



Investigating interactions of trust and interest similarity

Cai-Nicolas Ziegler^{a,*}, Jennifer Golbeck^b

^a Siemens AG, Corporate Research and Technology, Otto-Hahn-Ring 6, Geb. 31, D-81730 München, Germany

^b University of Maryland, College Park, A.V. Williams Building, College Park, MD 20740, USA

Received 6 September 2005; received in revised form 29 October 2006; accepted 3 November 2006

Abstract

Online communities that allow their users to express their personal preferences, such as the members they trust and the products they appreciate, are becoming increasingly popular. Exploiting these communities as playgrounds for sociological research we present two frameworks for analyzing the correlation between interpersonal trust and interest similarity. We obtain empirical results from applying the two frameworks on two real, operational communities, that suggest there is a strong correlation between both trust and interest similarity. We believe our findings particularly relevant for ongoing research in recommender systems and collaborative filtering, where people are suggested products based on their similarity with other customers, and propose ways in which trust models can be integrated into these systems.

© 2006 Elsevier B.V. All rights reserved.

Keywords: Trust; Similarity; Collaborative filtering; e-commerce; Distributed systems

1. Introduction

In today's networked worlds, uncertainty and anonymity are important factors that bear strong implications on decision-making. Several researchers have therefore proposed to incorporate the concept of interpersonal trust into electronic interactions [32], and proposed several computational trust models [37,43,41]. These models intend to predict which people are deemed trustworthy, based on the people one given agent trusts, the people the trusted agents trust, and so forth. Trust can be seen as an adequate decision-support tool that dissects relevant and reliable information sources from unreliable [28,26].

Recently, approaches incorporating trust models into online recommender systems have started gaining momentum [42,31,20], synthesizing recommendations based upon opinions from trusted peers rather than *most similar ones*.¹ Decentralized recommender systems cannot rely on generic collaborative filtering methods alone, because the latter have been shown to exhibit poor scalability [54,60]. Novel approaches that allow some sort of pre-filtering and scalable neighborhood formation (such as trust-based recommenders) are required. To this end, trust becomes a supplementary or even replacement filtering mechanism.

* Corresponding author.

E-mail addresses: cai.ziegler@siemens.com (C.-N. Ziegler), golbeck@cs.umd.edu (J. Golbeck).

¹ Recommender systems, or “collaborative filtering” systems, make product recommendations based upon the opinions of people bearing high similarity with the user for whom to make recommendations.

However, in order to provide *meaningful results*, trust must reflect user similarity to some extent; recommendations only make sense when obtained from like-minded people exhibiting similar taste.

Abdul-Rahman and Hailes [2] claim that given some predefined domain and context, e.g., communities of people reading books, people create ties of friendship and trust primarily with people resembling their own profile of interest. Jensen et al. [25] make likewise assumptions, supposing similarity as a strong predictor of friendship: “If I am a classic car enthusiast, for example, my friends will likely share my interests [...]. In other words, my circle of friends is likely to either share the same values as I do, or at least tolerate them.”

Reasons for that phenomenon are manifold and mostly sociologically motivated, like people’s striving for some sort of social affiliation [15]. For instance, Pescovitz [48] describes endeavors to identify trust networks for crime prevention and security. Its advocates operate “on the assumption that birds of a feather tend to flock together [...]”, an ancient and widely-known aphorism. However, though belief in the positive relation of trust and user similarity has been widely adopted and presupposed, thus constituting the foundations for trust-based recommender and rating systems, to our best knowledge, no endeavors have been made until now to provide “real-world” empirical evidence.

We want to investigate and analyze whether a positive correlation actually holds, relying upon data mined from two online communities, one focusing on books, the other on movies. Studies involve several hundreds of members indicating which books or movies they like and which other community members they trust. Demonstrating a correlation between trust and user similarity provides the foundation for filtering information based on the user’s social relationships. Collaborative filtering algorithms can then adopt some features of this social filtering by integrating results obtained through computational trust models. Our motivation mainly derives from incorporating trust models into decentralized recommender systems, exploiting trust not only for selecting small neighborhoods upon which to perform collaborative filtering, but also for weeding out irrelevant peers.

The paper is organized as follows. In Section 2, we briefly outline existing approaches dealing with the incorporation of trust into reputation systems and online recommenders. Next, we present a survey about results from socio-psychological research which bear relevance for our studies. Section 4 presents experiments we performed in order to investigate the presence or absence of interactions between trust and similarity in the context of the All Consuming community where implicit ratings

prevail. Complementing the evaluation framework, Section 5 presents an experimental study performed on FilmTrust, a community catering to movie lovers. Suggestions for the exploitation of correlation between trust and similarity are provided in Section 6, while Section 7 mentions open questions and possible future work.

2. Trust and recommender systems

Online recommender systems [49,33] intend to provide people with recommendations of products they might appreciate, taking into account their past ratings profile, purchase history, or interest. There are three main types of filtering systems [19]: collaborative, content-based, and economic filtering. While content-based filtering takes into account properties attributed to the nature of products themselves, collaborative filtering (CF) relies upon building “neighborhoods of like-minded customers” [53] whose rating history may then serve to generate new recommendations. Economic filtering, which filters items based on cost factors, has seen little practical application and exerts marginal impact only.

Sinha and Swearingen [57] have found that people prefer receiving recommendations from people they *know* and *trust*, i.e., friends and family-members, rather than from recommender systems. As a result, some researchers have started to focus on computational trust models as appropriate means to supplement or replace current collaborative filtering approaches.

Kautz et al. [29] mine social network structures in order to render expertise information exchange and collaboration feasible. Olsson [47] proposes an architecture combining trust, collaborative filtering and content-based filtering in one single framework, giving only vague information and insight, though. Another agent-based approach has been presented by Montaner et al. [42], who introduce “opinion-based” filtering. Montaner claims that trust should be *derived* from user similarity, implying that friends are exactly those people that resemble our very nature. However, Montaner’s model only extends to the agent world and does not reflect evidence acquired from real-world social studies concerning the formation of trust. Similar agent-based systems have been devised by Kinatader and Rothermel [31], Kinatader and Pearson [30], and Chen and Singh [11].

Apart from research in multi-agent systems, online communities also make use of trust networks. The well-known reviewers’ community Epinions (<http://www.epinions.com>) provides information filtering facilities based upon personalized webs of trust [20]. Guha, chief architect of Epinions, states that the trust-based filtering

approach has been greatly approved and appreciated by Epinions' members. However, empirical and statistical justifications underpinning these findings, like indications of a correlation between trust and interest similarity, have not been given so far. Epinions' rating and trust data have also been used by Massa and Avesani [38]. They propose to supersede CF-based neighborhoods by trust networks, making use of very basic propagation schemes. Initial empirical data has been provided in their work, indicating that precision does not decrease too much when using trust-based neighborhood formation schemes instead of common CF.

3. Evidence from social psychology

Research in social psychology offers some important results for investigating interactions between trust and similarity. However, most relevant studies primarily focus on *interpersonal attraction* rather than trust, and its possible coupling with similarity. Interpersonal attraction constitutes a major field of interest of social psychology, and the positive impact of attitudinal similarity on liking has effectively become one of its most reliable findings [4]. Studies have given extensive attention to three different types of interpersonal relationships, namely same-sex friendships, primarily among college students, cross-sex romantic relationships, again primarily among college students, and marriage [24]. Clearly, these three types of relationships also happen to be essential components of trust, though perfect equivalence does not hold. Moreover, the complex notion of interpersonal trust, already difficult to capture regarding the “lack of consensus” which has been pointed out by McKnight and Chervany [41], interacts with other sociological concepts not reflected through interpersonal attraction. These elusive components include reputation, skill, situational and dispositional aspects of interpersonal trust [37,36], and familiarity [13].

However, since explicit *trust* relationships have remained outside the scope of empirical analysis on the correlation with attitudinal similarity, we are forced to stick to interpersonal attraction instead. Clearly, results obtained must be treated with care before attributing them to interpersonal trust as well.

3.1. On interpersonal attraction and similarity

Early investigations date back until 1943, when Burgess and Wallin published their work about homogeneity of social attributes with respect to engaged couples [8]. Similarity could be established for nearly every

characteristic examined. However, according to Berscheid [4], these findings do not justify conclusions about positive effects of similarity on interpersonal attraction by themselves, since “part of the association between similarity and social choice undoubtedly is due not to personal preference, but to the fact that people tend to be thrown together in time and space with others similar to themselves”.

First large-scale experimental studies were conducted by Newcomb [44] and Byrne [9,10]. The former work focused on friendships between American college students and nowadays counts among the seminal works on friendship formation. By means of his longitudinal study, Newcomb could reveal a positive association between attraction and attitudinal value similarity. Byrne, doing extensive research and experiments in the area of attraction, conducted similar experiments, applying the now-famous “bogus stranger technique” [10].

