REACT: REcommending Access Control decisions To social media users

Gaurav Misra
Security Lancaster Research Center
Lancaster University
Lancaster, United Kingdom
Email: g.misra@lancaster.ac.uk

Jose M. Such
Department of Informatics
King's College London
London, United Kingdom
Email: jose.such@kcl.ac.uk

***Abstract*—The problems that social media users have in appropriately controlling access to their content has been well documented in previous research. A promising method of providing assistance to users is by learning from the access control decisions made by them and making future recommendations. In this paper, we present REACT, a learning mechanism which utilizes information available in the social network in conjunction with information about the content to be shared to provide users with access control recommendations. We demonstrate the highly accurate performance of REACT through a detailed empirical evaluation and also discuss ways of personalizing it for different users in order to improve performance even further.**

1. Introduction

Social media users interact with vast network of people representing various facets of their life such as work, family, education, etc. In such a scenario, it is essential for them to make informed access control decisions to preserve the "contextual integrity" of their information. Any user who discloses information on social media has a notion of the "intended recipients" and the context in which they would view that information and hence preservation of "contextual integrity" is essential in order to avoid a privacy breach [1]. Unfortunately, social media users have been found to be less than capable of accurately reasoning about the actual recipients of their content and even when they do, they struggle to appropriately control who can access their content on such platforms [2]. The inability of users to make appropriate access control decisions often results in "unintended disclosure" of information [3]. When users are aware of their audience, and are provided with finer grained access controls, they tend to be much more selective in granting access to their information [4]. Mainstream social media sites such as Facebook and Google+ have made an effort to assist users in managing their friend networks by providing them with "Lists" and "Circles" [4] respectively. However, recent research findings suggest that hardly any users employ these features when making access control decisions, arguably due to the effort this requires from them [5]. There have also been some efforts to improve visualization to enhance comprehension of access control policies among users [6] but even with these approaches, the burden of appropriately configuring access controls remains on the user. This burden can possibly be alleviated by providing accurate access control recommendation to users when they disclose information to their social network.

Providing recommendations to the user in a dynamic medium like social media is a sizeable challenge, and it is essential to consider and recognize the overall "context" of a disclosure to provide a "context-aware" recommendation which would safeguard the user's privacy [1]. In particular, the "social context" of any information disclosure in social media forms an essential part of the overall context and influences the access control decisions made by the user [7]. There are a lot of types of social network information which can be leveraged in order to appropriately define the social context and assist the user in configuring an appropriate access control policy. Different types of social relationships, commonly represented in terms of communities [8], [9], and the social relationship strength, commonly represented in terms of profile attribute similarity (or "closeness" exhibited by profile attributes) [10], [11], are considered important in influencing a user's access control policy. In addition to the "who", determined by social relationships, the "what", i.e. the content itself, also needs to be considered to make informed access control decisions [12]. As discussed later on in more detail and to the best of our knowledge, none of the previously proposed access control recommendation mechanisms consider both the "who" (relationship type and strength) and the "what".

In this paper, we bridge this gap and present REACT, which is an access control recommendation mechanism that considers the social context of information disclosure represented by the type and strength of social relationships in conjunction with information about the content in the form of annotations or "tags". Our results show that REACT provides highly accurate access control recommendations using these types of information. As privacy and access control behavior is very personal to individual users [13], we explore methods of personalizing REACT to go beyond a "one-size-fits-all" approach to improve performance even further. When the most favorable configuration is implemented for each user, REACT produces an **accuracy** of **93.2%** across our entire dataset.

2. Design

Figure 1 depicts REACT and the components it uses to make access control recommendations to the users. REACT is designed based on previous evidence which focuses on access control behaviors of social media users. For any such

ASONAM '17, July 31-August 03, 2017, Sydney, Australia

ACM ISBN 978-1-4503-4993-2/17/07...$31.00
http://dx.doi.org/10.1145/3110025.3110103

