

Henry Kautz, Bart Selman,
and Mehul Shah

Referral Web:

Combining Social Networks and Collaborative Filtering

*An interactive system for restructuring, visualizing, and
searching social networks on the Web.*

NUMEROUS STUDIES HAVE SHOWN THAT ONE OF THE MOST EFFECTIVE channels for disseminating of information and expertise within an organization is its informal network of collaborators, colleagues, and friends [1, 4, 7]. Indeed, the social network¹ is at least as important as the official organizational structure for tasks ranging from immediate, local problem-solving (for example, fixing a piece of equipment), to primary work functions, such as creating project teams.

Part of the success of social networks can be attributed to the “six degrees of separation” phenomena that means the distance between any two individuals in terms of direct personal relationships is relatively small. An equally important factor is there are limits to the amount and kinds of information a person is able or willing to make available to the public at large. For example, an expert in a particular field is almost certainly unable to write down all he knows about the topic, and is likely to be unwilling to make letters of recommendation he or she has written for various people publicly available. Thus, searching for a piece of information in this situation becomes a matter of searching the social network for an expert on the topic together with a chain of personal referrals from the searcher to the expert. The referral chain

serves two key functions: It provides a reason for the expert to agree to respond to the requester by making their relationship explicit (for example, they have a mutual collaborator), and it provides a criteria for the searcher to use in evaluating the trustworthiness of the expert.

Nonetheless, manually searching for a referral chain can be a frustrating and time-consuming task. One is faced with the trade-off of contacting a large number of individuals at each step, and thus straining both the time and goodwill of the possible respondents, or of contacting a smaller, more focused set, and being more likely to fail to locate an appropriate expert. In response to these problems we are building ReferralWeb, an interactive system for reconstructing, visualizing, and searching social networks on the World-Wide Web. Simulation experiments we ran before we began construction of ReferralWeb showed that automatically generated referrals can be highly

¹Note that by “social network” we explicitly include groups of people linked by professional activities. We do not use the word “social” in the more limited sense of “recreational.”

successful in locating experts in a large network [3]. The higher degree of responsiveness of an automated system can be effectively traded-off for the (potentially) lower accuracy of individual automatic referrals over those generated by a person.

Reconstructing and Querying Social Networks

A social network is modeled by a graph, where the nodes represent individuals, and an edge between nodes indicates that a direct relationship between the individuals has been discovered. There are many possible sources for determining direct relationships. At one extreme, which we reject as too burdensome, users could be required to enter lists of close colleagues. Analyzing of email logs provides a rich source of relationships, as shown by Schwartz and Wood [5]. In fact, the initial version of our system derived its network by analyzing mail archives [3]. However, the use of such information raises concerns of privacy and security that are hard to allay.

The current ReferralWeb system uses the co-occurrence of names in close proximity in any documents publicly available on the Web as evidence of a direct relationship. Such sources include:

- Links found on home pages
- Lists of co-authors in technical papers and citations of papers
- Exchanges between individuals recorded in net-news archives
- Organization charts (such as for university departments)

The network model is constructed incrementally. When a user first registers with the system, it uses a general search engine to retrieve Web documents that mention him or her. The names of other individuals are extracted from the documents. This can be done with high degree of accuracy (better than 90%), using techniques such as those described in [6]. This process is applied recursively for one or two levels, and the result merged into the global network model.

The network is then used to guide the search for people or documents in response to user queries. Most simply, a person may ask to find the chain between

himself/herself and a named individual (“What is my relationship to Marvin Minsky?”). To search for an expert, the user may specify both a topic and an effective social radius: for example, a query might be, “What colleagues of mine, or colleagues of colleagues of mine, know about simulated annealing?” Another kind of query takes advantage of a designated, known expert to control the search. For instance, one might ask to “List documents on the topic ‘annealing’ by people close to Scott Kirkpatrick.”

It is important to emphasize that ReferralWeb does not *replace* generic search engines such as AltaVista, but instead uses the social network to make their search more focused and effective. For example, in the previous scenario, the requirement of proximity to (the computer scientist) Kirkpatrick helps disambiguate the query between the computer science use of the term “annealing” and the use of the same term in metallurgy. The social network also prioritizes the answers, in that the user retrieves hits on people that are closest to him or herself, rather than simply a long list of hundreds of names or documents.

ReferralWeb has a number of features that sets it apart from many other collaborative filtering and recommender projects. Among those are the following:

- ReferralWeb attempts to uncover existing social networks, rather than provide a tool for creating new communities. While new communities may be appropriate for recreational uses of the Web, we emphasize helping individuals make more effective use of their large, existing networks of professional colleagues.
- While recommender systems are often designed to provide *anonymous* recommendations, ReferralWeb is based on providing referrals via chains of named individuals. This is critical because not all sources of information are equally desirable. For example, a user may choose to search for referrals to people who are closely associated with some known, trusted expert.
- Some recommender systems require the user to manually enter a personal profile of interests, preferences, or expertise. Recommendations are generated by matching profiles that exist within the

The ReferralWeb system is part of our effort to develop agent-based programs that address practical communication needs

system. By contrast, ReferralWeb primarily builds its model of its users' social network by data mining public documents found on the Web. This model includes many more individuals than those who explicitly register with the service.

- Users of ReferralWeb are not limited to any set of topic areas determined in advance. ReferralWeb uses a general full Web indexing engine (currently AltaVista) to match individuals to topic areas.

The ReferralWeb system is part of our effort to develop agent-based programs that address practical communication needs [2]. We have developed several internal prototypes of the system, and are currently building a version that will be made accessible externally on the Web.

Finally, the ReferralWeb system expands users awareness of their existing communities. Typically a user is only aware of a portion of the social network to which he or she belongs. By instantiating the larger community, the user can discover connections to people and information that would otherwise lay hidden over the horizon. **C**

REFERENCES

1. Granovetter, M. Strength of weak ties. *Amer. J. Sociology* 78 (1973), 1360–1380.
2. Kautz, H., Selman, B., and Coen, M. Bottom-up design of software agents. *Commun. ACM* 37, 7 (1994) 143–146.
3. Kautz, H., Selman, B., and Milewski, A. Agent amplified communication. In *Proceedings of AAAI-96* (Portland, Oreg.) MIT Press, Cambridge, Mass. 1996, 3–9.
4. Kraut, R., Galegher, J. and Edigo C. *Intellectual Teamwork: Social and Technological Bases for Cooperative Work*. Lawrence Erlbaum, Hillsdale, NJ, 1990.
5. Schwartz, M. F. and Wood, D. C. M. Discovering shared interests using graph analysis. *Comm. ACM* 36, 8 (1993), 78–89.
6. Sundheim, B. and Grishman, R., Eds. In *Proceedings of the Sixth Message Understanding Conference (MUC-6)*. Morgan Kaufmann, San Francisco, Calif., 1995.
7. Wasserman, S. and Galaskiewicz, J., Eds. *Advances in Social Network Analysis*. Sage, Thousand Oaks, Calif., 1994.

HENRY KAUTZ (kautz@research.att.com) is head of the Principles of Artificial Intelligence Research Department at AT&T Labs, Florham Park, N.J.

BART SELMAN (selman@research.att.com) is a research scientist at AT&T Labs, Florham Park, N.J.

MEHUL SHAH (shah@mit.edu) is a Masters student at M.I.T in the AT&T/M.I.T VI-A co-op program.

Permission to make digital/hard copy of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication and its date appear, and notice is given that copying is by permission of ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.

© ACM 0002-0782/97/0300 \$3.50