3.1.1. Analysis of interactions between similarity and attraction

The result of Byrne's experiment well aligned with Newcomb's findings and confirmed that attitude similarity is a determinant of attraction. Rather than further document this fact, which counts among the most reliable findings in social psychology today [4], researchers have ever since attempted to identify the factors that mediate and define the *limitations* of positive association between similarity and attraction.

For instance, Byrne conjectured that one's mere discovery of some other person holding similar attitudes is reinforcing in itself, arguing that “the expression of similar attitudes by a stranger serves as a positive reinforcement, because consensual validation for an individual's attitudes and opinions and beliefs is a major source of reward for the drive to be logical, consistent, and accurate in interpreting the stimulus world” [10]. We suppose likewise effects when forging bonds of trust. The sheer observation that some other peer holds interests similar to our own, e.g., reading the same kinds of books, intuitively renders the latter more trustworthy in our eyes and engenders sentiments of “familiarity”. In fact, automated collaborative filtering systems exploit this conjecture in order to make reliable predictions of product preference [57].

Social psychologists have identified some other likely factors accounting for the similarity-attraction association. For example, the information that another person possesses similar attitudes may suggest his sympathy towards the individual, and “it is known that the anticipation of being liked often generates attraction in return” [4]. Jones et al. [27] provided some large-scale

empirical analysis for reciprocation of attraction from similar others.

3.1.2. Limitations

While positive association was attested for attitudinal similarity and interpersonal attraction, evidence could not be expanded to similarity in general. Berscheid [4] therefore notes that despite “considerable efforts to find a relationship between friendship choice and personality (as opposed to attitude) similarity, for example, the evidence for this hypothesis remains unconvincing [...]”.

This inability to establish an association between personality similarity and attraction does not prove overly harmful to our hypothesis, since personal interests represent traits of *attitude* rather than *personality*. However, even attitude similarity fails to produce attraction under certain circumstances. Snyder and Fromkin [58] reveal that perceiving very high similarity with another individual may even evoke *negative* sentiments towards that respective person. Moreover, according to Heider [22], “similarity can evoke disliking when the similarity carries with it disagreeable implications”, which common sense anticipates, likewise. Take narcissist persons as an example.

3.2. Conclusion

The preceding literature survey has shown that interactions between similarity traits and *interpersonal attraction* are difficult to capture. Even though the tight coupling between both concepts counts among social psychology’s most reliable findings, there are numerous caveats to take into consideration, like subtle distinctions between various types of similarity, e.g., attitudinal similarity and personality similarity. Moreover, most studies argue that attitudinal similarity *implies* attraction, whereas the latter proposition’s inversion, i.e., positing that similarity *follows* from attraction, has been subject to sparse research only. Common sense supports this thesis, though, since people tend to adopt attitudes of friends, spouses, etc.

4. On correlations between trust and similarity

Evidence from socio-psychological research thus gives hints proposing the existence of positive interactions between *trust* and *interest similarity*. However, the theory does not provide exhaustive support that allows us to take this assumption for granted. Mind that trust and interpersonal attraction, though subsuming several common aspects, e.g., friendship, familiarity, etc., are *not* fully compliant notions.

We intend to establish a formal framework for investigating interactions between trust and similarity, believing that given an application domain, such as, for instance, the book-reading domain, people’s trusted peers are on average considerably more similar to their sources of trust than arbitrary peers. More formally, let A denote the set of all community members, $\text{trust}(a_i)$ the set of all users trusted by a_i , and $\text{sim}: A \times A \rightarrow [-1, +1]$ some similarity function:

$$\sum_{a_i \in A} \frac{\sum_{a_j \in \text{trust}(a_i)} \text{sim}(a_i, a_j)}{|\text{trust}(a_i)|} \gg \sum_{a_i \in A} \frac{\sum_{a_j \in A \setminus \{a_i\}} \text{sim}(a_i, a_j)}{|A| - 1} \quad (1)$$

For instance, given that agent a_i is interested in Science-Fiction and Artificial Intelligence, the chances that a_j , trusted by a_i , also likes these two topics are much higher than for peer a_e not explicitly trusted by a_i . Various social processes are involved, such as participation in those social groups that best reflect our own interests and desires.

4.1. Model and data acquisition

In order to verify or refute our hypothesis for specific domains, we need to define an information model, determine metrics and methods for evaluation, and apply our framework to real-world data.

4.1.1. Information model

In our information model proposed, we assume relationships of *binary* preference, i.e., for every $a_i \in A$, we divide the set of agents A into trusted and non-trusted ones. Moreover, ratings are assumed implicit and non-quantifiable with respect to the degree of appreciation. Though making the data more imprecise, the model becomes more general and bindings with real-world data, such as the All Consuming dataset, are facilitated.

- (a) **Set of agents** $A = \{a_1, a_2, \dots, a_n\}$. A contains all agents part of the community, identified uniquely through URLs etc.
- (b) **Set of products** $B = \{b_1, b_2, \dots, b_m\}$. All domain-relevant products are stored in set B . Unique identifiers either refer to proprietary product codes from an online store, such as Amazon.com’s ASINs, or represent globally accepted standard codes, like ISBNs.
- (c) **Trust rating function** $\text{trust}: A \rightarrow 2^A$. For every $a_i \in A$, $\text{trust}(a_i)$ returns the set of agents $A_i \subseteq A$ that a_i trusts.

- (d) **User ratings R_1, R_2, \dots, R_n .** Every agent a_i is assigned a set $R_i \subseteq B$ which contains his *implicit* product ratings. Implicit ratings, such as purchase data, product mentions, etc., are far more common in electronic commerce systems and online communities than explicit ratings [46].
- (e) **Taxonomy C over set $D = \{d_1, d_2, \dots, d_j\}$.** Set D contains categories for product classification. Each category $d_e \in D$ represents one specific topic that products $b_k \in B$ may fall into. Topics express broad or narrow categories. The partial taxonomic order $C: D \rightarrow 2^D$ retrieves all immediate sub-categories $C(d_e) \subseteq D$ for topics $d_e \in D$. We require that $C(d_e) \cap C(d_h) = \emptyset$ holds for all $d_e, d_h \in D, e \neq h$, and hence impose tree-like structuring, similar to single-inheritance class hierarchies known from object-oriented languages. Leaf topics d_e are topics with zero outdegree, formally $C(d_e) = \perp$, i.e., most specific categories. Furthermore, taxonomy C has exactly one top element, \top , which represents the most general topic and has zero indegree.
- (f) **Descriptor assignment function $f: B \rightarrow 2^D$.** Function f assigns a set $D_k \subseteq D$ of product topics to every product $b_k \in B$. Note that products may possess several descriptors, for classification into one single category may be too imprecise.

4.1.2. Data acquisition

All Consuming (<http://www.allconsuming.net>) represents one of the few communities that allow members to express which other agents they trust, as well as which items, in our case books, they appreciate. In fact, to the best of our knowledge, All Consuming, Epinions (<http://www.epinions.com>), and the FilmTrust community (discussed in detail later) are the only communities that feature trust *and* product ratings. We favored All Consuming over Epinions because of All Consuming restricting itself to the book domain where content information, needed for computing similarities between users, is readily available. This is not the case for most other domains.

Users may import their list of trusted persons from other applications like FOAF [12]. Moreover, All Consuming offers to automatically compile information about books its members have read from their personal *weblog*.² Members may explicitly assert trust statements and indicate books they own, have read, like most, and so forth.

Trust assertions from user a_i to a_j in All Consuming are boolean, either denoting *full* trust, i.e., a_i explicitly stating trust in a_j , or *no* trust, if a_i does not. Moreover, book mentions in All Consuming seldom reflect “real” ratings, like dislike or liking. They rather indicate that agent a_i has *read* or *purchased* book b_k . These statements therefore count among *implicit* ratings, which complies well with our model. Clearly, people tend to only buy and read books they expect to appreciate. In fact, numerous recommender systems are purely based upon implicit ratings [46].

Our tools mined data from about 2074 weblogs contributing to the All Consuming information base, and 527 users issuing 4.93 trust statements on average (see Fig. 1). These users have mentioned 6592 different books altogether. In order to obtain category descriptors $f(b_i)$ for all discovered books b_i , we collected classification information from the Amazon.com online bookshop (<http://www.amazon.com>). For each book, Amazon.com provides an average of about 4.1 classification topics. These topics represent leaf nodes relating to the huge Amazon.com book taxonomy, featuring 13,394 categories after duplicate removal and data cleansing. Note that the Amazon.com book taxonomy exhibits a tree structure, which complies with our formal model in Section 4.1.1. More statistical information describing the All Consuming and Amazon.com datasets is given in Ref. [64], Epinions is dealt with in Ref. [39].

