

Social Recommendations at Work

Tom Crecelius[†], Mouna Kacimi[‡], Sebastian Michel[‡], Thomas Neumann[‡],
Josiane Xavier Parreira[‡], Ralf Schenkel[‡], Gerhard Weikum[†]

[†]Max-Planck-Institut für Informatik
Saarbrücken, Germany

[‡]École Polytechnique Fédérale de Lausanne
Lausanne, Switzerland

socialsearch@mpi-inf.mpg.de

ABSTRACT

Online communities have become popular for publishing and searching content, and also for connecting to other users. User-generated content includes, for example, personal blogs, bookmarks, and digital photos. Items can be annotated and rated by different users, and users can connect to others that are usually friends and/or share common interests.

We demonstrate a social recommendation system that takes advantages of users connections and tagging behavior to compute recommendations of items in such communities. The advantages can be verified via comparison to a standard IR technique.

Categories and Subject Descriptors: H.4 [Information Systems Applications]: Miscellaneous; H.3.3 [Information Search and Retrieval]: Search Process

General Terms: Experimentation, Performance.

Keywords: social networks, scoring and ranking.

Introduction

With the advent of Web 2.0, new online services offer the possibility to store personal content, to share it with other people, and to explore other users' contents. Well-known examples of such social-community platforms are del.icio.us, YouTube, Flickr and LibraryThing. While differing in the type of content (e.g., videos, photos, bookmarks), most platforms follow a common pattern: *Users* join the community, then produce information by publishing their own *items* and by adding *tags*, ratings or comments to other content already available. Users additionally maintain a list of *friends*, which is often initially populated with people they already know from the offline world or other online communities, but over time they identify previously unknown users they share common interests with and add those users to their friends lists.

Social tagging offers an opportunity to exploit the “wisdom of the crowds” by identifying valuable content that is recommended by friends – either directly or indirectly (i.e., via friends of friends) and explicitly (e.g., ratings) or implicitly (e.g., by intensive tagging). To this end, the relationships among users should be taken into account for ranking, the intuition being that you trust the recommendations of your close friends more than those of your casual acquaintances. This situation resembles the paradigm of collaborative recommendation [1], which applies data min-

ing on customer-product and similar usage data to predict items that users are likely interested in. However, the high dynamics of social communities and the very high rate of content production and tagging calls for highly *efficient* and *scalable* methods.

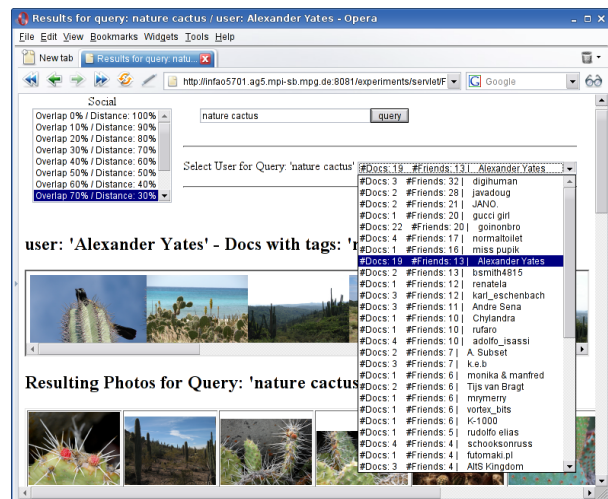


Figure 1: Screenshot of the user interface

Demonstration

The demo presents a social recommendation system that computes recommendations based on tags and items of related users efficiently at runtime, aggregating results from explicitly selected or implicitly defined friends of the querying user. The system provides different weighting schemes for social relations, including weights based on overlap of tag usage and distance in the social network. For leveraging the “social wisdom”, the system employs a new form of two-dimensional expansions: social expansion considers the strength of relations among users, and semantic expansion considers the relatedness of different tags. The usefulness of the generated recommendations is showcased with two different real-life datasets: a large crawl of Flickr with 10 million pictures and more than 50,000 users, and a large crawl of librarything.com with more than 6 million books and about 10,000 users.

REFERENCE

- [1] J. L. Herlocker et al. Evaluating collaborative filtering recommender systems. *ACM Transactions on Information Systems*, 22(1), 2004.