Factorization

Latent factor model

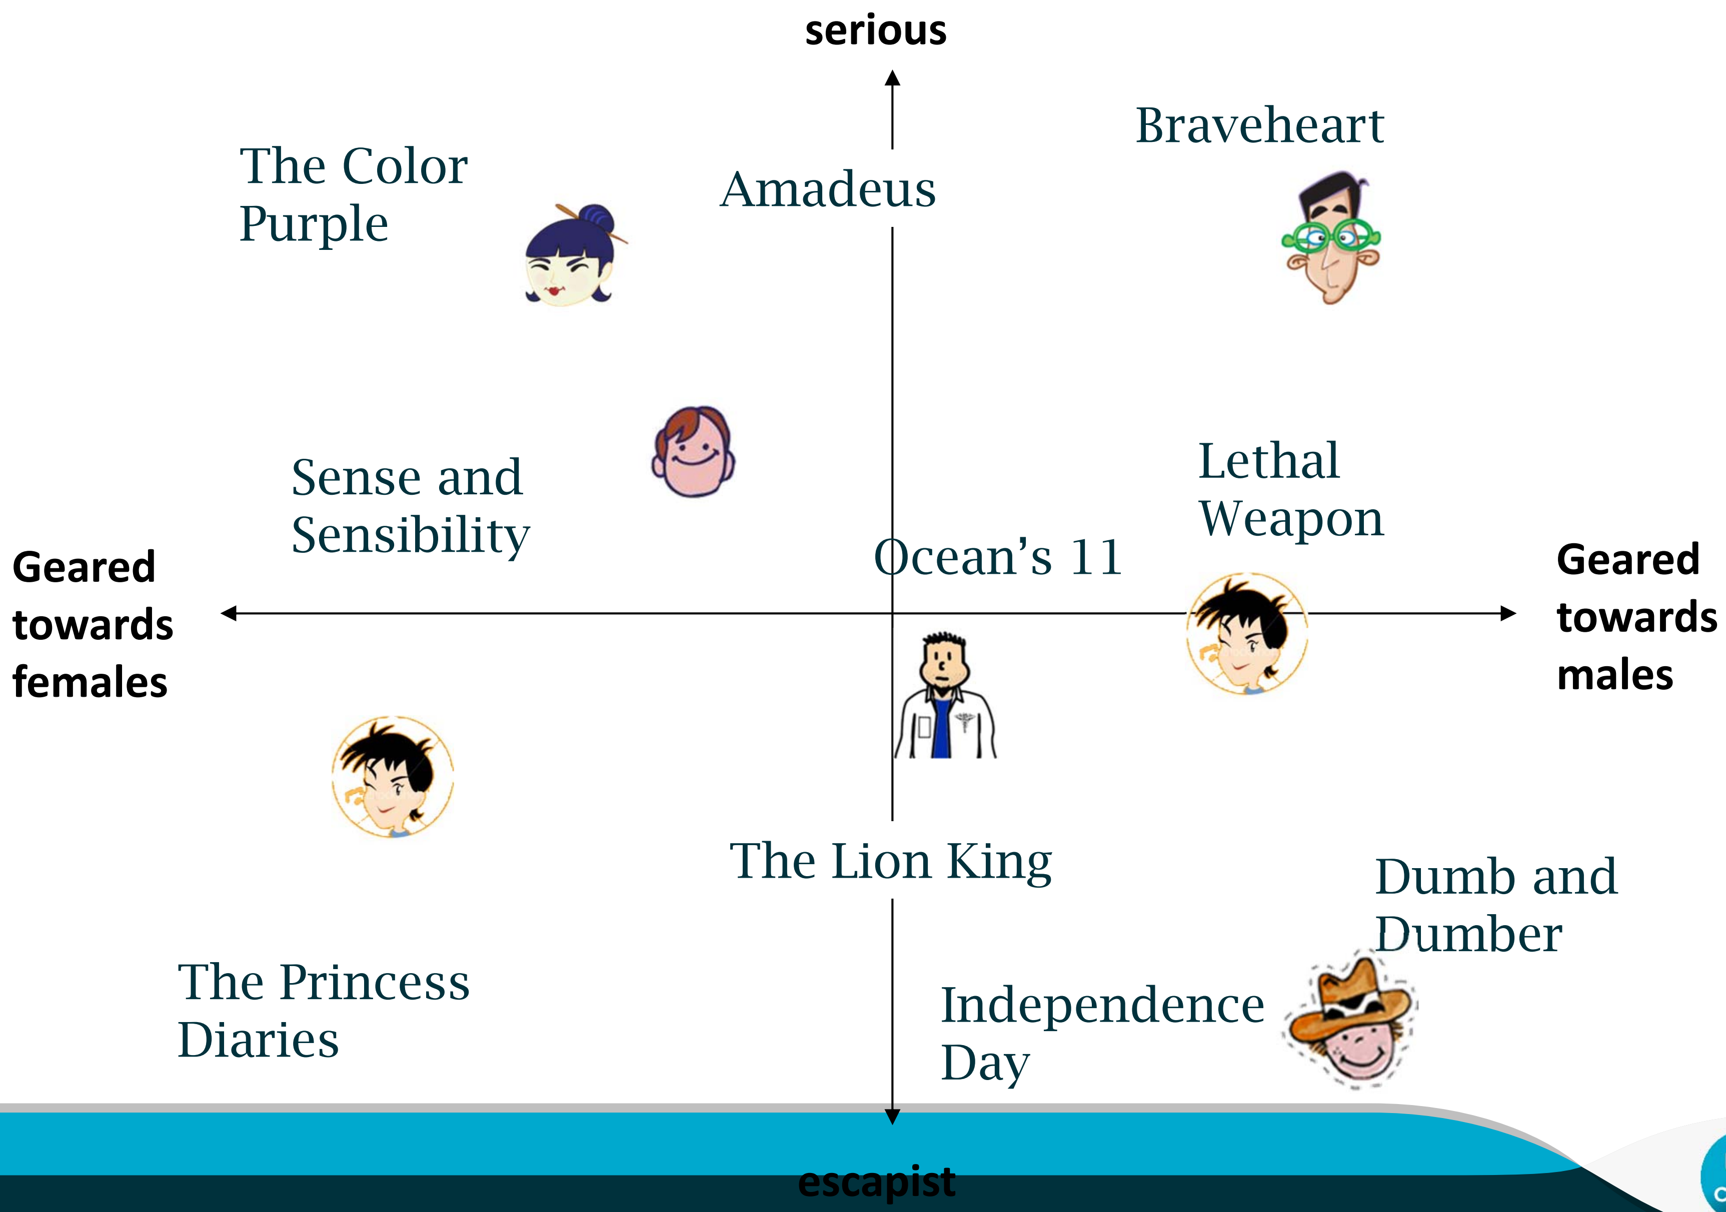

Latent factor model

users

items

1		3			5			5		4	
		5	4	?		4			2	1	3
2	4		1	2		3		4	3	5	
	2	4		5			4			2	
		4	3	4	2					2	5
1		3		3			2			4	

~

~

items

.1	-.4	.2
-.5	.6	.5
-.2	.3	.5
1.1	2.1	.3
-.7	2.1	-2
-1	.7	.3

●

users

1.1	-.2	.3	.5	-2	-.5	.8	-.4	.3	1.4	2.4	-.9
-.8	.7	.5	1.4	.3	-1	1.4	2.9	-.7	1.2	-.1	1.3
2.1	-.4	.6	1.7	2.4	.9	-.3	.4	.8	.7	-.6	.1

Estimate unknown ratings as an inner product of latent user and item factors

Latent factor model

users

items

1		3			5			5		4	
		5	4	?		4			2	1	3
2	4		1	2		3		4	3	5	
	2	4		5			4			2	
		4	3	4	2					2	5
1		3		3			2			4	

~

~

items

.1	-.4	.2
-.5	.6	.5
-.2	.3	.5
1.1	2.1	.3
-.7	2.1	-2
-1	.7	.3

●

users

1.1	-.2	.3	.5	-2	-.5	.8	-.4	.3	1.4	2.4	-.9
-.8	.7	.5	1.4	.3	-1	1.4	2.9	-.7	1.2	-.1	1.3
2.1	-.4	.6	1.7	2.4	.9	-.3	.4	.8	.7	-.6	.1