4.2. Similarity computation

In order to analyze interactions between trust and user similarity, we need mathematical models indicating how to *compute* similarity. The book domain bears some notable differences to most other domains like videos, computer games, and DVDs. First, every published book is uniquely identified by its ISBN, which makes it easy to ensure interoperability and gather supplementary information from various other sources, e.g., mentioned category descriptors from Amazon for any given ISBN. Second, the set of published books is vast and much larger than for videos or DVDs. Consequently, profile overlap, i.e., the number of books two given users $a_i, a_j \in A$ have both rated, is generally small. Common techniques used in collaborative filtering, such as computing Pearson’s correlation coefficient [55,50], are therefore bound to fail within our context. Even more advanced techniques, like Sarwar’s singular value decomposition [53], cannot reduce dimensionality satisfactorily for the underlying application domain.

² Weblogs are best described as personal online diaries.




Fig. 1. Visualization of the All Consuming network's largest connected component.

4.2.1. Profile generation

We propose another, more informed approach that does not represent users by their respective *product*-rating vectors of dimensionality $|B|$, but by vectors of interest scores assigned to *topics* taken from taxonomy C over product categories $d \in D$.

User profile vectors are thus made up of $|D|$ entries, which corresponds to the number of distinct classification topics. Moreover, making use of profile vectors representing interest in *topics* rather than product *instances*, we can exploit the hierarchical structure of taxonomy C in order to generate overlap and render the similarity computation more meaningful: for every topic $d_{k_e} \in f(b_k)$ of products b_k that agent a_i has implicitly rated, we also infer an interest score for all *super*-topics of d_{k_e} in user a_i 's profile vector. However, score assigned to super-topics decays with increasing distance from leaf node d_{k_e} . We furthermore normalize profile vectors with respect to the amount of

score assigned, according the arbitrarily fixed overall score s .

Suppose that $\vec{v}_i = (v_{i,1}, v_{i,2}, \dots, v_{i,|D|})^T$ represents the profile vector for user a_i , where $v_{i,k}$ gives the score for topic $d_k \in D$. Then we require the following equation to hold:

$$\forall a_i \in A : \sum_{k=1}^{|D|} v_{i,k} = s \quad (2)$$

By virtue of agent-wise normalization for a_i 's profile, the score for each product $b_k \in R_i$ amounts to $s/|R_i|$, inversely proportional to the number of distinct products that a_i has rated. Likewise, for each topic descriptor $d_{k_e} \in f(b_k)$ categorizing product b_k , we accord topic score $sc(d_{k_e}) = s/(|R_i| \cdot |f(b_k)|)$. Thus, the topic score for b_k is distributed evenly among its topic descriptors.

Let (p_0, p_1, \dots, p_q) denote the path from top element $p_0 = \top$ to descendant $p_q = d_{k_e}$ within the tree-structured

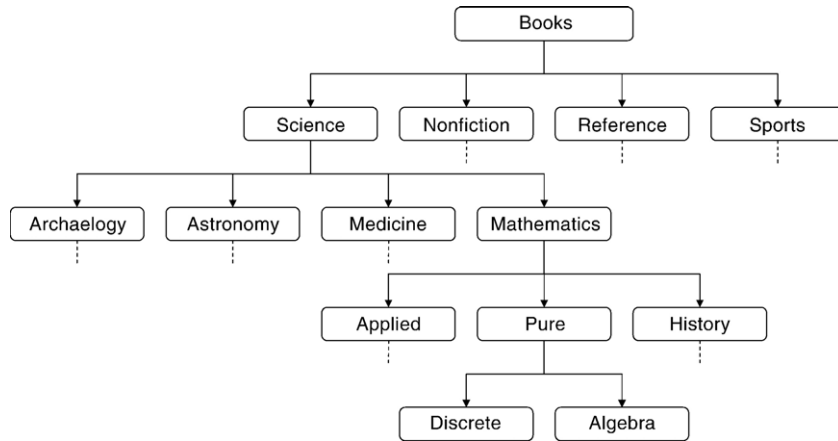


Fig. 2. Fragment from the Amazon.com book taxonomy.

taxonomy C for some given $d_{k_e} \in f(b_k)$. Then topic descriptor d_{k_e} has q super-topics. Score normalization and inference of fractional interest for super-topics imply that descriptor topic score $sc(d_{k_e})$ may not become *fully* assigned to d_{k_e} , but in part to all its ancestors p_{q-1}, \dots, p_0 , likewise. We therefore introduce another score function $sco(p_m)$ that represents the eventual assignment of score to topics p_m along the taxonomy path leading from $p_q = d_{k_e}$ to $p_0 = \top$:

$$\sum_{m=0}^q sco(p_m) = sc(d_{k_e}) \quad (3)$$

In addition, based on results obtained from research on semantic distance in taxonomies (e.g., see [7] and [52]), we require that interest score $sco(p_m)$ accorded to p_m , which is super-topic to p_{m+1} , depends on the number of siblings, denoted $sib(p_{m+1})$, of p_{m+1} : the fewer siblings p_m+1 possesses, the more interest score is accorded to its super-topic node p_m :

$$sco(p_m) = \kappa \cdot \frac{sco(p_{m+1})}{sib(p_{m+1}) + 1} \quad (4)$$

We assume that sub-topics have equal shares in their super-topic within taxonomy C . This assumption may imply several issues and raise concerns, e.g., when certain sub-taxonomies are considerably denser than others [51,52].

Propagation factor κ permits to fine-tune the profile generation process, depending on the underlying taxonomy's depth and granularity. For instance, we applied $\kappa=0.75$ for Amazon.com's book taxonomy, owing to its high level of nesting and "condensed" structure.³

³ As opposed to the Amazon movie taxonomy, which is comparatively shallow and broad. See [64].

Eqs. (3) and (4) describe conditions which have to hold for the computation of leaf node p_q 's profile score $sco(p_q)$ and the computation of scores for its taxonomy ancestors p_k , where $k \in \{0, 1, \dots, q-1\}$. We derive the following recursive definition for $sco(p_q)$:

$$sco(p_q) := \kappa \cdot \frac{sc(d_{k_e})}{g_q}, \quad (5)$$

where

$$g_0 := 1, g_1 := 1 + \frac{1}{sib(p_q) + 1},$$

and $\forall n \in \{2, \dots, q\}$

$$g_n := g_{n-1} + (g_{n-1} - g_{n-2}) \cdot \frac{1}{sib(p_{q-n+1}) + 1}$$

Computed scores $sco(p_m)$ are used to build a profile vector \vec{v}_i for user a_i , adding scores for topics in \vec{v}_i . The procedure is repeated for every product $b_k \in R_i$ and every $d_{k_e} \in f(b_k)$.

4.2.1.1. Example 1 (Profile computation). Suppose taxonomy C as depicted in Fig. 2, and propagation factor $\kappa=1$. Let a_i have implicitly rated four books, namely Matrix Analysis, Fermat's Enigma, Snow Crash, and Neuromancer. For Matrix Analysis, five topic descriptors are given, one of them pointing to leaf topic *Algebra* within our small taxonomy.

Suppose that $s=1000$ defines the overall accorded profile score. Then the score assigned to descriptor *Algebra* amounts to $s/(4 \cdot 5)=50$. Ancestors of leaf *Algebra* are *Pure*, *Mathematics*, *Science*, and top element *Books*. Therefore, score 50 must be distributed among these topics according to Eqs. (3) and (4). The application of Eq. (5) yields score 29.091 for topic *Algebra*.

Likewise, applying Eq. (4), we get 14.545 for topic *Pure*, 4.848 for *Mathematics*, 1.212 for *Science*, and 0.303 for top element *Books*.

The above procedure is repeated for all descriptors of all four books. The resulting values are then used to build profile vector \vec{v}_i for a_i .

4.2.2. Profile similarity computation

The presented approach computes flat profile vectors $\vec{v}_i \in [0, s]^{|D|}$ for agents a_i , assigning score values between 0 and maximum score s to every topic d from the set of product categories D . However, we still need to match these profile vectors against each other in order to come up with one single similarity metric value. Sarwar et al. [53] count nearest-neighbor techniques, like Pearson's correlation coefficient [55,50] and cosine similarity [3], widely known from information retrieval, among most popular approaches used for measuring profile proximity. We opt for Pearson correlation instead of cosine similarity since Pearson's correlation coefficient also allows for detecting negative correlation. Moreover, Pearson correlation represents the defacto standard for collaborative filtering [6,23].

