A Trust Framework for Social Participatory Sensing Systems

Haleh Amintoosi and Salil S. Kanhere

The University of New South Wales, Sydney, Australia
{haleha,salilk}@cse.unsw.edu.au

Abstract. The integration of participatory sensing with online social networks affords an effective means to generate a critical mass of participants, which is essential for the success of this new and exciting paradigm. An equally important issue is ascertaining the quality of the contributions made by the participants. In this paper, we propose an application-agnostic trust framework for social participatory sensing. Our framework not only considers an objective estimate of the quality of the raw readings contributed but also incorporates a measure of trust of the user within the social network. We adopt a fuzzy logic based approach to combine the associated metrics to arrive at a final trust score. Extensive simulations demonstrate the efficacy of our framework.

Keywords: trust framework, participatory sensing, online social networks, data quality, urban sensing, fuzzy logic

1 Introduction

The rapid improvement in mobile phone technology, in terms of storage, processing and sensing, has resulted in the emergence of a novel paradigm called participatory sensing [1]. The core idea is to empower ordinary citizens to collect and contribute sensor data (e.g., images, sound, etc) from their surrounding environment. This new paradigm has been effectively used to crowdsource information about road conditions [2], noise pollution [3], diet [4] and price auditing [5].

For participatory sensing to be a success, a key challenge is the recruitment of sufficient volunteers. Typically there is no explicit incentive for participation and people contribute altruistically. In the absence of adequate contributors, the application will very likely fail to gather meaningful data. Another challenge, particularly for tasks which require domain-specific knowledge (e.g., takings photos of rare plant species), is the suitability of the participants for the task at hand [6].

One potential solution to address these challenges is to leverage online social networks as an underlying publish-subscribe infrastructure for distributing tasks and recruiting suitable volunteers [7,8]. This new paradigm, referred to as *social participatory sensing*, offers the following advantages. First, it makes it easier to identify and select well-suited participants based on the information available in their public profile (e.g., interests, educational background, profession, etc). Second, social ties can motivate participants to contribute to tasks initiated by

K. Zheng, M. Li, and H. Jiang (Eds.): MOBIQUITOUS 2012, LNICST 120, pp. 237–249, 2013.

friends. Third, incentives can be offered in the form of reputation points or e-coins [9] and published on the contributors' profile. A real-world instantiation of social participatory sensing was recently presented in [10], wherein, Twitter was used as the underlying social network substrate.

The inherent openness of participatory sensing, while valuable for encouraging participation, also makes it easy for propagation of erroneous and untrusted contributions. When combined with social networks, other trust issues arise. People normally have more trust on contributions provided by their close friends than casual acquaintances, since interactions with close friends provides more emotional and informational support [11]. In particular, when data of the same quality is available from two social network contacts, one a close friend and the other a casual acquaintance, it is natural human tendency to put more credence in the data from the close friend. Hence, in social participatory sensing, it is crucial to consider both, the participant's social trust and the data quality as influential aspects in evaluating the trustworthiness of contributions. While there exist works that address the issue of data trustworthiness in participatory sensing (see Section.2), they do not provide means to include social trust and as such cannot be readily adopted for social participatory sensing.

In this paper, we present an application agnostic framework to evaluate trust in social participatory sensing systems. Our system independently assesses the quality of the data and the trustworthiness of the participants and combines these metrics using fuzzy logic to arrive at a comprehensive trust rating for each contribution. By adopting a fuzzy approach, our proposed system is able to concretely quantify uncertain and imprecise information, such as trust, which is normally expressed by linguistic terms rather than numerical values. We undertake extensive simulations to demonstrate the effectiveness of our trust framework and benchmark against the state-of-the-art. The results demonstrate that considering social relations makes trust evaluation more realistic, as it resembles human behaviour in establishing trustful social communications. We also show that our framework is able to quickly adapt to rapid changes in the participant's behaviour (transitioning from high to low quality contributions) by fast and correct detection and revocation of unreliable contributions. Moreover, we find that leveraging fuzzy logic provides considerable flexibility in combining the underlying components which leads to a better assessment of the trustworthiness of contributions. Our framework results in a considerable increases in the overall trust over a method which solely associates trust based on the quality of contribution.

The rest of the paper is organised as follows. Related work is discussed in Section 2. We present the details of our fuzzy system in Section 3. Simulation results are discussed in Section 4. Finally, Section 5 concludes the paper.

