Evaluating User Representations for the Trustworthiness of Interaction Partners

Sebastian Ries
Darmstadt University of Technology
Hochschulstraße 10
64289 Darmstadt, Germany
ries@tk.informatik.tu-darmstadt.de

Daniel Schreiber
Darmstadt University of Technology
Hochschulstraße 10
64289 Darmstadt, Germany
schreiber@tk.informatik.tu-darmstadt.de

ABSTRACT

Trust in Ubiquitous Computing is about finding trustworthy partners for risky interactions in presence of uncertainty about identity, motivation, and goals of the potential interactions partners. For the integration of users in the selection of interaction partners it is not sufficient to provide robust algorithms for trust establishment in service provision and recommendations, but there is a need for a user interface, which allows for intuitive interpretation of the collected information about the trustworthiness of the potential interaction partners. In this paper, we evaluate three representations for trust based on collected evidences in a user study. The results show that the graphical representation which we have developed for our trust model, called CertainTrust, fits to the needs of the users.

Author Keywords

Trust, User Interface, Recommendations

ACM Classification Keywords

H5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

INTRODUCTION

The main driving force behind the idea of ubiquitous computing is to support humans in their everyday life. Ubiquitous computing environments are expected to be made up by a huge number of heterogeneous, loosely coupled devices. Thus, collaboration between devices, e.g., sharing information or resources, is an indispensable enabler for the evolvement of the potential of ubiquitous computing. Frequent collaboration requires frequent decisions about with whom to interact demanding for a non-intrusive way of decision making. We favour the approach of selecting interaction partners based on trust, which is built on experiences from past interactions. Although, one may be uncertain about the identity, motivation or goals of its interaction partners, direct experiences from past interactions are a good indicator whether one should interact with this partner another time, or not. Since direct experiences may be rare, indirect experiences, as recommendations and reputation information, are considered as additional sources of trust (see Figure 1). As the main goal is to support users in their everyday life, the risk associated with those applications can include many aspects. This can reach from wasting time, while waiting for the transfer of an information to complete, to all kinds of financial and legal implications which the usual everyday life task is related to, e.g., when buying goods online.

In our opinion, a trust model has to support autonomous decision making, as well as the participation of the user in this process. For the integration of the user in the process of decision making, as well as the process of trust establishment, it is necessary to provide a representation of trust which allows for an intuitive access to the information about the trustworthiness of potential candidates for an interaction. In cases in which the user wants to intervene in the decision making process the collected information needs to be presented in away which allows rapid processing by the user, and supports him in his decision making process. In the case, that the user wants to take part in the trust establishment process, it is furthermore necessary, that the user is also capable of manipulating opinions about the trustworthiness in an easily accessible way.

The remainder of the paper is structured as follows: In section 2, we present a scenario motivating trust as a basis for collaboration. In section 3, we shortly discuss the related work. In section 4, we introduce requirements for a representation of trust which is suitable for agents and users. Furthermore, we present the three representations, which have been evaluated in the user study. Section 5 shows procedure and the results of the user study. The discussion of the results is in section 6. At last, we provide our conclusions.

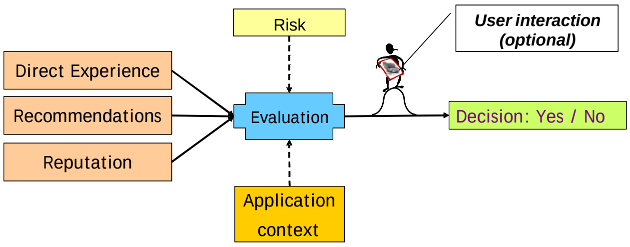

Figure 1: Trust as a well-founded basis for decisions. Main influence factors on trust evaluation and (optional) user intervention in the decision making process

SCENARIO

In this section we introduce a scenario in which trust enables non-intrusive collaboration in an opportunistic network. We present the scenario, as it has been presented to the participants of our user study.

You and your friends are on your way to a soccer match in the stadium in Frankfurt. You took your personal device – a next-generation mobile phone – with you, as usual. The week before, you informed your personal device, that you are looking for a certain song (mp3) and that you want to buy an mp3 player. While you are moving through the crowd in front of the stadium, your device searches (wireless) for potential interaction partners, which can offer the song or the mp3-player.

