Investigating Social Media Usage Patterns and Privacy Awareness with Composite Data Visualization

by

Ben Z. Yuan

B.S., California Institute of Technology (2014)

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Master of Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2017

Author ....
Department of Electrical Engineering and Computer Science
August 28, 2017

Certified by ....
Hal Abelson
Class of 1922 Professor of Computer Science and Engineering
Thesis Supervisor

Certified by ....
Ilaria Liccardi
Research Scientist, CSAIL
Thesis Supervisor

Accepted by ....
Professor Leslie A. Kolodziejski
Chair, Department Committee on Graduate Students

Investigating Social Media Usage Patterns and Privacy Awareness with Composite Data Visualization

by

Ben Z. Yuan

Submitted to the Department of Electrical Engineering and Computer Science
on August 28, 2017, in partial fulfillment of the
requirements for the degree of
Master of Science

Abstract

This thesis describes an investigation into the degree of awareness people have of their activity and audience on social media, and into the alignment of sharing expectations with actual sharing behavior. It is previously reported that people tend to share problematic posts on social media networks because they are not always aware of who can actually see their posts and other activity and do not always apply privacy settings effectively. We built a data collection tool that gathers social media data, like posts, connections, and private messages, from Facebook, Twitter, Instagram, and LinkedIn, and assembles a composite profile combining information from all four networks for visualization. We then conducted a user study evaluating people's data sharing patterns, audience perceptions, and data self-awareness on social media. We first surveyed participants to discover their own estimates of certain activity and visibility metrics like post type ratios, connection proportions by interaction frequency, and connections by presence on multiple networks; we then interviewed them with the aid of the tool's visualization to compare their answers with ones we computed from their collected data and gauge their reactions. Notably, we determined that participants tend to significantly overestimate the proportion of connections with whom they interact on social media, and we found that participants also have trouble recalling what types of posts they have made and how many people they share between networks; nevertheless, when presented with the actual computed information and a visualization of their social media activity and visibility, most participants reported being satisfied with their sharing strategy, although a minority did report a desire to change their behavior or re-examine their sharing settings. This document presents the methods used, the results from the user study, and suggestions and cautions for future work.

Thesis Supervisor: Hal Abelson
Title: Class of 1922 Professor of Computer Science and Engineering

Thesis Supervisor: Ilaria Liccardi
Title: Research Scientist, CSAIL

Acknowledgments

Completion of this thesis and the work leading to it would not have been possible without the support and guidance of several key people. Every one of them provided something essential for the realization, improvement, and finalization of this project and of the prose that describes it.

I would first of all like to thank Hal Abelson, who co-supervised this thesis, for providing motivation, wisdom, and careful comments during the progression of this work. His mentorship and timely input were crucial to its completion, and I am indebted to him for his guidance, feedback, and patience throughout the process.

I would also like to thank Ilaria Liccardi, who co-supervised this thesis as well, for providing feedback, ideas, and encouragement throughout the entire process, and for her assistance with the user study described in this work. No part of this project would have been possible without her knowledge and tireless support.

I would additionally like to offer my thanks to Fernanda Viégas and Martin Wattenberg for taking the time to review ideas and designs during the planning of this project. The feedback they offered helped improved the visualization ultimately developed for this work and made it much more effective.

I am especially grateful to the participants of the user study on which this work is based, for their participation and – in some cases – great patience and understanding when we encountered difficulties during the process.

Finally, I would like to offer thanks to all of the students, faculty, and staff in the Internet Policy Research Initiative, for making the lab a better place to work; and to my other friends and family, for making the world a better place to live.

Contents

1 Introduction **17**

2 Research aims and objectives **21**

2.1 Perceptions of audience composition 23

2.2 Information sharing patterns and perceptions 24

2.3 Effects of data visualization on perceptions of privacy risks 26

3 Related research **29**

3.1 Perceptions of audience composition 29

3.2 Information sharing perceptions and patterns 31

3.3 Visualization as a tool for understanding privacy issues 33

4 Polyhedron: data collection and visualization **35**

4.1 Data collection 36

4.1.1 Facebook 39

4.1.2 Twitter 43

4.1.3 LinkedIn 44

4.1.4 Instagram 47

4.1.5 Challenges and limitations 49

4.2 Data presentation 50

5 User study: data estimation and visualization **55**

5.1 Methods . . . 55

5.2 Results . . . 57

5.2.1 Participant demographics . . . 57

5.2.2 Connections by interaction frequency . . . 60

5.2.3 Connections by network grouping . . . 62

5.2.4 Content proportions by type . . . 63

5.2.5 Content breakdown by supplemental data presence . . . 68

5.2.6 Activity proportions . . . 71

5.2.7 Concern and control over visibility . . . 74

5.2.8 Subjective variations in social media usage strategies . . . 77

5.2.9 Effects of the Polyhedron visualization suite . . . 78

5.3 Discussion . . . 80

5.3.1 Perceptions of audience composition . . . 80

5.3.2 Information sharing patterns across networks . . . 81

5.3.3 Effects of data visualization on perceptions of privacy risks . . . 84

5.3.4 Summary . . . 85

6 Problems and challenges **87**

6.1 Evolution of target platforms . . . 87

6.1.1 Facebook . . . 88

6.1.2 Twitter . . . 89

6.1.3 LinkedIn . . . 90

6.1.4 Instagram . . . 90

6.2 Target platforms behaving differently for different people . . . 91

6.2.1 Facebook profile timeline . . . 93

6.3 Limitations imposed by data management policy . . . 94

6.4 Issues arising from participant data volume . . . 96

6.5 Reflections 97

7 Conclusions **99**

7.1 Summary of findings 99

7.2 Future work 102

7.2.1 Potential insights from text and photo content 102

7.2.2 Examination of additional networks 103

A Supplemental graphs **105**

A.1 Content proportions over time 105

A.2 Activity proportions over time 110

B Survey **115**

C Supplemental data tables **135**

Bibliography **147**

List of Figures

4.1 Overall architecture diagram of the Polyhedron add-on. 36

4.2 Overall architecture diagram of the Facebook page worker. Other page workers are similar. 38

4.3 Overview page from the visualization. 51

4.4 Post page from the visualization. 52

4.5 Location page from the visualization. 53

5.1 Participant estimation errors of connection proportions by activity level class. 61

5.2 P5 content proportion estimates for Twitter; note dotted line (survey result) and time trace (showing breakdown of computed content proportions within each year). 66

5.3 P2 and P11 content proportion estimates for Facebook; note dotted line (survey result) and time trace (showing breakdown of computed content proportions within each year). 67

5.4 Participant estimation errors of content proportions. 67

5.5 Participant estimation errors of activity proportions. 73

List of Tables

5.1 Participant demographics, summarized.
S = social, P = professional, O = other, blank = not used 58
5.2 Activity and connection total counts, summarized. 59
5.3 Mean error and standard deviation in percentage points for Part 17 (proportion estimates for connection types by interaction). Survey answers were linearly rescaled to add to 100. 60
5.4 Mean computed absolute set sizes and standard deviations for Part 17. . . . 61
5.5 Mean percentage error for Part 19 (cross-network connection count). 64
5.6 Mean percentages and standard deviations for content proportions (percentage points) by type. 64
5.7 Mean absolute counts and standard deviations for timeline content by type. . 65
5.8 Mean errors and standard deviations, in percentage points, for estimates of content proportions (percentage points) by type. 66
5.9 Mean errors and standard deviations, in percentage points, for estimates of content proportions by type, compared against numbers for the most recent calendar year containing activity. 68
5.10 Correlation coefficients between network timeline sizes and mean absolute error. 68
5.11 Mean percentages and standard deviations of "plain" posts (compared to all posts of the same general type, e.g. all photos), in percentage points. 69
5.12 Mean counts and standard deviations of "plain" posts. 69

5.13 Mean percentages and standard deviations of posts with location (compared to all posts of the same general type), including posts with both location and other data. 70
5.14 Mean counts and standard deviations of posts with location, including posts with both location and other data. 70
5.15 Mean percentages and standard deviations of posts with tagged people (compared to all posts of the same general type), including posts with both tagged people and other data. 71
5.16 Mean counts and standard deviations of posts with tagged people, including posts with both tagged people and other data. 71
5.17 Mean estimation errors and standard deviations, in percentage points, for estimates of activity frequency. Survey answers were renormalized to add to 100. 72
5.18 Mean absolute counts and standard deviations for activity frequency. 72
5.19 Mean estimation errors and standard deviations, in percentage points, for estimates of activity frequency, based on most recent calendar year of activity only. Survey answers were renormalized to add to 100. 73
5.20 Correlation coefficients between total activity count and mean absolute estimation error. 73
5.21 Mean answers (and standard deviations) for Part 14 (concern over visibility). 75
5.22 Mean answers (and standard deviations) for Part 15 (perceived control). . . 76
5.23 Correlation coefficients between mean absolute errors and mean Likert-scale concerns. 77

7.1 Summary of hypotheses investigated and what we learned about each (H1-H5). 101
7.2 Summary of hypotheses investigated and what we learned about each (H6-H8). 102

C.1 Part 10: Timeline content proportions (survey and computed, percentage points, to the nearest point) per participant, Facebook. 135
C.2 Part 10: Timeline content proportions (survey and computed, percentage points, to the nearest point) per participant, Twitter. 136
C.3 Part 10: Timeline content proportions (survey and computed, percentage points, to the nearest point) per participant, Instagram. 136
C.4 Part 13: Activity frequency percentages (survey, percentage points) per participant, Facebook. 136
C.5 Part 13: Activity frequency percentages (computed, percentage points, to the nearest point) per participant, Facebook. 137
C.6 Part 13: Activity frequency percentages (survey, percentage points) per participant, Twitter. 137
C.7 Part 13: Activity frequency percentages (computed, percentage points, to the nearest point) per participant, Twitter. 137
C.8 Part 14: Visibility concerns (5-point Likert scale, 5 = most concern), Facebook. 138
C.9 Part 14: Visibility concerns (5-point Likert scale, 5 = most concern), Twitter. 139
C.10 Part 14: Visibility concerns (5-point Likert scale, 5 = most concern), Instagram. 139
C.11 Part 14: Visibility concerns (5-point Likert scale, 5 = most concern), LinkedIn. 140
C.12 Part 15: Visibility control (5-point Likert scale, 5 = most control), Facebook. 140
C.13 Part 15: Visibility control (5-point Likert scale, 5 = most control), Twitter. . 141
C.14 Part 15: Visibility control (5-point Likert scale, 5 = most control), Instagram. 141
C.15 Part 15: Visibility control (5-point Likert scale, 5 = most control), LinkedIn. 142
C.16 Part 17: Connection proportion estimates by interaction level (survey, percentage points), Facebook. 142
C.17 Part 17: Connection proportion estimates by interaction level (computed, percentage points, to the nearest point), Facebook. 142

C.18 Part 17: Connection proportion estimates by interaction level (survey, percentage points), Twitter. 143
C.19 Part 17: Connection proportion estimates by interaction level (computed, percentage points, to the nearest point), Twitter. 143
C.20 Part 17: Connection proportion estimates by interaction level (survey, percentage points), Instagram. 143
C.21 Part 17: Connection proportion estimates by interaction level (computed, percentage points, to the nearest point), Instagram. 144
C.22 Part 17: Connection proportion estimates by interaction level (survey, percentage points), LinkedIn. 144
C.23 Part 17: Connection proportion estimates by interaction level (computed, percentage points, to the nearest point), LinkedIn. 144
C.24 Part 19: Connection count estimates by network grouping (survey). 145
C.25 Part 19: Connection counts by network grouping (computed). 145

Chapter 1

Introduction

Social media platforms are widely used by people to connect and share thoughts – to close friends, to old acquaintances, or to the whole world. Indeed, as of early 2017, nearly 70% of American adults use some type of social media [20]. People use these platforms for many different activities, like keeping up to date with the news, observing the lives of their friends and acquaintances, maintaining social and professional networks, and publishing thoughts, feelings, opinions, and life events to friends, family, and the public. However, people with long-term or high-profile presences on social media can have so many connections, and a record of posts and activities so large, that it is difficult or impossible to remember and reflect on all of them. Such a lack of awareness can have real consequences for one's social and professional life, if by consequence one makes a post that portrays oneself in an unfortunate light or reveals sensitive information to people with bad intentions; indeed, posts can cause problems for their authors even years after they were written [16].

While Facebook remains dominant in the social media landscape, people are increasingly turning to multiple networks to fulfill various types of interactions and needs [29]. Different networks lend themselves to different purposes – professional connection, social interaction, or other activity – and people may consider very different audiences when they interact with each one. Personal awareness of network activity could be very different from one network to another, based on purpose, interaction frequency and history: meticulous curation of one's public Twitter feed does not necessarily imply careful watch over one's Facebook activity, even though either one could be a source of social media liabilities. The social intersections between networks, groups of people among whom different purposes intersect, could potentially present their own problems, if they are primarily associated with a particular audience (e.g. work colleagues) but also have a view into material intended a more private audience (e.g. close friends). Examining how one interacts across the entire social media landscape, then, seems to become important for identifying potential privacy issues.

This project was interested in examining to what degree people are aware in practice of the types of connections, audiences, and information they share across social media platforms. To collect the quantitative data needed to answer these questions, we developed a tool that gathers a wide spectrum of data, including posting history, activity logs, and messaging history, from the social network profiles of willing participants that use multiple social media networks on a regular basis. We then surveyed our participants to gather estimates of certain metrics, including perceived interaction levels and cross-network presence of social media connections, types of data and metadata visible from participant posts, and types of posts and other activity. We also gathered information about participant expectations concerning social media – for what purposes they perceive themselves using each network, how concerned they are about posts and other activity being visible to unwanted people, and how much control they believe they have over who can see their activity. Following the survey, we interviewed participants in order to expand on their motivations and perceptions and to review discrepancies between their answers and the answers we computed from their data. We then presented them with an interactive visualization of their data, highlighting their activity, information, and connections across networks, and asked them to report on any surprises.

We found that people in our test population do tend to ascribe distinct purposes to each different social network, and do tend to behave differently on certain networks in ways consistent with these purposes – both in terms of how they make sharing decisions and in terms of how they interact with each network in practice. However, we also found that participants seemed to believe that they interacted more frequently with their social media contacts (and perhaps more broadly) than they actually did, significantly underestimating the set of people with whom they had no recent interaction at all; we also found that people had varying degrees of trouble estimating what types of content they had posted to social media, and what types of interactions they had conducted in the past. Despite these estimation errors, when presented with the computed data or the interactive visualization thereof, only a minority of participants (4 out of 12) registered surprise, although our visualization allowed them to identify potential issues with posts and contacts that they could then fix afterwards.

This document describes the project. It explains the research questions and hypotheses that the project aimed to address, and similar work previously conducted by others. It then describes the data collection and visualization methods used by our tool, which we named "Polyhedron" as a nod to its goal of examining the many faces people present on social media. The document follows with a description of the user study we performed with the assistance of our tool, the study's results, and our interpretation of the results, and describes the problems and challenges we faced as we progressed through the study. The document concludes with a summary of what was learned and potential directions for future work.

Chapter 2

Research aims and objectives

This project investigated the accuracy of people's self-perceptions of social network audiences across multiple platforms. We were especially interested in how well people can classify their connections by closeness based on interaction across multiple platforms, and how well their sharing intentions align with these classifications and the actual audience reached. This project additionally tried to determine whether people with mismatched perceptions, when provided with accurate information about their social exposure and visibility of personal information from different points of view, were motivated to change sharing behavior on the basis of recognition of oversharing or unintended disclosure.

The project aimed to address the following research questions:

RQ1. How accurately do people perceive the composition of their audiences on popular social networking platforms?

- How accurately can people estimate the distribution of relative engagement with their social networking connections?
- How accurately can people estimate the distribution of connections with presence on multiple social networks – and thus with presence in multiple distinct audiences?

RQ2. To what degree do sharing patterns differ across networks and contexts, and to what degree are people aware of these patterns?

RQ3. To what degree does the availability of composite profile information, generated from all available data, from multiple points of view affect the perception of privacy risks, and alter the importance of privacy in the decision making process governing information sharing?

- How much impact does the ability to investigate one's own composite profile from selected points of view, including from the perspectives of chosen friends, passing acquaintances, and total strangers, have on people's perceptions of privacy risk?
- Where expectations and reality mismatch regarding audience, to what degree do people plan to change their behavior when presented accurate information in an interactive format?

We ultimately tried to test the following hypotheses:

H1. People are not able to accurately estimate the distribution of their connections by frequency of social media engagement across all platforms. (We use frequency of engagement as a proxy for closeness.)

H2. People are not able to accurately estimate the distribution of their connections by which networks they are connected.

H3. There is a statistically significant difference between networks in the frequency with which people share text versus non-text posts, on the networks where both are possible.

H4. There is a statistically significant difference between networks in the frequency with which people tag other people relevant to a post, on the networks where this is possible.

H5. There is a statistically significant difference between networks in the frequency with which people attach location information to a post, on the networks where this is possible.

H6. People are not able to accurately estimate the composition of their own social media timelines by post type.

H7. People are not able to accurately estimate the proportions of activity types they undertake on each network.

H8. Viewing of one's own visualized composite profile data, along with information about audience and reach, with the ability to highlight the set of information visible to any particular connection or category of people, results in a measurable shift in one's opinions about privacy risk.

2.1 Perceptions of audience composition

As part of our work, we wanted to determine how well a person's mental model of their own audience on a social network matches up with the reality – that is, how well they can estimate the types of people they reach when posting. Previous work by Bernstein et. al. [1] suggests that a large portion of unintended sharing occurs because one forgets or does not realize the scope of audience to whom they are broadcasting information about themselves.

For our purposes, we separated connection types along a combination of interaction frequency and network purpose. We supposed that the frequency with which one interacts with a given connection could act as a measurable proxy for degree of social closeness, and so partitioned connections along these lines:

- "Close friends": those connections with whom a large number of interactions, corresponding to consistent daily activity, were recorded over the past month
- "Friends": those connections with whom at least a minimal number of interactions were recorded over the past month

- “Acquaintances”: those connections with whom at least one interaction was recorded over the past year
- “Mostly strangers”: everyone else

We treat connections connected on a professional basis as a special case, since the ramifications of unwanted sharing to this audience may be much more serious than unexpected sharing to other audiences. We apply the label “colleagues” to any connections who are not considered “friends” or “close friends”, but who are connected through LinkedIn[1].

We also examined awareness of social network connections who are connected on more than one network. Resharing by connections from one network to another is a potential reason why information may be shared to audiences other than intended; indeed, more conscientious social network participants do report being concerned about the possibility when making sharing decisions [29].

A result suggesting that people are not very accurate at estimating audience composition, relative engagement and interaction, and presence on multiple networks (H1 through H2) would allow us to attribute some amount of unintended sharing to such inaccuracy. On the other hand, a result suggesting that people can make accurate estimates of these quantities would suggest that some other phenomenon is mainly to blame.

2.2 Information sharing patterns and perceptions

We hypothesized that the nature of content shared on a given social network varies based on the perceived target audience reached by posting on that network. To test this hypothesis, we needed to determine what types of content are shared by people on each of these networks. The trouble with this goal is that many categorizations by which content can be classified, e.g. by political leaning, by target interest group constituency, or by overall sentiment

[1] As professional connections can be acquired independently of employment history, often in ways that are difficult to infer automatically, we use an inclusive metric here.

metrics, are difficult to determine in an automated fashion. While solutions exist to evaluate text-based content based on machine learning techniques, none appear to be available for local use, and we wanted to avoid sending participant data to remote servers for this work. Automatic interpretation of non-text content, like photos and videos, runs into similar issues.

Because of these limitations, our focus is limited to the high-level type of content (text vs. graphical) and the presence and absence of certain objectively measurable features. Specifically, we look at the following:

- Is the post primarily text / link, primarily a photo, or primarily a video?
- Does an external link, photo, video have a caption?
- Has the poster tagged anyone?
- Does the post have embedded location data?

We are further constrained by the fact that many networks do not implement all of these features; for instance, it is impossible to make a text-only post to Instagram, and it is impossible to attach location data to a LinkedIn post in a standard way. The hypotheses we are able to test are thus limited to comparisons among networks for which the relevant features exist and for which we are able to collect the necessary data.

Quantitative data affirming the hypotheses we are able to test (H3 through H5) would lend credence to prior work suggesting that perceived target audience affects sharing decisions [29]. On the other hand, failure to support these hypotheses would complicate the picture, suggesting that while people believe that target audience is important in making sharing decisions, there need not be a large impact in practice, at least on the quantities we are able to measure.

We additionally wanted to determine the degree to which people can estimate their own activity and content on social media. Posts can become problematic as a result of life events or world events, even years after they are initially contributed. For reasons similar to the above, we limited our evaluation to high-level content types – text, videos, and photo posts. In terms of activity metrics, we identified six common types of activity – posting, commenting on one's own posts, reacting to one's own posts, commenting on others' posts, reacting to others' posts, and resharing others' posts – that can be performed using timeline entries; we limited our evaluation to number of individual instances of these, rather than attempting to quantify time spent and/or looking at other activities like reading posts, as we found it impossible with our methods to capture historical values for the latter quantities.

Quantitative data affirming the hypothesis related to activity estimation (H6, H7) would suggest the potential existence of a source of problematic posts or other activity: latent posts made, or activity performed, and subsequently forgotten. On the other hand, failure to support this hypothesis would suggest that problematic disclosures likely arise for other reasons – either actualizing over short time spans, as investigated by Wang et. al. (2011) [26] and Sleeper et. al. (2013) [24], or resulting from risk miscalculations despite awareness.

2.3 Effects of data visualization on perceptions of privacy risks

Prior work, especially work by Bernstein et. al. [1] and by Dunbar [4], has highlighted the frequent mismatches between a person's perception of share audience and their actual information exposure. Presently, it is not always easy to estimate the visibility of information one shares on a given social network. One reason for this is that privacy settings themselves are usually de-emphasized visually. Facebook's privacy settings for posts are expressed as a single icon, requiring a mouse-over and/or click to display more information about any subleties that may have been specified beyond generic defaults; for other types of information, like education history items, privacy descriptors are hidden entirely unless the data element itself is moused over. On other platforms with less granularity over sharing, indicators are less conspicuous or even missing; for instance, Twitter indicates that a profile is public by the *absence* of a lock icon, and Instagram provides no visual indication whatsoever within its Web product. Hence, one might find it difficult to ascertain at a glance which privacy settings have been applied to a particular post, simply because the necessary visual indicators are not given prominence. Arguably, the difficulty of evaluating privacy exposure works in favor of the social media platforms, whose business model depends on people sharing information, as other work [23, 10] has shown that actively calling attention to privacy-related matters can discourage disclosure.

We hypothesized that providing a reasonable retrospective analysis capability for past posts, based on data harvested from multiple social media profiles, would help people better understand their privacy risks. Fundamentally, it can be difficult to recall posts made to social media in the past; being able to review the entire post stream at a glance, and identify specific points of concern, can be useful. Since many privacy risks may also arise due to cross-network interactions, we thought it would be useful to provide visibility information for contributed content, like posts and location data, as a single unified stream. Additionally, since specific networks and specific types of connections present different types of privacy risks based on potential consequences of oversharing, we wanted to provide an easy means of highlighting what information is visible from any chosen perspective, be it from the perspective of the public, total strangers, acquaintances, friends, connections on specific networks, or even specific contacts of interest.

As a concrete example, we describe a scenario based on concerns that Schoenebeck et. al. (2016) identified as being held by many young adults. Through interviews with college students using Facebook, they discovered that such people's usage patterns of Facebook changed significantly over time in accordance with personality changes during adolescence, and that very old posts can be a source of some embarrassment [22]. Given that control of presented image is a desirable goal in some circumstances, e.g. during application for employment [25], the presence of misrepresentational or compromising posts visible to certain audiences, e.g. a nascent professional network, would naturally be concerning. For the scenario of determining whether problematic posts are visible to potential employers, one should be able to select a viewpoint associated with this group – depending on the specific worry, this viewpoint may be the public, the set of LinkedIn connections, the set of connections on both LinkedIn and Facebook, or even a specific individual – and quickly determine what posts, from what points in time, are visible. If the set is relatively small, one should be able to inspect individual posts to see if they would present issues; if the set is large, one should be able to, at least, see what topics are discussed, and make a determination as to whether more careful evaluation is needed. When considering visibility from sets of connections, it should also be possible to see how large those sets are, and who the members are, to help evaluate whether the category presents a privacy risk in practice.

A result in favor of H8 would be in line with prior work suggesting that highlighting potential privacy problems can mitigate the occurrence of oversharing. A result not in favor of H8 would unfortunately be less informative: although it is certainly possible that providing easy visualization makes no useful difference, it is perhaps more likely that some deficiency exists in the type of visualization we are able to provide.

Chapter 3

Related research

This chapter describes previous findings in the following research areas aligned with our research objectives:

- Perceptions of audience composition: Researchers have examined the underlying reasons for oversharing of information on social media, and linked the phenomenon to misperceptions of who would see or be interested in particular posts.
- Information sharing perceptions and patterns: Researchers have examined people's sharing habits and intentions, and investigated the degree to which they align.
- Visualization as a tool for understanding privacy issues: Some work has been done on effective visualization techniques for helping people understand privacy issues and make more effective sharing decisions.

3.1 Perceptions of audience composition

Perception of audience is an important contributor to the phenomenon of oversharing; prior work has examined how people view their audiences on social media and how those perceptions can lead to problems. It is already known, due to Dunbar (2012), that people make many more social media connections than they could possibly associate with a genuine emotional connection, including connections with acquaintances who may be recognized by face but have little in common with oneself [3]; indeed, Dunbar (2016) finds by surveying people who use Facebook that while an average such person may have around 150 connections, only 14 connections or so can be expected to have an emotional association beyond general superficial awareness [4]. One can hypothesize that many posts made to social media are written mainly for the consumption of that "inner set", especially those posts with strong sentiments, secrets, or even falsehoods. Wang et. al. (2011) find through interviews of people using Facebook that, in many cases of oversharing, the unfortunate party has trouble remembering or knowing who the post will reach when they make it; while they may envision a particular audience, they often report not having been aware that their post could have spread beyond that audience [26]. Sleeper et. al. (2013) reach similar conclusions through surveys of people using Twitter [24].

Exactly who do people envision as their audience when they post? Litt and Hargittai (2016) attempt to answer this exact question, using a 2-month diary study and interview with posts collected over the study duration from Facebook, Twitter, and LinkedIn. They find that, in roughly half of cases where posts are shared to large sets of friends or the general public, people don't envision a specific audience at all; among the remainder of cases, people imagine audiences ranging from personal ties to specific communities to professional connections, though typically only one such grouping at a time [14]. Marwick and boyd (2011) further discover, through survey of people who use Twitter, that many actively reject the notion of audience altogether, believing that tailoring posts to audiences is somehow "inauthentic" [17].

Do these audiences vary by platform? Zhao et. al. (2016) find that they often do: through interviews, they discover that people, in fact, have a tendency to treat each network as a separate audience, and ascribe different norms and modes of interaction to each [29]. Given that many posts, even with specific audiences, are in practice shared to large groups, how widely do they spread in practice? Bernstein et. al. (2013) address this question on the Facebook network. By surveying people who use Facebook, as well as collecting information (with Facebook's cooperation) about how often posts are viewed, they found that people drastically underestimate their effective audience, believing the number of people who have seen a given post to be less than a quarter of the actual count [1]. Hence, while people may envision specific targeted audiences in some cases, it is quite common for posts to exceed that audience expectation, spreading beyond the set of people they thought would be interested.

We observe among this selection of prior work that comparatively little work has been done across multiple networks, with Litt and Hargittai (2016) being the only study (that we were aware of at the time we did the work) using actual data collected from more than one network. Bernstein et. al. (2013) effectively leverage access to an impressive amount of quantitative data but, due to the means by which it was gathered, are limited to making conclusions from a single specific network. Most other available research relies primarily on participant-reported survey and interview data, suggesting an opening for work combining both qualitative and quantitative data across multiple networks, as we attempt to achieve.

3.2 Information sharing perceptions and patterns

People share all kinds of information on social media, like biographical details, life events, opinions, and issues they find important. This information is shared through a variety of modes, like posts, reactions, comments, and reshares of others' posts, for a variety of reasons like eliciting reactions from others for self-validation [12], expressing solidarity (consciously or otherwise) with their friends [5], or "venting" about highly charged opinions or events as a means of catharsis [26]. A wide body of research has examined people's sharing habits and how they relate to privacy concerns.

How much of a factor is privacy in determining what to share? The answer to this question has evolved somewhat over time. Gross and Acquisti (2005) found through examination of

Facebook profiles that privacy concerns seemed to be largely absent, and that people freely shared all types of information without changing privacy controls from fairly permissive default settings [8]. On the other hand, Johnson et. al. (2012) found through survey questions, supplemented by individual post and connection information, that the picture had changed in many respects since the work by Gross: people were much more concerned about having information viewable by strangers, and the majority had taken steps to mitigate such unwanted disclosure; however, although some people were worried about unwanted disclosure to connections, few had taken effective steps to address the "insider threat" [10]. What kind of specific risks, if any, do people have in mind while sharing? Vitak and Kim (2014), through interview of people using Facebook, found a variety, including social rejection, presenting an unfavorable self-image, and disclosing information to the wrong parties [25]. How do people mitigate these risks? Vitak and Kim (2014) found that people used a variety of introspective strategies, including careful curation of friends, usage of the built-in friend grouping feature, and sometimes-lengthy consideration of possible collateral effects prior to posting [25], while Cho and Filippova (2016) discovered through focus-group interviews with people using Facebook that collaborative strategies, like discussing sharing norms out-of-band, are also employed [2].

