Privacy Wedges: Area-Based Audience Selection for Social Network Posts

Frederic Raber
DFKI GmbH
Saarland Informatics Campus
Saarland, Saarbrücken
frederic.raber@dfki.de

Alexander De Luca
DFKI GmbH
Saarland Informatics Campus
Saarland, Saarbrücken
alexander.de.luca@ifi.lmu.de

Moritz Graus
Saarland University
Saarland Informatics Campus
Saarland, Saarbrücken
s9mograu@stud.uni-saarland.de

ABSTRACT

We present Privacy Wedges, a user interface designed to allow users of online social networks to make meaningful decisions on who to share their posts with. By displaying the privacy settings for historical posts, it is possible to visualize them in a meaningful and comprehensive way. We conducted a user study with 26 participants that showed that unwanted disclosure could be significantly reduced compared to the current implementation of Facebook. That is, there were significantly fewer posts shown to friends they were not appropriate for or intended for.

1. INTRODUCTION

Online social networks like Facebook or Google+ hold a plethora of private information shared by their users. In many cases, this information is intended for a limited audience but actually shared with all friends of a user or even "everyone" [3]. This happens because often, users are not aware of who they are sharing their posts with, or the burden of dealing with complex privacy settings prevents them from (correctly) applying them [7].

As a big step towards a solution for oversharing in online social networks, we created Privacy Wedges, a user interface that enables users to easily define the audience they want to share a certain post with. The interface displays friends in wedges around the center (the user). The visualization is based on interpersonal distance to the respective friends, that is, less well-known friends are arranged further from the center in the interface.

To find out whether Privacy Wedges fulfills the requirement of reducing unwanted disclosure, we conducted a study with 26 participants comparing Privacy Wedges to the standard Facebook user interface for selecting recipients for a post.

The results of the study show that Privacy Wedges indeed significantly reduced false positives (FP) rates for sensitive social network posts when compared to Facebook's interface.

Symposium on Usable Privacy and Security (SOUPS) 2016, June 22–24, 2016, Denver, Colorado.

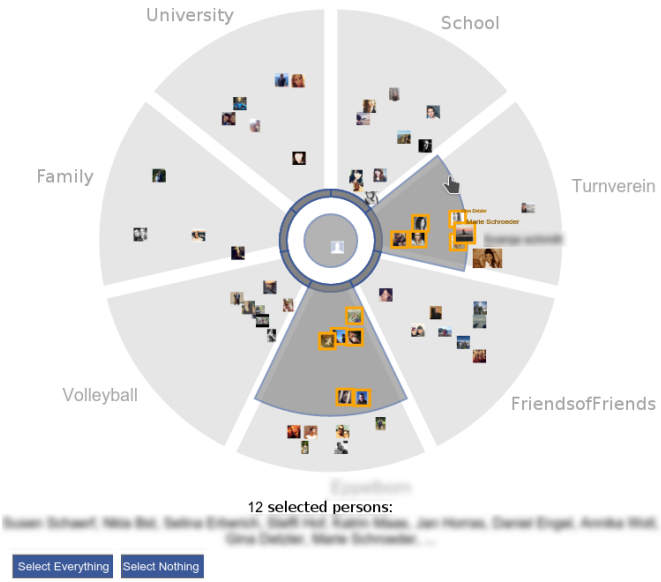

Figure 1: Main interface of Privacy Wedges: Friends are arranged in groups around the center depending on the interpersonal distance to the user. Selected friends are highlighted using a yellow border.

Summed up, this work contributes to the field of online social network privacy by providing a novel user interface for privacy-respectful posting on social networks, backed up by a preliminary user study that shows that the interface supports users in making better decisions on who to share their posts with.

2. RELATED WORK

In recent years, researchers have invested quite some effort to provide advanced user interfaces that support the users of online social networks in reducing (unwanted) information sharing.

A very simple yet efficient approach was presented by Wang et al. [8, 9]. Their so-called "privacy nudges" show examples of who the users are currently sharing a post with based on randomly selected photos of their friends. The users then have a few seconds to undo their post, that is, it is not posted. The results of their study showed that this simple nudge indeed decreased the amount of shared information.

