Personalization for Recurrent Activities

Eelco Herder

L3S Research Center / Leibniz University of Hanover
Hannover, Germany

April 14, 2016

What is personalization?

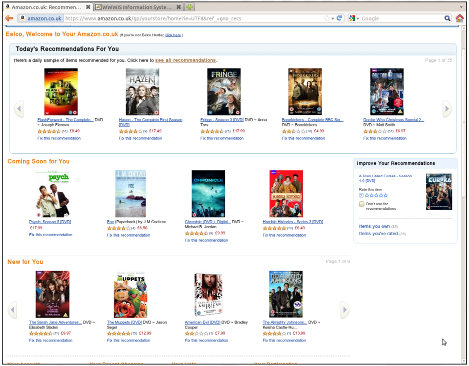

Product recommendations in Amazon. These recommendations are based on past purchases and past browsing behavior. The user can improve the recommendations by editing his or her user profile.

Personalization in the background

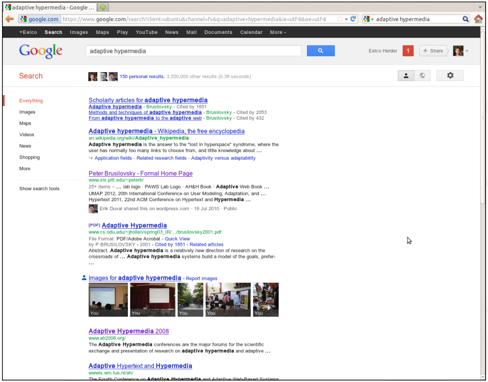

Google search results are personalized, based on past searches, current location, language settings (apparently 57 features in total).

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

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

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.

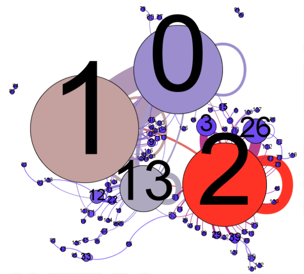

Recency and popularity in Web revisitation

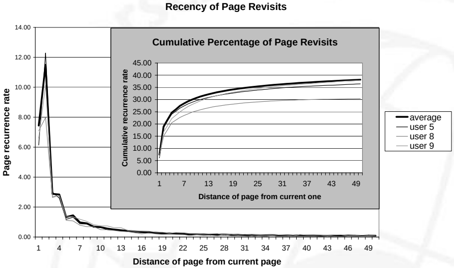

Most sites that we visit, are sites that we visited recently before.
Most users have a small number of sites that they visit frequently.

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

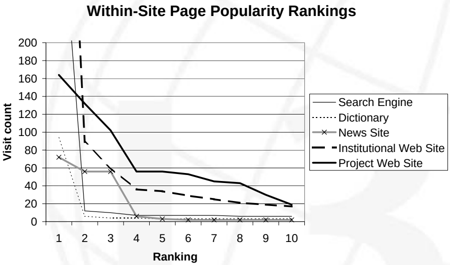

Routine behavior

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?

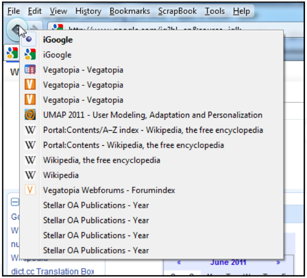

Yes: it makes routine tasks easier to carry out.

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.

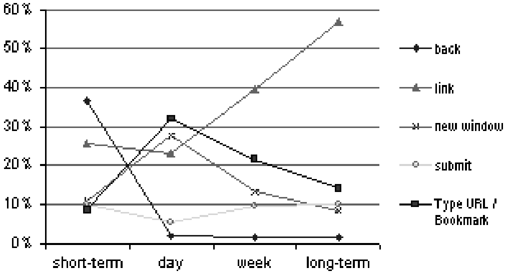

Hartmut Obendorf, Harald Weinreich, Eelco Herder and Matthias Mayer. Web Page Revisitation Revisited: Implications of a Long-term Click-stream Study of Browser Usage.. Proc. CHI 2007

How well can we predict routine Web revisits?

We experimented with the Web History Repository dataset.

- A Firefox plugin that invites the users to upload their usage data
- Browsing data of 180 users with more than 1000 page requests.
- Recorded between 2009 and 2011
- 1.5 million Web pages of which less than 15% were revisited.

George Papadakis, Ricardo Kawase, Eelco Herder, Wolfgang Nejdl. Methods for web revisitation prediction: survey and experimentation (2015). User Modeling and User Adapted Interaction (UMUAI) 25 (4), pp. 331-369.

