Personalization - Who needs what and why?

Eelco Herder

L3S Research Center / Leibniz University of Hanover
Hannover, Germany

7 September 2015

Outline

Who is Eelco Herder?

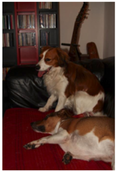

What is personalization?

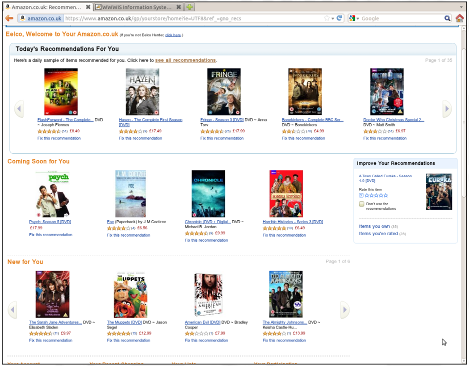

Product recommendations in Amazon. These recommendations are based on past purchases and past browsing behavior. The user can improve the recommendations by editing his or her user profile.

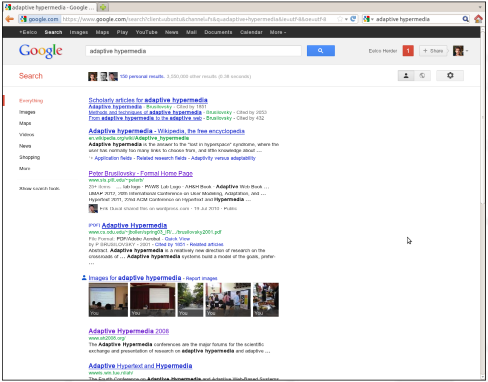

Google search results are personalized, based on past searches, current location, language settings (apparently 57 features in total).

A formal definition

Adaptive Hypermedia

By adaptive hypermedia systems we mean all hypertext and hypermedia systems which reflect some features of the user in a user model and apply this model to adapt various visible aspects of the system to the user.

Peter Brusilovsky: Methods and Techniques of Adaptive Hypermedia. User Modeling and User-Adapted Interaction 6 (2-3), 1996

Adaptive hyperwhat?

In 1945, Vanevar Bush envisioned a machine, the *memex*. By consulting several sources consecutively, a user builds an associative *trail* of related documents, which can be labeled and annotated with notes and comments.

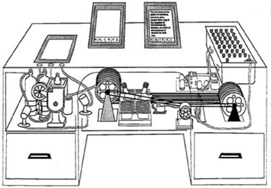

Similar to this idea, hypertext is a collection of documents that is connected by (associative) links.
The World Wide Web is the most common form of hypertext.

Does one size fit all?

In a library, a person looks for some books on China. What will the librarian recommend?

- Is the person a *small child* who saw a TV show about China and wants to learn about this exotic country?
- Or a *high school student* working on a paper?
- Perhaps a *prospective tourist*?
- A scholar interested in *Eastern philosophy*?
- Someone who can *read Chinese*?

Elaine Rich: User Modeling via Stereotypes. Cognitive Science 3, 329-354 (1979)

Most likely the librarian will make an educated guess, based on the person's appearance:

- age, style of clothing, accent, choice of words, ...

This initial guess might be confirmed or refuted by observations.

- It is assumed that a European cannot read Chinese, unless said otherwise
- Children are generally not (yet) interested in Eastern philosophy, but there are exceptions
- ....

The educated guess, a *stereotype* can be refined with follow-up questions.

Persons expect a *personalized* advice, even though the librarian does not know them.

And the same seems to yield for Web stores.

Jeff Bezos, amazon.com

If I have 3 million customers on the Web, I should have 3 million stores on the Web

When is personalization useful?

My supermarket is not personalized. Still, I can find all products that I need. Probably just because my needs are similar to everyone else's needs.

Personalization is deemed useful when:

- there are so many things to choose from that there is a need for guidance or recommendations
- the system is used by people with different goals and backgrounds

The ideal recommender

Your partner, your best friend or your mother probably knows a lot about you:

- the food you like, the books you read, the movies you watch
- things that interest you or that upset you
- your current needs, aspirations and goals
- dates of your birthday, your kids' birthdays, and holidays
- secret desires and phantasies

Still, this does not guarantee that your mother will buy you a present that you like.

It can be something that

- you already have
- you hate for some reason only known to you
- she bought to surprize you (sometimes this works out perfectly fine, though)

Goals of personalization

In the literature, the following goals are often mentioned:

- helping users to find information they need
- presenting information in the language of choice
- recommending products
- supporting collaboration
- taking over parts of routine tasks

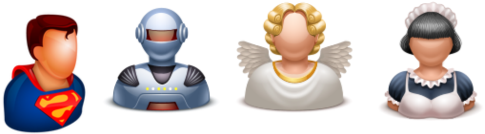

How does personalization work?