2 Related Work

To the best of our knowledge, the issue of trust in social participatory sensing hasn't been addressed in prior work. As such, we discuss about related research focussing on trust issues in participatory sensing.

In a participatory sensing system, trustworthiness can be viewed as the quality of the sensed data. In order to ascertain the data trustworthiness, it is highly desirable to ascertain that the sensor data has been captured from the said location and at the said time. [12] has proposed a secure service which allows participants to tag their content with a spatial timestamp indicating its physical location, which is later used by a co-located infrastructure for verification. A similar approach has been proposed in [13], in the form of a small piece of metadata issued by a wireless infrastructure which offers a timestamped signed location proof. Since these works rely on external infrastructure, they have limited scalability. Moreover, neither approach will work in situations where the infrastructure is not installed. In our proposed framework, we assume that sensor data is tagged with GPS coordinates/system time before being stored in phone memory, which is then used by trust server for verification. Data trustworthiness has been investigated from another point of view which tries to confirm that uploaded data preserves the characteristics of the original sensed data and has not been changed unintentionally or maliciously. In particular, there are several works which make use of Trusted Platform Module (TPM)[14], which is a micro-controller embedded in the mobile device and provides it with hardware-based cryptography as well as secure storage for sensitive credentials. In [15], each device has a trusted hardware element that implements cryptographic algorithms for content protection. [16] presents two TPM-based design alternatives: the first architecture relies on a piece of trusted code and the second design incorporates trusted computing primitives into sensors to enable them sign their readings. However, TPM chips are yet to be widely adopted in mobile devices. There is also recent work that does not require TPM. [17] proposes a reputation-based framework which makes use of Beta reputation [18] to assign a reputation score to each sensor node in a wireless sensor network. Beta reputation has simple updating rules as well as facilitates easy integration of ageing. However, it is less aggressive in penalizing users with poor quality contributions. A reputation framework for participatory sensing was proposed in [19]. A watchdog module computes a cooperative rating for each device according to its short-term behaviour which acts as input to the reputation module which utilizes Gompertz function [20] to build a long-term reputation score. Their results show an improvement over the non-trust aggregation based approaches and Beta reputation system. However, the parameters related to the participants' social accountability have not been considered. As such, their system cannot be readily used in our context.

3 Fuzzy Trust Framework

In this section, we explain the proposed framework for evaluating trust in social participatory sensing system. An overview of the architecture is presented in Section 3.1 followed by a detailed discussion of each component in Section 3.2.

3.1 Framework Architecture

Since our framework attempts to mimic how human's perceive trust, we first present a simple illustrative example. Suppose John is a member of an online social network (e.g., Facebook). He has made a profile and has friended several people. John is a vegetarian. He is also on a budget and is keen to spend the least possible amount for his weekly groceries. He decides to leverage his social circle to find out the cheapest stores where he can buy vegetarian products. Specifically, he asks his friends to capture geotagged photos of price labels of vegetarian food items when they are out shopping and send these back to him. One of his friends, Alex decides to help out and provides him with several photos of price labels. In order to decide whether to rely on Alex's contributions, John would naturally take into account two aspects: (i) his personal trust perception of Alex, which would depend on various aspects such as the nature of friendship (close vs. distant), Alex's awareness of vegetarian foods, Alex's location, etc and (ii) the quality of Alex's data which would depend on the quality of the pictures, relevance of products, etc. In other words, John in his mind computes a trust rating for Alex's contribution based on these two aspects. Our proposed trust framework provides a means to obtain such trust ratings by mimicking an approach similar to John's perception of trustworthiness in a scalable and automated manner. This trust rating helps John to select trustable contributions and accordingly plan for his weekend shopping. Our framework also affords a list of trustable friends for the data consumer (e.g., John) for future recruitment.

