POLITECNICO DI MILANO
Facoltà di Ingegneria dei Sistemi

POLO REGIONALE DI COMO

Master of Science in
Management, Economics and Industrial Engineering

"Management engineering"

Choosing the best information sources on the web With trusted social network

Supervisor: Prof. Cinzia Cappiello

Master Graduation Thesis by: Amir Salami
Student Id. number 736609

Academic Year 2010/2011

ACKNOWLEDGEMENT

I am heartily thankful to my supervisor, Prof. Cinzia Cappiello, whose encouragement, guidance and support from the initial to the final level enabled me to develop an understanding of the subject.

I would like to take the opportunity to thank those people who spent their time and shared their knowledge for helping me to complete my research with the best possible result.

Amir Salami

CONTENTS

List of Figures :

List of Tables :

Approach

This Research explores an approach to information seeking on the Web, based on the principles of word-of-mouth recommendation in social networks. Word-of-mouth through traditional channels (e.g. voice, face-to-face) provides access to new information that would not otherwise be available to the information seeker, and helps to filter out less relevant items from a broader pool of options (Granovetter, 1973, Kautz, Selman et al., 1997b, 1997a). This research explores these mechanisms in more detail and applies them in the virtual world.

The adopted approach is oriented around first identifying the most appropriate and trusted sources of recommendations and then using the knowledge held by these individuals to assist in the information seeking task. By combining technical systems that harness the knowledge and experience of users' social networks with their own knowledge of members of those networks, the goal is to reduce information overload and provide access to information that is more personally relevant and trustworthy.

This dissertation addresses the following research questions.

Research Questions

The research reported in this dissertation addresses the following principal question:

> 'To what extent can information- and recommendation-seeking within social networks be supported on the Web?'

This question can be broken down into a number of specific research questions:

1. How do people choose information and recommendation sources from among members of their social network?

2. Which factors influence judgements about the relevance and trustworthiness of these information and recommendation sources?

3. How do the characteristics of the task being performed affect these judgements?

4. To what extent can general principles derived from answers to the previous questions be operationalised as computational algorithms that replicate the process of seeking information and recommendations through social networks?

5. How feasible is the implementation of user-oriented systems that exploit such algorithms?

6. If such systems can be implemented, how do they perform relative to human performance of equivalent tasks?

Definition of Terms

Specialist terms, or those whose meaning may be open to interpretation, will mostly be defined in the body of the dissertation as required. However, for the sake of clarity a number of terms will be defined at this stage.

Information seeking is "a process in which humans purposefully engage in order to change their state of knowledge" (Marchionini, 1995) . This dissertation does not treat *information-seeking* and *recommendation-seeking* as distinct processes but as variations on the same theme. Recommendation-seeking is seen simply as a form of information-seeking in which the information seeker tries to obtain opinions or value judgements from trusted (or otherwise favoured) sources, as a means to distinguish between potentially relevant items and thereby reduce the search space.

In the context of this research, an individual's *social network* is defined in the first instance as the people they know personally and with whom they identify in some way, possibly through shared characteristics, socio-cultural identity, or other group membership. This may encompass family members, friends, colleagues, neighbours or other acquaintances. In the second instance the definition of the network may be extended to take in so-called 'friends-of-friends' (those people in the networks of members of one's own network), or even 'friends-of-friends-of-friends'.

This research views both social networks and *relevance* as being primarily a construction of the individual. On this basis, no assumptions are made from the outset about how the nature or origin of social relations may influence the information-seeking process. In particular, no assumptions are made about how particular classes of network members might contribute to the information-seeking process, as these issues will be examined by research presented later in this dissertation.

Literature Review

This Figure reviews literature in the fields of information-seeking, relevance, personalised search, recommender systems, social navigation, social search and trust. The structure of the review is outlined in Figure below, which also illustrates interrelations between these fields and how the limitations of work in one field motivate related work in another.

Figure 1: Conceptual structure of the literature review, showing limitations of and links between approaches

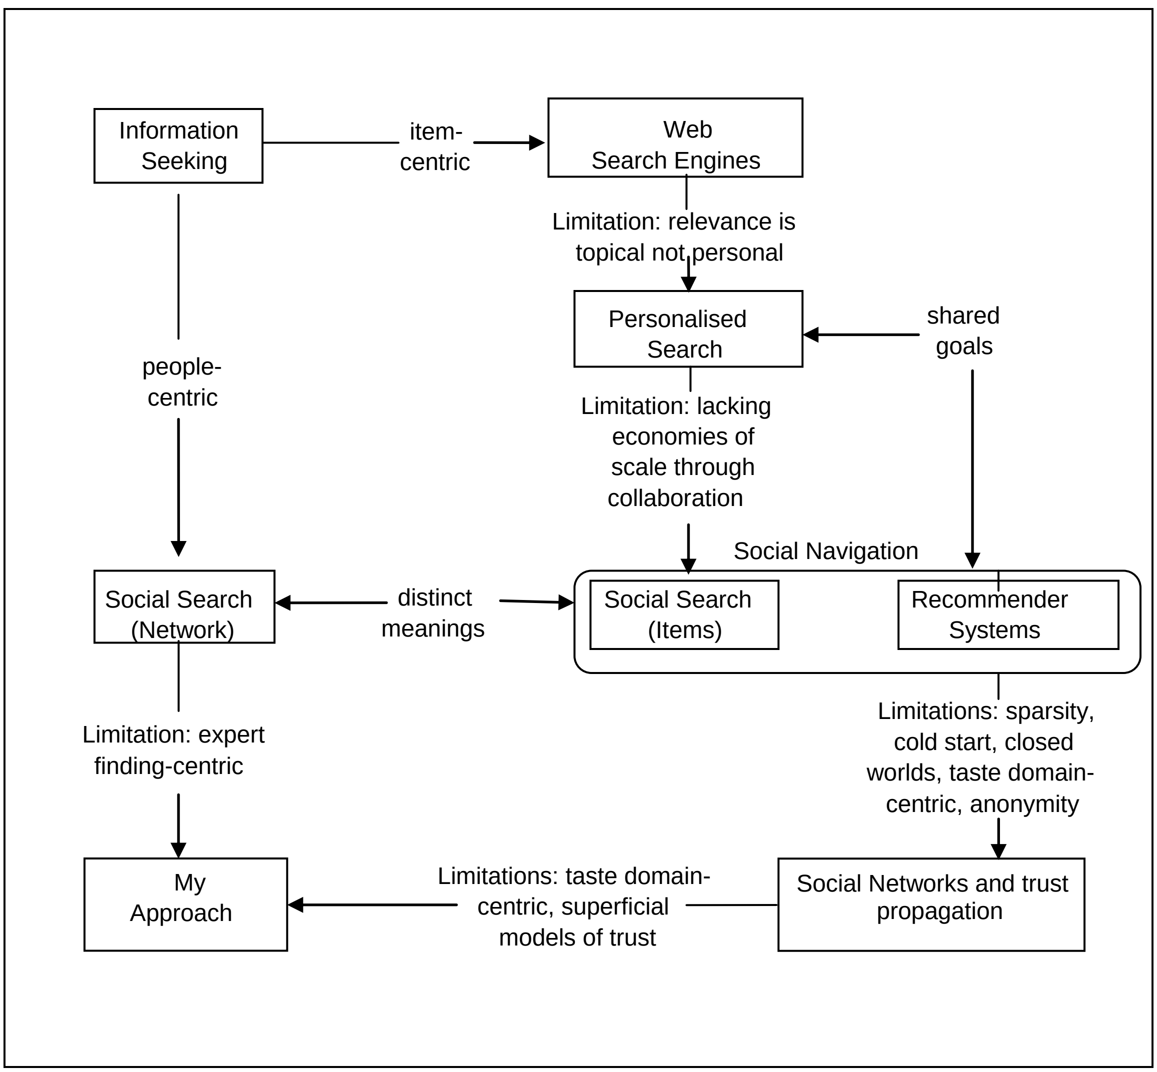

On the hypertext web, any person is allowed to make any statement with no requirements about its accuracy or truthfulness. When reading a web page, humans make many judgments based on the appearance of the page and the source of the information. Although someone could lie about their sources, it is relatively easy to generate at least some information about the source. On the Semantic Web, content is a series of statements that cannot be judged by appearance or professionalism. Since the underlying philosophy of the Semantic Web is to allow a computer to take distributed statements about the same resource and aggregate them, the source of information becomes removed one step from the presentation. The word “Trust” has come to have several definitions on the Semantic Web. Much research has focused on authentication of resources, including work on digital signatures and public keys. This provides confidence in the source or author of a statement, which is very important,

but trust in this sense ignores the credibility issue. Confirming the source of a statement does not have any explicit implication about the quality of the statement.

Reputation is more a social notion of trust. In our lives, we each maintain a set of reputations for people we know. When we need to work with a new, unknown person, we can ask people with whom we already have relationships for information about that person. Based on the information we gather, we form an opinion about the reputation of the new person. This system works well, even though there are a lot of people in the world, because communities tend to be highly interconnected, and the number of steps between any two people tends to be rather small. This is known as the Small World effect, and it has been shown to be true for a variety of social and web-based systems.I will talk about reputation method later after I explain and illustrate some important metrics.

Studying and analyzing Web 2.0 media, such as social networks, blogs, forums, wikis etc. has gained a big momentum, resulting in an increase of research in the related fields. Among the several facets of these social media, trust, influence, and ranking are receiving a lot of attention.

Several researchers have focused on trust prediction and propagation. Most researchers propose classification models, such as SVM-based methods to assign trust class labels using features such as user profile, user interactions, product reviews, and trust relations. A different approach is that of [Lim et], that employs the ”Trust Antecedent” framework proposed in management science and introduce quantitative - instead of qualitative - features, such as ability, benevolence and integrity in the prediction process.

A slightly different line of work focuses on how trust is propagated in a network of people . Whereas in our work we introduce the notion of trust in a social network, we assume that the trust between a pair of users is already known, either explicitly or implicitly. Moreover, trust propagation is thought to be covered by the more general notion of "influence" within such a network.

One common approach is to model the identification of influencers as a combinatorial optimization problem:

given a fixed number of nodes that can be initially activated or infected, find the set of nodes with maximum influence over the entire network - the one that generates the largest cascade of adoptions . Several works build on this

Information Cascade (IC) notion proposing various machine learning algorithms . Even though such approaches have been shown to improve over traditional social Network analysis metrics, they are solely based on the link structure of social networks, and do not take into consideration other important parameters, such as activity, rate of updates, and trust among users. In the same vein, researchers Have investigated the identification of likely influential users through link analysis techniques, as well as user activity-related parameters in order to identify influential users in blogs and social networks . Ranking on the web is primarily based on the analysis of the web graph as it is formulated by hyperlinks. In the case of blogs, several ranking algorithms have been suggested that exploit explicit (Eigen Rumor algorithm) and/or implicit (Blog Rank) hyperlinks between blogs. All these algorithms formulate a graph of blogs, based on hyperlinks and then apply Page Rank or a variation of it in order to provide an overall ranking of blogs. However, all these algorithms provide a static measure of blog importance that does not reflect the temporal aspects accompanying the evolution of the blogosphere. Recently, some effort has been done to also incorporate the content in the ranking process, when ranking twitterers (Twitter Rank).

To the best of our knowledge, this is the first extensive study of the effect of both overall "influence", as expressed by the analysis of the whole social graph, as well as by personalized aspects of "influence" such as trust, in ranking and recommending other users or content.

Social network analysis is the study of social entities (actors) and their interactions and relationships. The interaction and relationships are represented as a graph, where each node represents an actor (user), and the edge between two nodes represents their relationship. In my work, I employ social network analysis metrics such as centrality and rank prestige, in order to identify the “influential” actors in a social network, in terms of their position in the graph and their connections/interactions with other users . In addition to these global metrics, influence in a local scale is important for all actors. In this context, actors collaborate with the actors they trust and are influenced by their opinions. Moreover, trust and influence are reinforced for certain actors in the circle of trust and decrease for others.

In order to model the dynamics of trust and influence in the ”neighborhood” of a user, I employ my collaborative local scoring mechanism. In what follows, I provide a brief overview of the aforementioned metrics :

Social Network Analysis Metrics

The following represents the potential on-page factors:

1. Description Meta tag (A special HTML tag that provides information about a Webpage).
2. A website’s URL.
3. The title of a website.
4. Keyword Meta tags.
5. Density of a given keyword on a document.
6. Proximity of keywords defines how close keywords are in relation to each other.

7. Prominence of keywords defines where the keywords are on the HTML page. For example, a keyword with high prominence would be at the top of an HTML document.

8. Keywords using HTML bold and/or italics.

9. Overall size of a page.

10. Total number of pages within the website.

11. Number of outbound links.

12. Use of quotes text keywords.

13. Using underscores on text keywords.

14. The uniqueness of the content on your page relative to the other content on the web.

15. Content “freshness.” When was content last updated? Has it changed since the last time it was crawled?

16. Spelling and grammar.

A one-dimensional search algorithm might calculate the density of a keyword on a page, and use that keyword density as a measure of relevance. This type of search can quickly lead to text manipulation if the web authors are aware that they need simply to change the keyword density of their web document to indicate to a search engine what their document is about. Using only the on-page factors, web spam will be difficult to stop because the website optimizer can still control the parameters the search algorithm is using to determine ranking.

To this extent, off-page factors were introduced. These factors are difficult for the

Web page optimizer to control. Off-page metrics are more desirable in any ranking algorithm because they allow the search algorithm to determine which pages appear in search queries, rather than by webpage optimizers manipulating WebPages. The following represent the potential off-page factors:

1. Number of websites linking back to a website.
2. The page rank of a website.
3. The number and quality of directories a page is listed in. For example DMOZ or Yahoo.
4. How long a URL has been registered.
5. When a registered domain name will expire.
6. When the search engine spider last crawled the URL.
7. How many pages of the website were crawled (crawl depth).
8. How fast the pages can be crawled (crawl rate).

One reason for moving to metrics like those is that they are less obvious to the website optimizer. Major search engines like Google and Yahoo! have a majority of the world's search queries at their disposal. These search engines also have access to statistical data for how authoritative WebPages have evolved over time. Armed with this type of information, search engines can develop algorithms that can detect unnatural webpage behavior.

I love social network because it can positively affect almost all aspects of your business both directly and indirectly. Directly, it generates increased brand awareness, traffic, leads, sales, subscribers, and more followers on social media sites. Indirectly, it effects search rankings as a successful campaign attracts backlinks and mentions in content which leads to higher search engine rankings. Social network is so popular because it truly is that powerful, one campaign can affect your entire business.

Below is a list of metrics I use to track and prove ROI when analyzing my own site :

- **Traffic -** Inside the traffic umbrella there are multiple metrics equally important when analyzing success. Page Views to the campaign, for one. Pretty obvious, but it all starts with how much traffic the specific page generates throughout the “Viral” life span as well as each month afterward.
- **Unique Visitors -** Equally as obvious is the correlation between total visits and unique visits. This enables you to see what kind of new reach you gained as a result of the campaign.
- **Referring Urls -** This is by far one of the most important to track and watch. Knowing which social sites send you the most traffic and which blogs and news hubs are picking you up is very valuable. This helps you manage the entire conversation, the influence these sites have in your niche and which of the sites sending you traffic is sending visitors that are engaging in your content. Understanding the referring sites that send you the best traffic will help you in your future campaigns.
- **Conversions -** Not all campaigns generate direct sales/leads from the initial burst in traffic and referrals, but if your content does, make sure you are tracking where they came from. I have seen some sites which generally receive more leads from StumbleUpon, whereas other sites got leads from

Digg, Twitter or a mention on Mashable. Also, know that despite what critics say, it is possible to get direct sales, leads etc… from social media and viral campaigns. The more you understand this, the easier it is to improve your next viral campaign. When companies start out in viral marketing, a lot of it can be shoot first and aim later, but as you better understand your conversions you can become a "Black Ops" sniper that has the target in the cross hairs before you ever pull the trigger.

- **Micro-conversions -** Social Media can build upon small success to lead up to a reach of massive proportions. You may not get leads after your first influx in traffic, you may not sell crates full of product, but you can build a large audience and it is essential that you are tracking this. Set up outgoing click tracking with your analytics product to track RSS Subscribers, Twitter Followers, Facebook Fans and so forth – this will allow you to tie referral data to those clicks. Seeing the referrals that are generating those outbound clicks is essential, as it helps you understand again which funnel of traffic does best on your site. Sphinn might only bring you a couple hundred unique visitors, but out of those you might get 100 new Twitter followers. Diggmight send you 50,000 visits and generate nothing when it comes to new RSS or Facebook Fans. Again, the more you know about the hordes of traffic flooding to your site, the better feel you can get as to which communities might help you build your influence. I have had the most experience setting up outgoing link traffic with Google Analytics, Omniture and BLVD Status, it might be possible with other providers.

- **Loyalty/Avg Time on site -** You might know your bounce rate and average time on site for any specific date range, but are you narrowing down and looking at which social traffic tends to bring visitors that spend more time on your site? Do you know that visitors from Twitter are spending 43% more time on your site than Digg referrals, and Facebook visitors are spending more time than StumbleUpon? Understanding which referrers from social sites create more page views per visit and longer time on sites is invaluable.
- **Branded keywords and type-in traffic -** A big part of viral success is brand awareness, and a good way to track overall brand reach is to watch for surges in branded keyword traffic, either around your company as a whole, or if a specific product or aspect of the business was featured in the viral campaign. You should also watch whether the direct type-in traffic goes up during and after the viral campaign. If there is an increase in those typing in your brand or direct URL and coming to your site that is true word of mouth. (We noticed this with the launch of Social Media for Firefox a few years back, with huge surges in people typing in *www.97thfloor.com* or search things like "*97th Floor*," "97*th Floor Plugin*," "*Social Media for Firefox*," "*Chris Bennett Firefox*" there were literally hundreds of keyword phrases that brought tens of thousands of visitors all centered around the brand or the plugin.)
- **Short Urls -** Url shorteners with analytics tied to them is a great way to not only better track viral success throughout social sites like Twitter and other places your normal analytics can't reach, but it is also a good way to track

your influence and what kind of traffic your profile generates when you Tweet out a link.

- **Backlinks -** This is where the indirect portion of social media success comes in. It is common for a social media campaign to generate thousands of links, I have seen articles and graphics generate hundreds and often 1,000-2,000 links or so, and I have seen really good campaigns and pieces generate over 30,000 links alone. That kind of natural link growth is certain to help your search rankings increase. First, watch your backlink count to that specific url you were pushing from either Google Webmaster Tools or Yahoo Site Explorer or SEOMoz's tools. Pay attention for the big news, blog and hubs that pick your story up. These high profile links are going to help boost your rankings, *but – you also want to participate on those sites and help the success go further.* Stumble those pages, Tweet the articles mentioning you and help those 3rd party mentions get traction, this will throw lighter fluid on an already burning flame.
- **Keyword search -** Because of the viral success and the links that your url will incur, you will gain search rankings. This is specifically where the indirect results of social media come into play. Due to the word of mouth and links, your page ranks higher and for more unique terms – you may not have received direct sales from the initial burst in traffic or referrals, but you will more likely than not experience sales and leads due to the new rankings you are enjoying. Thus, social media has an indirect effect on lead gen and ROI. Segment your analytics to see all the keyword phrases that drive traffic to a specific viral piece and notice the amount of traffic they bring in. Track them all the way to the conversion to understand what kind

of impact they have on your bottom line. Segmenting only the words that point to the specific url you are tracking can be hard in most platforms, it can take a lot of tagging and extra work. Because of this, we created a feature in BLVD Status which lets you “filter by url” and see every word pointing to a specific page, as well as their rankings. We’ve seen social media campaigns get traffic from and rank for over 6,000 unique phrases which generated over 17,000 visits in one months’ time. This was two months after the viral traffic and word of mouth died down.

Social media and viral marketing is exciting and worthwhile, since you can experience positive ROI on your time and effort in many different ways, it can dramatically increase your bottom line whether through direct or indirect results. You just need to make sure you are tracking it from start to finish.

Centrality: The three centrality metrics, namely degree, closeness, and betweenness centrality, identify “key” users of the graph, in terms of information dissemination. Let n denote the size of the graph (i.e. the number of actors/ users).

Degree Centrality Gd(i) takes into consideration the node degree d(i) of a user i. The higher the node degree, the more central the user is:

$$Gd(i) = \frac{d(i)}{n-1} \tag{1}$$

Closeness Centrality Gc(i) of a user i signifies how easily this user interacts with all other users j (j 2 [1::n]). Let d(i; j) denote the distance of user i from user j, equal to the number of links in a shortest path. Then, according to closeness centrality, the shorter the distance of the user to all other actors, the more central the user is:

$$Gc(i) = \frac{n-1}{\sum_{j=1}^{n} d\left(i, j\right)} \tag{2}$$

Finally, Betweenness Centrality Gb(i) signifies the importance of user i with regards to the flow of information in the social network. If the user is between two non-adjucent users j and k then i has control over their interactions. If i is on the paths of many such interactions (i.e. between many users), then this is an important user, having a great amount of influence on what happens in the network. Let sp_{jk} be the number of shortest paths between j and k, and $sp_{jk}(i)$ (j$\neq$ i and k$\neq$ i) be the number of shortest paths that pass i. Betweenness centrality of a user i is defined as follows:

$$Gb(i) = \sum_{j<k} \frac{sp_{jk}(i)}{sp_{jk}} \tag{3}$$

Hubs and Authorities: Both terms were introduced as part of the well-known algorithm HITS . A hub is a node with many out-links and an authority is a node with many in-links. Let E be the set of directed edges (i.e. links) in the graph, then the authority Ga(i) and hub Gh(i) scores are iteratively calculated as follows:

$$Ga(i) = \sum_{(j,i)\in E} Gh(j) \tag{4}$$

$$Gh(i) = \sum_{(i,j)\in E} Ga(j) \tag{5}$$

PageRank: PageRank also identifies "authorities" in a graph. Transferring this notion to the social network paradigm, a user i is considered to be influential if

a) many other users endorse i (for example by "trusting" i, adding i's blog in their blogroll, or becoming i's followers), and

b) these users are in turn influential. The PageRank score Gp(i) of user i is iteratively computed as follows:

$$Gp(i) = (1-d) + d \sum_{(j,i)\in E} \frac{Gp(j)}{o_j} \qquad (6)$$

where Oj denotes the number of out-links of node j and d is the so-called damping factor.

Trust

Trust is a social phenomenon. As such, any artificial model of trust must be based on how trust works between people in society. To this end, we have carried out a survey of the social sciences and identified characteristics of trust that are relevant to our work. We outline them below. First, we must clarify the notion of trust.

Defining Trust

Trust is a complex notion whose study is usually narrowly scoped. This has given rise to an evident lack of coherence among researchers in the definition of trust . For our purposes, however, we find it instructive to use the following definition by Gambetta :

> *… trust (or, symmetrically, distrust) is a particular level of the subjective probability with which an agent will perform a particular action, both before [we] can monitor such action (or independently of his capacity of ever to be able to monitor it) and in a context in which it affects [our] own action.*

Mathematical probability has certain properties that make it unsuitable as a trust metric. For this reason, we will take Gambetta's use of the term 'subjective probability' above only as an indication of the existence of different levels of trust, which are dependent upon the truster.