Shortly before the security check, one of your friends meets some of his colleagues. Your personal device discovers that a member of this group (Dirk) offers the song you are looking for. To reduce the risk of getting a file which is damaged or contains a virus, your personal device collects recommendations from the mobile devices of your friends, who either know Dirk or have had a number of interactions with him. As the recommendations are positive, your personal device downloads the song from Dirk's device. Afterwards, it checks the file for noise and viruses. As the file is clean, your personal device generates a profile for Dirk, and notes that there has been a positive interaction in the context of mp3-exchange.

This all has been done without your interaction; only after the successful exchange, a short vibration of your personal device indicates the positive interaction.

In this scenario the personal device took all decisions by itself. But in the case that your personal device only collected a few number of evidences, not sufficient for a autonomous decision, it could have notified you. Then, you have to decide whether to interact or not. Although in this example the risk associated with the interaction is limited, in the case that someone offers you a used mp3-player for 10 EUR, you will be glad about all evidences about the trustworthiness of this potential interaction partner.

RELATED WORK

Our scenario is motivated by an application, called musicClouds [5], which allows for autonomous sharing of music files in opportunistic networks. While the basic idea of this application is quite similar to ours, this application focuses on defining filters for specifying the meta-information of the files of interest, but not on the selection of the interaction partner. In [6], Hui et al. argue for the relevance of pocket switched networking, focusing on opportunistic networking, since there are numerous scenarios in which local connectivity might be preferred over an internet-based connection, due to bandwidth, latency or costs.

There are numerous trust models which allow the evaluation of the trustworthiness of an interaction partner in distributed environments [1, 2, 6, 7, 9 - 13]. In this paper, we are not focusing on the strength and weaknesses of the algorithms proposed for trust establishment, weighting of recommendations, and autonomous decision making; instead we focus on the user interfaces of those models. The research here is much more restricted than the research on the models itself. For trust models, which focus only on trust relationships between agents [2, 10, 13], it is sufficient if the model has a number-based representation using a continuous domain. But as it is important for users to feel in control of the system, and to check and manipulate the parameters, there is a need for an intuitive user interface. Thus, as proposed in previous work [11], we need one trust model supporting both, a representation suitable for agents and autonomous decision making, and a representation for the users.

With FilmTrust [3], Golbeck proposed an application in which users are able to rate the trustworthiness of their friends in a one-dimensional number based interface. The main short-coming of this approach, is that users are not able to express how certain they are about their judgement. Furthermore, when they get a recommended rating for a movie, it is also represented as a single value, and thus, not stating on how many recommendations the rating is based.

An approach to model trust to enable collaboration in virtual communities and agent societies, mimicking a human notion of trust, is presented in [1]. The model uses 4 labels for a trust value (very trustworthy, trustworthy, untrustworthy, and very untrustworthy) and a parameter for uncertainty (u+, u°, u-). As the first set of labels seems intuitive, the second one is completely arbitrary. There is no expression, which allows to interpret both values together. A further problem of this approach is that the computational model, which is used for combining direct experiences and recommendations, is not based on a continuous domain, but on a discrete one, and thus has to be done based on table look-ups rather than on a sound mathematical model.

In [12], a trust model is presented which allows to map numbers to labels using fuzzy logic, but the model is not capable of expressing trust and certainty together.

As label-based user interfaces are restricted to a small set of labels, a general short-coming of those approaches is that they lead to very coarse-grained representations. One way to overcome this problem is using graphical representations, as they can be easily interpreted by users, and allow to for an almost continuous representation, leading to a smaller loss of information in the user interface.

A first approach for developing a two-dimensional graphical representation of opinions about trustworthiness, which is also capable of expressing uncertainty in an integrated way, is presented in [6]. We will compare this approach with ours in the experiment below.

A slightly different domain which also deals with representing recommendations based on evidences can be found in online recommender systems. Well-known examples are eBay (www.eBay.de) and Amazon (www.amazon.com). Although these approaches have a simpler way to cope the integration of direct experiences and recommendations as the models above, both have to cope with the problem of representing not only the average of all ratings, but also the amount of evidences from which this value was derived. Both approaches do not present an integrated way for presenting the rating and the amount of evidences in an integrated way.

EVALUATED REPRESENTATIONS

In this section we will shortly introduce the main aspects which we identified to be necessary for representing trust based on evidences and recommendations. Afterwards, we will explain the representations of trust and recommendations which have been evaluated in the user study.

