



HAL
open science

A Very Fast and Robust Trust Inference Algorithm in Weighted Signed Social Networks using Controversy, Eclecticism, and Reciprocity

Karim Akilal, Hachem Slimani, Mawloud Omar

► **To cite this version:**

Karim Akilal, Hachem Slimani, Mawloud Omar. A Very Fast and Robust Trust Inference Algorithm in Weighted Signed Social Networks using Controversy, Eclecticism, and Reciprocity. *Computers & Security*, 2019. hal-03034577

HAL Id: hal-03034577

<https://hal.science/hal-03034577>

Submitted on 1 Dec 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A Very Fast and Robust Trust Inference Algorithm in Weighted Signed Social Networks using Controversy, Eclecticism, and Reciprocity.

Karim Akilal^{a,*}, Hachem Slimani^a, Mawloud Omar^b,

^aLaboratoire d'Informatique Médicale, Faculté des Sciences Exactes, Université de Bejaia, 06000 Bejaia, Algérie

^bUnité de Recherche LAMOS, Faculté des Sciences Exactes, Université de Bejaia, 06000 Bejaia, Algérie

Abstract

The importance of trust in social networks instigated many research efforts to understand it and predict it. However, the complex nature of trust and its counterpart; distrust; makes these tasks challenging. While early trust inference approaches ignore distrust, it seems that this concept gained much attention in recent years. Surely, knowing whom to distrust is as important as knowing whom to trust. We show in this paper that trust and distrust can be quickly predicted using some social traits of the trustor and the trustee. Using a “*tug of war*” analogy involving these traits, we have devised an intuitive approach that uses only the direct neighbors of the trustor and those of the trustee to predict both trust and distrust. Experiments on four real-world social networks show that our algorithm is very fast, provides good predictions, and is robust to network sparsity.

Keywords: Online Social Network, Trust Inference, Distrust, Trust metric, Social Trait

1. Introduction

Online social networks are a prominent fact of our modern life. Billions of people all over the world interact, produce, and consume information everyday thanks to these networks. However, these advantages come with some downsides. Indeed, dangers exist, as noted by Huang et al. (2013), because of malicious actors who take advantage of the open nature of these social networks to mislead, or even harm, others. How, and why, should a user trust another one with information they may share with, or receive from, them? Organizational mechanisms, such as moderation, may be of some help, but they are prone to corruption (Shneiderman, 2015), and may become inefficient as data flows become voluminous. We thus need, as suggested by Matei et al. (2015), tools that can support timely, effective, and efficient knowledge extraction processes from such data, to help users by predicting and recommending how much should they trust each other.

Furthermore, as suggested by Massa and Avesani (2005), predicting distrust is as important as predicting trust. Indeed, social networks users should be able to discern whom

*Corresponding author

Email addresses: karim@akilal.com (Karim Akilal), haslimani@gmail.com (Hachem Slimani), mawloud.omar@gmail.com (Mawloud Omar)

Preprint submitted to Elsevier

February 7, 2019

to trust, and whom to —not passively ignore, but— *actively* distrust. Unfortunately, as noted by Ziegler (2013), most early approaches completely ignore distrust or consider it as the absence of trust (neutrality). In the real world, however, to distrust someone is definitely different from being neutral toward them. Distrust, as Cho (2006); Hawley (2013) put it, is not mere absence of trust. It is fundamentally characterized by a sentiment of unease, pessimism, and, to some extent, fear as explained by Lewicki et al. (1998); Lewicki and Brinsfield (2012). These aspects are not present when being neutral toward others. Distrust calls for precautions to be taken. Therefore, alerting users and recommending precaution by predicting distrust is important. In fact, we can argue, as did DuBois et al. (2011), that predicting distrust is often more important than predicting trust, because actually knowing whom to distrust —and to which extent— is the key to mitigate the very risk of trust: betrayal (and harm in some cases) (Jones, 1996; Mayer et al., 1995).

Unfortunately, incorporating distrust into most early approaches is often impractical or even impossible (Guha et al., 2004; Chiang et al., 2014; Tang et al., 2016b). Therefore the need for novel solutions that are able to process and predict both trust and distrust. Additionally, we consider that these solutions should satisfy three more conditions:

1. Accuracy: provide the best possible prediction as to how much a user should trust or distrust another one.
2. Robustness: the accuracy of the algorithm should not be very affected by the unavailability or invisibility of some, or most, regions of the network.
3. Speed: we expect the algorithm to provide a trust prediction very quickly as the user’s —often fast— (re)actions will depend on it.

Satisfying these conditions at once is a tough challenge. Indeed, algorithms may need to make costly computations to predict trust accurately, therefore lack the desired speed that we expect. They may also need to use most of the network edges, thus produce inaccurate predictions when these edges are not available. These are, in fact, some limitations of propagative trust prediction methods that also suffer from path dependences, trust decay on long paths, opinion conflict (Jiang et al., 2016a), and somewhat fail to handle distrust because of its lack of transitivity (Ziegler and Lausen, 2005; DuBois et al., 2011; Tang et al., 2014).

To tackle these issues, we think that trust prediction algorithms should opt for a more localized approach. They should extensively use the direct neighbors of the trustor and those of the trustee. By doing so, these algorithms will not have to use expensive graphs traversals, nor depend on edges that may be unavailable due to technical or privacy concerns.

Trust is often seen from the viewpoint of the trustor. For instance, in a quest for a cross-disciplinary definition of trust, Rousseau et al. (1998) concluded that a widely held definition is that “*Trust is a psychological state comprising the intention to accept vulnerability based upon positive expectations of the intentions or behavior of another*”. It is also stated as “*the evaluation by a trustor u of the willingness and/or ability of a trustee v to perform a task x* ” (Faulkner, 2014; Robbins, 2016; Bauer, 2017). This construct, while sound, understandably considers that u is the *subject* and v is the *object* of the *verb* *trust*. Such a characterization of trust gives the whole decision power to the trustor, and makes the trustee undergo it.

In this work, we take a different approach that considers that the *act* of trust (or distrust) is a joint decision made by both parties (the trustor and the trustee). Our reasoning was inspired by a recent TED talk by El Ghalid (2017). The speaker explains how fungi locate, grow toward, and infect a tomato plant using sensors that react to molecular signals emitted by the plant. One would think that the plant is a *victim* of fungi. But in light of this work, it is clear that the tomato plant is actually *responsible* of—or at least accomplice in—its own sad fate. Likewise, we think that trustees are not mere bystanders objects of trust; but real actors that take part in the decision to be (dis)trusted. Therefore, instead of seeing trust as a *thing* that flows from the trustor to the trustee (and thus *instinctively* try to propagate it), we see it as a thing that both individuals fight over.

Such a model makes the act of trust look like a *tug of war* game. A game where both opponents (the trustor and the trustee) try to pull the other to its side. We show that by considering three social traits, namely: controversy, eclecticism, and reciprocity, we are able to devise a simple and very fast algorithm that produces good predictions. Our algorithm operates only on the direct neighbors of the trustor and the trustee instead of propagating trust along paths of the network. Being *localized* is, indeed, what makes it fast and robust to network sparsity.

The remainder of this paper is organized as follows. In Section 2, we give a brief review of some related work. Next, in Section 3, we describe our approach, by first defining and formalizing the three social traits (controversy, eclecticism, and reciprocity), then explaining how these traits may be used in predicting trust. In Section 4, we conduct some experiments to evaluate the performances, the efficiency, and the robustness of the proposed approach. We discuss the results of these experiments in Section 5, and conclude this paper in Section 6 with a summary and some perspectives of future work.

