A Trust Model for Sharing Ratings of Information Providers on the Semantic Web

Jie Zhang, Robin Cohen

David R. Cheriton School of Computer Science, University of Waterloo
Waterloo, Ontario, N2L 3G1, Canada

Abstract. In the context of the Semantic Web, it may be beneficial for a user (consumer) to receive ratings from other users (advisors) regarding the reliability of an information source (provider). We offer a method for building more effective social networks of trust by critiquing the ratings provided by the advisors. Our approach models the consumer's private reputations of advisors based on ratings for providers whom the consumer has had experience with. It models public reputations of the advisors according to all ratings from these advisors for providers, including those that are unknown to the consumer. Our approach then combines private and public reputations by assigning weights for each of them. Experimental results demonstrate that our approach is robust even when there are large numbers of advisors providing large numbers of unfair ratings. As such, we present a framework for sharing ratings of possibly unreliable sources, of value as users on the Semantic Web attempt to critique the trustworthiness of the information they seek.

1 Introduction

The vision of the Semantic Web is to construct a common semantic interpretation for World Wide Web pages, in order to one day reliably run software to understand the information conveyed in any of its documents. In building the Semantic Web, however, information may be supplied by a wide selection of sources, with the result that a user seeking information will need to judge whether the content of any given source is in fact trustworthy. It is therefore important to develop models for trust in the context of the Semantic Web. Various approaches to date have been formulated about how best to form a Web of Trust, in order to share information and selectively choose trustworthy partners from whom information may be obtained. In our research, we are considering a problem that arises when social networks are formed in order to share trust ratings - that of unfair ratings. Dellarocas (Dellarocas 2000) distinguishes unfair ratings as unfairly high ratings and unfairly low ratings. Unfairly high ratings may be used to increase the trustworthiness of others and promote their services. They are often referred to as “ballot stuffing”. Unfairly low ratings of others are often referred to as “bad-mouthing”. In brief, the ratings of the trustworthiness of others, obtained from third parties, may in fact be suspect. What is required therefore is a mechanism for effectively adjusting the basis on which decisions of trust are made, to discount these possibly unfair ratings.

In this paper, we discuss our research in the context of sharing ratings of sources (called information providers) among users on the Semantic Web. We present an approach for modeling the trustworthiness of advisors - those users providing reputation ratings for potential providers from whom information may be obtained. We refer to the user seeking advice as the consumer. We first represent private reputation values, based on what is known about the advisors’ ratings for providers with whom the consumer has already had some experience. We then describe how to construct a public model of trustworthiness of advisors based on common, centrally held knowledge of providers and the ratings provided by advisors, including the reputation ratings of providers totally unknown to the consumer. We then outline how both private and public models can be combined, in order to obtain a value for the trustworthiness of each possible advisor. In summary, we offer a method for building more effective social networks of trust, by critiquing the advice provided by advisors.

In Section 2 we introduce the Semantic Web setting for sharing information about sources, and present some current research on modeling the trustworthiness of information sources based on ratings provided by advisors. Section 3 presents our approach for modeling the trustworthiness of advisors according to the ratings provided by them in the context of the Semantic Web. Section 4 provides an example that goes through each step of our approach. Section 5 includes some experimental results demonstrating what happens when there are large numbers of advisors providing large numbers of unfair ratings. Conclusions and future work are outlined in Section 6.

2 Background and Related Work

In this section, we discuss the setting of sharing information about sources, on the Semantic Web. We motivate the need to acquire information about the reliability of sources and then briefly outline some current research on modeling the trustworthiness of sources. This includes some discussion of approaches to communicate with other users to obtain advice about sources, sometimes referred to as a Web of Trust (Gil and Ratnakar 2002), as well as an approach for addressing the problem that some users may provide untruthful advice.

