

Sharing Policies in Multiuser Privacy Scenarios: Incorporating Context, Preferences, and Arguments in Decision Making

Ricard L. Fogues, Universitat Politècnica de Valencia
 Pradeep K. Murukannaiah, Rochester Institute of Technology
 Jose M. Such, King's College London
 Munindar P. Singh, North Carolina State University

Social network services enable users to conveniently share personal information. Often, the information shared concerns other people, especially other members of the social network service. In such situations, two or more people can have conflicting privacy preferences; thus, an appropriate sharing policy may not be apparent. We identify such situations as *multiuser privacy scenarios*. Current approaches propose finding a sharing policy through preference aggregation. However, studies suggest that users feel more confident in their decisions regarding sharing when they know the reasons behind each other's preferences. The goals of this paper are (1) understanding how people decide the appropriate sharing policy in multiuser scenarios where arguments are employed, and (2) developing a computational model to predict an appropriate sharing policy for a given scenario. We report on a study that involved a survey of 988 Amazon MTurk users about a variety of multiuser scenarios and the optimal sharing policy for each scenario. Our evaluation of the participants' responses reveals that contextual factors, user preferences, and arguments influence the optimal sharing policy in a multiuser scenario. We develop and evaluate an inference model that predicts the optimal sharing policy given the three types of features. We analyze the predictions of our inference model to uncover potential scenario types that lead to incorrect predictions, and to enhance our understanding of when multiuser scenarios are more or less prone to dispute.

CCS Concepts: •**Security and privacy** → **Social aspects of security and privacy**; *Social network security and privacy*;

General Terms: Experimentation; Human Factors

Additional Key Words and Phrases: Privacy; social media; multiuser; argumentation; crowdsourcing

ACM Reference Format:

Ricard L. Fogues, Pradeep K. Murukannaiah, Jose M. Such, Munindar P. Singh, 2016. Understanding Sharing Policies in Multiuser Scenarios. *ACM Trans. Comput.-Hum. Interact.* X, X, Article X (January 2016), 26 pages.
 DOI: 0000001.0000001

1. INTRODUCTION

A social network service (SNS) enables users to maintain social relationships via online interactions. SNS users share information with each other as they interact. Often, the information shared on an SNS involves more than one user. A natural example is a picture or video showing a group of people. Many SNSs enable users to connect the information they upload to other users so that the connected users can be notified of the uploaded information. Since the information shared varies depending on the SNS, these connections can take different forms, e.g., tags on a picture uploaded

Ricard Fogues worked on this project as part of his mobility stay at North Carolina State University, funded by the Erasmus Mundus Programme of the European Commission under the Transatlantic Partnership for Excellence in Engineering project. Pradeep Murukannaiah worked on this project when he was at North Carolina State University. Jose M. Such worked on this project while he was at Lancaster University. We thank the National Science Foundation (grant 0910868), US Department of Defense (the Science of Security Lablet grant), EPSRC (grant EP/M027805/1), Ministerio de Economía y Competitividad (grant TIN2014-55206-R), and the NC State College of Engineering for partial support. We also thank the associate editor, anonymous reviewers, and the members of NCSU Science of Security Lablet for useful comments.

Authors' addresses: Ricard L. Fogues, Departamento de Sistemas Informáticos, Universitat Politècnica de Valencia;

Pradeep K. Murukannaiah, Department of Software Engineering, Rochester Institute of Technology;

Munindar P. Singh, Department of Computer Science, North Carolina State University;

Jose M. Such, Department of Informatics, King's College London.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2016 Copyright held by the owner/author(s). 1073-0516/2016/01-ARTX \$15.00

DOI: 0000001.0000001

to Instagram or mentions in a tweet. Suppose Alice uploads a picture from last weekend's party in which she appears together with her friend Bob, and tags Bob in the picture. When these connections are created, the other users are linked to the uploaded information. Usually, a connection implies that the profile of the user can be accessed from the information or some personal information is shown in conjunction with the uploaded data. Although connections between information and users are widely employed by SNSs, they can pose a privacy threat. For example, Bob may find that the picture Alice uploaded is sensitive. However, Bob has no control over uploading that picture, and Alice's action can threaten Bob's privacy by revealing information about him from one setting or context into another. We identify a situation such as this as a *multiuser privacy scenario* or, for brevity, *multiuser scenario*.

Currently, SNSs do not provide mechanisms to handle multiuser scenarios [Fogués et al. 2015]. Thus, a user who did not upload a piece of information concerning him must deal with the privacy settings chosen by the uploader; at best, the user can remove the connection that links him to the shared information, but the information itself remains nevertheless. An ideal solution in a multiuser scenario is to respect each user's privacy. However, often such a solution may not be viable since the preferences of the users involved may conflict. For example, suppose Alice would like to share a picture in which Bob and she appear along with her friend Charlie. However, Bob would like to share this picture only with his common friends and he does not know Charlie. Here, no solution completely respects both Alice's and Bob's preferences.

Several researchers, e.g., [Besmer and Lipford 2010; Lampinen et al. 2011; Wisniewski et al. 2012; Such et al. 2014], have identified the lack of decision-support systems that help users resolve multiuser privacy conflicts as one of the biggest gaps in privacy management in social media. The main challenge that such decision-support systems would address is proposing an *optimal* solution: a solution most likely to be accepted by all those involved in the multiuser scenario. Although an optimal solution may not exist for each multiuser scenario, identifying one, when it exists, can minimize the burden on the users of having to resolve the conflict manually.

Other researchers (see [Squicciarini et al. 2009; Carminati and Ferrari 2011; Hu et al. 2013; Such and Criado 2016]) have proposed methods to automatically determine solutions to conflicts based on users' privacy preferences. These methods suffer from two main limitations. One, either they aggregate preferences in the same way regardless of the context or consider only a few, predetermined situations among the large number of potential situations. Two, they do not consider the reasons behind users' preferences. However, evidence from self-reported data [Lampinen et al. 2011; Wisniewski et al. 2012] suggests that users entertain the explanations provided by others and that the optimal solution may depend upon the particular context and reasons behind users' preferences. Following this idea, we empirically study three types of factors that potentially influence a privacy decision: the scenario's *context*, users' *preferences*, and their *arguments* about those preferences. An argument is a justification a user employs to convince the others involved that the user's expectations are reasonable and should be taken into account for making a decision.

Our key objective in this paper is to build an argumentation-based model that accurately represents a multiuser scenario. To this end, we (1) identify important factors that potentially influence the inference of sharing decisions; (2) evaluate the relative importance of these factors in inferring an optimal solution; and (3) develop a computational model that predicts an optimal solution for a given multiuser scenario.

We design a study where human participants are asked to choose what they think is the most appropriate sharing policy in a multiuser scenario we specified. Combining different values of the three types of factors we identified (context, preferences, and arguments), we generate 2,160 scenarios. Considering the sheer number of participants required, we conducted our study on Amazon Mechanical Turk (MTurk), collecting responses from 988 unique MTurk participants.

We are interested in understanding the norms of information sharing similar to those described by Nissenbaum [2004], Criado and Such [2015], and Murukannaiah et al. [2016] for multiuser scenarios. Therefore, our methodology crowdsources participants' opinions about what information sharing policies are optimal for a number of hypothetical scenarios involving third parties.

Shvartzshnaider et al. [2016] employ an approach similar to ours, where they crowdsource information sharing norms for an educational context generating hypothetical scenarios involving third parties. We later on show, however, that some scenarios lead to discrepancies in participants opinions and personalisation is the key to accurately recommending sharing policies for multiuser scenarios.

Contributions

- (1) We propose a novel model for representing and reasoning about multiuser scenarios, employing three types of features: contextual factors, user preferences, and user arguments.
- (2) We build a classifier, based on machine learning, as an exploratory attempt to predict the optimal sharing policy and evaluate the underlying privacy complexity of multiuser scenarios.
- (3) Employing the results obtained by the classifier, we identify what combinations of scenario-defining elements introduce variability in judgments on optimal sharing policy.

Organization

Section 2 identifies factors potentially influencing a multiuser sharing decision. Section 3 describes an inference model for predicting the optimal sharing policy and the MTurk study we conducted to build a dataset for training and testing our inference model. Section 4 describes our hypotheses and evaluation strategy. Section 5 presents our results and Section 6 describes how some of those results can be employed in a practical tool. Section 7 discusses our results, threats to their validity, and pointers for future work. Section 8 describes related works and Section 9 provides conclusions.

2. FACTORS INFLUENCING A MULTIUSER SHARING DECISION

We identify three important classes of factors that potentially influence the privacy decision in a multiuser scenario: context, user preferences, and arguments by users.

2.1. Context

Our definition of context reflects intuitions compatible with those of Nissenbaum's [2004] study of social settings that dictate the flow of personal information. Her definition of contextual integrity captures the idea that people share information of varying type and sensitivity in society, not simply as individuals in an undifferentiated social world, but as individuals in certain capacities (roles), in distinctive social contexts, such as health care, education, and employment. Since Nissenbaum formulates contextual integrity in the field of law and public policy, legal elements, such as obligation or entitlement, are prominent in the existing account. However, we posit that the legal elements do not adequately cover several aspects of sharing decisions in social networks.

Based on the literature, we focus on three elements of context as factors influencing sharing decisions in multiuser scenarios.

Relationships among the individuals. The roles of the individuals involved in a transaction are a defining element of context. Also, the relationship type is crucial when making *individual* decisions about privacy in social media [Marwick and Boyd 2011; Kairam et al. 2012]. Specifically, people share information differently with friends, family, and colleagues. We hypothesize that relationship types influence how a person perceives a multiuser scenario, because attributes of a relationship such as intimacy may influence the extent to which the parties involved in the scenario respect each other's opinion. For example, following Wisniewski et al. [2012], we imagine that a user's friend would respect the user's preferences.

Sensitivity of the information. The sensitivity of the information to be shared is known to influence individual sharing decisions in social media [Wang et al. 2011; Sleeper et al. 2013]. Besides, the judgment of appropriateness of sharing any information on an SNS is user and culture specific. For example, in some cultures, drinking alcohol is taboo. Thus, a picture showing a person drinking can be inappropriate in one culture, but normal in another. Each individual involved in a scenario has a perception of the sensitivity of the information. This perception may affect how important that person thinks his or her view is on the appropriate sharing decision.

Sentiment associated with the information. A user may employ social media for self-presentation and sharing as a means of maintaining that presentation [Hogan 2010; Kairam et al. 2012]. The sentiment conveyed by the personal information a user discloses on social media influences his or her self-presentation. For example, a user may share a picture of his recently broken leg with friends to receive emotional support. Similarly, a user can publicly post a picture of his graduation to obtain public recognition from everyone and congratulations from friends and family. In these examples, both pictures convey extremely different sentiments, nonetheless, both can be worth sharing, although with different audiences, depending on the intentions of the users. Since one of our goals is to observe the influence of arguments (which may convey the intentions of users) on the final privacy decision, we include sentiment as a factor contributing to context. Further, recent research on social media [Siersdorfer et al. 2010; You et al. 2015] presents methods to classify pictures based on sentiment, and some studies [Stieglitz and Dang-Xuan 2013] tackle the influence of the sentiment of the information on the level of its disclosure.

2.2. Preferences

The goal of each individual in a multiuser scenario is to persuade the sharer to apply the individual's preferred sharing policy to the information in question. The individuals' (including sharer's) preferences may be either compatible or conflicting. For example, Bob's preference of "I do not want my parents to see this picture" is compatible with Alice's preference of "I want only my friends to see it," as long as Bob's parents are not among Alice's friends. Optimally, the final decision to resolve a conflict should respect the preferences of every individual involved to the best possible extent. If all preferences are compatible with each other, the solution is trivial. However, in case of conflict, an acceptable solution may not be evident.

A sharing policy can imply no sharing, sharing publicly, or anything in between. Further, depending on the number of contacts and their types, sharing policies change from one SNS user to another. Given the large space of possibilities, for the sake of simplicity, we limit the sharing policies in our setting to three levels of disclosure.