However, their approach does not solve the problem that the respective post might have been appropriate for some of the users' friends, just not all of them.

Studies in the first days of Google+ [5] discovered that users select their audience according to the aspect of their life they belong to, as well as the tie strength of each of them. Kauer et al. [6] used this insight in their user interface for an audience selection approach based on interpersonal distance as a criterion: The interface shows all friends ordered by tie strength according to Xiang et al. [10] from left to right in a rectangle. The users grab and drag a slider from left to right in order to select friends to disclose the post to. Related to this, Reinhardt et al. [7] analyze interactions between users on social networks to determine how close they are to one another. Based on this ordering, they propose a privacy suggestions user interface shown next to the posting dialogue. Similarly, Goncalves et al. discuss a technique called "Narrowcasting", and present an interface that shows the users' friends grouped into different demographic categories and then share with this category only [2].

3. PRIVACY WEDGES

Privacy Wedges offers a graphical user interface, which allows audience selection based on interpersonal distance, for different groups. Figure 1 shows an example of audience selection in the UI. The interface contains the profile picture for each friend in the user's friend list, later denoted as "friend picture". Each of the user's friend groups is represented by a wedge in the UI.

The friend pictures are aligned around the center according to the *tie strength* between the friend and the user. The higher the tie strength to a friend, the closer the friend image is placed to the center. This rule holds for *all* friends in the UI, not only for each set of friends inside a wedge. Apart from the fixed distance to the center according to tie strength, the friend images have been arranged so that they do not overlap each other. The current implementation of Privacy Wedges does not yet incude a tie strength calculation. For the experiment, we let the participants create a friend list, grouped into friend groups and ordered by tie-strength.

Initially, no friend is selected. The user clicks and drags from the center of a wedge to the outer rim to select a subset of friends as recipients for a post. The selected area is colored grayish (see Figure 1). All friends which are inside the selected area of the wedge (from here on called "wedge area") will receive the post. The friend images within a wedge are highlighted with a yellow border. If the user wants to select all friends up to a certain tie strength, she clicks and drags the central circle in the graph to expand its radius. Apart from dragging wedges, Privacy Wedges supports the selection of single friends by clicking on the respective image.

All friend images start with the same size. When the user starts to hover over them, or drags the rim of the wedge area over a friend image, the image is magnified and the name of the friend is displayed below the image (see Figure 2). When the rim is dragged further away of the friend image, it is again miniaturized and the friend name is hidden. This additional feature highlights the friends that are currently on the cusp of being selected and helps the user to identify them.

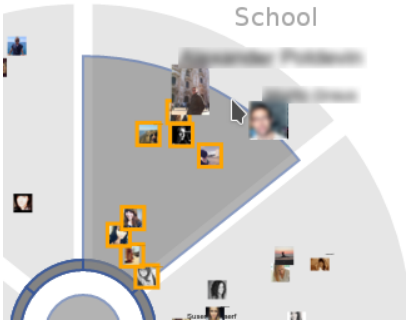

Figure 2: Magnifying effect when dragging a wedge area over friend images.

The number of selected friends and a list of the ten closest friends that are selected is displayed below the graph. In the bottom left corner, two buttons for selecting all friends and to reset the selection are placed, as shown in Figure 1.

All features of Privacy Wedges will be discussed in terms of typical use cases in a sharing situation in the following paragraphs. The different use cases and corresponding actions are depicted in Figure 3.

Select all friends up to a given tie strength (A):
Click and drag the central circle to expand its radius to the desired size. All friends inside the circle are selected as post recipients.
Typical use case: Only the friends up to a specific tie strength should receive the post, independent of the friend group. Example: A family visit to an amusement park. Although the information is suitable for all groups of friends, it is private information that might not be of interest for very distant friends.