Typology

Social scientists have collectively identified three types of trust. There is *Interpersonal Trust* which is the trust one agent has in another agent directly. This trust is agent and context specific . For example Alice may trust a specific agent Bob the Mechanic in the specific context of servicing her car but not in the context of babysitting her children.

The second type, *System Trust*, or *Impersonal Trust*, refers to trust that is not based on any property or state of the trustee but rather on the perceived properties or reliance on the system or institution within which that trust exists. The monetary system is one such example.

Finally, *Dispositional Trust*, or sometimes referred to as one's 'basic trust', describes the general trusting attitude of the truster. This is "a sense of basic trust, which is a pervasive attitude toward oneself and the world" .

Therefore, it is independent of any party or context.

Further subtypes of Dispositional Trust are defined by McKnight et al – Type A concerns the truster's belief on others' benevolence and Type B is the disposition that irrespective of the potential trustee's benevolence, a more positive outcome can be persuaded by acting 'as if' we trusted her.

Trust Characteristics

Trust is not an objective property of an agent but a *subjective degree of belief* about agents . The degrees of belief associated with trust range from complete distrust to complete trust. There is also a situation where an agent does not have an opinion of another's trustworthiness, i.e. the agent is said to be *ignorant* of the other agent's trustworthiness. A trusting action is taken despite *uncertainty* of outcome but in anticipation of a positive outcome .

This may draw some to conclude that trust is merely a game of chance, which is untrue. More than being a blind guess, a trusting decision is based on the truster's relevant prior *experiences* and *knowledge* . The experiences and knowledge forms the basis for trust in future familiar situations . In this sense, trust reasoning has an inductive form, rather than deductive. Furthermore, trust is *dynamic* and *non-monotonic* – additional evidence or experience at a later time may increase or decrease our degree of trust in another agent.

It may also seem intuitive to represent degrees of trust as some probability measurement. However, the problem with this is that the probability values will be meaningless unless it is based on well-defined repeatable experiments, which is an impossibility when dealing with most everyday real-life experiences. Another problem is that probability does not take the observers into account, merely their observations. Thus, probability is inherently transitive while trust is not necessarily so . If Alice trusts Bob and Bob trusts Cathy, it does not necessarily follow that Alice must trust Cathy by any degree. A formal argument for the non-transitiveness of trust is given in . Lastly, a trusting action may not follow the rules of rational choice theory . An agent may have reasons beyond the cognitive evaluation of risk and utility – a trust decision may be made “in terms of here and now” instead of pondering on future outcome.

In our current society it is more and more common to interact with strangers, people who are totally unknown to us. This happens for example when receiving an email asking for collaboration or advise from an unknown person, when we rely on reviews written by unknown people on sites such as Amazon.com, and also when browsing random profiles on social networking sites such as Facebook.
com or Linkedin.com. Even more surprising is the fact a huge amount of commercial exchanges happen now between strangers, facilitated by platforms such as Ebay.com. In all systems in which it is possible to interact with unknown people, it is important to have tools able to suggest which other users can be trustworthy enough for engaging with. Trust metrics and reputation systems have precisely this goal and become even more important, for instance, in systems where people are connected in the physical world such as carpooling systems or hospitality exchange networks (i.e. couchsurfing.com), in which users accept to have strangers into their car or their house.
A commonly cited definition of trust was proposed by Diego Gambetta:
“Trust (or, symmetrically, distrust) is a particular level of the subjective probability with which an agent will perform a particular action, both before [we] can monitor each action (or independently of his capacity of ever be able to monitor it) and in a context in which it affects [our] own action” . In all the previous example it is possible to consider the social relationship users can express as a trust statement, an explicit statement stating “I trust this person in this context” (for example as a pleasant guest in a house or as a reliable seller of items).

While research about trust issues spanned disciplines as diverse as economics, psychology, sociology, anthropology and political science for centuries, it is only recently that the widespread availability of modern communication technologies facilitated empirical research on large social networks, since it is now possible to collect real world datasets and analyze them . As a consequence, recently computer scientists and physicists started contributing to this research field as well .

Moreover we all start relying more and more on these social networks, for friendship, buying, working, ... As this field become more and more crucial, in the past few years many trust metrics have been proposed but there is a lack of comparisons and analysis of different trust metrics in the same conditions. As Sierra and Sabater put it in their complete “Review on Computational Trust and Reputation Models” : “Finally, analyzing the models presented in this article we found that there is a complete absence of test-beds and frameworks to evaluate and compare the models under a set of representative and common conditions. This situation is quite confusing, specially for the possible users of these trust and reputation models. It is thus urgent to define a set of test-beds that allow the research community to establish comparisons in a similar way to what happens in other areas (e.g. machine learning)” (emphasis added). Our goal is to fill this void and for this reason we set up Trustlet , a collaborative wiki in which we hope to aggregate researchers interested in trust and reputation and build together a lively test-bed and community for trust metrics evaluation. A project with similar goals is the Agent Reputation and Trust (ART) Testbed. However ART is more focused on evaluating different strategies for interactions in societies in which there is competition and the goal is to perform more successfully than other players, in a specific context. Our take with Trustlet is about evaluating performances of trust metrics in their ability to predict how much a user could trust another user, in every context. For this reason, we want also to support off-line evaluation of different trust metrics on social network datasets. The two testbeds are hence complementary.

In the following pages I will describe Trustlet, the reason behind its creation and its goals, I report the datasets I have collected and released and the trust metrics I have implemented and I present a first empirical evaluation of different trust metrics on the Advogato dataset.

The Advogato network-flow trust metric

Capacity constrained flow network. Capacities of nodes are set as a function of distance from seed.

Figure 2:

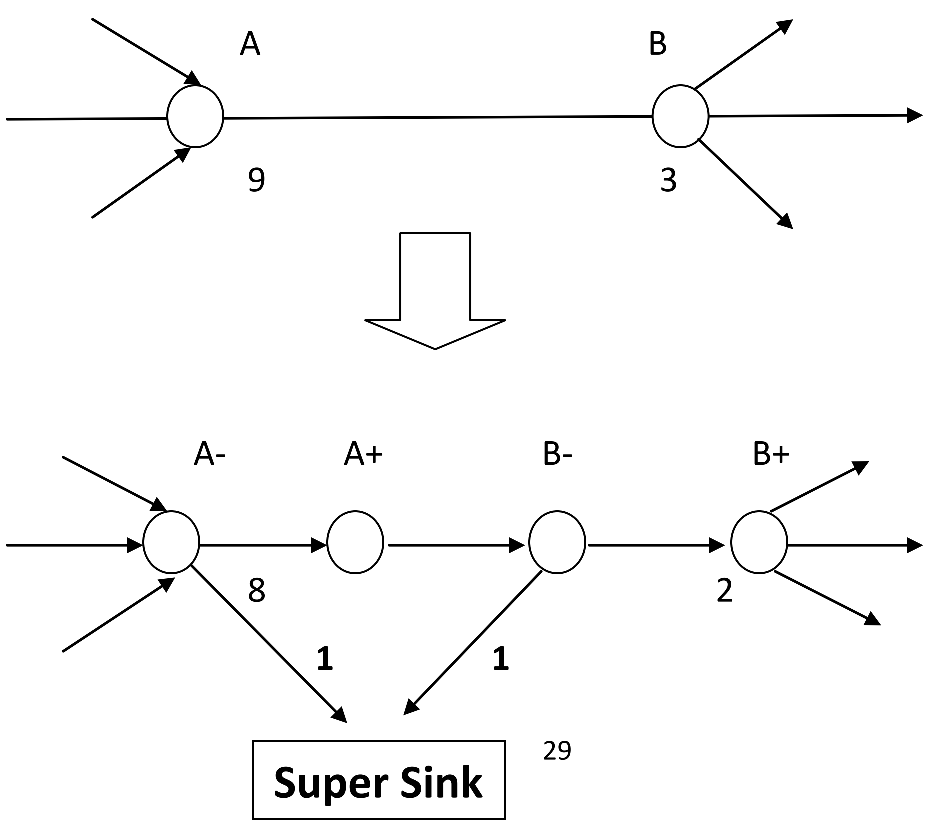

Trust metrics

Trust metrics are a way to measure trust one entity could place in another.
After a transaction user Alice on Ebay can explicitly express her subjective level of trust in user Bob. We model this as a trust statement from Alice to Bob. Trust statements can be weighted, for example on Advogato a user can certify another user as Master, Journeyer, Apprentice or Observer, based on the perceived level of involvement in the free software community. Trust statements are directed and not necessary symmetric: it's possible a user reciprocates with a different trust statement or simply not at all. By aggregating the trust statements expressed by all the members of the community it is possible to build the entire trust network (see for example Figure 1). A trust network is hence a directed, weighted graph. In fact trust can be considered as one of the possible social relationships between humans, and trust networks a subclass of social networks. Trust metrics are tools for predicting the trust a user could have in another user, by analyzing the trust network and assuming that trust can somehow be propagated. One of the assumptions is that people are more likely to trust a friend of a friend than a random stranger.
Trust metrics can either be local or global . A global trust metric is a trust metric where predicted trust values for nodes are not personalized.
On the other hand, with local trust metrics, the trust values a user sees for other users depend on her position in the network. In fact, a local trust metric predicts trust scores that are personalized from the point of view of every single user. For example a local trust metric might predict "Alice should trust Carol as 0.9" and "Bob should trust Carol as 0.1", or more formally trust (A,C)=0.9 and

trust(B,C)=0.1. Instead for global trust metrics, trust (A,B)=reputation(B) for every user A. This global value is sometimes called reputation. Currently most trust metrics used in web communities are global, mainly because they are simpler to understand for the users and faster to run on central servers since they have to be executed just once for the entire community. For example Ebay and Page rank are global. However we think that soon users will start asking for systems that take into account their own peculiar points of view and hence local trust metrics, possibly to be run in a decentralized fashion on their own devices.

While research on trust metrics is quite recent, there have been some proposals for trust metrics, although my goal is not to provide a complete review of trust metrics here.

Ebay web site shows the average of the feedbacks received by a certain user in his profile page. This can be considered as a simple global trust metric, which predicts, as trust of A in B, the average of all the trust statements received by B . In more advanced trust metrics trust can be extended beyond direct connections. The original Advogato trust metric is global, and uses network flow to let trust flow from a "seed" of 4 users, who are declared trustworthy a priori, towards the rest of the network. The network flow is first calculated on the network of trust statements whose value is Master (highest value) to find who classifies as Master. Then the Journeyer edges are added to this network and the network flow is calculated again to find users who classify as Journeyer. Finally the users with Apprentice status are found by calculating the flow on all but the Observer edges. The untrusted Observer status is given if no trust flow reached a node. By replacing the 4 seed users for an individual user A, Advogato can also be used as a local trust metrics predicting trust from the point of view of A.

The problem of ranking of web pages in the results of a search engine query can be regarded under a trust perspective. A link from page A to page B can be seen as a trust statement from A to B. This is the intuition behind the algorithm Page rank powering the search engine Google. Trust is propagated with a mechanism resembling a random walk over the trust network.

Moletrust is a local trust metric. Users are ordered based on their distance from the source user, and only trust edges that go from distance n to distance n+1 are regarded. The trust value of users at distance n only depend on the already calculated trust values at distance n-1. The scores that are lower than a specific threshold value are discarded, and the trust score is the average of the incoming trust statements weighted over the trust scores of the nodes at distance n-1. It is possible to control the locality by setting the trust propagation horizon, i.e. the maximum distance to which trust can be propagated.

Golbeck proposed a metric, TidalTrust , that is similar to Moletrust. It also works in a breadth first search fashion, but the maximum depth depends on the length of the first path found from the source to the destination. Another local trust metric is Ziegler's AppleSeed , based on spreading activation models, a concept from cognitive psychology.

Datasets and trust metrics evaluation

Research on trust metrics started a long time ago, but is somehow still in its infancy. The first trust metric could be even ascribed to the philosopher John Locke who in 1680 wrote: "Probability then being to supply the defect of our knowledge, the grounds of it are these two following: First, the conformity of anything with our own knowledge, observation and experience. Secondly, The testimony of others, vouching their observation and experience. In the testimony of others is to be considered: (1) The number. (2) The integrity. (3) The skill of the witnesses. (4) The design of the author, where it is a testimony out of a book cited. (5) The consistency of the parts and circumstances of the relation. (6) Contrary testimonies”. This quotation can give an idea of how many different models for representing and exploiting trust have been suggested over the centuries. However of course John Locke in 1680 didn’t have the technical means for empirically evaluating his “trust metric”. Even collecting the required data about social relationships and opinions was very hard in old times. The first contributions in analysis real social networks can be tracked down to the foundational work of Jacob Moreno and since then many sociologists, economists and anthropologists have researched on social networks and trust. But the advent of the information age has made it possible to collect, represent, analyze and even build networks way beyond that what is possible with pen and paper. Computer scientists and physicists have become interested in social networks, now that both huge amounts of data have become available and computing power has advanced considerably.

At Trustlet.org we have started a wiki to collect information about research on trust and trust metrics. We hope to attract a community of people with interest in trust metrics. We have chosen to use the Creative Commons Attribution license so that work can easily (and legally) be reused elsewhere. Our effort shares the vision of the Science Commons project which tries to remove unnecessary legal and technical barriers to scientific collaboration and innovation and to foster open access to data.

We believe the lack of generally available datasets is inhibiting scientific work.

It’s harder to test a hypothesis if it has been tested on a dataset that is not easily available. The other alternative is testing the hypothesis on synthesized datasets, which are hardly representative of real-world situations. Prior to the proliferation of digital networks data had to be acquired by running face-to-face surveys, which could take years to collect data of a mere couple of hundreds of nodes. The proliferation and popularity of on-line social networks has facilitated acquiring data, and the implementation of standards like XFN and common APIs like Open Social opens up new possibilities for research . A more widespread availability and controlled release of datasets would surely benefit research and this is one of the goal behind the creation of Trustlet.

Trust network datasets are directed, weighted graphs. Nodes are entities such as users, peers, servers, robots, etc. Directed edges are trust relationships, expressing the subjective level of trust an entity expresses in another entity .

We think it is important that research on trust metrics follows an empirical approach and it should be based on actual real-world data. Our goal with Trustlet is to collect as many datasets as possible in one single place and release them in standard formats under a reasonable license allowing redistribution and, at least, usage in a research context. At present, as part of our effort with Trustlet, we collected and released datasets derived from Advogato, people.

(See http://www.trustlet.org/wiki/Trust network datasets).

squeakfoundation.org, Robots.net and Epinions.com.

Advogato.org is an online community site dedicated to free software development, launched in November 1999. It was created by Raph Levien, who also used Advogato as a research testbed for testing his own attack-resistant trust metric, the Advogato trust metric . On Advogato users can certify each other as several levels: Observer, Apprentice, Journeyer or Master. The Advogato trust metric uses this information in order to assign a global certification level to every user. The goal is to be attack-resistant, i.e. to reduce the impact of attackers . Precise rules for giving out trust statements are specified on the Advogato site. Masters are supposed to be principal authors of an "important" free software project, excellent programmers who work full time on free software, Journeyers contribute significantly, but not necessarily full-time, Apprentices contribute in some way, but are still acquiring the skills needed to make more significant contributions. Observers are users without trust certification, and this is also the default. It is also the level a user certifies another user at to removea previously expressed trust certification.

For the purpose of this paper we consider these certifications as trust statements. T(A,B) denotes the certification expressed by user A about user B and we map the textual labels Observer, Apprentice, Journeyer and Master in the range [0,1] ,respectively in the values 0.4, 0.6, 0.8 and 1.0. This choice is arbitrary and considers all the certifications are positive judgments, except for "observer" which is used for expressing less-than-sufficient levels. For example, we model the fact raph certified federico as Journeyer as T(raph, federico)=0.8.

The Advogato social network has a peculiarly interesting characteristic: it is almost the only example of a real-world, directed, weighted, large social network. However, besides the leading work of Levien reported in his unfinished PhD thesis , I am just aware of another paper using the Advogato dataset which is focused on providing a trust mechanism for mobile devices .

There are other web communities using the same software powering Advogato.org and they have the same trust levels and certifications system:

robots.net, persone.softwarelibero.org, people.squeakfoundation.org, kaitiaki.org.nz.

We collected daily snapshots of all these datasets and made them available on Trustlet but we haven't used them for our analysis in this paper, mainly because they are much smaller than the Advogato dataset. Details about the characteristics of the Advogato trust network dataset are presented in following pages.

The other set of datasets we released is derived from Epinions.com, a website where users can leave reviews about products and maintain a list of users they trust and distrust based on the reviews they wrote .

Initial research outcomes

In the previous sections we highlighted the reasons for creating Trustlet and the way we hope it can develop into a collaborative environment for the research of trust metrics. As a first example of what we hope Trustlet will be able to bring to research on trust metrics, we report our first investigation and empirical findings. We chose to start studying the Advogato social network because of its almost unique characteristic. Trust statements (certifications) are weighted and this makes it a very useful dataset for researching trust metrics: most networks just exhibit a binary relationship (either trust is present or not) and the evaluation on trust metrics performances is less insightful.

The Advogato dataset we analyzed is a directed, weighted graph with 7294 nodes and 52981 trust relations. There are 17489 Master judgments, 21977 for Journeyer, 8817 for Apprentice and 4698 for Observers. The dataset is comprised of 1 large connected component, comprising 70.5% of the nodes, the second largest component contains 7 nodes. The mean in- and out-degree (number of incoming and outgoing edges per user) is 7.26. The mean shortest path length is 3.75. The average cluster coefficient is 0.116. The percentage of trust statements which are reciprocated (when there is a trust statement from A to B, there is also a trust statement from B to A) is 33%.

While a large part of research on social networks focuses on exploring the intrinsic characteristics of the network , on Trustlet we are interested in covering an area that received much less attention, analysis of trust metrics. We have compared several trust metrics through leave-one-out, a common technique in machine learning. The process is as follows: one trust edge (e.g. from node A to node B) is taken out of the graph and then the trust metric is used to predict the trust value A should place in B, i.e. the value on the missing edge.

We repeat this for all edges to obtain a prediction graph, in which some edges can contain an undefined trust value (where the trust metric could not predict the value). The real and the predicted values are then compared in several ways: the coverage, which is a measure of the edges that were predictable, the fraction of correctly predicted edges, the mean absolute error (MAE) and the root mean squared error (RMSE). Surely there are other ways of evaluating trust metrics: for example it can be argued that an important task for trust metrics is to suggest to a user which other still unknown users are more trustworthy, for example suggesting a user worth following on a social bookmarking site such as del.icio.us or on a music community such as Last.fm (for example because she is trusted by all the users the active user trusts). In this case the evaluation could just concentrate on the top 10 trustworthy users. But in this first work we considered only leave-one-out.

Evaluation of trust metrics on all trust edges

Table 1 reports our evaluation results of different trust metrics on the Advogato dataset. It is a computation of different evaluation measures on every edge present in the social network. The reported measures are fraction of wrong predictions, Mean Absolute Error, Root Mean Squared Error and coverage. We now describe the compared trust metrics. As already mentioned, we released the code and we plan to implement more trust metrics and release them and run the evaluations. We also applied a threshold function in case of trust metrics that can return values in a continuous interval, such as Moletrust and PageRank, so that for example a predicted trust of 0.746 becomes 0.8 (Apprentice). The compared trust metrics are some trivial ones used as baselines such as Random, which predicts simply a random trust score in the range [0.4, 1] thresholded in the normal way, or the metrics starting with “Always” which always return the corresponding value as predicted trust score. Other simple trust metrics are OutA which, in predicting the trust user A could have in user B, simply does the average of the trust statements outgoing from user A, and OutB which averages over the trust statements outgoing from user B.

The other trust metrics were already explained in Section of trust metrics, here we just report on how we thresholded and how we run them. Ebay refers to the trust metric that, in predicting the trust user A could have in user B, simply does the average of the trust statements incoming in user B, i.e. the average of what all the users think about user B. MoletrustX refers to Moletrust applied with a trust propagation horizon of value X. The values returned by PageRank as predicted trust follow a powerlaw distribution, there are few large PageRank scores and many tiny ones. So we decided to rescaled the results simply by sorting them and linearly mapping them in the range [0.4, 1], after this we thresholded the predicted trust scores. Our implementation of Advogato is based on Pymmetry1. AdvogatoGlobal refers to the Advogato trust metric run considering as seeds the original founders of Advogato community, namely the users “raph”, “federico”, “miguel” and “alan”. This is the version that is running on the Advogato web site for inferring global certifications for all the users. This version is global because it predicts a trust level for user B which it is the same for every user. AdvogatoLocal refers to the local version of Advogato trust metric. For example, when predicting the trust user A should place in user B, the trust flow starts from the single seed “user A”. This version is local because it produces personalized trust predictions which depends on the current source user and can be different for different users. AdvogatoLocal was run on a subset (8%) of all the edges since the current implementation is very slow. Due to the leave-one-out technique the network will be different for every evaluation and it has to be restarted from scratch for every single trust edge prediction.

The results of the evaluation are reported in Table 1.We start by commenting the column “fraction of wrong predictions”. Our baseline is the trust metric named “Random” which produces an incorrect predicted trust score 74% of the times. The best one is Ebay with an error as small as 35% followed by

Moletrust2 (36.57%), Moletrust3 (37.60%) and Moletrust4 (37.71%). Increasing the trust propagation horizon in Moletrust allows to increase the coverage but also increases the error. The reason is that users who are near-by in the trust network (distance 2) are better predictors than users further away in the social network (for example, users at distance 4).

Note that Moletrust is a local trust metric that only uses information located “near” the source node so it can be run on small devices such as mobiles which only need to fetch information from the (few) trust users and possibly the users trusted by them. This behaviour is tunable through setting the trust propagation horizon to specific values. On the other hand, Ebay, being a global trust metric, must aggregate the entire trust network, which can be costly both in term of bandwidth, memory and computation power. The AlwaysX metrics depend on the distributions of certifications and are mainly informative of the data distribution.

The fraction of wrong predictions of Advogato (both local and global) is high compared to Ebay and Moletrust. Advogato was not designed for predicting an accurate trust value, but to increase attack-resistance while accepting as many valid accounts as possible. A side effect is that it limits the amount of granted global certifications and assigns a lot of Observer certificates. In the case of AdvogatoGlobal, 45% of the predicted global certifications are marked as Observer which obviously has an impact on the leave-one-out evaluation. Different trust metrics might have different goals, that require different evaluation techniques.

Note that the local version of Advogato is more accurate than the global version.

The last metric shown in Table 1 is PageRank : the fraction of correct predictions is not too high but again the real intention of PageRank is to rank web pages and not to predict the correct value of assigned trust.

An alternative evaluation measure is the Mean Absolute Error (MAE). The MAE is computed by averaging the difference in absolute value between the real and the predicted trust statement on an edge. There is no need to threshold values because MAE computes a meaningful value for continuous values. The MAE computed for a certain thresholded trust metric is generally smaller than the MAE computed for the same trust metric when its trust score predictions are not thresholded. But in order to compare metrics that return real values and others that return already thresholded values, we consider the MAE only for thresholded trust metrics. The second column of Table 1 reports the MAE for the evaluated thresholded trust metrics. The baseline is given by the Random trust metric which incurs in a MAE of 0.2230. These results are the worst besides the trivial trust metrics that always predict the most unfrequent certification values.