In a nutshell, a personalized system tries to understand the user using

- information that he or she provides
- activities that a user carries out

And tries to guess what you currently want to do, by comparing

- you with people that are similar to you
- things that you own or like with similar things

Steps in the personalization process

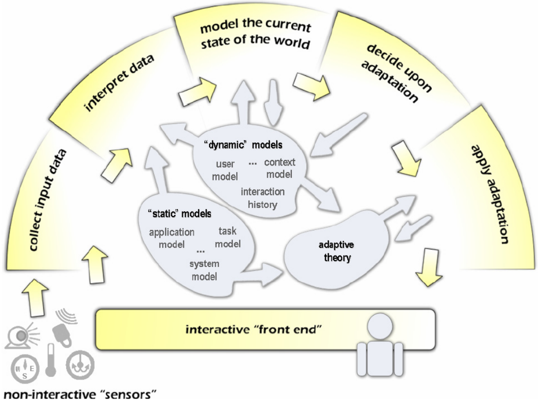

Alexandros Paramythis, Stephan Weibelzahl, Judith Masthoff. Layered evaluation of interactive adaptive systems: framework and formative methods. User Model. User-Adapt. Interact. 20(5): 383-453 (2010)

What can be adapted

The best known example of personalization are **recommendations** provided by stores like Amazon.

Other examples are **friend recommendations** in social media, such as Facebook.

Not everyone is aware of it, but Google **search results** are personalized as well.

But there are many more things that can be adapted.

Associative links

Associative links connect pages with similar topics and content. These links may be embedded in running text or displayed as related links.

Welcome To My Home Page

My name is Eelco Herder. I work as a postdoc at the L3S Research Center in Hannover, where I conduct research in the field of Human-Computer Interaction. My specialization is Adaptive Hypermedia.

I am currently involved in the European projects Grapple and Stellar, in which I work as local team leader and work package manager. I was general manager of the Prolean NoE, which finished successfully in 2008. I served as local chair of the Adaptive Hypermedia 2008 conference and as publicity chair of UMAP 2010 and UMAP 2011. Apart from project coordination, I supervise four Ph.D. students. Since 2007 I serve as the chair of the German SIG on Adaptive Systems and Interactive Systems ABIS.

My main research interests include Web personalization, user modeling, usability and HCI in general. In Grapple, we build a framework for user modeling, reasoning and profile enrichment using Web 2.0 services. Other research topics include Web usage analysis and the development of tools for Personal Information Management.

war. "Trotz der Kraft dieses katastrophalen Ereignisses konnten wir es nur sehen, weil unser Sonnensystem zufällig genau in die Mündung dieses Energie-Jets blickt", betonte Levan.

© dpa/C.Fehrholz

THEMA: Universum

SIEHE AUCH:

08/06/2011 Neues Mega-Teleskop liefert erste erstaunliche Bilder
08/06/2011 Google Earth: 'Space Station' auf dem Mars gesichtet
25/05/2011 Genaue Karte des erdnahen Universums veröffentlicht

Recommendations and other kinds of adaptive suggested links also fall in the category 'associative links'.

Local links

Step and page navigation allows people to move sequentially through pages. They provide local structure.

Step navigation typically consists of a text label (‘previous’, ‘next’) and an arrow. Step navigation can be regarded as ‘direct guidance’ through an information domain.

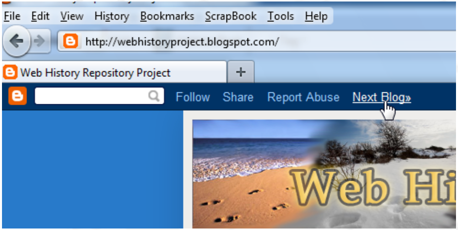

Page navigation provides additional information and options. Paging is a common method to divide large chunks of information (long texts or search result sets) in smaller pieces.

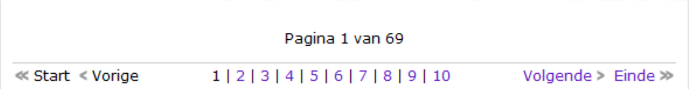

Structural links

Structural navigation connects one page to another; they impose a *hierarchy* on the hypertext structure.

The most common and most flexible form of structural navigation is the vertical or horizontal menu. The menu may be static or it may expand sub-items once a main item is chosen. There are also dynamic variations known as fly-out, pull-down or pop-up menus.

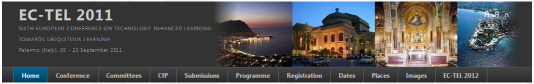

Tab navigation is a particular form of a horizontal menu. Amazon has experimented with many types of tab navigation.

Tabs have a scalability problem: there is a limited amount of horizontal space on a Web page.

Overview navigation tools

Similar to menus and navigation bars, overview navigation tools impose a (hierarchical) structure upon a Web site.

The most common overview is the site map.

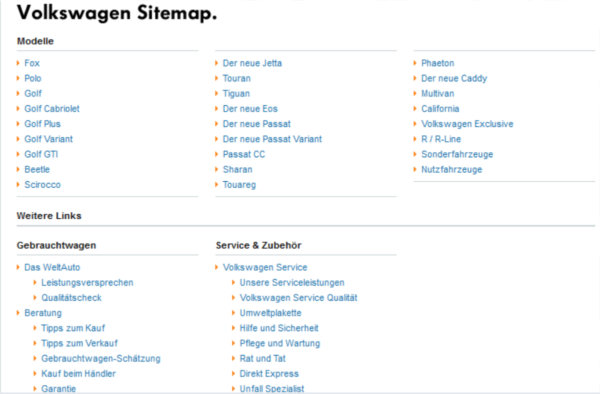

A more recent overview mechanism is the tag cloud, which lists keywords or topics related to a document or an item. The larger the tag is displayed, the more important or relevant it is.