Select friends of one or more friend groups up to a given tie strength (B):
Click and drag the inner rim of a wedge to create a wedge area and expand its size. As with (A), friends that are covered by the dragged wedge area are selected to receive the post. This process can be repeated with all available wedges. (A) and (B) can be combined.
Typical use case: Only the friends of one or several friend groups, up to a specific tie strength, should receive the post. Example: Pictures of a party at the university. The post is only suitable for a subset of the friends (most probably the friend group of fellow students). Additionally, the user might feel uncomfortable including distant fellow students in the recipients, and selects only friends up to a certain tie strength.

Exclude friends up to a given tie strength (C) and (D):
After a wedge area has been created in step (B), the user can click and drag the inner rim of the wedge area created in (B) to shrink it, excluding the inner friends of the wedge. As the shape of the wedge area becomes sickle-like by this process, this mode is later called "sickle mode"; the selected areas are denoted as "sickles".

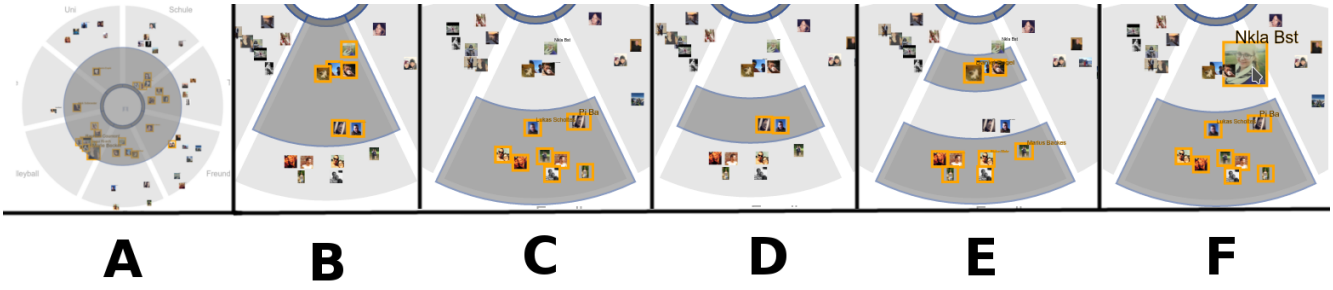

Figure 3: Actions for the different use cases in Privacy Wedges.

Typical use case: Only friends with intermediate tie strength of a friend group should receive the post. This is useful, for example, if you need feedback about something which is potentially embarrassing. Example: A rock band member asks for some nice techno events. He might not want his best friends to know of his "special" hobby, but still needs some friends who know him well in order to receive good recommendations.

Select multiple areas inside a friend group (E):
Following (C)/(D), the processes (B) and (C)/(D) can be repeated in order to create a second sickle/wedge area. This process can be repeated up to four times per wedge.
Typical use case: A clique of friends inside a friend group should not receive the post. Example: You post a picture of fireworks that you bought in a foreign country. A sub group of your friends would oppose such actions, and you want to exclude them.

Select/deselect single friends (F):
Independent of the selections in (A)-(E), a click on a profile picture selects or deselects the user from receiving the post
Typical use case: A single friend has to be added to or removed from the audience due to special reasons. Example: A user posts a picture of her new partner on Facebook. To avoid conflicts, she wants to exclude her ex-partner from the audience.

Select all, select nothing
Two buttons to select all or no friends can be found in the lower left corner of the UI (see Figure 1).
Typical use case: A post should be sent to all friends, or the user wants to reset the selection.

Privacy Wedges was implemented as a website. Its backend is based on Django, which is a web framework on top of Python 3. The interactive frontend is written in JavaScript. The required data, such as friend lists and tie strength, is stored in a mySQL database.

4. USER STUDY

We evaluated Privacy Wedges in a comparative study. Since Facebook is still the most frequently used online social network, we decided to use a mockup of the Facebook privacy settings dialog as a baseline condition, later referred to as the "reference interface". To be precise, we used the detailed privacy settings dialog, which can be accessed when clicking on "Custom Settings" in the recipient box of Facebook's posting dialog. The interface offers the user two different text boxes, which can be filled with one or more friends and friend groups. When the user clicks on one of the two text boxes, a dropdown list opens, containing different suggestions. Without additional input, only "friends" and "friends of friends" are suggested. After the user starts typing, friends and group of friends that match the entered text are suggested.

