

Analyzing and Predicting Privacy Settings in the Social Web

Kaweh Djafari Naini¹(✉), Ismail Sengor Altingovde², Ricardo Kawase¹,
Eelco Herder¹, and Claudia Niederée¹

¹ L3S Research Center, Hannover, Germany
{naini,kawase,herder,niederee}@L3S.de

² Middle East Technical University, Ankara, Turkey
altingovde@ceng.metu.edu.tr

Abstract. Social networks provide a platform for people to connect and share information and moments of their lives. With the increasing engagement of users in such platforms, the volume of personal information that is exposed online grows accordingly. Due to carelessness, unawareness or difficulties in defining adequate privacy settings, private or sensitive information may be exposed to a wider audience than intended or advisable, potentially with serious problems in the private and professional life of a user. Although these causes usually receive public attention when it involves companies' higher managing staff, athletes, politicians or artists, the general public is also subject to these issues. To address this problem, we envision a mechanism that can suggest users the appropriate privacy setting for their posts taking into account their profiles. In this paper, we present a thorough analysis of privacy settings in Facebook posts and evaluate prediction models that can anticipate the desired privacy settings with high accuracy, making use of the users' previous posts and preferences.

Keywords: Facebook · Privacy · Social networks

1 Introduction

Social networking sites such as Google+, Twitter and Facebook allow their users to post updates, tweets, pictures, links and videos to their circles of friends, their followers or to the whole world. By doing so, users generate a digital footprint that defines their 'online presence'.

Similar to the 'offline world', the various digital platforms provide means to define and structure a user's social network. Most of the services support a way to define different groups for information sharing within a user's social network, although each service provides its own implementation and terminology for this. Earlier social networks like Orkut had communities, for example. A member of a community could share posts, pictures and different sorts of information that were only visible to the members. In a similar fashion, current social network

platforms such as Facebook and Google+ provide analogous features with Facebook Groups and Google+ Circles. The common goal is to facilitate the social network users to manage the audience of their interactions and shared content.

The use of such structuring facilities becomes a necessity, because we increasingly share the same social networking sites with persons from different spheres of our real world social networks, such as colleagues, acquaintances, friends and family members. Since each of these groups represent different aspects of our lives, it is often desirable to also be able to maintain this separation in the digital world. For example, certain family affairs might better not be shared with colleagues or acquaintances. Maintaining the right balance of interaction and involvement within those social groups helps us to manage our different roles in life. For this purpose, social networks support their users to manage their groups and to keep control of their privacy settings. Unfortunately, in many cases, these settings are buried in menus, tabs and configurations that are notoriously hard to understand for the regular user [22]. Contributing to this discussion, Facebook founder Mark Zuckerberg claimed that the rise of social networking online means that people no longer have an expectation of privacy, adding ‘we decided that these would be the social norms now and we just went for it’ [13].

As a consequence, in the past years, there have been several cases of people who involuntarily, unknowingly ‘leaked’ information to the wrong audience. Common cases include public messages that were supposed to be privately sent to one particular recipient, and posts that are targeting a specific audience and are in fact publicly available (a recurrent issue on Twitter). The presence of inadequate or inappropriate information about a person in the public sphere can have serious impact, for example on employment opportunities. Exemplary cases are reported in [24]. An indication for the increasing awareness for this topic is the current legal discussion about the *right to be forgotten* now called the *right to erasure* in the European community. This discussion addresses the right of individuals not to be stigmatized as a consequence of a specific action performed in the past [23]. Although there is a distinction between the right to be forgotten and the right for privacy - the right for privacy constitutes information that is not publicly known, whereas the right to be forgotten involves removing information that was publicly known - there is a clear link: if people unintentionally share information to wrong audiences, they might later regret it and want the information to be ‘forgotten’. Ideally, it should be prevented that such information would be unintentionally publicly shared in the first place.

