

Accuracy of Metrics for Inferring Trust and Reputation in Semantic Web-based Social Networks

Jennifer Golbeck¹, James Hendler¹

¹University of Maryland, College Park
MIND Lab, 8400 Baltimore Ave
College Park, Maryland 20740
{golbeck, hendler}@cs.umd.edu

Abstract: While most research on the topic of trust on the semantic web has focused largely on digital signatures, certificates, and authentication, more social notions of trust which are reputation-based are starting to gain attention. In this paper, we describe an algorithm for generating locally-calculated reputation ratings from a Semantic Web Social Network. We present mathematical and experimental results that show the effectiveness of this algorithm to accurately infer the reputation of a node. We then describe TrustMail, an application that uses the network for rating email.

1 Introduction

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 [1,4,20,21].

Trust and reputation can be expressed on the semantic web using ontologies that provide a method for describing entities and the trust relationships between them. These ontological foundations allow for trust to be expressed in people, statements, or the content of information sources. They also facilitate expressing different trust relationships with respect to different topics. To make use of these ratings, the proper subset of entities and relationships are extracted and metrics to infer relationships are used on this set.

This work is motivated by the sudden explosion of social network data on the web. The Friend of a Friend (FOAF) project now comprises millions of files. We have created a trust extension to FOAF that allows people to create ratings for one another. This project, located at <http://trust.mindswap.org/>, provides tools and support for producing trust data that can be linked to an aggregator. The site also has a set of methods that provide data about the network and an interface for making trust inferences between two individuals. In two months, the network has grown to nearly 1,300 users (see figure 1) and they were used as the foundation for applying these metrics to the email application in section 5.




Fig 1: The trust network developed at <http://trust.mindswap.org/>

For methods that use the network to be successful, metrics over it must be accurate. In this paper, we present an algorithm for aggregating and inferring reputation ratings on the Semantic Web. We show mathematically and through experiments that this algorithm gives very accurate inferences, even in the face of high uncertainty. In section 5, we discuss applications that we have developed using this reputation system, and then present more areas where such an analysis can be useful.

2 Background and Related Work

This research covers several spaces of research that typically has involved different communities. Because of that, the following sections are broken up by subfield.

2.1 Small World Networks and Power Law Distributions

The concept of small world networks and the notion of six degrees of separation – the idea that any two people in the world are connected through only six intermediate people – began with Stanley Milgram’s seminal work, “The Small World Problem” [20] in 1967. He distributed letters to strangers in the Midwest, with the instruction that they pass the letters to friends with the ultimate goal of the letter reaching a specific person in Boston. Based on the number of people through whom the letters traveled, Milgram concluded that everyone in the country was connected through a chain of at most six people.

Since his paper was published, the concept of small worlds has been formalized. Small world networks are characterized by two main properties. First, there is a high degree of “connectance” compared with random graphs. Connectance is the property of clustering in neighborhoods. Given a node n , connectance is the fraction of edges between neighbors of n that actually exist compared to the total number of possible edges. Small world graphs have strong connectance; neighborhoods are usually very connected. The second property, which is computationally significant, states the average shortest path length between two nodes grows logarithmically with the size of the graph. This means that many computations can be expected to complete efficiently.

2.2 Defining Trust and Reputation

Trust and reputation are both socially loaded words. To use the terms as a computational concepts requires a strong definition. The first cut to make is to differentiate trust in this work from security. There has been quite a bit of research that uses trust in this sense, but this research focuses on trust as social concept, and the methods described here are not intended to be used as security measures.

There has been some work toward formalizing trust in the social sense as a computational concept. The work by Deutsch in 1962[8] contains a widely used definition. He states that trusting behavior occurs

when a person encounters a situation where he or she perceives an ambiguous path, the result of which can be good or bad, and the occurrence of the good or bad result is contingent on the action of another person. Additionally, the negative impact of the bad result is greater than the positive impact of the good result, so the person is motivated to make the correct choice. If this person chooses to go down the path, the person has made a trusting choice – the belief that the outcome will be good lies in the fact that the other person on whom the outcome depends is trusted to do the right thing.