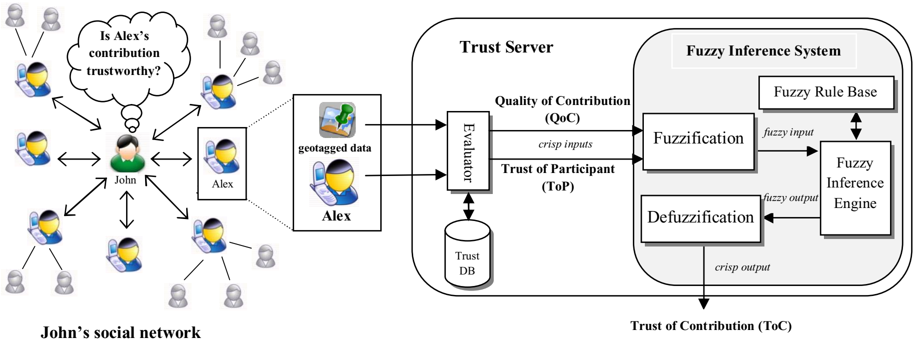

Fig. 1. Trust framework architecture

Fig. 1 illustrates the architecture of the proposed fuzzy framework. The social network serves as the underlying publish-subscribe substrate for recruiting friends as participants. In fact, the basic participatory sensing procedures (i.e., task distribution and uploading contributions) are performed by utilizing the social network communication primitives. A person wishing to start a participatory sensing campaign disseminates the tasks to his friends via email, message or by writing as a post on their profiles (e.g., Facebook wall). Friends upload their contributions via email or in the form of a message. We can also benefit from group construction facilities in Facebook or circles in Google Plus. The contributions received in response to a campaign are transferred (e.g., by using Facebook Graph API[1]) to a third party trust server, which incorporates the proposed fuzzy inference system and arrives at a trust rating for each contribution. This cumulative trust rating can be used as a criterion to accept/reject the contribution by comparing against a predefined threshold. Alternately, the ratings can be used as weights for computing summary statistics. Finally, ordered list of contributions according to their ToCs, or of participants according to their ToPs can be generated. ToP is also updated based on the quality of contributions. If below a specified threshold, participant's trust will be decremented by α; otherwise it will be incremented by β. Note that $\alpha>\beta$; since in typical social relations, trust in others is built up *gradually* after several trustworthy communications and torn down *rapidly* if dishonest behaviour is observed.

3.2 Framework Components

This section provides a detailed explanation of the framework components. In particular we focus on the trust sever and fuzzy inference system.

Trust Server. The trust server is responsible for maintaining and evaluating a comprehensive trust rating for each contribution. As discussed in Section 1, there are two aspects that need to be considered: (1) quality of contribution and (2) trust of participant. The server maintains a trust database, which contains the required information about participants and the history of their past contributions. When a contribution is received by the trust server, the effective parameters that contribute to the two aforementioned components are evaluated by the Evaluator and then combined to arrive at a single quantitative value for each. The two measures serve as inputs for the fuzzy inference system, which computes the final trust rating. In the following, we present a brief discussion about the underlying parameters and the evaluation methods.

Quality of Contribution (QoC)

In participatory sensing, contributions can be any form such as images or sounds. The quality of the data is affected not only but fidelity of the embedded sensor but also the sensing action initiated by the participant. The in-built sensors in mobile devices can vary significantly in precision. Moreover, they may not be correctly calibrated or even worse not functioning correctly, thus providing erroneous data. Participants may also use the sensors improperly while collecting data,(e.g., not focussing on the target when capturing images). Moreover, human-as-sensor applications such as weather radar in [10] are exposed to variability in the data quality due to subjectivity. For example, what is hot for one person may be comfortable for another. In order to quantify QoC, a group of parameters must be evaluated such as: relevance to the campaign (e.g., groceries in the above example), ability in determining a particular feature (e.g., price

[1] `http://developers.facebook.com/docs/reference/api/`

tag), fulfilment of task requirements (e.g., specified diet restrictions), etc. There already exists research that has proposed methods for evaluating the quality of data in participatory sensing. Examples include image processing algorithms proposed in [4] and outlier detection [23] for sound-based sensing tasks. Rather than reinventing the wheel, our system relies on the state-of-the-art methods for this evaluation. The result is a single value for QoC in the range of [0, 100].

Trust of Participant (ToP)

ToP is a combination of personal and social factors. Personal factors consist of the following parameters:

Expertise(E): It is defined as the measure of a participant's knowledge and is particularly important in tasks that require domain expertise. Greater credence is placed in contributions made by a participant who has expertise in the campaign. We employ expert finding systems for evaluating expertise. These systems employ social networks analysis and natural language processing (text mining, text classification, and semantic text similarity methods) to analyse explicit information such as public profile data and group memberships as well as implicit information such as textual posts to extract user interests and fields of expertise [21]. Dmoz[2] open directory project can be used for expertise classification. Expertise evaluation is done by incorporating text similarity analysis to find a match between the task keywords (e.g., vegetarian) and participant's expertise.

