

Computing Word-of-Mouth Trust Relationships in Social Networks from Semantic Web and Web2.0 Data Sources

Tom Heath¹, Enrico Motta¹, Marian Petre²

¹Knowledge Media Institute and Centre for Research in Computing

²Department of Computing and Centre for Research in Computing

The Open University, Walton Hall, Milton Keynes, MK7 6AA, United Kingdom

{t.heath, e.motta, m.petre}@open.ac.uk

Abstract. Social networks can serve as both a rich source of new information and as a filter to identify the information most relevant to our specific needs. In this paper we present a methodology and algorithms that, by exploiting existing Semantic Web and Web2.0 data sources, help individuals identify *who in their social network knows what*, and *who is the most trustworthy source of information on that topic*. Our approach improves upon previous work in a number of ways, such as incorporating topic-specific rather than global trust metrics. This is achieved by generating *topic experience* profiles for each network member, based on data from *Revyu* and *del.icio.us*, to indicate who knows what. Identification of the most trustworthy sources is enabled by a rich trust model of information and recommendation seeking in social networks. Reviews and ratings created on *Revyu* provide source data for algorithms that generate *topic expertise* and *person to person affinity* metrics. Combining these metrics, we are implementing a user-oriented application for searching and automated ranking of information sources within social networks.

1 Introduction

Social networks can serve as both a rich source of new information and as a filter to identify the information most relevant to our specific needs. Making optimal use of the knowledge within our social networks requires that we know firstly *who knows what*, and secondly *who is the most appropriate source of information on that topic*. In this paper we present a methodology and algorithms that address these issues by exploiting existing Semantic Web and Web2.0 data sources. Our approach supports an application that helps the user identify which members of their social networks may have knowledge on a particular topic, and of which topics each member of their network has knowledge. This is achieved by generating *topic-experience* profiles for each known person based on data from *Revyu* [4] reviews and ratings, and *del.icio.us*¹ social bookmarks.

¹ <http://del.icio.us/>

The second requirement is addressed by a rich trust model of information and recommendation seeking in social networks, based on previous empirical research. Reviews and ratings created on *Revyu* provide source data for algorithms that generate *topic expertise* and *person to person affinity* metrics. Combining all metrics derived in this fashion, we are implementing a user-oriented application for searching and automated ranking of information sources within social networks.

This paper describes in detail our methodology and algorithms for computing trust relationships, and briefly outlines the application we are developing that makes use of them. After reviewing related work in Section 2, Section 3 outlines the advantages of our approach. In Section 4 we summarize the findings of a previous study into how people choose sources for word of mouth recommendations. Section 5 introduces our technical approach, whilst Section 6 describes algorithms we have developed for computing trust relationships in word of mouth recommendation seeking scenarios, based on the findings of the previous study. Section 7 gives an overview of how these metrics are being used in applications that support information seeking using trust relationships in social networks. Section 8 concludes the paper with an outline of future work.

2 Related Work

The work of Granovetter [1] highlighted how social networks can serve as a source of new information to which an individual may not otherwise have access. In the context of job hunting, he found that weak, rather than strong, social ties are particularly useful, in that they are sufficiently well connected outside of the individual's immediate network (i.e. a sufficient proportion of acquaintances were not shared) as to provide valuable access to otherwise unavailable information about job opportunities.

In addition to this role of information source, our social networks can also serve as a filter, helping us identify the most relevant or appropriate information. At least two factors underpin this: firstly, the principle of homophily [5] states that we are likely to have more in common with members of our social networks than with other members of the population, and more likely to like what they like; secondly, we are better able to judge the appropriateness and trustworthiness (as information sources) of people we know, as we have greater background knowledge of their competence and trustworthiness in a particular domain.

These processes may be assisted by Web technologies in a number of ways. *Collaborative filtering* [6] recommender systems such as GroupLens [7] have typically sought to assist in information filtering by identifying others that share our preferences for newsgroup postings or some other type of item (such as items in an e-commerce site). Variations such as *Amazon* recommendations [8] perform a similar function but instead correlate item rather than people profiles. In the person-to-person approaches, collaborative filtering creates for each of us a social network of unknown others who nevertheless have shared tastes, and through whose preferences information can be filtered on our behalf. Whilst this can be of great value in informing decision-making, it does not allow us to use our own knowledge in

assessing the relevance or trustworthiness of a source, and does not address situations where we require recommendations from domain experts, irrespective of their likeness to ourselves.

