Delphine Reinhardt, Franziska Engelmann, Matthias Hollick: *Can I Help You Setting Your Privacy? A Survey-based Exploration of Users' Attitudes towards Privacy Suggestions*. In: Proceedings of the 13th ACM International Conference on Advances in Mobile Computing and Multimedia (MoMM), pp. 347-356, 2015.

# Can I Help You Setting Your Privacy? A Survey-based Exploration of Users' Attitudes towards Privacy Suggestions

Delphine Reinhardt
University of Bonn and Fraunhofer FKIE
Bonn, Germany
delphine.reinhardt@cs.uni-bonn.de

Franziska Engelmann
Secure Mobile Networking Lab
TU Darmstadt
Darmstadt, Germany
fengelmann@seemoo.tu-darmstadt.de

Matthias Hollick
Secure Mobile Networking Lab
TU Darmstadt
Darmstadt, Germany
mhollick@seemoo.tu-darmstadt.de

## ABSTRACT

Even avid users of mobile applications turn a blind eye to privacy settings. Still mobile applications remain the key means by which users share sensitive personal information. It is unclear if users just do not care, if they are missing the appropriate tools or user interfaces, or if they live in the delusion of being in control of their data. We argue that non-user-friendly design presents a key obstacle in making privacy controls work: it hinders users to effectively set up and maintain privacy settings. Our ultimate goal is to support the user by automatically suggesting access control lists based on an analysis of her communication metadata.To guide us in the design of such privacy suggestions, we perform an explorative questionnaire-based study with 42 participants. Our results confirm that users are overtaxed with existing schemes. We identify the expectations and preferences of users, thus facilitating the design of improved solutions.

## Categories and Subject Descriptors

K.6.5 [**Management of Computing and Information Systems**]: Security and Protection; H.1.2 [**Information Systems**]: User/Machine Systems—*Human factors*

## General Terms

Human Factors, Security, Experimentation

## 1. INTRODUCTION

The total number of mobile active Facebook users increased from 425 to 798 million between 2012 and 2015 [1, 19]. Already in 2011, 510 comments were estimated to be posted, 293,000 statuses to be updated, and 136,000 photos to be uploaded on average every minute by the Facebook community [18]. To control the access to these posted contents, social media users can often choose between predefined groups (e.g., everyone, friends) or define specific groups by manually selecting individuals from their contacts [10]. Several user studies, however, show that the existing mechanisms are not appropriate. For example, these mechanisms have been found to be time-consuming and complex for users [9], leading to sharing with unintended audience [7, 27]. Moreover, the proposed methods are static and thus do not consider the post context as well as the relationship dynamics between users [10]. As a result, users need to constantly maintain their access control lists in addition to the creation overhead. In practice, users however rarely reuse and update existing lists [29].

To improve the current state-of-the-art, the efforts required by users to configure their sharing settings should be reduced. To this end, we propose to support users by suggesting privacy settings based on the sensitivity of the content to be posted and the current strength of their social relationships to other users. For example, only socially close users would be suggested for a post rated as sensitive. By doing so, lists of contacts would be dynamically created and tailored to the post to be published. To cater for both awareness and control, users would be able to validate or disprove the suggestions before posting content. In order to create these suggestions, we first aim at analyzing communication data already available on the users' mobile phones (i.e., calls, SMS, MMS, and e-mails in our case) to identify and classify existing social relationships. We have conducted an explorative questionnaire-based study involving 42 participants in total. Among them, 19 participants tested our mobile app logging and processing both incoming and outgoing communication to extract statistics about their, e.g., duration, length, or the time of the day and contributed data during approximately one month. By conducting this deployment, we first have confirmed the feasibility of our approach by collecting real-world data for classifying social relationships as detailed in [22]. Secondly, we have examined the acceptance of potential users having run the app on their phone. Indeed, our approach relies on analyzing sensitive information about users to be able to provide suggestions, thus helping them to protect their privacy. Even if both collection and processing are conducted on the phones only, users might perceive them as an invasion to their pri-

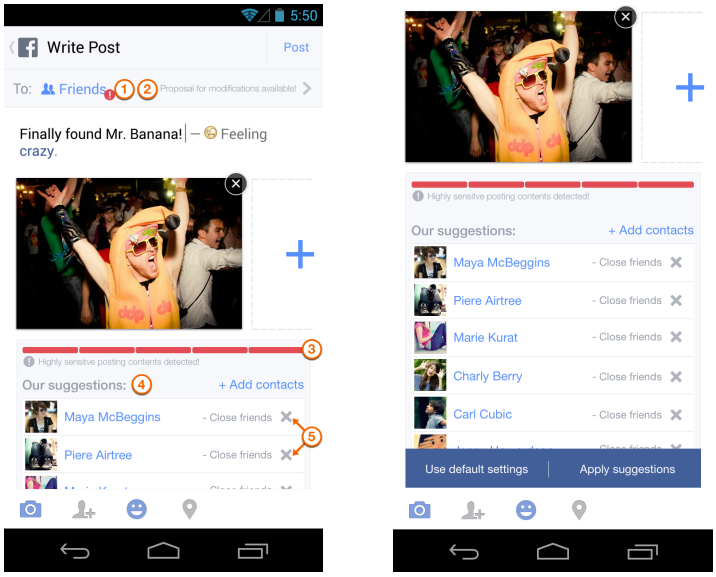


(a) Posting interface (scrolled up)

(b) Posting interface (scrolled down)

**Figure 1: Audience suggestions displayed during the creation of a new post**

vacy and be reluctant to use our proposed approach. Within the scope of this paper, we hence report the study results. Among others, our results confirm the difficulties users encounter in using existing mechanisms based on static lists, which is consistent with existing surveys such as [7, 9, 27]. In addition to provide us insights about the final design of our concept, the results are encouraging and show that a majority of the participants would be ready to use it. Even 7% would trust the system and let it automatically apply the generated suggestions.