4.1 Methodology

We compared Privacy Wedges to the conventional Facebook interface with respect to the following parameters in a repeated measures lab study: performance, user strategies, and the influence of the amount of friends shown in the interface on the aforementioned factors.

Each participant used Privacy Wedges as well as the Facebook interface and carried out each condition in two phases: one phase using explicit selection tasks (for example "select the six university friends to whom you feel closest"), later referred to as the "explicit task phase", and a second phase where the participants were given a post and had to decide themselves with whom to share, later called the "complex task phase". We tested the pragmatic and hedonic quality of each interface using the AttrakDiff questionnaire [4] after each condition.

The order of conditions as well as the set of tasks was randomized in order to minimize training effects. We made a recording of the screen as well as an audio recording to capture the interaction behavior with the interface and the verbal comments the participants gave while using the system. The interface stored the results of friend selection as well as the interaction times in a log file.

4.2 Procedure

The experiment started a few days before the first meeting. The participants were asked to hand in a list of their 60 closest Facebook friends, sorted by tie strength. Each of the friends had to be assigned to a friend group. The examiner fed the list into the database of both interfaces, and added the corresponding Facebook profile pictures to the Privacy Wedges interface.

In the lab, the participants first filled out an initial questionnaire regarding demographic data. The examiner then explained both interfaces and their functionalities, and gave the participant the chance to try out the interfaces. After a Q&A session, the subjects started with the first condition, and worked through both task phases.

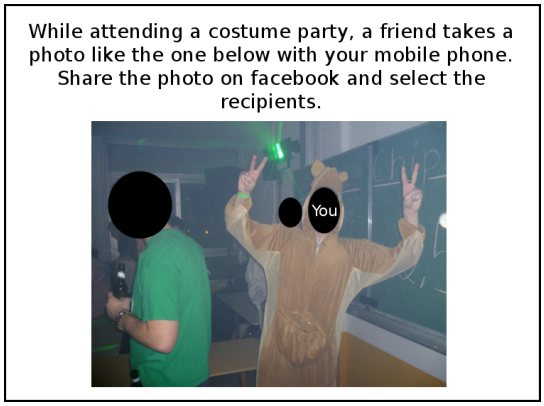

Figure 4: Example user study task with supplementary picture material.

As stated above, the first five tasks gave a definite description of which friends had to be selected, whereas the second six tasks consisted of hypothetical posts or picture examples, for which the participants had to determine the audience themselves. An example of a task description is given in Figure 4. In each task, the participant was first given the task description on the screen in textual form (for some tasks supplemented by the picture to post) and had to click "Next" to proceed to the selection interface.

The number of friend images was fixed to 50 persons for the explicit task phase, as this was the maximum amount of friends that users stated to be displayable while still being able to use the UI. As stated above, we also wanted to test on the influence of the number of friends on the correctness of the friend selection in the complex task phase. We started with 30 friend images, and increased this amount by increments of 5 friend images up to 55 friend images in this task phase.

When the last task of the condition was completed, we let the participants fill out the respective AttrakDiff questionnaire. This procedure was repeated for the second condition.

The participants were then invited to a second, final meeting seven days after the main experiment. We chose this timespan as we assumed that the participants' privacy preferences would not change significantly during this period of time. On the other hand, it was long enough to assume that they had forgotten which friends exactly they selected in the main experiment.

The meeting was used to identify the "ground truth" for each task. We went through the list of friends for each task in the main study and asked, for each friend in the list, whether he/she should receive the respective post. By this method, we were able to determine the desired privacy setting to which we compared the selected audience with the Facebook setting or Privacy Wedges. This part was followed by a questionnaire which asked for the criticality of a false positive (person received the post but should not have) and a false negative (person who should have received a post did not) on a ten-point scale (1=completely uncritical to 10=highly critical).