Requirements

In [8], reliability trust is defined as "the subjective probability by which an individual, A, expects that another individual, B, performs a given action on which its welfare depends". In [9] trust is defined as "a subjective expectation an agent has about another's future behavior based on the history of their encounters." The latter definition is a little narrower than the first one focusing on the interpretation and manipulation of trust by software agents; but both definitions go along with the idea to use trust as the basis for risky engagements. When trust is derived by interactions from past experiences, we see the following main parameters:

- the relation between the number of positive and negative evidences
- the total number of evidences
- the derived expectation value about the trustworthiness in the next interaction

There may be further parameters, but for not overloading the interface we focus on this small set. For the calculation of the expectation value, the user's attitude and expectations about the general behaviour of interaction partners in this context is a further relevant aspect. At last, when trust is not only based on direct experiences, but also on indirect ones, the sources of those recommendations become important. But this is beyond the scope of this paper. In general, we believe that it is necessary to have a trust model, which has a number representation based on a continuous domain for the interpretation and manipulation by software agents (as most of the models in section 3), and a graphical user interface which allows intuitive access to the trust information by users. In the remainder of this section, we present three representations which try to integrate with these aspects.

CertainTrust – Human Trust Interface

As we focus only on the usability of the representation of the trust model by humans, we do not present further details on the computational model – how recommendations are processed – and the exact mapping between evidences and opinions about trustworthiness (for details see [11]).

The graphical representation of opinions about trustworthiness of an interaction partner (short: opinions) for CertainTrust (CT) is called Human Trust Interface (HTI). It is basically a two-dimensional interface, with the parameters *trust (value)* and *certainty (value)*.

- *trust value t*:
  - meaning: How good are the evidences from past interactions (and recommendations) in average?
  - calculated as the relative frequency of positive interactions (in relation to the total amount of interactions)
  - trust value = 0: up to now there have been only bad interactions (very untrusted)
  - trust value = 1: up to now there have been only good interactions (very trusted)
- *certainty value c*:
  - meaning: How certain is the expectation?
  - increases with the ratio between the number of collected evidences (#ce) and the number of expected evidences (#ee) (until #ce = #ee)
  - certainty = 0: not any evidences available (uninformed / low certainty)
  - certainty = 1: the number of collected evidences is equal to the number of expected evidences

An additional assumption for this representation is that the users can define a number of maximal expected evidences per context (e.g., mp3-exchange) (see the number "20" in Figure 23). If the number of collected evidences is equal to the number of maximal expected evidences, the certainty value reaches its maximum, which means the user would expect the trust value to be a good estimate for the trustworthiness in the next interaction. In our user study we evaluated whether this assumption does reflect reality. The existence of such a number was supported by answers in our user study (see section: Discussion).

If there is less information available than the maximal number of expected evidences, the trust value based on this information is not expected to be representative and needs to be biased according to the users' preferences.

We propose to calculate the expectation value $E(t,c)$ as $E(t,c) = t * c + (1-c) * f$, where f depends on the users preferences or dispositional trust. In Figure 23, as well as in user study, we used the value $f = 0.5$. This corresponds to expecting unknown interaction partners to be trustworthy in 50% of the interactions. As long as the number of maximal expected evidences is in the range of 10 to 20, this also produces similar results as using the mean of a beta distribution $f(p|a,b)$ with a = 'number of positive evidences' +1 and b = 'number of negative evidences' +1. But when the number of maximal expected evidences is increased our expectation value gives more weight to the certainty of an opinion. In Figure 23 the expectation value is integrated via a red-yellow-green color-gradient: E(1,1) = 1 ⇔ green ⇔ "trustworthy, E(0.5,0) = 0.5 ⇔ yellow ⇔ undecided, E(0, 1) = 0 ⇔ red ⇔ untrustworthy.

If CertainTrust is to autonomously choose between two interaction partners, the choice would be for the one with the higher expectation value; in the case of equal expectation values for the one with higher certainty.

Example: The point "A" in Figure 23 represents an opinion which a based on a relative high amount of evidences, which in average have been very positive. For the next interaction, CertainTrust would suggest that the interaction candidate is "very trustworthy".

