Using Trust and Provenance for Content Filtering on the Semantic Web

Jennifer Golbeck
Maryland Information Network Dynamics Lab
University of Maryland
8400 Baltimore Avenue, Suite 200
College Park, Maryland, 20740
golbeck@cs.umd.edu

Aaron Mannes
Maryland Information Network Dynamics Lab
University of Maryland
8400 Baltimore Avenue, Suite 200
College Park, Maryland, 20740
awmannes@comcast.net

ABSTRACT

Social networks are a popular movement on the web. Trust can be used effectively on the Semantic Web as annotations to social relationships. In this paper, we present a two level approach to integrating trust, provenance, and annotations in Semantic Web systems. We describe an algorithm for inferring trust relationships using provenance information and trust annotations in Semantic Web-based social networks. Then, we present two applications that combine the computed trust values with the provenance of other annotations to personalize websites. The FilmTrust system uses trust to compute personalized recommended movie ratings and to order reviews. An open source intelligence portal, Profiles In Terror, also has a beta system that integrates social networks with trust annotations. We believe that these two systems illustrate a unique way of using trust annotations and provenance to process information on the Semantic Web.

1. INTRODUCTION

Tracking the provenance of Semantic Web metadata can be very useful for filtering and aggregation, especially when the trustworthiness of the statements is at issue. In this paper, we will present an entirely Semantic Web-based system of using social networks, annotations, provenance, and trust to control the way users see information.

Social Networks have become a popular movement on the web as a whole, and especially on the Semantic Web. The Friend of a Friend (FOAF) vocabulary is an OWL format for representing personal and social network information, and data using FOAF makes up a significant percentage of all data on the Semantic Web. Within these social networks, users can take advantage of other ontologies for annotating additional information about their social connections. This may include the type of relationship (e.g. "sibling", "significant other", or "long lost friend"), or how much they trust the person that they know. Annotations about trust are particularly useful, as they can be applied in two ways. First, using the annotations about trust and the provenance of those statements, we can compute personalized recommendations for how much one user (the source) should trust another unknown user (the sink) based on the paths that connect them in the social network and the trust values along those paths. Once those values can be computed, there is a second application of the trust values. In a system where users have made statements and we have the provenance information, we can filter the statements based on how much the individual user trusts the person who made the annotation. This allows for a common knowledge base that is personalized for each user according to who they trust.

In this paper, we will present a description of social networks and an algorithm for inferring trust relationships within them. Then, we will describe two systems where trust is used to filter, aggregate, and sort information: FilmTrust, a movie recommender system, and Profiles in Terror, a portal collecting open source intelligence on terrorist activities.

WWW2006, May 22–26, 2006, Edinburgh, UK.

2. SOCIAL NETWORKS AND TRUST ON THE SEMANTIC WEB

Social networks on the Semantic Web are generally created using the FOAF vocabulary [3]. There are over 10,000,000 people with FOAF files on the web, describing their personal information and their social connections [4]. There are several ontologies that extend FOAF, including the FOAF Relationship Module [2] and the FOAF Trust Module [4]. These ontologies provide a vocabulary for users to annotate their social relationships in the network. In this research, we are particularly interested in trust annotations.

Using the FOAF Trust Module, users can assign trust ratings on a scale from 1 (low trust) to 10 (high trust).There are currently around 3,000 known users with trust relationships included in their FOAF profile. These statements about trust are annotations of relationships. There are interesting steps that can be taken once that information is aggregated. We can choose a specific user, and look at all of the trust ratings assigned to that person. With that information, we can get an idea of the average opinion about the person's trustworthiness. Trust, however, is a subjective concept. Consider the simple example of asking whether the President is trustworthy. Some people believe very strongly that he is, and others believe very strongly that he is not. In this case, the average trust rating is not helpful to either group. However, since we have provenance information about the annotations, we can significantly improve on the average case. If someone (the *source*) wants to know how much to trust another person (the *sink*), we can look at the provenance information for the trust assertions, and combine that with the source's directly assigned trust ratings, producing a result that weights ratings from trusted people more highly

than those from untrusted people.

In this section, we present an algorithm for inferring trust relationships that combines provenance information with the user's direct trust ratings.

2.1 Background and Related Work

We present an algorithm for inferring trust relationships in social networks, but this problem has been approached in several ways before. Here, we highlight some of the major contributions from the literature and compare and contrast them with our approach.

There are several algorithms that output trust inferences ([14], [8]), but none of them produce values within the same scale that users assign ratings. For example, many rely on eigenvector based approaches that produce a ranking of the trustworthiness, but the rankings do not translate to trust values in the same scale.