Marsh[marsh94] addressed the issue of formalizing trust as a computational concept in his PhD dissertation. His model is complex and based on social and psychological factors. This work is widely cited, but the model is highly theoretical and often considered too difficult to practically implement. In this work, where trust will sometimes be derived implicitly, and, in the most expressive cases, be indicated with a numerical value, psychological factors cannot be considered.

In this work, trust is treated as a measure of uncertainty in a person or resource. Specifically, given an ambiguous path as described above, having trust in a person is defined as a measure of the confidence that the person will take the action that leads to the positive result. Reputation is synonymous with the measure of that trust.

2.3 Trust and Reputation on the Web

The issue of trust and reputation on the web has been around since the web itself began. How users could have trust in websites has taken many forms. In an unpublished masters thesis [9], Dudek studied user trust in e-commerce systems, and found that it closely related to the aesthetics of the site. More formal methods for rating the reputation of a site or of a user are also common. The eBay rating system tries to use customers positive and negative feedback ratings as a measure of a seller's reputation. Epinions, a consumer reviews website, also allows customers to rate the transactions with sellers, and maintains a more explicit trust rating system. The Epinions dataset is commonly used in the study of trust on the web. The PageRank algorithm, used by the Google search engine, also is a trust metric of sorts. It uses the number of links coming into a particular page as votes for that site. This rating, combined with other text processing, is used to score results. The PageRank algorithm is so effective at rating the relevance of pages, that its results are commonly used as a control for testing the effectiveness of trust metrics.

Explicitly creating networks and analyzing them to form opinions about the reputation of content and users has been an issue gaining more attention over the last few years, particularly with the advent of the semantic web. Unlike the hypertext web, where users had to maintain any reputation information in HTML pages, or in the backend database of a website, the semantic web is designed to let users make explicit statements about any resource, and maintain that data themselves in an open, distributed way.

The algorithms and applications presented in this work are designed to be used with semantic web based social networks, founded on the Friend-Of-A-Friend (FOAF) vocabulary[10,11]. The FOAF project defines a mechanism for describing people, and who they know. We extend that ontology by adding binary trust relations (trusts and distrusts). This work focuses on social networks, and the trust between people, but semantic trust relationships have been researched in several other contexts.

Gil and Ratnakar addressed the issue of trusting content and information sources [12]. They describe an approach to derive assessments about information sources based on individual feedback about the sources. As users add annotations, they can include measures of Credibility and Reliability about a statement, which are later averaged and presented to the viewer. Using the TRELLIS system, users can view information, annotations (including averages of credibility, reliability, and other ratings), and then make an analysis.

In the peer to peer context, the EigenTrust system [14] (based on PageRank) effectively computes global trust values for peers, based on their previous behavior. Individuals with poor performance will receive correspondingly low trust ratings. Their system was shown to be highly resistant to attack.

Raph Levin's Advogato project [19] also calculates a global reputation for individuals in the network, but from the perspective of designated *seeds* (authoritative nodes). His metric composes assertions from members to determine membership within a group. The Advogato website at <http://advogato.org>, for example, certifies users at three levels – apprentice, journeyer, and master. Access to post and edit

website information is controlled by these certifications. Like EigenTrust, the Advogato metric is quite attack resistant. By identifying individual nodes as "bad" and finding any nodes that certify the "bad" nodes, the metric cuts out an unreliable portion of the network. Calculations are based primarily on the good nodes, so the network as a whole remains secure.

Less centralized metrics are presented in work by Richardson, et al. [27]. That metric uses a local system for inferring reputations or trust relationships from a network. Richardson's work, like EigenTrust, presents a probabilistic interpretation of global belief combinations. The effectiveness of the system was shown in context with Epinions and BibServ.