The remaining of this paper is structured as follows. We first detail the current state-of-the-art in Sec. 2, before introducing our concept in Sec. 3. We present our questionnaire-based study and discuss the results in Sec. 4 and Sec. 5, respectively. We finally discuss related work in Sec. 6, before making concluding remarks in Sec. 7.

## 2. CURRENT FACEBOOK APPROACH

In what follows, we provide a brief overview of the current mechanisms available to users to control and review the audience of their posts. We employed these mechanisms in our questionnaire-based survey introduced in Sec. 4. We have selected Facebook as an example due to its large user base.

As of mid 2015, protecting the privacy of content in Facebook is cumbersome. When creating a post or uploading a picture, users can define its audience using the provided *audience selector tool*. They can decide to make it public, share it with all their contacts, or restrict the access to themselves only. They can also choose the *custom* option. In this case, they can manually select individual contacts or groups of contacts as well as explicitly exclude them. To create such a group, users can either use those predefined by Facebook (e.g., close friends, family, acquaintances, or colleagues) or specify a new one. Note that Facebook makes suggestions about potential members of the predefined groups but only based on the information manually provided by the users in their profiles (such as *work and education* or *relationships and family members*). To populate the different groups, users still need to individually select each member. Once a group is created, users can further define privacy restrictions to be applied to this specific group. Facebook finally remembers the user-selected sharing settings, which then become the default settings proposed for the next post. Note that the Facebook mobile app however does not provide all aforementioned options. By using it, users can neither edit and create new contact groups nor select the *custom* option.

After having selected the audience for their posts, users can review their settings using the proposed *View As* tool. It allows them to verify whether the selected settings correspond to their initial intention by adopting the perspective of the public or a particular contact. By using it, the users see their own timeline with the eyes of their selected audience.

## 3. PROPOSED APPROACH

Instead of manually populating the groups to share content with, we propose to support the users in this step by estimating the strength of their social relationships based on their interactions with other users. In addition, we aim at inferring the sensitivity of the post to be published to be able to suggest an appropriate audience. The users can either (partially or fully) adopt or ignore the displayed suggestions. In the latter case, they can use the existing mechanisms to select the post's audience, such as making it public or available to all their friends. In the following, we detail the different steps of our approach.

### 3.1 Mobile Data Collection and Processing

Our approach relies on an analysis of interactions between users to determine their degree of closeness. To this end, we consider communication data already available on the user phone. In our prototype, we consider contact names,

phone calls, SMS, MMS, and e-mails. Similarly to [14], we collect different features, such as call frequency and duration or e-mail length to classify the corresponding contacts into different categories. Indeed, a preliminary study has highlighted that users interact differently with their family, friends, or colleagues [3]. The data collection automatically runs in the background depending on the user's preferences, such as the time of the day or the frequency. Users can deactivate the collection function when they wish. Once the data are logged, they are leveraged for two purposes, namely transparency and classification. Since we are processing personal data, we believe that transparency is a key factor in the acceptance of our approach as already shown in orthogonal domains [8]. To this end, users can access and display different statistics about the collected data in dedicated interfaces. By doing so, users can visualize them and hence make an informed decision to continue using the app or not. Moreover, the collected data are primarily meant to serve in the classification of interacting users into several groups, such as friends, family, acquaintances, schoolmates, and colleagues. The classification results are however considered as out of scope of this paper.

## 3.2 Sensitivity Assessment

To suggest an appropriate audience, our approach further aims at estimating the sensitivity of posts by considering different factors, such as current time and location, tagged persons, and content type. For example, a picture taken at 2 am close to a nightclub might be more sensitive than one taken on a Sunday afternoon in a park. In our mock-up illustrated in Fig. 1, the estimated sensitivity is displayed using a colored bar (see Marker 3 in Fig. 1(a)). An non-critical post is coded using one green unit, while a very sensitive post is coded using five red units. If necessary, users can modify the estimated sensitivity degree.

## 3.3 Suggestion Display

When users create new posts, they are notified by an icon and a text that privacy suggestions are available and can be accessed (see Markers 1 and 2 in Fig. 1(a)). A list of suggested contacts is computed and displayed below the post to be published based on both the estimated sensitivity and inferred relationships strength (see Marker 4). For each contact, the profile picture, name, and the inferred relationship category are displayed. Users can remove individual contacts by deselecting them (see Marker 5) or add new ones using the *+ add contacts* option. Changes are taken into consideration to improve future suggestions. If the users are satisfied with the audience, they can apply these settings (see Fig. 1(b)). Otherwise, they can use the existing mechanisms using the *use default settings* option. The publication occurs when the users select the *post* option in Fig. 1(a).

# 4. QUESTIONNAIRE-BASED STUDY

We have distributed our questionnaire to 42 volunteers. 19 of them had experienced our app monitoring their communication patterns during one month. The questionnaire was in German and participants needed between 15 and 30 minutes in average to answer it. No incentives were provided. Note that the ethics committee of our university reviewed and approved our app deployment and the corresponding data processing. In the following, we first present demographic information about the participants, before surveying their use of current tools provided in Facebook. We finally study their attitudes towards privacy suggestions.