Do people share differently across networks? Zhao et. al. (2016) suggest that they do. Through interviews, they find that different networks are used for different modes of interaction. In their test population, they discover that Facebook is used mainly for personal interactions even among coworkers, is often used for highly curated content, and is used with the expectation of receiving feedback, while Instagram posts are made with much less selectivity and Twitter is used more as a broadcasting tool [29]. Even when they post about a particular idea or event on multiple networks, they tailor the message to each perceived audience or set of platform norms.

A review of the various studies in this space suggests that most address a single network at a time, with Facebook being a favorite – perhaps sensible in light of its large uptake in comparison to other networks [20]; few studies compare behavior across networks, and those that do rely on qualitative data like surveys and interview results. Thus, there is an opening for investigating quantitative differences in sharing intentions, perceptions, and patterns across different social networks.

3.3 Visualization as a tool for understanding privacy issues

Using data visualization as a tool for understanding privacy-related issues is not a new idea, and indeed, there are some noteworthy efforts to use visualization-based techniques as a way to convey understanding of privacy risks on social networks. Notably, Paul et. al. (2012) found that people who use Facebook find its privacy controls “confusing” [19]; their tool, C4PS, uses color coding and ready access to privacy settings in-line with data, and they determine by user study on synthetic Facebook data that their methods help people more effectively understand and set privacy constraints than Facebook’s own tools. Mazzia et. al. (2012) concentrate on the problem of group management on social networks; their tool, PViz, visualizes composition of friend networks and profile element visibility, allowing people to understand and measure what is visible to different groupings of Facebook friends, and they find through user study on synthetic and actual Facebook data that their methods work well for these purposes [18]. Both of these tools concentrate on Facebook specifically, and other efforts like Wang et. al. (2015)’s VeilMe [27] and Fang and LeFevre (2010)’s work on privacy wizards [6] likewise choose a single network on which to focus. There is room, then, for visualizations that provide understanding of information visibility from a cross-network perspective.

Chapter 4

Polyhedron: data collection and visualization

One of the difficulties of objectively answering any research question related to social media usage is that much of the data itself is not readily available in a conveniently consumable form. While all four of the social media platforms being targeted do provide APIs (application programming interfaces) of varying usefulness for third-party programs, there is some data that is not provided through this means; for instance, the Facebook API, when queried for a friends list, only provides those friends who have installed the same application, which was unacceptable for our needs. Additionally, usage of the API often requires review of the intended purpose by the API provider – which, while a reasonable safeguard against abuse, would cause problems for our intended use case should the provider withhold approval.

To collect data for our needs, we instead constructed a browser add-on, which we named Polyhedron, for the Mozilla Firefox web browser. Add-ons for Mozilla Firefox have several useful features: they can access and manipulate the contents of open web pages while using the active session's cookies; they can create 'page workers', which can load and manipulate pages invisibly; and they can interact with the local file system, allowing data captures to be saved for later processing and reloaded for updating. These capabilities permit us to collect information from the local session's perspective, and also present information by means of updating a browser tab shown by the add-on. We can then create tools that gather information from a person's social network accounts and present summaries and analysis for inspection.

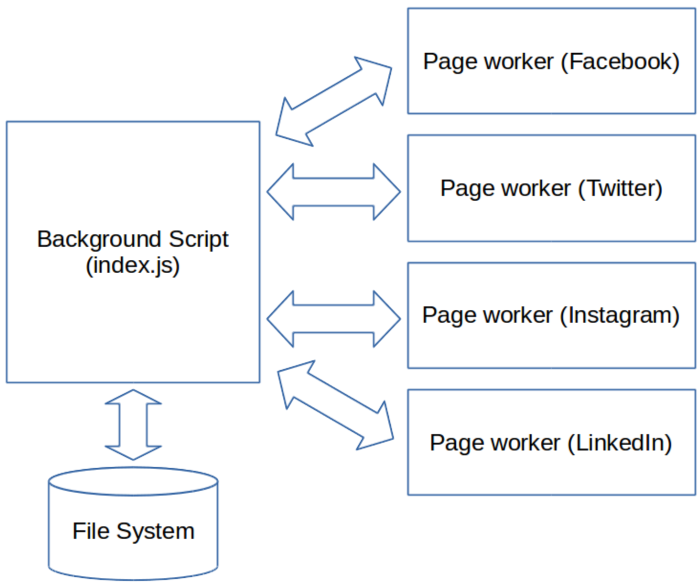

Figure 4.1: Overall architecture diagram of the Polyhedron add-on.

4.1 Data collection

There are several potential methods by which data may be gathered from each social media site targeted by our tool. Most simply, information may be extracted from the actual HTML markup shown during normal operation; three of our networks (Facebook, Twitter, and LinkedIn) are used primarily through a Web interface, with the fourth (Instagram) providing reasonably comprehensive interaction through the Web. It is also possible to reverse-engineer, through examination of the communication protocol, the means by which each Web page requests and presents data; the browser-based interfaces of all four networks all extensively use JavaScript to fetch partial data through Web-based APIs, and the protections on these APIs, for our purposes, are scant at best. Finally, the option exists to query some data from APIs provided specifically as interfaces for third-party programs. We did not explore this last option because of the need for approval from each platform and, in LinkedIn's case, the total inaccessibility of information beyond basic biographical data without a specific partnering agreement with LinkedIn itself.

For each platform, we use a combination of HTML scraping and direct API calls, depending on what is most convenient for a particular datum. Our interface to each platform begins with one or more 'page workers' set to navigate to pages of interest. Each page worker loads the target page in the background; a content script attached to each worker determines whether an active session exists, and if so, performs the necessary data collection. Where a known API call exists, we can perform it from the context of the loaded page using an 'XMLHTTPRequest', and parse the response accordingly. In the case where a usable API is missing or difficult to understand, we can parse the HTML markup directly and extract the necessary information. To avoid generating large amounts of traffic in a short period of time, and maintain API access within the bounds of human-like activity, we limit the number of API calls we make in a given time interval.

When a page worker retrieves a particular set of data, it sends a message to Polyhedron's main thread, containing the retrieved data in minimally processed form. The main thread then incorporates the data into a central model aggregating information from all available social networks into a single thread.

Many data has to be gathered by calling multiple HTTP endpoints in succession; for instance, many data are available only by navigating through a pagination scheme, and many other data had to be retrieved from an individual page per datum. We decided to limit the rate at which such data were gathered in an attempt to avoid rate-based "flagging" from the various service providers. In support of this, every page worker has an internal scheduler; for every datum or set of data needing a separate HTTP request, a new item is created on the scheduler, and the page worker itself consumes and processes these items at a rate of about 1 item per second. The scheduler periodically saves its state to disk

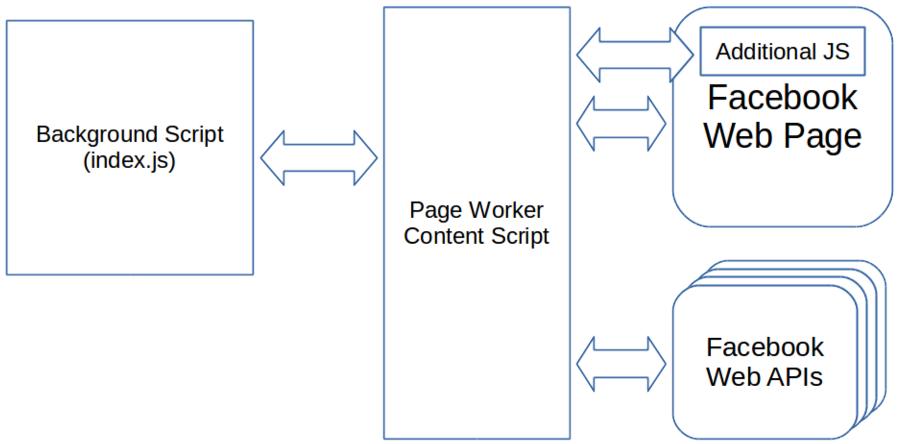

Figure 4.2: Overall architecture diagram of the Facebook page worker. Other page workers are similar.

and attempts to reload its state at tool startup, allowing data collection to resume from a reasonable point in time if interrupted.

Crucially, a page worker has the ability to use the active browser session's cookies. This property is essential to our data gathering methodology. As long as the browser has cookies for an active login session for a given social network, we are able to load pages and make API calls as if the person who created the session had done so themselves. Our positioning on a participant's browser thus allows us to gather information from each volunteered social network without needing to register ourselves explicitly with each platform.

In the following segments, we describe the data we collected from each network, as well as some specific remarks regarding collection for each network. Where we provide specific API details for particular networks, those details were current for data collection activities undertaken in the first half of 2017, and are likely to be out of date; replicating the work would require examining whether or not endpoints or mandatory parameters have changed in the meantime.

4.1.1 Facebook

For gathering data on Facebook, we wrote a page worker that interprets a combination of HTML and JSON information returned from Facebook's internal endpoints. We adopted this approach rather than use Facebook's normal API because many classes of information, like a participant's full friends list, are not available through the normal API, and because Facebook's normal API requires platform-level approval for meaningful use (which we believed might be difficult to secure for our use case, and thus did not pursue).

Facebook's primary interface relies heavily on a technology known as BigPipe, which delivers websites piecemeal in order to improve perceived responsiveness. A typical Facebook page consists of several UI containers, which may or may not be initially populated; when a container needs to be changed, the entire HTML markup needed for the alteration is pushed from the Facebook server and transplanted into the corresponding segment of the visible webpage. The response from the server also contains encoded calls to named procedures in Facebook's client-side JavaScript, which, among other tasks, updates the client state so that the next change can be requested correctly when needed. Additionally, the response contains a significant amount of information related to page styling and external resources, most of which is not interesting for our purposes. All of this data is sometimes delivered as fields in a single JSON object, and sometimes delivered as elements and `<script>` tags in a large (100 KB or more) HTML fragment, according to criteria not yet fully understood.

This architecture describes the operation of most of Facebook's website. Certain elements, however, do not quite operate as described. Some information, like comments and replies, is uniformly delivered as encoded JSON objects, containing the data necessary to reconstruct the HTML markup -- which is done client-side – rather than providing the markup itself. For instance, Facebook comments are delivered in the form of objects containing, among other information, comment text, author, and timestamp, rather than the corresponding HTML elements that are actually displayed.

Our Facebook page worker is able to handle both types of information. We parse JSON-encoded data directly, and add the necessary pieces to our own data model. Data encoded inside BigPipe pagelets can be retrieved, using the same API calls as the actual Facebook webpage, and parsed from the HTML markup. Because most of Facebook's generated HTML is hierarchical in nature, has stable – if often nondescriptive – element class names, and uses HTML tags in semantically appropriate ways (for instance, `<ul>` for unordered lists of friends), we can easily parse through the markup by using the jQuery library along with the standard DOM (Document Object Model) API provided by Firefox.

From Facebook, we are able to retrieve biographical information from the person being evaluated, as well as the person's social connections, posts, private messages, and Facebook activity. Additionally, we are able to get the privacy descriptors for many of these data. Facebook permits fine-grained permission controls to be applied to individual posts; one can specify individual connections or connection groups to include or exclude from the set of people who can view a given post.

To aid in our collection of data from Facebook, we embed a small amount of custom JavaScript into the context of the Facebook page content that lets us directly call some of the Facebook client library functions. We do this to retrieve the correct values for certain parameters that must be passed along with GET and POST requests to endpoints. We determined that this approach was necessary, as some required parameters cannot be retrieved purely by examining the page DOM.

Because most elements of Facebook are delivered in "cooked" form – that is, as rendered HTML rather than as raw data – it is not always easy to extract meaning from the information delivered. This is especially true when groups of data are present in a single constructed string with no additional semantic labeling and no clear indicator of which data, if any, are missing. For example, a workplace in Facebook is annotated with a job title, a start and end date, and a location, delivered in a single string without further labeling; if exactly one of the three are provided, it is not always easy to determine the difference between a job title and a location. There are means to circumvent this restriction – by, for example, examining the markup delivered in response to a request to edit the items of interest, and parsing through the delivered HTML form – that were not pursued due to lack of time and the forecasted additional complexity.

With our methods, starting from a page worker pointed at `https://www.facebook.com/me`, we are able to gather the following types of data from Facebook:

- Basic profile information, like full name, date of birth, and education and work history may be retrieved by making a POST request to `/profile/async/infopage/nav/` with the URL parameters '`profile_id`' and '`viewer_id`' set to the active profile's numerical ID and '`section`' set to one of "`overview`", "`education`", "`places`", "`contact_basic`", "`all_relationships`", "`about`", or "`year_overviews`". These sections correspond to the biographical tabs in '`/me/about`' describing an individual's basic information.

- Timeline posts may be retrieved by scanning '`/me`' for timeline data and making repeated API calls to an appropriate endpoint depending on which version of Facebook is being served to the current session. We identified three distinct endpoints: '`/ajax/pagelet/generic.php/ProfileTimelineSectionPagelet`', '`/timeline/jumper/async/`', and '`/profile/fig/timeline`'. In every case, post data may be recovered from the returned HTML content.

- For individual posts, privacy descriptors (visibility information) may be retrieved by making a POST request to '`/privacy/custom_dialog/`' with a number of different parameters, including the post ID. This API endpoint is normally called by Facebook to request a dialog for changing privacy settings. From the markup of the requested dialog, we can retrieve the list of people to whom a post is visible or not, as well as whether connections of tagged people should be able to see the post.

- The list of connections may be retrieved in two stages. Initially, we make a GET request to '`/{screenName}/friends`', where `{screenName}` is replaced with the actual "screen name" associated with the profile. The initial request requires certain URL parameters, like ‘dpr‘ set to 1 and ‘ajaxpipe_token‘ set to a token retrieved using a script injected into the page worker context. From the metadata returned along with the initial request, we can retrieve a collection token and cursor to pass along with GET requests to ‘`/ajax/pagelet/generic.php/AllFriendsAppCollectionPagelet`‘ for subsequent pages. In each case, data about individual connections can be extracted from the returned HTML content.

- Private messages may be retrieved by repeated POST requests to ‘`/ajax/mercury/thread_info.php`‘. This endpoint expects the connection profile’s ID, an offset, a timestamp, and a limit on retrieved message count (which the official Facebook client appears to always set to 20). Profile IDs may be discovered from metadata embedded in the connection list. The offset may initially be set to 0 and thereafter incremented by the number of posts retrieved, and the timestamp may initially be left blank and updated with a timestamp value included with the response.

- The Facebook activity log is presented in multiple sections, one per month. From a GET request to ‘`/{screenName}/allactivity`‘, blocks of metadata for all sections may be retrieved. For each section, we pass its metadata block to ‘`/ajax/pagelet/generic.php/TimelineEntStoryActivityLogPagelet`‘ to get the first batch of log entries for the section; if subsequent log entries exist, an HTML element with the class ‘`a.uiMorePagerPrimary`‘ contains the correct URL for the next batch of log entries in its ‘`ajaxify`‘ attribute. The activity log describes all active actions the participant has taken on Facebook, including posts, comments, reactions, and shares; we use this information to compile interaction statistics with particular connections as well as post type classification statistics.

4.1.2 Twitter

For Twitter, we wrote a page worker[1] that interprets a combination of HTML and JSON data returned from Twitter's internal API. While Twitter has an API for third-party applications with fairly generous terms of usage, we opted for this approach to maintain consistency of architecture and operation with the other platforms.

Although Twitter, like Facebook, pushes server-side-generated HTML fragments for incrementally loaded resources like the person's timeline and followers list, it does so in a much more concise fashion – as a JSON object usually containing only the necessary HTML fragment and some pagination information. The Twitter HTML markup is friendly to automated parsing: its class names are descriptive and, like Facebook, it uses HTML tags in semantically appropriate ways. Hence, we can use jQuery, along with the standard DOM API, to extract the desired data without too much difficulty. We gather most of the Twitter data we use from calls to the frontend API, with profile information lifted directly from the web page itself.

The information we gather from Twitter includes the limited set of biographical information (name, location, brief bio) provided in the profile, as well as a list of followers, accounts followed, and post history. While Twitter also allows location data to be attached to any specific post, we do not collect this data due to lack of time and required additional complexity: in the Web application, location data is displayed only in the detail view for an individual post, and not in the post stream overview, which would require an extra API call per post to retrieve.[2]

With our methods, starting with a page worker pointing at `https://twitter.com`, we perform the following steps, pulling out the following data:

[1] There are actually two separate page workers - one for collecting basic profile information only, and one for collecting everything else. There is no particular technical limitation that makes this necessary; in early versions of the tool, separate page workers were created for each API endpoint, and while most other page workers were merged together for performance reasons, these page workers were never merged.

[2] In practice, we added location annotations by using a postprocessing script fetching posts from the official API.

- The logged-in profile page is loaded by searching the Twitter front page for a link with CSS selector ‘`.DashboardProfileCard-name > a`‘ and sending it a click event.
- Basic profile information is retrieved directly from the profile page DOM; we extract the full name, screen name, biographical description, locality description, and whether or not the timeline is protected.
- We retrieve the post timeline through repeated calls to ‘`/i/profiles/show/{screenName}/timeline/with_replies`‘, where `{screenName}` is the screen name extracted earlier. Among other parameters, we provide a parameter ‘max_position‘ that functions as a cursor. This parameter is left blank for the most recent set of posts; for progressively earlier batches of posts, we take the ‘min_position‘ value provided with the returned data and use it as the ‘max_position‘ for the next call.
- The list of followers is retrieved through repeated calls to ‘`/{screenName}/followers/users`‘. This endpoint also requires a parameter ‘`max_position`‘ that functions analogously to the same parameter for the post timeline endpoint. For each follower entry in the HTML fragment returned with the API response, we are able to retrieve a numerical ID, a screen name, and a full display name.
- The list of people being followed is retrieved through repeated calls to ‘`/{screenName}/following/users`‘, analogously to the list of followers.

4.1.3 LinkedIn

The LinkedIn product, unlike Facebook or Twitter, provides only a basic third-party API for most developers; direct API access to their platform at the level afforded by other platforms requires an official partnership agreement. After reviewing the information LinkedIn provided about its partnership programs, and failing to discover any official mention of research usage[3], we elected to employ data collection techniques resembling those we used on the other platforms we investigated.

Before January 2017, LinkedIn's website most heavily resembled Facebook's website, in that information was primarily delivered in the form of rendered HTML markup. Hence, we extracted information from a person's LinkedIn account with means similar to that by which we did so from Facebook; that is, we queried well-known URLs and parsed the necessary data from a page DOM, adding them to our data model.

Sometime in January 2017, LinkedIn revealed a complete overhaul of their Web product, replacing the previous architecture with one centered largely on client-side rendering of JSON-encoded information. Data are accessible exclusively through calls to a set of endpoints exposed by LinkedIn's "Voyager" API, which was previously restricted to the platform's mobile product. Every object in the LinkedIn data model has a unique ID; by default, the LinkedIn JSON encoding encodes references to objects using this ID string, and provides all of the object definitions needed to interpret an API result as a list alongside the requested information in one HTTP response.

Because of the very large number of API endpoints often needed to render a single page, the LinkedIn Voyager API includes the ability to batch requests using a special endpoint '`/voyager/api/mux`'. This endpoint accepts a keyed collection of HTTP request descriptions, including endpoints, parameters, and additional headers, and returns a keyed collection of JSON responses. We use this feature in any case where the LinkedIn website uses it; we found that some of LinkedIn's APIs are only accessible via this batching method, and do not otherwise appear to be callable.

From LinkedIn, we fetch a participant's education and work history, skills and skill endorsements, and professional connections, along with the location of each connection. We are additionally able to pull some posts that the participant has made to LinkedIn;

[3]Indeed, anectodal evidence (e.g. `https://stackoverflow.com/questions/31141654/access-to-official-linkedin-data-for-an-academic-research-purpose`) suggests that LinkedIn does not generally provide access to data for academic research use.

however, there appears to be a time horizon associated with what can be retrieved, and some classes of information associated with posts (like location) are simply unavailable because they correspond to features that do not exist on the LinkedIn network. In more detail, this is what we retrieve with our methods, starting from a page worker pointed at ‘`https://www.linkedin.com/feed/`‘:

- Basic profile information, like full name and screen name, can be requested from the endpoint ‘`/voyager/api/me`‘, if not included in the metadata provided with ‘`/feed`‘.

- More detailed profile information is requested from the endpoint ‘`/voyager/api/identity/profiles/{screenName}/profileView`‘, where `{screenName}` is the screen name retrieved above. We specifically pull work experience information, education, and locality and industry description from this endpoint. Additionally, we retrieve the total number of skills associated with the profile.

- Connections can be requested from the endpoint ‘`/voyager/api/relationships/connections/`‘, with the parameters ‘`count`‘ set to 40, ‘`start`‘ set to the desired paging offset, and ‘`sortType`‘ set to ’`RECENTLY_ADDED`’. For each connection, we retrieve a full name and a screen name, and then make a separate call to ‘`/voyager/api/identity/profiles/{screenName}/profileView`‘ with the connection’s screen name to fetch location information. It is possible that fewer than 40 connections are returned even though more connections exist; the retrieval of fewer than 40 connections is no indication that there are no more connections to be found.

- The number of mutual connections a person has with some connection can be requested from the endpoint ‘`/voyager/api/identity/profiles/{longformID}/memberConnections?q=inCommon`‘, where ‘`longformID`‘ is an internal identifier specific to each connection.

- The list of skills can be retrieved by GET request to ‘`/voyager/api/identity/profiles/{screenName}/featuredSkills`‘, with URL parameters ‘`includeHiddenEndorsers`‘ set to `true` and ‘`count`‘ set to the total number of skills included in the profile. For each individual skill, the list of endorsers can be retrieved by one or more GET requests to ‘`/voyager/api/identity/profiles/{longformID}/endorsements`‘, where ‘`longformID`‘ is an internal identifier for the active profile; required URL parameters include ‘`q`‘ set to ’`findEndorsementsBySkillId`’, ‘`skillId`‘ set to the skill’s internal ID, ‘`pagingStart`‘ set to 0, ‘`count`‘ set to 20, and ‘`start`‘ set to the current paging offset (initially 0).

- Any posts made can be retrieved by GET request to ‘`/voyager/api/feed/updates`‘, with URL parameters ‘`count`‘ set to 5, ‘`moduleKey`‘ set to ‘`member-shares:phone`‘, ‘`profileID`‘ set to the active profile’s internal identifier, and ‘`q`‘ set to ‘`memberShareFeed`‘; if applicable, ‘`paginationToken`‘ is set to a token received in the response metadata, and ‘`start`‘ is set to the current offset in the feed.

4.1.4 Instagram

Instagram is a ’mobile-first’ application, so its Web product provides only a limited subset of functionality, and some classes of data that would otherwise be interesting are not available through the Web interface. The data accessible through the Web product can be accessed almost exclusively through calls to internal API endpoints, which serve plain JSON objects that client-side JavaScript would ordinarily reconstruct into HTML. We can usually add the retrieved information directly to the composite data model with minimal processing.

Our decision to use Instagram’s internal Website API does, as remarked, prevent us from capturing certain types of information. For instance, the list of photos that a person reacted to or commented upon is not accessible at all through the Web product, and would have to be determined through reverse-engineering the mobile API, which is protected by a secret key that would need to be determined by decompiling the mobile application. While this task is not impossible, doing so would have potentially added a large recurring workload, as the mobile application itself would have to be proactively monitored for updates, and any changes in the secret key disseminated to copies of the tool; usage of the mobile API, from what we determined from preliminary investigation, appears to require direct capture of login credentials, of which we had no wish to take possession; additionally, we supposed that accessing the mobile-only API from a browser would itself attract unwanted attention. (While Facebook, Instagram's current parent company, has granted privileged access to researchers in the past, we did not pursue this avenue, and cannot comment on the ease of receiving such access.)

From a page worker starting at '`https://www.instagram.com`', we can retrieve the following specific types of information:

- The Instagram screen name may be determined by finding the profile link on the page (with CSS selector '`.coreSpriteDesktopNavProfile`') and parsing it out from the link URL.

- Basic profile information may be retrieved by GET request to '`/{screenName}/?__a=1`', where screenName is the screen name determined above. The returned data includes an internal ID and the most recent few uploaded media items.

- Media items older than those returned with the basic profile may be retrieved by repeated GET requests to '`/graphql/query`' with URL parameters '`query_id`' set to a specific magic number (17880160963012870), '`id`' set to the internal profile ID, and '`first`' set to the required offset within the stream. Note that, while the provided magic number was accurate when we conducted the work, we have no evidence to suggest that it is stable.

- The list of followers may be retrieved similarly to media items with a different '`query_id`' magic number (17851374694183129).

- The list of people being followed may be retrieved analogously, with yet another '`query_id`' number (17874545323001329).

4.1.5 Challenges and limitations

Determining how to interpret a particular JSON result or DOM tree ended up being a much more laborious task than initially expected. The Firefox developer tools provide a network monitor allowing individual HTTP requests to be examined in detail, displaying request and response headers and showing JSON responses in tree view. However, this tool is often not as useful as one would hope: it is unable to provide insight into very long HTTP responses; it is unable to make sense of non-standard JSON; and when a JSON response contains a very long HTTP string as part of its content, the Firefox network monitor is unable to provide a further breakdown of its content. Hence, in most cases, the response text must be copied into a text editor and inspected manually – unfortunately, responses can be several hundred KB long and often lack line breaks, and currently available text editors, particularly editors with syntax highlighting, tend to perform very poorly with long lines. As a result, arriving at the correct interpretation of a HTTP response can take longer than one might expect given the actual useful information contained.

To make matters less favorable, the target platform operators have a propensity to enact changes to their platforms without prior notice. Social media products are constantly iterating on product features and design, and do not generally publish patch notes for minor front-end changes. Thus, breaking changes to the DOM or JSON schema of a particular page can arrive without notice to us, requiring us to repeat the interpretation process upon discovery of the change. One such change made by LinkedIn, involving a total replacement of the frontend product with a codebase built on completely different principles, required the re-tooling of the entire page worker for that network. The other networks have all made less drastic, but not much less inconvenient, changes to how their webpages are laid out and/or what information they choose to return in response to requests. Keeping up with these changes has proven a chronic challenge in practice; we expand on specific instances in Chapter 6.

Furthermore, there are challenges with interpreting the data itself to which we were unable to find good solutions. For instance, it is a relatively common phenomenon that people may use a different name on Twitter or Instagram, which are not strongly tied to a real-world identity, than on Facebook or LinkedIn. This is not a problem for assembling a composite profile for the people running the tool, but presents a challenge for deduplicating their connections. Thus, we underestimate the proportion of connections who are connected on sets of networks spanning both categories, and overestimate the proportion of connections who are not. Also, many of the metrics we were interested in evaluating are not explicitly available on all networks; for instance, while it might be possible to infer a person's religious or political views from inspecting their Instagram stream, doing so would require a degree of image processing ability beyond the scope of our project. Hence, we are not able to objectively evaluate participant awareness of exposure of certain classes of data on every network.

4.2 Data presentation

In addition to collecting information from the social media networks we chose, the Firefox add-on we created also presents that data in a browser tab for review. We present a summary of the participant's activity on each network, as well as a list of their connections. We also, to the extent we are able, permit the participant to filter the data by what is visible to a given connection or set of connections, as well as what is visible to the general public. Our aim is to allow the participant to examine the data they have shared on social media to a variety of audiences: the public, their colleagues, their friends, and even specific people of interest.

On an overview page, the tool displays summaries of activity on the social networks from which data was collected (Figure 4.3). Each network summary includes a real name and/or screen name, depending on what is available, as well as any basic biographical data presented by each network, and any relevant activity metrics like post count and connection count.

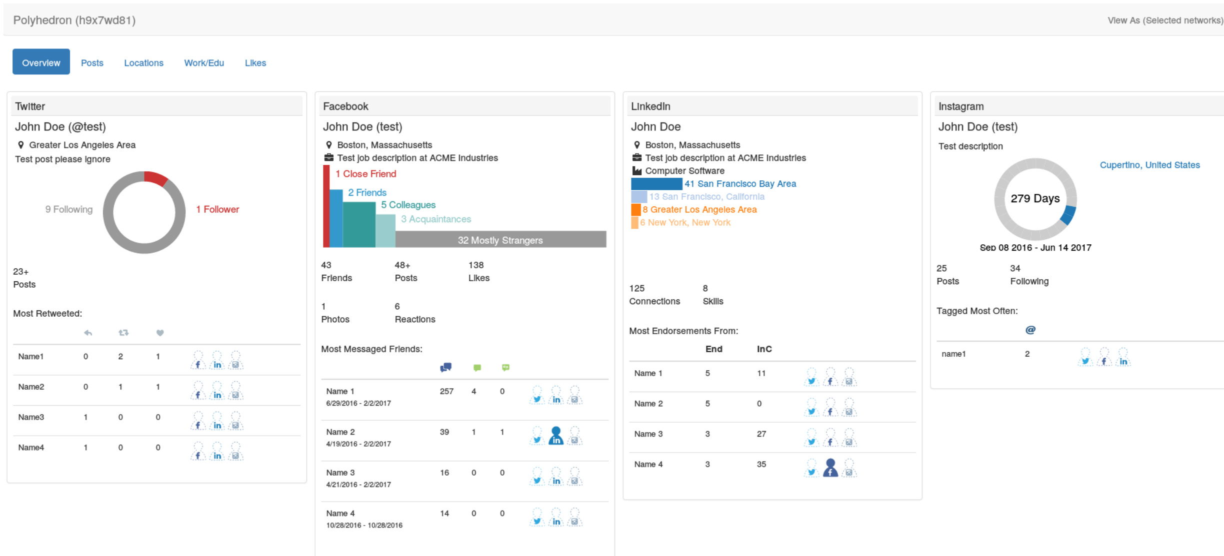

Figure 4.3: Overview page from the visualization.

Each network summary also presents an infographic describing some network-specific quality of interest: the follower-to-following ratio for Twitter, a count of connections by amount of engagement for Facebook, a count of connections by geographic location for LinkedIn, and a visualization of geographic location per post over time for Instagram.

