Adaptive Information-Sharing for Privacy-Aware Mobile Social Networks

Igor Bilogrevic
EPFL-IC-LCA1
Switzerland
igor.bilogrevic@epfl.ch

Kévin Huguenin
EPFL-IC-LCA1
Switzerland
kevin.huguenin@epfl.ch

Berker Agir
EPFL-IC-LCA1
Switzerland
berker.agir@epfl.ch

Murtuza Jadliwala
Wichita State University
USA
murtuza.jadliwala@wichita.edu

Jean-Pierre Hubaux
EPFL-IC-LCA1
Switzerland
jean-pierre.hubaux@epfl.ch

ABSTRACT

Personal and contextual information are increasingly shared via mobile social networks. Users' locations, activities and their co-presence can be shared easily with online "friends", as their smartphones already access such information from embedded sensors and storage. Yet, people usually exhibit selective sharing behavior depending on contextual attributes, thus showing that privacy, utility, and usability are paramount to the success of such online services. In this paper, we present SPISM, a novel information-sharing system that decides (semi-)automatically whether to share information with others, whenever they request it, and at what granularity. Based on active machine learning and context, SPISM adapts to each user's behavior and it predicts the level of detail for each sharing decision, without revealing any personal information to a third-party. Based on a personalized survey about information sharing involving 70 participants, our results provide insight into the most influential features behind a sharing decision. Moreover, we investigate the reasons for the users' decisions and their confidence in them. We show that SPISM outperforms other kinds of global and individual policies, by achieving up to 90% of correct decisions.

Author Keywords

Information-sharing; Decision-making; Machine Learning; User study; Privacy.

ACM Classification Keywords

H.5.2 Information interfaces and presentation (e.g., HCI): Miscellaneous.

INTRODUCTION

Mobile social networks are becoming extremely popular. As for 2013, more than 250 million people use their smartphones in order to get the latest updates from their favorite social networks[1]. Having access to users' personal data and physical context (through an increasing number of embedded sensors), mobile devices represent a simple means to quickly share information with others; location and photos are just two examples of data that can be easily shared. In addition to the user-triggered sharing decisions, applications such as FourSquare and the now-closed Gowalla enable users to configure their smartphones to share their location and co-presence automatically. With a small set of default information-sharing policies, users have the possibility to adjust the settings in order to match their sharing behaviors with their privacy concerns.

Prior studies on sharing behavior in mobile social networks have investigated the issues related to contextual information-sharing [4, 18, 19, 23]. By analyzing people's sharing behaviors in different contexts, they show that it is possible to determine the features that most influence users' sharing decisions, such as the identity of the person that is requesting the information and the current location [23]. For instance, tools such as the location-sharing systems Locaccino [24] and PeopleFinder [18] have been used to gain significant insight into the benefits of providing users with the ability to set personnal sharing policies. Two recurrent findings in UbiComp studies are that (i) users are not particularly good at effectively articulating their information-sharing policies (compared to their actual behavior) [18] and (ii) that sharing policies evolve over time [18, 24].

In order to overcome these two issues, machine learning techniques have been applied to improve to some extent the decision-making process [8, 9, 18]. The advantage of such systems is that they can decide in a (semi-)automatic fashion whether or not to share information. Most existing schemes, however, enable users to share only a specific kind of information (e.g., location). Moreover, they only make binary decisions on whether to share the requested information. In particular, this last issue is often mentioned as a crucial catalyst for overcoming concerns related to privacy [21] and to a more open, sharing behavior.

In our work, we perform a comprehensive study of

UbiComp'13, September 8–12, 2013, Zurich, Switzerland.

ACM 978-1-4503-1770-2/13/09...$15.00.
http://dx.doi.org/10.1145/2493432.2493510

[1]Social networking statistics, http://www.statisticbrain.com/social-networking-statistics/

information-sharing in mobile social networks, by tackling, all at once, the issues related to context, user-burden trade-offs and privacy. We introduce SPISM, a novel information-sharing system (implemented on Android) that decides, in a (semi-)automatic fashion, whether or not to share information (and the level of detail of the information to be shared) with other users or services, based on contextual features and past behavior. The decision-making core is supported by an active learning method that enables SPISM to either decide automatically – whenever the confidence in the decision is high enough – or to rely on the user's input otherwise. SPISM works with any existing (mobile) social network and can be used transparently by users, as it can operate at the operating system level, filtering all requests for personal information and replying according to the user's behavior.

The contribution of this work is three-fold. First, we develop a novel information-sharing system (SPISM) for (semi-)automatic decision-making in mobile social networks: It enables users to share different types of information (location, activity and co-presence of other people) with other users or services in a privacy-aware fashion. Second, we conduct a personalized online study involving 70 participants where, in addition to collecting data about their sharing behaviors, we provide insight into two other crucial factors in UbiComp studies [3]: The *reason* behind a decision to share and the confidence that the user has in her decision. Third, we evaluate SPISM with respect to the amount of training data (provided by the user) and its performance, and compare it against two policy-based mechanisms. Our results show that SPISM significantly outperforms both the individual user-privacy policies and several consolidated ones that are based on statistical analysis [2, 23], and it achieves up to 90% of correct sharing decisions. We also demonstrate the advantages of active learning techniques in our setting.

RELATED WORK

A substantial research effort has been made on the topic of privacy and information sharing in mobile social networks, notably with respect to the attitudes of people when sharing static and contextual data with other peers. The studies that are most related to our work can be grouped, from a high-level perspective, into two categories: (i) contextual information sharing and privacy [4, 19, 23] and (ii) machine learning for information sharing [1, 8, 9, 13, 17, 18].

Contextual Information Sharing and Privacy

Smith et al. [19] provide an early investigation on technologies that allow people to share their contextual information, such as location, in mobile social networks. In addition to allowing users to manually decide when to share their location with others, the authors implemented a system called *Reno* that can automate the process based on a set of pre-defined regions. By allowing Reno to automatically send notifications whenever the user entered or exited such regions, the authors show that there is both a value and a cost associated with automatic information disclosure. In particular, they show that static rules for location sharing in pre-defined regions are ineffective in accurately expressing the users' actual behavior when other contextual elements change, such as the time of the day or the day of the week. By taking into account such limitations in our work, we consider a wide set of contextual features (discussed in the "SPISM Information-Sharing Platform" section) in order to increase the flexibility of the decision-making process.

