

Operators for Propagating Trust and their Evaluation in Social Networks

Chung-Wei Hang Yonghong Wang
Munindar P. Singh
Department of Computer Science
North Carolina State University
Raleigh, NC 27695-8206, USA

chang@ncsu.edu, yhwang_y2k@yahoo.com, singh@ncsu.edu

December 13, 2008

Abstract

Trust is a crucial basis for interactions among parties in large, open systems. Yet, the scale and dynamism of such systems make it infeasible for each party to have a direct basis for trusting another party. For this reason, the participants in an open system must share information about trust. However, they should not automatically trust such shared information. This paper studies the problem of propagating trust in multiagent systems. It describes a new algebraic approach, shows some theoretical properties of it, and empirically evaluates it on two social network datasets. This evaluation incorporates a new methodology that involves dealing with opinions in an evidential setting.

1 Introduction

We consider autonomous parties such as people or businesses that interact with each other in modern information environments. Ultimately, whenever a party interacts with another, the two parties must trust each other sufficiently to be willing to carry out the desired interaction. In a general sense, a party, Alice, trusts another party Bob when Alice places her plans in the hands of Bob.

Importantly, trust between a trusting and trusted party must have a basis in some direct *relationship* (and with respect to a relevant purpose). The relationship in question could be based on or arise from a commercial or social transaction, or through mere participation in common groups, or through an assessment of certain attributes that apply to each party. No matter what exact form the direct relationship takes, the scale of real systems is such that a party would feature in direct

relationships with a relatively small number of others. Consequently, it would have reason to trust or distrust a relatively small subset of the parties with whom it might consider interacting.

The natural response to the above challenge is to enable the parties to share information with each other about whether and how much to trust others. It stands to reason, however, that trust need not propagate. For example, Alice may trust Bob and Bob may trust Charlie (and Alice may know this fact), but it may not be the case that Alice trusts Charlie. Although the above holds in general, in many practical settings, the propagation of trust is a reasonable description of what actually transpires. In real life, individuals and businesses give referrals and rely enormously on referrals to determine with whom to interact [Frenzen and Nakamoto, 1993]. Accordingly, we confine ourselves to the practical situations that do lend themselves to the propagation of trust. Such settings arise naturally whenever similarity in the needs of the various parties is a sufficient reason for the existence of trust. For example, to anticipate one of the datasets we study, if Alice trusts Bob’s opinions about movies and Bob trusts Charlie’s opinions about movies, then it stands to reason that Alice might trust Charlie’s opinions about movies.

Trust can be naturally multidimensional, including various aspects of competence and intent. Multidimensionality can explain the evidence pro and con regarding the propagation of trust. For example, if Alice trusts Bob as an IT specialist and Bob trusts Charlie as a plumber, Alice may yet not trust Charlie as a plumber because she doesn’t trust Bob to be a good judge of plumbers. The approach we describe below can accommodate multidimensionality but, for the sake of simplicity, we consider a single dimension and assume that it incorporates the aspect of trust in the sense of judging others and giving valuable referrals. Besides its simplicity, an important reason for this limited model is that the independently existing datasets upon which we empirically evaluate our approach do not lend themselves to multidimensional assessments.

We restrict our attention to situations where the various parties cannot willfully change their behavior to exploit others. For example, a hotel would treat all guests alike (modulo random variations). A hotel cannot install or remove a beautifying water fountain for different guests.

Further, we generalize the conceptualization of evidence here to mean not necessarily empirical but any basis of positive and negative opinions, respectively. This is important for our application domain of social networks, where the base relationships are based on opinions (even prejudices), not on empirical evidence.

There are three contributions in this paper. First, we define three operators, *aggregation*, *concatenation*, and *selection*, for efficiently and accurately propagating trust in social networks. Next, we prove some useful formal properties of the operators. Third, in order to evaluate our approach, we present a transformation from subjective opinions to objective (evidential) ratings.

2 Background

From the standpoint of trust, a system of interacting agents is naturally modeled as a weighted directed graph, each of whose vertices correspond to an agent and each of whose edges corresponds to a direct relationship of trust from the agent at the source vertex to the agent at the target vertex, the weight on the edge being a measure of the trust placed. Conventionally, this weight is a scalar

from the real interval $[0, 1]$. A path in such a graph thus corresponds to propagated trust and a weight on a path can be induced that reflects the measure of the trust propagated.

The earliest works on the propagation of trust assume that the propagation of trust along a path is multiplicative (e.g., the weighted average approach in [Golbeck, 2005]). For example, if Alice trusts Bob at 0.8 and Bob trusts Charlie at 0.7, then Alice is inferred as trusting Charlie at $0.8 \times 0.7 = 0.56$. This reflects the natural intuition that the measure of trust placed over a path falls as a path gets longer. (Our more nuanced approach captures this intuition as well.)

How trust is aggregated from more than one path is more challenging. An intuition is that the contributions of independent paths be added in some way, because they would reinforce each other. However, a simple addition of the path weights is generally inappropriate, because it would lead to double counting, and yield trust ratings greater than 1.0 unless somehow corrected.

Another challenge to trust propagation is the well-known *rumor problem*. What we would like to avoid are situations such as where Alice trusts Charlie because Bob told her he does and Bob trusts Charlie because Alice told him she does. We would like to make sure that there are no cycles in the flow of information. This is the reason that several trust propagation approaches represent only direct trust information in the graph models and use that information as the base for all their calculations of trust.