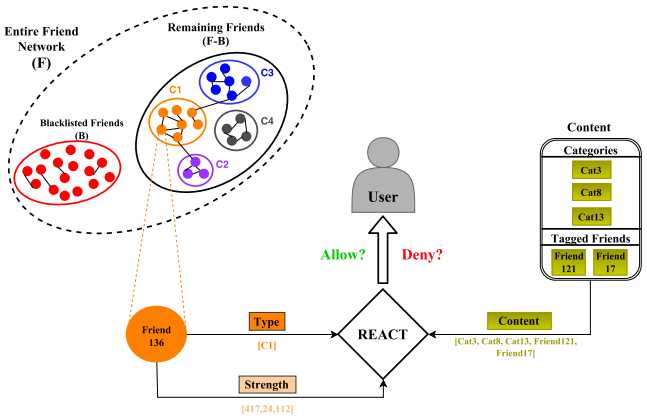

Figure 1: Figure showing the various components of REACT

recommendation mechanism to be usable, it is important to consider the social context of the information disclosure such that the "contextual integrity" of the information is preserved [7]. The social context can be derived from the interpersonal relationships, defined in terms of their "type" (friend, family, etc.) represented by partitions (or communities) [8], and the "strength" or closeness, represented by similarity of profile attributes [10], [11], all of which can be automatically retrieved and analyzed from the social network of the user. Additionally, the information about the content itself forms an integral part of the overall context of the disclosure. We describe the particular attributes of REACT which represent these types of information to represent the social context of disclosure in order to provide meaningful access control recommendations to the user.

2.1. Attributes Used

Relationship Type is represented by community membership in REACT. We considered network based community detection algorithms which require only the network connections of a user's friends in order to generate communities, which may or may not be overlapping as discussed later on. In our previous work, we empirically evaluated 8 well-known community finding algorithms and found that *Clique Percolation Method (CPM)* provides the best fit with respect to particular access control decisions made by social media users [14]. Even though CPM has a comparatively higher computational complexity as compared to other similar community detection algorithms, its impact on the timeliness of access control decisions would be minimal as it only needs to be executed to update communities when a new friend is added or removed from a user's network, but not every time a user would make an access control decision which happens much more often.

Relationship Strength can be calculated by fetching and analyzing the similarity of profile attributes and information in social networks [15], [16]. This has been also exploited as a basis to recommend access control decisions due to the known interplay between relationship strength and access control decisions [10]. However, there are four main problems associated with this approach: 1) many users leave profile attributes blank or set privacy controls so they cannot be retrieved; 2) there is a huge amount of potential attributes to be used and the associated time to fetch and analyze them; 3) some of the attributes could be seen as too privacy intrusive (e.g. entire conversations exchanged between users) to be used precisely to improve privacy and access control; 4) some of the attributes are dependent on the specific social media platform. Our previous research found that using only the size of network (*Total Friends*) and shared contacts (*Mutual Friends*) between individuals is sufficient in producing the same accuracy for access control recommendation as when using all attributes available while addressing the discussed problems [17]. In addition to these two attributes, REACT also uses the difference in network sizes of the user and each of their friends (*Friend Difference*) to represent relationship strength.

In addition to considering relationship defining attributes from the social network of a user, REACT also considers information about the **Content** being shared. This is essential in order to further define the context of any information disclosure, and previous research indeed showed that annotations of content by users employing "tags" can be used to create access control policies and that they are minimally disruptive for the user [18]. These tags can be about the topic of the content (e.g., Flickr has a list of potential topic categories for the photos users upload) or further information about the social context (e.g., friends being tagged in photos or mentions being made to them in text posts). The design of REACT is agnostic to the type of content being shared, and it considers all tags or information about the content available, whether manually provided by users or automatically inferred by tools such as [12].

2.2. Blacklisted Friends

Social media users often have vast friend networks consisting of many friends — as we will see later on users in our study had an average of 265 friends (s.d = 121). However, it has been empirically found that social media users often intend to share their content with a limited subset of their entire friend network [4]. REACT leverages this intuition by maintaining a "blacklist" of friends for each user. The access control decision corresponding to these friends would be a default "deny". Note that the specific number of blacklisted friends depends on various factors for individual users such as their access control behavior and their network size. For example, some users may want to connect to and share with as many people as possible. For these users, the number of friends in the blacklist would be minimal. Alternatively, users may choose to interact with only a small subset of their friends and the blacklist for these users may contain a majority of their friends. Thus, the blacklist construction in REACT is a personalized and dynamic process in which a blacklist is learned based on previous access controls decisions.