So how do people determine the trustworthiness, as information sources, of the people in their social networks? Various studies of information seeking in workplace settings [2, 3] found that people decide whom to ask for information based on what they know of the person, and how they value their knowledge and skills. Both studies found an effect of perceived source quality in determining the likelihood that an individual asks another for information.

In previous work [9] we extended these findings beyond workplace settings, and refined the notion of source quality or trustworthiness. These findings are summarized in Section 4 below. In this paper we will report on how we are using Semantic Web and Web2.0 data sources and social networks to calculate trust ratings between individuals, and how we are using these ratings to support information seeking from known and trusted sources.

Some existing work has been carried out in this area. For example, Massa and Avesani [10] use trust propagation mechanisms to increase the coverage of recommender systems without sacrificing the quality of recommendations to users. Perhaps the best known work in this area from a specifically Semantic Web perspective is that of Golbeck and colleagues. Golbeck and Mannes [11] use manual trust annotations between people (on a 1 to 10 scale) combined with provenance information about trust ratings and social network connections to infer trust ratings between unknown sources. Whilst this can be of value where insufficient annotations are provided by one's social network, it suffers a number of limitations. Firstly the trust ratings (either manual or computed) are not topic-specific; users are required to make global statements of their trust in another person, without further context being provided. This approach also requires sufficient manual trust annotations to bootstrap the process, without being able to rely on existing sources of information. In contrast, our approach aims to compute *person-person* and *person-topic* trust ratings according to a richer model of trust in word of mouth recommendation seeking, and based on existing data sources available in the Web.

3 Our Approach: Trusted Recommendations from a Social Network

We are investigating the use of social networks to provide relevant information and recommendations. In contrast to existing work, our approach aims to identify trusted sources from among known members of one's social network. This follows the principle that knowing the right person to ask is often the greatest challenge in seeking information or recommendations.

This *known person, source-centric* (rather than item-centric) approach has a number of advantages. It allows the user to employ existing knowledge of their social network to assess the quality and impartiality of recommendation sources, and follow-up enquiries with the source as they see fit. Therefore, in contrast to collaborative filtering our approach is less vulnerable to spamming, for the simple

reason that each user's exposure to the recommendations of others is limited in the first instance to those people they know. We proceed on the assumption that most users are unlikely to know others who wish to manipulate search indices on an ongoing, systematic basis. A recent investigation [12] (albeit journalistic, rather than scientific) demonstrated how easily ratings on travel review and recommendation sites such as *TripAdvisor*² can be skewed by those with a vested interest in promoting a particular establishment. Personally knowing those providing a review or recommendation acts as a safeguard against this form of manipulation.

Secondly, our source-centric approach does not assume completeness of the information in the system. For example, for a conventional recommender system to be able to recommend a hotel in Madrid to User A, some record of a hotel in Madrid must exist in the system. In contrast, whilst our approach can identify specific instances of recommended hotels in Madrid, simply identifying those known people with some knowledge of Madrid is sufficient to begin answering the user's information needs, without requiring substantial amounts of information. This is analogous to simply asking "who do I know that knows anything about Madrid?", and is in contrast to conventional collaborative filtering approaches, that whilst they may list "people like you", they are generally aimed towards informing the user that "people like you also liked X". In this sense they are item- rather than source-centric.

Thirdly, Linden, Smith, and York [8] outline limitations of traditional collaborative filtering that stem from its computational expense over large datasets. Computing the *co-preference*³ between all users of a system has been found not to scale where large numbers of users are concerned. By constraining recommendations to those coming from members of a user's social network, we reduce the number of co-preference relationships that must be computed in the system. We anticipate that such an architecture will allow the system to scale more readily.

Lastly, by using Semantic Web technologies we are able to exploit and integrate data from many different sources in computing trust relationships. Our approach uses FOAF-based definitions of users' social networks [13], allowing "friend lists" built up across different services to be reused. *Revyu* provides data about reviews and ratings in crawlable RDF and via a SPARQL endpoint. This brings practical benefits during development (such as query flexibility, ability to reuse common libraries) compared to the more restrictive data access allowed by *del.icio.us*. Crucially however, by being Semantic Web-aware, our approach allows for the generation or refinement of trust ratings based on additional Semantic Web data sources as they become available. This issue is discussed in Section 8.