More recently, Toch et al. [23] study the effect of the type of locations visited by the users on their willingness to share them with others. By considering simple statistical models that take into account factors other than the geographic location, the authors showed that the semantic category of the location being shared (such as a shopping center or a hospital) and the social group of the person asking for the location are significant factors in deciding whether to share the location. These results support earlier efforts [2, 11, 18] in providing a set of contextual features that have a statistically significant impact on the location-sharing behavior of mobile users. We use these results for our application when defining initial universal sharing policies, and will describe them in the "Evaluation" section.

In an attempt to capture the cost of mistakenly revealing a location due to ineffective sharing policies, in addition to sharing preferences, Benisch et al. [4] compare simple access control policies (white lists) to more sophistacated ones (based on time, day and location). They found out that (i) the accuracy of the sharing policies increases with their complexity (or flexibility), and that (ii) the accuracy benefits are the greatest for the highly sensitive information. This suggests that the notion of the cost of mistakenly revealing information to unauthorized parties (in particular contexts) is an important factor in designing and optimizing automated information-sharing mechanisms.

Wiese et al. [25] investigate the effect of physical and perceived social closeness on people's willingness to share information with others. Among the main results of the study, the authors show that social closeness and the frequency of communication are better predictors of sharing than physical proximity. Moreover, these two factors were also shown to have a capacity to predict sharing better than the social groups of the people asking for the information. Thus, the authors suggest that automatic methods for inferring social closeness could be suited for accurate information-sharing decisions more than physical co-location, in the case automated mechanisms (such as in [10, 14, 20, 22]) are envisaged.

Machine Learning and Information Sharing

Whereas studies on information-sharing attitudes and privacy shed light on the behavior of people and the factors that influence their decisions, they are mostly concerned about understanding the causes and effects of such behavior. Meanwhile, there has been a substantial effort in devising methods that help and nudge the users to make information-sharing decisions, or even make decisions on their behalf. We present some of these methods, including both supervised and unsupervised approaches for decision-making.

In [18], Sadeh et al. compare the accuracy of user-defined sharing policies with an automated mechanism (case-based reasoner) and a machine learning approach (random forests),

showing that these approaches have an accuracy better than the user-defined policies. Owing in part to the greater flexibility of the supervised machine-learning approaches compared to the more coarse-grained user-defined policies, the automated methods also benefited from the fact that users appeared to not be able to create sharing rules consistent with their own choices. On the contrary, the feedback provided by the users to the machine-learning methods did however appear to be consistent with their actual sharing behavior, which helped the automated methods to achieve better accuracy results. We include the user feedback in our learning mechanism and use it to adapt the automated decisions to the user behavior that can change over time.

Unsupervised or semi-supervised methods, which reduce the initial setup burden of the default sharing policies for each user, are investigated in [8, 9]. For instance, Danezis [8] proposes a method for automatically extracting privacy settings for online social networks; the method is based on the notion of a limited proliferation of information outside of a given social context. The proposed method, which determines cohesive groups of users where users belonging to a group have stronger ties to the users outside of the group, shows promising results on a limited set of evaluation samples. This study also shows that the social groups, and especially methods for their automated extraction, are a key factor to sharing private information in social networks. Our work uses both the Facebook social graph and our system's contacts list to automatically extract social groups or communities and uses them to relieve the user from the burden of manually assigning people to different social groups.

Fang and LeFevre [9] propose a novel approach to the inference and definition of access control policies for personal information on online social networks. They enable the supervised learning mechanism to learn the sharing preferences of a user by asking her a limited number of questions about her sharing behavior with some of her friends; these specific friends are the most "informative", i.e., those for which the classifier is most uncertain about. The authors show that their approach of iteratively asking questions about the most uncertain case (active learning with uncertainty sampling) reduces the effort required by the users and maintains a high accuracy compared to the ground truth (based on a 45-user study on Facebook). Active learning is a feature that we exploit in our application as well. Moreover, we allow users to update their sharing decision *a posteriori*, meaning that users are able to change their decision after it has been made; the application then learns from this new decision and takes it into account the next time the same context appears again.

Bigwood et al. [5] evaluate different machine learning algorithms for information sharing in terms of information over-exposure and correct decisions. Although their work is focused exclusively on binary (yes/no) location-sharing, the authors provide a machine-learning-based determination of the most influential features for the sharing decisions; moreover, they take into account cost-sensitive classifiers to reduce over-exposure. We believe this to be a promising direction to explore, and we will evaluate the effects of cost-sensitivity and granularity on the performance of our decision-making framework in the follow-up of this work.

THE SPISM INFORMATION-SHARING PLATFORM

In this section, we describe the functionality, the operating principle, the architecture and the design of the SPISM information-sharing platform.

In order to better understand the following, we need to distinguish between two different kinds of subscribers to SPISM: (i) the *requester*, who wants to know something about other subscribers by sending information requests, and (ii) the *target*, who receives requests for information.

The SPISM platform is composed of the *SPISM application*, that runs on mobile devices (as for now it is implemented only for the Android platform), and the *SPISM Information Sharing Directory (ISD)*, that runs on a dedicated server.

Overview

The SPISM application enables subscribers, who can be users, third-party online services or mobile apps, to request information about other subscribers. The information that can be requested includes contextual data (the geographic location and the wireless identifiers of physically co-located devices) and the time-schedule availability. The geographic location is determined by processing data obtained from the embedded GPS sensor (if available) or by WiFi tri-lateration (which relies on the Google localization service). The list of devices that are physically co-located with the target subscriber is obtained through periodic scans of the Bluetooth and WiFi interfaces. If a MAC address in the vicinity of the target is a known MAC address (there exist an entry associated with a subscriber in the contact list of the target), the name of the contact is displayed. Finally, the schedule availability is obtained from the subscriber's calendar (accessed through the on-device calendar application). Subscribers can specify a level of detail for the requested information: low, medium or high. The information sent by the target subscriber is provided with a level of detail lower or equal to the requested level. For the location, the coordinates are truncated; for the neighboring devices, the presence (i.e., some devices/no devices), the number, or the identifiers of the devices are provided; for the schedule availability, the availability (i.e., busy/available), the title or the detailed record of the calendar activity is provided. Figure 1 shows the main application windows, where subscribers can log in and register, request the location, the co-located devices and the availability of their contact, as well as enjoy additional features such as visualizing the past activity and their contacts' list.