2. Related work

As noted by Kramer and Cook (2004), contributions from various disciplines converge to the conclusion that trust often constitutes an important resource within social systems. However, while trust as a concept makes some consensus among scholars and disciplines (Rousseau et al., 1998), distrust still divides them. For instance, Guo et al. (2017) cite three major camps each with its own vision of distrust. The first camp considers trust and distrust as the opposite ends of the same continuum, with distrust and complete lack of trust being the same thing (Schoorman et al., 2007). The second camp also views distrust as the opposite of trust, but considers the existence of a middle region in the spectrum where an individual neither trusts nor distrusts another one, i.e., a state of *neutrality* (Robbins, 2016). The third camp considers distrust as a separate concept from trust. Their argument is that these two concepts have different antecedents and outcomes, and that both trust and distrust can—and often do—coexist in the same relation between two individuals (Lewicki et al., 1998; McKnight and Choudhury, 2006). A coexistence that, according to Schoorman et al. (2007), is not enough to separate the two concepts, and that they attribute to relations being multifaceted or multiplex. For their part, Cho (2006), Hawley (2013), Robbins (2016) among other scholars disagree with Schoorman et al. (2007) on the idea that distrust is the lack of trust. To quote Cho (2006): “Distrust is not just the absence of trust but the active expectation that the other party will behave in a way that violates one’s welfare and security”.

The debate is still ongoing, and so are efforts to predict trust and distrust. Indeed, according to Tang and Liu (2015), there are many ways to predict trust. Some are supervised, some are not. Some make use of interaction data (Jones et al., 2013; Huang et al., 2018) or emotions (Beigi et al., 2016), some make do with only trust relations in the network. For the sake of brevity, and to better position our work, we invite the reader to some interesting surveys like those by Jiang et al. (2016a) and Tang et al. (2016b), and focus hereafter on some unsupervised graph-based trust prediction approaches.

Unsupervised graph-based trust inference itself can be carried out in two different ways. First, using local metrics which answer how would a node u trust another one v . Second, using global metrics which answer how trustworthy, or leaning to trust, is a given node (Massa and Avesani, 2007; Tang and Liu, 2015). That is, the first family of methods operates at the edge-level, while the second family operates at the node-level. In what follows, we give a brief review of some works on the subject. We start with local and global approaches that were designed for unsigned networks (trust only), then extend our review to some other methods that apply for signed networks (trust and distrust).

2.1. Trust prediction in unsigned networks

As stated earlier, most early trust inference approaches tend to ignore distrust or simply consider it as the mere absence of trust. Among the most known graph-based local approaches there are TidalTrust (Golbeck, 2005a), MoleTrust (Massa and Avesani, 2007), SWTrust (Jiang et al., 2014), GFTrust (Jiang et al., 2016b), TiSoN (Hamdi et al., 2016), and many other algorithms all based on the transitive aspect of trust. A principle that states that if a node u trusts another node v that trusts a third one w , then the node u may trust w to some extent (Golbeck, 2005b). These algorithms are often two-phased processes. The first phase is that of *propagation* of which Guha et al. (2004) defined the four atomic operations: direct propagation (transitivity), transpose, co-citation, and trust coupling. This phase consists in exploring different paths from a source node to a sink node in order to estimate how would the source trust the sink using these propagation rules. The second phase consists in *aggregating* the different results obtained from different paths into a single final value. As noted by Jiang et al. (2016a), these approaches suffer from path dependence, trust decay, and opinion conflict. As for global approaches, many metrics have been proposed in the literature. For example, the PageRank (Page et al., 1999) and the HITS (Kleinberg, 1999) algorithms that were proposed to rank web pages, are often considered for social networks as well (Hu et al., 2018; Zhao et al., 2018). Similarly, the EigenTrust algorithm by Kamvar et al. (2003) has also known many variants (Chiluka et al., 2012; Kurdi, 2015). Note that, as argued by Tang and Liu (2015), trust is a subjective matter, and global metrics fail to describe how would a specific node u trust another one v , for v is likely to be trusted differently by nodes of the network. Moreover, these metrics were devised for unsigned networks, and can not be used *as-is* on signed networks, unless we choose to ignore the negative edges (Tang et al., 2016b). Such a choice means that the precious information conveyed by negative links will be lost, and accuracy would decline as such.

2.2. Trust prediction in signed networks

As empirically confirmed by Tang et al. (2014) and later by Gao et al. (2016), while trust may be transitive, distrust is often not. Indeed, Gao et al. (2016) have reported that in an Epinions dataset, if a node u distrusts another one v which also distrusts a third one w , then there are about 50% of cases where u trusts w , and about as much cases where u distrusts w . This uncertainty makes applying propagative approaches on signed networks even more challenging.

Still, the advent of signed social networks, and the benefits that negative links bring to network analytics (Papaoikonomou et al., 2013; Kunegis et al., 2013; Tang et al., 2016a) have instigated many efforts to predict trust in signed networks. In their work, Zolfaghar and Aghaie (2010) introduced some metrics such as the popularity and the gregariousness of a node. These metrics, while including negative links, do not consider edges' weights. For weighted and signed networks (WSNs), Mishra and Bhattacharya (2011) introduced two global metrics: BIAS and DESERVE. These two metrics respectively describe the bias of a trustor, and the prestige of a trustee. Some other efforts were devoted to adapt the PageRank and the HITS algorithms to take into account negative links (de Kerchove and Dooren, 2008; Shahriari and Jalili, 2014). More recently, Kumar et al. (2016) defined two new global metrics: FAIRNESS and GOODNESS, and a way to infer local trust simply by multiplying the fairness of the trustor by the goodness of trustee. Speaking of local trust, an interesting approach based on subjective logic algebra (Jøsang, 1999), built on opinions being represented by triplets of (trust, distrust, uncertainty), was proposed by Jøsang and Pope (2005), which also discussed trust and distrust transitivity and proposed a set of semantic requirements for trust transitivity that were later detailed in (Jøsang, 2016). Ziegler and Lausen (2005) proposed *Appleseed*, a propagative algorithm that is inspired by *spreading activation models* which roughly considers trust as an energy passing from nodes to their trustees. *Appleseed* handles distrust by passing it as a negative energy. Another interesting propagative approach was recently proposed by Gao et al. (2016). In their STAR algorithm, the authors defined a semiring (an algebraic structure) that operates on 2D values: trust and certainty. The trust that a node u puts in another one v is expressed by a pair of values (t_{uv}, c_{uv}) , where $t_{uv} \in [-1, +1]$ is the amount of trust/distrust, and $c_{uv} \in [0, +1]$ is the certainty value of t_{uv} . To assess the certainty dimension, the authors used the path length to a target node, and the degree of that target node. The semiring that they defined favors arcs with bigger certainty values and circumvent the intransitivity of distrust by simply ignoring paths with two successive negative links.

In summary, the dilemma of predicting trust and distrust is twofold: 1) In addition to their forecited downsides, propagative approaches also fail to completely embrace distrust because of its lack of transitivity. 2) Global metrics fail to describe how precisely should a given node u trust another node v , because v may in fact be controversial (Massa and Avesani, 2005).

Faced with this situation, we have decided to abandon the propagation route. We propose that local trust may still be predicted using three simple node-level metrics: controversy, eclecticism, and reciprocity. These metrics when used together, or more precisely one *against* the others, provide very fast, good, and robust trust and distrust predictions. Details of this approach are given in the next section.