Timeliness(T): Timeliness measures how promptly a participant performs prescribed tasks. It depends on the time taken to perform the task (t) and the task deadline (d). To evaluate this parameter, inverse Gompertz function defined as $T(t) = 1 - e^{-be^{-ct}}$ can be used because of its compatible with timeliness evolution. In the original inverse Gompertz function, the lower asymptote is zero; it means that the curve approaches to zero in infinity. In our case, timeliness rate will only be zero if contribution is received after the deadline; otherwise, a value between x and 1 is assigned to it. It means that the lowest timeliness rating will be x if contribution is received before the deadline, and is zero if received after the deadline. So, we modify the function as Eq.1 to calculate the timeliness(T):

$$T(t) = \begin{cases} 1 - [(1-x)e^{-be^{-ct}}] & \text{if } t < d \\ 0 & \text{otherwise} \end{cases} \tag{1}$$

Locality(L): Another significant parameter is locality, which is a measure of the participant's familiarity with the region where the task is to be performed. We argue that contributions received from people with high locality to the tasking region are more trustable than those received from participants who are not local, since the first group is more acquainted with and has better understanding of that region. According to the experimental results presented in [22], people tend to perform tasks that are near to their home or work place (places that they are considered 'local' to them). This implies that if we log the location of participants' contributions, we can estimate their locality. A participant's locality would be highest at locations from where they make maximum number

[2] `http://www.dmoz.org`

of contributions. In order to evaluate locality, we assume that the sensing area has been divided to n regions, and a vector V with the length equal to n is defined for each participant, where, *V(i)* is number of samples collected in region i. In this case, locality of a participant to region i is calculated by Eq. 2:

$$L(i) = V(i)/\sum_{i=0}^{n-1} V(i) \tag{2}$$

Next, we explain the social factors that affect ToP:

Friendship duration(F): In real as well as virtual communications, long lasting friendship relations normally translate to greater trust between two friends. So, friendship duration which is an estimation of friendship length is a prominent parameter in trust development. We use the Gompertz function to quantify friendship duration, since its shape is a perfect match for how friendships evolve. Slow growth at start resembles the friendship gestation stage. This is followed by a period of accumulation where the relationship strengthens culminating in a steady stage. As such, the friendship duration is evaluated according to Eq. 3, in which, b and c are system-defined constants and t is the time in months.

$$F(t) = e^{-be^{-ct}} \tag{3}$$

Interaction time gap(I): In every friendship relation, interactions happen in form of sending requests and receiving responses. Interaction time gap, measures the time between the consequent interactions and is a good indicator of the strength of friendship ties. If two individuals interact frequently, then it implies that they share a strong relationship, which translates to greater trust. We propose to use the inverse Gompertz function shown in Eq. 4, to quantify the interaction time gap, where, b and c are system-defined constants and t is the time in months.

$$I(t) = 1 - e^{-be^{-ct}} \tag{4}$$

The aforementioned parameters are combined by the Evaluator to arrive at a single value for ToP, as follows, $ToP = w_1 \times E + w_2 \times T + w_3 \times L + w_4 \times F + w_5 \times I$, where, w_i is the application specific weight of each parameter.

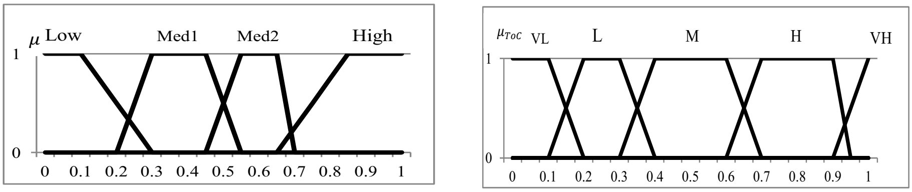

(a) Membership function for QoC and ToP

(b) Membership function for ToC

Fig. 2. Membership functions of input and output linguistic variables

Fuzzy Inference System. Our proposed framework employs fuzzy logic to calculate a comprehensive trust rating for each contribution, referred to as the Trust of Contribution (ToC). We cover all possible combinations of trust aspects and address them by leveraging fuzzy logic in mimicking the human decision-making process. The inputs to the fuzzy inference system are the crisp values of QoC and ToP. In the following, we describe the fuzzy inference system components. *Fuzzifier:* The fuzzifier converts the crisp values of input parameters into a linguistic variable according to their membership functions. In other words, it determines the degree to which these inputs belong to each of the corresponding fuzzy sets. The fuzzy sets for QoC, ToP and ToC are defined as:

T(QoC)=T(ToP)={Low, Med1, Med2, High}, T(ToC)= { VL, L, M, H, VH}

Fig.2(a) represents the membership function of QoC and ToP and Fig.2(b) depicts the ToC membership function. We used trapezoidal shaped membership functions since they provide adequate representation of the expert knowledge, and at the same time, significantly simplify the process of computation.

Inference Engine: The role of inference engine is to convert fuzzy inputs (QoC and ToP) to the fuzzy output (ToC) by leveraging If-Then type fuzzy rules. The combination of the above mentioned fuzzy sets create 4*4=16 different states which have been addressed by 16 fuzzy rules as shown in Table.1. Fuzzy rules help in describing how we balance the various trust aspects. The rule base design is based on our experience and beliefs on how the system should work. To define the output zone, we used *max-min* composition method as: $\mu_{T(ToC)}(ToC) = max[\underset{\substack{X\in T(ToP),\\ Y\in T(QoC)}}{min}(\mu_X(ToP), \mu_Y(QoC))]$, where $\mu_A(x)$ denotes the degree of x's membership to a fuzzy set A. The result of the inference engine is the ToC which is a linguistic fuzzy value.

Table 1. Fuzzy rule base for defining ToC according to QoC and ToP

Rule no.	if QoC	and ToP	Then ToC	Rule no.	if QoC	and ToP	Then ToC
1	Low	Low	VL	9	Med2	Low	M
2	Low	Med1	L	10	Med2	Med1	H
3	Low	Med2	L	11	Med2	Med2	H
4	Low	High	M	12	Med2	High	H
5	Med1	Low	L	14	High	Low	H
6	Med1	Med1	L	14	High	Med1	H
7	Med1	Med2	M	15	High	Med2	VH
8	Med1	High	M	16	High	High	VH

Defuzzifier: A defuzzifier converts the ToC fuzzy value to a crisp value in range of [0, 1] by employing the Centre of Gravity method (COG) [24], which computes the center of gravity of the area under ToC membership function.

To summarize, once a campaign is launched, participants begin to send a series of contributions. For each contribution, the Evaluator computes a value for QoC and ToP. These values are fed to fuzzy inference engine which calculates ToC for that contribution. The server utilizes the ToC to provide useful statistical results. For example, only contributions with a TOC greater than a certain threshold could be considered as trustable. Moreover, the ToP values could be used to select a list of trustable candidates for recruitment in future campaigns.

4 Experimental Evaluation

This section presents simulation-based evaluation of the proposed trust system. The simulation setup is outlined in Section 4.1 and the results are in Section 4.2.

4.1 Simulation Setup

To undertake the preliminary evaluations outlined herein, we chose to conduct simulations, since real experiments in social participatory sensing are difficult to organise. Simulations afford a controlled environment where we can carefully vary certain parameters and observe the impact on the system performance. We developed a custom Java simulator for this purpose. We simulate an online social network where 50 members participate in 200 campaigns, producing one contribution for each. In the ideal case, for each contribution, we would have computed the value of each of the underlying parameters discussed in Section 3.2 based on some typical probabilistic distributions. However, this would make the simulations quite complicated. Moreover, this exercise would digress from the primary objective of the evaluations: to evaluate if social trust is a useful contributor to the overall trust in social participatory sensing. For the sake of simplicity, we therefore, assign a random value of ToP to each participant and a random value of QoC for each contribution, both in the range of [0, 100], based on criteria specific to the scenarios and leave extra investigation for future work.