Raph Levin's Advogato project [9] also calculates a global reputation for individuals in the network, but from the perspective of designated seeds (authoritative nodes). His metric composes certifications between members to determine the trust level of a person, and thus their membership within a group. While the perspective used for making trust calculations is still global in the Advogato algorithm, it is much closer to the methods used in this research. Instead of using a set of global seeds, we let any individual be the starting point for calculations, so each calculated trust rating is given with respect to that person's view of the network.

Richardson et. al.[10] use social networks with trust to calculate the belief a user may have in a statement. This is done by finding paths (either through enumeration or probabilistic methods) from the source to any node which represents an opinion of the statement in question, concatenating trust values along the paths to come up with the recommended belief in the statement for that path, and aggregating those values to come up with a final trust value for the statement. Current social network systems on the Web, however, primarily focus on trust values between one user to another, and thus their aggregation function is not applicable in these systems.

2.2 Issues for Inferring Trust

When two individuals are directly connected in the network, they can have trust ratings for one another. Two people who are not directly connected do not have that trust information available by default. However, the paths connecting them in the network contain information that can be used to infer how much they may trust one another.

For example, consider that Alice trusts Bob, and Bob trust Charlie. Although Alice does not know Charlie, she knows and trusts Bob who, in turn, has information about how trustworthy he believes Charlie is. Alice can use information from Bob and her own knowledge about Bob's trustworthiness to infer how much she may trust Charlie. This is illustrated in Figure 1.

To accurately infer trust relationships within a social network, it is important to understand the properties of trust networks. Certainly, trust inferences will not be as accurate as a direct rating. There are two questions that arise which will help refine the algorithm for inferring trust: how will the trust values for intermeidate people affect the accuracy of the inferred value, and how will the length of the path affect it.

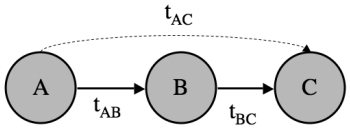

Figure 1: An illustration of direct trust values between nodes A and B (t_{AB}), and between nodes B and C (t_{BC}). Using a trust inference algorithm, it is possible to compute a value to recommend how much A may trust C (t_{AC}).

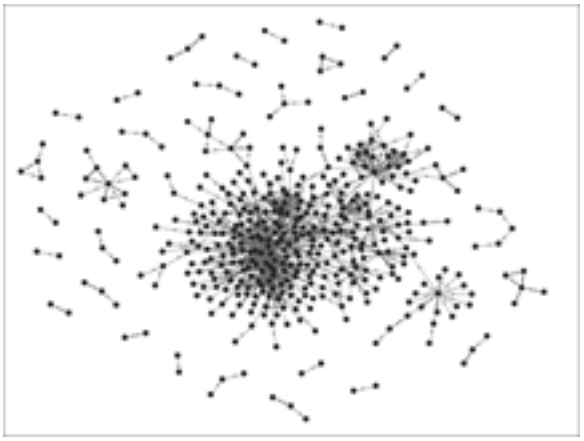

Figure 2: This figure illustrates the social network in the FilmTrust website. There is a large central cluster of about 450 connected users, with small, independent groups of users scattered around the edges.).

We expect that people who the user trusts highly will tend to agree with the user more about the trustworthiness of others than people who are less trusted. To make this comparison, we can select triangles in the network. Given nodes n_i, n_j, and n_k, where there is a triangle such that we have trust values t_{ij}, t_{ik}, and t_{kj}, we can get a measure of how trust of an intermediate person can affect accuracy. Call Δ the difference between the known trust value from n_i to n_k (t_{ik}) and the value from n_j to n_k (t_{ik}). Grouping the Δ values by the trust value for the intermediate node (t_{ij}) indicates on average how trust for the intermediate node affects the accuracy of the recommended value. Several studies [13],[4] have shown a strong correlation between trust and user similarity in several real-world networks.

It is also necessary to understand how the paths that connect the two individuals in the network affect the potential for accurately inferring trust relationships. The length of a path is determined by the number of edges the source must traverse before reaching the sink. For example, source-sink has length two. Does the length of a path affect the agreement between individuals? Specifically, should the source expect that neighbors who are connected more closely will give more accurate information than people who are further away in the network?

In previous work [4],[6] this question has been addresses

Table 1: Minimum $\overline{\Delta}$ for paths of various lengths containing the specified trust rating.

Trust Value	Path Length			
	2	3	4	5
10	0.953	1.52	1.92	2.44
9	1.054	1.588	1.969	2.51
8	1.251	1.698	2.048	2.52
7	1.5	1.958	2.287	2.79
6	1.702	2.076	2.369	2.92

using several real networks. The first network is part of the Trust Project, a Semantic Web-based network with trust values and approximately 2,000 users. The FilmTrust network[1], see Figure 2, is a network of approximately 700 users oriented around a movie rating and review website. We will use FilmTrust for several examples in this paper. Details of the analysis can be found in the referenced work, but we present an overview of the analysis here.

