Mozhgan Tavakolifard

On Some Challenges for Online Trust and Reputation Systems

Thesis for the degree of Philosophiae Doctor

Trondheim, August 2012

Norwegian University of Science and Technology
Faculty of Information Technology, Mathematics and Electrical Engineering
Department of Telematics

NTNU – Trondheim
Norwegian University of
Science and Technology

NTNU
Norwegian University of Science and Technology

Thesis for the degree of Philosophiae Doctor

Faculty of Information Technology, Mathematics and Electrical Engineering
Department of Telematics

© Mozhgan Tavakolifard

ISBN 978-82-471-3535-8 (printed ver.)
ISBN 978-82-471-3536-5 (electronic ver.)
ISSN 1503-8181

Doctoral theses at NTNU, 2012:127

Printed by NTNU-trykk

Dedication

I would like to dedicate this Doctoral dissertation to my husband, my parents, and my little daughter. There is no doubt in my mind that without their continued support and counsel I could not have completed this process.

دل گرچه در این بادیه بسیار شتافت یک موی ندانست ولی موی شکافت

اندر دل من هزار خورشید بتافت آخر به کمال ذره ای راه نیافت

ابوعلی سینا

Abstract

The Internet form a globally distributed network that provides a ubiquitous medium for interaction, the exchange of ideas, and commerce. The web is pervading our everyday lives in ways that were unimaginable even ten years ago. The evolving use of the web requires robust and efficient trust and reputation management mechanisms. During the past decade, online trust and reputation systems have provided cogent answers to emerging challenges in the global computing infrastructures relating to computer and network security, electronic commerce, virtual enterprises, social networks and cloud computing. The goal of these systems in such global computing infrastructures is to allow entities to reason about the trustworthiness of other entities and to make autonomous decisions on the basis of trust. This requires the development of computational trust models that enable entities to reason about trust and to verify the properties of a particular interaction. The robustness of these mechanisms, which is one of the critical factors for the success of this technology, is currently not being sufficiently addressed. The global computing infrastructure is highly dynamic with continuously appearing and disappearing entities and services. It is vital that the associated computational trust model is able to incorporate this dynamism and that equally flexible legislative and regulatory frameworks emerge. In this thesis, we present an overview of the characteristics of online trust and reputation models and systems through a multidimensional framework, which can serve as a basis to understand the current state of the art in the area. The critical open challenges that limit the effectiveness of today's trust and reputation systems are discussed by providing a comprehensive literature review. Furthermore, we present a set of our contributions as a way to address some of these challenges and propose prospectives for online trust and reputation systems.

Preface

This thesis is submitted to the Norwegian University of Science and Technology (NTNU) for partial fulfilment of the requirements for the degree of philosophiae doctor (Ph.d.).

This doctoral work has been performed at the Centre for Quantifiable Quality of Service in Communication Systems (Q2S), Centre of Excellence (CoE), with Professor Svein J. Knapskog as main supervisor and with co-supervisor Professor Peter Herrmann at the Department of Telematics.

Q2S is established and funded by the Research Council of Norway, NTNU, UNINETT, and Telenor.

Acknowledgements

First of all, I would like to thank my thesis advisor Professor Svein J. Knapskog. Your encouragement and support have been invaluable during my work. I would never have been able to finish this thesis without your help. I would like also to give my special appreciation to Professor Kevin C. Almeroth from University of California, at Santa Barbara for our collaboration and discussions that led my research during my visit there. Your contributions, encouragement, and insightful feedback are of very high value to me and working with you was very inspiring.

Many thanks to Anniken Skotvoll for providing a very pleasant and friendly work atmosphere. I felt like being at my home country from the first day because of you and Svein. Thanks also to Professor Peter Herrmann, my co-advisor, for our numerous valuable and fruitful discussions to the research presented in this thesis. Several other people also deserve thank for having directly or indirectly contributed to this thesis work through cooperation and/or co-authorship. A special thanks goes to two of the co-authors for their contributions: Professor Pinar Ozturk and Dr. Marie Elizabeth Gaup Moe.

Finally, but most importantly, I would like to thank my husband, my parents, my daughter Roxanna for all your understanding, support and love.

Contents

List of Papers

Publications Included in the Thesis

These papers are included as Part II of this thesis.

- PAPER A:
  Mozhgan Tavakolifard, Kevin C. Almeroth. *Social Computing: An Intersection of Recommender Systems, Trust/Reputation Systems, and Social Networks*. IEEE Network Magazine. In press, 2012.
- PAPER B:
  Mozhgan Tavakolifard, Kevin C. Almeroth, Pinar Ozturk. *Subjectivity handling of ratings for Trust and Reputation systems: An Abductive Reasoning Approach*. JDCTA: International Journal of Digital Content Technology and its Applications. Vol. 5, No. 11, 2011.
- PAPER C:
  Mozhgan Tavakolifard, Kevin C. Almeroth. *The Hidden Trust Network underlying Twitter*. Submitted to the IEEE Transactions on Information Forensics and Security. 2012.
- PAPER D:
  Mozhgan Tavakolifard, Pinar Ozturk. *Situation-based Trust Adjustment by Conditional Trust Reasoning*. Proceedings of the Networking and Electronic Commerce Research Conference (NAEC). Riva del Garda, Italy, Oct 2011.
- PAPER E:
  Mozhgan Tavakolifard, Svein J. Knapskog. *Trust Evaluation Initialization Using Contextual Information*. Proceedings of the The International Conference on Management of Emergent Digital EcoSystems (MEDES). San Francissco, USA, Nov 2011.
- PAPER F:
  Mozhgan Tavakolifard, Svein J. Knapskog. *A Probabilistic Reputation Algorithm for Decentralized Multi-Agent Environments*. Electronic Notes in Theoretical Computer Science. Vol. 244, 2009.
- PAPER G:
  Mozhgan Tavakolifard, Peter Herrmann, Svein J. Knapskog. *Inferring Trust*

based on Similarity with TILLIT. IFIP Advances in Information and Communication Technology. Vol. 300, 2009.

- PAPER H:
  Mozhgan Tavakolifard, Peter Herrmann, Pinar Ozturk. *Analogical Trust Reasoning*. IFIP Advances in Information and Communication Technology. Vol. 300, 2009.

Other Papers by the Author

These papers were also prepared while working with this thesis.

- Mozhgan Tavakolifard, Kevin C. Almeroth. *A Taxonomy to Express Open Challenges in Trust and Reputation Systems*. Journal of Communications. In press, 2012.

- Mozhgan Tavakolifard, Kevin C. Almeroth. *Trust 2.0: Who to Believe in the Flood of Online Data?*. Proceedings of the IEEE International Conference on Computing, Networking and Communications (ICNC 2012). Hawaii, USA, Jan 2012.

- Mozhgan Tavakolifard. *Conditional Trust Adjustment and Initialization*. Proceeding of the Third IEEE International Conference on Information Privacy, Security, Risk and Trust (PASSAT2011). Boston, USA, Oct 2011.

- Mozhgan Tavakolifard, Svein J. Knapskog, Peter Herrmann. *Trust Transferability Among Similar Contexts*. Proceedings of The 4th ACM International Symposium on QoS and Security for Wireless and Mobile Networks (Q2SWinet 2008). Vancouver, Canada, Oct 2008.

- Mozhgan Tavakolifard, Pinar Ozturk. *Situation-aware Trust Judgment*. Security Engineering Techniques and Solutions for Information Systems: Management and Implementation, Boudriga, Noureddine and Mohamed Hamdi (Ed.), ISBN: 9781615208036, IGI Global, USA, Feb 2010.

- Mozhgan Tavakolifard. *Situation-aware Trust Management*. Proceedings of the third ACM conference on Recommender systems. New York, USA, Oct 2009.

- Mozhgan Tavakolifard. *Similarity-based Techniques for Trust Management*. Web Intelligence and Intelligent Agents, Zeeshan-Ul-Hassan Usmani (Ed.), ISBN: 978-953-7619-85-5, InTech, Austria, Mar 2010.

- Marie Elisabeth Gaup Moe, Mozhgan Tavakolifard, Svein J. Knapskog. *Learning Trust in Dynamic Multiagent Environments using HMMs*. Proceedings of The 13th Nordic Workshop on Secure IT Systems (NordSec). Copenhagan, Denmark, Oct 2008.

Part I

THESIS INTRODUCTION

Introduction

The deployment of a global computing infrastructure raises new and difficult security and privacy issues. Traditional security mechanisms are of questionable effectiveness in the new global computing era. Part of the reason is that no common infrastructure can be assumed to enforce any notion of correct behavior, in part because even defining a common and acceptable standard is impossible. No single authority can define and enforce rules, and therefore, online interactions cannot be governed by common rules as before. Trust-based security mechanisms have emerged as a solution, significantly expanding the scope of traditional security models. Trust enables humans to accept risks and deal with uncertainty. These new mechanisms provide weaker security guarantees, but serve greater application areas.

However, the online environments such as the web, search engines, peer-to-peer networks, and new applications built on highly complex social networks introduce several challenges in the interpretation and use of online trust and reputation systems. For example, some of these challenges have their roots in the subjective nature of feedback and some of them are related to the ease with which online identities can be attacked. Before online reputation systems will be accepted as legitimate trust solutions, a better understanding is needed of how such systems can be compromised and how these problems can be solved.

Despite the promise of online trust and reputation systems, there remain significant challenges requiring further research and commercial development. In the work presented in this thesis, we describe the critical open challenges that limit the effectiveness of trust and reputation systems and have prevented their integration into large-scale distributed applications. Integrating reliable reputation solutions will contribute tremendously towards increasing user cooperation, thereby improving the performance of these applications. Our goal is to identify challenges that weaken trust and reputation systems, and to survey prominent strategies to overcome these challenges. In addition, we present our proposals to some of these problems.

The main part of this thesis, Part II, is a collection of eight papers. Part I shows the big picture view of the material covered in the papers.

The introduction is organized as follows. First, a general background on trust and reputation systems and a multidimensional framework for categorization and comparison of them are presented in Section 1. Then, the main problems are and solutions are overviewed in Section 2. Research goals and methodology are discussed in Section 3. An overview of the thesis contributions are presented in Section 5. A

state of the art survey is outlined in Section 4. Section 6 provides short summaries of the papers and identifies their specific contirbutions. Finally, concluding remarks and list of future work directions are provided in Section 7.

1. Background

Trust and reputation systems represent a significant evolution in support for Internet services, especially in helping users decide among a growing number of choices, from which movies to rent to which data sources to trust. In this section, we first describe the concept and nature of trust by indicating what is not trust [AR04]. Trust is not simply "confidence" because trust is about what the perception of what someone is willing to do, while confidence is about what another is capable of doing. Moreover, trust is not "reliability" since a person may not have a choice on whom she relies. Trust also differs from "hope" in terms of available choices. When a potentially risky action has to be taken, a person hopes that it will result in a satisfactory outcome.

A universally accepted definition of trust is still lacking despite extensive studies from philosophers, sociologists, and psychologists. One of the most commonly accepted definitions is from the sociologist Diego Gambetta [Gam00]: "... trust (or, symmetrically, distrust) is a particular level of the subjective probability with which an agent will perform a particular action, both before [we] can monitor such action (or independently of his capacity of ever be able to monitor it) and in a context in which it affects [our] own action". As stated in this definition, some of the characteristics of trust are: subjectivity, context-dependency, and dynamicity. It is easier to determine the properties of trust than to define exactly what trust itself is. The reason for this difficulty is that trust involves a combination of interrelated cognitive and non-cognitive constructs, some of which may or may not be called on depending on the entities and situations involved.

A *trust relationship* exists between two agents when one agent has an opinion about the other agent's trustworthiness and a *recommendation* is an opinion about the trustworthiness from a third party agent. If the referrer is not known by the recommendation requester, the requester can obtain recommendations about the unknown referrer as well. There is also the scenario where a recommendation requester may carry out a network search for a particular party and the received recommendation may be the result of the request being forwarded through a number of intermediary referrers. In both scenarios, when a referrer recommends another referrer, the result is a recommendation chain.

Reputation is defined as an "expectation about an agent's behavior based on information about or observations of his past actions." Therefore, reputation can be considered a collective measure of trustworthiness (in the sense of reliability) based on the referrals or ratings from members in a community. An individual's subjective trust can be derived from a combination of received referrals and personal experience.

The basic idea in existing online trust and reputation systems is to let parties generate feedback about each other after completion of a transaction, and aggregating the feedback to derive a reputation score. The reputation score is used to assist others in deciding whether or not to trust that party in the future. Resnick et al. identifies

three phases as being fundamental to any reputation system: (i) feedback generation, (ii) feedback distribution, and (iii) feedback aggregation [RKZF00].

In addition to reputation systems, some applications can use *collaborative filtering*. Collaborative filtering techniques calculate a personalized rating estimation of an item for a user as the weighted average of previous ratings given to that item by other users. The weights are proportional to the similarity between a current user and the previous users. The user similarity can be calculated using the users' profiles or as a function of the correlation between users' ratings assigned to a common set of items. For example, if User A likes Items X and Y, and User B likes Item Y, it is likely that User B will like Item X too.

There are similarities between collaborative filtering and reputation systems. Both types of systems collect ratings from members in a community/social network. The usefulness of the former arises when the emphasis is on the content, and the latter can be used when the source of information is a more important factor. Therefore, they are complimentary decision mechanisms for use in decision systems [JIB07].

In the following, we present several dimensions for classification of the current state of the art in trust and reputation systems. This classification can serve as a basis to understand the current state of the art in trust and reputation systems, to give an overview of research areas, and to help distinguish the areas that require more work.

Information type: Trust and reputation systems can use explicit or implicit information for decision making. Examples of implicit trust information can be found in social networks such as Facebook or LinkedIn. Entities within a social network can extract some degree of trust for information gathered through friends of friends. Although neither Facebook nor LinkedIn directly implement a reputation system, members of both systems are able to utilize reputable connections through friends within the environment. Another implicit form of trust information is the use of topological analysis in online social networks to determine reputation [PSD02]. In the Google search engine, reputation is determined by the number of links that point at a page, and from where the links originate. A link originating at a page with a high reputation is likely to mean that the target page has some value.

Trust value representation: Degrees of trust are represented as either discrete or continuous levels of trust. Humans are often better able to rate performance in the form of discrete verbal statements, than they are continuous measures. This limitation is also valid for determining trust measures. The discrete levels differ from one model to the next, with some using a bounded range [CY01] and others allowing the value to extend to infinity [Mau96]. Moreover, discrete values can be binary or multinomial. A binary trust representation allows complete trust in another agent or no trust at all. Binary trust representations are simple constructs and allow unambiguous implementations. The concept is, nevertheless, rather restrictive because users are forced to choose between trusting another agent completely or not at all. The ability to handle degrees of trust in multinomial form [JIB07, CY01] allows users to proceed in situations where the amount of trust in another agent is not complete, but sufficient for the situation concerned. The disadvantage of discrete measures is that they do not easily lend themselves to sound computational principles. Instead, heuristic mechanisms

like look-up tables must be used. On the other hand, continuous values [Mar94], represented as real numbers, are modeled as either objective or subjective probability values. The objective probabilities represent purely syntactic forms of trust values, *e.g.*, the beliefs of the agent does not influence the value, and subjective probabilities are intuitive "likelihood" measurements given by the agent depending on its current beliefs. All in all, it is better to maintain reputation values as multiple component scores. Applying different functions to the scores allows a rating best suited for the given situation to be calculated. Many proposed systems suggest maintaining multiple statistics about each user. For example, keeping separate ratings on a user's likelihood to defect on a transaction (its "trustworthiness") or user's likelihood to recommend malicious users (its "reliability" as a referral) [GKRT04].

Some proposals divide the span of trust into strata and assign qualitative labels to them [AR04]. For example, the stratification is given as the set of Very Trustworthy, Trustworthy, Untrustworthy, and Very Untrustworthy [AR04]. The use of strata with qualitative labels may initially be considered a good solution to the problem of subjectivity because it seems to provide a clear semantics and avoids the ambiguity associated with numerical values. Nevertheless, in order for it to have the claimed effect, a qualitative label such as "trustworthy" should hold the same meaning for one person as it does for another. This assumption is not necessarily the case because persons with different personality cultures may associate the same experience with different strata. For example, based on her own perception of trust, what is viewed by someone as "very trustworthy" may be judged as only "trustworthy" by another person. Previous work either only considered the positive values of trust or ignorance (absence of trust or no opinion about the trustworthiness) or considered distrust as well [Mar94].

Network architecture: The network architecture determines how feedback and reputation scores are communicated between participants in a reputation system. The two main types are *centralized* (or hierarchical) and *distributed* (or peer-based) architectures. In a centralized reputation system, a central authority (reputation center) collects feedback about a given participant from other members in the community who have had direct experience with that participant. The central authority derives a publicly available reputation score. Centralized structures work well within closed networks or where decentralized approaches are not suitable for management and control purposes. On the other hand, in a distributed reputation system [BFL96, CY01], each participant simply records an opinion about each experience with other parties and provides the information on-demand. Any user can compute a reputation score based on the received feedback from others and his/her own direct experiences.

Algorithm: A reputation system uses a specific method (*e.g.*, averaging, probabilistic-based or belief-based) to compute reputation values based on the collection of feedback from others. Some of the various methods for computing reputation and trust measures include.

1) Rank ordering: This method has no explicit reputation score and acts as an implicit indicator of reputation. For instance, in Slashdot, an online discussion board,

readers rate posted comments and postings are prioritized or filtered according to the ratings they receive from readers.

2) Simple summation or average of ratings: This method is the simplest form of computing reputation scores. The score is the sum of the number of positive ratings and negative ratings, for example, positive scores minus negative scores (*e.g.*, eBay) or the average (*e.g.*, Epinions and Amazon).

3) Probabilistic models: The reputation score is computed by updating Probability Density Functions (PDFs). The updated reputation score is computed as a combination of the previous reputation score and the new rating.

4) Fuzzy models: These methods represent trust and reputation as linguistically fuzzy concepts, where membership functions describe to what degree an agent can be described as trustworthy or not. Fuzzy logic provides rules for reasoning with fuzzy measures of this type.

5) Flow models: A participant's reputation increases as a function of incoming flow, and decreases as a function of outgoing flow (e.g., Google's PageRank and Advogato). In the case of Google, many hyperlinks to a web page contribute to increased PageRank whereas many hyperlinks from a web page contributes to a decreased PageRank for that web page.

6) Game theoretical models: Problematic social situations can be described as trust games with two players and two periods of play. A Trust Game is a one-sided Prisoners Dilemma Game. The restrictiveness of the social conditions under which problematic social situations have to be solved can be reduced by adding the notion of reputation (the possibility of obtaining or spreading information about a trustee's trustworthiness) and third parties. This can be explained by the fact that the principal effect of information from third parties is to reduce uncertainty about the behavior of the trustee.

7) Stochastic models: Events are modeled by Markov decision processes and reputation is aggregated using stochastic system theory [RCP05].

8) Belief models: Dempster-Shafer theory of evidence is an extension to probability theory with the advantage of being able to model uncertainty. It is a widely used model which provides the means for approximate reasoning under uncertainty. According to it, there is no direct relationship between a hypothesis and its negation and as a result the summation of probabilities of atomic elements may not necessarily result in a value of one. In this case, the remaining probability is interpreted as a state of uncertainty [JIB07].

Table 1. A comparison of existing online trust and reputation systems across several dimensions.

Model	Info type	Value representation	Architecture	Algorithm	Info source	Context awareness	Parameters
FIRE [HJS06]	E	Continuous	D	Custom-designed	Direct experiences, Referrals recommendation, Certifications	S	Role-based trust, Time, Reliability, Credibility
Confidant [RMVS05]	E	Real number in the range [0, 1]	D	Custom-designed	Direct experiences, Referrals recommendation	S	Time
Yu and Singh [YS02]	E	Real number in the range [0, 1]	D	Belief-based approach	Direct experiences, Referrals recommendation	S	Reilability, Time
TRAVOS [PTJL05]	E	Binary ratings and reputation value as a real number in [0, 1]	D	Probabilistic approach (Bayesian)	Direct experiences, Referrals recommendation	S	Time, Reliability of referrals
Conner et al. [CIM+09]	E	Continuous	D	Custom-designed	Direct experiences	M	
EigenTrust [KSGM03]	E	Simple summation	D	Flow model	Direct experiences, Referrals recommendation	S	
Gupta et al. [GJA03]	E	Continuous	C	Custom-designed	Direct experiences	M	Time
PeerTrust [XL04]	E	Continuous	D	Custom-desinged (normalization)	Direct experiences, Referrals recommendation	S	Context compatibility, reliability of referrals, Time
H-Trust [ZL09]	E	Bounded continuous	D	Custom-designed (inspired by H-index)		S	Credibility of referrals,
REGRET [SS01]	E	[0, 1]	D	Fuzzy approach	Direct experiences, Referrals recommendation	M	Criteria compatibility, Reliability of referrals, Time
HISTOS [ZMM99]	E	Continuous	D	?	Direct experiences, Referrals recommendation	S	Reliability of referrals, time
RATEWeb [MB09]	E	Continuous	D	Custom-designed (weighted-avarage)	Direct experiences, Referrals recommendation	M	Criteria compatibility, context compatibility, credibility of referrals, time
P-Grid [ACMD+03]	E	Binary	D	Custom-designed	Direct experiences, referral's recommendation	S	
PowerTrust [RY07]	E	Binary	D	Probabilistic-based approach (Bayesian method)	Direct experiences	S	Time
PRIDE	E	Discrete and unbounded	D	Custom-designed (certification)	Direct experiences	S	
TrustMe [SL03]	E	Continuous	D	Custom-designed	Direct experiences, referral's recommendation	S	Time
XRep [DdVP+02]	E	Binary, Discrete/Continuous	D	Custom-designed	Referral's recommendation	S	
Amazon	E	1-5 stars	C	Average of ratings	Direct experiences	S	
eBay	E	-1,0,1	C	Summation	Direct experience	S	
SlashDot	I	Strata & continuous	C	Rank ordering	Direct experiences	S	Credibility of feedback
ePinions	I	Strata	C	Average	Direct experiences	S	Credibility of feedback

E=Explicit, I=Implicit, D=Decentralized, C=Centralized, S=Single, M=Multiple

9) Semantic web and ontologies: This is a logical approach for trust formalization and mainly focuses on trust's semantic structure and its logical conditions and effects. As opposed to other approaches that focus on the uncertainty of trust, trust quantification, trust dynamics, and trust computings models and algorithms. The semantics of trust relationships are modeled using ontologies [GPH03].

10) Spread activation networks: This is an example of a cognitive science approach. Spread activation models simulate human comprehension through semantic memory, and are commonly described as "models of retrieval from long term memory in which activation subdivides among paths emanating from an activated mental representation" [ZL05].

11) Social network measures: These approaches attempt to find an answer to the question: in which way does a truster's level of trust in a trustee depend on his "local" network position and on the global network structure? In other words, they evaluate the effects of density, outdegree centrality, and centralization on the level of trust a trustor can have in a trustee [Bus98].

12) Custom-desinged models: In these models, trust values are calculated from handcrafted formula to yield the desired results. The flexibility of these approaches enables trust and reputation systems to define a composite trust metric to aggregate the essential parameters and factors they have been considered in their models. Basically, they include credibility of witnesses and time or a recency factor as the main variables. However, some advanced models may use other important variables such as the transitivity rate and context and criteria similarity into account as well.

Information source: The majority of trust models consider two types of knowledge in estimating the trustworthiness of a trustee in an interaction: direct experiences and referral's recommendations (or witness observations). Personal experience typically carries more weight than second hand recommendations or reputation, but in the absence of personal experience, trust often has to be based on recommendations from others. Furthermore, some trust and reputation systems designed particular information components for the situations when neither of these information sources is available. For the FIRE model [HJS06] uses *certifications* by target members. The other information source is called *role-based trust* [HJS06, SS05], in which agents trust each other based on the predefined roles and relationships that exist among them.

Context awareness: A single-context trust and reputation model is designed to associate a single trust value per partner without taking into account the context. These systems entail information being collected from a single method and being interpreted in a predefined way. This means all of the information collected about an entity is related to one explicit aspect of that entity's actions A multi-context model has the mechanisms to deal with several contexts at a time maintaining different trust values associated to these contexts for a single partner [CIM$^+$09]. Multi-context reputation systems can take advantage of numerous sources to gather information, or collect the information such that it can be used from different perspectives.

Parameters: This dimension addresses crucial parameters which may increase the accuracy of the expected reputation value.

1) Time: The time parameter is an essential parameter in any reputation calculation and indicates the recency and freshness of information. Older information should have less influence in the calculation.

2) Credibility of referrals: In a recommendation chain, recommendations from known referrals who already have had interaction with the requested party should have more weight as first-hand recommendations than those who are known but have not had any previous interactions with the requested party or those who are unknown.

3) Reliability of referrals: It is also important to consider the reliability and honesty of the referrals, even for well known referrals by the requesting party, when their recommendations are used to be able to calculate the confidence level of the generated recommendations and to alleviate the effect of dishonest information providers and spurious ratings. This can be measured by assessing the trend line and behavior of the referrer in the time interval [SS05].

4) Context compatibility: If the information used for calculation has not been generated in exactly the same context as the decision making, then information should be weighted based on the similarity between the current context and the context of the information in order to determine to what extent the received information should be taken into account.

5) Criteria compatibility: This factor determines the similarity between the criteria used to evaluate the outcome of an interaction.

Table 1 provides a comparison of some of the exiting trust and reputation systems against the aforementioned dimensions. Table 1 highlights the differences between the selected trust and reputation systems and compares them across the dimensions of the framework. Table 1 shows that some of the systems address a wider range of dimensions than others. This fact may not necessarily imply better quality and applicability of such systems. Instead, one should consider the context in which these systems are employed and evaluate how well they accomplish the goals and requirements of that particular environment.

2. Challenges

In this section, we identify challenges that weaken trust and reputation systems. The prominent strategies to overcome these challenges are also surveyed. We consider each phase of operation for such systems, namely: feedback generation, feedback distribution, and feedback aggregation. Each of these components needs safeguarding against a variety of adversarial threats. As a case in point, reliability in terms of reputation accuracy is a critical requirement for the aggregation component. This section, therefore, studies the extent to which existing research efforts counter these threats.

2.1 Feedback Generation

One of the most important tasks in a reputation system is generating accurate and representative feedback. Not only must a qualitative, opinion-based process be

reduced to quantitative facts, but also users will sometimes try to game the system. We have identified the following challenges.

Low incentive for providing feedback: There are two main reasons for this problem [Del03]. First, feedback constitutes a public good and once available, everyone can benefit, yet the provider benefits very little. Second, providing feedback presupposes that the provider will assume the risks of the transaction, risks that are typically higher for new products or users. To solve this problem, some models propose payments and financial rewards for honest feedback [JF03, MRZ02]. For example, Epinions provides incentives for reviewers, whereby they can earn money based on general use of reviews by consumers. Bizrate, a customer certified merchant, gives discounts as an incentive to fill out surveys. An alternative approach is to build incentives into the feedback aggregation equation. This goal can be accomplished by providing a small increase in reputation whenever a user provides reputation feedback to others [Mal01]. For instance, Amazon gives some members status as a top reviewer. Another approach is to use implicit feedback, where users' actions are recorded and the feedback is inferred from the recorded data [Del03]. For example, an assumption in Google's reputation score is that if enough people consider a page to be important enough to place links to it, and if the pointing pages are "reputable" themselves, then the information contained on the target page is likely to be valuable.

Bias toward positive feedback: It can be difficult to elicit negative feedback because of reciprocity. For example, the observed ratings on eBay are surprisingly positive. Of all ratings provided, less than 1% are negative, less than 0.5% are neutral and about 99% are positive [ZR02]. Providing anonymity may help to avoid this problem. It was also found that there is a high correlation between buyer and seller ratings, suggesting that there is a degree of reciprocation of positive ratings and retaliation for negative ratings. A possible remedy could be to not let sellers rate buyers.

Initialization and cold-start problem: Bootstrapping a reputation mechanism is not trivial. In many systems, users start with a neutral reputation. Newcomers are offered only a limited number of resources and so struggle initially to build their reputations. As other users in the system tend to interact with high reputable users, the chance of a new user being selected for interaction is generally rare (*e.g.*, in eBay, many users will not deal with individuals with a low reputation score [Mal01]). Hence, it is hard for a new user to raise her reputation score. This challenge may be a barrier to entry into the marketplace or community. Solutions include taking into consideration the interconnections among reputation systems and social networks. For example, the location of a given member of a community within a social network can be used to infer some properties about her degree of expertise, *i.e.*, her reputation [GH04].

Subjectivity: Feedback information is strongly influenced by subjectivity factors such as the feedback provider's taste and cultural background. One solution based on collaborative filtering, is to personalize the feedback by weighting it in inverse proportion to "taste distance" between the provider and the receiver of the feedback. Therefore, it will be easier for the receiver of the feedback to interpret it because it consists of opinions from like-minded people [Del03].

False feedback: When users incorrectly report their feedback, it creates errors in the system. We categorize the different forms of false feedback as follows:

1) Dishonest and unfair reports: This problem happens because of the low cost of submitting online feedback and the relative anonymity of the raters. Unfair ratings can be excluded using their statistical properties [CS01, Del00] or by using the rater's reputation [BLB04, CDdV$^+$02]. Slashdot addresses this issue by using the judgment of longstanding users as *a priori* trusted agents.

2) Collusion: Collusion occurs when two of more peers collectively boost each others reputations or conspire against one or more peers in the network. Dellarocas identifies three types of collusion misbehavior [Del03]:

a) *Ballot stuffing*: Parties engage in many fake transactions to artificially inflate their reputations and ratings. This problem is solved in eBay by only allowing participants to rate each other after the completion of a transaction, and charging a fee for each transaction. In the Sporas model [ZMM99], when a user rates another more than once, only the most recent rating is considered.

b) *Bad-mouthing*: This problem occurs when a malicious collective conspires against one or more users in the community and hurt their reputation by assigning unfairly low ratings to them.

c) *Positive and negative discrimination*: Discriminatory behavior can occur both when providing services and when providing feedback. A seller can, for example, provide good quality to all buyers except one in particular. Feedback about that particular seller will indicate that she is trustworthy except for the feedback from the victim buyer. Filtering techniques will give false positives, *i.e.*, judge the buyer victim unfairly in such situations. Only systems that are able to recognize the victim buyer as trustworthy would be able to handle this situation.

Cheap pseudonyms: In online environments where new identities may be created with minimal cost, these multiple identities create several problems, including the following:

1) Sybil-based collusion: Malicious entities may acquire multiple identities for the sole purpose of creating phantom feedback in the system. Proposed solutions to deal with Sybil attacks fall into centralized and decentralized approaches. In a centralized approach, a central authority issues and verifies credentials unique to each entity. To increase the cost of obtaining multiple identities, the central authority may require monetary or computational payment for each identity. In decentralized approaches, some proposed solutions include binding a unique identifier, such as IP addresses, to public keys or using network coordinates to detect nodes with multiple identities (*e.g.*, Kuro5hin allows only one rating from any single IP address). Other solutions take advantage of social knowledge to propagate reputations originating from trusted sources along the edges of a "web of trust". Thus, the effect of the attackers will be limited based on the expense of requiring social interactions [HZNR09].

2) Re-entry problem or churn attacks: In online communities, it is usually easy for members to disappear and re-register under a completely different online identity with zero or very low cost (*e.g.*, eBay). Models that treat unknown users and disreputable ones differently [Mar94, Gri05, TC04, TC04] are vulnerable to this problem, however, models that penalize newcomers are resistant [ZMM99]. There are two classes of approaches to this issue [Del03]: either making it more difficult to change online identities (*e.g.*, by using cryptographic authentication techniques), or making it unprofitable to exit and re-enter with a new identity (e.g., by imposing an upfront cost to each new entrant such as a fee or an implicit cost of having to go through an initial reputation-building phase with low or negative profits).

2.2 Feedback Distribution

Assuming reputation information can be collected and processed correctly and without malicious influence, the next challenge is to get the feedback to those who need it to make their decisions. Some of the challenges in this part of the process include:

Reputation lag problem: There is usually a time lag between an instance of a transaction and the corresponding effect on the reputation score (*e.g.*, in eBay, the buyer pays before the seller ships the item). A user has the opportunity to make use of this time lag to provide a large number of low quality services over a short period before the reputation score suffers any significant degradation [KC06]. Further, the re-entry problem can be combined with this problem in a way that a seller may re-enter the market each time a buyer learns of a dishonest seller. In this way, a seller can repeatedly take advantage of reputation lag.

Lack of portability between systems: The limited distribution of feedback limits its effectiveness. As a solution, Amazon allowed users to import their ratings from eBay [RKZF00]. Obviously only users with good reputations will take advantage of this feature, thereby diluting the value of the scores.

Inability to filter or search: Online communities run into several information overload problems due to the sheer size of many of these sites. The ability to filter and search by reputation would greatly improve their usability [Mal01].

Categorization: A reputation score is too general in most systems (*e.g.*, eBay) and there is little ability to use reputation scores in different categories. Reputation categories could enhance systems by providing better granularity. For example, a user might have a good reputation in one area (*e.g.*, quality of products) and a bad reputation in another area (*e.g.*, on-time delivery). This concept could work in conjunction with a search and filtering feature [Mal01].

2.3 Feedback Aggregation

Assuming reputation information can be collected and processed correctly and then delivered to a user, there is still the challenge in aggregating and displaying feedback so that it is truly useful in influencing future decisions about whom to trust. Some of the challenges in this part of the process include:

Inaccurate equations: Simple reputation schemes such as eBay's reputation score (*i.e.*, the sum of positive ratings minus the sum of negative ratings) can be misleading. For example, a user in eBay with 100 positive and 10 negative ratings would have the same total reputation score as a user with 90 positive and no negative ratings; however, the former should appear less reputable. This problem results in a vulnerability caused by "increased trust by increased volume." That is, a user can increase his/her trust value by increasing his/her transaction volume, thereby hiding the fact that she frequently misbehaves.

Value imbalance problem: In many reputation models (*e.g.*, eBay), all feedback is weighted equally regardless of the transaction value. This problem encourages Sybil attacks and collusion. A user can take advantage of this property to build a good reputation by honestly executing a number of small-value trades, and then using the accumulated reputation to cheat in a very high-value transaction [Del02].

Spread of false rumors: This problem occurs when the reputation of the feedback providers is not considered. One approach to this problem is to rely on pre-trusted identities. Another approach is to employ statistical methods to build robust formulations (*e.g*, a Bayesian framework) that can be reasoned about in a precise fashion [HZNR09].

Unlimited memory: Most reputation calculation algorithms use all transactions when calculating the overall score, thus, a new user might not understand how a site functions [Mal01]. Besides, a user can perform short duration malicious attacks with little risk of negative consequences because a lengthy previous history can heavily outweigh current actions. This problem can have a large impact on the system as the malicious users will continue to have a high reputation for a substantial period of time during which the system is slow to identify the malicious behavior and unable to sufficiently lower the user's reputation [HZNR09]. Therefore, the memory should be de-emphasized in some way, though this is not easy in practice. For example, a simple cut-off function handicaps the user by providing only the most recent information. Further, new users will likely require some time to become familiar with the mores of a site and they should not be penalized for initial bad behavior if the behavior is unintentional. One solution is to give less weight to negative feedback for new users and more weight for old users. Instead of a strict cut-off, this approach leads to a gradual change in the importance of more recent feedback [Mal01].

Dependence on profit margins: Reputation effects can induce users to accept short-term losses in order to realize larger long-term gains provided that the latter exceeds the former. In other words, the remaining horizon must be long enough and the profit per transaction must exceed a threshold. This result can have at least two potential interpretations. First, reputation mechanisms are not effective in highly competitive markets. Second, prices tend to be higher in markets where trust is based on reputation than in markets with perfect information [Del03].

Time sensitivity of reputation: treating old positive behavior equal to new negative behavior may result in attackers abusing the system by using previous altruism to hide current malicious behavior. Techniques have been proposed that use more aggressive short-term history and give more weight to recent negative behavior [HZNR09].

Denial of service attacks: Attackers may seek to subvert the mechanisms underlying the reputation system in centralized architectures, causing a denial of service. For instance, attackers can attempt to cause the central entity to become overloaded by attacking its network or computational resources. Attackers are then able to perform malicious actions without their negative reputation being known or without being punished for their negative behavior. Distributed architectures with enough redundancy are often less vulnerable to this attack. Techniques to cope with denial of service attacks are similar with the ones used by many routing protocols and include: use of acknowledgments, multi-path dissemination, gossip mechanisms, and forward error correction codes [HZNR09].

Playbooks: A playbook is a sequence of actions that maximizes profit of a participant according to certain criteria. A typical example is to act honestly and provide quality services over a period to gain a high reputation score, and then to subsequently milk the high reputation score by providing low quality services at a low production cost [KC06].

Exit problem: Since there is no incentive for a party leaving a system to maintain a good reputation, the entire accumulated reputation can be used for cheating (*e.g.*, in eBay). One solution to this problem is to introduce community membership rules that elicit good behavior. For example, online communities can levy a sufficiently high entrance fee that is refundable subject to maintaining a good reputation upon exit or reputation scores can be viewed as assets that can be bought and sold in a market [Del03].

While an innumerable variety of attacks can be devised by malicious peers, our above discussion identifies attack strategies most commonly observed in reputation systems.

3. Research goals and methodologies

For this thesis, we are following a mixed methods approach where different theories are brought together for the explicit purpose of solving particular practical problems that online trust and reputation systems design is faced with. The thesis is largely a theoretical examination of the issues at hand, but it proposes practical consequences as well. Some quantitative evaluation of particular aspects is also performed.

Due to the obvious time and spatial constraints, we cannot address all of the challenges that we face when designing online trust and reputation systems. In reducing the scope of this thesis, we have focused on the following problems:

- The problem of subjectivity in explicit ratings
- The initialization problem
- The categorization problem
- The lack of portability problem
- The Cold-start problem
- The time sensitivity problem

In our approach, we are mainly motivated by what humans do in traditional trust and reputation systems such as analogical, abductive, and inductive reasoning. The main idea is to consider contextual information, as a special kind of implicit feedback, in trust computations and the goal is to bring additional knowledge to the reasoning process by use of available auxiliary data or Meta-data (contextual data). Context qualifies a trust opinion, describing what the truster's belief in another's trustworthiness is really about. The introduction of context as an explicit notion may improve problem solving efficiency by better grounding what knowledge is used in decision making in the real world. With more information and better context, trust and reputation systems can help users to make more well-informed decisions. Our work improves the utility and accuracy of trust management systems by proposing methods on how to use *contextual information*.

This research is based on literature studies, group discussions and seminars, analytical modeling, implementations and evaluations using real data or simulation results. The models have been published and presented at international conferences and journals.

4. State of the Art

This section presents a survey of situation-aware approaches to trust management. The advantages and importance of using contextual information are also recognized by other researchers. For example, Neisse et al. [NWvS06] attempted to reduce the complexity in management of trust relationships. Neisse et al. [NWvSL07] and Gray et al. [GCJ03] focused on the improvement of the trust recommendation process. Holtmanns and Yan [HY06] investigated how to infer trust information in context hierarchies. Rehak et al. [RGPB06] improved the performance of trust management systems. They also provided protection against changes of identity and first time offenders in trust management systems. Bagheri and Ghorbani [BG06], Bagheri et al. [BBEZG08], and Gray et al. [GCJ03] provided methods that correlate trust information among various contexts.

In addition to differences in the main focuses and motivations, there have been differences also in the representation of the context information, as shown in Table 2.

Following, we shortly review the main contributions of these work:

Neisse et al. [NWvS06] proposed the idea of using the abstraction of context-aware domains to reduce the complexity in the management of trust relationships. In a large context-aware system, with thousand of components and users, trust relationships can not be associated with individual entities, as this can easily become unmanageable. Examples of context-aware management domain definitions are "nearby persons", "Personal devices", and "Working colleagues". The idea is to provide mechanisms to define and infer the trust degree of an entity based on the context information provided about that entity. According to their other work [NWvSL07], it is also possible to use context information to improve the recommendation process (i.e., to determine from whom to request recommendation). This will allow anonymous and still useful recommendation exchange.

Table 2. Context representation methods.

Context Representation Method	**Model**
Context-aware domains	Neisse et al., [NWvS06, NWvSL07]
Intensional Programming	Wan & Alagar [WA08]
Multi-dimensional goals	Gujral et al., [GDFB06]
Clustering	Rehak et al., [RGPB06]
Graph	Holtmanns & Yan, [HY06] Bagheri & Ghorbani [BG06]
Bayesian network	Bagheri et al. [BBEZG08]
Ontologies	Golbeck et al. [GPH03] Huang & Fox [HF06] Toivonen and Denker [TD04] Tavakolifard et al. [TKH08b, TKH08a]
Trust attributes	Caballero et al. [CBGS06a] Uddin et al., [UZA08] Gray et al., [GCJ03]
Case-based reasoning	Tavakolifard et al., [THÖ09]

Holtmanns and Yan [HY06] noted that context can often be structured hierarchically. For example, if you trust someone to drive your car, then you would most likely give him also your car keys or the keys to the garage. Therefore, it is necessary to identify possible hierarchical structures between different contexts in our model to be able to infer trust information from one into the other. In this work, entities that can be applications, other users or agents that act on behalf of users are structured into a context-based trust graph. Positions in this graph indicate the context-based trust level and changes based on events or over time. The structure of the trust graph reflects a certain hierarchy.

Alagar et al. [WA08] investigated the intensional programming paradigm for agents communication by introducing context as a first class object in the intensional programming language "Lucid". Intensional programming is a powerful and expressive paradigm based on Intensional Logic. Intensional logic is a branch of mathematical logic used to precisely describe context-dependent entities. In this paper definitions, syntax, and operators for context, and an operational semantic for evaluating expressions in extended Lucid are given. It is demonstrated that the extended Lucid language, called Agent Intensional Programming Language (AIPL), has the generality and the expressiveness for being an Agent Communication Language (ACL). Based on this work a context-specific trust model for multi-agent systems is introduced. The explicit introduction of context into the computation of trust, annotation of trust policies with context conditions, and definition of delegation through related contexts are some of the results given in this paper.

The context issue has also been viewed as multi-dimensional trust modeling for agents when goal requirements are multi-dimensional [GDFB06]. An agent's reward is determined by goal requirements and behavioral constraints of potential partners (e.g., quality, timeliness, availability, and cost).

Rehak et al. [RGPB06] defined a set of reference contexts in a metric space and associated truthfulness data with it. These data can be updated and queried with weight that decreases with distance between the current situation and the reference context. The model uses *Leader-Follower* clustering to identify the reference contexts to be representative of the data. The advantage of this clustering method is that it allows an on-line approach without pre-specifying the number of expected clusters, and requires only a single parameter as input. The biggest disadvantage is that it may easily under or over estimate the number of clusters. In an empirical test, it is shown that context-aware models easily outperform general trust models when the situation has an impact on partner trustfulness and that their performance and efficiency is comparable with general trust models where the trustfulness is independent of the situation. In this work, two advanced uses of context for multiagent trust modeling are proposed: (i) policy/norm learning at runtime by analyzing data regarding the performance of different agents in similar situations (e.g., when all agents fail in a certain situation, they may agree to introduce a policy that specifically prohibits such actions) (ii) reasoning based on uncertain identities by decomposing the single identity dimension into an identity subspace, where each agent is defined by one or more crucial properties. With this modification, the trust model can make predictions about the performance of agents by exploiting data characterizing a similar agent's performance in the past. The main advantages are that the extended model learns faster and once the new agent is categorized, its performance can be predicted. This is also a clear advantage in ad-hoc environments, where there is no agent platform to enforce unique identity.

Based on this model, Rehak et al. [RP07] concluded that the extension of a trust model with a context representation environment can be extended to encompass a more open situation (e.g., a wireless sensor network that is hard to identify and where the barriers of entry are quite low). In such environments it is not needed to have assumptions like: (i) proven identity, (ii) repetitive interactions, or (iii) similar trusting situations. The fact that two agents with presumably distinct identities can be considered identical by a context-sensitive trust model may provide protection against changes of identities as well. This approach is also effective against first time offenders; we can obtain a model with inductive properties, which is able to estimate the performance of new entrants using the experience with the similar partners in the past.

Golbeck et al. [GPH03] proposed an ontology for trust. Golbeck and Hendler [GH04] considered a model using context-specific reputation by assigning numeric ratings to different types of relations based on the context of the analysis. Toivonen and Denker [TD04] specified rules describing how certain context-sensitive information (trust factors) reduces or enhances the trust value for this trust ontology. The authors also argue that a specific advantage of making the context explicit in

message exchanges is that this information can be used in trust policies. For example, a policy can state that news information related to a particular location is to be trusted more if the reporting entity was at the location at the time when the event occurred. In this sense, policies define how to process context information to derive trustworthiness assertions. However, they have not answered how the context-sensitive trust factor should be determined. In addition, they have not addressed either the fact that the trust value might be different for different aspects of trust.

In the work by Huang and Fox [HF06], trust is formalized by using situation calculus in order to define a trust ontology. Situation calculus is a logic language specifically designed for representing dynamically changing worlds. It works in the following way: the changing world is represented by a set of fluents. A fluent is a property (of the world) whose value is dependent on situations. In other words, a fluent dynamically changes when the situation changes. The situation, in turn, changes when an action is performed by agent(s) in the world. Trust and context are represented as fluents.

Toivonen et al., [TLU06] used contextual information (context attributes) to adjust the output of a trust determination process. Each attribute can adjust the trust value positively or negatively according to a specified weight. For example, if t is the trust value and ω is the weight of the context property then the adjusting function can be t^{ω} for decrease or $\sqrt[\omega]{t}$ for increase. A context ontology connects the context attributes with each other in an appropriate manner, enabling the utilization of context attributes which do not exactly match the query, but are "close enough" to it. For example, the QoS properties of a network, over which some software component is downloaded, can be described in such an ontology.

In the example provided in Figure 1, we suppose that the current network (B_1) is not pre-evaluated with regard to its impact on trustworthiness. However, as its neighbors in the ontology are networks which have pre-evaluated trustworthiness values (B_2, U, and G). By using these values as well as their "semantic distance" to the current network, the resulting trustworthiness can be estimated. The semantic distance is calculated by taking into account the "upwards cotopy", that is, the distance between the currently investigated concept and a root-concept of the ontology. The upwards cotopy is calculated as the ratio between the number of shared nodes from the source node and the sink node to the root node, and the total number of nodes from the source and the sink to the root node. For example, in the case of B_1 and B_2, the numbers are $|Bluetooth, PacketSwitched, Wireless, Network| = 4$ and $|B1, B2, Bluetooth, PacketSwitched, Wireless, Network| = 6$ and the semantic distance between the source and the sink therefore is $\frac{4}{6} \approx 0.67$. If adjustment functions for B2, U, and G are $\sqrt[\omega_1]{t}$, $\sqrt[\omega_2]{t}$, and t^{ω_3} and their semantic distances to B_1 are d_1, d_2, and d_3 respectively then our estimate of adjusting function for B_1 will be $\sqrt[\omega_1 * d_1]{\sqrt[\omega_2 * d_2]{t^{(\omega_3 * d_3)}}}$.

In this work, the notion of context also has been applied to the reputations by emphasizing more the observations that have taken place under similar conditions as where the truster currently is. Two relationships have been considered between recommendations and context. First, as was the case with reputation, the contextual details at the time when the recommendation was made can be considered and compared with

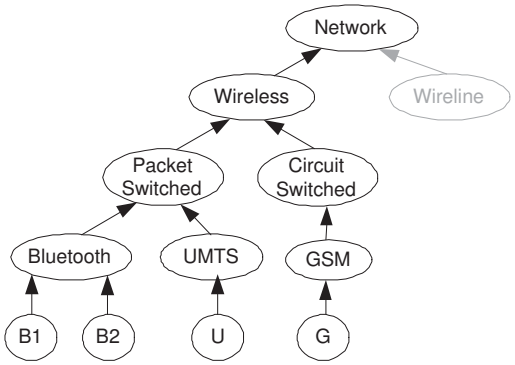

Figure 1. Concepts in the network ontology [TLU06].

the truster's current context. Note that considering this is not as straightforward as was the case with reputation, since recommendations come from others, not from the truster. Secondly, the recommendation content itself can be context-dependent.

Caballero et al., [CBGS06a] considered cases where an agent does not have enough information to produce a trust value for a given task, but she knows instead the previous partner behavior performing similar tasks. This model estimates trust using the information about similar tasks. The similarity ($D(s_1, s_2)$) between two tasks s_1 and s_2 is obtained from the comparison of the task attributes.

$$D(s_1, s_2) = 1 - \frac{1}{n}.\sum_{i=1}^{n} |s_{1_i} - s_{2_i}|$$

Where n is the number of task attributes, s_{1_i} is the $i-th$ attribute of task s_1, and s_{2_i} is the $i-th$ attribute of task s_2.

The same authors in another work [CBGS06b] obtain the similarity ($D(s_1, s_2)$) from the comparison of the task attributes in the ontology using formula below:

$$\frac{|S_1 \cap S_2|}{|S_1 \cap S_2| + \alpha(s_1, s_2)\,|S_1 \backslash S_2| + (1 - \alpha(s_1, s_2))\,|S_2 \backslash S_1|}$$

Where $0 < \alpha < 1$; S_1 and S_2 are the set of properties of concepts s_1 and s_2, respectively. Function α takes into account the depth of compared concepts in the ontology hierarchy.

Udding et al., [UZA08] proposed a model called CAT (A Context-Aware Trust Model), which uses some keywords to describe contexts. The similarity between two contexts with K_1 and K_2 as sets of keywords is calculated as $\frac{K_1 \cap K_2}{K_1 \cup K_2}$.

Bagheri and Ghorbani [BG06] proposed a framework for dynamically updating and inferring the unobserved reputation of environment participants in different contexts. This framework suggest the employment of a reputation tree structure to represent the relationship between the contexts of the environment. Reputation of a given identity in one context can be propagated to other contexts through two mechanisms,

namely: forward update and backward adjustment. This work does not mention how the reputation tree structure can be developed.

Bagheri and Ghorbani also proposed a framework for their previous proposal based on valuation networks. Global reputation is modeled as Dempster-Shafer belief functions on a Markov tree through which the relationship between various contexts of a unique environment is modeled by employing hyper-vertices of the Markov tree [BBEZG08]. Reputation of each identity in a given context is represented using a belief mass assignment function. The estimation of reputation in various contexts of the environment is performed by the employment of the message passing-based belief propagation model of the Shenoy-Shafer architecture.

Gray et al., [GCJ03] presented an initial investigation into addressing the issue of making trust-based security decisions in a given context. The authors considered several trust attributes for each context and proposed how trust is mapped across contexts based on common attributes among those contexts.

