Multi-Faceted Trust and Distrust Prediction for Recommender Systems

Hui Fang, Guibing Guo*, Jie Zhang

School of Computer Engineering, Nanyang Technological University, Singapore

Abstract

Many trust-aware recommender systems have explored the value of explicit trust, which is specified by users with binary values and simply treated as a concept with a single aspect. However, in social science, trust is known as a complex term with multiple facets, which have not been well exploited in prior recommender systems. In this paper, we attempt to address this issue by proposing a (dis)trust framework with considerations of both interpersonal and impersonal aspects of trust and distrust. Specifically, four interpersonal aspects (benevolence, competence, integrity and predictability) are computationally modelled based on users' historic ratings, while impersonal aspects are formulated from the perspective of user connections in trust networks. Two logistic regression models are developed and trained by accommodating these factors, and then applied to predict continuous values of users' trust and distrust, respectively. Trust information is further refined by corresponding predicted distrust information. The experimental results on real-world data sets demonstrate the effectiveness of our proposed model in further improving the performance of existing state-of-the-art trust-aware recommendation approaches.

Keywords: Trust, Distrust, Rating Behavior, Multi-facet, Recommender Systems

1. Introduction

Trust has been extensively exploited for improving the predictive accuracy of recommendations by ameliorating the issues such as *data sparsity* and *cold start* that recommender systems inherently suffer from [1, 18, 16, 3, 26, 9, 5]. In essence, trust provides additional information from which user preference can be better modelled, alternative or complementary to rating-based similarity. Both implicit [24] and explicit [18, 3, 16, 26, 9, 5] trust have been investigated in the literature. The former trust is usually inferred from user-item interactions (i.e., ratings) whereas the latter is directly specified by users indicating whom and to what extent they trust. In contrast, although distrust is recognized to play an equivalently important role as trust [22], the investigation of utilizing distrust in recommender systems is still in its infancy [30, 31]. To the best of our knowledge, no prior work has attempted to predict distrust for improving recommender systems.

Another issue of existent trust-aware recommender systems is the simplified modelling of trust as a concept with a single aspect, such as the ability to provide accurate ratings (known as *competence*) [24] or the probability of behaving maliciously. However, it is well acknowledged in social science that trust is a concept with multi-faceted properties [19, 21, 20]. One possible explanation is that only limited information is available in the few and publicly accessible data sets. Although some efforts have been made to capture multiple aspects (e.g. information credibility [12]) of raters (who give ratings) in recommender systems, they are essentially distinct concepts from trust. A generally agreed proposition states that people trusting each other may not always share similar preferences [10]. This statement leads to the following interesting research ques-

*Corresponding author. Address: Nanyang Technological University, 50 Nanyang Avenue, Singapore 639798. Tel.:+65-84005322

Email address: gguo1@e.ntu.edu.sg (Guibing Guo)

tion: *which aspects of (dis)trust reflect user preferences more and hence should be more considered for user preference modelling?* The answer would provide a guidance on whom and to what extent one can trust, especially given the fact that most available (i.e. explicit) trust scores are binary, i.e., either 1 (trust) or -1 (distrust) without specific degrees of trust or distrust.

In this paper, we aim to address the research question by proposing a framework of trust and distrust, taking into considerations both interpersonal and impersonal aspects of trust and distrust adapted from social science [20]. Specifically, four interpersonal aspects (i.e., benevolence, competence, integrity and predictability) are computationally modelled based on users' past ratings, while impersonal aspects (e.g., degree centrality) are formulated from the perspective of social links in trust networks. Note that the social links in a trust network consist of both trust and distrust connections among users. Two logistic regression models are developed and trained by accommodating these factors and then applied to predict continuous values of users' trust and distrust, respectively. We further refine the trust information using the predicted distrust information. These newly generated trust values can then be applied into the existing trust-aware recommender algorithms (i.e. TidalTrust, Merge and SocialMF). The experimental results on real-world data sets demonstrate the effectiveness of our proposed model for improving the performance of three representative trust-aware recommendation algorithms. In addition, the generality of our model is also empirically demonstrated. In all, our work is the first to comprehensively study the multiple aspects of trust and distrust in the context of recommender systems. The study results lead to refined trust and distrust predictions, and in consequence notable improvement on recommendation accuracy when the predicted trust and distrust are utilized in recommendation approaches.

The rest of the paper is organized as follows. Section 2 gives an overview of related research in the literature. Section 3 elaborates the proposed (dis)trust framework,

and Section 4 introduces the trust and distrust prediction models. The effectiveness of our approach is evaluated and discussed in Sections 5 and 6, respectively. Finally, the conclusion and future work are presented in Section 7.

2. Related Work

Both trust and distrust are well-known as heterogenous rather than homogenous concepts in the fields of social science and computational trust, each of which is composed of multiple aspects [19, 21]. Specifically, Mayer et al. [19] report that the trust relationship between a trustor (who specifies trust statements) and a trustee (who receives trust statements) is mainly influenced by the trustor's *propensity* to trust others in terms of three interpersonal aspects related with the trustee, namely *ability* (*competence*), *benevolence* and *integrity*. Mcknight and Chervany [21] enrich this model by adding one more aspect of the trustee—*predictability* as well as an impersonal aspect from the view of structural/institutional trust [20, 21]. Impersonal aspects are often utilized to predict positive or negative user links [14, 13] by virtue of the graph structures of social networks. We defer the formal definitions of these aspects till Section 3. These frameworks have been adopted as the underpinning of the socio-cognitive trust theory in the area of computational trust [2]. Consistently, in this work we employ both interpersonal and impersonal aspects of the trustee along with the trustor's propensity to formulate users' trust and distrust.

Trust is also applied in real applications, such as Epinions.com where users can explicitly specify other users as trustworthy or untrustworthy. The value of trust has been explored by many trust-aware recommender systems, given the strong and positive correlation between trust and preference [28]. For example, Donovan and Smyth [24] treat trust as a single aspect and equivalent with the expertise or competence of users. Massa and Anesani [18] replace user similarity with explicitly specified trust relation-

ships, and also allow trust relationships to propagate through the trust networks. They show that more robust recommendations can be produced without significant loss in accuracy. Golbeck [4] introduces a trust-flow-based method (called *TidalTrust*) to compute rating predictions for target items. She finds out that better accuracy can be achieved. Later works [3, 26] claim that better performance can be obtained by integrating both trust and similarity for recommendations. Jamali and Ester [8] design the *TrustWalker* approach to randomly select neighbors in the trust network formed by users and their trusted neighbors. TrustWalker combines trust information of the selected neighbors with an item-based technique, where both the ratings of the target item and similar items are considered. The recent work conducted by Guo et al. [5] focuses on the problems of data sparsity and cold start from which traditional recommender systems suffer. They empirically contend that by merging the ratings of trusted neighbors, the preferences of active users can be better modelled and hence the performance is improved.

Other than these neighborhood-based approaches, trust is also adopted in model-based approaches. For example, Ma et al. [17] design a latent factor model called *SoRec* based on probabilistic matrix factorization [23]. They fuse the user-item rating matrix with user-user trust matrix by sharing a common latent low dimensional user feature matrix. The two matrices are factorized by three sets of latent features: user vector and feature vector (for each user), and item vector. Experimental results demonstrate that SoRec outperforms the basic matrix factorization model and other trust related neighborhood models. However, although the trust information is considered, the real world recommendation processes are not reflected, where the two sets of latent features for each user cause the low interpretability of the model. To overcome this problem and model trust-aware recommender systems more realistically, they further propose *RSTE* [16], a linear combination of a basic matrix factorization technique and a trust-based approach. Jamali and Ester [9] later enhance this model by enabling trust propagation in their *SocialMF*

model. On the other hand, only very few works have been conducted to study the utility of distrust in recommender systems, although Victor et al. [31, 30] have shown that distrust is indeed helpful in trust-aware recommender systems.

