A Joint Evaluation Methodology for Service Quality and User Privacy in Location Based Systems

Luca Bedogni
luca.bedogni@unimore.it
University of Modena and Reggio Emilia
Italy

Chiara Franceschini
chiara.franceschini@unimore.it
University of Modena and Reggio Emilia
Italy

Federico Montori
federico.montori2@unibo.it
University of Bologna
Italy

ABSTRACT

Pervasive and ubiquitous applications provide novel and exciting services leveraging on a multitude of data obtained from people's devices, adapting the computation to the context in which the user currently is. This improves the service quality of these applications, which can provide a more tailored configuration of the application itself depending on the user context and needs. In these scenarios privacy is of paramount importance, since users must be also be protected against the misuse of their personal data. Analyzing ubiquitous systems in terms of service quality and privacy issues is however a challenging task, due to the heterogeneity of the possible attacks, which makes it difficult to compare two applications. In this paper we propose a novel methodology to jointly evaluate the service quality and the privacy issues in ubiquitous applications in an extensible and comparable way, building on the data available in each part of the system to be analyzed, and defining service qualities and privacy issues so that they can be easily re-used in other analyses. Our evaluation on a candidate application highlights the benefits of our proposal, showing the dependency between privacy levels and service quality, and paving the way for a novel methodology for the definition of these scenarios.

CCS CONCEPTS

• **Theory of computation** → *Theory of database privacy and security*; • **Security and privacy** → **Privacy-preserving protocols**; • **Information systems** → *Crowdsourcing*.

KEYWORDS

Location Based Services, Privacy, Performance evaluation

ACM Reference Format:
Luca Bedogni, Chiara Franceschini, and Federico Montori. 2023. A Joint Evaluation Methodology for Service Quality and User Privacy in Location Based Systems. In *ACM International Conference on Information Technology for Social Good (GoodIT '23), September 06–08, 2023, Lisbon, Portugal.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3582515.3609524

GoodIT '23, September 06–08, 2023, Lisbon, Portugal
© 2023 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0116-0/23/09.
https://doi.org/10.1145/3582515.3609524

1 INTRODUCTION

Ubiquitous applications and services are part of everyday lives, being them mobile applications, wearables, or with custom devices designed to provide specific, context aware services. These frameworks typically collect various data such as activities performed by the user, position, neighbors, environmental data and many others [9]. This enables applications to assess the context of the user and provide a multitude of services tailored specifically to the user owning the smart device, and to the particular scenario assessed through the data collection and analysis. Although these applications are now widespread, several concerns remain for what pertains the privacy issues related to the collection and storage of user data. In particular users are often unaware of the potential threats that hide behind the collection and management of specific data. In fact, understanding which data may raise a specific privacy issue is not always straightforward, hence seldom understood by users, as there are a multitude of potential privacy threats, raised by the collection of specific data with precise characteristics. There is a general understanding that the location of the user may reveal private aspects of the user's life, but issues may also arise from other types of data collected, as shown by [20] [4], making the scenario even more challenging. Clearly it exists a trade-off between data utility, in terms of quality of service received by the user, and privacy, with respect to possible attacks and information malicious entities may understand if such data is accessed. Moreover, users are willing to share their information in exchange for a service tailored to their needs, leading to a potential privacy paradox [15]. It is then important to balance the data collection process directly by the users, which must be able to decide which data to share and which to keep private. Although this is made possible by some applications, it is often limiting, as not giving the consent to collect any kind of data may result in the service to stop working, rather than just limiting its functions. More importantly, it is difficult to compare different applications providing a similar service, both in terms of data needed to accomplish the task, and the potential privacy vulnerabilities. A practical example was experienced around March 2020, when Mobile Contact Tracing applications were proposed to contrast the SARS-CoV2 pandemic. There was a number of different proposals, both in terms of architecture, data managed, communication protocols and so on. All of them claimed to be privacy preserving, and different studies tested each one of the variants against known attacks to check the system response. However, it was difficult if not impossible to directly compare two different proposals that aim to solve the same task, it was only possible to claim that systems were resilient specific attacks [1] [5].

In this research paper, we present an innovative and comprehensive methodology that aims to change the assessment of service qualities and privacy concerns in the realm of ubiquitous services. Our approach is centered around the development of a novel framework, which we refer to as the Data Meta Descriptor, enabling the establishment of a unified definition that seamlessly integrates both service qualities and privacy issues.

Furthermore, we also present a candidate implementation of our proposal, showing its feasibility and real-world applicability. We analyze it in detail, and provide performance evaluation on how our system can be seamlessly integrated into an existing ubiquitous application.

