Aggregation of Gradual Trust and Distrust

P. Victor
Ghent University
Belgium

C. Cornelis
Ghent University
Belgium

{Patricia.Victor,Chris.Cornelis}@UGent.be

M. De Cock*
University of Washington
USA
mdecock@u.washington.edu

E. Herrera-Viedma
University of Granada
Spain
viedma@decsai.ugr.es

Abstract

Trust and distrust are two increasingly important metrics in social networks, reflecting users' attitudes and relationships towards each other. In this paper, we study the indirect derivation of these metrics' values for users that do not know each other, but are connected through the network. In particular, we focus on the aggregation step that is required for fusing trust and distrust information obtained via multiple network paths.

Keywords: trust network, trust and distrust, aggregation, propagation, social network.

1 Introduction

A (virtual) trust network is a social network which contains trust scores assigned to each other by users (also called agents) in a given application. As an example, a trust network can be established among the users of an online Recommender System (RS); it has been shown that trust scores can be used to improve both the quality and quantity of generated recommendations (see e.g. [1, 5]).

Various representations of trust scores have been proposed, including probabilistic vs. gradual models, as well as models that capture only trust vs. those that explicitly distinguish trust from distrust. In this paper, we focus on the gradual, trust-and-distrust bilattice-based model introduced in [4], which orders trust scores both according to trustworthiness as to the amount of information they convey.

As many trust networks are large, it is unlikely that all agents know each other, in other words the network is unlikely to be fully connected. If an agent a wants to form a trust opinion about an unknown agent x, then a can inquire about x with one of its own trust relations, say b, who in turn might consult a trust connection, etc., until an agent connected to x is reached. The process of predicting the trust score along the thus constructed path from a to x is called *trust propagation*. On the other hand, it often happens that a has not one, but several trust connections that it can consult for an opinion on x. Combining several trust scores originating from different sources is called *trust aggregation*. Previously, a number of propagation operators for trust scores have been proposed [4]. In this paper, we embark on the problem of aggregating information originating from multiple trust score paths into an overall score.

The remainder of this paper is organized as follows: in Section 2, we recall preliminaries regarding the bilattice model and its trust propagation operators, while in Section 3, we propose a number of criteria that trust score aggregation operators need to fulfill, and also suggest one possible type of operators. Finally, in Section 4, we conclude and outline ideas for future investigation.

*On leave from Ghent University

2 Preliminaries

2.1 Bilattice-Based Trust Model

For our purposes, a trust network is represented as a directed graph (A, E, R) in which A is the set of agents (nodes), E is the set of trust connections (edges), and R is an $E \to [0,1]^2$ mapping that associates to each couple (a, b) of connected agents in E a trust score $R(a, b) = (R^+(a, b), R^-(a, b))$ in $[0,1]^2$, in which $R^+(a, b)$ is called the trust degree of a in b and $R^-(a, b)$ is called the distrust degree of a in b.

The set of trust scores can be endowed with a bilattice structure. In particular, the trust score space [5]

$$\mathcal{BL}^{\square} = ([0,1]^2, \leq_t, \leq_k, \neg)$$

consists of the set $[0,1]^2$ of trust scores, a trust ordering $\leq_t$, a knowledge ordering $\leq_k$, and a negation $\neg$ defined by

$$\begin{array}{rcl} (x_1, x_2) \leq_t (y_1, y_2) & \text{iff} & x_1 \leq y_1 \text{ and } x_2 \geq y_2 \\ (x_1, x_2) \leq_k (y_1, y_2) & \text{iff} & x_1 \leq y_1 \text{ and } x_2 \leq y_2 \\ \neg(x_1, x_2) & = & (x_2, x_1) \end{array}$$

for all (x_1, x_2) and (y_1, y_2) in $[0,1]^2$.

The "trust lattice" $([0,1]^2, \leq_t)$ orders the trust scores going from complete distrust $(0,1)$ to complete trust $(1,0)$. The "knowledge" lattice $([0,1]^2, \leq_k)$ evaluates the amount of available trust evidence, ranging from a "shortage of evidence", $x_1 + x_2 < 1$ (incomplete information), to an "excess of evidence", viz. $x_1 + x_2 > 1$ (inconsistent information). The boundary values of the $\leq_k$ ordering, $(0,0)$ and $(1,1)$, reflect ignorance, resp. contradiction.

2.2 Propagation Operators

In virtual trust networks, propagation operators are used to handle the problem of establishing trust information in an unknown agent by inquiring through other agents. The simplest case, atomic propagation, takes the trust score of agent a in agent b and the trust score of b in agent c, and uses this information to predict the trust score of a in c. In [4], four operators were proposed for this purpose, each reflecting a different strategy of dealing with the available trust information.