All the approaches mentioned above simply treat trust as a single-aspect term and adopt the explicit trust or distrust values without further adjustments. This simplification may work well when trust values can correctly refer to the trustworthiness of users. However, the exact fine-grained values of trust and distrust are often unavailable due to various concerns such as privacy issues. The most common form is simply the social links among users. In this case, the utility of trust and distrust may not be well exploited. Inaccurate or incomplete trust networks may further decline the performance of trust-aware recommender systems [29]. Therefore, we claim that it is important to infer and hence refine trust and distrust links for better recommendation performance.

Very few approaches for recommender systems have been proposed to capture the heterogenous property of (dis)trust. For example, Kwon et al. [12] adopt the source credibility theory to select credible neighbors by investigating multiple credibility attributes. The concept of "credibility" is essentially distinct from that of "trust" defined in our paper. Specifically, the former concept refers to the reliability of users' ratings for a given item, i.e. the reliability of the recommender. The attributes considered for selecting credible recommenders are mainly expertise, trustworthiness, similarity and attraction. However, the latter focuses on a better trust network which is most suitable for recommender systems. We only consider choosing trustworthy recommenders based on a set of (dis)trust antecedents. We intend to empirically reveal the correlations of each aspect with the trust relationship, and target better predictions of trust and distrust for recommender systems.

3. The (Dis)Trust Framework

In this section, we introduce the formal definitions of the interpersonal and impersonal aspects of trust and distrust from which they will be computationally modelled according to users' historic ratings and trust networks.

Trust in social science has been well recognized as a multi-faceted concept that consists of three major parts, namely dispositional trust, institution/structural-based trust, and interpersonal trust [21]. Dispositional trust, also known as a trustor's trust propensity, refers to the trustor's inherent propensity to trust other users. Mathematically, it could be treated as a continuous constant (in the range of $[0, 1]$) subject to each trustor. Institution/structual-based trust refers to a belief held by a trustor about impersonal things of a trustee such as environments and situation. Hence, in our framework, as all users are in the same environments, we differentiate this part of the trustee by regarding it as trustor's public view of the trustee's trustworthiness. This is mainly determined by impersonal aspects of the trustee such as her reputation and position in a trust network. The impersonal aspects also have an impact on trustor's perception and hence the trust evaluation [20]. Interpersonal trust mainly involves *benevolence*, *integrity*, *competence*, and *predictability*.

With respect to the original trust model in [19, 21], we make minor modification towards the connections between the aspects and trust as shown in Figure 1. Specifically, we regard the combination of each aspect of a trustee and the propensity of a trustor as *an aspect of the trustee perceived by the trustor*, or a *trusting belief* of the trustor that the trustee has the corresponding characteristic in her favor. Therefore, trust in our model is connected with four different trust beliefs (interpersonal aspects), each of which is regarded as a trust aspect of a trustee perceived by a trustor. Together with the trustor's trust propensity and impersonal aspects of the trustee, these aspects are known as the

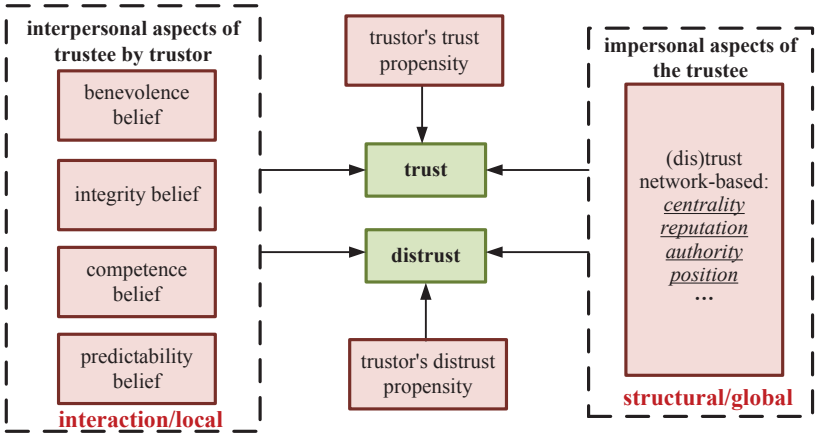

Figure 1: The proposed (dis)trust framework

antecedents of trust [19, 21], and elaborated as follows.

- **Benevolence** refers to the extent to which a trustee cares about the preferences of a trustor [21], i.e., the willingness of the trustee to do good deed for the trustor. For the user with whom the trustor has a high benevolence belief, her preferences are more likely to be similar with those of the trustor. In our case, it means that both users report similar ratings on many items.

- **Integrity** refers to the extent to which a trustee conforms to a norm or code of moral or artistic values [19]. It stresses the characteristic of the trustee to follow the norm or rules of an organization, and to have a core set of values to guide behaviors. To put it simply, the trustor believes that the trustee will always keep good-faith agreements, tell the truths, act ethically and fulfill the promises [21]. In contrast to benevolence, integrity is more concerned with the characteristic of the trustee than the trust relationship [21].

  The aspects of benevolence and integrity are somehow complementary to each other in evaluating the trustworthiness of a specific trustee. Specifically, although benevolence shows the honesty or willingness of a trustee towards a trustor, it may

fail to work in some scenarios where only limited interactions between the two users exist. This issue can be partially addressed by the integrity via considering the experience of all the users. Similarly, in the cases where integrity tends to be misleading, e.g., when most users are malicious, benevolence can help cope with this issue by relying more on personal experience between the two users.

- **Competence** refers to the ability or the power of a trustee to conduct the actions that are expected by a trustor in a specific domain [19]. Hence, competence is domain (context)-specific. For example, a user providing satisfying recommendations of purchasing cars may not be an expert of buying clothes. In other words, the user receiving a high competence belief from the trustor is capable of providing satisfactory recommendations to the trustor in a specific context. The more experience the trustee has in the specific context, the more competent she will be in the view of the trustor.

- **Predictability** refers to the consistency of a trustee's actions (good or bad, negative or positive) such that the trustor can make a prediction in a given situation [21]. Different from integrity, the value of predictability is neutral. Specifically, users' high predictability could mean that they always provide relatively high or low recommendations in need of the trustor, or consistently meet the trustor's preferences. Predictability is able to alleviate the problem of behaviors changing strategically, that is, a user may first act honestly but conduct dishonest behaviors later.

- **Impersonal aspects** represent different situations a trustor may encounter when interacting with a trustee. In our framework, they summarize the aspects of a trustee from the public view, which are independent of the interpersonal relationship between trustor and trustee. The representative information includes trustee's

reputation, position in the trustor network, degree centrality [25], authority, and even their profile information, etc.

As mentioned above, distrust is recognized as a distinct construct and opposed to trust. Trust and distrust may exist simultaneously between a trustor and a trustee. Distrust is also a multi-faceted concept, and is formalized as the mirror image of the trust concept [20]. Similarly, we connect the distrust with the aspects identified in the framework, which is illustrated in Figure 1.

4. Trust and Distrust Prediction

In this section, we firstly formulate the (dis)trust aspects based on users' historical experience. Then, we present two logistical regression models by accommodating these aspects to predict continuous values of users's trust and distrust, respectively. Finally, we further refine the trust links given the predicted trust and distrust values.

4.1. Formulations of Aspects

Given the formal definitions, we proceed to formulate the four aspects in the light of users' historical experience (i.e., ratings). For clarity, we first introduce a number of notations. Suppose there are two users: a trustor a and a trustee b, and each user has a set of experience denoted by E_a and E_b, respectively. A piece of experience is denoted by a 5-tuple $e_u = (u, j, r_{u,j}, t, c)$, indicating that a user u rated item j with a rating $r_{u,j}$ at time t under context c. Hence, users' experience can be represented as $E_a = \{e_{a1}, \ldots, e_{am}\}$ and $E_b = \{e_{b1}, \ldots, e_{bn}\}$, where m and n are the number of experience of users a and b, respectively.