3 The Reputation Inference Algorithm

3.1 Terms and Background

Different systems have used various ranges for trust and reputation ratings, including a three tiered system [19], a 1-9 rating system [13], and a continuous scale in the range [0,1] [27]. This work uses a $\{1,-1\}$ scale where 1 indicates trustworthiness, and -1 indicates untrustworthiness. This simple binary scale makes it easier to develop a theoretical understanding of the behavior of trust metrics.

If two nodes, say node A and node C, do not have a direct edge connecting them, the network can be used to generate an *inferred reputation rating* (see figure 2). If node A knows node B, and node B knows node C, then A can use the path to compose the inferred rating for C. The algorithm for inferring this value forms the basis of any trust or reputation system.




Fig 2. A reputation inference from node A to node C, can be made by following the path through node B

In real situations, people make evaluations of others based on local views of the world. To form an opinion about an unknown person, we turn to people we know and ask about the unknown person's reputation; in most cases, we do not rely on some centralized authority that determines the reputation each person deserves. One of the fundamental characteristics of this algorithm is that it is a purely local analysis. To infer a reputation rating starting at node A, the algorithm begins with A's neighbors and expands out. This means that node A may infer that node C has a very good reputation, while node D infers a poor reputation for node C. The difference will arise from the reputation ratings given along the paths from A to C and from D to C.

3.2 The Reputation Inference Algorithm

When performing an inference, the *sink* (s) is the node for which a rating is desired, and the *source* (i) is the node for whom the rating will be made. In this metric, the source polls each of the neighbors to which it has given a positive reputation rating. Neighbors with negative ratings are ignored, since their reputation means that they give unreliable information. Each of the source's trusted neighbors will return their rating for the sink. The source will then average these ratings and round the final value (i.e. a value between 0.5 and 1 rounds to a 1, while anything under 0.5 rounds down to 0). This rounded value is the inferred reputation rating from source to sink.

```

getRating(source, sink)
  numberOfNeighborsWithRatings = 0;
  sumOfRatings=0;

  mark source seen
  if sink is a direct neighbor of source
    source's rating of sink= rating(source, sink)
  else
    for each n adjacent to source
      if n is unseen and rating(source,n)>0
        if n has no ratingOfSink
          mark n seen
          inferredRating = getRating(n,sink)
          if inferredRating >= 0
            sumOfRatings +=inferredRating
            n's rating of sink = inferredRating
            numberOfNeighborsWithRatings++
          mark n unseen
        else
          if n's rating of sink >= 0
            sumOfRatings+=n's rating of sink
            numberOfNeighborsWithRatings++

    if numberOfNeighborsWithRatings > 0
      source's rating of sink=
        round(sumOfRatings/numberOfNeighborsWithRatings)

    else
      source's rating of sink = -1
  return source's rating of sink

```

Each of the source's neighbors will use this same process to come up with their reputation ratings for the sink– if there is a direct edge connecting them to the sink, the value of that edge is used; otherwise, the value is inferred. As shown in Figure 3, if a node is encountered in two paths from source to sink, it is considered in each path. Node B and C will both return ratings calculated through D. When a reputation rating from D is first requested, D will average the ratings from E and F. The value is then cached at D, so that the second time D's reputation rating is needed, no calculations are necessary.




Fig 3. An illustration of how nodes are used in the inference from node A to node G

Our system uses only 0 and 1 as reputation values. A value in the range [0,1] could be used, but the rounding in the algorithm will transform all of the values to 0 or 1 one iteration up from the nodes adjacent to the sink. Because of this, when a node polls its neighbors, it is essentially taking a majority

vote. If half or more neighbors say that the sink has a good reputation, the average works out to be half or more, which rounds up 1.