Subjective Logic (SL) – Opinion Triangle

The second representation uses a triple of parameters to represent opinions about the trustworthiness of an interactor (b, d, u) (b = *belief*, d = *disbelief*, u = *uncertainty*). For details on the model see [6]. The relation between these parameters and the collected evidences can be explained as follows:

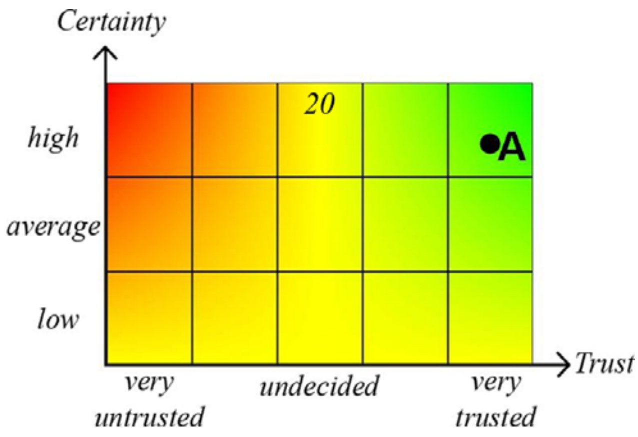

Figure 2: Example for CertainTrust - HTI

- *uncertainty*: depends on the number of collected evidences
  - uncertainty = 0: infinite number of evidences collected.
  - uncertainty = 1: not any evidence collected.
- *belief*: increases with the relative frequency of collected positive evidences:
  - belief = 1: infinite number of evidences collected, which have all been positive
  - belief = 0: not any positive evidence collected (only negative once, if any)
- *disbelief*: increases with the relative frequency of collected negative evidences:
  - disbelief = 1: infinite number of evidences collected, which all have been negative
  - disbelief = 0: not any negative evidence collected (only positive once, if any)

The axes for belief, disbelief, and uncertainty are indicated by the corresponding labels at the end of each axis. (see Figure 4). The interpretation of the opinion represented by A in Figure 4 is similar to the one explained in the example presented with CertainTrust.

This interface does not integrate a representation of the ex-

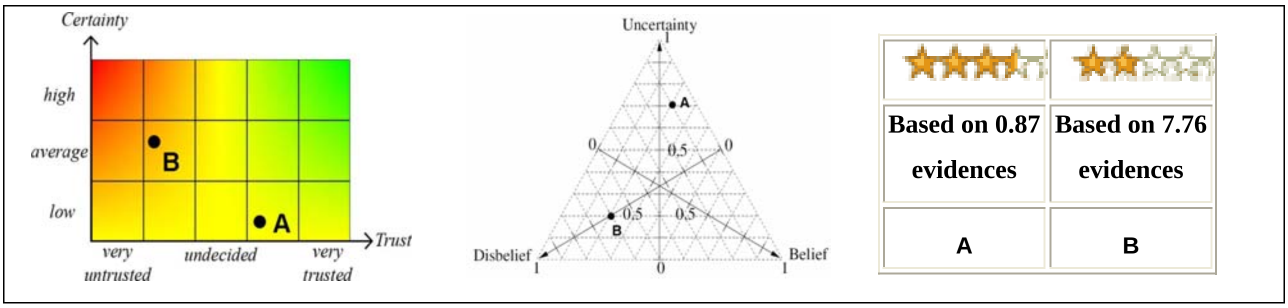

Figure 3: Sample of a setting: Each showing the opinion about the trustworthiness of 2 interactors based on the same evidences

pectation value in contrast to the interface of CertainTrust. Furthermore, the authors of this interface had not introduced a concept similar to the maximal number of expected evidences, but a static relation between collected evidences and uncertainty: u = 2 / (# pos. evid + # neg. evid. +2). As a result, uncertainty is below 0.1 in case there have been collected more than 18 evidences. In these cases all opinions are represented in the lower 10 percent of user interface. The uncertainty for an opinion based on 50 evidences is u = 0.0385, for an opinion based on 100 evidences it calculates to u = 0.0196. Due to this little differences in the uncertainty parameter and restricted resolution of the displays, the users will hardly be able to recognize which opinion is based on more evidences.