To see the relationship between path length and trust, we performed an experiment. We selected a node, n_i, and then selected an adjacent node, n_j. This gave us a known trust value t_{ij}. We then ignored the edge from n_i to n_j and looked for paths of varying lengths through the network that connected the two nodes. Using the trust values along the path, and the expected error for those trust values, as determined by the analysis of the correlation of trust and similarity determined in [4]. Call this measure of error Δ. This comparison is repeated for all neighbors of n_i, and for all n_i in the network.

For each path length, Table 1 shows the minimum averageΔ ($\overline{\Delta}$). These are grouped according to the minimum trust value along that path.

In Figure 3, the effect of path length can be compared to the effects of trust ratings. For example, consider the $\overline{\Delta}$ for trust values of 7 on paths of length 2. This is approximately the same as the $\overline{\Delta}$ for trust values of 10 on paths of length 3 (both are close to 1.5). The $\overline{\Delta}$ for trust values of 7 on paths of length 3 is about the same as the $\overline{\Delta}$ for trust values of 9 on paths of length 4. A precise rule cannot be derived from these values because there is not a perfect linear relationship, and also because the points in Figure 3 are only the minimum $\overline{\Delta}$ among paths with the given trust rating.

2.3 TidalTrust: An Algorithm for Inferring Trust

The effects of trust ratings and path length described in the previous section guided the development of TidalTrust, an algorithm for inferring trust in networks with continuous rating systems. The following guidelines can be extracted from the analysis of the previous sections: 1. For a fixed trust rating, shorter paths have a lower $\overline{\Delta}$. 2. For a fixed path length, higher trust ratings have a lower $\overline{\Delta}$. This section describes how these features are used in the TidalTrust algorithm.

2.3.1 Incorporating Path Length

The analysis in the previous section indicates that a limit on the depth of the search should lead to more accurate results, since the $\overline{\Delta}$ increases as depth increases. If accuracy decreases as path length increases, as the earlier analysis suggests, then shorter paths are more desirable. However, the tradeoff is that fewer nodes will be reachable if a limit is imposed on the path depth. To balance these factors, the path length can vary from one computation to another. Instead of a fixed depth, the shortest path length required to connect the source to the sink becomes the depth. This preserves the benefits of a shorter path length without limiting the number of inferences that can be made.

[1]Available at http://trust.mindswap.org/FilmTrust

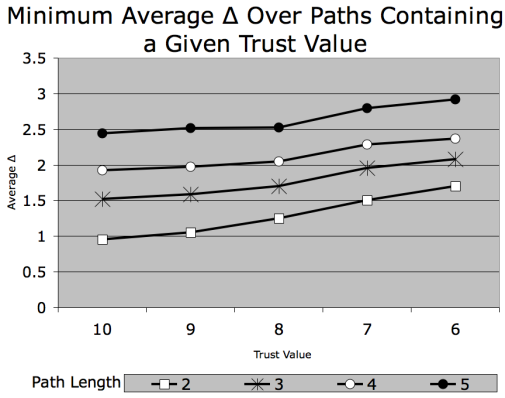

Figure 3: Minimum $\overline{\Delta}$ from all paths of a fixed length containing a given trust value. This relationship will be integrated into the algorithms for inferring trust presented in the next section.

2.3.2 Incorporating Trust Values

The previous results also indicate that the most accurate information will come from the highest trusted neighbors. As such, we may want the algorithm to limit the information it receives so that it comes from only the most trusted neighbors, essentially giving no weight to the information from neighbors with low trust. If the algorithm were to take information only from neighbors with the highest trusted neighbor, each node would look at its neighbors, select those with the highest trust rating, and average their results. However, since different nodes will have different maximum values, some may restrict themselves to returning information only from neighbors rated 10, while others may have a maximum assigned value of 6 and be returning information from neighbors with that lower rating. Since this mixes in various levels of trust, it is not an ideal approach. On the other end of possibilities, the source may find the maximum value it has assigned, and limit every node to returning information only from nodes with that rating or higher. However, if the source has assigned a high maximum rating, it is often the case that there is no path with that high rating to the sink. The inferences that are made may be quite accurate, but the number of cases where no inference is made will increase. To address this problem, we define a variable max that represents the largest trust value that can be used as a minimum threshold such that a path can be found from source to sink.

2.3.3 Full Algorithm for Inferring Trust

Incorporating the elements presented in the previous sections, the final TidalTrust algorithm can be assembled. The name was chosen because calculations sweep forward from

Table 2: $\overline{\Delta}$ for TidalTrust and Simple Average recommendations in both the Trust Project and FilmTrust networks. Numbers are absolute error on a 1-10 scale.

