A Privacy Preserving Framework for Rewarding Users in Opportunistic Mobile Crowdsensing

Federico Montori*, Luca Bedogni†
*Department of Computer Science and Engineering, University of Bologna, Italy
†Department of Physics, Informatics and Mathematics, University of Modena and Reggio Emilia, Italy
* federico.montori2@unibo.it, †luca.bedogni@unimore.it

***Abstract*—Crowdsensing is rapidly becoming an interesting approach for scenarios in which a significant amount of data is needed and a static infrastructure is not a viable option due to cost or other challenges. Although users collect data without any direct cost, it is common to reward them depending on the amount and quality of the data they provide. However, as this data also carries sensitive geolocation information, it also exposes the users to privacy concerns, if such data is accessed by a malicious entity. Geolocation information can disclose information about the habit of the user and his or her places of interest, however, in many cases, such information is crucial for the purpose of the application and cannot be omitted nor distorted. In this work, we present a novel framework for opportunistic MCS scenarios focused on maintaining the privacy of the users while rewarding them for their collected and geolocated data. We evaluate our proposal on real datasets, quantifying its benefits over other methodologies.**

I. Introduction

Modern Smart Cities rely on precise and timely data that aim to describe the properties of the city and its environmental status. Some of this data comes from manually installed sensors, such as induction loops or weather station, which report accurate data to the Smart City infrastructure. This data is then analyzed together with a multitude of other information, to eventually provide enhanced services for the community. Although this architecture enables novel and exciting possibilities, it also comes at a cost, both for installing the sensors and for maintaining the whole infrastructure. Moreover, it is challenging, if not impossible, to achieve a wide coverage with static sensors, as this would again raise the costs, and it would also be difficult to place them in hard to reach scenarios. It is also worth mentioning that nowadays many Smart Cities applications are on-demand and require the whole sensing infrastructure to be flexible. For these motivations, Mobile Crowdsensing (MCS) architectures have recently started to gain popularity, thanks to their ability to provide a large amount of data, directly reported by citizens, often through off-the-shelf devices. This is usually realized by utilizing a specific application, developed by the Smart City application manager or the crowdsourcer (i.e. the entity running the crowdsensing campaign). This application can sense different data from the user device, depending on the purpose of the crowdsensing campaign, and report this data back to a central entity for later processing. The cost of running the campaign for the crowdsourcer is then limited to maintaining the infrastructure and developing the application, since the cost for the sensors is inherently payed by the users when they buy the mobile devices. Therefore, although with reduced costs, still crowdsensing campaigns involve setting up an infrastructure capable of receiving data from users. Even though users may run the campaign virtually at no cost, providing them with incentives raises the possibility to gather participants for the campaign [1][2]. Typically, collected data is geolocalized, as it is needed to correlate measurements in the space domain. Clearly, this raises a privacy concern, since having access to the data reported by a user would enable a third party to understand her movements and daily routines. In this work we study the challenge of running crowdsensing campaigns with rewards, limiting the amount of private information that can be exploited accessing data sent by the users. It is straightforward to note that if measurements are timely correlated, it is then possible to reconstruct paths, which if repeated over different days may lead to understanding routines, which eventually results in possibly understanding political orientation, religious information and other private information. Guaranteeing higher privacy levels is of paramount importance for crowdsensing campaigns, as if users rely more on the platform collecting the measurements they may be more confident in providing it with their own data. The privacy leak is quantified in terms of possible traces reconstruction by raw GPS measurements, which is a topic actively studied [3][4]. Although for the purpose of this paper we designed a specific metric to quantify the privacy, we also note that it is possible to use other metrics, although those are out of the scope of this paper. For the purpose of this work, we focused particularly on opportunistic MCS scenarios, in which data is reported by the end devices in background without any specific action performed by the users. In fact, for participatory MCS, the topic of privacy has to be discussed differently, as, in such case, the central entity (e.g. the collector server) is typically aware of the location of users over time in order to issue them with tasks. We instead focus on frameworks in which the users contribute freely in a push-based policy, therefore the server is basically unaware of their behavior. We assume that users can select which of their performed measurements they want to use in order to obtain the reward, and we then propose an algorithm which out of N measurements collected by the users selects the best possible k needed to obtain the reward. We compare our proposal with other standard selection techniques, showing that the amount of private information that can be obtained by subsequent measurements is in general high. The evaluation is performed on two different, publicly available datasets, namely Gowalla and Brightkite [5]. The rest of this paper is structured as follows: Section II presents related works from literature; Section III details our framework and our proposal, in particular it explains the metric we used to quantify the privacy issues; Section IV presents the datasets, details the experiments and shows the results of the performance evaluation, and Section V concludes this study discussing future works.