System Model

The SPISM platform is composed of the ISD and the subscribers of the service, who can be either users or third-party online services. The roles of the ISD and of the subscribers are as follows:

- **ISD**: Its main purpose is to allow users to discover the current IP addresses of their contacts when they want to send them information requests. The ISD stores the list of registered SPISM subscribers, their credentials, their contact

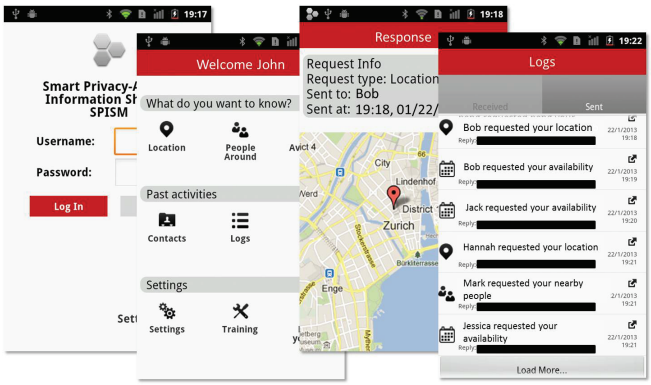

Figure 1. SPISM mobile application interfaces. From left to right, the different windows allow users to register and log in, check other subscribers' current location, the other devices around them, their availability. The subscribers can access other features such as the record of past activity and their contacts' lists.

lists and the MAC addresses of the Bluetooth interfaces of each user's mobile devices. The subscribers interact with the ISD in the registration phase (once per user), during the log-in phase (once per application start), when downloading the contacts lists, when periodically reporting their IP and updating their online status, and when sending information requests to one of their contacts.

- **Subscribers**: A subscriber, either an online service or a mobile user, can be a requester (when she sends queries to another subscriber) or a target (when she receives queries from other subscribers). In order to inform the ISD of her online status, each subscriber connected to the ISD sends periodic keep-alive messages. Requesters can see, at any time, the list of online and offline contacts, and they can choose to send queries to the online subscribers in their contacts list, in order to know their location, the devices around them and their availability. The requests that target subscribers receive and process are based on several features of their current physical and social contexts, including their current location, the time of the day and the people that are currently close by.

To enhance the security of the communications, all messages exchanged between the subscribers and the ISD are encrypted with a public-key certificate obtained from a trusted Certification Authority (CA). In order to protect users' privacy with respect to the ISD, no information requests or replies are tunneled through the ISD. This is a crucial aspect of our platform, as it prevents the service provider from learning the information sent by a subscriber about her location, physical context and availability. A shortcoming of this approach is that the requester knows the IP address of the target, and therefore she may be able to infer the target's coarse-grained location (based on IP-geolocation) and to infer the co-location of multiple targets if they share the same public IP (when connected to an access point using Network Address Translation for example). Conversely, the target may know the IP address of the requester. Note however, that a user can conceal her IP address by making use of proxies or anonymous networks such as Tor.[2]

Operating Principle

SPISM works as follows. A user first logs in to the ISD with her username and password. She can subsequently report her online status and obtain the online status (and IP addresses) of her contacts from the ISD. In a typical scenario, the user requests some information from one of her (connected) contacts. To do so, the user first chooses the type of information she wants to request, by selecting the corresponding icon in the main window (See Figure 1), and then she selects the target subscriber from the list of her connected contacts. Finally, the user specifies the level of detail for the requested information and the request is prepared and sent directly to the target subscriber's device. If the reply is received within a fixed amount of time (typically a few seconds) it is automatically showed to the user, together with the requested information if shared by the targeted requester (See Figure 1); otherwise, the user is redirected to the main window and she will be notified when the reply is received. At the targeted subscriber's device, the request is processed automatically when it is received: (1) The requested information is stored and (2) the information linked to the request (i.e., the time, the type of information requested and the requester) is combined with various contextual features (periodically collected in the background by SPISM from the various data sources and sensors available on the device) and fed to the information-sharing framework that we describe in detail in the next section. If SPISM can make the decision with enough confidence, based on the target subscriber's past decisions, the request is processed automatically. Otherwise, the target subscriber is notified and asked to decide; Her decision is then stored (note that the target subscriber can postpone her decision). Once a decision is made, it is sent back to the requester together with the requested information if the decision is positive. Before being sent, the requested information is processed to match the level of detail specified by the decision. All the sent and received requests are stored and can be accessed by the user by selecting the corresponding icon in the main window. In particular, the user can audit automatic decisions and correct those she disagrees with (to avoid similar errors in the future).

Decision Making

The SPISM information-sharing decision-making core processes each incoming information request. In order to make the decision, several contextual features are taken into account by the target device. Features such as the identity of and the social ties with the requester, the current location and the activity of the target, the people around the target and the time of the day were extensively studied in the past; several independent pieces of work show (with statistical significance) that they are strongly correlated with the information-sharing behavior of mobile users [2, 6, 7, 19, 23]. With these findings, we incorporated 18 such features in the SPISM decision-making core; the list of all the features we included is shown in Table 1. Due to the different natures of

[2]https://www.torproject.org/

	Feature	Type
Person	Familiarity	Float
	Social tie	Cat.
	User ID	Cat.
Service	Service category	Cat.
What?	Request type	Cat.
	Details	Float
Location	Latitude	Float
	Longitude	Float
	Semantic location	Cat.

	Feature	Type
When?	Time	Int.
	Weekday	Cat.
	Daytime	Cat.
	Activity	Cat.
With whom?	Neighbors	Int.
	Neighbors Type	Cat.
Last interact.	Time last request	Float
	Details last request	Float

Table 1. Features used by the SPISM machine learning framework to decide whether or not to share information and with what accuracy.

the features, some of them are defined as categorical (they are in a finite and pre-defined set of values, such as the social ties with the requester) or numerical (floating or integer values for the time and location coordinates).