Predicting always Journeyer (0.8) incurs in a small MAE because this value is frequent and central in the distribution. Ebay is the trust metric with the best performance, with a MAE of 0.0855. And it is again followed by Moletrust that in a similar way is more accurate with smaller trust propagation horizons.

A variant of MAE is Root Mean Squared Error (RMSE). RMSE is the root mean of the average of the squared differences. This evaluation measure tends to emphasize large errors, which favor trust metrics that remain within a small band of error and don’t have many outlying predictions that might undermine the confidence of the user in the system. For example, it penalizes a prediction as Journeyer when the real trust score should have been Master, or vice versa.

Table 1. Evaluation of trust metrics on all trust edges

	Fraction wrong pre-dictions	MAE	RMSE	Coverage
Random	0.737	0.223	0.284	1.00
AlwaysMaster	0.670	0.203	0.274	1.00
AlwaysJourneyer	0.585	0.135	0.185	1.00
AlwaysApprentice	0.834	0.233	0.270	1.00
AlwaysObserver	0.911	0.397	0.438	1.00
Ebay	0.350	0.086	0.156	0.98
OutA	0.486	0.106	0.158	0.98
OutB	0.543	0.139	0.205	0.92
Moletrust2	0.366	0.090	0.160	0.80
Moletrust3	0.376	0.091	0.161	0.93
Moletrust4	0.377	0.092	0.161	0.95
PageRank	0.501	0.124	0.191	1.00
AdvogatoLocal	0.550	0.186	0.273	1.00
AdvogatoGlobal	0.595	0.199	0.280	1.00

The baseline Random has an RMSE of 0.2839. With this evaluation measure too, Ebay is the best metric with an RMSE of 0.1563 and all the other performances exhibit a pattern similar to the one exposed for the other evaluation measures. However there is one unexpected result: the trivial trust metric OutA is the second best, close to Ebay. Remind that, when asked a prediction for the trust user A should place in user B, OutA simply returns the average of the trust statements going out of A, i.e. the average of how user A judged other users. This trust metric is just a trivial one that was used for comparison purposes. The good performance of OutA in this case is related to the distribution of the data in this particular social setting. The Observer certification has special semantics:

it is the default value attributed to a user unless the Advogato trust metric gives a user a higher global certification. So there is little point in certifying other users as Observer. In fact, the FAQ specifies that Observer is “the level to which you would certify someone to remove an existing trust certification”.

Observer certifications are only when a user changes its mind about another user and wants to downgrade her previously expressed certification as much as possible. This is also our reason for mapping it to 0.4, a less than sufficient level.

As a consequence of the special semantics of observer certifications, they are infrequently used. In fact only 638 users used the Observer certification at least once while, for instance, 2938 users used the Master certification at least once.

Trust metrics like Ebay and Moletrust work doing averages of the trust edges present in the network (from a global point of view for Ebay and only considering the ones expressed by trusted users for Moletrust) and, since the number of Observer edges is very small compared with the number of Master, Journeyer and Apprentice edges, these predicted average tend to be close to higher values of trust. This means that when predicting an Observer edge (0.4) they lead to a large error. This large error is weighted a lot by the RMSE formula. On the other hand, using the average of the outgoing trust edges (like OutA does) happens to be a successful technique for not incurring in large errors when predicting observer edges. The reason is that a user who used Observer edges tended to use it many times so the average of its outgoing edge certifications is a value that is closer to 0.4 and hence it incurs in lower errors on these critical edges and, as a consequence, in smaller RMSE. This effect can also be clearly seen when different trust metrics are restricted to predict only Observer edges and evaluated only on them. In this case (not shown in Tables), OutA gets the correct value for trust (Observer) 42% of times, while for instance, Ebay only 2.7% of times and Moletrust2 4%. The fact OutA exhibits a so small RMSE supports the intuition that evaluating which conditions a certain trust metric is more suited for than another one is not a trivial task. Generally knowledge about the domain and the patterns of social interaction is useful, if not required, for a proper selection of a trust metric for a specific application and context.

The last column of Table 1 reports the coverage of the different trust metrics on the Advogato dataset. Sometimes a trust metric might not be able to generate a prediction and the coverage refers to the number of edges that are predictable. The experiment shows that the coverage is always very high. Since local trust metrics use less information (only trust statements of trusted users) their coverage is smaller than the coverage of global trust metrics. Anyway, differently from other social networks , it is very high. The Advogato trust network is very dense, so there are many different paths from a user to another user. Even very local trust metrics such as Moletrust2, that only use information from users at distance 2 from the source user, are able to cover and predict almost all the edges.

Evaluation of trust metrics on controversial users

As a second step in the analysis we concentrated on controversial users . Controversial users are users which are judged in very diverse way by the members of a community. In the context of Advogato, they can be users who received many certifications as Master and many as Apprentice or Observer: the community does not have a single way of perceiving them. The intuition here is that a global average can be very effective when all the users of the community agree that “raph” is a Master, but there can be situations in which something more tailored and user specific is needed. With this in mind we define controversial users as Advogato users with at least 10 incoming edges and standard deviation in received certifications greater than 0.2. Table 2 shows the results of the evaluation of the different trust metrics when they are restricted to predicting the edges going into controversial users. In this way we reduce the number of predicted edges from 52981 to 2030, which is still a significant number of edges to evaluate trust metrics on.

In order to understand better the nature of trust edges under prediction in this second experiment, it is useful to note that, of edges going into controversial users, 1093 are of type Master, 403 of type Journeyer, 115 of type Apprentice and 419 of type Observer. The variance in the values of trust certificates is of course due to the fact that these users are controversial and it is also the reason for which predicting these edges should be more difficult.

We start by commenting the evaluation measures on AlwaysMaster (second row of Table 2) because it presents some peculiarities. Always Master predicts the correct trust value 53.84% (100% 46.16%) of times and, according to the evaluation measure “fraction of correctly predicted trust statements”, seems a good trust metric, actually the best one. However the same trust metric, AlwaysMaster, is one of the less precise when RMSE is considered. A similar pattern can be observerd for AdvogatoGlobal. In fact, since in general there is at least one flow of trust with Master certificates going to these controversial users, AdvogatoGlobal tends to predict almost always Master as trust value and

Table 2. Evaluation of trust metrics on trust edges going into controversial users

	Fraction wrong predictions	MAE	RMSE	Coverage
Random	0.799	0.266	0.325	1.00
AlwaysMaster	0.462	0.186	0.302	1.00
AlwaysJourneyer	0.801	0.202	0.238	1.00
AlwaysApprentice	0.943	0.296	0.320	1.00
AlwaysObserver	0.794	0.414	0.477	1.00
Ebay	0.778	0.197	0.240	0.98
OutA	0.614	0.147	0.199	0.98
OutB	0.724	0.215	0.280	0.92
Moletrust2	0.743	0.195	0.243	0.80
Moletrust3	0.746	0.194	0.241	0.93
Moletrust4	0.746	0.195	0.242	0.95
PageRank	0.564	0.186	0.275	1.00
AdvogatoLocal	0.518	0.215	0.324	1.00
AdvogatoGlobal	0.508	0.216	0.326	1.00

since almost half of the edges going into controversial users are of type Master, AdvogatoGlobal often predicts the correct one.

This means that the same trust metric might seem accurate or inaccurate depending on the evaluation measure. This fact once more highlights how evaluating trust metrics on real world datasets is a complicated task and a comparison of same metrics on many different datasets according to different evaluation methods would be highly beneficial for understanding the situation in which one trust metric is more appropriate and useful than another. We already previously explained why OutA is able to have a so small RMSE, the smallest one on controversial users: based on how Observer certifications are used in the system, OutA is the only metric that is able to avoid large errors when predicting the Observer edges.

Arriving at a comparison between a global trust metric such as Ebay and a local trust metric such as Moletrust, we were expecting the latter to be more accurate than the first on controversial users. While on the Epinions dataset, this is what was observed , the same is not true here. The reason is partly that in Epinions, the trust values were binary (either trust or distrust) and it was easier to discriminate. Another reason seems to be that on Advogato the user base is not divided in cliques of users such that users of one clique trust each other and distrust users of other cliques. In fact Advogato users are somehow similar and feel part of one single large community. It is future work to analyze if on a social network with a much higher polarization of opinions (such as for example essembly.com, a political site) the performances of local trust metrics are significantly better than global ones.

The trusty problems in the social network

Trust problems occur due to various reasons. Firstly users are not aware or don't care about the access control and privacy settings that might sometimes lead to unwanted situations.

Secondly, it's not too uncommon that people in the social networks friend with people who they have not ever met or don't really know at all. If you trust a person you don't know it can compromise and your personal information and give a false indication of trust for others.

Thirdly, sensitive private information could be compromised if an account is unauthorizedly accessed. It not only compromise that user's data but all of the friends are compromised as well. Access to account will also give an intruder the possibility to effect the trust relationships of the breached account.

The problem of joint fraud

A joint fraud in this context is a type of collusion where a group of people act cooperatively towards a common goal to deceive other users. Malicious users can form a group to gain trust that can be used to cheat a trust metric or actual individuals. Application of a joint fraud is for example boosting up a reputation rank in an online store. Joint fraud is important in this context because it's a result of system that relies on the trust mechanisms of social networks. A cleaver algorithm is needed to detect a joint fraud.

Policy on trusty mechanism of social network

Efficient policy for a trust mechanism needs to bring value to users by taking account the privacy of the users without compromising the usability too much. It also needs to take account trust exploiting possibilities of current solutions.

The general expression mode of users' identity

The most obvious solution here would be to demand a certification from the users of the social network to verify who they really are. This would solve the problems related to false identities and bogus persons. If all the users would be verified people could really trust that they are who they claim to be. At the moment the identities of the users is verified through email verification. Issue with email verification is that it’s trivial for a person to create bogus and false identities and it only merely blocks computer generated identities.

A new system that uses a stronger verification such as an E-bank account or national electronic identity could be generated to handle the verification process instead.

This approach still has major concerns and fallbacks. This would at first result in major costs. Implementation of such a system is not trivial because there is no universal way to verify a person reliably electronically. It’s likely that all users could not be verified especially those who live in developing countries where electronic verification methods rarely exists.

The most global form of verification would be verification by a credit card but this would again limit the user space to credit card holders only. The major problem with this type of verification is also that it will result in difficulties in case of identity theft. However, on the business perspective it would open up opportunities if the people could use their social network accounts for transactions. I see that this type of approach could happen in the near future when electronic verification develops enough.

Easier approach here could be the use of existing social network and instead of trusted 3rd party and let the members of the network decide who is trustworthy. Members of an user’s social network could verify that the user is who user claims to be. Users are likely to distinguish false identities and bogus users from the real ones. This could be done by polling all the users to verify their members to verify their friends or one might say that this is not even necessary because the bogus users don’t likely have a large social network. This type of identification method is easy to implement but it hides an internal problem of joint fraud in it. A group of people could verify a false identity to be a real one. At the moment identification usually works kind of other way around where user's can report malicious users to the maintenance. One might say that the identification of social networks are strong enough already.

It should be noted that none of the identification methods would resolve trust related problems such as joint-fraud where legitimate people maliciously co-operate. In addition to identification there should be a reliability metric to measure a reputation of an identity to avoid these type of issues.

The expression mode of trusty mechanism between users

Trust in social networks can be expressed by a trust metric. A trust metric can determine how trustworthy another person is to the user, but because of the nature of trust is subjective, a common algorithm for a trust mechanism is hard to derive.

User should have a control of the trust on social network like in the real life.

A trust metric in social network context can be based on the links between users. The previous works in this field have been concentrated on reputation of websites and P2P systems.

There are solutions like Eigentrust that evaluates trust between nodes in P2P networks and peer-trust. As P2P networks are somewhat similar to social networks - as they are connections between real people – these type of metrics might be relevant in social networks also.

Probably the most well-known trust metric in the world is Google's patented PageRank that evaluates the trustworthiness of website links. An algorithm similar to PageRank called NodeRank can be used when evaluating links in social networks. This approach has been criticized because connection based trust assessment hides the fact in human psychology trust is multi-dimensional and result of various parameters rather than connections. Trust evaluation should take account other activity related information such as profile information, comments and internet activity of the user because that's they way the users evaluate trust in psychological context.

Another approach is to evaluate the trust in social networks by parametrizing the connections in social networks by the activities on the network . This is done by evaluating trust by how often the peers in the network communicate with each other and the more they do the more they trust each other . A combining method has been researched that composes of trust relationships, influential and environmental factors. This model exists only in theory but it still is the most promising one.

It's possible to evaluate trust by comparing profile similarities in social networks. But this merely indicates the real psychological trust between users if they only have same interests.

Completely different approach is to redesign the social network architecture and instead of web-based centralized applications use a distributed P2P network to implement a social network and increase it's trustworthiness.

But this approach would need further research to increase performance and usability. Interaction based trustworthiness assessment can have applications in e-commerce communities where a reliable review and recommendation systems provide value to customers.

The expression mode of trusty mechanism of multi-hop users

Users should have control of privacy in social networks. Users should be able to decide whether and what information multi-hop users can see. In many cases it makes sense to share some information to peers that the peer don't know or trust but in many cases there are some information that peers only want to share within their own trust networks. In Facebook for example a user can very specifically determine what information is shared and who can access it. It's hard to determine whether a multi-hop user is trustworthy or not when that user is not in the user's social network.

A trivial solution for a metric here would be to count the mutual friends between the user and a multi-hop user and determine the trust that way. Such an algorithm exists an it's called TidalTrust. This could be possibly be applicable in economical sense but in sociological sense and using the common sense this is not however a good solution because in large communities there is usually people who have mutual friends but haven't never met and therefore a trust enabling social connection hasn't formed. People don't always cope well with other people and case can also be that even though two peers have many mutual friends they actually distrust each other or have had trust conflicts in the past.

This derives from the fact that trust is usually unidirectional and subjective.

How to solve the problem of joint fraud

A solution to the joint fraud problem can relay inside the social networks. If a strong identification is used - people use their real identities - and social network data is used a collusion detection mechanism can be built. This idea is based on a fact that when measuring reputation by excluding the nearest connections to friends. Another strong candidate for a solution of the joint fraud problem and in general a social network with malicious users is SocialTrust framework. This framework bases it's ideas to the connections and the qualities of these connection on social networks . SocialTrust is according to the research more efficient in precision compared to the PageRank and TrustRank trust assesment algorithms.

Classification dimensions

Trust and reputation can be analyzed from different perspectives and can be used in a wide range of situations. This makes the classification of trust and reputation models a difficult task. In this section we propose a set of aspects with which we classify the current computational trust and reputation models in a clear landscape. we focus our attention on computational models. Therefore, the classification dimensions have been selected considering the special characteristics of these kind of models and the environment where they have to evolve.

Conceptual model

According to the conceptual model of reference, trust and reputation models can be characterized as:

– **Cognitive**. As pointed out in (Esfandiari and Chandrasekharan, 2001), in models based on a cognitive approach 'trust and reputation are made up of underlying beliefs and are a function of the degree of these beliefs'. In the cognitive approach, the mental states that lead to trust another agent or assign a reputation, as well as the mental consequences of the decision and the act of relying on another agent, are an essential part of the model.

– **Game-theoretical**. Trust and reputation are considered 'subjective probabilities by which an individual, A, expects that another individual, B, performs a given action on which its welfare depends' (Gambetta, 1990). Trust and reputation are not the result of a mental state of the agent in a cognitive sense but the result of a more pragmatic game with utility functions, and numerical aggregation of past interactions.

Information sources

It is possible to classify trust and reputation models considering the information sources that they take into account to calculate trust and reputation values. Direct experiences and witness information are the "traditional" information sources used by computational trust and reputation models. In addition to that, a few models have recently started to use information associated to the sociological aspects of agents'behavior.

The kind of information available to an agent depends on its sensory capabilities. The use of several information sources, if they are taken into account in a smart way by the model, can increase the reliability of the calculated trust and reputation values but at the same time increases the complexity of the model. Moreover, scenarios that allow agents to obtain diverse information demand smarter (and, therefore, more complex) agents.

Direct experiences

This is, without doubt, the most relevant and reliable information source for a trust/reputation model. There are two types of direct experiences that an agent can include as part of its knowledge. The first, and used by all the trust and reputation models analyzed in this review, is the experience based on the direct interaction with the partner.
The second is the experience based on the observed interaction of other members of the community. This second type is not so common and restricted to scenarios that are prepared to allow it. Usually, in those models that consider the observation of other partners activity, a certain level of noise in the obtained information is assumed.

Witness information

Witness information (also called word-of-mouth or indirect information) is the information that comes from other members of the community.
That information can be based on their own direct experiences or it can be information that they gathered from others. If direct experience is the most reliable source of information for a trust/reputation model, witness information is usually the most abundant. However, it is far more complex for trust and reputation models to use it. The reason is the uncertainty that surrounds this kind of information. It is not strange that witnesses manipulate or hide pieces of information to their own benefit.

Sociological information

The base of this knowledge are the social relations between agents and the role that these agents are playing in the society. In real world, the individuals that belong to a given society establish different type of relations among them. Examples of these relations can be dependence, trade, competition, collaboration and so on. Also, each individual play one (or several) role(s)in that society. Both, the relations and the role or roles the individual play in the society influence his/her behavior and the interaction with the others.

The social relations established between agents in a multi-agent system are a simplified reflection of the more complex relations established between their human counterparts. This kind of information is only available in scenarios where there is a rich interaction between agents.

Currently, only a few trust and reputation models use this knowledge applied to agent communities to calculate or improve the calculation of trust and reputation values. These models use techniques like *social network analysis*. Social network analysis is the study of social relationships between individuals in a society that emerged as a set of methods for the analysis of social structures, methods that specifically allow an investigation of the relational aspects of these structures. The use of these methods, therefore, depends on the availability of relational data. Although currently the number of models that take into account this kind of information is reduced, we guess that the increase of complexity in multi-agent systems will make it more and more important in the near future.

Prejudice

The use of prejudice to calculate trust and reputation values is another mechanism not very common but present in current trust and reputation models. Prejudice is the mechanism of assigning properties (like for instance a reputation) to an individual, based on signs that identify the individual as member of a given group. These signs can be anything: a uniform, a concrete behavior, etc. A good analysis of the use of signs in trust is performed by Bacharach and Gambetta in (Bacharach and Gambetta, 2001).

As most people today use the word, “prejudice” refers to a negative or hostile attitude toward another social group, usually racially defined.

However, the negative connotations that prejudice has in human societies has to be revised when applied to agent communities. Differently from the signs used in human societies that range from skin color to sex, the set of signs used in computational trust and reputation models are usually out of ethical discussion.

Visibility types

Trust and reputation of an individual can either be seen as a global property shared by all the observers or as a subjective property assessed particularly by each individual.

In the first case, the trust/reputation value is calculated from the opinions of the individuals that in the past interacted with the individual being evaluated. This value is publicly available to all members of the community and updated each time a member issues a new evaluation of an individual. In the second case, each individual assigns a personalized trust/reputation value to each member of the community according to more personal elements like direct experiences, information gathered from witnesses, known relations between members of the community and so on. In the latter case, we cannot talk about the trust/reputation of an individual *x*, we have to talk about the trust/reputation of an individual *x* from the point of view of an individual *y*.

The position of taking trust and reputation as a global property is common in online reputation mechanisms. These systems are intended for scenarios with thousands or even millions of users. As pointed out by Dellarocas (Dellarocas, 2003), the size of these scenarios makes repeated interaction between the same set of players unlikely and, therefore, reduces the incentives for players to cooperate on the basis of hoping to develop a profitable relationship.

Take the example of an electronic auction house like those accessible nowadays through Internet. One day, the user wants to buy a book and the next day s/he wants to buy a computer. The intersection between users selling books and users selling computers is probably empty so the few personal experiences accumulated buying books are not useful in the computers' market. Computer sellers are unknown for the user so s/he has to rely on the information that people who bought computers in the past has left in the form of a reputation value. The robustness of these systems relies on the number of opinions available for a given partner. A great number of opinions minimize the risk of single individual biased perceptions.

In models that consider trust and reputation as a global property, the main problem is the lack of personalization of that value. Something that is bad for me could be acceptable for others and the other way around. Although this approach can be acceptable in simple scenarios where it is possible to assign a common

“way of thinking” to all members of the community, it is not useful when agents have to deal with more complex and subjective affairs.
The antithesis of these models are the models that consider trust and reputation as a subjective property. Each agent uses its personal experiences and what the other agents have said to it personally, among other things, to build the trust and reputation of each member of the community. These models are indicated for medium and small size environments where agents meet frequently and therefore it is possible to establish strong links among them.

Model’s granularity

Is trust/reputation context dependent? If we trust a doctor when she is recommending a medicine it does not mean we have to trust her when she is suggesting a bottle of wine. The reputation as a good sportsman does not help if we are looking for a competent scientist. It seems clear that the answer is yes: trust and reputation are context dependent properties. However, adding to computational trust and reputation models the capability to deal with several contexts has a cost in terms of complexity and adds some side effects that are not always necessary or desirable.
A single-context trust/reputation model is designed to associate a single trust/reputation value per partner without taking into account the context. A multi-context model has the mechanisms to deal with several contexts at a time maintaining different trust/reputation values associated to these contexts for a single partner.

One could argue that it is always possible to transform a single context model into a multi-context one just having different instances of the single-context model, one for each considered context. However, if there is something in trust and reputation environments that is usually scarce, that is the information used to calculate trust and reputation values. So what really gives to a model the category of being a multi context model is the capability of making a smart use of each piece of information to calculate different trust or reputation values associated to different activities. Identifying the right context for a piece of information or using the same information in several contexts when it is possible are two examples of the capabilities that define a real multi-context model.
Is this always necessary? Certainly not. Nowadays, there are very few computational trust and reputation models that care about the multi context nature of trust and reputation and even fewer that propose some kind of solution. This is because current models are focused on specific scenarios with very delimited tasks to be performed by the agents. In other words, it is possible to summarize all the agent activities in a single context without losing too much versatility. However, and similarly to what we have mentioned before about the use of sociological information, as the complexity of tasks to be performed by agents will increase in the near future, we may also expect an increase of the importance devoted to this aspect in trust modeling.

Agent behavior assumptions

The capacity to deal with agents showing different degrees of cheating behavior is the aspect considered here to establish a classification. We use three levels to categorize trust and reputation models from this point of view according to what we have observed in the analyzed trust and reputation models:

– **Level 0.** Cheating behavior is not considered. The model relies on a large number of agents who offer honest ratings to counteract the potential effect of the ratings provided by malicious agents.

– **Level 1.** The model assumes that agents can hide or bias the information but they never lie.

– **Level 2.** The model has specific mechanisms to deal with liars.

Type of exchanged information

The classification dimension here is the type of information expected from witnesses. We can establish two big groups. Those models that assume boolean information and those models that deal with continuous measures. Although it seems a simple difference choosing one approach or the other has a great influence in the design of the model.

Usually, models that rely on probabilistic methods work with Boolean information while those models based on aggregation mechanisms use continuous measures.

