Open Challenges in Relationship-Based Privacy Mechanisms for Social Network Services

Ricard Fogues[*1,2], Jose M. Such[†3], Agustin Espinosa[‡1], and Ana Garcia-Fornes[§1]

[1]Departamento de Sistemas Informáticos y Computación, Universitat Politècnica de València, Spain
[2]Computer Science Department, North Carolina State University, NC, USA
[3]School of Computing and Communications, Lancaster University, UK

Abstract

Social networking services (SNSs) such as Facebook or Twitter have experienced an explosive growth during the few past years. Millions of users have created their profiles on these services because they experience great benefits in terms of friendship. SNSs can help people to maintain their friendships, organize their social lives, start new friendships, or meet others that share their hobbies and interests. However, all these benefits can be eclipsed by the privacy hazards that affect people in SNSs. People expose intimate information of their lives on SNSs, and this information affects the way others think about them. It is crucial that users be able to control how their information is distributed through the SNSs and decide who can access it. This paper presents a list of privacy threats that can affect SNS users, and what requirements privacy mechanisms should fulfill to prevent this threats. Then, we review current approaches and analyze to what extent they cover the requirements.

Keywords. Social networks, privacy, access control, tie strength, interpersonal relationships, self-presentation.

1 Introduction

The advent of the Web 2.0 has supposed a revolution in how users interact with Web technologies. Social network services (SNS) are some of the most successful applications of this revolution [16]. Facebook with more than 900 million active users[1], Twitter with more than 500 million registered members[2], and Qzone with more than 51 millions of users are some of the biggest SNSs. The impact of these services on society, especially on young people, is unquestionable.

---

[*]rilopez@dsic.upv.es, rlopezf@ncsu.edu

[†]j.such@lancaster.ac.uk

[‡]aespinos@dsic.upv.es

[§]agarcia@dsic.upv.es

[1]Facebook statistics `http://newsroom.fb.com/`

[2]Twitter To Surpass 500 Million Registered Users On Wednesday. `http://www.mediabistro.com/alltwitter/500-million-registered-users`

Privacy problems associated with digital communication and network technologies have been a major concern among Internet users over the past decade [74]. The emergence of social networks has even increased these concerns. People register to these SNSs and share images, videos, and thoughts because they perceive a great payoff in terms of friendship, jobs, and other opportunities [20]. The popularity of SNSs attracts not only faithful users but third parties with adverse interest [2]. If we consider the huge amount of private information uploaded to those SNSs and the persistence of it in the social networks, the privacy of SNS users can be threatened [31]. Recent cases show that on-line thieves, stalkers, and bullies take advantage of the information available on SNSs and use it for purposes that were not the initially intended ones [34].

There are several definitions of privacy in the related literature. In the context of this survey, we use the definition of Alan Westin, who defined privacy as "the claim of individuals, groups, or institutions to determine for themselves when, how, and to what extent information about them is communicated" [69]. This definition implies that SNSs have to offer their users mechanisms that allow them to decide how their information is disclosed. Current SNSs have taken steps towards this objective, but there still exist several problems that make users feel they have lost control of their information and how it is shared among the SNS [65]. Users demand better privacy mechanisms, with richer and finer-grained privacy policies that take into account the way SNS users share information and interact with others. Moreover, privacy controls for these new access controls have to be easy to use, offering automatic suggestions and learning from the behavior of the users.

This article reviews studies that enhance privacy in social networks, as well as studies that explore human relationships over social networks and their behavior. Understanding how humans share and manage their friendships on SNSs is crucial so that researchers can adapt their models and methods to cope with the users' needs and expectations. Studies are classified according to the type of privacy risk they address. Each study is impartially presented and reviewed, showing the strengths and weaknesses of the proposals made in the study. The final objective of this paper is to survey research in the field of privacy in social networks as well as to promote and encourage future research and advances that overcome the current challenges that exist in this field.

1.1 Privacy and Social Networks

As pointed out more than a century ago by Warren and Brandeis [68], disclosure of private information and the misuse of it can damage people's feelings and cause considerable damage in people's lives. In SNS where intimate information of the users is managed, privacy is of paramount importance. A research of Gross and Acquisti [31] in the early days of Facebook showed that the majority of users were unconcerned about privacy risks. They tended to use default privacy configurations and personal data were generously provided. More recent studies, like the one from Boyd and Hargittai [8], show that the privacy awareness of SNS users has increased lately. The widespread media attention on SNS and on situations where the leakage of personal information of SNS users affected their lives has positively influenced the way SNS users manage their privacy [55]. Nevertheless, the high number of privacy risks that affect SNS users leaves room for improvement in this field of study.

The most important SNS users' privacy concerns are: identity theft [5], unauthorized access [55], misuse of personal information and stalking [31, 9, 59], and profiling [33]. In this study we focus on misuse of personal information. This threat refers to the possibility of a malicious dissemination of previously collected information. For instance, users may face blackmailing situations when embarrassing data is collected from a SNS by a third party. In the context of SNSs, misuse of personal

information usually occurs when users disclose inappropriate information due to a negligence during the configuration of their privacy settings or ignorance about how privacy is managed on the SNS. The rest of privacy threats affect different levels of privacy on SNSs and fall out of the scope of this study. Identity theft and unauthorized access are related to access control enforcement. For example, unauthorized access can occur if the authentication mechanisms of the SNS are not good enough or if the communication between the user and the SNS is not properly encrypted. Profiling is a threat when the party which owns the information on the SNS is not trustworthy. A typical case of profiling occurs when the party that manages the SNS sells the information available on the SNS to third parties that use it for marketing purposes.

It has been acknowledged that in order to properly minimize misuse of personal information, a new privacy mechanism is needed [26, 77]. In the next section, we detail the requirements for such a new privacy mechanism.

1.2 Requirements for a Social Network Service Privacy Mechanism

Any privacy mechanism has at its base an access control. Access controls dictate how permissions are given, what elements can be private, how access rules are defined, and so on. Access control models of current SNSs tend to be very simplistic. Nonetheless, recent improvements in facebook-like SNSs have enhanced the access control models. For example, now it is possible to define policies to deny access to groups of users, instead of individuals. Some SNSs allow the possibility to express a social distance of contacts that have access to the resource, for example, friends of friends (two hops), friends of friends of friends (three hops), and so on. Another addition that only a few SNSs have added is the possibility to choose the amount of information from a friend that we want to receive (Facebook). However, these models still lack key elements. One of the most important is the lack of diversity in the type of relationships. Most SNSs only employ "friend" as the only type of possible relationship. This lack of classification of contacts leads to privacy leaks to other members inside the social network. This is referred to by Johnson [36] as the *insider threat*. Gates [26] identifies the following requirements that an access control model for a SNS must fulfill:

- **Relationship-based**: People base their decision of sharing information on their relationship with others. Moreover, the properties of the relationship also affect the way people disclose their personal information. In social psychology, it is generally accepted that one discloses more of his/her personal information to someone in a strong relationship [19]. Hence, a control access model that tries to reflect the way people disclose and share in real life should be based on relationships.

- **Fine-grained**: The access control has to allow users to define access policies for single items. If the access control is available in a fine-grained format the privacy policies can be more flexible and they can express the user's preferences exactly. For example, a user should be able to define privacy policies for specific photos, individual blog entries, or even some words or phrases of a comment. In other words, users should be able to decide exactly to what extent others can access their information.

- **Interoperability**: Many SNSs have a specific objective; while Facebook aims to facilitate users contacting their friends, LinkedIn helps users to maintain their professional networks. Facebook and Linkedin have clearly different purposes. Because of this variety of purposes, users may have several multiple accounts in different SNSs, each one for a different social

objective. In this scenario, it is highly desirable for access controls to be interoperable and follow the users, so it is not necessary to define an entire new access control for each SNS.

- **Sticky policies**: Besides being interoperable, privacy policies should also follow the data to which they apply. For example, many SNS allow third party applications to access users' data. The privacy preferences assigned to that data should be respected by these third parties and in whatever context it might travel to. This idea was introduced by Karjoth et al. [41].

In the related literature, we have also identified additional requirements that play a crucial role in developing successful access control models for SNSs:

- **Content Type Management**: SNS enable users to share a variety of different pieces of information: photos, videos, comments, events, hobbies, and so on. Besides the miscellany in the format of the information, its content also matters when deciding who has access and who has not [32]. Flickr[3] employs tags so users can classify their pictures according to their type. A similar approach could be used so users could define permissions based on the type of the content.

- **Co-privacy**: SNS users like to upload items to their profiles, such as photographs and videos, where other users are depicted. Specially, SNSs that focus on helping users to maintain their friend relationships encourage users to upload publications of this kind. Items of this type can raise several privacy concerns. While the owner of the item is in charge of assigning a privacy policy to it, the other users related to the item can be affected if the privacy policy is not appropriate for their interests. It is possible to infer a great amount of information about an individual from information leaks that occur due to shared items and privacy preference conflicts [67]. Current access models do not consider these situations; thus, users are forced to use strategies like untagging, asking the owner to remove the photo or, in the most extreme situations, removing friendship links. Access controls should consider co-privacy management and offer mechanisms that allow every user involved in a single item to express their privacy preference so that the resulting privacy policy applied to that item maximizes the utility for everyone.

An access control acts as the base for a privacy mechanism, however, they require other elements to be functional. It is not realistic to assume that SNS users can understand access control models and use them intuitively. Powerful privacy models are useless if they lack usability and are not understood by the people that will use them [17]. Users need tools that guide them through the process of setting their privacy preferences. Users also require mechanisms that help them to understand their current privacy preferences and how their information is disseminated among other SNS users. According to related literature, a privacy mechanism for SNSs should fulfill these requirements:

- **Automatic relationship inferring**: If the access control has to be based on social relationships, these have to be accurately defined. SNS users tend to have a high number of friends. For example, according to the Facebook statistics, the average number of friends in that social network is 130. Hence, classifying every contact in a social network can represent a burden on the user. Privacy mechanisms should have the capacity to automatically infer the type of a relationship and make the whole process of friend classification easy and fast.

---

[3]`www.flickr.com`

4

- **Privacy setting recommendation**: While privacy is paramount on SNSs, users are focused on enjoying the functionality that these offer. For many users privacy settings represent a burden, for others privacy settings are difficult to manage and understand. Recommender tools can help users to set properly their privacy settings. While recommenders can help reduce the user's burden, they are rarely perfectly accurate. Thus, it is important for the user to be able to view, understand, and modify the recommended policy before it is applied

- **Privacy understandability**: Access controls can be complex and daunting for SNS users. Average SNS users do not have expertise on security, thus, it is difficult for them to accurately evaluate how their information is disclosed through the SNS [50]. Users require proper interfaces that show them how their privacy policies dictate their self-disclosure.

