Argumentation for Multi-party Privacy Management

(Position Paper)

Ricard Fogues
Universitat Politècnica de València
Camí de Vera, SN
Valencia, Spain
rilopez@dsic.upv.es

Pradeep Murukannaiah
North Carolina State University
Raleigh, NC, USA
pmuruka@ncsu.edu

Jose M. Such
Lancaster University
Lancaster, UK
j.such@lancaster.ac.uk

Agustin Espinosa
Universitat Politècnica de València
Camí de Vera, SN
Valencia, Spain
aespinos@dsic.upv.es

Ana Garcia-Fornes
Universitat Politècnica de València
Camí de Vera, SN
Valencia, Spain
agarcia@dsic.upv.es

Munindar Singh
North Carolina State University
Raleigh, NC, USA
mpsingh@ncsu.edu

ABSTRACT

Social network services enable users to share large quantities of private information. Often, the shared information concerns individuals who are members of the social network but did not upload the information to the service. In such situations, inappropriate sharing preferences can cause conflict and threaten users' privacy. Since related studies suggest that users prefer to solve multi-party privacy conflicts through negotiation, we introduce a novel approach based on negotiation through arguments. In our approach, users propose privacy settings and support their proposals with logical arguments. The final decision is based on a setting supported by sound arguments.

Categories and Subject Descriptors

K.4.1 [**Computers and Society**]: Privacy; I.2.11 [**Distributed Artificial Intelligence**]: Multiagent systems

Keywords

Privacy, Social Network, Argumentation

1. INTRODUCTION

A social network service (SNS) enables users to maintain social relationships via online interactions. As users on an SNS interact, they share information with each other. Often, the information shared on an SNS involves several users (e.g., a photo showing a group of people). Many SNSs enable their users to connect the information they upload to other users; these connections are usually employed to notify the concerned users that the information was uploaded. Since the information shared varies depending on the SNS, (e.g., Instagram is focused only on photos and Twitter on short textual messages), these connections can take different forms, e.g., tags on a photo or mentions in a tweet. For example, Alice uploads a photo from last weekend's party where she and her friend Bob appear together, and tags Bob on the picture. When these connections are created, the other users are also 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 SNS users, they can also pose a privacy threat. For example, Bob thinks that the photo uploaded by Alice is somewhat sensitive and he is not sure about uploading it to an SNS. However, since Bob has no control over uploading that photo, Alice's action can threat Bob's privacy. We identify situations like this as multi-party privacy conflicts.

Currently, SNSs do not have mechanisms to handle multi-party privacy conflicts [7]. Thus, a user who did not upload a piece of information concerning him has to either agree with the sharing preferences chosen by the uploader or remove the connection that links the user to the shared information. A trivial approach to solve such conflicts is to respect the sharing preferences of every party. However, the nature of conflicts in preferences can make this solution inviable. For example, according to their individual privacy preferences, Alice would like to share the photo with Charlie but Bob would like to share it only with common friends and he does not know Charlie. Here, there is no solution that completely respects both parties' preferences.

Decision support systems that help users resolve multi-party privacy conflicts have been identified as one of the biggest gaps in privacy management in social media [2, 11, 20, 15]. The main challenge for these systems is to propose solutions that can be accepted most of the time by all the users involved, minimising the burden on the users to resolve multi-party privacy conflicts.

Based on evidence that users negotiate over what privacy settings they should employ [11, 20], we hypothesize that, during the negotiation of setting a privacy preference, users employ arguments to convince the other parties that their demands are reasonable and should be taken into account. Therefore, we propose a new approach for managing multi-party privacy conflicts based on logical arguments. The forms of these arguments can be classified in a number of schemes, such as precedence or popular opinion. Besides, other variables, such as relationship types among the users, can play a key role during the negotiation and in the final decision.

2. RELATED WORK

One may think that the most direct approach to manage multi-party privacy policies is employing veto voting, as already suggested in [17]. That is, denying access takes precedence over granting access. Thus, if an individual wants to share the information with a given user, but another individual does not, the information is not shared. The obvious benefit of this approach is that it does not allow privacy breaches. However, there is a problem that advises against always employing this solution. Since denying access takes precedence, there may be cases in which veto voting leads to sharing utility loss. For example, Alice and Bob appear together in a photo. Bob initially opposes sharing the photo 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 the veto voting been applied, then the item would have not been shared with Charlie, being a missed opportunity to share.