Trust/Reputation reliability measure

Is the model providing a measure of how reliable is the trust/reputation value? Sometimes, as important as the trust/reputation value itself is to know how reliable is that value and the relevance it deserves in the final decision making process. Some models incorporate mechanisms that provide this kind of information. In the models we have analyzed, this measure is a single value associated to the trust or reputation value.

Depending on the model, the elements that are considered to calculate the reliability measure are different. Among them you can find elements like the number of experiences, the reliability of witnesses, how old is the information used to build trust and reputation, and so on.

Personalized Search

A number of researchers have attempted to offer personalised search, using a range of approaches for capturing broader information about the user from which to infer their information needs.

At the level of general Web search, Jeh and Widom (2003) present a modified version of the *PageRank* algorithm (Page, Brin et al., 1999). This approach takes as input a user's list of Web bookmarks, each of which is taken as an implicit endorsement of the relative importance of that document to the user. Based on this input, personalised *PageRank* scores can be calculated, which enables a personalised rather than a global view of the importance of Web documents, and can serve as a basis for ranking search results.

Specifically in the context of a job-seeking site, Bradley, Rafter and Smyth (2000) report on a system that filters search results by comparing these to a user profile based on jobs she has previously viewed and rated. This approach requires more extensive and ongoing input from the user compared to that of Jeh and Widom (2003), as the user must actively view and rate job advertisements in order to receive personalised results. The system is also domain-specific; however it could be extended to allow the capture of viewing and rating data for any collection of items. It remains to be seen whether explicit ratings such as these, compared to the implicit endorsement of bookmarking a Web site, provide more accurate data on which to base user profiles for personalised search.

Teevan, Dumais et al. (2005) report on an investigation into the potential value of personalised search results compared to those provided by current search engines whereby all users receive the same results. They found that search results currently reflect a broad range of different search intents, meaning that relevance to the intents of the group as a whole is generally high. However, the relevance of generic results to individual search intentions was considerably lower. Interestingly it was found that agreement about the relevance of results between individuals choosing the same search query was lower than found in previous studies. This finding is attributed to the study's emphasis on participants rating the relevance of results to their personal information needs rather than an abstract notion of the results' relevance to a topic.

Furthermore, it was found that inter-rater agreement on the relevance of results was relatively low (62%) even for those participants who used the same query and whose expressed intentions were the same. It was concluded that participants struggled to unambiguously describe their search intention, therefore the same description actually covered more than one intention and the relevance rating of results varied as a consequence. Based on these findings, Teevan et al. conclude that there may be value in personalising search results, and propose a technical approach based on re-ordering results retrieved through conventional search engines.

Many search personalisation approaches are limited by only exploiting information provided specifically by that user. For example, Bradley, Rafter et al.'s (2000) system only bases personalisation on viewing and rating data from the user themselves. Similarly Jeh and Widom's (2003) approach does not specifically address the use of other people's bookmark collections to aid one's own personalisation.

Using data about just one user does not allow for economies of scale through collaboration, whereby the knowledge and experience of other people could be used to aid the information-seeking process. Whilst search personalisation approaches go by a different name they share much in common with recommender systems, as both attempt to identify subsets of relevant items on the user's behalf. The next section will examine the two major classes of recommender systems, one of which takes an explicitly collaborative approach.

Content-Based Recommendation

Recommender systems generally follow one of two approaches, *content-based recommendation* or *collaborative filtering* (Balabanovic and Shoham, 1997). Content-based recommendations can be made in a number of ways: by matching the content of an item to some input such as a user profile (e.g. Balabanovic and Shoham, 1997) or keyword terms (in the case of a search engine); or by matching the content of an item to that of another item for which the user has already shown some preference. For example, action films may be recommended to the user where they have expressed an interest in this type of film, or satirical comedies where they have previously viewed or purchased items from this genre. Web search engines are an example of the content-based method, whereby results are returned based on content matches between documents and user-supplied keywords. Content-based recommendation is not limited to textual documents, and can be applied to other media formats. For example, Celma (2006) has used the approach to recommend musical artists based on the characteristics of their music.

Social Navigation

Social *navigation* is a design approach that aims to utilize the presence and actions of people in online environments as a means to assist others in navigating the same virtual spaces. Therefore users may be supported in locating and evaluating information and subsequent decision making, through mechanisms such as visualized traces of other peoples' activities, or direct communication channels (e.g. chat) with other users of a system (Dieberger, Höök, Svensson et al., 2001, Dieberger, Dourish et al., 2000).

The term was originally coined by Dourish and Chalmers (1994) in order to distinguish between social navigation (based on information from other people) and semantic navigation (based on the underlying structure of the information being navigated).

Whilst recommender systems, and in particular collaborative filtering applications that reuse the efforts or actions of other people to filter information, have been treated as one form of social navigation, there are many other avenues of research that fall under the same label. In fact, the nature of social navigation has been interpreted fairly broadly, giving rise to a wide range of applications and approaches. For example, the *Footprints* system (Wexelblat and Maes, 1999) uses visio-spatial metaphors such as maps and paths to indicate how previous users interacted with a Web site, whilst Svensson, Hook, Laaksolahti et al. (2001) bring together a number of social navigation features such as chat, recommendation, and avatars in a system for navigating food recipes.

Mobasher, Cooley and Srivastava (2000) describe a proof of concept system, called *WebPersonalizer*, that suggests potentially relevant pages to the user while they browse the site. Whilst not explicitly labelled as an example of social navigation, this application follows the same principles. Suggestions are made based on data about how previous users have navigated the site, obtained by analysing Web server logs. The analysis is performed anonymously therefore all users who follow the same navigation path on the site receive the same suggestions. This has the potential to be rather self-reinforcing, whereby all users are channelled along similar paths irrespective of their underlying information need or task.

In addition to social support for browsing, *social search* has been investigated. The term 'social search' can be interpreted in two ways. The first of these falls under the umbrella of social navigation and sees social search as supporting conventional search processes with information derived from the actions or preferences or other people.

This first interpretation of the term is adopted by researchers such as Ahn, Brusilovsky and Farzan (2005) who explore the use of page visit data and user annotations to supplement search results in their *Knowledge Sea* application. Search results are ranked using a conventional document ranking technique and then supplemented by displaying users' own visit frequency for particular documents alongside aggregate visit data from a wider group and indications of the degree of 'praise' the document has received. The use of page view data and endorsements (in the form of positive or negative praise) in the results interface bears many similarities to the use of customer purchase data or ratings in collaborative filtering recommender systems. However, the nature of the group from which aggregate statistics are drawn is not specified, and as with collaborative filtering performance is reliant on the behaviour of other anonymous users.

The second interpretation of the 'social search' label refers to searching a social network to identify particular individuals who may be able to assist with the current task.

The importance of maintaining privacy in social navigation systems has been raised (e.g. Dieberger, Höök et al., 2001). However, it is also argued in the same paper that a degree of *visibility* is essential in order for applications to retain utility, which certainly points towards pseudonymous and possibly even towards known identities in social navigation systems.

In support of the arguments put forward by Bonhard and colleagues described above, and counting against anonymous applications, Kautz, Selman et al. (1997b, 1997a) argue that not all information sources are equally desirable. Consequently, personal referrals between known individuals allow the information seeker to make judgements about the quality of the information they are receiving and may instil greater confidence in the information if the source is trusted.

How people select sources for information and recommendations will be reviewed in detail in following pages. However, before this, research and systems will be discussed that attempt to integrate more directly with known social networks, and utilize these to support the information-seeking process.

Social Networks and Trust

A number of attempts have been made to enhance social network-based approaches to information-seeking with notions of trust. In most cases trust is employed as a fairly broad, non-specific concept. These attempts are examined below, and can be analysed according to four dimensions:

1. **automation**: the degree to which trust ratings are automatically computed (versus provided manually)

2. **topicality**: the degree to which trust ratings are topical in nature (versus one global trust rating of an individual across all topics)

3. **individuality**: the degree to which trust ratings are personal (versus one global trust rating of an individual shared by all others)

4. **anonymity**: the degree to which the system operates over networks of known individuals (versus operating across systems of unknown individuals, or a mixture of the two)

Golbeck and Hendler reach beyond the network of personally known individuals by combining social networks and inferred trust/reputation relationships in an email filtering application (*TrustMail*) (Golbeck and Hendler, 2004) and film recommender system (*FilmTrust*) (Golbeck and Hendler, 2006).

The goal of *FilmTrust* is not to actively suggest items to the user unprompted, but to provide her with feedback on how likely she is to be interested in a film she has already found, based on direct or inferred trust relationships. Film reviews are also ranked on the same basis when displayed on the site. In a similar fashion, *TrustMail* annotates each email in the user's inbox with a trust rating, based on trust relationships computed through the network between sender and receiver.

To participate in the trust networks associated with these applications, and benefit from their filtering and ranking capabilities, the user must manually rate (on a 1-10 scale) the reputation of, or their trust in, people they know. In *TrustMail* these ratings are non-domain-specific 'reputation' ratings of the known person, whereas in *FilmTrust* the user is asked to rate her trust in the person in the context of films. These ratings then seed the algorithmic creation of trust scores for all other members of the wider network to whom the user is linked socially. Importantly, these scores are computed from the user's local perspective, rather than being global to the entire network. This work is characterised by a mixed approach to *automation*, no *topicality* in the *TrustMail* system but a limited amount in *FilmTrust*, a high level of *individuality*, and varying degrees of *anonymity*.

The approach is useful in that it enables trust ratings to be inferred between individuals who are connected to some degree, but do not know each other personally. This can be of value where insufficient information is available within one's immediate network, or one's immediate network is too small. In addition, there is some evidence (Golbeck and Hendler, 2006) to suggest that this approach can produce more accurate results than 'nearest neighbour' collaborative filtering techniques in situations where the user's tastes are divergent from the population as a whole.

However, the approach has a number of limitations. Firstly, the semantics of the trust relationships are often ambiguous or underspecified. In *TrustMail* users are asked to rate the general reputation of people they know. Whilst reputation may not be quite so context dependent as trust, this still appears to be a gross oversimplification. For example, a researcher may have an excellent academic reputation, but be known to be unreliable when repaying loans. In the context of email filtering the risks associated with this are small, however under-specifying relationships in this way does limit the value and reusability of the data.

The ontology Golbeck, Parsia and Hendler (2003) used to describe the trust ratings provided by users in *TrustMail* and *FilmTrust* does in fact allow specification of the topic or domain in which the trust is being asserted, and whilst users of the *FilmTrust* system are asked to provide trust ratings in the context of film reviews, this relationship is not explicitly encoded in output from the system.

Secondly, this approach does rely on provision of manual trust ratings between users to bootstrap the process. Whilst making just one social connection in the *FilmTrust* network does allow recommendations to be made for a user for 95% of films, it would be desirable to investigate existing sources of information from which trust relationships between known individuals could be inferred, in order to bypass this manual annotation process and lower the cost of participation for users.

Thirdly, the work of Golbeck and colleagues uses trust ratings as the basis for making similarity assessments between users. This is justified by reference to work by Ziegler and Lausen (2004) that found a correlation between trust and user similarity in the online community *All Consuming*. Whilst trust may serve as a valid proxy for similarity, this correlation may be due to a third factor which has not been accounted for, and as such the predictive validity of this relationship should be questioned.

Numerous other attempts have been made to integrate notions of trust with social networks. For example, Richardson, Agrawal and Domingos (2003) describe a distributed 'web of trust' approach, intended to support the assessment of 'belief' in assertions on the Semantic Web as a function of the user's subjective trust in the author of the statements. The approach assumes that no one entity will know the trustworthiness of every other, and therefore ratings cannot be assigned to an entity by a central source. On this basis, the authors propose that each user specifies a set of other trusted users, and a recursive propagation model is then used to compute a user's trust in all other connected members of the trust graph. This results in moderately *automated* trust ratings that are *individual* in nature, and therefore trust in the same entity may vary significantly between users. This user-centric model of trust is compatible with the perspective taken in this dissertation, as it gives a more personal view of the network. The approach of Richardson et al. does not support the specification of trust *topicality*, although this is raised by the authors as an area for future research.

Due to their statistical foundations, collaborative filtering systems require data sets of a significant size in order to perform at optimum levels. Massa and colleagues (Massa and Bhattacharjee, 2004, Massa and Avesani, 2004) use review and web of trust data from the reviewing site *Epinions3* to demonstrate how trust propagation techniques can be used to overcome the *cold-start*/*early-rater* and *sparsity* problems that can affect conventional collaborative filtering approaches.

The *cold-start* problem refers to situations in which items added to the catalogue of an e-commerce Web site can not be recommended using collaborative filtering until at least one customer purchases that item. Only at this point (and assuming that the customer already has purchases in common with other customers) can predictions be made of which other customers may be interested in the item. The extreme cold-start situation is that of a totally new recommender system where no data exists with which to correlate users or items.

Cold-start affects users in a similar fashion, as they must develop a profile that correlates them with other users (perhaps by rating or purchasing some items) before recommendations can be provided (Massa and Bhattacharjee, 2004). Early-rater problems (Dieberger, Dourish et al., 2000) describe one specific aspect of this situation, in which early adopters of a system gain little performance benefit in return for their input, as the system as a whole is not sufficiently populated with comparable users on which to base recommendations.

Sparsity is a measure of the degree to which items or users in a collaborative filtering system can be compared. Systems where users can on average be compared to a relatively low number of other users (due to a lack of overlap in profiles) are described as 'sparse', and will tend to provide lower quality recommendations (Massa and Avesani, 2004). These factors can all limit the ability to deploy recommender systems in settings where only small data sets are available on which to base recommendations.

Existing data from external sources is not commonly used to help bootstrap recommender systems. This is likely due to a lack of relevant data being available in an easily reusable form, either from the Web at large or from existing recommender systems. Issues such as privacy, data protection and maintaining competitive advantage reduce the incentives to share profile data, leading to duplication of effort by users who cannot benefit from using aggregate profiles of their own data across multiple recommender systems. If more data (such as reviews or broader profile information) were to be published online in an easily reusable form, this may provide a source of background data with which to bootstrap recommender systems, thereby reducing cold-start and sparsity issues.

Massa and colleagues (Massa and Bhattacharjee, 2004, Massa and Avesani, 2004) show that propagating trust through the network as a function of inverse network distance can provide systems with greater coverage of users and items on which to base recommendations, whilst keeping error relatively low. This is particularly useful when providing recommendations to new users who have not rated many items. Whilst these findings suggest there may be a role for this form of trust propagation, more sensitive trust metrics are required as the simplicity of the trust data on which it is based may be a limiting factor. See 2.7.6 for a full discussion of this issue.

In relation to the work of Golbeck and Hendler (but equally applicable to the related studies discussed above), O'Hara, Alani, Kalfoglou et al. (2004) observe that trust is not strictly transitive, and highlight this as a potential shortcoming of the work. This criticism applies to all the approaches described above that use trust propagation in order to compute metrics for indirectly connected (and therefore unknown) members of a social network. The results obtained by Golbeck and Hendler (2006) comparing their approach to collaborative filtering suggests that this may not significantly reduce the utility of the system in the context of film reviews. However, it may be that in domains less mediated by taste and where greater risk is involved, simple trust relationships such as these may not be so reliably propagated through an unknown network.

Personalized Relevance in Information-seeking through a Trusted Social Network

The Web has indisputably demonstrated its capabilities as an information sharing and dissemination platform. However, it is apparent that information-seeking applications on the Web would benefit from:

1) adopting more personalized notions of relevance

2) supporting a wider range of information-seeking tasks, which may vary in their characteristics

3) being sensitive to how variations in task characteristics may determine relevance

4) enabling greater involvement of the user's own knowledge in the information-seeking process

5) broadening their scope to include information that may not be available online

Social networks have long provided a powerful means for obtaining relevant and trustworthy information. This research proposes to address the shortcomings listed above by exploiting synergies between the Web and social networks. The outcome of the research will be approaches and systems that support information-seeking on the Web by harnessing the knowledge and experience of the user's social network, according to the principles of word-of-mouth recommendation. The aim is to increase personal relevance and facilitate greater use of trust, thereby improving the effectiveness of information-seeking and reducing information overload.

Numerous previous attempts have been made to support word-of-mouth in a Web environment, through, for example, collaborative filtering and online reviews. This research is not intended to replace these, but instead to develop complementary approaches and technologies that can overcome identified limitations in existing work. The factors outlined below distinguish this approach from previous work in the area.

Characteristics of this Approach

Source-centricity

In contrast to many of the search and recommendation approaches discussed before, this research takes a source-centric rather than item-centric approach to the information-seeking process; i.e. the emphasis is on identifying relevant sources before trying to identify relevant items.

The first challenge of this approach is *source identification*: finding out whom within a social network *knows* about topics relevant to the information need and therefore may be able to provide relevant information or recommendations. The second challenge is *source selection*: deciding which of these individuals to *trust* as sources of personally relevant information and recommendations. This research aims to develop approaches and systems that address both these challenges.

The reader may be interested to note that source identification and source selection can be seen as generalizations of McDonald and Ackerman's (1998) *expertise identification* and *expertise selection*. Regarding the concept of *trust*, many definitions have been proposed in the literature, of which Marsh (1994) provides a thorough review. For the purposes of this dissertation, and in the context of word-of-mouth recommendation-seeking, trust is defined here as '*confidence in another individual as a source of accurate and relevant information*'. This definition is deliberately neutral with respect to the source of evidence on which this confidence may be based.

Task-adaptivity

By definition, any information-seeking process must be aligned to the demands of the task by which it was originally motivated. This task will not only define the information need, but is also likely to have a number of other characteristics that will determine what constitutes an appropriate source of information or recommendations. This research aims to further understand these characteristics, and develop source identification and source selection processes that are sensitive and adaptive to them.

Social Networks and this Approach

The role of social networks in online environments, and online environments as reflections of social networks themselves, has received increasing attention in recent years. Garton, Haythornthwaite and Wellman (1997) emphasise the value of a social network perspective in the study of computer-mediated communication, and summarise some of the key units of analysis in the field of *social network analysis* (Scott, 2000).

Of particular relevance to the research presented here are the notions of *relations*, *ties* and *ego-centricity*. One or more *relations*, such as sharing information or being members of the same organization, create a *tie* (often classified as weak or strong) that connects a pair of actors. Research into the roles of strong and weak tie relationships is discussed in more detail later.

Garton et al. distinguish between *ego-centric* or *whole network* views of social networks. The ego-centric approach views the network from the perspective of a particular individual, whereas the whole network view considers an entire network comprised of individuals who meet a certain criterion. The former, ego-centric perspective on social networks is of greater relevance to this research.

Authors such as Mika (2004) have studied how information available on the Web reflects the structure of social networks in the offline world. By combining data harvested from the Semantic Web with conventional Web mining approaches, he is able to identify structural properties of the social network within the Semantic Web research community, such as various measures of each member's *centrality* within the network.

These metrics provide a basis for understanding some of the structural properties of a particular social network. As the research reported here is concerned primarily with the nature of one-to-one relationships in social networks, and the implications of these for information- and recommendation seeking, these measures of the structural attributes of social networks will not be considered in further detail.

In addition to ongoing work examining social networks themselves, whether online or offline, there has been an increasing interest in developing Web applications that include a social component. For example, the primary emphasis of sites such as *Facebook* and *LinkedIn* is in allowing people to express the connections in their social networks, forge new connections and engage in social interactions online.

In slight contrast, social annotation and bookmarking services, such as those summarized by Hammond, Hannay, Lund et al. (2005), allow individuals to store and annotate items for their own usage, but also share these resources with others through the social networking aspects of the sites.

Current trends in Web applications and ongoing research into social networks increase our understanding of the interaction between social factors and online environments, and provide a context for the research presented here. However, rather than looking at social networks purely from a structural/analytical viewpoint or from the perspective of technical applications, the research presented here requires a fuller understanding of how information and recommendations are sought within social networks, and the factors that shape this process.

Benefits of this Approach

Increased Personal Relevance

One fundamental premise of this approach is that members of one's social network are more likely to have knowledge relevant to one's own information needs than are people outside one's network.

Utility across a Range of Tasks

While similarity may provide a sound basis for increased personal relevance, the strength of this relationship is likely to vary according to the characteristics of the task that motivates the information-seeking, in which case additional factors will need to be taken into account in determining the relevance and trustworthiness of results. this research aims to be sensitive and adaptive to how peoples' information- and recommendation-seeking strategies may vary across tasks with different characteristics.

In cases where many potential information sources are identified within the user's social network, the approach presented here aims to help the user choose the most appropriate or trustworthy source of information given the characteristics of the information-seeking task. In doing so the aim is to be applicable and useful across a broader range of domains. This will be achieved by developing a detailed understanding of the source selection process in word-of-mouth recommendation.

Spam-resistant Information-seeking

A recent investigation (albeit journalistic, rather than scientific) (Walsh and Swinford, 2006) into 'review and rating spam' demonstrated how easily misleading reviews and ratings can be created on travel recommendation sites such as *TripAdvisor*, by those with a vested interest in promoting a particular establishment. The investigation suggested that this form of manipulation is widespread; consequently recommender systems that base recommendations on data that can be so easily falsified risk reducing the quality of their results (Josang, Ismail et al., 2007).

The use of social networks to support information-seeking makes the approach presented here less vulnerable to spamming, for the simple reason that each user is in the first instance only exposed to information or recommendations from people they know personally. This acts as a safeguard against manipulation of results, assuming that most users are unlikely to know others wishing to manipulate search indices on an ongoing basis, and at the expense of their acquaintances.

In the event that one individual persistently attempts to manipulate results, only those users who know the individual personally will be affected. These users will have the option of removing the individual from their social network (either virtually or in entirety!). The same benefits and safeguards do not apply to approaches based on social networks and trust propagation; by definition others beyond the immediate network will also be affected as trust relationships are propagated.

Openness to Additional Information

The approach presented here is oriented as much towards providing 'scaffolding' to support users in completing their information-seeking tasks, as it is toward providing solutions to their information needs. The aim is to augment rather than replace users' own assessments of members of their social networks as potential information sources. This is facilitated by the source-centricity of the approach and the use of social networks of known individuals.

Source Selection in Word-of-mouth Information-seeking

Background and Related Work

Word-of-mouth recommendation and referrals from others are powerful mechanisms for helping people acquire information and solve problems, in domains as diverse as finding piano teachers (Johnson Brown and Reingen, 1987) and successfully completing projects in the workplace (Cross, Parker, Prusak et al., 2001). Social networks of known individuals can serve as both a source of new information and as a filter to identify the information or items most relevant to one's specific needs (Granovetter, 1973),(Kautz, Selman et al., 1997a).

These processes have been extensively studied in a number of disciplines, particularly sociology, psychology, marketing and organizational sciences. In one of the earlier studies on the subject, Whyte (1954) provides an account of how interpersonal communication networks in local neighbourhoods can influence purchasing behaviour of domestic appliances. This study emphasised the existence of social networks that, through their role in information flow, can account for the non-random distribution of consumption patterns within the wider population. However, the work of Whyte (1954) was based on anecdotal evidence, and did not examine the nature of interpersonal relations between nodes in such networks or any effects these may have on the flow of information and subsequent purchasing decisions.

