

The promise of audience transparency. Exploring users' perceptions and behaviors towards visualizations of networked audiences on Facebook



Ralf De Wolf^{a,*}, Bo Gao^b, Bettina Berendt^b, Jo Pierson^a

^a Department of Media and Communication Studies, Vrije Universiteit Brussel, Brussels, Belgium

^b Department of Computer Science, KU Leuven, Heverlee, Belgium

ARTICLE INFO

Article history:

Received 26 February 2015

Accepted 10 April 2015

Available online 21 April 2015

Keywords:

Audience
Visualization
Affordances
Facebook
Evaluation

ABSTRACT

The presence of multiple audiences and the collapse of boundaries between them in Facebook make it difficult for users to know and to control who has access to their online contributions. Previous research has shown how visualizations of Facebook friends are useful, but mainly focused on the instrumental goal of controlling access. It is unclear, however, what value users themselves see in visualizations and whether knowledge and/or control are important to them. In this research, these questions were studied by evaluating FreeBu, a semi-automatic and interactive grouping technology that visualizes Facebook friends. The results indicate that audience visualizations are especially perceived useful for grouping and reflection purposes. Moreover, we found how users are attracted to larger groups, those with whom they communicate more, bridges and outliers in their network. The combined findings suggest that awareness is considered at least as important as control. Therefore, a shift from audience control to audience transparency is recommended.

© 2015 Elsevier Ltd. All rights reserved.

1. Introduction

Social network sites (SNSs) are widespread and adopted by diverse populations throughout the world. The online media environment increasingly becomes self-evident, and provides for new and open ways in communicating with others. The presence and collapse of audiences in SNSs, however, is also challenging for users (boyd, 2008; Marwick and boyd, 2010; Vitak, 2012). They are often unaware of employers scanning their profiles (Hargittai and Litt, 2013) and underestimate the size of their audience (Bernstein et al., 2013). Service providers, therefore, have invested in options that allow users to better control their personal information flow and audiences (e.g. Google circles, Facebook lists and groups). Audience visualizations have not been integrated in SNSs so far, even though visualizations have been proven useful for users to comprehend their online relationships (Lipford et al., 2008; Egelman et al., 2011; Wang et al., 2011, 2013; Mazzia et al., 2012).

In this research we contribute to the literature by studying the perceptions of users towards audience visualizations and which interface cues draw their attention to specific parts or aspects of audiences. Specifically, FreeBu, a semi-automatic and interactive grouping technology that visualizes Facebook friends is studied from a user perspective. In a first user study, a first version of FreeBu and audience visualization was studied. In a second user study, a second version of FreeBu – with three more audience visualizations – was researched. These visualizations are further described in Section 4.

* Corresponding author. Tel.: +32 2 629 16 33.

Our paper is organized as follows. We begin with outlining the problematic nature of networked audiences in SNSs and provide an overview of the proposed solutions and users' practices. Next, we give an overview of the mixed method research design, and describe the population of study, materials and procedure. In Section 6, the affordances of audience visualizations are discussed as well as how the users perceived and interacted with the different visualizations. Finally, in the discussion, we generalize our findings into recommendations for the further design of audience visualization technology.

2. Theoretical considerations

2.1. Networked audiences

Goffman (1959) described an individual's actions in everyday life as a performance on stage. A performance can be seen as "all activity of a given participant on a given occasion which serves to influence in any way any of the other participants" (p. 26). Goffman (1959) underlined the importance of knowing one's audience when presenting the self. On SNSs users have to imagine their audience, because they are physically absent. Moreover, the audience only becomes visible when others "like" or "comment" on one's performance. Research has indicated that the imagined audience can have a strong influence on behavior, just as the presence of the actual audience (Litt, 2012).

The situation on SNSs, especially, is challenging because of the exacerbation of context collapse and the co-presence of multiple audiences (Vitak, 2012). SNSs, such as Facebook, articulate one's offline network and not just one segment (boyd and Ellison, 2007; Ellison and boyd, 2013). Hence, users have to imagine multiple audiences from diverse contexts. Marwick and boyd (2010) have labeled the type of audiences in SNSs with the term "networked audiences". "(...) The networked audience is unidentified but contains familiar face; is both potentially public and personal" (p. 129).

The collapse of audiences, per se, should not be regarded as a problem. On the contrary, presenting the self towards multiple and diverse audiences can also facilitate the performance of an individual, free from offline constraints in time and space. It is, however, necessary that the imagined audience align with actual audience to adequately define and control the social situation (Litt, 2012).

2.2. Imagined' vs. 'actual' audiences: users' practices and proposed solutions

User studies have indicated how the alignment between the imagined and actual audiences is lacking. Survey data by Lampe et al. (2008) showed how peers and close connections are widely seen as one's audience. Strangers and casual acquaintances were seldom considered. Hargittai and Litt (2013) found that a considerable proportion of users in their sample lack the skills to present themselves optimally when looking for a job, not taking into account future employers as an audience. Bernstein et al. (2013) indicated that "(...) social media users consistently underestimate their audience size for their posts, guessing that their audience is just 27% of its true size." Wang et al. (2011) studied regrets on Facebook after posting content and found that users often do not remember who might see their disclosures.

Service providers have invested in audience management tools. Most noticeable are Google circles and Facebook lists. A user can manually create and edit Facebook friend lists. Facebook also has automatically generated lists, called "Smart Lists". These lists are based on the user's profile matching his friends' profile on work, school, family and city, and can be modified by the user.

Two other empty lists called "close friends" and "acquaintances" are also automatically generated, waiting for the user to fill them. Compared to Facebook's later integration of friend grouping into its platform, Google made the grouping an inherent feature of its platform from its conception. That is, when the user adds a new friend, he or she needs to add this friend into at least one circle. Google+ does not offer its users any automatic grouping features like Facebook smart lists.

Many researchers have criticized the settings of service providers. Jones and O'Neill (2010) asserted that Facebook's effort to group people automatically is inadequate. In their user study, they found the groups were incomplete and did not present the criteria users considered to group their friends. The user study of Johnson et al. (2012) found that Facebook access control models are effective for managing content with strangers (i.e. outsider threat). They also found that many users were concerned about insider threat, defined as "inappropriately sharing content with members of the friend network." However, the tools provided did not protect against insider threats. Mazzia et al. (2012) stated that the existing comprehension tools of Facebook do not align with the mental model of users.

2.3. The promise of audience visualizations

The settings, as described in previous section, are focused on audience control. Researchers, however, have proposed audience visualizations to make users more aware and provide for more informed decisions. Lipford et al. (2008) designed a prototype interface that provided visual feedback on one's audiences and found a 42% improvement in time to complete a task (e.g. answering questions about who gets to see what personal information). To help users in sharing information with different groups on Facebook, Egelman et al. (2011) designed an interface based on Venn diagrams, so users could see which group of friends overlapped. This visualization reduced errors with 55% in comparison to those using the Facebook interface. Mazzia et al. (2012) developed a privacy-policy comprehension tool (PViz) that shows the user the visibility of his/her profile

according to the grouping of friends. They compared PViz with Facebook's Custom Settings Page and Facebook's Audience View and found significant improvement in user accuracy for complex, group based tasks (e.g. limiting the visibility of content to a group of friends). Wang et al. (2011) and Wang et al. (2013) have researched the presence of audience cues when posting a status update on Facebook. Below the status-update box in Facebook they presented five random pictures of the user's Facebook friends, with following caption: "these people can see your post". This feature was well received by their participants and encouraged them to be cautious when disclosing.

3. Research questions and study overview

The literature review showed that it is necessary to align the imagined with the actual audience before adequate control to be possible. Moreover, the existing tools provided by SNSs are insufficient for this cause, while audience visualizations are put forward as a solution. It is, however, unclear how users value audience visualizations, and how such visualizations should look like to draw the users' attention.

The affordance theory states that in addition to the properties of certain objects, it is important to look at its possibilities of action (Norman, 1988). A handle of a window provides an affordance for opening it; a cord provides an affordance for pulling it or making a knot; a doorknob allows an affordance for turning it. "What are the affordances of audience visualizations for Facebook friends perceived by users?" is our first research question.

Many audience visualizations have been developed and studied (see above). Card sorting (Jones and O'Neill, 2010; De Wolf and Pierson, 2014) and friend tagging (McAuley and Leskovec, 2012) have been used to determine user's manual groupings and grouping criteria. "What audience visualizations do users prefer and what draws their attention?" remains to be studied and constitutes our second research question.