Strang and Linnho-Popien [SLP04] provided a survey of different approaches to model context for ubiquitous computing. In this work, numerous approaches are reviewed, classified relative to their core elements, and evaluated with respect to their appropriateness for ubiquitous computing. The authors concluded that the most promising assets for context modeling for ubiquitous computing environments can be found in the ontology category in comparison with other approaches like key-value models, mark-up scheme models, graphical models, object-oriented models, and logic based models. This selection is based on the six requirements dominant in pervasive environments: distributed composition, partial validation, richness and quality of information, incompleteness and ambiguity, level of formality, and applicability to existing environments.

5. Contributions

In this section, we outline our proposed solutions for some of the problems described in the previous sections. Our work improves the utility and accuracy of trust management systems by proposing methods on how to use *contextual information*.

We distinguish between *external* and *internal* context. External context is related either to the properties of the trustee or the object to be acted on (*e.g.*, information to be exchanged or something to be bought). These are the facts that exist independent of the reasoner. They are independent in the sense that they are there before and after the reasoner notices them. Internal context (*i.e.*, subjective/cognitive context), on the other hand, characterizes the mental and emotional state of the reasoner, the truster. A trust evaluation process is complicated by context in two ways: (1) trust is situation-specific (the effect of external context); a typical example is that a person may trust her financial advisor about investment analysis but does not trust the same advisor related to health-care issues, and (2) trust is person-specific (the effect of internal context); judgments of two persons on the same matter or event are often quite different.

We describe a holistic trust management approach that deals both with the situation-sensitivity of trust and the subjectivity problem. The impact of internal context

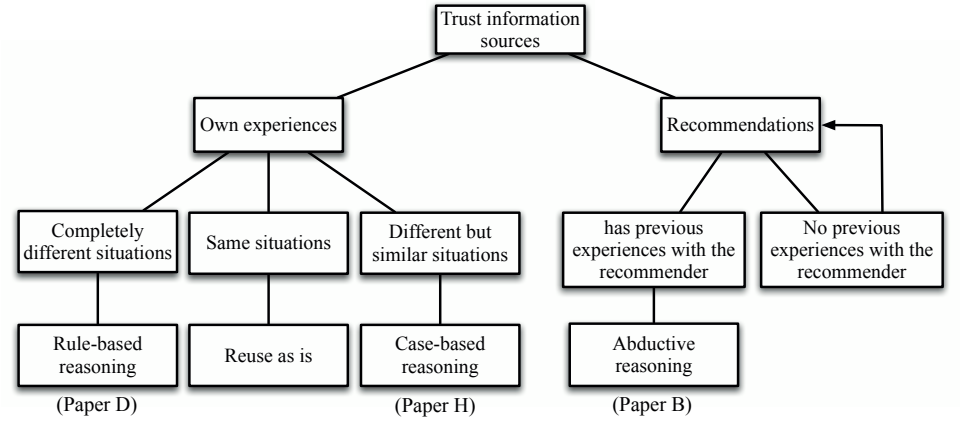

Figure 2. Different trust evaluation scenarios and their reasoning methods.

(subjectivity) and four types of external context: time, similarity, situation, and stereotypes are modeled and assessed. Our conception of different trust evaluation scenarios and the reasoning methods appropriate in each of them is illustrated in Figure 2. The two fundamental types of knowledge are acquired for trust evaluation through own experiences and from recommendations by third parties, i.e. the recommenders. If the truster previously has had interactions with the same trustee in the same situation, she can immediately use her past experiences in order to predict the outcome of the new interaction and make a decision on this basis. On the other hand, if the truster has had interactions with the trustee but in different situations, she can still use her past experiences, but should map the old and new situations and make necessary adaptations in order to draw a conclusion. We have used case-based reasoning (CBR) to handle such situation-specificity of trust. "Situation" in figure 2 refers to external context. The notion of internal context comes into play when recommendations from a third party is used to evaluate trust.

We now describe some of our work that attempts to solve the five problems identified in Section 3.

Unlimited memory and time sensitivity of reputation (Paper F): In an attempt to model the effect of time as a type of external context, we have proposed a formula for a dynamic longevity factor, $\lambda \in [0, 1]$, that make it possible to discount past ratings correctly . The longevity factor, λ, controls the rate at which past ratings are aged and discounted as a function of time. With $\lambda = 0$, past ratings are completely forgotten after a given time period. With $\lambda = 1$, past ratings are never forgotten. We propose to adjust λ after each interaction based on the similarity between the estimated and real outcome of the interaction. The higher the similarity, the larger the increase in the value of λ, and the larger the memory size (*i.e.*, time window). If the real outcome is not similar to what is expected, we are facing a change in behavior and a change in the value of λ. As a result, the size of memory for remembering the past rating/behavior

will be decreased. The amounts of increase and decrease are decided based on the application. For example, in risky applications, after a change in the behavior of the trustee, the value of λ should be decreased sharply. The initial value of λ should be zero. Only after a number of successful interactions is λ allowed to increase.

The sparsity and cold start problem (Paper G and Paper C): We consider the patterns of how individuals and groups use trust as another external context factor and groups trust as another external context factor. Based on the characteristics of how and what individuals and groups trust, we have proposed that the like-mindedness of individuals and groups can be utilized to identify other trust relationships. For instance, if one knows that, with respect to a specific property, two parties are trusted by a large number of different trusters, one can assume that the two parties have similar trust characteristics. Thus, if one has a certain degree of trust in the first party, one can safely assume a similar trustworthiness for the other party. In an attempt to provide high quality recommendations and proper initial trust values, even when no complete trust path or user profile exists, we propose `TILLIT`, a model based on a combination of trust inferences and user similarities. Similarity is derived from the structure of a trust graph and users' trust behavior as opposed to other collaborative-filtering-based approaches that use ratings of items or users' profiles. We describe an algorithm realizing the approach based on such a combination of trust inferences and user similarities, and validate the algorithm using a large, real-world data set.

Categorization (Paper D and Paper H): It is possible to draw information from feedback that is generated in a variety of situations, but situations in such a way that the feedback can be useful in other situations. For example, in online auctions, there are common factors between buying and selling activities that affect trust formation. Therefore, the feedback about a user as a buyer might be useful in calculating the reputation of the same user as a seller. We present a knowledge intensive and model-based, case-based reasoning framework that supports a system that can infer such information. The suggested method augments other work in environments where information is typically sparse (*e.g.*, there are many buyers and sellers, and it is unlikely that there is a previous transaction on which to calculate an accurate trust value). A trust rating can be calculated by inferring the lack of relationship information using other situational conditions. Such a solution allows better support for situation-aware trust and reputation management.

The CBR technique is particularly useful for tasks that are experience-intensive, involve plausible (i.e. plausible but not complete) reasoning and have incomplete rules to apply. The fundamental principle of the CBR technique is similar to that of the human analogical reasoning process which employs solutions of past problems to solve current ones. The reasoning process is generally composed of three stages: remembering, reusing, and learning. Remembering is the case-retrieval process, which recalls relevant and useful past cases. In the reusing step, the CBR system uses the recalled cases to find an effective solution to the current problem. Learning is the process of case base enhancement. At the end of each problem-solving session the new case and the problem-solving experiences are incorporated into the casebase [JHS99]. In our approach, the role of context (external) is to generate candidate cases. This

hypothesis-generation activity of the reasoner can be thought of as an instance of "cued recall" in cognitive psychology terminology. Context has been shown to have major influences on remembering cases and its inclusion in case-based problem solving empowers the case-based approach. The strong dependence between the context and a powerful memory-retrieval arise most probably from the role context plays in similarity assessment of two cases (i.e., the new and a past case). We proposed a rule-based reasoning model (far left in figure 2) for decision making when the truster does not have own similar past experiences or available recommendations about the trustee either (this we have addressed in a paper under preparation). The trust judgment then resorts to a set of domain-specific association rules.

Our framework can be coupled with existing models to make them situation-aware. Our model uses the underlying model of trust and reputation management to transfer information between situations and can also be used to transfer information from one system to another to provide more portability. We validate the proposed framework for the Subjective Logic Model [JIB07] and evaluate it by conducting experiments on a large, real-world data set.

Our second motivation for this work is trust transitivity. Trust is not always transitive in real life. For example, the fact that *A* trusts *B* to fix her car and *B* trusts *C* to look after his child does not imply that *A* trusts *C* to fix a car, or for child care. However, under certain semantic constraints, trust can be transitive and a trust referral system can be used to derive transitive trust. The semantic constraint is that the subject of trust should be the same along the entire path, for example all trust subjects should be "a good car mechanic" or "looking after a child". However, trust relations with the same subject are not always available. This constraint is relaxed in our work by introducing the notion of situation. We suggest that trust situations along a transitive trust path can be different but similar to each other. For instance, trust situations can be "to be a good car mechanic" or "to be a good motor mechanic". In this way, we are able to use trust information from available similar situations.

Initialization and low incentive for providing feedback (Paper E): When a user first comes into a system, there is little information available to use to build a trust recommendation. Further, gathering such information is difficult when there is little incentive to provide feedback. We categorize the decision making process with respect to these two factors based on the familiarity of the truster with the situation and the trustee. Different combinations of incomplete knowledge are:

1) Unfamiliar situation, familiar trustee: If the truster has had previous interactions with the trustee or other similar trustees, but in different situations, she can still use her past experiences. But in these new situations, the truster needs to map the old and new situations and make the necessary adaptations in order to draw a conclusion. As we mentioned earlier, case-based reasoning is used to handle such situation-specificity of trust.

2) Familiar situation, unfamiliar trustee: If the truster has had previous interactions in the same situations with other trustees (*i.e.*, a stereotype of the trustee), the trust judgment then resorts to a set of domain-specific association rules. We propose a rule-based reasoning algorithm to handle this situation. Past trustees are grouped

based on a common attribute with the current trustee and the general trustworthiness of the group can be summarized. Then, an opinion about those trustees as a group is formed and the current trustee is included in that group. In this way, the opinion about the group is effectively transformed into an opinion about the prospect.

3) Unfamiliar situation, unfamiliar trustee: If there is no situational or trustee information, the trust model uses a default trust value since there is no information to be used for the initialization of trust.

Subjective and unfair ratings (Paper B): Our approach to modeling the impact of subjectivity is based on the idea that a feedback provider's judgment method can be inferred and the target entity can be (re-)evaluated according to the value system of the receiver of the feedback. The judgment method is a function that maps an attribute value (*e.g.*, delivery time = late) to the value the feedback provider attached to that attribute (*e.g.*, unsatisfied). Thus, the receiver of the feedback will be able to translate (*i.e.*, eliminating subjectivity) for subsequent feedback from this particular user based on what she has learned. This method of extracting judgment information involves abductive reasoning.

Abductive inference is typically relied upon in imperfect domains, i.e., in the face of incomplete or inconsistent information as well as in cases where the domain does not provide a strong theory. Abductive reasoning in general is reasoning from consequences to antecedents and describes the process of discovering hypotheses (i.e., antecedent), and assesses the likelihood that a specific hypothesis entails a given as conclusion (i.e., consequence). Inference of 'it must have rained', upon seeing the grass wet is based on experiences: "when it rains, the grass gets wet. The grass is wet, then it must have rained". However, if there is a person near her car and a hose is on the grass, the inference would lend towards "when a person washes her car, the grass gets wet". This person may have washed her car (since she is near her car and there is a hose on the grass" [Har68].

We envisage that the truster can infer the *judgment method* of the recommender by observing the recommender's ratings and corresponding trustee's properties. For example, in cases where a recommender is known to consistently bias its ratings (e.g. always exaggerating positively or negatively, or always reporting the opposite of what it thinks), it is in fact possible to "re-interpret" the ratings. This can be done by extraction of the conditional relation between the trustee's properties as antecedents and the recommenders' ratings as consequences from the history of interactions[1]. Based on this information, in the future, the truster will be able to translate a new rating provided by the recommender into the actual properties of the trustee by employing abductive reasoning. Figure 1 (a) shows the trust value computation by the truster without considering the subjective difference. The recommender sends a rating about the trustee to the truster based on his own observations of the trustee's properties and the truster simply uses this rating as is in her own trust model (decision making

[1]The history contains two kinds of information for each interaction: the rating that the truster received from the recommender regarding the trustee before an interaction and the truster's own observation of the trustee's properties after the completion of the interaction.

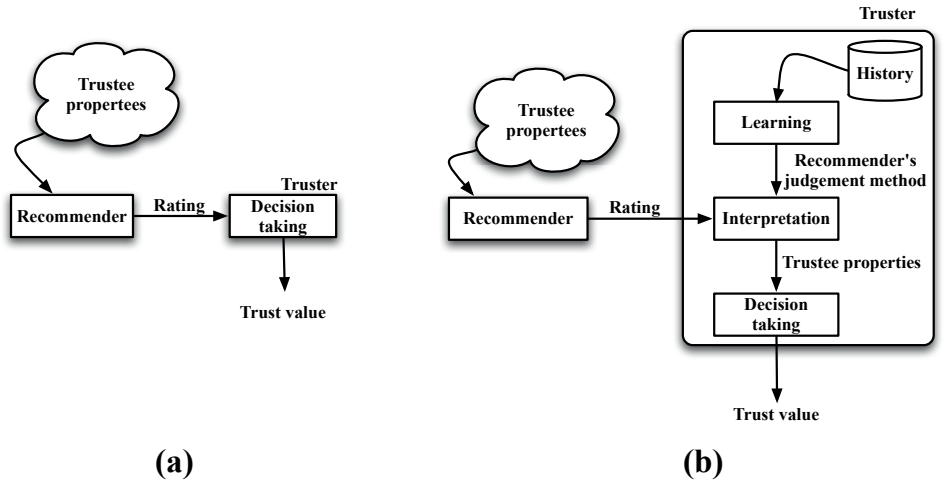

Figure 3. Trust value computation without (a) and with (b) subjectivity consideration.

model) as if she has generated this rating herself. Figure 1 (b) shows the same process, however, this time the truster considers the subjective differences and re-interprets the rating from the recommender by way of inferring the judgment method of the recommender from the historical data.

This proposal is implemented using subjective logic [JIB07]. This approach has been quantitatively compared with two other methods. The experiments show that an extended version of the "Beta trust model" [JIB07], a trust model without the elimination of subjectivity, with our method, in which subjectivity is eliminated, outperforms the original model. Although our method is not aimed at addressing the deception problem, it is able to cope with deception when a majority of feedback providers give deceptive, yet consistent ratings. In addition, our suggested method for trust and reputation systems may also be applied to other systems that include a rating mechanism such as recommender systems.

6. Contribution of Papers

This section presents the papers that have been published as part of this PhD project. Paper A is a survey and overview of the contributions of the other papers and which research goals they contribute to is given in Table 3.

For each paper, a short summary and a statement of the specific contributions of the thesis author are provided.

PAPER A

Social Computing: An Intersection of Recommender Systems, Trust/Reputation Systems, and Social Networks

Research goals/Papers	A	B	C	D	E	F	G	H
The problem of subjectivity in explicit ratings		X						
The initialization problem			X	X	X		X	X
The categorization problem				X				X
The lack of portability problem				X				X
The Cold-start problem			X	X			X	X
The time sensitivity problem						X		

Table 3. Summary of research goals and paper contributions.

The primary goal of this paper is to provide a brief survey of three popular social computing services: recommender systems, trust/reputation systems, and social networks. In this work, these services are approached from a data representation perspective and two of their main challenges: network sparsity and cold-start problems are disscussed. We also present a novel graph model, which provides an abstract taxonomy and a common data representation model for the three services. We are mainly motivated by the power of graph theory in data representation and analysis for social computing services. Through this model, we believe that it becomes more clear that data from different contexts can be related such that new solutions can be explored; thus, it may provide illumination of the aforementioned problems and stimulate new research.

This paper was the result of collaboration between the thesis author and Prof. Kevin C. Almeroth from University of California, at Santa Barbara (UCSB). Authors participated in discussions of the totality of the literature leading to the development of the model used in the paper. The paper was written mostly by the thesis author with input and editorial changes from the co-author.

PAPER B

Subjectivity handling of ratings for Trust and Reputation systems: An Abductive Reasoning Approach

This paper describes a missing part of the existing trust management models, which is handling subjectivity of recommendations. The approach is based on the idea that a recommender's judgment method can be inferred and the recommended entity can be (re)evaluated according to the value system of the truster who is about to make a decision related to an e-service. Extraction of the judgment method involves abductive reasoning which is implemented in the proposed account using subjective logic. This approach has been quantitatively compared with two other methods. The experiments show that an extended version of the 'Beta trust model', a trust model without subjectivity elimination, with our method outperform the original model with regard to dealing with subjectivity. Our suggested method for trust and reputation systems may also be applied for any other systems that includes a rating mechanism such as recommender systems.

This paper was the result of a joint work between the thesis author, Prof. Kevin C. Almeroth and Prof. Pinar Ozturk. It was written mostly by the thesis author with advice and comments by the co-authors.

PAPER C

The Hidden Trust Network underlying Twitter

In this paper, we study how to leverage Twitter activities and network structure to find a simple, efficient, but yet accurate method to infer implicit trust relationship among users. We derive hypotheses on the effects of using several different types of information such as micro-blogging activities and structure of the social network as the source of implicit trust inference and propose several methods based on them. We crawled an unbiased large data set from Twitter and we measured and compared the performance of the trust modeling strategies on this dataset. Our results confirm that the consideration of structural similarity in the network generated by users' behavior on retweeting messages can be a strong indication of implicit trust relationships among them.

This paper was written by the thesis author with advice and comments by Prof. Kevin C. Almeroth.

PAPER D

Situation-based Trust Adjustment by Conditional Trust Reasoning

This paper describes a context-sensitive trust management system that categorizes trust situations with respect to the experiences of a trustee. If the truster is familiar with the trustee, the trust judgment relies on case-based reasoning. Context-sensitivity is maintained in the description of the current and past situations that are compared. When the truster does not have any previous interaction with the trustee, a rule-based reasoner is used to assess the trustability of the trustee on the basis of available recommendations of third parties. The rules are automatically extracted from the history and encoded as conditions connecting contextual information to trust judgements. Through the use of subjective logic, this method explicitly incorporates uncertainty, thereby making it suitable in situations of partial ignorance and imperfect information. We evaluated our proposal using a large-scale real dataset.

Most parts of this paper, except the abstract and conclusion and some parts of the introduction were written by the thesis author with comments and editorial changes from Prof. Pinar Ozturk.

PAPER E

Trust Evaluation Initialization Using Contextual Information

In this paper, we propose to use contextual information for bootstrapping the reputation value. We use the Maximum Likelihood Estimation method for trust initial-

ization of probabilistic trust models. We show its implementation and effectiveness for a particular model called 'Beta reputation model' through simulations.

This paper was written by the thesis author, with advice and comments by the supervisor Prof. Svein J. Knapskog.

PAPER F

A Probabilistic Reputation Algorithm for Decentralized Multi-Agent Environments

The goal of this paper is to model trust and reputation in decentralized multi-agent systems. To achieve this, we have chosen the Ntropi model, among several other models, as a starting point, The efficiency of the model in such scenarios has been significantly improved by introducing a new probabilistic reputation algorithm for the Ntropi model.

This paper was written by the thesis author, with advice and comments by the supervisor Prof. Svein J. Knapskog.

PAPER G

Inferring Trust based on Similarity with TILLIT

In an attempt to provide high quality recommendations and proper initial trust values even when no complete trust propagation path or user profile exists, we propose TILLIT — a model based on combination of trust inferences and user similarity. The similarity is derived from the structure of the trust graph and users' trust behavior as opposed to other collaborative-filtering based approaches which use ratings of items or user's profile. We describe an algorithm realizing the approach based on a combination of trust inferences and user similarity, and validate the algorithm using a real large-scale data-set.

Most parts of this paper was written by the thesis author, with advice and comments by the supervisor Prof. Svein J. Knapskog and the co-supervisor Prof. Peter Herrmann. The abstract and some parts of the introduction was written by Prof. Peter Herrmann. The idea of the TILLIT model was the result of discussions between the thesis author and Prof. Peter Herrmann.

PAPER H

Analogical Trust Reasoning

In this paper, we present a knowledge-intensive and model-based case-based reasoning framework that supports the truster to infer such information. The suggested method augments the typically sparse trust information by inferring the missing information from other situational conditions, and can better support situation-aware trust management. Our framework can be coupled with existing trust management models to make them situation-aware. It uses the underlying model of trust management to transfer trust information between situations. We validate the

proposed framework for Subjective Logic trust management model and evaluate it by conducting experiments on a large real dataset.

This paper was written by the thesis author, with advice and comments by Prof. Peter Herrmann and Prof. Pinar Ozurk.

7. Conclusion and Possible Future Research Directions

Reputation systems confront many complex challenges, many of which yield no easy solutions. Efforts are underway to address these problems using a variety of approaches. This thesis examines current techniques used in reputation management systems and outlines a set of problems and proposed solutions. Furthermore, we present our proposals, which are motivated by what humans do in traditional trust and reputation systems. The main idea is to consider contextual information, as a special kind of implicit feedback, in trust computations and the goal is to bring additional knowledge to the reasoning process by use of available auxiliary data or Meta-data (contextual data). We believe that our work improves the utility and accuracy of trust management systems by proposing methods on how to use *contextual information*.

There are several unexplored areas for trust and reputation systems that present fertile opportunities for future research. The following list contains what we consider to be the most important open areas of research:

1) Most of the current trust and reputation mechanisms are centralized, resource-based, and personalized, which leaves space to research the suitability of many other types of system attributes. In addition, effective solutions need to be developed for the problems identified in Section 3 such as sufficient participation, easy identity changes, and strategic manipulation of online feedback.

2) Proposals from the academic community are not always deployable and are usually designed from scratch. Only in a very few cases do authors build on proposals from others. Hence, there is a need for a set of sound, standard principles for building trust and reputation systems. The design space and limitations of mediated trust and reputation mechanisms should be explored and a set of design parameters that work best in different settings should be understood. Formal models of those systems in both monopolistic and competitive settings should be developed.

3) Universal testbeds and evaluation metrics for comparison of the relative efficiency of trust and reputation mechanisms compared to that of more established systems are needed and theory-driven guidelines should be developed to decide which set of mechanisms to use.

4) A comprehensive set of robustness evaluation methods and criteria and a standardized set of attack types should be defined. Trust and reputation system robustness can be evaluated by implementing them in a real environment or from a theoretical perspective by third parties.

5) New domains where reputation mechanisms can be usefully applied need to be defined.

6) A calculated trust value should be presented to users in ways so that they can rely on this value. For example, the trust value should be accompanied with an explanation

of the estimation grounds and an uncertainty value, which shows how much data has been used for this estimation. The importance of explanation interfaces in providing system transparency and thus increasing user acceptance has been well recognized in a number of fields.

7) A decision to trust is a decision tied with risk. Even when the expectations are well grounded, there is an element of risk in trust, a chance that those who are trusted will not act as expected. The risk should be justified in order to confirm the current trust and to strengthen it, otherwise if the other party defects, trust decreases dramatically. The estimation of this risk remains a problematic area. Game theory is a powerful tool for this purpose.

8) There are fundamental differences between traditional and online environments. Therefore, adequate online substitutes for the traditional cues to trust and reputation that we are used to in the physical world should be found, and new information elements, specific to a particular online application, which are suitable for deriving measures of trust and reputation should be identified.

9) Social acceptance of trust and reputation systems is another critical factor. For the more widespread and general usage of these systems, social acceptance by all parties is an issue that needs to be considered.

Part II

INCLUDED PAPERS

PAPER A

Social Computing: An Intersection of Recommender Systems, Trust/Reputation Systems, and Social Networks

Mozhgan Tavakolifard, Kevin C. Almeroth

IEEE Network Magazine

In press, 2012

SOCIAL COMPUTING: AN INTERSECTION OF RECOMMENDER SYSTEMS, TRUST/REPUTATION SYSTEMS, AND SOCIAL NETWORKS

Mozhgan Tavakolifard,[1] Kevin C. Almeroth,[2]

[1] *Centre for Quantifiable Quality of Service in Communication Systems*
Norwegian University of Science and Technology
mozhgan@Q2S.ntnu.no

[2] *Department of Computer Science,*
University of California, Santa Barbara
almeroth@cs.ucsb.edu

Abstract Computational applications now go beyond personal computing, facilitating collaboration, and social interactions. Social computing is an area of information technology concerned with the intersection of human and social studies connected by computer networks. The primary goal of this paper is to provide a brief survey of three popular social computing services: recommender systems, trust/reputation systems, and social networks. We approach these services from a data representation perspective and discuss two of their main challenges: network sparsity and cold-start problems. We also present a novel graph model, which provides an abstract taxonomy and a common data representation model for the three services. We are mainly motivated by the power of graph theory in data representation and analysis for social computing services. Through this model, we believe that it becomes more clear that data from different contexts can be related such that new solutions can be explored; thus, it may provide illumination for the aforementioned problems and stimulate new research.

1. Introduction

Computing applications and technology have evolved rapidly over the past decade with the advance of Internet and Web technologies; the prevalence of computing resources and mobile devices; the accessibility of rich media content; and the resulting cultural and social changes. Computing is shifting to the edges of the network (i.e., networks are becoming more decentralized), and individual users are empowered with technology to use the Web for many purposes including engaging in social interaction, contributing their expertise, sharing content, and distributing information. Therefore, computer networks are inherently social networks, linking people, organizations,

and knowledge. Social computing is a novel and emerging computing paradigm at the intersection of computer science and the social sciences that involves a multi-disciplinary approach in analyzing and modeling social behaviors on different media and platforms to produce intelligent and interactive applications and results. Social computing is usually referred as a groups of *services* that are carried out by groups of people through, for example, recommender systems, trust/reputation systems, social networks, peer-to-peer networks, Wikis, and online auctions.

Three essential characteristics of computational social science are connectivity (forming relations among people within a group), collaboration (modeling the way people interact), and community (grouping or clustering of people through functional similarity and spatial closeness) [KLC09]. As social computing services become pervasive, many problems arise such as information overload and decision making problems. People are challenged to select products and reliable parties in transactions. As a solution, people seek advice from their friends or other trusted sources in *social networks* by using *trust/reputation systems* or using *recommender systems* to filter options according to their tastes. Thus, the focus of this tutorial paper is on these three social computing services. We give a brief overview of the services with the focus on data representation. The data in all of these services can be represented as a graph-based model. The major problem is that these graphs are, in reality, often too sparse. As a result, it is difficult to make predictions for new users. We believe that by using information available in a variety of different contexts, it might be possible to solve the problems of sparcity and cold-starts for new users; Thus, motivated by the power of graph theory in data representation and analysis of these services, we give an example of a common data representation as a graph-based model that exposes previously unexplored relationships among the various data elements.

Our example model neither emphasizes how the different algorithms for each service should work, nor the information that an algorithm should use (e.g., in the case of a recommender system, it does not address whether an algorithm should rely on others' ratings, on content-based features, or both). In addition, we make no claims about the results of algorithms being better or that they will be better received. By restricting its scope to exclude the actual aspect of social computing services, our framework provides a systematic and rigorous way to study these social computing services and stimulates new research directions on how to derive benefit from the interpretability among these services.

The remainder of the paper is organized as follows. Section 1.1 provides a brief overview of the three social social computing services with a focus on a graph-based data representation. We explore some of the current research challenges regarding these services, specifically the sparsity and cold-start problems, in Section 3. Then, we give an example of a common graph model that provides for all three services in Section 4. Finally, Section 5 offers concluding remarks and suggestions for future research.

Table 1. Example social computing services and methods.

Social Computing Services/Technologies	**Methods/Algorithms**
Recommender systems (e.g., Netfix)	Content-based Collaborative Hybrid
Trust/Reputation systems (e.g., eBay, Sporas, Histos)	Summation or Average Bayesian Systems Fuzzy models Flow models
Social networks (e.g., FaceBook, MySpace)	Node neighborhoods Ensemble of all paths

2. Background: Overview of Social Computing Services

In this section, we present an overview of the three selected social computing services: recommender systems, trust/reputation systems, and social networks with a focus on a graph-based representation of their underlying data. Each of these services has been implemented using several methods. Table 1 shows an example of each service and some of their methods. Other work provides a thorough survey on social computing [KLC09], recommender systems [AT05], trust/reputation systems [JIB07], and social networks [Sco06]. However, we present our brief survey with the goal of highlighting the common challenges of these services and possibility of designing a common framework as the solution. We believe that merging social networks, social trust relationships, and recommender systems can improve the accuracy of all of these services and improve a user's experience.

2.1 Social Networks

An online social network models connections among individuals or objects and facilitates information exchange between individuals or groups using relationships between users. Data is usually represented using graphs and matrices. Graph theory has been widely used to analyze social networks due to its representational capacity and simplicity [Kad12]. In general, the properties of social network graphs have been studied extensively. However, little is known in the research community about the properties of online social network graphs at scale, the factors that shape their structure, or the ways they can be leveraged in information systems.

In social networks, the representation by graphs is also called a "sociogram", where the nodes are called actors and the edges are called relationships. The relationship can be non-directional(e.g., marriage) or directional (e.g.,seller-buyer relationship). Characterizing the relationships that exist between a person's social group and his/her personal behavior has been a long standing goal of social network analysis.

Social networks are also known to be globally sparse and locally dense [Sco06]. Given a snapshot of a social network, inferring which new interactions among its members are likely to occur in the near future is formalized as a link prediction problem. The link prediction problem asks to what extent the evolution of a social network can be modeled using features intrinsic to the network itself. The link prediction problem is also related to the problem of inferring missing links from an observed network. In a number of domains, a network of interactions based on observable data is constructed and then other likely-to-exist links are inferred. All methods can be viewed as computing a measure of proximity or similarity between nodes relative to the network topology. In general, the methods are adapted from techniques used in graph theory and social network analysis; the dynamic power of graph theory lies not in its terminology but, like any other branch of mathematics, in its theorems. Two categories of link prediction methods are as follows:

- *Node neighborhood methods*: these approaches are based on the idea that two nodes are more likely to form a link in the future if their sets of neighbors have a large overlap.
- *Shortest paths methods*: these methods rank two nodes by the length of their shortest path. Such a measure follows the notion that collaboration networks are "small worlds," in which individuals are related through short chains. Some of these methods refine the notion of shortest-path distance by implicitly considering the ensemble of all paths between two nodes.

A special kind of social network is called an "affiliation network," in which nodes are actors and events to which the actors belong. Affiliation networks can also be described as collections of subsets of entities. Each event describes the subset of actors who are affiliated with it, and each actor describes the subset of events to which it belongs. Viewing an affiliation network this way is fundamental to the *hypergraph* approach.

As Figure 1 shows, hypergraph is a generalization of a graph, where an edge can connect any number of nodes. The nodes are actors and the edges are considered as the set of events. Furthermore, in some cases, the use of simple or directed graphs to represent the complex networks does not provide a complete description of the real-world systems under investigation. For example, in a collaboration network represented as a simple graph, we cannot know three or more users linked together in the network have collaborated on the same project or not.

2.2 Recommender Systems

The objective in a recommender system is to reduce information overload and retain customers by selecting a subset of items (e.g., movies or books) from a universal set based on user preferences. In its most common form, the recommendation problem is reduced to the problem of estimating ratings for the items that have not been seen by a user. Intuitively, this estimation is usually based on the ratings given by this user to other items and possibly other information as described below. Once we can estimate

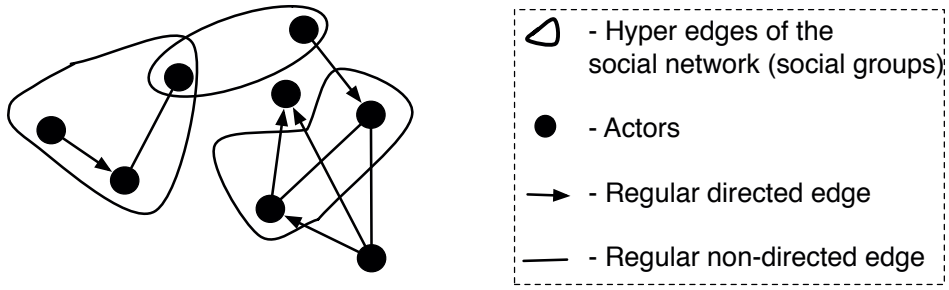

Figure 1. Hypergraph data representation for affiliation networks in social networks.

ratings for the yet unrated items, we can recommend to the user the items with the highest estimated ratings. The new ratings of the not-yet-rated items can be estimated in many different ways using methods from machine learning, approximation theory, and various heuristics. Recommender systems are usually classified according to their approach to rating estimation and have traditionally been studied from a content-based filtering versus collaborative design perspective [AT05]:

- *Content-based methods*: similar items to the ones the user preferred in the past will be recommended to the user. In particular, various candidate items are compared with items previously rated by the user and the best matching items are recommended. For example, if a particular user reads many online articles on the topic of nanotechnology, then content-based recommendation techniques will recommend other nanotechnology articles. This recommendation will be made because these articles will have more nanotechnology-related terms (e.g., "Nanooptics" and "Nanobiotechnology") than articles on other topics.
- *Collaborative filtering methods*: items that other people with similar tastes and preferences like will be recommended. For example, in a movie recommendation application, in order to recommend movies to a user, the collaborative recommender system tries to find other like-minded users, i.e., other users that have similar tastes in movies. Then, only the movies that are most liked by these like-minded users recommended.
- *Hybrid*: several recommendation systems use a hybrid approach by combining collaborative and content-based methods. This solution helps to avoid certain limitations of content-based and collaborative systems.

Current recommender systems use various kinds of data representations that usually capture three basic elements: user data (e.g., gender and address), item data (e.g., product category and price), and transaction data (e.g., user's rating, time and place of transaction). The research in recommender systems grew out of information retrieval and filtering; as a result, data is usually modeled as a user-item matrix.

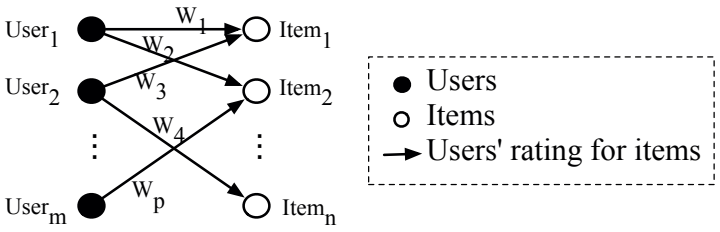

Figure 2. Graph-based data representation for recommender systems.

Another approach is a graph-theoretic model where a bipartite, directed and weighted graph with heterogeneous nodes (i.e., users and items) and homogeneous edges (i.e., purchases) can be used to represent the data. As Figure 2 shows, nodes represent users and items while edges represent users' ratings for items. Weights on the edges correspond to the rating values.

Despite significant research progress and growing acceptance in real-world applications, at least two major challenges limit the implementation of effective e-commerce recommendation applications. The first challenge is concerned with making recommendations based on sparse transaction data, also known as the sparsity problem. The second challenge is the lack of a unified framework to integrate multiple types of data and recommendation approaches. For better recommendation performance, a unified recommendation framework with the expressiveness to represent multiple types of input data and a generic computing mechanism to integrate different recommendation approaches is needed to fully exploit the rich information available at e-commerce sites. We explore these challenges in more detail in Section 3.

2.3 Trust/Reputation Systems

In the Web, where vast amounts of content is created by users, the question of whom to trust and what information to trust has become more important and more difficult. Trust/Reputation systems represent a significant trend in decision support for Internet services. The basic idea is to let parties rate each other, for example after the completion of a transaction, and to use the aggregated ratings to derive a trust or reputation score, which can assist others in deciding whether or not to transact with that party in the future [JIB07].

Jøsang distinguishes between two categories of trust: *reliability trust* and *decision trust* [JIB07]. Reliability trust is defined based on "the subjective probability by which an individual expects that another individual performs a given action on which its welfare depends." Decision trust is defined as "the extent to which one party is willing to depend on something or somebody in a given situation with a feeling of relative security, even though negative consequences are possible." A *trust relationship* exists between two agents when one agent has an opinion about the other agent's trustworthiness and a *recommendation* is a communicated opinion about the trustworthiness of

a third party. *Reputation* is defined as an "expectation about an agent's behavior based on information about or observations of his past actions." Therefore, reputation can be considered as a collective measure of trustworthiness (in the sense of reliability) based on the referrals or ratings from members in a community. An individual's subjective trust can be derived from a combination of received referrals and personal experience.

A reputation system uses a specific method (e.g., averaging, probabilistic-based or belief-based) to compute reputation values for a set of objects (e.g., users, goods, or services) within a community based on the collection of recommendations from others. These reputation values may be used by the entities in the community for decision-making purposes. Here, we describe some of the various methods for computing reputation and trust measures [JIB07].

- *Simple summation or average of ratings*: the simplest form of computing reputation scores is to sum the number of positive ratings and negative ratings separately, and to keep a total score as the positive score minus the negative score (e.g., eBay) or as the average of all ratings (e.g., Epinions and Amazon).
- *Bayesian systems*: a reputation score is computed by updating probability density functions (PDFs). The updated reputation score is computed by combining the previous reputation score with the new rating.
- *Fuzzy models*: these methods represent trust and reputation as linguistically fuzzy concepts where membership functions describe to what degree an agent can be described as trustworthy or not. Fuzzy logic provides rules for reasoning with fuzzy measures of this type.
- *Flow models*: A participant's reputation increases as a function of incoming flow, and decreases as a function of outgoing flow (e.g., Google's PageRank, Advogato). In the case of Google, many hyperlinks to a web page contributes to increased PageRank whereas many hyperlinks from a web page contributes to decreased PageRank for that web page.

Data for trust/reputation systems can be represented as a directed, weighted graph with homogenous nodes and edges. As shown in Figure 3, nodes are trustees and trusters (parties), edges are trust relationships, and the weights are trustworthiness values. The web of trust is often too sparse to predict trust values between non-familiar people, since in large online communities, a user has experience with only a very small fraction of the other community members. As a result, very often there will be no trust relation to an intended new partner of an e-commerce transaction.

3. Challenges

In this section, we provide further discussion on some of the major challenges for the three social computing services. In particular, we discuss the sparsity problem, which is one of the motivations for the graph-based representation model proposed in the next section.

The graphs in social networks, recommender systems, and trust/reputation systems are usually too *sparse*. In recommender systems, the numbers of users and items

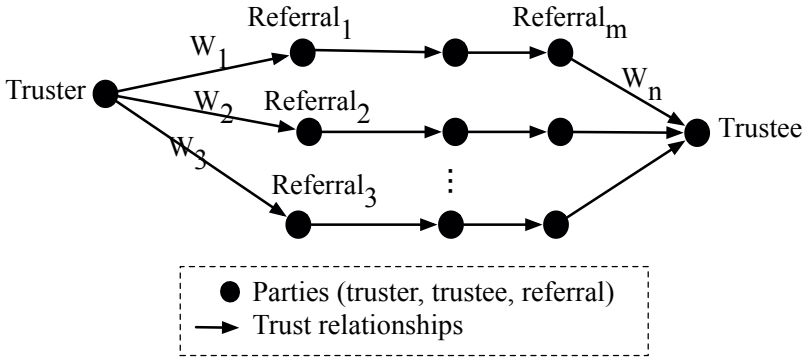

Figure 3. Graph-based representation for trust/reputation systems.

are very large. Even active users rate just a few of the total number of items and respectively, even very popular items are rated by only a few of the total number of users [AT05].

The *cold-start* problem emphasizes the importance of the sparsity problem. In recommender systems, this problem refers to the situation where an item cannot be recommended unless it has been rated by a substantial number of users. This problem applies to new and obscure items and it particularly effects users with eclectic taste. Likewise, a new user has to rate a sufficient number of items before the recommendation algorithm is able to provide reliable and accurate recommendations. In trust/reputation systems, a node must participate in interactions with others in order to raise its reputation score. As nodes in the system tend to interact with nodes with higher reputation scores, when a new node joins the system with a very low reputation score or no reputation score at all, its chance of being selected for interaction is generally rare. Hence, it is hard for a new user to raise his or her reputation score.

These problems may be alleviated by taking into consideration the interconnections among different services. As an illustration, recommendations in recommender systems are not delivered within a vacuum but rather cast within social networks. Thus, all recommender systems make connections among people either directly as a result of explicit user modeling or indirectly through the discovery of implicit relationships in data. Considering that a ratings dataset can be modeled as a bipartite graph rather than a matrix, social networks can also be formed by applying transformations on the bipartite graph, for example, two users are connected if they have rated a common item. As mentioned in the previous section, in social network theory this bipartite graph is referred to as an affiliation network.

Techniques to discover existing social networks from patterns embedded in interaction (transaction) data are analogous to collecting implicit declarations of preferences in recommender systems. Indeed, the use of social networks has expanded to many

diverse application domains of recommender systems such as digital libraries and community-based service location [PGF04].

Another example is the similarities between collaborative filtering and reputation systems. Both types of systems collect ratings from members in a community/social network. The usefulness of the former arises when the emphasis is on the content, and the latter can be used when the source of information is a more important factor. They are thus complimentary social mechanisms in global open distributed systems. There is significant potential to combine collaborative filtering and reputation systems [JIB07]. Another example is investigating Web-based social networks and its applicability to different tasks such as trust inferrencing within trust networks. In addition, the location of a given member of a community within a social network can be used to infer some properties about his or her degree of expertise, i.e., his or her reputation [GH04].

However, the methods used in these examples are application-specific. This fact limits the data inputs and representations that can be used. We believe that a model should be comprehensive to support diverse inputs and representations. Furthermore, it should be flexible to support a variety of different approaches. To this end, we propose a common representation model for all three services in the next section.

In addition, for the sparsity and cold-start problems, current approaches miss many desirable aspects such as explainability of their predictions in terms and constructs that are natural to the user/application domain, effusivity and subjectivity of ratings and feedback, and coping with easy name changing. In the next section, we present an example for a joint representation graph model that facilitates the collaboration among these services.

4. A Common Data Representation Model

The previous section showed that the field of social computing calls for a common taxonomy, data representation, and comprehensive model. This model should have the capability to represent different types of data inputs and to support different approaches using various methods. Motivated by these needs and the analysis power of graph theory, we take a connection-oriented approach toward social computing research and suggest an example common data representation model for the three services as a solution for the sparsity and cold-start problems. Our intuition is to seek for other contextual information when the data is sparse and there is no information available for prediction. In other words, the affiliation network in social networks may be used as an underlying context for recommender systems and trust/reputation systems. As a result, by merging the graphs of all of these services, it is possible to infer missing links of one using links from the others.

Our proposed model, as shown in Figure 4, is a heterogeneous two-layer weighted directed hypergraph in which the two layers of nodes represent users and items. Three types of links between nodes capture information about users, items, and transactions. Hyperlinks, shown as hyper edges in the figure, are social relations among users corresponding to affiliation networks in social networks. Other information about users, such as demographics, may also be added (grey edges). The links between items (dashed edges) captures the similarity between them. Different types of item

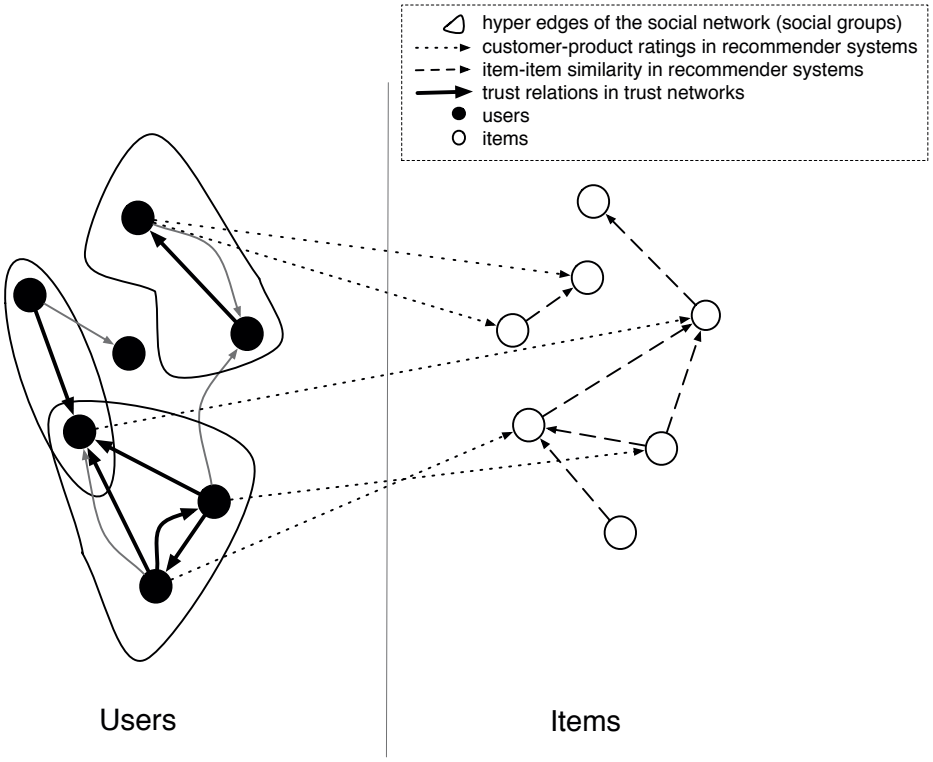

Figure 4. A common data representation model for recommender system, trust/reputation system, and social networks.

information can be used to compute similarity. For products like books and movies, the product description can also be used to compute product similarity. Inter-layer links (dotted edges) are formed based on the transaction information that captures the associations among users and items (e.g., purchase history, customers' rating, or browsing behavior). Different types of transaction information may be combined in the model by assigning different weights to reflect different association strengths. For instance, a high rating on a product may be weighted higher than browsing activity, because the former reflects the user's interest more directly.

We briefly describe the use of our graph model in solving various service-related problems. Our two-layer model captures all types of data inputs and covers all the data representations that were summarized in Section 1.1. The model is flexible because different combinations of edges can be activated at run time. A rich set of analytical tools developed in graph theory (e.g., random graph search, topological graph analysis, and link prediction) can be adopted to study properties of the model such as paths and clusters that may lead to improved methods for the services. As a case in point, the recommendation problem in recommender systems can be viewed as a link prediction

problem. For collaborative filtering, the non-present links (e.g., future transactions or potential interests) are predicted based on the links observed in the current graph. For the content-based and hybrid approaches, the links are predicted based on a graph that is enhanced by adding attributes about user and item nodes. In the following subsections, we explain the applicability of this model for each service in more detail.

4.1 Recommender Systems

As shown in Figure 4, closely related users, based on their relationship in the social network (hyperlinks) or people in the same trust network (thick solid edges), are clustered into groups. Users in the same group are potential neighbors for the collaborative filtering techniques that address the sparsity problem [MKR03, PPK05]. The cold start problem may also be addressed through explicit specification of a user's closest neighbors. For example, a new user joins an online book shop. There is no information available about the previous history of book purchases by this user. However, the books purchased by his/her close friends on the social network can be used as a basis for recommendations.

This representation satisfies all of the pertinent aspects for recommender systems outlined in the previous section. It utilizez a social network model, and thus, emphasizes connections rather than prediction. The nature of connections also aids in explaining the recommendations. The graph theoretic nature of connections allows the use of mathematical models (such as random graphs) to analyze the properties of the social networks in which recommender algorithms operate.

4.2 Trust/Reputation Systems

As shown in Figure 4, the sparsity and cold-start problems in trust networks may be improved by clustering users who are in the same social group (red hyperlinks) or users with similar historic ratings for products (dotted edges) in the same group. Then, the trust level is a common value for a group of users rather than individuals. As the groups can differ in purpose, one entity can be a member of more groups. Trust between two entities is then inferred based on their group memberships. Such models allow trust to be built between mutually unknown entities with less communication and computation load [Spa07]

Further, it is easier for the services to cope with the problem of multiple identities with this representation. In online communities, it is usually easy for members to disappear and re-register under a completely different online identity with zero or low cost. Community members can build a reputation, milk it by cheating other members, and then vanish and re-enter the community with a new identity and a clean record. In contrast, in an integrated system, it would be more costly for users to change identities in one service since they lose their current networks in the other services as well.

4.3 Social Networks

As shown in Figure 4, a social relationship between two users may be inferred based on a mutual or a transitive trust relation between them. In this way, the existence of

the trust network (thick solid edges) helps to bootstrap relations in the social network (hyperlinks) and results less sparsity.

Similar product rating patterns between two customers may also be used to induce a social relation between them. Therefore, item-item edges, which is the similarity between items in a recommender system (dashed and dotted edges), may be used to create a social relation (hyper edges) between the users who have have similar ratings for those items. In the simplest form, two users are connected if they have rated a common item. The cold-start problem is less of a problem in this approach as implicit ratings bootstrap the system [PGF04]. Perugini et. al [PGF04] posit that recommender systems have an inherently social element and is ultimately intended to connect people either directly as a result of explicit user modeling or indirectly through the discovery of relationships implicit in extant data.

5. Conclusions and Future Work

In this paper, we have described several challenges arising in social computing. Although these problems have been the focus of numerous papers, solutions to these problems in the context of the evolving Internet are still lacking. Specifically, in social computing, there exist the problem of sparsity, cold start users, multiple identities, and context insensitivity. We have shown through a novel example how the integration of the three social computing services can help to alleviate these problems.

For future work, effective solutions need to be developed for the problems identified in Section 3. We shortly discussed how link prediction in one service can help to reduce the cold-start and sparsity problem in the other services. However, future researchers can look to use our example graph-based model as a basis for solving a variety of important social computing problems and investigate further how graph theory tools and techniques such as random graph search and topological graph analysis can be applied using our model to help the propagation of data and knowledge from one service into another.

References

[AT05] 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.

[GH04] J. Golbeck and J. Hendler. Accuracy of metrics for inferring trust and reputation in semantic web-based social networks. In *International Conference on Knowledge Engineering and knowledge Management (EKAW), Northamptonshire*, pages 116–131, October 2004.

[JIB07] Audun Jøsang, Roslan Ismail, and Colin Boyd. A survey of trust and reputation systems for online service provision. *Decis. Support Syst.*, 43:618–644, March 2007.

[Kad12] C. Kadushin. *Understanding Social Networks: Theories, Concepts, and Findings*. Oxford Univ Pr, 2012.

[KLC09] Irwin King, Jiexing Li, and Kam Tong Chan. A brief survey of computational approaches in social computing. In *Proceedings of the International Joint Conference on Neural Networks (IJCNN)*, pages 2699–2706, Atlanta, Georgia, USA, June 2009.

[MKR03] B.J. Mirza, B.J. Keller, and N. Ramakrishnan. Studying recommendation algorithms by graph analysis. *Journal of Intelligent Information Systems*, 20(2):131–160, 2003.

[PGF04] S. Perugini, M.A. Gonçalves, and E.A. Fox. Recommender systems research: A connection-centric survey. *Journal of Intelligent Information Systems*, 23(2):107–143, 2004.

[PPK05] M. Papagelis, D. Plexousakis, and T. Kutsuras. Alleviating the Sparsity Problem of Collaborative Filtering Using Trust Inferences. In *Trust management: international conference (iTrust)*, page 224, Paris, France, May 2005.

[Sco06] J. Scott. *Social network analysis: A handbook*. Sage, 2006.

[Spa07] R. Spanek. Reputation system for large scale environments. In *International Conference on Digital Information Management (ICDIM)*, pages 621–626, Lyon, France, October 2007.