3. Our proposed approach

3.1. Notation and preliminaries

Our approach is graph-based. It treats social networks as weighted and signed directed networks (WSN). Nodes of these WSNs are individuals and arcs are trust relations between these individuals. We consider trust and distrust as two continuous and opposite states represented by real values in the interval $[-1, +1]$. The more a node u trusts another node v , the more this value is positive. And the more u distrusts v the more the value is negative. Table 1 summarizes the adopted notation.

Notation	Meaning
$\mathcal{G}(\mathcal{N}, \mathcal{E}, \mathcal{W})$	A weighted directed graph \mathcal{G} with nodes in \mathcal{N} connected by arcs in \mathcal{E} that are weighted using the mapping \mathcal{W} .
$\mathcal{W}(u, v)$	Weight of the arc going from node u to node v .
$\vec{\Gamma}(u)$	Set of the trustees of the node u .
$\overleftarrow{\Gamma}(v)$	Set of the trustors of the node v .
\mathcal{R}	Trust range. $\mathcal{R} = \max(\mathcal{W}) - \min(\mathcal{W})$.
$\vec{\mu}(v)$	Mean of trust received by v from its trustors.
$\overleftarrow{\mu}(u)$	Mean of trust put by u in its trustees.
$\vec{\sigma}(v)$	Controversy of the trustee v .
$\overleftarrow{\sigma}(u)$	Eclecticism of the trustor u .
$\rho(u)$	Reciprocity of the trustor u .

Table 1. Notation used throughout this paper.

3.1.1. Iverson brackets

To simplify mathematical writing, while keeping it rigorous, we use Iverson Brackets (Knuth, 1992). This notation makes an integer (0 or 1) from a logical statement P put between brackets as follows:

$$[P] = \begin{cases} 1 & \text{if } P \text{ is true,} \\ 0 & \text{if } P \text{ is false.} \end{cases}$$

Note that we use the “strong zero” convention adopted by Knuth (1992), Graham et al. (1994), and others; and that states that if P is false then $[P]f(x)$ would be equal to 0 even if $f(x)$ is undefined. To quote Donald E. Knuth:

In general, when an Iverson-bracketed statement is false, we want it to evaluate into a “very strong 0,” namely a zero so strong that it annihilates anything it is multiplied by—even if that other factor is undefined (Knuth, 1992).

3.2. Problem definition

Let $\mathcal{G} = (\mathcal{N}, \mathcal{E}, \mathcal{W})$ be a directed graph representing a social network, where \mathcal{N} is a set of nodes, \mathcal{E} a set of arcs between nodes of \mathcal{N} , and $\mathcal{W} : \mathcal{E} \mapsto [-1, +1]$ a mapping that associates to each arc (u, v) a weight $\mathcal{W}(u, v)$ that represents the trust ($\mathcal{W}(u, v) > 0$) or distrust ($\mathcal{W}(u, v) < 0$) that the node u puts in the node v . The present work aims to quickly predict how much would a node u (dis)trust another node v . In addition to accuracy and speed, our approach needs to be robust to network sparsity. To simplify our narrative, we will use *trust* to denote both states: trust (positive values) and distrust (negative values).

3.3. Metrics based on social traits

We define hereafter the three social traits that we use to characterize the nodes, namely: controversy, eclecticism, and reciprocity. Right after this, we describe how trust can be predicted using these traits.

Definition 3.1. A node v is said to be *controversial* if its trustors do not share the same opinion about it. That is, the more diversified the trust values that its trustors put in it, the more controversial v is. Formally, we define the *controversy* $\overleftarrow{\sigma}(v)$ of a node v using the standard deviation of the trust values that v receives from its trustors. That is:

$$\overleftarrow{\sigma}(v) = \frac{2 \left[\overleftarrow{\Gamma}(v) \neq \emptyset \right]}{\mathcal{R}} \sqrt{\frac{1}{|\overleftarrow{\Gamma}(v)|} \sum_{u \in \overleftarrow{\Gamma}(v)} (\mathcal{W}(u, v) - \overleftarrow{\mu}(v))^2}. \quad (1)$$

Definition 3.2. A node u is said to be *eclectic* if it is willing to evaluate different kinds of nodes. This difference being stated by u itself. Similarly to controversy, the *eclecticism* $\overrightarrow{\sigma}(u)$ of a node u is described by the standard deviation of the trust that u puts in its trustees. i.e,

$$\overrightarrow{\sigma}(u) = \frac{2 \left[\overrightarrow{\Gamma}(u) \neq \emptyset \right]}{\mathcal{R}} \sqrt{\frac{1}{|\overrightarrow{\Gamma}(u)|} \sum_{v \in \overrightarrow{\Gamma}(u)} (\mathcal{W}(u, v) - \overrightarrow{\mu}(u))^2}. \quad (2)$$

Remark. The values of $\overleftarrow{\sigma}(v)$ and $\overrightarrow{\sigma}(u)$ are in $[0, +1]$. This is because the standard deviation is always in the interval $[0, \mathcal{R}/2]$ (Al-Saleh and Yousif, 2016). The bigger $\overleftarrow{\sigma}(v)$ is, the more controversial v is. The bigger $\overrightarrow{\sigma}(u)$ is, the more eclectic u is.

Definition 3.3. A node u is said to be acting with *reciprocity* if it reciprocates the trust that it receives from its trustors. The *reciprocity* $\rho(u)$ of a node u is calculated as follows:

$$\rho(u) = \left[\overrightarrow{\Gamma}(u) \neq \emptyset \right] \left(1 - \sum_{v \in \overrightarrow{\Gamma}(u)} \frac{|\mathcal{W}(u, v) - \mathcal{W}(v, u)|}{\mathcal{R} |\overrightarrow{\Gamma}(u)|} \right). \quad (3)$$

The reader can easily verify that $\rho(u) \in [0, +1]$. The more u reciprocates trust, the more $\rho(u)$ converges toward $+1$. The less it does, the more $\rho(u)$ converges toward 0 .

3.4. Trust prediction as a tug of war game

We consider that the decision to trust is influenced by these social traits that we have defined. The rationale behind this idea is summarized by the following observations:

1. The less controversial v is, the more likely that $\mathcal{W}(u, v)$ will be equal to $\overleftarrow{\mu}(v)$.
2. The less eclectic u is, the more likely that $\mathcal{W}(u, v)$ will be equal to $\overrightarrow{\mu}(u)$.
3. The more u acts with reciprocity, the more likely that $\mathcal{W}(u, v)$ will be equal to $\mathcal{W}(v, u)$.