On every page, the tool displays a summary of connection count by network grouping. In addition to the total count of all connections discovered by the tool, this summary shows connection count by individual network, and provides a further breakdown of connection count by strict network set – for instance, one element of the summary shows a count of connections known through both Facebook and LinkedIn but not known through any other network. Additionally, on every page, a list of connections is shown; clicking on any of the summary counters filters this list to show only the connections included in the count. Every connection in the list has a "View As" button permitting the tool's point of view to be adjusted to that of the chosen connection, causing some of the displayed data to be highlighted or hidden depending on whether it is visible to the given viewpoint; a dropdown at the upper right corner of every page shows what viewpoint is currently active, and additionally allows certain special modes ("Myself", "Your friends", and "Public") to be

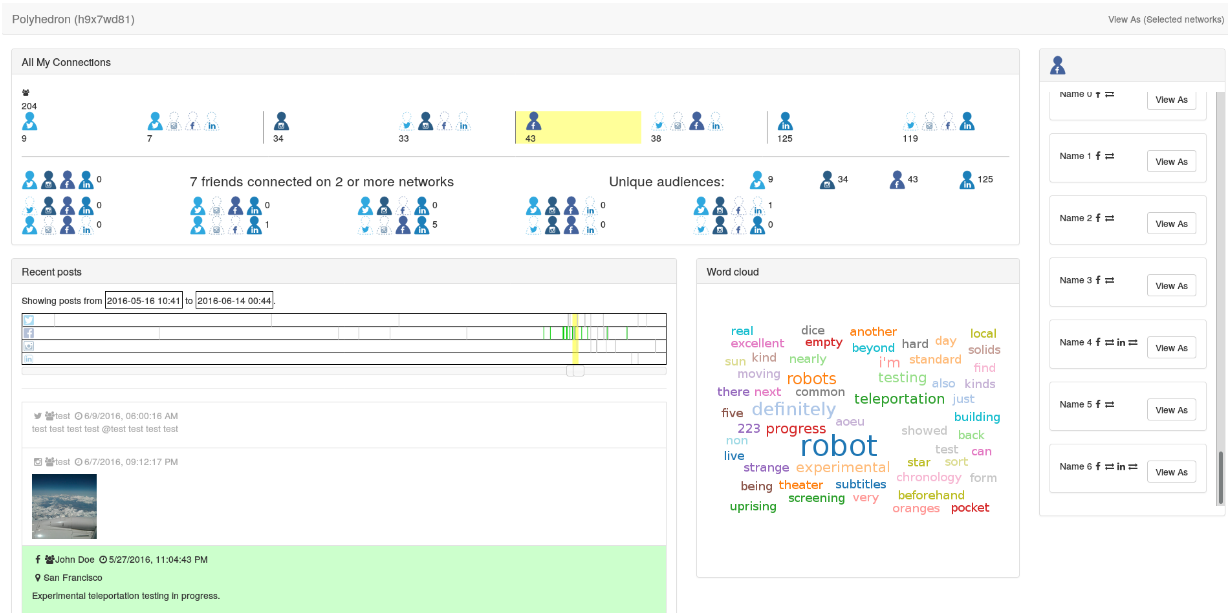

Figure 4.4: Post page from the visualization.

chosen.

The post page (Figure 4.4) contains a merged stream of all posts collected by the tool in reverse chronological order. An activity timeline shows post incidence over all time for each of three networks – Facebook, LinkedIn, and Instagram – and allows the post stream to be filtered by start and end time. The posts themselves are shown beneath, displaying for each post (where available) the post content, image or video thumbnail, external link title and description, post location descriptor, and post privacy descriptor. If a visibility filter is active, posts visible to the chosen viewpoint are highlighted with a green background, and posts hidden from this viewpoint are faded out. A "word cloud" displays words that are used frequently in the participant's post stream; the set of words displayed is taken from all posts visible to the current viewpoint.

For the scenario described in Section 2.3, the post page is the most useful one. By selecting a specific type of audience among the different network combinations (e.g. the LinkedIn set), one highlights the posts visible – both in the stream overview and in the actual list of posts – to that given audience, as well as when in time those posts were made; one also sees the words that are mentioned most frequently among the posts that are visible.

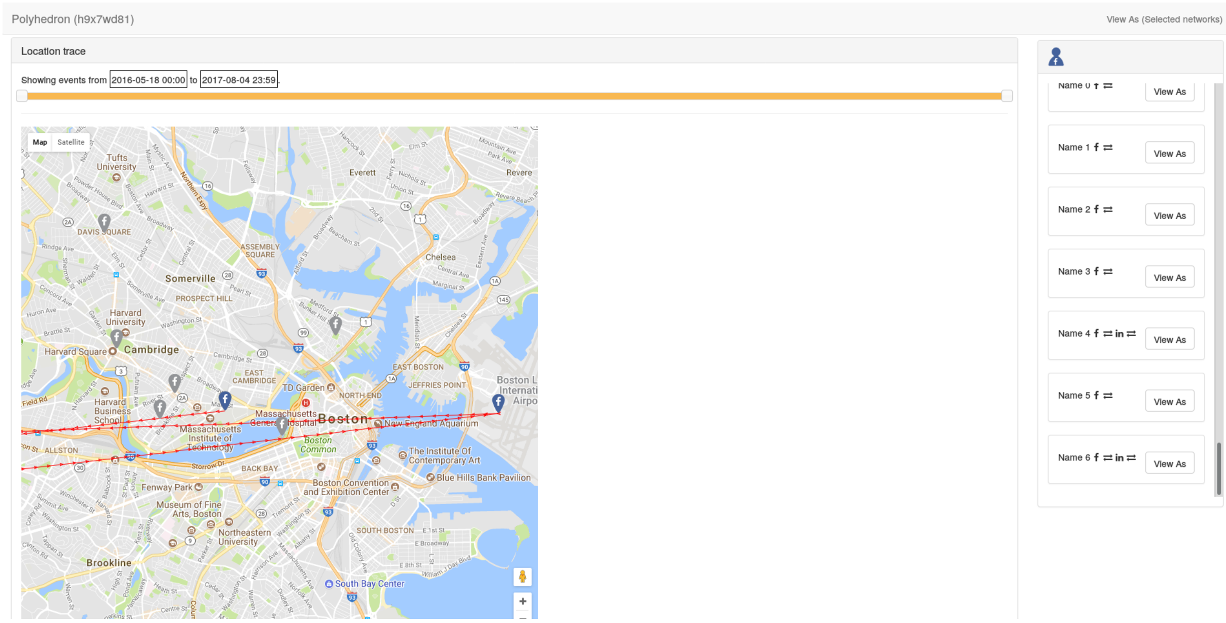

Figure 4.5: Location page from the visualization.

If an otherwise mostly private stream contains a collection of unexpectedly visible posts, one can focus attention on them and determine whether or not they are of concern. Additionally, if in the connections list one sees a person to whom disclosure is problematic, like a potential supervisor, or whom one no longer recognizes, one can re-evaluate whether the connection should still be retained, or if additional privacy settings should be applied in the future.

An additional page presents a map with pushpins labeling the recorded location of every post with a location descriptor, connected in order with directional arrows to present a linear sequence of locations visited. These pushpins respond to the active viewpoint, with locations corresponding to non-visible posts greyed out and disconnected from the sequence of arrows. Because of the large potential quantity of location descriptors, a slider allows the location stream to be filtered by start and end time. One more page presents any work and educational history explicitly provided to Facebook and LinkedIn.

While the tool supports gathering and displaying data at the same time, we found it impossible to maintain acceptable performance in this usage mode for the data volume generated by typical participants. Additionally, to preserve the effect of the user study on participants, we wanted to conceal the normal user interface during data collection, as it would have provided answers to many of the questions asked in the survey. Hence, we developed an alternative user interface that displays only data collection progress, to be displayed to participants as they ran the collection process on their own computers.

The initial design of the tool included a page for displaying certain predictive metrics computed from the collected data, like results from sentiment analysis and personality prediction. However, we were unable to secure the data necessary to reconstruct sensible prediction and analysis models for local use. Certain web services are able to compute such metrics given appropriate input, like Facebook Likes and post text content. However, for the purposes of our study, we did not wish to cause any participant data to be uploaded to any non-CSAIL servers. While the output of such analytics services could grant much more valuable insights than straight inspection of one's own post content, we leave an appropriate treatment of data analysis to future work.

Chapter 5

User study: data estimation and visualization

This chapter describes a user study conducted as part of this work to determine how well people are able to estimate the types of data they make visible on social media and the composition of their audiences across social networks, and how people evaluate their privacy strategies when presented with information about their activity and the data they have shared.

5.1 Methods

Participants in the study were asked to provide information from their social networking accounts – specifically, their profile, post, and activity information from Facebook, Twitter, Instagram, and LinkedIn, or for as many of these networks as they actually use. This information was collected using a modified version of the Polyhedron client, changed to hide the visualization interface and display only data collection progress, so as not to influence the outcome of later portions of the study involving the visualization.

After data collection was concluded, participants took a survey asking them to provide their own estimates of several quantities related to their own social media data. Among other information, participants are asked to estimate the visibility of certain biographical details on each network to the general public, the proportion of posts of particular types (text, photo, video) they make to each network, the proportion of frequency of various activities on each network (posting, commenting, reacting), and the proportion of connections by interaction frequency and by network. Additionally, the survey asks several subjective questions, like the perceived amount of control over visibility of certain attributes, as well as the participant's perceived concern over certain attributes becoming public.

We ran a separate post-processing script, specific to the study, on each bundle of collected data. This script computes estimates of ground truth based on the data collected from each participant's social media information, and compares these estimates to the answers provided in the survey. (Not all attributes are estimable; in particular, the visibility of several attributes on Twitter and Instagram can depend on the content of text posts or image posts, and so are difficult to determine automatically. Additionally, estimating some attributes proved impossible, as the necessary data could not be gathered using our methods.) The script produces a spreadsheet comparing the participant's answers to the computed estimates.

Immediately following data post-processing, each participant was interviewed on their usage of social media, using the spreadsheet as a guide. During each interview, we compared the participants' survey responses with our computed responses; where they differed, or where we determined that the answers were subjective or uncomputable with our methods, we asked the participants to describe how they arrived at their own answers.

Following the review of participant survey answers, we showed participants the full Polyhedron visualization using their saved data, and guided them through its various elements. During this process, we asked each participant some data-specific questions. For each grouping of connections by set of networks, we asked participants to characterize whether there were commonalities among the connections listed, and if so, what those commonalities were. Additionally, we demonstrated which sets of posts on all networks were visible to each set of people, and asked whether cross-network membership of connections was a major consideration when deciding what to post on a particular network. At the conclusion of the study, we asked participants to self-report whether they had been surprised by anything uncovered by the survey review or the visualization, and whether they planned to take any follow-up actions like reviewing connections or posts. Participants were compensated based on the degree of study completion, with those who completed both the survey and interview receiving \$25; an incentive of an additional \$30 was offered to the participant whose estimates were the most accurate overall.

5.2 Results

In this section we present findings related to our aims and research objectives from Chapter 2, along with supporting data figures. Additional data may be found in Appendix C, while additional graphs may be found in Appendix A.

5.2.1 Participant demographics

12 participants completed the entire study, including survey and interview. These participants were recruited by word-of-mouth from students and employees at a major university. We provide a demographic summary in Table 5.1, and total data size counts per participant in Table 5.2.

Participant age ranged from 20 to 46 ($\mu = 29.17$, $\sigma = 7.42$). 8 participants were male; 4 were female. 4 participants had completed up to a bachelor's degree; 4 had completed up to a master's degree; 3 had completed a doctoral degree.

11 participants used Facebook, 7 used Instagram, 7 used Twitter, and 12 used LinkedIn. All 11 participants using Facebook reported using it for social use, with 3 out of 11 reporting using it for professional use. 1 participant reported using Facebook for other use (as a news aggregator). All 12 participants reported using LinkedIn for professional use, with 2 out of

Participant #	Age range	Gender	Main Uses				Other networks
			FB	TW	IG	LI	
P1	30-40	F	S	S		P	
P2	20-30	M	S	P		P	
P3	20-30	F	SP			P	Snapchat
P4	30-40	M	SP	P	S	P	Google+, Snapchat
P5	40-50	M	S	P	S	P	WhatsApp
P6	30-40	M	S	O	S	P	
P7	30-40	M	S	SPO	S	SP	LiveJournal, Reddit, Tumblr
P8	20-30	M			SO	SPO	
P9	20-30	M	S			P	
P10	20-30	F	SP		S	P	
P11	20-30	F	S		S	P	Snapchat, WeChat, KakaoTalk, WhatsApp
P12	20-30	M	S	O		P	Google Hangouts

Table 5.1: Participant demographics, summarized.
S = social, P = professional, O = other, blank = not used

12 reporting using it for social use, and 1 out of 12 reporting use considered neither social nor professional (as a way to reflect on accomplishments). 2 out of 7 participants reported using Twitter for social use, 4 out of 7 reported using Twitter for professional use, and 3 out of 7 reported using Twitter for uses they considered neither social nor professional (as a news aggregator, as a log of read articles, as "a means of sharing experiences with conference-goers"). All 7 participants using Instagram reported using it for social use, with none reporting professional use, and 1 out of 7 reporting use that was considered neither social nor professional (to share personal creations).

All 12 participants used at least 2 of the social networks considered in our study. 4 participants used all four networks. 5 participants used 3 out of the 4 networks; of these, 3 did not use Instagram, and 2 did not use Twitter. 3 participants used 2 out of the 4 networks; of these, 2 used only Facebook and LinkedIn, and 1 used only Instagram and LinkedIn.

Participant #	Posts			Interactions[a]		Connections[b]				
	FB	TW	IG	FB	TW	Any	FB	TW	IG	LI
P1	351	3193[c]		4418	3251	621	365	151		147
P2	857	307		9607	515	1576	979	459		331
P3	128			8820		1091	958			341
P4[d]	820	3235[e]	356	43140	3682	1531	515	660	428	384
P5	827	911	9	7056	939	2418	844	658	85	1082
P6	727	994	64	4002	1003	4422	1886	717	681	1958
P7	259	519	29	1145	754	1027	394	383	96	310
P8			371			1052			178	919
P9	55			3015		538	522			62
P10	629		10	7041		1330	994		208	287
P11	587		105	5196		1035	750		363	44
P12[f]	75	700		156	700	352	156	103		142

[a]Including posts, reactions, and comments; not including private messages.

[b]Includes connections in either direction (follower or following) for Twitter and Instagram.

[c]Twitter does not permit retrieval of more than around 3200 Tweets by any means including official API; based on Tweets after March 2009.

[d]We encounted complications related to data volume while gathering this person's data. Hence, authoritative statistics about posts for P4 are only available after December 2015, and activity is only available after January 2012 (based on an estimate, computed from the Facebook activity log, of 2600 posts in this duration).

[e]Based on Tweets after December 2012.

[f]For this participant, we only have authoritative statistics about Facebook activity after September 2009; collection of activity before this time failed for unclear reasons.

Table 5.2: Activity and connection total counts, summarized.

	FB	TW	IG	LI
n	11	7	6^a	10^b
Family	+6.8 (7.5)	+0.4 (0.7)	+3.0 (3.4)	+3.7 (4.9)
Friends	+31.2 (25.9)	+21.3 (26.2)	+30.5 (14.2)	+17.7 (19.5)
Close Friends	+12.2 (12.3)	+2.4 (3.5)	+17.3 (18.2)	+6.2 (6.1)
Colleagues	-7.9 (6.4)	+5.4 (10.1)	-8.7 (9.1)	-26.8 (17.6)[c]
Acquaintances	-0.6 (27.2)	+22.0 (19.8)	+24.3 (30.2)	
Strangers	-39.2 (26.2)	-49.6 (23.3)	-64.4 (10.9)	

[a]One participant left this series of questions blank.
[b]Two participants left this series of questions blank.
[c]Survey answers in "Acquaintances" and "Strangers" were added to "Colleagues" when determining error.

Table 5.3: Mean error and standard deviation in percentage points for Part 17 (proportion estimates for connection types by interaction). Survey answers were linearly rescaled to add to 100.

5.2.2 Connections by interaction frequency

We asked participants to estimate what percentage of their connections fell within the interaction-based categories we provided. We considered "close friends" to be those with an average of 5 or more daily interactions (150 total) in the past 30 days, and "friends" to be an average of more than 3 interactions in the past 30 days. While computing answers from participant data, we marked connections as family members if they were listed as such on a Facebook profile, and marked connections as colleagues if we were able to determine that they were connected with the participant on LinkedIn.

Participant accuracy of proportion estimates varied. In Table 5.3, we provide the mean estimation error, in percentage points, for participant estimates of each category on each network; per-participant data is graphed in Figure 5.1. The size of the "strangers" category, containing people with whom no social interaction occurred within the prior 12 months, was consistently underestimated by every participant on every network except LinkedIn (on which such people would instead have been classified as 'colleagues'). Conversely, the size of the "close friends" and "friends" categories was consistently overestimated on every network.

Six participants noted specific cases where their own subjective interpretations of each category name did not match with the definitions we gave. Of these, two participants

	FB	TW	IG	LI
n	11	7	6	10
Family	4.6 (5.1)	0.4 (0.5)	2.8 (2.9)	1.0 (2.2)
Friends	28.5 (38.4)	5.9 (12.5)	11.8 (19.9)	6.6 (12.5)
Close Friends	2.5 (3.3)	0.1 (0.4)	0.3 (0.5)	0.6 (1.3)
Colleagues	109.6 (104.2)	40.9 (34.2)	43.2 (46.9)	553.2 (594.7)
Acquaintances	146.4 (97.1)	32.0 (35.8)	28.7 (25.9)	
Strangers	468.6 (367.9)	368.0 (206.2)	238.8 (162.3)	

Table 5.4: Mean computed absolute set sizes and standard deviations for Part 17.

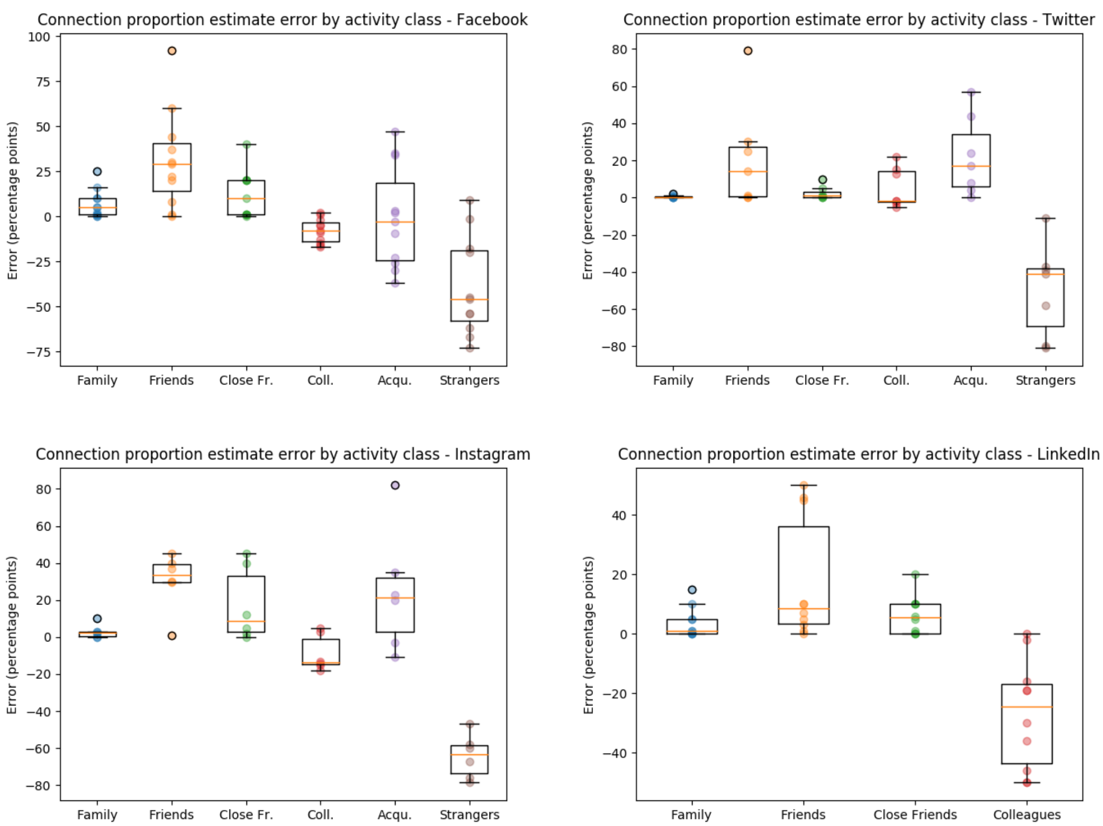

Figure 5.1: Participant estimation errors of connection proportions by activity level class.

(P4, P6) reported that interactions with close friends and family took place primarily by means other than the networks we considered. Additionally, two other participants (P1, P10) indicated that they actively avoided social media contact or exposure to close family members; P1 was specifically concerned about disapproval from family members, and P10 was specifically concerned about generating unwanted worry or attention. Two participants, P6 and P7, also noted that grouping all LinkedIn connections as "colleagues" is not necessarily accurate; P6 pointed out that LinkedIn connections are sometimes made with classmates or others who do not necessarily represent a persistent professional relationship, while P7 noted that he was happy to add people from a social context (e.g. Facebook) to LinkedIn even if a definite professional relationship was not present. P7 additionally reported having a habit of only adding connections to Twitter if he had met them in person at least once, and noted that the personal connection generated by doing so would prevent him from considering them "strangers" even in the absence of recent interaction.

5.2.3 Connections by network grouping

We asked participants, for each of 15 valid combinations of social networks, to estimate how many connections they had who were connections on *exactly* the requested set of networks and no others. For instance, if we asked about connections on Twitter and Facebook, we wanted the number of connections who were listed as friends on the Facebook network and had some follower or following relationship on Twitter, but were not listed as LinkedIn connections and had no relationship on Instagram. For this series of questions, we allowed participants to refer to their own social media profiles at will. In Table 5.5 we provide, for each category, the number of participants n with either non-zero estimated or computed value, the mean computed value μ (with standard deviation σ), mean percentage error (MPE %), and number of participants who overestimated or underestimated the category[1]; if a participant guessed more than 0 for an empty category, the percentage estimate is calculated as if the

[1]When computing these statistics, we removed 1 participant (P6), who had a mean category estimation error of 4351%, as an outlier.

category contained 1 member.

Looking at the results, we observe that most of the cross-network measures seem to be overestimated by participants on average compared to our answers. Additionally, with the exception of the Facebook-LinkedIn category, more participants tend to overestimate rather than underestimate in every multiple-network category.

We caution that some amount of error in the cross-connection measures is to be expected. Because connections can use different names on different networks, we are not always able to match connections across the different networks on which they are present. When we reviewed groupings of connections during the visualization phase, participants often pointed out connections that they knew were connected on more networks than the tool was able to list; in these cases, connections had used nicknames, alternate forms, mononyms, or pseudonyms on one or more networks. We note that this phenomenon was mostly absent between Facebook and LinkedIn, which mandate usage of a real name, but was often present for combinations of networks including Instagram and/or Twitter, which permit free choice of "real name" and indeed do not necessarily ask for a strict one-to-one mapping between an account and a real person.

5.2.4 Content proportions by type

When evaluating timeline content, we considered three general categories of posts: text, photos, and video. We asked participants to estimate which type of post they made most frequently on each network, and to estimate the relative proportion of each type per network.

For Facebook timeline posts, text was the most frequently used medium for the most participants (10 out of 11), followed by photos (1 out of 11). 4 out of 11 participants reported believing that photo posts outnumbered text posts, which was not borne out by the timeline; however, for all four individuals, the plurality of timeline posts made in 2017 were in fact photos.

For Twitter timeline posts, text was the most frequently used medium for all 7 partici-

	n	μ (σ)	MPE %	# Over	# Under
TW/IG/FB/LI	3	9.3 (9.0)	+2.5	1	1
TW/IG/FB	3	9.7 (9.5)	+150.4	2	1
IG/FB/LI	5	16.0 (9.4)	+170.5	3	2
TW/FB/LI	6	8.8 (6.2)	+141.8	6	0
TW/IG/LI	3	5.0 (0.0)	+145.5	2	1
TW/IG	3	13.7 (16.5)	+60.6	2	1
FB/LI	10	76.1 (60.0)	+14.2	4	6
TW/FB	6	6.2 (2.4)	+225.5	6	0
TW/LI	7	20.5 (17.8)	+197.2	4	2
IG/FB	5	57.2 (29.0)	+53.6	4	1
IG/LI	6	12.2 (16.5)	+339.6	4	2
IG	6	138.0 (82.3)	-8.2	2	4
TW	6	349.7 (184.6)	-16.5	0	6
FB	10	520.3 (246.3)	+12.1	8	2
LI	11	267.1 (296.8)	+10.5	6	5

Table 5.5: Mean percentage error for Part 19 (cross-network connection count).

	Facebook	Twitter	Instagram
n	11	7	7
Photos	22.4 (15.5)	7.7 (7.9)	96.4 (6.1)
Videos	1.8 (2.3)	0.3 (0.8)	1.8 (4.6)
Text	74.4 (16.1)	90.9 (8.4)	

Table 5.6: Mean percentages and standard deviations for content proportions (percentage points) by type.

pants with Twitter accounts. All 7 participants correctly identified this to be the case during the survey.

For Instagram timeline posts, photos were the most frequent post type for all 7 participants with Instagram accounts, agreeing with the survey results in each case.

We computed the mean percentages for the various post types, and present them in Table 5.6. Additionally, we computed the mean absolute counts for each type, and present them in Table 5.7. For Facebook and Twitter, we compared the ratios of text posts to total posts (text, photos, and video); we find using ANOVA that participants do appear to treat these differently ($p = 0.02$) at the $p = 0.05$ level.

Participant performance on estimating relative proportion varied significantly; some par-

	Facebook	Twitter	Instagram
n	11	7	7
Photos	110.6 (113.4)	96.4 (127.7)	124.9 (145.6)
Videos	11.7 (17.8)	6.6 (10.3)	10.0 (22.2)
Text	360.8 (244.0)	1305.4 (1163.1)	

Table 5.7: Mean absolute counts and standard deviations for timeline content by type.

ticipants were able to estimate these proportions relatively closely, while other participants' estimates were quite divergent from the computed result. We compute mean error, in percentage points, and standard deviation for each of the networks. Based on our data, we find that people overestimate count of photo posts on Facebook and Twitter, and underestimate count of text posts on the same networks, but the standard deviations on these measures are quite large indicating a wide variation of errors; examining the plot in Figure 5.4, we also find that most of the variance in the Twitter estimates is due to two significant outliers, with the other 5 values much closer to zero overall. We find that, on Instagram, people on average slightly underestimate photo count and slightly overestimate video count, but the effects are relatively small, consistent with Instagram primarily being seen as a photo-sharing platform. These values are presented in Table 5.8 and Figure 5.4.

There are a number of possible sources of participant error in this case. For example, P9 indicated in interview that he uploads many photo albums – the elements of which we do not count as separate photos, as we only tally photos present on the timeline – and that his answers reflect this reported behavior. Additionally, P5 reported that they had based their answers for Twitter on what he remembered from recent activity (see Figure 5.2), and had not necessarily taken all of timeline history into account; we observed similar effects for Facebook estimates for P2 and P11 (see Figures 5.3 and 5.3). Following this hypothesis, we repeated these computations while considering only data from the most recent calendar year containing activity. These mean errors are presented in Table 5.9. We observe for Facebook that, while mean errors move closer to zero, standard deviations increase; participant guesses are not conclusively better estimators of recent participant content proportions for this network.

	Facebook	Twitter	Instagram
n	11	7	7
Photos	+12.8 (16.6)	+9.2 (19.7)	-3.6 (3.6)
Video	+7.0 (12.3)	+1.5 (1.9)	+2.7 (2.9)
Text	-18.5 (24.2)	-9.6 (19.3)	+1.4 (3.5)

Table 5.8: Mean errors and standard deviations, in percentage points, for estimates of content proportions (percentage points) by type.

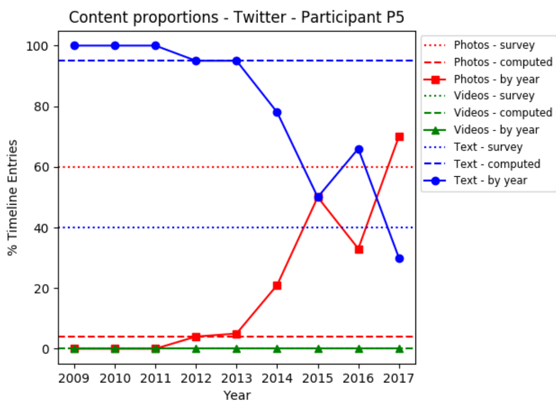

Figure 5.2: P5 content proportion estimates for Twitter; note dotted line (survey result) and time trace (showing breakdown of computed content proportions within each year).

However, mean errors and standard deviations do decrease for Twitter, suggesting that for this network, participant guesses are indeed influenced on average by recent activity consistent with P5's observation.

We also investigated whether network timeline size was related to the observed estimation error. For each network, we took correlation coefficients between the network size and the mean absolute error for each participant. We report these figures in Table 5.10. While the effect displayed is weak, there is some positive correlation for Facebook and Instagram between the network's timeline size and the absolute estimation error, suggesting that even tasks related to estimating content percentages does become slightly more difficult as more posts are made. We, however, found no such correlation for Twitter.