Based on user experience, we then model the four general trust aspects (i.e. beliefs, see Figure 1) of trustee b from the viewpoint of trustor a, as well as the trust value that

trustor a has towards trustee b. Note that belief could be modelled by evidence [27]. Following the definitions described in Section 3, we model the four aspects as follows.

- Benevolence, $Be(a, b)$. As benevolence refers to the closeness of shared experience between two users a and b, it is modelled as the user similarity which is usually used in collaborative filtering and computed by the Pearson correlation coefficient [1]:

$$Be(a,b) = \frac{\sum\limits_{j \in E_{a,b}} (r_{a,j} - \bar{r}_a)(r_{b,j} - \bar{r}_b)}{\sqrt{\sum\limits_{j \in E_{a,b}} (r_{a,j} - \bar{r}_a)^2} \sqrt{\sum\limits_{j \in E_{a,b}} (r_{b,j} - \bar{r}_b)^2}}, \tag{1}$$

  where $E_{a,b} = E_a \cap E_b$ is the set of shared experience on the commonly rated items between users a and b, and $\bar{r}_a, \bar{r}_b$ are the average of the ratings reported by users a and b, respectively. Alternative similarity measures such as cosine similarity [1] could also be applied.[1]

- Integrity, $In(b)$. As aforementioned, integrity is independent of the trustor-trustee relationship, hence it is formulated merely based on the past experience of the trustee regardless of the trustor's actions and evaluation. Specifically, the behaviors of the majority are treated as the norm or the code when evaluating the integrity of the trustee, i.e., the similarity between the trustee's behaviors and the majority's. Hence, integrity is computed by the similarity between the preferences of the trustee and the average:

$$In(b) = \frac{\sum\limits_{j \in E_b} (r_{b,j} - \bar{r}_b)(\bar{r}_j - \bar{r})}{\sqrt{\sum\limits_{j \in E_b} (r_{b,j} - \bar{r}_b)^2} \sqrt{\sum\limits_{j \in E_b} (\bar{r}_j - \bar{r})^2}}, \tag{2}$$

[1] Note that this also holds for the computation of the integrity (see Equation 2).

where $\bar{r}_j$ refers to the average of the ratings on item $j \in E_b$, and $\bar{r}$ is the average of the ratings on all items.

- Competence, $Co(a, b, c)$. The competence of the trustee b is described from the viewpoint of the trustor a under a specific context c. Two factors are taken into account, i.e., the number of user b's experience under context c (see Equation 4), and the ratio of correct recommendations given by user b to all the other users in the system (see Equation 3), employing the basic idea of O'Donovan and Smyth [24]. The competence is computed by integrating both factors:

$$Co(a, b, c) = \gamma \frac{\sum\limits_{j \in E_b} \underset{u \in U_j}{count}(|r_{b,j} - r_{u,j}| < \varepsilon)}{\sum\limits_{j \in E_b} \parallel U_j \parallel}, \tag{3}$$

where U_j represents the set of users who have a piece of experience about item j, and ε is a predefined error tolerance threshold below which a rating $r_{b,j}$ of the trustee b is treated as a correct recommendation for item j relative to the other's real preference $r_{u,j}$. And γ is defined by:

$$\gamma = \begin{cases} \frac{N_{b,c}}{N_c^a} & \text{if } N_{b,c} \leq N_c^a; \\ 1 & \text{otherwise;} \end{cases} \tag{4}$$

where $N_{b,c}$ is the number of experience under context c out of the total m experience that user b has, and N_c^a is the minimal number of experience under context c required by the trustor a such that a user can be regarded as a reliable recommender.

- Predictability, $Pr(a, b)$. Different from integrity, the predictability of trustee b is defined as the degree to which the (positive, neutral or negative) trend of b's rating

behaviors is distinct from that of trustor a. Formally, it is computed by:

$$\begin{aligned} n_u =& \underset{j \in E_{a,b}}{\text{count}}(|r_{a,j} - r_{b,j}| \leq \theta); \\ n_n =& \underset{j \in E_{a,b}}{\text{count}}(r_{a,j} - r_{b,j} > \theta); \\ n_p =& \underset{j \in E_{a,b}}{\text{count}}(r_{a,j} - r_{b,j} < -\theta); \\ Pr(a,b) =& \frac{max(n_u, n_p, n_n) - min(n_u, n_p, n_n)}{\parallel E_{a,b} \parallel}, \end{aligned} \tag{5}$$

where n_u, n_n and n_p refer to the neutral, negative and positive trends of user b's rating behaviors comparing to trustor a's behaviors, respectively; θ is a threshold predefined by trustor a. The intuition is that for a user who is highly predictable, the difference in trends should be significant. In case of $n_u = n_n = n_p$, we obtain the lowest predictability since it is difficult to predict the next behavior of the trustee.

- Impersonal aspects: due to the availability of (dis)trust links of each user, we specifically identify four kinds of impersonal aspects in our computational model on the basis of the degree of a trustee in the trust network. The degree, as one of the centrality measurements, essentially records the aggregate public relations of the trustee in the network. The four aspects based on degree of trustee b are trust indegree $d^+_{in}(b)$, trust outdegree $d^+_{out}(b)$, distrust indegree $d^-_{in}(b)$ and distrust outdegree $d^-_{out}(b)$, referring to trustee b's *incoming trust links*, *outgoing trust links*, *incoming distrust links* and *outgoing distrust links* respectively.

4.2. Trust Prediction

For trust prediction, we define $t_{a,b,c} \in [0, 1]$ as the trust value that trustor a has towards trustee b under context c, where 0 means completely not trust and 1 completely

trust. The trust value will be influenced by the set of eight aspects that we investigated, denoted by $A(a,b)$ =$\{Be(a,b),\ Co(a,b,c),\ In(b),\ Pr(a,b),\ d^{+}_{in}(b),\ d^{+}_{out}(b),\ d^{-}_{in}(b),\ d^{-}_{out}(b)\}$. In practice, users may specify other users as trusted neighbors ($t = 1$)[2], whereas if trustor a has no direct trust link to trustee b, we consider that a has no trust towards b ($t = 0$). The trust and absence of trust connections will help build a useful model of the trust aspects and the overall trust. Specifically, the expected probability[3] that trustor a completely trusts the trustee b under context c (denoted as $p^{+}(a,b,c)$) can be written as:

$$p^{+}(a,b,c) = E\big(t_{a,b,c} = 1 | A(a,b)\big). \tag{6}$$

We apply the logistic regression to classify trust from not trust, and obtain the importance weight of each aspect related with trust. To be specific, the logit of the probability is modelled as a linear combination of $A(a,b)$ [15]:

$$logit\big(p^{+}(a,b,c)\big) = log\left(\frac{p^{+}(a,b,c)}{1-p^{+}(a,b,c)}\right) = \alpha_0^{a+} + (\alpha_A^{a+})^T \cdot A(a,b), \tag{7}$$

where $\alpha_A^{a+} = \{\alpha_1^{a+}, \alpha_2^{a+}, \alpha_3^{a+}, \alpha_4^{a+}, \alpha_5^{a+}, \alpha_6^{a+}, \alpha_7^{a+}, \alpha_8^{a+}\}$, and α_0^{a+} is interpreted as the intrinsic trust propensity of trustor a. Then the probability $p^{+}(a,b,c)$ is derived by:

$$p^{+}(a,b,c) = \frac{1}{1+e^{-\left(\alpha_0^{a+} + (\alpha_A^{a+})^T \cdot A(a,b)\right)}}. \tag{8}$$

Based on the trust information directly specified by real users, we are able to train this model and learn the coefficients, i.e., the importance weight of each aspect related to trust. The weights α_A^{a+} can be used to compute implicit or refine explicit trust values

[2]Note that in some cases, user a might specify a real trust value ranged in [0,1] towards user b under context c. In this scenario, we treat $t_{a,b,c} = 1$ if the real value is bigger than 0, otherwise $t_{a,b,c} = 0$. We follow the same consideration for the distrust information.

[3]The probability is thus treated as the trust value.

from user experience.