3.3 Accuracy of Inferences

The quality of an inference can be measured by how accurate it is. Accuracy can be measured by choosing a source and sink that are connected, so the accrual rating that the source gives the sink is known. This edge from the source to the sink can be removed, and then the metric can be applied to infer a rating. Comparing the inferred rating to the actual rating will measure the accuracy of the metric. If the inferred values frequently match the actual values, the metric can be called highly accurate.

Accuracy will decrease when nodes along the path from source to sink return ratings different from the source's actual rating. Some of these differences will come from nodes with which the source would normally agree. Call these "good" nodes. Inaccuracy can also arise when nodes in the path occur that the source normally would disagree with. Call these "bad" nodes. The classification of "good" or "bad" will vary for each node, depending on which node is the source.

In reality, nodes will not be so easily categorized as "good" or "bad". However, the distinction makes this analysis clearer and does not diminish the quality of the results.

Bad nodes, which will produce many more inaccurate ratings than good nodes, are the largest threat to the quality of inferences in the network. At what points will results deteriorate? We will consider the worst case where "bad" nodes are always incorrect with their reputation ratings – they will always give the opposite rating that the source would give.

With this algorithm, the inference is accurate if a majority of the nodes return correct ratings. Since the bad nodes are always incorrect, the accuracy of the good nodes must compensate. Let b be the percentage of bad nodes in the network. Then g , the percentage of good nodes, is equal to $1-b$. Let p_a be the accuracy of the good nodes – how often the reputation ratings they give a node are consistent with what the source would give the node. To obtain a correct inference from a majority vote, we need

$$g * p_a \geq 0.5 \quad (1)$$

Let $a = g * p_a$. For a given graph with n nodes, the probability that the majority of the nodes will correctly rate the sink is given by the following formula:

$$\sum_{i=\lceil \frac{n}{2} \rceil}^n \binom{n}{i} a^i (1-a)^{n-i} \quad (2)$$

When $a \geq 0.5$, the probability that our inference is correct increases as the graph size increases.

$$\lim_{n \rightarrow \infty} \sum_{i=\lceil \frac{n}{2} \rceil}^n \binom{n}{i} a^i (1-a)^{n-i} \rightarrow 1 \quad (3)$$

The central limit theorem shows that for $a > 0.5$, the limit of the sum goes to 1 as the number of nodes goes to infinity. More specifically, if nodes are accurate at least half of the time, the probability that the summation will be 0.5 or greater increases as the population size increases. This is a critical point. As long as $g * p_a$ is greater than half, we can expect to have a highly accurate inference.

The formula above only gives the probability for the base level: nodes directly adjacent to the sink. As the algorithm moves up from the immediate neighbors of the sink toward the source, a will vary from node to node, but it will increase at each level. Figure 4 illustrates this point where the network starts with all good nodes, accurate in 70% of their classifications. After moving up three levels from the sink, the accuracy of the inference will be approximately 96%.

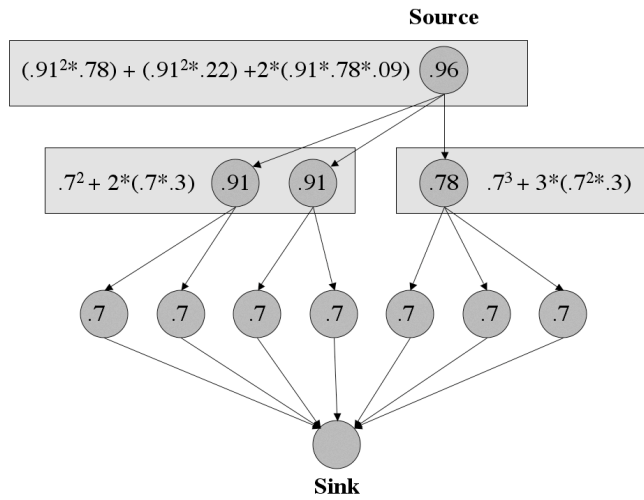


Fig 4. This figure shows a simple network and demonstrates the increasing probability of accuracy. Beginning with a uniform accuracy of 0.7, the probability of properly classifying the sink increases as we move up the search tree from sink to source. After only three levels, this source has a 96% chance of properly classifying the sink.