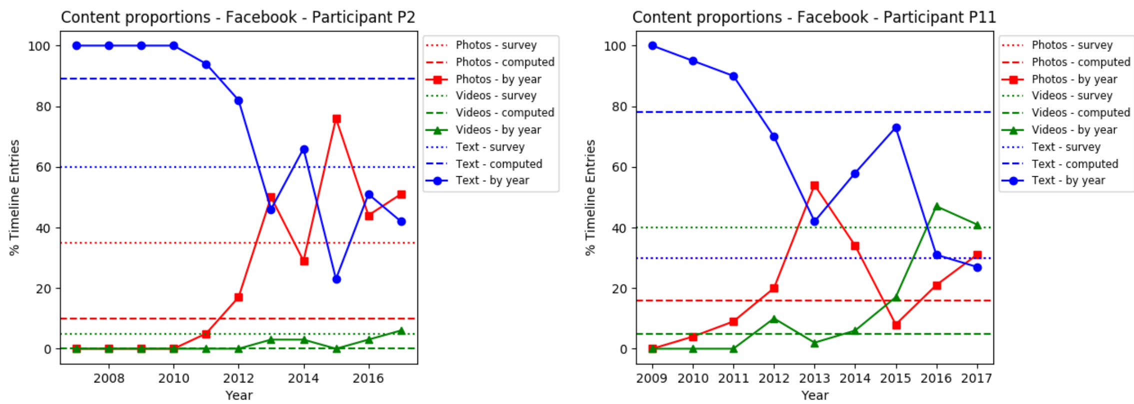

Figure 5.3: P2 and P11 content proportion estimates for Facebook; note dotted line (survey result) and time trace (showing breakdown of computed content proportions within each year).

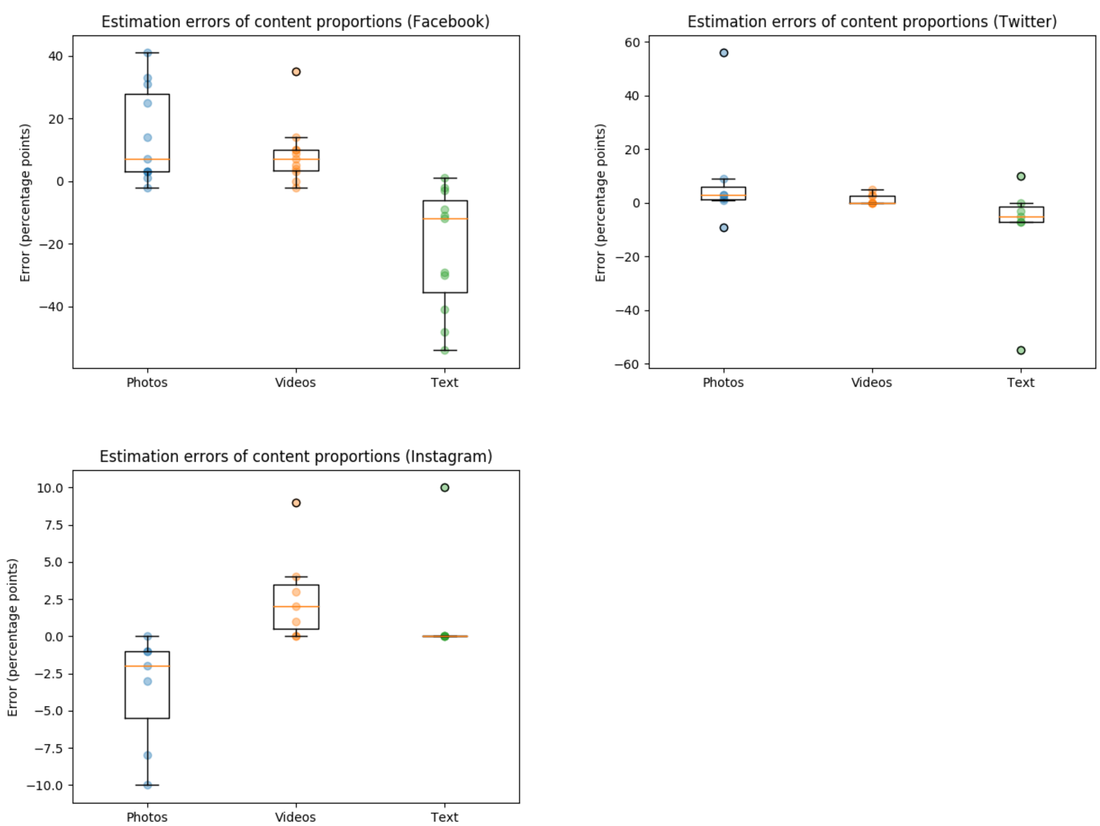

Figure 5.4: Participant estimation errors of content proportions.

	Facebook	Twitter	Instagram
n	11	7	7
Photos	-1.5 (27.1)	-7.9 (7.3)	-2.7 (4.1)
Video	-0.6 (11.5)	+1.4 (2.7)	+1.6 (2.7)
Text	+2.7 (26.0)	+7.1 (7.4)	+1.4 (3.5)

Table 5.9: Mean errors and standard deviations, in percentage points, for estimates of content proportions by type, compared against numbers for the most recent calendar year containing activity.

Facebook	Twitter	Instagram
0.304	0.032	0.244

Table 5.10: Correlation coefficients between network timeline sizes and mean absolute error.

5.2.5 Content breakdown by supplemental data presence

We asked participants, for each of the three types of posts, to consider three different types of additional material that may be associated with a post: a link (for text posts) or caption (for non-text posts), a location descriptor, and one or more tags of people. For each network and type of post, we asked participants to rank by relative frequency the presence of each of 8 possible combinations of such additional data on their own posts.

We also computed the actual relative frequencies of each subtype of post as percentages of total post volume of the overall type; for instance, we computed the percentage of plain text posts among all posts primarily containing text (including text posts with location and/or tagged people, but not photo posts with captions). We report figures in this section in percentage points, except for p-values.

We considered whether there might be differences in the presence of supplemental data between different networks by examining the fraction of posts that contain (or do not contain, respectively) supplemental information beyond their basic content. For text posts, where external links may be treated as part of the post content rather than as contextual data, one-way ANOVA does find a difference ($p = 3 \times 10^{-3}$) between Facebook ($\mu = 85.2$, $\sigma = 9.3$) and Twitter[2] ($\mu = 66.6$, $\sigma = 14.0$) in the fraction of posts without location or tag information.

[2]Observe that Instagram has no support for posts that are text-only.

	Facebook	Twitter	Instagram
n	11	7	7
Plain photos	26.5 (18.9)	0.0 (0.0)	4.9 (10.5)
Plain videos	25.9 (25.0)	3.6 (7.0)	8.6 (18.6)
Plain text (or text with link only)	85.2 (9.3)	66.6 (14.0)	

Table 5.11: Mean percentages and standard deviations of "plain" posts (compared to all posts of the same general type, e.g. all photos), in percentage points.

	Facebook	Twitter	Instagram
n	11	7	7
Plain photos	17.9 (14)	0.0 (0.0)	15.9 (38.9)
Plain videos	3.7 (6.6)	0.6 (1.0)	1.3 (2.4)
Plain text (or text with link only)	308.8 (209.1)	884.6 (883.2)	

Table 5.12: Mean counts and standard deviations of "plain" posts.

For photo posts, where captions may be considered contextual information, one-way ANOVA indicates that participants do behave differently ($p = 1 \times 10^{-3}$) on Facebook ($\mu = 26.5$, $\sigma = 18.9$), Twitter (no plain photo posts reported), and Instagram ($\mu = 4.9$, $\sigma = 10.5$) in terms of making "plain" rather than annotated posts. For video posts, though, one-way ANOVA does not show a statistically significant difference ($p = 0.06$) between Facebook ($\mu = 25.9$, $\sigma = 25.0$), Twitter ($\mu = 3.6$, $\sigma = 7.0$), and Instagram ($\mu = 8.6$, $\sigma = 18.6$) in this regard at the $p = 0.05$ level.

We also considered the prevalence of location information between networks by fractional share. For text posts, location tags are relatively rare on both Facebook ($\mu = 11.1$, $\sigma = 8.2$) and on Twitter ($\mu = 4.6$, $\sigma = 12.2$); in fact, 5 out of 7 participants with Twitter profiles have no location tags on Twitter for any posts. We are, however, unable to statistically separate the two networks on the basis of location tags on text posts ($p = 0.19$). For photo posts, location tags are still rare on Facebook ($\mu = 7.7$, $\sigma = 6.1$) and Twitter ($\mu = 2.5$, $\sigma = 6.2$), but significantly more common on Instagram ($\mu = 31.3$, $\sigma = 22.1$); one-way ANOVA does indicate that these three networks are statistically distinct in this case ($p = 6 \times 10^{-4}$). Location tag percentages on video posts are subject to fairly wide variation owing partially to the limited number of such posts on any network (meaning a single video post with a

	Facebook	Twitter	Instagram
n	11	7	7
Photos with location	7.7 (6.1)	2.5 (6.2)	31.3 (22.1)
Videos with location	14.2 (20.7)	1.0 (2.7)	22.6 (37.5)
Text with location	11.1 (8.2)	4.6 (12.2)	

Table 5.13: Mean percentages and standard deviations of posts with location (compared to all posts of the same general type), including posts with both location and other data.

	Facebook	Twitter	Instagram
n	11	7	7
Photos with location	11.5 (18.1)	9.0 (23.4)[a]	46.7 (74.4)
Videos with location	0.6 (0.8)	0.3 (0.8)[b]	4.0 (9.3)
Text with location	38.8 (38.6)	130.3 (344.7)[c]	

[a]Attributable entirely to 62 such posts by P4 and 1 such post by P7.
[b]Attributable entirely to 2 such posts by P4.
[c]Attributable entirely to 912 such posts by P4.

Table 5.14: Mean counts and standard deviations of posts with location, including posts with both location and other data.

location tag can change the percentages significantly), and we cannot statistically distinguish networks in this case ($p = 0.26$).

We repeated the comparison for fractional share of posts with tagged people by network. In the case of text posts, there is very clear separation ($p = 2 \times 10^{-6}$) between Facebook ($\mu = 4.5$, $\sigma = 3.5$) and Twitter ($\mu = 30.9$, $\sigma = 11.7$). This is not the case ($p = 0.46$) for photo posts between Facebook ($\mu = 6.2$, $\sigma = 8.7$), Twitter ($\mu = 12.2$, $\sigma = 15.5$), and Instagram ($\mu = 14.6$, $\sigma = 19.8$). It is also not demonstrably true ($p = 0.22$) that there is separation of tagging behavior for video posts between Facebook ($\mu = 17.1$, $\sigma = 22.4$), Twitter ($\mu = 23.0$, $\sigma = 38.6$), and Instagram (no video posts with tags).

We exclude LinkedIn because of the limited post history we were able to gather in general from this network, as well as the complete absence of certain features (like location tagging).

	Facebook	Twitter	Instagram
n	11	7	7
Photos with tagged people	6.2 (8.7)	12.2 (15.5)	14.6 (19.8)
Videos with tagged people	17.1 (22.4)	23.0 (38.6)	0.0 (0.0)
Text with tagged people	4.5 (3.5)	30.9 (11.7)	

Table 5.15: Mean percentages and standard deviations of posts with tagged people (compared to all posts of the same general type), including posts with both tagged people and other data.

	Facebook	Twitter	Instagram
n	11	7	7
Photos with tagged people	25.3 (19.6)	30.9 (40.1)	9.0 (13.8)
Videos with tagged people	0.9 (0.9)	1.0 (1.4)	0.0 (0.0)
Text with tagged people	17.1 (14.1)	349.6 (299.3)	

Table 5.16: Mean counts and standard deviations of posts with tagged people, including posts with both tagged people and other data.

5.2.6 Activity proportions

When evaluating participant activity on social media, we considered six types of actions: posting, reacting to one's own posts (e.g. by Like button), commenting on one's own posts, reacting to others' posts, commenting on others' posts, and sharing others' posts. For each social network, we asked participants to estimate which type of activity they performed most frequently, as well as to provide estimates of relative frequency for each type of activity. We then computed, using relevant action data, the actual proportions on the two networks (Facebook and Twitter) for which this was possible.

For Facebook, the most frequent type of activity was invariably reacting to other people's posts, and 9 out of 11 participants correctly estimated this in the survey. The remaining 2 participants estimated that sharing other people's posts was their most frequent Facebook activity.

For Twitter, posting was the most frequent activity type for 5 out of 7 participants, and sharing (in the form of retweeting) was the most frequent activity type for the remaining 2. All participants were able to correctly identify their most frequent activity type.

	Facebook	Twitter
n	11	7
Post	+0.9 (17.4)	-0.9 (9.6)
React own	+1.8 (2.4)	+1.4 (2.3)
Comment own	+2.1 (11.1)	+1.1 (2.5)
React other	-5.1 (24.2)	+2.1 (7.0)
Comment other	-9.0 (13.3)	-4.7 (5.6)
Share other	+12.1 (21.1)	+3.0 (14.0)

Table 5.17: Mean estimation errors and standard deviations, in percentage points, for estimates of activity frequency. Survey answers were renormalized to add to 100.

	Facebook	Twitter
n	11	7
Post	608.7 (723.5)	990.9 (941.3)
React own	7.7 (8.5)	0.0 (0.0)
Comment own	704.4 (1052.8)	9.4 (14.9)
React other	4379.3 (5565.8)	139.3 (164.7)
Comment other	2687.6 (4387.8)	178.9 (238.8)
Share other	121.1 (271.1)	229.3 (280.7)

Table 5.18: Mean absolute counts and standard deviations for activity frequency.

Activity proportion estimate accuracy varied. We present the mean errors and standard deviations in Table 5.17, and plot the errors in Figure 5.5. For the case where only the most recent calendar year of activity is considered, we present the mean errors and standard deviations in Table 5.19. We note that, in general, restricting ourselves to the most recent calendar year did not appear to improve estimates.

Similarly to what we did with content proportions, we investigated whether total number of actions taken is correlated with the absolute estimation error. We present those numbers in Table 5.20. In Twitter's case the errors appear to be completely uncorrelated. In Facebook's case, increasing total action count appears to have slight correlation with decreasing percentage estimation error.

We exclude LinkedIn and Instagram because information required to compute the necessary metrics is unavailable.

	Facebook	Twitter
n	11	7
Post	+2.0 (27.1)	-1.1 (7.9)
React own	+1.8 (2.4)	+1.4 (2.3)
Comment own	+5.8 (11.4)	+1.7 (2.2)
React other	-17.4 (30.1)	+14.3 (16.1)
Comment other	-2.1 (21.8)	+0.3 (6.0)
Share other	+12.3 (21.8)	-15.7 (16.9)

Table 5.19: Mean estimation errors and standard deviations, in percentage points, for estimates of activity frequency, based on most recent calendar year of activity only. Survey answers were renormalized to add to 100.

Facebook	Twitter
-0.323	0.018

Table 5.20: Correlation coefficients between total activity count and mean absolute estimation error.

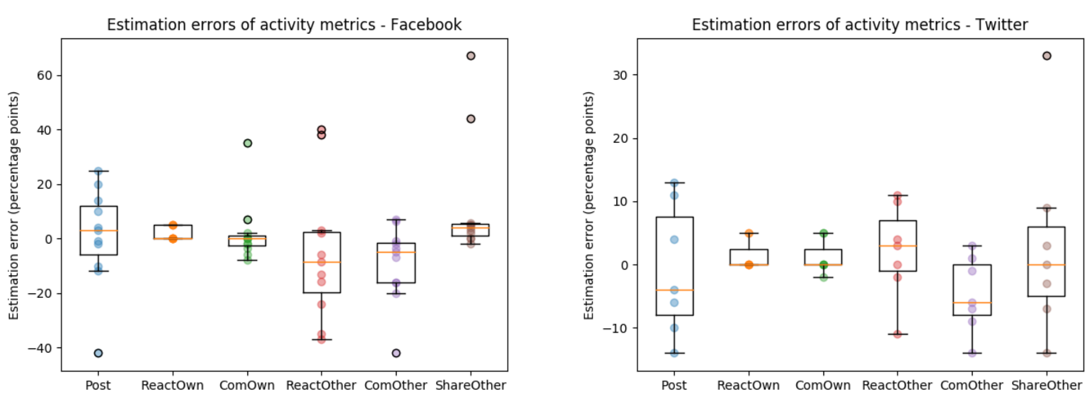

Figure 5.5: Participant estimation errors of activity proportions.

5.2.7 Concern and control over visibility

For this segment we were interested in 15 different types of events, like posting, being tagged, and reacting and commenting to posts by various groups of people (full list in Appendix B, Parts 14 and 15). For each type of event, we first asked participants how concerned they were that the information associated with each event (like a post or a comment) could be seen by strangers, and then asked how much control they believed they had over the visibility of each action. We asked participants to respond on a five-point Likert scale ranging from "no concern / no control" to "extreme concern / full control", with the option to leave blank any question about a feature the participant did not use. We list the mean participant responses for each category in Table 5.21 and Table 5.22.

One-way ANOVA by network suggests that there is a significant difference between networks for both concern ($p = 4.2 \times 10^{-13}$) and control ($p = 6.6 \times 10^{-3}$). Participants were most concerned about visibility of information on Facebook ($\mu = 2.48$) and least concerned about visibility on Twitter ($\mu = 1.22$). Participants believed that Instagram afforded the most control of visibility overall ($\mu = 2.24$) and that Twitter afforded the least ($\mu = 1.39$).

For every network, we found an inverse correlation between the amount of reported concern for each type of event and the amount of reported control. This was most pronounced for Facebook ($\rho = -0.75$) and least pronounced for LinkedIn ($\rho = -0.34$); for Twitter the value was $\rho = -0.52$ and for Instagram the value was $\rho = -0.58$.

We also hypothesized that the amount of expressed concern might have some correlation with individual error rates. We tested this idea by computing correlations between a participant's average Likert concern rating for a particular network and the following error metrics:

1. Mean absolute error across categories, activity proportions, percentage points (E_{ap}).

2. Mean absolute error across categories, content proportions, percentage points (E_{cp}).

Average answers, concern	FB	n	TW	n	IG	n	LI	n
Posts	2.09 (1.38)	11	1.14 (0.38)	7	1.86 (1.07)	7	1.67 (0.82)	6
Posts tagging friends	2.18 (0.87)	11	1.14 (0.38)	7	1.80 (0.45)	5		
Tags from friends	3.36 (1.21)	11	1.43 (0.79)	7	2.29 (1.38)	7		
React to own posts	1.78 (0.97)	9	1.14 (0.38)	7	1.33 (0.52)	6	1.33 (0.58)	3
React to own posts with tagged friends	1.89 (0.93)	9	1.14 (0.38)	7	1.33 (0.52)	6		
React to posts where tagged	2.55 (1.29)	11	1.43 (0.79)	7	1.71 (0.49)	7		
React to friend posts	2.50 (1.18)	10	1.14 (0.38)	7	1.83 (0.75)	6	1.40 (0.55)	5
React to posts where friends tagged	2.36 (1.21)	11	1.14 (0.38)	7	1.83 (0.75)	6		
React to non-friend posts	2.91 (1.58)	11	1.43 (0.79)	7	2.33 (1.51)	6	1.60 (0.89)	5
Comment on own posts	1.91 (0.94)	11	1.14 (0.38)	7	1.17 (0.41)	6	1.00 (0.00)	3
Comment on own posts with friends tagged	2.36 (0.81)	11	1.14 (0.38)	7	1.67 (0.52)	6		
Comment on posts where tagged	2.73 (1.19)	11	1.43 (0.79)	7	1.86 (0.90)	7		
Comment on posts where friends tagged	2.78 (1.09)	9	1.14 (0.38)	7	1.67 (0.82)	6		
Comment on friend posts	2.64 (1.12)	11	1.14 (0.38)	7	1.85 (0.90)	7	1.33 (0.58)	3
Comment on non-friend posts	3.30 (1.7)	10	1.14 (0.38)	7	2.14 (1.46)	7	1.50 (1.00	4
Average of all non-zero responses	2.50 (1.23)	157	1.22 (0.49)	105	1.79 (0.91)	95	1.45 (0.71)	29

Table 5.21: Mean answers (and standard deviations) for Part 14 (concern over visibility).

Average answers, control	FB	n	TW	n	IG	n	LI	n
Posts	3.45 (1.37)	11	2.29 (1.60)	7	3.14 (1.07)	7	2.00 (1.00)	5
Posts tagging friends	2.27 (1.10)	11	1.57 (0.79)	7	2.00 (0.89)	6		
Tags from friends	1.73 (0.79)	11	1.00 (0.00)	7	2.00 (0.89)	6		
React to own posts	3.45 (1.37)	11	1.86 (1.46)	7	3.17 (1.17)	6	2.00 (1.00)	3
React to own posts with tagged friends	2.18 (1.25)	11	1.57 (0.79)	7	2.33 (0.82)	6		
React to posts where tagged	1.64 (0.92)	11	1.14 (0.38)	7	2.17 (0.75)	6		
React to friend posts	1.36 (0.50)	11	1.29 (0.49)	7	2.17 (0.75)	6	1.50 (1.00)	4
React to posts where friends tagged	1.18 (0.40)	11	1.14 (0.38)	7	2.17 (0.75)	6		
React to non-friend posts	1.09 (0.30)	11	1.14 (0.38)	7	1.67 (1.03)	6	1.25 (0.50)	4
Comment on own posts	3.36 (1.36)	11	1.86 (1.46)	7	2.67 (1.03)	6	2.00 (1.00)	3
Comment on own posts with friends tagged	2.18 (1.25)	11	1.57 (0.79)	7	2.17 (0.98)	6		
Comment on posts where tagged	1.64 (0.81)	11	1.00 (0.00)	7	1.83 (0.98)	6		
Comment on posts where friends tagged	1.45 (0.52)	11	1.14 (0.38)	7	1.60 (0.89)	5		
Comment on friend posts	1.36 (0.67)	11	1.14 (0.38)	7	1.60 (0.89)	5	1.33 (0.58)	3
Comment on non-friend posts	1.00 (0.00)	11	1.14 (0.38)	7	1.40 (0.89)	5	1.33 (0.58)	3
Average of all non-zero responses	1.96 (1.22)	165	1.39 (0.84)	105	2.17 (1.00)	88	1.64 (0.78)	25

Table 5.22: Mean answers (and standard deviations) for Part 15 (perceived control).

Mean concern	(FB)	(TW)	(IG)
E_{ap}	-0.43	-0.27	
E_{cp}	0.48	-0.26	-0.09
E_{cif}	0.19	0.16	0.17

Table 5.23: Correlation coefficients between mean absolute errors and mean Likert-scale concerns.

3. Mean absolute error across categories, connections by interaction frequency, percentage points (E_{cif}).

The resulting correlation figures are reported in Table 5.23. They seem to indicate that people with higher self-reported concern do better on estimating activity proportions and non-Facebook content proportions, and worse on every other metric. However, except for the correlations between Facebook activity proportion error and content proportion error, the correlations are relatively weak.

5.2.8 Subjective variations in social media usage strategies

During the interview and as part of the survey, we asked participants to comment in more detail on how they use each network. We found that Facebook is used universally for social purposes – to catch up with friends and family members, especially people not often met in real life. Of the seven participants with Twitter accounts, we found that 6 out of 7 used it to share thoughts or interact in either a public or professional context, but 1 participant (P6) used it in an atypical private context as a repository for saving news articles. One participant using Instagram (P4) reported carefully curating content before sharing it, while the other participants did not remark on the level of curation. Almost all participants (11) using LinkedIn reported using it in a mostly passive capacity, with only one participant (P8) using it for any sort of active sharing; LinkedIn was universally treated as a space for professional interaction, with P5 specifically reporting that he is rarely even connected via LinkedIn to individuals outside the work context.

During the interview phase, we asked participants to describe how they curate their own appearances on social media. We found strategies to be split. 6 participants reported that they regulate, with varying degrees of precision, what particular connections or groups can see. P1 reported, for example, that she specifically excluded family members from viewing certain Facebook posts, like particular pictures or references to specific events, using the network's post privacy settings, out of concern that those family members would disapprove of those pictures or events. P10 reported that she also excluded family members specifically from many posts, as she was concerned that parents would express worry about certain posts with strong political sentiments. P4 reported trying to adopt a much more stratified division of information, making extensive use of the Facebook friends list feature and controlling what information each list receives; he expressed a desire to not "spam" his entire network with updates intended for specific people. On the other hand, the remaining 6 participants treated their entire social network persona as if it were public and posted accordingly, even if they used privacy settings to try to ensure otherwise. P6, in particular, expressed pessimism that the privacy settings presented by Facebook were effective – and, indeed, indicated on the survey that he felt he had no control over the visibility of any Facebook activity, regardless of the tools that the platform provides for the purpose.

5.2.9 Effects of the Polyhedron visualization suite

After reviewing survey answers with participants, we showed them the Polyhedron visualization of their collected data and stepped through each page, allowing them to interact and explore at each stage, and drawing attention to certain features of interest.

We started by examining the general overview for each network. For the most part, participants were unsurprised by the information presented in this section. Some of the elements displayed, like the distribution of connections by interaction frequency, had already been reviewed during the interview and were not necessarily "new" information. One participant showed mild surprise at a particular location descriptor displayed in the Instagram overview, but decided that this unexpected descriptor was not of serious relevance.

We proceeded to review the different groupings of connections by network presence, and asked participants whether or not the groupings represented any patterns or commonalities. Some smaller groupings contained fairly specific commonalities; for example, P10 was able to identify a set of co-workers at a specific previous place of employment, and a different set of people who share a particular common activity. Larger groupings tended to be ascribed with more generic commonalities if they received commonalities at all; for instance, P6 was unable to identify the FB-only set more specifically than noticing that many of them went to the same high school.

After reviewing connection groupings, we proceeded to examine participant post streams, highlighting posts by visibility to various sets of people (e.g. Facebook friends, LinkedIn connections, the general public). Where potential privacy concerns potentially existed – for example, Facebook posts visible to non-connections – we asked people to comment on the specific posts. Usually, participants remarked that such posts represented news or opinions that were indeed meant to be shared publicly. However, for one participant (P11), we discovered that posts before a certain date had been inadvertently left public; the participant immediately expressed concern because the highlighted set of posts conveyed a very different image than the participant wished to express on social media. A different participant (P3) showed mild surprise and concern at some of the contents of the "word cloud", which highlighted timeline words by frequency; the concerning words were judged after review to be harmless in context.

After reviewing the post stream, we examined the location stream, plotting geographical coordinates for those posts for which precise coordinates were provided. In most cases, participants were able to explain the origin of different clusters of location pins; often, they could identify specific places they had been, like restaurants, parks, shops, and places of worship, and relate them to different phases of their lives. In no case, however, did any participant express particular concern over the visibility of this data; even in a case where location pins were situated close to places of residence, the participant argued that the locations provided were sufficiently inexact as to provide little useful information.

Once participants had finished interacting with the visualization, we asked them whether they were surprised by anything they had seen and whether they would plan to make any changes based on what they had been shown. Out of 12 participants, 4 indicated that they would make some changes, ranging from removing old connections to changing privacy settings on old posts that had mistakenly been left public. The other 8 participants indicated that they were happy with the information presented and the techniques they were already using; indeed, 5 participants seemed to be reassured that their social media strategies appeared to be working well.

5.3 Discussion

5.3.1 Perceptions of audience composition

We observed that participants consistently overestimated the proportion of medium-to-high-interaction connections – “friends” and “close friends” – regardless of network. Indeed, when considered together, these categories are overestimated on average by between 24 percentage points (Twitter) and 48 percentage points (Instagram). Conversely, the proportion of zero-interaction connections – “strangers” – was consistently underestimated on average, by between 39 percentage points (Facebook) and 64 percentage points (Instagram). These numbers suggest that participants in our data set are biased towards considering themselves closely engaged with a larger portion of their social media audiences than they are in practice.

We caution that there are some possible explanations that do not necessarily suggest a simple deficiency in estimation. Six participants noted that the labels we used for each category carry pre-existing meanings that are not necessarily aligned with a metric based on interaction using the networks sampled. For example, P7 noted in interview that they primarily add to Twitter only those people they have met in person, and that they would not subjectively categorize such people as strangers despite lack of social media interaction since the participant has physically met each such connection at least once. Also, four participants (P4, P5, P11, P12) reported primarily using networks or methods other than those sampled, like Snapchat and WhatsApp, to keep in regular contact with family and close friends. We expect that, for these reasons, people might provide higher estimates of high-closeness categories like "friends" and "close friends" than are supported by the data we have. Nevertheless, the data we have seems to indicate a clear disparity, on average, between participant estimates and actual interaction counts, and so strongly suggests that H1, the hypothesis that people have trouble accurately estimating connection distribution by engagement frequency, applies to our population.

With regards to sorting connections by network grouping, it is more difficult to make firm conclusions. It is true that, looking only at the mean percentage errors, participants show large relative errors in estimating the number of cross-network connections of each possible type. However, we note that the absolute sizes of most categories are relatively small on average (with five categories averaging fewer than 10 connections). We also note that the categories with large relative differences between participant estimates and computed values are also those categories where we experienced issues with connection deduplication, and that the direction of these differences would be consistent with an underestimate on our side. Without certainty as to the size of the deduplication effect, we cannot conclusively say that H2, the hypothesis that people have trouble accurately estimating connection distribution by cross-network presence, applies to our population, even if some people clearly appear to have trouble with the task.

5.3.2 Information sharing patterns across networks

In terms of posting intent, we see some clear patterns. Facebook and Instagram are consistently viewed as more "personal" networks than Twitter or LinkedIn, in the sense that people report using Facebook and Instagram more to connect with family and friends, and more frequently share posts and photos in a more personal context. LinkedIn is used primarily as a tracking or self-advertising tool related to the professional network, while most participants use Twitter to further a public or professional image. Half of the participants we surveyed reported trying to exert a higher degree of control over which audience sees particular types of content; this is easier to do on Facebook (which allows post visibility to be fine-tuned) than on Twitter or Instagram.