To answer the research questions, we draw on the results of two user studies. In the first study, the first version of FreeBu was installed on the participants' computers. They interacted with FreeBu while thinking aloud. Meanwhile, a researcher asked the participant questions about his or her perceptions. In the second study, participants interacted with FreeBu online, with mouse movements being logged, and answered a questionnaire. Both studies took place in Belgium (08/2012–10/2012 respectively 11/2013–02/2014). The next section describes the audience visualizations and the design choices that were made in FreeBu. It is our intention to make the design choices explicit and testable with respect to our research questions.

4. Audience-visualization design choices in FreeBu

4.1. FreeBu #1

The modularity-based community detection algorithm has been shown to be useful for helping user group friends (Mazzia et al., 2012; Gao et al., 2012), and it was used for the algorithm in FreeBu #1. We used a hierarchical star-tree layout to show groups of friends (see Fig. 1).

The user is indicated with a blue circle, labeled with "self", in the middle of the image. On the first level of the "tree" are the red circles with various sizes. Each red circle represents a group. The more people contained in a group, the larger its size. On the second level, the child nodes (green) surrounding their parent nodes (red) are the user's Facebook friends. The labels are shown on top of the red circles if a community contains more than one person. The user can interact with the visualization by zooming and panning. The labels are typically school names, school years and work places, along with the group size. The user can customize the friend groups by moving the green nodes to different red nodes or creating new red nodes. Eventually the user can submit the friend groups as lists to their Facebook account.

4.2. FreeBu #2

FreeBu #2 (De Wolf and Pierson, 2014) is an extended version of FreeBu #1, with improved usability that includes three other types of audience visualizations. The design rationale is derived from experiences of users; the results of the interview study reported below and further research carried out after that study.

4.2.1. Architecture

FreeBu #1 is a desktop application written in Java, which requires users to install a Java Virtual Machine and therefore limits usage of the tool. FreeBu #2 is a web application, with the Facebook web login interface, which users can easily access.

4.2.2. Usability

The friend circles in the grouping visualization in FreeBu #1 sometimes overlap. We refined the layout algorithm and ways of representing the circles to avoid such overlap. We name the revised visualization the circle visualization (Fig. 2). In addition, we designed a new drag-and-drop interface for users to more easily compose individual lists that they can then upload to their Facebook account (rather than having to edit circles directly and then upload everything). This drag-and-drop interface, shown in Figs. 2–5, is consistent across the circle visualization and the additional new views, described below.

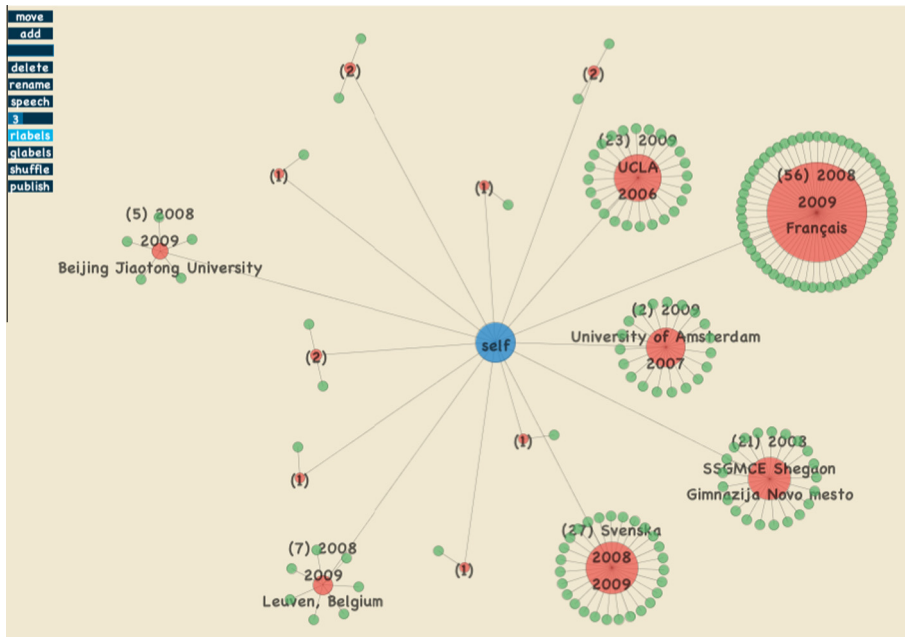


Fig. 1. The visualization in FreeBu #1.

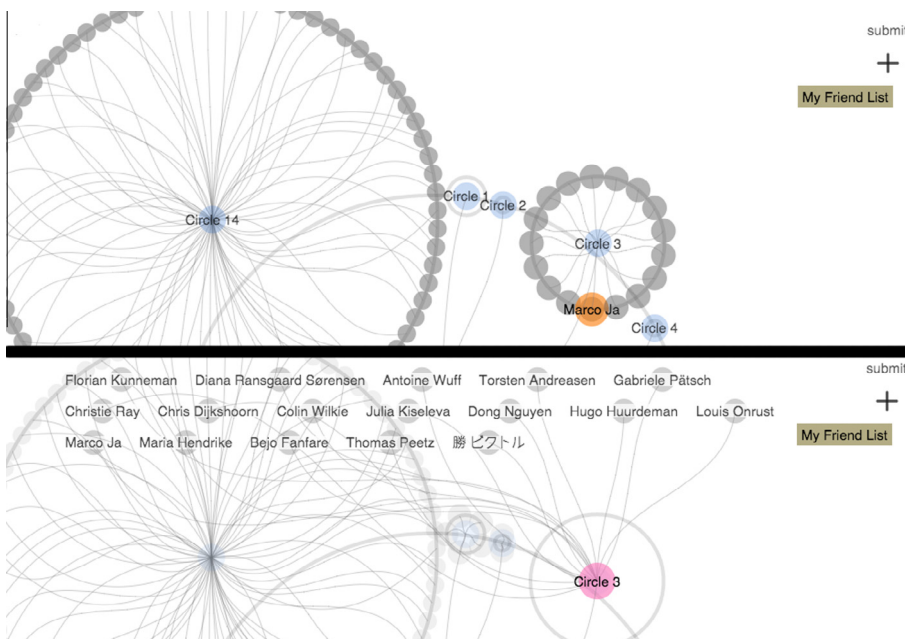


Fig. 2. The circle visualization in FreeBu #2.

4.2.3. Algorithm and interaction

In the circle visualization, we still use the same community detection algorithm as that of FreeBu #1. However, the algorithm sometimes has the resolution limit problem (Fortunato and Barthelemy, 2007), and large circles can be inconvenient for users to browse. We therefore made the circles “dividable”. The user can zoom into one circle, and the same algorithm will further divide this circle into sub-circles.

4.2.4. Other audience visualizations

Previous research found that besides users creating friend-groups based on the friends’ mutual connections (Figs. 1 and 2), users also employ other strategies (Jones and O’Neill, 2010; Gao and Berendt, 2013). These strategies were also translated into different audience visualizations in FreeBu #2.

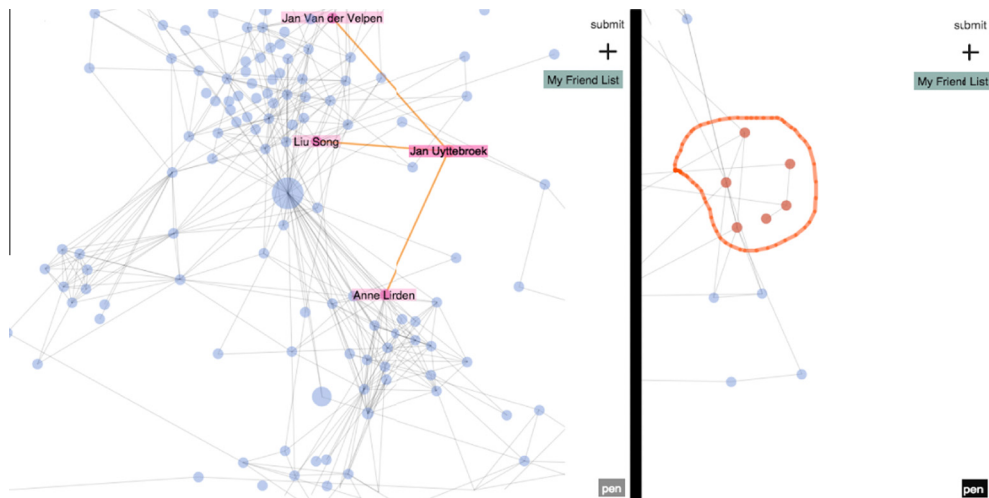


Fig. 3. The map visualization in FreeBu #2.