4.3. Distrust Prediction

Accordingly, following the process of trust prediction, the expected probability that the trustor a completely distrusts the trustee b under context c (denoted as $p^-(a,b,c)$) can be written as:

$$p^-(a,b,c) = E\big(d_{a,b,c} = 1 | A(a,b)\big), \tag{9}$$

where $d_{a,b,c} = 1$ represents that a completely distrusts b under context c. We also apply the logistic regression to classify distrust from not distrust, and obtain the importance weight of each aspect related with distrust:

$$logit\big(p^-(a,b,c)\big) = log\left(\frac{p^-(a,b,c)}{1-p^-(a,b,c)}\right) = \alpha_0^{a-} + (\alpha_A^{a-})^T \cdot A(a,b), \tag{10}$$

where $\alpha_A^{a-} = \{\alpha_1^{a-}, \alpha_2^{a-}, \alpha_3^{a-}, \alpha_4^{a-}, \alpha_5^{a-}, \alpha_6^{a-}, \alpha_7^{a-}, \alpha_8^{a-}\}$, and α_0^{a-} is interpreted as the intrinsic distrust propensity of the trustor a. Then the probability $p^-(a,b,c)$ is derived by:

$$p^-(a,b,c) = \frac{1}{1 + e^{-\left(\alpha_0^{a-} + (\alpha_A^{a-})^T \cdot A(a,b)\right)}}. \tag{11}$$

Based on the distrust information directly specified by real users, we are able to train this model and learn the importance weight of each aspect related to distrust. The weights can be used to compute implicit or refine explicit distrust values from user experience.

4.4. Trust Link Refinement

Given the predicted probability of complete trust $p^+(a,b,c)$ and distrust $p^-(a,b,c)$ according to Equations 8 and 11, we can further refine the trust link by filtering out the

possibly inaccurate trust or distrust link using the following rules:

$$\begin{aligned} &\text{if } p^+(a,b,c) > p^-(a,b,c), \text{ trust link from } a \text{ to } b; \\ &\text{if } p^+(a,b,c) < p^-(a,b,c), \text{ distrust link from } a \text{ to } b; \\ &\text{if } p^+(a,b,c) = p^-(a,b,c), \text{ no link from } a \text{ to } b. \end{aligned} \tag{12}$$

Furthermore, we could also refine the trust degree using Equation 13 for other specific purposes such as comparing the trust degrees between different user pairs.

$$t(a,b,c) = \begin{cases} p^+(a,b,c) - p^-(a,b,c) & \text{if } p^+(a,b,c) > p^-(a,b,c) \\ 0 & \text{otherwise.} \end{cases} \tag{13}$$

5. Evaluation

For evaluation, we aim to explore the effectiveness of our proposed (dis)trust framework by incorporating the generated trust information into three representative trust-aware recommender systems.

5.1. Data Sets

Three real-world data sets are used in the experiments, namely Epinions, FilmTrust and Flixster. Epinions enables users to review products by adding text comments and issuing numerical ratings in the range of $[1, 5]$. Besides, users can also explicitly specify other users as trust (to the trust list) or distrust (to the block list) based on whether the reviews and ratings of others are consistently valuable or useless for the user. We adopt the extended Epinions data set[4] where trust value is labeled as 1 and distrust as -1. We sample two subsets by randomly selecting $5,000$ and $10,000$ users, respectively. The

[4] http://www.trustlet.org/wiki/Epinions_datasets

other two data sets are FilmTrust (provided by Guo et al. [6]) and Flixster[5] where only trust exists and no distrust information is available. Users can only indicate others as trust, and provide item ratings scaled from 0.5 to 4.0 (5.0 in Flixster) with step 0.5. The statistics of the four data sets is presented in Table 1.

Table 1: The statistics of four data sets

Features	Epinions1	Epinions2	FilmTrust	Flixster
users	5,000	10,000	1,508	5,000
items	376,458	519,491	2,071	13,527
trust	744	3,443	2,853	2,898
distrust	424	1,398	n.a.	n.a.
ratings	968,467	2,017,158	70,998	264,540
avg rating	4.6964	4.6863	3.0028	3.6560

5.2. Experimental Settings

Since the two Epinions subsets are the only available collections that contain both trust and distrust information, we use them to train two logistic regression models for trust and distrust respectively. Specifically, the users who specify both trust and distrust statements to others are selected as the training data in order to learn the coefficients (i.e., the importance weights) of each trust and distrust aspect according to Equation 7 and Equation 10, respectively. Due to the limitation of data, we do not take into account the context information in the experiments. Besides, we empirically set $\varepsilon = 0.1$ for competence (see Equation 3) and $\theta = 0.1$ for predictability (see Equation 5) computations. Although the other data sets FilmTrust and Flixster do not contain distrust information (and hence cannot train a regression model independently), they may be useful in testing the effectiveness of these aspects by adopting the models learned from the Epinions data sets. The intuition is that although the exact or absolute coefficient values may vary in different data sets, the relative importance weights may follow the

[5] http://www.cs.sfu.ca/~sja25/personal/datasets/

same trends for key factors. In other words, as the coefficients learned from one data set A reflect the importance weights of the corresponding (dis)trust aspects related to (dis)trust, they capture the dependent relationships between these (dis)trust aspects with (dis)trust for the users in the data set sample A. In this case, if we assume that users in another data set B is sampled from the same user population as those in the data set A, the coefficients for these users in the data set B might have similar values as those in data set A. Under this assumption, to be specific, we apply the coefficients learned from Epinions1 to FilmTrust, and those learned from Epinions2 to Flixster according to the comparative sizes of the corresponding data sets.

After obtaining the aspect coefficients, we regenerate or predict the trust values in the light of different combinations of the two types of trust and distrust aspects, and in total we obtain 3 such different combinations and the corresponding trust values. Hence, the effectiveness of the new trust information (refined by the predicted distrust information) can be investigated by the recommendation performance in comparison with the original ones. Specifically, to demonstrate the effectiveness, we adopt three representative trust-aware algorithms to generate recommendations:

- **TidalTrust**, proposed by Golbeck [4], uses trust values to substitute user similarity to weigh user ratings when generating recommendations.

- **Merge**, proposed by Guo et al. [5], incorporates the ratings of trusted neighbors to form a more complete rating profile for active users, where the trust propagation length is 1.

- **SocialMF**, proposed by Jamali and Ester [9], considers the trust information and propagation of trust information into the matrix factorization model for recommender systems. In our experiments, we adopt the same settings of parameters as suggested in [9], and source code provided by MyMediaLite recommender system

library[6].

To have a better understanding of the effectiveness, we split each data set into three different views in terms of item-related properties as used in [5, 18]:

- **All** represents the whole data set.
- **Controversial Items** are those items which received ratings with standard deviation greater than 1.5.
- **Niche Items** are those items which received less than 5 ratings.

The experiments are conducted by applying the *leave-one-out* technique, that is, each rating is iteratively hidden whose value will be predicted by applying the TidalTrust, Merge, or SocialMF method until all ratings in the data sets are tested. The performance is evaluated by two commonly used measures: the root mean square errors (RMSE) and mean absolute errors (MAE). They both refer to the differences between the predictions and the ground truth, but differ from each other as indicated by their names. Generally, smaller RMSE and MAE values indicate better predictive accuracy.

6. Results and Analysis

The experimental results are presented in two-fold: (1) the importance weights of the trust and distrust aspects learned from logistic regression models; and (2) the effectiveness of the trust and distrust aspects applied in recommender systems in comparison with that of the original trust values.