The rest of this paper is organized as follows: Section 2 discusses the related works from literature; Section 3 highlights the main limitations of the scenario; Section 4 presents the Data Meta Descriptors; Section 5 describes the implementation of our framework on a candidate application, while Section 6 evaluates it; Section 7 concludes this works and discusses future works on this topic.

2 RELATED WORKS

As our lives become more entangled with ubiquitous services and devices, there is a continuous need to assess and prevent possible users' privacy breaks, which may expose information not meant to be shared with others. Depending on the device and on the scenario, there may be different vulnerabilities which arise from the use of personal data [16]. The possible issue may spark in different parts of the system, such as the communication network, due to data sent without encryption or protocols prone to known attacks, or while data is stored, due to possible identification of individuals. Other problems may also arise from the authentication schemes these devices use, which can lead to potentially huge problems [17]. In general, there is a gap between understanding potential privacy vulnerabilities of ubiquitous systems, and the user service such systems have to provide. It is well known that if users understand about potential privacy breaches, they may quit the ubiquitous service to be resilient against them, while also increasing their distrust towards digital systems [13]. On the other hand, they may also want to personalize it in terms of data collected, something which nowadays is difficult to do on a fine-grained basis, as it is rather an on-off choice, whether to run the service with the full data requested, or simply not using it [24]. Most privacy policies are written in natural language, detailing all the legal aspect needed to be compliant with the regulations, possibly slowing down the adoption of Ubiquitous and Pervasive applications [6]. Usable Privacy is a field which aims to make privacy policies more understandable by users, by making policies machine readable [22]. This enables the possibility to check how data is collected and used, and enables fine grained customization by the user, which is now able to clearly understand the data management process [14]. Nevertheless, users are well aware that many potential threats may hide beyond the data collection and analysis process, and to this end device manufacturers and applications developer have yet to provide global and satisfactory solutions [10]. In fact, there are still a number of challenges which need to be properly addressed [2]. Eventually it becomes a trade-off between privacy and utility, which however it would be preferable if the users themselves are able to configure it [19], while the stakeholder may also want to reward the user for the data they provide, such as in crowdsensing scenario, in which privacy has to be preserved [18]. Most of the research in this area is towards the development of models and frameworks which enable the possibility for application developers to quantify and balance the required service quality with the potential privacy leakage. The interaction between service utility and privacy issues is a well-known trade-off, target of many studies in literature [21] [8] [7], which show how the scenario can be heterogeneous and complex due to a number of different parameters, and due to many heterogeneous devices which have to interact among each other. An example of these frameworks is certainly [21], where the authors focus on Location Based Services and model the problem of service utility and privacy with a Stackelberg game. There are many interesting findings, such as the optimality of the solution leveraging their design, and the fact that they indeed show that there is a trade-off between service utility and privacy, but after once a certain degree of privacy has been reached, there is no need to restrict the quality requirements furthermore, as these does not directly translate into better privacy protection.

In [11] the authors build a model for user-centered privacy, but the use-case is limited to the medical scenario, hence challenging to be generalized to other domains, similarly to [12] which considers smart meters. Overall, studies often target a specific framework and propose user-centered methodologies which only apply to the domain of interest, without being able to generalize them also to other scenarios. It is instead much harder to find research which proceed towards the ability to compare different systems both on privacy and on utility. This depends on the fact that as it has already been stated, typically studies focus on very specific scenarios of interest, without considering data collection and management as a whole, together with the service utility.

3 SERVICE QUALITY AND PRIVACY ISSUES LIMITATIONS

In the current dynamic landscape concerning the trade-off between privacy and service quality, it is key to acknowledge and address the different challenges that exist. Within this context, we can identify and discuss four principal limitations that have to be analyzed, which we describe in detail below:

3.0.1 Limitation 1 – Understanding. : it is not always clear which attacks can be performed starting from the unauthorized access to such data and to what extent. Although there is no obligation for the stakeholder to inform the users about the possible consequences of the misuse of their personal data, the article 35 of the GDPR in Europe states that is mandatory to take into account possible threats and issues and assess their impact. However, to be compliant with articles 13 and 14 of the GDPR the data controller must only guarantee the adoption of appropriate measures and inform the users. This specific problem arises since not all users have the specific knowledge to address all the vulnerabilities and issues that the collection of a particular data triggers. Often, users may make uninformed decisions related to the practical threats related to data collection, or simply skip reading such information. This may also increase the distrust of users towards ubiquitous applications, eventually leaving them or not fully leveraging their benefits [13].