## 4.1 Demographics

A majority of our participants are male (69%). Their age ranges between 22 and 58 ($m$=30, $SD$=10). Most of them are students (64%) followed by employees (24%) and self-employed (10%). 2% are unemployed. 57% are studying or working in the fields of natural sciences, computer science, or engineering. Other working areas are accounting and management (7%), humanities (7%), social and healthcare (5%), architecture and civil engineering (2%), marketing and sales (2%), and production (2%). All participants use their phone multiple times a day and indicated to spend around 2.79 hours using it daily. 73% of our participants consult their Facebook profile at least once a day. Overall, the participants indicated to have an experience level of 4 on a scale from 1 (beginner) to 5 (expert) ($SD$=0.9). A Mann-Whitney U test shows that the participants' experience significantly differs based on their gender ($U$=262, $n$=42, $p$=0.046, $r$=40.4). In particular, male participants indicated to be globally more experienced than female participants. Based on a 5-point Likert scale, the participants rated the importance of the protection of their privacy on Facebook. The result shows that 79% consider the protection of their privacy as important to very important. Note that we only report statistically significant results within the scope of this paper. For example, we observe based on Kruskal-Wallis tests that the participants' age has no significant effect on their answers and therefore not comment further on it.

## 4.2 Experience with the Current Facebook Approach

We asked the participants about their experience with solutions currently available in Facebook (see Sec. 2 for a detailed description).

### *4.2.1 Visibility and Review*

Almost all participants (95%) use the existing configuration tools to restrict the visibility of their posts. In more detail, 45% always restrict the visibility to their friends, 7% to their friends plus friends-of-friends, and 5% do not share their posts with others. In comparison, 38% adapt the post visibility to the content to be published. Once the visibility of their posts is configured, 58% of our participants use the *View As* tool to review their configuration. In contrast, 12% rely on their friends to test their settings, while 27% do not review them. The remaining does not know. Similarly, 63% have already activated the timeline review function, while 23% have not used it yet and the remaining do not know this function.

While a majority of our participants claim to apply the provided tools to configure and review their settings, not all participants know them or use them.

### *4.2.2 List Creation and Management*

As detailed in Sec. 2, Facebook users also have the possibility to manually select individuals to create lists to share content with. In our sample, 85% know this option, but only 28% have already created at least one list. Among list creators, 18% indicated that they have never used their created list(s). Further 38% found that creating and managing lists

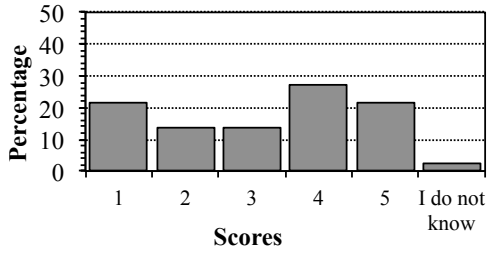


**Figure 2: Distribution of the scores attributed to the question "How often do you configure the visibility of a new post before publishing it". A score of 1 corresponds to never and 5 to always.**

is time-consuming. Note that the number of created lists ranges between two and ten for these participants.

Two Kruskal-Wallis H tests indicate a significant difference in the participants' perception of the effort demanded by the list creation and management depending on (1) their own experience ($\chi^2(2)$=4.91, $n$=34, $p$=0.027) and (2) how they review their privacy settings (see Sec. 4.2.1) ($\chi^2(2)$=9.57, $n$=34, $p$=0.023). In the former case, participants having already created and maintained lists find the process more cumbersome than those who have not. In the latter case, a pair-wise application of the Wilcoxon test with Bonferroni correction shows that participants relying on their friends to test their settings find this process significantly more cumbersome than participants who do not review them at all ($Z$=15.2, $n$=34, $p$=0.047, $r$=2.60).

Participants having not created lists yet, gave the following main reasons: (1) they do not see their utility (61%), (2) lists are confusing and require to remember who is in which list and which content is shared (22%), and (3) lists are time-consuming (13%) (multiple choices possible).

We hence observe that creating and managing lists is cumbersome for our participants. As a result, only a minority is actually using them on a regular basis.

### 4.2.3 Publishing Experience and Behaviors

93% of our participants have already published content on Facebook. As shown in Fig. 2, 22% never configure the visibility of their posts before publication. In this case, it is however difficult to interpret if these participants have, for instance, very restrictive privacy settings and therefore do not need to modify them, or if they are eager to share new content and thus do not take the time to change their settings. A Kruskal-Wallis H test reveals a relationship between the participants' reviewing behavior and the frequency at which they use lists to control the access to their posts ($\chi^2(2)$=4.23, $n$=11, $p$=0.040). In our sample, we observe that those, who rely on their friends to check their settings, use lists more often than others. Nevertheless, 11% revealed that they have already shared contents with an unintended audience and 19% do not know. Such result highlight that existing tools do not sufficiently support the users in their decision to either share contents or protect their privacy. To better understand the reason(s) of such unintentional sharing, we asked the participants to describe the underlying scenario. One participant mentioned that she had published a picture on her timeline, instead of sending it in a private message. Another one published a picture during a party and regretted it afterwards, as it was "not supposed to be seen by the rest of the world". As a result, 57% indicated that "they have already refrained from posting contents intended to one particular group of contacts" because (1) the individual configuration was too time-consuming (31%), (2) they were afraid to disclose them to untended audience (17%), (3) the individual configuration was too complex (9%), and (4) they did not know how to restrict the access to individuals (2%) (multiple choices possible). In the free-text field, two participants mentioned the time effort required by the current sharing lists: "[...] creating a new list is too time-consuming" and "[...]it is faster for me to send it using messages to the concerned persons." Another participant "did not published her post because she had tried to configure the visibility, but it somehow turned out to be too complicated".

In summary, the participants' answers show that the existing tools may not be optimal as they remain complex and time-consuming for several participants. As a result, our participants do not always know them or use them irregularly. For some of them, this has already resulted in sharing contents with unintended audience or even refraining from publishing contents.