A common feature of most of the popular approaches for trust is that they measure the extent of trust via a scalar [Gray et al., 2003; Katz and Golbeck, 2006; Levien, 2003; Quercia et al., 2007; Ziegler and Lausen, 2004]. Although the scalar representation is simple, it is not well-suited to the propagation of trust. Following Jøsang, we model one party's trust in another party in terms of a triple consisting of three scalars corresponding to belief (i.e., positive trust) b , disbelief (i.e., distrust or negative trust) d , and uncertainty u . However, each triple $\langle b, d, u \rangle$ satisfies the constraint that $b + d + u = 1$ —hence, this representation supports not three but two degrees of freedom. The additional degree of freedom compared to the traditional one-scalar representation helps us represent the certainty of information. In Jøsang's approach, the above triple representation is derived from a pair of numbers representing the positive evidence r and the negative evidence s , respectively, where $r + s > 0$. The $\langle b, d, u \rangle$ and $\langle r, s \rangle$ representations—called *belief* and *evidence* spaces, respectively—can be mapped to each other without loss of information. In notation, $\langle b, d, u \rangle = Z(\langle r, s \rangle)$ and $\langle r, s \rangle = Z^{-1}(\langle b, d, u \rangle)$, where Z is a suitable mapping [Wang and Singh, 2007]. Below, we use M_i to refer to a belief report of the form $\langle b, d, u \rangle$.

Wang and Singh (W&S) [2006], [2007] adopt Jøsang's framework [Jøsang, 1998], but deviate significantly from his approach. W&S base certainty on the strength, and not just the amount of evidence. First, for a given amount of evidence, increasing unanimity yields higher certainty. Second, holding the extent of the unanimity (or conflict) constant, an increasing amount of evidence yields increasing certainty.

Certainty provides a principled basis for propagating trust and distrust. For example, if Alice trusts Bob highly and Bob trusts Charlie highly, Alice would trust Charlie almost as highly. And, if Alice trusts Bob highly and Bob distrusts Charlie highly, Alice would distrust Charlie almost as highly as Bob does. In other words, even the propagation of distrust relies upon trust in the propagator. By contrast, if Alice distrusts Bob highly, then whether Bob trusts or distrusts Charlie, Alice would be uncertain about Charlie. It is impossible to capture the above nuances with a scalar

representation.

3 Approach

Building on the above model of a social network as a graph, we propose a model called *CertProp* that handles the propagation of trust. *CertProp* is based on three operators. Concatenation or \otimes deals with the propagation of trust ratings along a path. Selection or \textcircled{S} chooses the most trustworthy path to each witness, whereas aggregation of \oplus deals with the combination of trust ratings from paths between the same source and target. The aggregation operator is due to Jøsang [1998] as modified in [Wang and Singh, 2006]; the others are new here.

3.1 New Operator: Concatenation

Suppose Alice learns about Charlie from Bob. The extent of trust that Alice places in Charlie would be Bob’s trust in Charlie discounted by Alice’s trust in Bob.

Let agent A place trust $M_1 = \langle b_1, d_1, u_1 \rangle$ in agent B ’s references and B place trust $M_2 = \langle b_2, d_2, u_2 \rangle$ in agent C . Then A ’s trust in C due to the reference from B can be calculated by the concatenation $M_1 \otimes M_2$, defined as follows.

Definition 1 Concatenation operator \otimes . Suppose $M_1 = \langle b_1, d_1, u_1 \rangle$ and $M_2 = \langle b_2, d_2, u_2 \rangle$ are two belief functions, we define $M = M_1 \otimes M_2 = \langle b, d, u \rangle = Z(\langle b_1 r_2, b_1 s_2 \rangle)$, where $\langle r_2, s_2 \rangle = Z^{-1}(M_2)$.

3.2 Aggregation

The aggregation operator combines bodies of evidence. Consider a situation where Alice learns about David separately from both Bob and Charlie, and has factored in her trust for Bob and Charlie according to the above concatenation operator. What should be Alice’s combined trust in David? The answer depends upon the aggregation of the ratings obtained from two paths.

Definition 2 Aggregation operator \oplus . Let $M_1 = \langle b_1, d_1, u_1 \rangle$ and $M_2 = \langle b_2, d_2, u_2 \rangle$ describe ratings computed from two paths between the same source and target. Then $M_1 \oplus M_2 = Z(Z^{-1}(M_1) + Z^{-1}(M_2))$

3.3 Applying Concatenation and Aggregation

Consider an agent A with neighbors $\{B_1, \dots, B_m\}$, which A trusts M_1, \dots, M_m , respectively. Let the B_i trust the target agent C be M'_i . Then we can infer the trust placed by A in C , M as

$$M = (M_1 \otimes M'_1) \oplus (M_2 \otimes M'_2) \oplus \dots \oplus (M_m \otimes M'_m)$$

If a neighbor has not yet computed its trust in C , we can run the algorithm recursively to obtain the trust from merging and combining the trust from the neighbor’s neighbors, and so on, until we

get to the witnesses whose trust values in C are computed from their direct interactions with C . So the trust ratings can be merged in a bottom up fashion, from the leaves of the trust network up to its root A .

3.4 New Operator: Selection

The above approach in essence enumerates all paths from the originating agent to the target, and then uses the concatenation operator to combine beliefs along each path, the aggregate operator to combine beliefs from all paths. But there is a problem by doing it this way. For example, in Figure 1, witness W_i tells agent A that he has 1 positive and 1 negative experiences with the target. He tells the same information to agent B_2 . Both B_1 and B_2 are agent A 's neighbors. A should not use the aggregation operator \oplus to combine the beliefs from B_1 and B_2 , since their beliefs come from the same source.