3.0.2 Limitation 2 – Comparison. : if two separate systems A and B collect any kind of data, it is impossible to directly compare them with respect to the privacy vulnerabilities and the need for specific data to fulfill a specific task. What is typically done is a per-system analysis which aims to address the most common attacks and showcase whether such attacks are possible or not in any given framework. However, comparing two systems would require performing two separate analysis, and then compare the results, which cannot give a complete and extensive picture, as the comparison would be made only on the specific threats taken into account. Moreover, the comparison would be only on the basis of the tested attacks and vulnerabilities, without focusing on the practical issues for the users. It is also not possible to directly compare two applications based on the services they offer, in terms of data collected to provide any given service, and how this directly relates to the potential privacy issues such data collection raises.

3.0.3 Limitation 3 – Balance. : In certain cases, users may configure ubiquitous systems according to their preferences, in terms of data collection and active services. Typically, the denial of collecting specific data translates into part of the application to stop functioning, without the possibility for user to still access the service with reduced functionality and performance. This severely limits the possibility for the users to tailor any application according to their preferences, either accepting entirely how their data is handled or avoid using the system at all. It is also important to note that the balance between data collection and service provided to the end user is among the main ideas of the GDPR, which promotes a digital economy in which data has a key role while keeping users protected against their misuse.

3.0.4 Limitation 4 – Time consuming. : Analyzing privacy vulnerabilities is a time-consuming task, which needs to be repeated for each new system, increasing the possibility for errors. Moreover, in case a new vulnerability is found, existing systems would need to be tested from scratch against the novel attack, to check whether they may be at risk. It would be much more advantageous in this case to have automatic and semi-automatic controls available, which upon discovering a novel kind of attack, enable the seamless testing and reporting of other systems previously checked. Instead, this has to be done for each system separately, again requiring considerable time and specialized technical people. Such tools may also provide practical countermeasures to eliminate or limit the potential issues, by guiding stakeholders through the design and the development of their applications.

The aim of this work is to address all the aforementioned limitations leveraging on a novel methodology which builds upon the Data Meta Descriptor (DMD) component, described in Section 4.

Concerning the Understanding (Limitation 1), the DMD will make it possible for users to directly understand what threats hide behind specific data collection. Hence, they can make more informed decisions, and personally assess whether the permission to collect specific data should be given or not. Regarding instead the Comparison (Limitation 2), sharing an identical data definition, it will also be possible to directly compare 2 or more different systems with respect to the services they offer and the privacy vulnerabilities they may suffer. This would also allow to compare any system according to the amount of data they require to perform similar tasks. The Balance (Limitation 3) is then made possible since it would be allowed to directly assess the consequences of constraining the data collection process, not only as a binary on-off choice, but with more fine-grained possibilities, such as reducing frequency, constraining the collection based on external factors, and so on. Finally, the Time consuming (Limitation 4) can be addressed by developing automatic tools, which can be inherited by ubiquitous applications, to automatically assess potential privacy vulnerabilities and offering solutions to mitigate them.

Eventually, using the tools that the DMD enable, it would also be possible for ubiquitous applications developers to showcase an automatic assessment of their service, easily readable by anyone, which details all the data collection, management and analysis process, and the potential vulnerabilities which the user may experience when allowing the system to collect their data. It will also be made possible to directly understand how user preferences would affect the potential privacy vulnerabilities. Ultimately it is an equilibrium which every user may select between clear privacy vulnerabilities and ubiquitous utility, which is dependent on the scenario, the type of data accessed, and on the user's preferences.

4 DATA META DESCRIPTORS

In this section we present the Data Meta Descriptors (DMD), defined as a set of information related to a single chunk of data in terms of data type, frequency, precision and other parameters. A single ubiquitous system is composed by several DMDs, each one describing a part of the data available on the system at any stage of it. The overarching idea is to define DMD that can describe both privacy vulnerabilities as well as service quality. In this case, it would then be possible to jointly alter their parameters, and immediately understand how the system would react to the change.

In Figure 1 we show the overarching idea of our proposal. On the left it is possible to see the DMD repository, which contains the DMD definition for any Service Quality evaluated, as well as for privacy issues. Whenever a candidate ubiquitous system needs to be analyzed, the different DMD present in the system must be assessed. They will then be matched against the Service Qualities and Issues in the DMD Repository, eventually providing an evaluation of the whole system. In particular we note from Figure 1 that different Service Qualities and Issues may be found depending on the DMD available on each part of the system. For instance, consider a data collected at a certain frequency on the mobile device, which is then aggregated to be sent over the network, and anonymized once reaching the cloud. Clearly, even if it is the same raw data, its storage at different precisions, granularities and frequencies may offer different Service Quality and may raise different privacy issues.