Recall that, the goal of the trust framework is to assign a trust rating to each contribution which is further used as a criterion to accept/reject the contribution. As such, in the evaluations, we artificially create circumstances in which, some participants contribute poor quality data for a certain number of campaigns. We want to investigate if our trust framework is able to identify this behaviour and revoke untrusted contributions in a robust manner. In order to create all possible combinations of QoC and ToP, we assume that participants belong to one of the following four categories, each of which resembles one type of friend in a typical social participatory sensing system:

Category 1: Participants with high ToP (ToP$\geq$50) and high QoC (QoC$\geq$50).
Category 2: Participants with low ToP (ToP<50) but high QoC (QoC$\geq$50).
Category 3: Participants with high ToP (ToP$\geq$50) but low QoC (QoC<50).
Category 4: Participants with low ToP (ToP<50) and low QoC (QoC<50).

The threshold 50 used above for a trustworthy participant/contribution has also been used previously in [19,25]. Friends that belong to Category 1 would generally be more willing to volunteer and contribute data. As such, we assume that Category 1 contains more participants (20), in comparison with the other 3 categories, which contain 10 participants each. In the first scenario, we assume that participants do not alter their behaviour and thus QoCs follow the category settings throughout the entire simulation. In the second scenario, we assume that participants can transition from one category to another (details in Section 4.2).

As mentioned in Section 3, a ToC rating is calculated for each contribution and those with ToC lower than a predefined threshold are revoked from further calculations. The ToCs for the non-revoked contributions are then combined to form an overall trust for that campaign. In other words, $OverallTrust = \frac{\sum_{i=1}^{n} ToC}{n}$ in which, n is the number of non-revoked contributions. ToPs are also updated based related QoCs. We consider the overall trust as the evaluation metric. The greater the overall trust the better the ability of the system to revoke untrusted contributions. Overall trust has a value in the range of [0, 100].

We compare the performance of our framework against a baseline system, which only considers QoC for evaluating the trust of each contribution. In order to study the effect of other trust aspects, we incrementally add them to the baseline to see how considering each aspect influences trust. Specifically, we compare the following: (1) Baseline: where $ToC = QoC$ (2) Baseline-Rep: which follows the approach in [19] by calculating a reputation score for each participant according to the QoC of his successive contributions. This reputation score is used as a weight for QoC. In other words, $ToC = \sqrt{Rep * QoC}$ (3) Average: which includes ToP but computes the ToC simply as an average of ToP and QoC (4) Fuzzy: our proposed framework.

The revocation threshold is set to 50. Recall that, when ToP is updated, it is decremented by α if the QoC is below a threshold; otherwise it is incremented by β. We set the QoC threshold to 50 and α and β to 2 and 1, respectively.

4.2 Simulation Results

We first present the simulation results for the first scenario. Figure 3 depicts the evolution of the average overall trust as a function of the number of campaigns. As shown in the figure, our fuzzy trust method outperforms all the other methods. This confirms its success in mimicking the human trust establishing process by correct settings of fuzzy rules. In particular, we have set the rules in a way that results in early detection and severe punishment of untrusted contributions and also put greater emphasis on highly trusted contributions. The former has been done by assigning a very low(VL) value to ToC in case of low ToP and QoC (i.e., Rule no. 1 in Table. 1), whereas the latter has been obtained through assigning very high(VH) value to ToC in case of high QoC and above average ToP (i.e., Rule no. 15 and 16 in Table. 1).

Fig. 5 depicts two ordered lists provided by trust server. The first list sorts the participants in a descending order of their ToPs. This can be used as a suggestion list for data consumer for future recruitment of participants. The second list provides an ordered list of contributions according to the descending order of ToCs, which can help the data consumer to select the most trustable contributions based on a certain configurable threshold.

Next, we present results for the second scenario, wherein, the behaviour of the participants can change with time, which may result in a transition from one category to another. This scenario allows us to observe the performance of the schemes in the presence of noise. For example, consider a participant who is initially highly trusted and provides high quality data and thus belongs to