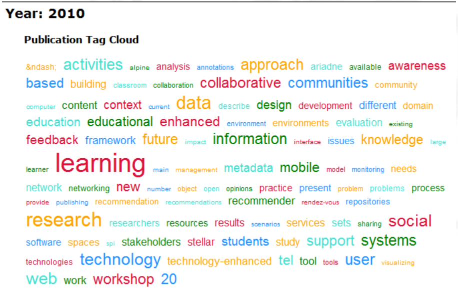

Other overviews include glossaries, categories, outlines and A-Z indexes.

Functional navigation

Functional navigation is a type of meta-navigation that brings the user to specific parts of a site or the Web that are not directly linked, or that are not part of the main content.

The most popular form of functional navigation is the search box.

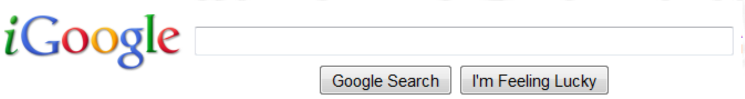

Submission forms allow people to submit information, to create an account, to provide a transaction or to obtain specific information based on their given input.

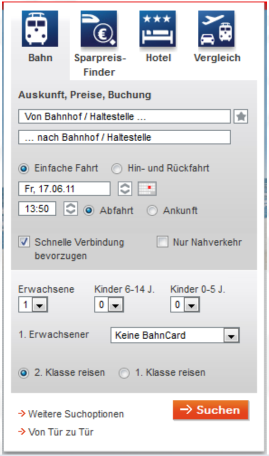

Temporal navigation

Temporal (or history) navigation tools allow users to return to pages that they have visited before.

Most temporal navigation tools are integrated in the browser: the back button, bookmarks, tabs and the history list.

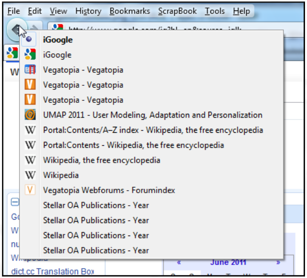

Adaptation techniques

- *Conditional text*: text that is only displayed when certain conditions are met
- *Page variants*: different versions (e.g. easy, intermediate, advanced) that can be chosen from
- *Direct guidance*: next page, next step, menu outline, trails
- *Adaptive sorting*: based on similarity, user background knowledge, relevance, ...
- *Adaptive hiding*
- *Adaptive annotation*: color coding, pop-up text
- *Adaptive overview maps*: site maps, history visualizations

Personalized e-learning

In e-learning situations the system usually knows what a learner wants to achieve.

The role of learner interest and background knowledge is also well-known.

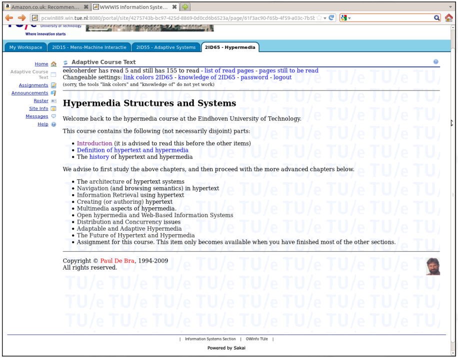

More examples of adaptive systems

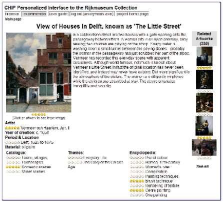

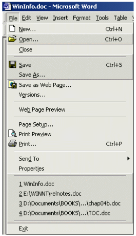

Summary

Personalization

- Personalization is an alternative to one-size-fits-all.
- Useful when there is much choice and many different needs.
- Support for routine tasks or for finding new information or products.
- Depends on knowledge about the users and their needs.
- There are many things that can be adapted.

Understanding the user

Tools such as Google Analytics show general trends:

- number of visits and users
- where do users come from, which systems do they use
- popular pages and keywords

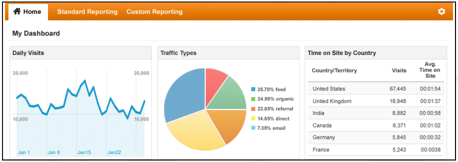

User modeling is about getting to know the *individual* user.

Which data can be of relevance?

- **Personal data, demographics**
  - Name, address, age, birthday, email address, gender, phone number, credit card information, ...
  - Education, profession, ...
- **Contacts and friends**
  - Friends' personal data, groups and group membership, chatlogs, ...
  - Social Media (Skype, Twitter, Facebook, LinkedIn, ...)

- **Device information**
  - System specs, display resolution, network speed and bandwidth, software and tools
- **Location**
  - Position, direction, speed, vehicle, . . .
- **Browsing history and bookmarks**
- *And much more*

Direct input from the user

User input (a user profile) is often gathered upon the first use of a system using forms or questionnaires. User input is also commonly gathered while the user interacts with the system.

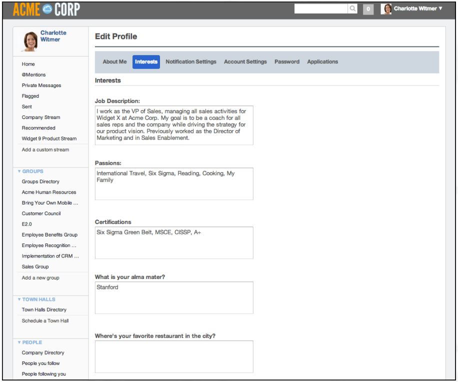

Another option is that the user gives *relevance feedback*. In recommendation systems such as Movielens, feedback is an essential part of the system, as the recommendation process mainly relies on user ratings and reviews of movies.