## 4.3 Suggestions of Privacy Settings

After having surveyed the experience of our participants with existing mechanisms in Facebook, we introduced our proposed approach by means of a textual description accompanied by mock-ups as those shown in Sec. 3. We then asked them to evaluate our approach based on different criteria. As shown in Fig. 3, a majority found our approach rather helpful (67%), time-saving (60%), protecting privacy (62%), and supportive (72%).

### 4.3.1 Data Collection

To infer the nature of social relationships and hence suggest privacy settings, our approach needs to monitor the users' communication patterns. While the processing is meant to be conducted on the phone, the data collection may be perceived by the users as an intrusion to their privacy. To gain insights about potential concerns, we asked our participants to rate the sensitivity of different data types based on a 5-point Likert scale. A score of 5 corresponds to highly sensitive. Fig. 4 summarizes the distribution of the participants' ratings. A Friedman test shows a significant difference in perceived sensitivity depending on the data type ($\chi^2(2)$=23.4, $n$=40, $p$=0.000). A pair-wise application of the Wilcoxon test with Bonferroni correction indicates that the participants rated the phone contacts as significantly more sensitive than e-mail ($Z$=1.30, $n$=40, $p$=0.028, $r$=0.20) and MMS metadata ($Z$=1.23, $n$=40, $p$=0.046, $r$=0.20). Differences between the others data types are statistically insignificant. Asked if the data collection would stop them to use our approach, the participants remain divided: 43% answered yes, 38% said no, and the remaining do not know. In particular, call metadata are the most selected data type followed by e-mail and SMS metadata that would stop participants from adopting our approach as illustrated in Fig. 5 (multiple choices possible).

### 4.3.2 Data Processing

After having investigated how participants perceive data collection, we next focus on data processing. To this end, we

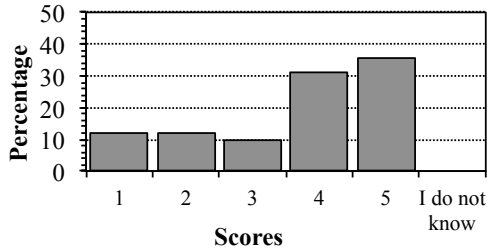


(a) Useless (score of 1) - helpful (5)

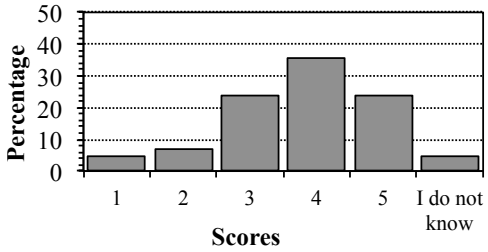


(b) Time-consuming (1) - time-saving (5)

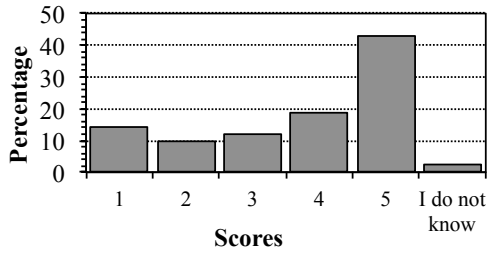


(c) Privacy threat (1) - protection (5)

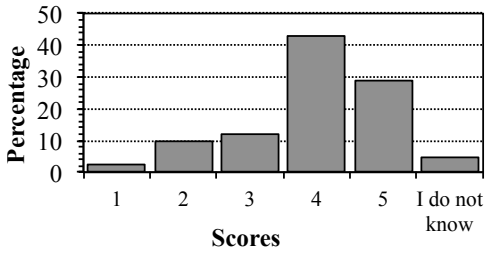


(d) Misleading (1) - supportive (5)

**Figure 3: Distribution of the participants' ratings for different evaluation criteria**

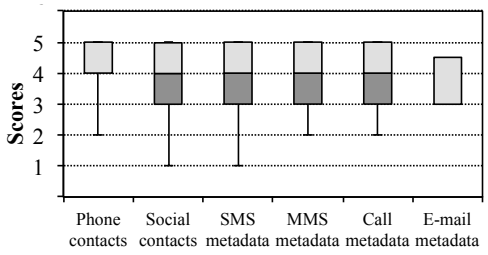


**Figure 4: Minimum, quartiles, and maximum score attributed to different data types collected by the app**

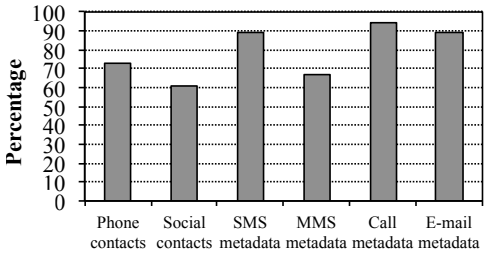


**Figure 5: Distribution of the participants' answer to the question "Indicate the data type(s) whose collection would stop you from using our approach"**

asked the participants to indicate which condition(s) need(s) to be fulfilled so that they would accept both the analysis of their data as well as the inference of their social relationships by our system. In both cases, more than 80% of the participants indicated that the collected data as well as inferred relationships should not leave their phone and be accessed by third parties. One participant especially commented that

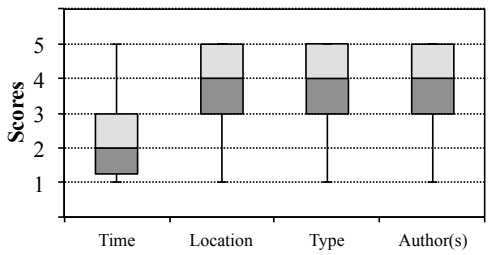


**Figure 6: Minimum, quartiles, and maximum score attributed to different factors potentially affecting the post's sensitivity**

"[f]or [him], it is important that the system does not spy on [him]. If using it, [he] want[s] to know exactly which data it obtains and how data is processed. Its actions, the data it processes, and the permissions it uses to do so must be clearly evident to [him]". As expected, confidentiality and transparency are confirmed to be key aspects in the participants' acceptance.