4 Experiments

The literature has demonstrated that both social and web-based systems exhibit small world behavior [4,6,7,17,21,23,28]. Because semantic reputation networks are essentially social networks that should have the small world properties, we used this model to automatically generate sample networks that were representative of what will occur in this domain. We used the β -graph model [30] to generate small world graphs. Our main results are on graphs with 400 nodes, but similar results were achieved for graphs with 1000 nodes.

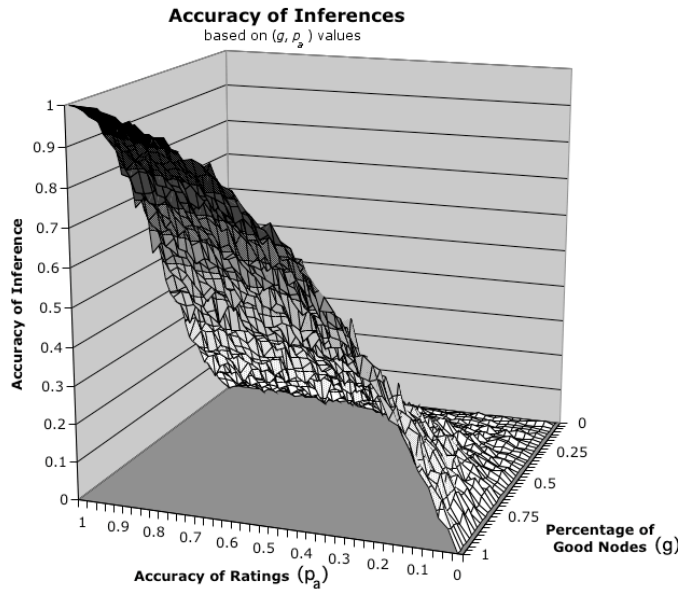



Fig 6: A surface map illustration how the accuracy of inferences changes with respect to changes in g and p_a

To test the effectiveness of this algorithm, we generated a series of small world graphs where one node was chosen as the source. We generated trust ratings from that source to every other node in the network for reference. To preserve the small world properties of the network, those ratings were not maintained as edges. A sink was randomly chosen in the graph and a rating from source to sink was inferred. The accuracy of that inference was determined by whether or not it matched the source's rating. For each edge in the graph, a reputation rating was assigned. To make this a worst case analysis,

every bad node (nodes that the sink rated a “0”) rated each neighbor opposite of how the source rated it. This makes the bad nodes’ ratings inaccurate 100% of the time. For the good nodes, ratings are correctly generated with probability p_a .

There are two variables in the network – percentage of good nodes, g , and the accuracy, p_a , of good nodes. If either variable were equal to 0, every inference would be inaccurate, so we did not include the 0 value for either variable in our experiments. Starting at 0.025 and using increments of 0.025, there are 1,600 pairs (g, p_a) . For each pair we generated 1,000 small-world graphs with 400 nodes, and one inference was made and checked for accuracy on each graph. This created a smooth picture of the surface of the space.

a (% of good nodes (g) * accuracy of good nodes (p_a))



Difference between Actual Accuracy and a




Fig 7. The top surface chart shows the value of $a (g * p_a)$, and the lower surface shows the difference between the inferred accuracy and the value predicted by a . Positive values indicate when the inferred value was higher than a , and negative values show when the inferred value was lower than a

As figure 6 shows, the accuracy of inferences remains very high, even as the accuracy of the initial ratings decreases. With no malicious nodes, inferences are 90% accurate for p_a as low as 65%. When attackers make up 10% of the population, inferences still remain over 90% accurate with p_a down to 0.75. With a population made up of 90% or more good nodes, the inferences actually remain more than 60% accurate, even when the initial ratings are below 40% accurate. Though initially this seems to contradict the result presented in the previous section – the inaccuracy should actually grow as it propagates up the network – the rounding gives us one additional benefit here. When looking at a network that comprises primarily good nodes, any ties in the voting will round up to a good rating. In a

random network, rounding up from 0.5 would be correct only 50% of the time, but since most nodes are good, rounding up is correct nearly all of the time. That means that after one level of inference, the probability of a correct inference actually increases to more than 50%. This probability will, in turn, increase in subsequent layers as described above.