The challenge of trusting information providers in a Web-based environment is discussed in (Paolucci and Sycara 2003). Paolucci et al. provide valuable insights into the need for trust on the Web, in the context of Web services, where Web sites dynamically exchange information using XML descriptions, but where it is difficult to ensure that the meaning of the messages being sent is well understood, without human intervention. The Semantic Web contributes by providing ontologies for Web services to interpret meanings in exchanged messages. According to (Paolucci and Sycara 2003), with the Semantic Web, the interaction between users and providers needs a process of capability matching to link users with providers of Web services. Specifically, providers advertise their capabilities, a user sends a request for the type of service she requires, a registry matches the capabilities of providers and the capabilities expected by the user, and finally the user selects the most suitable provider. However, in their advertisements, providers may lie about their capabilities in order to be selected by the user. To avoid selection of an untruthful provider, there is a need to properly model the trustworthiness of providers. In (Gil and Ratnakar 2002) this problem is reinforced for the Semantic Web: whether to trust the content of a Web resource, depending on the source. Richardson et al. (Richardson et al. 2003) explain further that due to the great diversity of the Web, it is difficult to expect the content to be consistent and of high quality. It then becomes important to decide how trustworthy each information source is.

Maximilien and Singh (Maximilien and Singh 2004, 2005) adopt an agent-based approach for modeling trust on the Semantic Web. Their work focuses on representing multiple qualities of services (QoS) for automatic runtime Web service selection. This trust model is based on a shared conceptualization of QoS and takes into account providers' quality advertisement, consumers' quality preferences, quality relationships, and consumers' quality tradeoffs. In order to select a Web service implementation, a consumer dynamically associates a trust value with each service implementation and selects the service implementation with the highest assigned level of trust. The trust value of each service implementation partially depends on its reputation value, which is determined by the set of quality values from other users who previously selected that provider.

Kagal et al. (Kagal et al. 2002) use a DAML+OIL trust ontology in a multi-agent system, which is based on a distributed trust and delegation mechanism verifying that a user's credentials are acceptable. The trust ontology is built for specifying credentials and checking if the credentials conform to policies. A policy maps credentials to a certain ability or right. The mechanism allows propagation of trust beliefs exchanged between users and avoids repeated checking of users' credentials.

The research of Gil and Ratnaker (Gil and Ratnakar 2002) provides a framework for users to express their trust about a source and the statements it contains, by annotating each part of a source to indicate their views. The focus of the work is on how to provide an effective interface for users to record their annotations. This TRELLIS system ultimately averages the ratings provided over many users and many analysis, to present a reflection of the trustworthiness of the source. A credibility-reliability pair emerges for each source-statement pair, to derive an overall rating of a single source, based on each of its associated statements.

Modeling trust on the Semantic Web, as discussed so far in this section, includes a reliance on the beliefs or ratings provided by third parties to be truthful. In fact, it is important to address the problem of possibly unfair or unreliable ratings. One approach that explores this possibility is that of Richardson et al. (Richardson et al. 2003). In this work, each user first explicitly specifies a small set of users whom she trusts, leading to a Web of Trust. This arrangement allows any user to compute the trustworthiness of a possible provider, based on the ratings supplied by others in her social network. The trust value of a provider is computed locally by combining the trust ratings provided by other users. One feature of this approach is to recursively propagate trust through the user's social network. In effect, trust in a provider is derived using some aggregating functions along each possible chain of trust from the user to the provider. One concern with this approach, however, is that this method of propagating trust may be computationally intractable, as there may be many different paths, of various lengths, which need to be aggregated.

In our own research, we are developing a model for representing the reliability of advisors from whom advice may be sought, when a user seeks to evaluate the trustworthiness of a provider. This framework is sufficiently general to operate in a variety of environments including electronic commerce, where buyers may make decisions about sellers by soliciting input on those sellers from other buyers in the marketplace.

In the context of the Semantic Web, our model is useful for the problem of determining the reliability of a provider being evaluated by a consumer by virtue of reputation ratings provided by advisors. Our focus is on addressing the problem of advisors who may be untrustworthy. The existence of malicious advisors is in fact acknowledged in (Richardson et al. 2003). But in contrast to the model of Richardson et al. (Richardson et al. 2003), we provide a more direct evaluation of each possible advisor in a Web of Trust, leading to an evaluation about how best to make use of that advisor's ratings of a possible provider being examined by a consumer.