So how may we combine beliefs which come from the same source, but propagate via different paths? We define a new operator, called *selection* operator, \textcircled{S} , to select one out of multiple paths that end at the same point. We select the path that offers the highest belief. Using any other path would lose valuable information. Adding a path would cause double counting. In this example, we have two paths: $P_1 = A \rightarrow B_1 \rightarrow W_i$, and $P_2 = A \rightarrow B_2 \rightarrow W_i$. Let A 's trust in B_1 and B_2 be M_1 and M_2 , respectively, and B_1 and B_2 's trust in W_i be M'_1 and M'_2 , respectively. Suppose $M_1 \otimes M'_1 = \langle 0.6, 0.3, 0.1 \rangle$, and $M_2 \otimes M'_2 = \langle 0.5, 0.25, 0.25 \rangle$. We pick the most reliable path from A to W_i . Since $M_1 \otimes M'_1 > M_2 \otimes M'_2$ ($0.6 > 0.5$), so P_1 is the most reliable path. So $(M_1 \otimes M'_1) \textcircled{S} (M_2 \otimes M'_2) = M_1 \otimes M'_1$.




Figure 1: Selection operator \textcircled{S}

Definition 3 *Selection operator \textcircled{S} .* Let P_1 and P_2 be two paths from A to W . Let $M_1 = \langle b_1, d_1, u_1 \rangle$ be the trust concatenated along P_1 , $M_2 = \langle b_2, d_2, u_2 \rangle$ be the trust concatenated along P_2 . Then if $b_1 \geq b_2$, then $M_1 \textcircled{S} M_2 = M_1$, otherwise $M_1 \textcircled{S} M_2 = M_2$.

3.5 Applying Concatenation, Aggregation, and Selection

For a given trust network, to propagate trust of A with C , we combine beliefs as following. Using selection, we find the best path (to a fixed depth) from A to each witness. Then we concatenate beliefs along this path. The resulting belief is the consolidated belief supported by that witness. Then we combine all beliefs supported by all witnesses by using aggregation.

For example, in Figure 2, by applying all three operators, the propagated trust is $M = ((M_1 \otimes M'_1) \textcircled{\text{S}} (M_2 \otimes M'_2) \textcircled{\text{S}} (M_3 \otimes M'_3)) \otimes M_5 \oplus (M_4 \otimes M'_4)$.

To illustrate the double counting problem without selection in Figure 2, let $M_1 \otimes M'_1 = \langle 0.8, 0.1, 0.1 \rangle$, $M_2 \otimes M'_2 = \langle 0.6, 0.2, 0.2 \rangle$, $M_3 \otimes M'_3 = \langle 0.6, 0.3, 0.1 \rangle$, $Z^{-1}(M_5) = \langle 50, 5 \rangle$, $M_4 = \langle 0.9, 0.05, 0.05 \rangle$, and $Z^{-1}(M'_4) = \langle 20, 0 \rangle$. The actual trust of C is M , where $Z^{-1}(M) = Z^{-1}(M_5) \oplus Z^{-1}(M'_4) = \langle 70, 5 \rangle$. However, by applying only concatenation and aggregation, the estimated trust $Z^{-1}(M) = Z^{-1}(((M_1 \otimes M'_1) \oplus (M_2 \otimes M'_2) \oplus (M_3 \otimes M'_3)) \otimes M_5) \oplus (M_4 \otimes M'_4) = \langle 118, 10 \rangle$, which double-counts M_5 . This shows that multiple paths via the same witness can lead to double-counting. This is especially obvious when the aggregated belief of the paths from the source to a witness is greater than one.




Figure 2: Trust propagation in a social network

4 Properties

To propagate trust meaningfully and efficiently, we desire certain properties of the three operators. CertProp follows Jøsang [1998] and W&S's[2006] aggregation (based on Jøsang [1998] but with a different Z), but different concatenation and selection. W&S showed that aggregation is associative and commutative. Now we establish some useful properties of concatenation and selection.

Theorem 1 *Selection $\textcircled{\text{S}}$ is commutative.*

Proof: We have to show $M_1 \textcircled{\text{S}} M_2 = M_2 \textcircled{\text{S}} M_1$. Let $M_1 = \langle b_1, d_1, u_1 \rangle$, $M_2 = \langle b_2, d_2, u_2 \rangle$, and $M_1 \textcircled{\text{S}} M_2 = M_1$, which means $b_1 > b_2$. Then $M_2 \textcircled{\text{S}} M_1 = M_1$.

Theorem 2 *Selection $\textcircled{\text{S}}$ is associative.*

Proof: We have to show $M_1 \textcircled{\text{S}} (M_2 \textcircled{\text{S}} M_3) = (M_1 \textcircled{\text{S}} M_2) \textcircled{\text{S}} M_3$. Let $M_1 = \langle b_1, d_1, u_1 \rangle$, $M_2 = \langle b_2, d_2, u_2 \rangle$, and $M_3 = \langle b_3, d_3, u_3 \rangle$.

Case 1 : Suppose b_1 is larger than b_2 and b_3 . If $b_2 \geq b_3$. Then $M_1 \textcircled{\text{S}} (M_2 \textcircled{\text{S}} M_3) = M_1 = M_1 \textcircled{\text{S}} M_3 = (M_1 \textcircled{\text{S}} M_2) \textcircled{\text{S}} M_3$. If $b_2 < b_3$, then $M_1 \textcircled{\text{S}} (M_2 \textcircled{\text{S}} M_3) = M_1 \textcircled{\text{S}} M_3 = M_1 = M_1 \textcircled{\text{S}} M_3 = (M_1 \textcircled{\text{S}} M_2) \textcircled{\text{S}} M_3$.

Case 2 : Suppose b_2 is larger than b_1 and b_3 . Then $M_1 \textcircled{S} (M_2 \textcircled{S} M_3)$
 $= M_1 \textcircled{S} M_2 = M_2 = M_2 \textcircled{S} M_3 = (M_1 \textcircled{S} M_2) \textcircled{S} M_3.$

Case 3 : Suppose b_3 is larger than b_1 and b_2 . Then $M_1 \textcircled{S} (M_2 \textcircled{S} M_3)$
 $= M_1 \textcircled{S} M_3 = M_3 = (M_1 \textcircled{S} M_2) \textcircled{S} M_3.$

Thus, $M_1 \textcircled{S} (M_2 \textcircled{S} M_3) = (M_1 \textcircled{S} M_2) \textcircled{S} M_3.$

Theorem 3 Concatenation \otimes distributes over aggregation \oplus .