A-priori prediction methods

These methods capture the *overall* chance that a user will revisit a page.

- Last Recently Used (recent items)
- Most Frequently Used (popular items)
- Frecency (frequency with bonus points for recent visits)
- Polynomial Decay (balance between recency and popularity)

$$DEC(m_i, I_{m_i}, i_n) = \sum_{j=1}^{|I_{m_i}|} \frac{1}{1 + (i_n - i_j)^{\alpha}},\ \alpha > 0$$

where i_k is the index of the chronologically last transaction in I_{m_i}, i_n is the index of the latest request of the system or user, and $|I_{m_i}|$ is the cardinality of I_{m_i}.

Propagation methods

These methods aim to identify pages that are often visited together.

- Simple transition matrices (preserves order)
- Symmetric association matrices (order-independent)
- Association rules (not practical for large numbers)

In addition, **drift methods** can be applied to minimize the influence of the more distant past.

- Decay-based (computationally expensive)
- Sliding window (limits computational requirements, can be of arbitrary length)

Experiment

We experimented with several single (a-priori, propagation and drift) methods and combinations of these methods.

We used two conditions:

- The full search space (all pages are candidates for revisits)
- An optimized search space (with an oracle that knows which pages will never be revisited)

We simulated the navigation of each user independently by 'replaying' their history and making predictions along the way.

Results in a nutshell

Of the individual methods, the *a-priori* method **Polynomial Decay** performed best (S@10=44%).

- A-priori methods also have low space and time requirements

Propagation methods capture other aspects of user navigation

- Combining Polynomial Decay with a **simple transition matrix** significantly improves performance (S@10=61%)

Combining the above methods with sliding-window *drift methods* does more harm than good.

- The benefits of a larger navigation history are higher

The *revisitation oracle* improved performance and decreased computation time.

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.

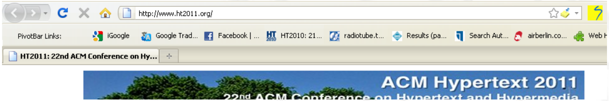

In a user evaluation we verified that better recommendations go hand-in-hand with better take-up

- The correlation between mouse clicks and blind hits was $r = 0.92; p < 0.01$.

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

Location-Based Services

Location-based services suggest new locations that match the user's inferred interests and preferences

- making use of content-based or collaborative recommendation techniques.
- distance is often used as the main criterion for inclusion in the recommendations.

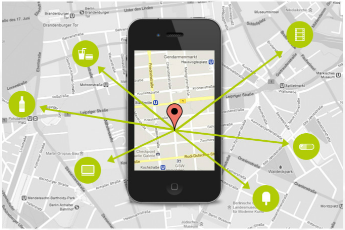

Eelco Herder, Patrick Siehndel and Ricardo Kawase (2014). Predicting User Locations and Trajectories. Proc. UMAP 2014

Daily Routines

Apart from visiting new locations, users often visit places that they visited before.

- Home and work
- Specialty stores
- Hiking areas
- Friends and relatives.

Support for Routine Activities

Locations most searched for: restaurants, stores, attractions, leisure.

- New locations, based on preferences and current location

Effective prediction of routine mobility patterns opens opportunities for recommendations and support of routine activities

- activities or locations to be included in your schedule
- minimize traveling time between destinations

Goals of Our Study

- We analyze, visualize and discuss patterns found in a dataset of GPS trajectories.
- We compare and analyze the performance of common prediction techniques that exploit the locations' popularity, recency, regularity, distance and connections with other locations.

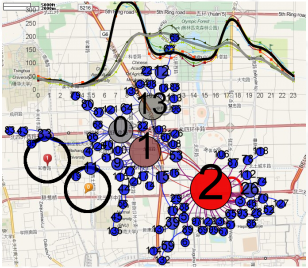

Datasets used

- The GeoLife GPS Trajectory Dataset (2012): a total of 17,621 trajectories from 178 users, mainly located in Beijing.
- The MSR GPS Privacy Dataset: 4,165 trajectories from 21 users, mainly located in and near Seattle.