The analysis of different privacy threats builds a repository of potential vulnerabilities modeled through DMDs, each one with a set of DMDs needed for the vulnerability to happen, in case they are accessed by a malicious entity. This has to be intended that "If a malicious entity has access to those DMDs, then the corresponding attack is possible". The repository also describes different ubiquitous services, with different possible configurations. For instance, a service may run with full functionality in case 4 different type

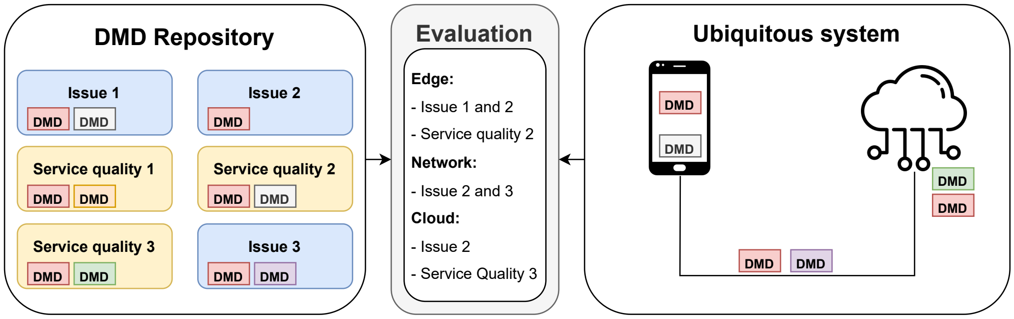

Figure 1: Data Meta Descriptors repository and example matching on a candidate ubiquitous system. The DMD repository contains both the privacy issues and the service qualities; the ubiquitous system is analyzed by surveying the data available in every part of it; joining the analysis and the DMD Repository it is then possible to perform an automated evaluation.

of data are collected, it may work with limited functionality and with only one type of data it provides basic functionality. Through this methodology it is also possible to evaluate different service qualities depending on the data. In fact, changing any DMD yields a different service quality, directly quantifiable, and eventually the new set of DMD is matched against the repository to look for privacy issues. This process can be repeated until the desired privacy level is reached, or until no further service degradation is accepted by the user. It is also known that typically there is a certain privacy level beyond which not additional privacy is granted, though the service quality degrades, as shown in [21].

An example of this behavior can be observed in location based services, for instance for the task "Find Point of Interest of Type X close to me". Even a simple service like the one just described can present a number of possible configurations and personalizations. The term "close to me" may be intended in different ways: close can be 100 meters, 1 kilometer, or close to another Point of Interest. Moreover, even the location from which the closeness of a Point of Interest is determined can vary: for instance, it may be possible to provide the location of a user with the best possible precision, uncovering the location of the user but getting in return the best possible service quality. Otherwise it is possible to reduce the granularity of it, still leveraging the service, though with a reduced quality. Other examples can be found for instance in the precision of the context obtained from inertial sensors such as accelerometers, gyroscope and magnetometers, which directly relates to the precision and frequency of sampling of data which again may uncover specific habits and routines of people [3] [4].

4.1 Overview

Data Meta Descriptors (DMD) describe a specific kind of data in any part of an ubiquitous system. It does not contain the data itself, but rather metadata which describe for instance the type of data, the frequency of collection and other attributes.

More formally, let $\Delta = \{dt, \overline{\kappa}\}$ be a generic DMD referring to data of type dt and with a set of attributes of variable length $\overline{\kappa}$, which contains a variable number of attributes, such as the precision of the data, the time granularity and so on.

In any section of a ubiquitous system such as a mobile device, the communication medium, the storage of data on the cloud, multiple DMD may exist. Let s be the s-th section of ubiquitous system U, hence

$$\overline{\Delta}_s^U = \{\Delta_{s,1}^U, \ldots, \Delta_{s,n}^U\}$$

is the set of DMDs of system U in its s-th section. In other words, $\overline{\Delta}_s^U$ describes all the data available in a specific part of the system, with its corresponding metadata.

4.2 DMD relations

In any ubiquitous system there exists a relation pertaining how data flows between its different sections. For instance, data can be collected on a mobile device and then sent to the cloud for storage. In this case, data collected on the mobile device has a set of attributes as defined in $\overline{\kappa}$, and when sending such data over the communication link the set $\overline{\kappa}$ may change.