For two given profile vectors $\vec{v}_i, \vec{v}_j \in [0, s]^{|D|}$, Pearson correlation is defined as below:

$$c(a_i, a_j) = \frac{\sum_{k=0}^{|D|} (v_{ik} - \bar{v}_i) \cdot (v_{jk} - \bar{v}_j)}{(\sum_{k=0}^{|D|} (v_{ik} - \bar{v}_i)^2 \cdot \sum_{k=0}^{|D|} (v_{jk} - \bar{v}_j)^2)^{\frac{1}{2}}} \quad (6)$$

Scalars \bar{v}_i and \bar{v}_j give mean values for vectors \vec{v}_i and \vec{v}_j . In our case, because of profile score normalization, both values are identical, i.e., $\bar{v}_i = \bar{v}_j = s/|D|$. Values for $c(a_i, a_j)$ range from -1 to $+1$, where negative values indicate negative correlation, and positive values positive correlation, respectively. Clearly, people who have read many books in common also have high similarity. For generic approaches to collaborative filtering, the proposition's negation also holds, i.e., people who have not read many books in common have low similarity. Our approach, on the other hand, may compute high similarity values even for pairs of agents that have little or even no books in common. Quality is highly dependent on the taxonomy's design and level of nesting. According to our scheme, the more score two profiles \vec{v}_i and \vec{v}_j have accumulated in same branches, the higher their computed similarity:

4.2.2.1. Example 2 (Positive correlation). Suppose the active user a_i has read only one single book b_m , bearing exactly one topic descriptor that classifies b_m into *Algebra*. User a_j has read another book b_n whose

topic descriptors point to diverse leaf nodes⁴ of History. Then $c(a_i, a_j)$ will still be reasonably high, for both profiles have significant overlap in categories *Mathematics* and *Science*.

Negative correlation occurs when users have completely diverging interests. For instance, in our information base mined from All Consuming, we had one user reading books mainly from the genres of Science Fiction, Fantasy, and Artificial Intelligence. The person in question was negatively correlated to another one reading books about American History, Politics, and Conspiracy Theories.

4.3. Experiment setup and analysis

We performed two experiments in order to analyze possible positive correlations between interest similarity and interpersonal trust for the book-reading domain. In both cases, experiments were run on data obtained from All Consuming (see Section 4.1.2). Considering the slightly different information makeup the two experiments were based upon, we expected the first to define an upper bound for the analysis, and the second one a lower bound. Results obtained confirm our conjecture.

4.3.1. Upper bound analysis

Before conducting the two experiments, we applied extensive data cleansing and duplicate removal to the All Consuming active user base of 527 members⁵. First, we pruned all users a_i having fewer than 3 books mentioned, removing them from user base A and from all sets $\text{trust}(a_j)$ where $a_i \in \text{trust}(a_j)$. Next, we discarded all users a_i who did not issue any trust assertions at all. Interestingly, some users had created several accounts. We discovered these "duplicates" by virtue of scanning through account names for similarity patterns and by tracking identical or highly similar profiles in terms of book mentions. Moreover, we stripped self-references, i.e., statements about users trusting themselves. Through application of data cleansing, 266 users were discarded from the initial test set, leaving 261 users for the upper bound experiment to run upon. We denote the reduced set of users by A' and corresponding trust functions by $\text{trust}'(a_i)$.

For every single user $a_i \in A'$, we generated his profile vector and computed the similarity score $c(a_i, a_j)$

⁴ Leaf nodes of topic History, representing history of mathematics, are not shown in Fig. 2.

⁵ All Consuming's crawled weblogs were not considered for the experiments, owing to their lack of trust web information.

for each trusted peer $a_j \in \text{trust}'(a_i)$. Next, we averaged these proximity measures, obtaining value z'_i :

$$z'_i := \frac{\sum_{a_j \in \text{trust}'(a_i)} c(a_i, a_j)}{|\text{trust}'(a_i)|} \quad (7)$$

Moreover, we computed a_i 's similarity with any other user from dataset A' , except a_i himself. Again, we took the average of these proximity measures and recorded the result s'_i :

$$s'_i := \frac{\sum_{a_j \in A' \setminus \{a_i\}} c(a_i, a_j)}{|A'| - 1} \quad (8)$$

A comparison of pairs (z'_i, s'_i) revealed that in 173 cases, users were more similar to their trusted peers than arbitrary ones. The opposite held for only 88 agents. Users had an average similarity score of 0.247 with respect to their trusted peers, while only exhibiting 0.163 with complete A' . In other words, users were more than 50% more similar to their trusted agents than arbitrary peers.

4.3.1.1. Distributions of z' and s' . Fig. 3(A) gives histogram representations for z' and s' , respectively. No agents have higher average similarity than 0.4, i.e., $s'_i \leq 0.4$ holds for all $a_i \in A'$. This is not the case for z' , as there remains a considerable amount of users a_i exhibiting an average trusted-peer similarity z' larger than 0.4. About 20 agents have $z'_i > 0.6$. Interestingly, while the overall peer similarity s' shows an almost perfect Gaussian distribution curve, its counterpart z' does not feature the typical bell shape. This observation raises some serious concerns when conducting analysis of statistical significance in Section 4.4.

4.3.1.2. Scatter plot. In order to directly match every user's overall similarity s' against his average trusted-peer similarity z' , Fig. 4(A) provides a scatter plot for the experiment at hand. The dashed line, dividing the scheme into an upper and lower region, models an agent a_i having identical similarity values, i.e., $s'_i = z'_i$. The plot exhibits a strong bias towards the upper region, which becomes particularly pronounced for agents a_i with $s'_i > 0.15$.

4.3.2. Lower bound analysis

The first experiment proposed that users tend to trust people that are significantly more similar to themselves than average users. However, we have to consider that All Consuming offers a feature that *suggests friends* to newbie users a_i . All Consuming chooses users who

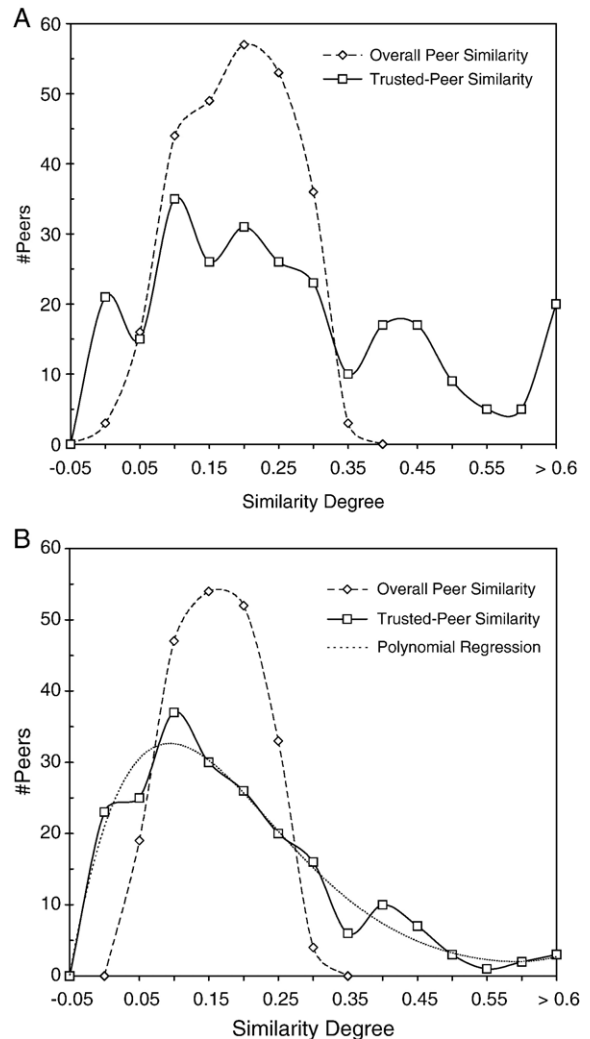


Fig. 3. Histograms of the upper bound (A) and lower bound (B) analysis.

have at least one book in common with a_i . We have reason to suspect that our first experiment was biased and too optimistic with respect to positive interactions between trust and similarity. Consequently, we pruned user set A' once again, eliminating trust statements whenever trusting and trusted user had at least one book in common. We denote the latter user base by A'' , now reduced to 210 trusting users, and indicate its respective trust functions by $\text{trust}''(a_i)$.

Clearly, our approach to eliminate All Consuming's intrusion into the natural process of trust formation entails the removal of many "real" trust relationships between users a_i and a_j , i.e., relationships which had been forged owing to a_i actually *knowing* and *trusting* a_j , and not because of All Consuming proposing a_j as an appropriate match for a_i .