- (1) *Share with all*: Anyone on the SNS can access the information. This sharing preference is at one end of the spectrum of sharing preferences.
- (2) *Share among themselves*: Only the individuals directly connected with the information can access the information. Since the scenarios presented in our study always include a number of individuals who are willing members of a group picture, the case of sharing among themselves equals the case of no sharing. That is, it does not matter whether the picture is shared among themselves on the SNS, because they shared the moment when the picture was taken, and sharing the moment can be more meaningful than sharing it online. Consequently, this preference is the direct opposite of *share with all*. Thus, we place it at the other end of the spectrum.
- (3) *Share with common friends*: Only common friends of the individuals involved in the scenario can access the information. As explained before, the space of possibilities is large and varies from user to user. Thus, there are no predefined privacy preferences that cover every possibility. Nonetheless, we consider that *share with common friends* is a reasonable compromise to present a privacy preference that lies in between the two ends of the spectrum represented by the preceding two preferences.

Two of the above sharing policies (share with all and share among themselves) represent the two ends of the spectrum of sharing possibilities and the other option is in between. We note that the sharing preferences Facebook employs in its basic privacy configuration are *only me*, *friends*, and *public*. The three sharing preferences we consider are the same ones Facebook employs, though adapted to sharing in multiuser scenarios. A benefit of choosing settings similar to Facebook, a popular SNS, is that many participants would be familiar with these privacy settings and would understand their privacy implications.

In addition to the three options above, we considered *share with all friends of all the parties* as a sharing preference to include in our study. First, we note that including an additional preference

would increase the total of number of scenarios to investigate. Second, a typical SNS user has a large number of friends (e.g., a recent study of two large samples of Facebook users found the mean values of a user's friends in the two samples to be 155.2 and 182.8 [Dunbar 2016]). As a result, sharing with all friends of all three parties involved is likely to include many strangers from each party's perspective. Thus, we concluded that, from a user's perspective, the risk-benefit tradeoff of *share with all friends of all the parties* may be similar to *share with all*. Consequently, we decided not to include sharing with all friends of all the parties as an additional preference.

2.3. Arguments

An individual involved in a multiuser scenario may employ arguments to persuade the others that his preferred sharing policy should be used, or at least considered for making the final decision. There are potentially many arguments one can employ to negotiate with or persuade another person. Following Walton et al. [2008], we understand arguments as instances of argumentation schemes, each scheme representing a form of inference from premises to a conclusion. Walton et al. show that arguments used in everyday conversation fall into a small number of schemes.

We identify four argumentation schemes that can be effective in deciding an optimal solution for a multiuser scenario: *argument from* (i) *good consequences*, (ii) *bad consequences*, (iii) *an exceptional case*, and (iv) *popular opinion*. It is important to note that we neither claim the foregoing as the only possible argumentation schemes applicable in resolving a multiuser privacy conflict nor do we seek to evaluate if these are the best possible schemes. Our objective is to evaluate if arguments, as instances of argumentation schemes, help in deciding the final policy in multiuser scenarios. Our motivation in adopting argumentation schemes is to restrict the arguments to be of a few well-defined types, instead of choosing the arguments arbitrarily.

The general structure and examples for the argumentation schemes we use is as follows.

Argument from good consequences. If A is brought about, then good consequences will occur. Therefore, A should be brought about.

An example of an argument from good consequences is: *We had a lot of fun during the party. Everybody's talking about how funny you were and they want to see your pictures. Let's share this picture with everybody.* An SNS user, who shares something, expects to receive some benefit, e.g., friendship, jobs, or other social opportunities [Ellison et al. 2007]. Thus, it is reasonable to argue that sharing certain information implies a good consequence.

Argument from bad consequences. If A is brought about, then bad consequences will occur. Therefore, A should not be brought about.

An example argument for bad consequences is: *It's a funny picture, but embarrassing since I appear drunk. I don't want strangers seeing it.* Sharing inappropriate information can hurt people's feelings and cause social tension. Thus, negative consequences can be valid arguments for not sharing.

Argument from an exceptional case. If this case is an exception, then the established rule can be waived in this case.

An example of an argument from an exceptional case is: *C'mon! it was our graduation party! something that we do only once in our lifetimes. We should show it to the world.* Although prior experience can guide future decisions, handling exceptions calls for a different approach. Instances of this scheme cover cases where an unusual privacy configuration may be called for: potentially, the opposite of the policy that might have been adopted if the arguments were not provided. Obviously, an individual must make a strong case to justify that the information is exceptional.

Argument from popular opinion. If a large majority (of the relevant group G) accepts a claim, then there is a presumption in favor of that claim. Similarly, if a large majority rejects a claim, then there is a presumption against that claim.

We do not consider explicit arguments from popular opinion. That is, in our setting, no user involved in a scenario provides an argument such as: “*the majority of us wants to share this picture only with common friends, hence, we should share it only with common friends.*” Instead, the popular opinion *emerges* when two or more users suggest the same sharing policy. Thus, although no individual explicitly employs an argument from popular opinion, the preferences of the individuals involved in a scenario may imply a popular opinion.

3. INFERENCE MODEL

We envision a computational model that incorporates the factors identified in Section 2 to recommend an optimal sharing policy for a multiuser scenario. Figure 1 shows an overview of the model.




Fig. 1. An overview of our envisioned inference model that predicts optimal sharing policy in a multiuser scenario

Given a piece of information to be shared in a multiuser scenario, we first identify all users involved in the scenario. This can be accomplished, for example, by enabling a tagging mechanism. We then ask each user identified to provide his or her preferred sharing policy for the information and an argument justifying that preference. As for the sharing context: (1) relationships can be obtained directly from the SNS; (2) users can be asked to rate the sensitivity of the information; and (3) several automated approaches exist for computing the sentiment of the information.

We note that all information above may not be available in a multiuser scenario. For example, some of the individuals involved in a multiuser scenario may not specify their preferences or arguments. Thus, we envision an inference model recommends a sharing policy based on available information and improves the recommendation as more information becomes available.

The core of our model is a *policy recommender*, which employs a machine learned *classifier* to recommend an optimal sharing policy for a given scenario. The recommender can employ a traditional classifier. However, two important challenges are (1) to engineer the features of the classifier based on the sharing context, and the preferences and arguments of the users involved, and (2) to build a dataset to train and test the classifier.

Next, we describe a crowdsourcing approach to build a dataset for training the policy recommender. This approach is useful not only for testing the recommender, but also for building a seed training set required for practical deployment of the recommender. Once the recommender is in use, additional instances can be added to the training set during deployment—as users share information via the model by adopting (or disregarding) the recommender’s suggestions.

3.1. Data Collection via Crowdsourcing

The dataset we seek must consist of a variety of multiuser scenarios and the optimal sharing policy in each of those scenarios. We expect that, in a multiuser scenario, one party’s preferred sharing policy will rarely be the sharing policy all the users involved adopt. Therefore, to collect actual first-party data, we would need to recruit cliques of participants that already know each other and ask them to recreate multiuser scenarios. This is nontrivial and challenging. Users are often reluctant to share sensitive information (one of the contextual factors in our model), biasing the study toward nonsensitive cases that users are willing to reveal [Wang et al. 2011]. An alternative is asking users

to self-report how they behave when they experience a multiuser scenario, but the results may not match participants' actual behavior because of the well-known dichotomy between users' stated privacy attitudes and their actual behavior [Acquisti and Gross 2006].

Considering the challenges above, we chose to create *situations* in which participants are *immersed* [Mancini et al. 2010] to improve behavior elicitation while avoiding biasing the study to nonsensitive situations. We present information about two or more individuals in a specific circumstance: a combination of context, preferences, and arguments. We ask participants to choose an optimal sharing policy for that circumstance.

An approach similar to ours is to ask participants to imagine themselves as involved in the presented situation. For example, showing a picture of some people and telling the participant that the picture depicts him or her with some friends. We chose against this approach for two reasons. First, participants may feel awkward imagining themselves in the presented situation (e.g., they may feel "that is silly, I would never do that!"), resulting in participants not immersing themselves effectively in the situation. Second, we have multiple users in each scenario and the optimal policy may depend on the particular user's role we ask the participant to play. This will require us to collect data from the perspective of each user involved in a scenario, increasing the number of scenarios to collect data from three fold.



We recruited participants for our study from Amazon MTurk [Paolacci et al. 2010]. We directed each participant to an external website that asked the participant to complete seven survey instruments: a presurvey questionnaire about demographics, five picture surveys (each involving a privacy conflict scenario and three sets of questionnaires), and a post-survey questionnaire about the participant's general opinions about resolving multiuser privacy conflicts. The study was approved by the IRB at North Carolina State University (details about participants and rewards are in Section 3.1.4).

3.1.1. Presurvey Questionnaire. We asked participants to report their age, gender, level of education, how frequently they use social media, and how often they share (multiuser) pictures online. Since some of the situations we presented to the participants could be inappropriate for young readers, we required participants to be older than 18 years of age and showed a disclaimer at the beginning that the survey may be inappropriate for some users.

3.1.2. Picture Survey. The picture survey is the core of our study. We first show a picture and describe a hypothetical scenario in which the picture was taken and next ask a series of questions. Table I shows two examples of picture survey. We generated these and several similar picture surveys by combining factors identified in Section 2, as described below.

- (1) Regarding context variables, we consider a predefined set of relation types, namely, *friends*, *family*, and *colleagues*. Squicciarini et al. [2011] and Toch et al. [2012] employ the same three types as an approximation to a user's relationships on a social network. Further, we assume that all individuals involved in a scenario have the same type of relationship with each other (i.e., all of them are either *friends*, *family*, or *colleagues*). Also, the pictures shown in the situations could be sensitive or nonsensitive and convey either a positive or a negative sentiment. This leads to 12 possible contexts. We found a representative picture for each of those combinations. Although we selected these 12 representative pictures, it is important to note that we ask participants to identify the contextual factors for each picture shown to them as indicated in Table I. Specifically, we ask participants to identify the type of relationship among the people involved in the scenario (family, colleagues, or friends), and rate sensitivity (Likert scale 1 = not sensitive at all, 5 = very sensitive), sentiment (1 = extremely positive, 5 = extremely negative). Appendix A includes all the 12 pictures we employed in our study along with participants' ratings of contextual factors for those pictures.
- (2) We assume that (1) an *argument from bad consequence* does not often support a *share with all* policy, and (2) an *argument from good consequence* does not often support a *share among themselves* policy. Although such combinations are conceivable, e.g., in scenarios where one of the users involved wishes a negative outcome for another, we posit that such scenarios are

Table I. Two example picture surveys (shortened version of those we used)

Picture		
Description	Aiko (C) took the picture above with her colleagues Ichiro and Katsu and, a French volunteer at the tsunami relief center	Three friends, Mark, Alex, and John, took the picture above during Mark's bachelor party on a boat in Ibiza
Rating	Identify the relationship between Aiko, Ichiro, and Katsu and rate the sensitivity and sentiment of the picture	Identify the relationship between Mark, Alex, and John and rate the sensitivity and sentiment of the picture
Context	Consider that Aiko wants to upload this picture to her social media account. What sharing policy should she apply for the picture?	Consider that Alex wants to upload this picture to his social media account. What sharing policy should he apply for the picture?
Preferences	Next, consider users' preferences as follows <i>Aiko</i> . Share among ourselves <i>Ichiro</i> . Share among ourselves <i>Katsu</i> . Share with all Considering the context and users' preferences, what sharing policy should Aiko apply for the picture?	Next, consider users' preferences as follows <i>Mark</i> . Share among ourselves <i>Alex</i> . Share with common friends <i>John</i> . Share with all Considering the context and users' preferences, what sharing policy should Alex apply for the picture?
Arguments	Finally, consider the users' preferences and arguments as follows <i>Aiko</i> . This was one of the worst natural disasters ever. Share among ourselves <i>Ichiro</i> . Tsunami was a disaster and our hand gestures are not appropriate; people may get the wrong idea. Share among ourselves <i>Katsu</i> . The picture shows the difficult situation of survivors; sharing this picture can encourage people to help. Share with all Considering the context, and users' preferences and arguments, what sharing policy should Aiko apply for the picture?	Finally, consider the users' preferences and arguments as follows <i>Mark</i> . There were some girls at the party; people might understand things the wrong way. Share among ourselves <i>Alex</i> . This was one of the best day of our lives. Share with common friends <i>John</i> . The is not like any other picture; it was from Mark's bachelor party! Share with all Considering the context, and users' preferences and arguments, what sharing policy should Alex apply for the picture?

not common in real life. We exclude these combinations for simplicity. Further, we restrict the *argument from exceptional case* to support only policies at either extreme: *share with all* or *share among themselves*. Thus, we consider six policy-argument combinations.