More formally, let i be a section of a ubiquitous system and j a subsequent section on which the data flows after i. We then have two DMD sets $\overline{\Delta}_i^U$ and $\overline{\Delta}_j^U$.

For any DMD, the following relation holds:

$$\kappa_i \geq \kappa_j, \forall i \prec j \wedge t_i^U \equiv t_j^U.$$

This relation between two set of attributes defines that attributes in κ_j cannot be less restrictive than attributes in κ_i. This is straightforward, as subsequent data flows cannot increase for instance the precision of data or its frequency, but only keeping it as it is or reducing it. For instance mobile devices can store locations with a precision of $10m$, and upon sending it they can keep it with the same precision or making it coarser such as $50m$, but not providing it with a more precise representation such as $5m$.

5 IMPLEMENTATION

In this section we provide the implementation of a candidate Service Quality metric and of a candidate Privacy metric. We recall that the proposed DMD framework can implement different metrics, for which it is impossible to give an all-encompassing implementation here. However, we believe that this example can, without loss of generality, showcase how the proposed framework works.

5.1 k-anonymity

To evaluate the privacy of users we leverage k-anonymity [23], which is defined as follows: a certain dataset is said to be k-anonymous if each person in the dataset is indistinguishable from $k-1$ users. In our scenario the ubiquitous service only collects four kind of information: a pseudonym p, a timestamp t, a generic measurement m and a location l.

Clearly t and m are non-identifiers of the user, while p is by definition a pseudonym. Therefore a possible attack is the re-identification of a user starting form its pseudonym. It is straightforward to note that in case there are few users in a certain area at a certain time, an attacker could easily associate the pseudonym to the user, hence performing the re-identification of her. On the inverse, in crowded areas the user can better hide herself among others, therefore making the attack more difficult.

We assume that in our area of interest users arrive with a Poisson process of parameter λ. Therefore, it is possible to compute the probability of having k users which report the exact same location, which is equal to

$$P(X=k) = \sum_{i=1}^{\Delta} \frac{(\lambda_i)^k}{k!} e^{-\lambda_i} \tag{1}$$

where Δ is the total area of interest, and λ_i is the parameter of the Poisson process in location i. Note that this definition allows to define more crowded places such as shops and points of interest, and less crowded ones such as roads.

5.2 Service Quality

To evaluate the service quality we count the number of POIs returned by a user query. The overarching idea is that the coarser the location, the higher the privacy, but at the same time it will also be returned a higher number of POIs, while for more precise locations the query will return closer POIs to the user position.

Let N be the total number of POIs in an area of size Δ. We compute the number of points inside an area of size δ as

$$\binom{N}{x} p^{\delta} (1-p)^{N-\delta} \tag{2}$$

where n is the number of POIs, and δ is the size of the area sent by the user. Without loss of generality we consider squared areas, such as in the popular MGRS geocoding. Assuming that the POIs are uniformly distributed in Δ, the probability of having one POI in the area sent by the user is given by

$$p = \frac{\delta}{\Delta}, \tag{3}$$

Figure 2 shows the probability of having a specific number of POIs in an area of a given size. This figure has been obtained with a total of 50 POIs in an area of $1km^2$. For sake of simplicity, we are representing δ as percentages with respect to Δ. Clearly for smaller areas the probability is shifted towards the left, meaning that there is a high probability of having a low number of POIs, while larger areas allow for more POIs in the same query.

We can then compute the expected value as

$$E[\delta] = \sum_{i=1}^{\infty} \delta_i p_i, \tag{4}$$

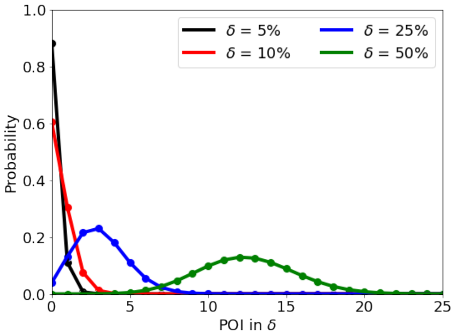

Figure 2: Probability of having a specific number of POIs in an area of size δ. Larger areas show higher probabilities of a larger number of POIs.