As illustrated in Figure 1, each of these three traits is trying to *pull* the value $\mathcal{W}(u, v)$ toward its respective extreme value. These *forces* apply only when their respective extreme values are meaningful. Therefore, we define:

$$\begin{cases} \overleftarrow{\zeta}(v) = (1 - \overleftarrow{\sigma}(v)) [\overleftarrow{\Gamma}(v) \neq \emptyset] \\ \overrightarrow{\zeta}(u) = (1 - \overrightarrow{\sigma}(u)) [\overrightarrow{\Gamma}(u) \neq \emptyset] \\ \varrho(u, v) = \rho(u) [\mathcal{W}(v, u) \neq 0] \end{cases} \quad (4)$$

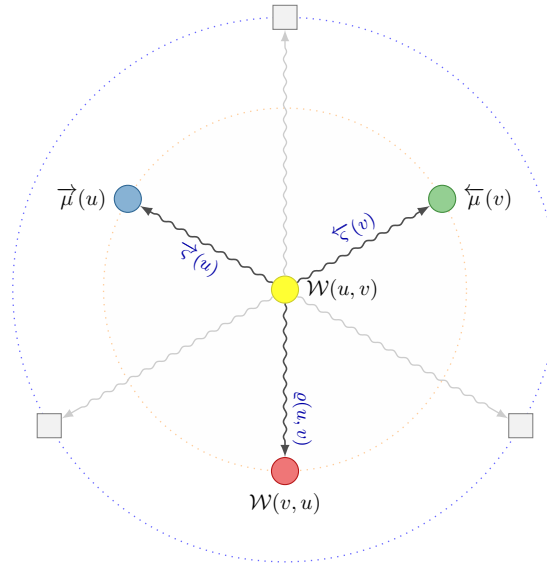


Figure 1. Trust as a three-way *tug of war*. The decision to trust/distrust is a *struggle* between three traits of the two actors: controversy (trustee), eclecticism (trustor), and reciprocity (trustor). Each one of these traits is a *force* that pulls the trust value $\mathcal{W}(u, v)$ toward its associated extreme value. For instance, the absence of controversy of v pulls $\mathcal{W}(u, v)$ toward $\overleftarrow{\mu}(v)$, the absence of eclecticism of u pulls this value toward $\overrightarrow{\mu}(u)$, and the reciprocity of u pulls it toward $\mathcal{W}(v, u)$. These forces and their associated extreme values can be simply and quickly computed. Other forces and extreme values (gray squares in the figure) may also apply but they would be more complex and slower to compute.

To predict the trust $\mathcal{W}(u, v)$ that the node u would put in v , we weight each of these extreme possible values ($\overleftarrow{\mu}(v)$, $\overrightarrow{\mu}(u)$, and $\mathcal{W}(v, u)$) by their respective forces: $\overleftarrow{\zeta}(v)$, $\overrightarrow{\zeta}(u)$, and $\varrho(u, v)$. That is:

$$\mathcal{W}(u, v) = \frac{\overleftarrow{\zeta}(v) \overleftarrow{\mu}(v) + \overrightarrow{\zeta}(u) \overrightarrow{\mu}(u) + \varrho(u, v) \mathcal{W}(v, u)}{\overleftarrow{\zeta}(v) + \overrightarrow{\zeta}(u) + \varrho(u, v)}. \quad (5)$$

The time complexity of this algorithm when trying to predict the trust that a node u would put in another one v is $\mathcal{O}(|\overrightarrow{\Gamma}(u)| + |\overleftarrow{\Gamma}(v)|)$.

4. Experimental evaluation

4.1. Datasets description

Following are the four datasets that we have used in our experiments. The first three originate from the *Stanford Large Network Dataset Collection*¹, and the last one from *Trustlet*². Table 2 summarizes some statistics of these datasets.

Bitcoin Alpha and OTC : Bitcoin is an anonymous cryptocurrency which is used by people around the globe to trade goods and services. Unfortunately, anonymity has some serious drawbacks such as the risk of fraud. These risks led to the emergence of some websites where users rate each other based on the trust that they put in them. We use two datasets from two sites: Bitcoin-Alpha and Bitcoin-OTC that were collected and scaled by Kumar et al. (2016) to fit in the interval $[-1, +1]$.

Wikipedia-Rfa : To be elected as a Wikipedia administrator, a user (or another member of the community on his behalf) submits a *Request for adminship* (Rfa). Members of the community cast their votes on the Rfa by a rating (+1 positive, 0 neutral, or -1 negative) and a comment. Kumar et al. (2016) analyzed these comments using the VADER sentiment engine (Gilbert, 2014) and generated a WSN with weights in $[-1, +1]$.

Robots.net : Robots.net is a community of robotics enthusiasts who certify each other using the levels *observer*, *apprentice*, *journeyer*, or *master*. We mapped these levels to real numbers (0.1, 0.4, 0.7, and 0.9 respectively). We ended up with a dataset, that while unsigned, allowed us to see how the studied algorithms behave with such data.

Network	Nodes	Arcs	Reciprocity Ratio
Bitcoin-Alpha	3783	24186	59.57%
Bitcoin-OTC	5881	35592	56.89%
Wikipedia-Rfa	9654	104554	0.07%
Robots.net	1725	3596	7.78%

Table 2. Statistics about the used datasets. The *reciprocity ratio* column indicates how much trust is reciprocated in the network. That is, the percentage of arcs such that $\mathcal{W}(u, v) = \mathcal{W}(v, u)$.

4.2. Evaluated algorithms

To evaluate the performances and the efficiency of the proposed approach, we conduct various tests using the following algorithms:

Reciprocal the easiest to implement for it relies on the assumption that if a node u trusts another node v , then v will probably trust u back as much as it trusts it. i.e., $\mathcal{W}(u, v) = \mathcal{W}(v, u)$ if the arc from v back to u exists, and 0 otherwise.

¹<http://snap.stanford.edu/data/>

²<http://www.trustlet.org/datasets/>

Bias and Deserve (BaD) we took $\text{DESERVE}(v)$ as the inferred trust value as described in Mishra and Bhattacharya (2011).

Fairness-Goodness (FxG) the trust from u to v is the product of the FAIRNESS of u by the GOODNESS of v as proposed by Kumar et al. (2016).

STAR we took the inferred trust value as proposed by Gao et al. (2016).

Tug of War (ToW) our own approach described in Section 3.

4.2.1. Performance evaluation metrics

Given a network $G(\mathcal{N}, \mathcal{E}, \mathcal{W})$ with $|\mathcal{N}| = N$. To compare the performances of the above algorithms, we have used the following metrics:

Mean Absolute Error (MAE) is the mean of the absolute differences between the actual trust values x_i and the inferred ones $y_i, i \in [1 \cdots N]$:

$$\text{MAE} = \frac{1}{N} \sum_{i=1}^N |x_i - y_i|.$$

Root Mean Squared Error (RMSE) is the root mean of the squared differences between the actual trust values x_i and the inferred ones $y_i, i \in [1 \cdots N]$:

$$\text{RMSE} = \sqrt{\frac{1}{N} \sum_{i=1}^N (x_i - y_i)^2}.$$

Pearson Correlation Coefficient (PCC) ranges between -1 and $+1$ and indicates how the actual trust values x_i correlate with the predicted ones y_i . The more the PCC converges toward $+1$, the more the two values are correlated:

$$\text{PCC} = \frac{\sum_{i=1}^N (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_{i=1}^N (x_i - \bar{x})^2} \sqrt{\sum_{i=1}^N (y_i - \bar{y})^2}},$$

where \bar{x} and \bar{y} are the arithmetic means of x_i and y_i respectively, and $i \in [1 \cdots N]$.

4.3. Experiments

Inspired by the work done by Kumar et al. (2016), we studied two scenarios during these experiments. First, we wanted to see how the evaluated algorithms would predict a yet-to-exist arc's weight given the rest of the network's arcs weights. Second, we studied how these algorithms behave when some, or most, of the network arcs are invisible (or unavailable). Note that unlike Kumar et al. (2016) who did *weight prediction* of edges that already exist, we are trying to predict weights of arcs that do not exist yet. Following are the details and the results of these experiments.

4.3.1. Leave-One-Out predictions

This is a classic trust prediction task in social networks. It consists in hiding an arc from the trust graph, and predicting its weight. It answers the question “How much would a node u (dis)trust another node v ” given all the other trust values in the network.