- **Self-presentation management**: In the beginning, social media was focused on establishing or maintaining friendship relationships through a digital channel. However, social media have increased the number of services offered and now users expect more benefits than friendship alone [20]. Some examples of social media use are to obtain fame[4], or for commercial brands to publicize their products, acquire recognition, and maintain contact with their customers[5]. In a nutshell, and as pointed out by Kairam et al. [40], social media users utilize the product to successfully tailor self-presentations for various parts of their network through selective information sharing. Self presentation is achieved by carefully tailoring self-disclosure. Privacy controls should help users to maintain their chosen self presentation on SNSs.

In the following sections of this paper, we will review studies that aim to totally or partially cover the requirements previously listed. First, Section 2 starts reviewing formal relationship-based access control models. Section 3 presents some prototypes of access controls that take into account the content type of the information being shared. Section 4 is centered on papers that deal with co-privacy. Section 5 presents papers that aim to accurately model and infer human relationships on SNSs. Section 6 presents works that propose privacy policy recommenders. Section 7 reviews papers that present methods to enhance the understandability of privacy settings. The remaining requirements (interoperability, sticky policies, and self-presentation management) are treated as open challenges and are covered further in section 8.

2   Relationship-based Access Control Models

This section reviews Relationship-based Access Control (ReBAC) models proposed for SNSs. The ReBAC paradigm provides users with better mechanisms for disclosure control, and they represent more naturally how humans decide what to share and with whom. All models studied in this section fulfill the two first requirements of access controls for SNSs. Specifically, all of them are based on relationships and their privacy policies are fine grained. However, the models differ in other features, table 1 shows a comparison of them. In the table, the features of the reviewed models have been divided into three subgroups.

The first group contains the properties that the models consider for the relationships. *Multiple relationship type* refers to the possibility that the model manages different kinds of relationship;

---

[4]http://www.wltx.com/news/tech/article/233056/378/Survey-Social-Media-Draws-Young-Fame-Seekers

[5]http://www.usatoday.com/story/money/personalfinance/2013/04/16/small-business--social-media-facebook/2075123/

	[25]	[24], [10]	[13], [12]	[11]	[15]
Relationship features					
Multiple relationship types		★	★	★	★
Tie-strength			★	★	
Directional relationship		★			★
User-to-user relationship	★	★	★	★	★
User-to-resource relationship				★	
Policy language features					
Policy individualization	★	★	★	★	★
Regular expression language					★
Based on ontology				★	
Arbitrary social distance	★	★	★	★	★
Social path specification		★	★	★	★

Table 1: Comparison of ReBAC models.

for example, a model could allow the differentiation between a friend and a family relationship. Another concept used to differentiate relationships, which is further explained in section 5, is tie strength. Some models allow the possibility of specifying a numerical value for the strength of the relationship. *Directional relationship* alludes to the possibility of defining asymmetric relationship; for example, user A can be related to user B but that does not imply that the opposite relationship exists. Finally, the last two features, *user-to-user relationship* and *user-to-resource relationship*, describe whether the models consider relationships among users and among users and items. For example, a user can somehow be related to a movie, which is not a user but a resource.

The second group of features includes those that affect the policy specification language of each model. As explained in the previous section, ReBAC models have to be fine-grained; the *policy individualization* refers to this characteristic. The following two features, *regular expression language* and *based on ontology*, specify whether the policy language uses any of these techniques. When specifying a privacy policy, a common resource is to define a maximum social distance, which is the number of hops between the owner of the resource and the accessor. *Arbitrary social distance* identifies what models allow this possibility. Besides the social distance, *social path specification* defines the type of the hops of that distance; for example, the family members of my friends represent a social distance of two, where the first hop is friends and the second family members.

Finally, the last subgroup specifies what models have implemented an *enforcement mechanism*. This mechanism has to evaluate the access rules attached to an item being accessed to determine whether the accessor satisfies those rules. In a centralized SNS like Facebook this mechanism can be integrated in the SNS; however in distributed social networks this mechanism can be more complex since it can require the use of third party trusted services.

Fong et al. [25] proposed a formal algebraic model for facebook-style social networks. The authors created this algebraic model to reflect how facebook-style SNSs model their access control. Even though this model cannot be classified as a true ReBAC, we considered this paper in this review since it formally shows the limitations of current SNS access control models. The model divides the authorization of access to a resource into two stages. Stage 1 is to reach the profile

6

of the resource's owner and Stage 2 is to access the resource. The model allows the definition of authorization policies for each stage independently. Since policy language considers the network topology properties, the model allows complex policies that are beyond what facebook-like SNSs offer (friends, friends of friends, no one, and public). Topology-based policies include: degree of separation, k common friends, k clique, trusted referral, and stranger. Since this work aims to model the access control of facebook-like SNSs, it also has the same limitations. Moreover, since users' profiles and resources are not treated the same, the authorization process is divided into two steps. This division can lead to unnecessary complexity of access policies.

Fong [24] proposes an access model for social networks based on relationships. In contrast to [25], this work uses social contexts and allows a generalization of relationships types (e.g. parent-child, employer-employee, etc). Social context is another dimension of relationships; some relationships have a different meaning depending on the context or they only exist in a given context. For example, a physician who is my treating physician in one medical case may very well be a consulting expert in a different medical case of mine. As a result, the physician may enjoy a different level of access in each case. This access model considers that, for each relationship, the social network defines its inverse. For example, if a social network has the relationships *parent* and *employer*, it must also contain the relationships *child* and *employee*, which are the inverse of *parent* and *employer* respectively. When a resource is being accessed, the evaluation of the authorization of access for that accessor is done based on an active context. This concept captures how people are willing to disclose different information depending on the context. The policy language defined by Fong allows the specification of unlimited sequences of relationships. For example, it is possible to express a policy that allows access to the father of a friend of a friend. This feature is an improvement with respect to [25] where the chain of relationship was restricted to friend and friend of a friend.

Bruns et al. [10] improved the previous work of Fong [24] adding *Hybrid Logic* to the model. The policies are divided into two sub-policies; one sub-policy is defined from the point of view of the owner of the resource and the other one from the point of view of the accessor. The improved model allows more flexible policies; for example, it is possible to grant access to the last four friends, or grant access if at least n friends of the owner fulfill a certain requirement.

These three works [25, 24, 10] share the same limitation. They do not consider the strength or intensity of the relationships (i.e., they only consider relationships as a boolean: either a relationship exist or not).

Carminati et al. [13, 12] propose a model that allows the specification of access rules that consider the type of the relationship, its depth, and its intensity. The proposal of Carminati et al. considers a distributed SNS; therefore, principals are in charge of specifying their access rules. The work also proposes a semidecentralized access control enforcement. When a principal wants to access a resource, she has to prove that she fulfills the requirements specified by the owner of the resource. A central and trusted server is responsible for storing all the relationships of the social network users. Thus, whenever the requester needs to prove to the resource owner the existence and the attributes of a given relationship, she requests this trusted server for this information. The policy language proposed by Carminati et al. allows the definition of policies that specify a type of relationship, a maximum depth, and a minimum strength. Policies can have several requirements, all of them have to be satisfied in order to obtain access. Several policies can be defined for a single resource. In this situation, only one of these policies have to be fulfilled in order to obtain access to the resource. One limitation of the policy language is that policies cannot refer to a chain of relationships with different types of relationships. For example, it is not possible to specify a policy

that grants access to the parents of the owner's friends.

Carminati et al. [11] propose an access control model based on semantic web technologies. This paper is an extension of the previous paper [13]; the main differences between the two proposals is that this model considers the user-to-resource relationship and it uses semantic technologies for the policy language. The model proposed by Carminati et al. considers the following five important elements of an online social network: (i) profiles, (ii) types of relationships among users (e.g. Bob and Alice are colleagues), (iii) resources, (iv) relationships between users and resources (e.g. Bob appears in a photo owned by Alice), and (v) actions. The use of semantic web technologies allows the model to infer about the relationships among users and resources. For example, it is possible to infer that a close friend is also a friend and anything that is accessible by friend could also be accessible by a close friend. The authors focus their article on the addition of semantic technologies to their previous access control model [13]. However, the authors did not evaluate how the addition of semantic technologies improved their previous work.

Cheng et al. [15] developed a ReBAC model using regular expression notation. Their model defines resources and users as the target of an action. The model permits a high generality of relationship paths in its policy specification, since the notation of the model is regular expression. The paths can be defined as patterns; for example, it is possible to define a policy that grants access to users that are connected to the resource owner by a path that contains at least one friend and a maximum of two coworkers. Since the users can be considered as targets of an action, it is possible to specify polices that hide the profile of the users and do not show them in searches performed by others that do not satisfy the specifications of the policy. Even when the model considers different types of relationships, it does not contemplate a value for the trust or strength of the relationships. This limitation restricts the power of expressiveness of the model because it is not feasible to define a type of relationship for each possible level of tie strength.

The models reviewed in this section are based on human relationships and depend on them to express privacy policies and define how the different elements can be accessed and by whom. These models assume that the social network provides a rich social model that is capable of representing different types of human relationships. Unfortunately, this is not true in current SNSs, as they usually only consider friend as the only type of possible relationship. Section 5 reviews studies that model human relationships and propose theoretical models and actual software tools to accurately predict and represent the type of a relationship and its intensity.

3   Content Type Management

None of the access control models reviewed above consider the type of the shared information. In other words, users cannot specify privacy policies for different types of content. For example, a user could not define a privacy policy that affects their family photos. A number of studies [32, 42, 75] show that content matters during privacy policies definition. Moreover, mechanisms to classify content, such as tags, improve the usability of privacy mechanisms.

Yeung et al. [75] prototyped the management of privacy for photos that considers content type. To classify the photos, the authors proposed the use of tags, which is the act of assigning descriptive keywords to resources. This is the same method that Flickr, the popular social network for photo sharing, employs so that users can classify their pictures according to their type. Yeung's system is based on OpenID as authentication protocol and the AIR policy language [38] which is based on RDF. The authors only proposed a prototype of the system and did not provide any evaluation.

Hart et al. [32] proposed a mechanism to manage privacy for blogs based on tags. The authors proposed a privacy language called Plog. This language is based on groups and post type. Users can define groups of users manually or group potential viewers by attributes that they all share (e.g., workplace or same school). The main focus of their study is to compare basic privacy policy mechanisms for blogs with a tag-based approach. To this aim, they developed a WordPress plugin and recruited twenty eight participants to evaluate their proposal. The authors did not use real data from the participants, instead, they created artificial data for imaginary users and asked the participants to manage that data as if it was theirs. Thus, they did not examine users' actual preferences. Their results showed that an approach that uses tags is more usable than one that does not.