II. Related work

MCS is a topic that has been investigated thoroughly over the very last years in all its different facets. Several architectures have been proposed over the years, which target specific campaigns in a plethora of use cases, among which we cite environmental monitoring [6], social trends detection [7] and traffic estimation [8]. However, it has been shown that users tend to be more willing to share their data with a third party organization, if that leads to a reward [1][2]. In general, rewards can be broadly categorized into two different classes: monetary or non-monetary. Monetary rewards refer to incentive mechanisms that eventually award the user a prize, which can be monetary or in the form of any good that the user is willing to accept. Non-monetary rewards refer instead to all those mechanism which do not give the user any money, such as gamification proposals.

In [9] the authors foresee a novel smart city architecture, in which the users are an active part of it. Basically, they leverage user gathered sensor data, in exchange of a reward which is no less than the sensing cost of the users. Other approaches such as [10] take into account gamification methodologies. Basically users share information about the quality of WiFi hotspots in a scenario. The more data they share, the higher the probability to virtually conquer WiFi territories. This also shows that the reward can also be non-monetary, though less appealing. Gamification is also used in [11], where the objective is to reduced expenses in heavy duty scenarios, in which the participation of the users require more effort. Interesting the approach described in [12], where users are rewarded according to how many friends and neighbors they convince to participate in the crowdsensing campaign. This has a double effect, as the users are willing to publicize the crowdsensing campaign to obtain a higher reward, while the entity running the campaign also achieves a higher number of participants. Besides the specific mechanism of the reward, [13] also proposes a method to quantify the quality of the data obtained through crowdsensing, and adapts the reward to the users based on that. It is well known that one of the major problems in crowdsensing scenarios is indeed the uncertainty on the data quality, as it is possibly sensed without taking proper attention to it by the users. Hence, [13] dynamically adapts the reward based on the quality itself, thus rewarding users with higher quality data more. Finally [14] presents two different approaches which reward users in crowdsensing scenarios: a crowdsourcer centric one and a user centric one. While the former leverages game theory through a Stackelberg game where the crowdsourcer is the leader and the users are the follower, the latter uses an auction based mechanisms, in which the users have more control over the data they send to the crowdsourcer, and can also decide based on the expected reward they can get. In relation with privacy, many existing works in literature leveraged independently private information enclosure in order to cope with common de-anonymization oriented attacks (e.g. collusion and eavesdropping). Notable works are LOCATE [15], which distributes the users trajectories across all the participants to make them anonymous, and PEPSI [16], which introduces a registration authority and obscures the sensitive data with an identity based encryption. More recently, notable works like [17] also leverage cloaking of the users location by increasing its granularity, however this is not feasible in applications where the precision of the geolocalization of the measurement is crucial. Concerning the privacy of mobile users, among different proposals rise k-anonimity [18] and differential privacy [19]. However, those are more oriented in preserving the privacy of a human being among a crowd of users, in datasets, while our work is focused towards users preserving themselves their own privacy by sharing loosely correlated locations and measurements.