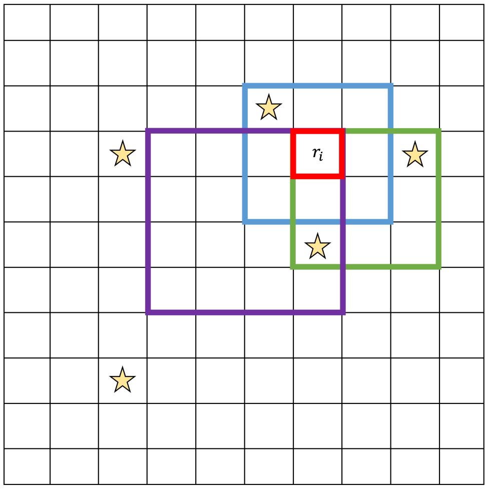

Figure 3: How measurement precision works. Red square is precision 0, green and blue squares are precision 1, and purple square is precision 2. Note that r_i can be anywhere inside the squared region.

which, given a certain x representing the area sent, contains the expected number of POIs in such area.

Starting from the expected value it is then possible to compute the service quality, simply defined as

$$SQ = 1 - \frac{E[\delta]}{N} \tag{5}$$

We show SQ in Figure 4 for a varying number of POIs. It is evident that when δ is small, the service quality is at maximum given the low number of POIs returned, while for larger δ the service quality drops. Clearly also the density of POIs matters, hence for less dense areas (i.e. 20 POIs) the service quality can be maintained higher with respect to denser areas.

Figure 3 shows how precision works. Smaller precision numbers report a more accurate location, while higher precision numbers refer to bigger areas. When users report bigger areas, the location of the user can be anywhere in such area.

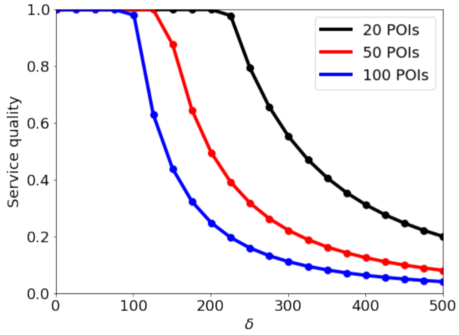

Figure 4: Evaluation of the sole Service Quality for three possible scenarios, versus the area sent by the user.

Given the square in which the user can be, the central server needs to report back the most appropriate POIs. We can decide how to deal with this, for instance it can report all the points in the area, or all the points which are located at most λ squares from any possible point in the square.

6 EVALUATION

In this section we provide performance evaluation of our system by considering our example metrics, and we highlight how starting from the single evaluation of either the Service Quality or the Privacy level it is possible to analyze them together. We remark again that although specific to a candidate service, the system can accommodate also other definitions.

In Figure 4 we show the evaluation of the Service Quality alone, for three reference scenarios. On the x-axis we vary the area size, while on the y-axis we plot the Service quality, defined as in Equation 5. We test our system in three reference scenarios, with a varying number of POIs which translates into a different density of them within the same area. With denser areas (i.e. with more POIs) the Service Quality drops more significantly and earlier than coarser scenarios. This happens even if the user queries for a smaller area, since the probability of having more POIs within that area increases for denser scenarios, hence the overall Service Quality is lower. It is also interesting to note that regardless of the scenario, there is always a maximum area size which is possible to use without affecting the Service Quality, which in our case it is around $100m^2$ for the 100 POIs scenario, $150m^2$ for the 50 POIs scenario and around $200m^2$ for the 20 POIs scenario. This suggests that for services which evaluates the Service Quality with the same methodology of our proposal, there is never the need to submit the precise location of the user, even for denser scenarios.

In Figure 5 we evaluate instead the Privacy level, accounting for the k-anonymity of the users. Specifically, we vary on the x-axis the k-anonymity to achieve, while on the y-axis we show the ratio of users which satisfy such constraint. The lines indicate groups of users which use different areas for their query. Clearly less conservatives users in terms of privacy (i.e. using a smaller area for their query) achieve a lower ratio of satisfied users as the k-anonymity requirement grows. More conservatives users instead can better hide within the crowd, since there are more other users which use similar locations when issuing their queries. It is also interesting to observe that depending on the area size selected for the query, there is a maximum k-anonymity level which can be achieved.

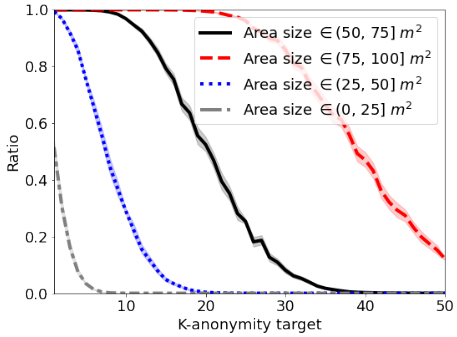

Figure 5: Ratio of users with a desired k-anonymity satisfied, depending on the area size they use to make queries.