Klemperer et al. [42] evaluated the usability of an access control based exclusively on category tags. The authors aimed at evaluating whether tags can be used to organize photos and define their privacy at the same time. Besides, they also studied if tags can decrease the number of privacy conflict. This is, privacy policies that are contradictory. Usually, this happens because users have problems building mental models of their privacy. For their study the authors developed an application that allowed participants to tag their pictures and assign privacy policies for their contacts. The authors asked the participants to use personal photos that were not necessarily uploaded to any SNS. Their results showed that the use of tags reduces the number of required privacy policies and also the number of privacy conflicts.

4   Co-privacy

As explained before, one of the requirements for ReBAC models is the management of co-privacy or, in other words, the management of the privacy settings of items shared among users of the social network that affect the intimacy of several individuals. An example of a common issue with items of this kind is a photograph with many users tagged in it. The user who took the photo uploads it to her SNS profile and tags every other user that appears in the photo. At this moment, SNSs leave the responsibility of setting a proper privacy setting for the shared item on the hands of the owner. This decision may suppose a threat to the privacy of the other involved users. The proposals reviewed in this section are focused on finding a proper privacy setting for items that involve several users.

Figure 1 depicts a conceptual map for all of the reviewed approaches that propose a co-privacy management mechanism for SNSs. As shown in the conceptual map, there are three different approaches for co-privacy management. *Condition preference sensitivity* refers to the possibility of the users involved to express how willing they are to allow violations of their preferences. The *Suggestions* approach is based on the idea that the owner of the item is the sole person responsible for the privacy management of that item; thus, other users are only allowed to suggest privacy configurations. Finally, the approaches that guarantee the *preferences of everybody* allow every user involved in the item to express their privacy preferences. Then a privacy policy is generated from the combination of every privacy preference.

Squicciarini et al. [63] propose collective privacy management based on the Clarke-Tax algorithm and incentives for users. Incentives are credits that are given to users to encourage them to assign co-owners of updated items. When the owner and co-owners specify their privacy preferences for a shared item, each one specifies their privacy preferences assigning a value of credits they are willing to spend in order to apply that policy. The Clarke-Tax algorithm ensures that the privacy policy
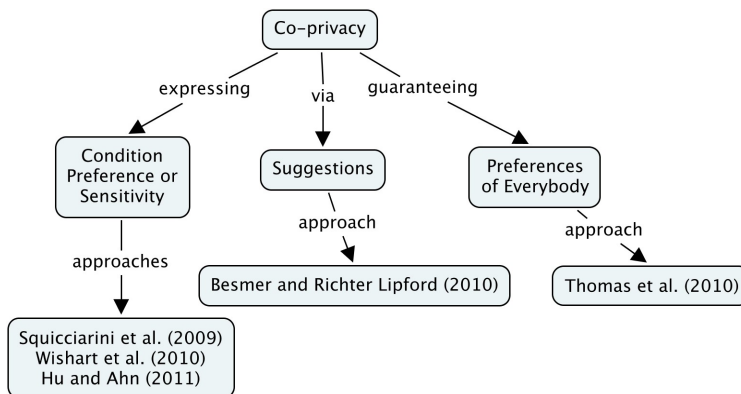
9

Figure 1: Co-privacy Conceptual Map

that maximizes the utility is chosen. The authors created a proof-of-concept application and tested its performance in terms of computing time. The authors propose a novel method that encourages users to participate in the collective managing of privacy by giving them rewards and finds the privacy configuration that has the highest utility according to the co-owners preferences. However, in the study, the privacy policies are very simple and owners only can specify if they prefer the item to remain private (only co-owners have access to it), if it is accessible to co-owners' friends, or if it is designated as public. These simplified privacy preferences avoid privacy conflicts, but they do not represent real privacy policies well.

Thomas et al. [67] study the risks of multi-party privacy in social networks. As a part of their research and as a way to show the dangers of unsuitable co-privacy management, the authors try to infer the information of Facebook users from two sources: links between friends and conversations with friends. The authors achieved an accuracy above 60% inferring information such as gender, political views or favorite TV shows. The authors propose a privacy framework to avoid these privacy conflicts. The framework is based on exposure policies instead of privacy policies. Each user referred to by a piece of information posted on any page of the SNS (e.g. Facebook wall) can define an exposure policy. For example, Alice posts on her Facebook wall a comment where Bob is referred to. Alice specifies a privacy policy because she is the owner of the information and Bob can define an exposure policy to limit the users that can access that comment where he appears. The exposure policies are defined in terms of the type of information and the page where it is posted. The authors prototyped their solution as a Facebook application that guarantees that every exposure policy is respected. The authors did not test their prototype, so experimental evaluation is lacking. Moreover, one of the main concerns of this proposal (also expressed by Thomas et al.) is that if several users are involved in the privacy management of an item, the group of users permitted to access that item tends towards the empty set.

Wishart et al. 2010 [71] propose a collaborative creation of privacy policies for shared items. The authors detect two roles, the owner of the item and the co-owners, which are designated by the owner and are individuals that are affected by the item. The main idea of the proposal is that owner and co-owners refine the privacy policy iteratively, and, hopefully, at the end, the preferences of everybody will be considered in the resulting policy. The authors define a model for privacy policies that contemplates the specification of strong and weak conditions. Weak conditions are overridden by strong conditions. In other words, weak and strong conditions establish a preference

order, where strong conditions are those that a user considers an accessor must fulfill and weak conditions are those that can be overridden by strong conditions specified by another co-owner. As a proof of concept, the authors developed a tool. The tool has not been tested or evaluated with real users. The concept of strong and weak condition introduces preferences when defining privacy policies. However, the proposal does not consider other problems that come from co-authoring. As the owner and co-owners refine the privacy policy, many condition conflicts can arise, for example, two strong conditions that contradict each other. Moreover, the process of refining can be virtually infinite; it is possible that at the end only the preferences of the most persistent user are considered, since there is no method to prevent an endless modification of the privacy policy.

Besmer and Lipford [4] propose a method where the owner of a photograph that involves several individuals is in charge of managing its privacy and the other involved users can only suggest privacy preferences. The authors performed an experiment to collect information about privacy concerns of SNS users about photographs and what strategies they use to control the leak of private information. According to their study, SNS users consider that the owner of a photo (the individual who uploaded the image to the SNS) is the one in charge of assigning a privacy policy. Therefore, the other users who appear in the photo may only suggest privacy policies relying on the responsibility of the photo owner. The authors developed a Facebook application that allows a user to send privacy suggestions to the owner of a photo where that user appears. The authors also performed an experimental evaluation of their approach and their application. According to their findings, the participants were comfortable using this approach. The main issue with this approach is that the owner of the photo may not responsible enough or cannot deal with the petitions of other users and the possible preference conflicts that can arise from these petitions. As a future improvement of the software proposed by the authors, it should help owners to decide which policy to apply in the case of interest conflicts.

Hu and Ahn [35] propose a multiparty authorization framework that enables collaborative management of shared data. The authors divide the users into three groups: *owner*, the user that uploaded the item to the SNS; *accessors*, the users that want to access the item; and *stakeholders*, the users that are affected by the item somehow, for example, being tagged in a photograph. The users in any of these groups have to specify their privacy preferences and assign a sensitivity score to the item being shared. As several preference conflicts can arise, the owner of the photo is in charge of specifying a conflict resolution strategy. The authors propose several conflict resolution strategies (for example: owner-overrides, full-consensus-permit, strong-majority-permit, and many more). Since the owner decides the conflict resolution strategy, she has greater control over the resulting privacy policy associated to the item. For example, if the owner considers that the item is very sensitive regarding her privacy, then she will assign a restrictive conflict resolution strategy like *strong-majority-permit*. The authors developed a prototype as a proof of concept and tested the policy evaluation performance of the prototype. They conclude that the prototype is fast evaluating policies. However, they did not test their proposal with users and did not evaluate the satisfaction of the users with their approach.

5 Modeling Human Relationships on Social Network Services

Current SNS make little effort to differentiate between users. Users are either friends or strangers, with nothing in between. This approximation does not represent human relationships well. As introduced in the paper by Granovetter [29], the concept of *tie strength* defines the relationship
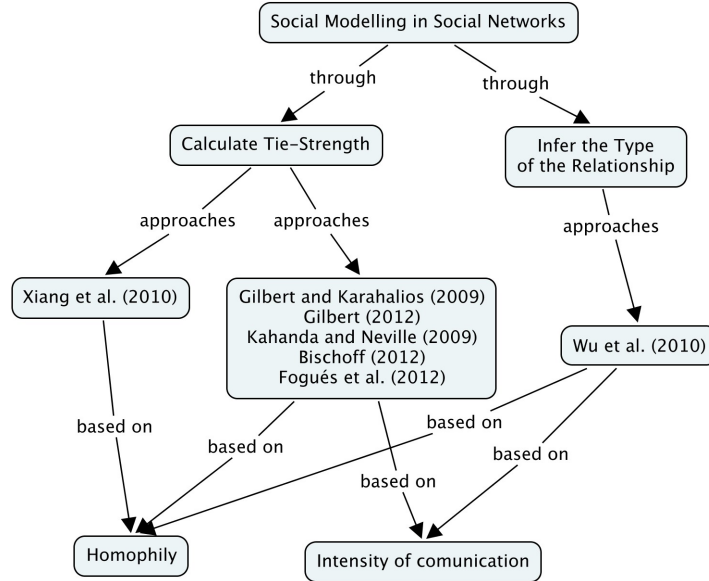
Figure 2: Concept Map of Human Relationship Models

between two individuals. In his work, Granovetter speaks about two different types of ties: *strong* and *weak*. On the one hand, strong ties usually include relations such as family and close friends. On the other hand, weak ties may refer, for example, to coworkers or less trusted friends.

As described in the previous sections, a rich relationship model can play a key role in privacy protection. SNSs aim at creating virtual versions of the real social networks of their users. An accurate representation of the real social graph of the users can help users to manage their privacy. Moreover, ReBAC models need relationships to be modeled truthfully. Wiese et al. [70] studied the correlation between information sharing willingness and tie strength. Their research proved that the strength of ties is even more significant than grouping (the current approach of most SNSs) for predicting sharing. They suggest that a mixture of grouping and tie strength could allow richer sharing policies. Several approaches to create a social model that is based on the concept of the tie strength have been proposed. Most of the works on this matter try to infer a value for the strength of the relationships.

Figure 2 depicts a conceptual map that sums up all of the reviewed approaches that propose models for human relationships on social networks. As seen in the figure, the models are based on two different concepts: homophily and intensity of communication. On the one hand, homophily states thats that the more two person have in common (job, friends, hobbies) the more likely it is that these two persons have a strong relationship. On the other hand, the idea of communication intensity is that if two persons interact frequently, then their relationship should be strong. The majority of models use both concepts and combine them.