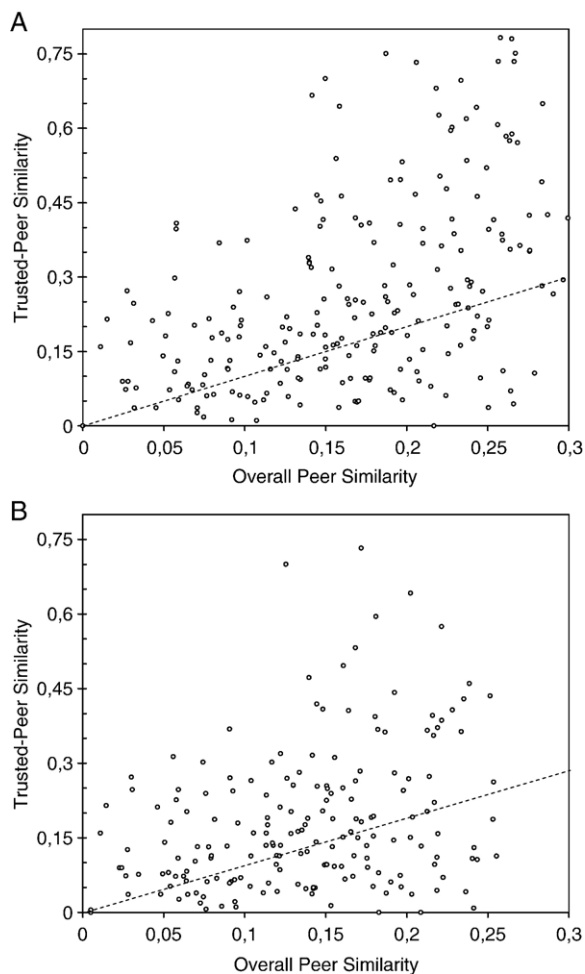


Fig. 4. Scatter plot for the upper bound (A) and lower bound (B) analysis.

For the second experiment, we computed values s_i'' and z_i'' for every $a_i \in A''$. We supposed results to be biased to the disadvantage of our conjecture, i.e., unduly lowering possible positive associations between trust and user similarity. Again, one should bear in mind that for set A'' , users did not have one single book in common with their trusted peers.

Results obtained from the second experiment corroborate our expectations, being less indicative for existing positive interactions between interpersonal trust and attitudinal similarity. Nevertheless, similarity values z_i'' still exceeded s_i'' : in 112 cases, people were more similar to their trusted fellows than arbitrary peers. The opposite held for 98 users. Mean values of z'' and s'' amounted to 0.164 and 0.134, respectively. Even for the lower bound experiment, users were still approximately 23% more similar to their trusted fellows than arbitrary agents.

4.3.2.1. Histogram curves. The bell-shaped distribution of s'' , depicted in Fig. 3(B), looks more condensed with respect to s' and has its peak slightly below the latter plot's curve. The differences between z'' and z' are even more pronounced, though, e.g., the shape of z'' 's histogram looks more "regular" than z' 's pendant. The approximation of z'' 's distribution, applying polynomial regression of degree 5, strongly resembles the Erlang- k distribution, supposing $k=2$. For similarity degrees above 0.35, peaks of z'' 's histogram are considerably less explicit than for z'' or have effectively disappeared, as is the case for degrees above 0.6.

4.3.2.2. Matching z_i'' and s_i'' . Fig. 4(B) gives the scatter plot of our lower bound analysis. The strong bias towards the upper region has become less articulate, though still clearly visible. Interestingly, the increase of ratio $z_i'' : s_i''$ for $s_i'' > 0.15$ still persists.

4.4. Statistical significance

We conclude our experimental analysis noticing that without exact knowledge of how much noise All Consuming's "friend recommender" adds to our obtained results, we expect the *true* correlation intensity between trust and interest similarity to reside somewhere within our computed upper and lower bound.

Moreover, we investigated whether the increase of mean values of z' with respect to s' , and z'' with respect to s'' , bears statistical significance or not.

For the analyses at hand, common parametrical one-factor ANOVA could not be applied to z' and s' , and z'' and s'' because of the following factors;

Gaussian distribution. The distributions of both samples have to be normal, even though small departures may be accommodated. While s' and s'' exhibit the latter Gaussian distribution property, z' and z'' obviously do not.

Equal variances. Data transformation, e.g., logarithmic, probits, etc., might be an option for z'' , bearing traits of Erlang-2. However, ANOVA also demands largely identical variances σ^2 . Since z'' 's variance is 5.33 times the variance of s'' , this criterion cannot be satisfied.

Owing to these two limitations, we opted for Kruskal–Wallis non-parametric ANOVA [56], which does not make any assumptions with respect to distribution and variance.

Table 1 shows result parameters obtained from analyzing the upper bound experiment. Since value p is

Table 1
Kruskal–Wallis ANOVA test results for the upper bound experiment

	<i>n</i>	Rank sum	Mean rank
z'	261	73702.0	284.56
s'	261	60719.0	234.44

Kruskal–Wallis Statistic 14.52.
 p 0.0001.

much smaller than 0.05, very high statistical significance holds, thus refuting the hypothesis that fluctuations between medians of s' and z' were caused by mere random.

For the lower bound experiment, on the other hand, no statistical significance was detected.

4.5. Conclusion

Both experiments suggest that the mean similarity of trusting and trusted peers exceeds the arbitrary user similarity. For the upper bound analysis, strong statistical significance was discovered, which was not the case for its lower bound pendant. However, assuming the true distribution curves to reside somewhere in between these bounds, and taking into account that both z' and z'' exhibit larger mean values than s' and s'' , respectively, the results we obtained bear strong indications towards positive interactions between interpersonal trust and interest similarity.

5. Trust and similarity in FilmTrust

FilmTrust (<http://trust.mindswap.org/FilmTrust/>) is a Web site that integrates social networks with movie ratings and reviews [16]. Users make connections to friends, and, more importantly for this work, rate how much they trust each friend's opinion about movies on a scale from 1 (low trust) to 10 (high trust). Users can also rate films on a scale from one half star (very bad) to four stars (excellent). Within the site, the social network is used as a recommender system to generate *predictive* ratings for movies based on the user's trust values for others in the network. The data from this Web site also provides an opportunity to study the relationship between user similarity and trust in a context where both trust values and opinions about movies have been *explicitly* rated on a scale. This section will introduce the FilmTrust network and use the data to illustrate a strong and significant correlation between trust and user similarity.

Note that by virtue of the different, prediction-oriented information model supposed for FilmTrust, the design of the study substantially differs from the evaluation frame-

work presented in previous sections, thus complementing our other approach.

5.1. FilmTrust introduction

The social networking component of the FilmTrust Web site allows users to make connections to friends. There are currently over 600 members of the FilmTrust community. Fig. 5 shows the structure of the social network. When adding a friend, users are required to provide a trust rating for that person. Because the context of the Web site is movie-specific, users are asked to rate how much they trust their friends' opinions of movies. Ratings are made on a 1 through 10 scale where 1 represents very low trust and 10 very high trust. There is no rating to reflect distrust in the system because the meaning of distrust is far less clear, both socially and computationally.

Part of the user's profile is a "Friends" page. In the FilmTrust network, relationships can be one-way, so the page displays a list of people the user has named as friends, and a second list of people who have named the user as a friend. An icon indicates reciprocal relationships and the trust ratings that the user assigned are shown next to each friend.

If trust ratings are visible to everyone, users can be discouraged from giving accurate ratings for fear of offending or upsetting people by giving them low ratings. Because honest trust ratings are important to the function of the system, these values are kept private and shown only to the user who assigned them. The ratings that people assigned to the user are not shown.

There are currently 672 members in the FilmTrust network. Of those, 241 are isolated nodes, they have no social connections and only participate for the movie




Fig. 5. Visualization of the FilmTrust network's largest connected component.

information. These disconnected nodes are not considered in our analysis. That leaves 431 remaining members. Because relationships are directional, that is, a person can list someone as a friend without that person listing them back, there are different in-and outdegrees for each node. The average outdegree (people the user has listed as a friend), is 3.72, while the average in-degree (people who listed a user as a friend) is 3.01. Thus, looking at all edges, the average degree is 6.73.

The average trust rating is 6.8 on a scale from 1 to 10, with a standard deviation of 2.23. The values are normally distributed with the exception of a spike for trust values of 10, see Fig. 6(A).

The other relevant feature of the Web site is a movie rating and review system. Users can choose any film and rate it on a scale of one half star to four stars. They can also write free-text reviews about movies. The user's "Movies" page displays data for every movie that he has rated or reviewed.

In FilmTrust, 1486 unique movies have been rated or reviewed by members. There are a total of 4833 reviews for 387 different movies. There are 13,997 ratings for 1479 different movies. This averages to 9.46 ratings per film. However, this distribution is skewed. 1437 movies have fewer than 50 ratings with about half (790) given only one rating. The average number of ratings given to this group of 1437 films is 2.66. On the other hand, fifty films have over 80 ratings, with an average of 202 ratings for each film in that group. This distribution is due to the fact that during the registration process, users are presented with and asked to rate the top fifty films from the AFI Top 100 Movies of All Time list. Thus, the movies on this list are much more likely to be rated. Of the 672 members, 563 have rated or reviewed films. On average, each person has rated or reviewed 24.9 films.