Table 2: The coefficients of trust aspects

Data set / Aspect	Epinions1		Epinions2	
	trust	distrust	trust	distrust
benevolence	0.772	-1.2295	0.6332	-1.4537
competence	2.3706	0.988	2.5458	1.6816
integrity	-0.5816	-0.1122	-0.6597	-0.461
predictability	-0.0471	0.313	-0.3666	0.5724
trust indegree	-0.055	0.0159	-0.0387	-0.0006
trust outdegree	0.0615	-0.0042	0.0677	0.011
distrust indegree	-0.0765	-0.3697	0.0016	-0.2533
distrust outdegree	-0.0125	0.2347	-0.0066	0.144

6.1. Importance of Trust and Distrust Aspects

We use the L2-regularized logistic regression provided by LIBLINEAR[7] to train the data of Epinions1 and Epinions2. The coefficients (i.e., the importance weights) of the trust and distrust aspects are illustrated in Table 2. Note that since the implementation of LIBLINEAR tends to minimize the bias part (to 0) during the model fitting process, we do not present the results of the aspect about trustor's propensity. In fact, its value is often equal to or very close to 0. Besides, since the logistic regression has a strict requirement on the sample size, we adopt a well-known rule of thumb, i.e., the 1 **in** 10 **rule** [7] to specify a minimum size of the sample for a reliable training. In particular, a minimum number 10 of trust (or distrust) links are required for each aspect, that is, at least 80 trust (or distrust) examples are required in order to obtain a reliable model. However, we find that only few users in the training sets could meet the requirement. Therefore, we turn out to train the logistic regression models based on the trust and distrust networks of all the users and adopt the learned coefficients for all the users. An alternative way is to divide users into different clusters according to user similarity and then the coefficients could be learned using all the users' experience within the same

[6] http://www.mymedialite.net
[7] http://www.csie.ntu.edu.tw/~cjlin/liblinear/

cluster. This may lead to more accurate coefficients for similar users. Although we do not conduct our experiments in this way in the current work, we demonstrate that our method based on the logistic regression models learned from all users (i.e. general knowledge) could already significantly improve the recommendation accuracy.

Table 2 shows that consistent results for the four interpersonal aspects[8] with trust and distrust are obtained in both Epinions1 and Epinions2 data sets. In general, benevolence and competence are both positively correlated with trust whereas integrity and predictability are negatively correlated. In other words, the first two aspects are more likely to increase the probability of trust, but the latter two decrease the probability. More specifically, competence shows the greatest correlation with trust, followed by benevolence. This may imply that users in recommender systems are more concerned with personal experience (e.g. benevolence) rather than collective opinions (e.g. integrity) when establishing trust. Further, a person whose behaviors are highly predictable does not guarantee high trustworthiness in trust building because the predictability is value-neutral. In contrast, competence and predictability present positive correlation with distrust whereas benevolence and integrity are negatively correlated with distrust. It should also be noted that the result for each individual impersonal aspect is not very consistent across the two data sets. This might be due to the fact that we only capture partial trust and distrust information for users in our data sets, as we only consider the trust and distrust information of each user to our sampled users. Overall, however, the coefficients for impersonal aspects could still be considered as consistent in the sense that the aggregated effect of the trust network related impersonal aspects (trust indegree and outdegree) is positive, while that of the distrust network related impersonal aspects is negative, for both trust and distrust. This could be partially explained as that a trustee with more

[8] Their values are in the range of [0,1], while the values for the four impersonal factors are integers (≥ 0).

trusted and trusting neighbors could have more far-reaching influence on other users, and thus are intended to be either more trusted or distrusted by others. In other words, a trustworthy user would be considered as more trustworthy by trustors, and further trusted by more people, and vice versa.

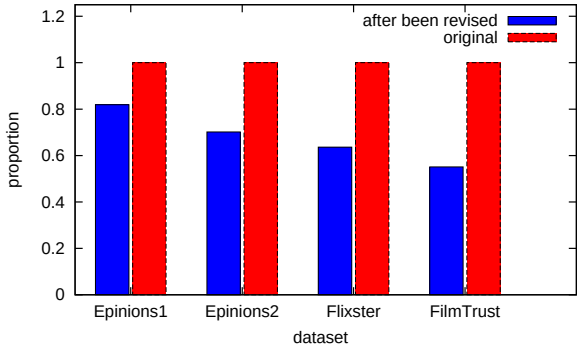

Figure 2: The comparison of refined trust links with the original ones

6.2. Effectiveness of the Proposed Model

Table 3: The comparison of performance based on Refined Trust using Epinions1

Methods	Aspects	**All**		**Controversial Items** 2,777 users 7,242 ratings		**Niche Items** 4,705 users 539,881 ratings		FilmTrust-All	
		RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE
TidalTrust	Original	**0.6759**	0.5756	1.7259	1.6045	**0.7287**	**0.6368**	0.9687	0.7564
	All	0.7432	0.6255	1.5910	**1.4226**	0.7653	0.6713	0.8355	0.6465
	Interpersonal	0.6833	**0.5675**	1.6996	1.5169	0.7674	0.6548	**0.7929**	**0.6177**
	Impersonal	0.7624	0.6469	**1.5873**	1.4574	0.7710	0.6787	0.9390	0.7315
Merge	Original	0.7441	0.5920	1.5490	1.3336	**0.7601**	**0.6103**	0.8788	0.6919
	All	0.7608	0.6140	1.5295	1.3490	0.7752	0.6324	0.8751	0.6892
	Interpersonal	**0.7234**	**0.5734**	1.5224	1.3270	0.7791	0.6384	**0.8748**	**0.6890**
	Impersonal	0.7890	0.6336	**1.4930**	**1.3172**	0.7811	0.6396	0.8766	0.6904
SocialMF	Original	1.4075	1.2177	–	–	–	–	1.0608	0.7760
	All	**1.3910**	**1.1820**	–	–	–	–	**0.9950**	**0.7310**
	Interpersonal	1.4455	1.2414	–	–	–	–	1.0639	0.7684
	Impersonal	1.6103	1.3370	–	–	–	–	1.081	0.7917

We predict the trust values based on the learned regression models for three scenarios: “All”, “Interpersonal” and “Impersonal”. “All” refers to considering both interpersonal and impersonal aspects, while the others refer to only considering interpersonal or impersonal aspects respectively. The effectiveness of these aspects in predicting

Table 4: The comparison of performance based on different trust aspects using Epinions2

Methods	Aspects	**All**		**Controversial Items** 2,653 users 12,775 ratings		**Niche Items** 8,922 users 731,116 ratings		Flixster-All	
		RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE
TidalTrust	Original	0.6420	0.5287	1.6467	1.5444	**0.7346**	**0.6399**	1.2449	0.9846
	All	0.7123	0.5857	1.7390	1.5215	0.7867	0.6764	**1.2129**	**0.9706**
	Interpersonal	**0.6255**	**0.5132**	1.6607	**1.4679**	0.7814	0.6627	1.2190	0.9789
	Impersonal	0.7194	0.5951	**1.6055**	1.4950	0.7766	0.6746	1.2449	0.9846
Merge	Original	0.6801	**0.5540**	1.4869	1.3235	**0.7298**	**0.5934**	1.0376	0.8163
	All	0.7032	0.5880	1.5467	1.3754	0.7432	0.6128	**1.0362**	**0.8150**
	Interpersonal	**0.6734**	0.5541	1.4868	**1.2830**	0.7503	0.6211	1.0366	0.8155
	Impersonal	0.7341	0.6043	**1.4669**	1.3235	0.7501	0.6143	1.0376	0.8163
SocialMF	Original	1.2799	1.1094	–	–	–	–	1.3747	**1.0440**
	All	**1.2559**	**1.0971**	–	–	–	–	**1.3716**	1.0450
	Interpersonal	1.2603	1.0999	–	–	–	–	1.3838	1.0506
	Impersonal	1.4474	1.2079	–	–	–	–	1.3747	1.0440

trust values would be investigated by applying the aforementioned three algorithms (i.e. TidalTrust, Merge1 and SocialMF) in terms of predictive accuracy for recommender systems. Besides, three different views[9] mentioned in Section 5.1 of data sets are studied. Lastly, we further employ the learned regression models from Epinions1 and Epinions2 to FilmTrust and Flixster where distrust information is unavailable.