Algorithm		
Network	TidalTrust	Simple Average
Trust Project	1.09	1.43
FilmTrust	1.35	1.93

source to sink in the network, and then pull back from the sink to return the final value to the source.

$$t_{is} = \frac{\sum_{j \in adj(j) \mid t_{ij} \geq max} t_{ij} t_{js}}{\sum_{j \in adj(j) \mid t_{ij} \geq max} t_{ij}} \quad (1)$$

The source node begins a search for the sink. It will poll each of its neighbors to obtain their rating of the sink. Each neighbor repeats this process, keeping track of the current depth from the source. Each node will also keep track of the strength of the path to it. Nodes adjacent to the source will record the source's rating assigned to them. Each of those nodes will poll their neighbors. The strength of the path to each neighbor is the minimum of the source's rating of the node and the node's rating of its neighbor. The neighbor records the maximum strength path leading to it. Once a path is found from the source to the sink, the depth is set at the maximum depth allowable. Since the search is proceeding in a Breadth First Search fashion, the first path found will be at the minimum depth. The search will continue to find any other paths at the minimum depth. Once this search is complete, the trust threshold (max) is established by taking the maximum of the trust paths leading to the sink. With the max value established, each node can complete the calculations of a weighted average by taking information from nodes that they have rated at or above the max threshold.

2.4 Accuracy of TidalTrust

As presented above, TidalTrust strictly adheres to the observed characteristics of trust: shorter paths and higher trust values lead to better accuracy. However, there are some things that should be kept in mind. The most important is that networks are different. Depending on the subject (or lack thereof) about which trust is being expressed, the user community, and the design of the network, the effect of these properties of trust can vary. While we should still expect the general principles to be the same−shorter paths will be better than longer ones, and higher trusted people will agree with us more than less trusted people−the proportions of those relationships may differ from what was observed in the sample networks used in this research.

There are several algorithms that output trust inferences, but none of them produce values within the same scale that users assign ratings. Some trust algorithms form the Public Key Infrastructure (PKI) are more appropriate for comparison. A comparison of this algorithm to PKI can be found in [1], but due to space limitations that comparison is not included here. One direct comparison to make is to compare the $\overline{\Delta}$ from TidalTrust to the $\overline{\Delta}$ from taking the simple average of all ratings assigned to the sink as the recommendation. As shown in Table 2, the TidalTrust recommendations outperform the simple average in both networks, and these results are statistically significant with p¡0.01. Even with these preliminary promising results, TidalTrust is not designed to be the optimal trust inference algorithm for every network in the state it is presented here. Rather, the algorithm presented here adheres to the observed rules of trust. When implementing this algorithm on a network, modifications should be made to the conditions of the algorithm that adjust the maximum depth of the search, or the trust threshold at which nodes are no longer considered. How and when to make those adjustments will depend on the specific features of a given network. These tweaks will not affect the complexity of implementation.

3. USING TRUST TO PERSONALIZE CONTENT

While the computation of trust values is in and of itself a user of provenance and annotations together, the resulting trust values are widely applicable for personalizing content. If we have provenance information for annotations found on the semantic web, and a social network with trust values such that a user can compute the trustworthiness of the person who asserted statement, then the information presented to the user can be sorted, ranked, aggregated, and filtered according to trust.

In this section we will present two applications that use trust in this way. The first, FilmTrust, is a movie recommendation website backed by a social network, that uses trust values to generate predictive recommendations and to sort reviews. The second, Profiles in Terror, is a web portal that collects open source intelligence on terrorist events.

3.1 FilmTrust

The social networking component of the website requires users to provide a trust rating for each person they add as a friend. When creating a trust rating on the site, users are advised to rate how much they trust their friend about movies. In the help section, when they ask for more help, they are advised to, "Think of this as if the person were to have rented a movie to watch, how likely it is that you would want to see that film."

Part of the user's profile is a "Friends" page,. In the FilmTrust network, relationships can be one-way, so users can see who they have listed as friends, and vice versa . If trust ratings are visible to everyone, users can be discouraged from giving accurate ratings for fear of offending or upsetting people by giving them low ratings. Because honest trust ratings are important to the function of the system, these values are kept private and shown only to the user who assigned them.

The other features of the website are movie ratings and reviews. Users can choose any film and rate it on a scale of a half star to four stars. They can also write free-text reviews about movies.