As will be seen in the sections that follow, we make various limiting assumptions (which are revisited as future work) in order to examine more clearly the need to adjust for possibly unfair ratings from advisors. In particular, we do not envisage entire chains of trust from advisor to advisor, instead evaluating independently the trustworthiness of each advisor, based in part on the user's own past experience. In addition, we represent the input from each advisor as a summary rating of a possible source as simply reliable or unreliable. We also allow an advisor to rate a source several times. In so doing, we are able to weight more heavily more recent evaluations of the source, allowing for dynamically varying trustworthiness of the source.

3 Modeling Trustworthiness of Advisors

In the discussion below, we use the following terminology:

- **User/Consumer:** person seeking information from various sources
- **Provider:** an information source, providing information
- **Advisor:** other users providing ratings of providers to consumers
- **Private reputation:** a determination of the reputation of an advisor by a user, based on commonly rated providers
- **Public reputation:** a determination of the reputation of an advisor by a user, based on a centrally held model of the advisor, from interactions with a whole set of providers

Our method for determining the trustworthiness of advisors is to employ a combination of what we refer to as private and public reputation values. To explain, the private reputation of an advisor is calculated by a consumer, based on ratings the advisor supplies of providers with whom the consumer has already had some experience. If the advisor is reputable and has similar preferences as the consumer, the consumer and advisor will likely have many ratings in common. This can then be used as the basis for assessing the trustworthiness of the advisor. In cases where the consumer has little private knowledge of the advisor, a public reputation will be elicited, reflecting the trustworthiness of that advisor, based on his ratings of all providers in the system. A weighted combination of private and public reputations is derived, based on the estimated reliability of the private reputation value. This combined value then represents the trustworthiness of the advisor.

3.1 Private Reputation

Our approach allows a consumer *C* to evaluate the private reputation of an advisor *A* by comparing their ratings for commonly rated providers $\{P_1, P_2, ..., P_m\}$. For one of the commonly rated providers P_i ($1 \leq i \leq m$ and $m \geq 1$), *A* has the rating vector R_{A,P_i} and *C* has the rating vector R_{C,P_i}. A rating for P_i from *C* and *A* is binary ("1" or "0", for example), in which "1" means that P_i is trustworthy and "0" means that P_i is untrustworthy. For the purpose of simplicity, we assume ratings for providers are binary. Possible ways of extending our approach to accept ratings in different ranges will be investigated as future work. Further discussion can be found in Section 6.

The ratings in R_{A,P_i} and R_{C,P_i}. are ordered according to the time when they are provided. The ratings are then partitioned into different elemental time windows. The length of an elemental time window may be fixed (e.g. three days) or adapted by the frequency of the ratings to the provider P_i, similar to the way proposed in (Dellarocas 2000). It should also be considerably small so that there is no need to worry about the changes of providers' behavior within each elemental time window. We define a pair of ratings (r_{A,P_i}, r_{C,P_i}), such that r_{A,P_i} is one of the ratings of R_{A,P_i}, r_{C,P_i} is one of the ratings of R_{C,P_i}, and r_{A,P_i} corresponds to r_{C,P_i}. The two ratings, r_{A,P_i} and r_{C,P_i}, are correspondent only if they are in the same elemental time window, the rating r_{C,P_i} is the most recent rating in its time window, and the rating r_{A,P_i} is the closest and prior to the rating r_{C,P_i}. We consider ratings provided by *C* after those by *A* in the same time window, in order to incorporate into *C*'s rating anything learned from *A* during that time window, before taking an action. According to the solution proposed by Zacharia et al. (Zacharia et al. 1999), by keeping only the most recent ratings, we can avoid the issue of advisors "flooding" the system. No matter how many ratings are provided by one advisor in a time window, we only keep the most recent one.