Some of these 18 features can be extracted from the request itself or the target mobile device, such as the time, the current schedule availability or the requester ID, whereas other features require more information, e.g., the social ties with the requester and the semantics of the current location of the target subscriber. To obtain such information, SPISM takes advantage of the existing social networks, such as Facebook, and other data available on the phone (e.g., call logs). In addition, other third-party services (such as Google Maps, OpenStreetMap and the Android application store, i.e., Google Play) are used to obtain more information about the location and type of application (in the case where the requester is a mobile application). In some cases, the extraction of the features requires access to the sensors embedded on the device; GPS and Bluetooth scans usually require a non-negligible amount of time and resources [16], and a per-request access to such sensors can drain the battery. For this reason, some time- or energy-consuming features (such as the GPS coordinates and Bluetooth MAC addresses of the nearby devices) are obtained periodically and cached, so that they can be polled by the device at any time instant without incurring resource-consuming operations. Note that the location, the list of nearby devices and the schedule availability are all used to make the decision and to be shared.

After all 18 features have been extracted from the request and determined from the context, they are aggregated into a feature vector and fed to a classifier. The output space of the classifier comprises four different classes that encode whether the information is shared and, if yes, the corresponding level of detail, specifically "No", "Yes (low)", Yes ("medium") and "Yes (high)". SPISM makes use of a Logistic classifier implemented in the WEKA[3] Android library.

STUDY AND DATA COLLECTION

In order to better understand how users share information and to evaluate the efficacy of the SPISM framework with respect to sharing decisions, we ran a user study in early 2013. The study consists of an online survey that puts the participants in realistic, personalized and contextual UbiComp sharing scenarios where they are asked to answer a set of questions regarding their willingness to share private information, the confidence in and reason for their decisions.

Participants and Remuneration

We recruited people directly from four large university campuses (in the US, Canada and Europe), and indirectly via the Amazon Mechanical Turk platform (MTurk)[4]. The latter allowed us to draw participants from a pool of non-student population, in order to limit the bias towards academic and student behaviors. To advertise our study, we used dedicated mailing-lists and we ran a media campaign through Facebook, LinkedIn, Google+ and official university websites, coordinated by our academic media office. We screened participants according to the following prerequisites: (i) aged between 18 and 80 years, (ii) with an active Facebook account with at least 50 friends and (iii) uses a smartphone. Such criteria were selected so as to sample people that are active in social networks and are aware of the information-sharing possibilities linked to the use of smartphones. Furthermore, we screened the MTurk workers who could access our survey based on their past Human Intelligence Task (HIT) approval rate (>95%) and the number of past approved HITs (>100). This was only a preliminary step for preventing non-serious and inexperienced MTurk workers from accessing our survey.

The survey requires access to private information of the participants (such as names of their friends on Facebook[5]) and it demands a significant amount of time (40 - 60 minutes). To provide incentives for the completion of the survey, we implemented two separate reward schemes: (i) the chance for one participant to win an Apple iPad and (ii) a fixed amount of money (US$4.5/HIT [12]). The first option was proposed to the participants recruited at the universities and through the academic media, whereas the second option was offered to the workers of the Amazon Mechanical Turk. We chose not to offer the second option to the academic participants due to our experience gained from previous on-campus studies: It appeared that the motivation for financial rewards was lower than for the possibility of winning a popular gadget.

Online Survey

We structured our survey in five parts: With a total of 94 questions, the first 19 are fixed (the same for each participant) and the last 75 are personalized (based on each participant's Facebook friends). In the very first part, the participants were required to log in to their Facebook account and grant our application access to their friend list.

[3]http://www.cs.waikato.ac.nz/ml/weka/

[4]https://www.mturk.com/mturk/welcome

[5]Before beginning the survey, the participants are informed that they would need to reveal the names of their Facebook friends for the purpose of this study. They approve a data retention and processing agreement, informing them that all data collected in our study is used solely for the purpose of our academic research project, and that we will not disclose or use it in any other way than what explicitly mentioned. Once the survey is completed, the name of the Facebook friends are replaced with anonymous identifiers.

In the first 15 questions, the participants were asked about their demographics, technology usage and privacy attitudes, in particular with respect to online social networks.
In the next question (16), the participants were asked to assign some of their friends to social groups, and we presented them with five distinct categories (based on [25]): (1) school colleagues, (2) friends, (3) family members, (4) work colleagues and (5) acquaintances. Each participant could assign one Facebook contact to at most one category. It is possible, however, that one such contact is a member of several categories (a school colleague that she works with currently). In this case, the participants were instructed to assign the contact to the most appropriate category.
In questions 17 through 19, the participants were asked to enter a set of information-sharing rules in free-text. The sharing rules are entered as a set of logical expressions that are based on the following *features*: (1) the participant's current location, (2) people nearby, (3) social group of the requester, (4) time of the day and (5) weekday/weekend. They can put *conditions* on these features (such as $=, <, >, \neq, \in$ or categorical values). For example, a location-sharing rule could be defined as:

"I am at a *friend's place* **AND** with *acquaintance* **AND** the requester is a *work colleague*: do not share"

In the last 75 questions, the participants were presented with sharing scenarios and they were asked to decide whether they want to share the specific information in the given context, their confidence in the decision and the level of detail. A typical scenario is "Would you share your *location* with *John* on *Saturday* at *11:PM*, assuming you are at *an event* with *work colleagues*?" (where the requester name is chosen from the participant's Facebook friends and the other features are chosen at random).

Depending on their answers ("Yes", "No" and "Uncertain") to the questions in this part, participants were presented with sub-questions. More specifically, "Yes" and "No" answers were followed by a set of additional questions asking the participants about the confidence in their decisions (i.e., "not so confident", "confident", "very confident") and the features that influenced the most their decision (i.e., "requester", "day of the week", "time", "location" or "neighboring people"). For "Yes" answers, the participants were also asked about the level of detail of the shared information ("low", "medium" or "high"). Similarly, "Uncertain" answers were followed by sub-questions regarding the reasons for being uncertain, such as a conflict between some features (in this case, the participant can specify the features that motivates her the most to share and to not share, and then specify in free text the reason they conflict) or simply a lack of information (in this case the participant can specify which information would have helped her reach a decision).