4 Previous Findings: Trust in Recommendation Seeking

In a previous paper [9] we presented the results of an empirical study examining how people select recommendation sources from among their social networks, and the factors that influence these decisions. Participants were presented with four recommendation seeking scenarios, asked to explain from whom they would seek

² <http://www.tripadvisor.com/>

³ The degree of preference two individuals share for an item

recommendations in each scenario, and to explain their reasons for these choices. Analysis of participants' responses identified five factors underlying the trust or confidence participants had in recommendations from specific sources: the *expertise*⁴, *experience*⁵, and *impartiality*⁶ of the source with regard to the topic of the recommendation seeking, the *affinity*⁷ between the source and recommendation seeker, and the *track record*⁸ of previous recommendations from the source.

These trust factors varied in their frequency of occurrence in participants' explanations for choosing a particular source. *Expertise*, *experience*, and *affinity* occurred most frequently, with relatively low occurrences of the *impartiality* and *track record* factors. Furthermore, the emphasis given to each of these factors was found to vary according to the characteristics of the recommendation seeking task.

Results suggested that the *criticality* of the task and the *subjectivity* of possible solutions were of primary importance in determining which trust factors were emphasised. In scenarios seen by participants as more critical, greater emphasis was placed on the recommendation source having relevant *expertise*. In contrast, in scenarios in which potential solutions were seen as more subjective, participants placed greater evidence on sources with which they shared a strong affinity.

A major shortcoming of the work of Golbeck and Mannes [11] is that trust relationships are represented as global traits between users, rather than being topical or domain-specific. A foundation for our work is the principle that trust can be topical, in that one person may be highly trusted for recommendations in one domain but trusted very little in others. For example, one may trust a friend who works in banking to give sound financial advice, but never trust her film recommendations. The findings of our previous study support our assertion of trust topicality, and suggest that any robust model of trust in word of mouth recommendation must take this into account.

It is worth noting that whilst the factors expertise, experience, and impartiality were clearly domain specific and therefore topical in nature, the study did not give a strong indication of affinity as a topical factor, but rather as a global construct. The range of responses that informed the *affinity* factor suggests that it represents more than simply shared tastes, encompassing instead similar outlooks on life, values, and expectations: "I would ask X, because we see the world in the same way".

5 Computing Knowledge and Trust Relationships

Based on the trust factors identified in this previous study, we have developed algorithms for computing *people-people* and *people-topic* trust metrics that signify

⁴ The source has relevant expertise, which may be formally validated through qualifications or acquired over time

⁵ The source has experience of solving similar scenarios, but without extensive expertise

⁶ The source does not have vested interests in a particular resolution to the scenario

⁷ The source has characteristics in common with the recommendation seeker such as shared tastes, standards, viewpoints, interests, or expectations

⁸ The source has previously provided successful recommendations to the recommendation seeker

respectively the affinity-based trust relationship between two individuals, and the expertise- and experience-based trustworthiness of an individual with regards to a topic. The metrics generated by these algorithms provide the foundations on which our system is built. An overview of the system is provided in Section 7.

We argue that auto-generating trust metrics from existing background data sources is crucial, for a number of reasons. Firstly, such an approach can help overcome the bootstrapping/cold-start problem, whereby a system is only useful to the user once they have provided a certain amount of data specifically to that system. We are exploiting a range of existing and widely used Web2.0 data sources, such as del.icio.us and Flickr, in the generation of our *experience* trust metrics. Initial weak metrics generated from these sources are then enhanced based on richer data from our *Revyu* Semantic Web reviewing and rating site. The integration of further sources into the trust metric generation process is technically feasible and highly desirable. Secondly, reuse of existing sources lessens the burden on the user, as they need not provide new data about their preferences to our system. Instead they can immediately reap the benefits of data they have provided in one system (such as bookmarks in del.icio.us, or reviews in *Revyu*), in the form of enhanced search results and personalization in our system.

Lastly, one additional mechanism for determining the trustworthiness of people's recommendations in a domain would be to ask them to rate their knowledge or expertise in a number of domains. However, such an approach would require a comprehensive yet manageable list of topics or domains, which by definition scales poorly to the full range of topics on which users might require recommendations. By reusing data from external sources that are themselves unconstrained in their coverage of topics (as users can use any tags they wish), we are not constraining the domains or topics in which trust metrics can be calculated.

In computing trust metrics for use within our system, we have given priority to the three trust factors arising most frequently in our previous study: *expertise*, *experience*, and *affinity*. Developing algorithms that directly represent the trust factors has not been possible in all cases. In particular, computing an expertise score in any one domain is problematic, as appropriate sources of background knowledge that indicate expertise are not widely available on the Web, are widely dispersed by topic, and are not generally available in structured, machine-readable form. For example, one's family doctor may have expertise in general healthcare. However, evidence of this in the form of a machine-readable certificate of qualification and competence from a recognised medical authority is not available on the Web. Consequently we have developed a metric (called *credibility*) that serves as a proxy for expertise. An individual is deemed credible with respect to a particular topic if their ratings of items related to that topic correlate highly with those of the community as a whole.