The Role of Weak Ties

When looking specifically at the relationship between the information seeker and an information source, one of the major themes in published work has been the notion of *strong* vs. *weak ties* in social networks, drawing on the work of Granovetter (1973). Whilst generally treated as discrete values of *strong*, *weak* or *absent*, *tie strength* is defined as a continuous variable stemming from a combination of amount of time, emotional intensity, intimacy and reciprocal services within a relationship. Importantly, it is posited that "the degree of overlap of two individuals' friendship networks varies directly with the strength of their tie to one another" (pp. 1360) (i.e. the stronger the tie between two individuals the greater the number of friends in common), and that a stronger tie correlates with greater similarity between two individuals.

Weak ties are considered more likely to act as 'bridges' between otherwise disconnected portions of the broader social network (supported empirically by Johnson Brown and Reingen, 1987). It is these weak ties that Granovetter found to play a key role in the diffusion of information to individuals who may not otherwise have been able to access it. Contrary to reasonable intuition, he found that weak rather than strong ties are more useful as sources of information about new jobs. This was attributed to the lower overlap between one's own social circle and those of others to whom one is weakly tied (i.e. a sufficient proportion of acquaintances were not shared). Consequently weak ties are more likely to be able to provide access to information about job opportunities that would be otherwise unavailable.

It is worth noting that Granovetter (1973) does not explicitly examine the way in which strong vs. weak ties affect the finding of a new job when elements of personal recommendation and referral are involved; the study is simply concerned with access to information about job vacancies.

Johnson Brown and Reingen (1987) identify a shortage of empirical evidence to support the importance of weak ties in communication flows in social networks. They argue that existing studies are insufficiently general, tending to focus on the role of weak ties in just one setting. Furthermore, they cite later work by Granovetter (1983) that highlights how the 'strength of weak ties' argument has often been used more as a post-hoc rationalization for empirical findings than as the focus of a systematic investigation.

The Role of Strong Ties

In addition to identifying shortcomings in the literature regarding the role of weak ties, Johnson Brown and Reingen (1987) also argue that there is potential for greater understanding of the role of strong ties in different aspects of word-of-mouth communication. The study they report seeks to provide empirical evidence for the 'strength of weak ties' argument of Granovetter (1973), whilst also examining the importance of strong ties in information-seeking and in influencing the decision-making of information recipients. Underpinning their work is a distinction between *relational form* and *relational content*. Relational form "refers to properties of the linkage between pairs of actors that exist independently of specific contents" ; tie strength is one of these properties that make up relational form. Word-of-mouth recommendation information is given as an example of relational content.

Johnson Brown and Reingen make a subtle distinction in their work between the activation of ties for the flow of information in general, and active information-seeking through ties. The former can be thought of as 'did information flow through this tie?'
whilst the latter can be conceptualized as 'was this tie actively sought out as an information source?'

From a study of word-of-mouth information flow regarding piano teachers in a metropolitan setting, Johnson Brown and Reingen found that: strong ties and ties between homophilous individuals (i.e. those who have characteristics in common) are more likely than weak or heterophilous ties to be activated for the flow of referral information.

However, the hypothesis that "active information-seeking is more likely to occur from strong-tie than weak tie sources of referrals" was not supported in the study. In fact, information was actively solicited from eighty six percent of weak ties used as sources, compared to active solicitation from only fifty percent of strong ties. This finding was attributed to the likelihood of incidental word-of-mouth communication increasing in line with communication frequency; therefore strong ties may be more likely to provide the required information in passing. It may be that where strong ties are unable to provide information in passing on a particular topic weak ties are actively sought instead.

Where referral information was provided by a strong tie it was perceived as more influential than referral information provided by weak ties. *Source credibility* is suggested as one explanation for the increase in perceived influence of information from strong ties, and a number of quotes are reported that suggest bases for this in factors such as *trusted opinions*, *valued recommendations*, and *knowledge of the field*. However, these factors are not investigated by Johnson Brown and Reingen, who do suggest that further investigation of how attributes such as credibility influence the choice of information source may complement the findings of relational analyses such as theirs.

Influences on Choice of Tie-Strength

Duhan, Johnson, Wilcox et al. (1997) investigate how attributes of the information seeker (*prior knowledge*) and the task (*difficulty*, role of *instrumental and affective evaluative cues*) impact upon the use of strong or weak ties as information sources. Their study used a scenario-based approach but focused solely on the domain of medical services, specifically the search for recommended obstetricians.

Duhan et al. found that the greater the perceived difficulty of the task, the greater the chance that strong-tie sources would be sought for recommendations; this finding supported their hypothesis of a positive relationship between task difficulty and the seeking of recommendations from strong ties. Contrary to another hypothesis, it was found that a greater importance of affective evaluative cues in decision-making did not correlate with a greater likelihood of seeking strongly-tied recommendation sources. However, as hypothesised, a greater importance of instrumental evaluative cues in decision-making was found to correlate with a greater likelihood of seeking weak ties for recommendations.

While the findings of Duhan et al. may appear to enhance our understanding of how task characteristics in particular impact upon the seeking of strong and weak ties as recommendation sources, their study has a number of limitations. The hypotheses investigated are based on a theoretical model formulated from previous research; however these hypotheses do not cover all possible relationships between factors present in the model, only certain relationships the authors predict to be of significance.

For example, the study predicts a relationship between task difficulty and recommendation-seeking from strong tie sources, but there is no comparable hypothesis testing a possible relationship between task difficulty and weak tie sources. In another example a positive relationship is predicted between the importance of instrumental cues and use of weak ties, without also examining possible relationships between instrumental cues and use of strong ties.

Consequently, it is not possible to conclude whether support for these latter two hypotheses was simply due to a greater chance of seeking recommendations *at all*, whether from weak or strong ties, as the design of the study is not sensitive to this. It is possible that other significant relationships exist that were not identified in the study but would invalidate the model. As a result, the study by Duhan et al. provides little evidence of the role of task characteristics in determining the use of strong or weak ties.

On this basis, it may be questioned whether relational form alone, and tie strength in particular, can provide an adequate, sufficiently granular, account of how people choose word-of-mouth information sources. In fact, attempting to explain source choice in terms of tie strength may represent a misapplication of the original research in this area. In Granovetter's (1973) work, tie strength is seen as a structural property that can influence information flow within networks, rather than a relational characteristic on which people base source selection decisions when actively seeking information. Consequently, tie strength may provide a rather blunt tool with which to understand source selection in information-seeking.

The Role of Source, Task and Individual Characteristics

A number of studies have moved beyond the broad strong/weak tie distinction and looked in more detail at the attributes of information sources that impact upon their selection by information seekers. Perhaps the largest body of work in this area concerns information-seeking within the workplace, from both human and non-human sources.

Workplace Studies

O'Reilly (1982) studied the frequency of use by welfare agency employees of a range of information sources, such as written documents, internal group members, and external sources. The impact of *source characteristics* (quality, accessibility), *task characteristics* (uncertainty, complexity) and *individual characteristics* (tenure, formal education, motivation) on frequency of use was investigated. In the context of this dissertation the most interesting findings relate to the source characteristics of quality and accessibility.

Accessibility of an information source was found to predict frequency of use for written documents (e.g. handbooks, procedures, memos, newsletters) and external sources but not human sources within the group. Further analysis found the frequency of use of group members to be a function of source quality, source accessibility, and the interaction between these factors. This interaction manifested itself in more frequent use of high quality, low accessibility sources than low quality, high accessibility sources, with sources of high quality and high accessibility being preferred.

O'Reilly acknowledges that *quality* is a subjective concept. He uses attributes such as *relevance*, *specificity*, *accuracy*, *reliability* and *timeliness* to define a more general notion of *information quality*, and it is at this higher level that the analysis is conducted.

Cross and Borgatti (2004) report a similar study that examined the impact of the source-seeker relationship on information-seeking. Through interviews with managers in a business consulting practice they identified four characteristics that were hypothesised to predict information-seeking: *awareness* of a potential source's expertise, *timely access* to the source, the *safety* of the relationship and willingness of the source to *cognitively engage* with the problem. A model based on these characteristics was then formulated and tested.

It was found that *awareness*, *timely access* and *engagement* were all predictors of source choice in information-seeking, however the same was not true for *safety*. These findings highlight that simply knowing who has knowledge or expertise in a topic is not sufficient in selecting an information source, as one must also be able to access a source who must also be willing to engage in problem solving. This study also provides some support for the findings of Borgatti and Cross (2003), as the *knowing/awareness* and *access* factors were found to be significant in both studies.

Morrison and Vancouver (2000) found that, in a sample of career-early aerospace engineers, *expertise* and *accessibility* of information sources both predict the likelihood of that source being used. Of these two factors, expertise was found to have the greatest impact. It is worth noting that the participants in Morrison and Vancouver's study were asked to rate information sources from a fixed list (supervisor, friend, colleague, mentor, documents) rather than sources they identified themselves. Despite this limitation, the results strongly support the findings of Borgatti and Cross (2003) and Cross and Borgatti (2004) relating to accessibility of information sources. The outcome related to perceived expertise of the source supports Borgatti and Cross's (2003) finding that perceived value of a source predicts use of that source.

Summary

The literature reviewed above provides indications of how the source selection process in information-seeking operates. A number of recurrent themes are present across the reviewed studies, such as the accessibility of sources and their perceived quality.

The diagram shown in Figure 3 below provides a representation of the information-seeking process from a source identification and source selection perspective, showing factors identified in studies reviewed before as having an impact on source selection in word-of-mouth information-seeking.

Figure 3. The information-seeking process from a source identification and source selection perspective.

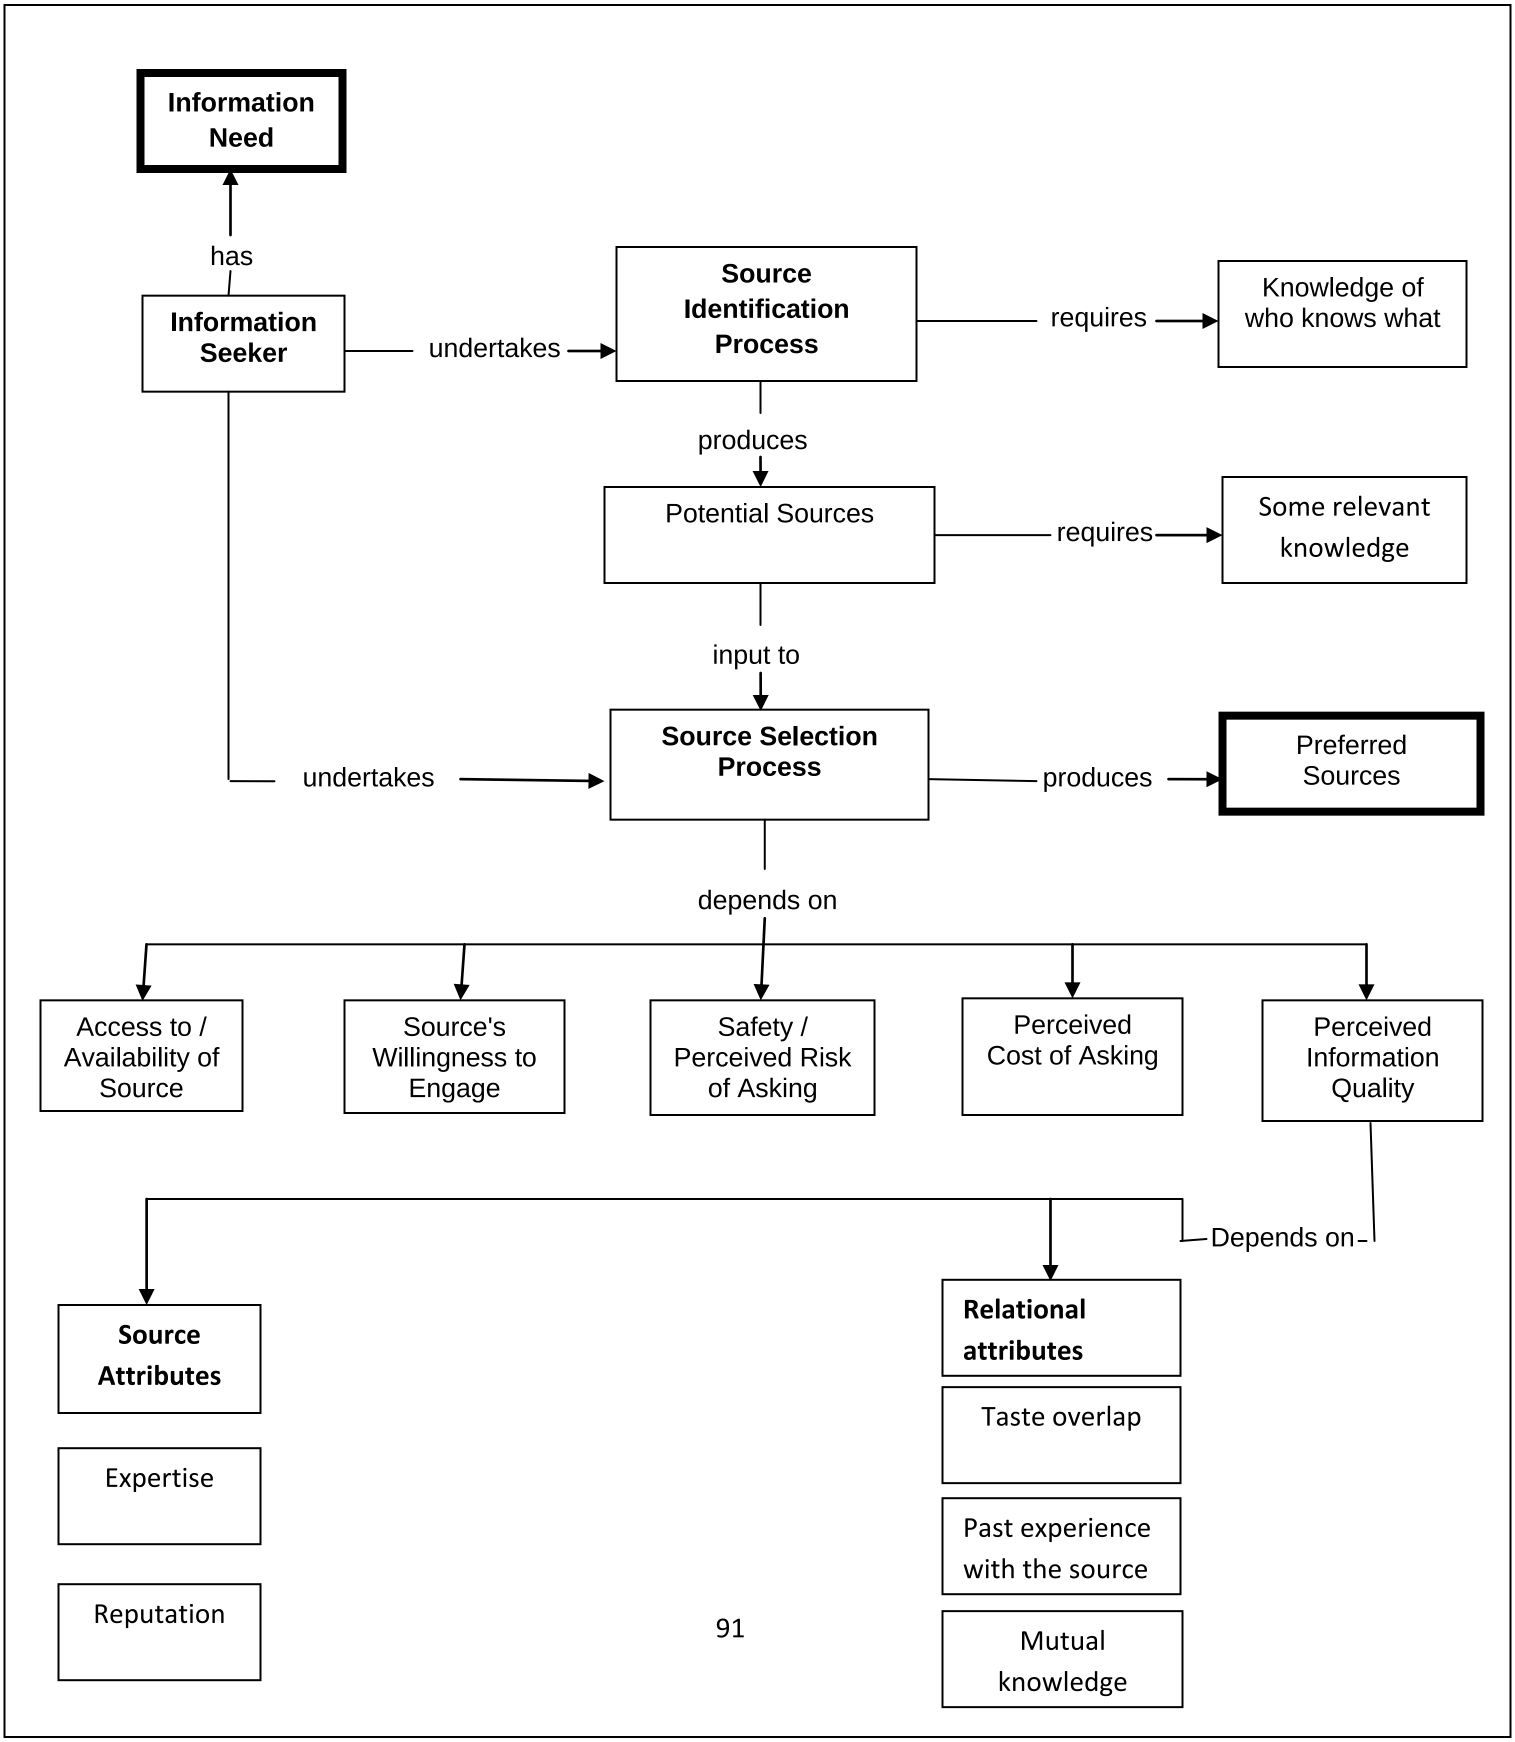

The studies in which the source and relational attributes shown were identified are shown in Table 3 below. Quality is included in this table to aid comparison between O'Reilly's work and that of others; however, Figure 3 above reflects the notion of quality as a higher-level construct that subsumes more specific factors.

	O'Reilly (1982)	Morrison and Vancouver (2000)	Bonhard and Sasse (2005)
Source Attributes			
"quality"	X		
expertise		x	x
Reputation			x
Relational Attributes			
taste overlap			x
past experience with the source			x
mutual knowledge			x

Table 3. Source and relational factors identified in existing literature as affecting perceived information quality

The literature on source selection in information-seeking is dominated by studies from workplace settings that deal primarily with information-seeking in job-related tasks. Studies investigating source selection in less informal and more taste oriented domains are less widespread. While Bonhard and Sasse's (2005) model does distinguish objective domains from taste domains, this factor is not systematically varied in the study on which the model is based and the findings remain oriented towards source selection in taste domains.

Overall, a picture of the source selection process does not emerge that is sufficiently consistent or generalisable to serve as a basis for implementing technical systems that support the selection of information sources within one's social network.

In order to establish some general principles from which the source selection process may be modelled, a further investigation is required that enhances our understanding of how people select information sources across a broader range of tasks, in domains not only mediated by taste. To address this need an empirical study was carried out to explore: from whom people seek information and recommendations in different scenarios; the factors that underlie their decisions about the trustworthiness of this information; and how the influence of these factors varies across different types of task.

Study of Source Selection in Word-of-mouth Information-seeking

This study addresses research Questions 1-3, introduced in previous pages:

1. How do people choose information and recommendation sources from among members of their social network?
2. Which factors influence judgements about the relevance and trustworthiness of these information and recommendation sources?

3. How do the characteristics of the task being performed affect these judgements?

Previous work in the area, as discussed above, does not provide a sufficiently comprehensive and consistent account of the information-seeking and recommendation-seeking process from which hypotheses can be derived and tested using quantitative methods. Therefore by necessity this study is exploratory in nature and qualitative in methodology. The aim is to identify central themes and factors in the decision-making process and gain insight into how the influence of these factors varies across different types of tasks, in order to identify general trends that can be operationalised in technical systems.

Design

The study consisted of semi-structured interviews in which participants were presented with a series of fictional recommendation-seeking scenarios and asked a number of open-ended questions exploring their decision-making process when selecting an information source.

Pilot

A pilot was conducted with three participants (who were not included in the main sample) to test the experimental protocol. This led to refinement of the interview script in order to ensure the results produced by the study would be sufficiently relevant to the research questions. In particular the open-ended questions used in the study were modified in order to be more structured, as the pilot had demonstrated that participants did not always understand how to respond to very open-ended questions.

Method

Participants

Twelve participants were recruited to the study using opportunistic sampling. Participation was voluntary, and no payment was received for taking part in the study. All participants were staff or students at The Open University, and varied in age from mid-20s to mid-50s. Seven participants were male and five were female.

Procedure

The study consisted of one semi-structured interview with each participant, on a one-to-one basis, in person. Interviews lasted between 16 and 60 minutes, varying according to the participant's engagement with the topic. After being given general instructions about how the interview would proceed, the participant was read in turn each of four hypothetical information- and recommendation-seeking scenarios (reproduced in Table 4 below) and asked to imagine themselves in this situation.

The scenarios used in the study were constructed by the researcher, and designed to closely represent everyday tasks and situations in which recommendations might be sought from members of one's social network. This contrasts with studies by authors such as O'Reilly (1982) where similar issues are investigated, but specifically in a workplace setting. It is not apparent how applicable such findings are outside that particular domain. The scenario-based approach bears some similarities to that used by Duhan, Johnson et al. (1997); however, in this case each participant was presented with multiple scenarios covering a range of domains, compared to Duhan et al's use of one scenario in a single domain.

Number	Domain	Text	Criticality	Modality
1	Plumber	"You move into a new house that requires renovation, including some substantial plumbing work. Who would you ask about recommended plumbers?"	High	Locating
2	Back Pain Treatments	"You are suffering from moderate and ongoing back pain and need to find some ways of getting it treated. Who would you ask about recommended ways of getting it treated?"	High	Exploring
3	Business Hotel	"You are travelling to Milan on business and need to find a hotel to stay in during your visit. Who would you ask about recommended hotels?"	Low	Locating
4	Holiday Activities	"You are planning a holiday to the east coast of the USA and need to find some information about how to spend your time there. Who would you ask about recommended activities?"	Low	Exploring

Table 4. Recommendation-seeking scenarios used in interviews with participants

The tasks described in the scenarios were varied along two dimensions: *task modality* and *task criticality*. Making up the task modality dimension, two of the scenarios (*plumber*, *business hotel*) described *locating* tasks, whilst two (*back pain*, *holiday activities*) described *exploring* tasks, as defined in Dzbor and Motta (2005). Locating tasks are those where the user is seeking a specific item or piece of information that is believed to exist, and the challenge is to identify an appropriate option or solution from among many. In contrast, exploring tasks are those where the user is attempting to develop a broad picture or understanding of a domain; the challenge in this case is to gather a representative range of perspectives from which later decisions may be taken.

Task criticality was defined as the degree of risk associated with a poorly chosen item or solution. This dimension was represented by two scenarios where the task was seen as low-criticality to the information seeker (*business hotel* and *holiday activities*), and two where the task was seen as highly critical (*plumber* and *back pain*).

The study was mindful of possible effects of domain (e.g. tourism, healthcare) and locality of task (for example, tasks based on information about the local area vs. information about distance locations), but these were not systematically varied in the study.