- (3) We limit the number of individuals involved in each scenario to three. This way, the implicit *argument from popular opinion* could work (when applied) without ties. Although some scenarios we employed showed pictures with more than three individuals, our scenarios discussed the preferences and arguments of only three of the individuals among the people involved with the picture.
- (4) We make sure that not all three individuals in a scenario use the same policy-argument combination. Our objective is to understand how a user decides a final policy given the scenario, and the preferences and arguments of others in the scenario. If all users thought the same way and wanted the same result, the solution would be trivial.

Putting the above together, we have: 12 pictures based on context, six policy-argument combinations each of first two individuals can employ, and five policy-argument combinations the last individual can employ (last restriction above). That is, we generated 2,160 scenarios. Each MTurk participant was shown five unique scenarios, making sure that no participant was shown the same picture twice. Scenarios were shown in random order to counter ordering bias. Further, we asked participants to immerse themselves in the particular scenario and ignore the resemblance or lack of resemblance of each scenario to other scenarios in which they might have seen that picture.

Following the picture and its description, we asked participants to identify the contextual factors for the scenario and answer three sets of questionnaires. We asked participants to answer these questionnaires sequentially and when answering a questionnaire, to consider only information provided to them up to that point.

Each of the three questionnaires tells participants that one of the individuals in the scenario wants to upload the picture to a social media account and asks participants what sharing policy should be applied. The participants choose one of the policies from *share with all*, *share with common friends*, and *share among themselves*. In the first questionnaire, participants know only the contextual attributes, but not the preferences or arguments of the individuals in the scenario. This case is similar to a real scenario where a user wants to upload and share information without asking others potentially concerned with the information. The second questionnaire introduces the preferences of all the users, but without their arguments supporting preferences. The third questionnaire employs all of the elements: the individuals in the scenario expose their preferences and support them with arguments. We keep the preferences fixed from the second questionnaire to the third, so we can observe the effect arguments have on the final decision.

3.1.3. Post-Survey Questionnaire. The post-survey questionnaire asks the following questions.

- (1) How important do you think the following factors are in choosing an appropriate policy when sharing information concerning multiple users on social media? (a) Relationship between stakeholders; (b) Sensitivity of the information shared; and (c) Sentiment of the information shared. The response to each factor was on a Likert scale (1 = not important at all, 5 = extremely important).
- (2) How confident will you be in choosing an appropriate policy for sharing information concerning multiple users on social media in the following cases? (a) You do not know the users' preferences or arguments; (b) You know the users' preferences, but not their arguments; and (c) You know the users' preferences and arguments. The response to each case was on a Likert scale (1 = not confident at all, 5 = extremely confident).

Responses to the above questions enable us to find correlations (or lack thereof) between participants' self-reported behavior and what they actually answered during the study.

3.1.4. Participants and Quality Control. We needed 432 participants to receive one response per scenario (each participant responds to five of the 2,160 scenarios described above). We intended to obtain two responses per scenario for completeness. However, we anticipated that some participants would begin a survey but not finish it, leaving gaps in the completed responses. To address this challenge, we launched the study on MTurk in multiple batches. For each batch, we checked if a particular survey needed additional responses and restricted the posted tasks accordingly. The final number of unique participants that completed the study was 988. At the end, each scenario had received at least two responses and some had received three responses. Compensation was provided for only those who completed all seven steps in the survey.

For quality control, we required participants to have completed at least 50 tasks on MTurk and to have had a success rate of at least 90% [Peer et al. 2014]. We included an attention check question [Gadiraju et al. 2015] in the ratings section of each picture survey, asking how many people (faces) were present in the picture, answering which requires counting from the picture. Participants answered the attention question incorrectly in a total of 38 instances (less than 1% of responses). If a participant incorrectly answered the attention check question in a picture survey, we excluded that

picture survey from analysis, and retained only those picture surveys where the participant answered the attention check questions correctly.

Table II. Demographics of MTurk participants of our study

Gender	Male: 46.3%, Female: 53.4%, Other: 0.3%
Age	18–20: 2%, 21–29: 36.6%, 30–39: 36%, 40–49: 13.7%, 50–59: 7.5%, 60 or more: 4.1%
Education	Graduate degree: 11.2%, Bachelor degree: 44.4%, College no degree: 30.9%, High school: 12.4%, Less than high school: 1%
Social media usage	Daily: 83.9%, Weekly: 12%, Monthly: 3.7%, Never: 0.4%
Pictures shared	Many (>5): 35.1%, Few (1–5): 45%, None: 18.1%, Not sure: 1.7%
Conflicts experienced	Many (>5): 2.8%, Few (1–5): 30.1%, None: 66%, Not sure: 1.1%

Table II summarizes our participants’ responses to the presurvey questionnaire. The question corresponding to the last row in the table was in the post-survey questionnaire, so that participants understand what multiuser conflicts look like before answering that question. As shown, the majority of our participants used social media on a daily basis. Over 80% of our participants had shared a picture showing multiple users and about one-third of them had experienced privacy conflicts.

3.2. Building a Training Set

We map the data collected in the MTurk study to a training set the policy recommender learns from. A training set consists of a set of data instances, where each data instance consists of a response variable (class) and a set of predictors (features). We set these as follows.

- A data instance corresponds to a participant’s response to a picture survey.
- The response variable is the final policy chosen by the participant.
- The predictors are divided into three cases based on the picture survey. The first case consists of contextual features; the second case consists of contextual and preference-based features; and the third case consists of contextual, preference-based, and argument-based features. For brevity, we refer to these cases as Context, Preferences, and Arguments, respectively.

3.2.1. Contextual Features. We compute the contextual features based on participants’ responses in the ratings section of the picture survey.

- (1) *Sensitivity* and *sentiment* each yield a feature with five levels corresponding to integer ratings from 1 to 5.
- (2) *Relationship* yields a feature with three levels: *family*, *friendship*, and *colleagues*.

3.2.2. Preference-Based Features. We compute the following features based on the preferences portrayed in the corresponding scenario (picture survey). For concreteness, we base our examples on Table I (left).

- (1) *Preference counts* represents three features corresponding to the number of participants preferring each of the three policies. In Table I (left), the preference counts are: *share with all* is 1, *share with common friends* is 0, and *share among themselves* is 2.
- (2) *Most and least restrictive policies* represent, among the preferred policies of the users in a scenario, the policy restricting sharing of information the most and the least, respectively. The order of policies from least to most restrictive is: *share with all*, *share with common friends*, and *share among themselves*. In Table I (left), the most restrictive policy is *share among themselves*, and least restrictive policy is *share with all*.
- (3) *Majority policy* represents the policy preferred by the majority of the users involved in the scenario. This feature can be *null* if there is no majority. The majority policy in Table I (left) is *share among themselves*.

3.2.3. *Argument-Based Features.* We compute argument-based features from the arguments users involved in a multiuser scenario provide for their corresponding preferences (also, recall from Section 2.3 that we restrict arguments to be instances of one of the argumentation schemes we consider).

- (1) *Argument counts* represent the number of times each type of argument is employed in the scenario. For Table I (left), the counts are: *argument from an exceptional case supporting share among themselves* is 1 (Aiko's), *argument from negative consequence supporting share among themselves* is 1 (Ichiro's), *argument from positive consequence supporting share with all* is 1 (Katsu's), and each of the remaining three argument-policy combinations is 0 (recall from Section 3.1.2 that there are six argument-policy combinations).
- (2) *Arguments supporting least restrictive policy* are the types of arguments that support the least restrictive policy. For Table I (left), the argument supporting least restrictive policy is *argument from positive consequence supporting share with all*.
- (3) *Arguments supporting most restrictive policy* are the types of arguments that support the most restrictive policy. For Table I (left), these are *argument from positive consequence* and *argument from an exceptional case* both supporting *share among themselves*.
- (4) *Arguments supporting majority policy* are the types of arguments that support the policy preferred by the majority of users. For Table I (left), these are *argument from positive consequence* and *argument from an exceptional case* both supporting *share among themselves*.

When arguments from distinct schemes support a policy, as in the arguments supporting *least restrictive policy* and *majority policy* cases above, we employ the combination of arguments as a distinct feature value. Also, if a majority policy does not exist, we set the corresponding argument-based feature to *null*.

4. EVALUATION

In this section, we describe our hypotheses and the techniques we use to evaluate those hypotheses.

4.1. Hypotheses

- (1) *H-Influence-Context:* Contextual factors, specifically, sensitivity and sentiment of the information shared, and the relationships among individuals involved influence the choice of optimal sharing policy in a multiuser scenario.
- (2) *H-Influence-Preferences:* Preferences of users involved in a multiuser scenario influence the choice of optimal policy for the scenario.
- (3) *H-Influence-Arguments:* Users' arguments for their preferred sharing policies influence the choice of optimal policy in a multiuser scenario.
- (4) *H-Prediction-Preferences:* Adding preferences to the contextual factors enhances the accuracy of an inference model predicting the optimal policy for a given multiuser scenario.
- (5) *H-Prediction-Arguments:* Adding arguments to preferences and contextual factors enhances the accuracy of an inference model predicting the optimal policy for a given multiuser scenario.
- (6) *H-Confidence-Preferences:* Adding preferences to the contextual factors enhances a user's confidence in choosing the optimal policy for a given multiuser scenario.
- (7) *H-Confidence-Arguments:* Adding arguments to preferences and contextual factors enhances a user's confidence in choosing the optimal policy for a given multiuser scenario.

4.2. Evaluation Strategy

4.2.1. *H-Influence-**. To evaluate our hypotheses about influences of different factors, we build *multinomial logistic regression* models (multiple predictors and one response variable). We adopt Akaike Information Criterion (AIC) as a measure of *goodness of fit* for these models.

$$AIC = 2k - 2 \ln L, \quad (1)$$

where L is the maximum value of the likelihood function for the model, and k the number of estimated parameters in the model. Lower values of AIC indicate better fit. AIC rewards goodness of fit (as assessed by the likelihood function), but includes a penalty that is an increasing function of the number of estimated parameters. The penalty discourages overfitting—increasing the number of parameters in the model almost always improves the goodness of the fit [Good and Hardin 2006].

Further, we focus on *coefficients* and their statistical significance. These coefficients help us understand each feature’s relative influence on the optimal policy. We employ *share among themselves* as the reference category for all models.

Linear regression models use the general linear equation $Y = b_0 + \sum b_i X_i$, where Y is a continuous response variable and b_i is the coefficient of predictor X_i . However, in logistic regression models, the coefficients are expressed in log-odds units. The coefficients show the effect of a predictor on the log odds of the optimal policy being in a given category ($Y = 1$) versus being in the reference category ($Y = 0$); that is, the odds of Y being a given category increase by a factor of e^{b_i} per unit change in X_i . The equation used by the logistic models is

$$P(Y = 1) = \frac{1}{1 + e^{-b_0 - \sum b_i X_i}} \quad (2)$$

We employ dummy variables for categorical predictors (e.g., relationship types). In these cases, the reference category is the one that is missing in the table reporting the coefficients of each model.

Since we consider several hypotheses in each model, the probability of a Type I error (false positive) is high. To reduce these errors, we apply the Holm-Bonferroni correction [Holm 1979] during the evaluation of the statistical significance of the coefficients yielded by the models.

We include all the features of a type (contextual factors, preferences, and arguments), when possible. However, some features of the same type can be highly correlated, causing *multicollinearity* [Gujarati and Porter 2009]. In that case, the coefficients may change erratically in response to small changes in the data. To counter this effect, we create independent models for highly correlated predictors. Note that this correction usually leads to higher coefficients.

Further, since each participant responded to a subset of scenarios (five out of 2,160), the personal privacy attitudes of participants may affect the estimated coefficients of the models. Thus, we employ mixed modeling [Robinson 1991] to create the logistic regression models. Mixed models offer the possibility of grouping samples of the data hierarchically. Specifically, we group responses by participants by employing the participant ID as a *random effect*. In this way, the participant ID captures the variability introduced by personal privacy attitudes in the responses. Introducing the ID enables us to obtain estimated coefficients for the other predictors (*fixed effects*) in a way that is less affected by variability in the participants’ responses.