III. Framework

In this section we detail our proposed privacy-aware framework for MCS scenarios. We recall that such framework is only applicable in contexts in which each observation is sent anonymously to the server (i.e. the server does not record any information about the connection with the participant's device). In other words, it must be supported by a data collection framework that implies the server being unaware of the identity of each user as well as their location over time. Furthermore, communication between the central entity and the participant is assumed to occur directly, however, if mediation is performed by other participants, as it happens in adhoc MCS architectures, the framework is still applicable as long as the data sent eventually reaches a common single repository. As highlighted in Section I, the main challenge in these architectures is the rewarding mechanism, as participants, in order to claim their compensations for their contributed data, would intuitively need to expose their identity and, consequently, the history of their location. This would hinder the privacy by construction guaranteed by the architecture, hence our proposal. An example of a suitable data collection framework can be found in [20], where, the rewarding mechanism was not taken into account.

Let P be a generic client participating in the considered MCS campaign (from now on we will use the term client and participant interchangeably). As the participant moves in the area of interest for the campaign, it sends sensor measurements – a.k.a. observations – to the central server, depending on the considered data collection framework. Let then I_i be the i-th observation sent by P to the central server and let us define I_i as follows:

$$I_i = \{l_i, t_i, d_i\},$$

where l_i is the location at which the observation has been performed, t_i is the time and d_i is the numerical value of the measurement. The whole set of chronologically ordered observations $\{I_1, \ldots, I_n\}$ performed by P is defined as the *trace* of P. As introduced earlier, even though the information I_i is reported by user P, no information about such association is found inside the tuple, which only contains data referring to the observation itself. Once the observation is received, the server binds the information tuple I_i to a randomly generated hash h_i which is a token representing the value of I_i to obtain rewards. Such token is then stored in the local memory of the server as well as sent back to P. At this time, the server keeps the connection to P alive only for the purpose of sending back the reward, after which the connection is assumed to be destroyed. This is crucial, because the connection itself obviously carries information that can expose the identity of P; with its destruction, subsequent instances of the connection from P do not carry any data that can make it amenable to other connections initiated by the same user. The client P then stores the hash received from the server in its local memory. We assume that, if all such hashes are associated to the same person, they can expose potentially sensitive data, such as the trajectory of P would be disclosed. All the hashes stored locally can then be used to collect the relative rewards from the entity running the crowdsensing campaign, by sending a set of them to the central server to collect the prize. Ideally, we assume that a reward can be released in exchange for a set of k hashes with $k > 1$, as we consider scenarios in which the number of measurements can be large and rewarding all of them singularly would be impractical. Once the hashes are received, the server erases the correspondent local copies and returns a single code in exchange, which can be used by the participant to collect the final prize (e.g. a QR code). More formally, we define each crowdsensing campaign with a set of rewards (or prizes) $\mathcal{R} = \{R_1, \ldots, R_M\}$. Each R_w, with $1 \leq w \leq M$ needs k_w distinct rewards to be obtained. For sake of simplicity, in this work we assume all the generated hashes $h_1, \ldots, h_n$ to have an atomic value. A client can then access the prize R_w by "spending" k_w hashes, provided that $n \geq k_w$. This is done by selecting the prize, selecting k_w hashes and sending them to the server. We assume for the sake of simplicity and without loss of generality that all the prizes need the same number k of hashes to be obtained. The way in which a participant selects the set of k hashes has an impact on the amount of information undisclosed, as then the respective k observations can be associated to the same person: note that the association between an information I_j and its hash h_j is always known to the central entity, though it is not stored the correlation between different measurements, hence locations. In the following subsections we will give the definition of metrics that can be used to select the k observations to minimize the risks.