Gilbert and Karahalios [28] proposed a model that predicts tie strength among users of Facebook. The authors selected a group of variables available at Facebook. These variables covered different tie-strength dimensions; for instance, the number of links shared correspond to the "services given" dimension. They evaluated their model with the participation of Facebook real users. The researchers asked the participants to assign a value of tie strength to their contacts. Besides, the researchers collected the values for the selected variables. Using the group of over 70 variables

Gilbert and Karahalios achieved an accuracy of 84%. One of the limitations of this model is the huge amount of information it requires to predict tie strength. This high volume of information has to be processed and can become a bottleneck in the performance of the model.

Gilbert [27] expands his previous work [28] and proposes a model to infer the tie strength of relationships in a different SNS, Twitter. Gilbert selects a set of variables that are similar to the variables chosen in [28]. Some variables that were selected in Facebook do not exist in Twitter. In this situation, the author chose analog variables; for example, the number of friends were replaced with the number of followers. In order to evaluate the new model, Gilbert developed a tool called *We Meddle* that predicted tie strength. Users were asked to try *We Meddle* and evaluate its predictions. This work showed that the model proposed in [28] can somehow be generalized and adapted to different SNSs.

In their work Kahanda and Neville [39] propose using transactional information to predict tie strength in social networks. The model is constructed using 50 variables. The variables are classified in 4 groups: (i) attribute-based Features, (ii) topological features, (iii) transactional features, and (iv) network-transactional features. The first three groups of variables are considered in other works [28, 73]; however, the fourth group, network-transactional features, represents a novel approach for tie strength prediction. The variables in this group capture the transaction information between nodes, but they represent it within the context of the larger network structure. For example, one of the variables in this group is the interaction among all the nodes in the network, not only between two pairs. According to the results of the Kahanda and Neville study, the most predictive variables are those in the fourth group. The study lacks an evaluation with humans; therefore, the accuracy of the results may have been affected as the tie strength is a purely human-dependent concept.

Xiang et al. [73] proposed a model to infer relationship strength based on profile similarity, with the goal of automatically distinguishing strong relationships from weak ones. The model proposed by Xiang et al. uses the concept of homophily to infer the tie strength between two individuals. Xiang et al. test their model with proprietary data of the SNS LinkedIn and data from students of Purdue University on Facebook. With the LinkedIn data they tested their tie strength prediction against other heuristics. The authors considered the number of times a user checked another user's profile page as the indicator of the tie strength between them. For the evaluation in Facebook, they calculated the ground truth tie strength between two users as a combination of the number of common networks, common groups, and common friends for those two users. The results of the evaluation show accurate results. However, the evaluation was synthetic; the ground truth tie strength values were calculated using heuristics and were not specified by the users themselves. Therefore, it lacked the corroboration of the results from the participants. The main difficulty of this model to work accurately is that it relies on profile information, which tends to be incomplete or of low quality. Few users specify their address, job, college, or other variables needed by the model to work [21].

Rana et al. [56] propose a theoretical framework for calculating social strength and a ranking of contacts sorted by their social strength. The algorithm proposed copes with the complexity of users using a wide variety of communication services. For example, a user can contact her friends using Twitter, Facebook, and SMS. Moreover, the algorithm also considers different ways of communication inside the same communication service. For example, in Facebook, a user can contact another sending a private message or through a comment anout a photo. These two ways of establishing contact are considered to be different *communication tools*. The algorithm counts the number of interactions on each communication service and assigns a level of importance to

the communications established on that service according to its ratio of usage. Finally, a value of tie strength for each contact is obtained adding the number of interactions with that contact and weighting the interactions according to the service and tool used. The main limitation of this algorithm is that it only considers the frequency of interaction between the user and each contact. However, as shown by other studies [43], a high number of interactions does not necessarily imply a strong tie.

Bischoff [6] analyzes online friendship in the Lastfm musical social network. Lastfm offers social features like friendship links, message exchange, and a personal profile. Moreover, Lastfm allows users to specify their musical tastes, favorite artists, and music event attendance. The author collected public information about 48,527 Lastfm users. The information collected contained variables such as: friend links, messages sent, tags assigned to artists, music recently heard, preferred albums, preferred artists, events attended, and demographic data (gender and country). The author used the gathered data to classify each pair of users as: no link between them, weak relationship, or strong relationship. The number of events coattended by both persons determines the strength of the relationship in the training set. For example, if two persons attended 2 events together, then they have a weak link; however, if they attended 11 or more, then they have a strong link. The author also used the data to create a friendship recommender system. In both experiments the results were promising. The most predictive variables in Bischoff's study were those related to homophily, such as coincidences in the preferred artists or same country. This study shows that, depending on the objective of the social network, different variables have to be considered for the task of inferring the tie strength.

Fogues et al. [23] introduce a tool called Best Friend Forever (BFF) that automatically groups and assigns a tie strength value for the contacts of a user. In order to infer tie strength values, BFF follows an approach similar to [28] and [39]. However, BFF uses a much smaller set variables, only 11. The reduction in the number of variables makes the variable collection task faster and less costly, thus increasing the utility of the tool. For automatic group creation, Fogus et al. used the algorithm proposed by [61]. This hierarchical diffusion algorithm is founded on the triadic closure principle, which suggests that, in a social network, there is an increased likelihood that two people will become friends if they have friends in common. The authors made an experimental evaluation of the tool and compared the results obtained by their tool with the preferences of 17 participants. Despite the reduction of variables for tie strength prediction, the tool performed accurately. On the matter of groups, the tool also worked with precision. Fogu'es et al. present a relationship model that focuses on a specific SNS, Facebook. However, although several relationship defining variables used for Facebook can also be found in other SNSs, a more general model that works in different social networks is needed. Moreover, the authors mention a correlation between tie strength and community creation, but they did not study this fact.

As explained in the previous section, some ReBAC models allow the specification of multiple types of relationships. Wu et al. [72] deal with the task of differentiating between personal and professional closeness. The authors performed a survey with users of a professional SNS called Beehive. This SNS was deployed at the IBM company in 2007. The authors asked the participants to assign a value of tie strength to their contacts in the social network from a professional and a friendship point of view. The results and the most predictive variables identified by the authors are close to the ones obtained by Gilbert and Karahalios [28]. To differentiate a professional relationship from a personal one, the authors identified a set of variables that worked as strong predictors of a professional relationship. This research shows that different types of relationships may need specific

predictors, making the relationship classification a complex problem.

This section reviewed studies that aim to deal with an essential ReBAC requisite, the modeling and definition of relationships. In Section 8, we identify future research paths that can improve the way relationships are modeled and classified. For example, considering how one person discloses information with another can help to define the type of relationship that exists between them.

A ReBAC model cannot exist without a suitable social model. However, a model that truthfully represents human relationships is still useless if its users cannot manipulate and understand it with ease. As explained in Section 1.2, another requirement for a ReBAC is usability. Sections 6 and 7 show studies that aim to increase the usability of privacy mechanisms for SNSs. The studies reviewed propose tools that help users to configure, adapt, and understand their privacy preferences in SNSs.

6 Recommender Systems

A first troublesome task that SNS users must face when dealing with their privacy preference configuration is the definition of privacy policies. As shown in section 2, a ReBAC model can consider a large number of variables and use a complex privacy policy definition language (e.g. a language based on regular expressions or ontologies). It is not realistic to assume that an average SNS user is familiar with these concepts and can effectively use them. A common approach to help users in this task is to suggest privacy policies to them or guide them during the process. Users should have the last call on what can be disclosed and what is private. While recommenders can help reduce the user's burden, they are rarely perfectly accurate. Thus, it is important for the user to be able to view, understand, and modify the recommended policy before it is applied; it is also important for the user to be able to maintain the policy over time.

Figure 3 depicts a conceptual map for all of the studied approaches that propose a privacy recommender system for SNSs. The following subsections review the studies according to the classification shown in the conceptual map.

6.1 Based on What Others Do

A first approach for privacy policy recommenders is based on how other users set their privacy policies. For example, if every friend of a user decides to hide their address information, then it is likely that a good recommendation for the user is to also hide her address information. In this section, we review studies that use this approach.

Bonneau et al. [7] propose a tool that suggests privacy policies based on expert users' configuration. Users can specify their privacy policies and then share them over the SNS. Any user can apply the privacy policies shared by another user, rate them, and recommend them to her friends. It can be expected that the best privacy policies and the experts who created them will have high rates and will be used by a high number of users. Moreover, users can subscribe to their favorite privacy expert. In this way, users' privacy policies will automatically update when their preferred expert updates her policies. This approximation only allows the automation of privacies with low granularity as only general policies can be shared by experts.

Squicciarini et al. [62] present a privacy manager named PriMa that suggests privacy policies taking into account the sensitivity of content according to what users tend to do on the SNS and tie strength. To use PriMa, first, the user expresses her concerns about disclosing each of the attributes
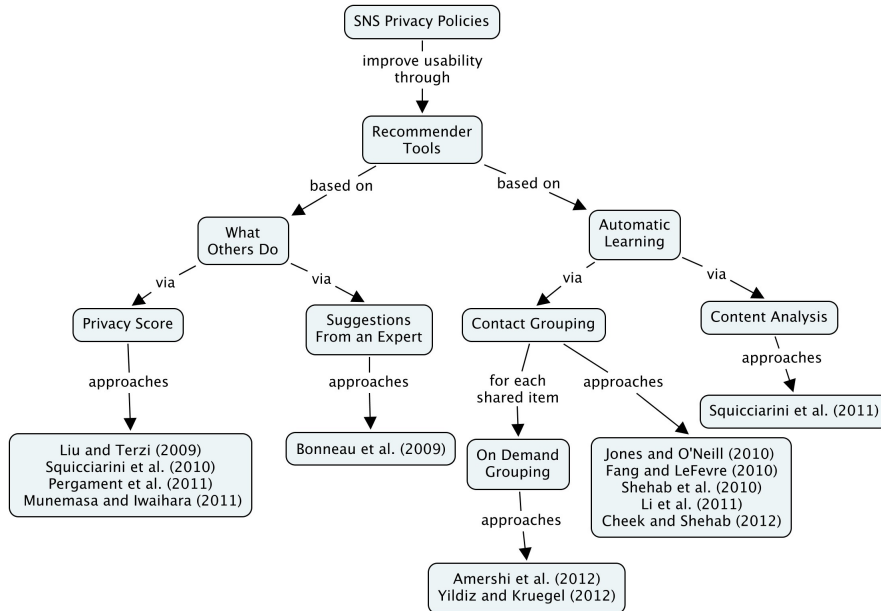
Figure 3: Conceptual Map of Privacy Recommender Systems

in her profile. If an attribute is left without a sensitivity value, PriMa infers this value from the values assigned to that trait by the contacts of the user. Once each trait has a sensitivity value, they are grouped in clusters depending on their sensitivity. In order to recommend policies, PriMa computes a user access score, which is a representation of the adequacy of a given target user to access a given cluster of attributes of the main user's profile. The score is based on the type of the relationship between the target user and the main user. The type of the relationship has an associated tie strength value that is predefined by the SNS provider. Finally, if the user access score is higher than a threshold, then that user has access to that attribute. The paper lacks an experimental evaluation also, the assumption of predefined types of relationships and tie strength values can lead to policy generation errors. For example, it is not possible to assume that every user agrees that a family relationship tie has a strength of 0.8 and a friend tie has a strength of 0.5.