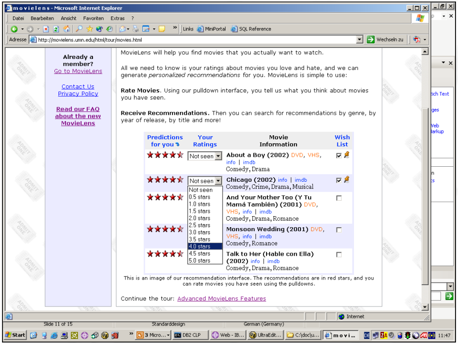

Another alternative: users make adaptations themselves by ordering lists, enabling or disabling options, dragging interface elements or by any other specific interaction with the system.

Completeness of user profiles

Users often do not fill out their profiles completely. For example, Twitter only asks 6 attributes, but these profiles are only completed up to 49%.

Would it be possible to complete user profiles by aggregating data from different sources?

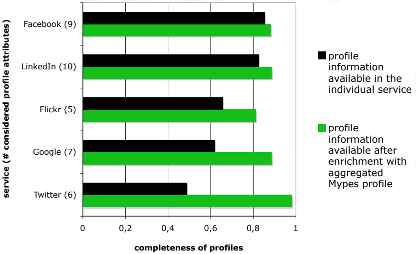

Fabian Abel, Eelco Herder, Geert-Jan Houben, Nicola Henze, Daniel Krause. Cross-system User Modeling and Personalization on the Social Web. UMUAI 23 (2-3), 2013, pp 169-209

Explicit ratings

- require additional work from the user
- in return, users get higher quality recommendations
- other motivations for rating include goodwill, having one's opinion's voiced and valued, and the ability to store their own likes and dislikes

In some domains - most notably hotel booking sites - users are particularly willing to express their opinions.

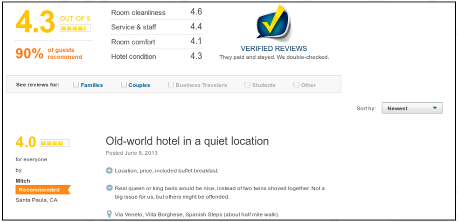

Rating scales

Different kinds of rating scales can be found on the Internet (and elsewhere), such as :

- A simple like-button: a boolean scale with two values (I like it or not)
- Five-star ratings: very popular in online stores and social networks
- Slider bars: allow for very fine-grained scales, such as 1-100

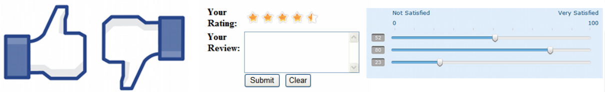

Boolean scales do not give the users sufficient possibilities to express their opinions. Boolean scales are also often too coarse-grained for collaborative filtering.

Very fine-grained scales may confuse the user: do I like 'The Hobbit' 44% or rather 47%? Fine-grained scales may make it unlikely to find users that give (exactly) the same rating to a particular item.

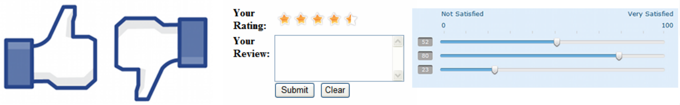

Implicit user profiling

In many cases users just want to start working on their tasks without first reading manuals, following an introductory tour or filling out forms.

Many adaptive systems attempt to infer knowledge directly by unobtrusively monitoring the user interactions with the system.

Implicit ratings

Implicit ratings are collected with little or no cost to the user

- may be based on the time spent reading information about a product
- or based on the products that the user actually bought, bookmarked or added to a wish list
- if implicit ratings are used, there is more uncertainty in the computation

Popular and frequently visited items

Many personalized systems exploit two very basic and common facts:

- A small number of products or other items is visited very frequently, a large number is visited very infrequently
- Most products or other items that are revisited have been visited quite recenty.

This distribution is called a *power law* and can be observed in many cases.

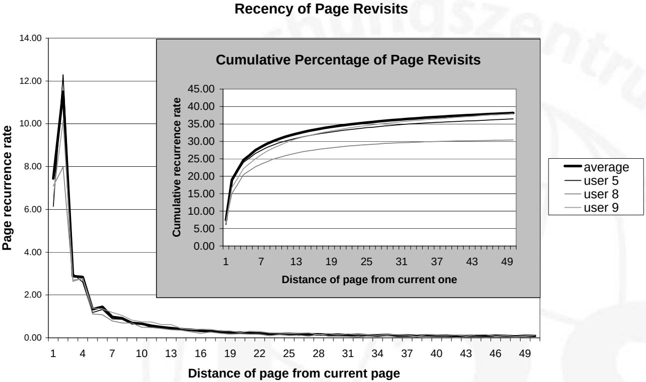

Percentage of page revisits as a function of distance from the current page. The popularity distribution shows a similar pattern.

From user data to a user model

Many interactions contain meaning in themselves, such as page visits, queries issued and items inspected or bought.

Other interactions need to be interpreted to become meaningful, such as key strokes, mouse clicks and eye gaze behavior.

The assumption for knowledge inference is that user interaction with a system is predictable to a certain extent.