3. Evaluation

In order to evaluate the performance of REACT, we conducted a user study in order to obtain ground truth access

control decisions to use for learning which is the standard way of evaluating automated access control mechanisms in the literature.

3.1. User Study

We created an application using Facebook Query Language (FQL) and the Facebook Graph API for participants to make access control decisions while disclosing 10 photos. Five of these photos were randomly downloaded from their Facebook profiles, and the participants were asked to select and bring 5 other photos which they had not yet uploaded on Facebook in order to avoid a scenario where a user makes access control decisions for all photos during the study for which they had already received comments and likes before, as that may have influenced their decisions. The participants were advised to bring photos which they considered to be personal (either included them or a family member) or considered sensitive so that they had a privacy implication. The participants logged into the application using their Facebook credentials and were then alerted about the data that would be accessed and asked for explicit permissions before moving on. Each participant was shown 10 photos sequentially on the screen, each on an individual page. They were asked to select categories for the photos from a predefined list of 15 popular Flickr categories, tag any friends in the photo, and make access control decisions for each photo. They were explicitly informed that any friend who was not selected would be denied access to the photo. Once the participants made the access control decisions and selected the categories for all 10 photos, their selections, friend lists and *Total Friends* and *Mutual Friends* profile attributes of all their friends were stored.

3.2. Participants

This research experiment was conducted at Lancaster University and participants were recruited primarily from among the staff and students. Additionally, we invited some participants who were external to the university through personal communication channels. All participants were compensated with £10 for their involvement.

We applied the typical pre- and post-experiment checks to maximize data quality. In particular, before the experiment, we screened participants and everyone who had a Facebook account and had uploaded at least 10 photos before the study was eligible to participate. After an initial registration phase, 31 participants were selected who took part. After completion of the user study, we checked all responses to make sure participants had correctly completed the experiment, finding 5 participants who did not (4 had randomly selected lists of alphabetically sorted friends, 1 had selected one single but different friend for each photo). The remaining 26 participants were considered for the analyses, including 15 males (57.7%) and 11 females (42.3%). The average age of the participants was 29 years (s.d = 6) and the average size of network was 265 friends (s.d = 121). In total, the ground truth dataset obtained during the experiment consisted of **68,840 access control decisions**.

3.3. Implementing REACT

We used Weka to implement REACT for the evaluation. We used 10 fold cross validation on the manually labeled dataset obtained from the user study. In our evaluations, Naive Bayes classification algorithm, Support Vector Machines (SVM) and Random Forest were tried and we found that Random Forest produced the best results. Thus, we do not report results corresponding to Naive Bayes and SVM in this paper due to lack of space.

We used CPM membership to represent **Relationship Type**, implemented using the SNAP library, to create communities for each user from their friend network (which was downloaded during the user study as a list of "nodes" representing each friend and "edges" representing links between them). The CPM algorithm can produce overlapping communities depending on the value of "k" in the implementation [19], i.e., if the value of "k" is kept at 2, it produces non-overlapping communities, but a value greater than 2 produces increasing number of potential overlapping communities. Given that we use community detection to represent relationship type, it is essential to consider overlapping communities as individuals on social media may share more than one relationship type (for e.g: a user's "co-worker" may also be a "family member"), so we changed the value of "k" to check whether creating overlapping communities has any effect on the performance of REACT. The CPM membership is represented as a binary vector of dimension n (where n is the number of communities CPM creates for a user given her social network), in which each element in the vector denotes whether a particular friend belongs to ("1") a particular community or not ("0"). For the example shown in Figure 1, the particular friend shown belongs to community "C1", so it would have a 1 for this community and a 0 for the rest of the communities CPM created from the friend network of the user.

The **Relationship Strength** was represented by *Total Friends, Mutual Friends*, which were both fetched directly from the users' and their friends' profiles, and *Friend Difference* which was calculated as a difference between the *Total Friends* attribute of the friend and the network size of the user. The friend in Figure 1 has has 417 total friends, 24 mutual friends with the user and a friend difference of 112.