Proof: We need to prove that $M_1 \otimes (M_2 \oplus M_3) = (M_1 \otimes M_2) \oplus (M_1 \otimes M_3)$. Assume the equivalent evidence corresponding to M_2 and M_3 is $\langle r_2, s_2 \rangle$ and $\langle r_3, s_3 \rangle$, respectively. Let $M_1 = \langle b_1, d_1, u_1 \rangle$. Then the equivalent evidence corresponding to $M_1 \otimes (M_2 \oplus M_3)$ is $\langle b_1(r_2 + r_3), b_1(s_2 + s_3) \rangle$. The evidence corresponding to $(M_1 \otimes M_2) \oplus (M_1 \otimes M_3)$ is $\langle b_1(r_2), b_1(s_2) \rangle \oplus \langle b_1(r_3), b_1(s_3) \rangle = \langle b_1(r_2 + r_3), b_1(s_2 + s_3) \rangle$, which completes the proof.

Conjecture 1 Concatenation \otimes distributes over selection \textcircled{S} .

5 Evaluating Trust Propagation

It is difficult to evaluate an approach on real data such as ours because large networks of agents interacting with one another do not exist. For this reason, we adapt two datasets of social networks, namely, FilmTrust and a PGP key ring. Each dataset is naturally modeled as a weighted directed graph.

We use the following strategy for evaluating a trust propagation algorithm over a graph. Typical networks carry a lot of redundancy, which we can exploit to evaluate the effectiveness of an approach for propagation [Katz and Golbeck, 2006]. The information associated with any specific edge may be induced from the other relevant edges, namely, those that fall on a sufficiently small path from the source to the target (of the given edge).

Specifically, let there be an edge from agent A to agent B of weight d_{AB} (denoting actual trust). We remove this edge temporarily and estimate the propagated trust i_{AB} between A and B based on paths from A to B . The difference between d_{AB} and i_{AB} reflects how effective an algorithm is at inferring the relationship between two agents. In essence, the elided direct relationship of weight d_{AB} yields the ground truth with which to evaluate the propagation.

5.1 Accuracy Metrics

To compare the propagated trust with the actual trust (the weight of the direct edge), we introduce two metrics: P -error and B -error, which are defined in evidence and belief space, respectively. Let $M_1 = \langle b_1, d_1, u_1 \rangle$, where $Z^{-1}(M_1) = \langle r_1, s_1 \rangle$, and $M_2 = \langle b_2, d_2, u_2 \rangle$, where $Z^{-1}(M_2) = \langle r_2, s_2 \rangle$. The P -error between M_1 and M_2 is $|\alpha_1 - \alpha_2|$, where $\alpha = \frac{r}{r+s}$. The B -error of M_1 and M_2 is defined as $|b_1 - b_2|$. Importantly P -error provides a metric for comparing with single-valued trust representations like TidalTrust [Katz and Golbeck, 2006], whereas B -error yields more accurate comparison because it considers the certainty of the trust ratings. B -error is more sensitive because it is low if either (a) two certainties are close, or (b) P -error is low.

5.2 Model Variants

Note that we find paths within a fixed length for the following reasons. First, it is computationally expensive to find all paths in a huge social graph. Besides, shorter paths yield better accuracy in general because longer chains are weaker [Katz and Golbeck, 2006; Yu et al., 2003]. However, a trade-off still exists. Deeper search may yield more evidence, but takes more time. Shallow search may give us accurate evidence quickly, but may not find any path. Thus, we propose three strategies. The *shortest* strategy first finds the shortest path from A to B , and then find all paths within that length. The *fixed* strategy searches all paths within a specified depth. In our experiments, we set this depth to seven to make sure we find at least one path for all connected pairs in the dataset. The *selection* strategy yields the most trusted paths to each witness found by *fixed*. We compare our approach, CertProp, to W&S [2006]. To show the influence of models and the strategies, we define three variants for both W&S and CertProp, which Table 1 summarizes.

Model Name	\oplus	\otimes	\textcircled{S}	Path
<i>W&S</i> (shortest)		W&S	No	Shortest
<i>W&S</i> (fixed)	W&S	W&S	No	Fixed
<i>W&S</i> (selection)		W&S	Yes	Fixed
<i>CertProp</i> (shortest)		New	No	Shortest
<i>CertProp</i> (fixed)	W&S	New	No	Fixed
<i>CertProp</i> (selection)		New	Yes	Fixed

Table 1: Variants of W&S and CertProp studied here

5.3 From Opinions to Evidence

Using the available social network datasets poses two challenges. One, the weights used in these networks are scalars whereas our approach needs two reals, namely, $\langle b, d, u \rangle$ (with $b + d + u = 1$). Two, the weights are subjective opinions and not evidence. Accordingly we propose heuristics for mapping opinions to evidence, so they become amenable to our approach. We propose two approaches to transform opinions into evidence: (1) *linear* and (2) *Weber-Fechner*. We consider subjective ratings drawn from a scale such as 1 to 10 (FilmTrust), 1 to 4 (PGP), and so on.