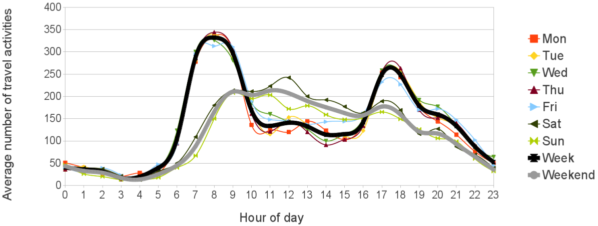

Early user models were hand-crafted and contained concrete things like user interests and personality traits. One might call them *mentalistic*:

- the user model explicitly represents the relevant aspects of the user as closely as possible.
- the model is built from human-like inferences.
- the data is easy to interpret.

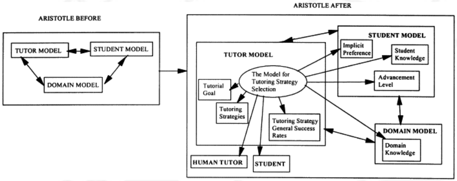

Currently, statistical models and machine learning technique have become more popular.

- implicit, statistical approaches are more flexible and better suitable for dealing with huge quantities of data.
- some approaches try to guess a user's *overall* interest in a page, document or other item.
- other techniques capture that 'if a user visits or likes X, then he might also visit or like Y'

Knowledge inference

In general, three approaches can be identified:

- **detecting patterns in user behavior**
  Useful when the aim of the adaptive system is to respond to recurrent behavior or to infer items that may be of the user's interest.
- **matching user behavior with the behavior of other users**
  Useful when a user behaves in a similar way to other users and is typically used for making recommendations involving items not seen before.
- **classifying users or products/content based on user behavior**
  Common applications include stereotyping and the modeling of user interests.

Enriching user profiles

Tim Berners-Lee

The Semantic Web is an extension of the current web in which information is given well-defined meaning, better enabling computers and people to work in cooperation.

A core building block of is the use of RDF triples: subject - predicate - object.

‘Peter is interested in Sweden’

subject: http://www.peter.de/foaf.rdf#me;
predicate: foaf:interest;
object: http://en.wikipedia.org/wiki/Sweden;

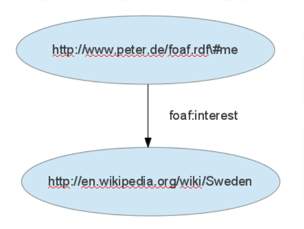

Ontologies

These simple subject-predicate-object statements can be used for describing ‘the world’.

There are several ontologies or vocabularies that can be used for these statements.
For example, Friend-of-a-Friend (FOAF) describes the relations between people and resources in the Web.

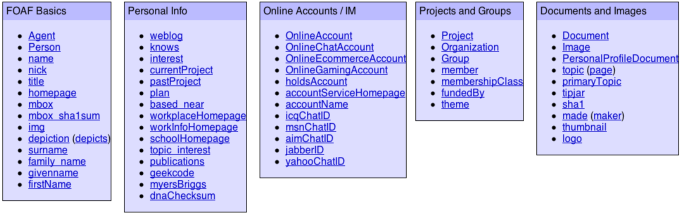

FOAF Basics

- Agent
- Person
- name
- nick
- title
- homepage
- mbox
- mbox_sha1sum
- img
- depiction (depicts)
- surname
- family_name
- givenname
- firstName

Personal Info

- weblog
- knows
- interest
- currentProject
- pastProject
- plan
- based_near
- workplaceHomepage
- workInfoHomepage
- schoolHomepage
- topic_interest
- publications
- geekcode
- myersBriggs
- dnaChecksum

Online Accounts / IM

- OnlineAccount
- OnlineChatAccount
- OnlineEcommerceAccount
- OnlineGamingAccount
- holdsAccount
- accountServiceHomepage
- accountName
- icqChatID
- msnChatID
- aimChatID
- jabberID
- yahooChatID

Projects and Groups

- Project
- Organization
- Group
- member
- membershipClass
- fundedBy
- theme

Documents and Images

- Document
- Image
- PersonalProfileDocument
- topic (page)
- primaryTopic
- tipjar
- sha1
- made (maker)
- thumbnail
- logo

Linked data

The Web enables us to link related documents. Similarly it enables us to link related data.

Linked Data refers to a set of best practices for publishing and connecting structured data on the Web.

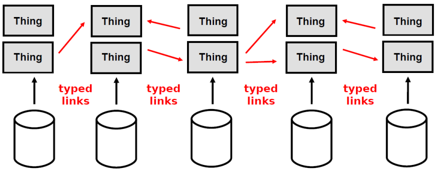

The Linking Open Data (LOD) project takes existing open data sets, makes them available on the Web in RDF and interlinks them with other data sets.

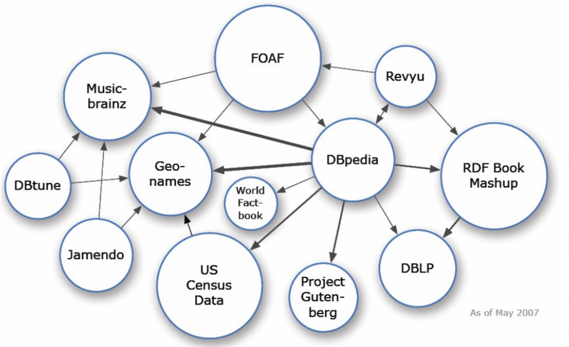
FOAF
Revyu
Music-brainz
Geo-names
DBpedia
RDF Book Mashup
DBtune
World Fact-book
Jamendo
US Census Data
Project Guten-berg
DBLP
As of May 2007