These explicit ratings of trust values and movies facilitate an analysis of the correlation between trust and user similarity.

5.2. Profile similarity computation

When users join the FilmTrust network, they are presented with a list of the top 50 films on the American Film Institute's top 100 movies list⁶ and asked to assign a star rating to any movies they have seen. Thus, there is a core set of films with ratings from many users that permit a better and more facile analysis.

Let $t_i(a_j)$ represent the trust rating that user $a_i \in A$ has assigned to user a_j and $r_j(b_k)$ represent the rating

⁶ In order to obtain the complete list, please refer to <http://www.afi.com/tvevents/100years/movies.aspx>.

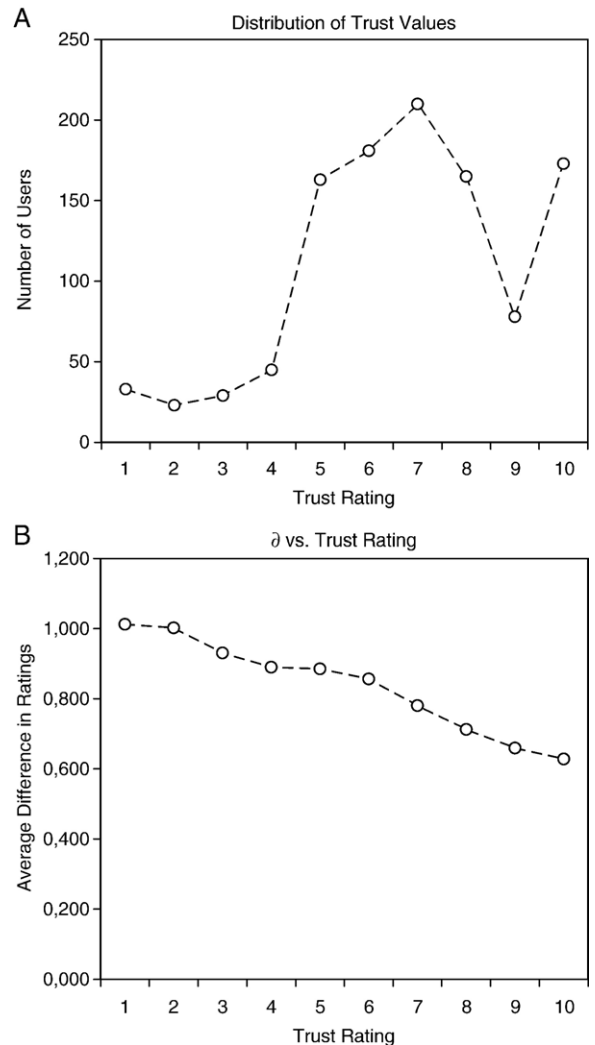


Fig. 6. Distribution of trust values (A) and mean difference $\bar{\delta}$ grouped by rating (B).

user a_j has assigned to movie b_k . Each neighbor of user a_j is contained in an adjacency list given by $\text{adj}(a_j)$. Each movie rated by user a_j is contained in the set $R_i = \{b \in B | r_i(b) \neq \perp\}$.

Similarity between users a_i and a_j on a given movie b_k is computed as the absolute difference $\delta = |r_i(b_k) - r_j(b_k)|$. To compute the correlation between trust and similarity, neighbors are grouped according to trust value. Since trust is allocated on an integer scale from 1 to 10, there are ten groups of neighbors. Let τ represent the trust value for which the similarity is being computed. Each movie that has been rated by user a_i and user a_j where $t_i(a_j) = \tau$ is compared using the absolute difference δ . The average $\bar{\delta}$, is used as the measure of similarity for trust value τ . That is, we want the average $|r_i(b_k) - r_j(b_k)|$, $\forall a_i, a_j \in A: t_i(a_j) = \tau, \forall b_k \in R_i \cap R_j$.

Table 2
Kruskal–Wallis ANOVA test results for FilmTrust data

Trust value	n	Rank sum	Mean rank
1	158	636,734.5	4029.97
2	203	820,833.5	4043.51
3	195	752,781.0	3860.42
4	296	1,129,921.0	3817.30
5	1235	4,681,499.5	3790.69
6	1043	3,874,654.5	3714.91
7	1217	4,233,673.5	3478.78
8	1142	3,771,554.0	3302.59
9	636	2,049,333.5	3222.22
10	898	2,713,791.0	3022.04

Kruskal–Wallis Statistic 17.96.
 p 0.0001.

For example, when $\tau=9$, we select all user pairs where user a_i has rated user a_j with a trust value of 9. Then, we select all movies b_k rated by both user a_i and user a_j , and take the absolute difference of their ratings for each b_k . The average of these differences over all movies for all $(a_i, a_j) \in A \times A$ where $t_i(a_j)=9$ is used as a measure of similarity for $\tau=9$.

As shown in Fig. 6(B), as the trust rating increases, $\bar{\delta}$ decreases. Recall that the difference in movie ratings is a measure of error, so a lower difference means higher similarity. Thus, we see that user similarity increases as trust increases.

5.3. Statistical significance

There are two questions as to the significance of these results. First, is there a statistical correlation between trust and similarity, and second, is the difference statistically significant.

Fig. 6(B) appears to show a *linear* relationship between $\bar{\delta}$ and trust. We compute this with the Pearson correlation [55]. It is important to point out that both trust rating and $\bar{\delta}$ values are on a scale of discrete values. Although the data are not continuous, they are not strictly categorical either. We choose to use the Pearson correlation, but because of the scalar nature of the values, the coefficient serves more as a general indicator than a precise measurement of the correlation.⁷ Nevertheless, the correlation between trust and $\bar{\delta}$ is strong and significant. For these data, the correlation coefficient $r=-0.987$ indicates an almost perfect negative linear relationship: as trust increases, the average

⁷ The analysis methodology presented herein resembles the approach adopted by Newman [45] to investigate the strength of assortative mixing in social networks.

difference in ratings decreases (ratings become more similar). For these sample sizes, this correlation is significant at $p<0.01$.

To test the significance of the change in $\bar{\delta}$ as trust values change, we used the same Kruskal–Wallis ANOVA as was used above. While the sample distributions are normal in this case, the variance is different in each group, making the standard ANOVA inappropriate. Table 2 shows the result parameters of the Kruskal–Wallis ANOVA. The p value is quite small, indicating that $\bar{\delta}$ decreases significantly as trust increases.

5.4. Conclusion

The data from the FilmTrust Web site give insight into how user similarity changes when there is a range of trust values. In the analysis presented here, we see that as the trust between users increases, the difference in the ratings they assign to movies decreases. This is a significant change, and the correlation between trust and similarity is strong. Our finding reinforces the results from Section 4 by showing that changes in similarity are correlated to changes in trust between users.

6. Exploiting correlations between trust and similarity

We envision trust to play an important role for *decentralized* recommender systems. These filtering systems suppose distributed data and control and currently face various problems inherent to their very nature:

6.1. Credibility and attack–resistance

The Semantic Web and other open systems lack dedicated mechanisms and facilities to verify user identity. These systems tend to encourage insincerity and fraudulent behavior. Moreover, penalization and banishment are hard to accomplish. Collaborative filtering becomes particularly susceptible to attack, for malicious users simply have to create profiles replicating the victim’s in order to obtain high similarity. Then they can lure the victim into buying items the purchase of which may provide some utility for the attacker.

6.2. Product–user matrix sparseness

Communities often limit the number of ratable products, therefore avoiding product–user matrices from becoming overly sparse. Besides, Ringo [55] and other systems require users to rate items from small product subsets to generate user profiles with sufficient overlap. However, decentralized recommender system cannot

suppose reduced item sets. Bear in mind that controlling product set contents and having users rate certain goods presupposes some central authority.

6.3. Computational complexity and scalability

Centralized systems are able to control and limit the number of members. Depending on the community's size, large-scale server clusters ensure proper operativeness and scalability. In general, recommender systems imply heavy computations. For instance, collaborative filtering systems compute Pearson correlation for users a_i offline rather than on-the-fly. Recall that coefficients $c(a_i, a_j)$ have to be computed for every other agent $a_j \in A$. Clearly, this approach does not work for large decentralized systems. Sensible prefiltering mechanisms which still ensure reasonable recall are needed.