We then count the number of such pairs for P_i, N_{P_i}. The total number of rating pairs for all commonly rated providers, N_{all} will be calculated by summing up the number of rating pairs for each commonly rated provider as follows:

$$N_{all} = \sum_{i=1}^{m} N_{P_i}$$

The private reputation of the advisor is estimated by examining rating pairs for all commonly rated providers. We define a rating pair (r_{A,P_i}, r_{C,P_i}) as a positive pair if r_{A,P_i} is the same value as r_{C,P_i}. Otherwise, the pair is a negative pair. Suppose there are N_f number of positive pairs. The number of negative pairs will be $N_{all} - N_f$. The private reputation of the advisor A is estimated as the probability that A will provide reliable ratings to C. Because there is only incomplete information about the advisor, the best way of estimating the probability is to use the expected value of the probability. The expected value of a continuous random variable is dependent on a probability density function, which is used to model the probability that a variable will have a certain value. Because of its flexibility and the fact that it is the conjugate prior for distributions of binary events (Russell and Norvig 2002), the beta family of probability density functions is commonly used to represent probability distributions of binary events (see, e.g. the generalized trust models BRS (Jøsang and Ismail 2002) and TRAVOS (Teacy et al. 2005)). Therefore, the private reputation of A can be calculated as follows:

$$\alpha = N_f + 1, \ \beta = N_{all} - N_f + 1$$

$$R_{pri}(A) = E(Pr(A)) = \frac{\alpha}{\alpha + \beta},$$

where $Pr(A)$ is the probability that A will provide fair ratings to C, and $E(Pr(A))$ is the expected value of the probability.

3.2 Public Reputation

When there are not enough rating pairs, the consumer C will also consider A's public reputation. The public reputation of A is estimated based on its ratings and other ratings for the providers rated by A. Each time A provides a rating $r_{A,P}$, the rating will be judged centrally as a fair or unfair rating. We define a rating for a provider as a fair rating if it is consistent with the majority of ratings to the provider up to the moment when the rating is provided.[1] As before, we consider only the ratings within a time window prior to the moment when the rating $r_{A,P}$ is provided, and we only consider the most recent rating from each advisor. In so doing, as providers change their behavior and become more or less reputable to each advisor, the majority of ratings will be able to change.

Suppose that the advisor A totally provides N'_{all} ratings. If there are N'_{f} number of fair ratings, the number of unfair ratings provided by A will be $N'_{all} - N'_{f}$. In a similar way as estimating the private reputation, the public reputation of the advisor A is estimated as the probability that A will provide fair ratings. It can be calculated as follows:

$$\alpha' = N'_{f} + 1,\ \beta' = N'_{all} - N'_{f} + 1$$

$$R_{pub}(A) = \frac{\alpha'}{\alpha' + \beta'},$$

which also indicates that the more the percentage of fair ratings advisor A provides, the more reputable it will be.

3.3 Trustworthiness

To estimate the trustworthiness of advisor A, we combine the private reputation and public reputation values together. The private reputation and public reputation values are assigned different weights. The weights are determined by the reliability of the estimated private reputation value.

We first determine the minimum number of pairs needed for C to be confident about the private reputation value it has of A. The Chernoff Bound theorem (Mui et al. 2002) provides a bound for the probability that the estimation error of private reputation exceeds a threshold, given the number of pairs. Accordingly, the minimum number of pairs can be determined by an acceptable level of error and a confidence measurement as follows:

$$N_{\min} = -\frac{1}{2\varepsilon^2} \ln \frac{1-\gamma}{2},$$

---

[1] Determining consistency with the majority of ratings can be achieved in a variety of ways, for instance averaging all the ratings and seeing if that is close to the advisor's rating.

where ε is the maximal level of error that can be accepted by *C*, and γ is the confidence measure. If the total number of pairs N_{all} is larger than or equal to $N_{\min}$, consumer *C* will be confident about the private reputation value estimated based on its ratings and the advisor *A*'s ratings for all commonly rated providers. Otherwise, there are not enough rating pairs, the consumer will not be confident about the private reputation value, and it will then also consider public reputation. The reliability of the private reputation value can be measured as follows:

$$w = \begin{cases} \dfrac{N_{all}}{N_{\min}} & if\ N_{all} < N_{\min}; \\ 1 & otherwise. \end{cases}$$

The trust value of *A* will be calculated by combining the weighted private reputation and public reputation values as follows:

$$Tr(A) = wR_{pri}(A) + (1 - w)R_{pub}(A)$$

It is obvious that the consumer will consider less the public reputation value when the private reputation value is more reliable. Note that when $w = 1$, the consumer relies only on private reputation.