LOD in 2007

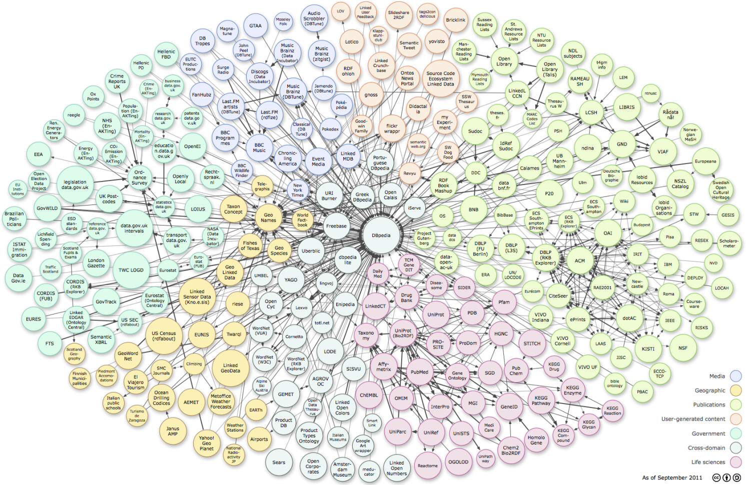
Media
Geographic
Publications
User-generated content
Government
Cross-domain
Life sciences
As of September 2011

LOD in 2011

Making use of linked data principles and the "LOD Cloud", several pieces of knowledge can be connected.

For example, making use of DBPedia, we can infer that Richard Cyganiak lives in a city with a population of 3.405.259.

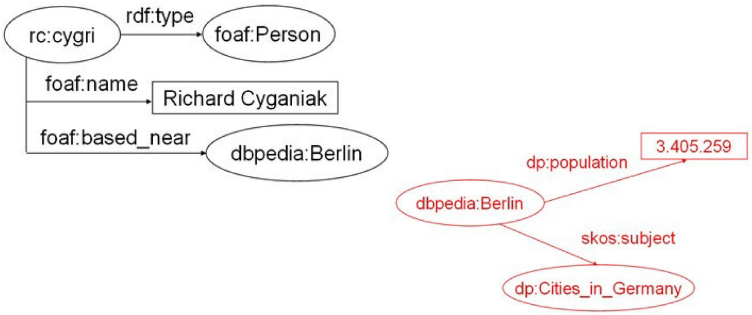

Summary

User models

- User models are based on personal data, friends, device information, usage data and much more.
- Direct input from the user: profiles and explicit ratings.
- Indirect profiling: usage logs, buying history.
- Indirect data and interpreted data contain *uncertainty*
- Linked data and ontologies can be used for enriching user profiles

Personalization techniques

Most information-oriented websites are not personalized. They provide the same kind of content to anyone.

In many cases, this is fine. Users find the information they need - as long as the navigation structure is understandable.

Semi-personalized systems

Semi-personalized systems do not make use of a user profile. They try to adapt the content to the (estimated) needs *of everyone*. For example:

- trending topics in Twitter
- most read items on a newssite

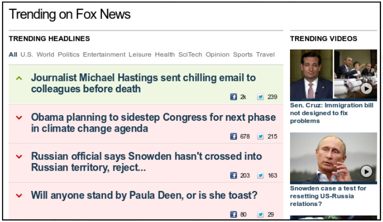

A personal touch and personal functionality

Particularly transaction sites (online stores) require users to log in or to create an account.
It is customary that these sites 'welcome' the user, which creates some basic 'mutual recognition'.

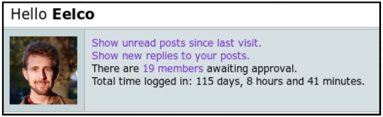

Why bother to log in

A personal greeting is not a very convincing reason for asking users to register.
More convincing - and practical:

- Easier, quicker checkout (less information to fill out when buying an item
- Access to previous orders, perhaps with the probability to re-order
- Saving of default location, default language, default values

I personally really appreciate this sort of functionality - most users expect (ecommerce) sites to offer such features.

Meine Übersicht

Hallo, Herr Eelco Herder!
Von Ihrer Benutzerkonto-Übersicht aus haben Sie die Möglichkeit, Ihre letzten Vorgänge einzusehen und Ihre Benutzerkonto-Daten zu bearbeiten. Wählen Sie dazu einen der untenstehenden Links, um Informationen anzusehen oder zu bearbeiten.

Letzte Bestellungen

Alle ansehen

Bestellung	Datum	Senden an	Bestell -summe	Status	
200156704	02.06.13	Herr Eelco Herder	40,86 €	*Bestellung komplett*	Bestellung ansehen | **Nachbestellen**
200144927	22.04.13	Herr Eelco Herder	40,94 €	*Bestellung komplett*	Bestellung ansehen
200134329	14.03.13	Herr Eelco Herder	73,44 €	*Bestellung komplett*	Bestellung ansehen
200122541	30.01.13	Herr Eelco Herder	49,96 €	*Bestellung komplett*	Bestellung ansehen
200109249	07.12.12	Herr Eelco Herder	47,38 €	*Bestellung komplett*	Bestellung ansehen

Personalization goes beyond this ‘basic’ functionality.