In terms of content proportions, we find that both Facebook and Twitter are used primarily for posting text, while Instagram is mostly used for posting photos. This is not so surprising given the primary purpose of each platform. We also find that, although both Facebook and Twitter are primarily used for text posts, the proportion of text posts on Twitter is statistically significantly higher than the proportion of text posts on Facebook. We thus find that H3, the hypothesis that there is a difference between networks in proportion of text to non-text posts, applies to our population. We note that all 11 participants with Facebook reported using the network in a social context at least some of the time, while only 2 of 7 participants (P1, P7) reported so for Twitter; posting photographs and other images may very well be incompatible with a professional use case for Twitter, which could account for the difference.

We observed that, in overall terms, posts on Facebook, Twitter, and Instagram tend to show different behavior when considering the percentage of "plain" posts (as opposed to posts tagging other people, containing captions, or providing location information): in particular, people are more likely to post such "plain" posts on Facebook than on Twitter or Instagram. There are a few reasons why this might be the case in practice. For instance, it is an accepted and frequent mode of interaction on Twitter to include mentions of Twitter accounts to which specific posts are addressed, or even accounts representing people, organizations, or ideas closely related to the post topic; indeed, we discover that, on Twitter, accounts are certainly "tagged" in this way more frequently for all post types than on Facebook. The distinction is statistically significant for text posts, thus validating H4, the hypothesis that there is a per-network difference in tagging frequency, for this type; however, for video and photo posts, we did not find the distinction to be statistically significant, thus failing to validate H4 for those types of posts.

Location information, on the other hand, was not a popular type of annotation on any network – except for Instagram, whose mobile app allows one-click geotagging of uploaded photos. Indeed, for photo posts, geotagging on Instagram is statistically significantly more frequent than on Facebook or Twitter, hence validating H5, the hypothesis that there is a per-network difference in location frequency, for this type. Video and text posts do not show similar levels of separation with regard to inclusion of location data, so H5 is not validated for those types.

Content proportion estimates seemed to vary most widely among those categories of data that were especially popular on each given network, so long as there was more than one such category – for instance, on Facebook, participant estimates varied greatly for text and photo posts. Participant estimates also exhibit strong bias in some cases: for instance, people consistently overestimate the number of photos they share on Facebook, and underestimate the amount of text. Two participants report a plausible explanation that may account for this particular bias – the mode with which they share photos, via albums rather than timeline posts. However, the estimates themselves seem to be spread widely enough for Facebook that we can accept H6, the hypothesis that people have trouble estimating content proportions, for this network. However, we do not accept H6 for Twitter, where the errors are generally smaller and show less compelling trends once outliers with very large estimation errors are removed; since our test population uses Twitter mainly for text posts and only occasionally to post photos, the smaller amount of variance does make some sense.

Activity proportion estimates seemed to vary most among popular types of activity, similarly to content proportion estimates. In particular, on Facebook, estimates of post and reaction frequency tended to err in either direction by up to 40 percentage points in the worst case, while estimates of less common modes of interaction like reacting to one's own posts were much less inaccurate. It appears that participants are fairly aware on average of the activities they do not perform on any given network, but that errors concerning the activities they do perform are potentially larger; interestingly, though, there is a slight trend for Facebook toward smaller overall errors with larger overall activity counts. Given that people's estimates of Facebook activity tended to be quite varied, we can accept H7, the hypothesis that people have trouble estimating activity metrics, for this network; however, for Twitter, the errors are generally smaller in size compared to overall activity, suggesting people are better estimators in this case, so we do not conclusively accept H7 in that case.

5.3.3 Effects of data visualization on perceptions of privacy risks

For most (8) of our participants, reviewing their social media activity seemed to be an exercise that yielded few surprises. When we reviewed their lists of multi-network connections, they expressed familiarity, or at least lack of overt concern, at the names highlighted; when we reviewed their post streams, we found that they were visibility-restricted to audiences that were in line with the general sharing intentions they expressed. For these participants, the tool-assisted audit we provided ended up being a simple validation of their existing policies and strategies, and did not seem to point to a meaningful behavioral shift – as the data presented suggested, to them, that none was necessary.

A handful (4) of participants, though, did identify some points of concern that they indicated wanting to work to change. Largely, this was surprise at the number of social media connections with whom they had not interacted in a while; as a result, these participants specifically expressed a desire to review their connections and prune the people they no longer recognized. In one specific case (P11), the concern was perhaps more overtly alarming: a group of posts from the participant's early social media history that were publicly viewable but, in the participant's view, should not have been. Hence, our tool was able to prompt reconsideration of post content, post visibility, and effective audience, in some cases – though this is necessarily subject to prior expectations.

For 8 of our participants – a majority of our sample – we find no evidence to support H8, the hypothesis that our visualization induces a measurable shift in privacy risk perception. However, for 4 participants, we find that our tool was able to uncover unexpected discrepancies between sharing expectations and practice, and in such scenarios, we find that H8 seems to apply.

Overall, it seems that while our sample population showed some difficulties estimating metrics about their own data, most considered their actual behavior to be reasonable when given a chance to review it in some detail, though doing so did cause a minority of participants to express desires to change certain aspects of their interactions with social media.

5.3.4 Summary

From a subjective viewpoint, we discover that people in our test population do ascribe specific purposes and audiences to different networks. Interviews suggest that Facebook usage primarily revolves around following and networking with family and friends, and Instagram is also used in more personal or insular contexts, while Twitter is used in a much more public context, as a medium for broadcast, information consumption, and interaction furthering a professional or public persona. Certain quantitative measures – the proportion of text vs. non-text between Facebook and Twitter, the greater prevalence of tagging other people on Twitter compared to other networks, and the much greater relative uptake of location on Instagram – suggest that, in addition to differing purposes, the separate networks promote different modes of communication.

We find additionally that people encounter varying degrees of difficulty recalling the types of content and activity on social media, particularly on networks like Facebook (and, to a lesser extent, Twitter), which allow various modes of interaction. We even find hints that the content estimation problem gets worse as the amount of content increases, and see specific instances where estimates are heavily influenced by recent activity. We also find that people have trouble estimating the proportions of interaction types they perform on social media, especially among those modes that they perform most often.

How do people handle choosing who to share information with, given that posts and activity tend to fade from awareness? Two main strategies expose themselves. One mode of thought, held by half of our test population, involves carefully scrutinizing the type of audience to which the post is to be made, and choosing a network or privacy level accordingly. The other strategy, adopted by the other half of our test population, is to behave as if all information is globally visible regardless of network or settings, and avoid posting information that does not belong in the public context.

When past sharing behavior and connection overlap is brought to light through visualization, people in our test population are occasionally surprised by specific elements, but largely show contentment or even reassurance after examining the results of their sharing strategies. However, in a handful (4) of cases, participants are motivated to change their behavior, either by reviewing posts or by re-evaluating connections. So, while our visualization does not appear to manufacture the appearance of problems where none can be perceived by the viewer to exist, we are nevertheless able to call attention to certain potential points of worry.

Chapter 6

Problems and challenges

This chapter describes some of the difficulties we encountered during our work, both in developing Polyhedron and in conducting the user study. The chapter presents a summary of lessons learned, as well as cautions for those looking to replicate or extend the work.

6.1 Evolution of target platforms

Modern social media websites, like the ones we targeted, regularly develop and deploy updates to their primary products, including both new features and updates to existing features. Because our methods rely both on a stable output format and a stable API, updates can cause parts of our tool to stop working until we can determine what has changed and adapt our methods accordingly. Every platform we targeted in this work has changed either their output format or API structure at least once during the course of the work, from early 2016 to mid-2017. We describe the changes observed during the period in which the user study was conducted, based on when we were able to identify and fix the resulting issues.

Our strategy for adapting to such changes could have used some improvement. Too often, we would discover the presence of changes "reactively", based on problems encountered with the data capture process when attempting the user study with actual participants. Even when we were proactive in determining the existence and effects of breaking changes, and adjusting our software to compensate, distribution of said changes was somewhat laborious due to our lack of an automatic update mechanism (omitted largely to avoid concerns that our software would upload malicious code to participant computers). Researchers performing similar work would have a much easier time if the following points were to be followed:

- Be proactive in testing any "data-scraping" tool on an adequate test set.
- Take advantage of any semantic information provided along with the returned data; it may not be necessary to match specific DOM structures if more descriptive sets of labels are provided.
- Obtain participant consent for automatic update deployment, where possible. The proportion of people who both actively demand hand-delivered updates and participate in such a study in the first place is likely quite small (none of our participants mentioned this as a concern).
- If the data scraping tool can be written to update itself and re-run collection of information without participants needing to restart the tool (or take any other action), a lot of effort can be saved.

It is true that these points are viewed by some as fairly elementary software engineering considerations; however, they are still worth repeating in this case. When collecting information from rendered HTML or internal API calls, it is necessary to treat the upstream provider as unstable, and architect accordingly.

6.1.1 Facebook

From December 2016 to July 2017, we adjusted our methods to account for the following unanticipated changes on the Facebook side:

- Sometime before 12 December 2016 the Facebook API for comments changed a specific variable '`feedbacktargets`' from an array of objects to a singular object ‘feedbacktarget‘.

- Sometime before 15 December 2016 Facebook changed the Work/Education profile tab’s internal name from ‘edu_work‘ to ‘education‘.

- Sometime before 5 June 2017 the CSS class of the full name link on a Facebook profile changed from ‘_8_2‘ to ‘_2nlw‘.

- Sometime before 20 June 2017 a group of Facebook profile timelines, which previously functioned by querying ‘/ajax/pagelet/generic.php/ProfileTimelineSectionPagelet‘ to retrieve older timeline entries, began to require queries to ‘/timeline/jumper/async‘ instead. These queries are of a completely different style and are initialized differently by the Facebook frontend.

6.1.2 Twitter

From December 2016 to July 2017, we adjusted our methods to account for the following unanticipated changes to the Twitter frontend:

- Sometime before 26 April 2017 the CSS class of the author name for a rendered post changed from ‘js-action-profile-name‘ to ‘fullname‘. Additionally, the correct CSS selector of the screen name for the logged-in profile on the home page changed from ‘.DashboardProfileCard-screennameLink‘ to ‘.DashboardProfileCard-name > a‘, and the screen name on the corresponding profile page was wrapped in an extra <b> tag.

- We determined on 5 June 2017 the correct CSS selector for a quoted tweet’s author was ‘div.QuoteTweet-originalAuthor span.username‘. In the past it had been ‘span.QuoteTweet-screenname‘, but this had not been true for some time.

- We determined around 22 June 2017 that at least some, if not all, posts indicated their existence as a retweet by showing an icon with CSS selector ‘span.Icon--retweeted‘.

In the past this had been ‘`div.retweeted`‘.

6.1.3 LinkedIn

Uniquely among all four platforms, LinkedIn experienced a complete revision of product in early January 2017. Prior to this time, LinkedIn provided most of its information as rendered HTML, requiring DOM parsing to extract useful information. After the change, which overhauled the Web product's entire frontend UI, LinkedIn served data using a JSON API previously reserved for its mobile product. While this particular change already required a total rework of the collection code we had finalized in late 2016, we note the following additional difficulties:

- Before the January 2017 product overhaul, a connection's geographical area could be fetched directly from the main connection listing. After the overhaul, this information was removed from the connection list, requiring queries to every single connection.
- Not all APIs use (or permit) usage with the `/voyager/api/mux` endpoint; some (for instance `/voyager/api/me`) must apparently be called directly if needed.
- In some cases, the results of API calls are provided as metadata entries when a LinkedIn page is first loaded. This was formerly true for `/voyager/api/me`, but stopped being true sometime before 10 May 2017.

6.1.4 Instagram

Perhaps uniquely among all four platforms, Instagram provides a Web product with limited functionality compared to the corresponding mobile product. As such, some data points we would have wanted to access, e.g. the list of posts one has "liked", are simply unavailable for analysis. Even this limited product, however, was still subject to changes, which required us to adapt our methods accordingly.

Between December 2016 and July 2017, we observed the following:

- Sometime before 26 April 2017 we found that the data schema returned by Instagram for posts had changed. Instagram accounts tagged in pictures were formerly represented as a list of objects under the path "`.usertags.nodes`", with each object containing an account profile descriptor at the path "`.user`". Afterwards, they were represented as lists of objects under the path "`.edge_media_to_tagged_user.edges`", with each object containing an account profile descriptor at the path "`.node.user`".

- Sometime before 5 June 2017 we found that the main API endpoint used by main Instagram features had changed from '`/query`' to '`/graphql/query`'. The query format had also changed significantly: instead of a string describing a query in an unknown query language, the new endpoint instead accepts a JSON object containing a numerical identifier for the query type, among other parameters.

6.2 Target platforms behaving differently for different people

We observed, when attempting to run our code on participant computers, that different people were often served different versions of the same website. We encountered this behavior primarily on the Facebook platform, where we found that individual features, though providing equivalent functionality, would behave very differently on the API level between participants. Because we did not anticipate the possibility when starting the work, each instance discovered required adaptation work on par with setting up collection for an entire newly discovered feature. In the following section we describe an instance where we were able to verify that this was happening.

There is, of course, a sensible reason why a platform like Facebook, given its resources, may wish to take such an approach. The Facebook product undergoes constant evolution, with new features being added and existing features being reworked to satisfy a need for improved performance and/or aesthetics. A well-accepted practice in determining the effectiveness of any particular approach to a feature is to present two or more versions to different groups of people, and then observe metrics of interest like load time and click patterns. This practice is known as "A-B testing". Facebook, apparently uniquely (at least based on our relatively small sample size), uses this approach extremely aggressively, segmenting its production website in this manner for multiple different features simultaneously.

There is at least one additional instance where we suspect this phenomenon to be at work, but were unable to verify due to time constraints: we observed while attempting to complete data collection on one potential participant that discovering private messages for Facebook friends was apparently impossible. This difficulty, coupled with lack of issues with other participants for the same data class, suggests that this participant was receiving a version of Facebook that returned private messages differently. Unfortunately, the time at which this was discovered precluded further investigation, so we do not expand on this phenomenon here.

Facebook's extensive use of A-B testing, and possible future trends toward highly personalized Web content delivery, exacerbate the concerns raised in Section 6.1: it is challenging, if not potentially impossible, to anticipate every possible change before it is encountered by some participant. For researchers who must use DOM scraping or API emulation rather than official platform APIs or privileged access to data (which tend to be more stable), we offer the following cautions, in addition to those raised in the prior section:

- Be proactive in examining returned participant data to review if it contains all the expected information. It is even better if the collection tool itself can do this, and flag potential anomalies while minimizing the amount of manual review researchers have to perform.

- Because not all "changes" (as expressed by the previous section) are necessarily uniform, data collection must determine whether or not a particular change applies to a given participant profile, and be able to handle every case without further input.

- Techniques for collecting data without relying on having to match any specific DOM structure would be immensely useful for this type of work. (In the limit, consider the case where every profile is rendered with a completely unique, possibly procedurally generated, possibly partially user-specified, page layout.)

6.2.1 Facebook profile timeline

Between December 2016 and July 2017, the period of time when we were working with study participants, we were able to identify three variants of the Facebook profile timeline being served to different groups. The Facebook timeline provides "infinite scrolling" – that is, a fragment of timeline is retrieved every time the page is scrolled to the bottom – and each variant accomplishes this task differently. We list the variants we encountered below, with the caution that other variants may exist that are not listed here. While none of them presented especially challenging problems from a collection perspective, their existence and method of discovery (observation of website behavior from debug logs provided along with data captures, which were manually retrieved from participant hardware), coupled with the long iteration cycle necessary to adapt to such changes (as it is often not possible to replicate the variant behavior on our own due to the nature of A-B testing), caused large delays in the data collection process in practice.

- The first variant we observed, and the one we initially developed against, initializes the timeline by first registering a number of time periods, each corresponding to a calendar year, and then placing a "scrolling pager" on the page. Each registered time period, as well as the initial scrolling pager, is associated with a bundle of metadata describing how to fetch the next segment of timeline. Whenever the bottom of the page is brought into view, an API call is made to '`/ajax/pagelet/generic.php/ProfileTimelineSectionPagelet`' with the metadata bundle corresponding to the missing segment. If the returned timeline fragment is not the earliest one in the time period, it contains metadata for a new scrolling pager corresponding to the next earliest fragment; otherwise, the stored metadata for the next earliest period is used.

- The second variant we observed, to which many people's Facebook timelines transitioned, does not segment the timeline by calendar year, but instead treats it as a single stream. This variant is initialized with a single metadata block describing initial timeline entries to render as well as a cursor to pass to the API for the next fragment. To fetch the following timeline segment, the cursor, among other parameters, is passed to '`/timeline/jumper/async/`' as a URL parameter; the returned JSON contains, among other information, a cursor for the following timeline segment, if it exists. While the cursor is passed as a URL parameter, the actual request made is a POST request; this fact is not cause for alarm.

- The third variant we observed among a handful of participants resembles the second variant in its use of a cursor, but the actual behavior is otherwise quite different. This variant tracks no paging metadata. Instead, an HTML element with the CSS selector '`a.uiMorePagerPrimary`' is embedded in the page markup, containing the URL for the next timeline segment in an attribute '`ajaxify`'. The endpoint called is '`/profile/fig/timeline`', with the cursor and other metadata included as URL parameters; the actual request made is a POST request, as in the second variant. The HTML included in the response may contain another pager element if there are timeline segments following the one returned.

6.3 Limitations imposed by data management policy

Because of the large volume of highly personal data collected as part of our study, we chose to take the strongest measures possible to ensure participant privacy. As a result, we were required to omit some features that we had originally planned for the Polyhedron tool, as they would have required transmission of this data to a third party. In particular, we originally wanted to include a visualization of inferred personality trait metrics computed using methods similar to the work by Kosinski et. al. [11]. However, the machine-learning models necessary for such methods are not readily available. There are Web-based services that provide interfaces to such models; for instance, Apply Magic Sauce[1] provides one such service generating predictions from lists of Facebook Likes. Usage of any such Web-based service, though, requires transmission of personal data by definition, and so we opted to not include this in the version of the Polyhedron tool used in the study. There also exist models providing similar metrics from analysis of actual text; for instance, IBM Watson [2] provides a number of sophisticated text analysis tools. However, relying on a Web-based third-party tool like the Watson cloud product for post text analysis would present even more privacy problems, since the text content of a person's social media timeline can be quite revealing – and, in many cases, reliably identifying. Hence, although the added information from services like these would have been of great use to us, we did not include them in the study.

The method of data collection that we used requires a logged-in session in order to gather information. We had no wish to take actual physical possession of any participant's login credentials or session tokens; in the event that a participant were to raise concerns, we believed it would be easier to argue the trustworthiness of our code, rather than of ourselves, through methods like code audits if necessary. As such, we chose to run the data collection process directly on participants' own hardware, ensuring that their login credentials and session tokens never left their own control. This decision, however, caused significant inconveniences for us as experimenters, and was a major contributor to the difficulty of completing this work. Whenever a participant's data contained omissions reflecting changes or variant behavior in the underlying social media platform, we were required to identify the cause of the variance, alter our software to handle the change properly, update the participant's local copy of the tool appropriately, and attempt to recollect the missing data. The inherent nature of software development often requiring multiple iterations to complete, combined with the relatively long time required to collect the necessary data (in some cases,

[1]https://applymagicsauce.com/
[2]https://www.ibm.com/watson/developercloud/

several hours per iteration), resulted in fairly severe delays for the entire data collection process; in extreme cases, one or more weeks could be required to collect a usable data sample. In all cases unwanted disclosure or loss of control of personal data was avoided, but at great cost in terms of time and efficiency.

Given the practical difficulties in operating under such a restricted data disclosure philosophy, we caution against insisting on such comprehensive constraints on data handling *a priori*. In particular, the following concessions and other actions should be investigated, some of which are also expressed in prior sections:

- An automatic update mechanism and an automatic data upload mechanism would significantly reduce the length of the software iteration cycle, especially if both can be designed to require no participant intervention. Making both mechanisms opt-out might be acceptable, and reduce the amount of researcher workload required.

- It may be possible to obtain participant consent for use of third-party services to perform some types data analysis on an opt-in basis, as long as doing so can provide some value to the participant (for instance, as part of an informative visualization aid with which the participant can interact).

6.4 Issues arising from participant data volume

To say that some people are quite heavily involved on social media is a large understatement. While we had some anticipation of the large volume of data of some of our participants, in many cases this great volume ended up causing complications for our collection process.

Web-based APIs will sometimes return an error result due to problems on the service provider side. In general, this is a low probability event per API call, and it is normally in the service provider's interest to minimise the probability to ensure a smooth experience for its customers. However, since some participants may have data requiring tens of thousands of API calls to retrieve in entirety, even low-probability events affecting data collection become highly probable or even near-certainties. In some cases, we encountered error conditions that did not successfully resolve even after retrying the failed call multiple times with delays, requiring us to restart collection of particular sets of data. Dealing with these situations gracefully was a challenge for us, and any extension of this work must take care to handle them properly and provide sufficient debugging information in the case of an unresolvable API failure.

In a handful of cases, managing the sheer amount of participant data was itself an issue. Firefox provides an intrinsic limit on the maximum length of any string, which made it difficult to serialize and save collected data if the JSON serialization ended up being larger than around 250 MB. Any similar work using JavaScript should be mindful of this limitation and be prepared to work around it. For instance, if the tool itself can condense information rather than relying on external postprocessing, significantly less information needs to be saved (though still enough to visualize later). In cases where even condensed information exceeds the bounds of available string serialization, data may be saved in pieces to different files.

6.5 Reflections

For data collection, we chose the route of relying primarily on directly making frontend API calls, emulating the communication behavior of the normal Web products, after extensive experimentation with directly manipulating a hidden tab by emulating human interaction, rejecting the latter due to performance issues associated with doing so when participants had large amounts of data and when too many such hidden tabs were being operated simultaneously. However, the route of using direct API calls ended up causing more than its own fair share of inconveniences. Primarily, there were serious delays resulting from the need to keep up with API changes; we found that, generally, the underlying API details tended to change somewhat more often than the resulting rendered HTML. Additionally, there were serious time costs incurred from attempting to adapt to the behavioural variance of the Facebook website among participants that we observed, due to the frequent need to iterate multiple times with affected participants. These issues could be mitigated or eliminated with a combination of proactive detection of changes, transparent software update, and transparent recollection of missing data, if the necessary participant consent can be obtained.

In cases where the necessary information can be reliably harvested from the DOM, it may very well be preferable to find some reasonable way to avoid the less-stable internal APIs and concentrate on pulling data directly from the rendered page, so long as the performance issues can be reliably overcome. One potential way to manage the runtime might be to remove DOM elements as they are processed; we were unable to make this work reliably, but by no means should this be taken as an assertion of impossibility. In cases where distinguishing data is not made impractical by ambiguous page layout, harvesting directly from the DOM may benefit from greater stability of rendered output. However, this does not eliminate the requirement to adapt quickly to any changes that do occur on the platform side. Given the rate at which active social networking websites deploy new features and changes, addressing the constant need to adapt to platform changes and distribute updates quickly – or, indeed, side-stepping the problem by using official API mechanisms or establishing research relationships directly with the target platform vendors – will be essential for replicating or extending this work.

Chapter 7

Conclusions

7.1 Summary of findings

The power of our investigation was ultimately restricted by its fairly small sample size, as – for reasons elaborated on earlier – participant data collection ended up being much more time-consuming than we had originally anticipated. Nevertheless, we were able to address most of our research questions, as applied to our sample population, using the methods presented. We summarize what we learned about each of our original hypotheses in Table 7.1 and Table 7.2.

We observed from our user study that people in our test set indeed have some difficulty estimating how frequently they interact with many of their contacts on social media – believing, by and large, that they interacted with more people than they did in practice (RQ1, H1). This is a telling result: it suggests that, when people do use social media, their mental image of interaction on every network – even those where they are supposedly posting for public consumption – tends to de-emphasize or even exclude the "silent audience". People also had some difficulty estimating the number of their cross-network connections with any degree of accuracy, though we found most of these sets to be small in practice (RQ1, H2).

We noticed that people tend to report interacting with different networks as if each one is a separate type of audience, and indeed, we find platform-specific differences both in the general types of posts made and in the contextual data – captions, location tags, and person tags – provided along with posts (RQ2, H3-H5). Some people are not completely aware, though, of how they have used each network in the past, or of what types of content they have previously posted, especially when they interact with each network frequently and in many different ways (H6, H7).

These look, in some ways, like conditions previously reported as conducive for problematic posts, although people report sharing strategies – careful consideration of audience on one hand, and treatment of spaces as public on the other hand – that one might expect to mitigate unwanted information disclosure. What happens when we show people their posting activity and their connections, grouped by network and grouped by activity? Most people (8) in our test set indeed found little reason to worry – they were fine with the post sharing decisions they had made in the past, and reported that the connections present on multiple networks were people who, indeed, they believed still belonged in every represented audience. The highlighting of connections with limited or no interaction, though, was enough to give a small number (3) of people pause, and prompt them to rethink whether all of those connections still made sense to retain. And, in one particular case (P11), our ability to highlight timeline content by visibility setting helped draw attention to a number of posts that, left alone, may have caused embarrassment or even reputational damage in the future. Hence, even though most of our population was comfortable with their sharing settings even after detailed review, we still reaffirm the value of data visualization in examining potential privacy issues on social media (RQ3, H8).

	Description	Conclusion
H1	People are not able to accurately estimate the distribution of their connections by frequency of social media engagement across all platforms.	Accepted: on all networks the category of least interaction is highly underestimated on average.
H2	People are not able to accurately estimate the distribution of their connections by which networks they are connected.	Promising but inconclusive: positive mean percentage error for all multiple-network sets (though more people underestimate the FB/LI set), but connection deduplication could account for much of the difference as sets are small in practice.
H3	There is a statistically significant difference between networks in the frequency with which people share text versus non-text posts, on the networks where both are possible.	Accepted: statistically significant difference between Twitter and Facebook.
H4	There is a statistically significant difference between networks in the frequency with which people tag other people relevant to a post, on the networks where this is possible.	Accepted for text; not accepted for photos or videos. Possibly related to differences in mode of interaction between networks (name mentions much more prevalent on Twitter than Facebook).
H5	There is a statistically significant difference between networks in the frequency with which people attach location information to a post, on the networks where this is possible.	Accepted for photos; not accepted for videos or text.

Table 7.1: Summary of hypotheses investigated and what we learned about each (H1-H5).

	Description	Conclusion
H6	People are not able to accurately estimate the composition of their own social media timelines by post type.	Accepted for Facebook: text content on timelines was underestimated compared to photo content. Not accepted for Instagram or Twitter, which show no decisive trends and relatively small error rate.
H7	People are not able to accurately estimate the proportions of activity types they undertake on each network.	Accepted for Facebook: people do make large estimation errors for Facebook activity, and overall people using Facebook do underestimate the frequency of reacting and commenting. Not accepted for Twitter, which shows no decisive trends and relatively smaller error rate.
H8	Viewing of one's own visualized composite profile data, along with information about audience and reach, with the ability to highlight the set of information visible to any particular connection or category of people, results in a measurable shift in one's opinions about privacy risk.	Promising but not generally accepted for this population: most (8) reported being content with their own sharing policies, with a minority (4) expressing a desire to change behavior.

Table 7.2: Summary of hypotheses investigated and what we learned about each (H6-H8).

7.2 Future work

7.2.1 Potential insights from text and photo content

A thread of investigation we considered but did not ultimately pursue was analysis of the actual text and photo content of participant post streams. Many different projects, both commercial and research, provide an ability to analyze a given corpus of text and predict certain psychometric qualities, like tone and personality traits. Additionally, researchers have attempted to use photographs on social media to determine both personality traits [15] and psychological well-being [21]. Including such information would increase the breadth of questions available for investigation: one could ask, for example, what type of personality one wants to present on social media to each potential audience, or what types of personality traits one wants to *hide*, and determine if that individual's sharing settings work to accomplish those desires.

The main difficulty with this task, in our view, is accomplishing it in such a way that the participants and experimenters can trust the integrity of the process. Because questions like these rely on processing highly sensitive personal data, we considered it a negative to have to send any information to servers outside our control, or to code beyond our understanding. Indeed, at every point in the study, we were fully prepared to conduct the entire data analysis process directly on participants' computers, and to explain in detail any portion of how the data collection process worked. In order to preserve this property, barring useful advancements in running machine learning algorithms on private data, we would need to acquire a predictive model of the right type in a form that could be run "locally" if needed (which is subject to the willingness of that model's owner to lend it for our use), or develop our own (which is subject to our ability to collect a reasonable quantity of training data). Otherwise, we would have to obtain participant consent to transmit significant portions of their own data to third parties, and adapt to any attendant change in the participation rate.

7.2.2 Examination of additional networks

We chose the networks we focused on – Facebook, Twitter, Instagram, and LinkedIn – based on high perceived usage share. However, during interviews, participants mentioned other social networks on which they spend a significant fraction of time and attention. Both Snapchat and WhatsApp received mentions from multiple participants, and indeed, both networks are quite popular: Snapchat had over 173 million active members as of mid-2017 [9], while WhatsApp had over 1.3 billion [28]. Participants in some cases reported interacting with their close connections more frequently on these networks (and by other means) than on any of the four networks we chose, which would cause connection interaction figures to be underreported; incorporating interaction data from these networks, and others, would allow us to construct a more complete picture of a participant's interactions with others on social media.

It should be noted that the Snapchat platform, in particular, would be difficult to address with the methods we used. The product is a mobile-only application without official developer APIs, and so would require a different approach to determine how best to extract data. While others in the past have done extensive reverse-engineering work on the Snapchat application [7], Snapchat itself has taken an aggressive stance against third-party applications, as unauthorized third-party access to its platform has been used in the past to contravene the "disappearing messages" property the product attempts to provide [13]. Hence, it can be anticipated that extending this work to gather quantitative data from the Snapchat platform might be tricky at best, given both the limited API access and the limited post history that appears to be available.