A. Correlation Metric

Given that the total number of hashes available to P is n, in this section we define a criterion for choosing k observations (and their respective hashes) out of n. Ideally, at each iteration, one would like to select a subset of data that yields the least possible information to a potential attacker. We choose to quantify the "sensitivity" of each piece of information by giving it a *correlation value*, which represents how much the information is related to the pool of measurements collected. Upon such premises, we define a correlation function $\Gamma(I_i)$ that returns a numerical value normalized by $[0, 1]$ (1 means that the element is maximally correlated, 0 means that the element is not correlated at all). This function is computed locally by the participant against all the elements of $\{I_1, \ldots, I_n\}$, then the least k correlated elements are selected to be converted in a reward. The definition of unique and universal metrics is out of the scope of this paper as the sensitivity of a piece of information can depend on a high number of factors with respect to different use cases. We recall that our framework is totally independent on the specific metrics used, hence it is flexible to accommodate other measures and methodologies to quantify privacy risks on top of the MCS environment taken into account. Nevertheless, we propose in this section the metric we use to address a major aspect of location-aware privacy: the disclosure of locations that are sensible for the user, such as the workplace, the home and any place that is visited frequently and periodically. In order to address it we make use of a *value function* calculated upon Markov chains. This method has been used and adapted extensively in literature for similar problems [21]. In our case we achieve our goal by defining our metricΓ as a function $\mathcal{V}^{\{\Delta t\}}$, where Δt defines a fixed time interval, which is our atomic unit of time. It is also necessary to define an additional attribute $\tau \in [0, 86400/\Delta t]$ for each observation, such that τ_i is the time of the day extracted by t_i divided by Δt. The process takes place through the following steps: first, we construct a graph such that each node is identified by a tuple $\langle l, \tau \rangle$, which represents the presence of the user in a location l at a defined time slot of the day τ. Next, for each couple of *consecutive* measurements within the trace I_i, I_j, we increase the weight of the edge $(\langle l_i, \tau_i \rangle, \langle l_j, \tau_j \rangle)$ by one, provided that they occur within the same day. Then, each node $\langle l_i, \tau_i \rangle$ is initialized with a value equals to its weighted indegree. A high value corresponds to a location and a time of major interest, since the user often traveled to such location during that time of the day. Subsequently, for each other node $\langle l_j, \tau_j \rangle$ reachable in at most γ steps, the initial value of $\langle l_j, \tau_j \rangle$ multiplied by the probability of the path connecting $\langle l_i, \tau_i \rangle$ and $\langle l_j, \tau_j \rangle$. The probability of a path of $\{E_1, \ldots, E_X\}$ ordered edges is defined as:

$$\prod_{x=1}^{X} weight(E_x).$$

At the end of the process, for each observation I_i, $\mathcal{V}^{\{\Delta t\}}(I_i)$ results in the value of the related node in the graph $\langle l_i, \tau_i \rangle$. Once values are calculated, we normalize them by their maximum value, in order to respect the metric properties.

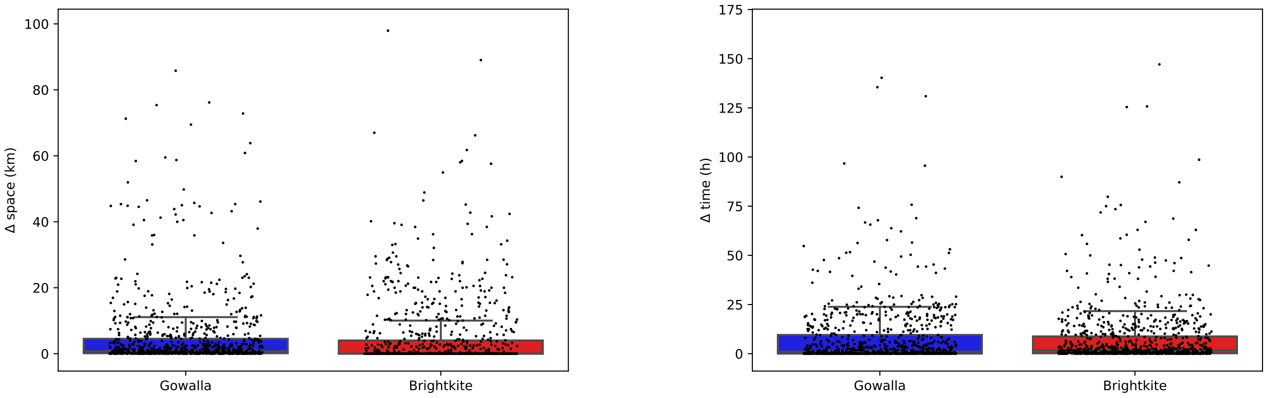

(a) Spatial distribution.

(b) Temporal distribution.