### *4.3.3 Content Analysis and Sensitivity*

As detailed in Sec. 3, our approach aims at proposing an appropriate audience based on the sensitivity of the content to be published. To be able to later match the content with different contact groups, we asked our participants to rate on a 5-point Likert scale how different factors influence a post's sensitivity. A score of 1 means no influence. Fig. 6 summarizes the results. A Friedman test shows a significant difference between the different factors ($\chi^2(2)$=32.656, $n$=37, $p$=0.000). A pair-wise application of the Wilcoxon test with Bonferroni correction shows that the participants perceived the time at which a content is created as less sensitive than the users' location ($Z$=-1.34, $n$=37, $p$=0.000, $r$=0.22), the post type ($Z$=-1.16, $n$=37, $p$=0.001, $r$=0.19),

**Table 1: Proposed content types categorized by sensitivity**

| # | Content | Sensitivity |
|---|---|---|
| a | Criticism or anger about job | Extreme |
| b | Personal issues | |
| c | Own relationships (e.g., status, statements to own relationship) | High |
| d | Family | |
| e | Health issues (e.g., disease, diet, etc.) | |
| f | Own contents that involve other contacts (e.g., through tagging) | |
| g | External contents involving the participant (e.g., through tagging) | Medium |
| h | Contents with negative slang (e.g., usage of swearwords) | |
| i | Job in general | |
| j | Party and night life | |
| k | Social or ethical statement (e.g., about animal testing) | Low |
| l | Religion | |
| m | Politics | |
| n | Hobby | |
| o | Entertainment (e.g., contents shared via 9GAG, Pinterest, or YouTube) | Very low |
| p | News | |
| q | Economics | |
| r | Culture | |
| s | Music | |
| t | Knowledge and science | |

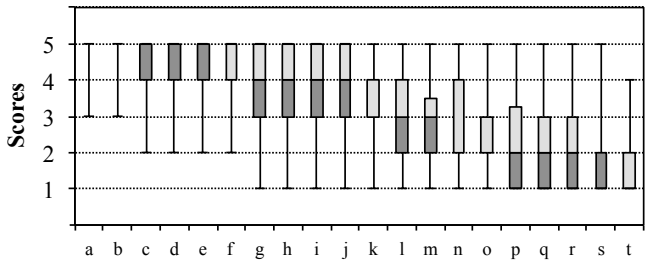


**Figure 7: Minimum, quartiles, and maximum score attributed to different content types detailed in Tab. 1**

or its author(s) ($Z$=-1.18, $n$=37, $p$=0.001, $r$=0.19). No statistically significant differences are observed between the other factors, but 75% of the participants indicated that these three factors have at least a moderate influence on the post's sensitivity. Regarding the participants' gender, we only observe a significant difference in the case of the post's author(s) ($U$=-79.0, $n$=39, $p$=0.018, $r$=12.6). Female participants especially rated this factor as more sensitive than male participants. Additional factors were also cited by the participants, such as potential persons tagged in the post or the type of devices from which the post is published.

We then further investigated how participants quantify the sensitivity of a post depending on its content. To this end, we submitted different content types partially derived from [27] to the participants. Again, a score of 1 means that the considered content type is rated as not sensitive at all, while a score of 5 means that it is highly sensitive. Based on the results displayed in Fig. 7, we categorize the proposed content types into different sensitivity categories presented in Tab. 1. Against our expectations, contents related to potential health issues are globally rated by the participants as less sensitive than other personal issues. Similarly, religion and politics are rated relatively low. While this classification is by no means exhaustive, it however allows us to determine trends in terms of content sensitivity and later map them to different generated contact groups.

In addition, different factors have a significant effect on the sensitivity scores attributed to certain content types. For example, iOS users find that sharing content involving other contacts is significantly more sensitive than Android users ($Z$=11.7, $n$=42, $p$=0.021, $r$=1.80). Participants sharing posts using the option "friends and friends of friends" rated the sensitivity of contents including negative slang higher than those sharing no posts ($Z$=-30.5, $n$=39, $p$=0.023, $r$=4.88) and those dynamically adapting the public to the post's content ($Z$=-20.6, $n$=39, $p$=0.031, $r$=3.31). This difference may be due to the fact that the latter subgroups of participants are more in control of the audience of the post than the former. Finally, the ratings of participants having already activated the *view as* tool (i.e., being potentially more aware of privacy issues than the others) are significantly higher than others regarding contents related to (1) family ($Z$=11.1, $n$=40, $p$=0.020, $r$=1.76), (2) criticism or anger about job ($Z$=7.30, $n$=39, $p$=0.030, $r$=1.17), and (3) health issues ($Z$=9.60, $n$=40, $p$=0.049, $r$=1.52).

### 4.3.4 Suggested Users

Instead of requesting users to manually populate the group(s) to share content with, our approach would suggest them potential users already classified into different groups. We therefore surveyed which groups would be appropriate. Except a 24% of the participants who would not make a distinction between different audience groups, the remaining named in average 4 groups mainly including friends, family, acquaintances, schoolmates, and colleagues. As a result, their answers converge with the groups proposed in our approach as defined in Sec. 3.1.