We have run the five algorithms on every arc of the used datasets. For every dataset, we have removed one arc at a time, and have predicted its weight. We then calculated the MAE, RMSE, and PCC metrics for every pair of algorithm and dataset. As shown in Table 3, the proposed approach outperforms all the other algorithms on every metric (MAE, RMSE, PCC), and on every dataset.

	Bitcoin-Alpha	Bitcoin-OTC	Wikipedia-Rfa	Robots.net
Reciprocal	(0.12, 0.27, 0.47)	(0.15, 0.32, 0.46)	(0.56, 0.63, 0.071)	(0.50, 0.60, 0.01)
FxG	(0.19, 0.33, 0.24)	(0.22, 0.38, 0.31)	(0.18, 0.24, 0.43)	(0.28, 0.33, -0.04)
BaD	(0.20, 0.34, 0.24)	(0.23, 0.40, 0.32)	(0.18, 0.23, 0.44)	(0.21, 0.31, 0.15)
STAR	(0.21, 0.32, 0.25)	(0.23, 0.35, 0.22)	(0.24, 0.32, 0.22)	(0.43, 0.50, 0.44)
ToW	(0.12, 0.24, 0.60)	(0.14, 0.27, 0.66)	(0.17, 0.22, 0.54)	(0.16, 0.22, 0.48)

Table 3. Results from the leave One-Out tests. Inside each cell of this table is a tuple (MAE, RMSE, PCC) of the output of an algorithm (row) on a dataset (column). Lower MAE and RMSE, and higher PCC, are better.

Efficiency comparison. While doing the previous experiment, we took the opportunity to study the efficiency of the five algorithms. We measured the time they take to do a *leave-one-out* prediction on every arc of the four used datasets. These tests were performed on an intel(R) i5-2450M CPU with 8GB of RAM. Figure 2 summarizes the time taken by these algorithms on the four datasets. It clearly demonstrates that the proposed approach is very close in terms of speed to the reciprocal algorithm which, obviously, is the fastest since its time complexity on a single arc is literally $\mathcal{O}(1)$. In fact, apart from the reciprocal algorithm, our approach is two orders of magnitude faster than the other three algorithms.

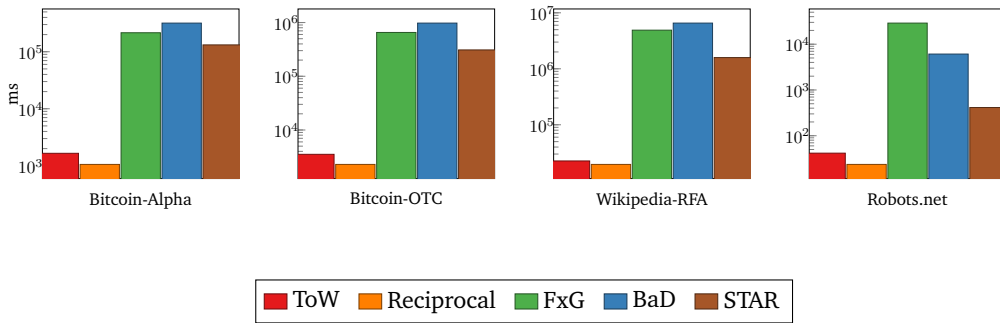


Figure 2. The time taken by each studied algorithm to perform a *leave-one-out* prediction on every arc of the used datasets. The y-axis is logarithmic and the durations are expressed in powers of 10 milliseconds.

4.3.2. Leave $\mathcal{N}\%$ Out predictions

How good would be trust predictions when some, or most, of the network’s arcs are not available? Reasons for such cases are various. They may range from privacy concerns to technical and performances limitations —especially for very large networks. Compared to the previous experiments, the quality of trust predictions in sparse networks may degrade as there is not much prior information to infer from.

In order to understand how these algorithms are affected by sparsity, we randomly remove $\mathcal{N}\%$ arcs at once, and try to predict them. Specifically, we have taken out 10%, 20%, and so on, up to 90% arcs of every dataset, and have tried to predict their weights. We have repeated these tests 100 times for every percent, dataset, and algorithm. We then calculated the average MAE, RMSE, and PCC for each algorithm, on each dataset with $\mathcal{N}\%$ hidden arcs. As reported in Figure 3, compared to the four other ones, the proposed approach provides the best MAE, RMSE, and PCC in every case, and is barely affected by the number of removed arcs. That is, our approach is sufficiently robust to networks sparsity.

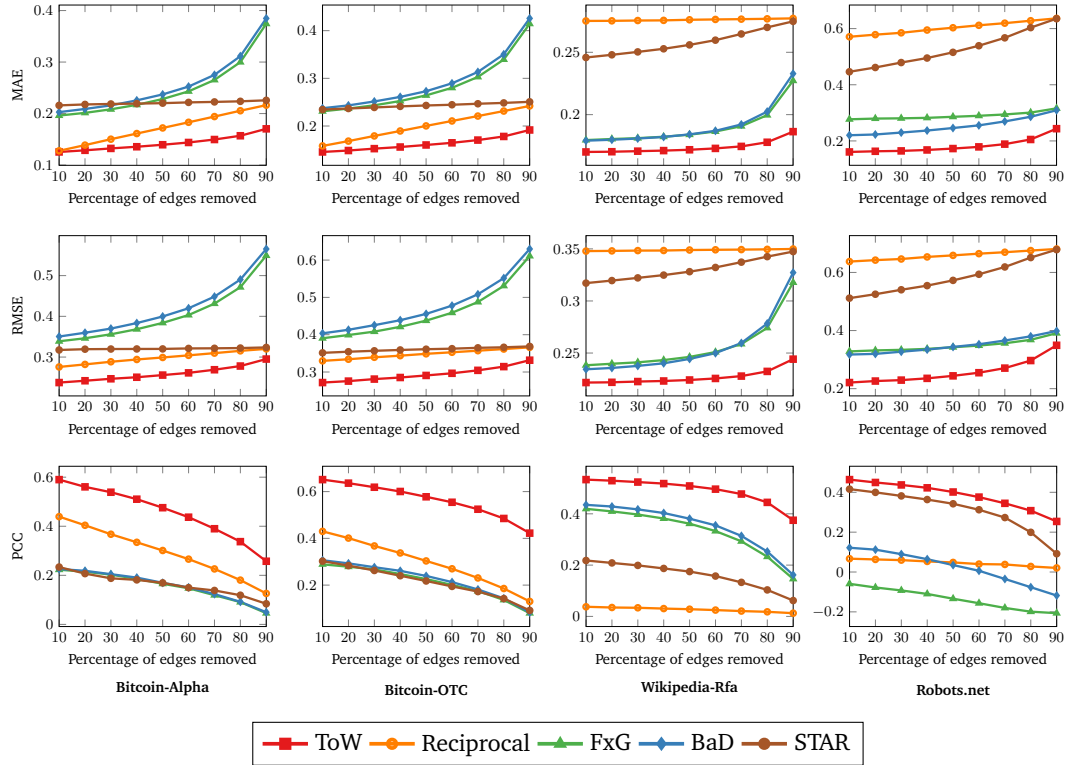


Figure 3. Leave- $\mathcal{N}\%$ -Out results. A grid of plots where: the x-axis are the percentage of removed edges (going from 10% up to 90% in steps of 10%). The three rows of plots describe how the MAE, RMSE, and PCC respectively change as we remove more edges from the datasets. Each column represents the results for a given dataset.