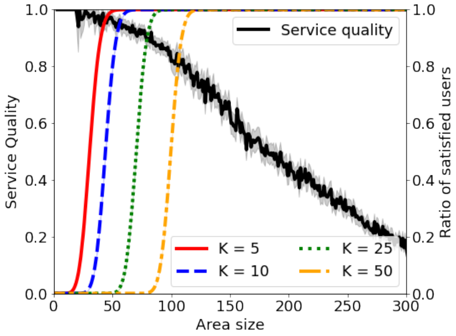

Figure 6: Joint evaluation of Service Quality and Privacy level. The evaluation leverages on the same independent variable, which is the area size selected by the user.

We finally show in Figure 6 the joint evaluation of Service Quality and Privacy. Having defined both on terms of common variables, it is then possible to assess and compare them at the same time, evaluating how one affects the other one. As we have shown earlier, the size of the area sent by the user is the key variable, and we leverage on it also in Figure 6 to highlight the differences between the two aspects taken into account. What it is then possible to do is to observe that in the considered scenario the maximum privacy level is achieved beyond $100m^2$ of area sent, while the Service Quality certainly drops even more as the area sent increases. In other words, evaluating such a system with our proposal allows the user to see that there is no benefit in sending too large areas in their queries, since the Service Quality drops significantly with no direct benefit on the Privacy level.

7 CONCLUSION

In this paper we have presented a novel methodology to jointly evaluate Service Quality and Privacy issues in Ubiquitous systems.

Our proposed methodology allows to describe Service Qualities and Privacy issues in terms of Data Meta Descriptors, which can then be matched together to evaluate them both at the same time. We have presented the overall structure which realizes this vision, and we have provided a candidate implementation for evaluating a specific Service Quality and the Privacy level of users. These implementation can then be later used to evaluate further Ubiquitous services, by simply defining the Data Meta Descriptors available in each part of the system. Our results also indicate that this evaluation can immediately highlight potential issues, and guide users towards selecting their own privacy preferences.

Future work on this topic include the implementation of novel Service Qualities and Privacy issues, to extend the system, and its evaluation on real candidate applications, as well as the definition of mitigation techniques which can be applied to the data to reduce the user exposure to privacy threats.

REFERENCES

[1] Nadeem Ahmed, Regio A. Michelin, Wanli Xue, Sushmita Ruj, Robert Malaney, Salil S. Kanhere, Aruna Seneviratne, Wen Hu, Helge Janicke, and Sanjay K. Jha. 2020. A Survey of COVID-19 Contact Tracing Apps. *IEEE Access* 8 (2020), 134577–134601.

[2] Florian Alt and Emanuel von Zezschwitz. 2019. Emerging Trends in Usable Security and Privacy. *i-com* 18, 3 (2019), 189–195.

[3] Luca Bedogni and Luciano Bononi. 2019. Vehicular Route Identification Using Mobile Devices Integrated Sensors. In *2019 IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops).* 820–825. https://doi.org/10.1109/PERCOMW.2019.8730753

[4] Luca Bedogni and Giacomo Cabri. 2020. Identification of Social Aspects by Means of Inertial Sensor Data. *Information* 11, 11 (2020).

[5] Luca Bedogni, Shakila Khan Rumi, and Flora D. Salim. 2021. Modelling Memory for Individual Re-Identification in Decentralised Mobile Contact Tracing Applications. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 5, 1, Article 4 (mar 2021), 21 pages.

[6] C. Bettini, S. Kanhere, M. Langheinrich, A. Misra, and D. Reinhardt. 2020. Is Privacy Regulation Slowing Down Research on Pervasive Computing? *Computer* 53, 06 (jun 2020), 44–52.

[7] Supriyo Chakraborty, Kasturi Rangan Raghavan, Matthew P. Johnson, and Mani B. Srivastava. 2013. A Framework for Context-Aware Privacy of Sensor Data on Mobile Systems. In *Proceedings of the 14th Workshop on Mobile Computing Systems and Applications* (Jekyll Island, Georgia) *(HotMobile ’13).* Association for Computing Machinery, New York, NY, USA, Article 11, 6 pages.

[8] Saksham Chitkara, Nishad Gothoskar, Suhas Harish, Jason I. Hong, and Yuvraj Agarwal. 2017. Does This App Really Need My Location? Context-Aware Privacy Management for Smartphones. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 1, 3, Article 42 (sep 2017), 22 pages.