Adaptability versus adaptivity

- *Adaptable* means that users can adapt system behavior themselves (e.g. iGoogle personalized start page)
- *Adaptive* means that the system adapts its own behavior on the user's behalf (e.g. Amazon recommendations)

Recommender systems

Definition

Recommender systems work from a specific type of information filtering system technique that attempts to recommend items that are likely to be of interest to the user.

Items may be: movies, tv programs, music, books, news, images, web pages, scientific literature, . . .

Collaborative filtering

For years, people have stood over the back fence or in the office break room and discussed books they have read, restaurants they have tried, and movies they have seen. And they used these discussions to form opinions.

At some point, you might observe that among your friends:

- Matt recommends the types of films that you like
- Paul typically recommends films that you despise
- And Margaret simply recommends everything.

Over time, you learn whose opinions you should listen to.

Schafer, J.B., Frankowski, D., Herlocker, J. and Sen, S. (2007) Collaborative Filtering Recommender Systems. The Adaptive Web, 291-324

Assumptions behind collaborative filtering

- There are other users with common needs or tastes
- People with similar tastes will rate things similarly
- Taste persists: CF has been successful for movies, books and electronics. If tastes change frequently, older ratings may be less useful (e.g. clothing)

- Item evaluation is personal: if objective (content-based) criteria for goodness are more relevant, collaborative filtering may not be very useful
- Items persist long enough to receive sufficient ratings: news stories are only important for a short time, which hinders CF
- Items are sufficiently homogeneous: for example music albums. Recommendations such as 'if you buy a hammer, you might also want to buy a refrigerator' are not very useful.

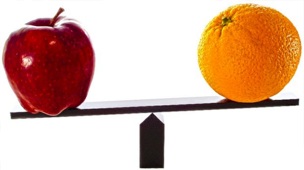

How does collaborative filtering work?

Traditional approach, in a nutshell:

- Identify the users who bought or liked the same items as you did (the *neighborhood*)
- Recommend items that the neighborhood bought or liked best

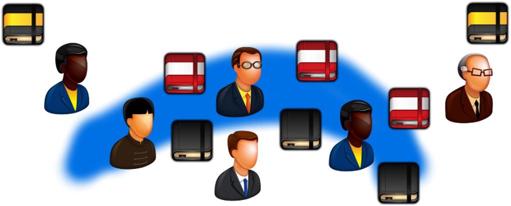

More scalable approach:

- Identify the items that are liked by all other users the same way as you like them
- Recommend items that are most similar to the ones that you like best

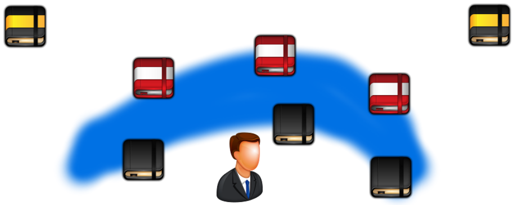

Content-based recommendation

Content-based recommendations are a useful alternative for collaborative filtering techniques, in various situations, such as:

- There are objective criteria for goodness (e.g. price and resolution of a camera)
- Items do not persist long enough to receive sufficient ratings (e.g. news stories)
- There are insufficient ratings in the system for collaborative filtering
- User ratings are not desirable or feasible (e.g. web site recommenders)

In summary: recommender systems

Pros and cons

- Collaborative filtering assumes that people with similar tastes will rate things similarly.
- Content-based recommendation assumes that there are objective criteria for ‘goodness’.
- Hybrid recommender systems combine both approaches.
- Although relatively successful, they only consider a user as someone who ‘likes’ things and ‘is like’ other people.

Returning users: the importance of the long tail

An important premise of personalization is that user behavior is relatively predictable. In about 80% of our actions that is the case:

- We have a couple of close friends that we visit very frequently.
- In the supermarket, we usually buy the same products.
- There are a couple of Web sites that we visit every day.

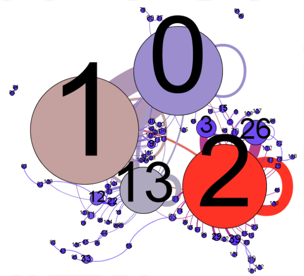

Recency and popularity in Web revisitation

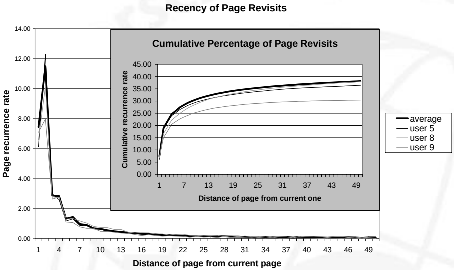

Most sites that we visit, are sites that we visited recently before.
Most users have a small number of sites that they visit frequently.

Hartmut Obendorf, Harald Weinreich, Eelco Herder and Matthias Mayer. Web Page Revisitation Revisited: Implications of a Long-term Click-stream Study of Browser Usage.. Proc. CHI 2007

Interestingly, there are some differences between types of sites.

- Search engines have just one most popular page
- Institutional web sites have several pages that are visited frequently
- News sites have a small number of frequently visited news categories

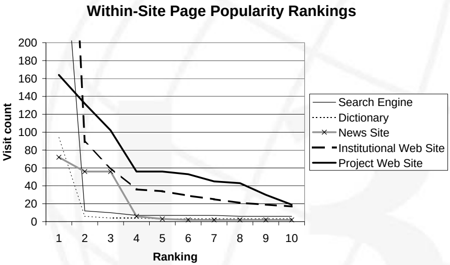

Routine behavior

If we visit an address or buy a product very frequently, we can usually find it very easily again.