4.3 Participants

We recruited 26 volunteers for the main study. The set of participants formed a convenience sample; they were recruited from among students and the examiner's acquaintances. Despite this, the age range covered a large interval from 18 to 54 years (mean: 24) and the participants had a variety of occupations, from students of different majors to engineers, consultants and persons working in the health or finance sector. Nine participants were female, 15 male. At the time of the study, each participant had been an active Facebook user for at least five years and a maximum of twelve years (mean: 8.76). On average, each user was member in 1.61 other additional social networks like Google+ or Twitter.

5. RESULTS

We will first explain the terminology for the data that was recorded during the experiment, and later present the measured results, followed by a discussion.

5.1 Terminology

The recorded data can be divided into two parts: First, the quantitative data that was logged by the interface, and second, the qualitative data resulting out of the questionnaire after each condition.

5.1.1 Quantitative data

We combined the list of intended shares of the second interview with the data logged by the UI in order to derive the following performance measure:

False Positives and False Negatives
For each task, the UI recorded the selected friends. We compared the list of selected friends with the list of intended recipients from the second interview. Friends that were in the list of selected but not in the list of intended friends were counted as false positives (FP) whereas friends were in the list of intended recipients but were not selected, were counted as false negatives (FN).

5.1.2 Qualitative data

Qualitative data was collected after each task set was finished: We gave the participants the AttrakDiff questionnaire in a paper version.

We asked the participants to judge the *severity* of a false positive and a false negative on a scale from 1 to 10 (1=very high severity, 10=very low severity).

5.2 Quantitative results

We conducted a paired t-test for the performance measure, namely the false positives and false negatives. To ensure the normal distribution of the samples, an important precondition for a t-test, we performed an F-test for each of the three samples beforehand. All F-tests were positive. As correctness of the answers can depend on the task type (explicit or complex), we analyzed the samples of the explicit and the complex task phases separately. The results of the statistics are shown in Table 1.

The false positives were significantly lower for Privacy Wedges in the complex task phase ($M_{wedges} = 1.92$; $M_{facebook} = 3.47$; $T = 2.26$; $p = .025$), and highly significantly lower in

Phase	Wedges		Facebook			
Measure	*Mean*	*SD*	*Mean*	*SD*	*T*	*p*
Explicit						
FP	.14	.43	1.25	2.81	4.50	$< .001$
FN	.15	.44	1.59	3.92	4.62	$< .001$
Complex						
FP	1.92	5.31	3.47	7.56	2.26	.025
FN	7.08	10.70	4.53	8.38	2.28	.024

Table 1: Results of the paired t-test for the performance measures.

the explicit task phase ($M_{wedges} = 0.14$; $M_{facebook} = 2.81$; $T = 4.50$; $p < .001$).

In addition, the false negatives were significantly different in the complex phase ($M_{wedges} = 7.08$; $M_{facebook} = 4.53$; $T = 2.28$; $p = .024$), but this time in favor of the reference interface. Privacy Wedges outperformed the Facebook interface again in the explicit phase with high significance ($M_{wedges} = 0.15$; $M_{facebook} = 1.59$; $T = 4.62$; $p < .001$).

As mentioned previously, we increased the number of friend images in the complex task phase for both interfaces. The mean values of false positives and false negatives for the different numbers of friend images is shown in Figure 6 in the appendix. We performed a univariate ANOVA to check for a significant trend. The results in Table 2 show that, for an increasing number of friend images, false negatives significantly increase with Privacy Wedges ($F_{5,150} = 2.30, p = .048$) whereas in the reference interface, the amount of false positives increases significantly ($F_{5,150} = 4.40, p = .001$). False positives did not increase significantly for Privacy Wedges, nor did false negatives for the reference interface.

Scale	*F*	*p*
Wedges		
FP	1.32	0.258
FN	2.30	0.048
Facebook		
FP	4.40	0.001
FN	0.63	0.68

Table 2: Statistical results for FP and FN depending on the amount of friend images.

5.3 Qualitative results

Privacy Wedges outperformed the Facebook interface in both hedonic and pragmatic quality, as depicted in Figure 5. The answers to the different questions can be found in Figure 7 in the appendix.