PAPER B

Subjectivity handling of ratings for Trust and Reputation systems: An Abductive Reasoning Approach

Mozhgan Tavakolifard, Kevin C. Almeroth, Pinar Ozturk

JDCTA: International Journal of Digital Content Technology and its Applications

Vol. 5, No. 11, 2011

Is not included due to copyright

PAPER C

The Hidden Trust Network underlying Twitter

Mozhgan Tavakolifard, Kevin C. Almeroth

Submitted to the IEEE Transactions on Information Forensics and Security.
2012

THE HIDDEN TRUST NETWORK UNDERLYING TWITTER

Mozhgan Tavakolifard,[1] Kevin C. Almeroth,[2]

[1] *Centre for Quantifiable Quality of Service in Communication Systems*
Norwegian University of Science and Technology
mozhgan@Q2S.ntnu.no

[2] *Department of Computer Science,*
University of California, Santa Barbara
almeroth@cs.ucsb.edu

Abstract Twitter is today the most prominent micro-blogging service available on the Web. Information overload is a major problem on most online social networks and particularly on Twitter. It is difficult to find the right people and content to focus on. Personalized and trust-based recommender systems have emerged as a solution to alleviate this problem on online social networks. However, there is no explicit trust network on Twitter. In this paper, we study how to leverage Twitter activities and network structure to find a simple, efficient, but yet accurate method to infer implicit trust relationship among users. we derive hypotheses on the effects of using several different types of information such as micro-blogging activities and structure of the social network as the source of implicit trust inference and propose several methods based on them. We crawled an unbiased large data set from Twitter and we measured and compared the performance of the trust modeling strategies on this dataset. Our results discover that the consideration of structural similarity in the network generated by users' behavior on retweeting messages can be the best indication of implicit trust relationships among them.

1. Introduction

Online social networks have emerged recently as the most popular application since the Web began and are considered by many groups such as scholars, advertisers, and political activists as an opportunity to study the propagation of ideas. For example, Twitter as a micro-blogging service counts with millions of users from all over the world and facilitates real-time propagation of information to a large group of users. The simplicity of Twitter and its real-time message streams are its most powerful features. These real-time message streams have greatly expanded the usage of social network sites from political campaigning to education, and from emergency news reporting to marketing and public relations. Particularly, Twitter is an ideal environment for the dissemination of breaking-news directly from the news source and/or geographical location of events; therefore, it has made interesting inroads into novel domains such

as emergency response and recovery under crisis situation (*e.g.*,, Twitter-based early warning systems [SOM10], help during a large-scale fire emergency[1], updates during riots in Kenya[2], and live traffic updates to track commuting delays[3]).

Social network sites have experienced an explosion in both the number of users and the amount of user contributed content in recent years, therefore, it is difficult to find the right people and content to focus on. Active Twitter users now face thousands of unread messages in their stream every day, as well as millions of other Twitter users that they could engage with if they wish. This challenge, which is called "attention scarcity", is a key challenge resulted from information overload and abundance of relationships among users. The linked structure of Twitter does not reveal actual interactions among people. In reality people interact with very few of those listed as part of their network. A study of social interactions within Twitter reveals that the driver of usage is a sparse and hidden network of connections underlying the declared set of friends and followers [HRW09].

To solve the attention scarsity problem, there is a need to present and suggest relevant data and contacts to the users of online social networks in a personalized and effortless way. By personalized, we mean that the help should be inherently personalized to individual users and by effortless, we mean that the help should be proactive without requiring any knowledge, skill or effort from the users. One way to do this is to make the hidden network visible. The underlying hidden network is often referred as the web-of-trust in the literature.

There is some work focused on exposing the hidden network or the web-of-trust on Twitter [AEG+10, AHTS10, NHL10] among the ongoing research on trust modeling for online social networks [GH06b, KG07, SS02, ZL05, Bus98, Gol06]. The main weakness is that they rely on explicit trust ratings. Trust inference can be based on explicit trust rating versus implicit information. There are many challenges associated with explicit ratings [TA12]; on the other hand, implicit feedback, where users' actions are recorded and the feedback is inferred from the recorded data can be used as well.

In our work, we aim to show that we are able to predict accurately and efficiently the hidden relationships in the Twitter by using the available knowledge about users' behavior. By exposing the hidden network we can solve the attention scarcity problem by making recommendations to users about the relevant data and contacts.

The main reason for why we use behaviors as indicator of trust in our approach is that we believe that the relationship among users in this hiddent network have to be intuitive. Humans must be able to comprehend why user *A* is related to user *B* and come to similar results when asked for a personal judgement.

More specifically, our solution uses **implicit** information that describe direct connections between people in Twitter and compose this information to infer in **real time** the links between two people who are not directly connected. We considered three different behaviors of users: follower-followee relationships, retweeting, and making

[1] http://factoryjoe.com/blog/2007/10/22/twitter-hashtags-for-emergency-coordination-and-disaster-relief/
[2] http://www.economist.com/node/10608764
[3] http://lifehacker.com/355453/track-commuting-delays-via-twitter-with-commuter-feed

tweets favorited as the indicators of trust. Our approach requires two main steps: inference and propagation. In the first step, inference, the hidden existing relationships from the Twitter network are extracted using the three trust indicating behaviors. Three webs of trust were generated in this step. In the second step, propagation, new potential links in the inferred network are predicted. We propose four different algorithms for trust propagation.

We conducted hypothesis driven experimentation and an in-depth analysis on a large Twitter dataset of more than 20000 users for evaluation. We compared the three webs of trust by applying the four trust propagation methods to see which yields the best results using different metrics to show the efficiency and accuracy of each method in prediction of the trust relationships. Our goal was to answer the following questions: which model among the three assumed behavior is a better indication of trust between linked users in terms of accuracy and efficiency? which of the four methods predict trust among unknown user pairs better? and whether there is any simple and efficient way to implicitly infer the trust relations?

The main advantage of our solutions is its efficiency, acceptable accuracy, and scalability. Although our proposal is based on the graph structure, the trust value between each two nodes can be calculated in a real time manner. Our algorithms are efficient and scale to million-node networks. Furthermore, we use implicit information as opposed to the related work. Nevertheless, we lost accuracy to some level for the sake of improving the algorithm's efficiency by using several heuristic methods.

The rest of this paper is organized as follows. Section 1.1 provides a brief background knowledge. We present the details of our various trust prediction methods for the Twitter network in Section 3 and describe the analysis and evaluation of these methods in Section 4. We explore the related work in Section 5. Finally, Section 5 provides concluding remarks and future research directions.

2. Background

In this section, we introduce the necessary background for successfully introducing trust to the Twitter network. Before we show how Twitter can be viewed through the lenses of our trust model, we would like to introduce a few salient concepts of Twitter. Twitter as an online social netowrk is an information sharing system, where users follow other users in order to receive information along the social links. Relationship links are directional, meaning that each user has followers and followees, instead of unidirectional friendship links. Twitter allows users to post and exchange 140-character-long messages, which are also known as "tweets". Tweets can be published by sending e-mails, SMS text-messages, and directly from smartphones using a wide array of Web-based services and can be repeated throughout the network, a process called re-tweeting. A retweeted message usually starts with RT @username, where the @ sign represents a reference to the one who originally posted the messages. The strength of Twitter as a medium for information diffusion stands out by the speed of retweets. Twitter users usually use hashtags (#) to identify certain topics. Hashtags are similar to a tag that is assigned to a tweet in its own body text.

Recommender systems have emerged as a promising solution for the aforementioned information overload problem in social network sites. For example, a new user to Tiwtter may follow a few college friends and the recommender may then help him by inferring from her existing social network a few other friends that she forgot to follow. As another example, an active user may find a majority of the many tweets he receives to be boring. To help him, a recommender may first identify his interest in arts by examining his previous tweets on the site, and then suggest to him a few interesting tweets on the topic of arts. All these helps can happen naturally without any extra user input or knowledge.

It has been shown in the prior research that incorporation of information about the web-of-trust in the recommendation prediction algorithm improve performance of the recommender systems [MA04, GH06a, WBS08]. The main challenge is that explicit trust information is not always available on online social networks. Furthermore, there are many problems associated with explicit trust ratings [TA12] such as *low incentive or no incentive for users to provide ratings*, *users' bias toward positive ratings*, *initialization and cold-start problem*, *subjectivity*, and *false ratings*.

On the other hand, users of social network sites interact with each other directly by making connections, sending messages, and sharing various contents. These direct interactions on social network sites can serve as a great data source, which may implicitly indicate trust. The main intuition is that trust between two users may result in certain typical behaviors. These behaviors are not only an expression of trust, but could also facilitate the development of further trust. Such behavioral expressions are not guaranteed expressions of trust; they are more noisy indicators of trust. The more often they occur, the more indicative they are of trust.

Motivated by the idea of using implicit users behaviors as indication of trust, we adopt the following as the trust definition in online social networks: "trust in a person is a commitment to an action based on a belief that the future actions of that person will lead to a good outcome [GH06b]". The action and commitment in this definition do not have to be significant. For example, in case of the Twitter social network, we could say user A trusts user B regarding semantics of the Twitter messages (tweets) if she chooses to read the message that B posts, retweet it, or make it favortied (commits to an action) based on her belief that B will not waste her time. Several properties of trust follow from this definition [GH06b], namely:

- Transitivity: The primary property of trust that is used in our work is transitivity. Trust is not perfectly transitive in the mathematical sense, that is, if A highly trusts B, and B highly trusts C, it does not always and exactly follow that A will highly trust C. There is, however, a notion that trust can be passed between people. For example, When we ask a trusted friend for an opinion about a babysitter, we are taking the friends opinion and incorporating that to help to form a preliminary opinion of the babysitter.

- Composability: There is more reasoning and justification for a belief with information from many people. if we look at trust recommendations as evidence

used to support the belief component of trust, then the trust values from many sources can be composed together to form a single opinion.

- Asymmetry and personalization: Trust is not necessarily identical in both directions for two people involved in a trust relationship. Because, individuals have different experiences, psychological backgrounds, and histories. For example, doctoral students typically say they trust their advisors more than the advisors trust the students.

An online social network can be modeled as a graph with users as the nodes and the relationship between users as the links. The web-of-trust is a directed weighted graph with the same nodes as the online social network graph. A link from user *A* to user *B* in the trust web shows user *A*'s trust in user *B* and the weight on the link shows how highly user *B* is trusted by user *A*. In the next section, we present the details of trust modeling on Twitter.

3. Inference and Prediction of the Hidden Web of Trust in Twitter

We aim to make out the web of trust from the interaction pattern that users have with each other. A two-step approach is taken in our solution. In the first step, trust inference, the hidden web of trust is extracted from the Twitter network using the implicit information about users behavior. Three sources of information are used to build three different webs of trust. In the second step, trust propagation, we propose four methods to predict trust links among users that are not connected directly to each other in the current web of trust. The goal is to make the web of trust less sparse. The details of these two steps are provided in the following.

3.1 Trust Inference

We consider three users behaviors as an expression of trust: follower-following relationships, retweeting, and the behavior of making tweet messages favoritedand. We assume trust links with weights in the range (0, 1]. The higher weight mean the higher trust.

There are also other possible indicating behaviors (e.g.,, direct conversation between users called *mentions* in Twitter), which we did not consider. We chose our trust indicating behaviors intuitively; and in our opinion, *direct conversations* are not necessarily an indication of trust between two persons. In following, we describe how each behavior can be considered as an indication of trust.

- Followee-follower: If user *A* trusts user *B*, then it is likely that user *A* follows user *B*.

- Retweet: Our second indication of trust is based on the propagation of information. If *B* retweets information from user *A* often, then we assume that *B* must be trusting *A*. The motivation behind this idea is that we observed people only

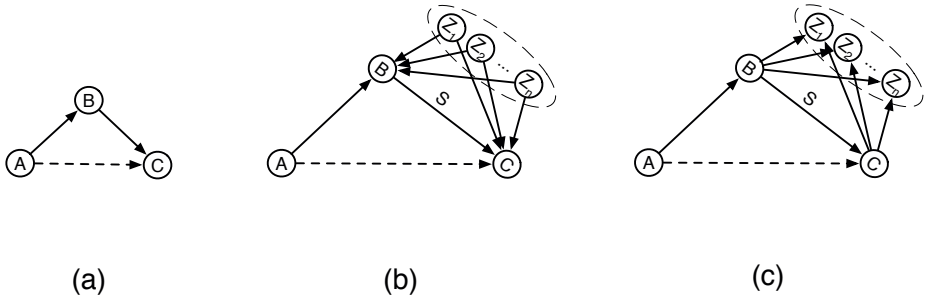

Figure 1. Different trust propagation methods: transitivity (a), similarity (b) and (c).

retweets from a small number of people and only a subset of a users followers actually retweet.

- Favorite: User *A* often making some of user *B* tweets as his/her own favorites and this indicates that user *A* trusts user *B*.

The webs of trust resulted from these inference methods are often too sparse to be helpful in practice (e.g.,, to be used by a recommender system) since a user has usually relationship or interaction with only a very small fraction of the total community members. Thus, very often there will be no trust link to an intended new user. Trust propagation methods can predict some of the missing links in the trust web to make it more dense and to alleviate the consequences of the sparsity and possible cold-start problems.

3.2 Trust Propagation

We suggest four different methods for trust propagation, prediction of the missing links, in the web of trust. Three of them are derived based on the *transitivity* feature of trust and the fourth one is based on the idea of *similarity* between users as a predictor of trust. By trust transitivity, we expect that people who the user trusts highly will tend to agree with the user more about the trustworthiness of others than people who are less trusted. For example, users are more likely to trust the "taste" of people they are following in the Twitter. That is, if *A* trusts *B* who trusts *C*, then *A* will also trusts *C*, as shown in Figure 1(a). On the other hand, this approach is helpful, provided that a complete transitive trust path exists between the truster and the trustee.

We propose an alternative approach based on similarity. One can exploit the like-mindedness resp. similarity of individuals based on collaborative filtering to infer trust to yet unknown parties. For instance, if one knows that with respect to a specific property, two parties are trusted alike by a large number of different trusters, one can assume that they are similar, as shown in Figure 1(b). Likewise, if two parties trust alike a large number of other users, they can be assumed to be similar, as shown in Figure 1(c).

If A has trust in B (there is a direct trust link from A to B in the web of trust) who is similar to C (they are *similar trustees*), then A can infer its trust to C. Two trustees are similar if they are both similarly trusted by other users Z_1, Z_2, ..., Z_n, as shown in Figure 1(b). This helps to predict new trust links, where it is not possible to predit any trust link from A towards B in a trust web using transitivity. The case of similar trusters is shown in Figure 1(c). We provide the details of each method's formulation in the following.

3.2.1 Trust Propagation through Transitivity

The simplest form of trust propagation is trust transitivity which is widely discussed in the literature [DKG+05, GKRT04, MBKM07, QHC07, YCB+02]. That is, if A trusts B who trusts C, then A will also trusts C, which we call it "simple-transitivity". This method has the time complexity of O(m*d) and the space complexity of O(m), where m is the number of links and d is the average degree of the graph.

It is important to consider the number of users like B that form a transitive path between user A and C. The higher the number of these users is, the stronger the predicted trust relationship between A and C will be. In an improved method called "weighted-transitivity", we assume user A trusts user C provided that there exists at least a number of other users like B that connects them in a transitive path $A \rightarrow B \rightarrow C$. For instance, the average number of paths of length two between users in the graph can be used as the threshold. The weighted-transitivity method predict a new trust link and assign its weight as T_{AC} by simply multiplying the weights of the links in the transitive path (T_{AB} and T_{BC}), as shown in 1. The intuition behind this multiplication is that we assume user B as a referral who sends his/her opinion of trust (T_{BC}) about user C to user A. Since it is also important to consider the user A's trust in user B (T_{AB}) as the referral in the formulation, the recommended trust by B (T_{BC}) is multiplied by users B's trustability (T_{AB}). As a consequence, If user B is not much trusted by user A, then his/her recommended trust value will be de-emphasized as well.

$$T_{AC} = T_{AB} * T_{BC} \tag{1}$$

Another formula is also proposed by Golbeck [Gol06] as the following:

$$T_{AC} = \frac{\sum T_{AB} * T_{BC}}{\sum T_{AB}}$$

We call this method as "golbeck-transitivity". The main difference in their method is that the trust value is not de-emphasized by the reliability of the referral. For example, the inferred trust value will be equal to T_{BC} in the case of one existing transitive path between A and C. Both weighted-transitivity and golbeck-transitivity methods have time complexity of O(m*d) and space complexity of O(m), where m is the number of links and d is the average degree of the graph.

In this paper, we just considered the paths of length two as the transitive paths. Theoretically, the transitive path between node A and node C can be of any length. However, previous work [Gol06] has addressed this issue and shown that, as expected, shorter paths lead to more accurate information. In our formulation, the predicted

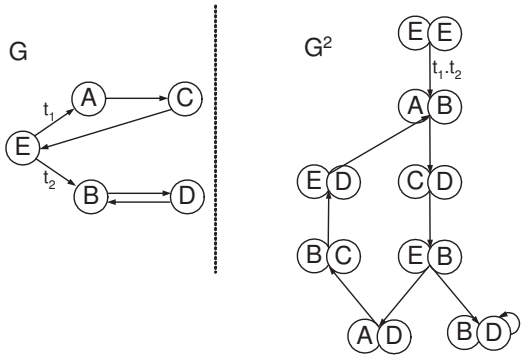

Figure 2. Similarity measurement.

value between two nodes over a very long path will be very small as the result of successive multiplications. Therefore, we considered only the paths of the length two in our study. It is intuitive since a user expect that neighbors who are connected more closely will give more accurate information than those who are further away in the network.

3.2.2 Trust Propagation Through Structural Similarity

This method addresses the way in which the level of trust in cooperative relations depends on similarity of nodes in the network structure. We measure similarity between each two users with respect of trusting other users or being trusted by other users. The intuition behind our algorithm is that, *similar* users are related to *similar* users! More precisely, users A and B are similar if they are related to users C and D, respectively, and C and D are themselves similar. The base case is that each user is similar to itself. If we call the web of trust G, then we can form a node-pair graph G^2 in which each node represents an ordered pair of nodes of G as depicted in Figure 5. A node (A,B) of G^2 points to a node (C,D) if, in G, A points to C and B points to D. Similarity scores are symmetric, so for clarity we draw (A,B) and (B,A) as a single node A,B (with the union of their associated links) [JW02].

SimRank [JW02] is a popular iterative fixed-point algorithm that computes similarity scores for node-pairs in G^2. The similarity score for a node υ of G^2 gives a measure of similarity between the two nodes of G represented by υ. Scores can be thought of as flowing from a node to its neighbors. Each iteration propagates scores one step forward along the direction of the links, until the system stabilizes (i.e.,, scores converge). Since nodes of G^2 represents pairs in G, similarity is propagated from pair to pair. Under this computation, two nodes are similar if they are linked by similar nodes.

For each iteration k, iterative similarity functions $sim_k(*,*)$ is introduced. The iterative computation is started with $sim_0(*,*)$ defined as

$$sim_0(A,B) = \begin{cases} 1, & if\ A = B \\ 0, & if\ A \neq B \end{cases} \tag{2}$$

On the $(k+1)$-th iteration, $sim_{k+1}(*,*)$ is defined in special cases as

$$\begin{aligned} sim_{k+1}(A,B) &= 1, && if\ A = B \\ sim_{k+1}(A,B) &= 0, && if\ I(A) = \emptyset\ or\ I(B) = \emptyset \\ sim_{k+1}(A,B) &= 0, && if\ O(A) = \emptyset\ or\ O(B) = \emptyset \end{aligned} \tag{3}$$

$I(A)$ is the set of in-neighbors of A while $O(A)$ specifies the set of A's out-neighbors. Individual in-neighbors are denoted as $I_i(A)$, for $1 \leq i \leq |I(A)|$, and individual out-neighbors are denoted as $O_i(A)$, for $1 \leq i \leq |O(A)|$. $sim_{k+1}(*,*)$ is computed from $sim_k(*,*)$ in the general case as follows:

$$sim_{k+1}(A,B) = \frac{w}{|I(A)||I(B)|} \sum_{A' \in I(A)} \sum_{B' \in I(B)} sim_k(A',B') \tag{4}$$

where $I(X)$ denotes the set of nodes linking to X (in-neighbors); if $I(A)$ or $I(B)$ is empty, then $sim_{k+1}(A,B) = 0$ by definition. For a node pair with $A = B$ we simply let $sim_{k+1}(A,B) = 1$. w is a constant between 0 and 1. can be thought of either as a confidence level or a decay factor. Consider a simple scenario where user X has two relations with users M and N, so we conclude some similarity between M and N. The similarity of X with itself is 1, but we probably do not want to conclude that $sim(M,N) = sim(X,X) = 1$. Rather, we let $sim(M,N) = w \times sim(X,X)$ meaning that we are less confident about the similarity between M and N than we are between X and itself. This formulas is alternately computed in iterations until the resulting similarity values converge. The structural similarity method has time complexity of $O(n^3)$ and space complexity of $O(n^2)$. We enhanced the algorithm to achieve the time complexity of $O(m*n)$ and space complexity of $O(n+m)$ by using the following heuristics [YLL10].

- As the similarity score can be seen as a random walker defined on a node-pair graph G^2 depicted in Figure 5 (b), the walker may wander into an enclosed subsection of the entire graph which has no out-link so that it will get stuck in the small subgraph with no possibility to return outside. The aforementioned scenario is associated with the fact that the graph is not strongly connected. A technique termed "teleportation" is used to make the graph irreducible and solve this problem.

- We represent similarity equations in a matrix form and employ a sparse storage scheme.

- The similarity matrix often contains an extremely large fraction of non-zeros entries whose values are almost 0 after several iterations. These small similarity

values require a significant amount of storage space with less practical information. We devised a pruning technique to eliminate impractical almost zero similarities by setting a threshold for each iteration. This dropping will also decrease the redundant similarity computations and space per iteration.

- For the similarity computation to be I/O-efficient, the adjacency matrix needs to be preordered, which requires off-line precomputation to minimize the bandwidth at query time. Therefore, A reordering technique is used, which not only speeds up the convergence rate, but achieves I/O efficiency as well.

The details of these techniques are far beyond the scope of this paper. In the next section, we describe the analysis and evaluation of the trust methods.

4. Analysis

The main challenge in this work is to quantify trust only on the basis of the observed communication behavior (a portion of the interactions between users). To understand how the different design choices perform in prediction of the trust links, we applied our methods to conduct an in-depth analysis on a large Twitter dataset. We compared the three webs of trust by applying the four trust propagation methods to see which yields the best results using the leave-one-out technique (a machine learning evaluation technique) and different metrics to show the efficiency and accuracy of each method in prediction of the trust relationships. We start by explanation of our data gathering approach.

4.1 Data Gathering and Crawling Algorithm

As a basis for evaluating our proposal, we first need data to evaluate. Extensive work has been conducted on top of online social networks and in many cases a partial data set is used. There are several reasons for this. First of all, it is hard to get a complete data set directly from the the online social network providers because social data is a very valuable asset and is protected by privacy regulations/laws. Secondly, it is a great challenge for crawlers to collect this huge amount of data from dynamic and customized pages. Moreover, rate limiting is enforced by most providers, preventing crawlers from making many requests within a short period of time. Finally, many users choose not to reveal their information to strangers because of privacy concerns.

An online social network can be modeled as a graph with users as nodes and the relationship between users as links. The crawling of the social graph starts from an initial node and proceeds iteratively. In every operation, we visit a node and discover all its neighbors. There are many ways, depending on the particular sampling method, in which we can proceed. The process for crawling this social graph and gathering a partial data set can be outlined as follows [YLW10]:

Algorithm 4.1: NODE SELECTION ALGORITHM(*seeds*)

```
1  Put seeds (starting nodes of the crawl) into a queue
2  Select a node from the queue
3  Crawl the node
4  Add the neighbors of the crawled node into the queue
5  Go to Step 2 or terminate if stop conditions are met
```

The gathered data set is decided by the following three factors: 1) choice of seeds as the starting point of a crawl, 2) node selection algorithm that decides which node to select from the crawling queue, and 3) size of the crawled subgraph, which is subject to real world resource constraints such as network bandwidth, time, machines, and the rate limits enforced by online social networks providers.

These factors may introduce biases towards high degree nodes and further contaminate or even skew the results. However, it has been widely documented that social networks have the properties of small world networks, where lots of nodes are tightly coupled together within a few hops of each other [Wat03]. The small world effect of online social networks makes the choice of seeds less critical [YLW10]. Therefore, node/link coverage (the number of nodes/links seen by the crawler versus the number of nodes/links in the graph) is not sensitive to the number of seeds neither to the degree of seeds. Moreover, crawling a small portion of the network is sufficient to reveal most nodes/links. It is a strong sign of the small world phenomenon.

There are several widely used node selection algorithms [YLW10], e.g., the *BFS (Breadth First Search)* algorithm, which simply selects the first item in the queue, the *Greedy* algorithm, which selects the node with the largest degree in the queue, or the *Random Walk* algorithm, which selects a node in the queue with probability proportional to its degree. Therefore, the probability of moving from a node u to its neighbor (the transition probability) is $\frac{1}{u_{degree}}$. However, these algorithms lead to samples that not only are biased towards high degree nodes, but also do not have provable statistical properties. We used an unbiased algorithm, the *Metropolis-Hastings Random Walk (MHRW)*, for our crawling purpose [GKBM10]. This algorithm obtains a uniformly distributed random sample of nodes by appropriately modifying the transition probabilities of the random walk. Pseudocode 4.2 shows the process. in every iteration of MHRW, at the current node u, the algorithm randomly selects a neighbor v and move there with probability $min(1, \frac{u_{degree}}{v_{degree}})$. It always accepts the move towards a node of smaller degree, and reject some of the moves towards higher degree nodes. As a result, the bias of RW towards high degree nodes is eliminated.

Algorithm 4.2: NODE SELECTION ALGORITHM(*seeds*)

```
queue ← seeds
while stopping criterion not met
do { u ← queue.GET()
     CRAWL(u)
     while true
     do { Select node v uniformly at random from
          neighbors of u
          Generate uniformly at random 0 ≤ p ≤ 1
          if p ≤ u_degree / v_degree
          then { queue.ADD(v)
                 return
```

In the rest of this section, we give the details about our crawling procedure on Twitter.

4.2 Twitter Crawling

Twitter API was used to gather the data, which allows developers to consume different types of data that Twitter exposes, such as user profiles, status updates and follower information. It should be noted that while the users that posted statuses are clearly currently active, the list of users obtained in successive steps may not have been active.

We selected the first twenty most active users among places where users have most tweeted [4] as the seeds.

The seeds were gathered in GMT:+1, +9, -8, -7, -6, -5, 0 corresponding to places, where the tweet counts were greater than one million. The public timeline command (API functions provided by Twitter) was used to sample the most active users. Twitter continually posts a series of twenty most recent status updates. The status updates in the timeline dataset are presumably a random snapshot of currently active users. Samples were made by retrieving the public timeline and extracting the set of users associated with the statuses in the timeline; then, details of these users were collected.

We gathered detailed information on the users and the list of users each of them were following. The constraint on the number of queries that we could issue in a day was the key-limiting artifact in the reach of our crawl. The use of Twitters API is rate limited. This means that every user is limited to perform a number of API calls per hour. The rate limit defaults to 150 public/unathenticated calls and 350 authenticated calls per hour and per IP. We used the authenticated calls for getting the list of followers, followees, and the one hundred last retweets for each user; while the unauthenticated used for checking the user's protection status and number of followers

[4] http://www.socialnetworkingsandiego.com/social-networking/twitter-as-a-marketing-media/

and followees to see if the user is a broadcaster or a celebrity. For this reason we were required to separate the crawling process across 8 different proxy servers as well as one main server, as a result, we were able to send 3150/hour authenticated API request and 1350 unauthenticated API request.A small portion of users were protected, which does not hurt the node/link coverage of crawlers especially for large social graphs.

Over a period of one month we crawled Twitter information streams of more than 20,000 users. Together, there are 144,962 followers-followees relations, 23,280 retweeted messages, and 50,713 favorited messages. Therefore, theses networks are very sparse. After gathering the data set, we applied several trust inference algorithms on it. In the next section, we explain about these algorithms in more details.

4.3 Evaluation

To validate our proposal and explore how the exploitation of various features influences the characteristics of the trust values generated by the different methods, we evaluated the trust propagation methods. The methods discussed in this paper vary in three design dimensions: (i) they use implicit indication of trust as opposed to explicit trust ratings, (ii) the emphasis is on efficiency and simplicity and (iii) the data sources exploited. In the first step, we generated three graphs (each one with 22,830 nodes) from the Twitter dataset based on the three assumptions about trust indicating behaviors, which are called Followers-followees, Retweet, and Favorite webs of trust. These graphs are represented as FF, RE, and FA adjacency matrices respectively that are defined below:

$$
\begin{aligned}
FF[i,j] &= \begin{cases} 1, & \text{if user } i \text{ follows user } j \\ 0, & \text{otherwise} \end{cases} \\
RE[i,j] &= \begin{cases} \frac{n_{re}}{n_t}, & \text{if user } i \text{ retweeted user } j \\ 0, & \text{otherwise} \end{cases} \\
FA[i,j] &= \begin{cases} \frac{n_{fa}}{n_f}, & \text{if user } i \text{ favorited user } j \\ 0, & \text{otherwise} \end{cases}
\end{aligned}
\tag{5}
$$

where n_{re} is the number of times user i has retweeted user j and n_t is total number of user i's retweets. Likewise, n_{fa} is the number of times user i has favorited user j's tweets and n_f is total number of user i's favorites. Matrices FF, RE, and FA have 144,962, 18,882, and 87,172 none-zero values and there are 470,928, 38,908, and 340,452 paths of length two in each one respectively.

Some of the statistical properties of the three webs of trust (Followers-Followees, Retweet, and Favortie) generated by the three different users' behavior (following, retweeting, and favoriting) are provided in Table 4.3. The total number of crawled users is 22830. Among them one user is both follower and followee of himself, 34 users were found that have retweeted themselves, and 46% of the favorite data was useless because 5,246 users have made their own tweets favorited. These cases were corrected first in the evaluation process.

Table 1. Some statistical properties of the dataset. The adjacancy matrices corresponding to the Followers-Followees, Retweet, and Favorite webs of trust are labeled as FF, RE, and FA respectively.

Statistics	**FF**	**RE**	**FA**
Number of elements	144962	23280	50713
Number of links	144962	12947	12301
Number of non-zero elements in the diameter	1	34	5246
Total number of paths of length two	2219461	25359	14881
Average number of paths of length two	4.81	1.54	1.56
Maximum number of paths of length two	62	9	6

Then, in the second step, we applied the four trust propagation methods: simple-transitivity, weighted-transitivity, golbeck-transitivity, and structural-similairty on these webs of trust to predict some missing links and make them less sparse. We used the leave-one-out technique (a machine learning evaluation technique) for the evaluation, which involves hiding one link in the graph and then trying to predict it using each of the propagation methods. The methods are compared according to the following metrics:

- Coverage: The percentage of currently available links that can be predicted using the propagation method.
- Triadic closure: The percentage of all the paths of length two in the graph that they would eventually close by the transitivity. In other words, the third closing link for them exists already in the graph.
- Mean absolute error (MAE): The weight of the newly predicted link is compared against the original weight of the hided link. The average of the prediction error over all links is then calucated.

The results of the evaluation are presented in following.

4.4 Results

We applied the four different methods for trust propagation: simple-transitivity, weighted-transitivity, golbeck-transitivity, and structural-similarity on the three trust graphs. The results are compared according to the three different metrics: triadic closure, coverage, and MAE and summarized in Table 3. We used the average number of paths of length two in each graph as the threshold value for the weighted-transitivity method. As it is shown in Figure 3, this is the best choice since the coverage decrease sharply when we increase the threshold value.

The results show that the weighted-transitivity method does not give a better result than the simple-transitivity method on any of the trust graphs. The coverage is less and the error is high. The weighted-transitivity method gives less error than the golbeck-transitivity method. The coverage is the same for both of them. This shows

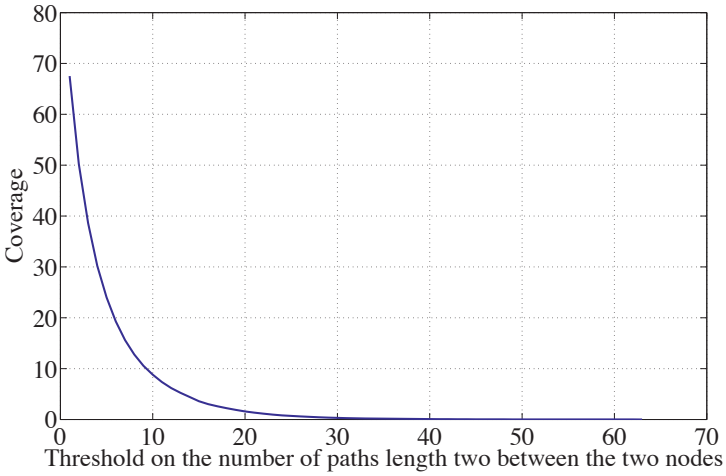

(a) Coverage vs. threshold in the Followers-followees web of trust

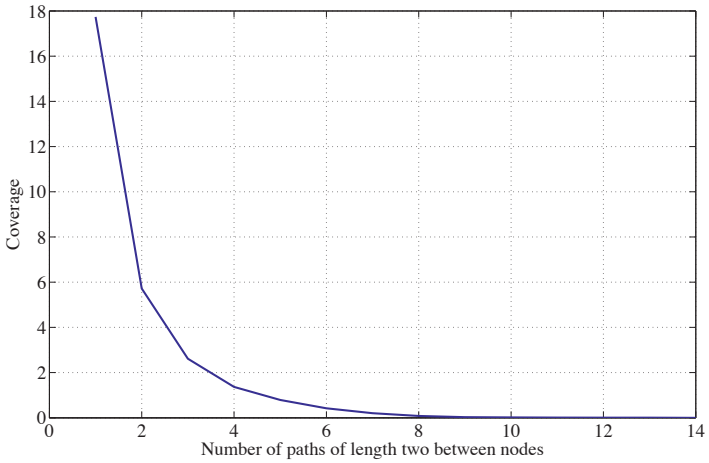

(b) Coverage vs. threshold in the Retweet web of trust

Figure 3. Some of the properties of the generate webs of trust

that consideration of the reliability of the recommender (the trust of truster in the recommender) is an important issue. The structural-similarity method on the trust graph generated by users' retweeting behavior gives the best result. The coverage is 99.96% meaning that we are able to predict almost every link in the graph and the error (MAE) is only 3.69% meaning that we are able to do the prediction accurately. As we discussed earlier, this method is very efficient as well with time complexity of O(m*n) and space complexity of O(n+m).

Table 2. tab:results

Method & Metric	**FF**	**RE**	**FA**
Triadic closure	21.22	12.26	12.16
Coverage for simple-transitivity	67.52	15.67	16.39
Coverage for weighted-transitivity	24.00	4.86	5.53
Coverage for golbeck-transitivity	24.00	4.86	5.53
Coverage for structural-similarity	90.87	99.96	99.98
MAE for weighted-transitivity	83.37	14.87	15.82
MAE for golbeck-transitivity	85.03	22.04	23.43
MAE for structural-similarity (C=0.8)	36.89	3.69	8.97

5. Related Work

There are many research work conducted on Twitter with different aims and scopes regarding the information overload problem. Most research initiatives study followee recommendations [AGHT11] or detecting spammers [BMRA10, YRS+09, SKV10, Wan10, LCW10, GAB10, CMP11]. Yet, little research has been done on modeling and inferring the trust relations among the users particularly using implicit information. The various approaches to trust inference can be summarized as: statistical patterns and clustering techniques [AEG+10], trust ontology [AHTS10], PageRank [NHL10], semantic web [Gol06], Bayesian networks [KG07], Fuzzy logic [SS02], spreading activation models [ZL05], game theory and social network measures [Bus98]. An important characteristic of Twitter is its real-time nature. Hence, the efficiency of such algorithms is very important, but has been ignored in the research. In this section, the existing work are compared according to the information source and algorithm that they use for trust prediction. We refer to following work among the vast amount of research work on trust modeling for social networks.

Huberman et al. [HRW09] find that Twitter users have a very small number of friends compared to the number of followers and followees they declare. This implies the existence of two different networks: a very dense one made up of followers and followees, and a sparser and simpler network of actual friends (those who have interactions with each other, e.g.,, they have sent direct messages to each other). The latter proves to be a more influential network in driving Twitter usage since users with many actual friends tend to post more updates than users with few actual friends.

Sibel et al. [AEG+10] present measures of trust based in social networks. The basis of their approach is an assumption that trust results in communication behavior patterns that are statistically different from communication between random members of a network. The proposed measure of who-trusts-whom relation in the network relies on detecting statistically significant patterns of the trust-like behavior and they validated these measures on a Twitter network data. Two types of trust were identified in this study: (i) conversational trust, the basis for measuring trust is the length and balance of conversations between two nodes and (ii) propagation trust, the metric is based on the percentage of tweets sent by node *A* that node *B* retweets. Theses measures are

based on the assumption that continued information exchange between members of a community can enhance trust in their relationships and receiving information that is believed to be true enhances the trust of the receiver in the sender. The conversational trust is symmetric, but the propagation trust is not, because node A may not trust node B, even if B retweets all tweets of A. The authors conjectured that trust is the foundation of communities, and that it should be possible to discover communities in the Twitter network by identifying clusters whose members trust each other. To test this conjecture, they analyzed the tweets and created communities based on conversational and propagation trust. The resulting two trust-graphs have similar structure, having roughly the same number of communities, as well as a very similar average community size. The trust-based communities created from conversational and propagation trust have a similarity higher modularity than could be expected for random graphs of the same size and node degree distribution. This result confirms that the trust-based communities capture similar relationships. Our work also confirms their findings that retweet behavior is a good indication of trust.

Anantharam et al. [AHTS10] developed a general ontology of trust that is independent of any specific domain and discussed how concepts in their ontology can be used in the context of Twitter as an application scenario. They define two types of trust called referral trust when one user sends another users tweet and functional trust when one user follow another user.

Noordhuis et al. [NHL10] applied PageRank (the Google's method for measuring the relative importance of a URL) to Twitters social graph of users and their followers to determine users of importance. A similar approach is introduced [MM10] by Moh and Murmann, which uses the calculated value as the credibility of users for detection of spammers.

Golbeck [Gol06] introduces an approach to integrate trust with annotations in Semantic Web systems. Then, she presents an application, FilmTrust, that combines the computed trust values with the provenance of other annotations to personalize the website. The FilmTrust system uses trust to compute personalized recommended movie ratings and to order reviews. In another paper, Kuter and Golbeck [KG07] propose to model the trust network as a Bayesian network and evaluate their proposal on the *FilmTrust* social network. Therefore, they also use explicit trust information.

The REGRET reputation system proposed by Sabater and Sierra [SS02] represent trust and reputation as linguistically fuzzy concepts, where membership functions describe to what degree an agent can be described as e.g., trustworthy or not trustworthy. Fuzzy logic provides rules for reasoning with fuzzy measures of this type. This model is a modular trust and reputation system oriented to complex small/mid-size e-commerce environments where social relations among individuals play an important role.

Ziegler and Lausen [ZL05] introduce Appleseed, a local group trust metric based on spreading activation models, designed for computing subjective neighborhoods of most trustworthy peers on the network. The basic intuition of Appleseed is motivated by *spreading activation models* from cognitive science. Spreading activation models simulate human comprehension through semantic memory, and are commonly de-

scribed as models of retrieval from long term memory in which activation subdivides among paths emanating from an activated mental representation.

Buskens [Bus98] in his interesting and unique paper propose a game-theoretic solution. Problematic social situations can be described as trust games with two players and two periods of play. A Trust Game is a one-sided Prisoners Dilemma Game. The restrictiveness of the social conditions under which problematic social situations have to be solved can be reduced by adding the notion of reputation (the possibility of obtaining or spreading information about trustee's trustworthiness) and third parties. This can be explained by the fact that the principal effect of information from third parties is to reduce uncertainty about the behavior of the trustee. This work is an attempt to find an answer to the question: In which way does a truster's level of trust in a trustee depend on his 'local' network position and on the global network structure? In other words, the author evaluates the effects of density, outdegree centrality, and centralization on the level of trust a trustor can have in a trustee using a simulated dataset. He concludes that higher density and outdegree induce more trust. Centralization increases trust if it is 'well organized,' i.e.,, actors who can place more trust are central in the network. Furthermore, he discusses theoretical evidence that the relative importance of density compared to outdegree increases if the trust problem at the dyadic level is large. Finally, he shows that, in many situations, a few simple network measures explain most of the effects of the network structure as a whole.

There are several other work that present trust modeling between a user and a statement on social networks. For example, Richardson et. al. [RAD03] use social networks with trust to calculate the belief a user may have in a statement. This is done by finding paths (either through enumeration or probabilistic methods) from the source to any node which represents an opinion of the statement in question, concatenating trust values along the paths to come up with the recommended belief in the statement for that path, and aggregating those values to come up with a final trust value for the statement. Current social network systems on the Web, however, primarily focus on trust values between one user to another, and thus their aggregation function is not applicable in these systems.

We used the definition of trust in social networks from the work by Goldbeck et al. [GH06b] and three behaviors as the indication of trust. The retweet behavior is also mentioned in the work by Sibel et al. [AEG+10] and Anantharam et al. [AHTS10]. Therefore, we confirm and complement their results in this way. Furthermore, the weighted-transitivity formula is inspired by the work proposed by Golbeck [Gol06].

6. Conclusion and Future work

In this paper, we present quantifiable measures for inferring trust based on users' communication behavior in Twitter and algorithms for predicting trust relationships between individuals that are not directly connected in the trust web. Moreover, we investigated how the different design alternatives influence the accuracy and sparsity of the predicted links in the trust web. Given a large dataset consisting of more than 20,000 user, we generated three different trust webs and applied four different trust prediction/propagation methods. We saw that retweet behavior is the best indication

of trust and structural similarity is the best trust propagation method among them. All in all, we conclude that the inherent and hidden underlying network in Twitter, which we call it Twitter's web of trust can be inferred efficiently and implicitly using the retweet behavior. Besides, This web of trust can be expanded efficiently for new predictions using the structural properties of this graph. The main advantages of our approach are accuracy, efficiency, scalability, and the use of implicit information. The disadvantage is the use of heuristics.

However, there is a lot more information in the behavioral trust graphs than is presented here, and so there are many directions for the future work. We may be able to improve the measures with simple semantic analysis. Efficient algorithms for statistically analyzing the tweets along different dimensions can considerably enhance the behavioral trust measures. In the future work we will further research the contextual information such as the semantics of tweets (e.g., topics and hashtags) and demographical information about users (location, age, ...) and their impact on trust inference. Therefore we plan to explore whether knowledge regarding the tweets' contextual information and users' demographical information can further leverage trust inference quality.

References

[AEG+10] S. Adali, R. Escriva, M.K. Goldberg, M. Hayvanovych, M. Magdon-Ismail, B.K. Szymanski, W.A. Wallace, and G. Williams. Measuring behavioral trust in social networks. In *IEEE International Conference on Intelligence and Security Informatics (ISI 2010)*, pages 150–152. IEEE, 2010.

[AGHT11] F. Abel, Q. Gao, G.J. Houben, and K. Tao. Analyzing user modeling on twitter for personalized news recommendations. *User Modeling, Adaption and Personalization*, pages 1–12, 2011.

[AHTS10] P. Anantharam, C.A. Henson, K. Thirunarayan, and A.P. Sheth. Trust model for semantic sensor and social networks: A preliminary report. In *Proceedings of the IEEE 2010 National Aerospace and Electronics Conference (NAECON)*, pages 1–5. IEEE, 2010.

[BMRA10] F. Benevenuto, G. Magno, T. Rodrigues, and V. Almeida. Detecting spammers on twitter. In *Proceedings of the 7th Annual Collaboration, Electronic messaging, Anti-Abuse and Spam Conference (CEAS)*, 2010.

[Bus98] V. Buskens. The social structure of trust. *Social Networks*, 20(3):265–289, 1998.

[CMP11] C. Castillo, M. Mendoza, and B. Poblete. Information credibility on twitter. In *Proceedings of the 20th international conference on World wide web*, pages 675–684. ACM, 2011.

[DKG+05] L. Ding, P. Kolari, S. Ganjugunte, T. Finin, and A. Joshi. Modeling and Evaluating Trust Network Inference. Technical report, Maryland Univ Baltimore, 2005.

[GAB10] D. Gayo-Avello and D.J. Brenes. Overcoming spammers in twitter-a tale of five algorithms. In *1st Spanish Conference on Information Retrieval, Madrid, Spain*, 2010.

[GH06a] J. Golbeck and J. Hendler. Filmtrust: Movie recommendations using trust in web-based social networks. In *Proceedings of the IEEE Consumer communications and networking conference*, volume 96. Citeseer, 2006.

[GH06b] Jennifer Golbeck and James Hendler. Inferring binary trust relationships in web-based social networks. *ACM Trans. Internet Technol.*, 6:497–529, November 2006.

[GKBM10] M. Gjoka, M. Kurant, C.T. Butts, and A. Markopoulou. Walking in facebook: A case study of unbiased sampling of osns. In *INFOCOM, 2010 Proceedings IEEE*, pages 1–9. Ieee, 2010.

[GKRT04] R. Guha, R. Kumar, P. Raghavan, and A. Tomkins. Propagation of trust and distrust. In *Proceedings of the 13th international conference on World Wide Web*, pages 403–412. ACM Press New York, NY, USA, 2004.

[Gol06] J. Golbeck. Combining provenance with trust in social networks for semantic web content filtering. *Provenance and Annotation of Data*, pages 101–108, 2006.

[HRW09] B.A. Huberman, D.M. Romero, and F. Wu. Social networks that matter: Twitter under the microscope. *First Monday*, 14(1):8, 2009.

[JW02] G. Jeh and J. Widom. SimRank: a measure of structural-context similarity. In *Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 538–543. ACM Press New York, NY, USA, 2002.

[KG07] U. Kuter and J. Golbeck. Sunny: A new algorithm for trust inference in social networks using probabilistic confidence models. In *Proceedings of the National Conference on Artificial Intelligence*. Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press; 1999, 2007.

[LCW10] K. Lee, J. Caverlee, and S. Webb. Uncovering social spammers: social honeypots+ machine learning. In *Proceeding of the 33rd international ACM SIGIR conference on Research and development in information retrieval*, pages 435–442. ACM, 2010.

[MA04] P. Massa and P. Avesani. Trust-aware collaborative filtering for recommender systems. *On the Move to Meaningful Internet Systems 2004: CoopIS, DOA, and ODBASE*, pages 492–508, 2004.

[MBKM07] R. Morselli, B. Bhattacharjee, J. Katz, and M. Marsh. Exploiting approximate transitivity of trust. In *Broadband Communications, Networks and Systems, 2007. BROADNETS 2007. Fourth International Conference on*, pages 515–524, 2007.

[MM10] T.S. Moh and A.J. Murmann. Can you judge a man by his friends?-enhancing spammer detection on the twitter microblogging platform using friends and followers. *Information Systems, Technology and Management*, pages 210–220, 2010.

[NHL10] P. Noordhuis, M. Heijkoop, and A. Lazovik. Mining twitter in the cloud: A case study. In *Cloud Computing (CLOUD), 2010 IEEE 3rd International Conference on*, pages 107–114. IEEE, 2010.

[QHC07] D. Quercia, S. Hailes, and L. Capra. Lightweight Distributed Trust Propagation. In *Data Mining, 2007. ICDM 2007. Seventh IEEE International Conference on*, pages 282–291, 2007.

[RAD03] M. Richardson, R. Agrawal, and P. Domingos. Trust management for the semantic web. *The Semantic Web-ISWC 2003*, pages 351–368, 2003.

[SKV10] G. Stringhini, C. Kruegel, and G. Vigna. Detecting spammers on social networks. In *Proceedings of the 26th Annual Computer Security Applications Conference*, pages 1–9. ACM, 2010.

[SOM10] Takeshi Sakaki, Makoto Okazaki, and Yutaka Matsuo. Earthquake shakes twitter users: real-time event detection by social sensors. In *Proceedings of the 19th international conference on World wide web*, WWW '10, pages 851–860, New York, NY, USA, 2010. ACM.

[SS02] J. Sabater and C. Sierra. Reputation and social network analysis in multi-agent systems. In *Proceedings of the first international joint conference on Autonomous agents and multiagent systems: part 1*, pages 475–482. ACM, 2002.

[TA12] M. Tavakolifard and K. C. Almeroth. Trust 2.0: Who to believe in the flood of online data? In *Proceedings of the International Conference on Computing, Networking and Communications (ICNC'12)*, 2012.

[Wan10] A.H. Wang. Don't follow me: Spam detection in twitter. In *Security and Cryptography (SECRYPT), Proceedings of the 2010 International Conference on*, pages 1–10. IEEE, 2010.

[Wat03] D.J. Watts. *Small worlds: the dynamics of networks between order and randomness*. Princeton Univ Pr, 2003.

[WBS08] F.E. Walter, S. Battiston, and F. Schweitzer. A model of a trust-based recommendation system on a social network. *Autonomous Agents and Multi-Agent Systems*, 16(1):57–74, 2008.

[YCB+02] Y. Yang, ACT Canberra, L. Brown, S. Wales, ACT ADFA, E. Lewis, and V.A. Melbourne. W3 Trust Model: Evaluating Trust and Transitivity of Trust of Online Services. In *International Conference on Internet Computing*, pages 354–362, 2002.

[YLL10] W. Yu, X. Lin, and J. Le. A space and time efficient algorithm for simrank computation. In *Web Conference (APWEB), 2010 12th International Asia-Pacific*, pages 164–170. IEEE, 2010.

[YLW10] S. Ye, J. Lang, and F. Wu. Crawling online social graphs. In *2010 12th International Asia-Pacific Web Conference*, pages 236–242. IEEE, 2010.

[YRS+09] S. Yardi, D. Romero, G. Schoenebeck, et al. Detecting spam in a twitter network. *First Monday*, 15(1), 2009.

[ZL05] C.N. Ziegler and G. Lausen. Propagation models for trust and distrust in social networks. *Information Systems Frontiers*, 7(4):337–358, 2005.