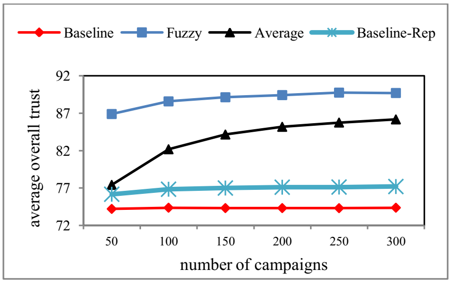

Fig. 3. Evolution of average overall trust for all methods, Scenario 1

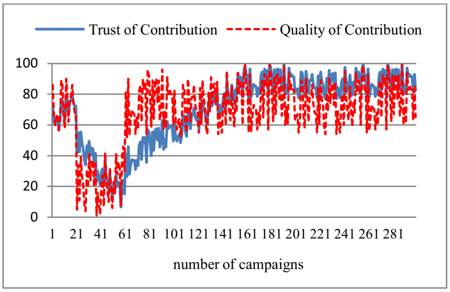

Fig. 4. Evolution of QoC & ToC for one participant, Fuzzy method, Scenario 2

Participant ID	0	8	13	9	1	14	. . . .	41	28	43	44	45
ToP	100	100	100	98	95	93	. . . .	20	16	13	10	5

Contribution ID	450	451	458	457	466	470	. . . .	495	494	490	498	496
ToC	96.6	96.6	96.6	96.4	95.7	87.2	. . . .	25	23.8	23.2	8.6	7.7

Fig. 5. Ranked lists provided by trust server with Fuzzy method, Scenario 1

category 1. After some time, this participant contributes low quality data for some campaigns. This may be because of incorrect operation of mobile device for the purpose of the sensing task (e.g., capturing unfocussed pictures). In this scenario, we assume that 15 participants transition from category 1 to category 3. In other words, the total population of the 4 categories changes from (20, 10, 10, 10) to (5, 10, 25, 10). The transitionary period lasts from the 20th to 60th campaign. Following this, the 15 participants transition back to category 1 and we return to the initial population distribution.

Fig. 6 shows the evolution of overall trust as a function of the number of campaigns in the Average and Fuzzy methods (the two Baseline methods are excluded, since we want to compare ToP related methods). There is a decrease in overall trust for both methods in the transition period, due to an increase in the number of category 3 participants, who produce low quality contributions. However, the fuzzy method is more robust at limiting the effect of these bad contributions and still achieves an acceptable level of trust. This is due to the correct adjustment of fuzzy rules such as rule no. 6 in Table. 1 which assigns a low trust rating to low quality contributions, which leads to their revocation.

As can be seen in this figure, there is a small decrease in overall trust after the transitionary period. The reason is that when participants transition to category 3, they begin providing low quality contributions, which in turn, results in low ToP for them (Recall that ToP is updated according to QoC). By transitioning back to category 1, they resume providing high quality contributions. But since ToP is still low, the obtained ToC is a value that is lower than before, but greater than revocation threshold. So, these contributions are not revoked and considered in overall trust calculation, which makes the aforementioned decrease.

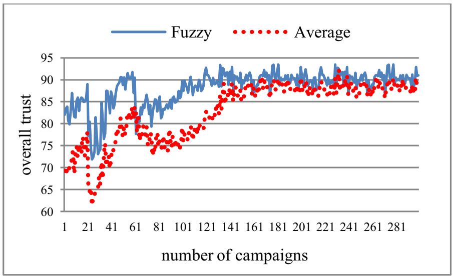

Fig. 6. Overall trust obtained in Fuzzy and Average methods in Scenario 2

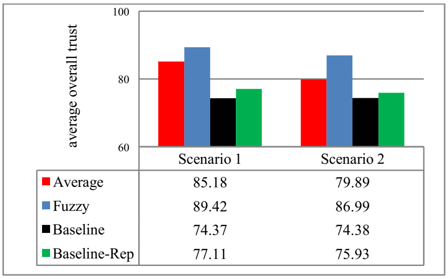

Fig. 7. Comparison of average overall trust for all methods for both scenarios

Fig. 7 presents summary results for both scenarios, averaged over 300 campaigns. Observe that the proposed fuzzy framework outperforms all other schemes in both scenarios. In particular, our scheme demonstrates high robustness to noisy contributions (scenario 2), as compared to the other schemes under consideration.

5 Conclusion

In this paper, we proposed an application agnostic trust framework for social participatory sensing system. Our system independently assesses the quality of the data and the trustworthiness of the participants and combines them via fuzzy inference engine to arrive at a comprehensive trust rating for each contribution. Simulations demonstrated that our scheme increases the overall trust by over 15% as compared to the Baseline method. As future work, we plan to extend the simulation scenarios to demonstrate the robustness of proposed framework.