Privacy Wedges received a positive score for all value pairs except for "undemanding - challenging" of the HQ-S scale. Other value pairs that could be improved are "cautious - bold" and "unprofessional - pofessional" of the HQ-I scale.

The AttrakDiff questionnaire was evaluated with an ANOVA using the four scales (PQ, HQ-S, HQ-I, ATT) as goal variables and the two conditions as factors. As shown in Table 3, Privacy Wedges outperformed the reference interface in all four dimensions with high significance.

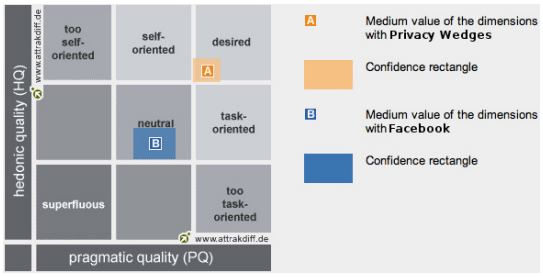

Figure 5: Hedonic and pragmatic quality for the Privacy Wedges and Facebook interfaces.

Scale	Mean		*F*	*p*
	Wedges	*Facebook*		
PQ	5.37	4.13	14.51	< 0.001
HQ-I	5.27	3.80	29.88	< 0.001
HQ-S	5.33	3.14	65.07	< 0.001
ATT	5.64	3.64	52.68	< 0.001

Table 3: Statistical results for the AttrakDiff questionnaire.

The mean values for a false negative are higher ($M_{FN} = 6.03, SD_{FP} = 3.06$) than the means of the false positives ($M_{FP} = 2.48, SD_{FP} = 1.97$), indicating a higher severity of false positives.

6. DISCUSSION

6.1 Performance

The main study results showed that for both the explicit and the complex tasks, users created significantly fewer false positives with Privacy Wedges, and a significantly higher number of false negatives were generated within the complex task phase. Furthermore, users tend to cause more false negatives with increasing numbers of friends in the interface, whereas the reference interface leads to more false positives.

These two factors lead to the assumption that, if errors are made with Privacy Wedges, the user tends to forget to include friends and selects an audience which is smaller than intended. With the current Facebook interface, a bigger audience than wanted is selected. Having a look at the perceived severity of false positives and false negatives, this gives Privacy Wedges an advantage over the reference interface, as the participants in our study were significantly more likely to accept false negatives, that is, missing recipients, than false positives.

6.2 User Experience

The AttrakDiff results clearly show that Privacy Wedges outperformed the reference interface in both pragmatic and hedonic quality. All four dimensions of the questionnaire had significantly different results, in favor of the proposed interface. The detailed results of the word pairs in Figure 7 contain two salient word pairs where Privacy Wedges still has potential to improve regarding the user stimulation (HQ-S): Users experience the interface as rather unchallenging and cautious. We are not sure whether these adjectives have a negative connotation in the application field of security

and privacy, but future research should keep these factors in mind when designing a UI like ours.

6.3 Application to historical posts

Privacy Wedges was designed to select the audience for a new social network post. Nevertheless, the same design can also be used to visualise and review the audience for historical user posts that have been published in the past, even if a different user interface was used for the selection: Selected users are initially selected as single friend selections (action F in Figure 3). If the UI detects two or more selected friends next to each other inside the same wedge, a selected area is created around them (actions B to E). By that means, Privacy Wedges is able to construct a selection that corresponds to the historical sharing setting using the wedge-based UI. We would like to examine on the usefulness and correctness of such a functionality in future research.

6.4 Limitations

This work has two main limitations: 1) the experimental setup with respect to the validity of our sampling and the number of participants; and 2) the scalability of our approach.

6.4.1 Experimental setup

As mentioned before, we relied on convenience sampling for our study, mainly to simplify the process of getting the participants' real friends data. Convenience sampling typically has the disadvantage of reducing generalizability. However, we argue that our sample was quite diverse for a convenience sample, and that our results are still valid to produce valuable insight into how the interface is used and whether it is able to reduce false positive rates when posting in online social networks.