First, the user may create a friend group based on a specific friend, e.g. 'My friend Bob introduced me to all the friends in this group.' For this strategy, we created the map visualization (Fig. 3), in which the user can select a specific friend and his/her immediate neighbors or encircle an arbitrary set of friends. Second, the user may create a friend group based on which location, school, or organization they share. For this strategy we created the column visualization (Fig. 4), in which friend names are grouped and stacked into columns according to the common profile attributes they share. Third, the user may create a friend group based on his/her emotional closeness towards them. As studies have shown that interaction frequency linearly corresponds to the strength of interpersonal ties (Dunbar and Sutcliffe, 2012; Sutcliffe et al., 2012), we visualize the chatting frequencies between the user and each of his/her friends on Facebook in the rank visualization (Fig. 5).

5. Research design

5.1. Participants in the studies

We selected adolescents and young adults as the population of our study. People in this age group typically go to work or go to college, which increases the presence of multiple types of audiences. Young users also tend to have more friends than the average Facebook user (Brandtzæg et al., 2010). For the two studies, participants were recruited by reaching out to different intermediaries who are in close contact with adolescents and young adults (e.g. youth organizations, teachers).

In the first user study, we contacted the head of different youth organizations to spread the call for participating in the initial testing of FreeBu #1. The users were compensated with a cinema ticket. Twelve people (7 male/5 female; 17–23 years old) participated in the first study. They all had a minimum of 300 Facebook friends.

For the second user study, the intermediaries distributed the survey and FreeBu #2 link via email and SNSs, such as Facebook and Twitter. By participating, the users had the opportunity to win cinema tickets. 98 Facebook users filled in

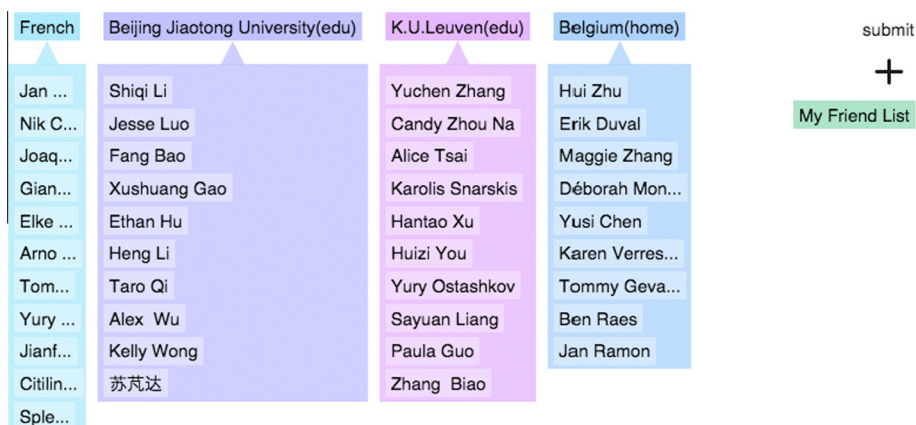


Fig. 4. The column visualization in FreeBu #2.




Fig. 5. The rank visualization in FreeBu #2.

the survey; of these, 49 also used FreeBu #2. In the following and in our report of the results in Section 6.2, we concentrate on these participants and their self-reported data from the questionnaire. Ages ranged from 16 to 34 ($M = 21.6$, $SD = 5.0$), 20 (29) participants were male (female), and they had between 100 and 1935 ($M = 465.3$, $SD = 327.1$) Facebook friends. 89.8% reported using Facebook on a daily basis.

Of the 49, 40 allowed us to collect their Facebook data, while nine chose to opt out of data collection. In Section 6.3, we report on these 40 users' behavioral data, obtained by logging their interactions with FreeBu #2 (set A). In addition, we include logging data from a further 39 users who used FreeBu #2 during the time of the study and allowed us to collect their data (4 had opted out of a total of 43), but without answering the questionnaire (set B). According to their Facebook-account data, these participants were similar to the other sample: ages ranged from 17 to 36 ($M = 23.9$, $SD = 5$), 24 (15) participants were male (female), and the numbers of friends ranged from 29 to 820 ($M = 349.3$, $SD = 169.0$).

5.2. Materials

In the first user study, participants interacted with FreeBu #1. The researcher used a semi-structured topic guide as the basis for the interview (See Appendix A.1). The core materials of the second study were FreeBu #2 and an online questionnaire. The questions were grouped into questions about usability, affordances, perceptions of the different visualization and privacy related questions¹. The affordances questions were operationalizations of the affordances discovered in the first study. All items were measured on a 7-point Likert scale. When needed, we included the option 'I don't know' or 'not applicable'. The negatively worded items are reverse-scored. The usability items were adapted from the Computer System Usability Questionnaire (Lewis, 1995). The visualizations of FreeBu #2 were enhanced, on its Web site, by tabs giving access to instruction videos and illustrations explaining FreeBu and the procedure of the study. In addition, an 'opt-out' button was available for participants who do not wish their data to be collected.

5.3. Procedure

5.3.1. General procedure

In the first study, the researcher installed FreeBu #1 on the computer of the participants and let them experiment whilst asking them questions from the semi-structured topic guide. We deliberately did not provide any task-based instructions but let participants interact with FreeBu in an open way in order to understand their perceptions. The interviews were audio-recorded and transcribed afterwards.

In the second study, participants were asked to test FreeBu #2 at home. To recruit participants, an invitation letter (e-mail) was sent, which included the link to FreeBu #2 together with the survey link. In the invitation letter, we advised

¹ In this paper, we focus exclusively on the research questions about audience visualizations. Hence, the relation between audience and privacy is not studied.

users to first check the tutorial. When a user logs into FreeBu #2, the “tutorial” button is distinctively shown on the home page, which directs the user to the tutorial page. This page contains detailed video instructions and illustrations that explain the purposes of the website and how to interact and make use of the visualizations. No further instructions were given to the participants to let them use the tool in the most natural state. Participants then used FreeBu #2 and/or answered the survey.

5.3.2. Ethical considerations

In both studies, we were confronted with an ethical issue: the collection and processing of personal data. We took different measures to respect user preferences in this process.

In the first study, the personal data were retrieved via the Facebook Graph API, using the participants' access tokens. Before users installed FreeBu #1, we explained the whole installation procedure and asked for their general consent. Under the guidance of one of the researchers, participants installed FreeBu #1 on their computers. In the first step, they logged into their Facebook account. In the second step, they retrieved their access token through the Facebook graph API explorer, where they gave permission to the FreeBu #1 to use user data, friend's data and extended permissions. We made it clear to participants that through the application we would have access to all information for which they gave their consent. In the third and final step, they copy-pasted the access token into the token submitter through which the selected personal data was downloaded on their computers and FreeBu #1 was started.

In order to use FreeBu #2, users first need to log in with their Facebook accounts. The standard Facebook authentication procedure was used to inform users what data the application would collect via Facebook Graph Application Programming Interface. Participants could also opt out if they only wanted to use FreeBu #2 without contributing their data. On the homepage and the tutorial page of FreeBu #2 we explicitly gave instructions on how to opt out.

5.4. Data, measures and analysis

In the first study, we used a grounded theory approach when coding the transcriptions (Strauss and Corbin, 1998). In the first phase of coding, which was completely data-driven, we maintained a close connection between the codes and data and coded the one word that appeared significant. In a second phase we categorized these codes with labels such as “functionality”, “appearance” and “usability”. In a third and last phase, we further organized and merged the codes and outlined different dimensions and properties where necessary. The coding process is included in Appendix A.2.