In order to detect sloppy answers (e.g., random answers or bots), we included a number of "dummy" questions that require human understanding to be correctly answered [12, 15]. These are questions such as simple computations (e.g., "$3 + 4$") or general-knowledge questions (e.g., "How many days are there in one week?"). Based on the answers to these questions and on the survey timing data (explained below), we ruled out dishonest participants from the dataset.

General Statistics and Validation

A total of 194 participants took part in our survey. 78 (40%) of them did not complete it, leaving 116 (60%) complete questionnaires. Out of these, 56 (48%) came from the university advertisement campaign (UNI) and 60 (52%) were recruited via MTurk. The average age of all the respondents is 27y±7 (Mturk avg. 31y±6, UNI avg. 25y±6), and 74% of them are male. 42% of all participants are students, 25% work in the IT industry and 8% in the education sector. It took 44±15 minutes on average to complete the survey (MTurk avg. 42 minutes, UNI avg. 47 minutes). We observed a sharp contrast, with respect to privacy concerns, between the two groups of participants: Most MTurk participants were not, or slightly, concerned about their privacy whereas most UNI participants were concerned about it.

Based on internal survey tests and detailed timing statistics, only the questionnaires that meet the following four validation criteria were retained.

- All answers to the dummy questions are correct;
- At least one different Facebook friend is assigned to each of the 5 social groups;
- The survey completion time is greater than 30 minutes.
- At least three of the following four timing conditions are met[6]: (1) Facebook friends assignment to groups time >5 minutes, (2) location sharing scenarios time >4 minutes, (3) activity sharing scenarios time >4 minutes, (4) nearby people sharing scenarios time >4 minutes.

All participants correctly answered the dummy questions. Based on timings, 46 (40%) of them were ruled out and 70 (60%) were kept for the analysis (33 MTurk and 37 UNI). The demographics remained mostly unaltered.

ANALYSIS AND EVALUATION

In this section, we present three sets of results. First, using descriptive statistics of the survey questionnaire, we discuss the effect on the sharing decisions of different contextual features, of the requester, of the information type, and the main reasons behind the decisions. Second, we compare the performance of the SPISM automated decision-making process against that of the users' own policies and an established default policy. Third, we discuss the effects of the increase of user-involvement on the performance of SPISM, by using active learning with different confidence thresholds.

Survey Statistics

Based on the survey data, we computed the proportion of "Yes/No/Uncertain" decisions for the different values of each contextual feature we considered, such as the participant's current location, the social group of the requester, the time

[6]These timing conditions were determined based on the observed timing distributions among all participants and on sample executions performed by test users.

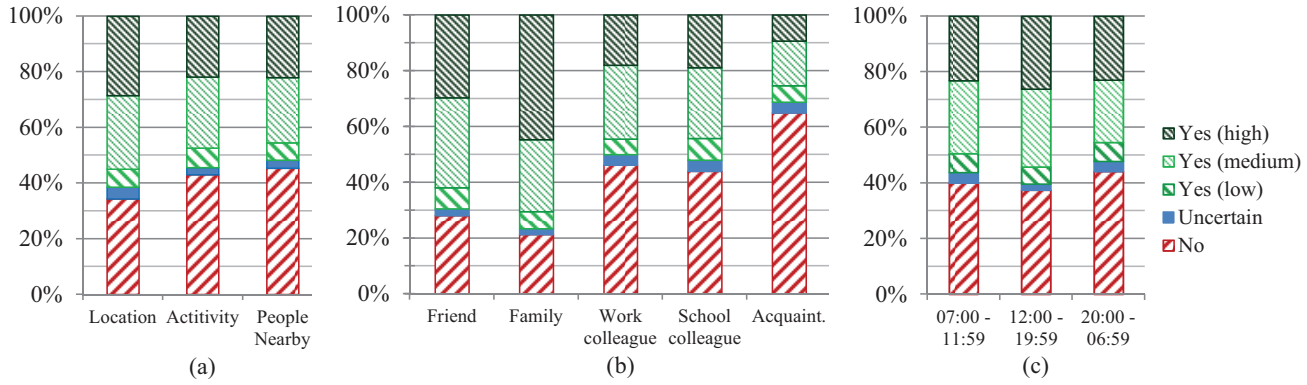

Figure 2. Histograms of the information-sharing decisions by (a) information type, (b) social group of the requester and (c) the time of the day.

of day, day of week, and the type of information requested. We found that the two that have the largest effect on the decision are the social group of the requester and the type of information that is being requested.

Regarding the type of information being asked, Figure 2a shows that users disclose their location in 64% of the cases (the sum of the "yes (low)", "yes (medium)" and "yes (high)" bars, aggregated over the 70 participants and for all the 25 location-sharing questions – out of the 75 questions – that is a total of 1,750 answers), and only 8% of the time at a coarse granularity ("Yes (low)"). The information about activity and people nearby is disclosed 50% of the time. People tend to be slightly more willing to share their location than to share other information[7]: Location, contrary to the activity and the co-presence of other people, is widely shared information in most mobile social networks. In addition, this was confirmed by self-reported privacy concerns about information sharing on OSNs (not shown in the paper).

Figure 2b shows the percentage of disclosure of information based on the social ties with the requester. We can see that, in accordance with previous UbiComp studies, there are substantial differences[7] between the close ties ("family" and "friend") and the more distant ones ("acquaintances" and "colleagues"). For instance, the close ties are granted access to any type of information (70%-80%) more than twice the times compared to the more distant ones (30%). Moreover, the level of detail of the shared information is much higher for the close ties (up to 45% of "yes (high)") compared to the distant ones (down to 8%). In fact, the proportion of "Yes (low)" and "Yes (medium)" does not vary significantly. Hence, the results indicate that users tend to exhibit a more tailored sharing behavior depending on the type of information, the social ties and closeness with the requester [25]. As illustrated in Figure 2c, the time at which the request is sent does not substantially influence the decision: users are slightly less willing to share in the evening but exhibit the same behavior in the morning as in the afternoon[7]. Our findings are aligned with those obtained in [5], where the time of day and the location do not appear to be influential factors when sharing personal information such as location, as opposed to the type of social ties with the requester.