PAPER D

Situation-based Trust Adjustment by Conditional Trust Reasoning

Mozhgan Tavakolifard, Pinar Ozturk

Proceedings of the Networking and Electronic Commerce Research Conference (NAEC)
Riva del Garda, Italy, Oct 2011

SITUATION-BASED TRUST ADJUSTMENT BY CONDITIONAL TRUST REASONING

Mozhgan Tavakolifard,[1] Pinar Ozturk,[2]

[1] *Centre for Quantifiable Quality of Service in Communication Systems*
Norwegian University of Science and Technology
mozhgan@Q2S.ntnu.no

[2] *Department of Computer and Information Science,*
Norwegian University of Science and Technology
pinar@idi.ntnu.no

Abstract World Wide Web a perfectly distributed and uncontrolled medium. However, the use of this technology to its utmost boundaries requires robust and efficient trust mechanisms. This paper describes a context-sensitive trust management system that categorizes trust situations with respect to the experiences of a trustee. If the trustee is familiar with the trustee, the trust judgment relies on case-based reasoning. Context- sensitivity is maintained in the description of the current and past situations that are compared. When the truster does not have any previous interaction with the trustee, a rule-based reasoner is used to assess the trustability of the trustee on the basis of available recommendations of third parties. The rules are automatically extracted from the history and encoded as conditions connecting contextual information to trust judgements. Through the use of subjective logic, this method explicitly incorporates uncertainty, thereby making it suitable in situations of partial ignorance and imperfect information. We evaluated our proposal using a large-scale real dataset.

1. Introduction

The World Wide Web is not only an information space, but also a medium for commerce and social interactions for citizens all over the world. However, despite that we all celebrate and enjoy the egalitarian and free nature of the WWW, problems regarding the security of information and services have started to manifest themselves. The human nature is biased to take advantage when the occasion offers itself. Hence, a major problem with such an open and distributed spaces as the WWW is that users lack sufficient information about the quality and security of the e-services and their providers. Conventional security mechanisms cannot handle the trust phenomenon in the way the new information systems would require. Therefore, the growth of services such as online transactions and information exchange is conditioned on the development and implementation of new trust management models.

Trust is context-locked, meaning that a trust value is associated with certain peculiarities pertinent to a situation. A typical example is that a person may trust her financial advisor about an investment analysis, but normally not in health-care related issues. Context is defined as "any information that can be used to characterize the situation of an entity which may be a person, a place, or an object which is considered relevant for the interaction between the user and application, including the user and the application themselves" [BD05]. A system is context-aware if "it uses context to provide relevant information and/or services to the user, where relevancy depends on the user's task" [DA00].

The need to consider a plurality of aspects as a basis for trust decisions has been recognized for a long time [RGPB06], yet the context issue has largely been neglected by the trust research community [Mar94, CF02, SS05, RHJ05]. A few exceptions exist. In [NWvSL07], it is shown that extension of trust models with context representations can reduce complexity in the management of trust relationships and improve the recommendation process. In [HY06], the possibility to infer trust information in context hierarchies is discussed, and in [RP07] and [TLU06], it is claimed that it is possible to learn policies/norms at runtime and provide protection against changes of identity and first time offenders. Hence, apart from the occasional work which is elaborated in section 4, the relationship between the notions of context and trust has not been given the attention it deserves.

TMMs set out with an initial trust value and modify it over time to provide a more accurate trust value. Typically TMMs use a default initial value which is context-neutral. CMF can provide a more informed (i.e. context-sensitive) initial value.

In our work a situation is defined in terms of a set of contextual attribute-value tuples $(X_1, X_2, \ldots, X_n)$ where X_i may be, for example, *location: Trondheim*. Our previous work [THO09] presents a context management framework (CMF) that can be combined with the existing trust management models (TMM) to extend their capabilities towards efficient modeling of the situation-aware trust through the following two functions:

1 Initialization of the trust values in unknown situations or for unknown trustees when there is no available information. In open systems such as ad-hoc networks, the agents are distributed across various platforms and can join or leave the system at their own will, which requires the assignment of estimated initial trust values. Most TMMs simply consider a default trust value for trust bootstrapping. In such a case, a high value is risky while a low value carries the risk that new agents might be ignored completely. CMF can help a TMM to bootstrap by providing an estimation of trust values based on similar situations or similar trustees previously observed (see figure 1-a).

2 Adjustment of the output of TMMs based on the situation, imposing situation awareness to the TMMs (see figure 1-b). The inability to take the situation into account limits the practical use of current trust models in domains where the agents perform diverse tasks in a highly dynamic environment.

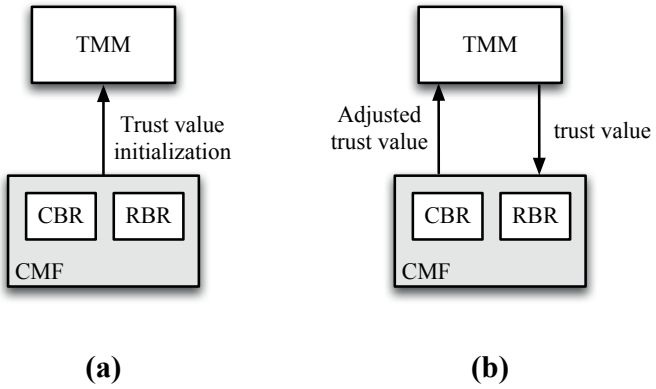

Figure 1. Scope and interconnection between the context management framework (CMF) and the trust management model (TMM). a) Initialization of the trust value in unknown situations or for unknown trustees. b) Adjustment of the output of a TMM based on the underlying situation.

The CMF is empowered by similarity-based and rule-based reasoning capabilities (depicted respectively as CBR and RBR modules in Figure 1). The CBR component is responsible for initialization of the trust values in unanticipated situations, while the RBR component is responsible for both initialization of the trust values for unknown trustees and adjustment of the trust judgements by TMMs according to the underlying situation. The CBR module was discussed thoroughly in [THO09] and was evaluated for a specific trust model that uses subjective logic. In that work, the trust model was extended to incorporate contextual factors. This paper focuses on the RBR module which uses subjective logic (described in section 3) for knowledge representation and reasoning, and explains how trust values are adjusted based on the underlying situation. Rules in the rule base are encoded as conditionals connecting contextual information to trust judgements. Through the use of subjective logic this method explicitly incorporates uncertainty, thereby making it suitable in situations of partial ignorance and imperfect information. We evaluated our proposal using a large-scale real dataset.

The rest of this paper is organized as follows. In section 2 we explain the CMF in more detail. Section 3 briefly explains the subjective logic. The proposed model for the rule-based trust inference is described in section 4. In section 5 we explain the application of our proposed model to recommender systems. Subsequently, in section 3, we present the evaluation plan and the obtained results. Section 4 provides an overview of the related research. Finally, we close by our concluding remarks and future research directions in section 5.

2. The Context Management Framework

We consider two approaches to the inference underlying the functionalities of the CMF: similarity-based reasoning and rule-based inference, depicted respectively as

		Trustee: Familar	Trustee: Unfamiliar
Situation	Familiar	None	RBR
	Unfamiliar	CBR	Default

Table 1. Initialization of the trust value.

case-based reasoner (CBR) and rule-based reasoner (RBR) modules in figure 1. The former provides the first role of the CMF, i.e. initialization of the trust values in unanticipated situations while the latter is responsible for both roles of CMF, i.e., initialization of the trust values for unknown trustees and adjustment of the trust values based on the underlying situation. We categorize the decision making on trust initialization in CMF (the initialization function, see figure 1-a) with regard to familiarity with the situation and the trustee, as shown in Table 1

- Familiar situation, familiar trustee: If the truster has previously had interactions with the same trustee in the same situation, then she can immediately use her past experiences to predict the outcome of the new interaction and take a decision on this basis. Therefore, there is no need for initialization.
- Unfamiliar situation, familiar trustee: If the truster has had previous interactions with the trustee, but in different situations, she can still use her past experiences, but should map the old and new situations and make necessary adaptations in order to draw a conclusion. For example, trusters trusting Bob as a good car mechanic would not automatically also trust him in undertaking heart surgeries. However, he may be capable of repairing motorcycles, since repairing cars and motorcycles demand similar knowledge and skills. We have used case-based reasoning (the CBR module) to handle such situation-specificity of trust [THO09], see figure 1 (a). The previous trust values can be revised and reused based on degree of similarity between the new and previous situations.
- Familiar situation, unfamiliar trustee (figure 1 (b)): The truster has previous experiences in the same situations with other trustees. The trust judgment then resorts to a set of domain-specific association rules. We propose a rule-based reasoning component (the RBR module) to handle this situation, which is elaborated in section 4.
- Unfamiliar situation, unfamiliar trustee: CMF does not provide any output to the TMM in this case. TMM uses a default trust value since there is no information to be used for initialization of trust.

Figure 2 shows the decision making process underlying the initialization of trust values.

The RBR module is also responsible for the second functionality of the CMF, i.e. adjustment of the trust value based on the situation. For example, someone has

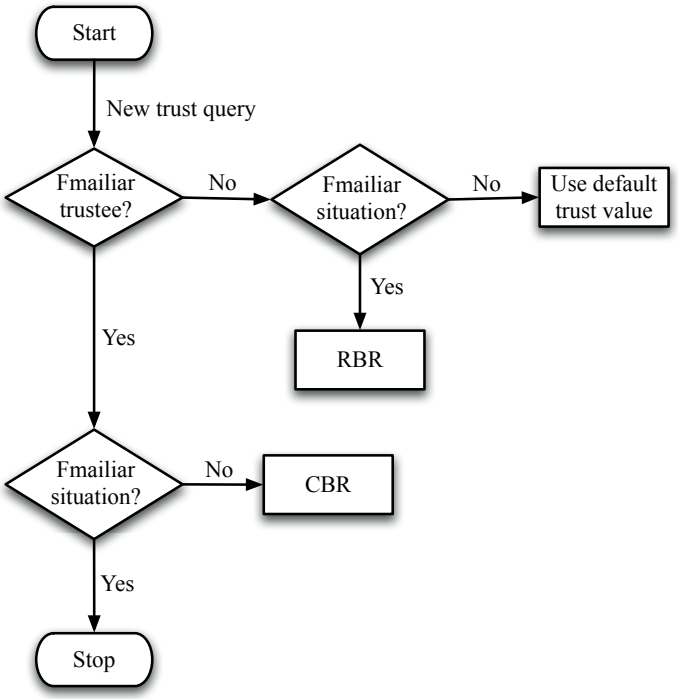

Figure 2. Decision making process in the context management framework for initialization of trust value.

reported an event of fire at university. The situational information has an impact on the trust we assign to this news in the following way. First of all, if the reporting person was at the site of accident, then our trust in the news will be increased. This is because we know that a person normally will have a more precise and correct perception of the situation if he is an eyewitness to the accident. Secondly, if we know that the the person was a journalist, this will increase our trust as well, since we know that journalists usually try to find out the true state of affairs in order to avoid reporting falsehoods.

2.1 Analogical Trust Judgment: Case-based Reasoning Module

The CBR technique [Kol93, AP94] is particularly useful in open and weak domains that lack the complete and certain knowledge and thus needs to exploit experience based knowledge. The fundamental principle of CBR is similar to human analogical reasoning [Gen83, HT97] in the sense of using solutions of past problems to solve the current similar ones. Two main components of a CBR system are the case base storing a number of previously solved cases and the CBR engine that finds and uses

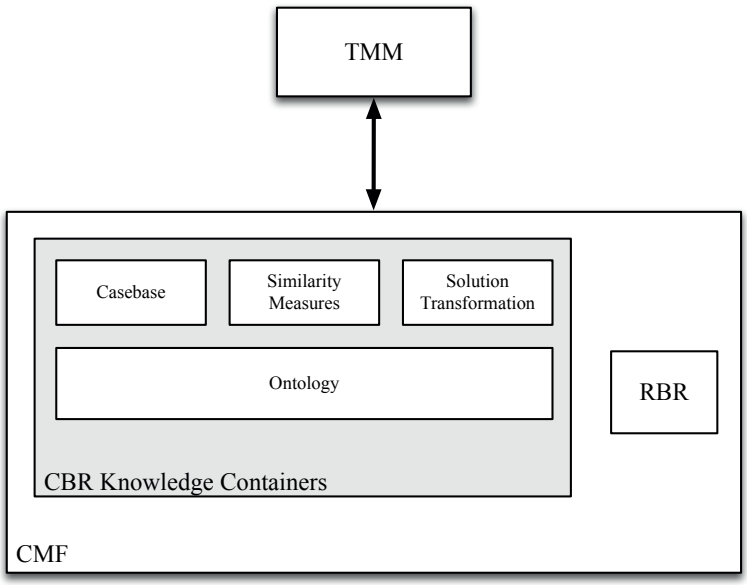

Figure 3. Knowledge containers in case-based reasoner (CBR). TMM: trust management model, RBR: rule-based reasoner, CMF: context management framework.

the previously solved cases (in the case base) in order to solve a new case. A case comprises two parts: a situation/problem description and a solution (only the past cases). In the presented work, a new case is a trust assessment query specifying the truster, trustee, and the other contextual information. Context has been shown to have major influences on remembering and comparing cases. The strong dependency between the context and a powerful memory-retrieval arises most probably from the role context plays in the similarity assessment of two cases (i.e., the new and a past case) [Özt98]. The query is matched with the problem description part of the cases in the case base and the cases are ranked according to their similarity with the query. The retrieved case provides a solution which is the trust value that the truster assigns to the trustee. In [THO09] we consider Subjective Logic as a representation language to represent the TMM and provide details for the solution transformation module.

3. Background: Subjective Logic

Subjective logic is a type of probabilistic logic that allows probability values to be expressed with degrees of uncertainty. Probabilistic logic combines the strengths of logic and probability calculus, meaning that it has the capacity of binary logic to express structured argument models, and it has the power of probabilities to express degrees of truth of those arguments. Subjective logic makes it possible to express uncertainty about the probability values themselves, meaning that it is possible to reason with argument models in presence of uncertain or partially incomplete evidence.

Subjective opinions express subjective beliefs about the truth of propositions,with degrees of uncertainty. A multinomial opinion is defined over $X = \{x_i | i = 1..k\}$, which is a set of exhaustive and mutually disjoint propositions x_i and is denoted by $\omega_X = (\vec{b}, u, \vec{a})$. $\vec{b}$ is a vector of belief masses over the propositions of X, u is the uncertainty mass, and $\vec{a}$ is a vector of base rate values over the propositions of X. These components satisfy $|\vec{b}| = |\vec{a}| = k$, $u + \sum \vec{b}(x_i) = 1$, and $\sum \vec{a}(x_i) = 1$ as well as $\vec{b}(x_i), u, \vec{a}(x_i) \in [0,1]$. $\vec{b}(x_i)$ denotes the belief mass over x_i that represents the amount of positive belief that x_i is true. The uncertainty mass u can be interpreted as the lack of committed belief mass in the truth of any of the propositions of X. In other words, uncertainty mass reflects that the belief owner does not know which of the propositions of X in particular is true, only that one of them must be true. The base rate vector $\vec{a}$ will play a role in determining probability expectation values over X and represents non-informative a priori probability over X before any evidence has been received. Given a frame of cardinality k, the default base rate for each element in the frame is $1/k$, but it is possible to define arbitrary base rates for all mutually exclusive elements of the frame, as long as the additivity constraint is satisfied.

Let $\vec{r}$ be a vector consisting of a number of observations over propositions of X. Then the corresponding opinion will be calculated as the following:

$$\begin{cases} \vec{b}(x_i) = \frac{\vec{r}(x_i)}{W + \sum_{i=1}^{k}} \\ \\ u = \frac{W}{W + \sum_{i=1}^{k} \vec{r}(x_i)} \end{cases} \tag{1}$$

W is a non-informative prior weight, a constant that is suggested to be equal to the cardinality of the frame for an a priori uniform distribution. The probability expectation of multinomial opinions is a vector expressed as a function of the belief vector, the uncertainty mass and the base rate vector. The function $\vec{E}_X$ from X to $[0,1]^k$ is the probability expectation vector over X and is expressed as

$$\vec{E}_X(x_i) = \vec{b}(x_i) + \vec{a}(x_i)u$$

$\vec{E}_X$ satisfies the additivity principle: $\vec{E}_X(\Phi) = 0$ and $\sum_{x \in X} \vec{E}_X(x) = 1$. The base rate vector expresses non-informative *a priori* probability, whereas the probability expectation function expresses the informative *a posteriori* probability.

Two operators are proposed in [Jøs07] to combine multinomial opinions: *cumulative fusion* (denoted by $\oplus$) and *averaging fusion* (denoted by $\underline{\oplus}$). The former is used in cases that opinions are independent (e.g. observations are made in disjoint time periods), while the latter is for dependent opinions (e.g. observations are in the same time period).

In [JG03, JPD05, Jøs08], Jøsang introduces the 'deduction' operator for the Subjective Logic denoted by $\odot$. Let $X = \{x_i | i = 1..k\}$ and $Y = \{y_i | i = 1..l\}$ be frames. Assume that an observer perceives a conditional relationship between the two frames X and Y. Where X plays the role of antecedent (what we have evidence about) and Y will play the role of consequent (about which we want to derive an opinion).

Let $\omega_{Y|X}$ be the set of conditional opinions on the consequent frame Y as a function of the opinion on the antecedent frame X expressed as

$$\omega_{Y|X} = \{\omega_{Y|x_i} | i = 1 \ldots k\}$$

where $\omega_{Y|x_i}$ is defined as opinion about Y given that x_i is TRUE. By using the notation $Y||X$ for conditional deduction, the expression for subjective logic conditional deduction can be defined as:

$$\omega_{Y||X} = \omega_X \odot \omega_{Y|X} \tag{2}$$

where $Y||X$ denotes the consequent opinion Y is derived as a function of the antecedent opinion X together with the conditional opinion $Y|X$. The expression $\omega_{Y||X}$ thus represents a derived value, whereas the expression $\omega_{Y|X}$ represents an input argument.

4. The Proposed Model: RBR module

The RBR module deals with situation-based trust reasoning, when the situation is familiar, but the trustee is unfamiliar. It has the following components [DD98]:

- **Knowledge base:** models the long-term memory as a set of rules.
- **Working memory:** models the short-term memory and contains facts related to the new problem, i. e. both the initially available ones and the ones inferred through firing of the rules.
- **Inference engine:** models reasoning by connecting the facts in the working memory with rules contained in the knowledge base to infer new information.

In this work, situations form the antecedents of the rules, while trust judgments (i.e. values) about unfamiliar trustees comprise the consequents.

In the RBR module (figure 4), the rules contained in the knowledge base represent the long-term memory and are either predefined or learned, based on the previous experiences (in section 5 we give an example of rule learning from past experiences). The rules are represented as conditional opinions. The facts contained in the working memory represent the current situation. A situation is composed of several contexts (e.g. time, location, etc.). Values of the contextual attributes form the antecedents of the conditional opinions. The inference engine compares the facts/situation in the working memory with the antecedents of the rules/conditional opinions to see which rules may fire. Those rules that can fire have their conclusions added to the working memory and the process continues until no other rule match the facts in the working memory. The conclusion part of a rule is represented as a consequent opinion which is an opinion about the trustworthiness of the trustee. Figure 4 shows the opinions in the rule-based model.

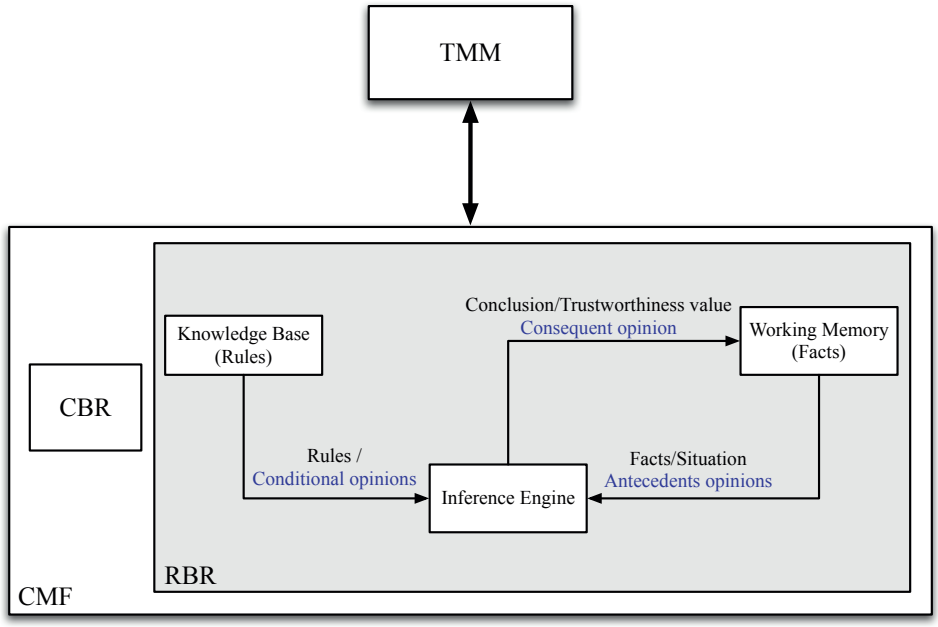

Figure 4. The rule-based reasoner. TMM: trust management model, CBR: case-based reasoner, CMF: context management framework.

4.1 Working Memory

Assume that the situation consists of n contexts $X1$, $X2$, ..., Xn and each context has $k_i, i = 1 \ldots n$ mutually disjoint propositions. Each antecedent opinion is denoted by ω_{Xi}. We assume a discrete trust model with l different trust values[1]. The consequent opinion is denoted by ω_Y where $|Y| = l$ (figure 5).

4.2 Knowledge Base

Conditional opinions ($\omega_{Y|Xi}$) for each context X_i represent the rules in the knowledge base. For a particular context X we have $\omega_{Y|X} = \{\omega_{Y|x_i} | i = 1..k\}$ where $\omega_{Y|x_i}$, represented in subjective logic as a triplet of $(\vec{b}_{Y|x_i}, u_{Y|x_i}, \vec{a}_{Y|x_i})$ where $|\vec{b}_{Y|x_i}| = l$, and $|\vec{a}_{Y|x_i}| = l$. The opinion $\omega_{Y|x_i}$ will be calculated from observation evidence vector $\vec{r}_{Y|x_i}$ according to (1). $\vec{r}_{Y|x_i}(j)$ represents the number of experiences which resulted in y_j when x_i was true and $|\vec{r}|$=l;

[1]Continues trust models can provide input ratings to our system based on the method proposed in [JLC08]

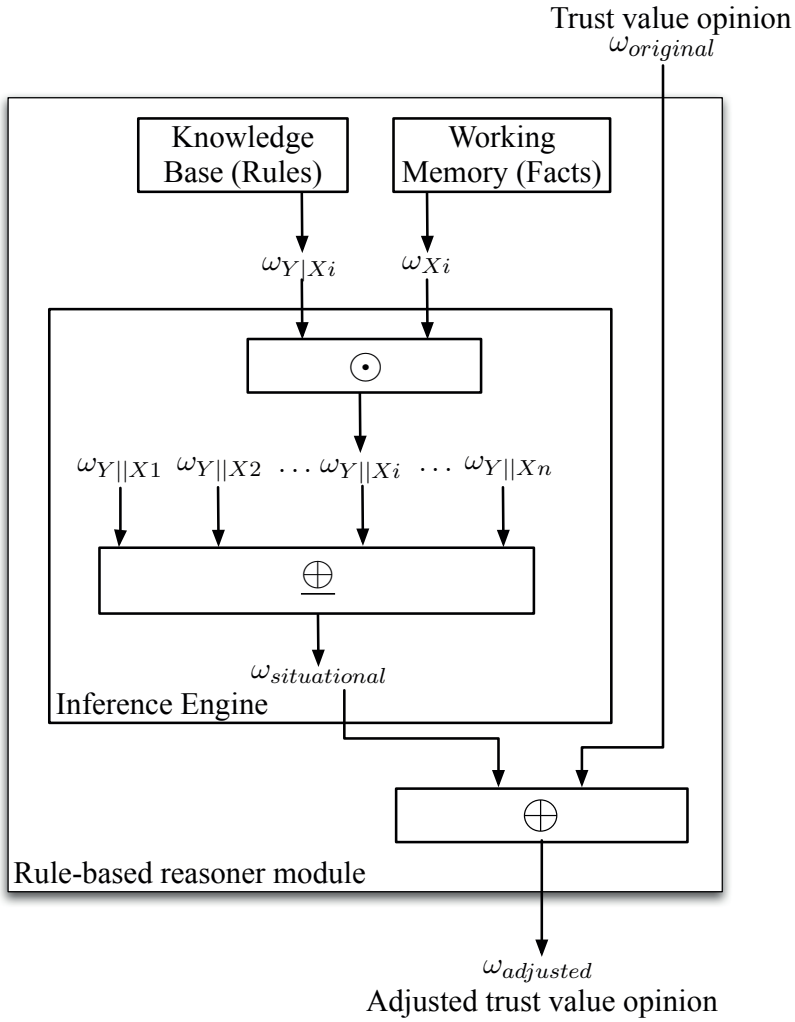

Figure 5. The rule-based reasoner module

$$\omega_{Y|x_i} = \begin{cases} \vec{b}_{Y|x_i}(y_j) = \frac{t_j}{m+c}, j = 1..l \\ u_{Y|x_i} = \frac{c}{m+c} \\ \vec{a}_{Y|x_i}(y_j) = \frac{1}{l}, j = 1..l \end{cases} \tag{3}$$

m is the number of previous experiences in which the proposition x_i for context X is true (that is, context X has the value of x_i) and t_j is the number of those which resulted in trustworthiness level j among them ($\sum_{j=1..l} t_j = m$).

4.3 Inference Engine

In the inference engine, the prediction about trustworthiness $\omega_{Y||Xi}$ in each context Xi is obtained by applying the deduction operator (2) on the antecedent opinion ω_{Xi} and the corresponding rule $\omega_{Y|Xi}$. In order to derive the final trustworthiness opinion, the opinions based on each context should be combined.

$$\omega_{Y||Xi} = \omega_{Xi} \odot \omega_{Y|Xi} \ , \quad i = 1..n$$

The derived opinion for each context X_i is considered as a dependent opinion since all of them are for one particular trustee. Therefore, the *averaging fusion* operator $\underline{\oplus}$ is used to combine them and the result would be the predicted opinion ($\omega_{\text{situational}}$) based on the situation.

$$\omega_{situational} = \omega_{Y||X1} \underline{\oplus} \omega_{Y||X2} \ \ldots \ \underline{\oplus} \ \omega_{Y||Xn}$$

This situational opinion is used as an initial trust opinion for an unknown trustee in situations which are familiar for the truster having had several experiences in those situations. However, if the situation is unfamiliar for the truster, then the CBR module should be used to derive an opinion based on other similar situations which are familiar for the truster, see figure 1(a).

The situational opinion may also be used to adjust the current trustworthy opinion based on the underlying situation, see figure 1(b). We use the *cumulative fusion* operator $\oplus$ to combine the situational opinion $\omega_{\text{situational}}$ with the original opinion ω_{original} to adjust the underlying situation, as they are independent opinions (one is calculated just based on the current situation and the other is calculated based on other factors such as previous experiences or recommendations about the trustee), see figure 1(b) and figure 5.

$$\omega_{ad\,justed} = \omega_{situational} \oplus \omega_{original}$$

5. Application Scenario: Rating Prediction in a Recommender System

In this section, we explain an application of the proposed model for recommender systems where users provide ratings (in a scale of l different levels) for objects. This example is not regarding trust about a trustee, however it is about rating about an object. We want to derive the opinion of a particular user about the rating of a particular object with several features, on the basis of historical ratings, i.e. ratings of the object previously provided by users. The situation consists of two context components: user attributes (e.g. age and gender) and object attributes. Each object attribute is coupled with a rating, the consequent opinion of the user, about that feature of the object. The whole process is described as a pseudo-code in algorithm 4.1. Each record in the history contains a *user*, an *object*, and the *rating* of the user for the object contexts. A query contains the *user* and the *object*.

BuildWorkingMemory procedure computes the antecedent opinions based on the users' attributes. In *BuildKnowledgeBase* procedure, a set of conditional opinions for

each consequent opinion and user attribute are calculated using available historical data. These conditional opinions constitues the ruels capturing the user's rating behavior about each object's feature based on personal characteristics (i.e. user's attributes such as age, gender, ...) and are learned from the history according to (3). In the *InferenceEngine* procedure, the deduction operator (2) is applied on each antecedent opinion and the related conditional opinion in order to predict the user's opinion about the object. The deduced opinions are combined using the *averaging fusion* operator $\underline{\oplus}$ to compute the final predicted opinion about the object.

The predicted opinion should then be converted to a single value in the set of rating levels (e.g. a value in the set $\{1,2,3,4,5\}$ for five-stars ratings). This can be done by assigning a point value v to each rating level i, and computing the normalized weighted point estimate score δ [JLC08]. Assume e.g. l different rating levels ($R = 5$ in our case) with point values evenly distributed in the range [0,1], so that $v(i) = \frac{i-1}{l-1}$. The point estimate rating score is then computed as:

$$\delta = \sum_{i=1}^{l} v(i) \times E_{\omega_{Y||X}}(i) \tag{4}$$

where $E_{\omega_{Y||X}}(i)$ is the probability expectation value of the predicted opinion ω for rating level i. It is calculated according to (3), then the point estimate in the range [0,1] can be mapped to a value in the scale of l levels.

$$predicted\_rating = \lfloor \delta \times (l-1) \rfloor + 1 \tag{5}$$

5.1 Example: MovieLens Recommender System

We explain the RBR module in detail for a particular recommender system called MovieLens [2]. In MovieLens the users provide 5-stars ratings for movies as objects. User attributes are age, gender, zipcode, and occupation [3], while movie attributes are film genres [4]. A film may be attributed to more than one genre. Much richer movie content can be obtained from the Internet Movie Database (IMDB) [5]. The user attributes consist of the antecedent opinions while the rating values for film genres forms the consequent opinions.

5.1.1 Working Memory

Four demographical context attributes (i.e. age, gender, occupation, zipcode) constitute a situation. The corresponding contexts are denoted as XA for age, XG for gender, XO for occupation, and XZ for zipcode (location). Following, we provide the

[2] http://www.grouplens.org/node/73

[3] Occupation list: administrator, artist, doctor, educator, engineer, entertainment, executive, healthcare, homemaker, lawyer, librarian, marketing, none, other, programmer, retired, salesman, scientist, student, technician, writer.

[4] Film genres: unknown, action, adventure, animation, children, comedy, crime, documentary, drama, fantasy, film-noir, horror, musical, mystery, romance, sci-fi, thriller, war, western.

[5] http://us.imdb.com

Algorithm 5.1: RULE BASED REASONER(*history*, *query*)

main
user ← *query.user*
object ← *query.object*
antecedents_set ← BUILDWORKINGMEMORY(*user.attributes*)
conditionals_set ← BUILDKNOWLEDGEBASE(*user.attributes*, *object.features*, *history*)
predicted_opinion ← INFERENCEENGINE(*antecedents_set*, *conditionals_set*)
comment: Conversion of the predicted_opinon to a rating level according to (4) and (5)
for $i \leftarrow 1$ **to** l
do $\begin{cases} \sigma \leftarrow \sigma + \frac{i-1}{l-1} \times \\ (deduced\_opinion.b[i] + deduced\_opinion.a[i] \times deduced\_opinion.u) \end{cases}$
predicted_rating ← $\lfloor \sigma \times (l-1) \rfloor + 1$
return (*predicted_rating*)

procedure BUILDWORKINGMEMORY(*user.attributes*)
for each *attribute* ∈ *user.attributes*
do {Compute antecedent opinions
return (*antecedents_set*)

procedure BUILDKNOWLEDGEBASE(*user.attributes*, *object.features*, *history*)
for each *record* ∈ *history*
do {Draw corresponding evidence vector according to (1)
Extract conditional opinions according to (3)
return (*conditionals_set*)

procedure INFERENCEENGINE(*antecedents_set*, *conditionals_set*)
for each *antecedent_opinion* ∈ *antecedents_set*
and *conditional_opinion* ∈ *conditionals_set*
do $\begin{cases} deduced\_opinion \leftarrow antecedent\_opinion \odot conditional\_opinion \\ predicted\_opinion \leftarrow predicted\_opinion \underline{\oplus} deduced\_opinion \end{cases}$
return (*predicted_opinion*)

	ω_{XG}			
Age	$\vec{b}(x_1)$	$\vec{b}(x_2)$	$\vec{b}(x_3)$	u
age is less than 25	1	0	0	0
age is between 26 and 49	0	1	0	0
age is more than 50	0	0	1	0

Table 2. The opinion about user's age ω_{XA}

	ω_{XG}		
Gender	$\vec{b}(x_1)$	$\vec{b}(x_2)$	u
gender is female	1	0	0
gender is male	0	1	0

Table 3. The opinion about gender ω_{XG}

definition for antecedent opinions (ω_{XA}, ω_{XG}, ω_{XO}, and ω_{XZ}) corresponding to the contexts.

ω_{XA} is the antecedent opinion about user's age. We consider three mutually disjoint propositions for context XA: x_1: *young*, x_2: *middle*, and x_3: *old*.

$$XA : \text{age} \begin{cases} x_1 : \text{young(13-25 yrs.)} \\ x_2 : \text{middle(26-49 yrs.)} \\ x_3 : \text{old(50+)} \end{cases}$$

Table 2 shows the belief masses and uncertainty values for the age opinion ω_{XA} based on the user's age. The default base rate $\vec{a}(x_1) = \vec{a}(x_2) = \vec{a}(x_3) = 1/3$ is used. In case that there is no information about user's age, we will have $\omega_{XA} = ((\vec{b}(x_1) = 0, \vec{b}(x_2) = 0, \vec{b}(x_3) = 0), u = 1)$ that indicates the complete uncertainty about the user's age.

ω_{XG} is the antecedent opinion about user's gender and there are two mutually disjoint propositions for the frame XG: x_1: *female* and x_2: *male*.

$$XG : \text{gender} \begin{cases} x_1 : \text{female} \\ x_2 : \text{male} \end{cases}$$

Table 3 gives the belief masses and uncertainty values for the opinion about users' gender (ω_{XG}) and the default base rate $\vec{a}(x_1) = \vec{a}(x_2) = 1/2$ is used. In case of complete uncertainty about $\omega_{XG} = ((\vec{b}(x_1) = 0, \vec{b}(x_2) = 0), u = 1)$.

Likewise, we define ω_{XO} as the antecedent opinion for the user's occupation that has 19 mutually disjoint propositions: x_1: *administrator*, x_2: *artist*, ..., $x_1 9$: *writer*. The default base rate of $\vec{a}(x_i) = 1/19$ is used.

$$XO: \text{occupation} \begin{cases} x_1 : \text{administrator} \\ x_2 : \text{artist} \\ \vdots \\ x_{19} : \text{writer} \end{cases}$$

The antecedent opinion for the user's zipcode is based on 10 mutually disjoint propositions, which correspond to 10 categories for the zipcode for the United States. Each category contains a set of states. For example, x_2 is a proposition corresponding to a category that contains DE, NY, and PA. $\vec{a}(x_i) = 1/10$

$$XZ: \text{zipcode} \begin{cases} x_1 : \text{CT, MA, ME, ...} \\ x_2 : \text{DE, NY, PA} \\ \vdots \\ x_{10} : \text{AK, AS, CA, ...} \end{cases}$$

5.1.2 The Consequent Opinions

We aim to derive the opinion of a particular user about the rating of a particular film, which may have several film genres [6]. The derived consequent opinions (i.e. the rating of a film genre) are combined using the *averaging fusion* operator $\underline{\oplus}$ to get the opinion of that user about the film.

Therefore, there are 19 consequent opinions ($Y1$, $Y2$, ..., $Y19$) for the 19 genres (Action, Adventure,..., Western). For each consequent opinion there are 5 mutually disjoint propositions equivalent to each rating level: y_1: *1 star*, ..., y_5: *5 stars*. The default base rate $\vec{a}(y_i) = 1/5, i = 1..5$ is used.

$$Y1: \text{rating as a "action" movie} \begin{cases} y_1 : 1 \text{ star} \\ y_2 : 2 \text{ stars} \\ \vdots \\ y_5 : 5 \text{ stars} \end{cases}$$

$$\vdots$$

$$Y19: \text{rating as a "western" movie} \begin{cases} y_1 : 1 \text{ star} \\ y_2 : 2 \text{ stars} \\ \vdots \\ y_5 : 5 \text{ stars} \end{cases}$$

[6] A meaningful conditional deduction requires that the antecedent is relevant to the consequent, or in other words that the consequent depends on the antecedents. In this case rating for a film is not dependent on user's attributes, nonetheless rating for a particular genre (e.g. romance) is relevant to the user's attributes. Thus, we consider rating for each film genre as a consequent opinion

5.1.3 Knowledge Base

Sets of conditional opinions for each consequent opinion (ω_{Y1}, ω_{Y2}, ..., ω_{Y19}) and each user's demographic attributes (XA, XG, XO, XZ) are learned from previous experiences in the history according to (3). These conditional opinions constitute the rules about the user's rating behavior for each film genre based on personal characteristics (age, gender, ...).

For instance, the conditional opinion $\omega_{Y1|XA}$ is a rule that tells us how users in different ages rate the *action* movies. $\omega_{Y1|XA} = \{\omega_{Y1|x_1}, \omega_{Y1|x_2}, \omega_{Y1|x_3}\}$ where $\omega_{Y1|x_1}$ tells us about the rating of young users for action movies, while $\omega_{Y1|x_2}$ represent the behavior of middle-aged users for action movies.

$$\omega_{Y1|x_1} = \begin{cases} \vec{b}(y_1) = 0, & \vec{a}(y_1) = 0.2 \\ \vec{b}(y_2) = 0, & \vec{a}(y_2) = 0.2 \\ \vec{b}(y_3) = 0.1667, & \vec{a}(y_3) = 0.2 \\ \vec{b}(y_4) = 0, & \vec{a}(y_4) = 0.2 \\ \vec{b}(y_5) = 0, & \vec{a}(y_5) = 0.2 \\ u = 0.8333 & \end{cases}$$

$$\omega_{Y1|x_2} = \begin{cases} \vec{b}(y_1) = 0.1429, & \vec{a}(y_1) = 0.2 \\ \vec{b}(y_2) = 0, & \vec{a}(y_2) = 0.2 \\ \vec{b}(y_3) = 0, & \vec{a}(y_3) = 0.2 \\ \vec{b}(y_4) = 0.1429, & \vec{a}(y_4) = 0.2 \\ \vec{b}(y_5) = 0, & \vec{a}(y_5) = 0.2 \\ u = 0.7143 & \end{cases}$$

$$\omega_{Y1|x_3} = \begin{cases} \vec{b}(y_1) = 0, & \vec{a}(y_1) = 0.2 \\ \vec{b}(y_2) = 0, & \vec{a}(y_2) = 0.2 \\ \vec{b}(y_3) = 0, & \vec{a}(y_3) = 0.2 \\ \vec{b}(y_4) = 0, & \vec{a}(y_4) = 0.2 \\ \vec{b}(y_5) = 0, & \vec{a}(y_5) = 0.2 \\ u = 1 & \end{cases}$$

These opinions are calculated according to (1). The observation vectors are:

$$\vec{r}_{Y1|x1} = (0,0,1,0,0)$$
$$\vec{r}_{Y1|x2} = (1,0,0,1,0)$$
$$\vec{r}_{Y1|x1} = (0,0,0,0,0)$$

For the sake of simplicity we used the default base rate $1/5$ in this example, however it is possible to use different base rates based on the common belief about that particular conditional opinion.

5.1.4 Inference Engine

The prediction about that user's opinion for *action* movie based on the user's age $\omega_{Y1||XA}$ is obtained by applying the deduction operator (2) on (ω_{XA}) and $\omega_{Y1|XA}$.

$$\omega_{Y1||XA} = \begin{cases} \vec{b}(y_1) = 0.1429, & \vec{a}(y_1) = 0.2 \\ \vec{b}(y_2) = 0, & \vec{a}(y_2) = 0.2 \\ \vec{b}(y_3) = 0, & \vec{a}(y_3) = 0.2 \\ \vec{b}(y_4) = 0.1429, & \vec{a}(y_4) = 0.2 \\ \vec{b}(y_5) = 0, & \vec{a}(y_5) = 0.2 \\ u = 0.7143 & \end{cases}$$

The *average fusion operator* $\underline{\oplus}$ is used to combine the derived opinions based on various user characteristics (age, gender, occupation, zipcode) together to obtain the user's opinion for *action* movies $\omega_{Y1||X}$.

$$\omega_{Y1||X} = \omega_{Y1||XA} \underline{\oplus} \omega_{Y1||XG} \underline{\oplus} \omega_{Y1||XO} \underline{\oplus} \omega_{Y1||XZ}$$

The derived opinions about several genres are combined using the *average fusion operator* $\underline{\oplus}$ to compute the user's opinion about a particular film that belongs to those genres $\omega_{Y||X}$. Appendix demonstrates the whole process of calculating the predicted ratings of a 36 years old, male, administrator user with a user zipcode=05201 for a movie belonging to the Sci-Fi and Fantasy as genres. Figure 6 illustrates the prediction of the rating of a user with four features: XA (age), XG (Gender), XO (Occupation), XZ (Zipcode) for a moive belonging to two genres: Y' and Y''.

6. Dataset and Experimentation

We have chosen the MovieLens dataset to evaluate our work. The GroupLens Research Project at the University of Minnesota has collected the MovieLens data [7]. The data consists of 100,000 ratings by 943 users on 1682 movies with every user having performed at 20 ratings. Simple demographic information for the users is included. There are 5 datasets, which are 80%/20% splits of the data into training and test data (training set of 80,000 ratings, and the test set of 20,000 ratings). Each of these data sets have disjoint test sets; this is for 5-fold cross validation (where we repeat our experiment with each training and test set and calculate the average of the results). The test sets are used as references for the accuracy of the predictions.

Our baseline is the Pearson algorithm [MA04], which relies on the Pearson correlation coefficient to produce a correlation metric between users. This correlation is then used to weigh the rating of each relevant user. The Pearson correlation between users A and B is defined as:

$$P_{A,B} = \frac{\sum_{i=1}^{m} (R_{A,i} - \bar{R}_A) \times (R_{B,i} - \bar{R}_B)}{\sigma_A \times \sigma_B}$$

where m is the number of movies that both users rated. $R_{A,i}$ is the rating, user A gave to movie i. $\bar{R}_A$ is the average rating user A gave to all movies, and σ_A is the corresponding standard deviation of those ratings. Once the Pearson correlation between a user and all other users is obtained, the predicted movie rating is calculated as:

[7] http://www.cs.umn.edu/Research/GroupLens/data/

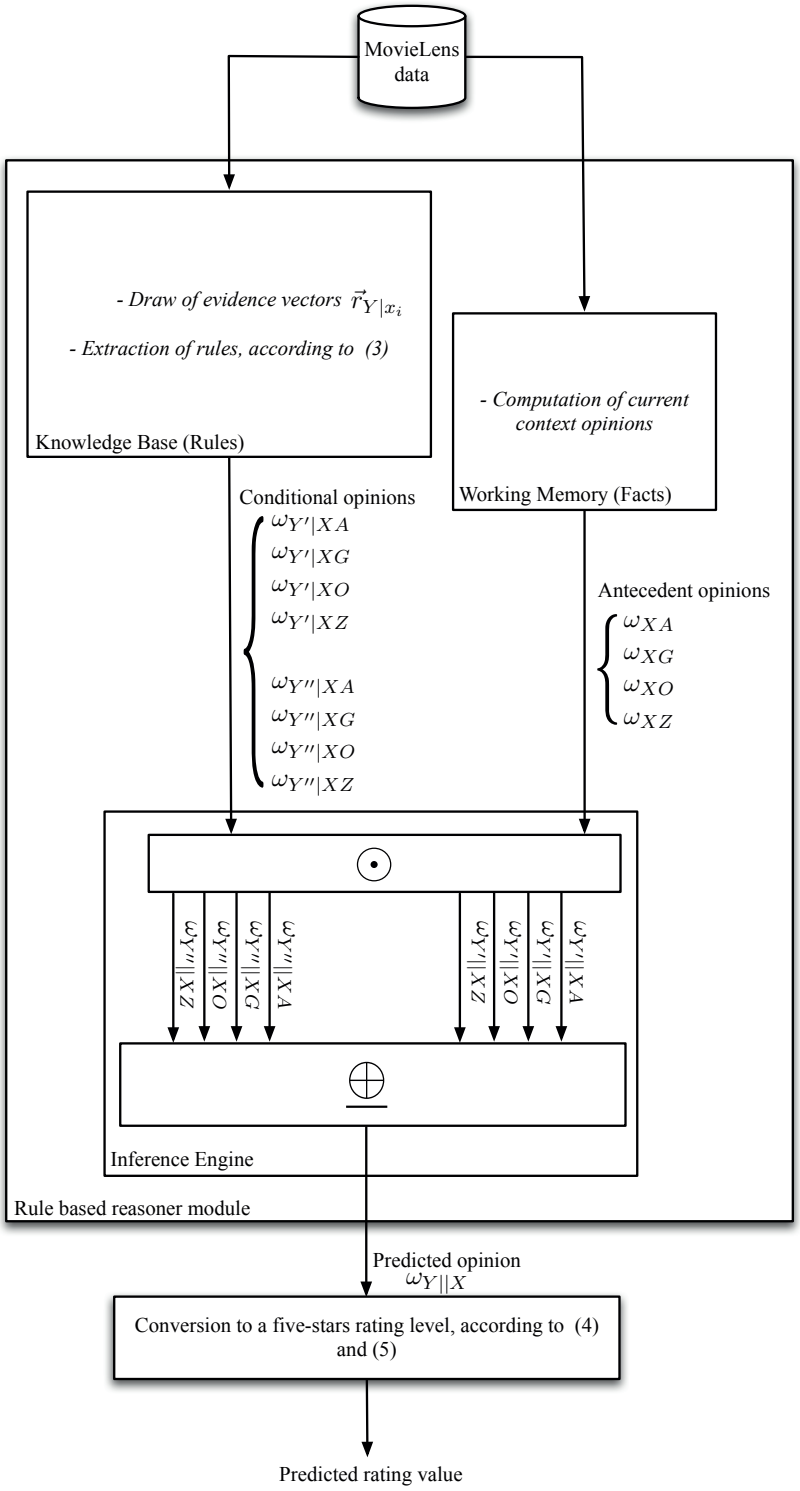

Figure 6. Prediction of the rating of a user with four features: XA (age), XG (Gender), XO (Occupation), XZ (Zipcode) for a movie with two genres: Y' and Y''.

	FCP	MAE	RMSE
Dataset1	0.3293	0.2371	0.2864
Dataset2	0.3107	0.2343	0.2834
Dataset3	0.3093	0.2291	0.2773
Dataset4	0.3020	0.2232	0.2706
Dataset5	0.3130	0.2347	0.2859
Average	0.3129	0.2317	0.2807
Pearson	0.1993	0.3049	0.3804

Table 4. Different Datasets

$$R_{A,i} = \bar{R}_A + \frac{\sum_{U=1}^{n} (R_{U,i} - \bar{R}_U) \times P_{A,U}}{\sum_{u=1}^{n} |P_{A,U}|}$$

The use of the Pearson correlation coefficient is quite common in the field of collaborative filtering, and results obtained with this method used to gauge the performance of other algorithms. The Pearson algorithm uses only the rating information whitout taking into account the situational information, while our method uses situational information to do the prediction.

Three types of evaluation criteria are used in this paper:

- FCP: fraction of correct predictions.
- MAE (Mean Absolute Error) : average of the prediction error (difference between probability expected values of predicted and real opinions) over all queries.
- RMSE (root mean squared error) : root mean of the average of the squared prediction error. RMSE tends to emphasize large errors.

6.1 Results

In table 4, we present the final results of the evaluation. We start by commenting the row fraction of correct predictions (FCP) that is approximately 0.31 and shows that from each 10 predicted ratings, three ratings are predicted with exact values. Further, the prediction errors (MAE and RMSE) for the other ratings that are not predicted exactly (seven ratings from each ten predicted ratings) are small in comparison with the Pearson method ($MAE \approx 0.23$ & $RMSE \approx 0.28$).

Table 5 gives the results for different values W in formula 1.

As illustrated in figure 7, the comparison of FCP, MAE and RMSE values for ten different values of W leads to the conclusion that $W = 5$ gives us the best results, i.e. lowest errors (MAE and RMSE) and highest FCP.

All-in-all, the results of the evaluation lead to the conclusion that our approach provides an improvement over the Pearson algorithm and this implies that contextual information is useful in making predictions.

Constant value W	FCP	MAE	RMSE
W=5	0.3195	0.2306	0.2780
W=5×10^1	0.3065	0.2463	0.3654
W=5×10^2	0.2847	0.2364	0.2828
W=5×10^3	0.2594	0.2489	0.3002
W=5×10^4	0.2502	0.2580	0.3167
W=5×10^5	0.2445	0.2598	0.3203
W=5×10^6	0.2445	0.2600	0.3207
W=5×10^7	0.2445	0.2601	0.3207
W=5×10^8	0.2445	0.2601	0.3207
W=5×10^9	0.2445	0.2601	0.3207

Table 5. Different constant values for W in the formula (1).

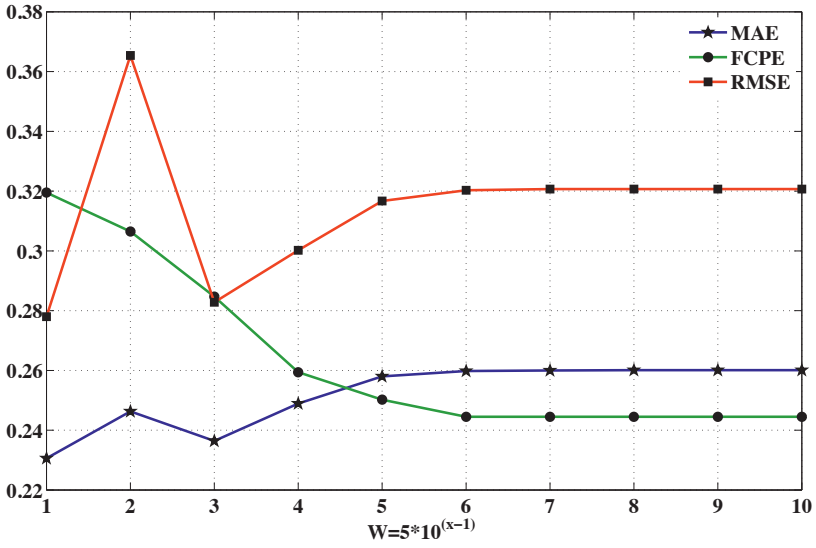

Figure 7. Comparison of FCP, MAE and RMSE values for 10 different values of W.

7. Related Work

Researchers have had different motivations to incorporate the notion of context into the trust management accounts. For example, [NWvS06] aims at reducing the complexity in management of trust relationships. [NWvSL07, GCJ03] focuces on the improvement of the trust recommendation process. [HY06] investigates how to infer trust information in context hierarchies. [RGPB06] improves the performance of trust management systems. [RGPB06, RP07] provide protection against changes of identity and

first time offenders in trust management systems. [TK08, TKH08b, TKH08a, THO09] propose a method to use the trust model to bootstrap in unanticipated situations. [BG06, BBEZG08, GCJ03, TKH08b, TKH08a, THO09] provide methods that correlate trust information among various contexts.