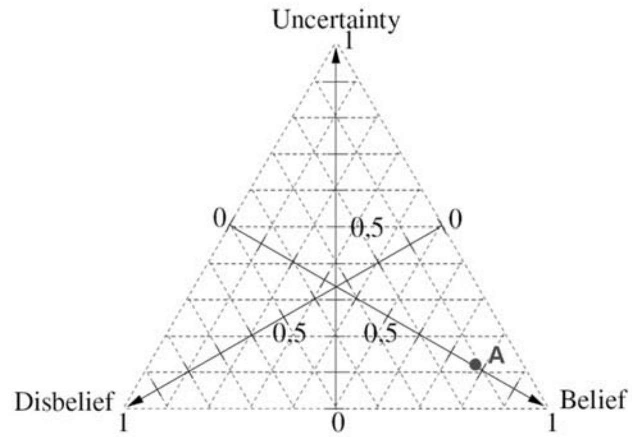

Figure 4: Example for Opinion Triangle (SL)

The Stars Representation

The interface "Stars" represents the average number of positive ratings of interactions as stars in the range of one to five stars (in half steps). Additionally, the interface shows the number evidences which contributed to the rating. For an example and interpretation of the stars see Figure 5 and Figure 6. The decimal places in the example of Figure 5 may originate from weighted evidences from recommendations.

Figure 5: Example for the Stars interface

In the end, this rating does not directly express the trustworthiness of an interaction partner, but it represent an average of the ratings, which have been given before. Thus, it can be support decisions as the approaches introduced above. A difference between the representations introduced before is that this interface does not integrate both values (stars and number of evidences) to a single value or point, but leaves the interpretation up to the user. This makes it harder to manipulate opinions, as both parameters have to be manipulated separately. As the Opinion Triangle, this representation does not integrate an expectation about the trustworthiness in the next interaction.

Figure 6: Interpretation of the stars is taken from Amazon (www.amazon.com)

USER STUDY

As pointed out above, we wanted to evaluate how intuitive users can interact with those representations. We conducted an experiment with three conditions (CT, SL, Stars), corresponding to the three interfaces described above.

This experiment allowed us to get insight into how good the interfaces do support the user in making a decision on the trustworthiness of a potential interaction partner.

We were also interested in how far the decision that would be taken by CertainTrust automatically matches the decisions the users took. A good prediction of the user's choice could make user interaction superfluous in cases as presented in the scenario.

We did not evaluate how the interfaces perform with regard to manipulating opinions, as this was not feasible with the Stars interface. We did include the Stars interface as a baseline as it is already widely used in internet sites on the web.

Specifically, our hypotheses for the experiment were as follows:

Hypothesis 1: The users will be faster in the CT Interface than in the Opinion Triangle (from Subjective Logic) to decide on the trustworthiness of interaction partners. In the Stars interface, we expected the user to be faster, as they were used to this interface.

Hypothesis 2: The participants' decisions in the Stars and SL interface will be the same as the decision suggested by CertainTrust.

Design

The experiment used a within-subject factorial design, with the interface type as primary factor. Every participant completed a series of eight tasks with every interface. The task was to pick one of two potential interaction partners which were displayed in the same interface. The same eight pairs in the same order were used for all conditions; however we do not expect any carry over effects, as the representations are sufficiently distinctive. We did however control for learning effects of the general setup by counterbalancing the order of interfaces between subjects. The within subject design does account for variability between subjects.

The study was conducted as an online survey, so that participants could take part remotely to prevent experimenter effects.

Participants

Thirty two subjects took part in the study. All but one did have prior experience with online shopping. 4 females and 27 males completed the study (one participant did not reveal the gender). They were in the age range of 21 to 40 years. Participants were not paid for taking part in the study.

(I) Model	(J) Model	Mean Difference (I-J)	Std. Error	Sig .(a)	99% Confidence Interval for Difference (a)	
					Lower Bound	Upper Bound
Stars	SL	-27838.000*	7057.201	.001	-50280.028	-5395.972
	CT	1851.156	3044.874	1.000	-7831.598	11533.910
SL	Stars	27838.000*	7057.201	.001	5395.972	50280.028
	CT	29689.156*	6967.931	.001	7531.001	51847.302
CT	Stars	-1851.156	3044.874	1.000	-11533.910	7831.598
	SL	-29689.156*	6967.931	.001	-51847.302	-7531.011

Based on estimated marginal means
*. The mean difference is significant at the .01 level.
a. Adjustment for multiple comparisons: Bonferroni.