4.2.2. H-Prediction-*. To evaluate our hypotheses about prediction, we build a *classification* model. We evaluate the prediction accuracy of these models via:

$$\begin{aligned} \text{precision} &= \frac{\text{TP}}{\text{TP} + \text{FP}}, \\ \text{recall} &= \frac{\text{TP}}{\text{TP} + \text{FN}}, \\ F_1\text{-measure} &= 2 \times \frac{\text{precision} \times \text{recall}}{\text{precision} + \text{recall}}, \end{aligned} \quad (3)$$

where TP, TN, FP, and FN refer to true and false positives and negatives. We implemented these models using Weka [Hall et al. 2009]. We perform ten-fold cross-validation and cross-validated paired *t*-test [Dietterich 1998] to test significance in differences, when required.

We note that the data we collect contains sharing policies that are considered optimal from the view points of third parties. Participants’ responses may have been different if they were personally involved in the scenarios shown during the study. Nonetheless, the data and results obtained from

the study are still useful to investigate what elements influence the optimal sharing policy and to build an initial dataset to train the policy recommender shown in Figure 1. However, in practice, to achieve a high prediction accuracy, this recommender would require adaptive capabilities to learn user-specific privacy behavior and predict sharing policies accordingly.

4.2.3. *H-Confidence.** To evaluate our hypotheses about users' confidence, we perform the Kruskal-Wallis test [Hollander and Wolfe 1999] on self-reported data from the post-survey questionnaire. The Kruskal-Wallis test is a nonparametric extension of one-way ANOVA and it relaxes the assumption that the populations underlying the samples compared are normally distributed. This test compares medians to determine if all ratings come from the same distribution. If the Kruskal-Wallis test determines that all the ratings do not come from the same distribution, we perform the multiple comparisons test [Hochberg and Tamhane 1987] to determine which variables are significantly different. As before, since we test multiple hypotheses, we employ the Holm-Bonferroni correction to reduce Type I errors.

5. RESULTS

In the following sections, we evaluate each of our hypotheses.

5.1. Context (*H-Influence-Context*)

Recall that we ask participants to identify the contextual factors for each picture shown to them. Contextual factors based on participants' responses are what we employ in all of our analyses. Figure 2 shows the distributions of the contextual variables computed from participants' responses. These distributions suggest that our picture surveys represent quite some variety in terms of relationship, sensitivity, and sentiment.




Fig. 2. Distributions of context variables computed from MTurk participants' responses

Table III shows the coefficients of the multinomial logistic regression models employing only contextual factors as predictors. The table shows the coefficients we obtain for models of optimal policies in three cases. Recall that in each case the participants provided the optimal sharing policy by considering a different set of factors: Case 1, considering only the contextual factors; Case 2, considering contextual factors and preferences; and Case 3, considering contextual factors, preferences, and arguments. The coefficients for each case are shown in two columns: one for the coefficients of the optimal policy being *share with all* versus the reference category *share among themselves*, and the other for the coefficients of the optimal policy being *share with common friends* versus the reference category. We highlight statistically significant differences at significance levels (α) of 5% and 1% with * and **, respectively. As mentioned before, where multiple hypotheses are tested, we adjust the significance levels via the Holm-Bonferroni correction.

In models for all three cases, we find that sensitivity is a significantly influential factor. Specifically, in the first case, where only contextual factors are considered, the estimated coefficient -5.485 indicates that the probability of the optimal sharing policy being *share with all* compared to the probability of being *share among themselves* decreases $e^{-5.485}$ times for each increase in the sensitivity level of the picture, assuming everything else remains unchanged.

The results also indicate that the significance of the type of the relationship decreases noticeably when preferences and arguments are considered (Cases 2 and 3). For example, in both Cases 2 and

3, *relationship = colleagues* is not statistically significant on the relative probability of the optimal policy being *share with all* versus *share among themselves*. Similarly, *relationship = friends* is not significant for the optimal policy being *share with common friends* versus the reference category. In addition, sentiment is not significant in any case.

Our intuition, based on these observations, is that, most of the times, users consider it appropriate to share a highly sensitive picture only among the individuals involved in that picture.

Table III. Regression coefficients for contextual variables

Variable	Coefficients					
	Context (Case 1)		Preferences (Case 2)		Arguments (Case 3)	
	All	Common	All	Common	All	Common
Sensitivity	-5.485**	-4.174**	-2.045**	-1.508**	-1.707**	-1.38**
Sentiment	-0.415	-0.424	-0.352	0.057	-0.193	0.048
Relationship = colleagues	-0.525**	-1.345**	-0.355	-0.457**	-0.245	-0.543**
Relationship = friends	1.162**	-0.464**	0.507**	-0.165	0.659**	-0.154
AIC	2481.1	3411.5	2471.0	4456.0	2599.0	4425.2

Further, we observe that as participants consider preferences and arguments, the coefficients of the contextual factors decrease and *AIC* increases. This indicates that, when users take preferences and arguments into account, the contextual factors' influence on the optimal sharing policy diminishes.

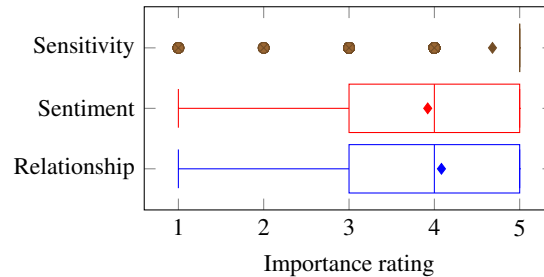


Fig. 3. Importance (from the post-survey questionnaire) of context variables

To further analyze the contextual factors, we compare these results with self-reported data collected in the post-survey questionnaire. In this way, we can observe if participants' opinions are consistent with the decisions they made in the picture survey. Figure 3 shows the boxplots for participants' importance ratings to each contextual factor. A diamond dot in the boxplot indicates the mean. Each dot outside a box indicates an outlier.

From Table III and Figure 3, we observe that the regression model and self-reported values follow a similar pattern: sensitivity is the most influential factor. However, an important distinction between the two is that although relationship and sentiment have similar ratings in the post-survey questionnaire, sentiment is not significantly influential in the regression model. Also, in the post-survey questionnaire relationship and sentiment have high (median of 4) importance ratings. However, their low regression coefficients suggest that participants did not consider relationship and sentiment quite as important as they reported in the post-survey questionnaire.

5.2. Preferences (*H-Influence-Preferences*)

The logistic models for the contextual factors suggest that employing preferences can influence the final sharing decision. Table IV shows the coefficients for models employing preference-based features. It shows only Case 2 (Preferences) and Case 3 (Arguments) since Case 1 (Context) does

not include preferences. We note that preference counts are highly correlated: when one count increases, the other two counts (naturally) decrease, causing multicollinearity. To counter this effect, we create independent models for the highly correlated predictors. Thus, the coefficient values for the preference counts features are obtained from different regression models: each of these models employs only one preference count as its predictor.

Table IV. Regression coefficients for preference features

Variable	Coefficients			
	Preferences		Arguments	
	All	Common	All	Common
# Share with all	1.32**	0.693**	1.374**	0.767**
# Share with common friends	0.136	1.029**	0.119	0.978**
# Share among themselves	-1.296**	-1.468**	-1.322**	-1.494**
Least restrictive policy = All	1.021**	0.013	1.197**	0.229
Most restrictive policy = Self	-1.421**	-2.461**	-1.31**	-2.304**
Majority policy = All	0.674**	-0.452*	0.766**	-0.538**
Majority policy = Common	-0.174	-0.195	-0.054	-0.253
Majority policy = Self	-0.584*	-1.07**	-0.536	-1.173**
AIC	2309.0	3849.2	2380.2	3820.1

First, we observe that the most restrictive policy has the most influence on the final policy, more so than even the majority policy. Second, we observe that preference counts have, in almost every case, a statistically significant influence on the final sharing policy. Third, the coefficients of the preference-based features do not change much from Preferences to Arguments. This lack of change indicates that preferences remain important even when both preferences and arguments are considered.

5.3. Arguments (*H-Influence-Arguments*)

Table V shows the regression coefficients for the argument-based features. Just as the policy counts are highly correlated, so are the argument counts. Thus, we use different models for those features.

Considering the absolute values of the coefficients (ignoring the sign), preferences supported by exceptional and positive arguments have the highest influence in the optimal sharing policy. In contrast, the arguments for negative consequences yield the lowest coefficients, suggesting that our participants weighed the benefits of sharing more than the potential risks of sharing a picture.

However, it is worth noting that *argument from bad consequences* supporting *self* (share among themselves) has the highest coefficient of all argument counts. This observation indicates that when a user makes a strong case for not sharing a picture, the other users respect that preference. This finding is consistent with those of Besmer and Lipford [2010] and Wisniewski et al. [2012].

5.4. Prediction (*H-Prediction-Preferences* and *H-Prediction-Arguments*)

The foregoing hypotheses concerned how various features influence a sharing policy. Now, we evaluate a predictive model that puts these features to use. Since our objective is to predict the actual policy (*all*, *common*, or *self*), we build a three-class classifier, which makes discrete predictions.

Figure 4 shows the distribution of sharing decisions opted by our participants in the entire dataset. Although participants' preferences were often conservative, we note that participants opted each preference on multiple occasions. This indicates that all three options we provided were useful.

We build *logistic regression* classifiers for different feature sets using the Weka implementation [Hall et al. 2009]. Table VI compares the precision, recall, and F_1 scores of the classifiers for the three cases. The size column in the table shows, for each class, the percentage of the actual (ground truth) instances belonging to that class.

Table V. Regression coefficients for argument features

Variable	Coefficients	
	All	Common
# Positive supporting all	1.883**	1.071**
# Positive supporting common	0.534**	1.475**
# Negative supporting common	-0.214	1.149**
# Negative supporting self	-2.713**	-3.436**
# Exceptional supporting all	1.274**	0.727**
# Exceptional supporting self	-1.066**	-1.197**
Positive supporting least restrictive policy	-1.612**	-1.257**
Negative supporting least restrictive policy	-1.728**	-0.168**
Exceptional supporting least restrictive policy	-1.334**	-1.043**
Positive supporting most restrictive policy	1.806**	1.354**
Negative supporting most restrictive policy	1.082*	1.269**
Exceptional supporting most restrictive policy	1.532**	1.446**
Positive supporting majority policy	2.641**	1.728**
Negative supporting majority policy	-0.111	0.671**
Exceptional supporting majority policy	-2.385**	-2.53**
AIC	2417.0	3975.6




Fig. 4. Distributions of sharing decisions computed from MTurk participants' responses

Table VI. Precision, recall, and F_1 scores of logistic regression classifiers for the three cases, considering all data instances

	Precision	Recall	F_1	Class	Size
Context	0.539	0.379	0.445	<i>all</i>	26.5%
	0.493	0.472	0.482	<i>common</i>	36.0%
	0.588	0.736	0.654	<i>self</i>	37.5%
	0.541	0.546	0.537	weighted mean	-
Preferences	0.329	0.047	0.082	<i>all</i>	15.4%
	0.606	0.622	0.614	<i>common</i>	40.8%
	0.610	0.779	0.684	<i>self</i>	43.8%
	0.565	0.602	0.563	weighted mean	-
Arguments	0.347	0.122	0.180	<i>all</i>	16.1%
	0.610	0.640	0.624	<i>common</i>	39.2%
	0.646	0.769	0.702	<i>self</i>	44.7%
	0.584	0.614	0.588	weighted mean	-

Analyzing the results obtained by the classifier, we observe that the F_1 measure increases (1) from Context to Preferences: contextual features to contextual and preference-based feature (*H-Prediction-Preferences*), and (2) from Preferences to Arguments: contextual and preference-based features to contextual, preference-based and argument-based features (*H-Prediction-Arguments*). Further, from a 10-fold cross-validated paired t -test [Dietterich 1998], we find that differences in the F_1 measures are significant ($p < 0.01$).

5.5. Confidence (*H-Confidence-Preferences* and *H-Confidence-Arguments*)

We compare a user’s confidence in choosing an optimal sharing policy given information corresponding to the three cases. Figure 5 shows the boxplots of the participants’ confidence ratings collected from the post-survey questionnaire. Further, from the Kruskal-Wallis and multiple comparison tests, we find that the median (\bar{x}) confidence rating for Preferences compared to Context, and Arguments compared to Preferences is significantly ($p < 0.01$) larger. This suggests that augmenting the contextual factors with preferences and arguments increases a user’s confidence in choosing an optimal sharing policy for a scenario.