Contextual information has been represented in several different forms such as Context-aware domains [NWvSL07], Intensional Programming [WA08], Multi-dimensional goals [GDFB06], Clustering [RP07], and Ontologies [TLU06].

The main contributions of these works are manifold: Neisse et al. [NWvS06] proposed the idea of using the abstraction of context-aware domains to reduce the complexity in the management of trust relationships. In a large context-aware system, with thousand of components and users, it is impractical to associate trust relationships with individual entities, as this can easily become unmanageable. Examples of context-aware management domain definitions are "nearby persons", "Personal devices", and "Working colleagues". This is the same as the common ground concept introduced earlier. The idea is to provide mechanisms to define and infer the trust degree of an entity based on the context information provided about that entity. According to [NWvSL07] it is also possible to use context information to improve the recommendation process (to determine from whom to request recommendation). This will allow anonymous and still useful recommendation exchange.

In [HY06] it is noted that context can often be structured hierarchically. For example, if you trust someone to drive your car, then you would most likely give him also your car keys or the keys to the garage. Therefore, it is necessary to identify possible hierarchical structures between different contexts in our model to be able to infer trust information from one into the other. In this work, entities, which can be applications, other users or agents that act on behalf of users, are structured into a context-based trust graph. Positions in this graph indicate the context-based trust level and changes based on events or over time. The structure of the trust graph reflects a certain hierarchy.

Alagar et al. [APW05] investigated the intensional programming paradigm for agent communication by introducing context as a first class object in the intensional programming language Lucid. Intensional programming is a powerful and expressive paradigm based on Intensional Logic. Intensional logic is a branch of mathematical logic used to precisely describe context-dependent entities. In their paper, definitions, syntax, and operators for context, and an operational semantic for evaluating expressions in extended Lucid, are given. It is demonstrated that the extended Lucid language, called Agent Intensional Programming Language (AIPL), has the generality and the expressiveness for being an Agent Communication Language (ACL). Based on this work, a context-specific trust model for multi-agent systems is introduced [WA08]. The explicit introduction of context into the computation of trust, annotation of trust policies with context conditions, and definition of delegation through related contexts are some of the new results given in this paper.

The context issue has also been viewed in multi-dimensional trust modeling for agents when goal requirements are multi-dimensional [GDFB06]. An agent's reward

is determined by goal requirements and behavioral constraints of potential partners (e.g. quality, timeliness, availability, and cost).

In [TK08], the authors propose an algorithm to estimate trust when truster and trustee are completely unfamiliar with each other. According to their algorithm the truster uses her past experiences that occurred in the same context as the current context to generate a training set. Then using maximum likelihood estimation, the trust value for a new trustee can be estimated.

[RGPB06] defines a set of reference contexts in a metric space and associate truthfulness data with it. These data are updated and queried with weight that decreases with distance between the current situation and the reference context. The model uses Leader-Follower clustering to identify the reference contexts to be representative of the data. The advantage of this clustering method is that it allows an on-line approach without pre-specifying the number of expected clusters, and requires only a single parameter as input. The biggest disadvantage is that it may easily under- or over-estimate the number of clusters. In an empirical test, it is shown that context-aware models easily outperform general trust models when the situation has an impact on partner trustfulness and that their performance and efficiency is comparable with general trust models where the trustfulness is independent of the situation. In this work two advanced uses of context for multiagent trust modeling is proposed: (i) policy/norm learning at runtime by analyzing data regarding the performance of different agents in similar situations (e.g. when all agents fail in a certain situation, they may agree to introduce a policy that specifically prohibits such actions) (ii) reasoning based on uncertain identities by decomposing the single identity dimension into an identity subspace, where each agent is defined by one or more crucial properties. With this modification, the trust model can make predictions about the performance of agents by exploiting data characterizing a similar agent's performance in the past. The main advantages are that the extended model learns faster and once the new agent is categorized, its performance can be predicted. This is also a clear advantage in ad-hoc environments, where there is no agent platform to enforce unique identity.

Based on this model, [RP07] conclude that the extension of a trust model with a context representation environment can be extended to encompass a more open situation (e.g. a wireless sensor network that is hard to identify and where the barriers of entry are quite low). In such environments it is not needed to have assumptions like: (i) proven identity, (ii) repetitive interactions and (iii)similar trusting situations. The fact that two agents with presumably distinct identities can be considered identical by a context-sensitive trust model may provide protection against changes of identities. This approach is also effective against first time offenders; we can obtain a model with inductive properties, which is able to estimate the performance of new entrants using the experience with the similar partners in the past.

[GPH03] proposes an ontology for trust. In [GH04] they have considered a model using context-specific reputation by assigning numeric ratings to different types of relations based on the context of the analysis. In [TD04] rules describing how certain context-sensitive information (trust factors) reduces or enhances the trust value have been specified for this trust ontology. The authors also argue that a specific advantage

of making the context explicit in message exchanges is that this information can be used in trust policies. For example, a policy can state that news information related to a particular location is to be trusted more if the reporting entity was at the location at the time when the event occurred. In this sense, policies define how to process context information to derive trustworthiness assertions. However, they have not answered how the context-sensitive trust factor should be determined. In addition, neither have they addressed the fact that the trust value might be different for different aspects of trust.

In [HF06], trust is formalized by using situation calculus in order to define a trust ontology. Situation calculus is a logic language specifically designed for representing dynamically changing worlds. It works in the following way: the changing world is represented by a set of fluents. A fluent is a property (of the world) whose value is dependent on situations. In other words, a fluent dynamically changes when the situation changes. The situation, in turn, changes when an action is performed by agent(s) in the world. Trust and context are represented as fluents.

In [TLU06] contextual information (context attributes) is used to adjust the output of a trust determination process. Each attribute can adjust the trust value positively or negatively according to a specified weight. For example, if it is the trust value and the weight of the context property, then the adjusting function can be for either decrease or increase. A context ontology connects the context attributes with each other in an appropriate manner, enabling the utilization of context attributes which do not exactly match the query, but are "close enough" to it. In this work, the notion of context also has been applied to the reputations by emphasizing more the observations that have taken place under similar conditions as where the truster currently is. Two relationships have been considered between recommendations and context. First, as was the case with reputation, the contextual details at the time when the recommendation was made can be considered and compared with the truster's current context. Note that considering this is not as straightforward as was the case with reputation, since recommendations come from others, not from the truster. Secondly, the recommendation content itself can be context-dependent.

In [CBGS07] cases are considered where an agent does not have enough information to produce a trust value for a given task, but she knows instead what the previous partner's behavior are when performing similar tasks. This model estimates trust using the information about similar tasks.

[BG06] proposes a framework for dynamically updating and inferring the unobserved reputation of environment participants in different contexts. This framework proposes the employment of a reputation structure tree to represent the relationship between the contexts of the environment. Reputation of a given identity in one context can be propagated to other contexts through two mechanisms, namely: forward update and backward adjustment. This work does not mention how to develop the reputation structure tree.

[BBEZG08] also proposes a framework for the author's previous proposal based on valuation networks. Global reputation is modeled as Dempster-Shafer belief functions on a Markov tree through which the relationship between various contexts of a unique

environment is modeled through hyper-vertices of the Markov tree. Reputation of each identity in a given context is represented using a belief mass assignment function. The estimation of reputation in various contexts of the environment is performed by the employment of the message passing-based belief propagation model of the Shenoy-Shafer architecture.

[GCJ03] presents an initial investigation into addressing the issue of making trust-based security decisions in a given context. The authors consider several trust attributes for each context and propose how to map trust across contexts based on common attributes among those contexts.

[SLP04] provides a survey of different approaches to model context for ubiquitous computing. In this work numerous approaches are reviewed, classified relative to their core elements, and evaluated with respect to their appropriateness for ubiquitous computing. The authors reach conclusion that the most promising assets for context modeling for ubiquitous computing environments can be found in the ontology category in comparison with other approaches like key-value models, mark-up scheme models, graphical models, object-oriented models, and logic based models. This selection is based on the six requirements dominant in pervasive environments: distributed composition, partial validation, richness and quality of information, incompleteness and ambiguity, level of formality, and applicability to existing environments.

The motivations of our work to incorporate the notion of context into the trust management are initialization of trust values (in unknown situations or for unknown trustees) and adjustment of the trust values based on the underlying situation. Two types of reasoning mechanisms collectively support the context-aware trust management process in our approach: case-based reasoning (CBR) and rule based reasoning (RBR). Among the related work, [TD04] resemble our approach. They formalize user-defined rules that take context-sensitive information into account on the basis of trust ontology. However, we extract the rules automatically from the history and incorporate the abductive reasoning paradigm to apply them.

8. Conclusion and Future Work

To wrap up, we have highlighted the focuses of attention within the trust management literature, and reviewed various accounts of situation-aware trust judgment. It seems that the importance of context-sensitivity has been recognized. However, mechanisms that reflect the situation-awareness on the quality of the trust judgment remains to be investigated in more details. Our framework based on the case-based reasoning rule-based reasoning is a step toward making the trust management models situation-aware. This framework has been validated for the Subjective Logic trust management model as an example and evaluated using a real large-scale dataset. The results of the evaluation lead to the expectation that our approach provides an improvement for the trust inference task and this implies that situational information is useful in making predictions.

In the future, we aim to add a risk management module (RMM) to this framework. Risk evaluation becomes important in inferring trust values among situations especially when the trustworthiness of some principal is completely unknown and

no recommendation information is available. The intuitive idea behind such a risk assessment can be to look up the in the casebase to see if there are any similar previous interactions, i.e., if we have previously encountered an entity with similar trust attributes and similar risk attributes in the same situation. The ontology part should be able to describe the level of situational risk, whereby the higher the risk of negative outcome, the higher the level of precision that must be captured.

APPENDIX

For a user with age=36, gender=M, occupation=administrator (Y1), and zip-code=05201 and a movie with Sci-Fi and Fantasy as genres we have

$$\begin{aligned}
\omega_{XA} &= ((0,1,0),0)\\
\omega_{XG} &= (0,1,0)\\
\omega_{XO} &= ((1,0,0,...,0),0)\\
\omega_{XZ} &= (0,0,0,0,0,1,0,...,0)
\end{aligned}$$

$$\begin{aligned}
\omega_{Y9|XA} = \{&((0.0565,0.0988,0.2372,0.3438,0.2615),0.0022),\\
&((0.0410,0.0891,0.2545,0.3722,0.2421),0.0011),\\
&((0.0224,0.0944,0.2587,0.3825,0.2375),0.0045)\}
\end{aligned}$$

$$\begin{aligned}
\omega_{Y9|XG} = \{&((0.0551,0.0940,0.2561,0.3369,0.2557),0.0022),\\
&((0.0381,0.0921,0.2481,0.3771,0.2438),0.0009)\}
\end{aligned}$$

$$\begin{aligned}
\omega_{Y9|XO} = \{&(0.0298,0.0795,0.2296,0.3924,0.2602,0.0085),\\
&((0.0462,0.0872,0.2321,0.3346,0.2744),0.0256),\\
&((0.0220,0.0604,0.2033,0.3956,0.2088),0.1099),\\
&((0.0287,0.0639,0.2311,0.4003,0.2701),0.0059),\\
&((0.0261,0.0912,0.2929,0.3492,0.2333),0.0074),\\
&((0.0497,0.1098,0.2693,0.3259,0.2110),0.0343),\\
&((0.1064,0.0868,0.2092,0.3165,0.2624),0.0187),\\
&((0.1651,0.1377,0.3430,0.2370,0.1001),0.0171),\\
&((0.0450,0.0901,0.2072,0.3604,0.1171),0.1802),\\
&((0.0315,0.0847,0.1961,0.3608,0.2785),0.0484),\\
&((0.0278,0.0955,0.2702,0.3626,0.2338),0.0101),\\
&((0.0415,0.1038,0.2859,0.3067,0.2300),0.0319),\\
&((0.0364,0.0820,0.2213,0.3920,0.2594),0.0089),\\
&((0.0292,0.1070,0.2938,0.4202,0.1109),0.0389),\\
&((0.0208,0.0311,0.1730,0.3772,0.3287),0.0692),\\
&((0.0188,0.0564,0.2657,0.3972,0.2368),0.0251),\\
&((0.0420,0.0975,0.2302,0.3657,0.2614),0.0031),\\
&((0.0280,0.0945,0.2469,0.4098,0.2015),0.0193),\\
&((0.0560,0.1276,0.2569,0.3371,0.2116),0.0108)\}
\end{aligned}$$

$$\omega_{Y9|XZ} = ((0.0233, 0.0857, 0.2405, 0.3956, 0.2485, 0.0064),$$
$$((0.0813, 0.0885, 0.2267, 0.3580, 0.2401), 0.0053),$$
$$((0.0599, 0.0945, 0.2439, 0.3630, 0.2331), 0.0056),$$
$$((0.0262, 0.0820, 0.2350, 0.3654, 0.2813), 0.0101),$$
$$((0.0391, 0.0918, 0.2445, 0.3614, 0.2547), 0.0085),$$
$$((0.0366, 0.0916, 0.2461, 0.3494, 0.2708), 0.0054),$$
$$((0.0214, 0.1059, 0.3148, 0.3528, 0.1985), 0.0065),$$
$$((0.0461, 0.0899, 0.2259, 0.3864, 0.2425), 0.0092),$$
$$((0.0435, 0.0994, 0.2660, 0.3599, 0.2220), 0.0093),$$
$$((0.0381, 0.0915, 0.2475, 0.3609, 0.2588), 0.0032)\}$$

$$\omega_{Y16|XA} = \{((0.0605, 0.1215, 0.2524, 0.3120, 0.2477), 0.0058),$$
$$((0.0519, 0.1186, 0.2677, 0.3466, 0.2118), 0.0033),$$
$$((0.0286, 0.1303, 0.2629, 0.3291, 0.2263), 0.0229)\}$$

$$\omega_{Y16|XG} = \{((0.0669, 0.1202, 0.2624, 0.3171, 0.2240), 0.0094),$$
$$((0.0492, 0.1210, 0.2628, 0.3387, 0.2258), 0.0025)\}$$

$$\omega_{Y16|XO} = \{((0.0457, 0.0874, 0.2823, 0.3091, 0.2487, 0.0269),$$
$$((0.0406, 0.0849, 0.2325, 0.2768, 0.2915), 0.0738),$$
$$((0.0426, 0.0638, 0.1064, 0.2766, 0.0851), 0.4255),$$
$$((0.0385, 0.1193, 0.2509, 0.3478, 0.2186), 0.0248),$$
$$((0.0464, 0.1014, 0.2749, 0.3412, 0.2171), 0.0190),$$
$$((0.0837, 0.1410, 0.1718, 0.2863, 0.2291), 0.0881),$$
$$((0.0583, 0.0828, 0.2270, 0.3129, 0.2577), 0.0613),$$
$$((0.1245, 0.1727, 0.2410, 0.2811, 0.1004), 0.0803),$$
$$((0.0682, 0.0909, 0.1364, 0.1136, 0.1364), 0.4545),$$
$$((0.0248, 0.0744, 0.2645, 0.2562, 0.2149), 0.1653),$$
$$((0.0260, 0.1172, 0.3073, 0.3255, 0.1719), 0.0521),$$
$$((0.0408, 0.1122, 0.2806, 0.2755, 0.1888), 0.1020),$$
$$((0.0408, 0.1119, 0.2228, 0.3337, 0.2699), 0.0209),$$
$$((0.0172, 0.1034, 0.2241, 0.3966, 0.0862), 0.1724),$$
$$((0.0673, 0.0865, 0.1827, 0.2308, 0.2404), 0.1923),$$
$$((0.0429, 0.1179, 0.2429, 0.3214, 0.2036), 0.0714),$$
$$((0.0526, 0.1261, 0.2647, 0.3319, 0.2168), 0.0079),$$
$$((0.0418, 0.1114, 0.2622, 0.3550, 0.1833), 0.0464),$$
$$((0.0988, 0.1490, 0.2184, 0.2877, 0.2114), 0.0347)\}$$

$$\begin{aligned}\omega_{Y16|XZ} = \{&,0.0402,0.0985,0.2714,0.3477,0.2221,0.0201),\\ &((0.0791,0.1196,0.2490,0.3103,0.2223),0.0198),\\ &((0.0596,0.1412,0.2535,0.3211,0.2070),0.0175),\\ &((0.0421,0.1154,0.2231,0.3229,0.2652),0.0312),\\ &((0.0470,0.1124,0.2489,0.3200,0.2489),0.0229),\\ &((0.0472,0.1192,0.2461,0.3646,0.2087),0.0141),\\ &((0.0280,0.1074,0.2908,0.3591,0.1924),0.0224),\\ &((0.0815,0.1541,0.2563,0.3007,0.1778),0.0296),\\ &((0.0555,0.1171,0.2651,0.3255,0.2121),0.0247),\\ &((0.0484,0.1140,0.2688,0.3122,0.2466),0.0101)\}\end{aligned}$$

The final opinion is the result of combination of all opinions using the average fusion operator $\underline{\oplus}$.

$$\begin{aligned}\omega_{Y||X} = &\left(\omega_{Y9||XA} \underline{\oplus} \omega_{Y9||XG} \underline{\oplus} \omega_{Y9||XO} \underline{\oplus} \omega_{Y9||XZ}\right)\\ \underline{\oplus} &\left(\omega_{Y16||XA} \underline{\oplus} \omega_{Y16||XG} \underline{\oplus} \omega_{Y16||XO} \underline{\oplus} \omega_{Y16||XZ}\right)\end{aligned}$$

where $\omega_{Y9||XA} = \omega_{XA} \odot \omega_{Y9|XA}$.

$$\begin{aligned}\omega_{Y9||XA} &= ((0.0410,0.0891,0.2545,0.3722,0.2421),0.0011)\\ \omega_{Y9||XG} &= ((0.0381,0.0921,0.2481,0.3771,0.2438),0.0009)\\ \omega_{Y9||XO} &= ((0.0298,0.0795,0.2296,0.3924,0.2602),0.0085)\\ \omega_{Y9||XZ} &= ((0.0233,0.0857,0.2405,0.3956,0.2485),0.0064)\\ \omega_{Y16||XA} &= ((0.0519,0.1186,0.2677,0.3466,0.2118),0.0033)\\ \omega_{Y16||XG} &= ((0.0492,0.1210,0.2628,0.3387,0.2258),0.0025)\\ \omega_{Y16||XO} &= ((0.0457,0.0874,0.2823,0.3091,0.2487),0.0269)\\ \omega_{Y16||XZ} &= ((0.0402,0.0985,0.2714,0.3477,0.2221),0.0201)\end{aligned}$$

The final predicted rating is $\omega_{Y||X} = ((0.0451,0.1079,0.2638,0.3471,0.2283),0.0079)$ that is equal to point estimate of 0.6514 and rating level of 4.

References

[AP94] A. Aamodt and E. Plaza. Case-based reasoning. *Proc. MLnet Summer School on Machine Learning and Knowledge Acquisition*, pages 1–58, 1994.

[APW05] V.S. Alagar, J. Paquet, and K. Wan. Intensional Programming for Agent Communication. In *Declarative agent languages and technologies II: second international workshop, DALT 2004, New York, NY, USA, July 19, 2004: revised selected papers*, page 239. Springer-Verlag New York Inc, 2005.

[BBEZG08] E. Bagheri, M. Barouni-Ebrahimi, R. Zafarani, and A.A. Ghorbani. A belief-theoretic reputation estimation model for multi-context communities. *Lecture Notes in Computer Science*, 5032:48–59, 2008.

[BD05] N.A. Bradley and M.D. Dunlop. Toward a multidisciplinary model of context to support context-aware computing. *Human-Computer Interaction*, 20(4):403–446, 2005.

[BG06] E. Bagheri and A.A. Ghorbani. Behavior analysis through reputation propagation in a multi-context environment. *International Conference on Privacy, Security and Trust (PST06)*, 2006.

[CBGS07] A. Caballero, JA Botia, and A. Gomez-Skarmeta. On the Behaviour of the TRSIM Model for Trust and Reputation. *LECTURE NOTES IN COMPUTER SCIENCE*, 4687:182, 2007.

[CF02] C. Castelfranchi and R. Falcone. Social trust: A cognitive approach. *Trust and deception in virtual societies*, pages 55–90, 2002.

[DA00] A.K. Dey and G.D. Abowd. Towards a better understanding of context and context-awareness. In *CHI 2000 workshop on the what, who, where, when, and how of context-awareness*, pages 304–307, 2000.

[DD98] J. Durkin and J. Durkin. *Expert systems: design and development*. Prentice Hall PTR Upper Saddle River, NJ, USA, 1998.

[GCJ03] E. Gray, Y. Chen, and C. Jensen. Initial Investigation into Cross-context Trust and Risk Assessment. In *IASTED International Conference on Communication, Network, and Information Security*, pages 56–61, 2003.

[GDFB06] N. Gujral, D. DeAngelis, K.K. Fullam, and K.S. Barber. Modeling Multi-Dimensional Trust. In *the Proceedings of the Workshop on Trust in Agent Societies*, pages 8–12, 2006.

[Gen83] D. Gentner. Structure-mapping: A theoretical framework for analogy. *Cognitive science*, 7(2):155–170, 1983.

[GH04] J. Golbeck and J. Hendler. Inferring Reputation on the Semantic Web. In *Proceedings of the 13th International World Wide Web Conference*, 2004.

[GPH03] J. Golbeck, B. Parsia, and J. Hendler. Trust Networks on the Semantic Web. In *Proceedings of Cooperative Intelligent Agents*, volume 2003. Springer, 2003.

[HF06] J. Huang and M.S. Fox. An ontology of trust: formal semantics and transitivity. In *Proceedings of the 8th international conference on Electronic commerce: The new e-commerce: innovations for conquering current barriers, obstacles and limitations to conducting successful business on the internet*, page 270. ACM, 2006.

[HT97] K.J. Holyoak and P. Thagard. The analogical mind. *American Psychologist*, 52:35–44, 1997.

[HY] S. Holtmanns and Z. Yan. Context-Aware Adaptive Trust.

[JG03] A. Jøsang and T. Grandison. Conditional inference in subjective logic. In *Information Fusion, 2003. Proceedings of the Sixth International inproceedings of*, volume 1, 2003.

[JLC08] A. Jøsang, X. Luo, and X. Chen. Continuous Ratings in Discrete Bayesian Reputation Systems. In *Proceedings of the Joint¡ Trust and PST Conferences on Privacy, Trust Management and Security (IFIPTM 2008), Trondheim*. Springer, 2008.

[Jøs07] A. Jøsang. Probabilistic logic under uncertainty. In *Proceedings of the thirteenth Australasian symposium on Theory of computing-Volume 65*, page 110. Australian Computer Society, Inc., 2007.

[Jøs08] A. Jøsang. Conditional Reasoning with Subjective Logic? *Journal of Multiple-Valued Logic and Soft Computing*, 15(1):5–38, 2008.

[JPD05] A. Jøsang, S. Pope, and M. Daniel. Conditional deduction under uncertainty. In *Proceedings of the 8th European inproceedings on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU 2005)*. Springer, 2005.

[Kol93] J. Kolodner. *Case-based reasoning*. Morgan Kaufmann Publishers Inc. San Francisco, CA, USA, 1993.

[MA04] P. Massa and P. Avesani. Trust-Aware Collaborative Filtering for Recommender Systems. *LECTURE NOTES IN COMPUTER SCIENCE*, pages 492–508, 2004.

[Mar94] S.P. Marsh. *Formalising trust as a computational concept*. Citeseer, 1994.

[NWvS06] R. Neisse, M. Wegdam, and M. van Sinderen. Context-Aware Trust Domains. In *1st European Conference on Smart Sensing and Context, Enschede, The Netherlands, Oct-2006*. Springer, 2006.

[NWvSL07] R. Neisse, M. Wegdam, M. van Sinderen, and G. Lenzini. Trust Management Model and Architecture for Context-Aware Service Platforms. *LECTURE NOTES IN COMPUTER SCIENCE*, 4804:1803, 2007.

[Özt98] Pinar Öztürk. A context model for knowledge-intensive case-based reasoning. *International Journal of Human Computer Studies*, 48:331–356, 1998.

[RGPB06] M. Rehak, M. Gregor, M. Pechoucek, and J.M. Bradshaw. Representing Context for Multiagent Trust Modeling. In *Proceedings of the IEEE/WIC/ACM International Conference on Intelligent Agent Technology (IAT 2006 Main Conference Proceedings)(IAT'06)-Volume 00*, pages 737–746. IEEE Computer Society Washington, DC, USA, 2006.

[RHJ05] S.D. Ramchurn, D. Huynh, and N.R. Jennings. Trust in multi-agent systems. *The Knowledge Engineering Review*, 19(01):1–25, 2005.

[RP07] M. Rehak and M. Pechoucek. Trust modeling with context representation and generalized identities. *Klusch, M., Hindriks, K., apazoglou, MP, Sterling, L.(eds.) CIA*, pages 298–312, 2007.

[SLP04] T. Strang and C. Linnhoff-Popien. A context modeling survey. In *Workshop on Advanced Context Modelling, Reasoning and Management as part of UbiComp*, 2004.

[SS05] J. Sabater and C. Sierra. Review on computational trust and reputation models. *Artificial Intelligence Review*, 24(1):33–60, 2005.

[TD04] S. Toivonen and G. Denker. The impact of context on the trustworthiness of communication: An ontological approach. In *Proceedings of the Trust, Security, and Reputation on the Semantic Web workshop, held in conjunction with the 3rd International Semantic Web Conference (ISWC 2004), Hiroshima, Japan*, volume 127, 2004.

[THÖ09] Mozhgan Tavakolifard, Peter Herrmann, and Pinar Öztürk. Analogical trust reasoning. *Trust Management III*, pages 149–163, 2009.

[TK08] M. Tavakolifard and S. Knapskog. A Probabilistic Reputation Algorithm for Decentralized Multi-Agent Environments. In *Proceedings of the 4th International Workshop on Security and Trust Management (STM 08)*. ENTCS, 2008.

[TKH08a] M. Tavakolifard, S.J. Knapskog, and P. Herrmann. Cross-situation trust reasoning. In *Proceedings of 2008 IEEE/WIC/ACM International Conference on Intelligent Agent Technology - IAT 2008*, volume 3, pages 67–71, Sydney, Australia, December 2008. IEEE Computer Society.

[TKH08b] M. Tavakolifard, S.J. Knapskog, and P. Herrmann. Trust transferability among similar contexts. In *Proceedings of the 4th ACM symposium on QoS and security for wireless and mobile networks*, pages 91–97. ACM New York, NY, USA, 2008.

[TLU06] S. Toivonen, G. Lenzini, and I. Uusitalo. Context-aware trust evaluation functions for dynamic reconfigurable systems. In *Proceedings of the Models of Trust for the Web workshop (MTW06), held in conjunction with the 15th International World Wide Web Conference (WWW2006) May*, volume 22. Citeseer, 2006.

[WA08] K. Wan and V. Alagar. An Intensional Functional Model of Trust. In *Proceedings of IFIPTM 2008 - Joint iTrust and PST Conferences on Privacy, Trust Management and Security*, pages 69–85. Springer, 2008.

PAPER E

Trust Evaluation Initialization Using Contextual Information

Mozhgan Tavakolifard, Svein J. Knapskog

Proceedings of the The International Conference on Management of Emergent Digital EcoSystems (MEDES)
San Francissco, USA, Nov 2011

TRUST EVALUATION INITIALIZATION USING CONTEXTUAL INFORMATION

Mozhgan Tavakolifard,[1] Svein J. Knapskog,[1]

[1] *Centre for Quantifiable Quality of Service in Communication Systems*
Norwegian University of Science and Technology
{mozhgan, knapskog}@Q2S.ntnu.no

Abstract The majority of existing trust and reputation models consider two types of knowledge in estimating the trustworthiness of a trustee in an interaction: personal direct experiences and recommendations from third parties. However, previous direct and recommended evidence is not available for new users. In addition, a new user joins the system with a neutral reputation value in most systems and must participate in interactions with others in order to raise its reputation score. Users usually tend to interact with high reputable ones; therefore, the chance of new-comers being selected for interaction is generally rare. As a result, it is hard for a new user to raise his or her reputation score. Furthermore, short-lived users preclude the others from gaining the necessary experiences to make an accurate evaluation. Even long-lived users might leave the system and rejoin with a new identity to lose their bad reputation and start with a neutral score. Hence, effective initialization mechanism is needed to avoid such problems in trust and reputation systems. We propose to use contextual information for bootstrapping the reputation value. We use the Maximum Likelihood Estimation method for trust initialization of probabilistic trust models. We show its implementation and effectiveness for a particular model called 'Beta reputation model' through simulations.

1. Introduction

The principle behind the World Wide Web is, without doubt, one of the most egalitarian inventions mankind has ever made in modern times. The Web is not only an information space but also a medium for human relationships; a plethora of web services spanning banking, shopping, health care and learning is becoming available for citizens all over the world. However, a steadily increasing number and variety of virtual social networks bring along some problems that cast a shadow on the huge advantages the Web may provide. A major problem with such an open and distributed space is that users lack sufficient information about the quality of the e-services and their providers. Conventional security mechanisms cannot handle the trust phenomenon in the way the new information systems would need. Therefore, the growth of services such as online transactions and information exchange is conditioned on the development of new trust and reputation management models.

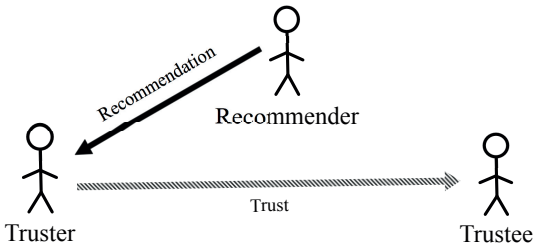

Figure 1. Relationship among truster, trustee, and recommender.

The recent *trust management systems* mimic the behavior that people exhibit in real life independent of the Internet. Then, if a person does not know about the person she is considering to make business with, she resorts to contact other people in her social networks to find out whether the candidate business partner has a good reputation. In a corresponding information system (i.e., referral systems), there are three agent roles: The *trustee* is the service provider, the *truster* is the agent interested in the provided service who needs to judge the trustworthiness of the provider, and the *recommender/referrer* can provide a rating to the truster about a trustee (see figure 1). An agent can play more than one role. For example, a truster often rates the trustee subsequently to a transaction she was involved in. The truster normally relies on her own personal direct experiences as long as they are thought to be sufficient and uses others' recommendations if she does not feel that she has enough experience with the trustee herself.

Hence, the majority of trust models consider two types of knowledge in estimating the trustworthiness of a trustee in the next interaction: experiences and recommendations. Recommendations about a trustee are derived from word-of-mouth and are frequently based on ratings about the trustee given by recommenders.

However, initial cases exist where previous direct and recommended evidence is unavailable. For example, in open systems such as ad-hoc networks, the agents are distributed across various platforms and can join or leave the system at their own will, which requires the assignment of estimated initial trust values. This case is called the cold-start problem. Moreover, newcomers are offered only a limited number of resources and so struggle initially to build their reputations. As other users in the system tend to interact with high reputable users, the chance of a new user being selected for interaction is generally rare (*e.g.*, in eBay, many users will not deal with individuals with a low reputation score [Mal01]). Hence, it is hard for a new user to raise his or her reputation score. This new-comers challenge may be a barrier to entry into the marketplace or community.

In both cases, the problem is one of how to minimize the risk inherent in "bootstrapping" trust evaluations when interacting with new, unknown users. The initialization problem may result in another problem in trust and reputation systems called "re-entry or churn attack". In online communities, it is usually easy for members to disappear

and re-register under a completely different online identity with zero or very low cost (*e.g.*, eBay). Therefore, users can hide their bad reputation in this way and start with a neutral one.

Most trust and reputation systems simply consider a default value for trust bootstrapping. In such a case, a high value is risky while a low value carries the risk that new-comers might be ignored completely. In order to address these issues, we propose that user can generalize their experiences with known partners in the same context in order to form tentative trust evaluations about unknown users.

The model we propose here can be applied to any probabilistic trust mechanism that uses numerical ratings to compare and exchange opinions, although, we demonstrate its use with a simple probabilistic model called 'Beta reputation model' [WJI05]. We add an initialization phase to the model using the contextual information and the Maximum Likelihood Estimation method. Our work is evaluated through simulations comparing an extended version of the Beta reputation model, which is, as a trust management model, enhanced with our proposed initialization phase.

The remainder of the paper is organized as follows. Section 1.1 provide a brief overview of the trust and reputation systems, our view of context, and short explanation of the Beta reputation model. The proposed method for trust bootstrapping is described in Section 2. Subsequently, in section 3, we present the evaluation plan and the obtained results. Section 4 provides an overview of the related research. Finally, we close by our concluding remarks in section 5.

1.1 Background: Trust and Reputation Systems

Trust and reputation systems represent a significant evolution in support for Internet services, especially in helping users decide among a growing number of choices. The basic idea of a trust system is to let parties generate feedback about each other after completion of a transaction, and aggregating the feedback to derive a reputation score. The reputation score is used to assist others in deciding whether or not to trust that party in the future. Jøsang et al. distinguishes between two categories of trust: *reliability trust* and *decision trust* [JIB07].

Reliability trust is defined based on "the subjective probability by which an individual expects that another individual performs a given action on which its welfare depends." Decision trust is defined as "the extent to which one party is willing to depend on something or somebody in a given situation with a feeling of relative security, even though negative consequences are possible". It is worth mentioning that the problem of trust evaluation that we address here is distinct from the problem of deciding to trust.

A *trust relationship* exists between two users when one has an opinion about the other's trustworthiness and a *recommendation* is a communicated opinion about the trustworthiness of a third party. *Reputation* is defined as an "expectation about an agent's behavior based on information about or observations of his past actions." Therefore, reputation can be considered as a collective measure of trustworthiness (in the sense of reliability) based on the referrals or ratings from members in a community.

An individual's subjective trust can be derived from a combination of received referrals and personal experience.

A reputation system uses a specific method (*e.g.*, averaging, probabilistic-based or belief-based) to compute reputation values based on the collection of feedback from others. Some of the various methods for computing reputation and trust measures include.

- *Rank ordering*: This method has no explicit reputation score and acts as an implicit indicator of reputation. For instance, in Slashdot [1], an online discussion board, readers rate posted comments and postings are prioritized or filtered according to the ratings they receive from readers.

- *Simple summation or average of ratings*: This method is the simplest form of computing reputation scores. The score is the sum of the number of positive ratings and negative ratings, for example, the positive score minus the negative score (*e.g.*, in eBay [2]) or the average of all ratings (*e.g.*, in Epinions [3] and in Amazon [4]).

- *Bayesian systems*: The reputation score is computed by updating Probability Density Functions (PDFs). The updated reputation score is computed as a combination of the previous reputation score and the new rating.

- *Fuzzy models*: These methods represent trust and reputation as linguistically fuzzy concepts, where membership functions describe to what degree an agent can be described as trustworthy or not. Fuzzy logic provides rules for reasoning with fuzzy measures of this type.

- *Flow models*: A participant's reputation increases as a function of incoming flow, and decreases as a function of outgoing flow (e.g., Google's PageRank and Advogato [5]). In the case of Google, many hyperlinks to a web page contribute to increased PageRank whereas many hyperlinks from a web page contributes to a decreased PageRank for that web page.

1.2 The Beta Reputation Model

We have chosen to study the use of a particular probabilistic trust and reputation system called the Beta reputation model (BRM). The BRM models the reputation formation for a truster as a sequence of observations, where each observation is the outcome of the rating done by a trustee, based on the outcome of an interaction. Ratings from all the users are gathered and each user reputation score will be updated based on them. The underlying mathematical model of the BRM considers the ratings

[1] http://slashdot.org/
[2] http://ebay.com/
[3] http://www.epinions.com/
[4] http://www.amazon.com/
[5] http://www.advogato.org/trust-metric.html

as a sequence of trials with binomial outcomes. For each trial there is a probability p of getting a *good* rating (recommendation) and a probability $(1/p)$ of getting a *bad* rating. The parameter p belonging to a truster is initially unknown, so due to lack of information it is assumed that it is drawn from a uniform distribution on [0,1] (The default trust value assumption). As ratings concerning this trustee start to arrive, there is more information available and we can update the distribution of p. In accordance with Bayesian inference we have *a prior* hypothesis X about the outcome of a trial, which is updated a *posteriori* to the actual outcome Y in accordance with *Bayes* Theorem:

$$P(X|Y) = \frac{P(X)P(Y|X)}{P(Y)}$$

The Beta distribution $Beta(\alpha,\beta) = \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} p^{\alpha-1}(1-p)^{\beta-1}$ is a conjugate prior for binomial trials (Bernoulli process). This means that if we assume that the prior X hypothesis is described by $Beta(\alpha,\beta)$, and Y is a sequence of ratings, then the posterior $P(X|Y)$ is also described by a Beta distribution. The initial prior is given by $Beta(1,1)$, which corresponds to the uniform distribution on [0,1] and can be considered as the default trust value. The reputation value is given as a function of the expectation value of the Beta distribution $E(p) = \alpha/(\alpha+\beta)$.

1.3 Context Sensitive Trust Management

Our work is mainly motivated by consideration of context in trust computations because of its ability to bring additional knowledge to the reasoning process. A trust evaluation process is complicated by two facts: *(i)* trust is situation-specific (e.g. a person may trust her financial advisor about investment analysis but does not trust the same advisor related to health-care issues), and *(ii)* trust is person-specific (e.g. judgments of two persons on the same matter or event are often quite different).

Context is defined as "any information that can be used to characterize the situation of an entity. An entity is a person, place, or object that is considered relevant to the interaction between the user and application, including the user and application themselves' ([BD05])'. A system is context-aware if "it uses context to provide relevant information and/or services to the user, where relevancy depends on the user's task" [DA00].

We build on a distinction between the external and the internal context. The term 'external' is used when the relevant context factors are external to the reasoner, relating either to the properties of the trustee or the object to be acted on (e.g., information to be exchanged or something to be bought), i.e. the external context is related to the 'situation'. The facts exist independently from the reasoner, in the sense that they are there before and after the reasoner notices them. The internal context, on the other hand, characterizes the mental and emotional state of the reasoner, i.e. the truster, and is the internal knowledge/mechanisms underlying the person's cognitive processes (e.g. mood and state-dependent effects). The 'state of mind' component of context emerges while the reasoner solves a problem, and captures various internal

parameters of human experience and activity: taste, personal standards, preferences, semantic differences, disposition, bias, and halo (tendency to rate according to a general impression).

2. The Proposed Method

We improve the BRM by the using the Maximum Likelihood Estimation (MLE) method to estimate the parameters α and β (which can be considered as the initial trust value) based on the history of transactions. This can be seen as a hierarchical Bayesian model for reputation, HBRM.

MLE is a popular statistical method used for fitting a statistical model to data, and provides estimates for the model's parameters. For a fixed set of data $x_1, x_2, \ldots, x_n$ in training set D and underlying probability model f_θ, maximum likelihood picks the values of the model parameters θ that make the data "more likely" than any other values of the parameters would make them, i.e. $\hat{\theta} = argmax_\theta f_\theta(x_1, \ldots, x_n)$.

In other words, we have a random variable and we know the form of its probability distribution, but we do not know the exact values of involved parameters. For instance, we might know that a given variable is Beta distributed but we may lack the exact value of the distributions parameters α and β. We are concerned with estimating these unknown values. Knowing the distribution type we can compute the likelihood of any sample set generated from the distribution for general values of the unknown parameters. Having a set of realizations of the considered random variable, we can simply fit the values of the unknown parameters that maximize the computed likelihood.

Therefore, we use the estimated values of α and β from MLE instead of the default values $\alpha = 1$ and $\beta = 1$ for initialization. This results in a hierarchical Bayesian method for trust and reputation evaluation, since we use another prior to estimate the prior for the original BRM.

We collect those past assigned trust values to form the training set D of experiences. There are two classes of generalized information, called classifiers, which we can use for forming the set D when estimating the parameters α and β: (i) *The external context* that chooses past trust values in the same situation and (ii) *The internal context* which selects past trust values with the same trustee properties. In the following, we present the details about these two generalization classifiers.

- The external context: all other trustees with whom the truster has had experiences in the same situation as the current situation are classified into a group as the training set D for the MLE method. In other words, the general trustworthiness of them will be a basis for a typical behavior in this situation.

- The internal context: all other trustees who have common attributes with the prospect are grouped together to form the training set D for the MLE method. That is, the general opinion about this group will be transformed into an opinion about the prospect.

When a user first comes into a system, there is little information available for building a trust evaluation. Further, gathering such information is difficult when there is little incentive to provide feedback. We categorize the decision making process with respect to these two factors based on the familiarity of the truster with the situation and the trustee. Different combinations of incomplete knowledge are:

- Familiar situation, familiar trustee: If the truster has previously had interactions with the same trustee in the same situation, then she can immediately use her past experiences to predict the outcome of the new interaction and take a decision on this basis. Therefore, there is no need for initialization.
- Unfamiliar situation, familiar trustee: If the truster has had previous interactions with the trustee or similar other trustees, but in different situations, he/she can still use his/her past experiences. In this case we group these previous trustees based on the situation (external context) similarities and use this set for the initialization phase.
- Familiar situation, unfamiliar trustee: The truster has previous interactions in the same situations with other trustees. Therefore, past trustees are grouped based on a common attribute with the current trustee (i.e. based on their internal context similarities) and this set will be used for the initialization phase.
- Unfamiliar situation, unfamiliar trustee: If there is no situational or trustee information available, the trust model uses a default trust value since there is no information to be used for the initialization of trust.

Figure 2 shows the decision making process underlying the initialization of trust values.

After forming the training set D, we can estimate α and β. Suppose $Beta(\alpha,\beta) = D = \{x_1, x_2, \ldots, x_n\}$. The likelihood function is

$$f(X|\alpha,\beta) = n[ln\Gamma(\alpha+\beta) - ln\Gamma(\alpha) - ln\Gamma(\beta)] + (\alpha-1)\sum lnx_i + (\beta-1)\sum ln(1-x_i)$$

This cannot be analytically maximized, so we use the *Newton Raphson* iteration method [Smi83] to find the maximum likelihood estimate. Because of space limitation, we do not give any detailed description of our use of this algorithm here.

3. Evaluation of HBRM vs. BRM

In evaluating our approach, we employed a simulated agent society where a set of truster agents interacts with a set of trustee agents over a number of rounds. Each trustee is assigned a performance profile that determines how it will behave and the ratings are binary. We compare performance of the original BRM with the model enhanced by the initialization phase under some common attack scenarios for trust and reputation systems. In the following we explain these scenarios in more detail:

- **Reputation lag problem and playbooks (case1)**: There is usually a time lag between an instance of a transaction and the corresponding effect on the

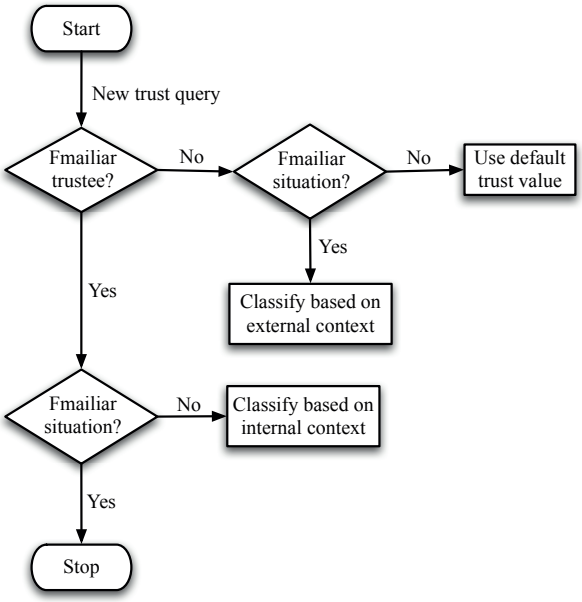

Figure 2. Decision making process for initialization of trust value.

reputation score (*e.g.*, in eBay, the buyer pays before the seller ships the item). A user has the opportunity to make use of this time lag to provide a large number of low quality services over a short period before the reputation score suffers any significant degradation [KC06]. Further, the re-entry problem can be combined with this problem in a way that a seller may re-enter the market each time a buyer learns of a dishonest seller. In this way, a seller can take advantage of a reputation lag repeatedly. A playbook is a sequence of actions that maximizes profit of a participant according to certain criteria. A typical example is to act honestly and provide quality services over a period to gain a high reputation score, and then to subsequently milk the high reputation score by providing low quality services at a low production cost [KC06].

- **Time sensitivity of reputation (case2)**: treating old positive behavior equal to new negative behavior may result in attackers abusing the system by using previous altruism to hide current malicious behavior.

- **Unlimited memory problem (case 3)**: Most reputation calculation algorithms use all transactions when calculating the overall score, thus, a new user might not understand how a site functions [Mal01]. Besides, a user can perform short duration malicious attacks with little risk of negative consequences because a lengthy previous history can heavily outweigh current actions. This problem can have a large impact on the system as the malicious users will continue to

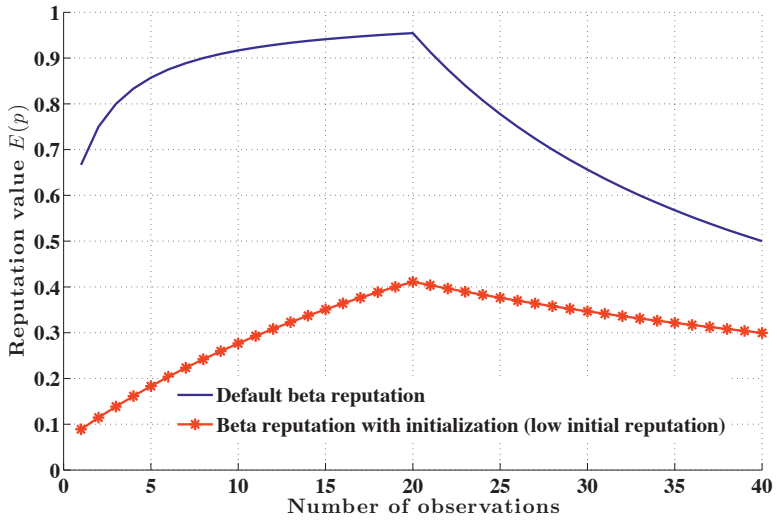

Figure 3. The BRM model compared to the HBRM model. The input is 20 good observations followed by 20 bad observations and a set of 10 low initial reputation values.

have a high reputation for a substantial period of time during which the system is slow to identify the malicious behavior and unable to sufficiently lower the user's reputation [HZNR09].

3.1 Results

Here we present the results of simulation for the three scenarios, namely case1, case2, and case3.

3.1.1 Case 1

We assume a scenario where we have 20 good observations followed by 20 bad observations. Such an input set of observations could come from a trust scenario where an agent builds its reputation value by behaving well for a certain amount of time, and then decides to take advantage of its good reputation by suddenly changing its behavior.

In Figure 3, we clearly see the difference between the results obtained when using the two models. Both simulations assume 10 low reputation values available by gathering information about the trustees in the same situation (external context) or similar trustees (internal context). These low reputation values are 10 randomly generated numbers in the range [0, 0.1]. We see the different slopes and different convergence of the two graphs. Both graphs have approximately the same increase in the reputation value after 20 good ratings; however, the BRM decreases more rapidly

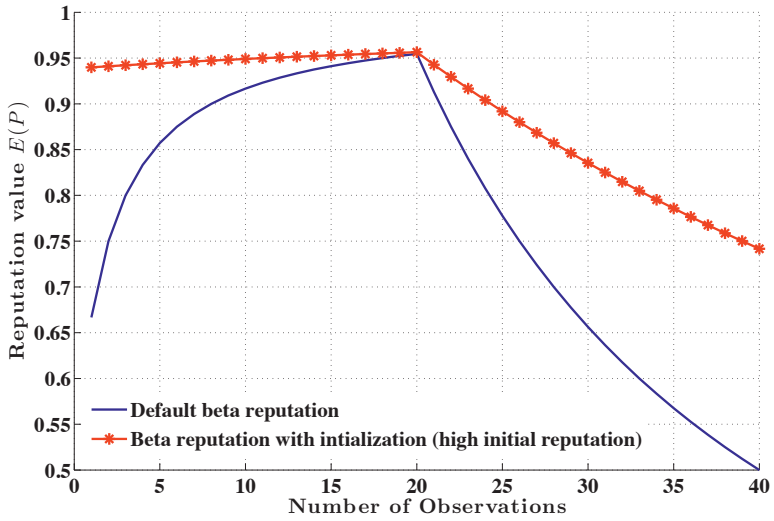

Figure 4. The BRM model compared to the HBRM model. The input is 20 good observations followed by 20 bad observations and a set of 10 high initial reputation values.

than the HBRM after 20 bad observations. This indicates that the HBRM estimation is more realistic than the BRM, and the BRM model is going to converge to the HBRM.

In Figure 4 we see the results for the same scenario with 10 high initial reputation values, which are 10 randomly generated numbers in the range [0.9, 1]. We observe the different slopes and convergence of the two graphs. It shows that HBRM is more stable with this kind of initialization. However, the decrease in reputation value is much stronger for the BRM after 20 bad observations. The effect of 20 bad observations for an agent from a well reputed community is often considered less damaging than for an unknown agent without any initial assumption.

In Figure 5 we see the results for the same scenario with 10 random initial reputation values, which are 10 randomly generated numbers in the range [0, 1]. We observe the same behavior of the two models and their convergence to the same reputation value, as one would expect.

3.1.2 Case 2

We assume a scenario wherein an agent has been compromised, i.e. taken over by a malicious agent. The agent then proceeds with a strategy of laying low, meaning that is waits for a long time without acting malicious, so that when it starts to show malicious behavior it can take full advantage of the good reputation that the previous owner of the agent had built up. An example of such a scenario will be having 9 good observations, then one bad observation at time $t = 10$, then no observations until time $t = 35$, followed by 5 bad observations.

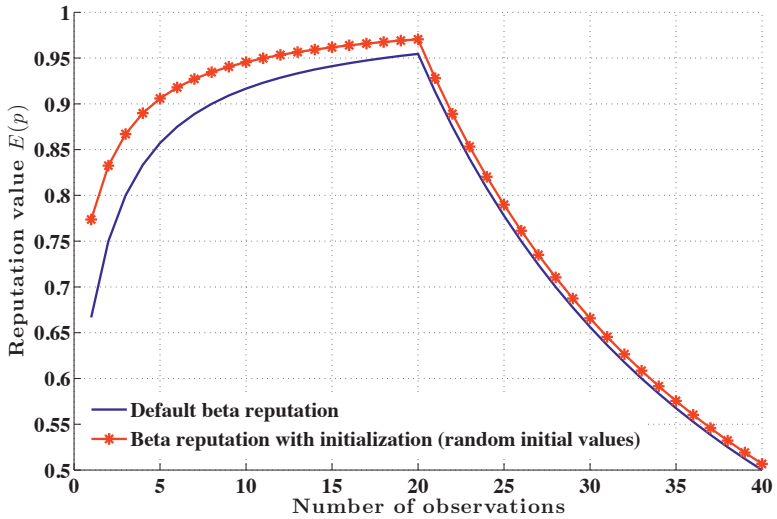

Figure 5. The BRM model compared to the HBRM model. The input is 20 good observations followed by 20 bad observations and a set of 10 random initial reputation values.