References

1. Burke, J., Estrin, D., Hansen, M., Parker, A., Ramanathan, N., Reddy, S., Srivastava, M.B.: Participatory sensing. In: WSW Workshop, ACM SenSys 2006 (2006)
2. Hull, B., Bychkovsky, V., Zhang, Y., et al.: Cartel: a distributed mobile sensor computing system. In: ACM SenSys 2006 (2006)
3. Rana, R.K., Chou, C.T., Kanhere, S.S., Bulusu, N., Hu, W.: Ear-phone: an end-to-end participatory urban noise mapping. In: ACM/IEEE IPSN 2010 (2010)
4. Reddy, S., Parker, A., et al.: Image browsing, processing, and clustering for participatory sensing: lessons from a dietsense prototype. In: ACM EmNets 2007 (2007)
5. Dong, Y., Kanhere, S.S., Chou, C.T., Liu, R.P.: Automatic image capturing and processing for petrolwatch. In: ICON 2011 (2011)
6. Reddy, S., Estrin, D., Srivastava, M.: Recruitment Framework for Participatory Sensing Data Collections. In: Floréen, P., Krüger, A., Spasojevic, M. (eds.) Pervasive 2010. LNCS, vol. 6030, pp. 138–155. Springer, Heidelberg (2010)
7. Krontiris, I., Freiling, F.: Integrating people-centric sensing with social networks: A privacy research agenda. In: IEEE PERCOM 2010 (2010)

8. Krontiris, I., Freiling, F.: Urban Sensing through Social Networks: The Tension between Participation and Privacy. In: ITWDC 2010 (2010)
9. Camenisch, J., Hohenberger, S., Lysyanskaya, A.: Balancing accountability and privacy using e-cash. In: De Prisco, R., Yung, M. (eds.) SCN 2006. LNCS, vol. 4116, pp. 141–155. Springer, Heidelberg (2006)
10. Demirbas, M., Bayir, M.A., Akcora, C.G., et al.: Crowd-sourced sensing and collaboration using twitter. In: IEEE WoWMoM 2010 (2010)
11. Hays, R.B.: The day-to-day functioning of close versus casual friendships. Journal of Social and Personal Relationships 6(1) (1989)
12. Lenders, V., et al.: Location-based trust for mobile user-generated content: applications, challenges and implementations. In: ACM HotMobile 2008 (2008)
13. Saroiu, S., Wolman, A.: Enabling new mobile applications with location proofs. In: ACM HotMobile 2009 (2009)
14. Trusted computing group, https://www.trustedcomputinggroup.org/home
15. Dua, A., Bulusu, N., Feng, W.C., Hu, W.: Towards trustworthy participatory sensing. In: HotSec 2009 (2009)
16. Saroiu, S., Wolman, A.: I am a sensor, and I approve this message. In: ACM HotMobile 2010 (2010)
17. Ganeriwal, S., Balzano, L.K., Srivastava, M.B.: Reputation-based framework for high integrity sensor networks. ACM TOSN 2008 4(3) (2008)
18. Commerce, B.E., Jøsang, A., Ismail, R.: The beta reputation system. In: Bled Electronic Commerce Conference (2002)
19. Huang, K.L., Kanhere, S.S., Hu, W.: On the need for a reputation system in mobile phone based sensing. In: Ad Hoc Networks (2011)
20. Kenney, J., Keeping, E.: Mathematics of statistics, part 1. Van Nostrand, Princeton (1962)
21. Alkouz, A., Luca, E.W.D., Albayrak, S.: Latent semantic social graph model for expert discovery in facebook. In: IICS 2011 (2011)
22. Alt, F., Shirazi, A.S., Schmidt, A., Kramer, U., Nawaz, Z.: Location-based crowdsourcing: extending crowdsourcing to the real world. In: ACM NordiCHI 2010 (2010)
23. Papadimitriou, S., Kitagawa, H., Gibbons, P., Faloutsos, C.: Loci: fast outlier detection using the local correlation integral. In: IEEE ICDE 2003 (2003)
24. Leekwijck, W., Kerre, E.E.: Defuzzification: criteria and classification. Fuzzy Sets and Systems 108(2), 159–178 (1999)
25. Shekarpour, S., Katebi, S.: Modeling and evaluation of trust with an extension in semantic web. Web Semantics: Science, Services and Agents on the World Wide Web 8(1) (2010)