Munemasa and Iwaihara [51] follow the line of Bonneau and Liu research [7]. These authors propose software that tells the user if her privacy settings are introvert or extrovert by computing a privacy score and comparing it with other SNS users' scores. The software created by the authors collects privacy settings of several SNS users and rates them. The rating of a privacy setting is based on the volume of information disclosed and how likely it is that others disclose the same information. The authors base the calculus of the privacy rating on the work of [47]. When a users configures her privacy settings, the software rates her configuration and compares it to the settings previously collected. From this information the software can determine whether the user's privacy setting discloses more or less information than the average user of the SNS. The authors performed an evaluation of their tool with 15 Facebook users. The evaluation participants only answered a questionnaire about the suitability of the tool. The article lacks experimentation that evaluates how accurate the recommendations made by the tool are (e.g. comparing them to the actual preferences of privacy concerned participants).

Pergament et al. [54] present a system named FORPS that helps user to decide what can be

accessed by a friend depending on the behavior of that friend on the SNS. FORPS calculates a privacy score for a target user for different themes. For example, a user can be very discreet with regard to religion and very indiscreet regarding political views. To calculate the score for each different theme, FORPS analyzes whether the target user discloses information on that matter with main the user's friends, her number of commentaries about that theme, and the sentiments on those commentaries. The system also analyzes the behavior of the friends of the main user in regard to the target user. Finally, FORPS recommends to the principal user a level of privacy for each theme towards the target user. The proposal lacks an experimental evaluation. Therefore, the utility of the system is not evaluated. Additionally, the system only recommends a numeric level of privacy for each theme. It is difficult to map a numeric value for what things is recommended to disclose and what are not. For example, on a 0 - 1 scale, what does a 0.3 degree of disclosure in politics mean? Should I disclose my preferred political party or not?

6.2   Based on Automatic Learning

Another approximation for recommenders is based on learning. Proposals in this section observe how the users interact with others and learn from their preferences. Once the recommender has learned enough, it is capable of automatically recommending privacy preferences. This section reviews studies that use the concept of automatic learning.

Jones and O'Neill [37] investigate what criteria humans use when they divide their social network in groups so they can share different information and avoid embarrassing situations with each group. The authors recruited 15 participants and asked them what factors they consider when creating groups in Facebook. According to their findings, the factors that affect grouping sorted from more important to less are: cliques (densely connected groups), tie strength, geographical location, organizational boundaries, temporal episodes, and functional roles. The authors also tested the clustering algorithm SCAN and compared its output with the groups manually created by the participants and they achieved 44.8% of similarity. In order to improve the accuracy of the algorithm the authors used some information from their interviews and added some information about tie strength to the algorithm. They could not add tie strength information for all participant data. With tie strength, the algorithm achieved an average similarity of 67%. This study shows that several factors are taken into account by users when creating groups. However, the authors did not collect enough data from participants' profiles and could not improve the clustering algorithm by modifying it, so it considers all the discovered factors.

Fang and LeFevre [21] propose a *wizard* software that suggests users privacy policies for different items on their profiles, like birthday, address, or telephone number. Since the proposal of Fang et al. uses supervised learning, participation of the user is needed. First, users' contacts are hierarchically grouped. To form the groups, the wizard considers community, profile, and activity features. For example, mutual friends is a community feature, while hobbies or fan pages that a user likes are considered to be activity features. Once the groups are created, the wizard asks the user to assign access grants to some of their contacts. The main idea of this process is that the user assigns access grants to the more representative users of the groups previously created. In this way, the user only needs to assign a low number of grants, and the process is much faster. According to the results of the proposal's evaluation, the wizard behaves better when only community features are considered. This can happen because many users do not specify their hobbies or demographic information. The privacy wizard is designed to protect only user's traits, like birth date, address, and telephone number. Other elements like images or videos are not considered by the wizard.

Moreover, to manage items like photos or videos, the wizard would need to be enhanced to consider specific features of the item, for example, tags on a photo.

Shehab et al. [60] introduce a privacy policy recommender system that is based on supervised learning. Their system works in 5 steps. In steps 1 and 2, the attributes of the main user's contacts are collected. Then, they are clustered according to their attribute similarity and a representative user is selected for each cluster. In step 3, the main user assigns access rights to the representative users. These labeled users are utilized by the classifier as the training set. In step 4, the rest of the contacts that are not labeled are classified and labeled accordingly. Finally, in the step 5, the main user's classifier look at the classifiers of the main user's neighbors and fuses itself with those that have similar clusters of users. One threat to privacy of using this recommender system is that it needs to access the privacy preferences of other users and these preferences should be also private.

Squicciarini et al. [64] propose a system called A3P that predicts a privacy policy for images in the context of social networks. A3P takes into account two variables when recommending a privacy policy for an image: social contexts and image content. A3P analyzes the content of the images and assigns a category to them. For the context, the authors predefine a set of social contexts (e.g. family, coworkers) and assign an intimacy value to each one of these contexts. Since A3P is based on supervised learning, it needs the user to specify some privacy policies before starting to predict policies. The policies can specify what contexts are allowed to access the image and what privileges each context has. Once A3P has learned enough from the user, it starts predicting policies using a policy mining algorithm. The policy mining algorithm considers the tendency in policy strictness of the user. Therefore, if the user tends to disclose more information, the suggested policies will be more extrovert and vice versa. Squicciarini et al. evaluated their proposal with humans; however, the photographs provided during the experimental evaluation were previously selected by the authors and do not correspond to real photos of the participants. An evaluation of how A3P performs with photographs taken by the experimental evaluation participants should be an extension of this paper.

Li et al. [44] present a privacy policy recommender based on semantics. Their approach is similar to the wizard presented in [21]. The main difference between the two approaches is that the one from Li uses semantics to find similarities among users. For example, a user can specify that he likes basketball and another that she likes NBA. Both hobbies are similar, but an approach without semantic knowledge will overlook this similarity. The authors present a k-Nearest Neighbors algorithm that uses semantic knowledge to compute the distance between two persons. When the user wants to specify access permissions for a new friend, the recommender suggests a policy that is already defined for the semantically closest k friends. The paper presents an experimental evaluation with 76 Facebook users and compares the performance of this approach against others. An interesting future improvement of this work would be the study how the use of techniques like uncertainty sampling can reduce the number of contacts to label so the recommender can accurately suggest privacy policies, thus reducing the effort required from the user.

Cheek and Shehab [14] introduce a privacy policy recommender that uses the similarity among friends to ease the process. Their approach is based on a graphical interface and offers two functionalities: an assistant to create groups of friends and *same-as* policy management. The group assistant guides users in the burdensome task of labeling their friends. The assistant presents a set of ten predefined social groups: Family, Close Friends, Graduate School, Under Graduate School, High School, Work, I do not know, Friends of Friend, Community, and Other. Each contact is presented to the user and the user is asked to select a group for that contact. In order to speed up the

process of labeling every friend, the assistant uses the Clasuet Newman Moore clustering algorithm to recommend groups for contacts that have not been labeled yet. Once every contact is labeled and belongs to a group, the *same-as* policy management asks the user to select a representative contact of each group and assign a privacy policy for that contact. The rest of the contacts on each group will be assigned the same privacy policy that was assigned to the representative contact of their group. The authors performed an extensive experimental evaluation where the users were asked to specify privacy policies for predefined groups of items: every album, demographic data, and educational data. The lack of granularity in privacy policies and the use of predefined groups of personal data limit the validity of the evaluation.

Amershi et al. [1] present a machine learning system called ReGroup that creates on-demand groups in social networks. The main difference between this proposal and others is that groups are not created once and used many times; instead, a new group is built for each item that is going to be shared. ReGroup uses a Naïve Bayes classifier to find similarities between users and make recommendations. When the user uploads an item to the SNS, ReGroup asks her to choose a contact to grant access to that item. After the first contact is chosen, the other contacts are sorted by similarity to the first one. Each time a similar contact is not selected, ReGroup adds a penalty to the similarity of that contact. In this way, this contact will appear in a later position the next time. The authors tested ReGroup and the standard application of Facebook, where users are sorted alphabetically in terms of group creation time and happiness of the user with the group created. The results obtained and the opinions of the participants highlight that each option works better depending on the size of the group. Alphabetical order is better for small groups, and ReGroup is better for large groups. This proposal alleviates the process of assigning a privacy policy for items; however, the user still has to select each user that is allowed to access the item separately. A future expansion of the article could be to compare what option is faster and easier for the user: correct a privacy policy suggestion or create a privacy policy from suggestions.

Yildiz and Kruegel [76] introduce a new algorithm that creates groups considering the users referenced by the item being shared and their social connections. The idea is that the groups created can help users to decide the privacy policy for the uploaded item. When a new item is uploaded, the algorithm creates a list of *participants*, who is the owner of the item, every user involved in the item (for example, users tagged in a photo), and a list of *candidates*, who are the friends of the participants. In each iteration of the algorithm, a candidate is inserted to the list of participants (users that are allowed to view the item). The inserted candidate is the one who maximizes a heuristic function. This function returns a high value when the analyzed candidate has many common friends with the users in the list of participants. Besides, the algorithm considers the tightness of the list of participants. If the participants are tightly connected, added candidates also have to be tightly connected to the participants; otherwise, the algorithm will look for loosely connected candidates. The algorithm keeps adding candidates until it cannot find an appropriate one. The authors performed an evaluation comparing the results of their algorithm with other local community finding algorithms. The authors considered for the best result for an algorithm to recover the entire group of the participants of the item while keeping to the minimum the number of members that do not belong to the participant group. In other words, they considered that the social cliques were entirely defined by the participants of the items. One problem with this assumption is that the privacy policies will probably tend to be very restrictive. Moreover, the article lacks an evaluation of the quality of the proposed groups made by human participants.