5. Discussion

Apart from handling distrust, we have considered in the introduction that trust prediction algorithms should be 1) as accurate as possible, 2) as robust to sparsity as possible, and 3) as fast as possible. In what follows, we discuss how these properties hold for the investigated methods.

5.1. Accuracy

Results from the two types of experiments show that the proposed approach outperforms the other studied ones. As shown in Table 3, it is the most accurate one among the investigated methods in a *leave-one-out* scenario. We should, however, highlight the fact that with the bitcoin datasets, the proposed and the reciprocal algorithms are very close in terms of MAE, and RMSE in a *leave-one-out* setting. This is explained by the fact that these two networks present a higher trust reciprocity compared to the Wikipedia-RFA and Robots.net networks (cf. Table 2). This difference may be due to the very nature of these networks. Indeed, the bitcoin networks are all about trust in trade. When a transaction goes as expected, both parties are generally satisfied and this leads to reciprocal ratings. On the other hand, networks like Wikipedia-RFA and Robots.net are about authority and skills. A highly rated expert has not to reciprocate the rating from a less-knowing node. With this said, in terms of PCC, our approach is clearly better. That is, the predicted trust values and the original ones are more correlated. A possible explanation of this result is that the proposed algorithm considers reciprocity *and* two other traits which the reciprocal algorithm does not take into account. When a node does not act with enough reciprocity, this trait will not weight a lot in Eq. (5). Therefore, the proposed algorithm will rely more on controversy and eclecticism.

5.2. Robustness

As revealed by the plots in Figure 3, our approach shows a very slow increase of MAE and RMSE and a slow decrease of PCC as we remove more arcs from the datasets. By contrast, other algorithms seem to be more sensitive to network sparsity. This may be explained by the design of these algorithms. For instance, the reciprocal algorithm relies on the hypothesis that a node likely reciprocates received trust. However, as we remove more links, reciprocal links are also removed and the inferred value is null. As for the STAR algorithm, one would think that links removal would break trust paths, but to our surprise, this algorithm shows a somewhat stable behavior. This may be attributed to the fact that arcs that were not considered when better ones (those with higher certainty values) were available are used when the better ones are removed. Still, its accuracy is inferior to that of the proposed approach. In fact, if we take a closer look at the plots in Figure 3, we notice that the proposed approach provides more accurate predictions at 90% removed arcs than what STAR provides at 10% removed arcs. Finally, approaches that rely on global metrics (BaD and FxG) need most, if not all, of the network to be visible to accurately calculate the nodes' characteristics. Indeed, the pairs of metrics in these algorithms (BIAS/DESERVE, FAIRNESS/GOODNESS for the BaD and FxG algorithms respectively) depend mutually on each other. Their incremental calculation spreads to large portions of the networks. On the other hand, the proposed metrics are more *local* since they depend only on the direct neighbors of the trustor and those of the trustee.

5.3. Speed

The previous observation leads us to the last point: speed. The reciprocal algorithm’s complexity for predicting a single arc is obviously $\mathcal{O}(1)$. The other algorithms are slower in comparison. However, while the global metrics ones spread to larger portions of the network, and while the STAR algorithm performs a graph traversal, our approach does not and is, hence, faster. Furthermore, knowing that trust fluctuates a lot in networks with high activity, it is worth mentioning that updates to global metrics will encompass most of the network. Updates to the proposed metrics, on the other hand, are limited to the trustor and the trustee neighborhoods, and are therefore faster.

Note also that the three proposed metrics can be independently calculated. As such, trust can be easily predicted using these metrics in parallel or in a distributed setting. By contrast, the other studied algorithms are harder to parallelize because of their design.

5.4. Final thoughts and perspectives

These experimental results prompt us to discuss trust prediction in general. Trust transitivity, as described by Guha et al. (2004) and others, states that if an individual u trusts another one v which itself trusts a third one w , then u *might* trust w to some extent. The emphasis on “*might*” and “*to some extent*” is important in this context and hints that this transitivity is hypothetical and not guaranteed. In fact, trust transitivity should adhere to some semantic requirements to be meaningful (Jøsang and Pope, 2005). This, and knowing that 1) trust itself is known to decay on long paths (which is understandable) (Liu et al., 2011), and 2) that multiple paths from a source to a sink may provide conflicting predictions (Jøsang and Pope, 2005), make transitivity-based predictions subject to a lot of uncertainty. In short, predicting trust by transitivity builds on intermediary and uncertain predictions to produce a final (and often more uncertain) one³. Actually, this aspect is why most propagative trust prediction approaches limit their processing to a small number of hops. It is not just about efficiency, but about accuracy (that decreases beyond some hops) (Golbeck, 2005a; Ziegler and Golbeck, 2015).

Now, compared to transitivity, local metrics such as the ones that we propose have more *truth* to them. They are computed locally using known relations from the inner circle of an individual. More specifically, the proposed approach has no dependence on transitivity nor on a hypothetical aspect of trust other than it being the result of *struggle* between social traits of the involved nodes in a trust relation. Such an aspect is not subject to decay nor conflicts, as opposed to transitivity. To put it briefly, we believe that there is more to know about, and from, the direct neighbors of an individual than we can gather from distant and not directly-related individuals. Naturally, that is not to say that we should *ditch* transitivity altogether, but that we should extensively explore the direct neighborhoods of the individuals in a trust relation, and look for more social traits and biases that affect trust. And *this* is, in fact, the essence of the proposed approach that, rather than propagating trust along paths from a trustor u to a trustee v , is built on the assumption that if we know *everything* trust-related about these individuals, then we should be able to know how one of them would trust the other.

³As argued by Jøsang (2016), if ω_X^B is the opinion of B about X , and ω_B^A is the opinion of A about B , then the opinion of A about X (denoted as $\omega_X^{[A;B]} = \omega_B^A \otimes \omega_X^B$) typically gets increased uncertainty mass, compared to the original opinion advised by B .

Despite the satisfying results, and the simplicity, of the proposed approach, there are some limitations that we should acknowledge, and that would be worth addressing in a future work. First and foremost, we considered only three traits in our work, however the *whole truth* about the involved individuals in a trust relation—and the relation itself—cannot be contained in these three traits. There should be more traits that decide on how we *give* and *receive* trust. Sure enough, the three traits that we have used gave competitively reliable, robust, and efficient predictions on the four used datasets, but we do think that there is still room for a lot of improvement in terms of accuracy by investigating other traits; all while keeping the balance between performances, robustness, and efficiency.

Another limitation of this work is the assumption that the three social traits that we have considered are equally important across all social networks. Indeed, finding out that the Bitcoin networks—by their nature—present a high reciprocity ratio compared to the other used two, makes us wonder whether associating various degrees of importance to each trait (depending on the network) would give even more accurate results. In other words, if social traits are *forces* that compete to affect trust, then these degrees of importance would be the properties of the *ground* where this struggle takes place, and that favors some traits over others. Statistical methods may be used to calculate these importance degrees, and thus describe the nature of the network more precisely than what we already did with the “reciprocity ratio” column in Table 2. In summary, characterizing a social network as a whole—not only its members—is, in our opinion, worth investigating.

6. Conclusion

We have explored in this paper the possibility of using some nodes traits such as controversy, eclecticism, and reciprocity as *forces* that affect the act of trust. Our intuition was to picture trust (and distrust) as a *tug of war* game where both the trustor and the trustee have their say. Where each one of them tries to pull the other opponent to their side by using their own traits.