As bad nodes are introduced into the network, the accuracy of inference drops off. For example, the inverse of the previous example – when 67% of the nodes are good, but the good nodes are accurate 100% of the time – does not translate to a high inference. Though the inferences are accurate 75% of the time in this case, it is nowhere close to the accuracy seen with the high number of good nodes.

These results also support the analysis described in the previous section and in Figure 4. Without that improvement, we would expect the probability of accurate inferences to be equal to the probability of an accurate initial rating – the variable a as presented above. As Figure 7 shows, for $p_a > 0.5$, the inferred value is always greater than a .

This system is *not* intended as a security system. Clearly, the error rates in this system would be unacceptable when sensitive information needs to be protected. However, its use as a system to make recommendations about the reputation of an information source is much better than most other recommendation systems [29]. In the recommendation context, our inferences are very accurate even in the face of exceptionally high uncertainty.

5 TrustMail: an Application

An ongoing project that has used these methods of reputation analysis is TrustMail [13]. TrustMail is an email client that looks up the mail sender in the reputation network and provides an inline rating for each email message. TrustMail can be configured to show trust levels for the mail sender either on a general level or with respect to a certain topic. Unlike spam filters that indicate which messages to ignore, reputation ratings in TrustMail can also tell users which messages are important to read.

Consider the case of two research groups working on a project together. The professors that head each group know one another, and each professor knows the students in her own group. However, neither is familiar with the students from the other group. If, as part of the project, a student sends an email to the other group's professor, how will the professor know that the message is from someone worth paying attention to? Since the name is unfamiliar, the message is not distinguishable from other, not-so-important mail in the inbox. This scenario is exactly the type of situation that TrustMail improves upon. The professors need only to rate their own students and the other professor. Since the reputation algorithm looks for *paths* in the graph (and not just direct edges), there will be a path from the professor one research group to students in the other through the direct professor to professor link. Thus, even though the student and professor have never met or exchanged correspondence, the student gets a high rating because of the intermediate relationship. If it turns out that one of the students is sending junk type messages, but the network is producing a high rating, the professor can simply add a direct rating for that sender, downgrading the reputation. That will not override anyone else's direct ratings, but will be factored into ratings where the professor is an intermediate step in the path.

The ratings alongside messages are useful, not only for their value, but because they basically replicate the way trust relationships and reputations work in social settings. For example, today, it would be sensible and polite for a student emailing a professor she has never met to start her email with some indication of the relationships between the student and the two professors, e.g., “My advisor has collaborated with you on this topic in the past and she suggested I contact you.” Upon receiving such a note, the professor might check with her colleague that the student's claims were correct, or just take those claims at face value, extending trust and attention to the student on the basis of the presumed relationship. The effort needed to verify the student by phone, email, or even walking down the hall weighed against the possible harm of taking the student seriously tends to make extending trust blindly worthwhile. In the context of mail, TrustMail lowers the cost of sharing trust judgments across widely dispersed and rarely interacting groups of people. It does so by gathering machine readably encoded assertions about people and their trustworthiness, reasoning about those assertions, and then presenting those augmented assertions in an end user friendly way.