In Section 8, we have identified future lines of research for privacy policy recommenders. A
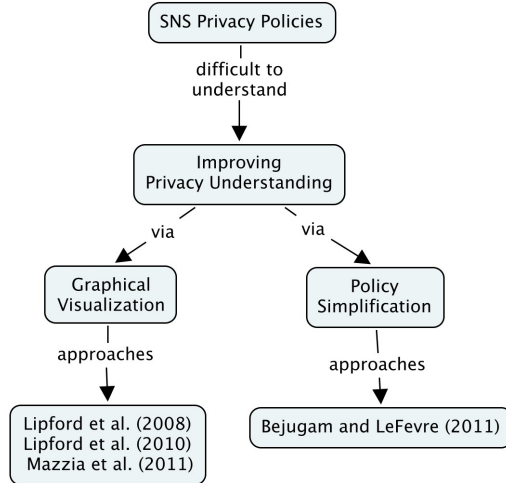
Figure 4: Conceptual Map for Improving Privacy Setting Understanding

crucial challenge is to combine privacy recommender tools with ReBAC models. As ReBAC models are key to the suitable management of information disclosure on SNSs, privacy policy recommenders should create policies according to the policy language defined by the access model.

The papers analyzed in this section focus on making privacy policies easy to configure. However, privacy policies should not only be manageable, they also have to be easy to maintain over time so that the users can adapt them as new relationships are established and previous ones change. Moreover, as the users set and refine their privacy policies, the global view of how their information is accessed by others grows in complexity. Ideally, users have to be perfectly aware of who has access to their private information and to what specific parts of that information. Section 7 expands the concept of usability by introducing the notion of privacy setting understanding and also reviews proposals that help users to create better mental models of their privacy settings in SNSs.

7 Improving Privacy Settings Understanding

A great obstacle that users find when dealing with privacy on SNSs is the difficulty of figuring out how and what personal information is disclosed over the SNS. SNSs have improved their user interfaces; for example, Google+ offers the Social Circles graphical tool that facilitates the laborious task of grouping contacts. These new interfaces facilitate the process of managing privacy settings and figuring out what others can see from our profile. However, there is still room for improvement, and if in the long term SNS developers aim to implement more complex access control models, they have to accompany them with easy-to-understand and easy-to-use interfaces. Currently, there are two different approaches to make privacy policies more understandable: graphical interfaces and privacy policy simplification. On the one hand, graphical interfaces use visual tools to enhance the understanding of the privacy policies. On the other hand, privacy policy simplification tries to remove unnecessary complexity of privacy policies and make them more legible for the user.

Figure 4 depicts a conceptual map for all of the reviewed approaches that propose tools or mechanisms that help users to understand their privacy settings in SNSs. The following subsections explain the approaches according to the order shown in the conceptual map.

7.1 Graphical Visualization

Lipford et al. [45] propose a new privacy management interface for Facebook called AudienceView. This new interface offers users the possibility to observe how their profiles are seen from different points of view. Each different point of view corresponds to one of the possible audiences: search, network, an individual friend, and friend group. To evaluate their proposal the authors recruited 16 participants and measured their performance in different tasks using the standard Facebook interface and the proposed new one. The results showed that users felt more comfortable with the new interface and it helped them to understand better the actual consequences of their privacy settings. Currently, the prototype is limited as it only shows how the profile data (name, hometown, hobbies...) is seen from different audience points of view. It would be very useful to expand the prototype so profile publications like photos or comments could be inspected by AudienceView. This improvement would probably require a new interface and a clever graphic design to make it usable and understandable for the user.

Lipford et al. [46] compare their previous proposal [45] with another kind of privacy policy representation, Expandable Grids. Expandable Grids show precisely what a policy allows or does not allow in a matrix using hierarchical axes that can be expanded or contracted to show more or less policy detail. Expandable Grids were initially proposed for access control policies in file systems. Therefore, a more extensive review of this representation is out of the scope of this paper. For further information of Expandable Grids, please refer to [57]. The experimental results obtained by Lipford et al. showed that both representations were highly usable, and they did not find clear advantages of one over the other. However, experiment participants did have clear and different preferences; some of them preferred the compact view of Expandable Grids and others the visual feedback of AudienceView.

Mazzia et al. [50] present PViz, an interface that is focused on helping users understand the visibility of their profiles. PViz presents the social structure of user's contacts in a graphical way. Contacts are automatically grouped in communities using the idea of modularity optimization. This methodology generates a hierarchical structure of groups. The higher levels of the hierarchy have groups of users that are loosely connected and share less attributes, the lower levels of the structure represent groups of users that share many common friends, tastes, and demographic data. This approach for automatically creating communities is also used in [21]. Moreover, PViz automatically labels each group, selecting the most common attribute of the contacts that conform the group. For example, if the majority of contacts in a group are from Washington, that group will be labeled as Washington. The graphical interface allows the user to select an attribute and see what communities are allowed to view that attribute. Communities are colored; the darker the color is, the larger the number of members of the community that can view that attribute. As the communities are hierarchically organized, the graphical interface allows users to zoom in or out of a group to see more or less detail of that group. The authors empirically evaluated PViz comparing it with AudienceView [45] and the Facebook standard interface. Their results showed that participants preferred PViz over the other two options; some participants even suggested that a combination of AudienceView and PViz could create a better solution.

7.2 Privacy Policy Simplification

Bejugam and LeFevre [3] present a policy simplification utility. Current privacy policy specification tools are based on rules that allow the user to specify positive rules ("Show this to") and negative

rules ("Hide this to"). These specifications can lead to long and complex policies, and even include redundancy, making them difficult to understand and adapt when new contacts or groups are created. The approach of Bejugam and LeFevre is based on two functionalities: automatic community creation and a policy consolidator. The automatic community creation algorithm is the same algorithm used in [50] and [21]. The policy consolidator transforms a privacy policy specification and produces the smallest equivalent privacy policy. The authors performed an experimental evaluation with nine subjects. Their test proved the utility of their approach and the hypothesis, which states that users will comprehend, remember, and maintain simplified policies easier than their verbose counterparts.

8 Open Challenges

Although the proposals shown in this survey cover some of the requisites that an access control for SNS should fulfill, we have identified many possible lines for future research. In this section, we outline some of the most challenging possible future directions in the research field of access control models for social networks and their usability. These possible paths of research are open challenges that were identified during the realization of this survey. The accomplishment of these open challenges will have a great impact in determining SNSs to be useful, entertaining, and privacy safe services available at the Web 2.0. On the one hand, SNS users will feel safer in the context of SNSs and will feel more inclined to register and use these services. On the other hand, we believe that the utility of SNSs will increase since users will upload and share much more information in their profiles.

8.1 ReBAC and Content Type

A number of works show that users take into account the type of the content when they are defining how that content will be disclosed [32, 42, 75]. However, current access control do not consider content type. On the one hand, including a content type as a new attribute of access control can improve the flexibility and expressiveness of privacy policies. On the other hand, more attributes can increase the complexity of access controls and privacy policies. Therefore, future research should evaluate the effect that the inclusion of the attribute content type can have over access controls. Specifically, new research should evaluate the complexity of the privacy policies created with the new access control, the required number of privacy policies that users need to express their privacy policies, the number of privacy conflicts generated by the policies, and the understandability of these policies.

8.2 Inferring Tie Strength from Different Sources

The studies about tie strength in SNSs use information available on the SNS; for instance, they consider the number of messages exchanged between users, the number of photographs shared or the number of common friends. A common pitfall of these works is that relationships that mainly occur outside the SNS are not well evaluated [28]. A possible solution for this problem is to search for information outside the SNS. Recent research [49], [18] infer tie strength values from data available on the Internet or from real life. Mixing data from the SNS and from other sources would create approaches that infer tie strength more accurately. However, the use of external data can create

some privacy concerns since it directly relates to re-identification and profiling privacy breaches (these breaches are defined in section 1.1).

The proposal of Murukannaiah and Singh [52] is a first attempt to use real world information to infer the social data of the user. The authors present a tool called Platys Social that runs on a mobile device. This software learns a user's social circles and the priority of the user's social connections from daily interactions. The software infers the interactions from information that is available on mobile devices, such as wi-fi networks, bluetooth connections, phone calls, and text messages. Combining all sources of information (real world through a mobile device, Internet, and social network) could positively improve the tie strength inference and classification of relationships.

8.3   Adaptive Relationship Models

In human relationships, the disclosure of private information represents an important part of these relationships [30]. One of the main reasons why people exchange private information is because they perceive that in the future this information disclosure may become a gain in tie strength. The way users share with others change the perception of both parts about the way they should interact in the future. For example, if user A constantly discloses information to user B, but user B only reveals a small portion of his information to user A, it can be assumed that in the mid/long term user A will reduce the amount of information disclosed to B, or even stop communication at all.

Such et al. [66] propose a self-disclosure decision-making mechanism for multiagent systems. The proposal is based on psychological findings regarding how humans disclose personal information in the building of their relationships. The implementation of this mechanism (or a similar one) to a relationship model would increase the accuracy of that model, especially over time. The decision of specifying a privacy policy would consider not only the current situation of the user's relationships, but also how the new specified policy itself would affect the relationships and the gain or loss of tie strength for those relationships.

8.4   Tie Strength Utility

The studies reviewed in section 5 infer the tie between two individuals in only one dimension, the strength. However, humans behave differently with their contacts depending not only on the strength of the relationship but also on its utility. Ties can be viewed not only for their strength but also how useful a tie is depending on the situation or the need. For example, a person can rely on a professionally well-connected friend to look for employment and ask another one for advice when going to have dinner, even when the tie strength with both friends is high.

A first attempt in this direction is the paper by Rosen and Chu [58]. The authors claim that the strong-weak tie dichotomy is conceptually misleading, and propose a multi-dimensional taxonomy of social network ties. The authors propose that a tie should not be evaluated by its strength but by its utility. In this sense, the authors study the possible social dimensions that matter during specific contexts. Future research should address to what degree each dimension affects the tie utility as well as how to identify the context an interaction belongs to.

8.5   Self-presentation Management

Kairam et al. [40] detected that the most common motivation to use SNS and share information is to create a self-presentation. Users care about what image they project on others and how they

affect others. Klout[6], which is a tool that helps users to see how they influence others in popular social networks, is a first step towards a self-presentation management system.

The studies reviewed in 7 aid SNS users to understand what other users can see in their profiles according to their privacy preferences. A similar approximation could be used so users could express their privacy policies based on facets. Users would find it more natural to specify which facet of their life they are willing to show to each contact, instead of a sequence of rules that specify who can see what and who cannot. For example, a user could choose to show a funny facet to her family members or friends, allowing them to see comical photos or posts, while showing a professional facet to colleagues or potential employers.

8.6 ReBAC: Usability and Visibility

A change in the access controls of SNSs will represent a change in their privacy controls as well. ReBAC models, such as the ones proposed by [11] or [15] use technologies like semantic web or regular expressions. Suitable privacy controls for SNSs should hide the complexity of such technologies and facilitate their use. Moreover, the inclusion of new attributes to the access control (tie strength and content type) also suppose a modification of privacy controls. Further studies should address how usable privacy controls can be created for new ReBAC controls for SNSs.