We can observe from Figure 6 that the BRM gives a steeper slope, while the HBRM with various initial values are just stretched at the x-axis. In particular, when there are no observations between time $t = 10$ and time $t = 35$, the BMR increases; in contrast to HBRM model which do not change. The latter behavior is considered more realistic.

3.1.3 Case 3

We want to see how the models react to a disruptive agent who changes its strategy. In particular, we consider an agent who follows a pattern of misbehavior adapted to a detection rule of three strikes and you are out. We have an example of this scenario, where an agent is showing good behavior for 10 observations to build up its reputation, and then proceeds with the disruptive behavior giving a pattern of 2 bad observations, one good observation, 2 bad observations, and so on.

In Figure 7, we can see that the BRM picks up this behavior with a decreasing reputation value, but the HBRM models with low and high initial values do not change considerably. Hence, the original model works equally well or better for this scenario.

We found that for each of the five conditions, the HBRM outperformed the BRM model after the first learning interval. In each case, the HBRM performs similarly to the standard model while training examples are gathered. However, once the first learning interval has passed, the outcome of the HBRM begin to improve, whereas BRM do not. The results we have presented show that an initialization mechanism based model can clearly help agents to make trust evaluations in situations where both

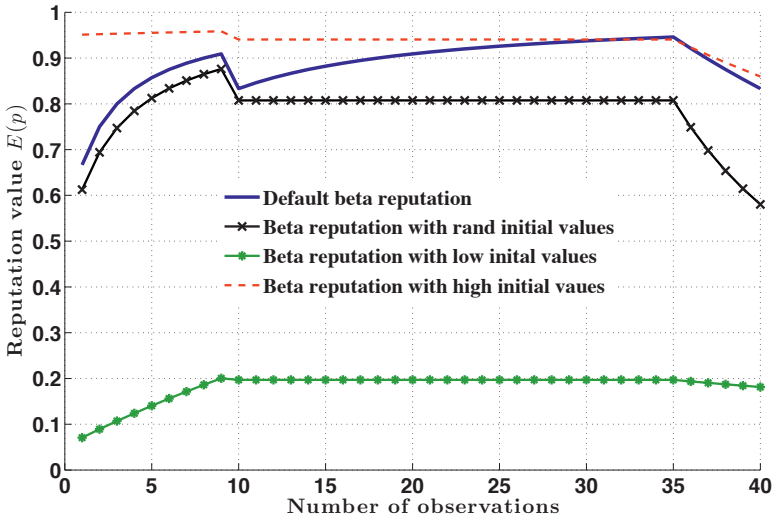

Figure 6. The BRM model compared to the HBRM models. The input is 9 good observations, 1 bad observation at time t=10, then no observation until time t=35 followed by 5 bad observations.

direct and reputational evidence is not forthcoming. One possible drawback to our approach is the use of the learning interval to control the formation of learning sets for the MLE. While we have referred to a number of trust evaluation models in this paper, it is worth highlighting here some related approaches, which attempt to address the issues of specific interest.

4. Related Work

This section gives a brief overview of the sparse but highly relevant research on trust initialization. The initialization problem may be alleviated by taking into consideration the interconnections among trust and reputation systems and social networks [JL98]. For example, the location of a given member of a community within a social network can be used to infer some properties about his or her degree of expertise, i.e., his or her reputation.

The FIRE system [HJS06] employs role-based trust to explicitly capture relationships between agents in certain roles. In this approach, rules specify an initial, predetermined degree of trust that will be conferred on partners for whom the rules match. This means that a degree of trust may be present even when no evidence is available. FIRE rules are explicitly specified in a domain specific manner by agent owners.

Burnett et al. [BNS10] propose a data-mining based categorical trust model. They propose that using the personal interactions with previously encountered trustees, the truster can derive some rules that allow him to characterize other trustees with

specific features as less or more trustworthy. These rules are learned using regression trees. Each rule maps trustees with specific features onto a trust value in the range [0,1]. The model adopts stereotypes as patterns for recognizing trustworthy agents. Following this structure, a Stereotrust agent has been implemented and performs categorical reasoning in three phases:(i) stereotypes rise from generalization of past experiences and are built using data mining and machine learning techniques; (ii) if direct experiences of past interaction with the same trustee are available, then trust is the average of the previous delegation results; (iii) otherwise, given trustee's manifesto and environmental conditions, stereotypes are applied as a filter to determine to which cluster the trustee belongs, thus finding the relative trust value. When a Stereotrust agent has stored an amount of experiences on the same task, it identifies some patterns for recognizing clusters of performers, thus associating them some appraised trust based on previous delegation results.

Tavakolifard et al. [THO09] present a context management framework that employs case-based reasoning [Mor94] to analyze the correlation between trust information among various situations and help to bootstrap in unanticipated situations using trust information available from similar situations. The case-based reasoning (CBR) technique is particularly useful for tasks that are experience-intensive, that involve plausible (i.e. not sound) reasoning and have incomplete rules to apply. The CBR technique [Kol93, AP94] is particularly useful in open and weak domains that lack the complete and certain knowledge and thus needs to exploit experience based knowledge. The fundamental principle of CBR is similar to human analogical reasoning [Gen83, HT97] in the sense of using solutions of past problems to solve the current similar ones. Two main components of a CBR system are the case base storing a number of previously solved cases and the CBR engine that finds and uses the previously solved cases (in the case base) in order to solve a new case. A case comprises two parts: a situation/problem description and a solution (only the past cases). In the presented work, a new case is a trust assessment query specifying the truster, trustee, and the other contextual information. Context has been shown to have major influences on remembering and comparing cases. The strong dependency between the context and a powerful memory-retrieval arises most probably from the role context plays in the similarity assessment of two cases (i.e., the new and a past case). The query is matched with the problem description part of the cases in the case base and the cases are ranked according to their similarity with the query. The retrieved case provides a solution, which is the trust value that the truster assigns to the trustee.

Rehak et al. [RGPB06] define a set of reference contexts in a metric space and associate truthfulness data with it. These data are updated and queried with weight that decreases with distance between the current situation and the reference context. The model uses Leader-Follower clustering to identify the reference contexts to be representative of the data. In this work, two advanced uses of context for multiagent trust modeling is proposed: (i) policy/norm learning at runtime by analyzing data regarding the performance of different agents in similar situations (e.g. when all agents fail in a certain situation, they may agree to introduce a policy that specifically prohibits such actions) (ii) reasoning based on uncertain identities by decomposing

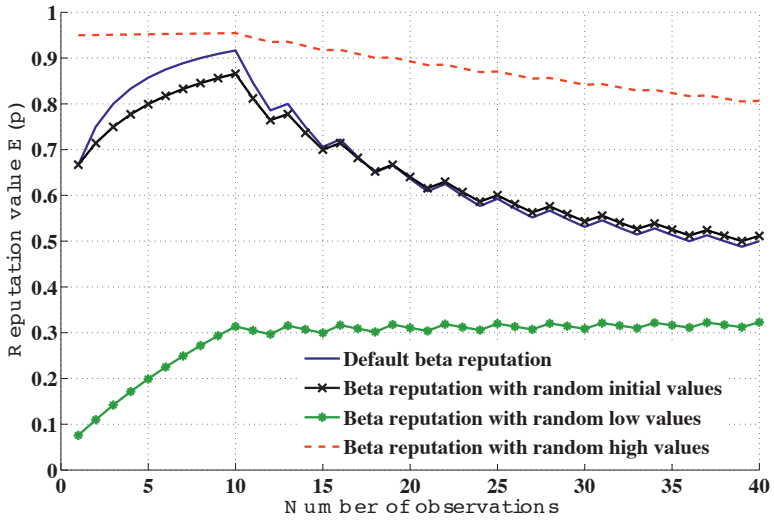

Figure 7. The Beta model (BRM) compared to the Hierarchical Beta models (HBRM). The input is 10 good observations, followed by a disruptive behavior giving a pattern of 2 bad observations, one good observation, 2 bad observations, and so on.

the single identity dimension into an identity subspace, where each agent is defined by one or more crucial properties. With this modification, the trust model can make predictions about the performance of agents by exploiting data characterizing a similar agent's performance in the past. Based on this model, authors [RP07] conclude that the extension of a trust model with a context representation can relax the existing models assumptions like: (i) proven identity, (ii) repetitive interactions and (iii) similar trusting situations. The fact that two agents with presumably distinct identities can be considered identical by a context-sensitive trust model may provide protection against changes of identities. This approach is also effective against first time offenders; we can obtain a model with inductive properties, which is able to estimate the performance of new entrants using the experience with the similar partners in the past.

It should be noted that Despotovi and Aberer [DA04, DA06] propose the MLE method as a feedback generation strategy for computing the reputation value, however, the BRM in our model is the feedback aggregation strategy and we propose the MLE for estimation of its parameters as a step towards initialization of the trust value, which results in a hierarchical Bayesian model for reputation.

5. Conclusion

In open systems, a number of situations can arise where a trust evaluation must be made, but no direct or recommended supporting evidence can be found. When a user is completely new in the system, direct evidence or recommendations can only obtained

(and subsequently propagated as reputation) when users takes a risk and interacts with the newcomer. We propose an approach for improving the performance of trust mechanisms in such initial cases by allowing trust evaluations to be "bootstrapped" by a priori assumptions based on contextual information. We have demonstrated how this approach can be used together with a relatively straightforward probabilistic trust model in order to significantly improve performance. The approach presented here can complement existing probabilistic trust evaluation techniques.

We have seen from the simulated examples that the BRM model and the HBRM model perform differently. In the hierarchical model, we assume another prior to estimate the prior for the original Bayesian process, which will be resulted in a hierarchical Bayesian process. In this way, we are able to estimate the disposition of the truster agent. We use information about reputation of agents in the similar context or reputation of other similar trustees as training data for the maximum likelihood estimator. We have shown the simulated results for the binary rating systems. Furthermore, our proposal can be extended for the multinomial systems with several levels of rating.

References

[AP94] A. Aamodt and E. Plaza. Case-based reasoning. *Proc. MLnet Summer School on Machine Learning and Knowledge Acquisition*, pages 1–58, 1994.

[BD05] N.A. Bradley and M.D. Dunlop. Toward a multidisciplinary model of context to support context-aware computing. *Human-Computer Interaction*, 20(4):403–446, 2005.

[BNS10] C. Burnett, T.J. Norman, and K. Sycara. Bootstrapping trust evaluations through stereotypes. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: volume 1-Volume 1*, pages 241–248. International Foundation for Autonomous Agents and Multiagent Systems, 2010.

[DA00] A.K. Dey and G.D. Abowd. Towards a better understanding of context and context-awareness. In *CHI 2000 workshop on the what, who, where, when, and how of context-awareness*, pages 304–307, 2000.

[DA04] Z. Despotovic and K. Aberer. Maximum likelihood estimation of peers' performance in p2p networks. In *The Second Workshop on the Economics of Peer-to-Peer Systems*, pages 1–6. Citeseer, 2004.

[DA06] Z. Despotovic and K. Aberer. P2p reputation management: Probabilistic estimation vs. social networks. *Computer Networks*, 50(4):485–500, 2006.

[Gen83] D. Gentner. Structure-mapping: A theoretical framework for analogy. *Cognitive science*, 7(2):155–170, 1983.

[HJS06] T.D. Huynh, N.R. Jennings, and N.R. Shadbolt. An integrated trust and reputation model for open multi-agent systems. *Autonomous Agents and Multi-Agent Systems*, 13(2):119–154, 2006.

[HT97] K.J. Holyoak and P. Thagard. The analogical mind. *American Psychologist*, 52:35–44, 1997.

[HZNR09] K. Hoffman, D. Zage, and C. Nita-Rotaru. A survey of attack and defense techniques for reputation systems. *ACM Computing Surveys*, 42(1):1–31, December 2009.

[JIB07] A. Jøsang, R. Ismail, and C. Boyd. A survey of trust and reputation systems for online service provision. *Decision Support Systems*, 43(2):618–644, March 2007.

[JL98] S.L. Jarvenpaa and D.E. Leidner. Communication and trust in global virtual teams. *Journal of Computer-Mediated Communication*, 3(4):0–0, 1998.

[KC06] R. Kerr and R. Cohen. Modeling trust using transactional, numerical units. In *Proceedings of the International Conference on Privacy, Security and Trust: Bridge the Gap Between PST Technologies and Business Services*, pages 21:1–21:11, October/November 2006.

[Kol93] J. Kolodner. *Case-based reasoning*. Morgan Kaufmann Publishers Inc. San Francisco, CA, USA, 1993.

[Mal01] R.A. Malaga. Web-based reputation management systems: Problems and suggested solutions. *Electronic Commerce Research*, 1(4):403–417, 2001.

[Mor94] B.W. Morris. SCAN: a case-based reasoning model for generating information system control recommendations. *International Journal of Intelligent Systems in Accounting, Finance and Management*, 3(1):47–63, 1994.

[RGPB06] M. Rehak, M. Gregor, M. Pechoucek, and J.M. Bradshaw. Representing Context for Multi-agent Trust Modeling. In *Proceedings of the IEEE/WIC/ACM International Conference on Intelligent Agent Technology (IAT 2006 Main Conference Proceedings)(IAT'06)-Volume 00*, pages 737–746. IEEE Computer Society Washington, DC, USA, 2006.

[RP07] M. Rehak and M. Pechoucek. Trust modeling with context representation and generalized identities. *Klusch, M., Hindriks, K., apazoglou, MP, Sterling, L.(eds.) CIA*, pages 298–312, 2007.

[Smi83] DM Smith. Maximum likelihood estimation of the parameters of the beta binomial distribution. *Appl. Stat*, 32:196–204, 1983.

[THO09] M. Tavakolifard, P. Herrmann, and P. Öztürk. Analogical Trust Reasoning. In *Trust Management III, IFIP Advances in Information and Communication Technology, Volume 300. ISBN 978-3-642-02055-1. Springer Berlin Heidelberg, 2009, p. 149*, page 149. Springer, 2009.

[WJI05] A. Whitby, A. Jøsang, and J. Indulska. Filtering out unfair ratings in bayesian reputation systems. *The Icfain Journal of Management Research*, 4(2):48–64, 2005.

PAPER F

A Probabilistic Reputation Algorithm for Decentralized Multi-Agent Environments

Mozhgan Tavakolifard, Svein J. Knapskog

Electronic Notes in Theoretical Computer Science
Vol. 244, 2009

A PROBABILISTIC REPUTATION ALGORITHM FOR DECENTRALIZED MULTI-AGENT ENVIRONMENTS

Mozhgan Tavakolifard,[1] Svein J. Knapskog,[1]

[1] *Centre for Quantifiable Quality of Service in Communication Systems*
Norwegian University of Science and Technology
{mozhgan, knapskog}@Q2S.ntnu.no

Abstract The importance of trust in electronic transactions is well understood. The majority of current trust models consist of a central entity that verifies compliance with the trust requirements, using standardized evaluation methods and criteria. In decentralized environments, the communication scenarios are more complex, and no universally accepted objective requirements or evaluation criteria exist. It should be noted that the situation would get even more complicated when agents are interacting with each other. The goal of this research is to model trust and reputation in decentralized multi-agent systems. To achieve this, we have chosen the Ntropi model, among several other models, as a starting point, The efficiency of the model in such scenarios has been significantly improved by introducing a new probabilistic reputation algorithm for the Ntropi model.

1. Introduction

The rapidly changing environments of the internet suffer from problems like fragile trustworthiness of millions active entities on the internet, e.g., humans and mobile agents. This problem is nontrivial, as more and more commercial transactions get carried out over the internet. Therefore, devising an effective approach for verification of trustworthiness in such complex environments is essential, since the trust mechanisms play a key role in the security of multi-agent systems. Also the trust establishment is nontrivial, since the traditional and social means of trust cannot be applied directly to virtual settings of these environments because in many cases the involved parties did not have any previous interaction. In such scenarios, reputation techniques may be used to stimulate service quality and acceptable user behavior in online markets and communities, and also sanction possible unacceptable user behavior. To this end, the Ntropi model [ARH98] was designed to facilitate the exchange of trust and reputation in information and/or business environments. The Ntropi classifies the trust into direct (explicit) and recommended classes. The direct class is based on the truster agent's previous personal experiences with the trustee agent. But the recommended trustworthy class is derived from word-of-mouth (e.g., opinions), which is called

reputation, and can be translated into direct or regular trust. This paper presents an automated and autonomous trust system using Bayesian inference along with improved Dirichlet distribution. Our main contributions are the application of maximum likelihood method in the trust/reputation model to estimate the parameters used in Dirichlet distribution, and also the introduction of a hierarchical Bayesian method in the proposed reputation management model. The maximum likelihood estimation method has been previously introduced in [DA04] as a feedback aggregation strategy. However, in this work the bootstrapping (when two unfamiliar agents face each other) is the main concern.

The rest of this paper is organized as follows: section 2 covers the relevant literature. In section 3 the Ntropi model and its analysis are presented. Section 4 discusses the proposed model in detail. Section 5 explains the experimental results and presents the evaluation process. Section 6 presents the conclusion and suggests future work.

2. Literature Review

We chose the Ntropi model [ARH98] among several other models because: 1) this model is mainly designed for decentralized multi-agent systems, 2) it covers more trust aspects in this area than other models, 3) it is a well received model in academia, 4) its proposed elements have been incorporated into Sun's JXTA framework [CY01] and Ericsson's trust model [QL04, QOL$^+$05]. The JXTA is an open source and a general purpose P2P framework currently available. Furthermore, the implementations have been analyzed in various popular P2P platforms such as Gnutella [Gen], Free Haven [Sni00] and Freenet [CMH$^+$02].

On the basis of recent surveys among existing reputation algorithms, the probabilistic algorithms, especially those with Bayesian inference seems to be more popular. Because these algorithms have a sound mathematical basis and are known to be suitable to formulate human characteristics, they are more flexible than the Ntropi's ad-hoc algorithm and need less interaction with users. Thus, the first feature in agent's definition, autonomy, seems more realistic.

The majority of Bayesian-based reputation algorithms are binomial (e.g. [BLB04]), allowing two-valued ratings, as either positive (e.g. good) or negative (e.g. bad). The main disadvantage of a binomial model is that it is not able to represent ratings with graded levels such as e.g. mediocre - bad - average - good excellent. In addition, the binomial models are in principle not able to distinguish between polarized ratings (i.e. many very bad and many very good ratings) and average ratings. The Ntropi offers graded multinomial ratings: for example "very trustworthy", "trustworthy", "moderate", "untrustworthy",and "very untrustworthy" which is more realistic. There are also several Bayesian based reputation models with graded ratings which seem more suitable. Some of these models have used Dirichlet as a priori distribution and multinomial models as likelihood distribution in their Bayesian inference.

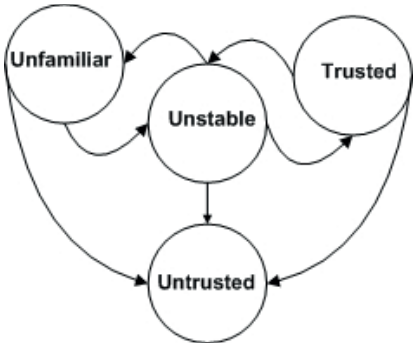

Figure 1. Phases of a trust relationship, with arrows indicating possible direction of phase transition.

3. Ntropi Model

In this section the Ntropi model by Farez Abdul-Rahman [ARH98] which forms the basis of our proposed model, is explained and an analysis of the model is given.

3.1 Model Description

The Ntropi is a trust model that is truly decentralized. It has no reliance on any third party and all entities can decide for themselves how to trust. It uses both reputational and experiential information. Recommendation which is a single opinion and reputation which is multiple opinions are combined. Trust values have a five-level scale: "very trustworthy", "trustworthy", "moderate", "untrustworthy",and "very untrustworthy". After receiving recommendations from recommenders about a prospect, a truster agent may decide to go ahead with the interaction. After that, he may give his experience a rating and notice the difference between his own rating and the recommended rating. This difference (called semantic distance in this model) shows the difference in rating standards.

These differences are recorded so that in the future the truster agent can adjust his trust values accordingly. Based on this history of differences, a translation table will be formed and recommendations will be translated. In order to turn "what he said" into "what we think he means" we get the most common semantic distances and add that to the recommended value. In order to combine more than one recommendation and calculating reputations, we need to know the trust in each recommender and give more weight to recommendations from more trustworthy recommenders. In Ntropi model, a trust relationship goes through phases. At any point in time, a trust relationship will exist in one of four phases, as shown in Fig.1. Recommendations in different phases may be considered in different ways.

We calculate our trust in the recommender based on the consistency of the recommenders' previous recommendations. If the distributions of semantic distances are more spread out, then there is less consistency. The less the spreading is, the more consistent the recommender is regarded to be. In this model the consistency is

obtained by first finding the semi-interquartile ranges (SIQR) of the ordered set of semantic distances for the active context, rounded to the nearest integer. Then a lookup table is used to convert the SIQR into a trustworthiness level. Then we assign weights to recommenders according to their trust value. For each recommended trust value the weights of those who recommended it will be summed up. The final reputation value is the trust value with the highest sum-of-weights.

If a recommender is not known by the recommendation requester, the requester can obtain recommendations about the unknown recommender. There is also the scenario where a recommendation requester may carry out a network search for a particular agent and the received recommendation may be the result of the request being forwarded through a number of intermediary recommenders. In both scenarios, when a recommender recommends another recommender, the result is a recommendation chain. The heads of the chain may contain more than one known recommender, all of which recommends the same first intermediary of the chain. an agent seeking recommendations about an unknown prospect will request recommendations from those recommenders that he already knows and trusts. Thus, a chain's heads should be a known recommender [ARH98].

3.2 Model Analysis

Our analysis of Ntropi model is as follows:

- The SIQR method is just one approach for finding the spread of semantic distances in this model. Other measures of dispersion in the data may be more appropriate for different applications, especially one where the requirement of unbounded, unimodal and symmetrical distribution (for which the SIQR is suitable for) does not exist.

- The SIQR, however, does not include all data points in the distribution, which may be another consideration when determining an appropriate spread measurement the standard deviation, for example, does include all data points.

- Furthermore, when converting the SIQR (or whatever the spread measure is) into a trust level, linear conversion need not be assumed. However, one may select different trust values for each SIQR value, depending on the weight one gives to the different SIQR/spread values.

- Another possible weakness of the approach taken in the Ntropi, where all chain heads must be known, is that it will not be possible to accept recommendations from chains with unknown heads, even if the requester is willing to use those recommendations. An example where unknown recommenders may be useful is when the alternative is to have no recommendation at all. This situation is analogous to asking for directions on the street, and demonstrate that at times one may successfully use advice from a stranger, i.e. when nothing is known about the recommending agent. This is particularly true in situations where possession of any information is better than no information, and, at the same

time, there is belief in the benevolence of the recommender as well as low perceived risk.

- We shall observe here that the weighted trust level approach provides a potential for customization and flexibility in weighting recommendations based on the trust levels of their recommenders. However, it also adds to the user's list of tasks to perform, namely that he must be able to define, and, if required, adjust the weights for each application that uses this model. In reality, this is not a very satisfactory situation since it will require additional help from the application itself in terms of either hardwiring the weighting based on well known properties of agents in the application domain, or employing some form of learning algorithm that can dynamically update the weights based on experience.
- Since its reputation algorithm results in the selection of recommendations with the highest weightings, it will potentially be ignoring other recommendations that also originated from trustworthy recommenders, albeit from those with lower comparative trustworthiness levels. Given a sufficiently high number of recommendations (for the same trust value) from lower trust recommenders, their recommended trust values may still be the winning value because their sum-of-weights will outweigh the recommenders with higher trust but with a lower population within the local set of recommendations. A better algorithm would be one where a new trust value is produced by the reputation/combination algorithm based on the recommendations received from all the trusted recommenders from the whole range of trust levels.

4. The Proposed Reputation Algorithm

The new reputation algorithm proposed in this paper is based on the Dirichlet reputation algorithm (as also) proposed by A. Jøsang [JH07]. We have improved this algorithm by using Maximum Likelihood Estimation method to estimate the parameters (for this algorithm) based on the observed data.

Jøsang's reputation algorithm is based on Bayes Theorem.

$$P(\Theta|X) \propto P(X|\Theta).P(\Theta) \tag{1}$$

Reading from left to right, the formula is interpreted as saying: the probability of the hypotheses Θ posterior to the outcome of experiment X is proportional to the likelihood of such outcome under the hypotheses multiplied by the probability of the hypotheses prior to the experiment. In the present context, the prior Θ will be an estimate of the probability of each potential outcome in our next interaction with principal p, whilst the posterior will be our amended estimate after one such interaction took place with outcome X.

It is important to observe here that $P(\Theta|X)$ is in a sense a second order notion, and we are not interested in computing it for any particular value of Θ. Indeed, as Θ is the unknown in our problem, we are interested in deriving the entire distribution in order to compute its expected value, and use it as our next estimate for trustworthiness.

In Ntropi model, trustworthiness of an agent can be referred as "very trustworthy", "trustworthy", "moderate", "untrustworthy",and "very untrustworthy". So the rating level is a discrete set. But in the Bayesian model the rating is a real number between 0 and 1. We should use a multinomial probability distribution for the likelihood in the Bayesian inference. Then the conjugate prior distribution will be Dirichlet distribution.

DEFINITION 1 *Agent A's trust in agent B is the accumulation of evaluations that agent A has of its past interactions with B. It reflects agent A's subjective viewpoint of B's capability. Trust value is denoted by θ_{dt} because it is direct trust.*

DEFINITION 2 *The reputation of agent B, from agent A's perspective, is the collective evaluation based on other agents' evaluations of B. It is an objective measure for agent B's capability, resulting from the evaluations of many other agents. Reputation value is denoted by θ_{rt} because it is recommendation trust.*

The estimator for successful cooperation is a combination of trust value and reputation value.

$$\hat{\theta} = w_1\hat{\theta}_{dt} + w_2\hat{\theta}_{rt} \tag{2}$$

Where w_1 and w_2 satisfy $w_1 + w_2 = 1$. They are weights to represent the importance of these two probabilities respectively and are decided by the personal characteristics of the agents.

4.1 The Unfamiliar Phase

The maximum-likelihood method estimates the parameters for the Dirichlet distribution. The parameters are not available in closed-form. We use a simple and efficient iterative scheme for obtaining the parameter estimates in this model from past experiences with other agents. This is our main contribution.

The Dirichlet distribution captures a sequence of observations of the k possible outcomes with k positive real parameters $\alpha(\theta_i)$, $i = 1...k$, each corresponding to one of the possible outcomes. The parameter α can be estimated from a training set with proportions: $D = \{p_1, p_2, ..., p_N\}$.

If agents A and B are complete strangers, i.e. B is in the unfamiliar phase with respect to A when these two strangers first meet, then A will need to collect those past experiences within the same context (context qualifies a trust opinion, describing what the truster's belief in another's trustworthiness is really about) as that in which he encounters B and summarize that set of experiences which will be the training set D. There are two classes of generalized information, called classifiers, which A can use for forming the set D when estimating the parameter α: *Context Experience* and *Stereotype*. Table 1 details these two generalization classifiers.

The maximum likelihood estimate of α maximizes $p(D|\alpha) = \prod_i p(P_i|\alpha)$. The log-likelihood can be written

$$\log p(D|\alpha) = N\log\Gamma(\sum_k \alpha_k) - N\sum_k \log\Gamma(\alpha_k) + N\sum_k (\alpha_k - 1)\log\bar{p}_k \tag{3}$$

Where $\log\bar{p}_k = \frac{1}{N}\sum_i \log p_{ik}$

Table 1. Generalized information classifier for first encounters

Description	**Classifier**
Context Experience	General trustworthiness of all other trustees we have experienced in the current context and use this as a basis for a typical behavior for trustees in this context
Stereotype	Groups past trustees based on a common attribute with the prospect and summarize the general trustworthiness of those trustees. We then form an opinion about those trustees as a group and include the prospect in that group, effectively transforming our opinion about the group into an opinion about the prospect

This objective is convex in in α since the Dirichlet distribution is the exponential family. This implies that the likelihood is unimodal and the maximum can be found by a simple search. The gradient of the log-likelihood with respect to one α_k is

$$g_k = \frac{d \log p(D|\alpha)}{d\alpha_k} = N\Psi(\sum_k \alpha_k) - N\Psi(\alpha_k) + N \log \bar{p}_k \tag{4}$$

$$\Psi(x) = \frac{d \log \Gamma(x)}{dx} \tag{5}$$

Ψ is known as the digamma function and is similar to the natural logarithm. As always with the exponential family, when the gradient is zero, the expected sufficient statistics are equal to the observed sufficient statistics. In this case, the expected sufficient statistics are

$$E[\log p_k] = \Psi(\alpha_k) - \Psi(\sum_k \alpha_k) \tag{6}$$

The observed sufficient statistics are a $\log p_k$.A fixed-point iteration for maximizing the likelihood, and can be derived as follows. Given an initial guess for α, we construct a simple lower bound on the likelihood which is tight at α. The maximum of this bound is computed in closed-form and it becomes the new guess. Such iteration is guaranteed to converge to a stationary point of the likelihood. For the Dirichlet, the maximum is the only stationary point.

As shown in [Ron89], a bound on $\Gamma(\sum_k \alpha_k)$ leads to the following fixed-point iteration:

$$\Psi(\alpha_k^{new}) = \Psi(\sum_k \alpha_k^{old}) + \log \bar{p}_k \tag{7}$$

This algorithm requires inverting the Ψ function a procedure which is described in [Ron89].

4.2 The Trusted and the Unstable Phases

Assume that A is the truster agent, B is the trustee agent and C is the recommender agent. Let there be k different discrete rating levels. This translates into having a state space of cardinality k for the Dirichlet distribution (in the case of our model k is 5). Let the rating level be indexed by i.

Each new rating of agent B by an agent C takes the form of a trivial vector where only one element has value 1, and all other vector elements have value 0. The index i of the vector element with value 1 refer to the specific rating level.

As a result of a new rating, the rating vector will be updated by adding the newly received rating vector $\vec{r}$ to the previously stored vector $\vec{R}$ (Bayesian inference). Agents may change their behavior over time, so it is desirable to give relatively greater weight to more recent ratings. This can be achieved by introducing a longevity factor $\lambda \in [0,1]$; which controls the rate at which old ratings are aged and discounted as a function of time. With $\lambda = 0$, ratings are completely forgotten after a single time period. With $\lambda = 1$, ratings are never forgotten. After encounters with other agents new $\vec{\alpha}$ will be calculated as follows:

$$\vec{\alpha}_{new} = \vec{\alpha}_{old}.\lambda + \vec{R} \qquad where\ 0 \leqslant \lambda \leqslant 1 \tag{8}$$

In order to adjust λ after each interaction (8) is used.

$$\lambda_{new} = \frac{\lambda_{old} + SIM}{n} \qquad where \quad SIM = 1 - \frac{\hat{\theta}_{rt} - outcome}{k-1} \ \ and\ n \geqslant 2 \tag{9}$$

In this formula, the similarity value (SIM) between our estimate and the outcome of the interaction is calculated first. If $\hat{\theta}_{rt}$ and $outcome$ are the same, then SIM will be equal to 1, otherwise will be less than 1 and greater than 0. The maximum value of their difference is $k-1$, and in this case SIM will be equal to 0. Based on the similarity between our estimation and the outcome of the interaction, the new value of λ will be calculated. In this formula, n is a natural number greater than or equal to 2 and is decided based on the application. For example, in risky applications, after a change in the behavior of the agent, the value of λ should be decreased sharply. Therefore greater value of n is needed.

Then we calculate the expected value for the Dirichlet distribution:

$$E(p(\theta_i)|\alpha) = \frac{\alpha(\theta_i)}{\sum_{i=1}^{k} \alpha(\theta_i)} \tag{10}$$

The reputation score can be expressed as a single value in some predefined interval. This can be done by assigning a point value $\hat{\theta}_{rt}$ to each rating level i (evenly distributed oint values in the range [0, 1] for k different rating levels), and computing the normalized weighted point estimate score. The point estimate reputation score is then computed as:

$$\hat{\theta}_{rt} = \sum_{i=1}^{k} \frac{i-1}{k-1} E(p(\theta_i)|\alpha) \tag{11}$$

5. Evaluation

To evaluate the performance and efficiency of the proposed algorithm a popular trust and reputation testbed for agent systems was used which is called ART. It is developed through a joint effort of Texas University, EMSE from France, ISTC from Italy and CWI from the Netherlands. The Agent Reputation and Trust (ART) Testbed [FKM+05] initiative has been launched with the goal of establishing a testbed for agent reputation- and trust related technologies. The ART Testbed serves two purposes: (1) as a competition forum in which researchers can compare their technologies against objective metrics, and (2) as a suite of tools with flexible parameters, allowing researchers to perform customizable, easily repeatable experiments. Annually, a workshop regarding ART's application is held in connection with the Autonomous Agent and Multi-agent Systems conference (AAMAS) which aims to bring together researchers who can contribute to a better understanding of trust and reputation in agent societies.

The reasons for this choice are: 1) as a versatile, universal experimentation site, the ART Testbed covers relevant trust research problems and unites researchers towards solutions via unified experimentation methods 2) Through objective, well-defined metrics, the testbed provides researchers with tools for comparing and validating their scientific models and the possibility of comparing a new model with previous models, 3) Standing on the shoulder of giants, and 4) reusability. We compared the proposed model with the Ntropi and have shown a considerable increased efficiency.

5.1 Metrics of Analysis in the testbed

In general, the most successful agent is selected as the appraiser with the highest bank account balance. In other words, the appraiser who is able to (1) estimate the value of its paintings most accurately and (2) purchase information most prudently, is deemed most successful. The Testbed also provides functionality to compute the average accuracy of the appraiser's final appraisals and the consistency of that accuracy, represented as its final appraisal error mean and standard deviation, respectively. In addition, the quantities of each type of message passed between appraisers are recorded [FKM+05].

5.2 Simulation Results

The Agent Skeleton is designed to allow researchers to implant within their appraiser agent-customized trust representations and algorithms while permitting standardized communication protocols with other entities. All appraiser agents participating in the ART Testbed are descendants of the same abstract class Agent. This class defines a set of abstract methods to be coded by the researcher to define the behavior of his/her appraiser agent, as well as a set of methods to facilitate the communication with other appraiser agents. The Agent class also provides methods for interacting with the Simulation Engine (for tasks such as verifying bank balances).

We used game rules similar to the rules in ART Testbed Competition 2007:

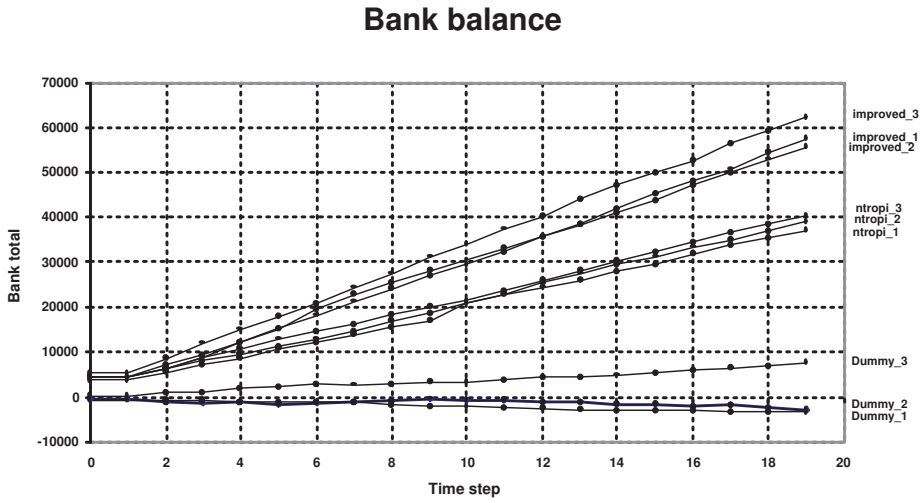

Figure 2. Results.

- Average-Clients-Per-Agent=20
- Client-Fee=100.0
- Opinion-Cost=10.0
- Reputation-Cost=0.1
- Timesteps-per-Session: 40

The game consisted of 12 agents from 4 different types:

- Simple: agents that do not use any model for trust related decisions.
- Ntropi: agents that use the Ntropi model for trust related decisions.
- Improved: agents that use the improved model for trust related decisions.
- Dummy: three dummy agents from the testbed itself which we used to have a more realistic environment.

Fig.2 shows the results. The horizontal axis in Fig.2 shows timesteps and the vertical axis is bank Total (the agent's bank balance). All of the improved agents have higher bank accounts during all time steps, and this shows their better performance.

6. Conclusion

The main contributions of this paper are that we have employed a second Bayesian algorithm in order to estimate the parameters for the priori trust, used a Dirichlet

distribution and introduced a new Hierarchical Bayesian-based reputation algorithm. In addition, we used the Maximum Likelihood Estimation algorithm to estimate the parameters of the Dirichlet distribution.

References

[ARH98] A. Abdul-Rahman and S. Hailes. A distributed trust model. *Proceedings of the 1997 workshop on New security paradigms*, pages 48–60, 1998.

[BLB04] S. Buchegger and J.Y. Le Boudec. A Robust Reputation System for Mobile Ad-hoc Networks. *Proceedings of P2PEcon, June*, 2004.

[CMH+02] I. Clarke, SG Miller, TW Hong, O. Sandberg, and B. Wiley. Protecting free expression online with Freenet. *Internet Computing, IEEE*, 6(1):40–49, 2002.

[CY01] R. Chen and W. Yeager. Poblano: A Distributed Trust Model for Peer-to-Peer Networks. *JXTA Security Project White Paper*, pages 1–26, 2001.

[DA04] Z. Despotovic and K. Aberer. Maximum Likelihood Estimation of Peers' Performance in P2P Networks. *2nd Workshop on Economics of Peer-to-Peer Systems, Cambridge, MA, USA, June*, 2004.

[FKM+05] K.K. Fullam, T.B. Klos, G. Muller, J. Sabater, A. Schlosser, Z. Topol, K.S. Barber, J.S. Rosenschein, L. Vercouter, and M. Voss. A specification of the Agent Reputation and Trust (ART) testbed: experimentation and competition for trust in agent societies. *Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems*, pages 512–518, 2005.

[Gen] Gnutella Protocol DevelopmentAn Overview of 0.6. Protocol Specification.

[JH07] A. Jøsang and J. Haller. Dirichlet reputation systems. *Proc of the Second International Conference on Availability, Reliability and Security (ARES)*, 2007.

[QL04] K. Quinn and A. Leddy. ebs/cd/a- 04:000737 uen. Technical report, Ericsson research, 2004.

[QOL+05] K. Quinn, D. OSullivan, D. Lewis, R. Brennan, and V.P. Wade. deepTrust Management Application for Discovery, Selection, and Composition of Trustworthy Services. *Proceedings of IDIP/IEEE 9th International Symposium on Integrated Network Management (IM 2005), Nice, France, May*, 2005.

[Ron89] G. Ronning. Maximum likelihood estimation of dirichlet distributions. *Journal of Statistical Computation and Simulation*, 32(4):215–221, 1989.

[Sni00] B.T. Sniffen. *Trust Economies in the Free Haven Project*. PhD thesis, MASSACHUSETTS INSTITUTE OF TECHNOLOGY, 2000.

PAPER G

Inferring Trust based on Similarity with TILLIT

Mozhgan Tavakolifard, Peter Herrmann, Svein J. Knapskog

IFIP Advances in Information and Communication Technology
Vol. 300, 2009

INFERRING TRUST BASED ON SIMILARITY WITH TILLIT

Mozhgan Tavakolifard,[1] Peter Herrmann,[2] Svein J. Knapskog,[1]

[1] *Centre for Quantifiable Quality of Service in Communication Systems*
Norwegian University of Science and Technology
{mozhgan, knapskog}@Q2S.ntnu.no

[2] *Department of Telematics,*
Norwegian University of Science and Technology
peter@item.ntnu.no

Abstract A network of people having established trust relations and a model for propagation of related trust scores are fundamental building blocks in many of todays most successful e-commerce and recommendation systems. However, the web of trust is often too sparse to predict trust values between non-familiar people with high accuracy. Trust inferences are transitive associations among users in the context of an underlying social network and may provide additional information to alleviate the consequences of the sparsity and possible cold-start problems. Such approaches are helpful, provided that a complete trust path exists between the two users. An alternative approach to the problem is advocated in this paper. Based on collaborative filtering one can exploit the like-mindedness resp. similarity of individuals to infer trust to yet unknown parties which increases the trust relations in the web. For instance, if one knows that with respect to a specific property, two parties are trusted alike by a large number of different trusters, one can assume that they are similar. Thus, if one has a certain degree of trust to the one party, one can safely assume a very similar trustworthiness of the other one. In an attempt to provide high quality recommendations and proper initial trust values even when no complete trust propagation path or user profile exists, we propose TILLIT — a model based on combination of trust inferences and user similarity. The similarity is derived from the structure of the trust graph and users' trust behavior as opposed to other collaborative-filtering based approaches which use ratings of items or user's profile. We describe an algorithm realizing the approach based on a combination of trust inferences and user similarity, and validate the algorithm using a real large-scale data-set.

1. Introduction

Many online communities are only successful if sufficient mutual trust between their members exists. Users want to know whom to trust and how much to trust in the competence and benevolence of other community members in a specific application domain. The process of building trust is hereby performed in two different ways. First, one can establish trust (or distrust) by gaining direct experience with another

party. Of course, every positive event increases the assumed trustworthiness of the trustee while every negative one reduces it. Second, one can gain trust based on recommendations of third parties. If, e.g., Alice has high trust in Bob's ability to assess the trustworthiness of other people, Bob has similar trust in Claire's recommendations, and Claire considers David trustable based on her personal experience with him, then Alice gains also trust in David even if she has no or very limited knowledge of him at all. This form of propagated trust is called trust transitivity.

Based on the two forms of trust, a so-called web of trust between community members is created which is often used in recommender systems helping users of e-commerce applications to get an idea about the trustworthiness of their mostly personally unknown cooperation partners. Unfortunately, however, these webs of trust are often too sparse to be helpful in practice since — at least in large online communities — a user has experience with only a very small fraction of the other community members. Thus, very often there will be no trust relation to an intended new partner of an e-commerce transaction at all [KLL+08].

As a model to increase the number of trust relations, we propose the method TILLIT[1] (Trust Inference Links based on Like-minded Interaction Transitions). It enables to derive trust not only from direct experience and by transitive propagation but also from the similarity between users and vice versa. In particular, two users are considered similar if they either built akin trust relations to other users or if they are trusted very similarly by others. This can be used to propagate already known trust to new trust relations encompassing people similar to those of the yet known relationships. Thus, the web of trust can be augmented significantly.

In our model, we measure similarity based on the existing web of trust in a community using an iterative fixed-point algorithm on node-pair graphs introduced later in this paper. As a method to describe the values of trust as well as its propagation we apply the TNA-SL model [JHP06] which is based on the Subjective Logic [Jøs01]. Our approach, however, would also work with other methods like [ARH00, GS02].

In comparison with other approaches based on similarity, our work has the following differences:

- It intends to alleviate the sparsity problem in the web of trust matrix itself instead of the matrix of users rating items in the system. Since users have usually few items rated in common, the classic recommender system techniques are often ineffective and are not able to compute a user similarity weight for many of the users. Instead, exploiting the web of trust, it is possible to propagate trust better and to infer additional trust information about other users.

- It calculates the similarity from the structure of the web of trust and trust relations (the trust graph structure and trust values) instead of user-item ratings.

- It proposes methods to convert trust values to similarity measures and vice versa based on the TNA-SL model.

[1] "Tillit" is the Norwegian word for trust.

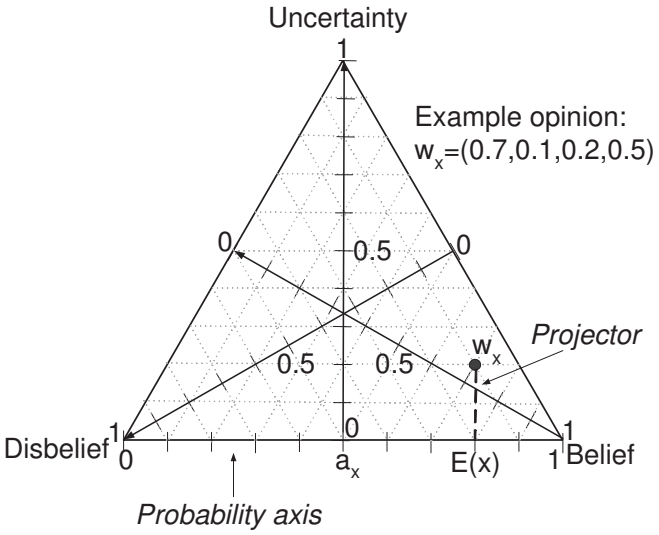

Figure 1. Opinion triangle with an example opinion [Jøs01].

We conducted experiments on a large real dataset showing how our proposed solution increases the coverage (number of trust relations that are predictable) while not reducing the accuracy (the error of predictions). This is especially true for users who have provided few ratings.

The rest of this paper is organized as follows: In section 2, we briefly explain the TNA-SL model as the background of our work. Our proposed model for trust inference is described in section 3. Next in section 4, we present the evaluation plan and results. Section 5 provides an overview of the related research. Finally, discussion and conclusion are given in section 6.

2. Trust Network Analysis with Subjective Logic

Our model is mainly based on TNA-SL [JHP06], a model for trust network analysis. TNA-SL uses the Subjective Logic [Jøs01] which enables to represent a specific belief calculus. There trust is expressed by a belief metric called opinion. An opinion is denoted by $\omega_B^A = (b,d,u,a)$ expressing the belief of a relying party A in the trustworthiness of another party B. The parameters b and d represent the belief resp. disbelief in B's trustworthiness while d expresses the uncertainty of A about to trust B or not. The three parameters are all probability values between 0 and 1 and fulfill the constraint $b+d+u=1$. The parameter a is called the base rate, and determines how uncertainty shall contribute to the opinion's probability expectation value which is calculated as $E(\omega_x^A) = b+au$. The opinion space can be mapped into the interior of an equal-sided triangle, where, the three parameters b, d, and u determine the position of the point in the triangle representing the opinion. Fig.1 illustrates an example where the opinion is $\omega_x = (0.7,0.1,0.2,0.5)$.

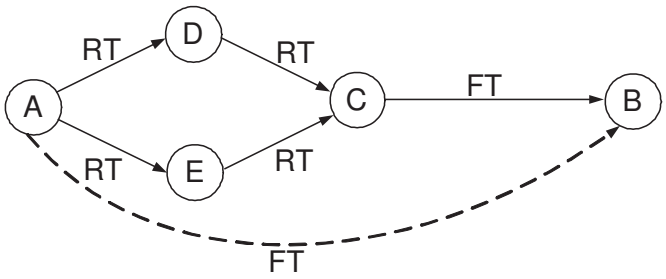

Figure 2. Referral trust transitivity and parallel combination of trust paths.

Based on TNA-SL, there are two different types of trust relations: *functional trust* (FT) and *referral trust* (RT). The former concerns A's direct trust in B performing a specific task; the latter concerns A's trust in B giving a recommendation about someone else doing a task or in other words is the trust in the ability to refer to a third party. As mentioned in the introduction, the simplest form of trust inference is trust transitivity which is widely discussed in literature [DKG+05, GKRT04, MBKM07, QHC07, YCB+02]. That is, if A trusts B who trusts C, then A will also trusts C. A valid transitive trust path requires that the last edge in the path represents functional trust and that all other edges in the path represents referral trust. Referral trust transitivity and parallel combination of trust paths are expressed as part of TNA-SL model (figure 2) [JHP06].

The discounting operator ($\otimes$) [Jøs02] is used to derive trust from transitive trust paths, and the consensus operator ($\oplus$) allows to combine parallel transitive trust paths. The trust network in figure 2 can then be expressed as

$$FT_B^A = ((RT_D^A \otimes RT_C^D) \oplus (RT_E^A \otimes RT_C^E)) \otimes FT_B^C$$

While we consider TNA-SL and the Subjective Logic as a suitable fundament for our similarity model, it can be, as already mentioned, adapted to all trust management models enabling to combine referral and functional trust (e.g., [ARH00, GS02]).

3. The Proposed Model

Our model for the estimation how much trust A can place in B considers not only direct experience and recommendations but also similarities between agents with respect of trusting other agents or being trusted by other parties. The two kinds of similarities between trusters resp. trustees can be gradually expressed by triples very similar to the first three operands of the opinion quadruples such that we can use the consensus operator of the subjective logic for the trust value computation.

3.1 Similar Trustees

If A has functional trust in C who is similar to B (they are *similar trustees*), then A can infer its functional trust to B ([DKG+05], see figure 3(a)). Two trustees are similar if they are both similarly trusted by other agents Z_1, Z_2, ..., Z_n (figure 3(b)).

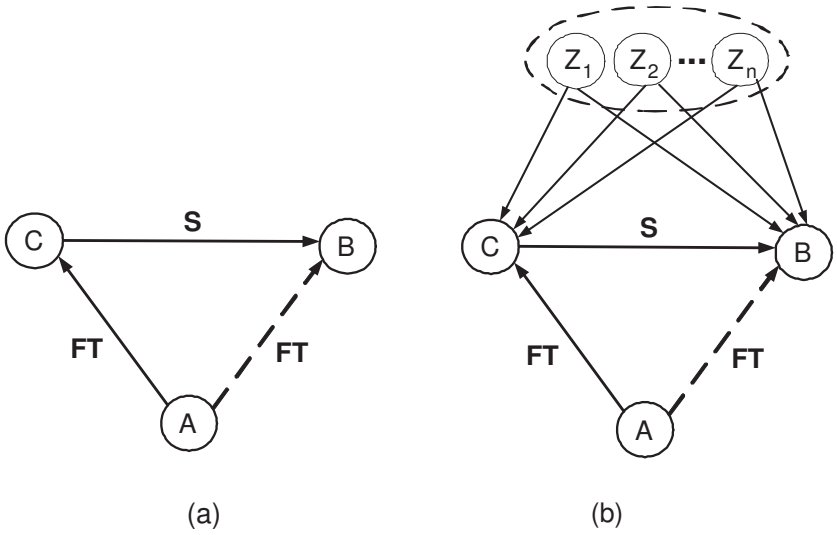

Figure 3. (a) Similar trustees (b) Similarly trusted.

This is an extension of TNA-SL in which it is not possible to infer any trust value of A towards B in a trust network.

Similarly to Jøsang's way to define opinions, we use triples to describe similarity which enables us to consider uncertainty. In particular, the degree of similarity depends on the number n of agents Z_1, Z_2, ..., Z_n used for the computation reflecting that we are more certain about the similarity of two parties if they are trusted by a significant large number of other agents in an akin way.