With our approach, we aim at providing suggestions to privacy settings and therefore support the users' in their configuration. Nevertheless, our objective is not to develop a system that would automatically decide for the users. On the contrary, we want that users maintain control over their privacy. This control however requires user interactions to, e.g., confirm that they agree with the proposed settings or

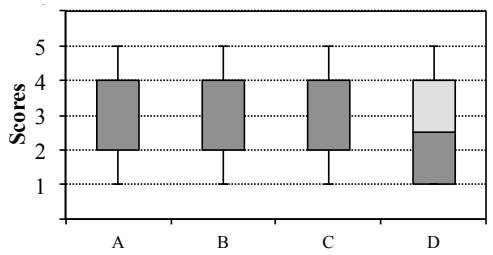


**Figure 8: Minimum, quartiles, and maximum score attributed to the following statements: (A) I would use this system, (B) This system would enhance my confidence in protecting my privacy when sharing content in online social networks, (C) This system would enhance my confidence in choosing the appropriate audience for my posts, and (D) I would share content more often using this system**

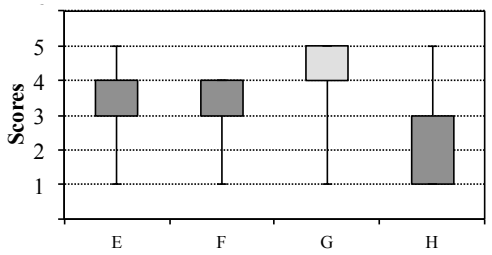


**Figure 9: Minimum, quartiles, and maximum score attributed to the following statements: (E) This system would help me to save time and effort when configuring my privacy preferences, (F) Applying and adapting the suggestions would help me to reconsider the potential risks and consequences of sharing content online, (G) I would trust the proposed system more if it takes into consideration my feedback, and (H) I would trust the proposed suggestions**

improve the suggestions. We therefore examined how often our participants would be ready to interact with our approach to improve the proposed suggestions. 17% claim to be ready to always provide feedback, while others are willing to do it either occasionally (43%) or only at the beginning (26%). The remaining do not want to provide feedback because they either would not want to use our approach or would not accept suggestions. Note that the answers only provide insights about the participants' claimed intention.

#### 4.3.5 Suggestion Presentation

We next surveyed the participants' preferences about the presentation of the suggestions. According to our expectations, most participants would prefer seeing them after having prepared their post and before publishing it (47%). In contrast, 24% would favor a presentation while they are composing their posts. Some 7% would trust the system and let it automatically apply the generated suggestions. For the latter users, the suggested audience would only be displayed in case of active user request. We find this result quite surprising, as we expected that users would like to keep the control on the audience and may not trust our system to such extend. By doing so, they would however reduce the number of interactions required, and hence the corresponding overhead. The remaining participants do not know which options they would prefer. One participant especially commented that she appreciated to be actively "reminded of the audience configuration" as in the first option. However, the user interactions required in this step should be kept to the minimum. Two participants mentioned that only suggestions that have not be configured, confirmed, or declined yet should be displayed. Another suggestion was to allow users to choose and switch between different modes, such as active display and background processing.

#### 4.3.6 Acceptance

We finally submitted to our participants multiple statements to be rated using a 5-point Likert scale. A score of 1 indicates a strong disagreement, while a score of 5 corresponds to a strong agreement. Figs. 8 and 9 compile the results. In particular, 61% of our participants indicated that they would use our approach. A majority thinks that our approach would enhance their confidence in both protecting their privacy and choosing the appropriate audience when sharing contents online (55% and 62%, respectively). The remaining users are either undecided or did not agree with the proposed statements. As a result of both latter results, 33% indicated that they would post more content online. Similarly, many participants believe that our approach would reduce the configuration overhead (64%) and increase their awareness about potential risks for their privacy (67%). However, they are overall doubtful about the accuracy of the suggestions and would trust the system more if they are able to validate or invalidate suggestions (81%).

In summary, our participants are globally positive about the proposed approach. Collecting and processing personal data would not stop them from testing our system as soon as all information remains on their devices. Participants are ready to interact with the system and consider it as an important factor to be able to trust the proposed suggestions.

## 5. DISCUSSION

The demographics of our participants may not be representative for the whole population. Indeed, a majority are students and the most represented field is "natural science, computer science, and engineering". As a result, our participants may have more expertise and understanding when interacting in online social networks than other population groups. However, their answers do not significantly differ from the other participants and the results show that most of them are also struggling with existing solutions. Moreover, no significant impact has been observed between the participants having tested our application and the others.

As all questionnaire-based studies, the answers provided by our participants reflect their claimed opinions and not necessarily their actual behavior. To investigate it under realistic conditions, a real-world deployment should be conducted and is considered as future work. At the exception of one question, the answers of the participants having installed our app however do not significantly differ from the ones having only answered our questionnaire as shown by Mann-Whitney U tests. In particular, the participants having installed our app rated the time overhead related to the creation and management of contact lists significantly higher than the others ($U$=74.5, $n$=34, $p$=0.015, $r$=12.8).

About the practical significance of the obtained results,

we observe two trends. Firstly, the differences observed between the perceived sensitivity of the different metadata collected as well as factors influencing the sensitivity of posted contents are statistically significant, but show only small to medium effect size. As a result, their practical significance is limited. These results however primarily aim at refining the design of our concept. In contrast, very large effect sizes are obtained for the remaining results, which are hence both statistically and practically significant.

In their free-text comments, several participants indicated that they encountered difficulties to understand the configuration of the current Facebook privacy controls. Moreover, the configuration overhead has already stopped several of them from publishing content online. In these cases, our approach may be appropriate and support them. Unlike [26, 27], the most cited reason why participants are not sharing content online is not the fear of sharing it with unintended audience and the potential consequences, but their lack of trust in the platform providers. Our approach provides no solutions in this case.