The idea of the *linear* transformation is normalization. In our representation, the ratings at the end points (intuitively, reflecting unanimity of various considerations) correspond to a lower uncertainty u than those in the middle. Further, the belief b derived from a rating of 10 should be the highest and that derived from a rating of 1 the lowest. Therefore, we translate a FilmTrust rating to our trust value $\langle r, s \rangle$ by simply interpreting the single number as the number of positive experiences r relative to a fixed total number of experiences of 10. For example, we translate an opinion rating of 4 to evidence $\langle r, s \rangle = \langle 4, 6 \rangle$. Likewise, we translate a PGP rating of 3 to $\langle 3, 1 \rangle$. Although this approach is simplistic, it provides us the metrics to compare with other trust propagation methods.

The *Weber-Fechner* transformation satisfies two observations. It is based on the *Weber-Fechner* “law” [Web], which says the relationship between stimulus and perception is logarithmic. If a stimulus (i.e., good experience) is tripled in strength, the corresponding perception (i.e., opinion ratings) will be three steps above the original value. For example, suppose agent A has 10 good transactions with agent B , and A ’s opinion rating about B is three. A ’s opinion rating about B will be four if A has 20 more good transactions. Second, as always, in our approach, the certainty of the object rating corresponding to the average opinion should be the lowest. Based on these observations, we define a transformation function to transform opinions into evidence.

Definition 4 Weber-Fechner Transformation. *Suppose Σ is the set of normalized opinion ratings. Let A , P , U be the average, the most popular, and the least popular opinion in Σ , respectively. Then $W(\sigma \in \Sigma) = \langle k^{\sigma/A}, k \rangle$ is the transformation from opinions into evidence, where $k = \ln \left(\frac{\#P - \#U}{|P - U|} + e \right)$, $\#P$ and $\#U$ are the counts of P and U .*

The intuition behind k is to capture the slope between the numbers of the most and least popular opinions. If k is larger, the difference of the probabilities between the transformed ratings corresponding to two successive opinions will be bigger. For example, the number of the most and least popular opinions in FilmTrust and PGP are 199, 308 and 202, respectively, which means it requires more positive evidence to increase the opinions by one in PGP than in FilmTrust. We add a correction e to make sure k is greater than one. The exponent σ/A captures the logarithmic relationship between evidence and ratings.

For example, in the FilmTrust dataset, the average of normalized opinions is 0.68, the most and least popular normalized opinions are 0.7 and 0.2, which have 228 and 26 counts, respectively. Then, $W(0.3) = \langle 2.21, 6.01 \rangle$, and $W(0.7) = \langle 6.33, 6.01 \rangle$.

Figure 3 compares the certainties of *linear* and *Weber-Fechner* transformed ratings in the FilmTrust and PGP dataset. With *Weber-Fechner*, the average rating has the lowest certainty, but not with *linear* does not (in PGP). In FilmTrust, the average normalized rating is close to 0.7, while the average rating in the PGP dataset is close to 1, which is the default value in PGP setting. Sections 5.4 and 5.5 show how these transformations affect the results.

5.4 FilmTrust

Our first evaluation is based on *FilmTrust* [Katz and Golbeck, 2006; Kuter and Golbeck, 2007], a small social network of film buffs. In this dataset, the 538 vertices represent agents (users) and the 1,234 directed edges represent their trust relationships with other users. The weight of an edge is an integer in $[1, 10]$, which reflects the strength of the source agent’s qualitative rating or opinion of the target agent.

Katz and Golbeck’s [2006] algorithm, TidalTrust, collects trust data from all referral paths with the shortest length from a source to a sink. It selects referral paths with strength above a threshold and uses them to compute the overall trust value.

Kuter and Golbeck’s [2007] trust inference model, *Sunny*, provides a confidence measurement based on probabilistic sampling. *Sunny* exhaustively finds all possible paths from a source to




Figure 3: Rating distributions and associated certainties

a sink. The confidence measurements in Sunny, W&S, and CertProp, reduce the influence of path length—longer paths contribute less to the overall result than shorter ones. But considering additional paths provides access to more opinions, and can yield more accurate inference.

Figure 4 shows the paths found by the *fixed* variant, which exhaustively searches all paths within the fixed length three from the source to a sink (diamond), after the directed edge (dashed arrow) is removed. In this case, we are left with two paths of length two and 13 paths of length three. The *shortest* variant only finds the two paths in the shortest length, which is two. Conversely, the *fixed* strategy considers all 15 paths of up to a fixed length of three. Our experiments below show how search strategies affect the accuracy of the propagation.




Figure 4: Paths of length less than or equal to three from a source (solid rectangle) to a sink

Figure 5 shows the average *P-errors* of TidalTrust, Sunny, W&S, and CertProp. The *shortest* variant has similar performance to Sunny, which outperforms TidalTrust. The *fixed* variant has the best performance with both W&S and CertProp.

To provide more insightful comparisons among the variants of W&S and CertProp, Figure 6 shows both the *P-error* and *B-error* with different search strategies. We draw three conclusions. First, the *fixed* variant also yields higher certainty (reflecting the additional evidence found) than the *shortest*, which indicates the former provides better estimates by considering more evidence (paths). Next, although CertProp and W&S have similar *P-errors* in both *linear* and *Weber-Fechner*




Figure 5: Average P -errors evaluated on FilmTrust

cases, CertProp has a lower B -error (more accurate). This is because CertProp has better concatenation operator, which is further discussed in Section 5.6. Third, both CertProp and W&S yield better performance in the *Weber-Fechner* case. This result shows that the *Weber-Fechner* transformation successfully reduces the subjectivity in FilmTrust. The subjectivity of the opinions decreases the accuracy of CertProp and W&S.




Figure 6: Average P -errors and B -errors with different search strategies, evaluated on FilmTrust with *linear* and *Weber-Fechner* transformations

To show more evidence of how the amount of evidence affects prediction accuracy, we use random walk sampling [Leskovec and Faloutsos, 2006] to generate five 25% samples. The sampling method preserves the properties of the graph, such as degree distribution, clustering coefficient, and so on, with a moderate sampling size of 15% to 25%. As Figure 7 shows, the confidence of

the prediction increases with the number of paths found. Also the average ratings, the depth of the paths found, and the P -errors from the sample data are similar to the results from the whole FilmTrust data. We apply the same sampling method in Section 5.5 to reduce the size of the PGP dataset.