Before evaluating our performance, we first present the ratio of "reliable" trust links to the original ones according to Equation 12 based on our model. As illustrated in Figure 2, a substantial ratio of the original trust links are filtered out as "unreliable" ones by our model. The results in Tables 3 and 4 of our method are based on these "reliable" trust links. Later we will show whether this difference would lead to the performance improvement of the three recommendation algorithms.

In the view of *All*, Tables 3 and 4 show that our model could almost achieve the best performance with regard to RMSE and MAE for all three algorithms on the three data sets. Our method could achieve similar results with the original trust values on Epinions and Flixster, but demonstrate significant differences on FilmTrust (the t-test verifies its

[9]In this work, we did not investigate views of controversial items and niche items for the SocialMF algorithm due to the inconvenience of implementation.

statistical significance at the 5% level, i.e. p-value$= 0.0415 < 0.05$). This may be explained by the fact that most ratings on Epinions and Flixster data sets are highly skewed. Specifically, the average ratings are 4.6964, 4.6863 (out of 5) and 3.6560 (out of 4) in Epinions1, Epinions2 and Flixster, respectively (see Table 1). In contrast, the average rating in FilmTrust is 3.0028 out of 4. The same trends could be observed in the view of *Controversial Items* due to less skewed distributed ratings, where our approach obtains much better performance than that with original trust value. It should be noted that in the views of *Niche Items*, the performance of our method is worse than that with the original trust. This is mainly because niche items are defined as those which received less than 5 ratings. In that case, the problem of data sparsity becomes more serious as we filter out some recommenders who might provide ratings to niche items. This problem could be addressed by predicting more implicit trust links with our model. The improvements on the three methods over those with the original trust are remarkable (around 0.13 in RMSE and 0.18 in MAE at most), as Koren [11] points out that small improvements in RMSE may lead to significant improvements in real applications.

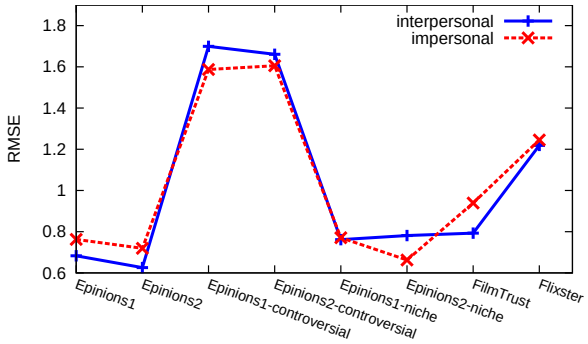

Figure 3: Performance comparison of TidalTrust method by considering interpersonal and impersonal aspects

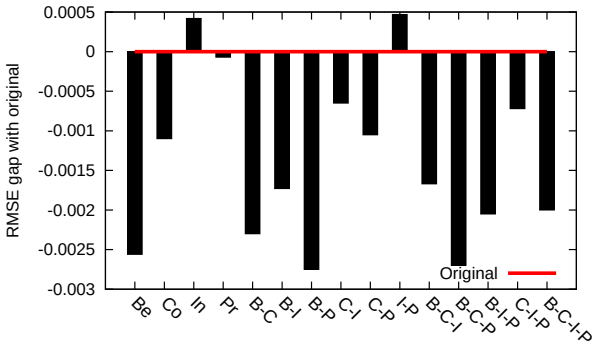

Figure 4: Performance comparison of TidalTrust method in All View On Epinions1

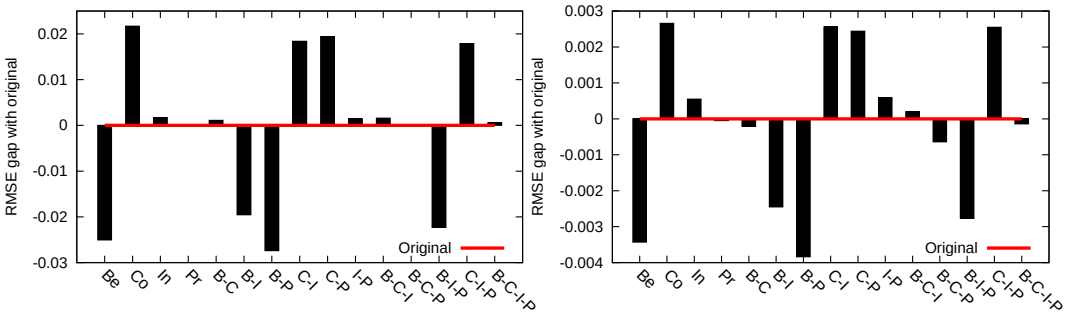

(a) Controversial items **(b) Niche items**

Figure 5: The performance comparison of TidalTrust method in different views on Epinions1

6.2.1. Interpersonal and Impersonal Aspects

Figure 3 presents the performance of TidalTrust algorithm by considering only interpersonal or impersonal aspects. As demonstrated in Figure 3, we can see that the trust derived from interpersonal aspects (i.e. on rating history) is more effective than that from impersonal factors (i.e. on trust and distrust network) in terms of RMSE. However, for controversial items, the algorithms depending on trust values derived from impersonal aspects perform saliently better than those from interpersonal aspects. This is due to the fact that users' ratings of controversial items are quite dissimilar, increasing the difficulty on extracting valuable information for personalized recommendation according to rating history. On the contrary, the impersonal aspects, modelled based on the trust and distrust networks, would not be affected by those controversial ratings. Hence, they might

infer more reliable trust and distrust values. Besides, we also explore the effectiveness of each interpersonal aspect as well as their combinations without considering the impersonal aspects. The results are presented in Table 5 and Figures 4 and 5, where B, C, I and P denote benevolence, competence, integrity and predictability, respectively. Hence all the combinations of trust aspects can be represented by concatenating letters. For example, B-C refers to the combination of the benevolence and competence. As can be seen in Table 5[10], overall, the performance increases as more aspects are involved in. We thus could conclude that all the four interpersonal aspects are reasonable and each of them contributes to the success of our trust and distrust prediction.

Table 5: The performance comparison of Tidaltrust based on interpersonal aspects on Flixster

# Aspects	RMSE	Improvement	MAE	Improvement
1	1.2415±0.0031	-	0.9828±0.0036	-
2	1.2383±0.0055	0.26%	0.9805±0.0043	0.23%
3	1.2352±0.0028	0.25%	0.9783±0.0034	0.22%
4	1.2129	1.81%	0.9706	0.78%

A clearer and more detailed demonstration is illustrated in Figures 4 and 5 which present the comparison of different interpersonal aspects in terms of performance gaps in different views of data sets. The histogram under the horizontal solid line (representing the original trust performance) means a better performance than the baseline in terms of RMSE. More specifically, for single aspect, benevolence achieves the best performance than the other three aspects. In contrast, in the view of *Controversial* or *Niche* items (see Figure 5), competence obtains the worse performance than predictability or integrity whose performance is equivalent to that of original trust values. Besides, in the view of *All* (see Figure 4), the performance gap between competence and integrity or predictability is not so significant since all the RMSE gaps are smaller than 0.005. Hence, although competence is an important aspect for trust modelling (see Table 2), it is not that useful

[10]The improvements are computed as $\frac{\text{RMSE (or MAE) value of n-1 aspects - that of n aspects}}{\text{RMSE (or MAE) value of n-1 aspects}}$, $n = 2, 3, 4$.

in recommender systems as a single aspect. Furthermore, the performance of $B - P$ is better than that of B (the best of single aspect) and that of $B - I - P$ or $B - C - P$ (the best of the combinations of three aspects). This implies that predictability, modelled in a different way and providing additional information, can complement benevolence in building the trust relationship in recommender systems. However, integrating with other aspects (e.g., competence or integrity) may not result in better performance. For the best combination $B - P$, benevolence is closely related to individuals' similarity, and predictability, on the contrary, provides indications of the consistency of the similarity trend. In this sense, the two aspects are complementary to each other, and capable of generating better trust values for recommender systems. However, when other aspects are incorporated, redundant and even noisy information could be brought in, and thus deteriorates the performance.