Social networks meet movie information on the "Ratings and Reviews" page shown in Figure 4. Users are shown two ratings for each movie. The first is the simple average of all ratings given to the film. The "Recommended Rating" uses the inferred trust values, computed with TidalTrust on the social network, for the users who rated the film as

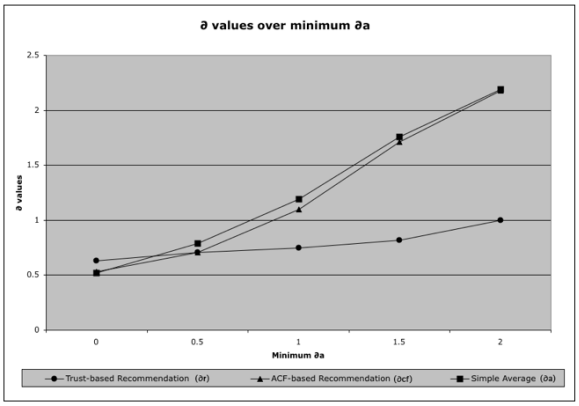

Figure 4: A user's view of the page for "A Clockwork Orange," where the recommended rating matches the user's rating, even though δa is very high (δa = 2.5).).

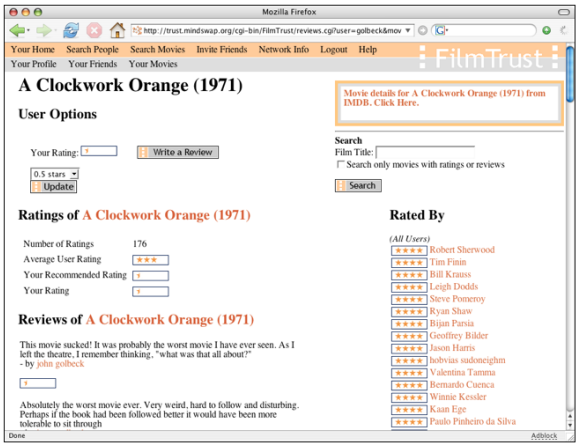

Figure 5: The increase in δ as the minimum δa is increased. Notice that the ACF-based recommendation (δcf) closely follows the average (δa). The more accurate Trust-based recommendation (δr) significantly outperforms both other methods.

weights to calculate a weighted average rating. Because the inferred trust values reflect how much the user should trust the opinions of the person rating the movie, the weighted average of movie ratings should reflect the user's opinion. If the user has an opinion that is different from the average, the rating calculated from trusted friends - who should have similar opinions - should reflect that difference. Similarly, if a movie has multiple reviews, they are sorted according to the inferred trust rating of the author. This presents the reviews authored by the most trusted people first to assist the user in finding information that will be most relevant.

3.1.1 Site Personalization: Movie Ratings

One of the features of the FilmTrust site that uses the social network is the "Recommended Rating" feature. As figure 4 shows, users will see this in addition to the average rating given to a particular movie.

The trust values are used in conjunction with the TidalTrust algorithm to present personalized views of movie pages. When the user chooses a film, they are presented with basic film data, the average rating of the movie, a personalized recommended rating, and the reviews written by users. The personalized recommended rating is computed by first selecting a set of people who rated the movie. The selection process considers trust and path length; details on how this set of people are chosen are provided in [5]. Using the trust values (direct or inferred) for each person in the set who rated the movie as a weight, and computing the weighted average rating. For the set of selected nodes S, the recommended rating r from node s to movie m is the average of the movie ratings from nodes in S weighted by the trust value t from s to each node:

$$r_{sm} = \frac{\sum_{i \in S} t_{si} r_{im}}{\sum_{i \in S} t_{si}} \quad (2)$$

This average is rounded to the nearest half-star, and that value becomes the "Recommended Rating" that is personalized for each user.

As a simple example, consider the following: Alice trusts Bob 9 Alice trusts Chuck 3 Bob rates the movie "Jaws" with 4 stars Chuck rates the movie "Jaws" with 2 stars

Then Alice's recommended rating for "Jaws" is calculated as follows:

$$\frac{t_{Alice->Bob} r_{Bob->Jaws} + t_{Alice->Chuck} r_{Chuck->Jaws}}{t_{Alice->Bob} + t_{Alice->Chuck}}$$
$$= \frac{(9*4+3*2}{9+3} = \frac{42}{12} = 3.5$$

For each movie the user has rated, the recommended rating can be compared to the actual rating that the user assigned. In this analysis, we also compare the user's rating with the average rating for the movie, and with a recommended rating generated by an automatic collaborative filtering (ACF) algorithm. There are many ACF algorithms, and one that has been well tested, and which is used here, is the classic user-to-user nearest neighbor prediction algorithm based on Pearson Correlation [7]. If the trust-based method of calculating ratings is best, the difference between the personalized rating and the user's actual rating should be significantly smaller than the difference between the actual rating and the average rating.