Our mathematical description of these social traits allowed us to design a very fast and robust algorithm to predict trust and distrust in weighted signed social networks. Indeed, experimental evaluation of two scenarios (leave-one-out, and leave- $\mathcal{N}\%$ -out) show that our approach presents good and stable performances. Being localized, rather than propagative, turned out to be an advantage both in terms of speed and robustness. In addition to these benefits, it is worth noting that the proposed algorithm is absurdly easy to implement since it consists in calculating means and standard deviations.

As a future work, we would like to explore more traits that can affect trust. Or, to carry on with the adopted *tug of war* analogy, we would like to add more *ropes* to the game as illustrated in Figure 1. Some of these new traits may require complex and slower computations, but we think that they are still worth exploring in order to improve the accuracy and the robustness of our approach. Also worth investigating are the properties that characterize social networks as a whole. These properties might indeed improve the accuracy of the proposed approach, by giving various degrees of importance to the used social traits that affect trust relations.

7. References

- Al-Saleh, M.F., Yousif, A.E., 2016. Properties of the standard deviation that are rarely mentioned in classrooms. *Austrian Journal of Statistics* 38, 193–202. doi:<https://doi.org/10.17713/ajs.v38i3.272>.
- Bauer, P.C., 2017. Conceptualizing trust and trustworthiness doi:10.2139/ssrn.2325989.
- Beigi, G., Tang, J., Wang, S., Liu, H., 2016. Exploiting emotional information for trust/distrust prediction, in: *Proceedings of the 2016 SIAM International Conference on Data Mining*, SIAM. pp. 81–89. doi:10.1137/1.9781611974348.10.
- Chiang, K.Y., Hsieh, C.J., Natarajan, N., Dhillon, I.S., Tewari, A., 2014. Prediction and clustering in signed networks: a local to global perspective. *The Journal of Machine Learning Research* 15, 1177–1213.
- Chiluka, N., Andrade, N., Gkorou, D., Pouwelse, J., 2012. Personalizing eigentrust in the face of communities and centrality attack, in: *2012 IEEE 26th International Conference on Advanced Information Networking and Applications*, pp. 503–510. doi:10.1109/AINA.2012.48.
- Cho, J., 2006. The mechanism of trust and distrust formation and their relational outcomes. *Journal of retailing* 82, 25–35. doi:10.1016/j.jretai.2005.11.002.
- DuBois, T., Golbeck, J., Srinivasan, A., 2011. Predicting trust and distrust in social networks, in: *Privacy, Security, Risk and Trust (PASSAT) and 2011 IEEE Third International Conference on Social Computing (Social-Com)*, 2011 IEEE Third International Conference on, IEEE. pp. 418–424. doi:10.1109/PASSAT/SocialCom.2011.56.
- El Ghalid, M., 2017. How fungi recognize (and infect) plants. URL: https://www.ted.com/talks/mennat_el_ghalid_how_fungi_recognize_and_infect_plants. last visited: 2018-07-30.
- Faulkner, P., 2014. The practical rationality of trust. *Synthese* 191, 1975–1989. doi:10.1007/s11229-012-0103-1.
- Gao, P., Miao, H., Baras, J.S., Golbeck, J., 2016. Star: semiring trust inference for trust-aware social recommenders, in: *Proceedings of the 10th ACM Conference on Recommender Systems*, ACM. pp. 301–308. doi:10.1145/2959100.2959148.
- Gilbert, C.H.E., 2014. Vader: A parsimonious rule-based model for sentiment analysis of social media text.
- Golbeck, J., 2005a. Personalizing applications through integration of inferred trust values in semantic web-based social networks, in: *Semantic Network Analysis Workshop at the 4th International Semantic Web Conference*, p. 30.
- Golbeck, J.A., 2005b. *Computing and Applying Trust in Web-based Social Networks*. Ph.D. thesis. College Park, MD, USA. AAI3178583.
- Graham, R.L., Knuth, D.E., Patashnik, O., 1994. *Concrete mathematics a foundation for computer science*. Addison-Wesley.
- Guha, R., Kumar, R., Raghavan, P., Tomkins, A., 2004. Propagation of trust and distrust, in: *Proceedings of the 13th International Conference on World Wide Web*, ACM, New York, NY, USA. pp. 403–412. doi:10.1145/988672.988727.
- Guo, S.L., Lumineau, F., Lewicki, R.J., et al., 2017. Revisiting the foundations of organizational distrust. *Foundations and Trends® in Management* 1, 1–88. doi:10.1561/34000000001.
- Hamdi, S., Gancarski, A.L., Bouzeghoub, A., Yahia, S.B., 2016. Tison: Trust inference in trust-oriented social networks. *ACM Transactions on Information Systems (TOIS)* 34, 17. doi:10.1145/2858791.
- Hawley, K., 2013. Trust, distrust and commitment. *Noûs* 48, 1–20. doi:10.1111/nous.12000.
- Hu, Y., Wang, S., Ren, Y., Choo, K.K.R., 2018. User influence analysis for github developer social networks. *Expert Systems with Applications* 108, 108 – 118. doi:10.1016/j.eswa.2018.05.002.
- Huang, H., Dong, Y., Tang, J., Yang, H., Chawla, N.V., Fu, X., 2018. Will triadic closure strengthen ties in social networks? *ACM Trans. Knowl. Discov. Data* 12, 30:1–30:25. doi:10.1145/3154399.
- Huang, Z., Olteanu, A., Aberer, K., 2013. Credibleweb: A platform for web credibility evaluation, in: *CHI '13 Extended Abstracts on Human Factors in Computing Systems*, ACM, New York, NY, USA. pp. 1887–1892. doi:10.1145/2468356.2468694.
- Jiang, W., Wang, G., Bhuiyan, M.Z.A., Wu, J., 2016a. Understanding graph-based trust evaluation in online social networks: Methodologies and challenges. *ACM Computing Surveys* 49, 10:1–10:35. doi:10.1145/2906151.
- Jiang, W., Wang, G., Wu, J., 2014. Generating trusted graphs for trust evaluation in online social networks. *Future Generation Computer Systems* 31, 48 – 58. doi:10.1016/j.future.2012.06.010. special Section: *Advances in Computer Supported Collaboration: Systems and Technologies*.
- Jiang, W., Wu, J., Li, F., Wang, G., Zheng, H., 2016b. Trust evaluation in online social networks using generalized network flow. *IEEE Transactions on Computers* 65, 952–963. doi:10.1109/TC.2015.2435785.
- Jones, J.J., Settle, J.E., Bond, R.M., Fariss, C.J., Marlow, C., Fowler, J.H., 2013. Inferring tie strength from online directed behavior. *PLOS ONE* 8, 1–6. doi:10.1371/journal.pone.0052168.