Let T be a t-norm, S a t-conorm and N a negator. The propagation operators P_1, P_2, P_3 and P_4 are defined by, for (t_1, d_1) and (t_2, d_2) in $[0,1]^2$,

$$\begin{array}{l} P_1((t_1, d_1), (t_2, d_2)) = (T(t_1, t_2), T(t_1, d_2)) \\ P_2((t_1, d_1), (t_2, d_2)) = (T(t_1, t_2), T(N(d_1), d_2)) \\ P_3((t_1, d_1), (t_2, d_2)) = \\ (S(T(t_1, t_2), T(d_1, d_2)), S(T(t_1, d_2), T(d_1, t_2))) \\ P_4((t_1, d_1), (t_2, d_2)) = \\ (T(t_1, t_2), S(T(t_1, d_2), T(d_1, t_2))) \end{array}$$

P_1 reflects the basic strategy of taking over information only from trusted sources, while P_2 exhibits a paranoid behaviour by taking over distrust information even from an unknown party. P_3, on the other hand, also takes into account distrusted sources, reversing their opinion rather than ignoring or copying it; P_4 mitigates P_3's behaviour by reversing only distrust information and ignoring trust information coming from a distrusted party.

Note that only P_1 is associative. The fact that the other operators are not associative means that we need to fix a particular evaluation order to propagate trust scores over paths with more than 2 edges. In a network with a central authority (CA) that maintains all trust information, this order is arbitrary. On the other hand, if there is no CA, and each agent has access only to the trust scores it has issued, it is necessary to perform the propagation in a right-to-left direction. With this order, at each node in the propagation path, an agent combines its trust score in its successor, with the propagated trust score it receives from this successor. This is illustrated below for a path containing three edges, and a generic propagation operator P:

$$\begin{array}{l} P((t_1, d_1), (t_2, d_2), (t_3, d_3)) = \\ P((t_1, d_1), P((t_2, d_2), (t_3, d_3))) \end{array}$$

In the remainder, we assume that the right-to-left evaluation order is used for propagation.

3 Trust Score Aggregation

When an agent a needs to establish an opinion about another agent x, and there is more than one path linking them, we require a way of combining the information provided by each of those paths. In this section, before discussing possible trust score aggregation strategies, we first establish a number of basic requirements for an adequate trust score aggregation operator.

3.1 Characterization of Trust Score Aggregation Operators

Let $\mathcal{BL}^{\square} = ([0,1]^2, \leq_t, \leq_k, \neg)$ be the trust score space previously introduced. In this space, we look for a trust score aggregation operator $\Omega : ([0,1]^2)^n \to [0,1]^2$ $(n \geq 1)$ satisfying the following requirements or properties:

1. *Idempotence property* (**RQ1**). If an agent a receives the same trust score $(t,d) \in [0,1]^2$ from different paths linked to the agent x, i.e., $[(t_1,d_1) = (t,d), \ldots, (t_n,d_n) = (t,d)]$, then the opinion of a about x should be that same trust score, i.e,

$$\Omega((t_1,d_1),\ldots,(t_n,d_n)) = (t,d).$$

2. *Monotonicity property* (**RQ2**). If through a given trust score path to x an agent a receives more knowledge (resp., more trust), then the overall trust (resp., knowledge) of a's opinion about x cannot decrease. Therefore, a trust score aggregation operator Ω should be monotonously increasing with respect to both $\leq_t$ and $\leq_k$, such that if $(t_j,d_j) \leq_k (t_j^{'},d_j^{'})$, then $(p,q) \leq_k (p^{'},q^{'})$; and if $(t_j,d_j) \leq_t (t_j^{'},d_j^{'})$, then $(p,q) \leq_t (p^{'},q^{'})$, with

$$\Omega((t_1,d_1),\ldots,(t_j,d_j),\ldots,(t_n,d_n)) = (p,q)$$
$$\Omega((t_1,d_1),\ldots,(t_j^{'},d_j^{'}),\ldots,(t_n,d_n)) = (p^{'},q^{'})$$

3. Commutativity property (**RQ3**). To compute the opinion of agent a about agent x the order of the trust scores should not matter. That is,

$$\Omega((t_1,d_1),\ldots,(t_n,d_n)) =$$
$$\Omega(\pi(t_1,d_1),\ldots,\pi(t_n,d_n))$$

where π is a permutation over the set of trust scores.