# 6. RELATED WORK

The results of our questionnaire-based study detailed in Sec. 4 highlight that existing solutions are not optimal for users. This observation is in line with [9, 10, 13, 28]. Indeed, [13] shows that users encounter difficulties to translate their privacy conception into settings in online social networks. Predefined lists of contacts are especially shown to be inappropriate, as they do not consider the context [10] and the content of the post [9]. Additionally, [28] demonstrates that users have issues controlling what they share with overlapping contact lists. The aforementioned works however focus on establishing the limitations of existing solutions, but do not propose alternatives such as our approach.

One key component of our approach is the inference of the nature of social relationships between users. To reach this goal, different methods and data types have already been exploited. At a network level, e-mail traffic is analyzed in [6], while interactions in an online social network are considered in [12, 23, 24, 30]. In particular, [23, 24, 30] focus on analyzing the closeness between users to identify common interests and hence make better content recommendations. However, these works do not consider privacy issues. Moreover, calls between mobile phone users are leveraged in [15, 17]. In contrast, our approach relies on the analysis of local data collected on each mobile phone. For example, [4, 20] base their analysis on proximity data collected using Bluetooth-enabled phones. These information are further completed by locations, calls, and SMS in [5] to identify friends and determine the nature of their relationships between symmetric and asymmetric friendships. Similar data are considered in [14] to distinguish relationships between family, colleagues, and social contacts. These works however do not consider using the inferred relationships to suggest privacy settings to the users.

Based on the inferred relationships, our approach aims at suggesting contacts with which users would like to share their post. A similar idea is proposed in [2], in which only features extracted from Facebook are considered to suggest new members in order to complete the group currently created by the users. However, the proposed solution does not take into account the posts' sensitivity, and hence the users' privacy protection. While existing solutions also work towards assisting users in protecting their privacy, they adopt different approaches. For example, several solutions focus on defining general privacy policies and preferences, such as [7, 11, 16, 21]. Our proposal differs from these solutions, as we aim at providing suggestions tailored to the content to be published instead of general policies. [25] also adopts this idea by considering the sensitivity of the contents to be posted, but additionally leverages the manually defined users' privacy preferences as well as the risk of exposure and disclosure for each profile data. In contrast, we aim at relying on communication data available on the users' mobile phone to infer the nature of their social relationships rather than only considering users' inputs. Our work finally shares similarities with [26], in which privacy nudges are introduced to invite users to consider the content and audience of their posts more carefully before publishing them. The nudges however only aim at increasing the users' awareness about potential privacy issues, but do not support the users in the selection of their sharing settings.

# 7. CONCLUSIONS

Online social networks and communities such as Facebook, Google+, Twitter, etc. are publishing user-generated content with an steadily increasing pace. The use of mobile applications to share content has evolved from a niche use case to the dominant form of content sharing for billions of users. Controlling privacy settings is already challenging on a desktop computer, hence, even avid users turn a blind eye to privacy settings in mobile applications. Within our work, we propose to support the user with suggestions for appropriate access control list settings, thus facilitating to remain in control of her privacy even on mobile platforms with limited controls. These suggestions are derived from the communication metadata of the user, which we propose to analyze locally on her device. We have demonstrated the feasibility of the extraction of this metadata by developing an app that collects the data in a privacy-conscious manner and deployed it in a real-world experiment involving 19 users during one month. Our main contribution is an explorative questionnaire-based study with 42 participants to guide us in the design of the aforementioned privacy setting suggestion scheme that utilizes the gathered metadata. Our results confirm related work in that existing privacy controls overtax the user. We additionally find that a substantial set of users mistrust the platform provider to respect their privacy. While the majority of users embraces the idea of privacy suggestions (and perceive it as helpful, time-saving, supportive, and privacy-enhancing), there is also scepticism to have a system scrutinizing communication metadata to mine the necessary social relationships—even for a fully local implementation residing exclusively on the mobile system. At the same time, keeping the data local was one of the most important reasons to accept the analysis of communication metadata at all. We further are able to identify some of the most privacy-sensitive content categories for user generated content. Altogether, we can show that user-friendly suggestions would be highly sought after by end users and have the potential to tip the privacy scale in favor of the users.

# 8. ACKNOWLEDGMENTS

Our thanks go to the participants of the user study and A. García Bouso for her feedback.

# 9. REFERENCES


[1] Facebook Statistics. Online: http://newsroom.fb.com/company-info/ (accessed 04.15), 2015.

[2] S. Amershi, J. Fogarty, and D. Weld. Regroup: Interactive Machine Learning for On-demand Group Creation in Social Networks. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI)*, 2012.

[3] D. Christin, A. Bentolila, and M. Hollick. Friend is Calling: Exploiting Mobile Phone Data to Help Users in Setting their Privacy Preferences. In *Proceedings of the 4th International Workshop on Security and Privacy in Spontaneous Interaction and Mobile Phone Use (IWSSI/SPMU)*, 2012.

[4] N. Eagle and A. S. Pentland. Reality Mining: Sensing Complex Social Systems. *Personal Ubiquitous Computing*, 2006.

[5] N. Eagle, A. S. Pentland, and D. Lazer. Inferring Friendship Network Structure by Using Mobile Phone Data. *Proceedings of the National Academy of Sciences*, 2009.

[6] H. Ebel, L.-I. Mielsch, and S. Bornholdt. Scale-free Topology of E-mail Networks. *Physical review E*, 66(3), 2002.

[7] L. Fang and K. LeFevre. Privacy Wizards for Social Networking Sites. In *Proceedings of the 19th ACM International Conference on World Wide Web (WWW)*, 2010.