After being read each scenario, the participant was asked a series of questions, which can be paraphrased as:

* *From whom they would seek a recommendation?*

* *Was there anyone they would not ask?*

* *What were the reasons for these decisions*?

These questions made up a common script used by the experimenter , which provided a general structure for the interviews. This structure was broadly followed, however in line with the exploratory nature of the study deviation by participants was permitted in order to capture as rich an account of the decision making process as possible. Participants often provided lengthy responses which rendered later questions irrelevant, in which cases these questions were skipped by the experimenter. Asking participants if there was anyone they would not ask provided an opportunity for participants to elaborate on their source selection rationale, and often provided a richer picture of their decision-making process.

It was emphasised to each participant that there were no right or wrong answers to the questions asked by the interviewer, but simply that the research was interested in how they approach the problems presented in the interview.

Participants were not limited to specifying information sources within a certain proximity in their social network. Some did ask for clarification regarding whether they could cite sources not known to them personally, and some actively cited other sources such as the Web, however these cases were rare. Participants were also not constrained to citing sources with any particular tie-strength, as this was not a variable in the study. This allowed for examination of the salient properties of the information source or the interpersonal relationship as these impacted on the task in the scenario, without this being obscured by questions of tie-strength.

Participants were also asked to describe any analogous recommendation-seeking scenarios from their own experiences which came to mind in the course of the interview, and describe to their decision-making process on these occasions. Data from these accounts was included in the analysis.

Audio recordings of the interviews were made and transcribed to form the basis for the analysis.

Analysis

Following the methodology described in Smith (1995), inductive analysis of the transcripts was carried out to identify themes in respondents' decision-making. Each transcript was systematically analysed to identify factors that determined from whom respondents would seek recommendations. The factors identified across all transcripts were aggregated into a master list, from where they were grouped into a list of initial themes which was grouped again to produce the super-ordinate themes described below.

Results and Discussion

Five factors were identified that influenced participants' choice of sources for word-of-mouth recommendations, and the trust and confidence they had in information from these sources. Definitions of these factors are provided below, followed by frequency data and illustrative quotes taken from transcripts of the interviews. From now onwards these factors will be referred to as 'trust factors'. Factors related to practical aspects and diversity of responses were also raised, however, these were not included in the analysis as they do not relate to trust and relevance issues.

Trust Factors: Definitions

- **Expertise**: the source has relevant expertise of the domain of the recommendation-seeking; this may be formally validated through qualifications or acquired over time.
- **Experience**: the source has experience of solving similar scenarios in this domain, but without extensive expertise.
- **Impartiality**: the source does not have vested interests in a particular resolution to the scenario.
- **Affinity**: the source has characteristics in common with the recommendation seeker, such as shared tastes, standards, values, viewpoints, interests, or expectations.
- **Track Record**: the source has previously provided successful recommendations to the recommendation seeker.

Note that *expertise*, *experience* and *impartiality* relate to relationships between an information source and the topic of the recommendation-seeking (these are person → topic factors), whereas *affinity* and *track record* capture a relationship between the source and recommendation seeker (these are person → person factors).

Trust Factors: Illustrative Quotes

The following quotes from participants in the study illustrate the five trust factors:

Expertise

> "*I would probably go and ask my friend who is a plumber or my friend who is a gas fitter, working on the principle that their domain expertise, their knowledge, is in a similar area.*"

Quote 1. Participant ID 16, Plumber scenario

> "*Maybe I would immediately approach my doctor in the surgery where I'm registered, and ask his advice. …I wouldn't be confident that the advice is reliable…from the people who I don't know as specialists in the area.*"

Quote 2. Participant ID 10, Back Pain scenario

Experience

> "*I guess it depends on the location of the flat where I lived. If it was somewhere near to my parents I'd probably ask them first, for their advice, because they've got more experience, they've met people in the past who've done good jobs for them etc. etc.*"

Quote 3. Participant ID 05, Plumber scenario

"*People I know in the area, it's good to have word-of-mouth, you know they've got experience good or bad.*"

Quote 4. Participant ID 14, Plumber scenario

Impartiality

"*…with travel agents you'd have to question what they were promoting to you - is it because they get commission?*"

Quote 5. Participant ID 08, Holiday Activities scenario

"*Who wouldn't I ask? [I have] no specific examples. Actually its travel agents, as they're trying to sell you something; people who have no personal relationship to me and are interested in selling a product.*"

Quote 6. Participant ID 16, Holiday Activities scenario

Affinity

"*There is someone I would not ask [for] recommendations, who it would probably help to speak with... they have been to the States this summer and previous times... but ... because we're different persons she cares about different details than me... and adding to is that I don't think we have the same style in things we are after, so I wouldn't be urged to ask her advice.*"

Quote 7. Participant ID 17, Holiday Activities scenario

"*[I] may not ask people who I don't feel comfortable with, who haven't got the same values as me, or have a completely different lifestyle that I don't relate to.*"

Quote 8. Participant ID 12, Plumber scenario

Track Record

"*I looked on the internet yesterday about going to see a masseur, but they were too expensive so I'll go back to [ask] my sister as I had a good experience with [recommendations from] her before.*"

Quote 9. Participant ID 07, Back Pain scenario

"*Like the plumbing one [I wouldn't ask] someone who'd given me bad recommendations of hotels in the past.*"

Quote 10. Participant ID 16, Hotel scenario

Trust Factors: Occurrence Frequencies

Whilst the goal of the analysis was not to produce quantitative results for statistical analysis, it is useful to examine the frequencies of occurrence of the different trust factors in participants' explanations for choosing a particular recommendation source. As shown in following Figure 4 , expertise, experience, and affinity occurred most frequently, with relatively low occurrences of the impartiality and track record factors.

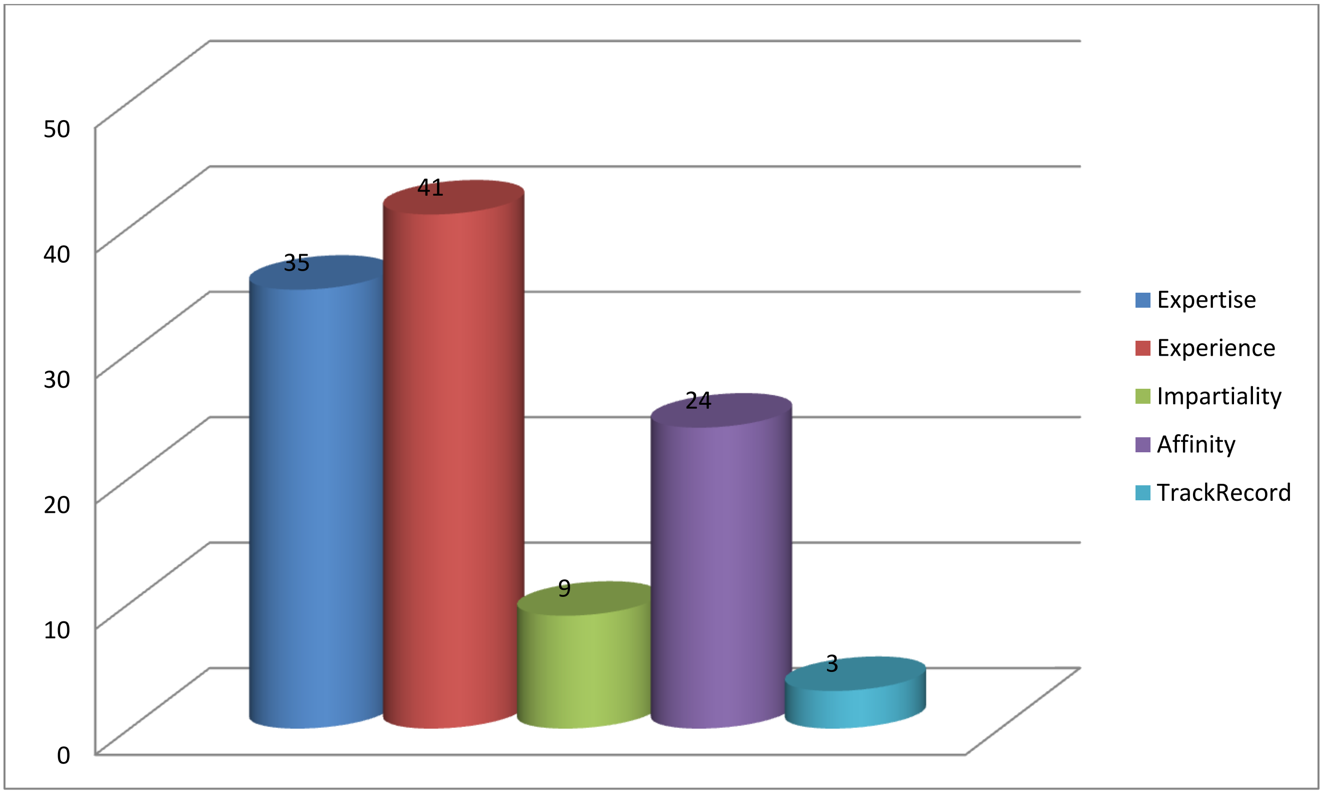

Figure 4: Response frequencies for each factor, summed across 4 scenarios in each of 12 interviews

Trust Topicality

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 more general construct. This may seem counter-intuitive at first as this aspect of affinity contrasts with taste, which is generally treated as a domain-specific characteristic. The relationship between affinity and taste is explored in the following section, along with a general discussion of how the findings relate to previous work in the area.

Relation of Trust Factors to Previous Work

Comparing the results of this study to the findings of previous research it is apparent that whilst some commonalities exist some novel trust factors have been identified. *Expertise* was identified as a factor in source selection by Morrison and Vancouver (2000) and Bonhard and Sasse (2005). Bonhard and Sasse also identified how *past experience* with a source can affect future use of that source for information-seeking, as can taste overlap between the information source and information seeker.

The affinity factor identified in this study appears to be relatively novel. One reason for this not having been previously identified may be that outside the formal roles and structures of the workplace there may be greater potential for exercising personal discretion in selection of sources, increasing the use of affinity relative to other trust factors. Furthermore, the tendency for existing studies to examine either taste domains or workplace expert finding may explain why the more universal notion of affinity has not been previously recognized.

It appears that affinity may be crucial where subjective recommendations are sought rather than simply factual information, a conclusion consistent with the findings of Bonhard and Sasse (2005) regarding taste domains. However, the data obtained in the study (e.g. Quotes 7 and 8 above) indicated that affinity represents more than simply shared tastes and is in fact domain-independent. In addition to style and taste, affinity appears to encompass more universal traits such as similar outlooks on life, values, lifestyle, expectations and attention to detail. Whilst affinity and taste are no doubt related in some way, the results of this study suggest that they are not interchangeable. In fact, shared tastes would appear to be one sub-component of the broader notion of affinity, which can be thought of as 'taste++'.

The study reported here did not identify a specific role for mutual knowledge or reputation, both of which were identified by Bonhard and Sasse. However, Bonhard and Sasse do not adequately define the concept of reputation, which may simply reflect a personal or social perception of the quality of information from a particular source. Reputation may in fact represent an aggregate measure of factors identified in this study, particularly expertise, experience, impartiality and track record. Any role of affinity would likely depend on whether a personal or group level definition of reputation was adopted.

In contrast to previous research, the study reported here identified relevant experience as a key factor in determining the trustworthiness of or confidence in an information source, along with the source's impartiality with respect to the domain of the task.

General Trends in Application of Trust Factors

Individuals did vary in the source selection strategies they reported, however some general trends emerged, most significantly that the emphasis given to each of the trust factors varied according to the characteristics of the recommendation-seeking task. These trends are examined in the following sections.

Effects of Criticality, Subjectivity

In tasks perceived as highly *critical* (e.g. the back pain scenario), emphasis was placed on externally validated '*expertise*', as illustrated by Quote 1 and Quote 2 above. This finding is consistent with the claims of Dieberger, Dourish et al. (2000) that "some domains depend more heavily on expert recommendations" . In less critical tasks respondents were less selective. Some participants indicated a particular willingness to seek information from a broad range of sources in less critical situations, on the basis that information from less trusted sources could be filtered or disregarded later if necessary, as illustrated by the quote below:

> *"My view is I gather everything from everybody and filter it, so I wouldn't be averse to asking people who maybe wouldn't like the same holiday, I'd still be prepared to take on board what they recommended, because I'd then filter it out, rather than not taking it."*

Quote 11. Participant ID 12, Holiday Activities scenario

Where tasks were perceived to have an objectively correct solution, respondents also widely cited '*expertise*' or '*experience*' of the recommender as influencing their choice. However, where suitable solutions were more *subjective* (such as in the holiday activities scenario), respondents emphasized the '*affinity*' factor. Some participants indicated that they would reject sources with highly relevant experience if there was not an affinity between themselves and that source, as illustrated by Quote 7 above.

These results suggest that the criticality of the task and the subjectivity of possible solutions were of primary importance in determining which trust factors were emphasized. 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.

Effects of Task Modality

Effects of task modality (i.e. locating vs. exploring) were not readily apparent in the data. This may indicate that sources are chosen in the same way irrespective of modality. However, it is also possible that variation in criticality of the tasks and subjectivity of solutions masked any such effects in this study.

Domain of Task and Nature of Relationship

Respondents indicated that they would choose information sources with '*expertise*' or '*experience*' appropriate to the domain of the task (e.g. a doctor in the back pain scenario). However, any variation in how the trust factors are employed across domains such as tourism and healthcare is attributable to factors such as the criticality and subjectivity of the task, not to differences in strategy that are specific to particular domains.

Close family and friends were often cited as sources. Whilst trust factors such as 'affinity' and 'track record' likely contribute to this finding, it is also probable that respondents cited these sources for practical reasons; they are easily accessible, and the seeker can better assess their suitability to give recommendations in a particular domain. The precise nature of the relationship between respondent and the source they chose did not appear of great importance. Practical factors such as the source being a gatekeeper to others (as a family doctor may be), and the social acceptability of asking someone were also mentioned.

Conclusions

This part reports on an empirical study examining how people seek recommendations from members of their social networks, across a range of scenarios. The study demonstrates that people make detailed and complex decisions when identifying sources of recommendations, and assessing the trustworthiness of such sources. Furthermore, these decisions take into account a detailed knowledge of potential recommendation sources.

Analysis of the data identified five factors that influenced from whom participants would seek recommendations, and how trustworthy these sources would be perceived to be: *expertise*, *experience*, *impartiality*, *affinity*, *track record*.

The specific factors on which source selection decisions were based varied according to the characteristics of the task. In particular the criticality and subjectivity of the task were found to influence the factors most attended to in a given scenario.

Whilst providing support for a number of findings from existing research, the results of this study make a number of novel contributions: they provide results that may generalize more readily, as a range of scenarios were used beyond purely workplace or taste domains, and these were supplemented by participants own accounts; they expand upon previous research by identifying new factors that influence source selection, thereby further unpacking the notion of source quality.

These findings address Research Questions 1-3, by identifying how people choose information and recommendation sources from among their social network, the factors that influence judgments of the relevance and trustworthiness of these sources, and how source selection decisions vary according to the characteristics of the task.

Data Acquisition from Distributed Sources

In order to compute experience, expertise and affinity trust metrics two basic types of data are required: data that connects people to domains or topics, from which experience and expertise metrics can be computed; and data that connects people to other people, from which affinity metrics can be computed.

The APIs of many so-called 'Web2.0' O'Reilly (2005) services such as *Amazon*, *Del.icio.us*, *Flickr* and *Face book* provide data that may address some of these requirements. For example, keyword tags that people have used to annotate photos or bookmarks may indicate domains in which they have experience, whilst reviews of items on *Amazon* may provide a basis for computing affinity scores between users using collaborative filtering-style approaches.

Some use is made of data from these services, such as tagging data from the social bookmarking site *Del.icio.us*. Because tagging is unconstrained in the terms that can be used, tagging data has the potential to provide evidence of an individual's experience across an infinite number of domains, which would not be possible if a fixed topic list and manual ratings were used.

However, the data available through services such as *Del.icio.us* and *Amazon* is limited in a number of ways that affects its utility in this research, particularly in the extent to which reviews, tags and social network data can be integrated.

For example, Code Fragment 1 shows anonymised review and user data retrieved from the *Amazon Associates Web Service API* in response to a CustomerContentLookup operation. The operation takes an *Amazon* CustomerID as input and in this case returns all information the customer has made public about themselves (the Customer Full response group was requested).

CustomerIDs can be obtained by querying the API for reviews of a known item; these CustomerIDs can then be used in CustomerContentLookup operations to obtain additional data about the user. As Code Fragment 1 shows, personal information such as name, location and nickname are returned in such queries. Occasionally a user's email address is used as the value of the nickname field, however this is not consistent and in most cases no data is available that can be used to uniquely identify a user as a basis for integration with other types of data from external data sources, such as social network information.


```
<Customer>
<CustomerId>A89YPC0B3HML7X</CustomerId>
<Nickname>joebloggs</Nickname>
<WishListId>7DCW9CVSFW7RI</WishListId>
<Location>
<UserDefinedLocation>Bloggsville, Arizona, United
States</UserDefinedLocation>
</Location>
<CustomerReviews>
```

```xml
<TotalReviews>18</TotalReviews>
<TotalReviewPages>2</TotalReviewPages>
<Review>
<ASIN>1234567890</ASIN>
<Rating>4</Rating>
<HelpfulVotes>2</HelpfulVotes>
<Reviewer>
<CustomerId>A89YPC0B3HML7X </CustomerId>
<Name>Joe Bloggs</Name>
<Location>Bloggsville, Arizona, United States</Location>
</Reviewer>
<TotalVotes>2</TotalVotes>
<Date>1998-08-29</Date>
<Summary>A great account of a tricky situation</Summary>
<Content>In this witty book author Joe Bloggs recounts the
challenges of being given one of the world's most common names.
</Content>
</Review>
</CustomerReviews>
</Customer>
```

Code Fragment 1. Example of the structure of a CustomerFull Response Group to a CustomerContentLookup Operation on the Amazon Associates Web Service API.

Resource Description Framework (RDF)

Data on the Semantic Web is not published as tables or lists in HTML documents, but as '*triples*' according to the '*Resource Description Framework*' (Klyne and Carroll, 2004). RDF defines both a graph-based data model based on *subject*, *predicate*, *object* triples, and the RDF/XML format (Beckett, 2004) through which an RDF graph of one or more triples can be serialized as an XML document14. The subject of any RDF triple must be a URI or a 'blank node', the predicate must be a URI, and the object can be either a URI, a Literal or a blank node (Klyne and Carroll, 2004).

Publishing data in RDF conveys a number of benefits: data is machine-readable, easily integrated for querying or other forms of processing, and easily linked across disparate sources. Traditional data formats such as Comma Separated Variables (CSV), 'vanilla' XML and even HTML can all be described as machine-readable, as data can be represented in these formats and parsed reliably by software applications. However, data represented in RDF is machine-readable in a different way. Not only is it machine-readable at a syntactic level (i.e. it can be parsed reliably) but also at a semantic level, in that the meaning of RDF data is made explicit.

Sources of Social Network Data

Many potential sources of social network data exist on the Web, particularly in social networking sites such as *Facebook* and *LinkedIn*. However, as discussed above and despite the availability of the *Facebook Platform*, the data held by these sites is not published in formats that afford easy integrating and linking with data from other sources. Consequently this research uses social network data published in RDF using the 'Friend of a Friend' (FOAF) vocabulary (Brickley and Miller, 2007).

Taking an RDF-based approach affords users greater choice and flexibility in how their personal information is managed and published, as data can be made available in locations of their choosing and under their control, from where it can be shared with third party applications.

The FOAF vocabulary provides properties and classes for describing common features of people and their social networks. The basic unit for defining social relationships in FOAF is the knows property, simply used to state that Person A knows Person B. This degree of semantics is sufficient for many application scenarios, and avoids potentially awkward social situations arising from individuals having different perceptions of the nature of a relationship.

Other vocabularies, such as the Relationship vocabulary (Davis and Vitiello, 2005), have been proposed that go beyond the shallow semantics of foaf:knows to describe greater subtleties in the relationships between individuals. The greater specificity provided by such vocabularies may be beneficial for certain applications, but is unlikely to enhance this research as it is not apparent how different relationship types may predict trust relationships between individuals in the domains with which this research is concerned.

REVYU : A Semantic Web Reviewing and Rating Site

Introduction

Revyu was developed to enable the collection of data from which trust metrics could be derived and integrated with social network data. Revyu is a reviewing and rating site in the mould often associated with Web2.0 but which has been present on the Web for some time. Prominent examples of such sites include *Epinions* and the reviewing functionality of *Amazon*.

Revyu was launched as a live, publicly accessible Web site at http://revyu.com/ in November 2006. As of November 2007 more than 650 reviews have been created by more than 150 reviewers. The reviews in the system cover a range of types of items including books, films, concerts, hotels, restaurants and academic papers. The Revyu homepage is shown in Figure 5 below.

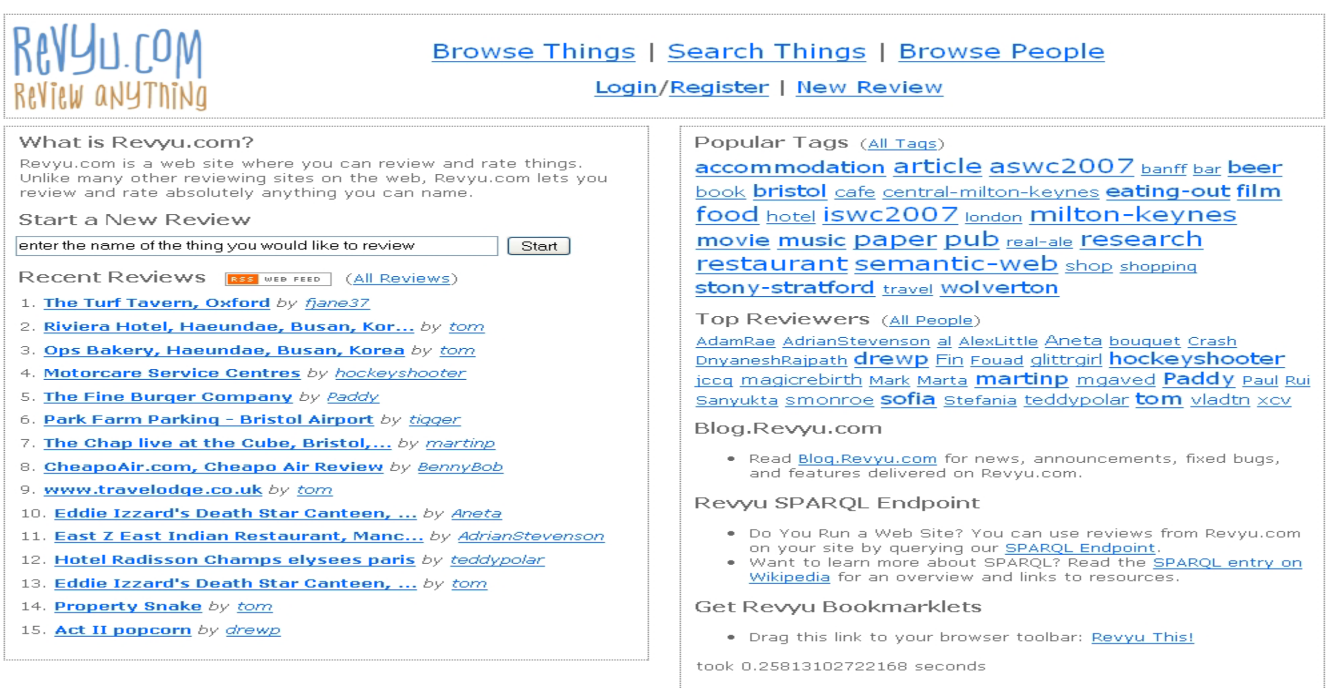

Novel Features of Revyu

Revyu differs from existing Web-based review and rating systems in a number of significant ways. Firstly, users of the site are not restricted by the closed worlds of conventional reviewing sites that limit reviews to items from a specific domain, sold by a particular company, or catalogued in an existing database. Instead Revyu takes a more open-world approach where users are free to review anything they choose. In addition to giving the user flexibility this has the benefit of not requiring a database to be maintained of items suitable for review, as is the case with existing cross-domain review sites such Secondly, reviewing sites that provide data for reuse via an API are not widespread. As a result, sites such as *Epinions* and *TripAdvisor* become closed world silos of reviews available on the Web but not well interlinked with other relevant data. Even where APIs are provided, by *Amazon* for example, these reviews are generally made available in formats such as XML that do not afford interlinking at the data level. This hinders the interlinking and aggregation of all reviews of a particular item from across the Web, because without the use of universal identifiers such as URIs it is not easy to determine if two reviews refer to the same item.