Fig. 1: Distribution of spatial and temporal intervals between each pair of consecutive measurements in Gowalla and Brightkite datasets. The strip plot shows only about 10% of the points for displaying purposes.

IV. Performance Evaluation

In this section we evaluate the performance of our framework. In Section IV-A we describe the datasets we used for our evaluation through and ad-hoc simulation, and in Section IV-B we detail the parameters and the environment used.

A. Datasets

To evaluate our work, we used 2 different datasets retrievable on the web. The datasets are check-in based, meaning that they record, for each user, a list of location visited in GPS coordinates alongside with the timestamp. We chose such datasets because they have been used extensively in the literature for a multitude of related problems. A brief description of the datasets is reported below:

- **Gowalla**: Gowalla[1] [5] is a location-based social network and its dataset has been used in various studies on the privacy.
- **Brightkite**: Brightkite[2] [5] is a location-based social network and its dataset has been used in various studies on the privacy. Brightkite and Gowalla are often used in comparison studies as they have a similar structure.

Each datasethas been adapted in order to contain data about 100 users, with each trace being exactly 90 observations long, for the sake of the comparison. Due to the extreme heterogeneity of spatial and temporal differences within each datasets – certain users may turn off their sensing for many days and/or travel significant distances – we extracted only traces containing 90 subsequent observations such that the time difference between two consecutive observations is less than a week and the spatial distance is less than 100km. Subsequently, location data has been normalized in order to identify areas instead of GPS locations. In order to have a standardized area denotation, we made use of the Military Grid Reference System (MGRS) [22] with different area size depending on the datasets. As Gowalla and Brightkite have sparsely distributed observations (both in time and space), then we set the precision of the MGRS squares to be 10km. The experiments can be scaled to different precision in datasets displaying a finer granularity of data. Figure 1 shows the distribution of space and time differences in the two datasets by means of a boxed strip plot. Regarding the spatial difference we can observe that there are for both datasets some cases in which two consecutive observations are performed by the same user at a distance up to 100km. However, these are mostly outliers, as the vast majority of the spatial distances are concentrated in the lowermost section of the plots. Indeed, the median values are ∼780m for Gowalla and ∼0m for Brightkite (meaning that Brightkite users are mostly still). A similar behavior is observable in the time domain, in which some time intervals can be as large as 100 hours, however the median values are around 60min for Gowalla and around 100min for Brightkite.

B. Experimental Setup

In order to evaluate our experiments we performed simulations with 100 users, each of them performing 90 measurements. The purpose of the experiment is showing that, as the number of measurement increases, potentially selecting k observations among them makes the correlation values to drop. In order to show this, we performed simulations with the value of k fixed to a certain amount for all users. In real deployments, the value of k might be different across users and even across prizes for the same user over time, however, fixing it to one value does not prejudicate our experiments to be representative of the reality and it facilitates the display of results. In particular, we have three different values of k: 10, 20 and 30. The simulation consists in calculating, for each user, the correlation of the k values selected over N, where N ranges from 0 to 90, which is the maximum length of the