Similarly, large volumes of data are available on the Web that may indicate an individual's experience with regard to a particular topic. However, automatically validating with any degree of confidence that this is the case may not be feasible. Therefore a proxy metric (*usage*) has been developed that suggests an individual has experience in a particular topic. Comparing ratings between individuals allows us to compute affinity metrics with some degree of confidence, without resorting to proxy measures.

6 Algorithms for Generating Trust Metrics

The algorithms used to compute trust metrics in our system are detailed below. The algorithms rely primarily on data from Revyu, however *usage (experience)* metrics are also computed based on del.icio.us tagging data. Tags used in Revyu and del.icio.us seed the list of topics for which individuals may have usage or credibility scores. In Section 8 we discuss further potential Semantic Web data sources on which to base trust calculations.

6.1 Credibility (Expertise) Algorithm

```
for each tag in Revyu
  get all items tagged with that tag, by anyone
  for each item
    find the mean item rating
    for each review of the item
      subtract rating from mean rating to
      give a rating distance
      adjust sign of the rating distance to
      ensure it is positive
      divide rating distance by highest
      possible rating minus 1 to give
      normalized rating distance
      subtract normalized rating distance
      from 1 to give credibility score for
      that review in the range 0-1
      sum each reviewer's credibility
      scores for the current tag to give a
      credibility total for this tag
  for each reviewer with a credibility total for this tag
    divide the credibility total by the number of
    reviews from which it is gained, giving a
    reviewer's credibility score for that tag, in the
    range 0-1
```

Fig. 1. Credibility (Expertise) algorithm in pseudo-code

At present the algorithm does not take into account tags for which only one item exists, or tags for which multiple items exist but where all have only been reviewed by the same person. This can lead to the situation where an individual is assigned a

credibility rating of 1 for a particular topic, by virtue of being the only reviewer of things tagged with that tag. It could be argued that within the scope of the knowledge currently held within the system, this person is justifiably credible and expert on the topic, as no contradictory information exists. However, we do not accept this argument, and anticipate some negative effects of this artifact when we evaluate the algorithms. Methods for mediating this effect are being sought in ongoing research.

6.2 Usage (Experience) Algorithm

This algorithm calculates the prevalence of an individual in the reviews of items that have been tagged with a particular tag, thereby providing a relative measure of their experience with the topic.

```
-----  
for each tag in Revyu  
    count how many times each reviewer has reviewed an  
    item tagged with that tag (by anyone); this gives a  
    reviewer's tag count  
    find the highest of these tag counts  
    divide each reviewer's tag count by the highest tag  
    count to give a usage score in the range 0-1  
-----
```

Fig. 2. Usage (Experience) algorithm in pseudo-code

Catching all people who have reviewed something that has ever been tagged with the target tag helps ensure that people are credited with experience in a relevant domain, even if they haven't used a particular keyword tag themselves. This helps ensure a broader spread of experience scores across related topics.

One consequence of this algorithm is that the individual with the highest *tag count* will be assigned a usage (experience) score of 1 for that topic, by virtue of having reviewed the greatest number of things tagged with a particular tag, and irrespective of the overall number of reviews of items tagged with that tag. Following evaluation we may modify this algorithm to ensure no scores of 1 can be assigned, and also to adjust scores relative to the total number of reviews.

6.3 Affinity Algorithm

The following algorithm computes an affinity score between an individual and another person they know, based on analysis of their reviews in Revyu. In addition to Revyu review data, the algorithm must be seeded with some basic details of the known person. This is supplied to the algorithm in the form of a FOAF description of the user's social network.

```

get all reviews by the user (User A)
get all reviews by the known person (User B)
count the number of items that both users have reviewed
divide this by the highest number of total reviews by
either user, to give an item overlap ratio in the range
0-1

where both users have reviewed the same item
    subtract the rating of User B from that of
    User A, to give a rating distance

    adjust the sign of the rating distance to
    ensure it is positive

    divide rating distance by highest possible
    rating minus 1, to give a normalized rating
    distance in the range 0-1

    subtract the normalized rating distance from 1
    to give a rating overlap for that review

    sum all item-level rating overlaps between
    users A and B, then divide by the number of
    items that both users have reviewed, to give a
    mean rating overlap

combine the item overlap ratio and mean rating overlap
to produce a measure of the affinity between User A and
User B