This implies that there is a need for better support for selecting adequate privacy settings in social networks. With that in mind, in this work, we investigate to what extent it is possible to predict the privacy setting of posts. We build our work on top of Facebook’s privacy settings. Facebook is arguably the most popular social network and it provides its users a range of privacy options. In order to understand the users’ privacy behavior, we first provide an analysis of privacy settings for Facebook posts. Subsequently, we present a method for predicting privacy settings by employing classification based on a small but effective set of features that are available at post creation time. Evaluations show that

privacy settings can be predicted with high accuracy, which may allow automatic privacy-setting assistance for the end users and third party apps.

The remainder of this paper is structured as follows. In Section 2, we summarize the relevant related work. In Section 3, we describe our efforts to collect the sensitive data, followed by a data analysis in Section 3.1. In Section 4, we describe the experiments and results towards a privacy prediction method. We finally discuss and conclude our work in Section 5.

2 Related Work

Social media sites, such as Twitter, Facebook and Google+, are designed to share information - and other content, such as pictures, videos and links - with other users. Studies on the usage of social media platforms focus, among others, on usage motivations [14], user behavior [16], and relations and social capital [4,5]. A recent study on Facebook [33] shows that the dynamic and temporal changes of the relationships between users lead to conflicting privacy needs of the user.

Apart from relatively harmless updates, such as sharing a link or other types of public content, messages on Twitter and Facebook may contain highly personal information such as the user's location or email address. For this reason, social media sites typically offer their users several ways for indicating the intended audience of shared messages. First of all, there are *default* settings, which can be adapted by the user. Second, users can overrule these default settings for specific messages. Third, in many cases it is possible to delete, hide or edit a message post hoc.

However, as indicated by several studies (e.g. [20]), users often do not inspect or adapt the default settings offered by the system; thus, most messages are sent with the default settings. Due to this behavior, messages often have a wider audience than intended or expected by the user. According to a recent report from the Pew Internet & American Life Project [21], particularly males and young adults have posted content that they regret. Not surprisingly, these are also the users with the least restricted privacy settings. However, due to the raising awareness of privacy issues and their implications, more and more users actively manage their privacy settings and prune their profiles.

Other studies on Facebook privacy analyze user concerns regarding sharing personal content with a public audience or with third-party applications [9,11]. Similarly, in YouTube it has been observed [17] that users follow different strategies for balancing the pros and cons of sharing with privacy. As an example, users do share videos with private content, but can ensure that their faces are not displayed and their identities are not disclosed. In the context of mobile apps, it is again reported that users' privacy settings are diverse, yet can be represented via a relatively small number of privacy profiles [18,19].

Still, research has shown that users typically disclose more personal information online than they would do in face-to-face situations. There are many risks associated with content that is unknowingly disclosed to the public. Some of these risks - including mobbing, loss of reputation, family problems and lost

career opportunities - are summarized in [25,29]. A remarkable initiative to raise attention for these issues is the site PleaseRobMe¹, which aggregates and shows tweets of users who report to be away from home. In addition, the user is informed via a (public) tweet.

With the goal of raising awareness for the problems related to sensitive information leaks and privacy settings, Kawase et al. [15] introduced FireMe!, a website that contains live streams of people who publicly tweet offensive comments towards their working environment, bosses and coworkers. In their work, they built a system that, once an offensive message was detected in the twittersphere, the author of the offending tweet was sent an alert message. Their results show that only 5% of the users who were alerted by the system later on deleted the compromising tweet. The authors called for the deployment of an alert system that prompts users before a compromising tweet is sent. In fact, our work goes into this direction. By understanding and predicting privacy settings, we might be able to advise (suggest) users the appropriate privacy settings for a given post, before it is effectively out there.

There are other works in the literature that aim to recommend privacy settings. Fang et al. suggest building models that can predict whether a user's friends should be allowed to see certain attributes (such as the birth date or relationship status) in the Facebook profile of the user [7,8]. Similarly, Ghazinour et al. build a classifier to predict the privacy-preference category of a user (such as "pragmatic" or "unconcerned"). Furthermore, they employ a simple kNN approach to determine the similar users to a given user, and based on the preferences of these similar users, they suggest privacy settings, again, for the attributes in the user's Facebook profile [10]. Our work differs from those in that we do not address such general attributes but we aim to recommend a privacy setting for every post (be it a status update, a video link or a photo) made by the user. Machine learning and/or collaborative filtering methods are further employed to recommend privacy settings for the location-sharing services [28,31] and mobile apps [18]. The latter domains involve different dynamics and/or features for setting privacy options than those in Facebook; the social network addressed in our study.