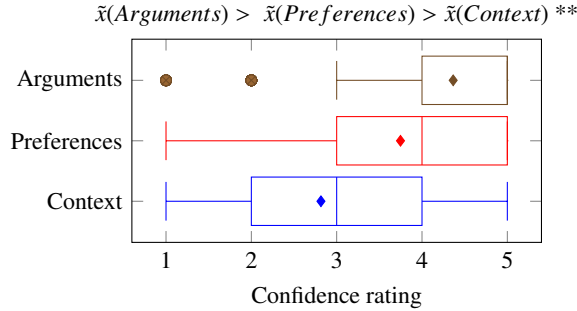


Fig. 5. Users’ confidence (from the post-survey questionnaire) in choosing the optimal policy for different cases

5.6. Analysis of Misclassified Instances

We investigate the data instances our prediction model (Section 5.4) misclassified to find out whether some combinations of scenario-defining elements make the prediction of optimal sharing policy difficult. Because manually identifying such combinations is nontrivial, since our prediction model employs multiple features, we employ two automated techniques for this purpose.

First, we cluster the misclassified instances via the expectation-maximization (EM) algorithm [Dempster et al. 1977]. Our intuition in clustering is that there can be multiple types of scenarios that make prediction difficult. The EM algorithm yields six clusters. We observe that the distribution of scenarios with different ground truth policies (*all*, *common*, and *self*) is almost uniform across the six clusters—none of the six cluster contains more than a 40% scenarios with a given optimal sharing policy as ground truth. Thus, the scenarios in each cluster are not linked to a specific optimal sharing policy (and so are our following observations about the clusters).

Next, to report the clusters visually, in a human-readable format, we build a decision tree with the six clusters as decision targets. The decision tree yields nearly perfect accuracy (99.93%). Figure 6 shows the resulting decision tree, where each leaf represents a cluster and shows its name and size. For brevity, we employ the following notation for representing the nonleaf nodes in the tree:

- $\#share[All|Common|Self]$ represents the number of users expressing that specific sharing preference in a scenario. For instance, in a scenario, if two users prefer the picture to be shared with common friends, $\#shareCommon$ equals 2.
- $\#arg[Positive|Negative|Exceptional]$, represents the number of users employing a given argument to support their preferences.
- $[most|least]RestrictivePolicy$ represents the most or least restrictive policy employed in a given scenario. For example, in a scenario, if $\#shareAll = 2$ and $\#shareCommon = 1$, then $mostRestrictivePolicy = common$.
- $argTypeFor[MostRestrictive|LeastRestrictive|Majority]Policy$ represent the types of argument employed in a given scenario to support the most restrictive, least restrictive, or majority policy, respectively.

- *sensitivity*, *sentiment*, and *relationship* represent the values of these contextual factors employed in a given scenario.

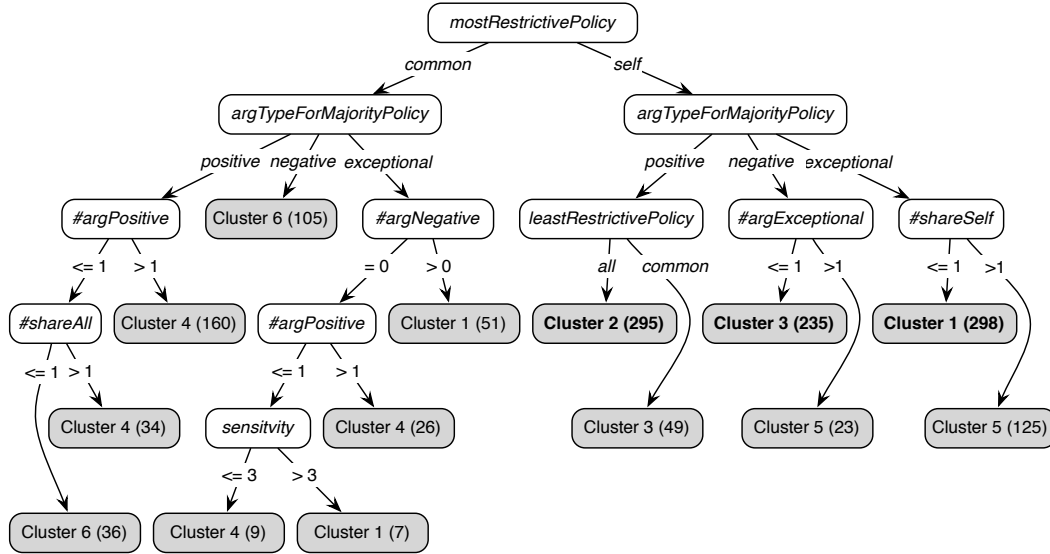


Fig. 6. A decision tree with the six clusters of misclassified scenarios as decision targets (leaf nodes; shaded).

Our objective in training the decision tree is to interpret the paths in the tree. Our intuition is that the path from the root to a cluster-leaf in the tree describes the combinations of feature values that lead to misclassifications corresponding to the cluster.

On the one hand, we observe that few nodes in the tree represent contextual factors. Specifically, there are only two paths that contain *sensitivity*; further, the clusters corresponding to these two paths represent only 16 scenarios. Their rarity suggests that combinations of contextual factors defining a scenario have little effect on whether the scenario is misclassified or not.

On the other hand, we observe that most nodes in the tree (including nodes at the top two levels) are related to users' preferences and arguments in a scenario. This suggests that preferences and arguments can help explain several misclassified instances. By observing multiple paths in the tree, we infer that misclassified scenarios often present some conflict in terms of preferences and arguments. To understand this better, consider three leaf nodes representing Clusters 1, 2, and 3 (highlighted as bold in Figure 6), which account for the majority of misclassified instances (61.4%).

- In Cluster 1, the most restrictive policy is *self*, and the majority policy is either *all* or *common* (since $\#shareSelf \leq 1$).
- In Cluster 2, the most restrictive policy is *self* and the least restrictive policy is *all*.
- In Cluster 3, a *negative consequence* argument supports the majority policy, whereas an *exceptional case* argument seems to support the minority policy.

In each of these clusters, a participant must resolve the conflict somehow to choose a final policy. We conjecture that a participant's choice of final policy for a scenario posing conflicts depends on how the participant interprets and more importantly, prioritizes the preferences and arguments involved. Further, since such prioritization is likely to be subjective, two participants may choose different policies as optimal for the same scenario, which can eventually lead to misclassification.

6. MUPPET: A TOOL FOR MULTIUSER PRIVACY DECISION ASSISTANCE

Despite achieving only moderate accuracy, our prediction model can be a valuable tool for deciding sharing policies in multiuser scenarios. To illustrate this potential benefit, consider how one would employ our model in practice. Imagine a tool, Muppet, for multiuser privacy decision assistance.

Step 1. Given a piece of information to be shared, the Muppet tool identifies users concerned with the sharing of the information. In a simple use case, one of the users involved can manually tag the information with identifiers for all the concerned users. However, this step (or parts of it) can potentially be automated. For example, many picture editors automatically detect and tag faces in a picture; if the information to be shared is a social media post, concerned users are often mentioned (e.g., by including “@Alice”) within the post.

Step 2. Muppet identifies the context of the information. Again, one of the users involved with the information can manually identify the context of the information. However, some contextual information can be automatically identified. Specifically, (1) relationship information can be obtained from an SNS provided users organize their connections on the SNS based on relationships, and (2) several automated approaches exist for computing the sentiment for both text and pictures, e.g., [You et al. 2015].

Step 3. Each relevant user specifies their preferred sharing policy for the information and provides an argument supporting their preference. In our current model, a user must choose one of the predefined sharing policies and argument types. For example, a user may choose the preferred policy as *all* and the argument as *exceptional case*.

Step 4. Given the information above (context of the information, users’ preferences, and argument), the Muppet tool recommends a sharing policy for the scenario. The concerned user may adopt or disregard the recommendation. However, if the user’s final decision can be observed, Muppet model can be retrained in an online fashion, e.g., as in Murukannaiah and Singh [2015].

Considering the difficulty of collecting data from first parties, an initial training of our model can be performed on a dataset synthesized from third parties choosing optimal sharing policies in a variety of scenarios (such as the dataset we built from our MTurk study). A downside of this approach is that such a dataset may not accurately reflect users’ sharing behavior in multiuser scenarios (since third parties do not share information for real).

An advantage of training with a dataset of third-party views, though, is that the resulting model would be functional off-the-shelf. As explained above, the privacy settings selected by the participants represent socially accepted information-flow norms. This is strengthened by the fact that there was a consensus between participants on the optimal policy for a scenario in the majority of scenarios employed in our study (specifically, for 63% of the scenarios we employed, all participants that rated a given scenario provided the same optimal sharing policy). Therefore, our model can recommend sharing policies starting from the first piece of information a user wants to share via a tool based on our model. Although the model may yield moderately accurate recommendations in the beginning, we posit that its accuracy would increase if the training set is updated with first-party data collected from groups of users, and the recommender retrained accordingly.

Personalizing Muppet

As shown in Table VI, a machine learned classifier that employs contextual features, preferences, and arguments achieves a moderate accuracy. However, this classifier ignores the characteristics of the users involved in a scenario. Since sharing norms may vary depending on individuals and groups, we envision personalizing Muppet’s recommendations depending on characteristics of the specific users involved in a multiuser scenario. We posit that such a recommender, which learns from and adapts to users, would yield a higher accuracy than the classifier we presented in this paper.

As an initial investigation to understand the influence of personal features on the optimal sharing policy in a multiuser scenario, we build a multinomial regression model employing demographic

variables (collected in the presurvey) as predictors and the optimal sharing preference provided by the participants as the response variable.

As shown in Table VII, the coefficients for most demographic variables are not statistically significant. However, two main exceptions are gender and sharing experience. Specifically, women are much more reluctant than men to share information openly once they consider the preferences and arguments of the parties involved in the multiuser scenario. This indicates that, in many cases, women dismiss the least restrictive policy proposed in the scenario. Next, the greater the sharing experience a user has, the more inclined that user is to share openly. Based on these results, Muppet can include heuristics to personalize recommendations. For example, Muppet may recommend a more conservative policy for a scenario which involves a majority of women compared to a similar scenario with a majority of men.

Table VII. Regression coefficients for demographic variables

Variable	Coefficients					
	Context (Case 1)		Preferences (Case 2)		Arguments (Case 3)	
	All	Common	All	Common	All	Common
Gender = male	0.486**	0.064**	0.515**	0.121**	0.463**	0.137**
Gender = female	0.541	0.036	-11.74**	-0.463	-11.53**	0.215
Age	-0.101*	-0.067	-0.136**	-0.026	-0.121**	0.003
Education level	-0.051	-0.042	0.057	-0.009	0.068	-0.033
Socialmedia use frequency	0.142*	0.061	0.181*	0.055	0.195*	0.027
Sharing frequency	-0.285*	-0.121	-0.309*	-0.103	-0.324*	-0.061
Sharing experience	0.281**	0.246**	0.294**	0.288**	0.296**	0.273**
Conflict experience	-0.068	-0.102	-0.029	-0.172	-0.031	-0.118
AIC	6127.15	9087.8	5872.8	8493.28	5988.5	8549.93

7. DISCUSSION

Our results show that contextual factors, preferences, and arguments influence the optimal sharing policy. However, we find a disparity between the importance ratings participants assigned to contextual factors and how these factors actually influence the optimal policy. Participants considered that the sentiment a picture conveys is an important element (rating 4 out of 5). However, according to our regression model, the influence of sentiment on the optimal policy is not statistically significant.

Our regression and inference models show that arguments help in predicting the optimal policy. Further, the majority of participants reported that they feel much more confident about what policy to apply to a shared item if they know others' arguments. These results indicate that approaches based only on aggregating user preferences, e.g., [Thomas et al. 2010; Carminati and Ferrari 2011; Hu et al. 2011] (as discussed below), may not be effective. Our results show that users are willing to accommodate others' preferences if they know their reasons. Therefore, knowing the preferences alone is not sufficient to set a policy that is satisfactory for the majority of the users involved.

The inference model achieves the F_1 score of 58.6% when trained with features based on contextual factors, preferences, and arguments. Although the accuracy of our model is not high in absolute terms, it performs noticeably better than what some naive inference models can achieve. Specifically, (1) a *random classifier* that randomly chooses one of three sharing policies as the optimal policy would yield a 33.3% accuracy, and (2) a *majority-class classifier* that always chooses the policy corresponding to the majority-class as the optimal policy would yield a 44.6% accuracy (Table VI). We are not aware of other models (that predict sharing policies in multiuser scenarios) with which we can compare our model. Our model can serve as a baseline for future works on this topic.