4 An Example

To illustrate how our approach models trustworthiness of advisors, this section provides an example that goes through each step of the approach.

In the setting of sharing information on the Semantic Web, a provider P_0, which is an information source, provides some statements. Whether a consumer *C* can trust these statements depends on how much *C* trusts P_0. To model the trustworthiness of the provider P_0, the consumer *C* seeks advice from two advisors A_x and A_y who have experience with P_0. The advice about P_0 from A_x and A_y are ratings representing the trustworthiness of P_0 in terms of providing reliable content. Before aggregating the ratings provided by A_x and A_y, the consumer *C* needs to evaluate the reliability of those ratings, which depends on the trustworthiness of the advisors A_x and A_y. Our approach effectively models the trustworthiness of advisors based on how reliable the previous ratings provided by them are.

To demonstrate what ratings provided by advisors may look like, we assume both the advisors A_x and A_y have rated one of the providers, P_i. We are in fact interested in all P_i's for which A_x or A_y has supplied ratings and *C* has had experience. Table 1 lists some of the ratings provided by A_x and

A_y for P_i. The symbol "T" represents a sequence of time windows, in which T_1 is the most recent time window. To simplify the demonstration, we assume that each advisor provides at most one rating within each time window. Some advisors might have not provided any ratings for the provider within some time window. For example, the advisor A_y has not provided any ratings for P_i within the time window T_{n-1}. As can be seen from Table 1, the consumer C also provides some ratings for P_i; some of the ratings are within the same time windows as the ratings provided by A_x and A_y. We assume that the ratings provided by C are after those provided by A_x and A_y if they are within the same time window.

Table 1. Ratings Provided by A_x, A_y and C for P_i

	P_i						
T	T_1	T_2	…	T_j	…	T_{n-1}	T_n
A_x	1	1	…	1	…	1	1
A_y	1	0	…	1	…	-	0
C	1	-	…	0	…	1	1

Suppose that A_x and A_y each provides 40 ratings in total for providers. In this case, $N'_{all}(A_x) = N'_{all}(A_y) = 40$. The advisor A_x provides 35 fair ratings ($N'_f(A_x) = 35$), and A_y provides 20 fair ratings ($N'_f(A_y) = 20$). A rating here is considered as a fair rating when it is consistent with the majority of ratings for the provider within a same time window. Then the public reputation values of A_x and A_y are calculated as follows:

$$R_{pub}(A_x) = \frac{35+1}{35+1+(40-35)+1} = 0.86;$$

$$R_{pub}(A_y) = \frac{20+1}{20+1+(40-20)+1} = 0.5,$$

which means that A_x is more likely to provide fair ratings.

Suppose that the consumer C provides 30 ratings that are within the same time windows of the same providers with A_x and A_y. Therefore, $N_{all}(A_x) = N_{all}(A_y) = 30$. Within those 30 ratings pairs, 25 of ratings provided by A_x are same as the ratings provided by C ($N_f(A_x) = 25$), and A_y provides only 20 same ratings ($N_f(A_y) = 20$). Then the private reputation values of A_x and A_y are calculated as follows:

$$R_{pri}(A_x) = \frac{25+1}{25+1+(30-25)+1} = 0.81;$$

$$R_{pri}(A_y) = \frac{20+1}{20+1+(30-20)+1} = 0.66,$$

which means that A_x is more likely to provide fair ratings and have similar preferences with C.