Figure 7: The results of CertProp (fixed) on the sampling data, showing the certainty increases with the number of the paths found, while the average rating, the depth of paths found, and P -error are preserved by the random-walk sampling method

5.5 PGP Web of Trust

The *web of trust* is a concept used in Pretty Good Privacy (PGP). The main idea is that instead of relying on centralized certificate authorities, the web of trust establishes a decentralized trust model of public keys, in order to verify the authenticity of the bindings between users and their public keys.

For example, user C receives a digitally signed email from user S . C needs S 's public key to verify the digital signature. However, a fake public key can be easily created with S 's name. To verify the authenticity of the public key, one simple way is to find another trustworthy user who can confirm the public key belongs to S . In other words, the authenticity of the public key of S can be verified if C can find a user who signs the public key of S . This is called a *signature relation*. The web of trust is a directed graph representing the signature relation among users. Each vertex represents a user, and an edge from user A to user B means the public key of B is signed by user A . Besides, an integer trust value in $[1, 4]$ is associated with each edge to indicate the strength of the relation. To verify the authenticity of S , C needs to find a path (a confirmation chain of signature relations) to reach S . There could be more than one path between two users. The more disjoint paths exist the harder it is to fake the confirmation chain.

We repeat the experiment of Section 5.4 on the PGP dataset. The web of trust data is collected from the PGP server snapshot [wot] on June 5, 2008. It consists of 39,246 vertices and 317,979

edges. The average distance (i.e., the shortest path) of all vertex pairs is around 6. We sample the graph to scale down the size of the data using Random Walk sampling [Leskovec and Faloutsos, 2006].

We compare CertProp with W&S and a trust propagation approach, *Naive*. In Naive trust propagation, concatenation is defined as the multiplication of the probabilities $\alpha = \frac{r}{r+s}$. Aggregation is defined by the average of the probabilities. There is no certainty concept in Naive.

Figure 8 compares CertProp with W&S and Naive. Just as for FilmTrust, CertProp has similar performance with respect to *P-errors*, but outperforms W&S in *B-errors*. Naive is the worst among the three approaches. Also, in *Weber-Fechner* case, both CertProp and W&S improve, but the concatenation of CertProp yields greater improvement than W&S.

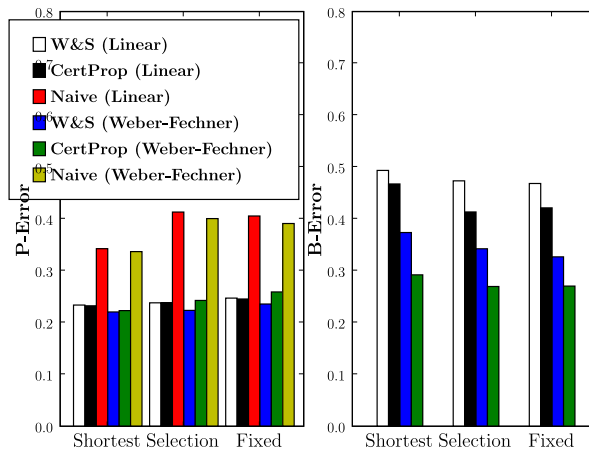


Figure 8: Average *P-errors* and *B-errors* of W&S, Naive, and CertProp, with different search strategies, evaluated on PGP with *linear* and *Weber-Fechner* transformations

5.6 Discussion

We can draw the following conclusions from the evaluation in Sections 5.4 and 5.5.

Overall Performance of CertProp and W&S CertProp (*fixed*) is the best trust propagation approach in both datasets. In general, CertProp is more accurate and efficient than W&S in terms of *B-error*, although they have similar *P-errors*. As mentioned in Section 5.1, the accuracy of probability-certainty trust model is better estimated by *B-error*. Besides, CertProp and W&S both have better *P-errors* than Sunny and TidalTrust in FilmTrust, and Naive in PGP.

Performance of Search Strategies The *fixed* variant is generally more accurate than *shortest*, whereas the *selection* variant is outperformed by *fixed* in FilmTrust but has similar performance as *fixed* in PGP, depending on the properties of the networks. In most cases, *selection* does not perform well because it ignores too much evidence.

Performance of \otimes in CertProp and W&S The main difference between CertProp and W&S is in their concatenation operators. CertProp concatenation discounts the evidence of the referral by the belief of the referrer, whereas W&S concatenation discounts the belief of the referrer by the belief of the referrer. The certainty of concatenated trust in CertProp is generally higher than in W&S, because the belief can be viewed as certainty-discounted evidence. The results in Sections 5.4 and 5.5 indicate the concatenation in CertProp has better accuracy than W&S.

Further, CertProp concatenation works evidence space. In other words, it requires no Z^{-1} mapping [Wang and Singh, 2006] (as described in Section 2). To show how Z^{-1} mapping affects the efficiency, we compute the concatenation of two trust ratings, M_1 and M_2 , where $Z^{-1}(M_1) = \langle r_1, s_1 \rangle$ and $Z^{-1}(M_2) = \langle r_2, s_2 \rangle$. Each of r_1, s_1, r_2, s_2 is an integer ranging over $[1, 100]$. We consider all combinations represented by *Cartesian product* $r_1 \times s_1 \times r_2 \times s_2$. Thus, there are total 100 million combinations. Generally, the Z^{-1} mapping can be approximated using binary search or Newton’s method [Wang and Singh, 2007]. Instead of computing it in real-time, we build a look-up table in advance. Despite the speed up of Z^{-1} , we find that each CertProp concatenation takes on average 0.12 milliseconds, whereas W&S takes 133.37 milliseconds.¹

Performance of \textcircled{S} in CertProp and W&S The selection operator \textcircled{S} does not help in either W&S or CertProp. We observed that, in datasets like FilmTrust, opinions are assigned by users, who trust others not based on evidence, but primarily based on impressions. For example, A may give B a high opinion just because everybody trusts B . Selection is helpful only in evidence-based datasets where double counting is undesirable. In other cases, Selection is overly cautious and ignores too much evidence when a witness can be reached by many paths.