Our analysis of the misclassified instances reflects the complexity of considerations that go into argument evaluation. Our results suggest that arguments are crucial in helping participants conceive the information-flow norms for a given context. However, as our analysis suggests, these norms may

be subjective. That is, although some users may consider an information flow in a specific context as a norm, others may consider the same flow in the same context as a privacy violation.

7.1. Threats to Validity

We identify two threats to validity of our findings. First, any study based on surveys is susceptible to participant confusion. We mitigated this threat via explicit quality control measures commonly used in MTurk studies, such as selecting workers based on number of previous tasks completed and task success rate, and employing attention-check questions, as detailed in Section 3.1.4.

Second, to avoid the unreliability of self-reported attitudes, the well-known reluctance of participants to provide sensitive information, and the logistical challenges of finding participants in sets of three friends, we employ scenarios in which users immerse themselves and provide their assessment of the policies, preferences, and arguments under consideration. This is based on methodologies known to work well in other domains, such as [Mancini et al. 2010], as well as methodologies used in other social media privacy studies, such as [Hart et al. 2009; Klemperer et al. 2012; Squicciarini et al. 2011; Wiese et al. 2011]. Generalizations, however, should be made with caution, as participants' decisions in such immersive scenarios may not necessarily match those in real scenarios.

7.2. Limitations and Directions

We identify a number of directions for future work. For logistical reasons, we employ predefined types of arguments, preferences, and contexts. Thus, our findings are specific to these values. Nonetheless, other relationship types, sharing preferences, and argumentation schemes can be appropriate or even desirable in a multiuser scenario. A future study could attempt to understand what arguments users would employ to support their preferences during a multiuser scenario—the prospective benefit being to discover argumentation schemes geared to multiuser sharing. Automated techniques for identifying arguments [Lippi and Torroni 2016] can be valuable in this regard.

Similarly, the influence of other elements suggested in the literature on the optimal policy is worthy of investigation. For example, Nissenbaum [2004] identifies entitlement as a context-defining factor. In SNSs, users may consider that the owner of the information item is more entitled than anyone else to define the final sharing policy for that item.

We restrict the number of sharing policies to three options (share with all, common, themselves) so as to limit the possible scenarios to a manageable number in our user study. Therefore, another path for future research is to investigate how users solve privacy conflicts when fine-grained sharing preferences are used. Additionally, researchers should look into whether modifying the potential viewers can help users reach a satisfactory agreement. For example, a user who cannot be persuaded through arguments to share a picture with common friends may waive her objection if the other parties agree to exclude a specific common friend from the target audience.

In our study design, to limit the complexity of the multiuser scenarios, we restricted the negotiation among the parties to one round. Wisniewski et al. [2012] and Besmer and Lipford [2010] identify several social aspects that people employ to figure out an optimal sharing policy (e.g., trust and group norms). A future direction is to incorporate such social aspects in the model presented in this paper to determine how they evolve during successive rounds of a negotiation. For example, arguments can change from round to round depending upon the preferences expressed by others. And, group norms can dictate what arguments can be employed in a negotiation.

Our work contributes by informing the development of decision-support systems for multiuser scenarios by modeling context, preferences and arguments to find an optimal sharing policy. However, additional challenges need to be addressed in order to develop a usable user interface for any such decision-support system, such as finding an adequate trade-off between user intervention and automation to avoid burdening users. To this aim, suitable defaults, e.g., [Watson et al. 2015], or recommender systems based on machine-learning approaches, e.g., [Fang and LeFevre 2010], could be applied. As future work, we plan on building a recommender tool and evaluating it with users, so as to collect information about its accuracy and usability.

A user may employ a multiuser decision-support system such as Muppet frequently. For example, sharing group photos may be a daily routine for a user. Therefore, automating the process to a good extent is an important consideration. To this end, we alluded to automated techniques for recognizing sensitivity (e.g., [Peddinti et al. 2014]), sentiment (e.g., [You et al. 2015]), and relationships (e.g., [Murukannaiah and Singh 2012; Fogués et al. 2014]). However, this is still an active research area. Thus, a key challenge for future work on multiuser decision-support systems is evaluating not only the recommendations of such systems but also the automated techniques they build on.

In training and evaluating our policy recommender, we consider overshares (e.g, recommending *share with all* when the optimal is *share with common*) and undershares (e.g., recommending *share among themselves* when the optimal is *share with common*) as equally bad. However, in practice, recommending a conservative sharing policy may be less harmful than recommending a liberal policy when there is doubt on the optimal. Thus, a direction for future work is to consider cost-sensitive models that would weigh overshares and undershares differently.

The findings presented in this paper can be employed in a number of ways. The logistic regression models show how elements of context, preferences, and arguments influence information-flow norms that dictate sharing decisions in multiuser scenarios. Therefore, formal models of norms, e.g., [Barth et al. 2006], [Singh 2013], and [Criado and Such 2016], can be adapted for multiuser scenarios by including those elements. Also, the reported accuracy of our classifier can serve as a benchmark for future classifiers.

Our study is cross-sectional (one time) and in the scenarios of the study, decisions are to be made in one round. Thus, our study does not provide data to understand how users deal with situations where their arguments are countered by others' arguments.

Finally, each participant in our study responded to five scenarios. Although we employed random effects in building the logistic regression models, personal privacy attitudes may bias the results. A future study could require participants to engage in negotiation and persuasion to decide an optimal sharing policy for a scenario. Such a study requires coordinating multiple participants; thus, conducting such a study on MTurk is nontrivial. Employing automated agents (built based on the data we collected in our study) against one or a few real participants is a viable and interesting direction.

8. RELATED WORK

As Wisniewski et al. [2012] identify, the lack of tools to manage multiuser scenarios forces SNS users to employ a number of *coping strategies*, which are of limited effectiveness in practice. Such strategies include blocking other users, filtering friendship requests by the number of common friends, and self censorship. In a similar previous study, Lampinen et al. [2011] explore SNS-users' perceptions of control over online disclosure through a series of interviews. The qualitative data obtained in their study suggests that users manage interpersonal boundaries both individually and collaboratively. Lampinen et al. further classify the strategies that SNS users employ to manage their privacy into two super classes: preventive and corrective. Strategies in both classes can be accomplished through individual or collaborative means. For example, a collaborative preventive strategy is asking for approval before disclosing content from those involved. These two studies show that SNS users need functionalities to manage multiuser privacy. Our work is a stepping stone toward offering such functionality.

Besmer and Lipford [2010] propose a method where the owner of a picture in a multiuser scenario is in charge of deciding the sharing policy and the other users involved can suggest privacy preferences. They developed a Facebook application that enables a user to send privacy suggestions to the owner of a picture where that user appears. Besmer and Lipford show that the owners of pictures usually entertain and accept suggestions from others when they decide sharing policies—which is consistent with our findings. A shortcoming of their approach is that it is *manual*, i.e., the owner of a picture must determine an optimal solution based on suggestions from others. Such manual effort may overload the owner of the picture.

Thomas et al. [2010] propose *veto voting* as a direct approach to manage multiuser sharing policies. That is, denying access takes precedence over granting access. Thus, if an individual wants to

share some information with a given user, but another individual does not, the information is not shared. Whereas this approach protects privacy, it may lead to utility loss. For example, suppose Alice and Bob appear together in a picture. Bob initially opposes sharing the picture with Charlie as he does not know him. However, if Alice tells him that Charlie is her friend and that everything is OK, then Bob may accept sharing with Charlie. Had veto voting been applied, the picture would have not been shared with Charlie, thereby missing, for example, a potential opportunity for Bob to be friends with Charlie.

Other approaches too tackle multiuser privacy conflicts. However, they exhibit various limitations, which we describe next. Some of these approaches require too much human intervention during conflict resolution: Wishart et al. [2010] require users to solve the conflicts *manually*. Likewise, Squicciarini et al. [2009] requires users to resolve conflicts nearly *manually* by participating in difficult-to-comprehend auctions with fake money to handle every possible conflict.

Approaches that provide automated support [Carminati and Ferrari 2011; Hu et al. 2011; Thomas et al. 2010] help resolve multiuser conflicts, but they usually consider only one fixed way of aggregating user preferences, without considering how users would compromise and the concessions they might be willing to make in a specific situation. Hu et al. [2013] consider more than one way of aggregating user preferences, but the user who uploads an item chooses the aggregation method to be applied, which becomes a unilateral decision without considering any input from others. Clearly, solutions that do not consider input from the other users involved may lead to solutions that are far from what some users would be willing to accept. This may lead users to manually resolve conflicts most of the time. Such and Criado [2014; 2016] provide an improvement over the fixed ways of aggregating user preferences by automatically inferring the particular situation for the conflict and applying the *concessions* that are known to happen during offline negotiations in those situations [Lampinen et al. 2011; Wisniewski et al. 2012]. Situations are modelled considering the individual preferences of each user involved, the sensitivity of the content, and the relationships to the potential audience. While this approach captures and adapts to known situations, it may not capture opportunistic concessions or agreements that may arise in potentially unknown situations.

Some recent works propose game-theoretic mechanisms to tackle multiuser privacy conflicts. They define negotiation protocols, which are a means of standardising the communication between participants in the process of negotiating a solution to a multiuser privacy conflict by defining how the participants can interact with each other. These protocols are then enacted by users themselves manually [Hu et al. 2014] or automatically by software agents [Such and Rovatsos 2016] to negotiate an agreed sharing decision for a particular item. Participants can follow different strategies when enacting the negotiation protocols, and these strategies are analysed using well-known game-theoretic solution concepts such as the Nash equilibrium. However, such proposals may not work well in practice since they assume users are perfectly rational and do not capture the social idiosyncrasies that users consider in real life [Lampinen et al. 2011; Wisniewski et al. 2012].

Iliia et al. [2015] present a mechanism to enforce fine-grained access control in pictures by blurring the faces of the users depicted in the picture based on each users' access control list. This approach can limit the utility of sharing information. However, used in conjunction with a decision-support mechanism based our findings, Iliia et al.'s approach could enforce access control differently for different users in case they cannot agree on a sharing policy.

Arguments are commonly used to resolve conflicts in other domains. Murukannaiah et al. [2015] employ arguments in requirements engineering to resolve conflicts in stakeholders' goals. Their findings that arguments lead to consistent responses aligns with our observation. Williams and Williamson [2006] incorporate arguments and Bayesian networks for breast cancer prognosis, where they exploit arguments to develop an explanation for the prognosis. A similar application in the privacy domain is to explain to a user, via arguments, the consequences of a particular sharing decision.

9. CONCLUSIONS

Sharing all kinds of information on SNSs is routine for many people. Examples of such information are a picture from the Christmas party and a tweet about your imminent trip with friends. The

information shared often involves multiple users. When the privacy preferences of two or more users do not align, they should be able to negotiate so as to balance privacy and utility for each user. The related literature proposes methods based on aggregating privacy preferences. However, aggregation does not capture how people align with others' preferences. In contrast, we propose employing arguments to support preferences. We present an inference model that employs sharing context, and users' preferences and arguments to predict an optimal sharing policy in a multiuser scenario. We conduct a crowdsourcing study to train and test our model.

Via a series of multinomial logistic regression models, we show that all three feature types we consider influence the optimal sharing policy in a multiuser scenario. We find that (1) among the contextual variables, sensitivity has the highest influence on the optimal policy; (2) among preference-based features, the most restrictive policy has the highest influence on the optimal policy and, in particular, not the majority policy; and (3) users may value arguments for sharing more than arguments for not sharing; however, if the argument for not sharing is an exceptional case argument, users usually support not sharing.

By training an inference model for each feature type, we find that a model employing argument-based features predicts optimal policy with higher accuracy than those not employing arguments. Further, from self reported data, we find that introducing arguments increases a user's confidence in choosing the final sharing policy. Further, we investigate the data instances our prediction model misclassified to find out whether some combinations of scenario-defining elements make the prediction of optimal sharing policy difficult. We find that conflicts between arguments and preferences generate the majority of misclassification. This indicates that users interpret and prioritize arguments subjectively, which suggests that a tool that aims at helping users deal with multiuser conflicts must be adaptive and learn from the user.