Appendix A

Supplemental graphs

This section contains supplemental graphs, including graphs of content proportions and activity proportions.

A.1 Content proportions over time

Each of these graphs shows, for each type of content, the proportion estimate provided in the survey (dotted line), the actual computed value over the entire timeline (dashed line), and the computed value considering only content within each individual timeline year (time series with solid lines).

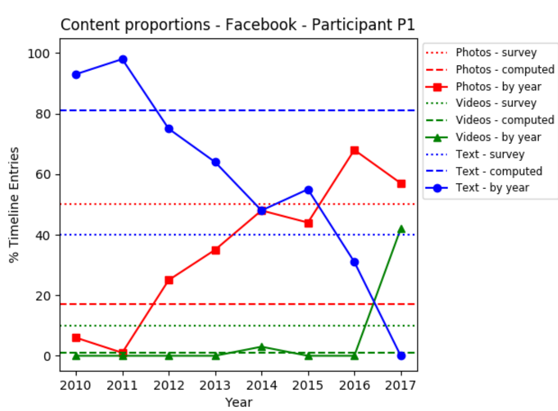

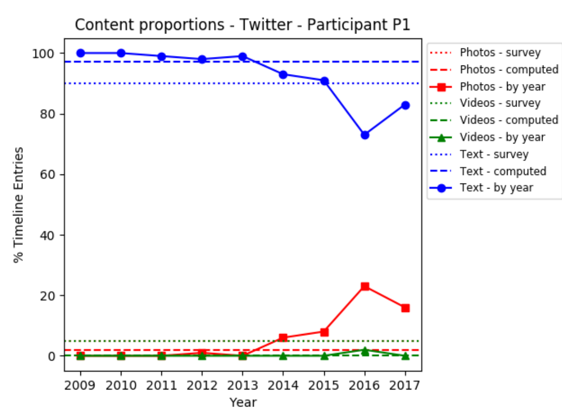

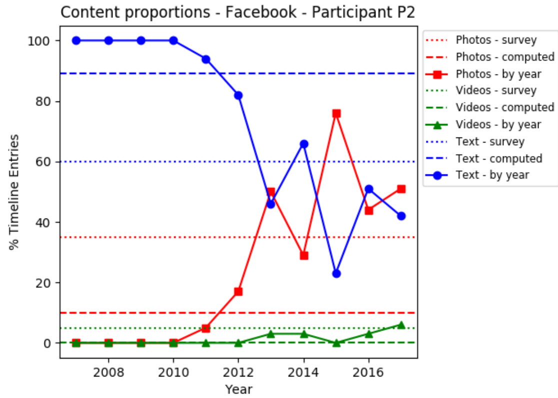
Content proportions - Facebook - Participant P2
% Timeline Entries
Year
0
20
40
60
80
100
2008
2010
2012
2014
2016
Photos - survey
Photos - computed
Photos - by year
Videos - survey
Videos - computed
Videos - by year
Text - survey
Text - computed
Text - by year

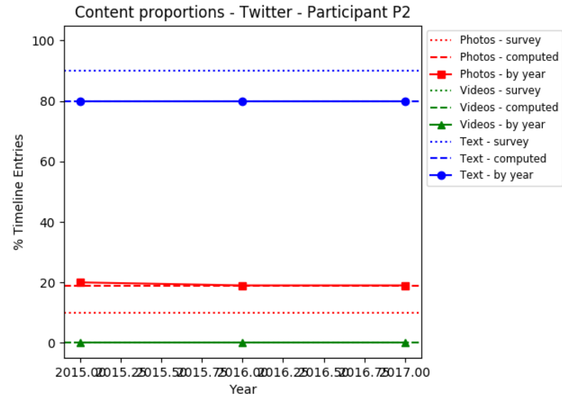
Content proportions - Twitter - Participant P2
% Timeline Entries
Year
0
20
40
60
80
100
2015.00
2015.25
2015.50
2015.75
2016.00
2016.25
2016.50
2016.75
2017.00
Photos - survey
Photos - computed
Photos - by year
Videos - survey
Videos - computed
Videos - by year
Text - survey
Text - computed
Text - by year

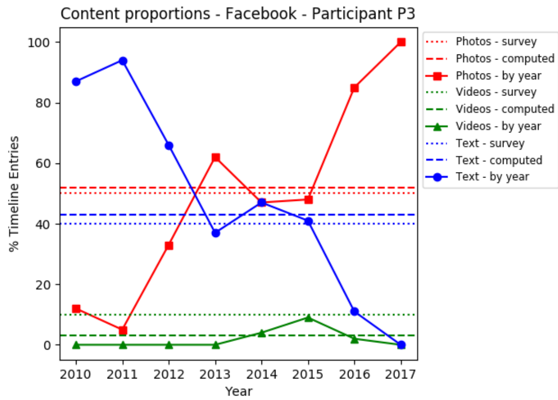
Content proportions - Facebook - Participant P3
% Timeline Entries
Year
0
20
40
60
80
100
2010
2011
2012
2013
2014
2015
2016
2017
Photos - survey
Photos - computed
Photos - by year
Videos - survey
Videos - computed
Videos - by year
Text - survey
Text - computed
Text - by year

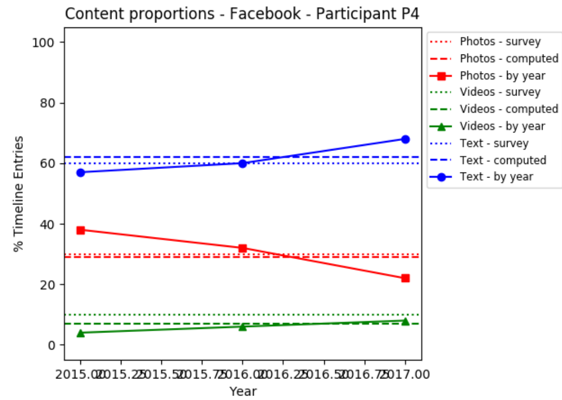
Content proportions - Facebook - Participant P4
% Timeline Entries
Year
0
20
40
60
80
100
2015.00
2015.25
2015.50
2015.75
2016.00
2016.25
2016.50
2016.75
2017.00
Photos - survey
Photos - computed
Photos - by year
Videos - survey
Videos - computed
Videos - by year
Text - survey
Text - computed
Text - by year

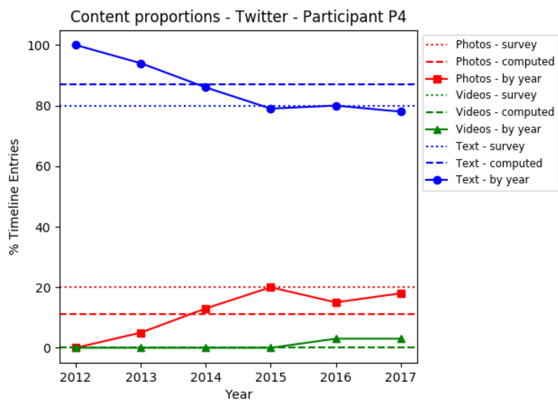
Content proportions - Twitter - Participant P4
% Timeline Entries
Year
0
20
40
60
80
100
2012
2013
2014
2015
2016
2017
Photos - survey
Photos - computed
Photos - by year
Videos - survey
Videos - computed
Videos - by year
Text - survey
Text - computed
Text - by year

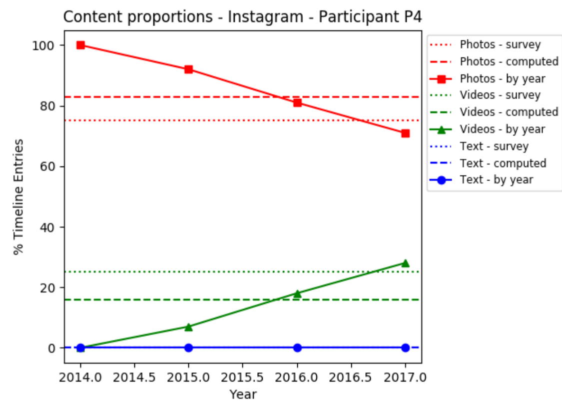
Content proportions - Instagram - Participant P4
% Timeline Entries
Year
0
20
40
60
80
100
2014.0
2014.5
2015.0
2015.5
2016.0
2016.5
2017.0
Photos - survey
Photos - computed
Photos - by year
Videos - survey
Videos - computed
Videos - by year
Text - survey
Text - computed
Text - by year

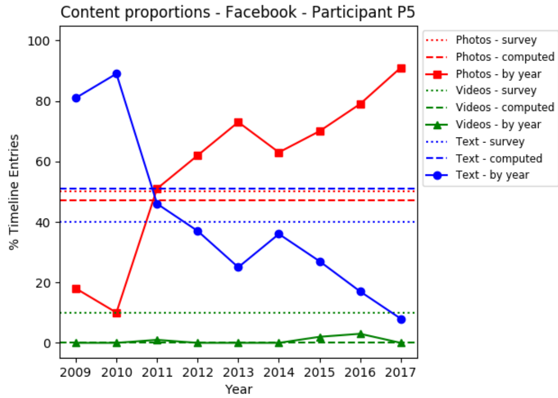
Content proportions - Facebook - Participant P5
% Timeline Entries
Year
2009 2010 2011 2012 2013 2014 2015 2016 2017
0 20 40 60 80 100
Photos - survey
Photos - computed
Photos - by year
Videos - survey
Videos - computed
Videos - by year
Text - survey
Text - computed
Text - by year

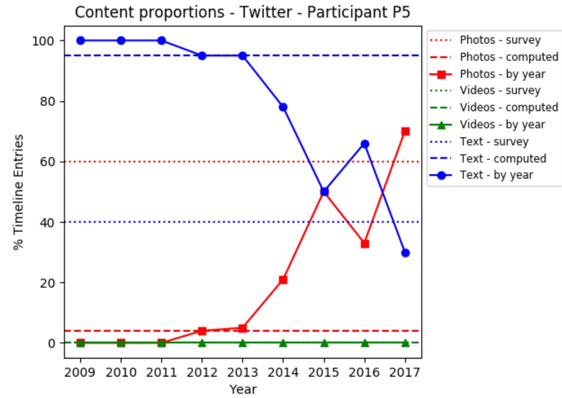
Content proportions - Twitter - Participant P5
% Timeline Entries
Year
2009 2010 2011 2012 2013 2014 2015 2016 2017
0 20 40 60 80 100
Photos - survey
Photos - computed
Photos - by year
Videos - survey
Videos - computed
Videos - by year
Text - survey
Text - computed
Text - by year

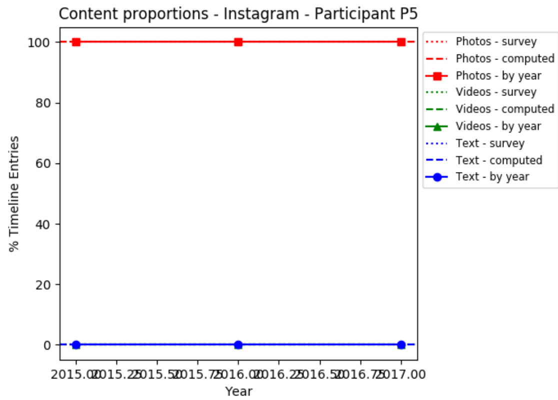
Content proportions - Instagram - Participant P5
% Timeline Entries
Year
0 20 40 60 80 100
Photos - survey
Photos - computed
Photos - by year
Videos - survey
Videos - computed
Videos - by year
Text - survey
Text - computed
Text - by year

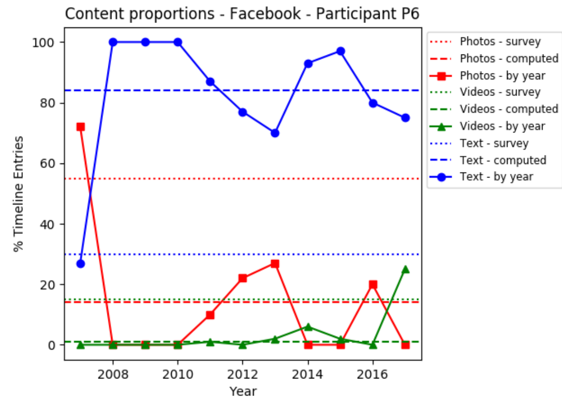
Content proportions - Facebook - Participant P6
% Timeline Entries
Year
2008 2010 2012 2014 2016
0 20 40 60 80 100
Photos - survey
Photos - computed
Photos - by year
Videos - survey
Videos - computed
Videos - by year
Text - survey
Text - computed
Text - by year

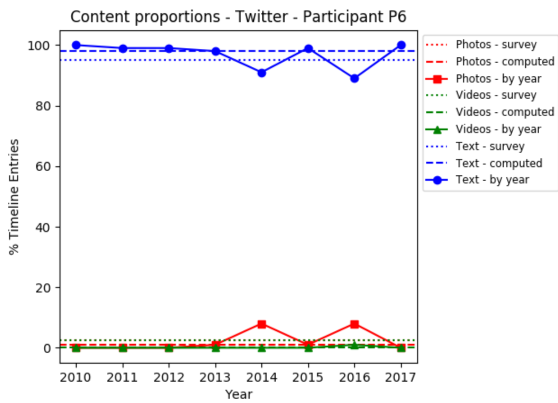
Content proportions - Twitter - Participant P6
% Timeline Entries
Year
2010 2011 2012 2013 2014 2015 2016 2017
0 20 40 60 80 100
Photos - survey
Photos - computed
Photos - by year
Videos - survey
Videos - computed
Videos - by year
Text - survey
Text - computed
Text - by year

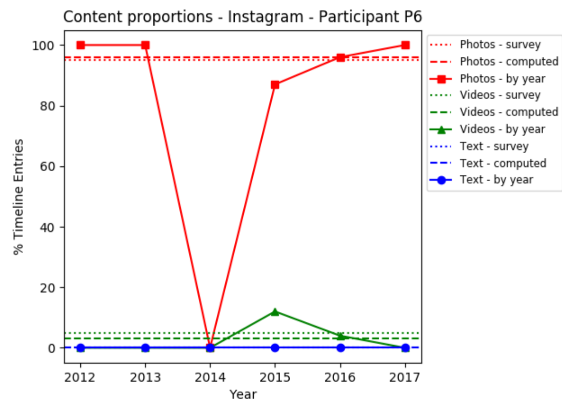
Content proportions - Instagram - Participant P6
% Timeline Entries
Year
2012 2013 2014 2015 2016 2017
0 20 40 60 80 100
Photos - survey
Photos - computed
Photos - by year
Videos - survey
Videos - computed
Videos - by year
Text - survey
Text - computed
Text - by year

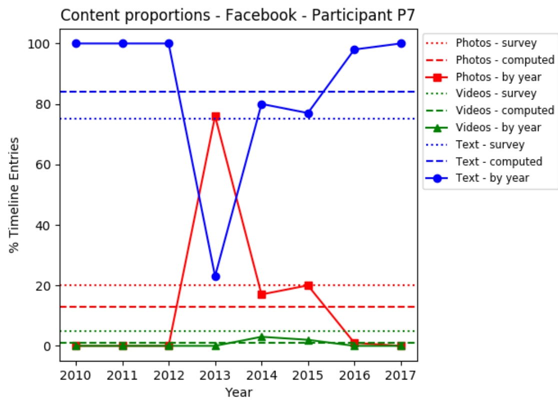
Content proportions - Facebook - Participant P7
% Timeline Entries
Year
Photos - survey
Photos - computed
Photos - by year
Videos - survey
Videos - computed
Videos - by year
Text - survey
Text - computed
Text - by year

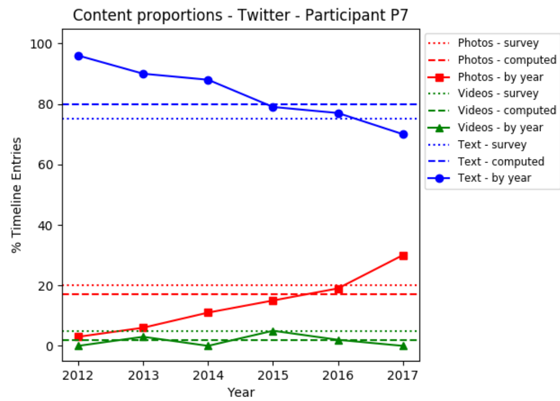
Content proportions - Twitter - Participant P7
% Timeline Entries
Year
Photos - survey
Photos - computed
Photos - by year
Videos - survey
Videos - computed
Videos - by year
Text - survey
Text - computed
Text - by year

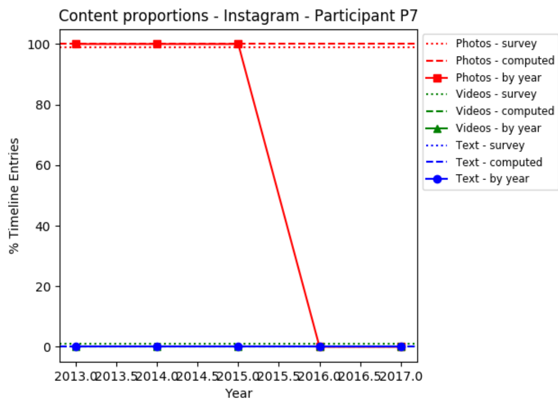
Content proportions - Instagram - Participant P7
% Timeline Entries
Year
Photos - survey
Photos - computed
Photos - by year
Videos - survey
Videos - computed
Videos - by year
Text - survey
Text - computed
Text - by year

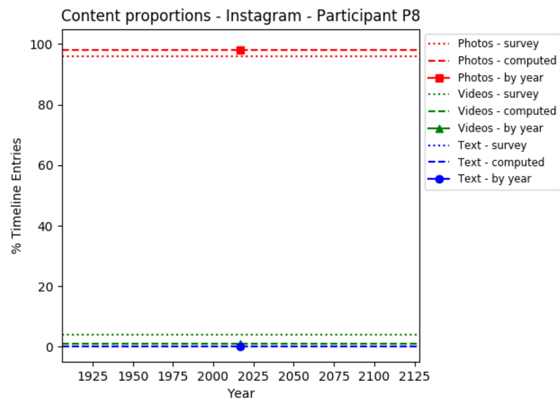
Content proportions - Instagram - Participant P8
% Timeline Entries
Year
Photos - survey
Photos - computed
Photos - by year
Videos - survey
Videos - computed
Videos - by year
Text - survey
Text - computed
Text - by year

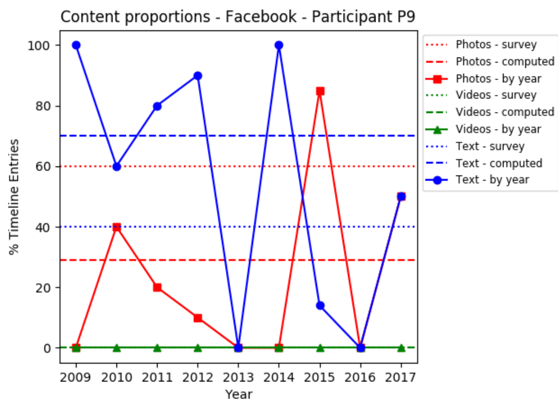
Content proportions - Facebook - Participant P9
% Timeline Entries
Year
Photos - survey
Photos - computed
Photos - by year
Videos - survey
Videos - computed
Videos - by year
Text - survey
Text - computed
Text - by year

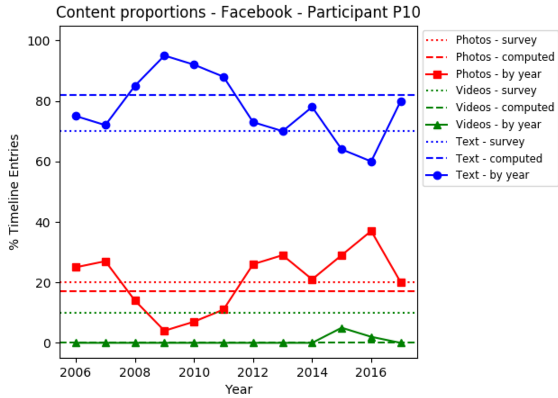
Content proportions - Facebook - Participant P10
% Timeline Entries
Year
0
20
40
60
80
100
2006
2008
2010
2012
2014
2016
Photos - survey
Photos - computed
Photos - by year
Videos - survey
Videos - computed
Videos - by year
Text - survey
Text - computed
Text - by year

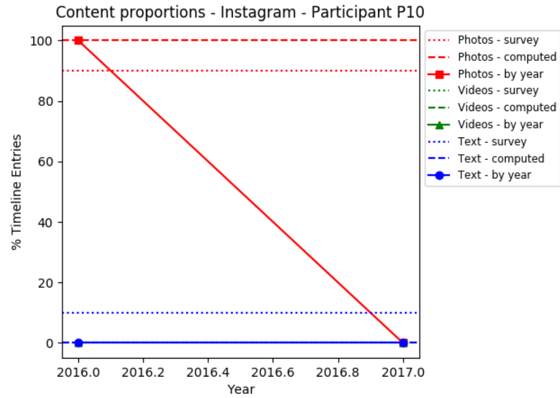
Content proportions - Instagram - Participant P10
% Timeline Entries
Year
0
20
40
60
80
100
2016.0
2016.2
2016.4
2016.6
2016.8
2017.0
Photos - survey
Photos - computed
Photos - by year
Videos - survey
Videos - computed
Videos - by year
Text - survey
Text - computed
Text - by year

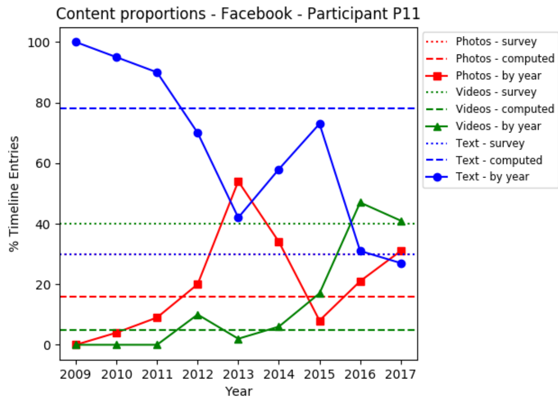
Content proportions - Facebook - Participant P11
% Timeline Entries
Year
0
20
40
60
80
100
2009
2010
2011
2012
2013
2014
2015
2016
2017
Photos - survey
Photos - computed
Photos - by year
Videos - survey
Videos - computed
Videos - by year
Text - survey
Text - computed
Text - by year

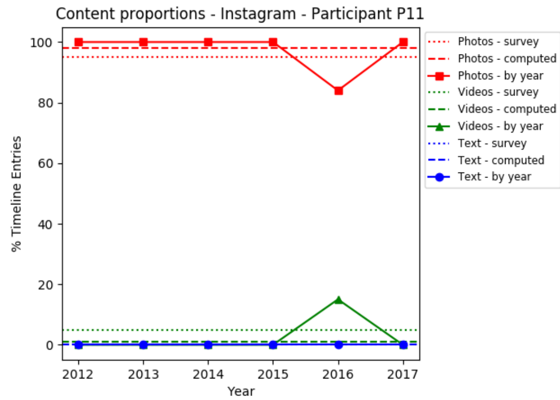
Content proportions - Instagram - Participant P11
% Timeline Entries
Year
0
20
40
60
80
100
2012
2013
2014
2015
2016
2017
Photos - survey
Photos - computed
Photos - by year
Videos - survey
Videos - computed
Videos - by year
Text - survey
Text - computed
Text - by year

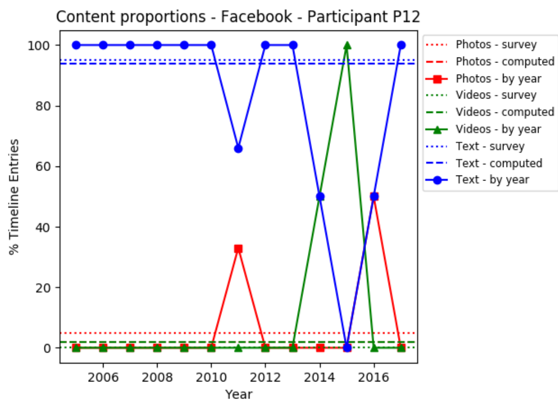
Content proportions - Facebook - Participant P12
% Timeline Entries
Year
0
20
40
60
80
100
2006
2008
2010
2012
2014
2016
Photos - survey
Photos - computed
Photos - by year
Videos - survey
Videos - computed
Videos - by year
Text - survey
Text - computed
Text - by year

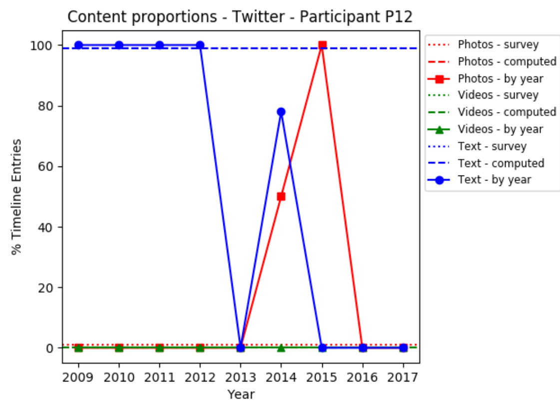
Content proportions - Twitter - Participant P12
% Timeline Entries
Year
0
20
40
60
80
100
2009
2010
2011
2012
2013
2014
2015
2016
2017
Photos - survey
Photos - computed
Photos - by year
Videos - survey
Videos - computed
Videos - by year
Text - survey
Text - computed
Text - by year

A.2 Activity proportions over time

Each of these graphs shows, for each type of activity, the proportion estimate provided in the survey (dotted line), the actual computed value over the entire timeline (dashed line), and the computed value considering only activity metrics within each individual timeline year (time series with solid lines).