```

Fig. 3. Affinity algorithm in pseudo-code

At present several aspects of the affinity computation process are subject to variation pending the outcome of evaluations into the effectiveness of the algorithms. Firstly, the relative importance of *item overlap ratio* and *mean rating overlap* in computing affinity is not fully clear, and may vary according to the item overlap ratio. For example, a high mean rating overlap based on few overlapping items may be of less value as a measure of affinity than a slightly lower mean rating overlap based on a large number of overlapping items. The most reliable means for combining these measures is an ongoing question for our research. One option may be to base affinity scores purely on *mean rating overlap*, weighted according to the number of overlapping items. An alternative may be to introduce confidence measures whereby affinity scores are based solely on mean rating overlap, but the confidence of this measure is expressed based on the item overlap ratio.

6.4 Generating Usage (Experience) Scores from del.icio.us Data

In order to increase the range of topics for which users in the system have usage/experience scores, we have extended the *usage* (experience) algorithm to take into account users' tags on *del.icio.us*. Where a user of the system has a *del.icio.us* account, their most used tags are retrieved. For each tag that has received a certain amount of usage (above an arbitrary threshold), the user is recorded as having some *experience* of that topic. A standard nominal experience score (currently 0.1) is assigned irrespective of the frequency of usage of the tag above the threshold, in recognition that tag usage is not necessarily strongly correlated with real experience of the topic. For example, in the course of researching possible holiday destinations a user may bookmark many resources using the tag *hawaii*, but eventually choose Mexico instead for their holiday. In contrast, where a user has reviewed an item we can be reasonably confident that they have some experience of the topics denoted by that item's tags.

Where a user has an existing experience score for a particular topic that exceeds the nominal score derived from their *del.icio.us* tags, the existing score stands unchanged. Where they have an existing score lower than the nominal score, this is increased in line with the nominal score for *del.icio.us*-derived experience. No attempt is made to supplement Revyu-derived *credibility* and *affinity* metrics based on *del.icio.us* data, as bookmarks do not carry ratings, endorsements, or other value judgments from which these may be derived.

6.5 Representing Computed Trust Relationships

Once computed, trust relationships based on these metrics are stored in a triplestore, according to a simple ontology that models the relationships between people and topics identified in our earlier study. This triplestore provides the data for the application outlined below. Trust relationships will also be republished on the Web for potential reuse in other applications.

7 Supporting Information Seeking with Trusted Social Networks

Using trust relationship data computed according to the algorithms detailed above, we are currently completing the implementation of a system that enables people to locate and explore trusted information sources within their social networks, and access items rated highly by these sources. An example of output from the system is shown in Figure 4 below.

Explore film

film Experts in Your Network

1 enrico (Revyu)

[the consequences of love](#)

1 mgaved (Revyu)

[Perfume: The Story of a Murderer](#)

1 slowman (Revyu)

[Pan's Labyrinth \(El Laberinto Del Fauno\)](#)

1 Fin (Revyu)

[Perfume:The Story of a Murderer](#)

0.96 martinp (Revyu)

[The Departed](#)

[The Prestige](#)

[Babel](#)

[Idiocracy](#)

0.95 Paddy (Revyu)

[casino royale](#)

Fig. 4. System output showing *film* experts in the first author's social network, ranked according to *expertise*

As discussed above, the role of trust in information seeking is not constant, but varied and situational, depending on characteristics of the task such as its *criticality* and *subjectivity*. Consequently, in our approach the relative importance of *topic expertise* and *person to person affinity* in ranking of potential information sources is varied according to the *criticality* and *subjectivity* of the information seeking task. We intend to carry out user evaluations to assess the relative merits of different mechanisms for representing *criticality* and *subjectivity* in the system. Current approaches being considered include allowing the user to select *criticality* and *subjectivity* measures in the interface, and pre-categorizing the domains of queries according to their *criticality* and *subjectivity* profiles.

8 Conclusions and Future Work

In this paper we have presented our approach to generating trust profiles for members of a user's social network, in the context of word of mouth recommendation seeking. This approach is based on algorithms for computing *person-topic* (expertise, experience) and *person-person* (affinity) trust metrics, that have been developed based on previous research. By utilizing people's social networks, and employing a rich model of trust in recommendation seeking, our approach overcomes the limitations of previous work in the field.