In the same fashion, privacy visualization tools will have to adapt to ReBAC models. For example, visualization tools should explain to the users in an understandable way how their information is disseminated according to a specific type of relationship. Moreover, complex models that allow the specification of a privacy policy hierarchy, user-to-resource relationship, or social paths will require visualization tools with clever designs. Google+ and the friend circles application is a good example of how privacy visualization tools can be designed to be usable and engaging for users.

8.7 ReBAC and Co-privacy

As stated in Section 4, co-privacy can cause several privacy conflicts among the users involved. SNSs should offer integrated solutions for situations where a shared item can cause tension among users. This tension and the complete lack of control of these situations can lead to users adopting strategies like un-friending the user who uploaded the controversial item or deleting their profile from the SNS. Hence, it is necessary to add shared item or co-privacy management to the the access control models for SNSs. However, there is no proposal that brings together a formal ReBAC model and shared content privacy management.

Some of the proposals reviewed in section 4 could be merged with a formal ReBAC. Users should be able to specify how strict their privacy preferences are on a shared item. Moreover, the ReBAC should have a privacy conflict resolution that guarantees that the resulting privacy policy maximizes the utility for every used involved. For example, a ReBAC policy language could allow users to specify a value that indicates how strong or weak their conditions are for each item or with respect to a certain type of relationship.

---

[6]http://klout.com/home

24

8.8 Privacy Settings Interoperability

There are many different SNSs, and they accomplish different objectives; for example, Facebook and Google+ are focused on maintaining friendships; LinkedIn[7] is focused on professional life; Meetic[8] aims to help users to find dates with similar people; and Flickr[9] is a social network that is centered on photograph sharing. Besides, the wider use of mobile devices that work together with SNSs also increases the variety and diversity of SNSs and their uses.

It is a common practice for SNS users to have profiles on different SNSs and to use each one for different purposes according to the general objective of the SNS. As pointed out in this paper, setting privacy settings is a burdensome task and users struggle doing it. Therefore, if a user has to configure her privacy settings separately for each one of the SNS she is using, this task becomes even more tiresome. Users need privacy settings that are interoperable among every SNS that they are members of. A first step in achieving interoperability for privacy settings could be that access control models of SNS use a universal and well-defined privacy ontology for SNSs. This ontology should be able to differentiate among the different contexts that each SNS belongs to. To improve the interoperability of privacy settings, personal privacy agents could be developed. These agents will take care of their principals' privacy, adapting their behavior to the context their principal is at that moment.

8.9 Sticky Policies

Private information stored in the profile of a SNS user can move or be moved to different contexts. In order to maintain the privacy of that information, it is necessary for the privacy policies associated to the information move along with it. Policies of this kind are known as sticky policies and they enable users to improve control over their personal information as it moves across multiple parties. Sticky policies are orthogonal to ReBAC models, however, they become more necessary as mobile and pervasive systems gain in popularity. Users can share their information using a wide variety of devices and applications. Sticky policies can guarantee self-disclosure consistency among every device and application.

A common situation (in the context of SNSs) where personal information moves to a different party is the use of applications created by third parties and integrated in the SNSs. In [22] the authors propose a method to keep the maximum amount of personal information hidden to third party applications in Facebook. However, to the best of our knowledge there is no study that proposes how information is managed after the third party acquires it. The recent study of Pearson and Mont [53] presents a general framework for using sticky policies called EnCoRe. A similar approach could be used for SNSs, where the SNS itself would work as a trusted authority in charge of controlling that the information is disclosed by third party applications according to the owner's preferences and also of regulating the access to such data.

9   Conclusions

With the constant increase in the use of SNSs, mobile devices that connect us with others at all times, and in general, the spread of pervasive computing, we consider that privacy will be of

---

[7]http://www.linkedin.com/
[8]http://www.meetic.com
[9]http://www.flickr.com/

paramount importance during the next few years. In a connected world, controlling our privacy and what others are able to see about us will be more difficult and complex. Hence, powerful, easy-to-use, and intuitive privacy solutions will be the subject of many research efforts during next years.

In this paper, we have reviewed approaches that offer partial solutions to the most critical problems of privacy management on SNSs. However, current SNSs have not adopted them and still lack the suitable privacy management tools. Approaches like Google+, where the control of information dissemination has been given great visibility [48], are first steps towards SNSs that are more respectful of privacy. In the not-so-distant future we envision a SNS that offers a privacy mechanism that satisfies every requisites mentioned in this paper and provides the features that users demand. In order to develop this ideal SNS, developers and researchers will have to deal with several challenges. The inclusion of ReBAC models in popular SNSs will improve the control of privacy for the users. However, ReBAC models are complex and SNSs will require a thorough design that guarantees the usability of the privacy management. Moreover, sticky policies and privacy settings interoperability will represent a technological challenge. We believe that universal privacy policy languages and access control models will be required to ensure these two requisites. The model and the language proposed would need to be flexible enough to allow different SNSs that focus on different social aspects.

References

[1] S. Amershi, J. Fogarty, and D. Weld. Regroup: Interactive machine learning for on-demand group creation in social networks. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. ACM, 2012.

[2] S. Barnes. A privacy paradox: Social networking in the united states. *First Monday*, 11(9):11–15, 2006.

[3] R. Bejugam and K. LeFevre. enlist: Automatically simplifying privacy policies. In *Data Mining Workshops (ICDMW), 2011 IEEE 11th International Conference on*, pages 620–627. IEEE, 2011.

[4] A. Besmer and H. Richter Lipford. Moving beyond untagging: photo privacy in a tagged world. In *Proceedings of the 28th international conference on Human factors in computing systems*, pages 1563–1572. ACM, 2010.

[5] L. Bilge, T. Strufe, D. Balzarotti, and E. Kirda. All your contacts are belong to us: automated identity theft attacks on social networks. In *Proceedings of the 18th international conference on World wide web*, pages 551–560. ACM, 2009.

[6] K. Bischoff. We love rock'n'roll: Analyzing and predicting friendship links in lastfm. In *Proceedings ACM Web Science Conference, Web Science 2012, WebSci'12, Evanston, Illinois, USA, June 22-24, 2012. to appear*, 2012.

[7] J. Bonneau, J. Anderson, and L. Church. Privacy suites: Shared privacy for social networks. In *Symposium on Usable Privacy and Security (SOUPS)*. Citeseer, 2009.

[8] D. Boyd and E. Hargittai. Facebook privacy settings: Who cares? *First Monday*, 15(8), 2010.

[9] D. Boyd and J. Heer. Profiles as conversation: Networked identity performance on friendster. In *System Sciences, 2006. HICSS'06. Proceedings of the 39th Annual Hawaii International Conference on*, volume 3, pages 59c–59c. IEEE, 2006.

[10] G. Bruns, P. Fong, I. Siahaan, and M. Huth. Relationship-based access control: its expression and enforcement through hybrid logic. In *Proceedings of the second ACM conference on Data and Application Security and Privacy*, pages 117–124. ACM, 2012.

[11] B. Carminati, E. Ferrari, R. Heatherly, M. Kantarcioglu, and B. Thuraisingham. Semantic web-based social network access control. *computers & security*, 30(2-3):108–115, 2011.

[12] B. Carminati, E. Ferrari, and A. Perego. Rule-based access control for social networks. In *On the Move to Meaningful Internet Systems 2006: OTM 2006 Workshops*, pages 1734–1744. Springer, 2006.

[13] B. Carminati, E. Ferrari, and A. Perego. Enforcing access control in web-based social networks. *ACM Transactions on Information and System Security (TISSEC)*, 13(1):6, 2009.

[14] G. P. Cheek and M. Shehab. Policy-by-example for online social networks. In *Proceedings of the 17th ACM symposium on Access Control Models and Technologies*, SACMAT '12, pages 23–32, New York, NY, USA, 2012. ACM.

[15] Y. Cheng, J. Park, and R. Sandhu. A user-to-user relationship-based access control model for online social networks. In N. Cuppens-Boulahia, F. Cuppens, and J. Garcia-Alfaro, editors, *Data and Applications Security and Privacy XXVI*, volume 7371 of *Lecture Notes in Computer Science*, pages 8–24. Springer Berlin Heidelberg, 2012.

[16] M. Cooke and N. Buckley. Web 2.0, social networks and the future of market research. *International Journal of Market Research*, 50(2):267–292, 2008.

[17] L. Cranor and S. Garfinkel. *Security and Usability*. O'Reilly Media, Inc., 2005.

[18] L. Ding, D. Steil, B. Dixon, A. Parrish, and D. Brown. A relation context oriented approach to identify strong ties in social networks. *Knowledge-Based Systems*, 24(8):1187–1195, 2011.

[19] S. Duck. *Human relationships*. Sage Publications Ltd, 2007.

[20] N. Ellison, C. Steinfield, and C. Lampe. The benefits of facebook friends: Social capital and college students' use of online social network sites. *Journal of Computer-Mediated Communication*, 12(4):1143–1168, 2007.

[21] L. Fang and K. LeFevre. Privacy wizards for social networking sites. In *Proceedings of the 19th international conference on World wide web*, pages 351–360. ACM, 2010.

[22] A. Felt and D. Evans. Privacy protection for social networking apis. *2008 Web 2.0 Security and Privacy (W2SP'08)*, –, 2008.

[23] R. Fogues, J. Such, A. Espinosa, and A. Garcia-Fornes. Bff: A tool for eliciting tie strength and user communities in social networking services. *Information Systems Frontiers*, –:1–13, 2013.

[24] P. Fong. Relationship-based access control: protection model and policy language. In *Proceedings of the first ACM conference on Data and application security and privacy*, pages 191–202. ACM, 2011.

[25] P. Fong, M. Anwar, and Z. Zhao. A privacy preservation model for facebook-style social network systems. In M. Backes and P. Ning, editors, *Computer Security – ESORICS 2009*, volume 5789 of *Lecture Notes in Computer Science*, pages 303–320. Springer Berlin Heidelberg, 2009.

[26] C. Gates. Access control requirements for web 2.0 security and privacy. *IEEE Web*, 2(0), 2007.

[27] E. Gilbert. Predicting tie strength in a new medium. In *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work*, pages 1047–1056. ACM, 2012.

[28] E. Gilbert and K. Karahalios. Predicting tie strength with social media. In *Proceedings of the 27th international conference on Human factors in computing systems*, pages 211–220. ACM, 2009.

[29] M. Granovetter. The strength of weak ties. *American journal of sociology*, 78(6):l, 1973.