In the second study, we analyzed the answers given to the survey questions as well as mouse-movement data logged by FreeBu #2. To compare between affordances and visualizations in the survey we made use of a Wilcoxon signed-ranked test. Mouse positions correlate highly with eye-gaze positions (e.g. Chen et al., 2001; Huang et al., 2012), a well-known indicator of attention (e.g. Geiselman and Bellezza, 1977). For the logging results, in line with the exploratory nature of our studies, we focused on descriptive statistics, we used *#checks*, the count of how often a user clicks on or hovers over a visual object (node, nametag, bar), to measure attention to the entity (Facebook friend) represented by the object. We complement this by measures that take the effect of the visual context of the object into account: in the circle visualization, we use the average percentage of friends checked by the user in a circle (%checked Friends) and the average number of checks per checked friend per circle (*#checks/friend/circle*) over all circles from all the users, to indicate the extent to which a user interacts with a circle. For example, a circle contains 5 friends, when the user checks 3 of them, the percentage of checked friends in this circle is 60%. When the total number of checks on this circle is 18, the number of checks per checked friend in this circle is 6. In the map visualization, we use the average number of checks on an individual friend node (*#checks/friend*) to indicate the extent to which the user interacts with nodes of certain size. Furthermore, we use Mahalanobis distance (De Maesschalck et al., 2000) to quantify the relative position of a node in the map visualization. The node that is far away from the other nodes on the screen (i.e. the outliers) will be assigned a relatively high score, while the node that “belongs to the crowd” will be assigned a relatively low score. The scores are normalized for inter-group comparison. We also use Betweenness centrality (Newman, 2005) scores to scale the sizes of the nodes. A node's Betweenness is equal to the number of shortest paths from all vertices to all others that pass through that node. A node with large Betweenness score indicates its central role in connecting different parts of the user's social network. We linearly map the nodes' Betweenness centrality scores to their diameters ranging from 10 pixels to 30 pixels. In the column visualization, we use the average number of checks per column (*#checks/column*) over all the columns from all the users, to measure the extent to which a user interacts with a column. In the rank visualization, similar to the map visualization, we use *#checks/friend*, only that each friend is represented as the friend's photo with bar chart.

6. Results

6.1. Study 1: interview results

We identified three affordances as perceived by the participants. To illustrate these, we make use of the quotes of the participants.²

² The participants' names are pseudonymized.

During the interviews many participants indicated how difficult it was to remove friends using the settings provided by Facebook. Several indicated that FreeBu #1 would make it easier to identify the friends they want to unfriend. The visualization provided by FreeBu #1 showed people with whom the participant did not have any affiliation, or in some cases were not familiar with at all. The conversation we had with Jonas (M; 17) serves as an example. *“I think I know this guy through playing an online game, but I do not really know him”*. When we asked him whether or not he would unfriend this person from his Facebook account, he replied *“If I could see him like this I would”*. The last sentence indicates that users need an extra push for managing their audience. Scanning the different categories, Len (M; 21) noticed that he had added two persons with the same name, and that only one of the other was familiar to him. Like Jonas, he would also unfriend this person from his Facebook account.

Participants indicated that it would be easy to share certain information with a limited audience. The quote of Lizzy (F; 18) illustrates this line of thought. *“These people are all involved in my youth movement. This would make it easy to post something just for them, such as things we do in my youth movement”*.

Some participants indicated that they would use the groups provided by FreeBu #1 to decrease the visibility for certain audiences. For Lara (F; 19), the functionality was rather straightforward when we asked her why she would use this technology. *“Ah, that’s a simple question. To better categorize your friends into different groups”*. We noticed that users found the visualization of the different categories relevant in further grouping their audiences and not starting on something without any preparation. When analyzing a large group of friends, Peter (M; 17) claimed that he could easily subdivide further. *“I could certainly make further categorizations. This guy used to be a member of our movement. This one is still a member. This one I got to know on Expies [a camp for youth movements]. So I would make a differentiation between all of these. Many people in this category apparently have people in it from KSA-KSJ-VKSJ [the overarching organization]”*.

Apart from unfriending and grouping audiences, most participants also identified the overview of one’s audiences provided by FreeBu #1 as valuable. The conversation with Stella (F; 18) illustrates this. She made clear that she did not like the thought on deleting or grouping friends. *“I don’t actually think that I would use it for that purpose, because I consider all these people as my friends. They can know everything about me”*. Lizzy (F; 18) also indicated the value of the people in her friends list, but preferred minimizing visibility to deleting persons. *“I find it too harsh to just delete friends. It would be a shame if I started remove friends”*. When asked why, she replied with a laugh, *“Because then I would have fewer friends. I know that may sound foolish”*. Most participants liked FreeBu #1 because it gave them an overview of their Facebook friends. For example, when asking Uschi (F; 19) whether she would use the tool to limit information access towards certain audiences she replied: *“no, it is just fun to see your friends like this”*.

We noticed that the participants often mentioned that the groups were ‘correct’ or ‘wrong’, indicating that they were reflecting on the different groups. The reaction of Ted (M; 17) illustrates this. *“This one is spot on. I know all of these people via my girlfriend. This one is her cousin. My girlfriend herself, however, is not in this category”*. Although participants did not always consider the groups to be ‘correct’, they were looking for links. When scanning the different groups, Lara (F; 19) mentioned the following: *“This group consists of people who were together with me in elementary school. This group consists of people of my youth movement and other youth movements. But this group is not correct. Different people are just thrown together. Maybe they are all residents of Ghent.”* Seth (M; 20) also remembered that he once added someone who has the same surname as he does. Most of the time the small groups were considered “correct”, while the larger groups were perceived as a merging of different groups. Although this pushed participants to reflect on who was in the group and make further connections, several participants indicated they wanted to know “how” FreeBu #1 made these categorizations.

Together, the interview data differentiated between three affordances: unfriending, grouping and reflecting on one’s Facebook friends. These affordances were further studied in the survey study.

6.2. Study 2: survey results

The affordances of FreeBu #2 in general and those of the different visualizations in particular were further studied in a survey study. In [Appendix B.1](#) the average scores and standard deviations are provided. Because of the small sample size the emphasis is on descriptive statistics.³ In [Appendix B.2](#), however, exploratory factor analysis and reliability testing of perceived value are included for further study.

The interview study clarified three different affordances: unfriending, grouping and reflection. The different affordances were translated into 7-point Likert scale items measuring the perceived value of FreeBu #2. [Appendix B.1](#) provides the items and scores. The first three items can be considered a direct translation of the three affordances (PV1–PV3), whereas the other items are further elaborations. The data shows how every perceived value item is scored higher than the neutral 4 on a 7-point Likert scale. The results show how the participants value FreeBu for different reasons, especially for grouping and reflection purposes. We used a Wilcoxon Signed Rank test to study what is valued the most by the participants. The value for grouping or reflecting on Facebook friends is significantly greater than unfriending audiences. 18 of 38 participants scored higher on grouping than unfriending audiences ($z = -2.820$, 14 ties, $p < 0.01$). 20 of 38 participants scored higher

³ Initially, it was our intention to study what influences the perceived value of audience visualizations (i.e. measuring the predictors using regression analyses). In the survey study privacy variables were also integrated (privacy concern, privacy management, turbulence) to explore the relationship between privacy and the perceived value of audience visualizations.

Table 1

The number of users who used each of the four visualizations, in A, B and the merged set.

User set	Circle visualization	Map visualization	Column visualization	Rank visualization
A	30	29	26	31
B	27	30	24	21
A∪B	57	59	50	52

on reflecting than unfriending audiences ($z = -2.699$, 10 ties, $p < 0.01$). No significant results were found when comparing grouping with reflecting.

The usability items were adapted from the Computer System Usability Questionnaire (Lewis, 1995). The average score on each item are shown in Appendix B.1. Overall, the scores on usability are rather high. 15 participants also left comments about the usability in an open question. Many indicated that not all visualizations were displayed or loaded very slowly, e.g. “zooming was tiresome and slow.” Another point of critique was that it was very hard to learn how the technology worked, e.g. “I found it very hard to learn how to use this.”

The mean score for all questions regarding the visualizations was never lower than 4 on a 7-point Likert scale, with the exception of the fifth and sixth item of the column visualization. (“The relation between the characteristics in the columns made clear who my Facebook friends are”, $M = 3.83$; “The columns provided me with a clear image of my Facebook friends”, $M = 3.81$).

6.3. Study 2: logging results

In this section, we report the analysis results of the log data of the user interactions with the four visualizations. Users from set A and set B are differentiated, as described in Section 5.1. Table 1 shows the respective numbers of A and B users in the four visualizations, from which we can see that the circle and map visualizations attract slightly more users than the other two.

We perform four two-tailed Mann–Whitney U -Tests ($\alpha = 0.05$) on A–B pairs for the number of mouse checks in the four visualizations, and find that the difference between A and B in each visualization is insignificant. Therefore we merge the two sets of users, and summarize our explorative analysis results in Fig. 6. From this figure we can see that the users in the circle visualization generally have the most mouse checks, while those of the column visualization have the least. The median #checks for the map and the rank visualizations are similar –43 and 40, respectively. But the rank visualization has a broader upper quartile, Whisker, and outlier ranges than those of the map visualization, and is comparable to those of the circle visualization. This indicates that the rank visualization also attracts more mouse checks compared to the map and the column visualizations.