[1] https://snap.stanford.edu/data/loc-gowalla.html

[2] https://snap.stanford.edu/data/loc-Brightkite.html

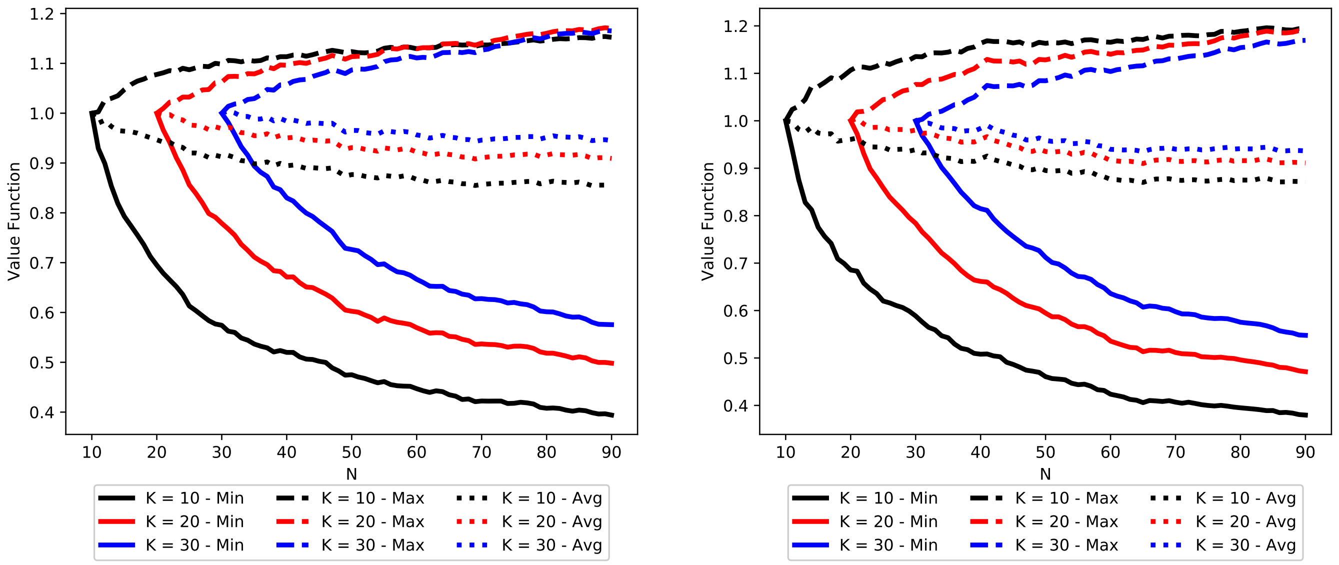

(a) Gowalla dataset

(b) Brightkite dataset

Fig. 2: Average correlation values of the selection of k elements N ranging from 0 to 90 for 100 users.

trace. For each value of N, we extract k elements in three different ways:

- We sort the N measurements by their correlation value, as shown in Section III-A and pick the bottom k values. This ensures the best combination of unrelated k values according to our metric. The average correlation value of the k-sized subset of measurements selected this way is the **Min** correlation value and reflects our goal.
- As a term of comparison, we pick the top k values from the N sorted measurements. This represents the worst combination, i.e. the k values that yield the maximum information according to our metric. The average correlation of the k-sized subset of measurements selected this way is the **Max** correlation value.
- As a further term of comparison, we also pick the mean correlation value over all the N measurements as a fair average case. Such correlation value is the **Avg** correlation value.

These three values are shown in Figure 2 for both Gowalla and Brightkite datasets for all the values of k. Logically lines start from the relative value of k as it would not make sense to select k values over a quantity that is less than k. For a better readability, values are normalized by the selection of k elements over k, for which the correlation value is set to 1. By observing the plots we observe what is expected: selecting the k elements with minimum correlation leads to a low correlation of the subset, whereas it happens the opposite when selecting the maximum. It is also noticeable how Max and Min lines tend to be almost monotone even though every time a new observation is added (i.e. N increases) the values in the graph can change. It is also worth mentioning that the three Avg lines have the same absolute values, however, they are shifted due to normalization. Figure 3 shows two area plots, one for each dataset, representing the difference between the Max and the Min correlation values for an increasing value of N. When $N = k$ Max and Min are the same value, therefore the area is 0. We can appreciate how, differently from the plots in Figure 2, the shape of these ones in monotone, meaning that, with our metric, by adding a new observation, the difference between Max and Min cannot decrease. This also leads to the concept that the more a participant "waits" before sending its measurements, the higher the chance for the k measurements to be unrelated with his or her habits and, therefore, to yield sensitive information.