There are other proposals in the related literature that aim to help users resolve multi-party privacy conflicts [19, 12, 3, 9, 8]. However, some of them [19, 12] need too much human intervention during the conflict resolution process to be practicable, by requiring users to solve the conflicts *manually* [19] or very close to *manually* [12], e.g., participating in difficult-to-comprehend auctions with fake money for each and every possible conflict. Other approaches to resolve multi-party privacy conflicts are more automated [17, 3, 9], but they only consider one fixed way of aggregating user's privacy preferences without considering how users would actually achieve compromise and the concessions they might be willing to make to achieve it depending on the specific situation. Only [8] considers more than one way of aggregating users' privacy preferences, but the user that uploads the 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 all the users involved may lead to solutions that are far from what some users would be willing to accept. All of this causes these more automated mechanisms to have difficulties to adapt to different situations that may motivate different users' concessions, which has the potential to cause these mechanisms to suggest solutions that may not be acceptable by all users, so that users may need to end up resolving multi-party conflicts manually most of the time. The work presented in [14] provides an improvement over these fixed ways of aggregating privacy preferences from users by suggesting 3 different methods that would be selected depending on the particular situation, but again, only a limited number of aggregation methods is considered.

Finally, some very recent works propose game-theoretic negotiation mechanisms to tackle the multi-party privacy conflict resolution problem [10, 16]. These proposals provide an elegant analytic framework to study the problem and the kind of solutions that can be obtained based on well-known solution concepts such as the Nash equilibrium. However, as shown in [10], these proposals may not work well in practice since it seems they may not be able to capture well the social idiosyncrasies users actually consider in the real life when they face multi-party privacy conflicts [11, 20].

3. PROBLEM STATEMENT

In this paper, we propose a novel approach to resolve multi-party privacy conflicts based on argumentation. As in [2, 19], users negotiate over what sharing preferences should be applied. However, in our approach, users do not classify their preferences by their strength. Instead, their preferences have to be supported by logical arguments. As in a real discussion, the sound arguments are the ones considered in the conclusion of the negotiation, while weak or unsound arguments are discarded. This means that the preferences supported by the sound arguments will be applied, or at least, taken into consideration.

We model the multi-party privacy management problem as a multi-agent scenario. Each individual linked to a piece of information to be shared on an SNS employs a personal agent that acts on his or her behalf during the negotiation. These personal agents consider the following variables during the negotiation:

Ownership : An individual can play one of two roles, the owner of the information, or a stakeholder—an individual who is somehow linked to the data but is not the owner. As concluded by Besmer et al. [2], SNS users believe that being the owner of the information implies some authority over the final sharing settings. Thus, during the negotiation, agents must be aware of this position of power of the owner and behave accordingly.

Relationships among the individuals : An individual has specific types of social relationship with other individuals, e.g., close friend or sibling. The closeness and authority relationships can modify how a person negotiates. Closeness can affect on how much an opinion is taken into account. For example, it is very likely that a user's sibling respects the user's preferences. On the other hand, a person with authority over others can impose her opinion even when all the others have different views.

Sensitivity of the information : The appropriateness to share something on an SNS is subjective. For example, in some cultures drinking alcohol is a taboo, thus, a photo showing a person drinking can be inappropriate, where as it can normal in other cultures. Each individual involved in the negotiation has a perception of the sensitivity of the information. This perception will affect how that person tries to impose her view.

The goal of each agent during the negotiation is to apply a sharing preference to the information as close as possible to its principal's preference. The preferences of two agents can be compatible or conflicting with each other. For example, the preference of Bob: "I do not want my parents to see this photo" is compatible with the preference of Alice: "I want only my friends to see it", as long as Alice's friends do not include Bob's parents.

The negotiation consists of a predetermined number of rounds. During each of these rounds the agents propose privacy settings supported by arguments. An agent's arguments and settings can change and adapt to what other agents propose from round to round. After the preset number of rounds of negotiation, a final decision is made. This decision 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 conflicts, the preferences supported by the sound arguments should take precedence over the others. Since creating such sharing preferences can be difficult for an average SNS user, this process should be automated, so that a sharing preference recommender suggests the final sharing

configuration.

Consider an example negotiation between Alice and Bob:

1. Bob: It's a funny photo, but embarrassing since I appear drunk. I don't want strangers seeing it.
2. Alice: We had a lot of fun during the party. Everybody's talking about how funny you were and they want to see your photos. Let's share it with everybody.
3. Bob: Well, if you insist... we should share it like we always share photos like this, only with common friends.
4. Alice: C'mon! it was our graduation party! It's something that we only do once in our lifetime. We should show it to the world.

In this example, Alice is the owner. She and Bob are close friends, thus, they negotiate using informal language. Alice thinks that the photo is not very sensitive, while Bob thinks the contrary. In this example, there are two rounds of negotiation and in each round, Alice and Bob give arguments to support why they should share the photo in a specific way.