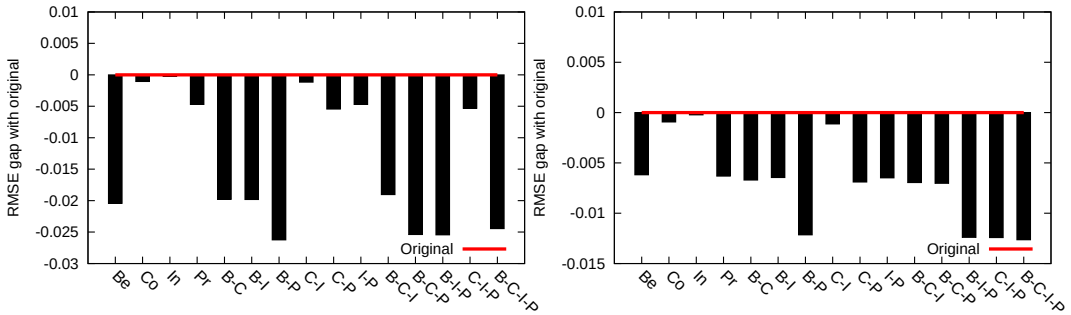

(a) FilmTrust **(b) Flixster**

Figure 6: (a) Performance comparison of TidalTrust method in All view.

6.2.2. Generalization

It is observed that similar trends of performance are obtained on FilmTrust and Flixster using the coefficients learned from Epinions1 and Epinions2, respectively. As illustrated in Tables 3 and 4, TidalTrust, Merge and SocialMF could achieve better performance with the trust information learned by using the trained logistic regression mod-

els (for trust and distrust) on Epinions. Moreover, as shown in Figures 6(a) and 6(b), benevolence consistently shows better performance than other single aspects, and the combination of benevolence and predictability reaches the best performance among the overall 15 combinations of impersonal aspects. Hence, we conclude that the trust model learned from one data set can be applied to other data sets where distrust information is unavailable. It is important because most real-world data sets do not contain such information due to various reasons such as privacy concern. In other words, the knowledge learned from one community can be (partially) reused to model the trust and distrust in other communities.

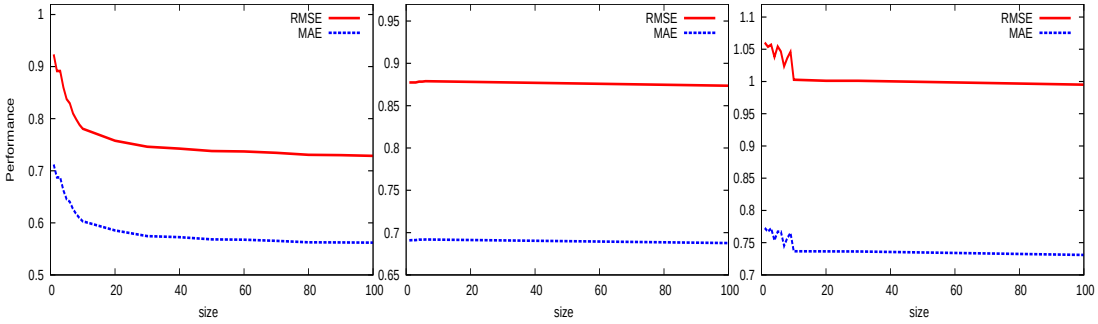

(a) TidalTrust **(b) Merge** **(c) SocialMF**

Figure 7: The performance by varying the size of predicted trust network on FilmTrust

6.2.3. The Size of the Predicted Trust Network.

Figure 7 shows the performance of the three algorithms on FilmTrust data set by varying the size of the predicted trust network. Here, the X-axis refers to that the predicted trust network is certain times as large as the original network. As demonstrated in the figure, we can see that all three algorithms obtain better recommendation accuracy as the size of the trust network increases up to a certain point, verifying the effectiveness of our model in predicting implicit trust and distrust values (or relationships). Note that both Merge and SocialMF incorporate the mechanism of trust propagation to improve the recommendation accuracy. Therefore, the corresponding performance of incorporat-

ing our model could be further improved if the propagation length is made bigger than 1 (especially for the Merge method [5]).

6.2.4. Predicted Distrust Information.

In our model, we employ the predicted distrust value to refine the trust value according to Equations 12 and 13. Figure 8 pictures the performance comparison of the TidalTrust algorithm on three data sets by differentiating between considering and not considering the predicted distrust information. As illustrated, by considering all aspects, i.e. both interpersonal and impersonal ones, we can see that the performance of TidalTrust has been saliently improved if incorporating predicted implicit distrust information into trust value prediction. This demonstrates that the noisy (less reliable) trust links could be validly removed by our model.

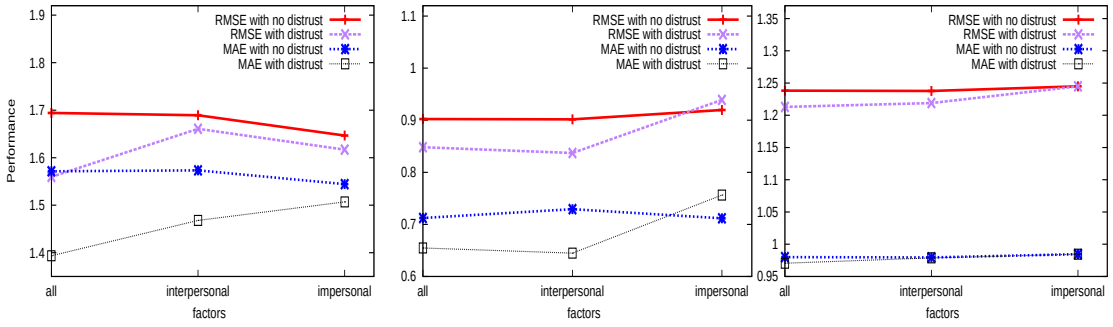

(a) Epinions-Controversial Items **(b) FilmTrust** **(c) Flixster**

Figure 8: The performance comparison of TidalTrust between "without" and "with" distrust

7. Conclusion and Future Work

This paper explored the multiple facets of trust and distrust predictions for recommender systems. Specifically, we identified both the interpersonal and impersonal aspects according to the trust theory from social science. The four interpersonal aspects, namely benevolence, competence, integrity and predictability, were formally defined in

the trust theory based on which they were computationally modelled in the light of user experience (i.e. ratings) in the systems, while the impersonal aspects are computed on the basis of users' trust and distrust network. Then the importance of each aspect to trust or distrust was learned by applying a corresponding logistic regression model trained by real-world data sets that contained trust or distrust information.

After learning the two logistic regression models, we predicted the (implicit) trust and distrust values, where the trust values were further refined by the distrust values. These newly generated trust values were taken as input to three representative trust-based recommendation algorithms (i.e TidalTrust, Merge and SocialMF) in order to validate the effectiveness of our proposed model. The experimental results showed that: (1) benevolence and competence were positively correlated with trust whereas the integrity and predictability were negatively correlated. On the other hand, competence and predictability were positively correlated with distrust whereas the benevolence and integrity were negatively correlated. All the four interpersonal aspects were useful for the existing trust-based recommendation algorithms in that each individual aspect can achieve comparable performance derived from the original trust values. The combination of benevolence and predictability can achieve the best performance among all the 15 combinations made by the four aspects; (2) the learned trust models can be applied to other communities where distrust information is not available for evaluating both the trust and distrust relationships. Our results could serve as a guidance to effectively build implicit trust or distrust networks (competitive to robust explicit networks) based on our proposed framework when users had no explicit trust or distrust information. Incorporating distrust information could effectively remove noisy and redundant data in the original explicit trust network. Therefore, when encountering a data set A without distrust information, an alternative way is to learn the coefficients from sampled Epinions data set which has comparable size with the data set A; (3) incorporating impersonal as-

pects can further improve the performance of the existing trust-based recommendation algorithms. In addition, the interpersonal aspects would take greater effect when there were lots of rating data of users, whereas the impersonal aspects would contribute more to the controversial items; and (4) our ability of predicting the implicit trust values could complement the trust network, which could further improve the performance of trust-aware recommender systems. In other words, if we want to improve the performance of a specific trust-aware recommender system, we can predict more possible trust links using our approach to increase the size of the existing trust network.