Furthermore, we can see that the results in Table 1 and Fig. 6 are consistent with the survey results of Study 2 in the sense that the circle and the rank visualizations are generally preferred while the column visualization is less favored by the users.




Fig. 6. The boxplot for #checks for each visualization, the cross signs are the outliers, the numbers in the brackets indicate the counts of non-outliers.

Moreover, we measured the session time that a user stays within the FreeBu, and found that the average session time per user is 13.8 min, with standard deviation 15.2 min (excluding an outlier that exceeds 10 h). The median time is 9.3 min. We did not record the session times for different visualizations of each user, but indeed this more-fine grained logging would provide us more insight into users' behavior with the FreeBu in future work.

6.3.1. Interactions with the circle visualization

The circle visualization shows an initial set of circles as friend groups for the user to interact. Due to the nature of this type of visualization, the differences in the sizes of the circles may influence the user's subsequent interactions. We use two metrics “%checked Friends” and “#checks/friend/circle” to measure the extent to which a user interacts with a circle, as described in Section 5.4.

Figs. 7 and 8 show our observations on these two metrics, with grouped circle size, for A users, B users and the two sets merged. Each dot represents an averaged value corresponding to the group of circles with a size range (e.g. [1–10]). The number in the brackets is the total number of circles with that size range. For example, [11–20] (39) indicates that there are 39 circles with size between 11 and 20 friends.

From Fig. 7 we can see that, when the circle size is not extreme, between 11 and 130 friends, the percentage of checked friends remains stable around 30% for the majority of the circles, regardless of their sizes. This pattern is similar in A and B users. This suggests a “checking-threshold” for a user to “comprehend” a group of visual objects. From Fig. 8 we can see that a friend node in a circle is typically checked 1.3 times by a user (for both A and B) across the whole range of circle sizes.

6.3.2. Interactions with the map visualization

In the graph-layout of the map visualization, we look into how different sizes and positions of the nodes affect users' interactions with the visualization.

It is expected to see that larger nodes attract more repeated checks from the users, as shown in Fig. 9. This pattern is consistent in both set A and B users. Fig. 10 groups the normalized Mahalanobis distances with a 0.2 interval, from the corresponding #checks/friend, we see that the nodes with large Mahalanobis scores attract more user interactions, which is especially the case for set A users. Recall that the nodes with large Mahalanobis scores are the ones that are isolated from the rest of the nodes, i.e. outliers. Fig. 11 shows two example map visualizations, in which if a node is checked, it is colored.




Fig. 7. The percentage of checked friends, grouped by circle size.

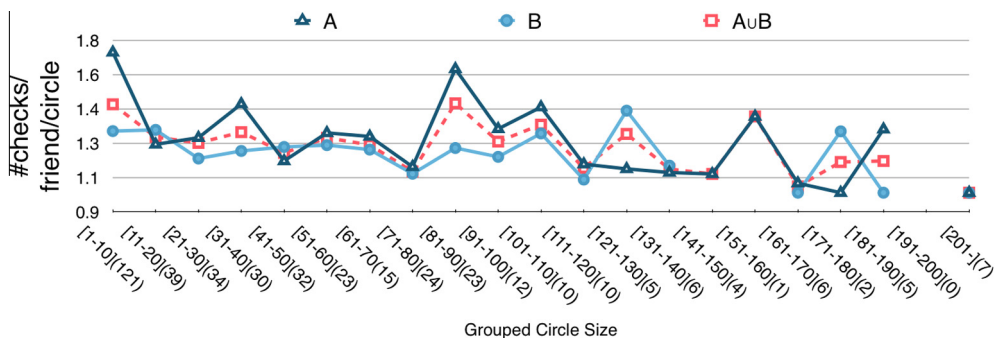


Fig. 8. The number of checks per friend within a circle, grouped by circle size.

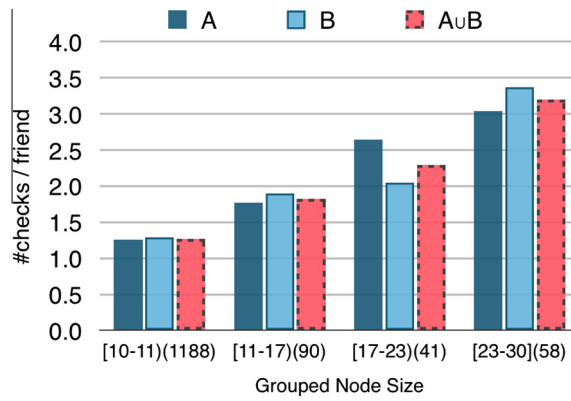


Fig. 9. The average number of checks per friend node with grouped node size in the map visualization. To read the X-axis, e.g. [11–17] (90) indicates all the nodes with [11,17] pixel size, and there are 90 such nodes in total.

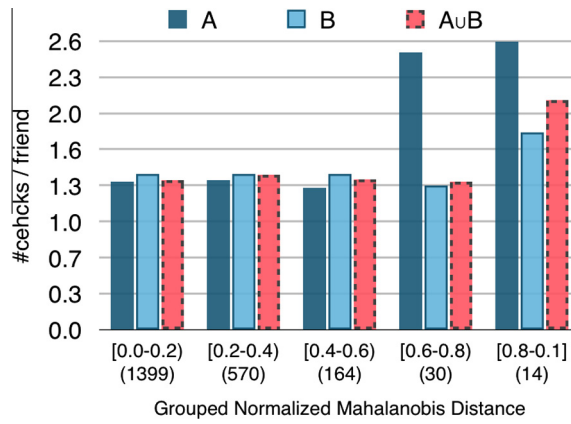


Fig. 10. The average number of checks per friend node with grouped Mahalanobis distance in the map visualization. To read the X-axis, e.g. [0.2–0.4](570) indicates all the nodes with their normalized Mahalanobis distance in [0.2, 0.4], and there are 570 such nodes in total.

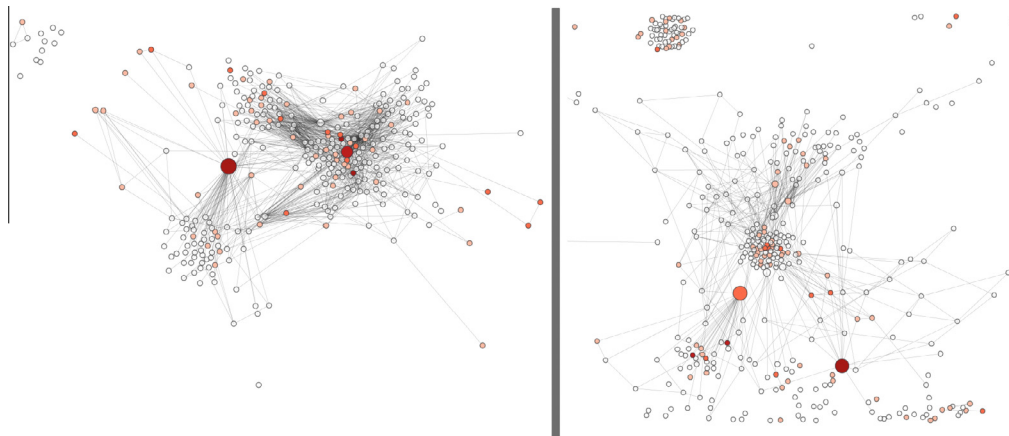


Fig. 11. Two examples showing users' interaction focus in the map visualization.