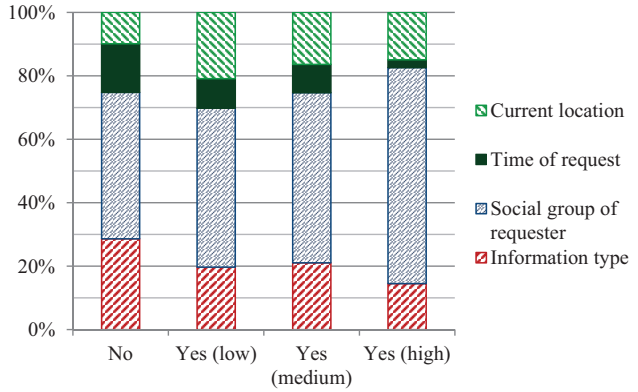

Figure 3. Histograms of the main reasons for (not) sharing.

We also looked at the reasons for (not) disclosing information and at the users' confidence in their decisions. First we observe that the social ties with the requester is by far the most frequent reason for sharing (or not) information (45%-67%), followed by the type of information (15%-28%) and the current location (11%-21%). Second, we see again that the higher the level of detail (Figure 3), the more important the social ties with the requester (on average). Unsurprisingly, the confidence that the participants have in their decision (Figure 4) is lower for the intermediate level of detail: It can be observed that the proportion of "Very confident" is significantly lower for "low" and "medium" levels of detail than for "No" and "Yes (high)". In addition, the proportion of "Not so confident" is more than doubled for the most borderline decision, i.e., "Yes (low)". This could be explained by the fact that users try to minimize the risk by limiting the level of detail when their confidence is low.

STATIC POLICIES

We compared the performance of our SPISM decision framework with two other policy-based approaches. For the following comparisons, we used 10-fold cross validation and a logistic regression binary classifier. In order to be consistent with the policy-based approaches, we only compare the

[7]With statistical significance, based on global and pair-wise χ^2 homogeneity tests with $p < 0.01$.

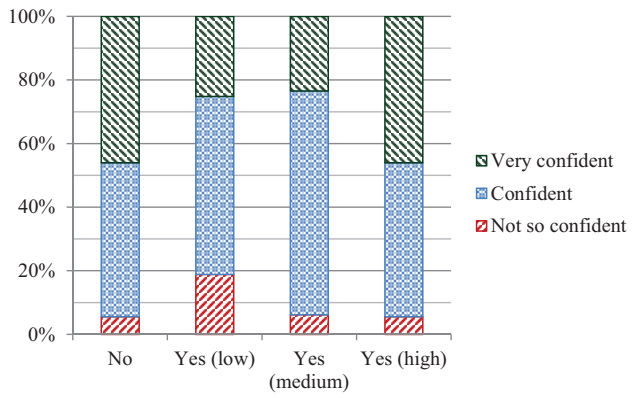

Figure 4. Histograms of the users' confidence in their decisions.

binary ("Yes/No") decisions here as the participants were instructed to only specify share/not share policies in the survey. The first policy-based approach, called *AT studies*, is inspired from the findings presented in [2, 23], and is derived by the following two rules:

1. Do not share any information while sleeping (12 AM - 6 AM) or eating (12 PM - 1 PM).
2. Do not share any information when you are around people that are not your family members or friends, except when you are at an event.

The second policy-based approach is derived from the individual policies that each participant specified in free text in the survey. We selected a random sample of 19 participants and we manually transposed their free-text policies to a format suitable to be evaluated against their own decisions. The participants specified between 1 and 15 policies (avg. 6.9).

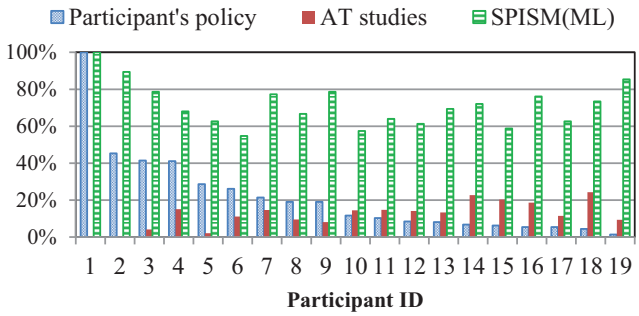

Figure 5. Histograms of the proportion of correct sharing decisions for three different sharing policy approaches. The AT studies' policies are derived from [2, 23], the participants' individual policies are derived from their free text answer in the survey and the SPISM approach is based on machine learning (without active learning).

The results of the three-way comparison are shown in Figure 5 where the results are sorted in descending order, based on the performance of the participant's individual policies. First, we can observe that the SPISM machine-learning approach consistently outperforms the other two approaches (this holds for all users when compared only to the AT policies defined earlier). The SPISM performance rate is between 53% and 100%, with an average of 71%. Compared to the participant's policy (avg. 22%) and the AT studies (avg. 12%), SPISM is significantly better at adapting itself to the user's sharing behavior. We also observe that usually where the participants' own policies correctly represent their actual behavior, the AT policies exhibit the worst performance (left side of Figure 5). The inverse appears to be true as well, as the policies inspired by the AT studies perform better for the participants whose own policies do not particularly match their actual behavior. This points out an interesting question, which is outside of the scope of this work: Are people who are not able to articulate well their sharing policy better suited to not trying to modify the default policies at all?

For the individual policies, we also observed the correctness of the decisions as a function of the number of policies, and found that a small number of policies (1-5) achieved up to 41% of correct decisions, followed by a slightly better performance for the number of policies between 6 and 9 (up to 45%), and then a much worse performance (up to 28% of correct decisions) for the highest number of policies (10 - 15). This suggests that there is an advantage in having a moderate number of sharing policies (up to 9) but not higher; With a larger number of policies, the risk of having overlapping but contradicting policies is higher, which could result in a worse overall performance.