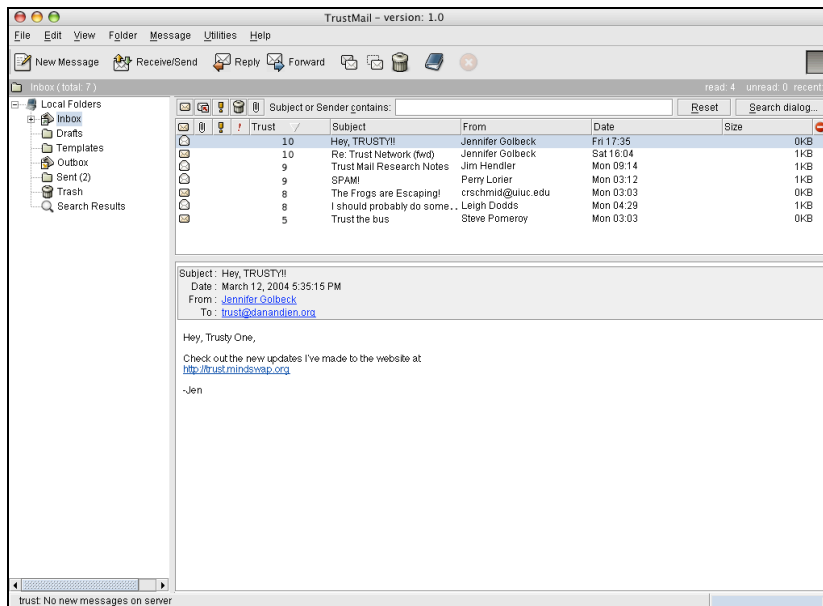


Fig 8. The TrustMail Interface

6 Future Work

We have quite a bit of work in progress to extend the results presented here. The Trust Project, at <http://trust.mindswap.org>, provides a trust ontology that does not use the $\{-1,1\}$ scale discussed here. Since users generally want a broader range of values, we provide a rating system from 1-10. Analyzing metrics for a range of values requires a variation on the theoretical approach described in this work. We are currently investigating mechanisms for performing a similar analysis, using the actual trust network built as part of the project, to measure the accuracy of new metrics and help refine them.

One of the practical limitations of this work and the Trust Project is that they require explicit trust ratings to infer a trust rating. The data available for explicit trust networks is very limited, while the FOAF project's simple social network (people connected only by a "knows" relationship) has millions of people. We are currently developing an implicit trust metric that will make reputation recommendations based on the likelihood that two individuals may be connected. This implicit metric will open trust analysis to a much wider space of data. A longer-term project will integrate the implicit metric with explicit trust metrics to use both datasets together.

7 Acknowledgements

This work was supported in part by grants from DARPA, the Air Force Research Laboratory, and the Navy Warfare Development Command. The Maryland Information and Network Dynamics Laboratory is supported by Industrial Affiliates including Fujitsu Laboratories of America, NTT, Lockheed Martin, and the Aerospace Corporation. The authors thank Aditya Kalyanpur for his comments on this paper.

The applications described in this paper are available from the Maryland Information and Dynamics Laboratory, Semantic Web Agents Project (MIND SWAP) at <http://www.mindswap.org>. You can add your own ratings to the trust network at <http://trust.mindswap.org/>.

References

1. Adamic, L., "The Small World Web". *Proceedings of ECDL*, pages 443-- 452, 1999.
2. Adding SVG Paths to Co-Depiction RDF, <http://Jibbering.com/svg/codepiction.html>
3. The Advogato Website: <http://www.advogato.org>
4. Albert, R., Jeong, H. AND Barabasi, A.-L. "Diameter of the world-wide web." *Nature* **401**, 130-131, 1999.
5. Bharat, K and M.R. Henzinger. "Improved algorithms for topic distillation in a hyperlinked environment," *Proc. ACM SIGIR*, 1998.