In our approach, we aim to directly support users in choosing privacy settings at the moment that they submit a post. This goal is similar to the work presented in [32], although addressing a different media type. In their work, Zerr et al. propose a method for detecting private photos in Flickr that are posted publicly by extracting a set of visual features. Their results show that a combination of visual and textual features achieves a considerable performance for classifying and ranking private photos. While following a similar goal, we operate in a different setting: we use social network specific features and we aim to predict more fine-granular privacy settings. In our previous work [3], we have also used different types of social network features, but for different prediction tasks, namely for suggesting Facebook posts for content retention and summarization.

¹ <http://pleaserobme.com/>

3 Dataset

In this section we present the two datasets that we used for our experiments. Both datasets have been collected using an experimental Facebook App². This app has been developed in the context of a different work [3], where all the participants authorized us to use their Facebook data (i.e., the content and privacy settings of the posts as well as the basic user profile) for research purposes. Further, to comply with Facebook’s Platform Policies³, we took extra care regarding the participants’ privacy. Most importantly, the data will not be disclosed to third parties and the data collected represent the minimal amount of information needed in order to perform the experiments.

Dataset 1. The first dataset contains 45 users from 10 different countries. The users are all researchers and/or students in the field of computer science from the first authors’ institution. We expect the data from these users to be trustworthy; the users are presumably more knowledgeable in using such digital platforms. From these 45 users, we collected all their posts, summing up to 26,528 posts (posts per user varies from 13 up to 3,176 posts). This dataset has been collected during February and March 2014.

Table 1. Datasets

	Dataset 1	Dataset 2
No. of users	45	649
No. of posts	26,528	769,205
Avg. no. of posts per user	602.431	1,185.215
Variance no. of posts per user	545,343	5,484,176
Min no. of posts per user	13	100
Max no. of posts per user	3,176	30,715

Dataset 2. The second dataset has been collected using the CrowdFlower crowdsourcing platform, where the workers were asked to use the same Facebook app mentioned above. In this case, the authors are not personally familiar with the participants and therefore we have no a priori information on their knowledge of using social networking sites. Especially the former issue raised the concern of reliability, as there could be some workers who use fake profiles to finish the task and get paid. As a remedy, we considered only the data from workers who have a Facebook account that exists for at least 4 years and who have posted at least 25 posts each year. Using the CrowdFlower platform, we ended up with a much larger dataset, including 649 users and 769,205 posts in total. The crowdsourcing task has been running only for one day on November 27, 2014. In Table 1, we summarize the characteristics of both datasets.

² <http://www.l3s.de/~kawase/forgetit/evaluation2015/>

³ <https://developers.facebook.com/policy/>

3.1 Data Analysis

The privacy settings in Facebook regarding the audience of a post can be one of the following five main alternatives:

- **EVERYONE**: This setting means that the post is public. Even non-Facebook users are able to see these posts.
- **SELF**: Only the user who created the post can see it.
- **ALL_FRIENDS**: Posts with this setting are visible to users who are friends with the post creator, and to the friends of those tagged in the post.
- **FRIENDS_OF_FRIENDS**: In addition to the friends, posts with this setting are also visible to friends of the poster’s friends, and to friends of friends of those tagged in the post.
- **CUSTOM**: In this setting the user deliberately specifies a customized privacy setting that includes or excludes specific users or groups from the audience. This option is usually accompanied by the fields *privacy_allow* and/or *privacy_deny*. These fields list users or group ids. CUSTOM includes three sub-values:
 - **ALL_FRIENDS**: Posts with this setting are visible to all friends of the post creator, except to some users or groups that are manually chosen by the creator.
 - **FRIENDS_OF_FRIENDS**: Posts with this setting are visible to all friends of friends of the post creator, except to some users or groups that are manually chosen by the creator.
 - **SOME_FRIENDS**: Posts with this setting are only to specific users or groups that were manually chosen by the post creator.