[9] Anind K. Dey. 2001. Understanding and Using Context. *Personal and Ubiquitous Computing* (2001).

[10] Pardis Emami-Naeini, Sruti Bhagavatula, Hana Habib, Martin Degeling, Lujo Bauer, Lorrie Faith Cranor, and Norman Sadeh. 2017. Privacy Expectations and Preferences in an IoT World. In *Proceedings of the Thirteenth USENIX Conference on Usable Privacy and Security* (Santa Clara, CA, USA). USENIX Association, USA, 14 pages.

[11] Paul Grace and Mike Surridge. 2017. Towards a Model of User-Centered Privacy Preservation. In *Proceedings of the 12th International Conference on Availability, Reliability and Security* (Reggio Calabria, Italy) *(ARES ’17).* Association for Computing Machinery, New York, NY, USA, Article 91, 8 pages.

[12] Timo Jakobi, Sameer Patil, Dave Randall, Gunnar Stevens, and Volker Wulf. 2019. It Is About What They Could Do with the Data: A User Perspective on Privacy in Smart Metering. *ACM Trans. Comput.-Hum. Interact.* 26, 1, Article 2 (jan 2019), 44 pages.

[13] Xiaodong Jiang and J.A. Landay. 2002. Modeling privacy control in context-aware systems. *IEEE Pervasive Computing* 1, 3 (2002), 59–63.

[14] Johanna Johansen. 2019. Making GDPR Usable: A Model to Support Usability Evaluations of Privacy. In *Privacy and Identity Management.*

[15] Bin Liu, Jialiu Lin, and Norman Sadeh. 2014. Reconciling Mobile App Privacy and Usability on Smartphones: Could User Privacy Profiles Help?. In *Proceedings of the 23rd International Conference on World Wide Web* (Seoul, Korea) *(WWW ’14).* Association for Computing Machinery, New York, NY, USA, 201–212.

[16] Franco Loi, Arunan Sivanathan, Hassan Habibi Gharakheili, Adam Radford, and Vijay Sivaraman. 2017. Systematically Evaluating Security and Privacy for Consumer IoT Devices. In *Proceedings of the 2017 Workshop on Internet of Things Security and Privacy* (Dallas, Texas, USA) *(IoTS&P ’17).* Association for Computing Machinery, New York, NY, USA, 1–6.

[17] A. Machanavajjhala, J. Gehrke, D. Kifer, and M. Venkitasubramaniam. 2006. L-diversity: privacy beyond k-anonymity. In *22nd International Conference on Data Engineering (ICDE’06).* 24–24.

[18] Federico Montori and Luca Bedogni. 2020. A Privacy Preserving Framework for Rewarding Users in Opportunistic Mobile Crowdsensing. In *2020 IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops).* 1–6.

[19] Elham Naghizade, Lars Kulik, Egemen Tanin, and James Bailey. 2020. Privacy- and Context-Aware Release of Trajectory Data. *ACM Trans. Spatial Algorithms Syst.* 6, 1, Article 3 (jan 2020), 25 pages.

[20] Sashank Narain, Triet D. Vo-Huu, Kenneth Block, and Guevara Noubir. 2016. Inferring User Routes and Locations Using Zero-Permission Mobile Sensors. In *2016 IEEE Symposium on Security and Privacy (SP).* 397–413.

[21] Reza Shokri, George Theodorakopoulos, Carmela Troncoso, Jean-Pierre Hubaux, and Jean-Yves Le Boudec. 2012. Protecting Location Privacy: Optimal Strategy against Localization Attacks. In *Proceedings of the 2012 ACM Conference on Computer and Communications Security* (Raleigh, North Carolina, USA) *(CCS ’12).* Association for Computing Machinery, New York, NY, USA, 617–627.

[22] Garfinkel Simson and Lipford Heather Richter. 2014. Usable Security: History, Themes, and Challenges. *Synthesis Lectures on Information Security, Privacy, and Trust* 5, 2 (2014), 1–124.

[23] Latanya Sweeney. 2002. k-Anonymity: A Model for Protecting Privacy. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems* 10, 05 (oct 2002), 557–570.

[24] Tidiane Sylla, Mohamed Aymen Chalouf, Francine Krief, and Karim Samaké. 2020. Towards a Context-Aware Security and Privacy as a Service in the Internet of Things. In *Information Security Theory and Practice*, Maryline Laurent and Thanassis Giannetsos (Eds.). Springer International Publishing, Cham, 240–252.