The number of participants was mainly chosen to get a useful amount of data to produce valuable insights into the interface and at the same time to keep the workload managable. For instance, due to the lack of a tie strength calculation in the current implementation of Privacy Wedges, we let the users do the ordering with respect to tie strength for the experiment, and had to manually check the results. For future versions, we are planning to implement the tie strength calculation by Reinhardt et al. [7].

6.4.2 Scalability

We discovered in the pilot study that the user interface of Privacy Wedges is limited to a certain amount of friend images that can be shown while still maintaining the functionality of the UI. As visualized in Figure 6, the number of false negatives increases linearly with the number of displayed friends. In other words, with an increasing amount of displayed friends, the users tend to overlook more friends and share the post with a smaller audience than intended.

In contrast, the Facebook interface has an increasing false positive rate when more friends are added. This means users tend to forget to exclude friends and share the post with a bigger audience than intended.

The qualitative results of the main study indicate that false positives are perceived as more severe on average than false negatives, which favors the behavior of Privacy Wedges given the tested number of friends. Nevertheless we can only see a rough trend within the experimental results, and cannot foresee the amount of false negatives for a larger amount of friends, as is common in a real Facebook account. Future work should further explore this issue.

The proposed approach also suffers from an additional scalability issue due to the limited space available to display friend images. We already implemented a magnifying functionality to cope with the problem. It allows us to overlap friend images initially. We enlarge the friend images that are of interest when a selection is done, and foreground those images that are on the margin of being selected.

We propose the following two features, which can reduce the space problem within Privacy Wedges:

First, we further improve the clarity of the interface by grouping together clusters of friends. The clusters can be formed by friends which have a high tie-strength between each other or which are almost always selected or deselected together.

Second, according to Christakis and Fowler [1], only a small subset of Facebook friends are real friends that are of importance when sharing a post. The interface can therefore concentrate on displaying of the portion of friends with highest tie strength. All additional friends can be placed at the outer rim of the wedges as other friends. Similarly, the size of the friend images inside a wedge can be changed. At the moment, all friend images have the same size, unless a selection is done ("magnifying functionality"). In addition to this, we can scale the friend images according to their tie strength. The closest friends in the center of the wedge are displayed relatively large, whereas the size decreases with a decreasing tie strength. Finally, the friends that are most distant are only denoted as small dots on the outer rim of a wedge. Still, if the user hovers over them during a selection process, these dots are enlarged so that the friend images can be recognized.

7. CONCLUSION

In this paper, we presented Privacy Wedges, a user interface to support privacy-respectful posting on online social networks. It was designed to reduce the amount of unwanted data sharing, i.e. sharing a post with an audience it was not intended for.

Privacy Wedges indeed significantly reduced the amounts of wrong recipients at the cost of increased false negative rates, meaning that some friends who should have been part of the recipient list did not receive a post. Based on our qualitative findings, this was behavior that was highly favored by the study participants over creating false positives, i.e. sharing a post with friends who should not have received it.

By visualizing the privacy settings that have been made using the facebook interface in the past, Privacy Wedges can be used to visualize these settings and indicate the privacy of the user for the already published posts.

8. APPENDIX

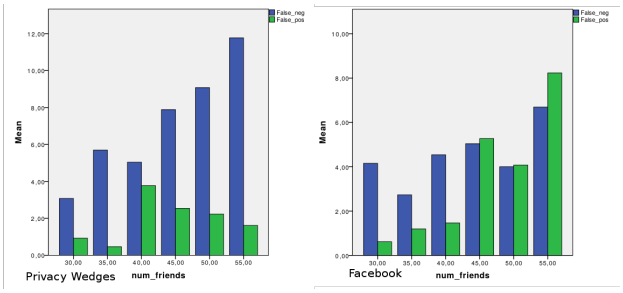

Figure 6: Mean values for FP and FN depending on the number of friend images.