For **Content**, we used the tags, provided by the users during the study, about photo categories and friends appearing in the photos. The photo categories selected by the users were coded as a binary vector which denotes whether the particular category was selected ("1") or not ("0") by the user. In Figure 1, the user has selected 3 photo categories, namely, "Cat3", "Cat8" and "Cat13". In addition to the mandatory selection of photo categories, users were also given the option of "tagging" their friends in photos. This "tag" was also a binary variable for each of the user's friends where "1" represents the case where a particular friend was tagged in a particular photo and a "0" represented case where the particular friend was not tagged by the user. For the photo in Figure 1, the user has tagged two friends: "Friend121" and "Friend17".

TABLE 1: Confusion matrix for evaluating performance

		Recommendation	
		Allow	Deny
Access Control Decisions	Allow	TP	FN
	Deny	FP	TN

REACT also benefits from isolating the **Blacklisted Friends** and assigning a default "deny" decision to them. For our evaluation, blacklisted friends were identified by looking at the ground truth access control decisions made by users during the user study and all friends who were never granted access by the user in any of the 10 photos were added to the blacklist. In reality, it may require even less disclosure decisions to identify blacklisted friends for each user. Our data shows that the median percentage of blacklisted friends for a user identified after the first photo is 76.8% and reaches 93.1% after only the fourth photo. This suggests that the blacklist can be learned by REACT very quickly for all users with very few previous access control decisions made by the user.

3.4. Metrics

We evaluate the performance of REACT using the ground truth access control decisions made by users during the user study. In order to compare the actual access control behavior with the recommendations made by REACT, we use several established evaluation metrics for machine learning classifiers [20] to give a broad picture of the performance of REACT. We measure *Specificity (i.e., true negative rate)* as a proportion of "deny" instances (from ground truth) that are correctly recommended as such. *Sensitivity (i.e., true positive rate or recall)* is the proportion of "allow" instances that are correctly recommended as "allow" by the REACT. *Precision* is the proportion of "allow" recommendations which were actually "allow" in the ground truth access control decisions. Finally, *Accuracy* measures the proportion of correct, both "allow" and "deny", recommendations [20].

Table 1 shows the confusion matrix for evaluating the performance of REACT. The described metrics are calculated using the following equations:

$$Specificity = \frac{TN}{TN + FP} \tag{1}$$

$$Sensitivity = \frac{TP}{TP + FN} \tag{2}$$

$$Precision = \frac{TP}{TP + FP} \tag{3}$$

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN} \tag{4}$$

where, TP = True Positives, FP = False Positives, TN = True Negatives, FN = False Negatives

4. Results

4.1. Overall Results

Table 2 shows there is a rather important difference between specificity, sensitivity, and precision; with specificity being much higher than the other two for the mode without any balancing. One of the reasons for this may be that, despite the use of a blacklisting approach by REACT, there is still a difference between the number of friends being recommended allow and deny, which is perfectly consistent with the expected access behavior exhibited by users when confronted with individual access control decisions [4]. In cases like this, the machine learning literature recommends the use of class balancing techniques [21] to produce more balanced results across the three metrics. In particular, and as shown in Table 2, in addition to the normal way of training a classifier (first 4 rows), we considered the three most used and well-known class-balancing techniques that are implemented in Weka: 1) *Class Balancer*, in which synthetic instances of the rarer class ("allow" in our case) are introduced; 2) *Spread Subsampling*, in which the instances from the most common class are randomly removed; and 3) *Cost Sensitive Learning*, which penalizes misclassifications of the rare class over the other.

TABLE 2: Overall results produced by REACT for all metrics, for all k-values of CPM, and for all learning methods

Metric	k-value	Specificity	Sensitivity	Precision	Accuracy
No Balancing	k=2	97.8%	64.4%	74.5%	94.7%
	k=3	97.8%	65.6%	75.3%	94.9%
	k=4	97.8%	64.6%	74.6%	94.8%
	k=5	97.8%	64.7%	74.4%	94.7%
Class Balancer	k=2	97.1%	70.1%	84.8%	92.1%
	k=3	**97.2%**	**72.7%**	**85.7%**	**92.6%**
	k=4	97.2%	70.7%	85.2%	92.3%
	k=5	97.2%	70.9%	85.2%	92.3%
Spread Subsample	k=2	97.8%	64.5%	74.7%	94.8%
	k=3	97.8%	64.8%	74.3%	94.8%
	k=4	97.7%	65.2%	74.2%	94.8%
	k=5	97.8%	64.5%	74.4%	94.7%
Cost Sensitive	k=2	95.0%	75.9%	60.2%	93.2%
	k=3	95.1%	76.6%	61.0%	93.4%
	k=4	95.0%	76.6%	60.6%	93.3%
	k=5	94.7%	77.5%	59.7%	93.2%