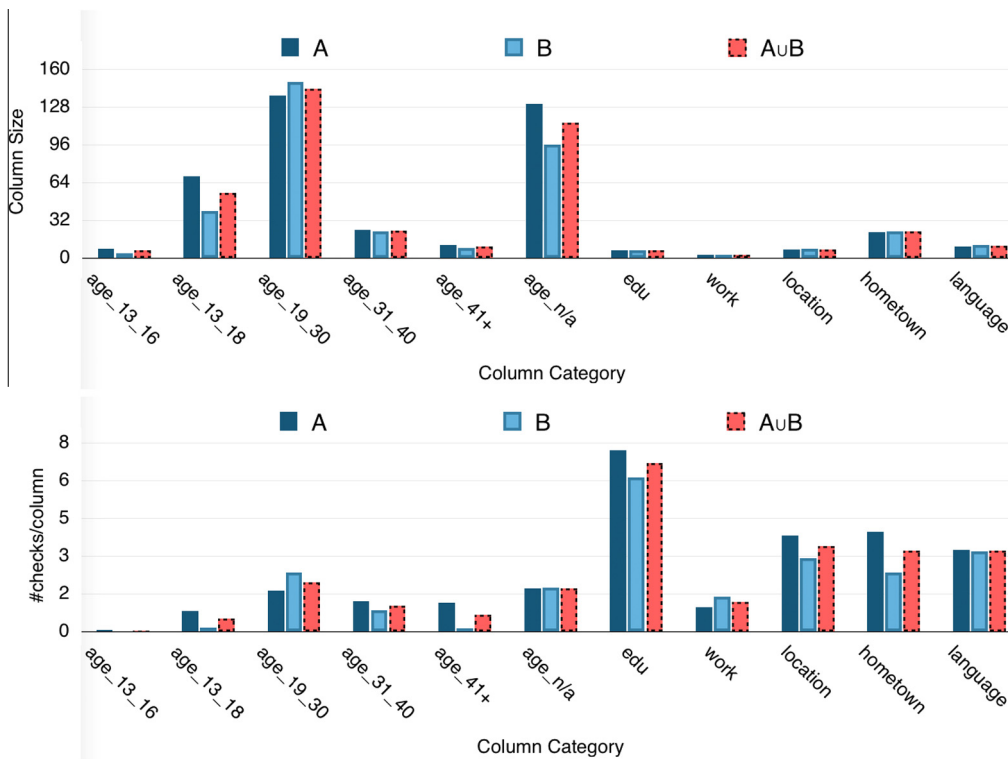


Fig. 12. The column sizes (top) and #checks/column (bottom) on different types of columns in the column visualization.

The more number of checks, the darker the color becomes. We can see that not only users focus on the larger nodes, but also focus on the “bridging” and “isolated” parts of a map visualization.

6.3.3. Interactions with the column visualization

In the column visualization, large columns (the ones containing more friends) are placed to the left of the screen, so that the user will firstly see. However, we find that users seek to interact with specific columns in which they are interested, despite the initial positioning of the columns. As shown in Fig. 12-top, large columns are age-related, while the other types of columns such as “hometown”, “edu” (education), “work” are much smaller. But from Fig. 12-bottom, we see that the average numbers of checks per type of column by the users are concentrated on these smaller columns, especially education-related ones. We can also observe this pattern is similar between A and B users.

6.3.4. Interactions with the rank visualization

In the rank visualization, the friends with whom the user has chatted most frequently online are put to the left of the screen. When visualization starts, the user sees the first 16 of these friends, as shown in Fig. 5. Mouse scrolling to the right reveals more friends. We find that coincidentally, 16 is also the average number of friends with whom a user has chatted on Facebook in both set A and B. Fig. 13 shows users’ (A, B and the merged set) average number of checks per friend (#checks/friend) on the top 20 friends, as they attract more than 90% of the total user interactions. From Fig. 13, we can see that users check the top two friends very frequently, then the number of checks starts to decrease, but increases again around the 15th friend. We can observe this pattern in both A and B users.

7. Discussion

7.1. Results and reflections

The presence and collapse of multiple audiences on SNSs challenges users. Although presenting one’s own identity towards multiple audiences can be liberating and free from certain constraints, knowing and controlling that audience remains necessary. The literature review indicated that users do not fully know their audience or are unable to control them, while audience visualizations were put forward as a promising solution. In this research we contributed by studying how users perceive audience visualizations and what draws their attention. Both are necessary requirements to align the imagined with the actual audiences.




Fig. 13. The average number of checks per friend on the top 20 friends in the rank visualization.

Studying the initial circle visualization, the interview data pointed out that the participants perceive the technology useful for unfriending, grouping and reflection purposes. When adding and studying three other visualizations (mapping, columns and ranking) in a second study, the survey data indicated that reflecting and grouping are significantly greater valued than unfriending.

Comparing the different values items of all types of visualizations in the survey data it is not clear what the participants preferred. Only the column visualizations received lower scores on providing a clear image of one's Facebook friends. The logging data showed how the participants paid attention to different aspects of their networked audiences. For example, often participants explored a similar percentage of a community of friends regardless of its size, or larger nodes positioned as bridges in the map visualization attracted users' attention. The rank visualization gives partially visualized individual friends, that is, the friends with whom the user has chatted. The results show how the participants mainly interacted with the top 20 of their Facebook friends, arguably their intimate friends. Finally, the logging data made clear that the column visualization received relatively less user-attention than the others. This could indicate that the profile-based friend groups presented in the form of columns of stacked names are less clear or interesting to the users. This finding is also supported by the survey results, considering the low score on providing a clear image of Facebook friends.

The existing audiences settings in SNSs are focused on grouping. For example, Facebook lists make it possible to group "close friends" from "colleagues" or "acquaintances". Based on the findings, however, we argue that users value the "networkedness" that defines their networked audiences. The existing settings do not provide an overview of one's entire network in SNSs. Previous research has indicated the value of audience visualizations for access control and privacy management. Our findings further indicate that users value audience visualizations for reflection, next to grouping, and, when interacting with the visualizations, are drawn to different parts of their network. We therefore propose a shift from "audience control" to "audience transparency", and propose to move beyond an approach that is solely oriented on controlling audiences.

7.2. Strengths, limitations and future research

Previous research has proven audiences visualizations to be useful for users of SNSs in making more aware decisions (Lipford et al., 2008; Egelman et al., 2011; Mazzia et al., 2012), using a task based research design. We did not specify any particular tasks for the participants to focus on how they perceived and interacted with the technology in their most natural state. We argue that this open approach allows for a better understanding of users' perceptions and behaviors.

A research design that combines an open method with an evaluation of a technology under development is also challenging, and it may have influenced the response rate in both studies. The user studies are based on “snapshots” of FreeBu use. It is plausible that users might have a different perception towards FreeBu over time. Although the second user study allowed for using FreeBu #2 multiple times, only four participants chose to do so.

Future research should focus on studying audience visualizations in other platforms. It might be that users prefer other audience visualizations and perceive other affordances when operationalized in other types of SNSs (e.g. LinkedIn, Twitter, etc.). Moreover, Wang et al. (2011) indicated that the perception of one's audiences is dependent of their life stage. This study is limited to the population of adolescents and young adults users.

7.3. Design implications

FreeBu #1 did not explain to its users how it establishes the different groups of friends. It thus remains a black box. Then again, we noticed during the interview that not knowing the underlying algorithm provoked users to reflect on their audiences and how they are connected. We advise designers to take into the account the trade-off between motivating users to reflect on their networked audience through obscurity on the one hand, and providing them with a clear explanation of the algorithm's working on the other hand. For example, by selectively visualizing part of the “black box”.

The column visualization received low scores on providing a clear image of one's Facebook friends. Although scrolling through or combining different attributes of Facebook friends' profile data can help in knowing one's online audience, it can also be overwhelming. We therefore recommend being selective in choosing certain attributes when developing an attribute-based visualization. That said, we suggest the usage of education-related attributes because they attracted more attention.

Though the first user study indicated that FreeBu #1 would be useful for unfriending, the Wilcoxon signed-rank test showed that, after adding three more visualizations, the participants valued the grouping and awareness option more. This is not surprising. Restricting access to those who once were allowed to be a Facebook friend can be a hard decision to make. We therefore argue to focus on reflection and grouping options when designing audience visualizations.

8. Conclusion

Goffman (1959) made an analogy between everyday life and dramaturgy to study social reality where an individual is a performer and his or her environment the audience. With the emergence of SNSs it seems that the front stage has expanded and the audience collapsed. Over the years, service providers have invested in options to better control the information flow and group audiences. This approach is limited. We argue that it is necessary to acknowledge networked audiences on SNSs and make users aware of how exactly they are networked. The latter requires a shift from “audience control” to “audience transparency”.

Appendix A. (interview study)

A.1. Interview questions FreeBu #1

What do you think the grouping is based on?
 Does this grouping make sense to you?
 What is missing in the grouping structure?
 What would you like to see in the grouping structure?
 What do you think about the interface?
 What did you not expect?
 Would you also group your friends like this?
 Which features do you appreciate?
 Do you find tool easy to use?
 Is there something you find annoying?