MACHINE LEARNING

In order to assess the potential of (semi-)automatic information-sharing decision making, which constitutes the core of SPISM, we evaluate the performance of a logistic classifier in predicting the users' sharing decisions. To do so, we use the survey data comprised of 75 scenarios for each of the 70 participants: Each scenario corresponds to a feature vector and the decision made by the participant constitutes the ground truth. We considered only the "Yes" and "No" decisions. We evaluate the performance of the classifier in terms of the proportion of correct predictions (i.e., that match the user's decision), the proportion of cases where the information is incorrectly shared (whereas the user would have not shared it), thus compromising the user's privacy, and the proportion of cases where the information is incorrectly not shared (whereas the user would have shared it), thus reducing the utility of the system.

Firstly, we consider the case where the users first manually make n decisions to train the classifier, and then the classifier makes the remaining decisions automatically. For several values of n, and for each participant, we compute the average proportions of correct and incorrect decisions following a 20-fold cross-validation approach. For each value of n, we obtain one data point (i.e., a proportion of "correct", "share less", and "share more" decisions) for each user and each fold, that is 1,400 data points. We represent the results across the different users and folds by showing the median, the first and third quartiles, and the 5 and 95-percentiles, as depicted in Figure 6a. It can be observed that the median proportion of correct decisions increases from 60% and reaches 70% for a training set of only 30% of the data, which correspond to $\sim$25 scenarios. The proportion of correct decisions then quickly stabilizes around 74% after approximately 40 decisions (i.e.,

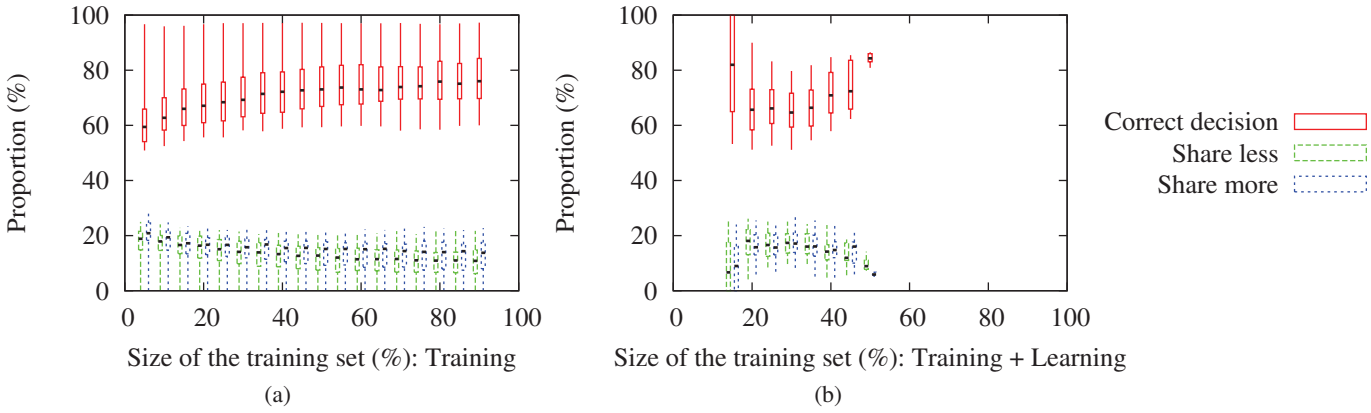

Figure 6. Performance of the machine learning-based decision making algorithm.

∼50% of the data). The third quartile and the 95-percentile show that for more than 25% of the users, the proportion of correct decisions goes up to 80% and for some of them, it is consistently higher than 96%. The proportion of incorrect decisions is evenly distributed between sharing and not sharing the information yet slightly biased towards incorrectly sharing the information. Should a user favor her privacy over the utility of the system, she could assign a higher error-penalty to this type of errors in order to make decisions in a conservative way. Without penalties and active learning, over-sharing happens in 10-20% of the cases, in line with the results reported in [5] and obtained with different classifiers. Note that the size of the training set (represented on the x-axis) represents the burden of the user as she has to manually make the corresponding decisions.

Secondly, we consider the case of active learning in which the user is asked to manually make the decision when the confidence of the classifier is low. The classifier outputs a distribution over the possible decisions; we define the confidence as the normalized entropy of this distribution. The classifier is first initialized with 10% of the data. For each user, we run the active learning-based classifier for several values of the confidence threshold (under which the user is asked to make the decision). Each experiment gives one data point comprised of (1) the proportion of decisions (including the first 10%) the user has to manually make and (2) the proportions of correct and incorrect decisions (among the decisions that are made automatically). In order to represent the data in a form that is comparable to that of Figure 6a, we group the data points in bins of size 5% (on the x-axis as represented in the figure) based on the proportion of manual decisions. Note that the number of data points varies across the different bins. Within each bin, we compute the median and the relevant percentiles. The result are depicted in Figure 6b. It can be observed that active learning outperforms training-only learning in most cases (i.e., for a given number of manual decisions, it provides a higher proportion of correct decisions). The proportion of manual decisions remains lower than 50% which shows that the classifier can make the decision with very high-confidence for at least half of the scenarios. For some users, the proportion of manual decisions remains low (∼20%), regardless of the confidence threshold, and the proportion of correct decisions is high (∼80%). This corresponds to the users whose decisions are highly predictable. With active learning, we observe a significantly improved performance in terms of over-sharing compared to the absence of active learning. We posit that, coupled with cost-sensitive classifiers, active learning can be used to improve the correctness of the sharing decisions while maintaining a significantly lower over-sharing rate.

CONCLUSION AND FUTURE WORK

Mobile social networks allow users to share an increasing number of contextual information, such as their location, their activity and their co-presence with others. To simplify the sharing process and improve usability, the research community has been studying sharing preferences and developing applications that, based on several contextual features, can automate to some extent the sharing process. Machine-learning approaches have been developed and evaluated for specific instances of information (mostly location) or for online social network (without the notion of context).

In this paper, we present and evaluate a novel privacy-preserving information-sharing system (SPISM) that decides in a (semi-)automated fashion whether or not to share different types of contextual information and to what level of detail. Using a personalized online user-study involving 70 participants, we show that SPISM significantly outperforms both individual and general user-defined sharing policies, achieving up to 90% of correct sharing decisions, with only a limited cost for the user in terms of initial setup thanks to active learning. We also show that the system has a slight bias towards incorrectly sharing the information, which could be mitigated by introducing a penalty for those kind of errors. Furthermore, our results provide significant insight into two other crucial aspects of UbiComp studies: The reasons behind a sharing decisions and the participants' confidence in them. We show that the type of the requested information, in addition to the social ties of the requester, is an influential feature in the decision process.