To combine the private and public reputation values, the weight w should be determined. Suppose $\varepsilon = 0.1$ and $\gamma = 0.9$, then $N_{\min} = -\frac{1}{2\times 0.1^2}\ln\frac{1-0.9}{0.2} = 150$. Since N_{all} is less than $N_{\min}$, $w = \frac{30}{150} = 0.2$. The trust values of A_x and A_y will be calculated as follows:

$$Tr(A_x) = 0.2\times 0.81 + (1-0.2)\times 0.86 = 0.85;$$
$$Tr(A_y) = 0.2\times 0.66 + (1-0.2)\times 0.5 = 0.53,$$

which clearly indicates that A_x is more trustworthy than A_y. As a result, the consumer C will place more trust in the advice provided by A_x. It will consider the advice provided by A_x more heavily when aggregating the advice provided by A_x and A_y for modeling the trustworthiness of the information provider P_0. Discussion of possible aggregation functions is necessary when employing our model to reach final decisions about which sources to trust. A brief summary of some aggregation functions and references of some others can be found in (Richardson et al. 2003). We leave the topic of selecting effective aggregation functions to future work. Our framework serves the purpose of representing the trustworthiness of advisors, so that this may be taken into account, when determining how heavily to rely on their advice.

5 Experimental Results

Our approach models the trustworthiness of advisors according to how reliable the ratings provided by them are. To demonstrate the effectiveness of the approach, we carry out some modest preliminary experiments involving advisors who provide different percentages of unfair ratings. The expectation is that trustworthy advisors will be less likely to provide unfair ratings, and vice versa. We also examine how large numbers of dishonest advisors (i.e. advisors that provide unfair ratings) will affect the estimation of advisors' trustworthiness. Results indicate that our approach is still effective by making adjustments to rely more heavily on private reputations of advisors, in this case.

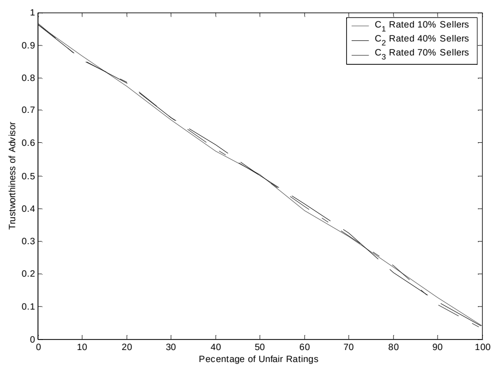

Fig. 1. Trustworthiness of Advisor

The first experiment involves 100 providers, 3 consumers, and one advisor. The 3 consumers, C_1, C_2 and C_3, rate 10, 40 and 70 randomly selected providers, respectively. The advisor totally rates 40 randomly selected providers. [2] We examine how the trust values the consumers have of the advisor change when different percentages (from 0% to 100%) of its ratings are unfair. As illustrated in Figure 1, the trust values the consumers have of the advisor decrease when more percentages of the advisor's ratings are unfair. From this figure, we can also see that our approach is still effective when the consumer C_1 does not have much experience with providers, in the sense that C_1 can still reduce the reputation of the advisor when it provides more unfair ratings.

The second experiment involves 100 provides, 80 advisors, and one consumer. The consumer and each advisor rate 80 of the randomly selected providers. We model the trust value the consumer has of one of the advisors, *A*. The trustworthiness of the advisor will be modeled as the combination of its private and public reputations (referred to as the CR approach) and as only its public reputation (referred to as the PR approach), respectively. The advisor *A* will provide different percentages (from 10% to 100%) of unfair ratings. Figure 2 illustrates the trustworthiness of *A* when 24 (30% of all) advisors are dishonest. Those dishonest advisors provide the same percentage of unfair ratings as the advisor *A* does. Re-

[2] Note that we simplify the experiments by limiting each consumer or advisor to provide at most one rating for each provider.