Estimate unknown ratings as an inner product of latent user and item factors

Latent factor model

users

items

1		3			5			5		4	
		5	4	4		4			2	1	3
2	4		1	2		3		4	3	5	
	2	4		5			4			2	
		4	3	4	2					2	5
1		3		3			2			4	

~

~

items

.1	-.4	.2
-.5	.6	.5
-.2	.3	.5
1.1	2.1	.3
-.7	2.1	-2
-1	.7	.3

●

users

1.1	-.2	.3	.5	-2	-.5	.8	-.4	.3	1.4	2.4	-.9
-.8	.7	.5	1.4	.3	-1	1.4	2.9	-.7	1.2	-.1	1.3
2.1	-.4	.6	1.7	2.4	.9	-.3	.4	.8	.7	-.6	.1

Estimate unknown ratings as an inner product of latent user and item factors

Matrix Factorization

- Pros
  - Well evaluated in data mining
  - Very strong and accurate model
  - Can scale to Web-size datasets
  - Can incorporate contextual dependency
  - Many variants and open implementations
- Cons
  - Can easily overfit
  - Requires optimization of parameters
  - Requires regularization
  - Meaningless latent factors

Hybrid recommendations

- Each core method has its own pros and cons
- Combine core methods for recommendations
  - Leverage the advantages and hide shortcoming
  - Recall the Netflix winning ensemble!
- Lots of hybrid methods – no standard

Hybridization method	Description
Weighted	The scores (or votes) of several recommendation techniques are combined together to produce a single recommendation.
Switching	The system switches between recommendation techniques depending on the current situation.
Mixed	Recommendations from several different recommenders are presented at the same time
Feature combination	Features from different recommendation data sources are 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.

Weighted, Switching, and Mixed

- The prediction is computed from the outputs of individual methods
  - Linear combination of recommendations
  - What is the weight of each recommender?
- Switching: the system uses some criterion to switch between recommendation techniques
  - CB technique is applied first and then CF
  - When to switch? What is the switching criterion?
- Mixed: recommendations generated by individual techniques are presented to users
  - The user has to decide
    - Decision support tool rather than recommender system

Feature combination and Meta-Level

- Feature combination: features used by one technique, are also used by other techniques
  - Content features are used by collaborative filtering to compute similarity
  - Machine learning technique uses ratings and content features to predict new ratings
  - Plenty of options for combination. Which are beneficial?
- Meta-level: use the user models generated by one technique as input for other techniques
  - Mediation of user modeling data
  - Plenty of options for user model interoperability. Which are beneficial?

Cascade and Feature augmentation

- Cascade: one method produce a coarse list of recommendations, which is refined by another
  - Utility-based technique places items into buckets of equal preference
  - Collaborative technique is applied to break ties
  - Which methods can be cascaded? What is their best ordering?
- Feature augmentation: output of one method is incorporated by another method
  - Content-based book recommendations
  - Recommend “related authors” and “related titles”
  - Which features of which methods can be augmented? What is their best ordering?

Hybrid Recommendations

- Hybrid methods are the state-of-the-art
  - Most powerful and most popular
  - Leverage the advantages of the individual methods
  - Generate recommendations superior to individual methods
- Plenty of unexplored options for hybridization
  - The most simple and widely used methods are weighted, switching, and mixed hybridizations
  - Several focused studies of cascade and feature augmentation hybridizations
  - Very few studies on feature combination and meta-level hybridizations

Evaluating Recommender Systems

- Algorithmic evaluation
  - Offline datasets, statistic evaluations
    1. Measure how good is the system in predicting the exact rating value (value comparison)
    2. Measure how well the system can predict whether the item is relevant or not (relevant vs. not relevant)
    3. Measure how close the predicted ranking of items is to the user's true ranking (ordering comparison).
- User studies
  - Let users play with the system
  - Collect and analyze feedback
  - Compare with non-personalized system

Algorithmic Evaluation

- Split the data into training and test sets
- Build a model using training data
- Compare the predicted ratings for test set items with the actual rating stored in the test set
- N-fold validation is often applied