Table 1: Pairwise Comparisons (time in milliseconds)

Procedure

Participants were presented the same scenario as described in section 2. Afterwards they answered a questionnaire regarding their opinion about the relationship between trust and evidences.

The procedure for the experiment was as follows. The interface of the first condition was explained to the subjects. They were then sequentially presented with the 8 pairs of potential interaction partners (named A and B) using the same trust representation. The time they took to decide between the two marked by a click on a button and the interaction partner they preferred was logged for analysis.

The same was done for the other two conditions, so that every user judged on every pair of interaction partners three times, once in every interface. In Figure 2 there is an example how the same example (setting) is presented using different interfaces.

Setting	Evidences for Interactor A		Evidences for Interactor B	
	pos.	neg.	pos.	neg.
1	0.5903	0.2883	1.868	5.8968
2	1.9294	1.9294	3.8816	3.8816
3	9.8276	2.0319	0.9623	0.0385
4	8.7332	2.036	1.9467	8.8225
5	1.9089	1.9499	8.7332	2.036
6	2.0119	9.4476	0.4129	0.4328

Table 2: # Evidences per interaction partner and setting

For the evaluation we use only 6 of 8 settings, as the number of evidences presented in the three interfaces was identical for those (for the numbers see Table 2). As it hard for users to distinguish the uncertainty of opinions which are based on more than 20 evidences in the Opinion Triangle, we restricted ourselves to opinions based on less evidences.

Results

Hypothesis 1:

For the analysis of the first hypothesis we aggregated the time the users took for their decisions over the 6 settings per model (CT, SL, Stars). The mean times (in ms) per model are given in ascending order: 44635.094 (CT), 46486.250 (Stars), and 74324.250 (SL). The Kolmogorov-Smirnov test indicated that the data is normally distributed. For the further analysis we did one-way repeated measures ANOVA:

Mauchly's test indicated that the assumption of sphericity had been violated, $\chi^2(2)=24.0$), $p<.01$, therefore degrees of freedom were corrected using Greenhouse-Geisser estimates of sphericity ($\varepsilon=.645$). The results show that the choice of the representation has significantly affected the time a user needs to select an interaction partner, $F(1.3, 40.0)$, $p<.01$, $\omega=0.56$. Using the benchmarks for effect size this represents a strong effect.

Bonferroni post hoc tests revealed that the mean time of the participants was significantly higher than when using the Opinion Triangle of Subjective Logic as in both other interfaces. For details see Table 1.

Hypothesis 2:

As the selected interaction partner in an experiment leads to nominal, non-parametric data – a participant either selected to interact with A or B – we first counted the frequencies of the selection of A and B. (Figure 7 shows the corresponding the percentage). We can say that the majority of the participants has selected the same interaction partner independently of the model. Furthermore, the interaction partner which has been selected by the majority of the users was in all settings the one which would have been autonomously selected by CertainTrust. Thus, we consider this hypothesis to be true.

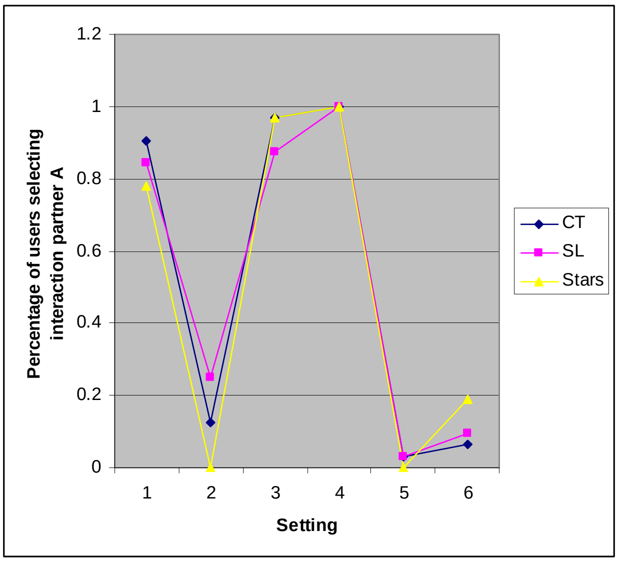

Figure 7: Percentage of participants selection interaction partner A per setting and per model (the lines are only for ease of reading)

For a further analysis, whether the representation has an effect on the choice of the interaction partner, we calculate the percentage of participants which have chosen the same interaction partner as the one proposed by CertainTrust per setting and per model. The results show that there are not any significant differences (p>.05) between the participants choices in the different models. The mean values are: .943 (CT), .891 (SL), and .927 (Stars); the confidence intervals are given in Figure 8.