A.2. Coding process interviews FreeBu #1

The coding process consisted of three phases to structure the qualitative data. This process helped with summarizing and analyzing the way the participants valued FreeBu #1.

Phase 1: Open coding	Phase 2: Labeling the codes	Phase 3: Organizing the codes and making further categories
Concept	Appearance	Appearance
Zooming	Usability	Color
Ugly	Appearance	Ugly
Period	Management	Labels
Place	Small group	Static
Alone	Small group	
Context	Mistakes	Functionality
Separately	Small group	Overview
Circles	Usability	Reflecting
Contact	Management	Grouping
Control	Management	Deleting
Visualization	Appearance	Announcements
Scary	Emotions	New
Enthusiastic	Emotions	
Feedback	Functionality	Emotions
Wrong group	Mistakes	Surprised
Mistakes	Mistakes	Scary
Using	Functionality	
Integrated	Usability	Small group
Mixed	Big group	
Big groups	Big group	Management
Small groups	Small groups	Period
Color	Appearance	Contact
Correct	Functionality	Context
Labels	Functionality	
Lists	Functionality	Usability
Mixed	Big group	Dragging
Mind map	Appearance	Zooming
Names	Appearance	Integrated
Not handy	Functionality	Searching
Unfamiliar	Functionality	
Divided	Functionality	Big group
Unnecessary	Functionality	
Remarkable	Functionality	Mistakes
Overview	Functionality	
Posting	Functionality	
Privacy	Functionality	
Searching	Usability	
Relevant	Functionality	
Scratch	Functionality	
Shift	Functionality	
Unchanged	Appearance	
Changes	Appearance	
Mixed	Big group	
Dragging	Usability	
Deleting	Functionality	
Mess	Appearance	
See	Functionality	
Meaningful	Functionality	

Appendix B. (survey study)*B.1. Means scores for perceived value, usability, and the four visualizations*

	M (S.D.)	Valid responses
<i>Perceived value</i>		
PV1 – FreeBu helps me remove unwanted audiences on Facebook	4.5 (1.29)	38
PV2 – FreeBu helps me group my Facebook friends	5.23 (1.25)	39
PV3 – FreeBu lets me think about who my Facebook friends are	5.56 (1.27)	39
PV4 – FreeBu clarifies the relationships with others I am not fully aware of	5.40 (1.31)	38
PV5 – FreeBu indicates close and distant friends	4.97 (1.48)	39
PV6 – FreeBu helps me share information with specific groups of friends	4.79 (1.44)	38
PV7 – FreeBu gives a wrong image of my Facebook friends (reverse coded)	5 (1.15)	39
PV8 – FreeBu helps me with making Facebook lists	5 (1.39)	38
PV9 – FreeBu does not provide a good overview of my Facebook friends (reverse coded)	4.05 (1.69)	39
PV10 – Overall, I find FreeBu a useful tool	4.69 (1.44)	39
<i>Usability</i>		
USA 1 – FreeBu was difficult to use (reverse coded)	4.07 (1.64)	42
USA 2 – It was easy to learn to use FreeBu	5.45 (1.56)	42
USA 3 – Whenever I made a mistake using FreeBu, I recovered easily and quickly	4.84 (1.46)	37
USA 4 – The videos with instructions were clear	6.05 (1.28)	40
USA 5 – The interface of FreeBu was pleasant	5.2 (1.36)	40
USA 6 – The videos and instructions were helpful when testing FreeBu	5.87 (1.51)	39
USA 7 – The interface was easy to use	5.38 (1.39)	40
USA 8 – FreeBu had all the functions and capabilities I expect it to have	5.1 (1.28)	40
USA 9 – FreeBu helps with making Facebook lists	5.46 (1.24)	37
USA 10 – It was difficult to find the information I needed (reverse coded)	4.63 (1.41)	40
USA 11 – Whenever something didn't work it was easy to solve	4.17 (1.63)	36
USA 12 – Overall, I am satisfied with FreeBu	5.07 (1.33)	42
USA 13 – Overall, I am satisfied with the usability of FreeBu	5 (1.40)	42
<i>Circle visualization</i>		
CIRC 1 – The circles were well-arranged	4.94 (1.63)	35
CIRC 2 – The circles provided me with a clear image of who my Facebook friends are	4.86 (1.72)	35
CIRC 3 – The circles were pleasant to see	5.09	35

(continued on next page)

Appendix B (continued)

	(1.62)	
CIRC 4 – The circles would match with a grouping I would make	4.62	34
	(1.78)	
CIRC 5 – The Facebook friends who were grouped together also belonged together	4.57	35
	(1.67)	
<i>Map visualization</i>		
MAP 1 – The map was well-arranged	4.14	35
	(1.7)	
MAP 2 – The provided me with a clear image of my Facebook relations	4.66	35
	(1.85)	
MAP 3 – The map was pleasant to see	4.43	35
	(1.80)	
MAP 4 – The map indicated how groups of friends are connected with each other	5.14	35
	(1.68)	
MAP 5 – The map indicated which groups are completely segregated	5.14	35
	(1.87)	
MAP 6 – The map provided me with a clear image of who my Facebook friends are	4.68	34
	(1.82)	
<i>Column visualization</i>		
COL 1 – The columns were well-arranged	4.81	31
	(1.76)	
COL 2 – The columns provided me with the characteristics of my Facebook friends	4.45	31
	(1.84)	
COL 3 – The columns were pleasant to see	4.87	31
	(1.71)	
COL 4 – The characteristics in the columns were relevant	4.1	30
	(1.73)	
COL 5 – The relation between the characteristics in the columns made clear who my Facebook friends are	3.83	30
	(1.7)	
COL 6 – The columns provided me with a clear image of who my Facebook friends are	3.81	31
	(1.64)	
<i>Ranking visualization</i>		
RANK 1 – The ranking was well-arranged	5.83	30
	(1.29)	
RANK 2 – The ranking provided me with a clear image with who I chat	4.7	30
	(1.84)	
RANK 3 – The ranking was pleasant to see	5.6	30
	(1.13)	
RANK 4 – The ranking provided me with a clear image of who my Facebook friends are	4.3	30
	(1.86)	
RANK 5 – The ranking provided me with how much contact I have with my Facebook friends	4.3	30
	(1.92)	

B.2. Exploring the dimensions in the perceived value of FreeBu #2

Exploratory factor analysis (EFA) gives insight into the underlying dimensions of a variable that consists of multiple items. Considering the small sample size, the outcome should be approached with care. Nonetheless, the results may be used in the further development of scales that measure the perceived value of audience visualizations. We especially recommend exploring the relationship between privacy related variables (e.g. privacy concern, privacy turbulence) and the perceived value of audience visualizations.

The factor analysis of perceived value of FreeBu reveals a three-factor solution (see table below). The first factor contains items that refer to demarcating audiences and sharing information with a particular group of people, e.g. “FreeBu helps me with making Facebook lists”. The first factor aligns with the grouping affordance found in the interview study. The second factor contains items that clarify one’s relationship with other Facebook friends, e.g. “FreeBu lets me think about who my Facebook friends are”. The second factor contains both the reflection and unfriending affordance found in the interview

study. Finally, the third factor contains two items that refer to an overview of one's Facebook friends, e.g. FreeBu gives a wrong image of my Facebook friends. Table 2 presents the factor loadings (after varimax rotation) of the factors that were retained and left out values 0.5 or less.

Further analyses proves factor 1 ($\alpha = 0.86$, $M = 4.92$, $SD = 1.31$) to be one-dimensional and consistent. The Cronbach's alpha of the three items in factor 2 is above 0.7, which is promising. However, when removing item 1 the Cronbach's alpha increases from 0.71 to 0.78, with a $M = 5.41$ and $SD = 1.22$. Factor 3, however, has an unacceptable low internal consistency ($\alpha = 0.52$) and is not considered to be adequate for further analyses.

Items	Factor 1 Controlling FB friends	Factor 2 Relationship with FB friends	Factor 3 Overview of FB friends
PV 2	0.746		
PV 6	0.816		
PV 8	0.836		
PV 1		0.614	
PV 3		0.764	
PV 4		0.716	
PV 7			0.581
PV 9			0.632
PV 5			
Eigen value	3.298	1.657	1.270
% Of variance	36.64	18.42	14.11
KMO			0.68
Bartlett's test (Approx. Chi Square)			109.339***

Factor analysis with varimax rotations of measures of perceived value of FreeBu #2 (***) $P < 0.001$.