Besides these basic properties, we also propose a number of additional requirements to guarantee the intuitive behaviour of the trust score aggregation process; we motivate these requirements by examples. In the scenario in Figure 1, b and c are both fully trusted acquaintances of a that are connected to x. Propagation with any of the four operators from the previous section results in the two trust scores (t,d) and $(0,0)$. However, it can be argued that c's opinion of x (ignorance) should not contribute to the final outcome; indeed, a $(0,0)$ edge can be considered as no edge at all, as it carries no information. In other words, $(0,0)$ should act as a neutral element of the aggregation operator **(RQ4)**:

$$\Omega((t_1,d_1),\ldots,(t_{n-1},d_{n-1}),(0,0)) =$$
$$\Omega((t_1,d_1),\ldots,(t_{n-1},d_{n-1}))$$

This scenario also shows why a naive average of trust and distrust degrees, leading to $\left(\frac{t}{2},\frac{d}{2}\right)$, would be a poor aggregation strategy in this case.

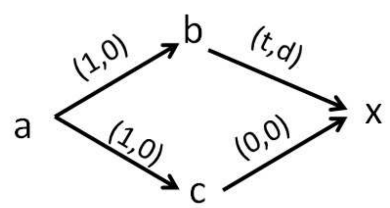

Figure 1: Scenario with ignorance

In Figure 2, two fully trusted acquaintances of a express completely opposite trust opinions of x. Again, a simple average of trust and distrust degrees, yielding $(0.5,0.5)$, is unsuitable, since it does away with the conflicting information a receives, and cannot be distinguished from a scenario in which a receives information that x is half to be trusted and half to be distrusted. A more intuitive result seems to be $(1, 1)$, reflecting the inconsistency a faces in its assessment of x. This brings us to a final requirement for trust score aggregation operators: if the arguments consist of an equal number of $(1, 0)$ and $(0, 1)$ trust scores, and all other arguments are $(0, 0)$, then the result should be $(1, 1)$. We denote this condition by **(RQ5)**.

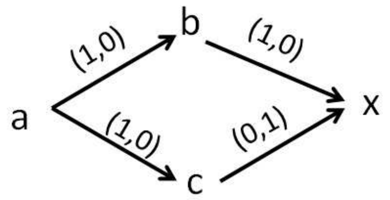

Figure 2: Scenario with total inconsistency

3.2 Trust Score Aggregation Operators

Both examples from the previous section reveal that aggregation should, in a way, respect the *amount of information* carried by the different trust scores to be merged. Such a behaviour is exhibited by $\oplus$, the join of the information lattice $([0, 1]^2, \leq_k)$:

$$(t_1, d_1) \oplus (t_2, d_2) = (\max(t_1, t_2), \max(d_1, d_2))$$

It can be checked that $\oplus$ satisfies **(RQ1)**–**(RQ5)**, making it a very attractive operator at first glance. However, when there are more than two agents contributing to the aggregated trust score, $\oplus$ may not be so suitable. To see this, Figure 3 shows a variation of the previous scenario, in which again the contributing agents express conflicting opinions, but in this case a majority of the sources vote against x. Applying the join yields $(1, 1)$, which does not reflect the fact that a has more evidence to distrust than to trust x. In fact, we expect that the aggregated trust score (t, d) contains more distrust than trust information (i.e., $t < d$).

The example shows that $\oplus$ does not allow for compensation in merging (possibly conflicting) information; a weighted aggregation operator might therefore be more suitable to model the trust score aggregation process.

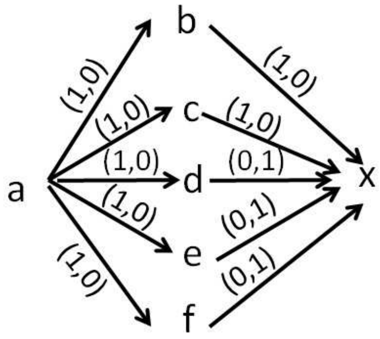

Figure 3: Scenario with partial inconsistency

One of the best-known examples is the *Ordered Weighted Averaging aggregation operator* [6], or OWA for short. This operator models a weighted average process in which n scalar values are first (decreasingly) ordered and then weighted according to their ordered position by means of the weighting vector $W = \langle w_i \rangle$, such that $w_i \in [0, 1]$ and $\Sigma_i^n w_i = 1$. The OWA's main strength is its flexibility, since it enables us to model a whole range of aggregation strategies: e.g., if the weighting vector equals $W_1 = \langle 0, \cdots, 0, 1 \rangle$, then the corresponding OWA operator equals min; the average is modeled by $W_2 = \langle 1/n, \cdots, 1/n \rangle$, etc. The OWA operator can be analyzed by several measures, among which the orness-degree that computes how similar its behaviour is to that of max:

$$orness(W) = \frac{1}{n-1} \sum_{i=1}^{n} ((n - i) \cdot w_i).$$

$orness(W_1)$ is 0, while the weighting vector $\langle 0.8, 0.1, 0.1, 0, 0 \rangle$ yields an orness-degree of 0.925, which means that its behaviour is very maximum-like. OWA operators with an orness degree less than 0.5 approach the minimum operator.

Remark that $\oplus$ can be modeled by two separate OWAs, one for the trust and one for the distrust degrees, both using the weighting vector $\langle 1, 0, \cdots, 0 \rangle$. Unlike the join however, the more general OWA operators can also be used to handle partial inconsistencies. Let us reconsider Figure 3, in which a received information from 5 trusted agents. A possible way to handle this scenario is to compute the trust degree in x as the degree to which the majority of the trusted agents trust x; this can be achieved, for instance, by means of the weighting vector $W_T = \langle 3/6, 2/6, 1/6, 0, 0\rangle$.

An analogous strategy can be pursued for computing the distrust degree. The OWA used in this case does not necessarily need to equal the one used for the trust degree; for instance, one might argue that the presence of a few trusted agents distrusting x is already sufficient to decide that x should be distrusted. This behaviour can e.g. be modeled by the weights $W_D = \langle 4/5, 1/5, 0, 0, 0\rangle$. Applying the OWA operators OWA_{W_T} and OWA_{W_D} yields the aggregated trust score of $(5/6, 1)$ for the scenario in Figure 3, which better reflects the partial inconsistency and the tendency to distrust x. Note that both weighting vectors clearly show a maximum-like behaviour, since the orness-degrees are $\frac{5}{6}$ and $\frac{19}{20}$, for trust and distrust, respectively.

The above example can be generalized into a general class of trust score aggregation operators.

Definition 1 *The trust score OWA operator $T\text{-}OWA_{W,V}$ associated with weighting vectors W and V is defined as*

$$T\!-OWA_{W,V}\left((t_1, d_1), \cdots, (t_n, d_n)\right) = \left(OWA_W\left(t_1, \cdots, t_n\right), OWA_V\left(d_1, \cdots, d_n\right)\right),$$

with $W = \langle w_1, \cdots, w_n\rangle$ and $V = \langle v_1, \cdots, v_n\rangle$ such that

1. *$orness(W) \geq 0.5$, $orness(V) \geq 0.5$*
2. *$w_i = 0$ when $(t_{\tau(i)}, d_{\tau(i)}) = (0, 0)$, with $\tau(i)$ denoting the index of the argument with the i-th largest trust degree.*
3. *$v_i = 0$ when $(t_{\delta(i)}, d_{\delta(i)}) = (0, 0)$, with $\delta(i)$ denoting the index of the argument with the i-th largest distrust degree.*

The first condition ensures that both components of a $T\text{-}OWA$ operator have a maximum-like behaviour, and hence that it remains possible to model the join operation. The second and third conditions ensure that $(0, 0)$ acts as a neutral element for any $T\text{-}OWA_{W,V}$. As an example, reconsider the scenario in Figure 3, but this time imagine x to be unknown to agent f, i.e., $R(f, x) = (0, 0)$, and suppose moreover that a also completely trusts agents g, h, i, and j, all of whom have no information about x. In this scenario, only 4 out of 9 acquaintances (b, c, d, and e) can deliver useful information, and their opinions are fully conflicting; hence, we expect the final aggregated result to be $(1, 1)$. However, if we compute the weights along the same lines as before and use e.g. $W = \left\langle \frac{5}{15}, \frac{4}{15}, \frac{3}{15}, \frac{2}{15}, \frac{1}{15}, 0, 0, 0, 0\right\rangle$, then the trust degree of the final result will be $\frac{9}{15}$ instead of 1; this is due to the fact that some of the ignorance trust scores are also weighted. Conditions 2 and 3 prevent this: the trust and distrust weights for the ignorant agents will be zero, and the other weights can be determined as if there were no edges from a to f–j. E.g., in our scenario, since only 4 agents have an opinion about x, W can be constructed as $\langle 2/3, 1/3, 0, 0, 0, 0, 0, 0, 0\rangle$, so that x is trusted to the degree that the majority of the *knowledgeable* agents trusts x. Following a similar strategy for the distrust part generates a final trust score of $(1, 1)$, which does correspond to our intuition.