DEFINITION 1 *The similarity opinion S_B^C from C towards B is the triple*[2] *(*similarity*,* non-similarity*,* uncertainty*). If $C = B$, the similarity opinion is defined to be* $(1,0,0)$*. Otherwise, it is calculated based on the measure $sim_{te}(C,B)$ of similarity between the two trustees C and B which is introduced in subsection 3.3:*

$$S_B^C = (\frac{n \cdot sim_{te}(C,B)}{c+n}, \frac{n \cdot (1 - sim_{te}(C,B))}{c+n}, \frac{c}{c+n}) \quad (1)$$

c is a constant determining how fast uncertainty is replaced by assurance. As higher its value is, as more agents are needed to reduce the uncertainty value in favor of the similarity and non-similarity values. The similarity opinion fulfills the constraints that the sum of all three values is equal to 1.

Our similarity opinion is a special form of referral trust. It reflects that the akin trust evaluations of B and C by several other trusters are a kind of recommendation by these agents to A to treat B and C similarly. Thus, we see the discounting operator $\otimes$

[2]This metric is inferred from a metric for the trust value computation [JK98] by Jøsang and Knapskog.

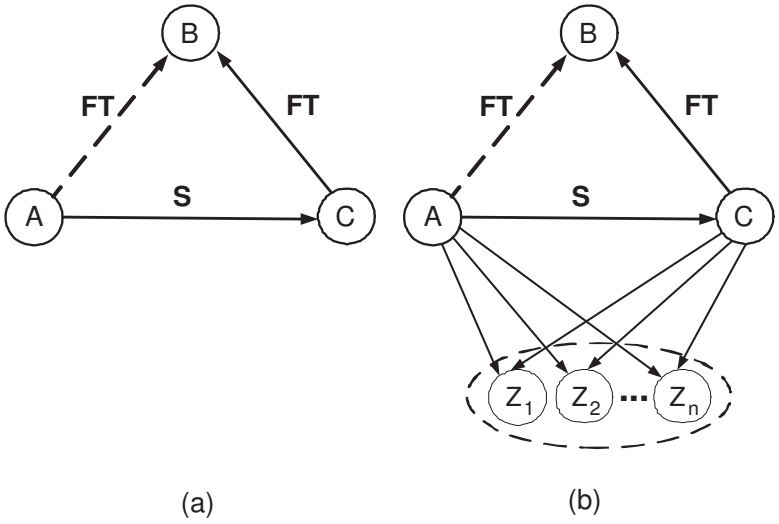

Figure 4. (a) Similar trusters (b) Similarly trusting.

as the correct mechanism to combine the similarity opinion between B and C with the functional trust of A in C in order to infer the functional trust of A in B:

$$FT_B^A = S_B^C \otimes FT_C^A \tag{2}$$

As higher the similarity between B and C is, as closer the trust of A to B will equal to that between A and C. As lower this similarity is, as more uncertain A will be about whether to trust B or not.

3.2 Similar Trusters

If C has functional trust to B and A is similar to C (they are *similar trusters*), then A can also infer functional trust towards B ([DKG+05], see figure 4(a)). We call C and A similar trusters if they have alike trust in several other agents Z_1, Z_2, ..., Z_n. In this case, if C has functional trust to a new agent B, then A can infer a functional trust to B (figure 4(b)). Again using TNA-SL alone, there is no way to infer a new trust value.

Like (1), the similarity opinion S_C^A from A to C is calculated using the measure of similarity $sim_{tr}(C,A)$ between trusters which is also introduced in subsection 3.3:

$$S_C^A = (\frac{n \cdot sim_{tr}(C,A)}{c+n}, \frac{n \cdot (1 - sim_{tr}(C,A))}{c+n}, \frac{c}{c+n}) \tag{3}$$

This similarity opinion is discounted by the functional trust FT_B^C from C to B to form the new trust value.

$$FT_B^A = S_C^A \otimes FT_B^C \tag{4}$$

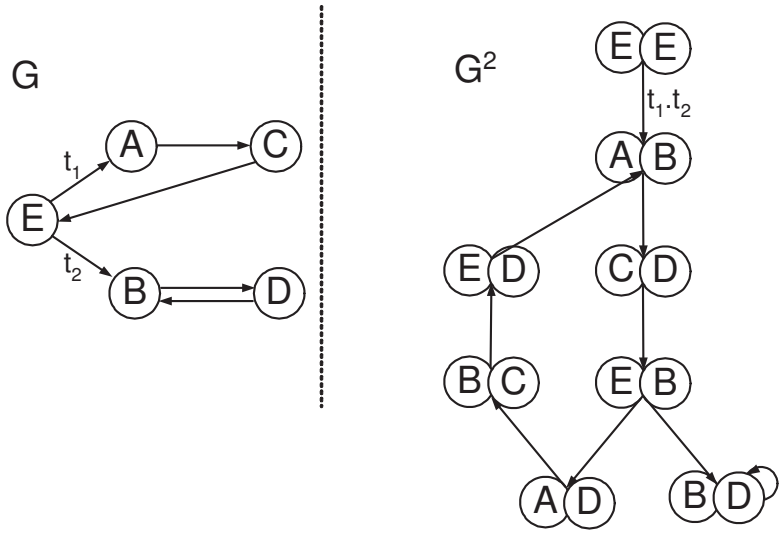

Figure 5. Similarity measurement.

3.3 Similarity Calculation

In order to measure similarities, we model trusters, trustees, and trust relationships as a graph with nodes representing trusters and trustees and edges representing trust relations. The intuition behind our algorithm is that, *similar* trustees are related to *similar* trusters. More precisely, trusters A and B are similar if they are related to trustees C and D, respectively, and C and D are themselves similar. The base case is that each node is similar to itself. If we call this graph G, then we can form a node-pair graph G^2 in which each node represents an ordered pair of nodes of G as depicted in figure 5. A node (A,B) of G^2 points to a node (C,D) if, in G, A points to C and B points to D. Similarity scores are symmetric, so for clarity we draw (A,B) and (B,A) as a single node A,B (with the union of their associated edges) [JW02].

We propose an iterative fixed-point algorithm on G^2 to compute similarity scores[3] for node-pairs in G^2. The similarity score for a node υ of G^2 gives a measure of similarity between the two nodes of G represented by υ. Scores can be thought of as flowing from a node to its neighbors. Each iteration propagates scores one step forward along the direction of the edges, until the system stabilizes (i.e., scores converge). Since nodes of G^2 represents pairs in G, similarity is propagated from pair to pair. Under this computation, two trustees are similar if they are trusted by similar trusters.

For each iteration k, iterative similarity functions $sim_{te,k}(*,*)$ for trustees and $sim_{tr,k}(*,*)$ for trusters are introduced. The iterative computation is started with

[3] An alternative approach to measure this similarity is to model an agent's mental structure as an ontology and using various methods proposed in our previous work [TKH08a, TKH08b]

$sim_{0,*}(*,*)$ defined as

$$sim_{0,*}(A,B) = \begin{cases} 1, & if A = B \\ 0, & if A \neq B \end{cases} \tag{5}$$

On the $(k+1)$-th iteration, $sim_{*,k+1}(*,*)$ is defined in special cases as

$$\begin{gathered} sim_{*,k+1}(A,B) = 1, \quad if A = B \\ sim_{te,k+1}(A,B) = 0, \quad if\ I(A) = \emptyset\ or\ I(B) = \emptyset \\ sim_{te,k+1}(A,B) = 0, \quad if\ O(A) = \emptyset\ or\ O(B) = \emptyset \end{gathered} \tag{6}$$

$I(A)$ is the set of in-neighbors of A while $O(A)$ specifies the set of A's out-neighbors. Individual in-neighbors are denoted as $I_i(A)$, for $1 \leq i \leq |I(A)|$, and individual out-neighbors are denoted as $O_i(A)$, for $1 \leq i \leq |O(A)|$. $sim_{te,k+1}(*,*)$ is computed from $sim_{tr,k}(*,*)$ in the general case as follows:

$$sim_{te,k+1}(A,B) = \frac{\sum_{i=1}^{n}\sum_{j=i}^{n} sim_{tr,k}(I_i(A),I_j(B)) \cdot (1 - \text{distance}(I_i(A),I_j(B),A,B))}{\sum_{i=1}^{n}\sum_{j=i}^{n} sim_{tr,k}(I_i(A),I_j(B))} \tag{7}$$

and $sim_{tr,k+1}(*,*)$ is computed from $sim_{te,k}(*,*)$ in the general case as:

$$sim_{tr,k+1}(A,B) = \frac{\sum_{i=1}^{n}\sum_{j=i}^{n} sim_{te,k}(O_i(A),O_j(B)) \cdot (1 - \text{distance}(A,B,O_i(A),O_j(B)))}{\sum_{i=1}^{n}\sum_{j=i}^{n} sim_{te,k}(O_i(A),O_j(B))} \tag{8}$$

Formulas (7) and (8) are alternately computed in iterations until the resulting similarity values sim_{tr} and sim_{te} converge. The corresponding algorithm is sketched as the procedure *CalculateSimilarity* in figure 4.1.

The *distance* function is used to compare trust relations. $distance(A,B,C,D)$ expresses the difference between the trust from A, B to C, D. It averages the Euclidean distances between the trust values of A and C resp. B and D on the opinion triangle (see figure 1):

$$\begin{aligned} \text{distance}(A,A,C,D) &= \sqrt{(b_{AC} + \tfrac{1}{2}u_{AC} - b_{AD} - \tfrac{1}{2}u_{AD})^2 + \tfrac{3}{4}(u_{AC} - u_{AD})^2} \\ \text{distance}(A,B,C,C) &= \sqrt{(b_{AC} + \tfrac{1}{2}u_{AC} - b_{BC} - \tfrac{1}{2}u_{BC})^2 + \tfrac{3}{4}(u_{AC} - u_{BC})^2} \\ \text{distance}(A,B,C,D) &= \begin{cases} \tfrac{1}{2}(\sqrt{(b_{AC} + \tfrac{1}{2}u_{AC} - b_{BD} - \tfrac{1}{2}u_{BD})^2 + \tfrac{3}{4}(u_{AC} - u_{BD})^2} \\ +\sqrt{(b_{AD} + \tfrac{1}{2}u_{AD} - b_{BC} - \tfrac{1}{2}u_{BC})^2 + \tfrac{3}{4}(u_{AD} - u_{BC})^2}) \end{cases} \end{aligned} \tag{9}$$

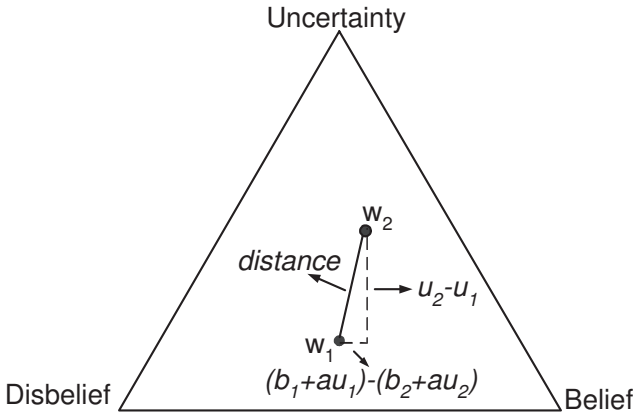

Figure 6. The distance between opinions.

For the sake of simplicity, all base rate values (a_{AD}, a_{AC}, a_{BD}, a_{BC}) are assumed to be $\frac{1}{2}$. The factor $\frac{3}{2}$ is used for the vertical axis to adapt the measures. Otherwise, the opinion triangle would be compressed and the distance between the points (0,1,0) and (0,0,1) would not be equal to one. Figure 6 illustrates the *distance* function graphically.

4. Evaluation

We chose a publicly available dataset taken from a real system known as Advogato [adv]. Advogato (`http://advogato.org`) is an online community site dedicated to free software development. On Advogato a user can certify another user as "Master", "Journeyer", "Apprentice" or "Observer", based on the perceived level of involvement in the free software community. The Advogato social network is an example of a real-world, directed, weighted, large social network. There are indeed other web communities using the same software powering Advogato.org and they also have reached similar trust levels and use the same certifications system, but we do not use them for our analysis in this paper, mainly because:

- Our model is based on user-user trust matrix and not the user-item rating matrix.
- They are much smaller than the Advogato dataset.

4.1 Dataset

Precise rules for giving out trust statements are specified on the Advogato site. *Masters* are supposed to be principal authors of an "important" free software project, excellent programmers who work full time on free software. *Journeyers* contribute significantly, but not necessarily full-time. *Apprentices* contribute in some way, but are still acquiring the skills needed to make more significant contributions. *Observers* are users without trust certification, and this is the default. It is also the level at which a user certifies another user to remove previously expressed trust certifications.

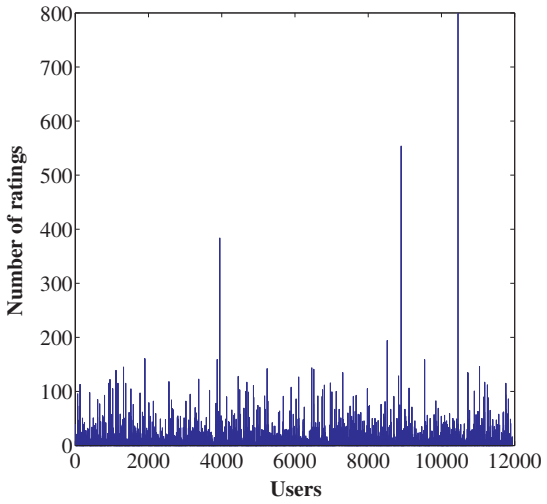

Figure 7. Users' rating activity.

The Advogato dataset is a directed, weighted graph with 11934 nodes and 57610 trust relations. There are 18053 Master judgments, 23091 for Journeyer, 10708 for Apprentice and 5758 for Observers. Figure 7 illustrates the allocation of ratings that correspond to each user. In our tests, we apply our model to 3 different datasets and the results are averaged. Each 3000 users built a trust graph of approximately 4000 relations.

For the purpose of this paper, we consider these certifications as trust statements. Trust statements are directed and not necessarily symmetric. By aggregating the trust statements expressed by all the members of the community it is possible to build the entire trust network. A trust network is hence a directed, weighted graph. Arbitrarily, we map the textual labels Observer, Apprentice, Journeyer and Master respectively to rating values 0, 1, 2, 3. which have to be yet converted to subjective logic opinions. In general, with n-level rating values (in our case $n = 3$) in which the number of ratings of level i is described by function $f(i)$, we can use the following conversion method in which c is a constant:

$$b = \frac{\sum_{i=1}^{n} i \cdot f(i)}{c + n \cdot \sum_{i=0}^{n} f(i)}, \quad d = \frac{\sum_{i=0}^{n-1} (n-i) \cdot f(i)}{c + n \cdot \sum_{i=0}^{n} f(i)}, \quad u = \frac{c}{c + n \cdot \sum_{i=0}^{n} f(i)} \tag{10}$$

In this formula, the highest rating value 3 is mapped to three positive valuations, while 2 corresponds to two positive valuations and a negative one, etc.

4.2 Plan

We use the leave-one-out technique [FH89] (a machine learning evaluation technique) to show the performance of our approach. Leave one out involves hiding one trust edge and then trying to predict it. The predicted trust edge is then compared with the real edge (using the distance function) and the difference is the prediction error. This procedure is repeated for all edges in the trust graph. The real and the predicted values are then compared in several ways: the coverage, which refers to the fraction of edges for which, after being hidden, the algorithm is able to produce a predicted edge, FCPE which is the fraction of correctly predicted edges, MAE (mean absolute error) which is average of the prediction error over all edges, and RMSE (root mean squared error) which is the root mean of the average of the squared prediction error. RMSE tends to emphasize large errors.

The evaluation can be described in pseudo-code as in algorithm 4.1. First, the similarity matrix is calculated by calling the procedure *CalculateSimilarity* from the main procedure. Since similarity is symmetric, the similarity of trustees is stored in the lower triangle of the similarity matrix and the similarity of trusters in the upper triangle. Next, for each edge in the real trust graph, an equivalent trust edge is calculated by calling procedure *PredictTrustEdge*. This procedure takes the real trust graph without that edge as an input. The predicted edges form the predicted trust graph. Finally, the real and predicted trust graph are compared according to the four metrics (coverage, FCPE, MAE, and RSME) by calling procedure *DoEvaluation*.

4.3 Results Summary

Figure 8 depicts the similarity measures among the first 150 users. For each two users, their similarity as trustees is in the lower triangle of the similarity matrix and their similarity as trusters is in the upper triangle of the similarity matrix.

In table 1 we present the final results of the evaluation. We start by commenting the column "coverage". The coverage becomes an important issue on a very sparse dataset that contains a large portion of cold start users since many trust values become hardly predictable [MA07]. Our baseline is a method called "Random" which randomly generates trust edges. Results ($coverage \approx 0.6$) indicate that our model is able to predicate approximately one edge from each two existing edges. The second important result is the fraction of correctly predicted edges (FCPE) which is 0.8. It shows that from each 10 predicted edge 8 edges are predicted correctly. Further, prediction errors (MAE and RMSE) computed are small in comparison with the Random method ($MAE \approx 0.14$ & $RMSE \approx 0.18$).

Figure 9 shows the sparsity of the trust graph before and after prediction for the first dataset. The sparseness has been decreased significantly. All-in-all, the results of the evaluation lead to the expectation that the method TILLIT will increase the coverage of trust relationships significantly, and that the accuracy of the predicted additional will be fairly high as well.

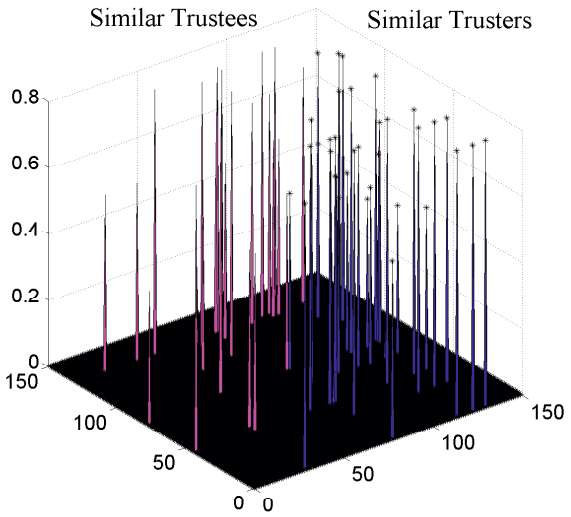

Figure 8. Similarity measures among the first 150 users

Table 1. Final evaluation results

Metric	Dataset1	Dataset2	Dataset3	Average	Random
Coverage	0.5783	0.5678	0.6520	0.5994	1
FCPE	0.8169	0.8299	0.8227	0.8232	0.3068
MAE	0.1389	0.1427	0.1409	0.1408	0.4570
RMSE	0.1823	0.1828	0.1864	0.1838	0.5036

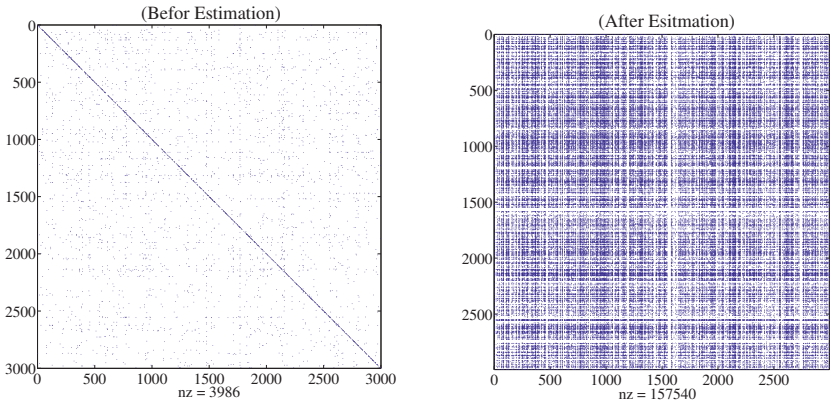

Figure 9. Sparsity of the trust graph before and after prediction for the first dataset

5. Related Research

Most popular approaches proposed to deal with the sparsity problem include dimensionality reduction of the user-item matrix, application of associative retrieval techniques in the bipartite graph of items and users, item-based similarity instead of user-based similarity, and content-boosted collaborative filtering (see [PPK05]). The dimensionality reduction approach addresses the sparsity problem by removing unrepresentative or insignificant users or items so as to condense the user-item matrix. We briefly explain those which are based on similarity measurement and thus more closely resemble our work, see Table 5. These approaches can be categorized in two groups: rating-based similarity and profile-based similarity.

Paper/Technique	**Web of trust**	**User-item ratings**	**rating-based similarity**	**profile-based similarity**
Papagelis et al. [PPK05]		X	X	
Massa and Avesani [MA04]	X		X	
Massa and Bhattacharjee [MB04]	X		X	
Massa and Avesani [MA07]	X			
Avesani et al., [AMT04]	X		X	
Avesani et al., [AMT05]	X		X	
Weng et al., [WMG06]	X			
Lathia et al., [LHC08]	X			
Gal-Oz et al., [GOGH08]	X			
O'Donovan and Smyth [OS05]		X	X	X
Ziegler and Golbeck [ZG07]		X		X
Ziegler and Lausen [ZL04]		X		X
Golbeck [Gol06]		X		X
Golbeck and Hendler [GH04]		X		X
Golbeck [Gol05]	X			
Bedi and Kaur [BK06]	X			
Bedi et al., [BKM07]	X			
Hwang and Chen [HC07]		X		
Kitisin and Neuman [KN06]	X			
Fu-guo and Sheng-hua [FgSh07]		X		X
Peng and Seng-cho [PSc09]	X			X
Victor et al., [VDCCT08]	X			
Victor et al., [VCDCPdS09]	X			

Recently, several researches have suggested that the incorporation of a notion of trust into the standard CF model can effectively solve the sparsity problem and thus

provide better recommendations. A user can build his personalized web of trust by specifying those friends or users he trusts. The trust web can be constructed through the explicit trust ratings provided by users [HC07].

In [PPK05], authors explain how similarity can benefit from special characteristics of trust such as the ability to propagate along chains of trusted users; in this way similarity can support transitivity. They develop a model to establish trust between users by exploiting the transitive nature of trust. In their model they use ordinary measures of similarity taken from collaborative filtering to form the potential trust between the users which would be propagated in a similar way to the word-of-mouth scheme through a trust graph. Finally, by transforming the value back into similarity measure terms, it could be made appropriate for use in collaborative filtering algorithms. More specifically, for each pair of users they first calculate how similar they are, applying Pearsons correlation coefficient formula over the user-item ratings, and then they calculate the indirect trust between them. Next, this trust value is converted to a similarity metric using their formula. However, their model simply adopts similarity as trustworthiness. Hence, it still possesses the limitations of similarity-based CF as discussed. The main contribution of this work is that a trust metric has been designed, which helps a user to quantify the degrees of trust it should place on others.

Massa et al. present in [MB04, MA04] evidence that, by incorporating trust, recommender systems can be more effective than systems based on traditional techniques like collaborative filtering. In [MA04], the authors analyze the potential contribution of Trust Metrics in increasing the performances of Recommender Systems and proposed an architecture for trust-aware Recommender Systems. In this paper, it is proposed that a peer can establish trust on other peers through explicit trust statements and trust propagation. A trust model is built directly from users' direct feedbacks. This trust model is incorporated into the recommendation process for recommending various items (such as books, movie, music, software etc.) to on-line users. Users can express their personal web of trust by identifying those reviewers whose reviews and ratings are consistently found to be valuable. Massa et al. argue that it is possible to predict trust in unknown users by propagating trust even there were no direct connection between them. However, it is not clear how a user quantify the degrees of trust when making trust statements. The authors show how the similarity measure, on average, is computable only against a very small portion of the user base and is, in most cases, a noisy and unreliable value because computed on few items rated in common by two users. Instead, trust-aware techniques can produce a trust score for a very high number of other users; the trust score of a user estimates the relevance of that users' preferences. In this paper, similarity is measured using Pearsons correlation coefficient on user-item ratings.

They also show, in their subsequent experiment [MA04], that the incorporation of trust metric and similarity metric can increase the coverage of recommender systems while maintaining the recommendation accuracy. Due to the limitation on trust value representation, in their experiments, the webs of trust are built on binary relationships among users and the propagating trusts are computed simply based on the distances between them.

The work of [MB04] builds a trust model directly from trust data provided by users as part of the popular epinions.com service. A big limitation of the work in [MA04] and [MB04] is that they require some explicit trust ratings in order to infer further trust rating.

Avesain et al. in [AMT04, AMT05] apply the trust model into the ski mountaineering domain. They present a community-based website in which users can share their opinions about the snow conditions of different ski routes and also express their trust on others opinions. The trust score of a user depends on the trust statements of other users on him/her and their trust scores. However, the trust model requires the direct feedback of users and the effectiveness of the trust model on the skiing community has not been validated.

In [WMG06] have proposed that peers predict the new items' ratings based on the recommendations of the peers that are trusted directly or indirectly. A trust metrics has been designed to help peers to determine the degrees of trust should be placed on others. The design of trust metrics also stimulates a novel method to make prediction, which is featured by the recommendation adjustment and pseudo-recommendation. It has been shown by the experimental results that the trust metrics and corresponding prediction making approach do improve the performance of traditional similarity-based CF in terms of coverage, prediction accuracy and robustness.

A number of techniques for performing collaborative filtering from the point of view of a trust-management problem are outlined in [LHC08]. In this work authors propose a variation of k-nearest neighbor collaborative filtering algorithm for trusted k-nearest recommenders. This algorithm allows users to learn who and how much to trust one another by evaluating the utility of the rating information they receive. They mainly address the problem of learning how much to trust rating information that is received from other users in a recommender system.

A model for computing trust-based reputation for communities of strangers is proposed in [GOGH08]. The model uses the concept of knots, which are sets of members having high levels of trust in each other. Different knots typically represent different view points and preferences. The assumption underlying this knot-aware reputation model is that use of relatively small, but carefully selected, subsets of the overall community's reputation data yields better results than those represented by the full dataset.

In [OS05], O'Donovan and Smyth argue that profile similarity is just one of a number of possible factors that might be used to influence recommendation and prediction, and the reliability of a partner profile to deliver accurate recommendations in the past is another important factor, if a profile has made lots of accurate recommendation predictions in the past it can be viewed as more trustworthy than another profile that has made many poor predictions. They claim that the reliability of a user profile to deliver accurate recommendation in the past is an important factor for influencing recommendation and prediction. A user is viewed as more trustworthy if he has made more accurate predictions in the past than other users. The trust metrics are calculated at both the Item and Profile levels. trust values are calculated both the Item and Profile levels. Item Level trust is a representation for a producer's trustworthiness

with respect to the recommendation of a specific item. Profile Level trust is a less fine-grained metric, representing a recommendation producers trust as a whole, without respect to one specific item. For example, we might wish to refer to John's overall trustworthiness based on a series of different past recommendations. This score is simply an average over the Item Level trust scores for every item in the users profile. Essentially these metrics summarize the relative number of correct recommendations that a given user has made, according to a predefined error bound. They propose to modify the way that recommendation partners are generally selected or weighted during the recommendation process. They argue that profile similarity on its own may not be sufficient, that other factors might also have an important role to play. Specifically they introduce the notion of trust in reference to the degree to which one might trust a specific profile when it comes to make a specific rating prediction. They develop two different trust models, one that operates at level of the profile and one at level of the items within a profile. In both of these models trust is estimated by monitoring the accuracy of a profile at making predictions over an extended period of time. Trust then is the percentage of correct predictions that a profile has made in general (profile-level trust) or with respect to a particular item (item-level trust). They describe how this trust information can be incorporated into the recommendation process and demonstrate that it has a positive impact on recommendation quality. However, this system only uses a global trust metric and provides neither any personalization nor trust propagation.

Ziegler and Golbeck in [ZG07] experimentally prove that there exists a significant correlation between the trust expressed by the users and their profile similarity based on the recommendations they made in the system. This correlation is further studied as survey-based experiments in [Gol06]. Ziegler and Lausen in [ZL04] mention that in order to provide meaningful results for recommender system applications, they expect notions of trust to clearly reflect user similarity. In this work, they provide empirical results obtained from one real, operational community and verify latter hypothesis for the domain of book recommendations.

Golbeck et al. in [GH04] describe an E mail filtering system based on trust ratings. The predicted trust of a user is given by a weighted average of her neighbors trust ratings. They have shown that the weighted average metric can provide better results than other metrics. However they still need the explicit trust ratings from users and do not use any mail ratings information.

Golbeck in [Gol05] present FilmTrust, a website that uses trust in Semantic Web-based social networks, to create predictive movie recommendations. She show how these recommendations are more accurate than other techniques in certain cases, and discuss this as a mechanism of Semantic Web interaction. Within the FilmTrust website, trust in social networks has been used to personalized the user experience. Trust took on the role of a recommender system forming the core of an algorithm to create predictive rating recommendations for movies. The accuracy of the trust-based predicted ratings in this system is significantly better than the accuracy of a simple average of the ratings assigned to a movie and also the recommended ratings from a Person-correlation based recommender system.

In [BK06] a model that incorporates the social recommendation process is proposed. The trustworthy peers of the user become the recommender agents and suggest movies to the user according to the tastes of the user. The agents in our system also learn from their experience in dealing with the trustworthy peers and update the degree of trust on them. In the proposed system, they have tried to merge the advantages of the mechanical recommender system with the more humane recommendation process to make their recommendations trustworthy and useful for the user.

In [HC07] an improved mechanism to the standard CF techniques by incorporating trust into CF recommendation process is presented. They derive the trust score directly from the user rating data based on users' prediction accuracy in the past and exploit the trust propagation in the trust web. We investigate the effects of both the local trust metric and the global trust metric in the standard CF recommendation. The global metric has shown to have an advantage over other approaches in prediction coverage. The local metrics provide more accurate recommendations than those provided by standard CF technique. The overall performance of their trust-based recommender system is presented and favorably compared to other approaches. Experimental results verify that the incorporation of trust into CF process can indeed improve the prediction accuracy while maintain satisfactory prediction coverage.

[KN06] propose an approach to include the social etc) offer a process for collecting and distributing reputation/trust factors e.g. user's past behaviors and reputation together as an information rating from a user to another user. Some systems element of trust that can be incorporated into the current which allows anonymity collect feedback on their users' past recommender system framework and show their experiments in behaviors.

In [FgSh07] authors argue that items belonging to different topics need different trustworthy users to make recommendation, so topic-level trust will be more effective than profile-level trust in incorporating into the recommendation process. Based on this idea, they design a topic-level trust model which helps a user to quantify the trustworthy degree on a specific topic, and propose a new recommender algorithm by incorporating the new model into the mechanics of a standard collaborative filtering recommender system. Their proposed algorithm combines topic trust with profile similarity. The results from experiments based on Movielens dataset show that the new method can improve the recommendation accuracy of recommender systems.

[BKM07] proposes the design of a recommender system that uses knowledge stored in the form of ontologies. The interactions amongst the peer agents for generating recommendations are based on the trust network that exists between them. Recommendations about a product given by peer agents are in the form of Intuitionistic Fuzzy Sets specified using degree of membership, non membership and uncertainty. The presented design uses ontologies, a knowledge representation technique, instead of databases for creating annotated content for Semantic Web. Seeing the potential and popularity of ontologies among researchers, they believe that ontologies will be build and maintained in numerous knowledge domains for the Semantic Web and future applications. The presented recommender system uses temporal ontologies that absorb the effect of changes in the ontologies due to the dynamic nature of domains,

in addition to the benefits of ontologies. A case study of tourism recommender system is chosen to generate the recommendations for the selection of destination, travel agents and the flight schedule. A comparison of the generated recommendations with the manual recommendations by peers establishes the validity of the presented recommender system.

[PSc09] is motivated by the need to provide recommendations about blog articles, so that bloggers/readers can find desired articles easily. Accordingly, this study proposes to exploit the trust relationships between bloggers and readers via explicit trust ratings to generate recommendations in a reliable and satisfactory way. Furthermore, rather than only using a single trust rating, this work presents a multi-faceted model that considers trust by dividing a general trust rating into multiple trust ratings for different types of blog articles, thus enabling trust relationships to be evaluated in a fine-grained manner. To help ease information overload in the blogosphere, this work proposes a trust-enhanced collaborative filtering approach that integrates multi-faceted trust based on article type and user similarity. An online blog article recommender system, called iTrustU, is also designed to evaluate the effectiveness of the proposed approach in terms of accuracy and quality of recommendations. Results of a 45-day online experiment with 179 participants from the Internet demonstrate that the proposed integrated approach yields a significantly higher accuracy than traditional approaches, especially for cold-start users. Analysis results indicate that trust and similarity among bloggers/readers have a significantly positive correlation in the blogosphere. Effective recommender systems can be achieved by exploiting trust relationships in a trust network. The proposed approach is applicable not only to the blogosphere, but also to online social communities when trust relationships already exist between users.

[VDCCT08] examines the problem of cold-start users in recommender systems and propose to connect the newcomer to an underlying trust network among the users of the recommender system which alleviates the so-called cold start problem. In this paper, they study the effect of guiding the new user through the connection process, and in particular the influence this has on the amount of generated recommendations. Experiments on a dataset from Epinions.com support the claim that it is more beneficial for a newcomer to connect to an identified key figure instead of to a random user.

In [VCDCPdS09] the authors advocate the use of a trust model in which trust scores are (trust,distrust)-couples, drawn from a bilattice that preserves valuable trust provenance information including gradual trust, distrust, ignorance, and inconsistency. They pay particular attention to deriving trust information through a trusted third party, which becomes especially challenging when also distrust is involved. In our work we provide an alternative approach to deal with the sparsity problem.

In our work we provide an alternative approach to deal with the sparsity problem. We measure similarity based on the users' trust relationships, i.e. trust graph structure and trust values (in contrast to the other approaches which have used user-item ratings or profile similarity), and propose novel formulas to convert it to subjective logic opinions. The consideration of these similarities leads to extra information accessible for trust inferences.

Algorithm 5.1: EVALUATION($users, trust\_graph$)

procedure CALCULATESIMILARITY($users, trust\_graph$)
 repeat
  for each $i, j \in users$
   do if $i = j$
   then $similarity\_matrix[i, j] \leftarrow (1, 0, 0)$
   else
    if $i < j$
     then $neighbors \leftarrow$ common in-neighbors of i and j
    comment: similarity of trustees
     else $neighbors \leftarrow$ common out-neighbors of i and j
    comment: similarity of trusters
    if $number\_of\_neighbors == 0$
     then $sim \leftarrow 0$
     else $sim \leftarrow$ GETSIMILARITY($neighbors$)
    comment: According to (7) and (8)
    $similarity\_matrix[i, j] \leftarrow$ GETOPINION($sim, number\_of\_neighbors$)
    comment: According to (1)
 until *converge*
 return ($similarity\_matrix$)

procedure PREDICTTRUSTEDGE($(i, j), trust\_graph$)
 $opinion \leftarrow (0, 0, 1)$
 for each $k \in users - \{i, j\}$
  do
   $similarity\_trustee(k, j) \leftarrow similarity\_matrix[min(k, j), max(k, j)]$
   $similarity\_truster(i, k) \leftarrow similarity\_matrix[max(i, k), min(i, k)]$
   $predicted\_opinion\_te \leftarrow trust\_opinion(i, k) \otimes similarity\_trustee(k, j)$
   $predicted\_opinion\_tr \leftarrow trust\_opinion(k, j) \otimes similarity\_truster(i, k)$
   $opinion \leftarrow (opinion \oplus predicted\_opinion\_te \oplus predicted\_opinion\_tr)$
 return ($opinion$)

procedure DOEVALUATION($trust\_graph, predicted\_trust\_graph$)
 $coverage \leftarrow$ number of predicted edges in predicted_trust_graph
 $fcpe \leftarrow$ fraction of correctly predicted edges
 $mae \leftarrow$ mean absolute error of predicted values
 $rmse \leftarrow$ root mean squared error of predicted values
 output ($coverage, fcpe, mae, rmse$)

main
 global $similarity\_matrix \leftarrow$ CALCULATESIMILARITY($users, trust\_graph$)
 for each $edge \in trust\_graph$
  do
   $predicted\_edge \leftarrow$ PREDICTTRUSTEDGE($edge, trust\_graph - edge$)
   $predicted\_trust\_graph \leftarrow predicted\_trust\_graph \cup predicted\_edge$
 DOEVALUATION($trust\_graph, predicted\_trust\_graph$)

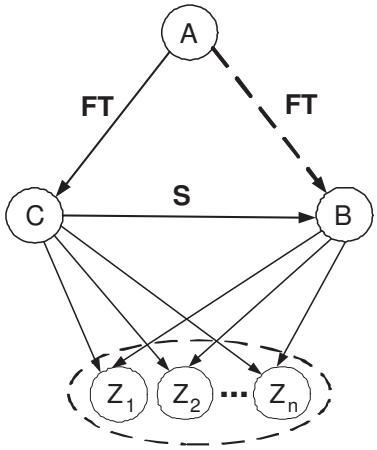

Figure 10. Coupling: a trust propagation method.

6. Discussion and Conclusion

In order to overcome sparseness of the web of trust, we consider users' similarity as a factor to derive trust connectivity and trust values. The main idea is that we account two persons similar if either a fair number of others have akin trust in them or if they themselves trust several other people alike. In the first case, every person who has trust in one of them can infer similar trust to the other one, at least as an estimated starting value. In the second case, a person may infer the trust value of a third party from other trusters similar to her.

We consider a similarity-based recommendation system for singers and songs as a good application example for our model. Normally, in systems like iTunes only the most popular songs or other songs of artists, of whom one already has bought songs, are advertised without any guarantee that one likes these songs as well. Using our approach, it is possible to find other customers who have an akin taste about music as the customer Alice reading the advertisements. Songs rated positively by these customers but not bought yet by Alice can be advertised to her since she will like them probably as well. This will make Alice more receptive to the advertisements.

In the future, we aim to evaluate the accuracy of a whole recommender system that employs our proposed model. Furthermore, we assess the possibility of modeling some of other trust propagation methods using our approach. An example is transposition resp. reciprocity [GKRT04] assuming that A's trust in B causes B to develop also some level of trust towards A. Another propagation method is Coupling, in which A's trust in C propagates to B because C and B trust people in common [GKRT04]. This propagation rule is depicted in figure 10. According to this rule we can use the similarity between trusters to propagate the trust in one trustee to another.

Moreover, one can use similarity in a complete different way. Trust is very specific and nobody trusting Bob as a good car mechanic will automatically trust him also in undertaking heart surgeries. But probably, he will be capable in repairing motorcycles.

Thus, there is a large similarity between the domains of repairing cars and motorcycles but a very low one between both of these and medical surgery. We think to use trust relations in one domain to infer ones in similar domains and consider ontologies describing the degrees of similarity between the domains as a useful means. All-in-all, we are convinced, that the various forms of similarity are good vehicles to tackle the major problem of too sparse webs of trust in online communities.

References

[adv] `http://www.trustlet.org/wiki/Advogato_dataset`.

[AMT04] P. Avesani, P. Massa, and R. Tiella. Moleskiing: trust-aware decentralized recommender system. In *1st Workshop on Friend of a Friend, Social Networking and the Semantic Web. Galway, Ireland*, 2004.

[AMT05] P. Avesani, P. Massa, and R. Tiella. A trust-enhanced recommender system application: Moleskiing. In *Proceedings of the 2005 ACM symposium on Applied computing*, pages 1589–1593. ACM New York, NY, USA, 2005.

[ARH00] A. Abdul-Rahman and S. Hailes. Supporting trust in virtual communities. In *Proceedings of the 33rd Hawaii International Conference, Volume 6*, Maui, Hawaii, 2000. IEEE Computer Society Press.

[BK06] P. Bedi and H. Kaur. Trust based personalized recommender system. *INFOCOM Journal of Computer Science*, 5(1):19–26, 2006.

[BKM07] P. Bedi, H. Kaur, and S. Marwaha. Trust based recommender system for the semantic web. *Proc. of the IJCAI07*, pages 2677–2682, 2007.

[DKG+05] L. Ding, P. Kolari, S. Ganjugunte, T. Finin, and A. Joshi. Modeling and Evaluating Trust Network Inference. Technical report, MARYLAND UNIV BALTIMORE DEPT OF COMPUTER SCIENCE AND ELECTRICAL ENGINEERING, 2005.

[FgSh07] Z. Fu-guo and X. Sheng-hua. Topic-level Trust in Recommender Systems. In *Management Science and Engineering, 2007. ICMSE 2007. International Conference on*, pages 156–161, 2007.

[FH89] K. Fukunaga and D.M. Hummels. Leave-one-out procedures for nonparametric error estimates. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(4):421–423, 1989.

[GH04] J. Golbeck and J. Hendler. Reputation network analysis for email filtering. In *Proceedings of the First Conference on Email and Anti-Spam*, volume 44, pages 54–58, 2004.

[GKRT04] R. Guha, R. Kumar, P. Raghavan, and A. Tomkins. Propagation of trust and distrust. In *Proceedings of the 13th international conference on World Wide Web*, pages 403–412. ACM Press New York, NY, USA, 2004.

[GOGH08] N. Gal-Oz, E. Gudes, and D. Hendler. A Robust and Knot-Aware Trust-Based Reputation Model. In *Proceedings of IFIPTM 2008 - Joint iTrust and PST Conferences on Privacy, Trust Management and Security*, pages 167–182. Springer, 2008.

[Gol05] J. Golbeck. Semantic Web Interaction through Trust Network Recommender Systems. In *Proc. of the ISWC05 Workshop on End User Semantic Web Interaction*, page 2005, 2005.

[Gol06] J. Golbeck. Trust and nuanced profile similarity in online social networks. *Journal of Artificial Intelligence Research*, 2006.

[GS02] T. Grandison and M. Sloman. Specifying and analysing trust for internet applications. In *Proceedings of the 2nd IFIP Conference on E-Commerce, E-Business & E-Government (I3E)*, pages 145–157, Lisbon, 2002. Kluwer Academic Publisher.

[HC07] C. Hwang and Y. Chen. Using Trust in Collaborative Filtering Recommendation. *Lecture Notes in Computer Science*, 4570:1052, 2007.

[JHP06] A. Jøsang, R. Hayward, and S. Pope. Trust network analysis with subjective logic. In *Proceedings of the 29th Australasian Computer Science Conference-Volume 48*, pages 85–94. Australian Computer Society, 2006.

[JK98] A. Jøsang and S. J. Knapskog. A metric for trusted systems. In *Proceedings of the 21st National Security Conference*. NSA, 1998.

[Jøs01] A. Jøsang. A Logic for Uncertain Probabilities. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 9(3):279–311, 2001.

[Jøs02] A. Jøsang. The consensus operator for combining beliefs. *Artificial Intelligence*, 141(1-2):157–170, 2002.

[JW02] G. Jeh and J. Widom. SimRank: a measure of structural-context similarity. In *Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 538–543. ACM Press New York, NY, USA, 2002.

[KLL+08] Y.A. Kim, M.T. Le, H.W. Lauw, E.P. Lim, H. Liu, and J. Srivastava. Building a web of trust without explicit trust ratings. In *Data Engineering Workshop, 2008. ICDEW 2008. IEEE 24th International Conference on*, pages 531–536, 2008.

[KN06] S. Kitisin and C. Neuman. Reputation-based Trust-Aware Recommender System. *Securecomm and Workshops, 2006*, pages 1–7, 2006.

[LHC08] N. Lathia, S. Hailes, and L. Capra. Trust-Based Collaborative Filtering. In *Proceedings of IFIPTM 2008 - Joint iTrust and PST Conferences on Privacy, Trust Management and Security*, pages 119–134. Springer, 2008.

[MA04] P. Massa and P. Avesani. Trust-Aware Collaborative Filtering for Recommender Systems. *LECTURE NOTES IN COMPUTER SCIENCE*, pages 492–508, 2004.

[MA07] P. Massa and P. Avesani. Trust-aware recommender systems. In *Proceedings of the 2007 ACM conference on Recommender systems*, pages 17–24. ACM Press New York, NY, USA, 2007.

[MB04] P. Massa and B. Bhattacharjee. Using Trust in Recommender Systems: An Experimental Analysis. In *Trust Management: Second International Conference, ITrust 2004, Oxford, UK, March 29-April 1, 2004: Proceedings*. Springer, 2004.

[MBKM07] R. Morselli, B. Bhattacharjee, J. Katz, and M. Marsh. Exploiting approximate transitivity of trust. In *Broadband Communications, Networks and Systems, 2007. BROADNETS 2007. Fourth International Conference on*, pages 515–524, 2007.

[OS05] J. O'Donovan and B. Smyth. Trust in recommender systems. In *Proceedings of the 10th international conference on Intelligent user interfaces*, pages 167–174. ACM New York, NY, USA, 2005.

[PPK05] M. Papagelis, D. Plexousakis, and T. Kutsuras. Alleviating the sparsity problem of collaborative filtering using trust inferences. In *Proceedings of iTrust*, pages 224–239. Springer, 2005.

[PSc09] T.C. Peng and T.C. Seng-cho. iTrustU: a blog recommender system based on multi-faceted trust and collaborative filtering. In *Proceedings of the 2009 ACM symposium on Applied Computing*, pages 1278–1285. ACM New York, NY, USA, 2009.

[QHC07] D. Quercia, S. Hailes, and L. Capra. Lightweight Distributed Trust Propagation. In *Data Mining, 2007. ICDM 2007. Seventh IEEE International Conference on*, pages 282–291, 2007.

[TKH08a] M. Tavakolifard, S. Knapskog, and P. Herrmann. Cross-Situation Trust Reasoning. In *Proceedings of The Workshop on Web Personalization, Reputation and Recommender Systems (WPRRS08)*. IEEE Computer Society Press, 2008.

[TKH08b] M. Tavakolifard, S. Knapskog, and P. Herrmann. Trust Transferability Among Similar Contexts. In *Proceedings of The 4th ACM International Workshop on QoS and Security for Wireless and Mobile Networks (Q2SWinet 2008)*. ACM, 2008.

[VCDCPdS09] P. Victor, C. Cornelis, M. De Cock, and P. Pinheiro da Silva. Gradual trust and distrust in recommender systems. *Fuzzy Sets and Systems*, 160(10):1367–1382, 2009.

[VDCCT08] P. Victor, M. De Cock, C. Cornelis, and A. Teredesai. Getting cold start users connected in a recommender systems trust network. *Computational Intelligence in Decision and Control*, 1:877–882, 2008.

[WMG06] J. Weng, C. Miao, and A. Goh. Improving collaborative filtering with trust-based metrics. In *Proceedings of the 2006 ACM symposium on Applied computing*, pages 1860–1864. ACM New York, NY, USA, 2006.

[YCB+02] Y. Yang, ACT Canberra, L. Brown, S. Wales, ACT ADFA, E. Lewis, and V.A. Melbourne. W3 Trust Model: Evaluating Trust and Transitivity of Trust of Online Services. In *International Conference on Internet Computing*, pages 354–362, 2002.

[ZG07] C.N. Ziegler and J. Golbeck. Investigating interactions of trust and interest similarity. *Decision Support Systems*, 43(2):460–475, 2007.

[ZL04] C.N. Ziegler and G. Lausen. Analyzing correlation between trust and user similarity in online communities. *Lecture notes in computer science*, pages 251–265, 2004.

PAPER H

Analogical Trust Reasoning

Mozhgan Tavakolifard, Peter Herrmann, Pinar Ozturk

IFIP Advances in Information and Communication Technology
Vol. 300, 2009

ANALOGICAL TRUST REASONING

Mozhgan Tavakolifard,[1] Peter Herrmann,[2] Pinar Ozturk,[3]

[1] *Centre for Quantifiable Quality of Service in Communication Systems*
Norwegian University of Science and Technology
mozhgan@Q2S.ntnu.no

[2] *Department of Telematics,*
Norwegian University of Science and Technology
peter@item.ntnu.no

[3] *Department of Computer and Information Science,*
Norwegian University of Science and Technology
pinar@idi.ntnu.no

Abstract Trust is situation-specific and the trust judgment problem with which the truster is confronted might be, in some ways, similar but not identical to some problems the truster has previously encountered. The truster then may draw information from these past experiences useful for the current situation. We present a knowledge-intensive and model-based case-based reasoning framework that supports the truster to infer such information. The suggested method augments the typically sparse trust information by inferring the missing information from other situational conditions, and can better support situation-aware trust management. Our framework can be coupled with existing trust management models to make them situation-aware. It uses the underlying model of trust management to transfer trust information between situations. We validate the proposed framework for Subjective Logic trust management model and evaluate it by conducting experiments on a large real dataset.

1. Introduction

This paper presents a context management framework (CMF) that employs case-based reasoning [Mor94] to analyze the correlation between trust information among various situations and help to bootstrap in unanticipated situations using trust information available from similar situations. The case-based reasoning technique is particularly useful for tasks that are experience-intensive, that involve plausible (i.e. not sound) reasoning and have incomplete rules to apply.

The fundamental principle of the case-based reasoning technique is similar to that of the human analogical reasoning process which employs solutions of past problems to solve current ones. The reasoning process is generally composed of three stages: remembering, reusing, and learning. Remembering is the case-retrieval process, which retrieves relevant and useful past cases. In the reusing step, the case-based reasoning

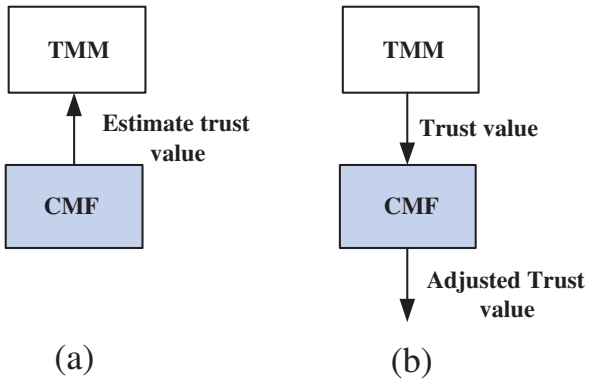

Figure 1. Scope and interconnection of context management framework (CMF) and trust management model (TMM). a) Estimation of the trust value in unknown situations. b) Adjustment of the output of TMM (trust value) based on the underlying situation.

system applies the cases that have been retrieved to find an effective solution to the current problem. Learning is the process of casebase enhancement. At the end of each problem-solving session the new case and problem-solving experiences incorporated into the casebase [JHS99].

We present a universal mechanism (called CMF) that can be combined with existing trust management models (TMM) to extend their capabilities towards efficient modeling of the situation-aware trust by

- estimating the trust values based on similar situations, in unknown situations or for unknown trustees when there is no information available. Therefore, CMF can help TMM to bootstrap (Figure 1(a)).
- adjusting the output of TMM (trust value) based on the underlying situation, thus, providing situation-awareness for TMM (Figure 1(b)).

In our approach TMM is implemented using the Subjective Logic [JHP06]. One of our main contributions is the extension of the Subjective Logic with a context-sensitive domain model.