References

- Bernstein, M.S., Bakshy, E., Burke, M., Karrer, B., 2013. Quantifying the invisible audience in social networks. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. ACM, New York, NY, USA, pp. 21–30.
- boyd, 2008. Taken Out of Context: American Teen Sociality in Networked Publics (PhD). ProQuest.
- boyd, d., Ellison, N.B., 2007. Social network sites: definition, history, and scholarship. *J. Comput. Mediat. Commun.* 13 (1), 210–230.
- Brandtzæg, P.B., Lüders, M., Skjetne, J.H., 2010. Too many Facebook "Friends"? content sharing and sociability versus the need for privacy in social network sites. *Int. J. Hum. Comput. Interact.* 26 (11–12), 1006–1030.
- Chen, M.C., Anderson, J.R., Sohn, M.H., 2001. What can a mouse cursor tell us more?: correlation of eye/mouse movements on web browsing. In: CHI '01 Extended Abstracts on Human Factors in Computing Systems. ACM, New York, NY, USA, pp. 281–282.
- De Maesschalck, R., Jouan-Rimbaud, D., Massart, D.L., 2000. The mahalanobis distance. *Chemom. Intell. Lab. Sys.* 50 (1), 1–18.
- De Wolf, R., Pierson, J., 2014. Who's my audience again? Understanding audience management strategies for designing privacy management technologies. *Telemat. Informat.* 31 (4), 607–616. <http://dx.doi.org/10.1016/j.tele.2013.11.004>.
- Dunbar, R., Sutcliffe, A., 2012. Social complexity and intelligence. In: Shackelford, T.K., Vonk, J. (Eds.), *The Oxford Handbook of Internet Studies*. Oxford University Press, New York.
- Egelman, S., Oates, A., Krishnamurthi, S., 2011. Oops, I did it again: mitigating repeated access control errors on Facebook. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. ACM, New York, NY, USA, pp. 2295–2304.
- Ellison, N.B., boyd, d., 2013. Sociality through social network sites. In: D., W. (Ed.), *The Oxford Handbook of Internet Studies*. Oxford University Press, Oxford, pp. 151–172.
- Fortunato, S., Barthelemy, M., 2007. Resolution limit in community detection. *Proc. Natl. Acad. Sci. USA* 104 (1), 36–41.
- Gao, B., Berendt, B., Clarke, D., de Wolf, R., Peetz, T., Pierson, J., Sayaf, R., 2012. Interactive grouping of friends in OSN: towards online context management. In: 2012 IEEE 12th International Conference on Data Mining Workshops (ICDMW), pp. 555–562. <http://dx.doi.org/10.1109/ICDMW.2012.88>.
- Gao, B., Berendt, B., 2013. Circles, posts and privacy in egocentric social networks: an exploratory visualization approach. In: 2013 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM), pp. 792–796.
- Geiselman, R.E., Bellezza, F.S., 1977. Eye movements and overt rehearsal in word recall. *J. Exp. Psychol.: Hum. Learn. Memory* 3 (3), 305–315.
- Goffman, E., 1959. *The Presentation of Self in Everyday Life*. Penguin Books, Harmondsworth.
- Hargittai, E., Litt, E., 2013. New strategies for employment? internet skills and online privacy practices during people's job search. *IEEE Secur. Privacy* 11 (3), 38–45.
- Huang, J., White, R., Buscher, G., 2012. User see, user point: gaze and cursor alignment in web search. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. ACM, New York, NY, USA, pp. 1341–1350.
- Johnson, M., Egelman, S., Bellovin, S.M., 2012. Facebook and privacy: it's complicated. In: Proceedings of the Eighth Symposium on Usable Privacy and Security. ACM, New York, NY, USA, pp. 9:1–9:15.
- Jones, S., O'Neill, E., 2010. Feasibility of structural network clustering for group-based privacy control in social networks. In: Proceedings of the Sixth Symposium on Usable Privacy and Security. ACM, New York, NY, USA, pp. 9:1–9:13.
- Lampe, C., Ellison, N.B., Steinfield, C., 2008. Changes in use and perception of Facebook. In: Proceedings of the 2008 ACM Conference on Computer Supported Cooperative Work. ACM, New York, NY, USA, pp. 721–730.
- Lewis, J.R., 1995. IBM computer usability satisfaction questionnaires: psychometric evaluation and instructions for use. *Int. J. Hum.-Comput. Int.* 7 (1), 57–78. <http://dx.doi.org/10.1080/10447319509526110>.
- Lipford, H.R., Besmer, A., Watson, J., 2008. Understanding privacy settings in Facebook with an audience view. In: Proceedings of the 1st Conference on Usability, Psychology, and Security. USENIX Association, Berkeley, CA, USA, pp. 2:1–2:8.
- Litt, E., 2012. Knock, knock. Who's there? The imagined audience. *J. Broadcast. Electron. Media* 56 (3), 330–345.
- Marwick, A.E., boyd, d., 2010. I tweet honestly, I: tweet passionately: twitter users, context collapse, and the imagined audience. *New Media Soc.*

- Mazzia, A., LeFevre, K., Adar, E., 2012. The PViz comprehension tool for social network privacy settings. In: Proceedings of the Eighth Symposium on Usable Privacy and Security. ACM, New York, NY, USA, pp. 13:1–13:12.
- McAuley, J., Leskovec, J., 2012. Image labeling on a network: using social-network metadata for image classification. In: Fitzgibbon, A., Lazebnik, S., Perona, P., Sato, Y., Schmid, C. (Eds.), Computer Vision – ECCV 2012. Springer, Berlin, Heidelberg, pp. 828–841.
- Newman, M.E., 2005. A measure of betweenness centrality based on random walks. *Soc. Netw.* 27 (1), 39–54.
- Norman, D., 1988. *The Design of Everyday Things*. Doubleday, New York.
- Strauss, A., Corbin, J.M., 1998. *Basics of Qualitative Research: Techniques and Procedures for Developing Grounded Theory*. SAGE Publications.
- Sutcliffe, A., Dunbar, R., Binder, J., Arrow, H., 2012. Relationships and the social brain: integrating psychological and evolutionary perspectives. *Br. J. Psychol.* 103 (2), 149–168.
- Vitak, J., 2012. The impact of context collapse and privacy on social network site disclosures. *J. Broadcast. Electron. Media* 56 (4), 451–470.
- Wang, Y., Norcie, G., Komanduri, S., Acquisti, A., Leon, P.G., Cranor, L.F., 2011. I regretted the minute I pressed share: a qualitative study of regrets on Facebook. In: Proceedings of the Seventh Symposium on Usable Privacy and Security. ACM, New York, NY, USA, pp. 10:1–10:16.
- Wang, Y., Leon, P.G., Scott, K., Chen, X., Acquisti, A., Cranor, L.F., 2013. Privacy nudges for social media: an exploratory Facebook study. In: Proceedings of the 22nd International Conference on World Wide Web Companion. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland, pp. 763–770.

Ralf De Wolf is a PhD candidate within the user research unit of iMinds-SMIT at Vrije Universiteit Brussel. He holds a masters degree in sociology from the University of Ghent, Belgium, where he also completed a degree in teaching political and social sciences. His PhD focuses on the social aspects of security and privacy for social network sites. Key issues are the relationships between identity and privacy, and contextual privacy problems.

Bo Gao has been a PhD student since March 2011 in the Machine Learning research group from the department of computer science in KU LEUVEN in Belgium, under Prof. Bettina Berendt's supervision. His research interests include social media data analysis; privacy feedback/awareness web applications and information visualization.

Bettina Berendt is a professor in the Declarative Languages and Artificial Intelligence (DTAI) group at the Department of Computer Science of KU Leuven, Belgium. In her research, teaching and project activities, she aims at defining and developing human-centric data science. Specific research interests include web, social-web, and text mining, semantic web, how these technologies affect privacy and social fairness (especially whether and how they can be made to serve them), and feedback and awareness tools for information literacy.

Jo Pierson is Associate Professor in the Department of Media and Communication Studies at the Vrije Universiteit Brussel, Belgium. He is also Senior Researcher and staff member at the research center iMinds-SMIT (Studies on Media, Information and Telecommunication) since 1996, in charge of the research on privacy in social, mobile and ubiquitous media. He lectures on socio-economic issues relating to the information society, digital media marketing, and user innovation in new media. Drawing upon science and technology studies and media-sociological approaches, his research focus is on privacy, surveillance, datafication, and user empowerment in digital media.