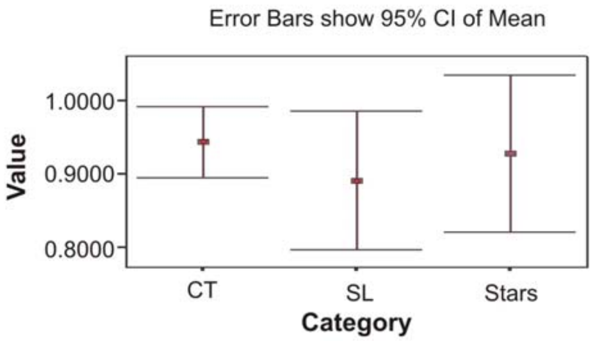

Figure 8: Mean values: average percentage of participants selecting the same interaction partner as proposed by CT.

DISCUSSION

Effects of the representation

We can learn from the results, that representation has a significant effect on the time the participants needed to select an interaction partner. Although, we expected most participants to be faster in the Stars interface as they are probably used to this representation (all, but one had experiences with online shopping), the participants were slightly, non-significantly slower than in the CertainTrust HTI. In both interfaces the participants were significantly faster than in the Opinion Triangle. We identified two reasons to explain these effects. First, we believe that users are not used to interpret triangular representations with three axes, as orthogonal two axis layouts are more common in everyday life. On the website on Subjective Logic (http://sky.fit.qut.edu.au/~josang/sl/demo/BV.html), we can find the statement that the Opinion Triangle is a more mathematical representation. That page offers further representations for opinions, but as to our knowledge the details have neither been evaluated nor published, we decided to use the representation which is usually used as graphical representation of opinions for Subjective Logic. We believe that the integration of trust and certainty in one graphical representation, as opposed to the Stars interface, has further benefit. Additionally, we tried to support the users' decisions with the green-yellow-red color-gradient. The answers by the participants on the questions at the end of the experiment also indicated that they felt comfortable with the interface of CertainTrust (small sample from the questions):

Question A: Which interface would you prefer in the described scenario: (CT 68.8%, SL 3.1%, Stars 28.1%)

Question B: The color-gradient supported your decision for an interaction partner. Do you agree? The mean value of the answers on a scale from 0 (strongly disagree), 1 (disagree), 2 (slightly disagree), 3 (slightly agree), 4 (agree), 5 (strongly agree) is 3.81.

Question C: Can you imagine using the Human Trust Interface without the labels after a short time of familiarization? mean value (scale as above): 3.19

Rationality of the choices

At this point, we have still to discuss, if the decisions of the participants were reasonable. The definition of the "best choice" is not trivial, as the choice is influenced by the relative frequency of positive interactions, as well as the total number of evidences, and furthermore the risk which is associated with an interaction. In settings 1, 3, and 5 the decision goes along with the relative frequency of positive and negative interactions and the amount of collected evidences. The more interesting cases are the settings 2, 4, and 6. For the exact numbers of evidences see Table 2.

In setting 2, there are two candidates for the interaction having the same of ratio of positive to total evidences, but the opinion about B is based on more evidences. Using the Stars interface (3 stars for both) all users decided to interact with B – which is the candidate about whom more information is available. In both remaining interfaces, a small number of participants selected the candidate about whom less information is available. In both cases, the expected trustworthiness calculated by CertainTrust is 0.5. Although in this case, the expectation values are identical, the selection algorithm favours the one with the higher number of evidences; thus, going along with the majority of the participants.

In setting 4, there is candidate A about whom we have very little, but very positive information (mean: 0.96) information; and candidate B about whom we have higher amount of information, mostly positive (mean: 0.81). In all three interfaces the majority of the participants selected user B. In the face that interactions are associated with a certain risk and a little amount of evidences may be quite misleading, we believe this choice is rational.