sults indicate that the trustworthiness of *A* modeled by using the CR and

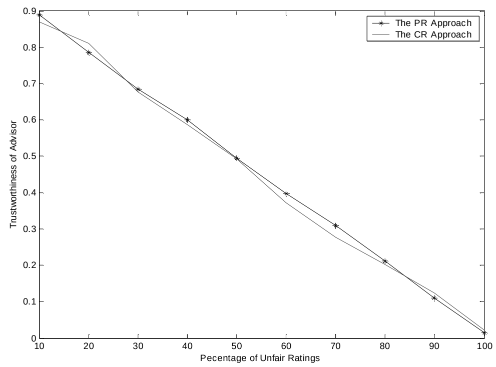

Fig. 2. Trustworthiness of *A* When Majority of Advisors are Honest

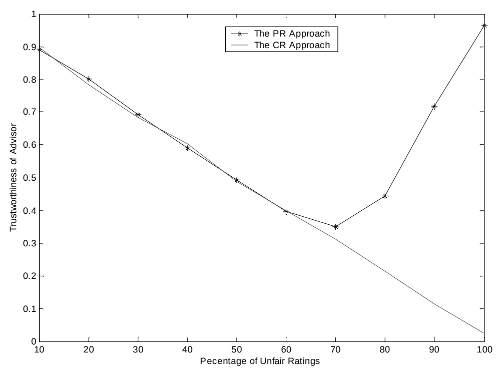

Fig. 3. Comparison of the CR and PR Approaches

PR approaches decreases when more percentages of ratings provided by *A* are unfair. Therefore, these two approaches are not affected when only a small number of advisors are dishonest. Figure 3 represents the trustworthiness of *A* when 40 (60% of all) advisors are dishonest. In this figure, the trustworthiness of *A* modeled by using the CR approach still decreases when more percentages of ratings provided by *A* are unfair, which indicates that our approach is still effective when the majority of advisors provide large numbers of unfair ratings. In contrast, the trustworthiness modeled by using the PR approach increases when more than 60% of ratings provided by the dishonest advisors are unfair, which indicates that the PR approach is only effective when the majority of ratings are fair.

6 Conclusions and Future Work

In this paper, we first introduce the Semantic Web setting for sharing information about sources. Due to the fact that any user on the Web can become an information source, there is a need to form a Web of Trust. Current research on modeling the trustworthiness of information sources on the Semantic Web relies on the unrealistic assumption that advice provided by advisors about an information source is truthful. A typical approach to address this problem is to critique advisors' advice based on their trustworthiness. We present an approach for modeling the trustworthiness of advisors. Our approach allows a consumer to estimate the trustworthiness of an advisor based on the advisor's ratings for providers with whom the consumer has already had some experience. It also models the trustworthiness of the advisor based on all its ratings and common knowledge of providers who might be totally unknown to the consumer. The above results are finally combined by our approach. The experiments are carried out in the setting where advisors might provide different numbers of unfair ratings. Experimental results indicate that our approach can effectively model the trustworthiness of advisors even when consumers do not have much experience with providers. Furthermore, our approach is still effective when the majority of advisors provide large numbers of unfair ratings.

Our approach of combining both private and public reputation values offers useful improvement for the modeling of the trustworthiness of advisors. A model such as BRS ((Jøsang and Ismail 2002) that relies on public reputation has the problem that it is only effective when the majority of ratings are fair, whereas a model like TRAVOS (Teacy et al. 2005) that uses private reputation has difficulty when a consumer is new to the system.

For the purpose of simplicity, the current approach limits ratings for providers to be binary. In future work, we will extend our approach to accept ratings in different ranges. Instead of using the numerical difference of two ratings, comparison of the two ratings could take into account the semantics of rating levels (Chen and Singh 2001). For example, although the numerical differences of the pairs are same, the difference between "5" (very trustworthy) and "3" (neutral) is smaller than that between "4" (trustworthy) and "2" (untrustworthy). In consequence, the similarity between "5" and "3", say 0.2, should be set to be larger than the similarity between "4" and "2", say 0. When these extensions are made, the Dirichlet family of probability density functions (Gelman et al. 2004), which is the multivariate generalization of the beta family, can be used to represent probability distributions of discrete similarity values. Our model will evaluate private and public reputation values based on aggregation of those discrete similarity values

Our approach represents trustworthiness of providers using a single rating provided by consumers or advisors. For future work, as in the research of (Richardson et al. 2003), we will also extend our approach to accept multiple ratings representing different dimensions of trustworthiness of providers. We could for example, examine credibility and reliability of providers as used by Gil and Ratnakar (Gil and Ratnakar 2002) or a quality of service ontology used by Maximilien and Singh (Maximilien and Singh 2004, 2005). We would then need to explore methods to combine the different kinds of ratings provided by advisors, for example whether to weight one dimension more heavily than another.