4. ARGUMENTS

In our approach, agents employ arguments to convince the other parties that their privacy preferences should be used, or at least considered for the final decision. Different types of arguments can be employed during a negotiation. Each of these arguments corresponds to an argumentation scheme [18]. These schemes are argument forms that represent inferential structures of arguments used in everyday discourse. Although, given the appropriate situation, almost every argumentation scheme can be used during a negotiation, we hypothesize that in our context only four schemes fit: argument (i) from consequences, (ii) from analogy, (iii) for an exceptional case, and (iv) from popular opinion. Argumentation schemes work as classes of arguments and each argument employed by an agent is an instance of a class.

We consider that individuals value their relationship with others, thus, they are well intended. Consequently, individuals do not employ fallacies during the negotiation and the arguments used are based on actual facts. The goal of an agent is to guarantee that its privacy preferences are respected as much as possible. Thus, the agents are likely to employ the arguments that are the most convincing.

Argument from Consequences: If A is brought about, then good (bad) consequences will occur. Therefore, A should (not) be brought about. An example of good consequences in our context is: *We had a lot of fun during the party. Everybody's talking about how funny you were and they want to see your photos. Let's share it with everybody.* An example of argument from bad consequences is: *It's a funny photo, but embarrasing since I appear drunk. I don't want strangers seeing it.*

When an SNS user shares something, she expects to obtain some kind of benefit in terms of friendship, jobs, and other social opportunities [5]. Therefore, it is reasonable to argue that sharing certain information implies a good consequence. On the other hand, sharing inappropriate information can harm people's feelings and cause social tensions. Thus, negative consequences can also form a valid argument.

Argument from Analogy: Generally, case C_1 is similar to case C_2 and A is true (false) in case C_1. Then, A is true (false) in C_2. An example of an argument using this scheme is: *We should share it like we always share photos like this, only with common friends.*

We find a number of approaches for managing privacy on SNS based on tools that automatically suggest privacy preferences [1, 4, 6, 13]. These tools employ past privacy settings employed by the user to infer new configurations. Many of these tools have been developed and evaluated satisfactorily with users. Thus, past decisions can be exploited for suggesting new privacy settings.

Argument for an Exceptional Case: If the case of x is an exception, then the established rule can be waived in the case of x. An example in our domain is: *C'mon! it was our graduation party! It's something that we only do once in our lifetime. We should show it to the world.*

Although previous privacy configurations can act as a guide for future elements, exceptions need a different approach. The scheme for an exceptional case is, at some level, the opposite to the scheme of argument from analogy. The arguments created from this scheme cover cases where privacy recommending tools would fail. Obviously, an individual has to make a hard case to justify why the new element is such an exception that needs a different consideration to the other previous and similar elements.

Argument from Popular Opinion: If the large majority in a particular reference group G accepts A as true (false), then there exists a presumption in favor of (against) A. An example is: the majority of the people that appear in the photo think that it should be kept private. Therefore, we should not share it with anyone.

Argument created from this scheme can take two forms: (i) they can be explicitly employed in an utterance, or (ii) they can emerge from the suggested privacy settings supported by other arguments. The first form can only be used in the following round of negotiation. It is not possible to use it in the first round as individuals still do not know what others' opinions are. The second form for an argument from popular opinion automatically emerges when two or more individuals suggest the same sharing preferences. Thus, although no individual explicitly employed an argument from popular opinion, it is considered for the final outcome.

5. OPEN CHALLENGES

The goal of our research is to build a privacy recommender tool that helps users to decide what sharing configuration they should apply to a piece of information that concerns several individuals. This goal entails a number of challenges that future work should address.

First, since the recommender must provide suggestions that are similar to what humans do, we need to collect data from human participants. However, generating real scenarios where privacy conflicts arise is nontrivial. Hence, we plan to survey SNS users using hypothetical situations that present such conflicts. The variety of situations must be sufficient to collect several instances of every possible combination of variable values and arguments. Figure 1 shows an example of a possible hypothetical scenario presented to the participants.

The data collected will be used to generate a predictive model that, given a set of arguments, privacy settings, relationship types, sensitivity values, and roles as input, provides a privacy setting as output. The predictive model can be based on a machine learning technique (e.g., decision trees) or heuristics. The model will be trained and tested with the data collected from the real SNS users.

Besides the variables proposed, it is worth noting that social pressure might play a role. That is, if the majority of

During a wedding, the groom (A) takes the photo below with three friends of his (B,C, and D). The photo was taken during the dance after the ceremony. Two weeks after the wedding, A uploads the photo to Facebook. They talk about what the privacy policy for the photo should be. These are their arguments.
- A: I know we always share photos of us only with common friends. However, this is different, it is from my wedding! I want everybody to see this photo.
- B: I appear drinking, I don't want anyone but us to see it.
- C: The photo is great, let's share it with everyone.
- D: The photo is great, let's share it with everyone.