Acknowledgements

We would like to thank Peng Gao, Jens Grossklags, Urs Hengartner, Mathias Humbert, Aylin Jarrah Nezhad, Bradley Ma-

lin, Xin Tang and Romain Tavenard for the fruitful discussions and the anonymous reviewers, as well as our shepherd, for helping us to improve the quality and presentation of the paper.

REFERENCES

1. An, X., Jutla, D., Cercone, N., Pluempitiwiriyawej, C., and Wang, H. Uncertain Inference Control in Privacy Protection. *International Journal of Information Security 8*, 6 (2009), 423431.
2. Anthony, D., Henderson, T., and Kotz, D. Privacy in Location-Aware Computing Environments. *IEEE Pervasive Computing 6*, 4 (2007), 64.
3. Barkhuus, L. The mismeasurement of privacy: using contextual integrity to reconsider privacy in HCI. In *Proc. of ACM CHI* (2012), 367–376.
4. Benisch, M., Kelley, P. G., Sadeh, N., and Cranor, L. F. Capturing Location-Privacy Preferences: Quantifying Accuracy and User-Burden Tradeoffs. *Personal and Ubiquitous Computing 15*, 7 (2011), 679–694.
5. Bigwood, G., Abdesslem, F. B., and Henderson, T. Predicting Location-Sharing Privacy Preferences in Social Network Applications. In *Proc. of AwareCast* (2012).
6. Brush, A. J., Krumm, J., and Scott, J. Exploring End User Preferences for Location Obfuscation, Location-Based Services, and the Value of Location. In *Proc. of ACM UbiComp* (2010).
7. Consolvo, S., Smith, I., Matthews, T., LaMarca, A., Tabert, J., and Powledge, P. Location Disclosure to Social Relations: Why, When, & What People Want to Share. In *Proc. of ACM CHI* (2005).
8. Danezis, G. Inferring Privacy Policies for Social Networking Services. In *Proc. of ACM AISEC* (2009).
9. Fang, L., and Lefevre, K. Privacy Wizards for Social Networking Sites. In *Proc. of ACM WWW* (2010).
10. Hsieh, G., Tang, K. P., Low, W. Y., and Hong, J. I. Field Deployment of IMBuddy: A Study of Privacy Control and Feedback Mechanisms for Contextual IM. In *Proc. of ACM UbiComp* (2007).
11. Mancini, C., Thomas, K., Rogers, Y., Price, B. A., Jedrzejczyk, L., Bandara, A. K., Joinson, A. N., and Nuseibeh, B. From Spaces to Places: Emerging Contexts in Mobile Privacy. In *Proc. of ACM UbiComp* (2009).
12. Mason, W., and Suri, S. Conducting Behavioral Research on Amazons Mechanical Turk. *Behavior Research Methods 44*, 1 (2012), 1–23.
13. Miettinen, M., and Asokan, N. Towards Security Policy Decisions Based on Context Profiling. In *Proc. of ACM AISEC* (2010).
14. Miluzzo, E., Lane, N. D., Fodor, K., Peterson, R., Lu, H., Musolesi, M., Eisenman, S. B., Zheng, X., and Campbell, A. T. Sensing Meets Mobile Social Networks: The Design, Implementation and Evaluation of the CenceMe Application. In *Proc. of ACM SenSys* (2008).
15. Patil, S., Gall, Y., Lee, A., and Kapadia, A. My Privacy Policy: Exploring End-User Specification of Freeform Location Access Rules. In *Proc. of USEC* (2012).
16. Priyantha, B., Lymberopoulos, D., and Liu, J. Littlerock: Enabling Energy-Efficient Continuous Sensing on Mobile Phones. *Pervasive Computing, IEEE 10*, 2 (2011), 12–15.
17. Riva, O., Qin, C., Strauss, K., and Lymberopoulos, D. Progressive Authentication: Deciding When to Authenticate on Mobile Phones. In *Proc. of USENIX Security* (2012).
18. Sadeh, N., Hong, J., Cranor, L., Fette, I., Kelley, P., Prabaker, M., and Rao, J. Understanding and Capturing People's Privacy Policies in a Mobile Social Networking Application. *Personal and Ubiquitous Computing 13*, 6 (2009), 401–412.
19. Smith, I., Consolvo, S., Lamarca, A., Hightower, J., Scott, J., Sohn, T., Hughes, J., Iachello, G., and Abowd, G. Social Disclosure of Place: From Location Technology to Communication Practices. *Pervasive Computing* (2005), 151–164.
20. Tang, J. C., Yankelovich, N., Begole, J., Van Kleek, M., Li, F., and Bhalodia, J. ConNexus to Awarenex: Extending Awareness to Mobile Users. In *Proc. of ACM CHI* (2001).
21. Tang, K., Hong, J., and Siewiorek, D. The Implications of Offering More Disclosure Choices for Social Location Sharing. In *Proc. of ACM CHI* (2012).
22. Tang, K. P., Keyani, P., Fogarty, J., and Hong, J. I. Putting People in Their Place: An Anonymous and Privacy-Sensitive Approach to Collecting Sensed Data in Location-Based Applications. In *Proc. of ACM CHI* (2006).
23. Toch, E., Cranshaw, J., Drielsma, P. H., Tsai, J. Y., Kelley, P. G., Springfield, J., Cranor, L., Hong, J., and Sadeh, N. Empirical Models of Privacy in Location Sharing. In *Proc. of ACM UbiComp* (2010).
24. Toch, E., Cranshaw, J., Hankes-Drielsma, P., Springfield, J., Kelley, P. G., Cranor, L., Hong, J., and Sadeh, N. Locaccino: A Privacy-centric Location Sharing Application. In *Proc. of ACM UbiComp (adjunct papers)* (2010).
25. Wiese, J., Kelley, P., Cranor, L., Dabbish, L., Hong, J., and Zimmerman, J. Are You Close with Me? Are You Nearby?: Investigating Social Groups, Closeness, and Willingness to Share. In *Proc. of ACM UbiComp* (2011).