- Jones, K., 1996. Trust as an affective attitude. *Ethics* 107, 4–25.
- Jøsang, A., 2016. Computational trust, in: *Subjective Logic*. Springer, pp. 243–270.
- Jøsang, A., Pope, S., 2005. Semantic constraints for trust transitivity, in: *Proceedings of the 2nd Asia-Pacific conference on Conceptual modelling-Volume 43*, Australian Computer Society, Inc.. Australian Computer Society, Inc., Darlinghurst, Australia, Australia. pp. 59–68.
- Jøsang, A., 1999. An algebra for assessing trust in certification chains, in: *Proceedings of the Network and Distributed Systems Security Symposium (NDSS'99)*. The Internet Society, p. 80.
- Kamvar, S.D., Schlosser, M.T., Garcia-Molina, H., 2003. The eigentrust algorithm for reputation management in p2p networks, in: *Proceedings of the 12th International Conference on World Wide Web*, ACM, New York, NY, USA. pp. 640–651. doi:10.1145/775152.775242.
- de Kerchove, C., Dooren, P.V., 2008. The PageTrust algorithm: How to rank web pages when negative links are allowed? pp. 346–352. doi:10.1137/1.9781611972788.31.
- Kleinberg, J.M., 1999. Authoritative sources in a hyperlinked environment. *Journal of the ACM (JACM)* 46, 604–632. doi:10.1145/324133.324140.
- Knuth, D.E., 1992. Two notes on notation. *The American Mathematical Monthly* 99, 403–422. doi:10.2307/2325085.
- Kramer, R.M., Cook, K.S., 2004. *Trust and distrust in organizations: Dilemmas and approaches*. Russell Sage Foundation.
- Kumar, S., Spezzano, F., Subrahmanian, V., Faloutsos, C., 2016. Edge weight prediction in weighted signed networks, in: *Data Mining (ICDM), 2016 IEEE 16th International Conference on*, IEEE. pp. 221–230. doi:10.1109/ICDM.2016.0033.
- Kunegis, J., Preusse, J., Schwagereit, F., 2013. What is the added value of negative links in online social networks?, in: *Proceedings of the 22Nd International Conference on World Wide Web*, ACM, New York, NY, USA. pp. 727–736. URL: <http://doi.acm.org/10.1145/2488388.2488452>, doi:10.1145/2488388.2488452.
- Kurdi, H.A., 2015. Honestpeer: An enhanced eigentrust algorithm for reputation management in p2p systems. *Journal of King Saud University - Computer and Information Sciences* 27, 315 – 322. doi:10.1016/j.jksuci.2014.10.002.
- Lewicki, R.J., Brinsfield, C., 2012. Measuring trust beliefs and behaviours. *Handbook of research methods on trust* 29.
- Lewicki, R.J., McAllister, D.J., Bies, R.J., 1998. Trust and distrust: New relationships and realities. *Academy of management Review* 23, 438–458. doi:10.5465/amr.1998.926620.
- Liu, G., Wang, Y., Orgun, M.A., et al., 2011. Trust transitivity in complex social networks., in: *AAAI*, pp. 1222–1229.
- Massa, P., Avesani, P., 2005. Controversial users demand local trust metrics: An experimental study on epinions.com community, in: *Proceedings of the 20th National Conference on Artificial Intelligence - Volume 1*, AAAI Press. pp. 121–126.
- Massa, P., Avesani, P., 2007. Trust metrics on controversial users: Balancing between tyranny of the majority. *International Journal on Semantic Web and Information Systems (IJSWIS)* 3, 39–64.
- Matei, S.A., Bertino, E., Zhu, M., Liu, C., Si, L., Britt, B., 2015. *A Research Agenda for the Study of Entropic Social Structural Evolution, Functional Roles, Adhocratic Leadership Styles, and Credibility in Online Organizations and Knowledge Markets*. Springer International Publishing, Cham. pp. 3–33. URL: https://doi.org/10.1007/978-3-319-05467-4_1, doi:10.1007/978-3-319-05467-4_1.
- Mayer, R.C., Davis, J.H., Schoorman, F.D., 1995. An integrative model of organizational trust. *Academy of management review* 20, 709–734.
- McKnight, D.H., Choudhury, V., 2006. Distrust and trust in b2c e-commerce: Do they differ?, in: *Proceedings of the 8th international conference on Electronic commerce: The new e-commerce: innovations for conquering current barriers, obstacles and limitations to conducting successful business on the internet*, ACM. pp. 482–491. doi:10.1145/1151454.1151527.
- Mishra, A., Bhattacharya, A., 2011. Finding the bias and prestige of nodes in networks based on trust scores, in: *Proceedings of the 20th international conference on World wide web*, ACM. pp. 567–576. doi:10.1145/1963405.1963485.
- Page, L., Brin, S., Motwani, R., Winograd, T., 1999. *The PageRank citation ranking: Bringing order to the web*. Technical Report. Stanford InfoLab.
- Papaoikonomou, T., Kardara, M., Tserpes, K., Varvarigou, T., 2013. The strength of negative opinions, in: *Iliadis, L., Papadopoulos, H., Jayne, C. (Eds.), Engineering Applications of Neural Networks*, Springer Berlin Heidelberg, Berlin, Heidelberg. pp. 90–99. doi:10.1007/978-3-642-41016-1_10.
- Robbins, B.G., 2016. What is trust? a multidisciplinary review, critique, and synthesis. *Sociology Compass* 10, 972–986. doi:10.1111/soc4.12391.

- Rousseau, D.M., Sitkin, S.B., Burt, R.S., Camerer, C., 1998. Not so different after all: A cross-discipline view of trust. *Academy of Management Review* 23, 393–404. doi:10.5465/amr.1998.926617.
- Schoorman, F.D., Mayer, R.C., Davis, J.H., 2007. An integrative model of organizational trust: Past, present, and future. *Academy of Management review* 32, 344–354. doi:10.5465/amr.2007.24348410.
- Shahriari, M., Jalili, M., 2014. Ranking nodes in signed social networks. *Social Network Analysis and Mining* 4, 172. doi:10.1007/s13278-014-0172-x.
- Shneiderman, B., 2015. Building trusted social media communities: A research roadmap for promoting credible content, in: *Roles, trust, and reputation in social media knowledge markets*. Springer, pp. 35–43. doi:10.1007/978-3-319-05467-4_2.
- Tang, J., Aggarwal, C., Liu, H., 2016a. Node classification in signed social networks, in: *Proceedings of the 2016 SIAM International Conference on Data Mining, SIAM*. pp. 54–62. doi:10.1137/1.9781611974348.7.
- Tang, J., Chang, Y., Aggarwal, C., Liu, H., 2016b. A survey of signed network mining in social media. *ACM Computing Surveys* 49, 42:1–42:37. doi:10.1145/2956185.
- Tang, J., Hu, X., Liu, H., 2014. Is distrust the negation of trust?: The value of distrust in social media, in: *Proceedings of the 25th ACM Conference on Hypertext and Social Media, ACM, New York, NY, USA*. pp. 148–157. doi:10.1145/2631775.2631793.
- Tang, J., Liu, H., 2015. Trust in social media. *Synthesis Lectures on Information Security, Privacy, & Trust* 10, 1–129. doi:10.2200/S00657ED1V01Y201507SPT013.
- Zhao, H., Xu, X., Song, Y., Lee, D.L., Chen, Z., Gao, H., 2018. Ranking users in social networks with higher-order structures, in: *AAAI Conference on Artificial Intelligence*.
- Ziegler, C.N., 2013. *Trust Propagation Models*. Springer International Publishing, Cham. pp. 99–131. doi:10.1007/978-3-319-00527-0_7.
- Ziegler, C.N., Golbeck, J., 2015. Models for trust inference in social networks, in: *Propagation Phenomena in Real World Networks*. Springer, pp. 53–89.
- Ziegler, C.N., Lausen, G., 2005. Propagation models for trust and distrust in social networks. *Information Systems Frontiers* 7, 337–358. doi:10.1007/s10796-005-4807-3.
- Zolfaghar, K., Aghaie, A., 2010. Mining trust and distrust relationships in social web applications, in: *Intelligent Computer Communication and Processing (ICCP), 2010 IEEE International Conference on, IEEE*. pp. 73–80. doi:10.1109/ICCP.2010.5606460.