On first analysis, it did not appear that that the personalized ratings from the social network offered any benefit over the average. The difference between the actual rating and the recommended rating (call this δr) was not statistically different than the difference between the user's actual rating and the average rating (call this δa). The difference between a user's actual rating of a film and the ACF calculated rating (δcf) also was not better than δa in the general case. A close look at the data suggested why. Most of the time, the majority of users actual ratings are close to the average. This is most likely due to the fact that the users in the FilmTrust system had all rated the AFI Top 50 movies, which received disproportionately high ratings. A random sampling of movies showed that about 50% of all ratings were within the range of the mean +/- a half star (the smallest possible increment). For users who gave these near-mean rating, a personalized rating could not offer much benefit over the average.

However, the point of the recommended rating is more to provide useful information to people who disagree with the average. In those cases, the personalized rating should give the user a better recommendation, because we expect the people they trust will have tastes similar to their own [13].

To see this effect, δa, δcf, and δr were calculated with various minimum thresholds on the δa value. If the recommended ratings do not offer a benefit over the average rating, the δr values will increase at the same rate the δa values do. The experiment was conducted by limiting δa in increments of 0.5. The first set of comparisons was taken with no threshold, where the difference between δa and δr was not significant. As the minimum δa value was raised it selected a smaller group of user-film pairs where the users made ratings that differed increasingly with the average. Obviously, we expect the average δa value will increase by about 0.5 at each increment, and that it will be somewhat higher than the minimum threshold. The real question is how the δr will be impacted. If it increases at the same rate, then the recommended ratings do not offer much benefit over the simple average. If it increases at a slower rate, that means that, as the user strays from the average, the recommended rating more closely reflects their opinions. Figure 5 illustrates the results of these comparisons.

Notice that the δa value increases about as expected. The δr, however, is clearly increasing at a slower rate than δa. At each step, as the lower threshold for δa is increased by 0.5, δr increases by an average of less than 0.1. A two-tailed t-test shows that at each step where the minimum δa threshold is greater than or equal to 0.5, the recommended rating is significantly closer to the actual rating than the average rating is, with p¡0.01. For about 25% of the ratings assigned, δa¡0.5, and the user's ratings are about the same as the mean. For the other 75% of the ratings, δa¿0.5, and the recommended rating significantly outperforms the average.

As is shown in Figure 5, δcf closely follows δa. For δa¡1, there was no significant difference between the accuracy of the ACF ratings and the trust-based recommended rating. However, when the gap between the actual rating and the average increases, for δa¿=1, the trust-based recommendation outperforms the ACF as well as the average, with p¡0.01. Because the ACF algorithm is only capturing overall correlation, it is tracking the average because most users' ratings are close to the average.

Figure 4 illustrates one of the examples where the recommended value reflects the user's tastes. "A Clockwork Orange" is one of the films in the database that has a strong collective of users who hated the movie, even though the average rating was 3 stars and many users gave it a full 4-star rating. For the user shown, δa=2.5 - a very high value - while the recommended rating exactly matches the user's low rating of 0.5 stars. These are precisely the type of cases that the recommended rating is designed to address.

Thus, when the user's rating of a movie is different than the average rating, it is likely that the recommended rating will more closely reflect the user's tastes. When the user has different tastes than the population at large, the recommended rating reflects that. When the user has tastes that align with the mean, the recommended rating also aligns with the mean. Based on these findings, the recommended ratings should be useful when people have never seen a movie. Since they accurately reflect the users' opinions of movies they have already. Because the rating is personalized, originating from a social network, it is also in line with other results [11][12] that show users prefer recommendations from friends and trusted systems.

One potential drawback to creating recommendations based solely on relationships in the social network is that a recommendation cannot be calculated when there are no paths from the source to any people who have rated a movie. This case is rare, though, because as long as just one path can be found, a recommendation can be made. In the FilmTrust network, when the user has made at least one social connection, a recommendation can be made for 95% of the user-movie pairs.

The purpose of this work is not necessarily to replace more traditional methods of collaborative filtering. It is very possible that a combined approach of trust with correlation weighting or another form of collaborative filtering may offer equal or better accuracy, and it will certainly allow for higher coverage. However, these results clearly show that, in the FilmTrust network, basing recommendations on the expressed trust for other people in the network offers significant benefits for accuracy.

3.1.2 *Presenting Ordered Reviews*

In addition to presenting personalized ratings, the experience of reading reviews is also personalized. The reviews are presented in order of the trust value of the author, with the reviews from the most trustworthy people appearing at the top, and those from the least trustworthy at the bottom. The expectation is that the most relevant reviews will come from more trusted users, and thus they will be shown first.