To overcome these issues, Revyu is built natively on Semantic Web technologies. As a result, the site identifies reviews (and all other types of objects in the system) with URIs and exposes these on the Web in RDF according to the principles of Linked Data (Berners-Lee, 2007), and via a SPARQL endpoint. This enables reuse of data from Revyu in third party applications, more flexible querying via SPARQL, and easier integration and linking of data across different sources.

Thirdly, Revyu exploits this ease of data integration to enhance the site with data from external sources without requiring this data to be replicated at Revyu.

Lastly, the majority of conventional reviewing and rating sites only identify reviewers by nicknames or unique identifiers that have only local rather than global scope. As a result one can rarely base decisions about the trustworthiness or value of a review on pre- existing knowledge of the reviewers, as nicknames obscure their true identity and prevent one from identifying all reviews by known and trusted individuals. Instead, characteristics such as writing style must be relied upon in judging the suitability or trustworthiness of a review.

To overcome this and enable integration of reviews with social network data Revyu includes a SHA1 hash (Eastlake and Jones, 2001) of the reviewer's mailbox URI in its RDF output of reviews, using the mbox_sha1sum property from the FOAF vocabulary. This serves to uniquely identify a reviewer without disclosing his identity to those who do not already know his email address, as the SHA1 algorithm makes it "computationally infeasible to find a message which corresponds to a given message digest, or to find two different messages which produce the same message digest" (Eastlake and Jones, 2001) .

User Walkthrough

Users can search or browse the site to read existing reviews, descriptions of things reviewed on the site, and profiles of reviewers. To the non-specialist Revyu appears like any regular Web site: little indication is given that it is based on Semantic Web technologies. All site content is published in HTML and RDF/XML, however users viewing the site with a conventional Web browser will never be exposed to the underlying RDF data unless they explicitly request it, either by clicking a link in HTML pages on the site or by sending appropriate Accept headers in their HTTP request . Figure 6 shows a review created on Revyu, as it appears in a conventional Web browser.

Figure 6 : The HTML view of a review on Revyu

Generating Semantic Web Content by Completing Web Forms

Users who wish to create reviews and ratings can do so simply by registering with the site and filling in a Web form as shown in Figure 7 . The reviewing form can be accessed by following a link on the Revyu site or using the Revyu 'bookmarklet', a browser widget that redirects the user from the site they are currently viewing to the reviewing page on Revyu; this can be helpful where the user wants to review a certain Web page or a thing described by the Web page, as a relationship between the reviewed item and the origin Web page is recorded by Revyu .

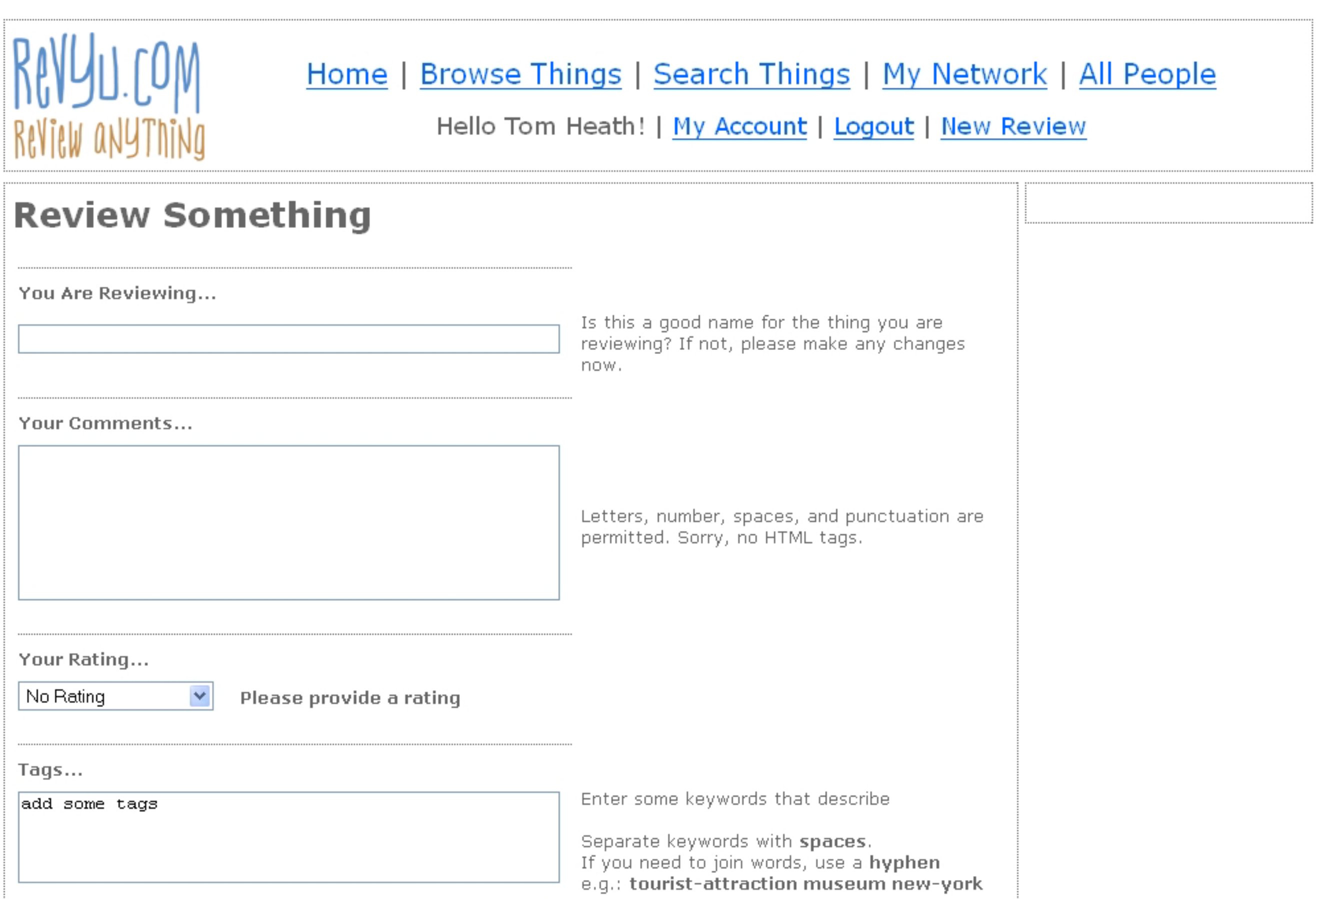

Figure 7 : The upper half of the Revyu Reviewing Form

The Revyu reviewing form in Figure 7 simply asks users to provide a name for the thing they wish to review, the text of their review, a numerical rating (on a scale of 1-5, where 1 represents Very Bad, and 5 represents Very Good), some keyword tags related to the thing being reviewed, and one or more links to related Web resources.

This mode of interacting will be familiar to those who have written reviews at sites such as *Epinions* or *Amazon*, and is designed to enable novice users to contribute reviews through a Web2.0-style interface, but make these reviews available online in the appropriate Semantic Web format.

Web2.0 applications and services such as *Wikipedia*, *Flickr* and *Del.icio.us* have enabled non-specialist users to contribute to the Web on a scale that is inline with the original vision of a 'read-write Web' (Berners-Lee and Fischetti, 2000), but had not previously been achieved. This has been made possible by providing simple, well-structured interfaces based on Web forms, through which users can, for example, edit wiki entries or tag photos and bookmarks. Such interfaces lower the cost of adding content and annotations to the Web compared to traditional publishing techniques that involve specialist skills and software.

Following a similar approach, Revyu is designed to be usable by humans whilst transparently generating machine-readable RDF metadata based on their input. By adhering to this well established interaction pattern, Revyu allows users to create Semantic Web data that can be used in computing trust metrics for this research, without requiring any knowledge of RDF.

In an evaluation of Semantic Web applications deployed to members of the Semantic Web community (Heath, Domingue and Shabajee, 2006) it was found that the usability of applications hindered their uptake, even by those knowledgeable in the field. In the light of these findings, tools that make semantic annotation feasible for specialists and non-specialists alike are required if user-generated Semantic Web data is to be created on a significant scale.

To date users of Revyu have created over 20,000 RDF triples which are publicly available on the Semantic Web. Whilst not a large figure by some standards, it is significant that these triples have been generated primarily from direct user input, rather than by data mining or extraction from natural language.

Reviews submitted through the reviewing form are converted to RDF and stored as persistent triples in the Revyu triplestore . From there they are immediately available on the site in HTML and RDF/XML, and via the Revyu SPARQL endpoint.

The Role of Tagging in Revyu

Tagging versus Classification

A decision was made when designing and implementing Revyu to not require users to classify reviewed items according to an existing taxonomy, but instead allow them to tag with one or more descriptive keywords an item being reviewed. This decision was made for both user-oriented and implementation-related reasons: classifying reviewed items would require the user to identify an appropriate category in an existing, fixed taxonomy to which not all reviewers could subscribe. Furthermore, if users were to be given complete flexibility in what they reviewed then such a classification would by definition be large and therefore complex. A sufficiently comprehensive classification was not readily available, and even the entire range of ontologies available on the Web were not seen to provide adequate coverage of all types of items that users might wish to review. Even were this was not the case, developing a sufficiently usable interface with which users could easily categorize any item was As a result, keyword tagging was chosen in favour of classification, as this can aid other users of the site in browsing or searching for reviews, whilst not creating barriers to the contribution of reviews and allowing for reviewing of items that might be not be easily categorized but can be described with a few keywords.

When users start entering tags in the Tags field of the Revyu reviewing form, suggestions are displayed of tags they may want to use based on those already present in the system. This helps avoid spelling mistakes, aids convergence on particular syntactic forms, and ensures consistency of tag usage.

A less desirable consequence of the use of tagging in Revyu is that machine-readable statements regarding the nature of reviewed items cannot be made with any confidence from tagging data alone. For example, the tag book not may refer to a volume of reading material but to a service for booking concert tickets. Similarly, an item tagged film may not be a movie film but a particular brand of photographic film. Therefore, by default Revyu makes no assumptions about the type of reviewed items based on how they have been tagged.

By allowing less structured input from users the burden of identifying the 'type' of reviewed items is transferred to Revyu if the site is to provide additional functionality based on this information. Derivation of type information from tagging data is currently undertaken in two domains, books and films, using external data sources to help ensure accurate results. Similar heuristics may feasibly be implemented for items such as music albums, pubs, restaurants and hotels.

Inferring the Type of Reviewed Items

Identifying Films on Revyu

The majority of contemporary films have homepages, which are generally provided by the film studio but carry little if any machine-readable data about the picture. However, coverage of films is very high in *Wikipedia*, which provides an external source against which Revyu data can be verified by querying the *DBpedia* (Auer, Bizer, Lehmann et al., 2007) SPARQL endpoint. The following heuristic is used to identify films: for each reviewed item tagged 'film' or 'movie', look for items in *DBpedia* of type 'film' and with the same name. For any items for which this heuristic returns a match, an rdf:type statement is added to the Revyu triplestore asserting that this item is a film. This type information is exposed in the RDF descriptions of items on the Revyu site and also used to trigger retrieval of additional information about the reviewed item for display on the site.

Identifying Books on Revyu

Whilst *Wikipedia* (and thus *DBpedia*) has extensive coverage of films, the coverage of books is less comprehensive; therefore a different heuristic is used to identify books reviewed on Revyu. When reviewing books, reviewers often place links to an *Amazon* Web page about the book in one of the Links fields of the reviewing form (generally the Other Links' field, as described below).

Where these links exist they are parsed and analyzed to extract ISBN numbers. If a valid ISBN is identified then an rdf:type statement is added to the Revyu triplestore asserting that this item is a book. Again, this type of information is used to retrieve additional information about the item, also as described below. Parsing links to external resources in this way is preferred over simply looking up all items tagged 'book', due to the potential for books and other items with the same name to cause false positives.

Identifying Related Tags

Many tags are used together when reviewing items, presumably because they are related in some way. An algorithm is used to identify tags that frequently co-occur (above a certain threshold of co-occurrence, to avoid identifying spurious connections) from tagging data in Revyu. For example, the algorithm finds that 'pub' is related to 'beer' and 'food'.

These relations are then logged in the Revyu triplestore and republished in both HTML and RDF. In the HTML pages about each tag, tags that co-occur above a certain threshold are displayed to the user. This threshold is set low for HTML output, as human readers of the page are unlikely to infer erroneous information based on these relationships. The RDF output uses the skos:related property of the SKOS vocabulary (Miles and Brickley, 2005), asserting that these two concepts are related. This makes these conceptual relationships accessible to other applications wishing to find information about connections between tags. In contrast to the HTML output, relationships exposed in RDF descriptions of tags are based on a more conservative threshold, in order to avoid erroneous inferences based on these assertions.

Finding co-occurrence relationships between tags is certainly not unique to Revyu; what makes this work more noteworthy is the republishing of these relationships to the Web in RDF. At present no attempt is made to link tags to other concepts in e.g. WordNet (van Assem, Gangemi and Schreiber, 2006), as sufficient accuracy cannot be guaranteed, especially when dealing with homonyms. However, techniques described by Specia and Motta (2007) suggest how Revyu tags may be better integrated with the Semantic Web.

The Role of Links in Identifying Reviewed Items

As discussed earlier in this section, Revyu takes an open world view of the reviewing process by not constraining users to reviewing items from a fixed database; anything that the user can name can be reviewed. This has the potential to create a situation where an item has been reviewed, but the exact 'identity' of the item is not apparent from the content of the review. To minimize the occurrence of such situations Revyu allows reviewers to specify a number of links that are associated with the item being reviewed in one of three ways: the home page of the item, a page that contains additional information about the item but is not the home page, or the actual location of the item where it exists on the Web Figure 8 below shows the three Link fields on the Revyu reviewing form.

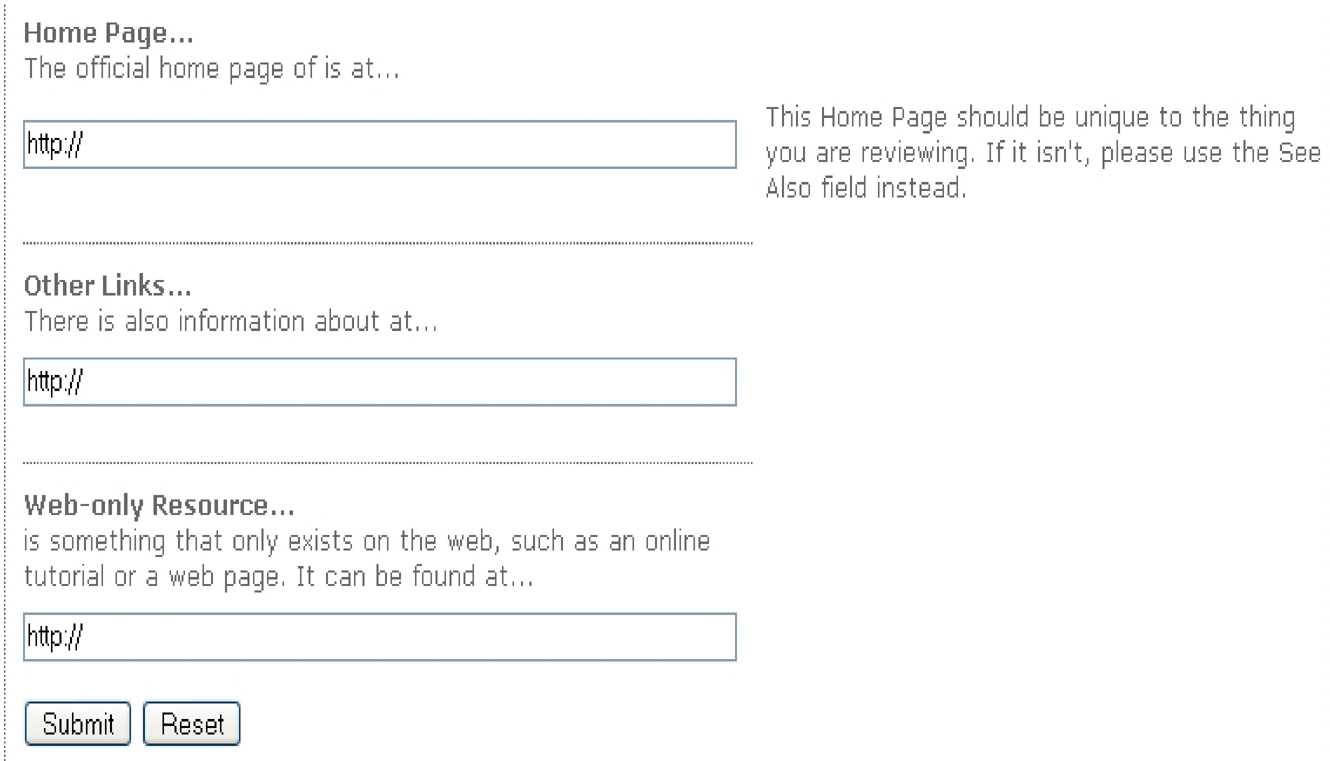

Figure 8 : The lower half of the Revyu reviewing form

These external links provide a way for human users of the site to disambiguate reviewed items in cases where there is any ambiguity. Disambiguation can also be carried out by applications that use Revyu's machine-readable RDF output, as the contents of these fields are saved as RDF triples when the review is submitted, using the foaf:homepage, rdfs:seeAlso and owl:sameAs predicates respectively.

The owl:sameAs property indicates that two URIs identify the same item, thereby linking a thing's representation on Revyu to its true location on the Web. RDF-aware users can also enter URIs that represent things other than Web documents ('non-information resources') into this 'Location' field in order to link Revyu-generated URIs to equivalent URIs minted by other data providers.

Links made using rdfs:seeAlso are of less value for these purposes, however the homepage property is defined in the FOAF ontology as 'Inverse Functional', meaning that the object of a foaf:homepage triple uniquely identifies the subject of the triple. Consequently it can be inferred that two resources that have the same foaf:homepage are in fact the same resource. This feature opens up the possibility of using Semantic Web lookup services such as *Sindice* (Tummarello, Oren and Delbru, 2007) to identify other sources of information about items reviewed on Revyu.

Links to other Data Sets

Where possible, links are made between Revyu data and items in external data sets (see Figure 9) in order to avoid Revyu data becoming an isolated island of RDF. Publishing these links in RDF connects Revyu in to a growing Web of Linked Data that is signified in particular by initiatives such as the Linking Open Data community project (Bizer, Heath, Ayers et al., 2007).

Many of these links are created during the same processes described above that attempt to derive type information from tagging data by validating against external sources. For example, where a reviewed film or book is found to exist in *DBpedia* or the *RDF Book Mashup* (Bizer, Cyganiak and Gauss, 2007), owl:sameAs statements are added to the Revyu triplestore to record that both URIs identify the same item. Likewise, where a user provides the URI of their FOAF file at registration time, owl:sameAs statements are made between the reviewer's Revyu URI and the URI they use to identify themselves in their FOAF description. These statements are then republished in the reviewer's RDF description on Revyu.

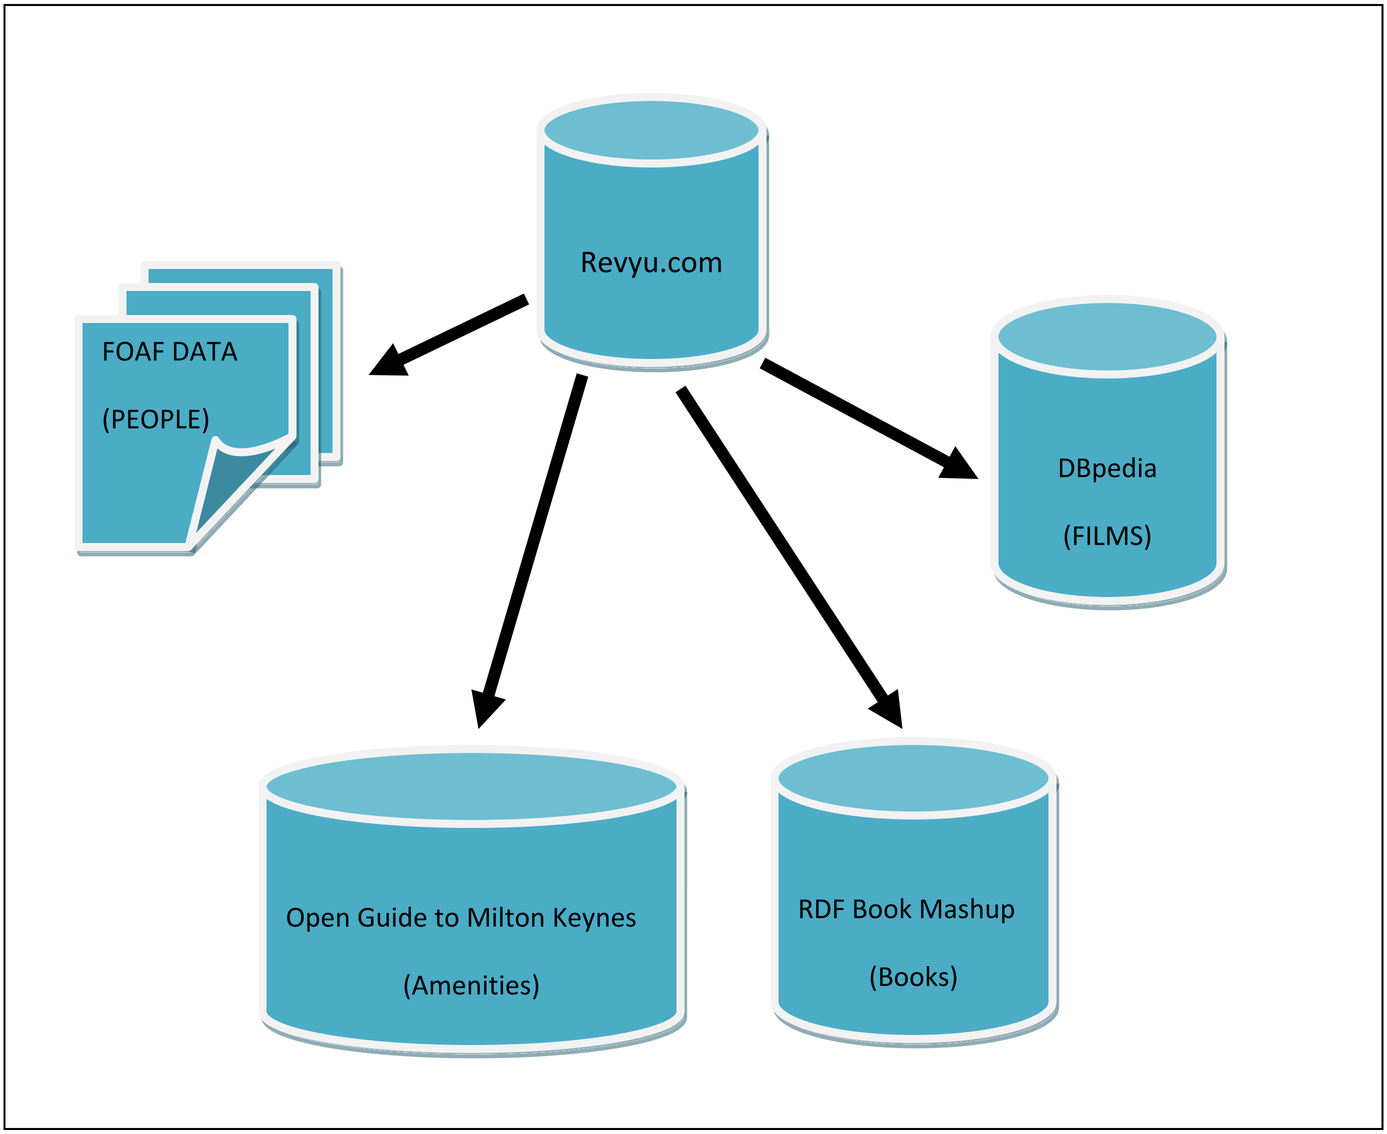

Figure 9 : Links from Revyu to external data sets

Consuming Linked Data

Links between Revyu and external data sources are used as the basis for retrieving additional information about reviewed items from external Semantic Web data sources, without requiring the reviewer to provide this information. This information is shown alongside review data from Revyu in the HTML pages about an item, thereby enhancing the experience provided to users of the site without placing an additional burden on reviewers.