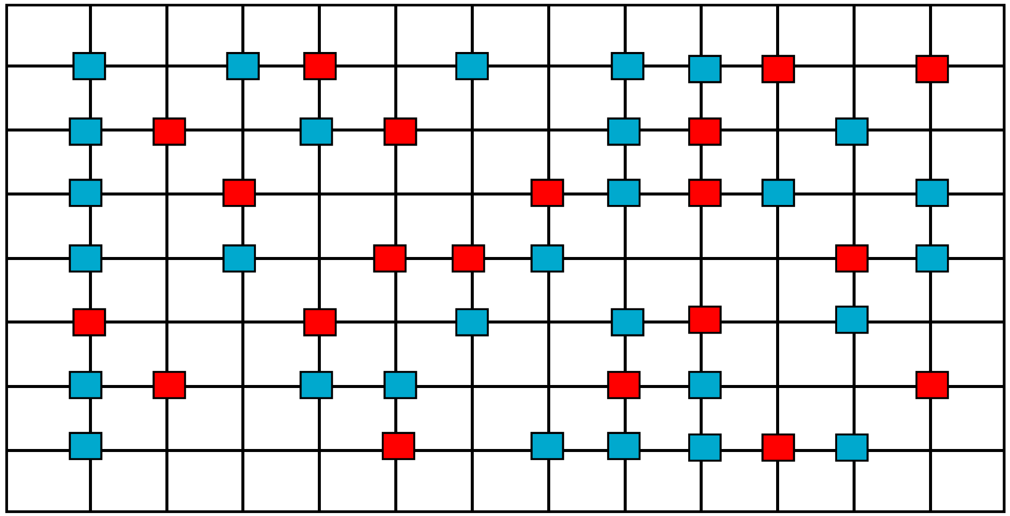

Evaluation metric: predictive accuracy

- Measure whether the predicted ratings are close to the true user ratings
  - Mean Absolute Error (MAE)

$$MAE = \frac{\sum_{i=1}^{N} |p_i - r_i|}{N}$$

- Less appropriate for tasks like “find good items”
  - Users examine only top rated items
- Mean squared error can be computed as well
  - Netflix Prize competition’s RMSE
  - Does 3.9 or 4.1 stars really matter?

Evaluation metric: classification accuracy

- Measure if item classification is correct
- Discretize the rating scale
  - Precision is the ratio of relevant items selected by the recommender to the number of items selected
  - Recall is the ratio of relevant items selected to the number of relevant
  - F-measure is used as well
- Is the accuracy of recommending the first and the last item in the list equally important?

	Selected	Not Selected	Total
Relevant	N_{rs}	N_{rn}	N_r
Irrelevant	N_{is}	N_{in}	N_i
Total	N_s	N_n	N

Evaluation metric: ranking accuracy

- Measure whether the recommended items were ordered accurately
  - Spearman correlation
  - Kendall’s Tau
- Important when logical dependencies exist between the recommended items
- Merge several accuracy metrics
  - Normalized discounted cumulative gain (NDCG)