V. Conclusions and Future Work

In this paper we have presented a novel framework for privacy preservation in opportunistic MCS scenarios for the purpose of rewarding users. This framework ensures that measurements reported by participants are not associated directly to the person of interest and provides a way for the user to expose a subset of the measurements in order to obtain the related reward. We also provided a mechanism thanks to which each subset of measurements can be released without any direct association of such subsets with their owner. Furthermore, we defined a policy of choosing the subset of measurements by assigning a correlation metric to each of them so that the subset with the least average correlation can be chosen at each iteration; in this way even the risk of indirect association can be lowered. We provided simulation results that demonstrate the effectiveness of our solution by using our metric, however, several other metrics can be used with this framework depending on the application domain and requirements. We consider this as a pioneering effort in the field of rewarding users in opportunistic MCS scenarios, in fact we envision a consistent number of improvements as a

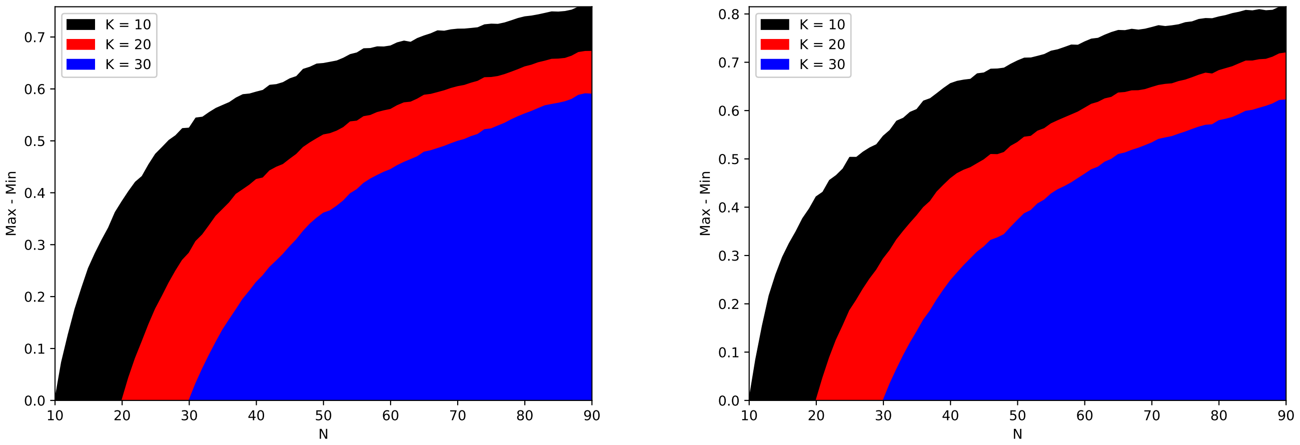

(a) Gowalla dataset

(b) Brightkite dataset

Fig. 3: Area plot of the differences between Max and Min correlation values for both the Gowalla (Figure 3a) and Brightkite (Figure 3b) dataset, for different K and N values.

future work. First of all, this framework will be tested with a larger and differentiated number of metrics that can give us more details on which of them has a higher information gain in different datasets. Indeed, we also plan to experiment our framework on a higher number of well-known datasets with substantial differences in terms of data granularity and distribution. Lastly, we also envision how to include metrics that do not imply an absolute correlation – i.e. the correlation of an information with all the others –, but also a correlation between pairs of measurements or subsets. This boosts the computational complexity of the research for the minimum, as it becomes an NP-hard problem, therefore we will investigate further on possible solutions.

References

[1] F. Montori, L. Bedogni, and L. Bononi, "A collaborative internet of things architecture for smart cities and environmental monitoring," *IEEE Internet of Things Journal*, vol. 5, no. 2, pp. 592–605, April 2018.

[2] X. Zhang, Z. Yang, W. Sun, Y. Liu, S. Tang, K. Xing, and X. Mao, "Incentives for mobile crowd sensing: A survey," *IEEE Communications Surveys & Tutorials*, 2016.

[3] M. Gramaglia, M. Fiore, A. Tarable, and A. Banchs, "Preserving mobile subscriber privacy in open datasets of spatiotemporal trajectories," in *IEEE INFOCOM 2017 - IEEE Conference on Computer Communications*, May 2017, pp. 1–9.

[4] L. Bedogni, M. Fiore, and C. Glacet, "Temporal reachability in vehicular networks," in *IEEE INFOCOM 2018 - IEEE Conference on Computer Communications*, April 2018.