TABLE 3: Number of users for whom each attribute subset provided best results

Attribute Subset	No. of Users
Principal Components	8
Information Gain	9
All Attributes	9

Overall, we found that CPM with "k" value of 3 is the most suitable as it provides the best results for all cases and hence it can be said that considering overlapping CPM communities benefits REACT. Looking at all the 4 metrics, it is clear that *Class Balancer* provides the best overall trade-off among the class balancing techniques. In the remainder of the paper, we only report results corresponding to Class Balancer and using $k = 3$ for CPM unless specified otherwise. For this configuration of REACT, the metrics calculated using the entire dataset of 68840 access control decision are: **Specificity = 97.2%, Sensitivity = 72.7%, Precision = 85.7%** and **Accuracy = 92.6%**.

4.2. Personalizing REACT

So far in this paper, we have seen that REACT produces highly accurate access control recommendations according to our evaluations. It uses the attributes discussed earlier in this paper to provide a holistic representation of the social context and content of a disclosure. However, privacy and

TABLE 4: All metrics, calculated over entire dataset, when using best attributes for each user compared to all attributes

	Specificity	Sensitivity	Precision	Accuracy
All Attributes	**97.2%**	72.7%	**85.7%**	92.6%
Best Attributes	96.8%	**77.4%**	84.5%	**93.2%**

TABLE 5: User characteristics to identify suitable subset of attributes to configure REACT

Attribute Subset		T.F	B.R	Audience		Communities		Allow Ratio
				Avg*	SD*	Total	Used	
PCA (8)	**Avg**	214.8	72.1%	13.5	15.2	6.1	5.2	6.1%
	SD	115.2	29%	18.3	19.4	4.3	4.5	6.3%
I. Gain (9)	**Avg**	322.4	61.1%	33.4	38.1	9.8	8.7	14%
	SD	150.8	26.8%	22.7	27.9	6.6	6	15.1%
All (9)	**Avg**	251.6	51.1%	42.9	43.1	9.6	8.9	17.6%
	SD	69.1	27%	32.1	28.3	6.1	6.6	15.4%

T.F - Total Friends
B.R - Blacklist Rate (ratio of friends who are blacklisted)
*Average and Std. deviation calculated across 10 photos for an average user. Significant difference with Kruskal-Wallis Test ($p < 0.05$)

access control behavior is very personal to individual users [13], so there is the potential that a personalized approach would render even better performance than a"one-size-fits-all" solution. In order to personalize REACT, we tried two established attribute selection techniques to identify subsets of attributes that could be used to configure REACT for individual users:

- *Principal Components*: Principal Components Analysis (PCA) is an established method of attribute selection which was used to identify the appropriate attributes for each individual user.
- *Information Gain*: In this approach, we only included those attributes in the classifier which were contributing in terms of information gain.

We tried these attribute selection techniques for each individual user and compared the results with those produced by including all available attributes. We did not focus on accuracy alone and wanted to find the best trade-off between all metrics described earlier for each individual user when selecting the best mechanism to select a subset of attributes.

Table 3 shows that using PCA to select optimal set of attributes produces best results for 8 users while using information gain is best for 9 users and 9 users for whom using the entire set of attributes produces the best trade-off. It is worth noting, however, that there were 17 users for whom all three types of attributes (type and strength of relationship as well as content) were represented in the particular attributes that contributed to the classifier (either PCA or Information Gain) and 9 other users for whom at least two types of attributes were represented which highlights the importance of considering all three types of attributes while designing REACT.

Table 4 shows the results of REACT calculated over the entire dataset of 68840 access control decisions made by all users when: 1) All attributes were used for each of the 26 users, and 2) Only the best subset of attributes was used for each individual user. It is evident from the table that using the best attribute subset provides a better overall trade-off as compared to using all attributes for all users as it provides a big improvement in terms of sensitivity at the expense of very little change for the other metrics.