REFERENCES

- Alessandro Acquisti and Ralph Gross. 2006. Imagined Communities: Awareness, Information Sharing, and Privacy on the Facebook. In *Proceedings of the 6th International Conference on Privacy Enhancing Technologies (PET)*. Springer-Verlag, Cambridge, UK, 36–58.
- Adam Barth, Anupam Datta, John C. Mitchell, and Helen Nissenbaum. 2006. Privacy and Contextual Integrity: Framework and Applications. In *Proceedings of the 2006 IEEE Symposium on Security and Privacy (SP)*. IEEE Computer Society, Oakland, CA, 184–198.
- Andrew Besmer and Heather Lipford. 2010. Moving Beyond Untagging: Photo Privacy in a Tagged World. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI)*. ACM, Atlanta, 1563–1572.
- Barbara Carminati and Elena Ferrari. 2011. Collaborative access control in on-line social networks. In *Proceedings of the 7th International Conference on Collaborative Computing: Networking, Applications and Worksharing (CollaborateCom)*. 231–240.
- Natalia Criado and Jose M Such. 2015. Implicit contextual integrity in online social networks. *Information Sciences* 325 (2015), 48–69.
- Natalia Criado and Jose M. Such. 2016. Selective Norm Monitoring. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9-15 July 2016*. 208–214.
- Arthur P. Dempster, Nan M. Laird, and Donald B. Rubin. 1977. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)* 39, 1 (1977), 1–38.
- Thomas G. Dietterich. 1998. Approximate Statistical Tests for Comparing Supervised Classification Learning Algorithms. *Neural Computation* 10, 7 (Oct. 1998), 1895–1923.
- Robin I. M. Dunbar. 2016. Do online social media cut through the constraints that limit the size of offline social networks? *Royal Society Open Science* 3, 1 (2016). DOI: <http://dx.doi.org/10.1098/rsos.150292>
- Nicole B. Ellison, Charles Steinfield, and Cliff Lampe. 2007. The benefits of Facebook “friends:” Social capital and college students' use of online social network sites. *Journal of Computer-Mediated Communication* 12, 4 (2007), 1143–1168.
- Lujun Fang and Kristen LeFevre. 2010. Privacy Wizards for Social Networking Sites. In *Proceedings of the 19th International Conference on World Wide Web (WWW)*. ACM, Raleigh, North Carolina, 351–360.
- Ricard L Fogués, Jose M Such, Agustin Espinosa, and Ana Garcia-Fornes. 2014. BFF: A tool for eliciting tie strength and user communities in social networking services. *Information Systems Frontiers* 16, 2 (2014), 225–237.
- Ricard L. Fogués, Jose M. Such, Agustín Espinosa Minguet, and Ana García-Fornes. 2015. Open Challenges in Relationship-Based Privacy Mechanisms for Social Network Services. *International Journal of Human-Computer Interaction* 31, 5 (2015), 350–370.

- Ujwal Gadiraju, Ricardo Kawase, Stefan Dietze, and Gianluca Demartini. 2015. Understanding Malicious Behavior in Crowdsourcing Platforms: The Case of Online Surveys. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI)*. ACM, Seoul, 1631–1640.
- Phillip I. Good and James W. Hardin. 2006. *Common Errors in Statistics (And How to Avoid Them)*. Wiley, New York.
- Damodar N. Gujarati and Dawn C. Porter. 2009. *Basic Econometrics* (Fifth ed.). McGraw-Hill/Irwin, New York, Chapter 10, 120–150. Multicollinearity: What happens if the regressors are correlated?
- Mark Hall, Eibe Frank, Geoffrey Holmes, Bernhard Pfahringer, Peter Reutemann, and Ian H. Witten. 2009. The WEKA Data Mining Software: An Update. *SIGKDD Explorations Newsletter* 11, 1 (Nov. 2009), 10–18.
- Michael Hart, Claude Castille, Rob Johnson, and Amanda Stent. 2009. Usable Privacy Controls for Blogs. In *Proceedings of the International Conference on Computational Science and Engineering - Volume 04 (CSE)*. IEEE Computer Society, 401–408.
- Yosef Hochberg and Ajit C. Tamhane. 1987. *Multiple Comparison Procedures*. John Wiley & Sons, New York.
- Bernie Hogan. 2010. The Presentation of Self in the Age of Social Media: Distinguishing Performances and Exhibitions Online. *Bulletin of Science, Technology and Society* 30, 6 (2010), 377–386.
- Myles Hollander and Douglas A. Wolfe. 1999. *Nonparametric Statistical Methods*. Wiley, New York.
- Sture Holm. 1979. A Simple Sequentially Rejective Multiple Test Procedure. *Scandinavian Journal of Statistics* 6, 2 (1979), 65–70. <http://www.jstor.org/stable/4615733>
- Hongxin Hu, Gail-Joon Ahn, and Jan Jorgensen. 2011. Detecting and resolving privacy conflicts for collaborative data sharing in online social networks. In *Proceedings of the 27th Annual Computer Security Applications Conference (ACSAC)*. ACM, Orlando, 103–112.
- Hongxin Hu, Gail-Joon Ahn, and Jan Jorgensen. 2013. Multiparty Access Control for Online Social Networks: Model and Mechanisms. *IEEE Transactions on Knowledge and Data Engineering* 25, 7 (July 2013), 1614–1627.
- Hongxin Hu, Gail-Joon Ahn, Ziming Zhao, and Dejun Yang. 2014. Game Theoretic Analysis of Multiparty Access Control in Online Social Networks. In *Proceedings of the 19th ACM Symposium on Access Control Models and Technologies (SACMAT)*. ACM, London, Ontario, 93–102.
- Panagiotis Iliia, Iasonas Polakis, Elias Athanasopoulos, Federico Maggi, and Sotiris Ioannidis. 2015. Face/Off: Preventing Privacy Leakage From Photos in Social Networks. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security (CCS)*. ACM, Denver, Colorado, 781–792.
- Sanjay Kairam, Mike Brzozowski, David Huffaker, and Ed Chi. 2012. Talking in Circles: Selective Sharing in Google+. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI)*. ACM, Austin, 1065–1074.
- Peter Klemperer, Yuan Liang, Michelle Mazurek, Manya Sleeper, Blase Ur, Lujo Bauer, Lorrie Faith Cranor, Nitin Gupta, and Michael Reiter. 2012. Tag, you can see it!: Using tags for access control in photo sharing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI)*. Austin, 377–386.
- Airi Lampinen, Vilma Lehtinen, Asko Lehmuskallio, and Sakari Tamminen. 2011. We’re in it Together: Interpersonal Management of Disclosure in Social Network Services. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI)*. ACM, Vancouver, 3217–3226.
- Marco Lippi and Paolo Torroni. 2016. Argumentation Mining: State of the Art and Emerging Trends. *ACM Transactions on Internet Technology* 16, 2, Article 10 (March 2016), 25 pages.
- Clara Mancini, Yvonne Rogers, Arosha K. Bandara, Tony Coe, Lukasz Jedrzejczyk, Adam N. Joinson, Blaine A. Price, Keerthi Thomas, and Bashar Nuseibeh. 2010. ContraVision: Exploring Users’ Reactions to Futuristic Technology. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI)*. ACM, 153–162.
- Alice Marwick and Danah Boyd. 2011. I tweet honestly, I tweet passionately: Twitter users, context collapse, and the imagined audience. *New Media & Society* 13, 1 (September 2011), 114–133.
- Pradeep K. Murukannaiah, Nirav Ajmeri, and Munindar P. Singh. 2016. Engineering Privacy in Social Applications. *IEEE Internet Computing* 20, 2 (March 2016), 72–76.
- Pradeep K. Murukannaiah, Anup K. Kalia, Pankaj R. Telang, and Munindar P. Singh. 2015. Resolving Goal Conflicts via Argumentation-Based Analysis of Competing Hypotheses. In *Proceedings of the 23rd IEEE International Requirements Engineering Conference*. Ottawa, 156–165.
- Pradeep K. Murukannaiah and Munindar P. Singh. 2012. Platys Social: Relating Shared Places and Private Social Circles. *IEEE Internet Computing* 16, 3 (May 2012), 53–59. DOI: <http://dx.doi.org/10.1109/MIC.2011.106>
- Pradeep K. Murukannaiah and Munindar P. Singh. 2015. Platys: An Active Learning Framework for Place-Aware Application Development and Its Evaluation. *ACM Transactions on Software Engineering and Methodology* 24, 3, Article 19 (May 2015), 33 pages.
- Helen Nissenbaum. 2004. Privacy as contextual integrity. *Washington Law Review* 79 (2004), 119.
- Gabriele Paolacci, Jesse Chandler, and Panagiotis G. Ipeirotis. 2010. Running experiments on Amazon Mechanical Turk. *Judgment and Decision Making* 5, 5 (2010), 411–419.

- Sai Teja Peddinti, Aleksandra Korolova, Elie Bursztein, and Geetanjali Sampemane. 2014. Cloak and swagger: Understanding data sensitivity through the lens of user anonymity. In *2014 IEEE Symposium on Security and Privacy*. IEEE, 493–508.
- Eyal Peer, Joachim Vosgerau, and Alessandro Acquisti. 2014. Reputation as a sufficient condition for data quality on Amazon Mechanical Turk. *Behavior Research Methods* 46, 4 (2014), 1023–1031.
- George K. Robinson. 1991. That BLUP is a good thing: The estimation of random effects. *Statist. Sci.* 6, 1 (1991), 15–32. <http://www.jstor.org/stable/2245695>
- Yan Shvartzshnaider, Schrasing Tong, Thomas Wies, Paula Kift, Helen Nissenbaum, Lakshminarayanan Subramanian, and Prateek Mittal. 2016. Learning Privacy Expectations by Crowdsourcing Contextual Informational Norms. In *Proceedings of the Fourth AAI Conference on Human Computation and Crowdsourcing (HCOMP)*.
- Stefan Siersdorfer, Enrico Minack, Fan Deng, and Jonathon S. Hare. 2010. Analyzing and Predicting Sentiment of Images on the Social Web. In *Proceedings of the 18th ACM International Conference on Multimedia (MM)*. ACM, Firenze, Italy, 715–718.
- Munindar P. Singh. 2013. Norms as a Basis for Governing Sociotechnical Systems. *ACM Transactions on Intelligent Systems and Technology* 5, 1, Article 21 (Dec. 2013), 23 pages. DOI: <http://dx.doi.org/10.1145/2542182.2542203>
- Manya Sleeper, Rebecca Balebako, Sauvik Das, Amber Lynn McConahy, Jason Wiese, and Lorrie Faith Cranor. 2013. The post that wasn't: Exploring self-censorship on Facebook. In *Proceedings of the conference on Computer supported cooperative work (CSCW)*. ACM, 793–802.
- Anna Cinzia Squicciarini, Mohamed Shehab, and Federica Paci. 2009. Collective Privacy Management in Social Networks. In *Proceedings of the 18th International Conference on World Wide Web (WWW)*. ACM, Madrid, 521–530.
- Anna Cinzia Squicciarini, Smitha Sundareswaran, Dan Lin, and Josh Wede. 2011. A3P: Adaptive Policy Prediction for Shared Images over Popular Content Sharing Sites. In *Proceedings of the 22nd ACM Conference on Hypertext and Hypermedia (HT)*. ACM, New York, 261–270.
- Stefan Stieglitz and Linh Dang-Xuan. 2013. Emotions and information diffusion in social mediasentiment of microblogs and sharing behavior. *Journal of Management Information Systems* 29, 4 (2013), 217–248.
- Jose M. Such and Natalia Criado. 2014. Adaptive Conflict Resolution Mechanism for Multi-party Privacy Management in Social Media. In *Proceedings of the 13th Workshop on Privacy in the Electronic Society*. ACM, 69–72.
- Jose M. Such and Natalia Criado. 2016. Resolving Multi-party Privacy Conflicts in Social Media. *IEEE Transactions on Knowledge and Data Engineering* 28, 7 (2016), 1851–1863.
- Jose M. Such, Agustín Espinosa, and Ana García-Fornes. 2014. A survey of privacy in multi-agent systems. *Knowledge Engineering Review* 29, 03 (2014), 314–344.
- Jose M. Such and Michael Rovatsos. 2016. Privacy Policy Negotiation in Social Media. *ACM Transactions on Autonomous and Adaptive Systems* 11, 4 (2016), 1–29. Issue 1.
- Kurt Thomas, Chris Grier, and David M. Nicol. 2010. unFriendly: Multi-party Privacy Risks in Social Networks. In *Proceedings of the 10th International Conference on Privacy Enhancing Technologies (PET)*. Springer-Verlag, Berlin, 236–252.
- Eran Toch, Yang Wang, and Lorrie Faith Cranor. 2012. Personalization and Privacy: A Survey of Privacy Risks and Remedies in Personalization-based Systems. *User Modeling and User-Adapted Interaction* 22, 1-2 (April 2012), 203–220.
- Douglas Walton, Christopher Reed, and Fabrizio Macagno. 2008. *Argumentation Schemes*. Cambridge University Press.
- Yang Wang, Gregory Norcie, Saranga Komanduri, Alessandro Acquisti, Pedro Giovanni Leon, and Lorrie Faith Cranor. 2011. 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*. ACM, 10.
- Jason Watson, Heather Lipford, and Andrew Besmer. 2015. Mapping User Preference to Privacy Default Settings. *ACM Transactions on Computer-Human Interaction* 22, 6, Article 32 (Nov. 2015), 20 pages.
- Jason Wiese, Patrick Gage Kelley, Lorrie Faith Cranor, Laura Dabbish, Jason I. Hong, and John Zimmerman. 2011. Are You Close with Me? Are You Nearby?: Investigating Social Groups, Closeness, and Willingness to Share. In *Proceedings of the 13th International Conference on Ubiquitous Computing (UbiComp)*. ACM, New York, 197–206.
- Matt Williams and Jon Williamson. 2006. Combining Argumentation and Bayesian Nets for Breast Cancer Prognosis. *Journal of Logic, Language, and Information* 15, 1–2 (2006), 155–178.
- Ryan Wishart, Domenico Corapi, Srdjan Marinovic, and Morris Sloman. 2010. Collaborative Privacy Policy Authoring in a Social Networking Context. In *Proceedings of the IEEE International Symposium on Policies for Distributed Systems and Networks (POLICY)*. 1–8.
- Pamela Wisniewski, Heather Lipford, and David Wilson. 2012. Fighting for My Space: Coping Mechanisms for SNS Boundary Regulation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI)*. ACM, Austin, 609–618.
- Quanzeng You, Jiebo Luo, Hailin Jin, and Jianchao Yang. 2015. Robust Image Sentiment Analysis Using Progressively Trained and Domain Transferred Deep Networks. In *Proceedings of the Twenty-Ninth AAI Conference on Artificial Intelligence (AAAI)*. AAAI Press, 381–388.