Note that every trust score OWA operator fulfills **(RQ1)**–**(RQ5)**; this follows from the properties of the classic OWA operator (idempotency, commutativity, monotonicity), and the construction of the weighting vectors. The orness degrees of the OWAs used to construct the particular $T\text{-}OWA_{W,V}$ operator have a lot of impact on its interpretation. As an example, suppose agent a receives information about agent x from 6 fully trusted acquaintances, namely $(1, 0)$, $(0.9, 0)$, $(0.8, 0)$, $(0, 0.8)$, $(0, 0.9)$, and $(0, 1)$. Choosing $W_1 = \left\langle \frac{3}{6}, \frac{2}{6}, \frac{1}{6}, 0, 0, 0\right\rangle$ yields $t = \frac{14}{15}$, while $W_2 = \left\langle \frac{2}{6}, \frac{2}{6}, \frac{2}{6}, 0, 0, 0\right\rangle$ results in $t = \frac{9}{10}$. Both operations expose a maximum-like behaviour, but the ways in which the operation is implemented clearly differ: in the latter scenario, the opinion of every agent in the trusted majority is equally important, while in the former scenario a is of a more gullible nature, illustrated by the higher orness-degree. The same discussion applies to OWA_V. E.g., the distrust weights $V = \langle \frac{4}{5}, \frac{1}{5}, 0, 0, 0, 0 \rangle$ result in $d = \frac{49}{50}$, reflecting the fact that a is more prudent when distrust is involved (note the higher orness-degree of OWA_V). Applying $T\text{-}OWA_{W_1,V}$ then yields $(t, d) = (0.93, 0.98)$, which tells us that a had to deal with highly inconsistent sources. Remark that, due to the effect of the W/V-choice, $t < d$, although there is as much information contained in the distrust opinions as in the trust opinions.

4 Conclusion and Future Work

Research in trust networks is still in its infancy, in particular when it comes down to the representation, propagation and aggregation of distrust. In this paper, we have built upon previous work in which trust scores are modeled as (trust,distrust)-couples that are drawn from a bilattice, and investigated which requirements a trust score aggregator needs to fulfill. We introduced a class of trust score aggregation operators, called *T-OWA*, which has its origin in the classic ordered weighting average operator, and which meets all the proposed criteria. The main asset of *T-OWA* is its versatility: the operator is employable in a wide range of agent networks because it allows to model specific behaviours when dealing with trust and distrust.

A first step in our future research involves the further exploration of the weighting vectors. For instance, it makes sense to generate weights which not only depend on the orness-degree, but also on the amount of knowledge that the corresponding trust scores contain; one can argue that the opinion of agents who do not show much doubt (i.e, not much ignorance or inconsistency) should be weighted more in the aggregation. In other words, an agent's trust score (t, d) could be evaluated according to the proximity of $t + d$ to 1.

Secondly, we also want to investigate the suitability of other existing and/or new operators for the trust score aggregation process. Possible avenues include the weighted ordered weighted aggregation operator (WOWA, [3]) and the linguistic OWA operator [2], which allow for even more flexibility.

Acknowledgements

Patricia Victor would like to thank the Institute for the Promotion of Innovation through Science and Technology in Flanders, and Chris Cornelis the Research Foundation–Flanders, for funding their research.

References

[1] J. Golbeck (2005). Computing and Applying Trust in Web-Based Social Networks. PhD thesis.

[2] F. Herrera, E. Herrera-Viedma (1997). Aggregation Operators for Linguistic Weighted Information, IEEE Transactions on Systems, Man and Cybernetics—Part A: Systems and Humans 27(5), p. 646-656.

[3] V. Torra (1996). The WOWA operator and the interpolation function W^*, International Journal of Intelligent Systems 12, p. 153-166.

[4] P. Victor, C. Cornelis, M. De Cock, P. Pinheiro Da Silva (2009). Gradual Trust and Distrust in Recommender Systems. Fuzzy Sets and Systems 160(10), p. 1367-1382.

[5] P. Victor, C. Cornelis, M. De Cock, A. Teredesai (2009). Trust- and Distrust-Based Recommendations for Controversial Reviews. In: *Proceedings of Web Science Conference (WebSci'09).*

[6] R. R. Yager. On Ordered Weighted Averaging Aggregation Operators in Multicriteria Decision making (1988). IEEE Transactions on Systems, Man, and Cybernetics, 18, p. 183-190.