From a practitioner's point of view, the above issues need to be addressed before decentralized recommender systems can be crafted. The knowledge that trust and similarity positively correlate can be exploited in the following way, solving the Gordian knot without resorting to swords:

Trust addresses the credibility problem. Every agent builds his own neighborhood of trusted peers, relying upon direct trust statements and those from trusted peers, likewise [14]. For deriving trust, numerous metrics have been proposed during the mid-nineties, among those [40,1,5,35], and more recently [17,18]. However, we believe that local group trust metrics like Levien's Advogato [34] and Appleaseed [61,62] best fit neighborhood formation in decentralized systems [59]. Unfortunately, trust cannot handle product–user matrix sparseness, nor substantially reduce dimensionality. Supplementary approaches are needed, e.g., taxonomy-based filtering techniques [63] similar to the one proposed.

Increased computational complexity and loss of scalability are mitigated and may even be eliminated when supposing positive correlation between trust and user similarity. Note that the complexity issue *per se* does not require the latter correlation to hold: limiting collaborative filtering to selected peers part of agent a_i 's trust neighborhood entails complexity reduction, too. However, when supposing that trust does not reflect similarity, serious tradeoffs are implied, because scalability comes at the expense of neighborhood quality. The trust neighborhood A_i of agent a_i only represents one tiny fraction of the overall system A . Moreover, this fraction does not necessarily contain similar peers. Instead, trusted agents are on average no more similar than arbitrary ones. The number of agents $a_j \in A_i$ with $c(a_i, a_j)$ above some threshold t , found by the filtering process, degrades *pro-*

portionally with the neighborhood's size. On the other hand, when assuming that trust does correlate with similarity, the degradation does not take place as fast, thus ensuring reasonable neighborhood quality.

The approach pursued by Epinions [20,21] relies upon trust networks as only filtering mechanism, clearly exploiting the correlation. Positive user feedback backs the design decision. Nevertheless, we believe that trust should *supplement* rather than replace existing filtering techniques. For instance, application of collaborative filtering to computed trust neighborhoods A_i might boost precision significantly.

7. Discussion and outlook

We articulated our hypothesis that dependencies between trust and user similarity exists when the community's trust network is tightly bound to some particular application. Empirical evidence has been provided based upon data obtained from the All Consuming book-readers' community and the FilmTrust community of movie aficionados. To our best knowledge, similar experiments have not been performed before, since communities incorporating explicit trust models are still very sparse.

We believe that our results will have substantial impact for ongoing research in recommender systems, where discovering user similarity plays an important role. Decentralized approaches will especially benefit from trust network leverage. The outstanding feature of trust networks refers to sensible prefiltering of like-minded peers and credibility of recommendations. Arbitrary social networks, on the other hand, only allow for computation complexity reduction.

Though backing our experiments with information involving several hundreds of people, studies for distinct interest domains are required. We would also like to run our analysis on communities larger than All Consuming and FilmTrust.

References

- [1] Alfarez Abdul-Rahman, Stephen Hailles, A distributed trust model, New Security Paradigms Workshop, September 1997, pp. 48–60, Cumbria, UK.
- [2] Alfarez Abdul-Rahman, Stephen Hailles, Supporting trust in virtual communities, Proceedings of the 33rd Hawaii International Conference on System Sciences, Maui, HI, USA, January 2000.
- [3] Ricardo Baeza-Yates, Berthier Ribeiro-Neto, Modern Information Retrieval, Addison-Wesley, Reading, MA, USA, ISBN: 0-20-139829-X, May 1999.
- [4] Ellen Berscheid, Interpersonal attraction, in: Daniel Gilbert, Susan Fiske, Gardner Lindzey (Eds.), 4th edition, The Handbook of Social Psychology, vol. II, McGraw-Hill, New York, NY, USA, ISBN: 0-19-521376-9, 1998.

- [5] Thomas Beth, Malte Borchering, Birgit Klein. Valuation of trust in open networks, Proceedings of the 1994 European Symposium on Research in Computer Security, 1994, pp. 3–18, Brighton, UK.
- [6] John Breese, David Heckerman, Carl Kadie, Empirical analysis of predictive algorithms for collaborative filtering, Proceedings of the Fourteenth Annual Conference on Uncertainty in Artificial Intelligence, Morgan Kaufmann, Madison, WI, USA, July 1998, pp. 43–52.
- [7] Alexander Budanitsky, Graeme Hirst, Semantic distance in WordNet: an experimental, application-oriented evaluation of five measures, Proceedings of the Workshop on WordNet and Other Lexical Resources, Pittsburgh, PA, USA, June 2000.
- [8] Ernest Burgess, Paul Wallin, Homogamy in social characteristics, *American Journal of Sociology* 2 (49) (1943) 109–124.
- [9] Donn Byrne, Interpersonal attraction and attitude similarity, *Journal of Abnormal and Social Psychology* 62 (1961) 713–715.
- [10] Donn Byrne, *The Attraction Paradigm*, Academic Press, New York, NY, USA, 1971.
- [11] Mao Chen, Jaswinder Singh, Computing and using reputations for internet ratings, Proceedings of the 3rd ACM Conference on Electronic Commerce, ACM Press, Tampa, FL, USA, 2001, pp. 154–162.
- [12] Edd Dumbill, Finding friends with XML and RDF, IBM's XML Watch, June 2002.
- [13] Sabine Einwiller, The significance of reputation and brand in creating trust between an online vendor and its customers, in: Otto Petrovic, Markus Fallenböck, Christian Kittl (Eds.), *Trust in the Network Economy*, Springer-Verlag, Heidelberg, Germany, ISBN: 3-211-06853-8, 2003, pp. 113–127.
- [14] Andrew Fernandes, Risking trust in a public key infrastructure: old techniques of managing risk applied to new technology, *Decision Support Systems* 31 (3) (2001) 303–322.
- [15] Barbara Given, Teaching to the brain's natural learning systems, Association for Supervision and Curriculum Development, Alexandria, VA, USA, May 2002.
- [16] Jennifer Golbeck, Computing and Applying Trust in Web-based Social Networks, PhD thesis, University of Maryland, College Park, MD, USA, April 2005.
- [17] Jennifer Golbeck, James Hendler, Accuracy of metrics for inferring trust and reputation in Semantic Web-based social networks, Proceedings of the 14th International Conference on Knowledge Engineering and Knowledge Management, Northamptonshire, UK, October 2004.
- [18] Jennifer Golbeck, Bijan Parsia, James Hendler, Trust networks on the Semantic Web, Proceedings of Cooperative Intelligent Agents, Helsinki, Finland, August 2003.
- [19] David Goldberg, David Nichols, Brian Oki, Douglas Terry, Using collaborative filtering to weave an information tapestry, *Communications of the ACM*, ISSN: 0001-0782 35 (12) (1992) 61–70.
- [20] Ramanathan Guha, Open rating systems, Technical Report, Stanford Knowledge Systems Laboratory, Stanford, CA, USA, 2003.
- [21] Ramanathan Guha, Ravi Kumar, Prabhakar Raghavan, Andrew Tomkins, Propagation of trust and distrust, Proceedings of the Thirteenth International World Wide Web Conference, ACM Press, New York, NY, USA, May 2004.
- [22] Fritz Heider, *The Psychology of Interpersonal Relations*, Wiley, New York, NY, USA, 1958.
- [23] Jonathan Herlocker, Joseph Konstan, Al Borchers, John Riedl, An algorithmic framework for performing collaborative filtering, Proceedings of the 22nd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, ACM Press, Berkeley, CA, USA, ISBN: 1-58113-096-1, 1999, pp. 230–237.
- [24] Ted Huston, George Levinger, Interpersonal attraction and relationships, *Annual Review of Psychology* 29 (1978) 115–156.
- [25] Carlos Jensen, John Davis, Shelly Farnham, Finding others online: reputation systems for social online spaces, Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, ACM Press, Minneapolis, MN, USA, ISBN: 1-58113-453-3, 2002, pp. 447–454.
- [26] Andrew Jones, On the concept of trust, *Decision Support Systems* 33 (3) (2002) 225–232.
- [27] Edward Jones, Linda Bell, Elliot Aronson, The reciprocation of attraction from similar and dissimilar others, in: Charles McClintock (Ed.), *Experimental Social Psychology*, Holt, Rinehart, and Winston, New York, NY, USA, 1972.
- [28] Audun Jøsang, Rolan Ismail, Colin Boyd, Survey of trust and reputation systems for online service provision, *Decision Support Systems* (2005) (To appear).
- [29] Henry Kautz, Bart Selman, Mehul Shah, Referral Web: combining social networks and collaborative filtering, *Communications of the ACM*, ISSN: 0001-0782 40 (3) (March 1997) 63–65.
- [30] Michael Kinatader, Siani Pearson, A privacy-enhanced peer-to-peer reputation system, Proceedings of the 4th International Conference on Electronic Commerce and Web Technologies LNCS, vol. 2378, Springer-Verlag, Prague, Czech Republic, September 2003.
- [31] Michael Kinatader, Kurt Rothermel, Architecture and algorithms for a distributed reputation system, in: Paddy Nixon, Sotirios Terzis (Eds.), Proceedings of the First International Conference on Trust Management LNCS, vol. 2692, Springer-Verlag, Crete, Greece, April 2003, pp. 1–16.
- [32] Mathias Klang, Who do you trust? Beyond encryption, secure e-business, *Decision Support Systems* 31 (3) (2001) 293–301.
- [33] Joseph Konstan, Introduction to recommender systems: algorithms and evaluation, *ACM Transactions on Information Systems*, ISSN: 1046-8188 22 (1) (2004) 1–4.
- [34] Raph Levien, Attack-resistant Trust Metrics, PhD thesis, University of California at Berkeley, Berkeley, CA, USA, 2004. To appear.
- [35] Raph Levien, Alexander Aiken, Attack-resistant trust metrics for public key certification, Proceedings of the 7th USENIX Security Symposium, San Antonio, TX, USA, January 1998.
- [36] Stephen Marsh, Optimism and pessimism in trust, in: Jose Ramirez (Ed.), Proceedings of the Ibero-American Conference on Artificial Intelligence, McGraw-Hill, Caracas, Venezuela, 1994.
- [37] Stephen Marsh, Formalising Trust as a Computational Concept, PhD thesis, Department of Mathematics and Computer Science, University of Stirling, Stirling, UK, 1994.
- [38] Paolo Massa, Paolo Avesani, Trust-aware collaborative filtering for recommender systems, in: Robert Meersman, Zahir Tari (Eds.), Proceedings of the DOA/CoopIS/ODBASE Confederated International Conferences (1), LNCS, vol. 3290, Springer-Verlag, Larnaca, Cyprus, ISBN: 3-540-23663-5, October 2004, pp. 492–508.
- [39] Paolo Massa, Bobby Bhattacharjee, Using trust in recommender systems: an experimental analysis, in: Christian Jensen, Stefan Poslad, Theodosios Dimitrakos (Eds.), Proceedings of the 2nd International Conference on Trust Management, LNCS, vol. 2995, Springer-Verlag, Oxford, UK, March 2004.
- [40] Ueli Maurer, Modelling a public key infrastructure, in: Elisa Bertino (Ed.), Proceedings of the 1996 European Symposium on