In Figure 1, we present the distribution of the alternative privacy settings for the posts in Datasets 1 and 2. We can observe some interesting patterns regarding the usage of Facebook privacy settings. First of all, for both datasets, we see a clear dominance of posts that are visible to all of the users’ friends (45 to 50%) and of public posts (around 30%). We also observe a rather high demand for the option that denies access to specific users or groups (around 10%)




Fig. 1. Distribution of the privacy settings for Datasets 1 and 2

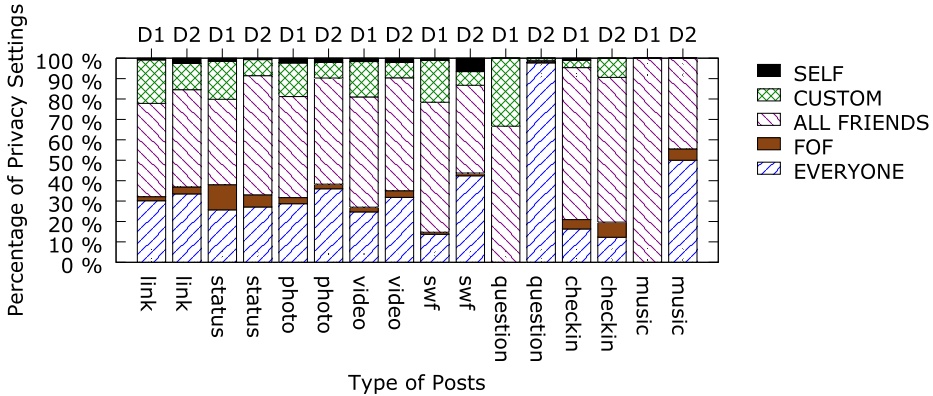


Fig. 2. Distribution of posts normalized by post type for Dataset 1 (D1) and Dataset 2 (D1)

for Dataset 1 and 20% for Datasets 2), which is in line with our expectations from Section 1. This shows that quite often, users carefully ‘hide’ posts from particular users in their social networks.

Next, we provide a more detailed analysis taking into account the types of the posts. In total, there are eight different types that we identified in our datasets. Each type has unique characteristics that may influence the users in choosing the appropriate privacy setting. In Figure 2, we plot the distribution of privacy setting over post types. We see that especially post types of *music* and *Flash content* (shown as *swf*) are the ones that are shared with more general audiences (i.e., EVERYONE and FRIENDS_OF_FRIENDS), whereas post types like *link*, *status*, *video* and *photo* are more likely to be visible to restricted audiences (e.g., with privacy setting CUSTOM). As another interesting observation, almost all the posts of type the *question* (97.79%) are public in Dataset 2 (note that, while the situation is different for Dataset 1, we notice that there are only two posts of type *question* in this dataset, and hence findings are not representative). This is because of the fact that the privacy settings for *Questions* are not directly chosen by the user. In Facebook, *Questions* can only be posted in *Groups* or in *Events*, and the privacy settings are inherited from them. If a question is posted in a public group or in a public event, question is considered public (EVERYONE), and the creator of the post is not given the option to change it.

We also make an analysis of our datasets from the perspective of users. While doing so, we report the findings for Dataset 2, as the number of users is considerably smaller in Dataset 1 (though similar trends are observed for Dataset 1 as well). In Figure 3, we report the percentage of users who have posts with certain combinations of privacy settings. For instance, almost 43% of the users have posts from four different privacy settings, EVERYONE, ALL_FRIENDS, CUSTOM and SELF. Similarly, another 34% of users have posts from all the privacy settings. In general, the distribution in Figure 3 implies that these users




Fig. 3. Distribution of users by their privacy settings combination (for Dataset 2)




Fig. 4. Distribution of the privacy changes by each user (for Dataset 2)

are not unaware of the privacy setting, and indeed they intentionally use different privacy settings for their different posts.