Does it make sense to recommend very frequently visited items?

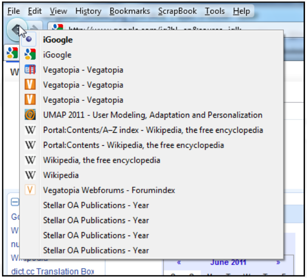

Yes: it makes routine tasks easier to carry out.

But what about less frequent things?

I know the URLs of the Web sites that I visit frequently by heart. They are also in my bookmarks, my history, and the Web browser suggests it too when I type.

But what about that hotel that I found about a year ago?
What was its name? In which city exactly? Did I find it on booking.com or somewhere else? Was it in July or in September?

The irony is: once the time since the last visit was long enough for me to forget about it, it is also long enough for the browser to 'forget' it.
I have to dig into my email archives or search for it again.

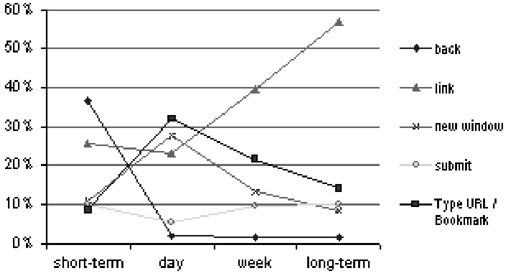

Hartmut Obendorf, Harald Weinreich, Eelco Herder and Matthias Mayer. Web Page Revisitation Revisited: Implications of a Long-term Click-stream Study of Browser Usage.. Proc. CHI 2007

The Pivotbar: a dynamic browser toolbar

The Pivotbar provides personalized recommendations for Web pages or sites to be revisited. If a user would visit travel-related sites, it would recommend hotels visited in the past.

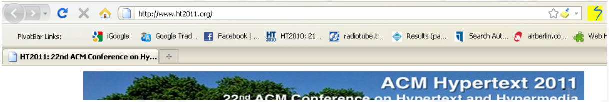

In a user evaluation we verified that better recommendations go hand-in-hand with better take-up.

Ricardo Kawase, Georgios Papadakis, Eelco Herder, Wolfgang Nejdl. Beyond the Usual Suspects: Context-Aware Revisitation Support. Proc. Hypertext 2011

Summary

Personalization techniques

- Personalization is not always needed, but in many cases very useful.
- Users expect access to previous orders and other account-related data.
- Recommendation techniques are used for finding or suggesting new items.
- Better support is needed for infrequent actions, such as products that one orders only once in a while.

Different voices on personalization

The ambition of adaptivity is that not only ‘everyone should be computer literate’, but also that ‘computers should be user literate’ *(Browne, 1990)*

Personalization is a designers’ approach to achieve harmony between users, tasks, environments and the system *(Benyon, 1993)*

Personalization is an overrated concept. Rather than investing time and energy on trying to predict individual users’ needs it would be better to enhance the overall site design. *(Nielsen, 1998)*

How personalized is the web?

Transaction-oriented sites - such as online stores and travel planners - usually maintain profiles of their users and offer access to previous transactions. *Users expect this to be standard functionality.*

Support for routine behavior is generally focused on tasks that are carried out very frequently, *the long tail of infrequent activities is largely ignored*.

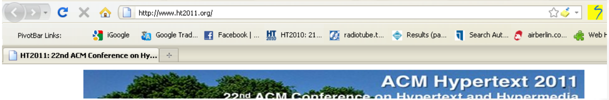

Many sites attempt to provide recommendations, but the quality of the recommendations varies.

- Often due to lack of sufficient data about their users.
- User interest in some items, such as news stories, is hard to predict.
- Even Amazon, a pioneer in recommendations, still often provides irrelevant recommendations.

Current recommenders are definitely not yet like your partner, your best friend or your mother, who usually know a lot about you.

The filter bubble

Personalization in the background, such as Google search results, may be harmful.
If search results or news articles are always personalized, this means that no two persons will see the same results.

Users will only get results that match their own viewpoints, isolating them in their own *filter bubbles*.
(To make things worse, the personalization process still needs improvement)

Possible consequences:

- Search results may be relevant, but may not challenge your current thinking.
- Online stores recommend books and movies that are entertaining, but not necessarily informative or useful ones.
- One may loose contact with people with other (political or ideological) views.

Eli Pariser. Beware online "filter bubbles". TED Talk, March 2011.
http://www.ted.com/talks/eli_pariser_beware_online_filter_bubbles.html

What about our personal data?

Personalization depends on the availability of data about the user. Data collection practices and privacy policies are often critized in the media.

Chris Hoofnagle

"We now have ten years of experience with privacy selfregulation online, and the evidence points to a sustained failure of business to provide reasonable privacy protections"

A particular problem is that user data is owned and controlled by the site or company that has collected it, not by the users themselves.

Michael Zimmer. The Externalities of Search 2.0: The Emerging Privacy Threats when the Drive for the Perfect Search Engine meets Web 2.0. First Monday 13 (3), March 2008

Future directions - a wishlist

Open user data and the adoption of linked data practices would facilitate personalization on smaller sites with less usage data.

User control is essential. Users should be able to *scrutinize* their model in one way or another.

Transparency in personalization is much needed. One should be able to turn it off.

Personalization comprises more than generating and displaying lists of recommendations. Many concepts are still waiting for adoption on the web.