Online Appendix to: Sharing Policies in Multiuser Privacy Scenarios: Incorporating Context, Preferences, and Arguments in Decision Making


Ricard L. Fogues, Universitat Politècnica de València
Pradeep K. Murukannaiah, Rochester Institute of Technology
Jose M. Such, King's College London
Munindar P. Singh, North Carolina State University

A. PICTURE SURVEY SCENARIOS

Table VIII shows the 12 pictures we employ in the picture surveys (Section 3.1.2). For each picture, the table provides a description, context, and arguments.

- The description includes the scenario in which the picture was taken and the people involved in its sharing.
- The context includes relationship, sensitivity rating, and sentiment rating of the picture. Note that the context was identified by MTurk participants who answered picture surveys corresponding to the picture.
- The arguments are of types: positive consequence, negative consequence, and exceptional case. Note that (1) each argument is employed along with a sharing policy, and (2) different combinations of arguments are employed in each picture survey (examples in Table I).

Table VIII: A list of pictures we employ in the picture surveys

Picture and Context		Relationship: Friends (92.2%) Sensitivity rating: $\mu = 1.56$ ($\sigma = 0.96$) Sentiment rating: $\mu = 1.77$ ($\sigma = 1.46$)
Description	Tim, Ashley, and Jerry just graduated. Tim's father took the picture above after the graduation ceremony. Tim wants to upload the picture to his social media account.	
Arguments	<p><i>Positive consequence argument.</i> People we know will be happy to see that we are finally done with college.</p> <p><i>Negative consequence argument.</i> Our gestures are not appropriate for a moment like this; people might think that we did not take our college time seriously.</p> <p><i>Exceptional case argument.</i> This is not like any of our other pictures. It was our graduation, which happens only once in our lifetimes.</p>	

Picture and Context

Relationship: Friends (98.3%)
 Sensitivity rating: $\mu = 3.29$ ($\sigma = 1.16$)
 Sentiment rating: $\mu = 3.82$ ($\sigma = 1.11$)

Description


Three friends, Santosh, Arun, and Nitin, decided to perform some stunts on a motorcycle. Unfortunately, while performing a stunt, Arun and Nitin had a minor accident. Santosh took the picture below at that very moment. Santosh wants to upload the picture to his social media account.

Arguments

Positive consequence argument. Fortunately, none of us got hurt. This picture makes anyone who sees it laugh out loud.

Negative consequence argument. People looking at this picture may think that we are reckless drivers, which is not true.

Exceptional case argument. Motorbike stunts are not something we do every-day.

Picture and Context

Relationship: Friends (96.5%)
 Sensitivity rating: $\mu = 3.78$ ($\sigma = 1.16$)
 Sentiment rating: $\mu = 2.99$ ($\sigma = 1.26$)

Description

Three friends, Mark (the groom), Alex, and John, go on a boat in Ibiza during a bachelor party. They get drunk and meet some girls. This is one of the pictures Alex took during that party. Alex wants to upload the picture to his social media account, the day after the party.

Arguments

Positive consequence argument. This was one of the best day of our lives; we should share these good memories with others.

Negative consequence argument. There were some girls in the party; people might understand things the wrong way.

Exceptional case argument. This is not like any of our other pictures. This was Mark's bachelor party!

Picture and Context

Relationship: Friends (94.5%)
 Sensitivity rating: $\mu = 4.11$ ($\sigma = 1.13$)
 Sentiment rating: $\mu = 4.24$ ($\sigma = 1.20$)

Description The picture above is from an end-of-semester party Vanessa (the girl wearing dark shirt with the US flag on it) hosted. The party went wild and a neighbor called the police. There was some tension after the police arrived and a few people invited to the party were arrested. The picture also shows Vanessa's friends Natasha (the girl on the grass) and Jessica (the girl at the very left of the picture). Vanessa and her friends felt that the police abused their authority. Vanessa wants to upload the picture to her social media account, a few days after the incident.

Arguments

Positive consequence argument. Police arrested some of us in the party. This picture shows that police abused their power and arrested some of us for a silly reason.

Negative consequence argument. We do not deserve what happened and this picture brings back those traumatic memories.

Exceptional case argument. This pictures reminds us of one of the wildest party we have had.

Picture and Context



Relationship: Family (97.7%)
Sensitivity rating: $\mu = 1.90$ ($\sigma = 1.25$)
Sentiment rating: $\mu = 2.01$ ($\sigma = 1.61$)

Description The Moore brothers and their parents, wives, and children took part in a photoshoot. The following is the best picture from the photo shoot. Frank wants to upload the picture to his social media account.

Arguments

Positive consequence argument. We took this picture so that our loved ones can see it an remember us.

Negative consequence argument. Our children appear in this picture; what others can do with this picture concerns me.

Exceptional case argument. This is not like any other picture; we hired a photographer to take it.

Picture and Context



Relationship: Family (96.7%)
Sensitivity rating: $\mu = 2.52$ ($\sigma = 1.28$)
Sentiment rating: $\mu = 3.09$ ($\sigma = 1.06$)

Description Jerry (the grandfather dressed as Santa), Timmy (Jerry's grandson), and April (Jerry's granddaughter) took a picture during Christmas night. Jerry wants to upload the picture to his social media account.

Arguments

Positive consequence argument. It was a lovely evening and this picture brings back good memories.

Negative consequence argument. Poor Timmy got scared. I do not think it is fair to Timmy to share a picture of him crying and scared of his grandpa.

Exceptional case argument. This was the first time Jerry and Timmy took a picture together.

Picture and Context

Relationship: Family (95.3%)
Sensitivity rating: $\mu = 4.51$ ($\sigma = 0.76$)
Sentiment rating: $\mu = 2.45$ ($\sigma = 1.35$)

Description

Dolores and Philip decide to have their baby, Rose, at home with the help of Ann, who is Dolores' sister and a doula. They took the picture below during the labor. Philip wants to upload the picture to his social media account, a few days after Rose was born.

Arguments

Positive consequence argument. This picture shows that Rose was born peacefully at our home surrounded by those who love us.

Negative consequence argument. The idea of giving birth at home was to be at a private place surrounded by only the people we love.

Exceptional case argument. This is not like any other family picture; this shows that Rose is coming to our lives.

Picture and Context

Relationship: Family (96.5%)
Sensitivity rating: $\mu = 4.21$ ($\sigma = 1.01$)
Sentiment rating: $\mu = 3.93$ ($\sigma = 1.16$)

Description

Mary, Sophia, and Charles attend their mother's funeral. Another family member takes some pictures and circulates them among the family members. Mary wants to upload the picture to her social media account.

Arguments

Positive consequence argument. Many people knew our mother and loved her, including our friends. When people see this picture they will remember her and see that we all were there to say goodbye.

Negative consequence argument. This picture may appear highly inappropriate to many people.

Exceptional case argument. This is not like any other picture, we were saying goodbye to mom.

Picture and Context

Relationship: Colleagues (94.4%)
 Sensitivity rating: $\mu = 1.77$ ($\sigma = 1.10$)
 Sentiment rating: $\mu = 2.83$ ($\sigma = 0.92$)

Description

Maria, Bonita, and Felipe, three junior employees in a company, attend a business lunch in which they meet their seniors. One of the other employees took the following picture and sent it to Maria. Maria wants to upload the picture to her social media account.

Arguments

Positive consequence argument. This picture shows that we are making good progress in our careers.

Negative consequence argument. This was a professional event and our seniors might want to keep it private.

Exceptional case argument. This is an exceptional event since we attended a professional party for the first time.

Picture and Context

Relationship: Colleagues (86.7%)
 Sensitivity rating: $\mu = 2.67$ ($\sigma = 1.32$)
 Sentiment rating: $\mu = 2.69$ ($\sigma = 1.24$)

Description

Aiko (C) took the picture above with her colleagues Ichiro and Katsu and, a French volunteer at the tsunami relief center. Aiko wants to upload this picture to her social media account.

Arguments

Positive consequence argument. The picture shows the difficult situation in which the survivors live. Sharing this can encourage people to help.

Negative consequence argument. Tsunami was a disaster and our gestures are not appropriate. People may get the wrong idea.

Exceptional case argument. This was one of the worst natural disasters.

Picture and Context

Relationship: Colleagues (92.9%)
 Sensitivity rating: $\mu = 3.26$ ($\sigma = 1.41$)
 Sentiment rating: $\mu = 2.46$ ($\sigma = 1.50$)

Description

Jerry, Laura, and Sabrina work together in a company. They were asked to attend the Christmas party dressed. However, a guy in their company (the one in pink dress) brought the whole dressing to a new level. They took the following picture at the party. Jerry wants to upload the picture to his social media account, a few days after the party.

Arguments

Positive consequence argument. People think that I have a boring life because I work at a boring place; this will prove them wrong.

Negative consequence argument. This is embarrassing; people will pick on us because of this picture.

Exceptional case argument. This is an exceptional event since a Christmas party happens only once a year.

Picture and Context

Relationship: Colleagues (98.0%)
 Sensitivity rating: $\mu = 3.52$ ($\sigma = 1.26$)
 Sentiment rating: $\mu = 3.50$ ($\sigma = 1.00$)

Description

The hospital where Bryan, Martin, and Sophia work has recently changed its shift policy making shifts much longer. Doctors complain that these shifts leave them exhausted. During one such long shifts, at 4am, Bryan takes a picture of his two colleagues Martin and Sophia sleeping while they wait for another patient to come to emergencies. Bryan wants to upload the picture to his social media account, a few days after the picture was taken.

Arguments

Positive consequence argument. This new shift policy is too demanding. A picture like this can convince the management to change the policy.

Negative consequence argument. If people think that we sleep on our job, they won't trust the hospital.

Exceptional case argument. A doctor sleeping on a chair is exceptional.