However, the *selection* result verifies the observation made by Katz and Golbeck [2006]—shorter paths have more accurate trust inference for a fixed trust rating. As Figure 9 shows, in FilmTrust, shortest paths tend to be more trustworthy because they are selected more by the selection \textcircled{S} operator. For example, over 90% of paths of length of 2 are selected. The PGP evaluation also has a similar result.

Performance of Linear and Weber-Fechner Transformation Our evaluation shows that, for opinion networks, evidence-based trust propagation approaches, like CertProp and W&S, are more accurate in *Weber-Fechner* transformed graphs than *linear* transformed ones. This result verifies our assumption that the opinions in the datasets are subjective. Also it indicates *Weber-Fechner* transformation can reduce such subjectivity, whereas *linear* transformation fails to reduce the subjective because it cannot adapt different properties of the datasets.

¹The experiment is conducted on MacBook Pro with 2.16 GHz Intel Core 2 Duo, and 2 GB 667 MHz DDR2 SDRAM, in Java Run-time Environment 5.0.




Figure 9: The distribution of the length of the paths found by CertProp (selection). Shorter paths tend to be more trustworthy and hence are selected more by selection.

6 Literature

Trust models are widely studied in various domains. Trust propagation methods can be categorized as the models that consider the witness information sources [Sabater and Sierra, 2005], or retrieve and aggregate ratings from the social network [Ramchurn et al., 2004]. Here, we focus on trust propagation approaches not covered in above surveys.

Richardson *et al.* [2003] discuss an abstract framework for trust propagation. Each user maintains trust in a small number of other users. A user’s trust in any other user can be computed by using the existing web of trust recursively. Richardson *et al.* first enumerate all paths between the user and every other user who has a local belief in a given statement. Next they calculate the belief associated with each path by using a concatenation operator along each path, and combine the beliefs associated with all paths using a predefined aggregation operator.

Yu and Singh [2002; 2003] study distributed reputation management in a social network whose member agents cooperate with each other to find the trustworthiness of other agents. When an agent wants to find the trustworthiness of a service provider, it uses the referral network to find witnesses, and combines the beliefs of those witnesses regarding the service provider. Yu and Singh’s approach is based on the Dempster-Shafer theory of evidence. It treats a service with medium quality as of unknown quality, not as a known medium quality. More importantly, it does not completely address double counting in that multiple paths through the same agent are treated as if they were independent. Although the approach identifies witnesses who have direct experience, so there is no double counting at that level.

Gray *et al.* [2003] solve trust propagation problem in mobile ad hoc networks. In their approach, the concatenation is calculated as the average of trust of each edge along the referral path, where the trust of edges is discounted by the depth in the referral path. Trust from different paths is aggregated by choosing the most trusted of the available paths. Quercia *et al.* [2007] also prop-

agates trust in mobile network context. Based on the web of trust, they build *relationship graphs* where nodes are relationships. Two *related* nodes are linked if they are either from the same rater or rating the same person. At the beginning, some of nodes are rated at the first place. Then trust is propagated from rated nodes to unrated ones by computing a *predictive* function. However, the number of relationships in real-world networks is much larger than the number of nodes. For example, there are 1, 234 edges (538 nodes) in FilmTrust, and 317, 979 edges (39, 246 nodes) in PGP. Propagating trust in the corresponding relationship graphs would be more computationally expensive.

Trust management in peer-to-peer systems have been also widely studied [Kamvar et al., 2003; Zhou et al., 2008]. Spectral decomposition is used on the adjacency matrix of the network graph to estimate global reputation, which is also called the *global group* trust metric [Ziegler and Lausen, 2004]. Guha *et al.* [2004] associate trust relations to matrix operations. For example, the commutativity of trust is associated to matrix transpose, while direct propagation is associated to matrix multiplication. Different from global group methods, our approach models trust from a personal perspective. Propagated trust takes into account personal bias. Besides, matrix approaches require several iterations to converge.

Advocate [Levien, 2003] and Appleseed [Ziegler and Lausen, 2004] are two *local group* trust metrics, which propagate global trust only on the local subgraph. Advocate applies network flow on a modified graph, where capacities are assigned to edges based on the depth in the referral path. The deeper the edge is the less capacity it has. Conversely, Appleseed adopts *spreading activation*. It spreads *energy* across the graph, and, when propagating through a node, divides energy among successors based on the edge weights. The idea of Appleseed is similar to the spectral decomposition, which requires several iterations to converge. CertProp propagates trust in local subgraphs. It provides *local* trust, rather than group-level (local group) or global trust. In other words, the propagated trust of Alice from Bob is different from the trust of Alice from Charlie.

Wang and Singh [2006] present path algebra for trust propagation, and define aggregation and concatenation. We use their aggregation but our concatenation operator is different (and faster) than theirs without loss of accuracy. We also introduce a new selection operator to avoid possible double counting in trust networks.

7 Conclusion

Trust propagation is a natural concept in settings where distributed agents must cooperate to identify the most trustworthy parties with whom to deal. Doing so helps the agents leverage each other's knowledge, benefiting the most from each other's efforts in establishing and evaluating trust relationships, and thereby increasing the social good.