In addition to completing implementation of the system outlined above, a number of outstanding issues remain which are the subject of ongoing research. Firstly we are investigating the integration of additional sources of data. The contents of users' FOAF files, when combined with other Semantic Web datasets, provide a potentially rich source of information about users' experience of particular topics. For example, where a user states in their FOAF file that they are `based_near` a particular location, we can assume they have some experience of this location, and consequently increase their experience rating for this topic. Use of the Geonames service⁹ may allow us to locate other nearby locations, and assume the user also has some (although likely less) experience of these.

Amongst Web2.0 data sources, *Flickr*¹⁰ in particular may provide a good basis for assessing people experience of particular locations or activities, as photos are likely to be tagged with a location name. In contrast however, it may also lead to significant noise in the system where people have tagged items using words that whilst representing some aspects of the contents of the picture, do not indicate particular experience of a topic. Whilst sources of reviews such as Amazon and Yahoo Reviews are potentially rich in terms of quantity of reviews, they do not provide information from known sources, as reviewers are rarely reliably identifiable.

Regarding the trust relationship algorithms, we aim to investigate how trust relationships may decay over time, and how any rate of decay may vary across different domains. For example, the trustworthiness of a person as a source of knowledge on ancient history may decay very slowly, whereas trust in another individual as a source of restaurant recommendations in London may quickly decay if it isn't regularly updated. Representing these issues in our algorithms is an area of future investigations.

Lastly we aim to use patterns in tag co-occurrence to disambiguate topics, and also as a means to propagate trust scores in one topic to others that are related. Throughout these processes we will continue to evaluate the techniques we develop to ensure that they reliably address user needs.

⁹ <http://www.geonames.org/>

¹⁰ <http://flickr.com/>

Acknowledgements

This research was partially supported by the Advanced Knowledge Technologies (AKT) project. AKT is an Interdisciplinary Research Collaboration (IRC), which is sponsored by the UK Engineering and Physical Sciences Research Council under grant number GR/N15764/01. The AKT IRC comprises the Universities of Aberdeen, Edinburgh, Sheffield, Southampton and the Open University.

References

1. Granovetter, M. S.: The Strength of Weak Ties. *The American Journal of Sociology* 78 (1973) 1360-1380
2. Borgatti, S. P., Cross, R.: A Relational View of Information Seeking and Learning in Social Networks. *Management Science* 49 (2003) 432-445
3. O'Reilly, C. A.: Variations in Decision Makers' Use of Information Sources: The Impact of Quality and Accessibility of Information. *The Academy of Management Journal* 25 (1982) 756-771
4. Heath, T., Motta, E.: Reviews and Ratings on the Semantic Web. In: Proc. Poster Track. 5th International Semantic Web Conference (ISWC2006) (2006)
5. McPherson, M., Smith-Lovin, L., Cook, J. M.: Birds of a Feather: Homophily in Social Networks. *Annual Review of Sociology* 27 (2001) 415-44
6. Goldberg, D., Nichols, D., Oki, B. M., Terry, D.: Using Collaborative Filtering to Weave an Information Tapestry. *Communications of the ACM* 35 (1992) 61-70
7. Konstan, J. A., Miller, B. N., Maltz, D., Herlocker, J. L., Gordon, L. R., Riedl, J.: GroupLens: applying collaborative filtering to Usenet news. *Communications of the ACM* 40 (1997) 77-87
8. Linden, G., Smith, B., York, J.: Amazon.com Recommendations: Item-to-Item Collaborative Filtering. *IEEE Data Engineering Bulletin* (2003) 76-80
9. Heath, T., Motta, E., Petre, M.: Person to Person Trust Factors in Word of Mouth Recommendation. In: Proc. CHI2006 Workshop on Reinventing Trust, Collaboration, and Compliance in Social Systems (Reinvent06) (2006)
10. Massa, P., Avesani, P.: Trust-aware Collaborative Filtering for Recommender Systems. In: Proc. International Conference on Cooperative Information Systems (CoopIS) (2004)
11. Golbeck, J., Mannes, A.: Using Trust and Provenance for Content Filtering on the Semantic Web. In: Proc. WWW2006 Workshop on Models of Trust for the Web (2006)
12. Walsh, G., Swinford, S.: Hotel review websites: a five-star scam. (2006) http://travel.timesonline.co.uk/tol/life_and_style/travel/article634136.ece
13. Brickley, D., Miller, L.: FOAF Vocabulary Specification. <http://xmlns.com/foaf/0.1/> (2005)