Unlike the personalized ratings, measuring the accuracy of the review sort is not possible without requiring users to list the order in which they suggest the reviews appear. Without performing that sort of analysis, much of the evidence presented so far supports this ordering. Trust with respect to movies means that the user believes that the trusted person will give good and useful information about the movies. The analysis also suggests that more trusted individuals will give more accurate information. It was shown there that trust correlates with the accuracy of ratings. Reviews will be written in line with ratings (i.e. a user will not give a high rating to a movie and then write a poor review of it), and since ratings from highly trusted users are more accurate, it follows that reviews should also be more accurate.

A small user study with 9 subjects was run on the FilmTrust network. Preliminary results show a strong user preference for reviews ordered by the trustworthiness of the rater, but this study must be extended and refined in the future to validate these results.

The positive results achieved in the FilmTrust system were encouraging from the perspective of creating intelligent user interfaces. However, in other applications, filtering and rating information based on its provenance is even more critical. In the next section, we introduce the Profiles In Terror portal and present a beta version of a system that integrates trust with the provenance of information to help the user see results from the most trusted perspective.

3.2 Profiles In Terror

In the wake of the major intelligence failures of the last decade, intelligence reformers have pointed to group-think and failure of imagination as a recurring problem for intelligence agencies. A Trust Network could be an important

asset to help intelligence agencies avoid this pitfall. A trust analysis network would be an asset both to teams focused on specific problems and for the broader intelligence community. A trust network would be useful both for facilitating communication and for evaluating internal communication. Since the intelligence community of even a medium-sized nation-state could have several thousand intelligence community stake-holders (agents, collectors, policy-makers, analysts, and other intelligence consumers), all of these stakeholders cannot possibly know each other and need some means to evaluate the veracity of the information they receive. A trust network would help stakeholders identify other intelligence community members with relevant knowledge for advice and counsel. A trust network could also provide broader insight into the functioning of the intelligence community. In addition to helping stakeholders, trust systems can be useful for those doing meta-analysis on the performance of the intelligence community as a whole.

As intelligence communities are changing to face new challenges they are embracing a model of competitive collaboration. In this model divergent analyses are brought before policy-makers rather than attempting to forge a consensus. A trust network could be used to help identify and understand the data different sub-communities relied on to come to their conclusions and look at how different elements of the intelligence community view one another and their work.

In the murky world of intelligence, virtually every piece of data can be subject to dispute. Even seemingly certain information, such as date and place of birth may not be known with confidence. This problem is even more severe when more complex phenomena are being interpreted. Different units may become attached to particular theories and uninterested in alternate explanations.

The intelligence trust network would allow various stakeholders to enter a numerical rating as to their confidence in another stakeholders work, with the possibility of giving subratings for particular issues or topics (such as a particular nation or organization.) Raters would have the option of including comments. In a smaller-scale portal provenance would be assigned to the ratings and openly visible. In a large-scale portal that encompassed an entire intelligence agency, or even several agencies semi-anonymity might be necessary so that raters would feel free to contribute comments without potential repercussions. However, it would be important for stakeholders to be able contact specific raters.

For example, an analyst is assessing the stability of a regime. He comes across a report that men in the ruling family have a genetic heart defect. This was previously unknown and there is no confirmation. If it is true it has a substantial impact on the regimes stability. The analyst does not have any prior knowledge of the source, but sees that while the source has a range of ratings, there is a cluster of analysts who consistently trust this source on issues involving the regime in question. She does not know these analysts but sees from her network that some of them are well regarded by people she trusts. She contacts these analysts and learns that the source is a case officer who has recruited a high-level source within the regime who has consistently provided solid and unique information. The analyst writes her report taking this new information to account.

The trust network would allow multiple users to enter different ratings and their rationale. Within an intelligence community's trust network certain analysts and sources will gain reputations, and other stakeholders can search databases by their ratings. While the system will be able to tally and average the results, these totals may not always be strong indicators of the reliability of information or the validity of a hypothesis. In general, in trust networks, most ratings cluster together and the interesting results will be found with the outliers.

For example, tracking the movements of an individual suspected to be a major terrorist leader, an analyst comes to the conclusion that a major attack is in the works. His argument persuades several other analysts and he is given a high trust rating. When policy-makers begin examining options to capture the individual the situation become more complex. It will require substantial diplomatic efforts and could reveal sensitive sources. The policy-makers are being pressed by the analysts to move against the individual, but know that such a move will come at a high cost. While the key analyst has numerous high ratings, particularly on terrorist travel issues the policy-makers find an analyst who does not particularly trust the key analyst. The second analyst is called in to review the situation. He brings up several weaknesses in the report. The key analyst responds effectively to these points and the policy-makers move ahead with confidence to intercept the suspected terrorist.