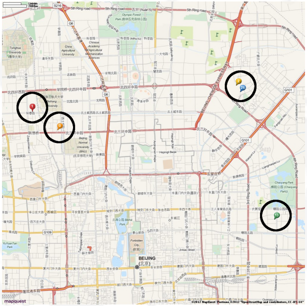

Overall Travel Activity

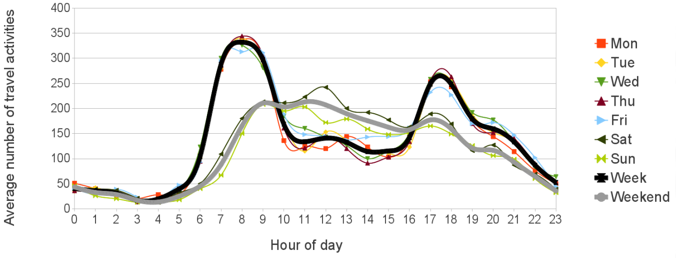

- Morning and evening rush hour on weekdays.
- Small travel peak during lunchtime.
- Weekend traffic starts later and remains stable during the day.

Connections between locations

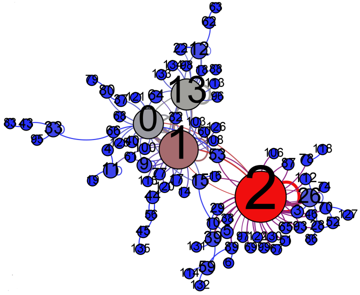

- Graph of one exemplary user, made using Gephi.
- A small number of frequently visited locations.
- Other locations connected to one of these main locations.

Predicting Future Locations

We compare five basic methods for predicting when a person will revisit a particular location.

We only consider basic methods and do not attempt to find optimal combinations of these methods.

Our purpose is to verify the performance of each method and to what extent these prediction methods are correlated.

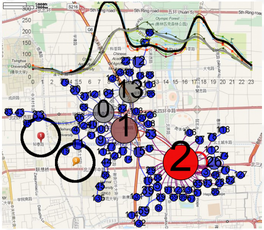

Prediction Methods

- *Top-N locations* - most popular places (baseline).
- *Last-N locations* - last visited places.
- *Hour top-N locations* - top-n locations for a particular hour.
- *Top-N closest locations* - closest to the current location (often used in location-based services).
- *Simple Markov Model* - the probability that a user will travel to some location starting from the current location.

Evaluation Measures

As we are interested in predicting locations that people will revisit, we apply the above-mentioned methods to each user individually.

We use the success rates $S@1$ and $S@5$:

- for many applications it is sufficient if the next location is included in a small set of recommendations

We also report the *Shannon entropy*

- indicates to what extent the predictions cover the whole set of frequently and less frequently visited locations,

Success Rates

	Top-N	Last-N	Hour	Distance	Markov
$S@1$	0.286	0.204	0.467	0.275	**0.626**
$S@5$	0.612	0.546	0.829	0.49	**0.931**

- *Distance-based* (closest locations) performs worst - most location-based services consider distance as an important factor.
- *Top-N* and *Last-N* have moderate performance.
- *Hour-based* performs significantly better.
- A simple *Markov Model* predicts the correct location in 62% of the cases - $S@5$ is 93%.

Entropy

Low entropy measures indicate that a method often suggests the same locations.

	Top-N	Last-N	Hour	Distance	Markov	Actual
Top-1	0	4.142	1.803	4.165	2.852	4.139
Top-5	2.322	4.635	4.201	4.896	3.157	-

- Not surprisingly, *Top-N* has a low entropy.
- *Last-N* and *Distance* reach the highest entropy, but had rather low success rates.
- The *Markov model* has reasonable entropy values (and good success rates).

Weekdays versus Weekend

We repeated our experiments with separate models for weekdays and weekend, with more or less similar results..

$S@5$ for *Top-N* and *Last-N* was *7% higher* during weekdays and about *7% lower* during the weekend

- confirms our observation that weekend patterns are *less stable* than weekday patterns.

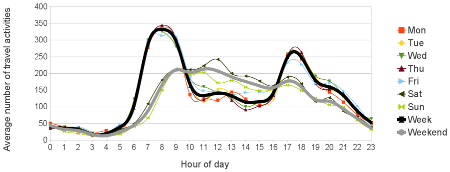

Summary of Findings

Human mobility patterns contain strong regularities:

- most time is spent on a small number of (popular) locations.
- these popular locations serve as *starting points* for visits to other locations.
- the *purpose* of locations depends on the time of day.

Basic predictive methods have reasonable performance.

- *Markov models* have the best performance.
- *Distance* seems less important (even though often used by location-based services).

Entropy and correlation measures provide indications on how they can be combined in more complex models.

Design Implications

Most location-based services focus on the recommendation of *new locations*.

Strong regularities in human behavior form a basis for supporting *everyday activities* involving already visited locations.

Most locations can be connected to one *base location*. This can be exploited in various ways:

- recommendations for regular stops on the way back home
- targeted advertisements on a Saturday-morning shopping trip