For a deeper insight, we investigate how often users change their privacy settings in different posts. We performed a temporal analysis on the posts of each user to compute how many times she changed her settings; i.e., the number of times a user selects a different privacy setting for her post than that of the preceding post in chronological order. Figure 4 depicts the percentages of such changes for each user as shown on the x-axis. On the average, the users choose different privacy settings in 10.8% of the posts. This further supports our previous finding showing that at least some users deliberately choose different privacy settings. Given the fact that choosing the privacy settings is a task that is frequently triggered, and that the decision is quite often varying, we believe that users could benefit from tools that suggest the appropriate settings. Therefore, in the next section, we present a first step in the direction of predicting privacy settings.

4 Privacy Prediction Experiments

The data analysis in the previous section suggests that there might be dependencies between the privacy settings of a post and some characteristics of the the post or the user who wrote the post. In this section, we investigate whether it is

possible to automatically predict a privacy setting for a post. Such a predictor can be used for recommending the most appropriate privacy setting to the user at the time of posting, and hence help to avoid cases of information leaking as exemplified before.

4.1 Experimental Setup

Target Classes. To build a predictor with reasonable accuracy that can be employed in a practical setting, we opt for building a binary classifier and predicting whether a post has low or high privacy at an abstract level, rather than assigning each post to one of the privacy levels described in Section 3. We assume that posts that have the privacy setting EVERYONE or FRIENDS_OF_FRIENDS are in the class *Low_Privacy*, as they are visible to a very general audience. In contrast, the posts with the setting ALL_FRIENDS, SELF and CUSTOM are said to be in the class *High_Privacy*, as the user has the intention of sharing the post with a specific audience, i.e.; with only her friends, which can be the most typical case in a social platform, or even with a certain subset of them.

Dataset. For our classification experiments, we employ the crowd-sourced dataset (Dataset 2) that includes a reasonably large number of users and postings and, hence, can yield generalizable results. From the latter dataset, we discard all non-English posts using the language detector tool provided by [26], as we aim to construct features based on the post content. Furthermore, for each user, we label the posts as Low and High Privacy; and get all the posts in the class with smaller number of instances, and undersample the posts from the other class. This is to obtain a balanced dataset (as the dataset is otherwise skewed in various ways; some users have a large number of posts, and furthermore, they are biased for a certain privacy class only). At the end, our dataset includes a total of 93,460 posts from 469 users; with an average of approximately 100 posts from each class, per user.

Features. In our experiments, we use features from six different categories (see Table 2). First, we have *metadata* features obtained from a post, such as the type of the post (e.g., link, photo, status, video, etc.), whether the post includes one or more of the predefined Facebook fields (such as message, story or description) and its length, and number of tagged users in the post. The *context* features capture the platform and time related information. From the post content, we first extract *sentiment* features, i.e., the positivity, negativity and objectivity scores computed using a vocabulary based sentiment analysis tool, namely, SentiWordNet [6]. The *keyword* feature category captures whether a post includes a keyword that might be related to a certain concept like family, friends, work, travel, etc. Note that, for each of the latter concepts, we manually compiled a small list (up to 20 words) of representative words. Another feature category is the *word vector*, i.e., the entire content of the post as a bag of words, as typical in text classification. We keep top-1000 most words with the highest tf-idf scores in the word vector. Finally, we have the *user* features, such as the number of posts

Table 2. The list of features used for the privacy prediction task

Feature	Description	Feature	Description
Post metadata		Context	
has(message)	post has a message	sendFromMobile	post sent from an mobile application
length(message)	length of the message	dayTimes	(morning, afternoon, evening, night)
norm(length(message))	length normalized per user	sendAtWeekend	post sent during weekend
has(story)	has a story	Sentiment	
length(story)	length of the story	negative	the negativity score of a post
norm(length(story))	length normalized per user	positive	the positivity score of a post
has(description)	has a description	objective	the objectivity score of a post
length(description)	length of the description	Users	
norm(length(description))	length normalized per user	no_posts	total number of posts of a user
has(link)	post includes a link	no_friends	total number of friends of a user
has(icon)	post has an icon	gender	gender of the user
has(caption)	post has an caption	age	age of the user
type	type of post	country	country of the user
status_type	status_type of a post	education	the education level of the user
icons	describes user activity	Keywords	
tagged users	users tagged in a post	words_family	contains word from the list
Word vector		words_friends	contains word from the list
bag of words	top-1000 words using tf/Idf	words_work	contains word from the list
		words_holiday	contains word from the list
		words_travel	contains word from the list