Supplementing Reviewer Information with FOAF Data

Users registering with the site are not asked to provide copious information to populate their user profile, only an email address, screenname and password (real name can optionally be provided). Instead, where a reviewer maintains their own RDF (i.e. FOAF) description in another location they may also provide its URL. In this case Revyu dereferences this URI and queries the resulting graph for relevant information the reviewer chooses to share about themselves, such as photographs, homepage links, interests, and locations. This information is then used to enhance the reviewer's (HTML) profile page (figure 10) , thereby exploiting the data integration capabilities of a Semantic Web to provide the kind of rich user profiles often associated with Web2.0 applications without the information needing to be duplicated in Revyu. This approach reduces the burden on the user by not requiring them to manage multiple redundant sets of personal information stored in different locations, as one central set of personal information can be maintained in their FOAF file.

Figure 10 : The author's Revyu profile page, showing review data from the site (left) alongside information from his external FOAF file (right)

In addition, where a user knows another reviewer they can choose to add this person to their social network (as recorded on Revyu). This relationship is then recorded in the triplestore using the foaf:knows property. All such triples are exposed in the user's RDF description on the site, allowing them to be combined with other FOAF data from the Web to provide an integrated definition of the user's social network.

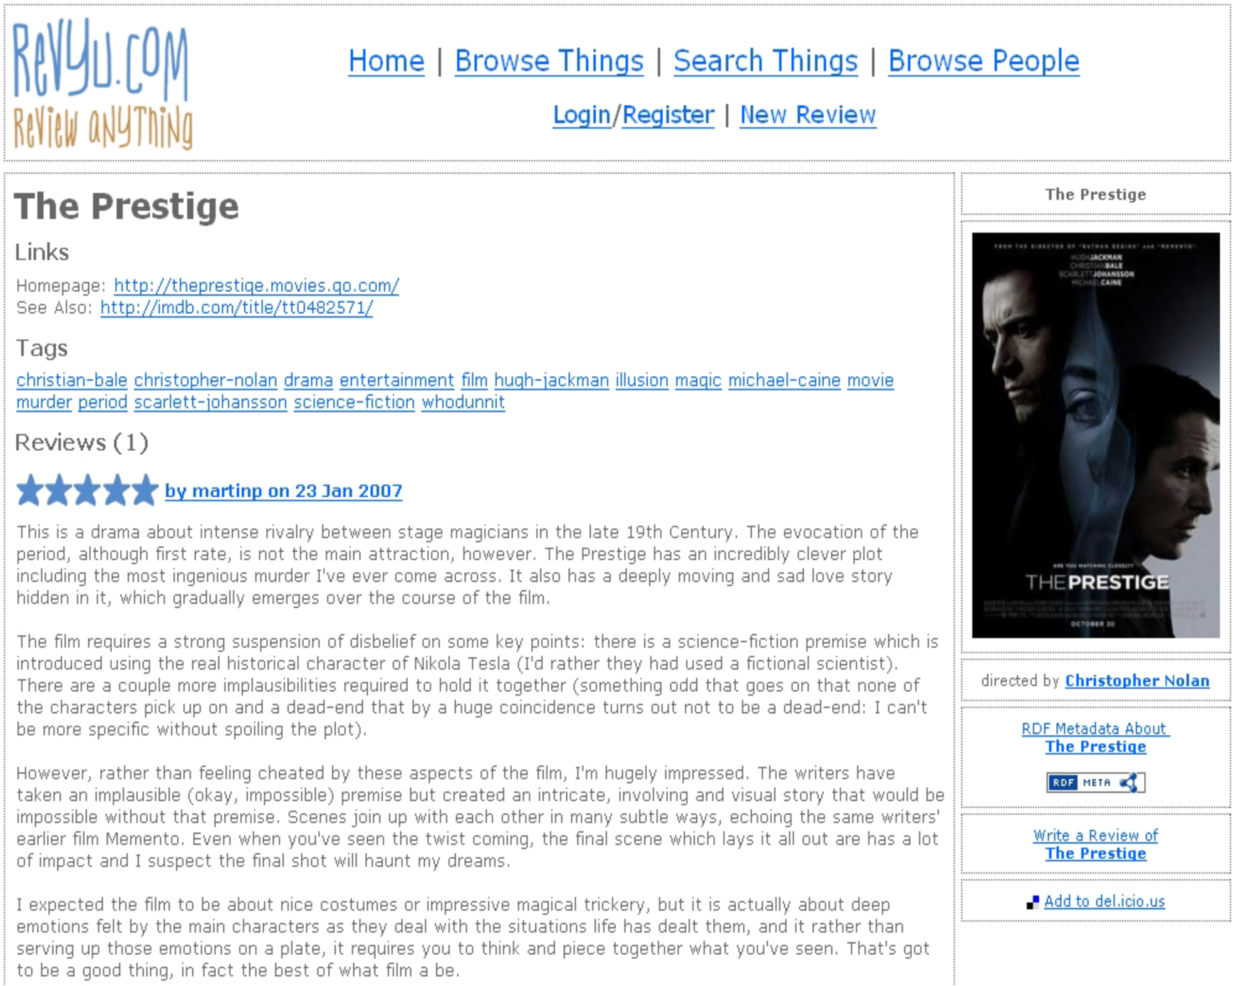

Figure 11 : A film review on Revyu (left) shown alongside film data from DBpedia (right)

Similarly links between Revyu and the *RDF Book Mashup* (Bizer, Cyganiak and Gauss, 2007) are exploited as the basis for retrieving book cover and author information which is also then displayed on the Revyu HTML page about the book, as shown in next Figure .

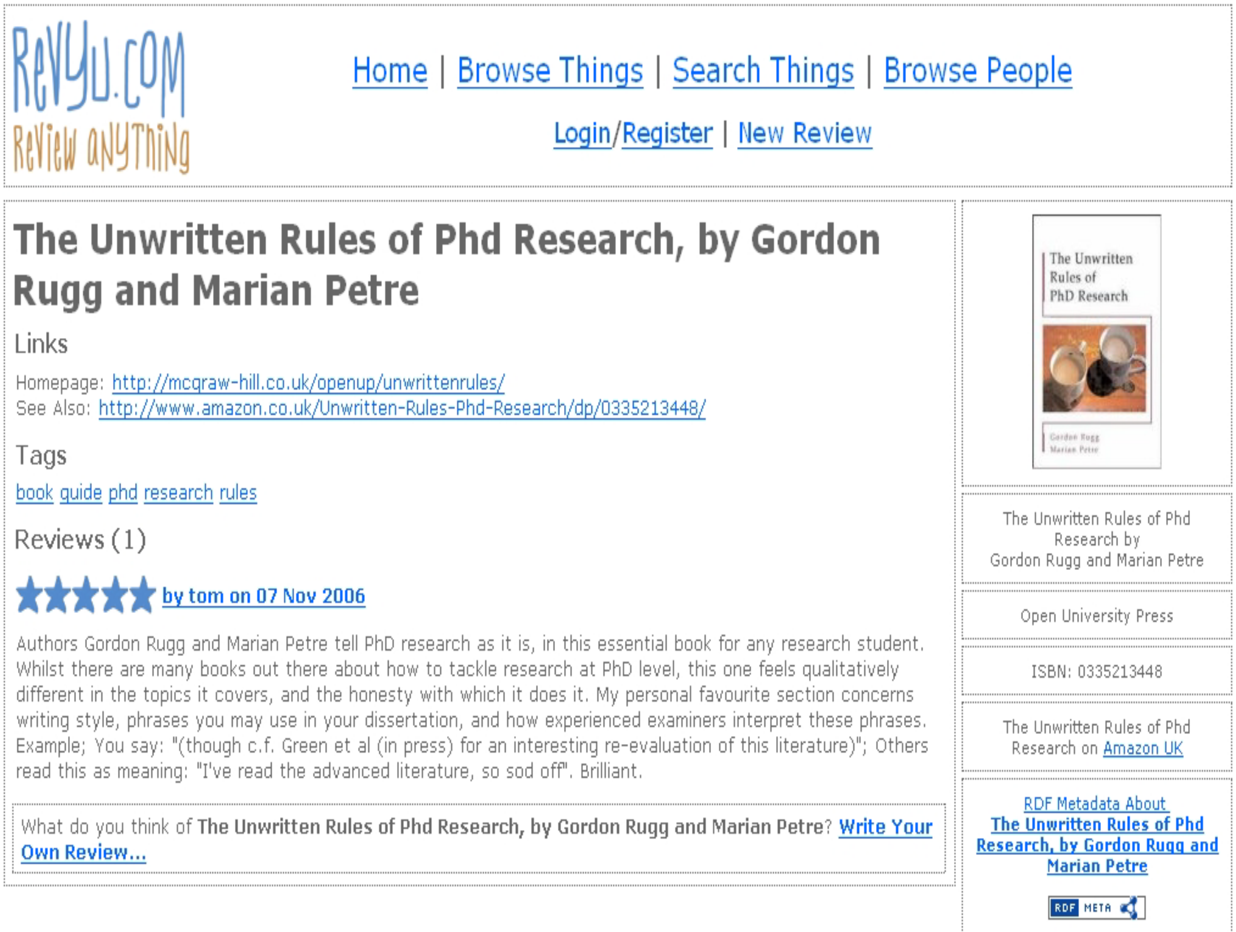

Figure 12 : A book review on Revyu (left) shown alongside film data from Amazon/RDF Book Mashup (right)

This approach could be described as using Semantic Web data to produce Web2.0-style mashups at the human-readable, HTML level, whilst also creating linked data mashups at the RDF level. Not only does this linked data approach to mashups reduce issues with licensing of data for republication, it is also a more Web-like approach; duplicating data is of much lesser value than linking to it, and the user agent of the future should be able to 'look ahead' to linked items and merge data accordingly.

It should be noted that no claims are being made that this form of human-oriented mashup represents something that could not have been achieved using conventional Web2.0 approaches, or provides immediate user benefits over conventional Web2.0 mashups. What distinguishes this approach however is the simultaneous publishing of data and human-oriented mashups, which brings several significant benefits for the developer, for the Semantic Web at large and ultimately for future Web users.

Supplementing Reviewed Items by Pre-population

Whilst links from films and books on Revyu to corresponding items in external data sets are created heuristically, a different approach has been followed when linking Revyu to data from the Open Guide to Milton Keynes (Gaved, Heath and Eisenstadt, 2006) and papers from the 6th International Semantic Web Conference (ISWC+ASWC 2007).

The Open Guide to Milton Keynes is a member of the Open Guides family of wiki-based city guides that publish data in RDF. Milton Keynes is a town in south east England, and home of The Open University. Whilst some amenities in the locality, such as pubs and restaurants, were already reviewed on Revyu, many more were listed in the Open Guide due to its longer history.

Therefore, after identifying items existing in both locations and making the appropriate mappings to avoid duplication, skeleton records were created in Revyu for the remaining items, setting links back to their Open Guide URIs. These skeleton records provide a basic representation of items within Revyu (a title, rdf:type statement, keyword tags and links back to the item in the original data source). This serves to encourage users to review items they recognize, ensures greater coverage and consistency of entries than is possible through organic growth, and ensures that items are properly linked across data sources.

These links enable latitude and longitude data for many items to be retrieved from RDF exposed by the Open Guide, and used to show a Google Map of the items location, as shown in Figure 13 . The same approach can also be used to expose address, telephone, and opening time information held in the Open Guide, and can be extended to Open Guides for other locations, such as London and Boston.

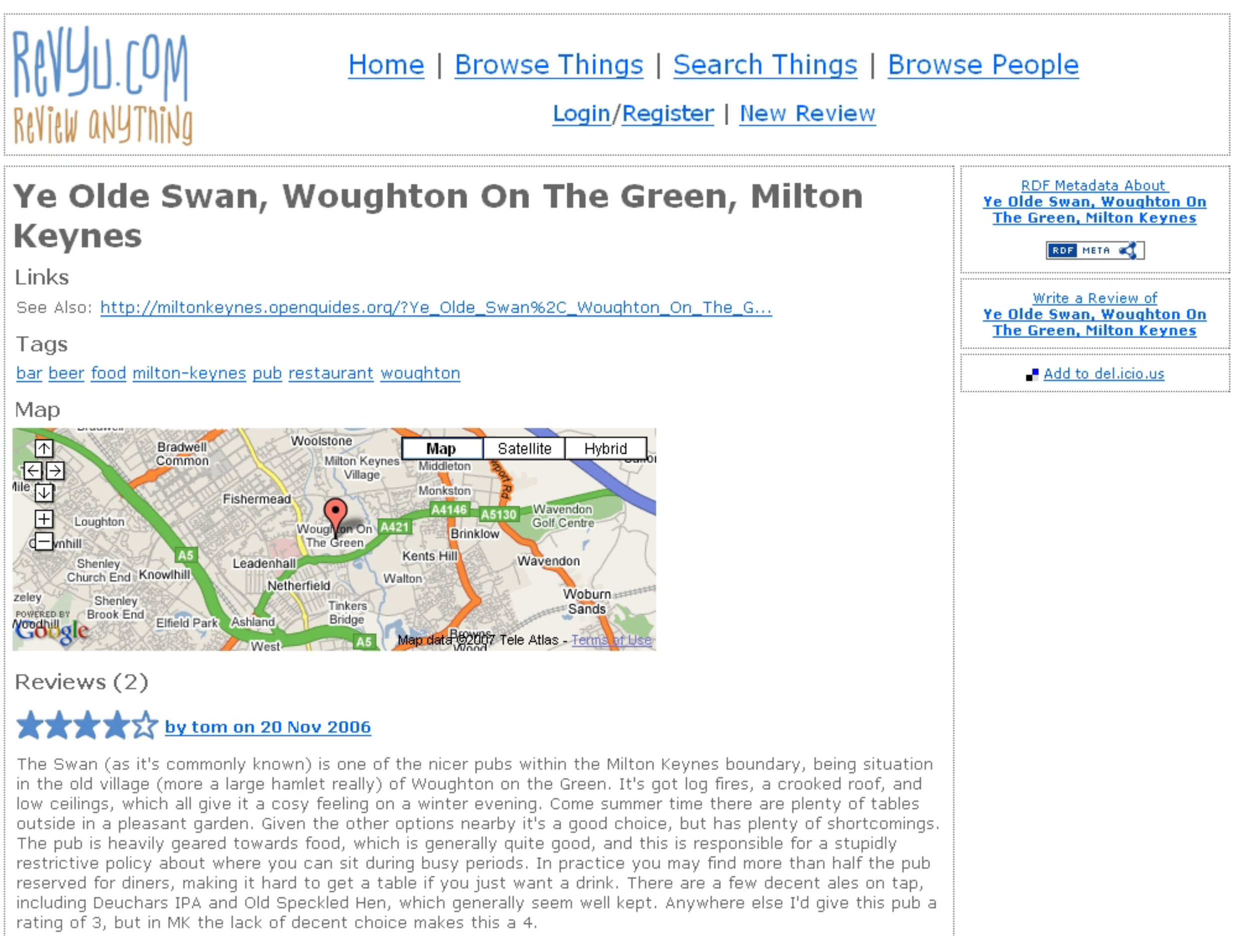

Figure 13 : Geodata from the Open Guide to Milton Keynes used to display a Map of a reviewed item's location

It should be noted that the goal of pre-population from external datasets is not to constrain, but merely to seed users' conceptions of what can be reviewed, where well-defined external data sets exist describing items that may usefully be reviewed in Revyu.

Reusability of Revyu Data

By making content available in standard formats, Revyu reviews can be syndicated and reused by reviewers who use the site and administrators of third party sites who wish to add value to their existing content by adding review information, or combined with reviews from other sources that are also published in RDF. This can be particularly valuable in overcoming the scenario where an item may not have been reviewed many times on one particular site, but reviews exist elsewhere on the Web.

Multiple routes are provided for accessing and reusing Revyu data. With one line of JavaScript code a user's ten latest reviews can be displayed on a remote Web site. This provides a simple mechanism for syndication of reviews by users who are less technically proficient. More sophisticated syndication options are available via RSS feeds of the latest reviews across the entire site and from each individual user.

Third parties interested in data integration rather than simple syndication have two options: retrieving RDF data from the site by crawling or making one-off HTTP requests; or accessing the data they require via queries to the Revyu SPARQL endpoint.

Revyu exposes data about things, reviews, people, and tags via its SPARQL endpoint, which relies on the RAP SPARQL engine operating against the same MySQL-based triplestore. Providing such a query interface allows third parties to retrieve reviews and related data in a flexible fashion, for reuse in their own applications. Whilst in some ways analogous to Web2.0 APIs which provide remote query capabilities, SPARQL endpoints afford many advantages to the developer: for example, common libraries can be used to query multiple RDF graphs yet return the results as one resultset, effectively allowing joins over multiple data sources.

Availability of Data

The issue of available data has shaped many aspects of this research, and continues to be a limiting factor. A lack of review data that was readily available and in a form that enabled integration with social network data led to the creation of Revyu. Whilst Revyu has provided a substantial amount of data with which to test the ideas in this research, a significant increase in available data is required if the benefits of my approach are to be fully investigated. A number of approaches are being considered in order to address this.

Further Pre-population of Revyu

One approach to increasing the amount of available data in the system is the pre-population of Revyu with skeleton records describing things that people may wish to review, in order to attract potential reviewers to the site. Use of this technique with data from the Open Guide to Milton Keynes was described before.

This approach has been considered with a number of significant data sets, such as descriptions of roughly 12,000 films from *DBpedia* (Auer, Bizer et al., 2007) and 70,000 hotels from *Geonames*. Being Semantic Web data sets, integration of Revyu with data from these sources would enable a number of linking opportunities that could greatly enhance the site at a user and data level.

However, initial investigations have identified a number of issues with this approach. The amount of data cleaning required with external data sources can be substantial, in order to address issues such as encoding of foreign characters and removal of bogus data generated by automated methods. Translation of the cleaned data into a format suitable for consumption can also be very resource intensive. This process involves taking the source data as input and generating new RDF graphs that are suitably structured for import into Revyu. Much of this can be achieved using SPARQL CONSTRUCT queries (Prud'hommeaux and Seaborne, 2007) for graph transformations; however except through use of property functions SPARQL does not provide string manipulation functions essential for this kind of data processing, such as when minting URIs. As a result, much of the processing must be carried out programmatically, which in turn increases the resource requirements.

Conclusions

Few mechanisms currently exist that allow non-specialist users to contribute to the Semantic Web. This is in stark contrast to both the conventional Web and Web2.0. Early growth of the Web is widely attributed to individuals creating personal sites by copying and pasting HTML code. Whilst this approach may not be appropriate to a Semantic Web (novice users may not understand the semantics of statements contained in copied code), Web2.0 applications have demonstrated that regular users can contribute content without specialist skills. With few exceptions, similar tools enabling grassroots publishing on the Semantic Web are not currently available. Revyu is one exception.

Revyu is rare in its status as a publicly available service in daily use that is oriented towards human users but also embodies current best practices in developing for the Semantic Web. By adhering to the well established interaction pattern of completing forms in a Web browser, Revyu allows users to create review data that is immediately usable on the Semantic Web. This occurs without any user knowledge of RDF, ontologies, or even the principles of the Semantic Web.

By providing reviews in a reusable format that is easily integrated and linked with other data, Revyu provides source data that is in a format suitable for computing trust metrics that can be integrated with social networks, as discussed previously.

Contributions of the Research

Contribution 1

In order to maximize the value and effectiveness of word-of-mouth recommendation it is important to select the most appropriate information sources. Existing literature has much to say on the matter, however this is mostly confined to either workplace settings or taste domains, as discussed . The first three research questions address issues raised by the shortcomings of previous work on source selection in word-of-mouth information-seeking. It is in addressing these questions, and providing a richer understanding of the source selection process and at a more general level, that this research makes its first major contribution :

- ► An empirical study of decision making in recommendation-seeking identified five trust factors that influence the choice of information sources and their perceived trustworthiness. Variations were identified in how these factors are applied across situations with varying levels of criticality and subjectivity.

These findings provide a basis for systems that may support the source selection process across a range of different tasks. Those that are more critical in nature, and poorly served by current recommender approaches, may benefit greatly from the support of trusted social networks, especially where trust is defined in a task-appropriate fashion.

Contribution 2

Shortcomings were identified in the data available on the Web with which to investigate these questions (4,5,6) . These shortcomings are outlined before , and resulted in the second major contribution of this research:

- Revyu, a live, public reviewing and rating Web site. The site is built on Semantic Web technologies to enable integration of review data with social networks, and easy reuse of the data in deriving word-of-mouth related trust metrics.

Providing review data that is more easily reusable has tangible technical benefits. It also opens review data up to a wider range of systems and service providers who may not otherwise have had access to such information. This may in turn lead to a greater number of systems that develop functionality based on reviews in order to better serve their users.