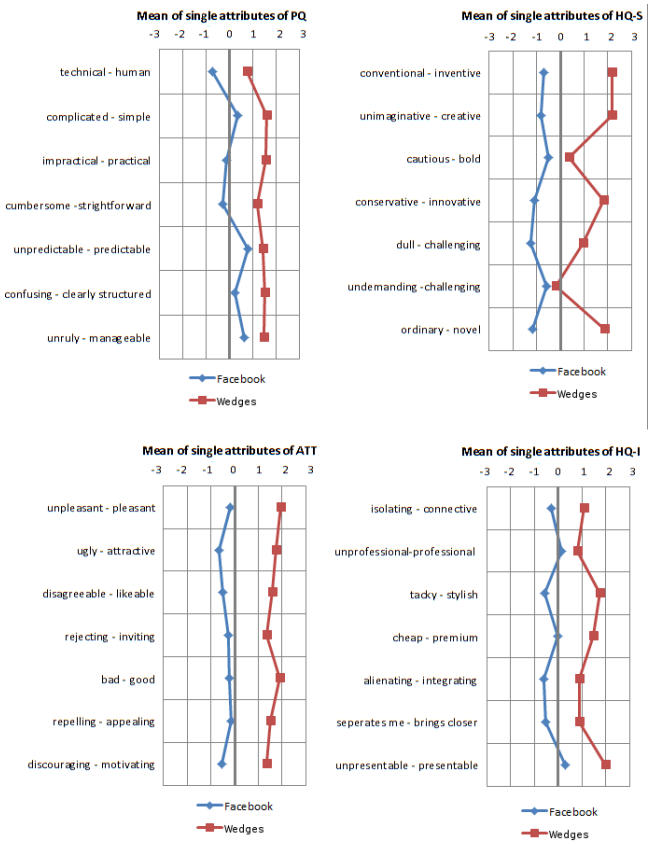

Figure 7: Participants' answers to the AttrakDiff questionnaire.

9. REFERENCES

[1] N. A. Christakis and J. H. Fowler. *Connected: The Surprising Power of Our Social Networks and How They Shape Our Lives*. Back Bay Books, 2011.

[2] J. Goncalves, V. Kostakos, and J. Venkatanathan. Narrowcasting in social media: Effects and perceptions. In *Advances in Social Networks Analysis and Mining (ASONAM), 2013 IEEE/ACM International Conference on*, pages 502–509, Aug 2013.

[3] 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*, WPES '05, pages 71–80, New York, NY, USA, 2005. ACM.

[4] M. Hassenzahl, M. Burmester, and F. Koller. Attrakdiff: Ein fragebogen zur messung wahrgenommener hedonischer und pragmatischer qualitaet. In G. Szwillus and J. Ziegler, editors, *Mensch & Computer 2003: Interaktion in Bewegung*, pages 187–196, Stuttgart, 2003. B. G. Teubner.

[5] S. Kairam, M. J. Brzozowski, D. Huffaker, and E. H. Chi. Talking in circles: Selective sharing in google+. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '12)*, pages 1065–1074, New York, NY, 2012.

[6] M. Kauer, B. Franz, T. Pfeiffer, M. Heine, and D. Christin. Improving privacy settings for facebook by using interpersonal distance as criterion. In *CHI '13 Extended Abstracts on Human Factors in Computing Systems*, pages 793–798, New York, NY, USA, 2013.

[7] D. Reinhardt, F. Engelmann, and M. Hollick. Can i help you setting your privacy? a survey-based exploration of users' attitudes towards privacy suggestions. In *Proceedings of the 13th ACM International Conference on Advances in Mobile Computing and Multimedia (MoMM)*, 2015.

[8] Y. Wang, P. G. Leon, A. Acquisti, L. F. Cranor, A. Forget, and N. Sadeh. A field trial of privacy nudges for facebook. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '14, pages 2367–2376, New York, NY, USA, 2014. ACM.

[9] Y. Wang, P. G. Leon, K. Scott, X. Chen, A. Acquisti, and L. F. Cranor. Privacy nudges for social media: An exploratory facebook study. In *Proceedings of the 22Nd International Conference on World Wide Web*, WWW '13 Companion, pages 763–770, New York, NY, USA, 2013. ACM.

[10] 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*, WWW '10, pages 981–990, New York, NY, USA, 2010. ACM.