We also looked at the individual user characteristics in order to try and recommend a particular attribute subset (PCA, Information Gain or All Attributes) for an individual user by looking at their characteristics. Looking at Table 5, we find that PCA can be used to select attributes for users who have smaller friend networks and who select from a smaller section of their network as shown by their higher blacklist rate (ratio of total friends who were blacklisted), lower allow ratio (percentage of total access control decisions which were "allow") and smaller number of communities used.

5. Related Work

REACT uses a conjunction of relationship type, relationship strength and information about the content to provide users with access control recommendations. As we list in Table 6, none of the previous works, to the best of our knowledge, have used a conjunction of all these three types of information to create a holistic solution.

There have been a number of works that proposed using community detection algorithms to facilitate the definition of access control policies [8], [9], [22], [23]. Although community detection could indeed help in learning the access control policies of individual users [23], it does not provide enough goodness of fit with actual access control decisions by users to be used as a single source of information [14].

Another method for access control recommendation has been to use profile attributes [10], [11], [16], [17], [24]. In particular, recommendations in this line of work are based on measures of similarity or closeness between users. Although the closeness and strength of relationships is known to influence access control decisions, it is not the only aspect considered by users when making access control decisions [4], [25].

Finally, previous research showed that tags can be used to create access control policies [18], some methods have been proposed to classify or categorize the nature of the content and leverage this classification to inform access control decisions [12], and the nature of the content was also used to detect cliques in a user's network which can then be used to enhance access control mechanisms [26]. However, the context of a disclosure cannot be completely defined only with the content, and most privacy theories like contextual integrity also recognize the crucial role of the individuals involved and their social relationships [1].

6. Conclusion and Future Work

In this paper, we presented REACT, an access control recommendation mechanism which leverages information representing interpersonal relationships between social media users in conjunction with information about the content. Considering these types of information together enables REACT to represent the overall social context of the information disclosure and make relevant access control recommendations to the users. Our empirical evaluation of REACT shows that it performs very well with respect to all the

TABLE 6: The types of attributes which were considered by previous research

	Relationship Type	Relationship Strength	Content
Amershi et al. [10]	✖	✔	✖
Cheek et al. [9]	✔	✖	✖
Danezis [23]	✔	✖	✖
Fang et al. [8]	✔	✔	✖
Jones et al. [22]	✔	✖	✖
McAuley et al. [11]	✖	✔	✖
Misra et al. [17]	✖	✔	✖
Squicciarini et al. [24]	✖	✔	✔
Squicciarini et al. [12]	✖	✖	✔
Yildiz et al. [26]	✖	✖	✔

metrics including accuracy. We further explore the potential of personalizing REACT by observing that using Principal Component Analysis (PCA) to select the most appropriate attribute set for users who grant access to a smaller section of their friend network may improve results even more. The design of REACT is agnostic to the type of information being shared and can be easily adapted to other types such as text.

An interesting future work may be to investigate the best method of obtaining information about the content. In our evaluation, we used the annotations provided by the users as these are usually considered minimally intrusive [18]. While this is shown to produce good results and the mechanism does not need to know the content itself (which could bear privacy implications) but just metadata about it, other options such as automatic analysis of the content [12] may be more beneficial if the users' effort in annotating content is deemed to be prohibitively large. As REACT is agnostic to the type of information being shared, any such mechanism can be easily adapted to the particular type of content (such as natural language processing for text). Another future work could be extending REACT to consider the cases where one user is not the only one potentially affected by the content to be shared. This could be achieved by implementing conflict detection and resolution mechanisms for when the access control decisions of the users involved differ [27].

References

[1] H. Nissenbaum, "Privacy as contextual integrity," *Washington Law Review*, vol. 79, p. 119, 2004.

[2] G. Hull, H. R. Lipford, and C. Latulipe, "Contextual gaps: Privacy issues on facebook," *Ethics and information technology*, 2011.

[3] M. Johnson, S. Egelman, and S. M. Bellovin, "Facebook and privacy: it's complicated," in *Proc. of SOUPS*. ACM, 2012, p. 9.