6. Brin, S and L. Page, "The anatomy of a large-scale hypertextual Web search engine," *Proc. 7th WWW Conf.*, 1998.
7. Broder, R Kumar, F. Maghoul, P. Raghavan, S. Rajagopalan, R. Stata, A. Tomkins, and J. Wiener. "Graph structure in the web." *Proc. 9th International World Wide Web Conference*, 2000.
8. Deutsch, Morton. 1962. "Cooperation and Trust. Some Theoretical Notes." in Jones, M.R. (ed) *Nebraska Symposium on Motivation*. Nebraska University Press.
9. Dudek, C. (2003). "Visual Appeal and the Formation of Trust in E-commerce Web Sites." Unpublished Masters Thesis, Carleton University, Ottawa, Canada.
10. Dumbill, Ed, "XML Watch: Finding friends with XML and RDF." IBM Developer Works, <http://www-106.ibm.com/developerworks/xml/library/x-foaf.html>, June 2002.
11. FOAFBot: IRC Community Support Agent: <http://usefulinc.com/foaf/foafbot>
12. Gil, Yolanda and Varun Ratnakar, "Trusting Information Sources One Citizen at a Time," *Proceedings of the First International Semantic Web Conference (ISWC)*, Sardinia, Italy, June 2002.
13. Golbeck, Jennifer, Bijan Parsia, James Hendler, "Trust Networks on the Semantic Web," *Proceedings of Cooperative Intelligent Agents 2003*, August 27-29, Helsinki, Finland.
14. Kamvar, Sepandar D. Mario T. Schlosser, Hector Garcia-Molina, "The EigenTrust Algorithm for Reputation Management in P2P Networks", *Proceedings of the 12th International World Wide Web Conference*, May 20-24, 2003, Budapest, Hungary.
15. Kleczkowski, A. and Grenfell, B. T. "Mean-fieldtype equations for spread of epidemics: The 'small-world' model." *Physica A* **274**, 355–360, 1999.
16. Kleinberg, J, "Authoritative sources in a hyperlinked environment," *Journal of the ACM*, 1999.
17. Kumar, Ravi, Prabhakar Raghavan, Sridhar Rajagopalan, D. Sivakumar, Andrew Tomkins, and Eli Upfal. "The web as a graph". *Proceedings of the Nineteenth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*, May 15-17, 2000.
18. Labalme, Fen, Kevin Burton, "Enhancing the Internet with Reputations: An OpenPrivacy Whitepaper," <http://www.openprivacy.org/papers/200103-white.html>, March 2001.
19. Levin, Raph and Alexander Aiken. "Attack resistant trust metrics for public key certification." *7th USENIX Security Symposium*, San Antonio, Texas, January 1998.
20. Milgram, S. "The small world problem." *Psychology Today* **2**, 60–67, 1967.
21. Moore, C. and Newman, M. E. J. "Epidemics and percolation in small-world networks." *Physical Review E* **61**, 5678–5682, 2000.
22. Mutton, Paul and Jennifer Golbeck, "Visualization of Semantic Metadata and Ontologies," *Proceedings of Information Visualization 2003*, July 16-18, 2003, London, UK.
23. Newman, Mark, "Models of the small world", *J. Stat. Phys.* 101, 819-841 (2000).
24. Open Privacy Initiative: <http://www.openprivacy.org/>
25. RDFWeb: FOAF: 'the friend of a friend vocabulary', <http://rdfweb.org/foaf/>
26. RDFWeb: Co-depiction Photo Meta Data: <http://rdfweb.org/2002/01/photo/>
27. Richardson, Matthew, Rakesh Agrawal, Pedro Domingos. "Trust Management for the Semantic Web," *Proceedings of the Second International Semantic Web Conference*, 2003. Sanibel Island, Florida.
28. Spertus, E, "ParaSite: Mining structural information on the Web," *Proc. 6th WWW Conf.*, 1997.
29. Swearingen, Kristen and R. Sinha. "Beyond algorithms: An HCI perspective on recommender systems," *ACM SIGIR 2001 Workshop on Recommender Systems*, New Orleans, LA, 2001
30. Watts, D. Small Worlds: The Dynamics of Networks between Order and Randomness. Princeton, NJ: Princeton University Press, 1999.