A trust network may also help understand organizational and inter-organizational communication. This is where the ability to tally results can be useful. If a particular unit is consistently giving particularly high or low ratings to individuals in another unit it may indicate a breakdown in communications. It is possible that the two units are increasingly overlapping, but are not in direct contact, or do not understand the other group's work. The data from the trust network could indicate this deficiency and managers could take steps to correct it - by holding joint meetings or assigning the groups to joint projects. Alternately, high-ratings for the same information across several linked units might indicate group think and be a warning to management to bring in an alternate unit to "red-team" the situation.

Whether shared by a small team, an agency, or several agencies, a trust network can be a useful tool for the intelligence community. It will serve a valuable role in bringing alternate views to the attention of intelligence community stakeholders and facilitating communication between specialists in disparate agencies. Finally, it can provide an analytical basis for understanding how the intelligence community itself disseminates and analyzes information.

In the Profiles In Terror web portal, we have begun the steps to integrate trust information into the presentation of the metadata. We track provenance for each statement asserted to the portal (see figure 6. The portal also tracks probabilities associated with each statements. This means if an analyst has a piece of information, but he or she is not confident in the quality of it, they can associate a probability. In figure 6, we see a probability of 0.5 associated with the statement that Abu Mazen participated in the event Munich Olympics Massacre. We are currently integrating a trust network to the system which will combine the trust inferences discussed earlier in this paper, with provenance and probabilities in the Profiles in Terror system. This will allow statements to be filtered and ranked according to the personal trust preferences of the individual analyst.

Figure 6: A sample page from the PIT portal illustrating provenance information for a statement, as well as probabilities.

4. CONCLUSIONS AND FUTURE WORK

In this paper, we have presented a two level approach to integrating trust, provenance, and annotations in Semantic Web systems. First, we presented an algorithm for computing personalized trust recommendations using the provenance of existing trust annotations in social networks. Then, we introduced two applications that combine the computed trust values with the provenance of other annotations to personalize websites. In FilmTrust, the trust values were used to compute personalized recommended movie ratings and to order reviews. Profiles In Terror also has a beta system that integrates social networks with trust annotations and provenance information for the intelligence information that is part of the site. We believe that these two systems illustrate a unique way of using trust annotations and provenance to process information on the Semantic Web.

5. ACKNOWLEDGMENTS

This work, conducted at the Maryland Information and Network Dynamics Laboratory Semantic Web Agents Project, was funded by Fujitsu Laboratories of America – College Park, Lockheed Martin Advanced Technology Laboratory, NTT Corp., Kevric Corp., SAIC, the National Science Foundation, the National Geospatial-Intelligence Agency, DARPA, US Army Research Laboratory, NIST, and other DoD sources.

6. REFERENCES

[1] T. Beth, M. Borcherding, and B. Klein. Valuation of trust in open networks. *Proceedings of ESORICS 94.*, 1994.

[2] I. Davis and E. V. Jr. Relationship: A vocabulary for describing relationships between people. 2004.

[3] J. P. Delgrande and T. Schaub. Expressing preferences in default logic. *Artif. Intell.*, 123(1-2):41–87, 2000.

[4] J. Golbeck. *Computing and Applying Trust in Web-based Social Networks*. Ph.D. Dissertation, University of Maryland, College Park, 2005.

[5] J. Golbeck. Filmtrust: Movie recommendations using trust in web-based social networks. *Proceedings of the Consumer Communication and Networking Conference*, 2006.

[6] J. Golbeck. Generating Predictive Movie Recommendations from Trust in Social Networks. *Proceedings of The Fourth International Conference on Trust Management*, 2006.

[7] J. Herlocker, J. A. Konstan, and J. Riedl. Explaining collaborative filtering recommendations. *Proceedings of the 2000 ACM conference on Computer supported cooperative work*, 2000.

[8] S. D. Kamvar, M. T. Schlosser, and H. Garcia-Molina. The eigentrust algorithm for reputation management in p2p networks. *Proceedings of the 12th International World Wide Web Conference*, May 20-24, 2004.

[9] R. Levin and A. Aiken. Attack resistant trust metrics for public key certification. *7th USENIX Security Symposium*, 1998.

[10] M. Richardson, R. Agrawal, and P. Domingos. Trust management for the semantic web. *Proceedings of the*

Second International Semantic Web Conference, 2003.

[11] R. Sinha and K. Swearingen. Comparing recommendations made by online systems and friends. *Proceedings of the DELOS-NSF Workshop on Personalization and Recommender Systems in Digital Libraries*, 2001.

[12] K. Swearingen and R. Sinha. Beyond algorithms: An hci perspective on recommender systems. *Proceedings of the ACM SIGIR 2001 Workshop on Recommender Systems*, 2001.

[13] C.-N. Ziegler and J. Golbeck. Investigating Correlations of Trust and Interest Similarity. *Decision Support Services*, 2006.

[14] C.-N. Ziegler and G. Lausen. Spreading activation models for trust propagation. March 2004.