[4] S. Kairam, M. Brzozowski, D. Huffaker, and E. Chi, "Talking in circles: selective sharing in google+," in *Proc. of the SIGCHI*. ACM, 2012, pp. 1065–1074.

[5] P. Wisniewski, B. P. Knijnenburg, and H. Richter Lipford, "Profiling facebook users privacy behaviors," in *SOUPS2014 Workshop on Privacy Personas and Segmentation*, 2014.

[6] 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*. ACM, 2012, p. 13.

[7] G. Misra and J. M. Such, "How socially aware are social media privacy controls?" *IEEE Computer*, vol. 49, no. 3, pp. 96–99, 2016.

[8] L. Fang and K. LeFevre, "Privacy wizards for social networking sites," in *Proc. of the 19th international conference on World wide web*. ACM, 2010, pp. 351–360.

[9] G. P. Cheek and M. Shehab, "Human effects of enhanced privacy management models," *Dependable and Secure Computing, IEEE Transactions on*, vol. 11, no. 2, pp. 142–154, 2014.

[10] S. Amershi, J. Fogarty, and D. Weld, "Regroup: Interactive machine learning for on-demand group creation in social networks," in *Proc. of the SIGCHI*. ACM, 2012, pp. 21–30.

[11] J. J. McAuley and J. Leskovec, "Learning to discover social circles in ego networks." in *NIPS*, vol. 272, 2012, pp. 548–556.

[12] A. C. 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*. ACM, 2011, pp. 261–270.

[13] A. Acquisti, L. Brandimarte, and G. Loewenstein, "Privacy and human behavior in the age of information," *Science*, vol. 347, no. 6221, pp. 509–514, 2015.

[14] G. Misra, J. M. Such, and H. Balogun, "Non-sharing communities? an empirical study of community detection for access control decisions," in *IEEE/ACM ASONAM, 2016*. IEEE, 2016, pp. 49–56.

[15] E. Gilbert and K. Karahalios, "Predicting tie strength with social media," in *Proceedings of the SIGCHI conference on human factors in computing systems*. ACM, 2009, pp. 211–220.

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

[17] G. Misra, J. M. Such, and H. Balogun, "Improve: Identifying minimal profile vectors for similarity based access control," in *IEEE TrustCom*. IEEE, 2016, pp. 868–875.

[18] C.-m. A. Yeung, L. Kagal, N. Gibbins, and N. Shadbolt, "Providing access control to online photo albums based on tags and linked data." in *AAAI Spring Symposium: Social Semantic Web: Where Web 2.0 Meets Web 3.0*, 2009, pp. 9–14.

[19] I. Derényi, G. Palla, and T. Vicsek, "Clique percolation in random networks," *Physical review letters*, vol. 94, no. 16, p. 160202, 2005.

[20] G. Calikli, M. Law, A. K. Bandara, A. Russo, L. Dickens, B. A. Price, A. Stuart, M. Levine, and B. Nuseibeh, "Privacy dynamics: Learning privacy norms for social software," in *Proceedings of the 11th International Symposium on Software Engineering for Adaptive and Self-Managing Systems*. ACM, 2016, pp. 47–56.

[21] N. Japkowicz and S. Stephen, "The class imbalance problem: A systematic study," *Intelligent data analysis*, 2002.

[22] S. Jones and E. O'Neill, "Feasibility of structural network clustering for group-based privacy control in social networks," in *Proceedings of the SOUPS*. ACM, 2010, p. 9.

[23] G. Danezis, "Inferring privacy policies for social networking services," in *Proc of the 2nd ACM workshop on Security and artificial intelligence*, 2009.

[24] A. Squicciarini, S. Karumanchi, D. Lin, and N. DeSisto, "Identifying hidden social circles for advanced privacy configuration," *Computers & Security*, vol. 41, pp. 40–51, 2014.

[25] A. E. Marwick *et al.*, "I tweet honestly, i tweet passionately: Twitter users, context collapse, and the imagined audience," *New Media & Society*, vol. 13, no. 1, pp. 114–133, 2011.

[26] 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*. IEEE, 2012, pp. 353–359.

[27] J. M. Such and N. Criado, "Resolving multi-party privacy conflicts in social media," *IEEE Transactions on Knowledge and Data Engineering*, vol. 28, no. 7, pp. 1851–1863, 2016.