NAME	REFERENCES	DOMAIN
ACR News	Mobasher et al. 2000	Netnews filtering
Amazon	Amazon 2001	E-commerce
Amalthaea	Moukas 1997	Web recommender
Anatagonomy	Skagami et al. 1997	Personalized newspaper
Beehive	Huberman and Kaminsky 1996	Sharing news
Bellcore Video Recom	Hill 1995	Movie recommender
Casmir	Berney and Ferneley 1999	Document recommender
CDNow	CDnow	E-commerce
Fab	Balabanovic and Shokam 1997	Web recommender
GroupLens	Resnick et al. 1994	Netnews recommender
ifWeb	Minio and Tasso 1996; Asnicar and Tasso 1997	Web recommender
InfoFinder	Krulwich and Burkey 1995, 1996	Information recommender
INFOrmer	Riordan and Sorensen 1995; Sorensen et al. 1997	Netnews filtering
Krakatoa Chronicle	Kamba et al. 1995	Personalized newspaper
LaboUr	Schwab et al. 2001	Document recommender
Let's Browse	Lieberman et al. 1999	Web recommender
Letizia	Lieberman et al. 1995	Web recommender
LifeStyle Finder	Krulwich 1997	Purchase, travel and store recommender
MovieLens	Good et al. 1999	Movie recommender
News Dude	Billsus and Pazzani 1999	Netnews recommender
NewsWeeder	Lang 1995	Netnews recommender
NewT	Sheth and Maes 1993	Netnews filtering
Personal WebWatcher	Mladenic 1996	Web recommender
PSUN	Sorensen and McElligot 1995	Netnews recommender
Re:Agent	Boone 1998	E-mail filtering
Recommender	Basu et al. 1998	Movie recommender
Ringo/FireFly	Shardanand 1994; Shardanand and Maes 1995	Music recommender
SIFT Netnews	Yan and Garcia-Molina 1995	Netnews filtering
SiteIF	Stefani and Strappavara 1998	Web recommender
Smart Radio	Hayes and Cunningham 1999, 2000	Music lists recommender
Syskill & Webert	Pazzani et al. 1996; Pazzani and Billsus 1997	Web recommender
Tapestry	Goldberg 1992	E-mail filtering
Webmate	Chen and Sycara 1998	Web recommender
WebSail	Chen et al. 2000	Web search filtering
WebSell	Cunningham et al. 2001	Purchase recommender
Websift	Cooley 1999	Web recommender
WebWatcher	Armstrong et al. 1995; Joachims et al. 1997	Web recommender

Domains

- News
- Movies
- Web pages
- Documents
- Travel
- Email
- Music
- Web search
- Social media
- People
- eCommerce
- eHealth
- … more and more …

Challenge: Data sparsity

- Personalized systems succeed only if sufficient information about users is available
  - No User Model = No Personalization
- How to gather enough user modeling data in unobtrusive manner?
- If the required data is not available
  - Web of trust to identify "similar users"
  - Use external data sources
    - – Web mining
  - The output is always an approximation
- Similarly: new item problem

Challenge: : Contextualization

- Systems should adapt to user context
  - Some methods cannot cope with this
- Largely depends on the definition of context but in practice this includes
  - Short term preferences (“tomorrow I want …”)
  - Information related to the specific space-time position of the user (“less than 5 mins walking)
  - Motivations of search (“present to my wife”)
  - Circumstances (“some time to spend here”)
  - Emotions and mood (“I feel adventurous”)
  - …

Challenge: : Privacy

- Personalization is based on personal data
  - Privacy vs. personalization tradeoff
    - More user information = more accurate personalization
    - More user information = less user privacy
- Laws that impose stringent restrictions on the usage and distribution of personal data
  - Systems must cope with these legislation
    - e.g., personalization systems exchanging user profiles could be impossible for legal reasons
- Personalization systems must be developed in a way that limits the possibility of an attacker learning/accessing personal data

Challenge: : Robusteness

- Recommender systems should be robust against attacks aiming at modifying the system such that it will recommend an item more often than others
  - Shilling
  - Nuking
- Some algorithms may be more robust than others

Challenge: : Scalability

- Personalization techniques rely on extensive user/item descriptions
  - Many of them are hardly scalable
- Techniques that can overcome this
  - Feature selection
  - Matrix factorization
  - Latent factors
  - Clustering and partitioning
  - Distributed computing
  - P2P architectures
  - Parallel computing
  - …

Open Challenges:

- Generic user models and personalization
- Portable and mobile personalization
- Emotional and value aware personalization
- User trust and recommendations
- Persuasive personalized technologies
- Group-based personalization
- Interactive sequential personalization
- Complex and bundle recommendations
- Robustness of business recommenders systems
- Semantically enhanced personalization
- Personalization in social applications
- Personalization in the Internet of Things
- People recommender systems
- Personalization or information bubble
- … more and more …

Resources

- The Adaptive Web – Methods and Strategies of Web Personalization
- Recommender Systems – An Introduction
- Recommender Systems Handbook
- Persuasive Recommender Systems – Conceptual Background and Implications
- More detail at www.recommenderbook.net

Thank you

CSIRO Digital Productivity Flagship
Jill Freyne

Senior Research Scientist

t +61 2 9372 4453
e jill.freyne@csiro.au