and friends, gender, age, country and education (the latter is obtained from the crowdsourcing platform). All the features in these categories are concatenated to obtain a single instance vector, i.e., applying the early fusion approach for different types of features (e.g., see [27]). Note that since our predictor is to be employed during the post creation time, it is not possible to use typical social network features based on community feedback (e.g no. of likes, no. of comments etc.) employed in other contexts [2].

Classifiers and Evaluation Metrics. We apply the well-known classification algorithms NaiveBayes[12] as well as a fast decision tree learner, REPTree [30][1]. For both algorithms, we use the implementation provided by the WEKA library⁴. For the evaluation, we use well-known measures from the literature: the true positive rate (TPR), false positive rate (FPR), precision, recall, F-Measure, and area under the ROC curve (AUC). All the reported results are obtained via 5-fold cross-validation. Remarkably, this implies that the posts of a particular user are distributed to training and test sets at each fold; and hence, the model will learn to predict the privacy based on not only other users previous decisions, but the user’s own decisions, as well.

Results and Discussions. In Table 3, we compare the prediction performance for NaiveBayes and RepTree classifiers. The average TPR (i.e., accuracy) of the NaiveBayes predictor is 0.692, which is better than the random baseline with 0.5 accuracy (as we have a balanced dataset). Moreover, when predicting the High Privacy class, the classifier has a higher TPR (i.e., 0.745). This is useful in practice, as predicting a highly private post as public is more dangerous (as these are the cases where the information is exposed to a larger audience

⁴ <http://www.cs.waikato.ac.nz/ml/weka/>

Table 3. Classification results using all the features

Naive Bayes						
TP Rate	FP Rate	Precision	Recall	F-Measure	AUC	Class
0.640	0.255	0.715	0.640	0.675	0.780	LOW_PRIVACY
0.745	0.360	0.674	0.745	0.708	0.780	HIGH_PRIVACY
0.692	0.308	0.694	0.692	0.691	0.780	Avg.
REPTree						
TP Rate	FP Rate	Precision	Recall	F-Measure	AUC	Class
0.810	0.191	0.809	0.810	0.810	0.887	LOW_PRIVACY
0.809	0.190	0.810	0.809	0.809	0.887	HIGH_PRIVACY
0.809	0.191	0.809	0.809	0.809	0.887	Avg.

Table 4. Classification results for each category of features

REPTree						
Feature category	TP Rate	FP Rate	Precision	Recall	F-Measure	AUC
Word vector	0.715	0.285	0.719	0.715	0.714	0.793
Post	0.641	0.358	0.642	0.641	0.640	0.709
Users	0.600	0.400	0.601	0.600	0.598	0.673
Sentiment	0.591	0.408	0.593	0.591	0.588	0.652
Context	0.583	0.417	0.588	0.583	0.577	0.634
Keywords	0.553	0.446	0.553	0.553	0.553	0.592

than intended) than vice versa. The overall performance of the RepTree classifier is even more impressive, as it yields an accuracy of 0.809 for both classes (and, on the average). For this classifier, average F-measure and AUC metrics are also over 0.80. These findings reveal that it is possible to predict the privacy class of a post with good accuracy, and such a predictor can serve in suggesting the privacy setting of a post when it is first created.

Note that, since our dataset includes different numbers of posts from each user (but with the same number of instances from each class), it is also interesting to investigate whether classification performance is biased for the users who have more posts than the average. To this end, we filtered the dataset used in previous experiments, so that each user remaining in the dataset now has exactly 100 posts (50 from each class). In this new setup, the accuracy of the RepTree classifier is still 0.788, which implies that the accuracy can improve with more training instances from a particular user. Nevertheless, even for the case of 100 posts per user, the prediction accuracy is high (note that the scores for the other evaluation metrics are also similar and not reported here for brevity).