Our evidence-based approach, CertProp, provides efficient operators, concatenation, aggregation, and selection, that can propagate trust accurately. These operators satisfy useful algebraic properties.

Besides, we motivate a new way to transform subjective opinions into objective evidence based

on *Weber-Fechner* law. This transformation also follows the idea that the average opinion yields the lower certainty of transformed trust. It helps reduce the subjectivity in opinion-based datasets so that the evidence-based approaches like CertProp and W&S can apply.

Our evaluation is conducted in two social network datasets, FilmTrust and PGP. The results show that CertProp provides accurate propagated trust. It indicates that *Weber-Fechner* transformation successfully reduces the subjectivity in two quite different datasets. However, although the selection operator has solid theoretical basis, the subjective datasets fail to show the effectiveness of it. Also, the selection operator ignores too much evidence while it reduces double-counting.

In the future, we would study additional trust propagation datasets, for example, Epinions [Guha et al., 2004] and Advocate [Levien, 2003; Ziegler and Lausen, 2004]. Besides, we would like to enhance the capability of attack-resistance and rumor-resistance. Finally, we will refine the selection operator in order to find balance between gathering evidence and reducing double-counting.

Acknowledgments

Thanks to Jennifer Golbeck for sharing the FilmTrust dataset. We thank Chris Hazard for his helpful comments.

References

Weber-fechner law. http://en.wikipedia.org/wiki/weber-fechner_law.

Web of trust. <http://www.lysator.liu.se/jc/wotsap/wots2/>.

Jonathan Frenzen and Kent Nakamoto. Structure, cooperation, and the flow of market information. *Journal of Consumer Research*, 20:360–375, December 1993.

Jennifer Golbeck. *Computing and Applying Trust in Web-based Social Networks*. PhD thesis, University of Maryland, College Park, 2005.

Elizabeth Gray, , Jean-Marc Seigneur, Yong Chen, and Christian Jensen. Trust propagation in small worlds. *Lecture Notes in Computer Science*, 2692:239–254, 2003.

R. Guha, Ravi Kumar, Prabhakar Raghavan, and Andrew Tomkins. Propagation of trust and distrust. In *WWW '04: Proceedings of the 13th international conference on World Wide Web*, pages 403–412. ACM Press, 2004.

Audun Jøsang. A subjective metric of authentication. In *Proceedings of the 5th European Symposium on Research in Computer Security (ESORICS)*, 1998.

Sepandar D. Kamvar, Mario T. Schlosser, and Hector Garcia-Molina. The EigenTrust algorithm for reputation management in p2p networks. In *WWW '03: Proceedings of the 12th international conference on World Wide Web*, pages 640–651. ACM Press, 2003.

- Yarden Katz and Jennifer Golbeck. Social network-based trust in prioritized default logic. In *Proceedings of the 21st National Conference on Artificial Intelligence (AAAI)*, pages 1345–1350, Menlo Park, 2006. AAAI Press.
- Ugur Kuter and Jennifer Golbeck. Sunny: A new algorithm for trust inference in social networks using probabilistic confidence models. In *Proceedings of the 22st National Conference on Artificial Intelligence (AAAI)*, pages 1377–1382, Menlo Park, 2007. AAAI Press.
- Jure Leskovec and Christos Faloutsos. Sampling from large graphs. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining (KDD '06)*, pages 631–636, New York, NY, USA, 2006. ACM Press.
- Raph Levien. *Attack Resistant Trust Metrics*. PhD thesis, UC Berkeley, 2003.
- Daniele Quercia, Stephen Hailes, and Licia Capra. Lightweight distributed trust propagation. In *Seventh IEEE International Conference on Data Mining (ICDM 2007)*, pages 282–291, 2007.
- Sarvapali D. Ramchurn, Dong Huynh, and Nicholas R. Jennings. Trust in multi-agent systems. *Knowledge Engineering Review*, 19(1):1–25, 2004.
- Matthew Richardson, Rakesh Agrawal, and Pedro Domingos. Trust management for the semantic Web. In *The Semantic Web: Proceedings of the 2nd International Semantic Web Conference (ISWC)*, volume 2870 of *LNCS*, pages 351–368. Springer-Verlag, 2003.
- Jordi Sabater and Carles Sierra. Review on computational trust and reputation models. *Artificial Intelligence Review*, 24(1):33–60, 2005.
- Yonghong Wang and Munindar P. Singh. Trust representation and aggregation in a distributed agent system. In *Proceedings of the 21st National Conference on Artificial Intelligence (AAAI)*, pages 1425–1430, Menlo Park, 2006. AAAI Press.
- Yonghong Wang and Munindar P. Singh. Formal trust model for multiagent systems. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1551–1556, Detroit, 2007. IJCAI.
- Bin Yu and Munindar P. Singh. Searching social networks. In *Proceedings of the 2nd International Joint Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, pages 65–72, New York, July 2003. ACM Press.
- Bin Yu and Munindar P. Singh. Distributed reputation management for electronic commerce. *Computational Intelligence*, 18(4):535–549, November 2002.
- Bin Yu, Mahadevan Venkatraman, and Munindar P. Singh. An adaptive social network for information access: Theoretical and experimental results. *Applied Artificial Intelligence*, 2003. To appear.

Runfang Zhou, Kai Hwang, and Min Cai. GossipTrust for fast reputation aggregation in peer-to-peer networks. *IEEE Transactions on Knowledge and Data Engineering*, 20(9):1282–1295, 2008.

Cai-Nicolas Ziegler and Georg Lausen. Spreading activation models for trust propagation. In *EEE '04: Proceedings of the 2004 IEEE International Conference on e-Technology, e-Commerce and e-Service*, pages 83–97, Washington, DC, USA, 2004. IEEE Computer Society.