In setting 6, there are six participants who selected candidate A in the Stars interface, while in the other interfaces only one (CT) or two (SL) participants selected A. In short, about candidate A, there is very little information available (similar amount of positive and negative evidences), while about candidate B, there is more information available, which is mostly negative. In this case, we can say that the participants made a better choice in the CertainTrust and Subjective Logic interface, than in the Stars interface.

At last, we can ask, if there is a number of maximal expected evidences, which allows to interpret the trust value as representative expectation for future interactions. At the beginning of the experiment, after the participants were given the scenario and before they were introduced to the user interfaces, we presented the participants this statement: “Having collected a certain number of evidences, you are able to properly estimate the trustworthiness of your interaction partner.” Most users agreed (all, but 5). On a scale from 0 to 5 (as above) the mean was 3.41. For the mp3-exchange the majority of users expected 6-10 evidences, for buying the mp3-player (used, 10€) the majority expected 21-50 evidences.

CONCLUSION

We have presented a representation of trust, which allows for representing the main parameters – total amount of evidences, relative frequency of positive evidences, and expectation value – in an integrated way. Furthermore, we have shown that users can intuitively interpret our interface, and make good decisions. A dispositional trust component has already been integrated in the representation, and will simply affect the color-gradient, which expresses the expectation value. As orthogonal two-axis layouts are well-known, we believe that users also will be able to easily manipulate opinions. In the next steps the integration of the sources of recommendations, as well as the risk, which is associated to with an engagement, will be evaluated.

ACKNOWLEDGEMENTS

The authors’ work was supported by the German National Science Foundation (DFG) as part of the PhD program "Enabling Technologies for Electronic Commerce" (first author) and by SAP CEC Darmstadt in the AUGUR project (second author).

REFERENCES

1. A. Abdul-Rahman and S. Hailes. Supporting trust in virtual communities. In *Proceedings of Hawaii International Conference on System Sciences*, 2000.
2. S. Buchegger and J.-Y. Le Boudec. A Robust Reputation System for Peer-to-Peer and Mobile Ad-hoc Networks. In *P2PEcon*, 2004
3. J. Golbeck. *Computing and Applying Trust in Web-Based Social Networks*. PhD thesis, University of Maryland, College Park, 2005.
4. A. Heinemann. *Collaboration in Opportunistic Networks*. PhD thesis, Technische Universität Darmstadt, 2007.
5. P. Hui, A. Chaintreau, J. Scott, R. Gass, J. Crowcroft, and C. Diot. Pocket switched networks and human mobility in conference environments. In WDTN ’05*: Proceeding of the 2005 ACM SIGCOMM workshop on Delay-tolerant networking*, pages 244–251, New York, NY, USA, 2005. ACM Press.
6. A. Jøsang. A logic for uncertain probabilities. International Journal of Uncertainty, *Fuzziness and Knowledge-Based Systems*, 9(3):279–212, 2001.
7. A. Jøsang and R. Ismail. The beta reputation system. In *Proceedings of the 15th Bled Conference on Electronic Commerce*, 2002.
8. A. Jøsang and R. Ismail, and C. Boyd. A survey of trust and reputation systems for online service provision. In *Decision Support Systems*, 43(2), 2007
9. L. Mui, M. Mohtashemi, and A. Halberstadt. A computational model of trust and reputation for e-businesses. In *Proceedings of the 35th Annual Hawaii International Conference on System Sciences (HICSS’02)*-Volume 7, Washington, DC, USA, 2002. IEEE Computer Society.
10. D. Quercia, S. Hailes, and L. Capra. B-Trust: Bayesian trust framework for pervasive computing. In K. Stølen, W. H. Winsborough, F. Martinelli, and F. Massacci, editors, *4th International Conference on Trust Management* (iTrust), volume 3986 of Lecture Notes in Computer Science, pages 298–312. Springer, 2006.
11. S. Ries. CertainTrust: A trust model for users and agents. In *Proceedings of the 2007 ACM Symposium on Applied Computing*, pages 1599 – 1604. ACM, 2007.
12. J. Sabater and C. Sierra. Reputation and social network analysis in multi-agent systems. *In Proceedings of the 1st International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 475–482, New York, NY, USA, 2002. ACM Press.
13. W. T. Teacy, J. Patel, N.R. Jennings, and M. Luck. TRAVOS: Trust and reputation in the context of inaccurate information sources. *Autonomous Agents and Multi-Agent Systems*, 2006