[30] K. Greene, V. Derlega, and A. Mathews. Self-disclosure in personal relationships. *The Cambridge handbook of personal relationships*, –:409–427, 2006.

[31] R. Gross and A. Acquisti. Information revelation and privacy in online social networks. In *Proceedings of the 2005 ACM workshop on Privacy in the electronic society*, pages 71–80. ACM, 2005.

[32] M. Hart, C. Castille, R. Johnson, and A. Stent. Usable privacy controls for blogs. In *Computational Science and Engineering, 2009. CSE'09. International Conference on*, volume 4, pages 401–408. IEEE, 2009.

[33] M. Hildebrandt. Defining profiling: A new type of knowledge? In M. Hildebrandt and S. Gutwirth, editors, *Profiling the European Citizen*, pages 17–45. Springer Netherlands, 2008.

[34] G. Hogben. Security issues and recommendations for online social networks. *Position Paper ENISA European Network and Information Security Agency*, 80211(1), 2007.

[35] H. Hu and G.-J. Ahn. Multiparty authorization framework for data sharing in online social networks. In Y. Li, editor, *Data and Applications Security and Privacy XXV*, volume 6818 of *Lecture Notes in Computer Science*, pages 29–43. Springer Berlin Heidelberg, 2011.

[36] M. Johnson, S. Egelman, and S. M. Bellovin. Facebook and privacy: it's complicated. In *Proceedings of the Eighth Symposium on Usable Privacy and Security*, SOUPS '12, pages 9:1–9:15, New York, NY, USA, 2012. ACM.

[37] S. Jones and E. O'Neill. Feasibility of structural network clustering for group-based privacy control in social networks. In *Proceedings of the Sixth Symposium on Usable Privacy and Security*, page 9. ACM, 2010.

[38] L. Kagal, C. Hanson, and D. Weitzner. Using dependency tracking to provide explanations for policy management. In *Policies for Distributed Systems and Networks, 2008. POLICY 2008. IEEE Workshop on*, pages 54–61, June 2008.

[39] I. Kahanda and J. Neville. Using transactional information to predict link strength in online social networks. In *Proceedings of the Third International Conference on Weblogs and Social Media (ICWSM)*, 2009.

[40] S. Kairam, M. Brzozowski, D. Huffaker, and E. Chi. Talking in circles: selective sharing in google+. In *Proceedings of the 2012 ACM annual conference on Human Factors in Computing Systems*, CHI '12, pages 1065–1074, New York, NY, USA, 2012. ACM.

[41] G. Karjoth, M. Schunter, and M. Waidner. Platform for enterprise privacy practices: Privacy-enabled management of customer data. In R. Dingledine and P. Syverson, editors, *Privacy Enhancing Technologies*, volume 2482 of *Lecture Notes in Computer Science*, pages 69–84. Springer Berlin Heidelberg, 2003.

[42] P. Klemperer, Y. Liang, M. Mazurek, M. Sleeper, B. Ur, L. Bauer, L. F. Cranor, N. Gupta, and M. Reiter. Tag, you can see it!: using tags for access control in photo sharing. In *Proc. CHI*, pages 377–386, 2012.

[43] K. Koroleva and A. Bolufé Röhler. Reducing information overload: Design and evaluation of filtering & ranking algorithms for social networking sites. In *Proceedings of the ECIS 2012 European Conference on Information Systems*, 2012.

[44] Q. Li, J. Li, H. Wang, and A. Ginjala. Semantics-enhanced privacy recommendation for social networking sites. In *Trust, Security and Privacy in Computing and Communications (Trust-Com), 2011 IEEE 10th International Conference on*, pages 226–233. IEEE, 2011.

[45] H. Lipford, A. Besmer, and J. Watson. Understanding privacy settings in facebook with an audience view. In *Proceedings of the 1st Conference on Usability, Psychology, and Security*, pages 1–8. USENIX Association Berkeley, CA, USA, 2008.

[46] H. Lipford, J. Watson, M. Whitney, K. Froiland, and R. Reeder. Visual vs. compact: A comparison of privacy policy interfaces. In *Proceedings of the 28th international conference on Human factors in computing systems*, pages 1111–1114. ACM, 2010.

[47] K. Liu and E. Terzi. A framework for computing the privacy scores of users in online social networks. In *Data Mining, 2009. ICDM'09. Ninth IEEE International Conference on*, pages 288–297. IEEE, 2009.

[48] S. Mahmood and Y. Desmedt. Poster: preliminary analysis of google+'s privacy. In *Proceedings of the 18th ACM conference on Computer and communications security*, pages 809–812. ACM, 2011.

[49] Y. Matsuo, J. Mori, M. Hamasaki, T. Nishimura, H. Takeda, K. Hasida, and M. Ishizuka. Polyphonet: An advanced social network extraction system from the web. *Web Semantics: Science, Services and Agents on the World Wide Web*, 5(4):262 – 278, 2007. World Wide Web Conference 2006Semantic Web Track.

[50] A. Mazzia, K. LeFevre, and E. Adar. The pviz comprehension tool for social network privacy settings. In *Proceedings of the Eighth Symposium on Usable Privacy and Security*, SOUPS '12, pages 13:1–13:12, New York, NY, USA, 2012. ACM.

[51] T. Munemasa and M. Iwaihara. Trend analysis and recommendation of users' privacy settings on social networking services. In A. Datta, S. Shulman, B. Zheng, S.-D. Lin, A. Sun, and E.-P. Lim, editors, *Social Informatics*, volume 6984 of *Lecture Notes in Computer Science*, pages 184–197. Springer Berlin Heidelberg, 2011.

[52] P. Murukannaiah and M. Singh. Platys social: Relating shared places and private social circles. *Internet Computing, IEEE*, 16(3):53–59, May 2012.

[53] S. Pearson and M. Mont. Sticky policies: An approach for managing privacy across multiple parties. *Computer*, 44(9):60–68, 2011.

[54] D. Pergament, A. Aghasaryan, J. Ganascia, and S. Betgé-Brezetz. Forps: friends-oriented reputation privacy score. In *Proceedings of the First International Workshop on Security and Privacy Preserving in e-Societies*, pages 19–25. ACM, 2011.

[55] G. Pike. Fired over facebook. *Information Today*, 28(4):26–26, 2011.

[56] J. Rana, J. Kristiansson, and K. Synnes. Enriching and simplifying communication by social prioritization. In *Advances in Social Networks Analysis and Mining (ASONAM), 2010 International Conference on*, pages 336–340. IEEE, 2010.

[57] R. Reeder, L. Bauer, L. Cranor, M. Reiter, K. Bacon, K. How, and H. Strong. Expandable grids for visualizing and authoring computer security policies. In *Proceedings of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*, pages 1473–1482. ACM, 2008.

[58] D. Rosen and K. Chu. The utility of communication network ties: Reconceptualizing the social network tie measure. In *System Sciences (HICSS), 2011 44th Hawaii International Conference on*, pages 1–8. IEEE, 2011.

[59] P. Samarati and L. Sweeney. Protecting privacy when disclosing information: k-anonymity and its enforcement through generalization and suppression. Technical report, Technical report, SRI International, 1998.

[60] M. Shehab, G. Cheek, H. Touati, A. Squicciarini, and P. Cheng. Learning based access control in online social networks. In *Proceedings of the 19th international conference on World wide web*, pages 1179–1180. ACM, 2010.

[61] K. Shen, L. Song, X. Yang, and W. Zhang. A hierarchical diffusion algorithm for community detection in social networks. In *Cyber-Enabled Distributed Computing and Knowledge Discovery (CyberC), 2010 International Conference on*, pages 276–283. IEEE, 2010.

[62] A. Squicciarini, F. Paci, and S. Sundareswaran. Prima: An effective privacy protection mechanism for social networks. In *Proceedings of the 5th ACM Symposium on Information, Computer and Communications Security*, pages 320–323. ACM, 2010.

[63] A. Squicciarini, M. Shehab, and F. Paci. Collective privacy management in social networks. In *Proceedings of the 18th international conference on World wide web*, pages 521–530. ACM, 2009.

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

[65] K. Strater and H. Lipford. Strategies and struggles with privacy in an online social networking community. In *Proceedings of the 22nd British HCI Group Annual Conference on People and Computers: Culture, Creativity, Interaction-Volume 1*, pages 111–119. British Computer Society, 2008.

[66] J. M. Such, A. Espinosa, A. García-Fornes, and C. Sierra. Self-disclosure decision making based on intimacy and privacy. *Information Sciences*, 211:93 – 111, 2012.

[67] K. Thomas, C. Grier, and D. Nicol. unfriendly: Multi-party privacy risks in social networks. In *Privacy Enhancing Technologies*, pages 236–252. Springer, 2010.

[68] S. D. Warren and L. D. Brandeis. The right to privacy. *Harvard Law Review*, 4(5):pp. 193–220, 1890.

[69] A. Westin. Privacy and freedom. *Washington and Lee Law Review*, 25(1):166, 1968.

[70] J. Wiese, P. Kelley, L. Cranor, L. Dabbish, J. Hong, and J. Zimmerman. Are you close with me? are you nearby? investigating social groups, closeness, and willingness to share. In *Proceedings of the 13th international conference on Ubiquitous computing*, pages 197–206. ACM, 2011.

[71] R. Wishart, D. Corapi, S. Marinovic, and M. Sloman. Collaborative privacy policy authoring in a social networking context. In *Policies for Distributed Systems and Networks (POLICY), 2010 IEEE International Symposium on*, pages 1–8. IEEE, 2010.

[72] A. Wu, J. DiMicco, and D. Millen. Detecting professional versus personal closeness using an enterprise social network site. In *Proceedings of the 28th international conference on Human factors in computing systems*, pages 1955–1964. ACM, 2010.

[73] R. Xiang, J. Neville, and M. Rogati. Modeling relationship strength in online social networks. In *Proceedings of the 19th international conference on World wide web*, pages 981–990. ACM, 2010.

[74] M. Yao, R. Rice, and K. Wallis. Predicting user concerns about online privacy. *Journal of the American Society for Information Science and Technology*, 58(5):710–722, 2007.

[75] C. Yeung, L. Kagal, N. Gibbins, and N. Shadbolt. Providing access control to online photo albums based on tags and linked data. In *Proceedings of the AAAI Spring Symposium on Social Semantic Web: Where Web*, volume 2, 2009.

[76] H. Yildiz and C. Kruegel. Detecting social cliques for automated privacy control in online social networks. In *Pervasive Computing and Communications Workshops (PERCOM Workshops), 2012 IEEE International Conference on*, pages 353 –359, march 2012.

[77] C. Zhang, J. Sun, X. Zhu, and Y. Fang. Privacy and security for online social networks: challenges and opportunities. *Network, IEEE*, 24(4):13–18, 2010.