[5] E. Cho, S. A. Myers, and J. Leskovec, "Friendship and mobility: user movement in location-based social networks," in *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 2011, pp. 1082–1090.

[6] R. K. Rana, C. T. Chou, S. S. Kanhere, N. Bulusu, and W. Hu, "Ear-Phone : An End-to-End Participatory Urban Noise Mapping System," *Proceedings of the International Conference on Information Processing in Sensor Networks IPSN*, pp. 105–116, 2010.

[7] B. Guo, H. Chen, Z. Yu, X. Xie, S. Huangfu, and D. Zhang, "FlierMeet: A Mobile Crowdsensing System for Cross-Space Public Information Reposting, Tagging, and Sharing," *IEEE Transactions on Mobile Computing*, 2015.

[8] B. Pan, Y. Zheng, D. Wilkie, and C. Shahabi, "Crowd Sensing of Traffic Anomalies Based on Human Mobility and Social Media," *Proceedings of the 21st ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, pp. 344–353, 2013.

[9] M. Pouryazdan, B. Kantarci, T. Soyata, and H. Song, "Anchor-assisted and vote-based trustworthiness assurance in smart city crowdsensing," *IEEE Access*, vol. 4, pp. 529–541, 2016.

[10] F. Wu and T. Luo, "Wifiscout: A crowdsensing wifi advisory system with gamification-based incentive," in *2014 IEEE 11th International Conference on Mobile Ad Hoc and Sensor Systems*, Oct 2014, pp. 533–534.

[11] Y. Ueyama, M. Tamai, Y. Arakawa, and K. Yasumoto, "Gamification-based incentive mechanism for participatory sensing," in *2014 IEEE International Conference on Pervasive Computing and Communication Workshops (PERCOM WORKSHOPS)*, March 2014, pp. 98–103.

[12] G. Yang, S. He, Z. Shi, and J. Chen, "Promoting cooperation by the social incentive mechanism in mobile crowdsensing," *IEEE Communications Magazine*, vol. 55, no. 3, pp. 86–92, March 2017.

[13] D. Peng, F. Wu, and G. Chen, "Pay as how well you do: A quality based incentive mechanism for crowdsensing," in *Proceedings of the 16th ACM International Symposium on Mobile Ad Hoc Networking and Computing*, ser. MobiHoc '15. New York, NY, USA: ACM, 2015, pp. 177–186.

[14] D. Yang, G. Xue, X. Fang, and J. Tang, "Incentive mechanisms for crowdsensing: Crowdsourcing with smartphones," *IEEE/ACM Trans. Netw.*, 2016.

[15] I. Boutsis and V. Kalogeraki, "Privacy preservation for participatory sensing data," in *Pervasive Computing and Communications (PerCom), 2013 IEEE International Conference on*. IEEE, 2013, pp. 103–113.

[16] E. De Cristofaro and C. Soriente, "Participatory privacy: Enabling privacy in participatory sensing," *IEEE Network*, vol. 27, no. 1, pp. 32–36, 2013.

[17] T. Peng, Q. Liu, D. Meng, and G. Wang, "Collaborative trajectory privacy preserving scheme in location-based services," *Information Sciences*, vol. 387, pp. 165–179, 2017.

[18] L. Sweeney, "k-anonymity: A model for protecting privacy," *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, vol. 10, no. 05, pp. 557–570, 2002.

[19] C. Dwork, *Differential Privacy*. Boston, MA: Springer US, 2011, pp. 338–340.

[20] F. Montori, L. Bedogni, and L. Bononi, "Distributed data collection control in opportunistic mobile crowdsensing," in *Proceedings of the 3rd Workshop on Experiences with the Design and Implementation of Smart Objects*. ACM, 2017.

[21] L. Bedogni and M. Levorato, "Rising user privacy against predictive context awareness through adversarial information injection," in *2018 IEEE Global Communications Conference (GLOBECOM)*. IEEE, 2018, pp. 1–6.

[22] R. Lampinen, "Universal transverse mercator (utm) and military grid reference system (mgrs)," 2001.