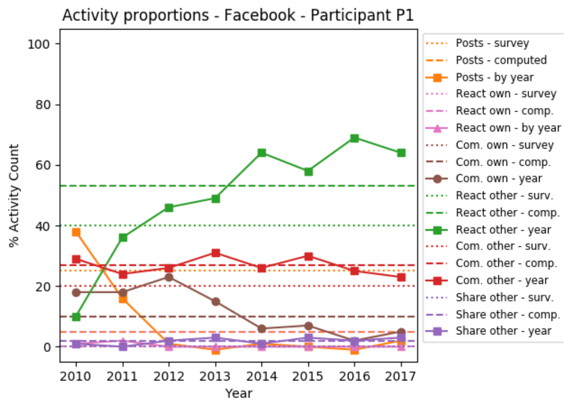

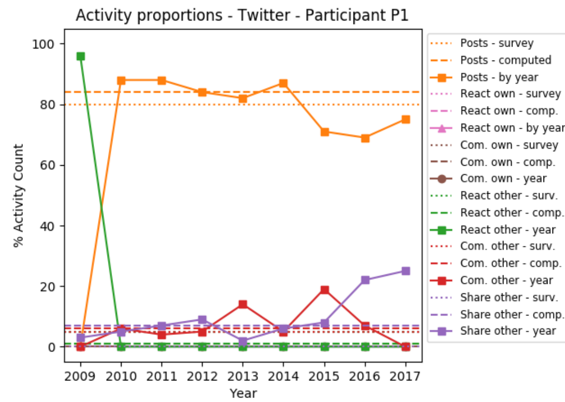

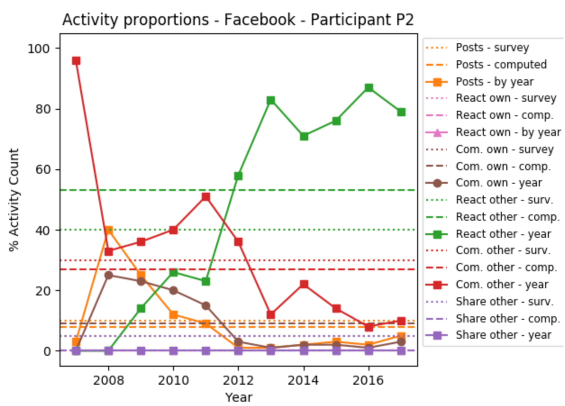

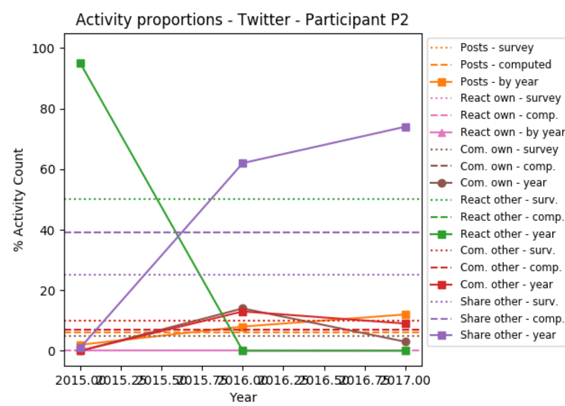

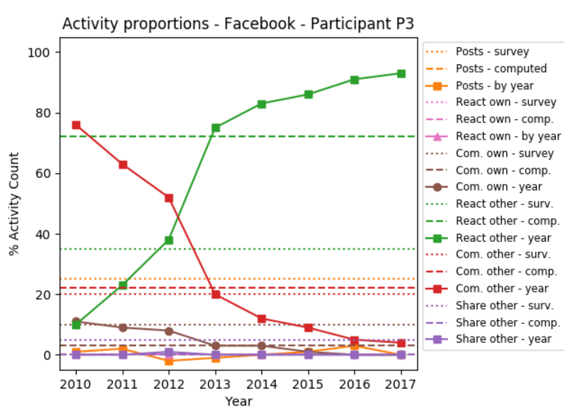

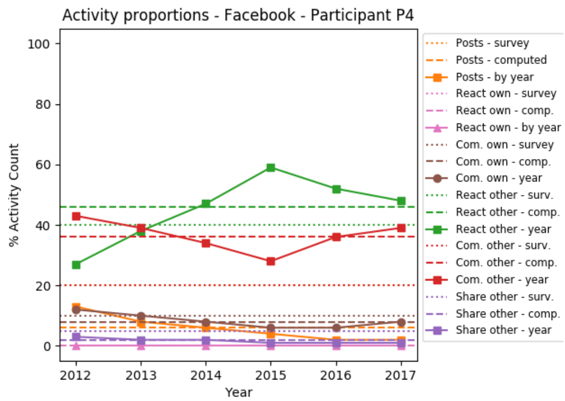
Activity proportions - Facebook - Participant P4
% Activity Count
Year
Posts - survey
Posts - computed
Posts - by year
React own - survey
React own - comp.
React own - by year
Com. own - survey
Com. own - comp.
Com. own - year
React other - surv.
React other - comp.
React other - year
Com. other - surv.
Com. other - comp.
Com. other - year
Share other - surv.
Share other - comp.
Share other - year

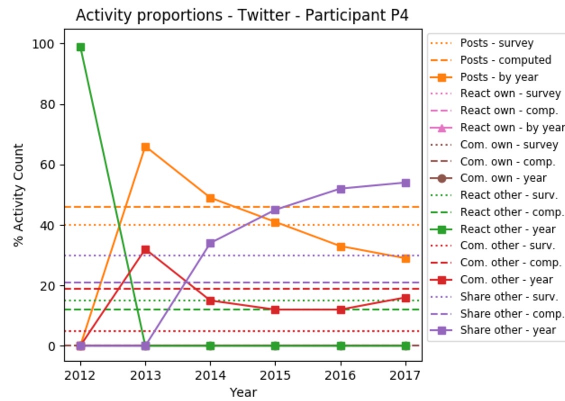
Activity proportions - Twitter - Participant P4
% Activity Count
Year
Posts - survey
Posts - computed
Posts - by year
React own - survey
React own - comp.
React own - by year
Com. own - survey
Com. own - comp.
Com. own - year
React other - surv.
React other - comp.
React other - year
Com. other - surv.
Com. other - comp.
Com. other - year
Share other - surv.
Share other - comp.
Share other - year

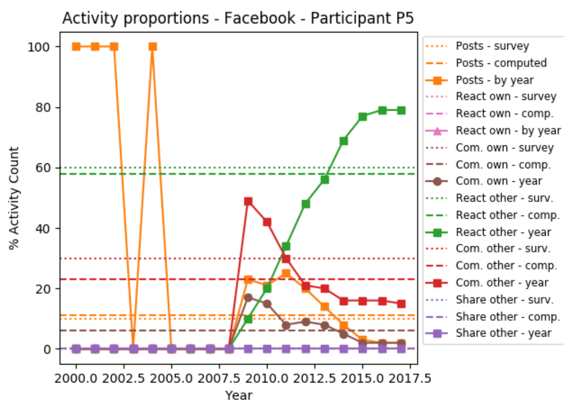
Activity proportions - Facebook - Participant P5
% Activity Count
Year
Posts - survey
Posts - computed
Posts - by year
React own - survey
React own - comp.
React own - by year
Com. own - survey
Com. own - comp.
Com. own - year
React other - surv.
React other - comp.
React other - year
Com. other - surv.
Com. other - comp.
Com. other - year
Share other - surv.
Share other - comp.
Share other - year

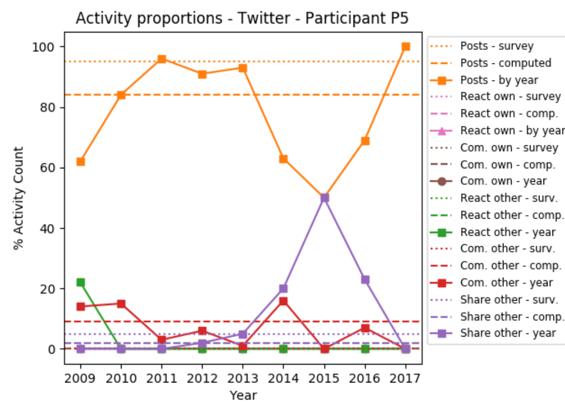
Activity proportions - Twitter - Participant P5
% Activity Count
Year
Posts - survey
Posts - computed
Posts - by year
React own - survey
React own - comp.
React own - by year
Com. own - survey
Com. own - comp.
Com. own - year
React other - surv.
React other - comp.
React other - year
Com. other - surv.
Com. other - comp.
Com. other - year
Share other - surv.
Share other - comp.
Share other - year

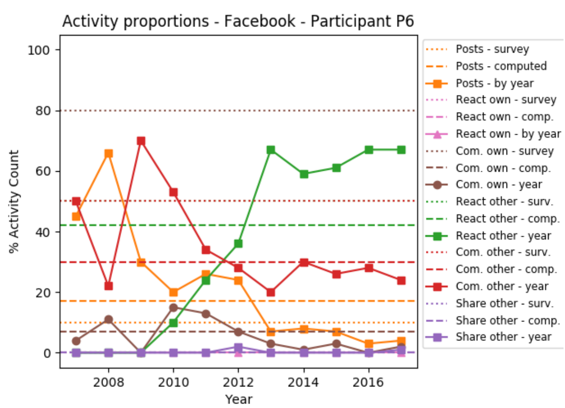
Activity proportions - Facebook - Participant P6
% Activity Count
Year
Posts - survey
Posts - computed
Posts - by year
React own - survey
React own - comp.
React own - by year
Com. own - survey
Com. own - comp.
Com. own - year
React other - surv.
React other - comp.
React other - year
Com. other - surv.
Com. other - comp.
Com. other - year
Share other - surv.
Share other - comp.
Share other - year

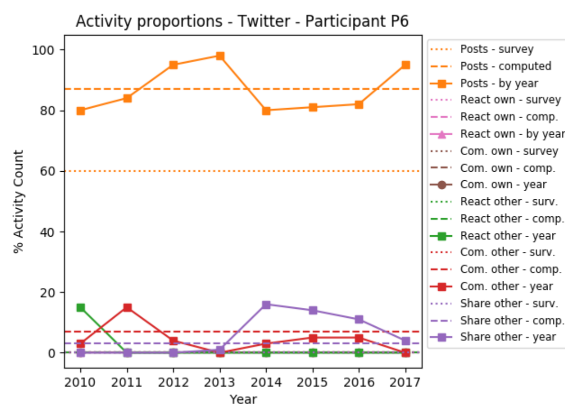
Activity proportions - Twitter - Participant P6
% Activity Count
Year
Posts - survey
Posts - computed
Posts - by year
React own - survey
React own - comp.
React own - by year
Com. own - survey
Com. own - comp.
Com. own - year
React other - surv.
React other - comp.
React other - year
Com. other - surv.
Com. other - comp.
Com. other - year
Share other - surv.
Share other - comp.
Share other - year

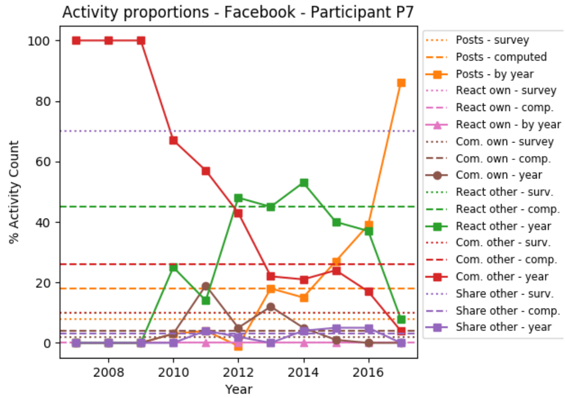
Activity proportions - Facebook - Participant P7
% Activity Count
Year
Posts - survey
Posts - computed
Posts - by year
React own - survey
React own - comp.
React own - by year
Com. own - survey
Com. own - comp.
Com. own - year
React other - surv.
React other - comp.
React other - year
Com. other - surv.
Com. other - comp.
Com. other - year
Share other - surv.
Share other - comp.
Share other - year

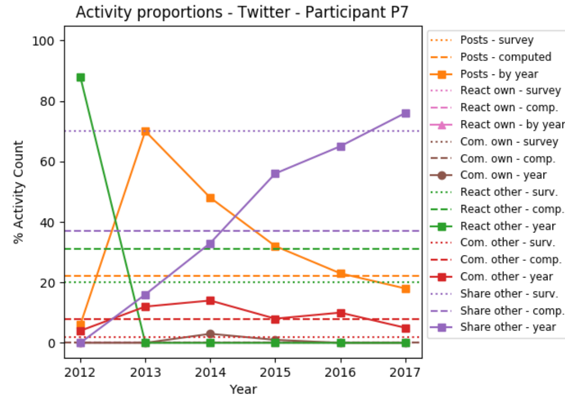
Activity proportions - Twitter - Participant P7
% Activity Count
Year
Posts - survey
Posts - computed
Posts - by year
React own - survey
React own - comp.
React own - by year
Com. own - survey
Com. own - comp.
Com. own - year
React other - surv.
React other - comp.
React other - year
Com. other - surv.
Com. other - comp.
Com. other - year
Share other - surv.
Share other - comp.
Share other - year

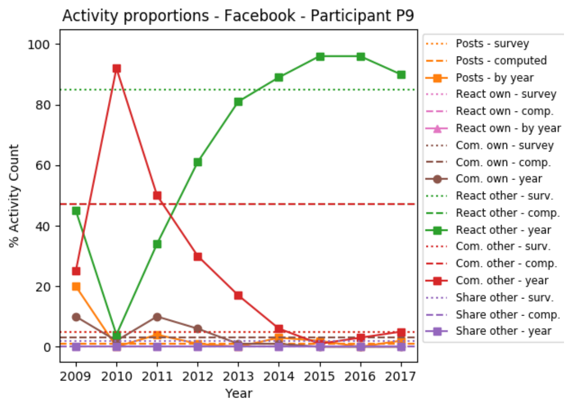
Activity proportions - Facebook - Participant P9
% Activity Count
Year
Posts - survey
Posts - computed
Posts - by year
React own - survey
React own - comp.
React own - by year
Com. own - survey
Com. own - comp.
Com. own - year
React other - surv.
React other - comp.
React other - year
Com. other - surv.
Com. other - comp.
Com. other - year
Share other - surv.
Share other - comp.
Share other - year

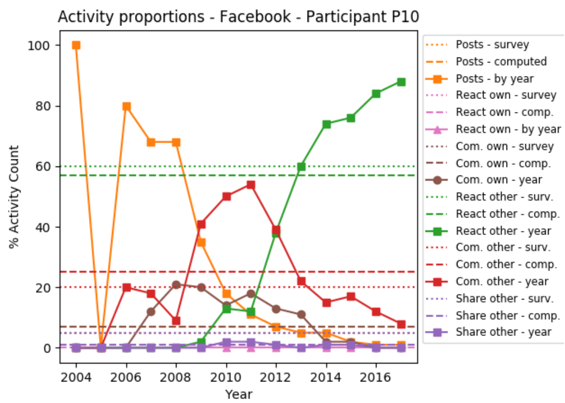
Activity proportions - Facebook - Participant P10
% Activity Count
Year
Posts - survey
Posts - computed
Posts - by year
React own - survey
React own - comp.
React own - by year
Com. own - survey
Com. own - comp.
Com. own - year
React other - surv.
React other - comp.
React other - year
Com. other - surv.
Com. other - comp.
Com. other - year
Share other - surv.
Share other - comp.
Share other - year

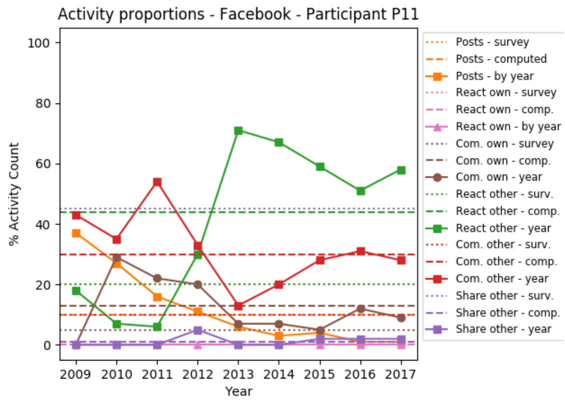
Activity proportions - Facebook - Participant P11
% Activity Count
Year
0
20
40
60
80
100
2009
2010
2011
2012
2013
2014
2015
2016
2017
Posts - survey
Posts - computed
Posts - by year
React own - survey
React own - comp.
React own - by year
Com. own - survey
Com. own - comp.
Com. own - year
React other - surv.
React other - comp.
React other - year
Com. other - surv.
Com. other - comp.
Com. other - year
Share other - surv.
Share other - comp.
Share other - year

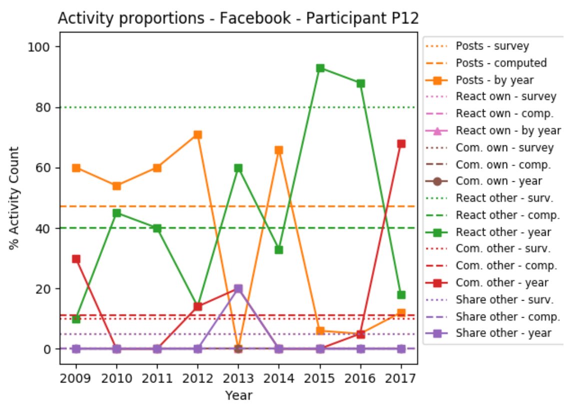
Activity proportions - Facebook - Participant P12
% Activity Count
Year
0
20
40
60
80
100
2009
2010
2011
2012
2013
2014
2015
2016
2017
Posts - survey
Posts - computed
Posts - by year
React own - survey
React own - comp.
React own - by year
Com. own - survey
Com. own - comp.
Com. own - year
React other - surv.
React other - comp.
React other - year
Com. other - surv.
Com. other - comp.
Com. other - year
Share other - surv.
Share other - comp.
Share other - year

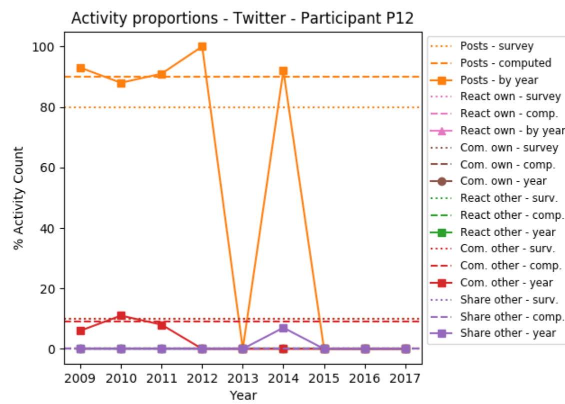
Activity proportions - Twitter - Participant P12
% Activity Count
Year
0
20
40
60
80
100
2009
2010
2011
2012
2013
2014
2015
2016
2017
Posts - survey
Posts - computed
Posts - by year
React own - survey
React own - comp.
React own - by year
Com. own - survey
Com. own - comp.
Com. own - year
React other - surv.
React other - comp.
React other - year
Com. other - surv.
Com. other - comp.
Com. other - year
Share other - surv.
Share other - comp.
Share other - year

Appendix B

Survey

This section contains a reproduction of the survey as given to the participants during the user study for this work. Note that the following survey items were not used for the analysis presented in this work:

- Part 18 was omitted in its entirety.
- No analysis was performed on data from Part 5, Parts 7-9, or Part 12.

Part 1: Demographic information

Age

Gender ---------

Education

Employment

Part 2: Sharing behaviors

Describe, in your own words, which social networking applications you use (including networks not mentioned elsewhere in this survey) and your reasons for using each one. Be as specific as possible.

Potential motivations may include feature set (for instance, the ability to share pictures or add location information to posts), audience (public vs. private), content type (social, professional, other), and modes of interaction (messaging, blogging, etc).

Part 3: Social network participation

Check off the social networks among the following choices for which you have a profile and with which you have interacted within the last two weeks (even by merely logging in).

Facebook	Twitter	Instagram	LinkedIn
☐	☐	☐	☐

Part 4: Main uses

What do you use each social network for primarily? (Choose all that apply.)

	Facebook	Twitter	Instagram	LinkedIn
Social	☐	☐	☐	☐

	Facebook	Twitter	Instagram	LinkedIn
Professional	☐	☐	☐	☐
Other (explain below)	☐	☐	☐	☐

Part 5: Information visible to others

What kinds of information about you do you think can be determined by the general public through each of your social networks - even if you have set privacy controls? Choose all that apply.

	Facebook	Twitter	Instagram	LinkedIn
Name	☐	☐	☐	☐
Type of work	☐	☐	☐	☐
Place of work	☐	☐	☐	☐
Home address	☐	☐	☐	☐
City currently located	☐	☐	☐	☐
Gender	☐	☐	☐	☐
Age	☐	☐	☐	☐
Religion	☐	☐	☐	☐
Sexual orientation	☐	☐	☐	☐

	Facebook	Twitter	Instagram	LinkedIn
Spouse / partner name				
Likes / dislikes				
Places where I have been (based on location tags)				
Places where I have been (based on photos)				
What matters to me				
My educational history				
My interests				
My hobbies				
My skills				
My close friends				
How can people infer who your close friends are?				

Part 6: Most used medium

Select the medium you use most frequently on each network.

	Facebook	Twitter	Instagram	LinkedIn
Photos				

	Facebook	Twitter	Instagram	LinkedIn
Video	○	○	○	○
Text	○	○	○	○
I do not share content	●	●	●	●

Part 7: Plain text posts

Provide the relevant frequency order of how you share text posts within each network. For each type with which you share, assign a number from 1 to 8, where 1 is most common and 8 is least common. If you don't share with a particular type, leave it as 0. Don't repeat non-zero numbers within the same column (social network).

	Facebook	Twitter	Instagram	LinkedIn
Plain text	0	0	0	0
Plain text with link	0	0	0	0
Plain text with tagged people	0	0	0	0
Plain text with location	0	0	0	0
Plain text with location and tagged people	0	0	0	0
Plain text with location and link	0	0	0	0
Plain text with link and tagged people	0	0	0	0
Plain text with location, link, and tagged people	0	0	0	0

Part 8: Photo posts

Provide the relevant frequency order of how you share photos within each network. For each type with which you share, assign a number from 1 to 8, where 1 is most common and 8 is least common. If you don't share with a particular type, leave it as 0. Don't repeat non-zero numbers within the same column (social network).

	Facebook	Twitter	Instagram	LinkedIn
Plain photo(s)	0	0	0	0
Photo(s) with caption	0	0	0	0

	Facebook	Twitter	Instagram	LinkedIn
Photo(s) with tagged people	0	0	0	0
Photo(s) with location	0	0	0	0
Photo(s) with location and tagged people	0	0	0	0
Photo(s) with location and caption	0	0	0	0
Photo(s) with caption and tagged people	0	0	0	0
Photo(s) with location, caption, and tagged people	0	0	0	0

Part 9: Video posts

Provide the relevant frequency order of how you share videos within each network. For each type with which you share, assign a number from 1 to 8, where 1 is most common and 8 is least common. If you don't share with a particular type, leave it as 0. Don't repeat non-zero numbers within the same column (social network).

	Facebook	Twitter	Instagram	LinkedIn
Plain video(s)	0	0	0	0
Video(s) with caption	0	0	0	0
Video(s) with tagged people	0	0	0	0
Video(s) with location	0	0	0	0
Video(s) with location and tagged people	0	0	0	0
Video(s) with location and caption	0	0	0	0
Video(s) with caption and tagged people	0	0	0	0
Video(s) with location, caption, and tagged people	0	0	0	0

Part 10: Content type proportion estimates

Estimate the percentages of posts of each type that you post on each social network. These should add up to 100%. If you do not share on a network, leave all values for that network at 0.

	Facebook	Twitter	Instagram	LinkedIn
Photos	0	0	0	0
Video	0	0	0	0
Text	0	0	0	0

Part 11: Most frequent action types

Among the following actions, which do you perform most frequently on each social network?

	Facebook	Twitter	Instagram	LinkedIn
I post material	○	○	○	○
I react to my own posted material	○	○	○	○
I comment on my own posted material	○	○	○	○
I react to people's posted material	○	○	○	○
I comment on people's posted material	○	○	○	○
I re-share people's posted material	○	○	○	○
I do not post on this network	◉	◉	◉	◉

Part 12: Action types by relative frequency

Provide a frequency ordering for each type of action you perform on each social network. Rate each action from 1 to 6, where 1 is most frequent, and 6 is least frequent. Do not reuse numbers within a single social network. If you do not perform an action on a given social network, give it a rating of 0.

	Facebook	Twitter	Instagram	LinkedIn
Posting material	0	0	0	0
Reacting to material you posted	0	0	0	0
Commenting on material you posted	0	0	0	0

	Facebook	Twitter	Instagram	LinkedIn
Reacting to material others posted	0	0	0	0
Commenting on material others posted	0	0	0	0
Resharing material others posted	0	0	0	0

Part 13: Action type proportions

Provide a percentage estimate for each type of action you perform on each social network. Try to make sure that the numbers add up to 100. If you do not perform a given action on a given network, write 0.

	Facebook	Twitter	Instagram	LinkedIn
Posting material	0	0	0	0
Reacting to material you posted	0	0	0	0
Commenting on material you posted	0	0	0	0
Reacting to material others posted	0	0	0	0
Commenting on material others posted	0	0	0	0
Resharing material others posted	0	0	0	0

Part 14: Visibility concerns

How concerned are you that your information can be viewed by strangers, i.e. by people you have not approved even if your network is set to private?

Choose the best answer that fits for each row.

Facebook

	Not used	Not at all concerned	Slightly concerned	Somewhat concerned	Moderately concerned	Extremely concerned
I post material (without tagging anyone)	(•)	( )	( )	( )	( )	( )
I post material where I tag friends	(•)	( )	( )	( )	( )	( )

	Not used	Not at all concerned	Slightly concerned	Somewhat concerned	Moderately concerned	Extremely concerned
I am tagged in a friend's post	◉	○	○	○	○	○
I react to my own posts (where I haven't tagged anyone)	◉	○	○	○	○	○
I react to my own posts (where I've tagged friends)	◉	○	○	○	○	○
I react to posts where I am tagged	◉	○	○	○	○	○
I react to my connections' posts	◉	○	○	○	○	○
I react to someone else's post in which a friend is tagged	◉	○	○	○	○	○
I react to a post made by a non-connection	◉	○	○	○	○	○
I comment on my own posts (where I haven't tagged anyone)	◉	○	○	○	○	○
I comment on my own posts (where I've tagged friends)	◉	○	○	○	○	○
I comment on posts where I am tagged	◉	○	○	○	○	○
I comment on my connections' posts	◉	○	○	○	○	○
I comment on someone else's post in which a friend is tagged	◉	○	○	○	○	○
I comment on a post made by a non-connection	◉	○	○	○	○	○

Twitter

	Not used	Not at all concerned	Slightly concerned	Somewhat concerned	Moderately concerned	Extremely concerned

	Not used	Not at all concerned	Slightly concerned	Somewhat concerned	Moderately concerned	Extremely concerned
I post material (without tagging anyone)	◉	○	○	○	○	○
I post material where I tag friends	◉	○	○	○	○	○
I am tagged in a friend's post	◉	○	○	○	○	○
I react to my own posts (where I haven't tagged anyone)	◉	○	○	○	○	○
I react to my own posts (where I've tagged friends)	◉	○	○	○	○	○
I react to posts where I am tagged	◉	○	○	○	○	○
I react to my connections' posts	◉	○	○	○	○	○
I react to someone else's post in which a friend is tagged	◉	○	○	○	○	○
I react to a post made by a non-connection	◉	○	○	○	○	○
I comment on my own posts (where I haven't tagged anyone)	◉	○	○	○	○	○
I comment on my own posts (where I've tagged friends)	◉	○	○	○	○	○
I comment on posts where I am tagged	◉	○	○	○	○	○
I comment on my connections' posts	◉	○	○	○	○	○
I comment on someone else's post in which a friend is tagged	◉	○	○	○	○	○
I comment on a post made by a non-connection	◉	○	○	○	○	○

Instagram

	Not used	Not at all concerned	Slightly concerned	Somewhat concerned	Moderately concerned	Extremely concerned
I post material (without tagging anyone)	◉	○	○	○	○	○
I post material where I tag friends	◉	○	○	○	○	○
I am tagged in a friend's post	◉	○	○	○	○	○
I react to my own posts (where I haven't tagged anyone)	◉	○	○	○	○	○
I react to my own posts (where I've tagged friends)	◉	○	○	○	○	○
I react to posts where I am tagged	◉	○	○	○	○	○
I react to my connections' posts	◉	○	○	○	○	○
I react to someone else's post in which a friend is tagged	◉	○	○	○	○	○
I react to a post made by a non-connection	◉	○	○	○	○	○
I comment on my own posts (where I haven't tagged anyone)	◉	○	○	○	○	○
I comment on my own posts (where I've tagged friends)	◉	○	○	○	○	○
I comment on posts where I am tagged	◉	○	○	○	○	○
I comment on my connections' posts	◉	○	○	○	○	○
I comment on someone else's post in which a friend is tagged	◉	○	○	○	○	○

	Not used	Not at all concerned	Slightly concerned	Somewhat concerned	Moderately concerned	Extremely concerned
I comment on a post made by a non-connection	◉	○	○	○	○	○

LinkedIn

	Not used	Not at all concerned	Slightly concerned	Somewhat concerned	Moderately concerned	Extremely concerned
I post material (without tagging anyone)	◉	○	○	○	○	○
I react to my own posts (where I haven't tagged anyone)	◉	○	○	○	○	○
I react to my connections' posts	◉	○	○	○	○	○
I react to a post made by a non-connection	◉	○	○	○	○	○
I comment on my own posts (where I haven't tagged anyone)	◉	○	○	○	○	○
I comment on my connections' posts	◉	○	○	○	○	○
I comment on a post made by a non-connection	◉	○	○	○	○	○

Part 15: Control over visibility

How much control do you believe you have over the visibility of your actions on each social network?

Choose the best answer that fits for each row.

Facebook

	Not used	No control	Low control	Medium control	High control	Full control
I post material (without tagging anyone)	◉	○	○	○	○	○
I post material where I tag friends	◉	○	○	○	○	○

	Not used	No control	Low control	Medium control	High control	Full control
I am tagged in a friend's post	●	○	○	○	○	○
I react to my own posts (where I haven't tagged anyone)	●	○	○	○	○	○
I react to my own posts (where I've tagged friends)	●	○	○	○	○	○
I react to posts where I am tagged	●	○	○	○	○	○
I react to my connections' posts	●	○	○	○	○	○
I react to someone else's post in which a friend is tagged	●	○	○	○	○	○
I react to a post made by a non-connection	●	○	○	○	○	○
I comment on my own posts (where I haven't tagged anyone)	●	○	○	○	○	○
I comment on my own posts (where I've tagged friends)	●	○	○	○	○	○
I comment on posts where I am tagged	●	○	○	○	○	○
I comment on my connections' posts	●	○	○	○	○	○
I comment on someone else's post in which a friend is tagged	●	○	○	○	○	○
I comment on a post made by a non-connection	●	○	○	○	○	○

Twitter

	Not used	No control	Low control	Medium control	High control	Full control
I post material (without tagging anyone)	●	○	○	○	○	○
I post material where I tag friends	●	○	○	○	○	○
I am tagged in a friend's post	●	○	○	○	○	○

	Not used	No control	Low control	Medium control	High control	Full control
I react to my own posts (where I haven't tagged anyone)	◉	○	○	○	○	○
I react to my own posts (where I've tagged friends)	◉	○	○	○	○	○
I react to posts where I am tagged	◉	○	○	○	○	○
I react to my connections' posts	◉	○	○	○	○	○
I react to someone else's post in which a friend is tagged	◉	○	○	○	○	○
I react to a post made by a non-connection	◉	○	○	○	○	○
I comment on my own posts (where I haven't tagged anyone)	◉	○	○	○	○	○
I comment on my own posts (where I've tagged friends)	◉	○	○	○	○	○
I comment on posts where I am tagged	◉	○	○	○	○	○
I comment on my connections' posts	◉	○	○	○	○	○
I comment on someone else's post in which a friend is tagged	◉	○	○	○	○	○
I comment on a post made by a non-connection	◉	○	○	○	○	○

Instagram

	Not used	No control	Low control	Medium control	High control	Full control
I post material (without tagging anyone)	◉	○	○	○	○	○
I post material where I tag friends	◉	○	○	○	○	○
I am tagged in a friend's post	◉	○	○	○	○	○
I react to my own posts (where I haven't tagged anyone)	◉	○	○	○	○	○

	Not used	No control	Low control	Medium control	High control	Full control
I react to my own posts (where I've tagged friends)	●	○	○	○	○	○
I react to posts where I am tagged	●	○	○	○	○	○
I react to my connections' posts	●	○	○	○	○	○
I react to someone else's post in which a friend is tagged	●	○	○	○	○	○
I react to a post made by a non-connection	●	○	○	○	○	○
I comment on my own posts (where I haven't tagged anyone)	●	○	○	○	○	○
I comment on my own posts (where I've tagged friends)	●	○	○	○	○	○
I comment on posts where I am tagged	●	○	○	○	○	○
I comment on my connections' posts	●	○	○	○	○	○
I comment on someone else's post in which a friend is tagged	●	○	○	○	○	○
I comment on a post made by a non-connection	●	○	○	○	○	○

LinkedIn

	Not used	No control	Low control	Medium control	High control	Full control
I post material (without tagging anyone)	●	○	○	○	○	○
I react to my own posts (where I haven't tagged anyone)	●	○	○	○	○	○
I react to my connections' posts	●	○	○	○	○	○
I react to a post made by a non-connection	●	○	○	○	○	○
I comment on my own posts (where I haven't tagged anyone)	●	○	○	○	○	○

	Not used	No control	Low control	Medium control	High control	Full control
I comment on my connections' posts	◉	○	○	○	○	○
I comment on a post made by a non-connection	◉	○	○	○	○	○

Part 16: Relationship types

What different types of social relationships do you have with the connections on each of your social networks? Select all that apply.