The contributions of our current work can be mainly summarized by the following two aspects: (1) our study serves as the initial step aiming to fill in the gap between trust and distrust as multi-aspect concepts and the relatively simple usage of trust and (especially) distrust in recommender systems. The newly predicted trust and distrust values can effectively enhance the performance of the existing trust-aware recommender systems. Given the relatively large base of this kind of recommender systems, the influence is considerably significant for the area of recommender systems. With the increased accuracy of recommendation, users will be able to achieve more informed decision making; and (2) we introduce the formal definitions of the interpersonal and impersonal aspects of trust and distrust from which they will be computationally modeled according to users' historic ratings and trust networks. It can inspire and lead the research in the computational trust area to build more robust and practical trust models, which well support users' decisions on which others to trust.

The future work is discussed as follows: (1) in the current work, we assume the trust aspects are independent with each other, and linearly correlated with trust. In the future, we might employ more complex machine learning techniques to capture the possible dependency among trust aspects for more powerful (dis)trust prediction algorithm; (2) the present work focuses on predicting (implicit) trust and distrust values according to

our proposed trust framework for recommender systems by comparing the performance using refined trust values (or links) relative to the original ones. In the future, we will further verify our research framework by exploring other candidate impersonal aspects such as reputation and closeness centrality of trustees, etc.; (3) this study formalizes distrust as the mirror image of the trust concept. In the future, we could consider another interesting case in which trust and distrust are not predicted by the same aspects/antecedents, but by different ones; and (4) we plan to design a trust-aware recommender system by relatively equally considering both the predicted trust and distrust values instead of using the predicted distrust values to refine trust values.

References

[1] G. Adomavicius and A. Tuzhilin. Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge and Data Engineering*, 17(6):734–749, 2005.

[2] C. Castelfranchi and R. Falcone. Socio-cognitive theory of trust. *J. Pitt. London: Wiley*, 2005.

[3] M. Chowdhury, A. Thomo, and B. Wadge. Trust-based infinitesimals for enhanced collaborative filtering. In *Proceedings of the 15th International Conference on Management of Data (COMAD)*, 2009.

[4] J. Golbeck. Generating predictive movie recommendations from trust in social networks. In *Trust Management*, pages 93–104. 2006.

[5] G. Guo, J. Zhang, and D. Thalmann. A simple but effective method to incorporate trusted neighbors in recommender systems. In *Proceeding of the 20th Conference on User Modeling, Adaptation, and Personalization (UMAP)*, pages 114–125. 2012.

[6] G. Guo, J. Zhang, and N. Yorke-Smith. A novel bayesian similarity measure for recommender systems. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI)*, 2013.

[7] F. E. Harrell, K. L. Lee, R. M. Califf, D. B. Pryor, and R. A. Rosati. Regression modelling strategies for improved prognostic prediction. *Statistics in Medicine*, 3(2):143–152, 1984.

[8] M. Jamali and M. Ester. Trustwalker: a random walk model for combining trust-based and item-based

recommendation. In *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 397–406. ACM, 2009.

[9] M. Jamali and M. Ester. A matrix factorization technique with trust propagation for recommendation in social networks. In *Proceedings of the fourth ACM Conference on Recommender Systems*, pages 135–142. ACM, 2010.

[10] A. Jøsang, W. Quattrociocchi, and D. Karabeg. Taste and trust. In *Trust Management V*, pages 312–322. Springer, 2011.

[11] Y. Koren. Factor in the neighbors: Scalable and accurate collaborative filtering. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 4(1):1–24, 2010.

[12] K. Kwon, J. Cho, and Y. Park. Multidimensional credibility model for neighbor selection in collaborative recommendation. *Expert Systems with Applications*, 36(3):7114–7122, 2009.

[13] J. Leskovec, D. Huttenlocher, and J. Kleinberg. Predicting positive and negative links in online social networks. In *Proceedings of the 19th International Conference on World Wide Web*, pages 641–650. ACM, 2010.

[14] D. Liben-Nowell and J. Kleinberg. The link-prediction problem for social networks. *Journal of the American Society for Information Science and Technology*, 58(7):1019–1031, 2007.

[15] X. Liu, A. Datta, H. Fang, and J. Zhang. Detecting imprudence of 'reliable'sellers in online auction sites. In *Proceedings of the 11th International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom)*, pages 246–253, 2012.

[16] H. Ma, M. R. Lyu, and I. King. Learning to recommend with trust and distrust relationships. In *Proceedings of the third ACM Conference on Recommender Systems*, pages 189–196. ACM, 2009.

[17] H. Ma, H. Yang, M. R. Lyu, and I. King. Sorec: social recommendation using probabilistic matrix factorization. In *Proceedings of the 17th ACM conference on Information and Knowledge Management*, pages 931–940. ACM, 2008.

[18] P. Massa and P. Avesani. Trust-aware recommender systems. In *Proceedings of the 2007 ACM Conference on Recommender Systems (Recsys)*, pages 17–24, 2007.

[19] R. C. Mayer, J. H. Davis, and F. D. Schoorman. An integrative model of organizational trust. *The Academy of Management Review*, 20(3):709–734, 1995.

[20] D. H. McKnight and N. L. Chervany. Trust and distrust definitions: One bite at a time. In *Trust in Cyber-Societies*, pages 27–54. Springer, 2001.

[21] D. H. McKnight and N. L. Chervany. What trust means in e-commerce customer relationships: An

interdisciplinary conceptual typology. *International Journal of Electronic Commerce*, 6(2):35–59, 2001.

[22] D. H. McKnight and V. Choudhury. Distrust and trust in b2c e-commerce: Do they differ? In *Proceedings of the 8th International Conference on Electronic Commerce*, pages 482–491. ACM, 2006.

[23] A. Mnih and R. Salakhutdinov. Probabilistic matrix factorization. In *Proceedings of the Advances in Neural Information Processing Systems*, pages 1257–1264, 2007.

[24] J. O'Donovan and B. Smyth. Trust in recommender systems. In *Proceedings of the 10th International Conference on Intelligent User Interfaces (IUI)*, pages 167–174, 2005.

[25] T. Opsahl, F. Agneessens, and J. Skvoretz. Node centrality in weighted networks: Generalizing degree and shortest paths. *Social Networks*, 32(3):245–251, 2010.

[26] S. Ray and A. Mahanti. Improving prediction accuracy in trust-aware recommender systems. In *Proceedings of the 43rd Hawaii International Conference on System Sciences (HICSS)*, pages 1–9, 2010.

[27] G. Shafer. *A mathematical theory of evidence*, volume 1. Princeton university press Princeton, 1976.

[28] P. Singla and M. Richardson. Yes, there is a correlation:-from social networks to personal behavior on the web. In *Proceedings of the 17th International Conference on World Wide Web (WWW)*, pages 655–664, 2008.

[29] M. Srivatsa and M. Hicks. Deanonymizing mobility traces: Using social network as a side-channel. In *Proceedings of the 2012 ACM Conference on Computer and Communications Security*, pages 628–637, 2012.

[30] P. Victor, C. Cornelis, M. De Cock, and A. Teredesai. Trust-and distrust-based recommendations for controversial reviews. *IEEE Intelligent Systems*, 26(1):48–55, 2011.

[31] P. Victor, N. Verbiest, C. Cornelis, and M. D. Cock. Enhancing the trust-based recommendation process with explicit distrust. *ACM Transactions on the Web (TWEB)*, 7(2):42–59, 2013.