The rest of this paper is organized as follows: In section 2, we briefly explain the Subjective Logic as an example of the trust management model. Our proposed model for trust inference is described in section 3. Next in section 4, we present the evaluation plan and results. Section 5 provides an overview of the related research. Finally, conclusion and some ideas for future work are given in section 6.

2. Subjective Logic Trust Management Model

In this section, we briefly explain the Subjective Logic fundamentals and give reasons why it needs to be extended with a *situation* dimension. Subjective Logic [Jøs01] enables the representation of a specific belief calculus in which trust is expressed by a belief metric called opinion. An opinion is denoted by $\omega_B^A = (b, d, u, a)$ expressing the

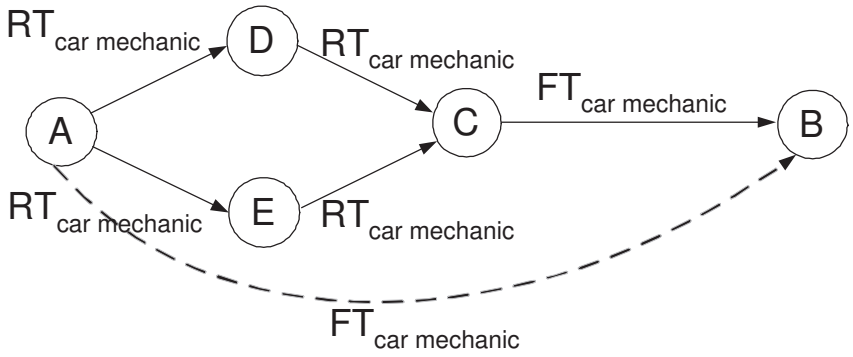

Figure 2. Trust transitivity and parallel combination of trust paths. FT is functional trust and RT is referral trust.

belief of a relying party A in the trustworthiness of another party B. The parameters b and d represent the belief respectively. disbelief in B's trustworthiness while u expresses the uncertainty in A's trust in B. All the three parameters are probability values between 0 and 1, and fulfill the constraint $b+d+u=1$. The parameter a is called the base rate and determines how uncertainty contributes to the opinion's probability expected value which is calculated as $E(\omega_x^A) = b + au$. The opinion space can be mapped into the interior of an equal-sided triangle, where the three parameters b, d, and u determine the position of the point in the triangle representing the opinion.

Based on the Subjective Logic, there are two different types of trust relations: *functional trust* (FT_B^A) and *referral trust* (RT_B^A). The former concerns A's direct trust in B performing a specific task ,while the latter concerns A's trust in B giving a recommendation about someone else doing a task. In other words, it is the trust in the ability to refer to a suitable third party. The simplest form of trust inference is trust transitivity which is widely discussed in literature [DKG+05, GKRT04, QHC07]. That is, if A trusts B who trusts C, then A will also trusts in C. A valid transitive trust path requires that the last edge in the path represents functional trust and that all other edges in the path represents referral trust. Referral trust transitivity and parallel combination of trust paths are expressed as part of the Subjective Logic model (figure 2) [JHP06].

The discounting operator ($\otimes$) [Jøs02] is used to derive trust from transitive trust paths, and the consensus operator ($\oplus$) allows to combine parallel transitive trust paths. The trust network in figure 2 can then be expressed as

$$FT_B^A = ((RT_D^A \otimes RT_C^D) \oplus (RT_E^A \otimes RT_C^E)) \otimes FT_B^C \quad (1)$$

There are two reasons for extension of the Subjective Logic with situation representation. First, It has been shown [CH96] that trust is not always transitive in real life. For example, the fact that A trusts B to fix her car and B trusts C to look after his child does not imply that A trusts C for fixing the car, or for looking after her child. However, under certain semantic constraints, trust can be transitive and a trust referral system can be used to derive transitive trust. The semantic constraint in the Subjective

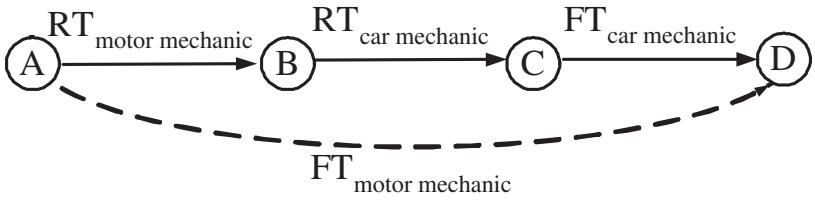

Figure 3. Trust transferability among similar situations.

Logic is that the subject of trust should be the same along the entire path, for example all trust subjects should be "to be a good car mechanic" (figure 2) or "looking after her child". On the other hand, this constraint is relaxed in our proposal by introducing the notion of situation. We suggest that trust situations along a transitive trust path can be different but similar to each other. For instance, trust situations can be "to be a good car mechanic" or "to be a good motor mechanic" (figure 3). In this way, we are able to use trust information from available similar situations (section 6 provides the details).

Second, Jøsang introduces three different versions of the consensus operator (denoted by $\oplus$, $\underline{\oplus}$, $\tilde{\oplus}$ respectively) for fusion of independent, dependent, and partially dependent trust opinions [JMP06]. If *A* and *B* have simultaneously observed the same event in the situation then their opinions are dependent. If *A* and *B* observed the same event during two partially overlapping situations then their opinions are partially dependent (e.g. *A* and *B* observed the same event of fire at the same time. *A* was in the place of fire, while *B* saw it on TV). Jøsang assumes that fraction of the overlapping observations is known and proposes formulas to estimate dependent and independent parts of the two observations to define the consensus operator of partially dependent opinions ($\tilde{\oplus}$). We propose to calculate the fraction of overlapping observations as the similarity measure between the two situations.

3. The Proposed Framework

We consider two approaches for the inference task among situations: rule-based inference and similarity-based reasoning, depicted respectively as case-based reasoner (CBR) and rule-based reasoner (RBR) modules in figure 4. The former provides the first role (Figure 1(a)), estimation of the trust value in unanticipated situations and the latter is responsible for the second role (Figure 1(b)) of CMF, adjustment of the trust values based on underlying situation. The gray box in figure 4 shows the focus of this paper.

3.1 Case-based Reasoner Module

In the case-based reasoning approach, knowledge is distributed among the four knowledge containers: ontology, casebase, similarity measures, and solution transformation.

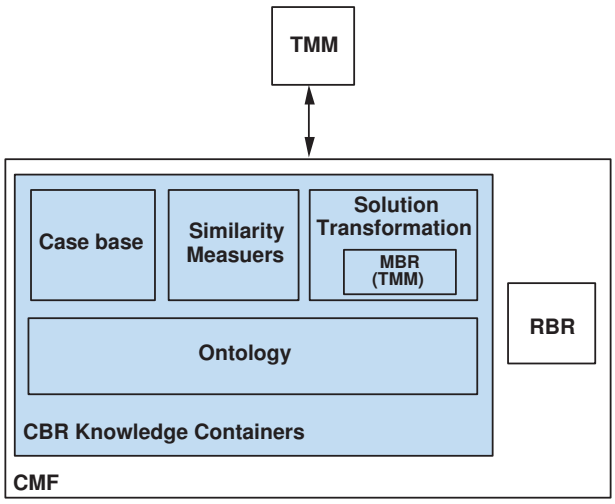

Figure 4. Knowledge containers in case-based reasoner (CBR). TMM: trust management model, MBR: Model-based reasoner, RBR: rule-based reasoner, CMF: context management framework.

- *Ontology*: We represent the situations in the pertinant domain in form of an ontology. A situation consists of set of *contexts* which are captured as nodes of the ontology. Figure 5 depicts the ontology related to user-movie ratings. In this example, a situation has two main contexts: *User* and *Movie*. Demographic information for the users (age, occupation, sex, and zip code) are *local contexts* for the *User* context and movie genres are *local contexts* for the *Movie* context.

- *Casebase*: The characterizations of the previous experiences and the recommendations (trust information including truster, trustee, trust value, and situation) are stored as elements of cases in the casebase. Cases are represented as attribute-value pairs.

- *Similarity* [1]: The similarity between situations is a weighted sum of the similarity between their contexts. Similarity between contexts, in turn, are computed as the wighted sum of the similarity between the underlying local contexts. According to the Tverskys formula [T+77], the similarity between two concepts A and B can be determined in the following way:

$$S(A,B) = \frac{|U(A) \cap U(B)|}{|U(A) \cap U(B)| + \alpha\,|U(A) \backslash U(B)| + (1-\alpha)\,|U(B) \backslash U(A)|} \qquad (2)$$

  $U(A)$ and $U(B)$ are the sets of properties of concepts A and B, respectively. The function U takes into account the depth of compared concepts in the ontology hierarchy. α is a value in the range $[0, 0.5]$. The value of 0 implies that the

[1] In [TKH08a] we provide a comprehensive set of similarity measurement algorithms.

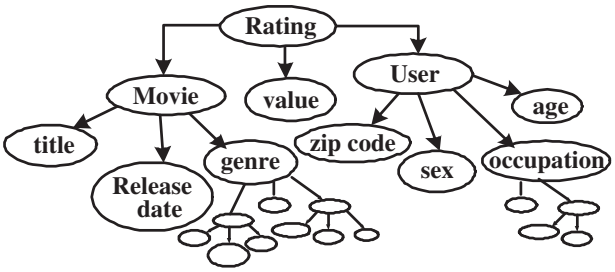

Figure 5. The ontology example for user-movie ratings.

differences of *A* with respect to *B* are not sufficient to conclude that they are similar, and the value of 0.5 means that the differences are necessary and sufficient to conclude such an assumption. Figure 6 illustrates an example of the similarity calculation.

In our approach, equation (2) is used to compare the attributes with each other, while the comparison between the *values* of an attribute is performed using the following general comparasion guidelines:

- *Categorical*: values in the same category are similar (e.g., weather).
- *Continuous*: closer values are alike (e.g., time).
- *Hierarchical*: values in the same hierarchy are similar (e.g., location).

Attributes which do not have these characteristics may require a custom comparator to be defined for them.

- *Solution transformation*: The model-based reasoner (MBR) is responsible for adaptation or transformation of a solution (trust value) from previous experiences to the current problem of trust judgment. It uses TMM to estimate trust value for the current situation based on trust values of the similar situations (see figure 4). In section 3.2.1, we consider the Subjective Logic model as TMM and provide details for the solution transformation module.

3.2 Processes

CMF is generally composed of three processes: Remembering, Reusing, and Learning.

- *Remembering*: The query (the current trust assessment question) is compared to cases (past trust assessment experiences) in the casebase and N most similar cases are retrieved (N nearest neighbors). This process uses the ontology to measure the similarity between the query and each case in the casebase.
- *Reusing*: A trust value is predicted for the query using the solution transformation module.

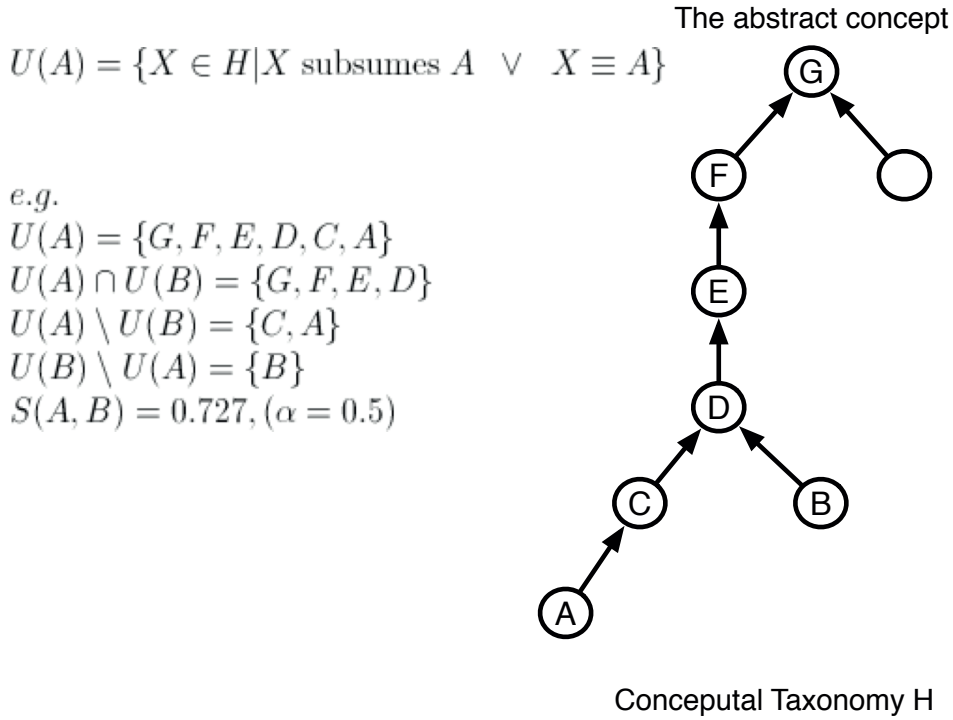

Figure 6. Relations taxonomy.

- *Learning*: A new case is built from the query and the predicted value and is added to the casebase for future uses.

In following, we explain the details for solution transformation module considering the Subjective Logic as TMM.

3.2.1 Solution Transformation in Case of the Subjective Logic

We explain the functionality of the model-based reasoner through extension of the Subjective Logic model as TMM. If A has functional trust in B in situation C_1, then A can infer its functional trust to B in situation C_2 which is a similar situation. For example, if A trusts B as a good car mechanic then A will probably trust B in repairing motorcycles since there is a large similarity between the domains of repairing cars and motorcycles.

Similarly to Jøsang's way to define opinions, we use triples to describe similarity which enables us to use the Subjective Logic operators.

DEFINITION 1 *The similarity opinion $S_{C_1}^{C_2}$ from C_1 towards C_2 is the triple* [2] *(*similarity*,* non-similarity*,* uncertainty*) and fulfills the constraints that the sum of all three values is equal to* 1*. If $C_1 = C_2$, the similarity opinion is defined to be* $(1,0,0)$*. Otherwise, it is calculated based on the measure of similarity ($S(C_1,C2)$) between the two situations C_1 and C_2 and the depth of concepts in the ontology (see (2)):*

$$S_{C_1}^{C_2} = (\frac{S(C_1,C_2)\cdot UN(C_1,C_2)}{k+UN(C_1,C_2)}, \frac{(1-S(C_1,C_2))\cdot UN(C_1,C_2)}{k+UN(C_1,C_2)}, \frac{k}{k+UN(C_1,C_2)}) \quad (3)$$

Here, k is a constant and $UN(C_1,C_2) = |U(C_1) \cup U(C_2)|$ defining the number of properties in play at all. In general, the higher the similarity value is, the less uncertain we are, and the uncertainty will be lower as more details ($UN(C_1,C_2)$) are available in comparison of the two situations C_1 and C_2.

Our similarity opinion is a special form of referral trust. It reflects that the akin situations of C_1 and C_2 is a kind of recommendation (reminding) to A to treat in situations C_1 and C_2 similarly. Thus, we see the consensus operator $\otimes$ as the correct mechanism to combine the similarity opinion between C_1 and C_2 with the functional trust of A in B in order to infer the functional trust of A in B:

$$FT_{B,C_1}^{A} = S_{C_1}^{C_2} \otimes FT_{B,C_2}^{A} \quad (4)$$

$FT_{B,X}^{A}$ is extended notation for A's functional trust to B which considers the underlying situation X. The higher the similarity between C_1 and C_2 is, the closer the trust of A to B in situation C_1 will be equal to that of between A and B in situation C_2. The lower this similarity is, the more uncertain A will be about whether to trust B or not in the second situation.

The same conversion formula can be used for Referral Trust.

$$RT_{B,C_1}^{A} = S_{C_1}^{C_2} \otimes RT_{B,C_2}^{A} \quad (5)$$

4. Evaluation

We chose MovieLens data [3] in view of the fact that we needed a context-enriched data to evaluate our work. The MovieLens data has been collected by the GroupLens Research Project at the University of Minnesota [4]. The data consists of 100,000 ratings from 943 users on 1682 movies with every user having at least 20 ratings and simple demographic information for the users is included. Figure 5 depicts the ontology which corresponds to the MovieLens data.

[2] This metric is inferred from a metric for the trust value computation [JK98] by Jøsang and Knapskog.
[3] http://www.grouplens.org/node/73
[4] http://www.cs.umn.edu/Research/GroupLens/data/

User attributes are age, sex and 19 occupation categories [5], zipcode, and movie attributes are 19 film genres [6]. Much richer movie content can be obtained from the Internet Movie Database (IMDB) [7]. We consider user and movie concepts as contexts and user and movie attributes as local contexts to form the situation for each rating.

4.1 Data Setup

There are 5 datasets which are 80%/20% splits of the data into training and test data (training set of 80,000 ratings, and the test set of 20,000 ratings). Each of these datasets have disjoint test sets; this is for 5 fold cross validation (where we repeat our experiment with each training and test set and average the results). The test sets are used as references for the accuracy of the predictions.

In the MovieLens data, rating values 1 and 2 represent negative ratings, 4 and 5 represent positive ratings, and 3 indicates ambivalence (we consider them as -2,-1,0,+1,+2). In order to convert these rating values to the Subjective Logic opinions (the triple $(b,d,u), b+d+u=1$) we can use the following conversion method:

$$b = \frac{\sum_{i=2}^{n}(i-1)\cdot f(i)}{c+(n-1)\cdot\sum_{i=1}^{n} f(i)}, \quad d = \frac{\sum_{i=1}^{n-1}(n-i)\cdot f(i)}{c+(n-1)\cdot\sum_{i=1}^{n} f(i)}, \quad u = \frac{c}{c+(n-1)\cdot\sum_{i=1}^{n} f(i)} \tag{6}$$

where the number of ratings at level i is described by function $f(i)$ and c is a constant.

4.2 Experimental Setup

The casebase is built up from the ratings in the training set. Each case is composed of four parts: user identifier, movie identifier, rating value, and situation including user and movie information. Ratings in the test set forms queries to CMF and each query is composed of three parts: user identifier, movie identifier, and the situation (the rating value is removed). The rating value in the query is predicted by CMF using the casebase, and then consequently compared with the removed value in the test set.

Four types of evaluation criteria are used in this paper:

- Coverage: measure of the percentage of movies in the test dataset that can be predicted.
- FCP: fraction of correct predictions.
- MAE (Mean Absolute Error) : average of the prediction error (difference between probability expected values of predicted and real opinions) over all queries.

[5] Occupation list: administrator, artist, doctor, educator, engineer, entertainment, executive, healthcare, homemaker, lawyer, librarian, marketing, none, other, programmer, retired, salesman, scientist, student, technician, writer.
[6] Film genres: unknown, action, adventure, animation, children, comedy, crime, documentary, drama, fantasy, film-noir, horror, musical, mystery, romance, sci-fi, thriller, war, western.
[7] http://us.imdb.com

- RMSE (root mean squared error) : root mean of the average of the squared prediction error. RMSE tends to emphasize large errors.

The evaluation is described as a pseudo-code in algorithm 4.1. First, the casebase and the set of queries are built from training and test sets, respectively. Second, the *Remember* procedure is called for each query computes the similarity between each case in the casebase and the query. Cases with a similarity less than a threshold are ignored and the ten most similar cases among the remainings are retrieved. Next, by calling the *Reuse* procedure, a rating value is predicted for the query (R_q) based on the rating values of the retrieved cases ($R_i, i = 1..10$) and their similarity measures (S_i) which are calculated by the *Similarity* procedure.

$$R_q = (S_1 \otimes R_1) \oplus (S_2 \otimes R_2) \oplus \ldots \oplus (S_{10} \otimes R_{10}) \tag{7}$$

Then, a new case is built which contains user and movie information of the query and the predicted rating value is added to the casebase by calling the *Learn* procedure. The predicted ratings form the *predicted set*. Finally, the *test* and *predicted* sets are compared according to the four metrics (Coverage, FCP, MAE, and RSME) by calling the *Evaluate* procedure.

The *Similarity* procedure (see algorithm 4.2) calculates weighted average of similarity measures of local contexts (age, sex, occupation, and zipcode for users and genres for movies) to determine the similarity between situations. In our implementation these weights are 0.2, 0.15, 0.1, 0.05, 0.5 respectively and are determined based on the fact that how much the local context can affect the rating decision. The comparator for each local context are:

- Age: Closer values are more similar.
- Sex: The similarity value is 1 for identical sex values and 0 otherwise.
- Occupation: The similarity is calculated according to (2) for similarity measurement on the ontology.
- Zipcode: ZIP codes are numbered with the first digit representing a certain group of U.S. states, the second and third digits together representing a region in that group (or perhaps a large city) and the fourth and fifth digits representing a group of delivery addresses within that region. We assign similarity values of 1, 0.75, 0.5 to the same delivery address, region, and state group respectively.
- Movie genre: The similarity is calculated using (2) to measure similarity on the ontology.

Our baseline is the Pearson algorithm [MA04] which relies on Pearson correlation coefficient to produce a correlation metric between users. This correlation is then used to weigh the rating of each relevant user. The Pearson correlation between users A and B is defined as:

$$P_{A,B} = \frac{\sum_{i=1}^{m} (R_{A,i} - \bar{R}_A) \times (R_{B,i} - \bar{R}_B)}{\sigma_A \times \sigma_B} \tag{8}$$

Algorithm 4.1: CONTEXT MANAGEMENT FRAMEWORK($test\_set, training\_set$)

main
global $casebase, similarity$
comment: Build "casebase" from the training set and "queries" from the test set
$similarity[1..size(casebase)] \leftarrow 0$
comment: "similarity" array stores similarity measures between the query and the cases
for each $query \in queries$
do $\begin{cases} neighbors \leftarrow \text{REMEMBER}(query, casebase) \\ predicted\_rating \leftarrow \text{REUSE}(neighbors) \\ \text{LEARN}(query, predicted\_rating) \\ predicted\_set \leftarrow predicted\_set \cup predicted\_rating \end{cases}$
EVALUATE($test\_set, predicted\_set$)

procedure REMEMBER($query$)
for each $case \in casebase$
do $\begin{cases} sim \leftarrow \text{SIMILARITY}(query, case) \\ \textbf{if } sim >= THRESHOLD \\ \quad \textbf{then } similarity[case] \leftarrow sim \end{cases}$
return (ten most similar cases)

procedure REUSE($neighbors$)
$predicated\_opinion \leftarrow (0,0,1)$
for each $ncase \in neighbors$
do $\begin{cases} similarity\_opinion \leftarrow (similarity[ncase], 0, 1 - similarity[ncase]) \\ new\_opinion \leftarrow similarity\_opinion \otimes ncase.rating \\ predicted\_opinion \leftarrow predicted\_opinion \oplus new\_opinion \end{cases}$
return ($predicted\_opinion$)

procedure LEARN($query, predicted\_rating$)
$new\_case \leftarrow query.user \cup query.movie \cup predicted\_rating$
$casebase \leftarrow casebase \cup new\_case$

procedure EVALUATE($test\_set, predicted\_set$)
$coverage \leftarrow$ fraction of predicted ratings
$fcp \leftarrow$ fraction of correct predictions
$mae \leftarrow$ mean absolute error of predictions
$rmse \leftarrow$ root mean squared error of predictions
output ($coverage, fcp, mae, rmse$)

Algorithm 4.2: SIMILARITY$(query, case)$

procedure SIMILARITY$(query, case)$
$user_q \leftarrow query.user$
$user_c \leftarrow case.user$
$age\_sim \leftarrow 1 - \frac{age_q - age_c}{age_{max} - age_{min}}$
if $sex_q == sex_c$
  then $sex\_sim \leftarrow 1$
  else $sex\_sim \leftarrow 0$
$occupation\_sim \leftarrow$ ONTOLOGYSIM$(occupation_q, occupation_c)$
comment: "OntologySim" calculates contextual similarity according to (2)
if $zipcode_q(1) == zipecode_c(1)$
  then
    if $zipcode_q(2,3) == zipecode_c(2,3)$
      then
        if $zipcode_q(4,5) == zipecode_c(4,5)$
          then $zipcode\_sim \leftarrow 1$; **comment:** the same delivery address
          else $zipcode\_sim \leftarrow 0.75$; **comment:** the same region
      else $zipcode\_sim \leftarrow 0.5$; **comment:** the same state group
  else $zipcode\_sim \leftarrow 0$
$movie\_sim \leftarrow$ ONTOLOGYSIM$(movie_q.genre, movie_c.genre)$
$total\_sim \leftarrow 0.2 \cdot age\_sim + 0.15 \cdot sex\_sim + 0.1 \cdot occupation\_sim$
$+0.05 \cdot zipcode\_sim + 0.5 \cdot movie\_sim$
return $(total\_sim)$

where m is the number of movies that both users rated. $R_{A,i}$ is the rating, user A gave to movie i. $\bar{R}_A$ is the average rating user A gave to all movies, and σ_A is the standard deviation of those ratings. Once the Pearson correlation between a user and all other users is obtained, the predicted movie rating is calculated as:

$$R_{A,i} = \bar{R}_A + \frac{\sum_{U=1}^{n} (R_{U,i} - \bar{R}_U) \times P_{A,U}}{\sum_{u=1}^{n} |P_{A,U}|} \tag{9}$$

Use of the Pearson correlation coefficient is quite common in the field of collaborative filtering, and results obtained with this method will be used to gauge the performance of other algorithms. Moreover, the Pearson algorithm uses only the rating information while our method use situational information to do the prediction.

Table 1. Final evaluation results

Metric	DS1	DS2	DS3	DS4	DS5	Average	Pearson
Coverage	43.82	43.88	44.94	45.42	45.06	44.62	99.83
FCP	0.3629	0.3497	0.3299	0.3345	0.3417	0.3437	0.1993
MAE	0.1605	0.1600	0.1656	0.1648	0.1626	0.1627	0.3049
RMSE	0.2742	0.2717	0.2757	0.2739	0.2724	0.2736	0.3804

DS=Dataset

4.3 Discussion of the Obtained Results

In table 1, we present the final results of the evaluation. We start by commenting the row "Coverage". The coverage becomes an important issue on a very sparse dataset that contains a large portion of cold-start users since many trust values become hardly predictable [MA07]. The results ($Coverage \approx 0.45\%$) indicate that our model is able to predicate approximately one rating from each two ratings. For the Pearson algorithm the coverage is not perfect merely because not all movies in the test dataset have a rating in the training dataset. The second important result is the fraction of correct predictions (FCP) is 0.34 which shows that from each 10 predicted ratings between 3 and 4 ratings are predicted with exact values. Further, the prediction errors (MAE and RMSE) for the other ratings that are not predicted exactly (between 6 and 7 ratings from each 10 predicted ratings) are small in comparison with the Pearson method ($MAE \approx 0.12$ & $RMSE \approx 0.20$).

All-in-all, the results of the evaluation lead to the expectation that our approach provides an improvement over the Pearson algorithm and this implies that situational information is useful in making predictions.

5. Related Research

CMF is a *knowledge-intensive CBR* which is designed to extend situational inference capabilities of *trust* management models. More precisely, the aim is to reuse the available trust information (direct experiences and recommendations) in similar situations for the current problem and we use semantic (ontology-based) similarity measures. Although CBR techniques are extensively used for recommender systems [AAM02, RGBDAGC08] and there are some works which use CBR to build more trust through providing explanations [PC06, PC07, Lea96], to the best of our knowledge this proposal is quite new. In this section, we briefly explain the related researches which are based on context-aware trust management and thus more closely resemble our goal.

According to the literature, the extension of a trust model with context representation can reduce complexity in the management of trust relationships [NWvSL07], improve the recommendation process [NWvSL07], help to infer trust information in context hierarchies [HY06], improve performance [RP07], help to learn policies/norms at runtime [RP07, TLU06], and provide protection against changes of identity and first

time offenders [RP07]. Context related information has been represented as Context-aware domains [NWvSL07], Intensional Programming [WA08], Multi-dimensional goals [GDFB06], Clustering [RP07], and Ontologies [TLU06].

[SLP04] provides a survey of different approaches to model context for ubiquitous computing. In particular, numerous approaches are reviewed, classified relative to their core elements and evaluated with respect to their appropriateness for ubiquitous computing. The authors conclude that the most promising assets for context modeling of ubiquitous computing environments can be found in the ontology category in comparison with other approaches like key-value models, mark-up scheme models, graphical models, object-oriented models, and logic based models. This selection is based on the six requirements dominant in pervasive environments: distributed composition, partial validation, richness and quality of information, incompleteness and ambiguity, level of formality, and applicability to existing environments.

We present a state-of-the-art survey of context representation for trust management in [TKH08b]. In the rest of this section ontology-based approaches to this problem are examined in more details.

Golbeck et al. [GPH03] propose an ontology for trust. In [GH04] the authors consider a model using context-specific reputation by assigning numeric ratings to different types of connections based on context of the analysis. In [TLU06] rules to describe how certain context-sensitive information (trust factors) reduces or enhances the trust value have been specified for this trust ontology.

In [TLU06] contextual information (i.e., context attributes) is used to adjust the output of a trust determination process. Each attribute can adjust the trust value positively or negatively according to a specified weight. As an illustration, if t is the trust value and ω is the weight of the context property then the adjusting function can be t^{ω} for decrease or $\sqrt[\omega]{t}$ for increase. A context ontology connects the context attributes with each other in an appropriate manner, enabling the utilization of context attributes which do not exactly match the query, but are "close enough" to it.

In [CBGS07], cases where a truster does not have enough information to produce a trust value for a given task, but she knows instead the previous partner behavior performing similar tasks, are considered. This model estimates trust using the information about similar tasks. The similarity between two tasks is obtained from the comparison of the task attributes.

6. Conclusion and Future Directions

To sum up, we propose a framework based on the case-based reasoning paradigm and the representation of deep knowledge to make existing trust management models situation-aware. This framework has been validated for the Subjective Logic trust management model as an example and evaluated using a real large-scale dataset. It can also be considered as an inference mechanism which deals with the sparsity and cold-start problems of a web of trust.

The original Subjective Logic can be applied to determine transitivity only if the subject of the trust relations along the entire path is the same. However, trust relations with the same subject are not always available. Our proposal opens up the possibility

to draw transitivity also when the subject (situation) of the available trust relations are not the same but are similar. First, the trust relations with similar situations with the current problem are retrieved from the casebase using the ontology and the similarity measurement algorithm (remembering past similar trust experiences). Next, they are converted (using (4) and (5)) to equivalent trust relations in the current problem by solution transformation module (reusing the trust information from the past similar trust experiences). Then, the transitive trust path is formed and final trust is calculated according to the Subjective Logic (1). Solution of the current problem is stored as a new case in the casebase (the learning process of CBR).

In the future, we aim to add a Risk Management Module to this framework. Risk evaluation becomes important in inferring trust values among situations especially when the trustworthiness of some principal is completely unknown and no recommendation information is available. The intuitive idea behind such a risk assessment can be to look up the in the casebase to see if there are any similar previous interactions, i.e., if we have previously encountered an entity with similar trust attributes and similar risk attributes in the same situation. The ontology part should be able to describe the level of situational risk, whereby the higher the risk of negative outcome, the higher the level of precision that must be captured.

References

[AAM02] S. Aguzzoli, P. Avesani, and P. Massa. Collaborative Case-Based Recommender Systems. *LECTURE NOTES IN COMPUTER SCIENCE*, pages 460–474, 2002.

[APW05] V.S. Alagar, J. Paquet, and K. Wan. Intensional Programming for Agent Communication. In *Declarative agent languages and technologies II: second international workshop, DALT 2004, New York, NY, USA, July 19, 2004: revised selected papers*, page 239. Springer-Verlag New York Inc, 2005.

[CBGS07] A. Caballero, JA Botia, and A. Gomez-Skarmeta. On the Behaviour of the TRSIM Model for Trust and Reputation. *LECTURE NOTES IN COMPUTER SCIENCE*, 4687:182, 2007.

[CH96] B. Christianson and W.S. Harbison. Why Isn't Trust Transitive? In *Proceedings of the International Workshop on Security Protocols*, pages 171–176. Springer-Verlag London, UK, 1996.

[DKG+05] L. Ding, P. Kolari, S. Ganjugunte, T. Finin, and A. Joshi. Modeling and Evaluating Trust Network Inference. Technical report, MARYLAND UNIV BALTIMORE DEPT OF COMPUTER SCIENCE AND ELECTRICAL ENGINEERING, 2005.

[GDFB06] N. Gujral, D. DeAngelis, K.K. Fullam, and K.S. Barber. Modeling Multi-Dimensional Trust. In *the Proceedings of the Workshop on Trust in Agent Societies*, pages 8–12, 2006.

[GH04] J. Golbeck and J. Hendler. Inferring Reputation on the Semantic Web. In *Proceedings of the 13th International World Wide Web Conference*, 2004.

[GKRT04] R. Guha, R. Kumar, P. Raghavan, and A. Tomkins. Propagation of trust and distrust. In *Proceedings of the 13th international conference on World Wide Web*, pages 403–412. ACM Press New York, NY, USA, 2004.

[GPH03] J. Golbeck, B. Parsia, and J. Hendler. Trust Networks on the Semantic Web. In *Proceedings of Cooperative Intelligent Agents*, volume 2003. Springer, 2003.

[HY] S. Holtmanns and Z. Yan. Context-Aware Adaptive Trust.

[JHP06] A. Jøsang, R. Hayward, and S. Pope. Trust network analysis with subjective logic. In *Proceedings of the 29th Australasian Computer Science Conference-Volume 48*, pages 85–94. Australian Computer Society, 2006.

[JHS99] C. Jung, I. Han, and B. Suh. Risk Analysis for Electronic Commerce Using Case-Based Reasoning. *Int. J. Intell. Sys. Acc. Fin. Mgmt*, 8:61–73, 1999.

[JK98] A. Jøsang and S. J. Knapskog. A metric for trusted systems. In *Proceedings of the 21st National Security Conference*. NSA, 1998.

[JMP06] A. Josang, S. Marsh, and S. Pope. Exploring Different Types of Trust Propagation. *LECTURE NOTES IN COMPUTER SCIENCE*, 3986:179, 2006.

[Jøs01] A. Jøsang. A Logic for Uncertain Probabilities. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 9(3):279–311, 2001.

[Jøs02] A. Jøsang. The consensus operator for combining beliefs. *Artificial Intelligence*, 141(1-2):157–170, 2002.

[Lea96] D.B. Leake. CBR in Context: The Present and Future. *Case-Based Reasoning: Experiences, Lessons, and Future Directions*, pages 3–30, 1996.

[MA04] P. Massa and P. Avesani. Trust-Aware Collaborative Filtering for Recommender Systems. *LECTURE NOTES IN COMPUTER SCIENCE*, pages 492–508, 2004.

[MA07] P. Massa and P. Avesani. Trust-aware recommender systems. In *Proceedings of the 2007 ACM conference on Recommender systems*, pages 17–24. ACM Press New York, NY, USA, 2007.

[Mor94] B.W. Morris. SCAN: a case-based reasoning model for generating information system control recommendations. *International Journal of Intelligent Systems in Accounting, Finance and Management*, 3(1):47–63, 1994.

[NWvSL07] R. Neisse, M. Wegdam, M. van Sinderen, and G. Lenzini. Trust Management Model and Architecture for Context-Aware Service Platforms. *LECTURE NOTES IN COMPUTER SCIENCE*, 4804:1803, 2007.

[PC06] P. Pu and L. Chen. Trust building with explanation interfaces. In *Proceedings of the 11th international conference on Intelligent user interfaces*, pages 93–100. ACM New York, NY, USA, 2006.

[PC07] P. Pu and L. Chen. Trust-inspiring explanation interfaces for recommender systems. *Knowledge-Based Systems*, 20(6):542–556, 2007.

[QHC07] D. Quercia, S. Hailes, and L. Capra. Lightweight Distributed Trust Propagation. In *Data Mining, 2007. ICDM 2007. Seventh IEEE International Conference on*, pages 282–291, 2007.

[RGBDAGC08] JA Recio-García, D. Bridge, B. Díaz-Agudo, and PA González-Calero. CBR for CBR: A Case-Based Template Recommender System. In *Advances in Case-Based Reasoning, 9th European Conference, ECCBR*, 2008.

[RP07] M. Rehak and M. Pechoucek. Trust modeling with context representation and generalized identities. *Klusch, M., Hindriks, K., apazoglou, MP, Sterling, L.(eds.) CIA*, pages 298–312, 2007.

[SLP04] T. Strang and C. Linnhoff-Popien. A context modeling survey. In *Workshop on Advanced Context Modelling, Reasoning and Management as part of UbiComp*, 2004.

[T+77] A. Tversky et al. Features of similarity. *Psychological Review*, 84(4):327–352, 1977.

[TKH08a] M. Tavakolifard, S. Knapskog, and P. Herrmann. Cross-Situation Trust Reasoning. In *Proceedings of The Workshop on Web Personalization, Reputation and Recommender Systems (WPRRS08)*. IEEE Computer Society Press, 2008.

[TKH08b] M. Tavakolifard, S. Knapskog, and P. Herrmann. Trust Transferability Among Similar Contexts. In *Proceedings of The 4th ACM International Workshop on QoS and Security for Wireless and Mobile Networks (Q2SWinet 2008)*. ACM, 2008.

[TLU06] S. Toivonen, G. Lenzini, and I. Uusitalo. Context-aware trust evaluation functions for dynamic reconfigurable systems. In *Proceedings of the Models of Trust for the Web Workshop (MTW06), held in conjunction with the 15th International World Wide Web Conference (WWW2006) May*, volume 22, 2006.

Bibliography

[ACMD+03] K. Aberer, P. Cudré-Mauroux, A. Datta, Z. Despotovic, M. Hauswirth, M. Punceva, and R. Schmidt. P-grid: a self-organizing structured p2p system. *ACM SIGMOD Record*, 32(3):29–33, 2003.

[AR04] A. Abdul-Rahman. *A framework for decentralised trust reasoning*. PhD thesis, University College London, 2004.

[BBEZG08] E. Bagheri, M. Barouni-Ebrahimi, R. Zafarani, and A.A. Ghorbani. A Belief-Theoretic Reputation Estimation Model for Multi-context Communities. *Lecture Notes in Computer Science*, 5032:48, 2008.

[BFL96] M. Blaze, J. Feigenbaum, and J. Lacy. Decentralized trust management. In *Proceedings of the IEEE Symposium on Security and Privacy*, pages 164–173, 1996.

[BG06] Ebrahim Bagheri and Ali A. Ghorbani. Behavior analysis through reputation propagation in a multi-context environment. In *Proceedings of the International Conference on Privacy, Security and Trust: Bridge the Gap Between PST Technologies and Business Services*, pages 40:1–40:7, 2006.

[BLB04] S. Buchegger and J.Y. Le Boudec. A robust reputation system for mobile ad-hoc networks. In *Proceedings of the Workshop on the Economics of Peer-to-Peer Systems*, June 2004.

[Bus98] V. Buskens. The social structure of trust. *Social Networks*, 20(3):265–289, 1998.

[CBGS06a] A. Caballero, J. Botia, and A. Gomez-Skarmeta. A new model for trust and reputation management with an ontology based approach for similarity between tasks. *Multiagent System Technologies*, pages 172–183, 2006.

[CBGS06b] Alberto Caballero, Juan A. Botía Blaya, and Antonio F. Gómez-Skarmeta. A new model for trust and reputation management with an ontology based approach for similarity between tasks. In *MATES*, pages 172–183, 2006.

[CDdV+02] F. Cornelli, E. Damiani, S.D.C. di Vimercati, S. Paraboschi, and P. Samarati. Choosing reputable servents in a P2P network. In *Proceedings of the International Conference on the World Wide Web*, pages 376–386, May 2002.

[CIM+09] W. Conner, A. Iyengar, T. Mikalsen, I. Rouvellou, and K. Nahrstedt. A trust management framework for service-oriented environments. In *Proceedings of the International Conference on the World Wide Web*, pages 891–900, 2009.

[CS01] M. Chen and J.P. Singh. Computing and using reputations for internet ratings. In *Proceedings of the ACM Conference on Electronic Commerce*, pages 154–162, October 2001.

[CY01] R. Chen and W. Yeager. Poblano: A distributed trust model for peer-to-peer networks. Sun Microsystems, inc. White Paper, 2001.

[DdVP+02] E. Damiani, D.C. di Vimercati, S. Paraboschi, P. Samarati, and F. Violante. A reputation-based approach for choosing reliable resources in peer-to-peer networks. In *Proceedings of the ACM Conference on Computer and Communications Security*, pages 207–216, 2002.

[Del00] C. Dellarocas. Immunizing online reputation reporting systems against unfair ratings and discriminatory behavior. In *Proceedings of the ACM Conference on Electronic Commerce*, pages 157–164, October 2000.

[Del02] C. Dellarocas. Goodwill hunting: An economically efficient online feedback mechanism for environments with variable product quality. In *Workshop on Agent Mediated Electronic Commerce IV: Designing Mechanisms and Systems*, pages 93–112, July 2002.

[Del03] C. Dellarocas. The digitization of word-of-mouth: Promise and challenges of online reputation systems. *Management Science*, 49(10):1407–1424, October 2003.

[Gam00] D. Gambetta. Can we trust trust. *Trust: Making and breaking cooperative relations*, pages 213–237, 2000.

[GCJ03] E. Gray, Y. Chen, and C. Jensen. Initial Investigation into Cross-context Trust and Risk Assessment. In *Proceedings of the IASTED International Conference on Communication, Network, and Information Security*, pages 56–61, 2003.

[GDFB06] N. Gujral, D. DeAngelis, K. Fullam, and K.S. Barber. Modeling multi-dimensional trust. In *the Proceedings of the Workshop on Trust in Agent Societies*, pages 8–12, 2006.

[GH04] J. Golbeck and J. Hendler. Accuracy of metrics for inferring trust and reputation in semantic web-based social networks. In *Proceedings of the International Conference on Knowledge Engineering and Knowledge Management*, October 2004.

[GJA03] M. Gupta, P. Judge, and M. Ammar. A reputation system for peer-to-peer networks. In *Proceedings of the International Workshop on Network and Operating Systems Support for Digital Audio and Video (NOSSDAV)*, pages 144–152, 2003.

[GKRT04] R. Guha, R. Kumar, P. Raghavan, and A. Tomkins. Propagation of trust and distrust. In *Proceedings of the International Conference on the World Wide Web*, pages 403–412, 2004.

[GPH03] J. Golbeck, B. Parsia, and J. Hendler. Trust networks on the semantic web. *Cooperative Information Agents VII*, pages 238–249, 2003.

[Gri05] N. Griffiths. Task delegation using experience-based multi-dimensional trust. In *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 489–496, 2005.

[Har68] G.H. Harman. Enumerative induction as inference to the best explanation. *The Journal of Philosophy*, 65(18):529–533, 1968.

[HF06] J. Huang and M.S. Fox. An ontology of trust: formal semantics and transitivity. In *ICEC06: Proceedings of the 8th international conference on Electronic Commerce*, pages 259–270, 2006.

[HJS06] T.D. Huynh, N.R. Jennings, and N.R. Shadbolt. An integrated trust and reputation model for open multi-agent systems. *Autonomous Agents and Multi-Agent Systems*, 13(2):119–154, 2006.

[HY06] S. Holtmanns and Z. Yan. Context-Aware Adaptive Trust. In *Proceedings of the Ambient Intelligence Developments Conference*, 2006.

[HZNR09] K. Hoffman, D. Zage, and C. Nita-Rotaru. A survey of attack and defense techniques for reputation systems. *ACM Computing Surveys*, 42(1):1–31, December 2009.

[JF03] R. Jurca and B. Faltings. An incentive compatible reputation mechanism. In *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1026–1027, July 2003.

[JHS99] C. Jung, I. Han, and B. Suh. Risk Analysis for Electronic Commerce Using Case-Based Reasoning. *Int. J. Intell. Sys. Acc. Fin. Mgmt*, 8:61–73, 1999.

[JIB07] A. Jøsang, R. Ismail, and C. Boyd. A survey of trust and reputation systems for online service provision. *Decision Support Systems*, 43(2):618–644, March 2007.

[KC06] R. Kerr and R. Cohen. Modeling trust using transactional, numerical units. In *Proceedings of the International Conference on Privacy, Security and Trust: Bridge the Gap Between PST Technologies and Business Services*, pages 21:1–21:11, October/November 2006.

[KSGM03] S.D. Kamvar, M.T. Schlosser, and H. Garcia-Molina. The eigentrust algorithm for reputation management in p2p networks. In *Proceedings of the International Conference on the World Wide Web*, pages 640–651, 2003.

[Mal01] R.A. Malaga. Web-based reputation management systems: Problems and suggested solutions. *Electronic Commerce Research*, 1(4):403–417, 2001.

[Mar94] S.P. Marsh. *Formalising trust as a computational concept*. PhD thesis, Dept. of Computing Science and Mathematics, University of Stirling, April 1994.

[Mau96] U. Maurer. Modelling a public-key infrastructure. In *Proceedings of the European Symposium on Research in Computer Security (ESORICS)*, pages 325–350, 1996.

[MB09] Z. Malik and A. Bouguettaya. Rateweb: Reputation assessment for trust establishment among web services. *The VLDB Journal*, 18(4):885–911, 2009.

[MRZ02] N. Miller, P. Resnick, and R. Zeckhauser. Eliciting Honest Feedback in Electronic Markets. *Working Paper Series*, August 2002.

[NWvS06] R. Neisse, M. Wegdam, and M. van Sinderen. Context-Aware Trust Domains. In *Proceedings of the European Conference on Smart Sensing and Context*, 2006.

[NWvSL07] R. Neisse, M. Wegdam, M. van Sinderen, and G. Lenzini. Trust Management Model and Architecture for Context-Aware Service Platforms. *Lecture Notes in Computer Science*, 4804:1803, 2007.

[PSD02] J.M. Pujol, R. Sangüesa, and J. Delgado. Extracting reputation in multi agent systems by means of social network topology. In *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 467–474, 2002.

[PTJL05] J. Patel, W. Teacy, N. Jennings, and M. Luck. A probabilistic trust model for handling inaccurate reputation sources. *Trust Management*, pages 413–419, 2005.

[RCP05] K. Regan, R. Cohen, and P. Poupart. The advisor-pomdp: A principled approach to trust through reputation in electronic markets. In *Proceedings of the Conference on Privacy Security and Trust*, 2005.

[RGPB06] M. Rehak, M. Gregor, M. Pechoucek, and J.M. Bradshaw. Representing Context for Multiagent Trust Modeling. In *Proceedings of the IEEE/WIC/ACM International Conference on Intelligent Agent Technology (IAT)*, pages 737–746, 2006.

[RKZF00] P. Resnick, K. Kuwabara, R. Zeckhauser, and E. Friedman. Reputation systems. *Communications of the ACM*, 43(12):45–48, December 2000.

[RMVS05] Y. Rebahi, V.E. Mujica-V, and D. Sisalem. A reputation-based trust mechanism for ad hoc networks. In *Proceedings of the IEEE Symposium on Computers and Communications (ISCC)*, pages 37–42, 2005.

[RP07] M. Rehak and M. Pechoucek. Trust modeling with context representation and generalized identities. In *Proceedings of the International Workshop on Cooperative Information Agents XI (CIA)*, pages 298–312, 2007.

[RY07] AGP Rahbar and O. Yang. Powertrust: A robust and scalable reputation system for trusted peer-to-peer computing. *Parallel and Distributed Systems, IEEE Transactions on*, 18(4):460–473, 2007.

[SL03] A. Singh and L. Liu. Trustme: anonymous management of trust relationships in decentralized p2p systems. In *Proceedings of the International Conference on Peer-to-Peer Computing*, pages 142–149, 2003.

[SLP04] Thomas Strang and Claudia Linnhoff-Popien. A context modeling survey. In *In: Workshop on Advanced Context Modelling, Reasoning and Management, UbiComp 2004 - The Sixth International Conference on Ubiquitous Computing, Nottingham/England*, 2004.

[SS01] J. Sabater and C. Sierra. Social regret, a reputation model based on social relations. *ACM SIGecom Exchanges*, 3(1):44–56, 2001.

[SS05] J. Sabater and C. Sierra. Review on computational trust and reputation models. *Artificial Intelligence Review*, 24(1):33–60, 2005.

[TC04] T. Tran and R. Cohen. Improving user satisfaction in agent-based electronic marketplaces by reputation modelling and adjustable product quality. In *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 828–835, July 2004.

[TD04] S. Toivonen and G. Denker. The impact of context on the trustworthiness of communication: An ontological approach. In *Proceedings of the Trust, Security, and Reputation on the Semantic Web workshop, held in conjunction with the 3rd International Semantic Web Conference (ISWC 2004), Hiroshima, Japan*, volume 127, page 17, 2004.

[THÖ09] M. Tavakolifard, P. Herrmann, and P. Öztürk. Analogical trust reasoning. *Trust Management III*, pages 149–163, 2009.

[TKH08a] Mozhgan Tavakolifard, Svein J. Knapskog, and Peter Herrmann. Cross-situation trust reasoning. In *WI-IAT '08: Proceedings of the 2008 IEEE/WIC/ACM International*

Conference on Web Intelligence and Intelligent Agent Technology, pages 67–71, Washington, DC, USA, 2008. IEEE Computer Society.

[TKH08b] Mozhgan Tavakolifard, Svein Johan Knapskog, and Peter Herrmann. Trust transferability among similar contexts. In *Q2SWinet '08: Proceedings of the 4th ACM symposium on QoS and security for wireless and mobile networks*, pages 91–97, New York, NY, USA, 2008. ACM.

[TLU06] S. Toivonen, G. Lenzini, and I. Uusitalo. Context-aware trust evaluation functions for dynamic reconfigurable systems. In *Proceedings of the Models of Trust for the Web Workshop*, May 2006.

[UZA08] M.G. Uddin, M. Zulkernine, and S.I. Ahamed. Cat: a context-aware trust model for open and dynamic systems. In *Proceedings of the 2008 ACM symposium on Applied computing*, pages 2024–2029. ACM, 2008.

[WA08] K. Wan and V. Alagar. An intensional functional model of trust. *Trust Management II*, pages 69–85, 2008.

[XL04] L. Xiong and L. Liu. Peertrust: Supporting reputation-based trust for peer-to-peer electronic communities. *IEEE Transactions on Knowledge and Data Engineering*, 16(7):843–857, 2004.

[YS02] B. Yu and M.P. Singh. An evidential model of distributed reputation management. In *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 294–301, 2002.

[ZL05] C.N. Ziegler and G. Lausen. Propagation models for trust and distrust in social networks. *Information Systems Frontiers*, 7(4):337–358, 2005.

[ZL09] H. Zhao and X. Li. H-trust: A group trust management system for peer-to-peer desktop grid. *Journal of Computer Science and Technology*, 24(5):833–843, 2009.

[ZMM99] G. Zacharia, A. Moukas, and P. Maes. Collaborative reputation mechanisms in electronic marketplaces. In *Proceedings of the Annual Hawaii International Conference on System Sciences*, January 1999.

[ZR02] R. Zeckhauser and P. Resnick. Trust among strangers in Internet transactions: Empirical analysis of eBay's reputation system. *The Economics of the Internet and E-commerce*, pages 127–157, 2002.