Another valuable direction for future work is to go beyond a generalized reputation rating for an information source, to one that determines whether to trust a source on a particular topic or segment of its information. In this case, we would want to model the advisors' trustworthiness with respect to these segments of the provider, as well. This may result in the design of a more elaborate private reputation model or a method of determining what weight to place on this private reputation, when advisors have only currently rated different segments of the source. It would also be valuable to learn which advisors to rely on, for which different elements of a source.

For future work, we will also carry out further experiments to continue to compare our model with competing approaches. It is important to note that we are focused in this paper on the question of judging the trustworthiness of advisors, as part of the process of evaluating how much to trust the content of an information source. In fact, we would like to see our approach integrated into a full scale decision-theoretic framework for selecting reputable sources. The performance of the overall system would then need to be evaluated, as well.

References

Chen M, Singh JP (2001) Computing and using reputations for internet ratings. In: Proceedings of the 3nd ACM Conference on Electronic Commerce.

Dellarocas C (2000) Immunizing online reputation reporting systems against unfair ratings and discriminatory behavior. In: Proceedings of the 2nd ACM Conference on Electronic Commerce, Minneapolis, MN.

Gelman A, Carlin JB, Stern HS, Rubin DB (2004) Bayesian Data Analysis. Chapman & Hall/CRC, Boca Raton.

Gil Y, Ratnakar V (2002) Trusting information sources one citizen at a time. In Proceedings of the First International Semantic Web Conference (ISWC), Sardinia, Italy.

Jøsang A, Ismail R (2002) The beta reputation system. In: Proceedings of the 15th Bled Electronic Commerce Conference.

Kagal L, Finin T, Joshi A (2002) Developing secure agent systems using delegation based trust management. In: Proceedings of the First International Autonomous Agents and Multi Agent Systems (AAMAS), Bologna, Italy.

Maximilien EM, Singh MP (2004) Toward autonomic web services trust and selection. In: Proceedings of the 2th International Conference on Service Oriented Computing (ICSOC), pp 212–221.

Maximilien EM, Singh MP (2005) Agent-based trust model using multiple qualities. In: Proceedings of 4th International Autonomous Agents and Multi Agent Systems (AAMAS 2005).

Mui L, Mohtashemi M, Halberstadt A (2002) A computational model of trust and reputation. In: Proceedings of the 35th Hawaii International Conference on System Science (HICSS).

Paolucci M, Sycara K, Nishimura T, Srinivasan N (2003) Toward a semantic web ecommerce. In: Proceedings of 6th Conference on Business Information Systems (BIS2003), Colorado Springs, Co, USA, pp 153–161.

Richardson M, Agrawal R, Domingos P (2003) Trust management for the semantic web. In: Proceedings of the 2nd International Semantic Web Conference (ISWC).

Russell S, Norvig P (2002) Artificial Intelligence: A Modern Approach. Second Edition, Prentice Hall, Englewood Cliffs, New Jersey.

Teacy WTL, Patel J, Jennings NR, Luck M (2005) Coping with inaccurate reputation sources: Experimental analysis of a probabilistic trust model. In: Proceedings of 4th International Autonomous Agents and Multi Agent Systems (AAMAS 2005).

Zacharia G, Moukas A, Maes P (1999) Collaborative reputation mechanisms in electronic marketplaces. In: Proceedings of the 32nd Hawaii International Conference on System Sciences (HICSS-32), Maui, Hawaii.