Finally, for the RepTree classifier reported in Table 3, we further investigate the performance of each feature category in isolation. Table 4 reveals that

keyword features and word vectors are the least and most useful features, respectively. It is further remarkable that the classifier that use all features in combination perform considerably better than those based on a single feature category.

5 Conclusion

In this paper, we presented an approach for supporting users in selecting adequate privacy settings for their posts. This work is based on a thorough analysis on privacy settings on social networks, particularly in Facebook. Our analysis shows that users customize their privacy settings quite often: for roughly one out of ten posts, a new privacy setting is chosen over time. The data also has shown that the type of post has a significant impact on the the choice of privacy settings. While posts of the type ‘music’ and ‘question’ tend to have a larger (less restricted) audience, ‘status’, ‘photo’ and ‘video’ are more often restricted to a smaller audience.

Targeting a supporting tool that could suggest users preferable privacy settings, we performed experiments for the privacy settings prediction task. By relying on different categories of features that can already be identified at the time of post composition, we were able to achieve a very good prediction performance with a recall and precision of more than 80% on average.

Additionally, our analysis demonstrated clear differences in users’ behavior with respect to privacy settings. We observed that there are some users who are very sloppy regarding privacy settings, having most of their posts publicly available and not changing the settings. We also observed users who very often customized their settings, and users who prefer sharing data mostly with their friends. This difference in behavior indicates that a personalized model for privacy prediction might improve the already good results of the experiments presented in this paper. This is part of our future work, where we plan to collect more contributors (users) willing to collaborate with our research, and on top of a bigger user base, we plan to explore the best personalized methods for privacy settings prediction.

Acknowledgments. This work was partially supported by the K3 project funded by the German Federal Ministry of Education and the European Commission Seventh Framework Program under grant agreement No.600826 for the ForgetIT project.

References

1. Breiman, L.: Bagging predictors. *Machine Learning* **24**(2), 123–140 (1996)
2. Chelaru, S., Orellana-Rodriguez, C., Altingovde, I.S.: How useful is social feedback for learning to rank youtube videos? *World Wide Web* **17**(5), 997–1025 (2014)
3. Djafari Naini, K., Kawase, R., Kanhabua, N., Niederée, C.: Characterizing high-impact features for content retention in social web applications. In: *Proc. of WWW 2014*, pp. 559–560 (2014)