Use the following definitions. These may be different from your own subjective definitions.

- Family: Anyone with whom you have a familial or family-like relationship in the offline world, regardless of social media engagement.
- Close friends: People with whom you've interacted frequently on any network in the recent past - a daily average of more than 5 private messages, comments, and/or post tags within the past month.
- Friends: People with whom you've interacted at least 3 times on any network, in some capacity, in the past month.
- Colleagues: People who are not in the above categories, and with whom you have some professional connection.
- Strangers: People with whom you haven't interacted at all on any network within the last 12 months, and who are not colleagues.
- Acquaintances: People with whom you've interacted at least once on any network within the last 12 months, and are not in any of the above categories.

	Facebook	Twitter	Instagram	LinkedIn
Family	☐	☐	☐	☐
Close friends	☐	☐	☐	☐
Friends	☐	☐	☐	☐
Colleagues	☐	☐	☐	☐
Acquaintances	☐	☐	☐	☐
Strangers	☐	☐	☐	☐

Part 17: Relationship count estimates

Estimate the percentage of connections you have on each social network that fit in each category. The total percentage should add to 100%.

Use the following definitions. These may be different from your own subjective definitions.

- Family: Anyone with whom you have a familial or family-like relationship in the offline world, regardless of social media engagement.
- Close friends: People with whom you've interacted frequently on any network in the recent past - a daily average of more than 5 private messages, comments, and/or post tags within the past month.
- Friends: People with whom you've interacted at least 3 times, on any network, in some capacity, in the past month.
- Colleagues: People who are not in the above categories, and with whom you have some professional connection.
- Strangers: People with whom you haven't interacted at all on social media within the last 12 months, and who are not colleagues.
- Acquaintances: People with whom you've interacted at least once on any network within the last 12 months, and are not in any of the above categories.

Use the "Other Relation" fields to describe any important social relationship types you have that are not defined above.

	Facebook	Twitter	Instagram	LinkedIn
Family	0	0	0	0
Close friends	0	0	0	0
Friends	0	0	0	0
Colleagues	0	0	0	0
Acquaintances	0	0	0	0
Strangers	0	0	0	0
Other Relation 1	0	0	0	0
Describe this relation				
Other Relation 2	0	0	0	0
Describe this relation				
Other Relation 3	0	0	0	0
Describe this relation				

Part 18: Distinctive friends

Provide the names (and usernames, where requested and known) of three people who fit the following criteria most closely:

- **Facebook**: People you've commmunicated / interacted with most frequently.
- **Twitter**: People whose posts you've liked / retweeted most frequently.
- **Instagram**: People whom you have tagged most frequently on your own feed.
- **LinkedIn**: People who have endorsed you for the broadest range of skills.

		Facebook	Twitter	Instagram	LinkedIn
1	Name				
	Username				
2	Name				
	Username				
3	Name				
	Username				

Part 19: Network estimation

For each social network combination listed below, estimate the number of connections you have in common on *exactly* that set of networks. (For instance, if the question asks for Twitter and Instagram, estimate the number of people who are connections on both Twitter and Instagram but *not* on either Facebook or LinkedIn.)

For this question, feel free to reference your social media profiles.

Network combination	Estimated shared connections
	0
	0
	0
	0
	0
	0
	0

Network combination	Estimated shared connections
	0
	0
	0
	0
	0
	0
	0
	0

Submit

Appendix C

Supplemental data tables

This section contains supplemental data tables, including per-participant survey responses (and corresponding computed responses) as used in the analysis in Chapter 5.

	Photos (s)	Photos (c)	Videos (s)	Videos (c)	Text (s)	Text (c)
P1	50	17	10	1	40	81
P2	35	10	5	0	60	89
P3	50	52	10	3	40	43
P4	30	29	10	7	60	62
P5	50	47	10	0	40	51
P6	55	14	15	1	30	84
P7	20	13	5	1	75	84
P9	60	29	0	0	40	70
P10	20	17	10	0	70	82
P11	30	16	40	5	30	78
P12	5	2	0	2	95	94

Table C.1: Part 10: Timeline content proportions (survey and computed, percentage points, to the nearest point) per participant, Facebook.

	Photos (s)	Photos (c)	Videos (s)	Videos (c)	Text (s)	Text (c)
P1	5	2	5	0	90	97
P2	10	19	0	0	90	80
P4	20	11	0	0	80	87
P5	60	4	0	0	40	95
P6	2.5	1	2.5	0	95	98
P7	20	17	5	2	75	80
P12	1	0	0	0	99	99

Table C.2: Part 10: Timeline content proportions (survey and computed, percentage points, to the nearest point) per participant, Twitter.

	Photos (s)	Photos (c)	Videos (s)	Videos (c)	Text (s)	Text (c)
P4	75	83	25	16	0	0
P5	100	100	0	0	0	0
P6	95	96	5	3	0	0
P7	99	100	1	0	0	0
P8	96	98	4	1	0	0
P10	90	100	0	0	10	0
P11	95	98	5	1	0	0

Table C.3: Part 10: Timeline content proportions (survey and computed, percentage points, to the nearest point) per participant, Instagram.

	Post	React own	Comment own	React other	Comment other	Share other
P1	25	5	10	40	20	0
P2	10	0	5	40	30	5
P3	25	5	10	35	20	5
P4	20	5	10	40	20	5
P5	10	0	0	60	30	0
P6	10	0	80	50	50	0
P7	8	0	2	10	10	70
P9	5	0	3	85	5	2
P10	5	5	5	60	20	5
P11	20	0	5	20	10	45
P12	5	0	0	80	10	5

Table C.4: Part 13: Activity frequency percentages (survey, percentage points) per participant, Facebook.

	Post	React own	Comment own	React other	Comment other	Share other
P1	5	0	10	53	27	2
P2	8	0	9	53	27	0
P3	0	0	3	72	22	0
P4	6	0	9	47	36	2
P5	11	0	6	58	23	0
P6	17	0	7	42	30	0
P7	18	0	4	45	26	3
P9	1	0	3	47	47	0
P10	7	0	7	57	25	1
P11	10	0	13	44	30	1
P12	9	0	0	67	23	0

Table C.5: Part 13: Activity frequency percentages (computed, percentage points, to the nearest point) per participant, Facebook.

	Post	React own	Comment own	React other	Comment other	Share other
P1	80	5	5	5	5	0
P2	10	0	5	50	10	25
P4	40	5	5	15	5	30
P5	95	0	0	0	0	5
P6	60	0	0	0	0	0
P7	8	0	0	20	2	70
P12	80	0	0	10	10	0

Table C.6: Part 13: Activity frequency percentages (survey, percentage points) per participant, Twitter.

	Post	React own	Comment own	React other	Comment other	Share other
P1	84	0	0	1	6	7
P2	6	0	7	39	7	39
P4	46	0	0	12	19	21
P5	84	0	0	2	9	2
P6	87	0	0	0	7	3
P7	22	0	0	31	8	37
P12	90	0	0	0	9	0

Table C.7: Part 13: Activity frequency percentages (computed, percentage points, to the nearest point) per participant, Twitter.

	P1	P2	P3	P4	P5	P6	P7	P9	P10	P11	P12
Posts	2	1	1	5	4	3	1	2	2	1	1
Posts tagging friends	3	2	2	4	2	3	1	2	2	2	1
Tags from friends	3	4	3	5	5	3	2	4	3	4	1
React own posts	2		1	1	3	3	1		1	1	3
React own posts with tags	3		1	1	3	3	1		2	1	2
React other posts with self tagged	4	2	2	4	5	3	1	2	1	2	2
React friends' posts	4	2	2	4	2	3	1	1		4	2
React other posts with friends tagged	4	3	3	3	1	3	1	1	1	4	2
React non-connection posts	5	4	4	4	1	3	1	4	1	4	1
Comment own posts	2	1	1	1	3	3	1	2	3	1	3
Comment own posts with tags	3	2	2	3	4	3	1	2	2	2	2
Comment other posts with self tagged	4	3	2	5	3	3	1	3	1	3	2
Comment friends' posts	4		2	3	4	3	1	2		4	2
Comment other posts with friends tagged	4	3	3	4	1	3	1	2	2	4	2
Comment non-connection posts	5	4	5	5	1	3	1	4		4	1

Table C.8: Part 14: Visibility concerns (5-point Likert scale, 5 = most concern), Facebook.

	P1	P2	P4	P5	P6	P7	P12
Posts	1	1	1	1	2	1	1
Posts tagging friends	1	1	1	1	2	1	1
Tags from friends	1	3	1	1	2	1	1
React own posts	1	1	1	1	2	1	1
React own posts with tags	1	1	1	1	2	1	1
React other posts with self tagged	1	3	1	1	2	1	1
React friends' posts	1	1	1	1	2	1	1
React other posts with friends tagged	1	1	1	1	2	1	1
React non-connection posts	1	1	3	1	2	1	1
Comment own posts	1	1	1	1	2	1	1
Comment own posts with tags	1	1	1	1	2	1	1
Comment other posts with self tagged	1	3	1	1	2	1	1
Comment friends' posts	1	1	1	1	2	1	1
Comment other posts with friends tagged	1	1	1	1	2	1	1
Comment non-connection posts	1	1	1	1	2	1	1

Table C.9: Part 14: Visibility concerns (5-point Likert scale, 5 = most concern), Twitter.

	P4	P5	P6	P7	P8	P10	P11
Posts	2	4	2	1	2	1	1
Posts tagging friends	2		2	1	2		2
Tags from friends	2	5	2	1	2	1	3
React own posts	1		2	1	2	1	1
React own posts with tags	1		2	1	2	1	1
React other posts with self tagged	2	2	2	1	2	1	2
React friends' posts	2		2	1	2	1	3
React other posts with friends tagged	3		2	1	2	1	2
React non-connection posts	5		2	1	2	1	3
Comment own posts	1		2	1	1	1	1
Comment own posts with tags	2	1	2	1	2		2
Comment other posts with self tagged	3	1	2	1	3	1	2
Comment friends' posts	1		2	1	2	1	3
Comment other posts with friends tagged	3	1	2	1	2	1	3
Comment non-connection posts	5	1	2	1	2	1	3

Table C.10: Part 14: Visibility concerns (5-point Likert scale, 5 = most concern), Instagram.

	P1	P2	P3	P4	P5	P6	P7	P8	P9	P10	P11	P12
Posts			2			1		2	3	1		1
React own posts						1		2				1
React friends' posts			2			1		2		1		1
React non-connection posts			2			1		3		1		1
Comment own posts						1		1				1
Comment friends' posts						1		2		1		1
Comment non-connection posts						1		3		1		1

Table C.11: Part 14: Visibility concerns (5-point Likert scale, 5 = most concern), LinkedIn.

	P1	P2	P3	P4	P5	P6	P7	P9	P10	P11	P12
Posts	4	5	4	5	4	1	1	4	4	3	3
Posts tagging friends	3	3	1	3	1	1	1	3	4	3	2
Tags from friends	3	1	3	2	1	1	1	2	2	2	1
React own posts	4	5	4	5	4	1	1	4	4	3	3
React own posts with tags	3	3	1	3	1	1	1	4	4	2	1
React other posts with self tagged	2	1	1	2	1	1	1	2	4	2	1
React friends' posts	1	2	1	2	1	1	1	2	2	1	1
React other posts with friends tagged	1	1	1	2	1	1	1	1	2	1	1
React non-connection posts	1	1	1	1	1	1	1	1	2	1	1
Comment own posts	4	5	4	5	4	1	1	4	3	3	3
Comment own posts with tags	3	3	1	3	1	1	1	4	4	2	1
Comment other posts with self tagged	3	1	1	3	1	1	1	2	2	2	1
Comment friends' posts	2	2	1	2	1	1	1	2	2	1	1
Comment other posts with friends tagged	2	1	1	2	1	1	1	1	3	1	1
Comment non-connection posts	1	1	1	1	1	1	1	1	1	1	1

Table C.12: Part 15: Visibility control (5-point Likert scale, 5 = most control), Facebook.

	P1	P2	P4	P5	P6	P7	P12
Posts	2	5	2	4	1	1	1
Posts tagging friends	2	3	2	1	1	1	1
Tags from friends	1	1	1	1	1	1	1
React own posts	2	5	2	1	1	1	1
React own posts with tags	2	3	2	1	1	1	1
React other posts with self tagged	1	2	1	1	1	1	1
React friends' posts	1	2	2	1	1	1	1
React other posts with friends tagged	1	2	1	1	1	1	1
React non-connection posts	1	2	1	1	1	1	1
Comment own posts	2	5	2	1	1	1	1
Comment own posts with tags	2	3	2	1	1	1	1
Comment other posts with self tagged	1	1	1	1	1	1	1
Comment friends' posts	1	2	1	1	1	1	1
Comment other posts with friends tagged	1	2	1	1	1	1	1
Comment non-connection posts	1	2	1	1	1	1	1

Table C.13: Part 15: Visibility control (5-point Likert scale, 5 = most control), Twitter.

	P4	P5	P6	P7	P8	P10	P11
Posts	3	4	3	1	3	4	4
Posts tagging friends	2	1	3	1	2		3
Tags from friends	2		3	1	3	1	2
React own posts	3		3	1	4	4	4
React own posts with tags	3		3	1	2	2	3
React other posts with self tagged	2		3	1	2	3	2
React friends' posts	2		3	1	2	3	2
React other posts with friends tagged	2		3	1	2	3	2
React non-connection posts	1		3	1	1	3	1
Comment own posts	3		3	1	2	3	4
Comment own posts with tags	2		3	1	1	3	3
Comment other posts with self tagged	1		3	1	1	3	2
Comment friends' posts	1		3	1	1		2
Comment other posts with friends tagged	1		3	1	1		2
Comment non-connection posts	1		3	1	1		1

Table C.14: Part 15: Visibility control (5-point Likert scale, 5 = most control), Instagram.

	P1	P2	P3	P4	P5	P6	P7	P8	P9	P10	P11	P12
Posts			3			1		3	1			2
React own posts						1		3				2
React friends' posts			1			1		3				1
React non-connection posts			1			1		2				1
Comment own posts						1		3				2
Comment friends' posts						1		2				1
Comment non-connection posts						1		2				1

Table C.15: Part 15: Visibility control (5-point Likert scale, 5 = most control), LinkedIn.

	Family	Friends	Close friends	Colleagues	Acquaintances	Strangers
P1	1	92	1	1	5	0
P2	0	40	40	0	20	0
P3	1	10	1	20	9	50
P4	10	50	20	10	10	0
P5	25	30	20	10	10	5
P6	16	1	1	2	50	30
P7	5	30	20	2	38	5
P9	10	20	1	3	15	51
P10	5	60	20	5	5	5
P11	2	23	10	3	60	2
P12	3	0	0	7	15	75

Table C.16: Part 17: Connection proportion estimates by interaction level (survey, percentage points), Facebook.

	Family	Friends	Close friends	Colleagues	Acquaintances	Strangers
P1	0	0	0	9	35	54
P2	0	3	0	15	17	62
P3	0	3	0	20	19	56
P4	0	6	0	26	47	18
P5	0	1	0	15	33	50
P6	0	0	0	19	3	76
P7	0	0	0	15	4	78
P9	0	0	0	8	18	71
P10	0	0	0	7	31	59
P11	0	1	0	2	25	69
P12	3	0	0	16	13	66

Table C.17: Part 17: Connection proportion estimates by interaction level (computed, percentage points, to the nearest point), Facebook.

	Family	Friends	Close friends	Colleagues	Acquaintances	Strangers
P1	0	79	1	0	10	10
P2	0	30	0	0	70	0
P4	2	15	5	30	33	15
P5	0	0	0	25	25	50
P6	1	1	1	8	40	40
P7	0	25	10	5	20	40
P12	0	0	0	30	0	70

Table C.18: Part 17: Connection proportion estimates by interaction level (survey, percentage points), Twitter.

	Family	Friends	Close friends	Colleagues	Acquaintances	Strangers
P1	0	0	0	2	6	90
P2	0	0	0	3	13	81
P4	0	1	0	8	16	73
P5	0	0	0	10	1	87
P6	0	0	0	14	0	85
P7	0	0	0	7	12	79
P12	0	0	0	17	0	81

Table C.19: Part 17: Connection proportion estimates by interaction level (computed, percentage points, to the nearest point), Twitter.

	Family	Friends	Close friends	Colleagues	Acquaintances	Strangers
P4	10	40	40	0	10	0
P5	0	0	0	0	0	0
P6	3	1	2	1	83	10
P7	5	45	45	0	5	0
P8	3	40	0	10	20	27
P10	0	30	5	10	40	15
P11	0	25	20	5	42	2

Table C.20: Part 17: Connection proportion estimates by interaction level (survey, percentage points), Instagram.

	Family	Friends	Close friends	Colleagues	Acquaintances	Strangers
P4	0	3	0	14	21	60
P5	1	2	0	25	15	55
P6	1	0	0	19	1	77
P7	2	0	0	13	8	76
P8	0	0	0	25	0	74
P10	0	0	0	7	17	73
P11	0	0	0	1	15	81

Table C.21: Part 17: Connection proportion estimates by interaction level (computed, percentage points, to the nearest point), Instagram.

	Family	Friends	Close friends	Colleagues	Acquaintances	Strangers
P1	0	50	0	40	5	5
P2	0	0	0	0	0	0
P3	0	10	10	30	30	20
P4	15	5	20	30	30	0
P5	5	5	10	30	50	0
P6	1	1	1	90	4	4
P7	5	10	5	30	20	30
P8	1	45	0	25	9	20
P9	0	0	0	0	0	0
P10	10	10	10	50	10	10
P11	0	50	10	0	40	0
P12	0	0	0	100	0	0

Table C.22: Part 17: Connection proportion estimates by interaction level (survey, percentage points), LinkedIn.

	Family	Friends	Close friends	Colleagues
P1	0	0	0	100
P2	0	2	0	96
P3	0	3	0	96
P4	0	2	0	96
P5	0	0	0	99
P6	0	0	0	99
P7	0	0	0	99
P8	0	0	0	100
P9	1	0	0	98
P10	0	0	0	100
P11	0	4	4	90
P12	0	0	0	100

Table C.23: Part 17: Connection proportion estimates by interaction level (computed, percentage points, to the nearest point), LinkedIn.

	P1	P2	P3	P4	P5	P6	P7	P8	P9	P10	P11	P12
TIFL	0	0	0	50	0	420	2	0	0	0	0	0
TIF	0	0	0	300	12	600	2	0	0	0	0	0
IFL	0	0	0	200	3	578	2	0	0	150	35	0
TFL	20	10	0	100	15	650	15	0	0	0	0	10
TIL	0	0	0	0	15	450	2	0	0	0	0	0
TI	0	0	0	0	8	600	5	0	0	0	0	0
FL	20	20	150	200	40	1200	10	0	56	150	40	10
IF	6	10	0	50	70	700	50	0	0	0	0	20
IL	4	2	0	50	60	675	15	0	0	20	0	10
IF	0	0	0	200	50	750	10	0	0	136	250	0
IL	0	0	0	0	10	500	2	16	0	50	25	0
I	0	0	0	30	76	175	75	50	0	20	250	0
T	116	350	0	150	397	360	300	0	0	0	0	57
F	372	1000	813	200	1000	300	402	0	523	700	800	132
L	147	350	192	80	1000	1100	308	870	6	50	40	60

Table C.24: Part 19: Connection count estimates by network grouping (survey).

	P1	P2	P3	P4	P5	P6	P7	P8	P9	P10	P11	P12
TIFL	0	0	0	22	4	29	2	0	0	0	0	0
TIF	0	0	0	23	2	2	4	0	0	0	0	0
IFL	0	0	0	34	16	84	11	0	0	12	7	0
TFL	2	7	0	22	6	37	8	0	0	0	0	8
TIL	0	0	0	5	0	4	0	0	0	0	0	0
TI	0	0	0	37	1	10	3	0	0	0	0	0
FL	34	158	208	70	106	233	42	0	46	63	16	18
IF	2	6	0	10	7	7	7	0	0	0	0	5
IL	2	15	0	18	58	34	20	0	0	0	0	10
IF	0	0	0	80	15	175	32	0	0	68	91	0
IL	0	0	0	7	4	20	0	45	0	4	1	0
I	0	0	0	220	43	357	44	133	0	124	264	0
T	145	431	0	523	580	594	339	0	0	0	0	80
F	327	808	750	254	688	1319	288	0	476	851	636	125
L	109	151	133	206	888	1517	227	874	16	208	20	106

Table C.25: Part 19: Connection counts by network grouping (computed).

Bibliography

[1] Michael S. Bernstein, Eytan Bakshy, Moira Burke, and Brian Karrer. Quantifying the invisible audience in social networks. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '13, pages 21–30, New York, NY, USA, 2013. ACM. ISBN 978-1-4503-1899-0. doi: 10.1145/2470654.2470658. URL `http://doi.acm.org/10.1145/2470654.2470658`.

[2] Hichang Cho and Anna Filippova. Networked privacy management in facebook: A mixed-methods and multinational study. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing*, CSCW '16, pages 503–514, New York, NY, USA, 2016. ACM. ISBN 978-1-4503-3592-8. doi: 10.1145/2818048.2819996. URL `http://doi.acm.org/10.1145/2818048.2819996`.

[3] R. I. M. Dunbar. Social cognition on the internet: testing constraints on social network size. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 367(1599):2192–2201, 2012. doi: 10.1098/rstb.2012.0121.

[4] R. I. M. Dunbar. Do online social media cut through the constraints that limit the size of offline social networks? *Royal Society Open Science*, 3(1), 2016. doi: 10.1098/rsos.150292. URL `http://rsos.royalsocietypublishing.org/content/3/1/150292`.

[5] Veikko Eranti and Markku Lonkila. The social significance of the facebook like button. *First Monday*, 20(6), 2015. ISSN 13960466. doi: 10.5210/fm.v20i6.5505. URL `http://firstmonday.org/ojs/index.php/fm/article/view/5505`.

[6] Lujun Fang and Kristen LeFevre. Privacy wizards for social networking sites. In *Proceedings of the 19th International Conference on World Wide Web*, WWW '10, pages 351–360, New York, NY, USA, 2010. ACM. ISBN 978-1-60558-799-8. doi: 10.1145/1772690.1772727. URL http://doi.acm.org/10.1145/1772690.1772727.

[7] GibsonSec. Snapchat security advisory, 2013. URL http://gibsonsec.org/snapchat/.

[8] Ralph Gross and Alessandro s 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. ISBN 1-59593-228-3. doi: 10.1145/1102199.1102214. URL http://doi.acm.org/10.1145/1102199.1102214.

[9] Alex Heath. Snap misses across the board for q2 earnings, stock gets whacked, 2017. URL http://www.businessinsider.com/snap-q2-earnings-results-2017-8.

[10] Maritza Johnson, Serge Egelman, and Steven M. Bellovin. Facebook and privacy: It's complicated. In *Proceedings of the Eighth Symposium on Usable Privacy and Security*, SOUPS '12, pages 9:1–9:15, New York, NY, USA, 2012. ACM. ISBN 978-1-4503-1532-6. doi: 10.1145/2335356.2335369. URL http://doi.acm.org/10.1145/2335356.2335369.

[11] Michal Kosinski, David Stillwell, and Thore Graepel. Private traits and attributes are predictable from digital records of human behavior. *Proceedings of the National Academy of Sciences*, 110(15):5802–5805, 2013. doi: 10.1073/pnas.1218772110. URL http://www.pnas.org/content/110/15/5802.abstract.

[12] Mounia Lalmas, Heather O'Brien, and Elad Yom-Tov. Measuring user engagement. *Synthesis Lectures on Information Concepts, Retrieval, and Services*, 6(4):1–132, 2014.

doi: 10.2200/S00605ED1V01Y201410ICR038. URL http://dx.doi.org/10.2200/S00605ED1V01Y201410ICR038.

[13] Steven Levy. Snapchatś non-vanishing message: you can trust us, 2015. URL https://www.wired.com/2015/04/snapchats-non-vanishing-message-you-can-trust-us/.

[14] Eden Litt and Eszter Hargittai. The imagined audience on social network sites. *Social Media + Society*, 2(1):2056305116633482, 2016. doi: 10.1177/2056305116633482. URL http://dx.doi.org/10.1177/2056305116633482.

[15] Leqi Liu, Daniel Preotiuc-Pietro, Zahra Riahi Samani, Mohsen Ebrahimi Moghaddam, and Lyle H Ungar. Analyzing personality through social media profile picture choice. In *ICWSM*, pages 211–220, 2016.

[16] David Lumb. Twitter fires a vr manager after his past comes to light, 2016. URL https://www.engadget.com/2016/10/19/twitter-fires-its-head-of-vr-after-24-hours/.

[17] Alice E. Marwick and danah boyd. I tweet honestly, i tweet passionately: Twitter users, context collapse, and the imagined audience. *New Media & Society*, 13(1): 114–133, 2011. doi: 10.1177/1461444810365313. URL http://dx.doi.org/10.1177/1461444810365313.

[18] Alessandra Mazzia, Kristen LeFevre, and Eytan Adar. The pviz comprehension tool for social network privacy settings. In *Proceedings of the Eighth Symposium on Usable Privacy and Security*, SOUPS '12, pages 13:1–13:12, New York, NY, USA, 2012. ACM. ISBN 978-1-4503-1532-6. doi: 10.1145/2335356.2335374. URL http://doi.acm.org/10.1145/2335356.2335374.

[19] Thomas Paul, Martin Stopczynski, Daniel Puscher, Melanie Volkamer, and Thorsten Strufe. *C4PS - Helping Facebookers Manage Their Privacy Settings*, pages 188–201.

Springer Berlin Heidelberg, Berlin, Heidelberg, 2012. ISBN 978-3-642-35386-4. doi: 10.1007/978-3-642-35386-4_15. URL https://doi.org/10.1007/978-3-642-35386-4_15.

[20] Pew Research Center. Social media fact sheet, 2017. URL http://www.pewinternet.org/fact-sheet/social-media/.

[21] Matthew Pittman and Brandon Reich. Social media and loneliness: Why an instagram picture may be worth more than a thousand twitter words. *Computers in Human Behavior*, 62:155 – 167, 2016. ISSN 0747-5632. doi: http://dx.doi.org/10.1016/j.chb.2016.03.084. URL http://www.sciencedirect.com/science/article/pii/S0747563216302552.

[22] Sarita Schoenebeck, Nicole B. Ellison, Lindsay Blackwell, Joseph B. Bayer, and Emily B. Falk. Playful backstalking and serious impression management: How young adults reflect on their past identities on facebook. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing*, CSCW '16, pages 1475–1487, New York, NY, USA, 2016. ACM. ISBN 978-1-4503-3592-8. doi: 10.1145/2818048.2819923. URL http://doi.acm.org/10.1145/2818048.2819923.

[23] Fuming Shih, Ilaria Liccardi, and Daniel Weitzner. Privacy tipping points in smartphones privacy preferences. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, CHI '15, pages 807–816, New York, NY, USA, 2015. ACM. ISBN 978-1-4503-3145-6. doi: 10.1145/2702123.2702404. URL http://doi.acm.org/10.1145/2702123.2702404.

[24] Manya Sleeper, Justin Cranshaw, Patrick Gage Kelley, Blase Ur, Alessandro Acquisti, Lorrie Faith Cranor, and Norman Sadeh. "i read my twitter the next morning and was astonished": A conversational perspective on twitter regrets. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '13, pages 3277–

3286, New York, NY, USA, 2013. ACM. ISBN 978-1-4503-1899-0. doi: 10.1145/2470654.2466448. URL http://doi.acm.org/10.1145/2470654.2466448.

[25] Jessica Vitak and Jinyoung Kim. "you can't block people offline": Examining how facebook's affordances shape the disclosure process. In *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing*, CSCW '14, pages 461–474, New York, NY, USA, 2014. ACM. ISBN 978-1-4503-2540-0. doi: 10.1145/2531602.2531672. URL http://doi.acm.org/10.1145/2531602.2531672.

[26] Yang Wang, Gregory Norcie, Saranga Komanduri, Alessandro Acquisti, Pedro Giovanni Leon, and Lorrie Faith Cranor. "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*, SOUPS '11, pages 10:1–10:16, New York, NY, USA, 2011. ACM. ISBN 978-1-4503-0911-0. doi: 10.1145/2078827.2078841. URL http://doi.acm.org/10.1145/2078827.2078841.

[27] Yang Wang, Liang Gou, Anbang Xu, Michelle X. Zhou, Huahai Yang, and Hernan Badenes. Veilme: An interactive visualization tool for privacy configuration of using personality traits. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, CHI '15, pages 817–826, New York, NY, USA, 2015. ACM. ISBN 978-1-4503-3145-6. doi: 10.1145/2702123.2702293. URL http://doi.acm.org/10.1145/2702123.2702293.

[28] WhatsApp. Connecting one billion users every day, 2017. URL https://blog.whatsapp.com/10000631/Connecting-One-Billion-Users-Every-Day.

[29] Xuan Zhao, Cliff Lampe, and Nicole B. Ellison. The social media ecology: User perceptions, strategies and challenges. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, CHI '16, pages 89–100, New York, NY,

USA, 2016. ACM. ISBN 978-1-4503-3362-7. doi: 10.1145/2858036.2858333. URL http://doi.acm.org/10.1145/2858036.2858333.