- Research in Computer Security, LNCS, vol. 1146, Springer-Verlag, Rome, Italy, 1996, pp. 325–350.
- [41] Harrison McKnight, Norman Chervany, The meaning of trust, Technical Report MISRC 96-04, Management Information Systems Research Center, University of Minnesota, MN, USA, 1996.
- [42] Miquel Montaner, Beatrix López, Josep de la Rosa, Opinion-based filtering through trust, in: Sascha Ossowski, Onn Shehory (Eds.), Proceedings of the Sixth International Workshop on Cooperative Information Agents, LNAI, vol. 2446, Springer-Verlag, Madrid, Spain, September 2002, pp. 164–178.
- [43] Lik Mui, Mojdeh Mohtashemi, Ari Halberstadt, A computational model of trust and reputation, Proceedings of the 35th Hawaii International Conference on System Sciences, Big Island, HI, USA, January 2002, pp. 188–196.
- [44] Theodore Newcomb, *The Acquaintance Process*, Holt, Rinehart, and Winston, New York, NY, USA, 1961.
- [45] Mark Newman, The structure and function of complex networks, *SIAM Review* 45 (2) (2003) 167–256.
- [46] David Nichols, Implicit rating and filtering, Proceedings of the Fifth DELOS Workshop on Filtering and Collaborative Filtering, ERCIM, Budapest, Hungary, 1998, pp. 31–36.
- [47] Tomas Olsson, Decentralized social filtering based on trust, Working Notes of the AAAI-98 Recommender Systems Workshop, Madison, WI, USA, 1998.
- [48] David Pescovitz, The best new technologies of 2003, *Business 2.0*, vol. 11, Time Inc. Publishing, November 2003.
- [49] Paul Resnick, Hal Varian, Recommender systems, *Communications of the ACM*, ISSN: 0001-0782 40 (3) (1997) 56–58.
- [50] Paul Resnick, Neophytos Iacovou, Mitesh Suchak, Peter Bergstorm, John Riedl, GroupLens: an open architecture for collaborative filtering of netnews, Proceedings of the ACM 1994 Conference on Computer-Supported Cooperative Work, ACM, Chapel Hill, NC, USA, 1994, pp. 175–186.
- [51] Philip Resnik, Using information content to evaluate semantic similarity in a taxonomy, Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence, Montreal, Canada, 1995, pp. 448–453.
- [52] Philip Resnik, Semantic similarity in a taxonomy: an information-based measure and its application to problems of ambiguity in natural language, *Journal of Artificial Intelligence Research* 11 (1999) 95–130.
- [53] Badrul Sarwar, George Karypis, Joseph Konstan, John Riedl, Application of dimensionality reduction in recommender systems, ACM WebKDD Workshop, Boston, MA, USA, August 2000.
- [54] Badrul Sarwar, George Karypis, Joseph Konstan, John Riedl, Item-based collaborative filtering recommendation algorithms, Proceedings of the Tenth International World Wide Web Conference, Hong Kong, China, May 2001.
- [55] Upendra Shardanand, Patti Maes, Social information filtering: algorithms for automating “word of mouth”, Proceedings of the ACM CHI Conference on Human Factors in Computing Systems, ACM Press, Denver, CO, USA, May 1995, pp. 210–217.
- [56] Sidney Siegel, John Castellan, *Non-parametric Statistics for the Behavioral Sciences* (2nd edition), McGraw-Hill, New York, NY, USA, 1988.
- [57] Rashmi Sinha, Kirsten Swearingen, Comparing recommendations made by online systems and friends, Proceedings of the DELOS-NSF Workshop on Personalization and Recommender Systems in Digital Libraries, Dublin, Ireland, June 2001.
- [58] Charles Snyder, Howard Fromkin, *Uniqueness: the Human Pursuit of Difference*, Plenum, New York, NY, USA, 1980.
- [59] Cai-Nicolas Ziegler, Semantic Web recommender systems. In Wolfgang Lindner, Mauro Mesiti, Can Türker, Yannis Tzitzikas, and Athena Vakali, editors, *EDBT 2004 Workshops (PhD, DataX, PIM, P2P and DB, and ClustWeb)*, LNCS, vol. 3268, pages 78–89. Heraklion Greece, revised selected papers, Springer-Verlag, November 2004.
- [60] Cai-Nicolas Ziegler, Towards Decentralized Recommender Systems. PhD thesis, Albert-Ludwigs-Universität Freiburg, Freiburg i.Br., Germany, June 2005.
- [61] Cai-Nicolas Ziegler, Georg Lausen, Spreading Activation Models for Trust Propagation, Proceedings of the IEEE International Conference on e-Technology, e-Commerce, and e-Service, Taipei, Taiwan, IEEE Computer Society Press, March 2004.
- [62] Cai-Nicolas Ziegler, Georg Lausen, Propagation models for trust and distrust in social networks, *Information Systems Frontiers* 7 (4–5) (2005) 337–358.
- [63] Cai-Nicolas Ziegler, Georg Lausen, Lars Schmidt-Thieme, Taxonomy-driven computation of product recommendations, Proceedings of the 2004 ACM CIKM Conference on Information and Knowledge Management, pages 406–415, Washington, D.C., USA, ACM Press, ISBN: 1-58113-874-1, November 2004.
- [64] Cai-Nicolas Ziegler, Sean McNee, Joseph Konstan, Georg Lausen, Improving recommendation lists through topic diversification, Proceedings of the 14th International World Wide Web Conference, Chiba, Japan, ACM Press, May 2005.



Cai-Nicolas Ziegler is consultant and researcher in the domain of text mining and natural language processing (NLP) at Siemens AG, Corporate Research and Technology (CT) in Munich, Germany. Before joining Siemens AG, he has been a post-doctoral researcher at DBIS, the Databases and Information Systems group of the University of Freiburg, Germany. Cai-Nicolas received his PhD in Computer Science from the University of Freiburg in 2005, and his Diploma (equivalent to MSc) from the University of Passau.



Jennifer Golbeck is a post-doctoral researcher at MINDSWAP and the University of Maryland Institute for Advanced Computer Studies (UMIACS) at the University of Maryland, College Park. Her research interests include social network analysis, trust models, the Semantic Web, and their relationships to intelligent systems. Jennifer received her PhD in Computer Science from the University of Maryland in 2005, and her Master of Science and Bachelor's degrees from the University of Chicago.