4. Ellison, N.B., Gray, R., Vitak, J., Lampe, C., Fiore, A.T.: Calling all facebook friends: exploring requests for help on facebook. In: Proc. of ICWSM 2013 (2013)
5. Ellison, N.B., Steinfield, C., Lampe, C.: Connection strategies: Social capital implications of facebook-enabled communication practices. *New Media & Society* **13**(6), 873–892 (2011)
6. A. Esuli and F. Sebastiani. Sentiwordnet: a publicly available lexical resource for opinion mining. In: Proc. of LREC 2006, pp. 417–422 (2006)
7. Fang, L., Kim, H., LeFevre, K., Tami, A.: A privacy recommendation wizard for users of social networking sites. In: Proc. of the 17th ACM Conference on Computer and Communications Security, CCS 2010, pp. 630–632 (2010)
8. Fang, L., LeFevre, K.: Privacy wizards for social networking sites. In: Proc. of WWW 2010, pp. 351–360 (2010)
9. Felt, A., Evans, D.: Privacy protection for social networking platforms. In: Proc. of Web 2.0 Security and Privacy (2008)
10. Ghazinour, K., Matwin, S., Sokolova, M.: Monitoring and recommending privacy settings in social networks. In: Proc. of the Joint EDBT/ICDT Workshops, pp. 164–168 (2013)
11. R. Gross and A. Acquisti. Information revelation and privacy in online social networks. In: Proc. of the ACM Workshop on Privacy in the Electronic Society, pp. 71–80 (2005)
12. John, G.H., Langley, P.: Estimating continuous distributions in bayesian classifiers. In: Proc. of the 11th Annual Conference on Uncertainty in Artificial Intelligence (UAI), pp. 338–345. Morgan Kaufmann (1995)
13. Johnson, B.: Privacy no longer a social norm, says facebook founder (2010). <http://www.theguardian.com/technology/2010/jan/11/facebook-privacy/>
14. Joinson, A.N.: Looking at, looking up or keeping up with people?: motives and use of facebook. In: Proc. of CHI 2008 (2008)
15. Kawase, R., Nunes, B.P., Herder, E., Nejdil, W., Casanova, M.A.: Who wants to get fired? In: Proc. of WebSci 2013, pp. 191–194 (2013)
16. Lampe, C., Ellison, N.B., Steinfield, C.: Changes in use and perception of facebook. In: Proc. of CSCW 2008 (2008)
17. Lange, P.G.: Publicly private and privately public: Social networking on youtube. *Journal of Computer-Mediated Communication* **13**(1), 361–380 (2007)
18. Lin, J., Liu, B., Sadeh, N.M., Hong, J.I.: Modeling users' mobile app privacy preferences: restoring usability in a sea of permission settings. In: Proc. of the 10th Symposium on Usable Privacy and Security (SOUPS 2014), pp. 199–212 (2014)
19. B. Liu, J. Lin, and N. M. Sadeh. Reconciling mobile app privacy and usability on smartphones: could user privacy profiles help? In: Proc. of WWW 2014, pp. 201–212 (2014)
20. Liu, Y., Gummadi, K.P., Krishnamurthy, B., Mislove, A.: Analyzing facebook privacy settings: user expectations vs. reality. In: Proc. of the 11th ACM SIGCOMM Conference on Internet Measurement (IMC 2011), pp. 61–70 (2011)
21. Madden, M.: Privacy management on social media sites. Technical report, Pew Internet and American Life Project (2012)
22. Madejski, M., Johnson, M., Bellovin, S.M.: A study of privacy settings errors in an online social network. In: Proc. of PerCom 2012 Workshops, pp. 340–345 (2012)
23. Mantelero, A.: The eu proposal for a general data protection regulation and the roots of the 'right to be forgotten'. *Computer Law & Security Review* **29**, 229–235 (2013)
24. Mayer-Schönberger, V.: Delete - The Virtue of Forgetting in the Digital Age. Morgan Kaufmann Publishers (2009)

25. Rose, C.: The security implications of ubiquitous social media. *International Journal of Management and Information Systems* **15**(1) (2011)
26. Shuyo, N.: Language detection library for java (2010)
27. Snoek, C., Worring, M., Smeulders, A.W.M.: Early versus late fusion in semantic video analysis. In *Proc. of the 13th ACM Int'l Conference on Multimedia*, pp. 399–402 (2005)
28. Toch, E., Sadeh, N.M., Hong, J.I.: Generating default privacy policies for online social networks. In: *Proc. of CHI 2010 (Extended Abstracts Volume)*, pp. 4243–4248 (2010)
29. Toch, E., Wang, Y., Cranor, L.: Personalization and privacy: a survey of privacy risks and remedies in personalization-based systems. *User Modeling and User-Adapted Interaction* **22**, 203–220 (2012)
30. Witten, I.H., Frank, E.: *Data Mining: Practical Machine Learning Tools and Techniques*, 2nd edn. Morgan Kaufmann Publishers Inc., San Francisco (2005)
31. Xie, J., Knijnenburg, B.P., Jin, H.: Location sharing privacy preference: analysis and personalized recommendation. In: *Proc. of the 19th Int'l Conference on Intelligent User Interfaces*, pp. 189–198 (2014)
32. Zerr, S., Siersdorfer, S., Hare, J., Demidova, E.: Privacy-aware image classification and search. In: *Proc. of SIGIR 2012*, pp. 35–44 (2012)
33. Zhao, X., Salehi, N., Naranjit, S., Alwaalan, S., Voida, S., Cosley, D.: The many faces of facebook: experiencing social media as performance, exhibition, and personal archive. In: *Proc. of CHI 2013* (2013)