[8] J. L. Herlocker, J. A. Konstan, and J. Riedl. Explaining Collaborative Filtering Recommendations. In *Proceedings of the 3rd ACM Conference on Computer Supported Cooperative Work (CSCW)*, 2000.

[9] M. Johnson, S. Egelman, and S. M. Bellovin. Facebook and Privacy: It's Complicated. In *Proceedings of the 8th ACM Symposium on Usable Privacy and Security (SOUPS)*, 2012.

[10] P. G. Kelley, R. Brewer, Y. Mayer, L. F. Cranor, and N. Sadeh. An Investigation into Facebook Friend Grouping. In *Proceedings of the 13th IFIP TC 13 International Conference on Human-computer Interaction (INTERACT)*, 2011.

[11] P. G. Kelley, P. Hankes Drielsma, N. Sadeh, and L. F. Cranor. User-controllable Learning of Security and Privacy Policies. In *Proceedings of the 1st ACM Workshop on AISec*, 2008.

[12] G. Kossinets and D. J. Watts. Empirical Analysis of an Evolving Social Network. *Science*, 2006.

[13] M. Madejski, M. Johnson, and S. M. Bellovin. A Study of Privacy Settings Errors in an Online Social Network. In *Proceedings of the 10th IEEE International Conference on Pervasive Computing and Communications (PERCOM Workshops)*, 2012.

[14] J.-K. Min, J. Wiese, J. I. Hong, and J. Zimmerman. Mining Smartphone Data to Classify Life-facets of Social Relationships. In *Proceedings of the 16th ACM Conference on Computer Supported Cooperative Work (CSCW)*, 2013.

[15] S. H. Mirisaee, S. Noorzadeh, A. Sami, and R. Sameni. Mining Friendship from Cell-Phone Switch Data. In *Proceedings of the IEEE 3rd International Conference on Human-Centric Computing (HumanCom)*, 2010.

[16] J. Mugan, T. Sharma, and N. Sadeh. Understandable Learning of Privacy Preferences Through Default Personas and Suggestions. Technical Report CMU-ISR-11-112, Carnegie Mellon University, http://ra.adm.cs.cmu.edu/anon/anon/isr2011/CMU-ISR-11-112.pdf (accessed 04.15), 2011.

[17] J.-P. Onnela, J. Saramäki, J. Hyvönen, G. Szabó, D. Lazer, K. Kaski, J. Kertész, and A.-L. Barabási. Structure and Tie Strengths in Mobile Communication Networks. *Proceedings of the National Academy of Sciences*, 2007.

[18] C. Pring. 100 Social Media Statistics for 2012. Online: http://thesocialskinny.com/100-social-media-statistics-for-2012/ (accessed 04.15), 2012.

[19] E. Protalinski. Facebook has over 425 Million Mobile Users. Online: http://www.zdnet.com/article/facebook-has-over-425-million-mobile-users/ (accessed 04.15), 2012.

[20] D. Quercia and L. Capra. FriendSensing: Recommending Friends Using Mobile Phones. In *Proceedings of the 3rd ACM Conference on Recommender Systems (RecSys)*, 2009.

[21] R. Ravichandran, M. Benisch, P. G. Kelley, and N. Sadeh. Capturing Social Networking Privacy Preferences: Can Default Policies Help Alleviate Tradeoffs Between Expressiveness and User Burden? In *Proceedings of the 5th ACM Symposium on Usable Privacy and Security (SOUPS)*, 2009.

[22] D. Reinhardt, F. Engelmann, A. Moerov, and M. Hollick. Show Me Your Phone, I Will Tell You Who Your Friends Are: Analyzing Smartphone Data To Identify Social Relationships. In *Proceedings of the 14th ACM International Conference on Mobile and Ubiquitous Multimedia (MUM)*, 2015.

[23] Y. Shen and R. Jin. Learning Personal + Social Latent Factor Model for Social Recommendation. In *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, 2012.

[24] Y. Song, L. Zhang, and C. L. Giles. Automatic Tag Recommendation Algorithms for Social Recommender Systems. *ACM Transactions on the Web (TWEB)*, 5(1), 2011.

[25] A. Squicciarini, F. Paci, and S. Sundareswaran. PriMa: An Effective Privacy Protection Mechanism for Social Networks. In *Proceedings of the 5th ACM Symposium on Information, Computer and Communications Security (ASIACCS)*, 2010.

[26] Y. Wang, P. G. Leon, K. Scott, X. Chen, A. Acquisti, and L. F. Cranor. Privacy Nudges for Social Media: An Exploratory Facebook Study. In *Proceedings of the 22nd International Conference on World Wide Web (WWW)*, 2013.

[27] Y. Wang, G. Norcie, S. Komanduri, A. Acquisti, P. G. Leon, and L. F. Cranor. "I Regretted the Minute I Pressed Share": A Qualitative Study of Regrets on Facebook. In *Proceedings of the Seventh Symposium on Usable Privacy and Security (SOUPS)*, 2011.

[28] J. Watson, A. Besmer, and H. R. Lipford. +Your Circles: Sharing Behavior on Google+. In *Proceedings of the 8th ACM Symposium on Usable Privacy and*

*Security (SOUPS)*, 2012.

[29] J. Wiese, P. G. Kelley, L. F. Cranor, L. Dabbish, J. I. Hong, and J. Zimmerman. Are You Close with Me? Are You Nearby?: Investigating Social Groups, Closeness, and Willingness to Share. In *Proceedings of the 13th ACM International Conference on Ubiquitous Computing (UbiComp)*, 2011.

[30] X. Yang, Y. Guo, and Y. Liu. Bayesian-Inference based Recommendation in Online Social Networks. In *Proceedings of the 30th IEEE Conference on Computer Communications (INFOCOM)*, 2011.