What privacy policy do you think should be applied to the photo?

- ○ Public photo (A, C, and D)
- ○ Private, only the four friends (B)
- ○ Other:

Figure 1: Example of hypothetical scenario.

users involved advocate for a privacy policy, this might have an effect on the final decision. Indeed, most of current approaches to this problem apply some kind of voting mechanism: majority voting, veto voting, and uploader overwrites. To evaluate the effect of social pressure, it is necessary to create scenarios where a majority of users want something and there is a user that opposes strongly.

In our approach, we assume that agents do not employ fallacies and that arguments are always valid. This assumption reduces the complexity. However, future work should look into ensuring the validity of arguments and punishing or applying penalties to the agents that use invalid arguments.

REFERENCES

[1] S. Amershi, J. Fogarty, and D. Weld. Regroup: Interactive machine learning for on-demand group creation in social networks. In *Proc. of CHI 30th*. ACM, 2012.

[2] A. Besmer and H. Richter Lipford. Moving beyond untagging: photo privacy in a tagged world. In *Proc. of the CHI 28th*, pages 1563–1572. ACM, 2010.

[3] B. Carminati and E. Ferrari. Collaborative access control in on-line social networks. In *IEEE CollaborateCom*, pages 231–240, 2011.

[4] G. P. Cheek and M. Shehab. Policy-by-example for online social networks. In *SACMAT '12*, pages 23–32, New York, NY, USA, 2012. ACM.

[5] N. Ellison, C. Steinfield, and C. Lampe. The benefits of facebook friends: Social capital and college students' use of online social network sites. *Computer-Mediated Communication*, 12(4), 2007.

[6] L. Fang and K. LeFevre. Privacy wizards for social networking sites. In *Proc. of WWW 19th*, pages 351–360. ACM, 2010.

[7] R. Fogues, J. M. Such, A. Espinosa, and A. Garcia-Fornes. Open challenges in relationship-based privacy mechanisms for social network services. *International Journal of Human-Computer Interaction*, (just-accepted), 2015.

[8] H. Hu, G. Ahn, and J. Jorgensen. Multiparty access control for online social networks: model and mechanisms. *IEEE TKDE*, 2013.

[9] H. Hu, G.-J. Ahn, and J. Jorgensen. Detecting and resolving privacy conflicts for collaborative data sharing in online social networks. In *Proc. ACSAC*, pages 103–112. ACM, 2011.

[10] H. Hu, G.-J. Ahn, Z. Zhao, and D. Yang. Game theoretic analysis of multiparty access control in online social networks. In *Proceedings of SACMAT '14*, pages 93–102, New York, NY, USA, 2014. ACM.

[11] A. Lampinen, V. Lehtinen, A. Lehmuskallio, and S. Tamminen. We're in it together: interpersonal management of disclosure in social network services. In *Proc. CHI*, pages 3217–3226. ACM, 2011.

[12] A. Squicciarini, M. Shehab, and F. Paci. Collective privacy management in social networks. In *Proceedings of the WWW 18th*, pages 521–530. ACM, 2009.

[13] A. Squicciarini, S. Sundareswaran, D. Lin, and J. Wede. A3p: adaptive policy prediction for shared images over popular content sharing sites. In *Proc. of Hypertext 22nd*, pages 261–270. ACM, 2011.

[14] J. M. Such and N. Criado. Adaptive conflict resolution mechanism for multi-party privacy management in social media. In *Proceedings of the 13th Workshop on Privacy in the Electronic Society*, pages 69–72. ACM, 2014.

[15] J. M. Such, A. Espinosa, and A. García-Fornes. A survey of privacy in multi-agent systems. *The Knowledge Engineering Review*, 29(03):314–344, 2014.

[16] J. M. Such and M. Rovatsos. Privacy policy negotiation in social media. *arXiv preprint arXiv:1412.5278*, 2014.

[17] K. Thomas, C. Grier, and D. Nicol. unfriendly: Multi-party privacy risks in social networks. In *Privacy Enhancing Technologies*, volume 6205 of *Lecture Notes in Computer Science*, pages 236–252. Springer, 2010.

[18] D. Walton, C. Reed, and F. Macagno. *Argumentation Schemes*. Cambridge University Press, 2008. Cambridge Books Online.

[19] R. Wishart, D. Corapi, S. Marinovic, and M. Sloman. Collaborative privacy policy authoring in a social networking context. In *Proc. of POLICY '10*, 2010.

[20] P. Wisniewski, H. Lipford, and D. Wilson. Fighting for my space: Coping mechanisms for sns boundary regulation. In *Proc. CHI*, pages 609–618. ACM, 2012.