

Photo Privacy Conflicts in Social Media: A Large-scale Empirical Study

Jose M. Such
Department of Informatics
King's College London
jose.such@kcl.ac.uk

Joel Porter
Security Lancaster
Lancaster University
j.porter2@lancaster.ac.uk

Sören Preibusch
Google
Mountain View, CA
preibusch@google.com

Adam Joinson
School of Management
University of Bath
A.Joinson@bath.ac.uk

ABSTRACT

Items in social media such as photos may be co-owned by multiple users, i.e., the sharing decisions of the ones who upload them have the potential to harm the privacy of the others. Previous works uncovered coping strategies by co-owners to manage their privacy, but mainly focused on general practices and experiences. We establish an empirical base for the prevalence, context and severity of privacy conflicts over co-owned photos. To this aim, a parallel survey of pre-screened 496 uploaders and 537 co-owners collected occurrences and type of conflicts over co-owned photos, and any actions taken towards resolving them. We uncover nuances and complexities not known before, including co-ownership types, and divergences in the assessment of photo audiences. We also find that an all-or-nothing approach seems to dominate conflict resolution, even when parties actually interact and talk about the conflict. Finally, we derive key insights for designing systems to mitigate these divergences and facilitate consensus.

Author Keywords

Privacy; photo sharing; online social networks; co-ownership; conflicts; social media

ACM Classification Keywords

K.4.1. Computers and Society: Privacy; H.5.m. Information Interfaces and Presentation (e.g. HCI): Miscellaneous

INTRODUCTION

The interpersonal and multiparty nature of privacy was already identified in foundational privacy theories such as those of Altman [2], Petronio [40], and Nissenbaum [37]. This very nature was also acknowledged as critical to successfully unpack privacy in computer-mediated communications [38, 52]. Social media are computer-mediated communication infrastructures where *multiparty* privacy manifests, as many of the hundreds of billions of items that are uploaded are *co-owned* by multiple users [58, 23]. One particular manifestation is what we call *multiparty privacy conflicts* (MPCs), which arise

when the privacy preferences of the uploader and co-owners of an item do not align. In this paper, we focus on *photo MPCs*, i.e., MPCs caused by photos shared over social media. Take a very simple but illustrative example: Alice takes a photo of her and her friends Bob and Charlie and then shares it with others through social media. What if Alice (the uploader) sets the photo visible to all in the social media site but Bob and Charlie (the co-owners) would have preferred to just share among common friends?

Handling MPCs is an important part of users' interpersonal boundary regulation processes, and there are a number of strategies users are known to utilise [31, 56, 13, 24]. These strategies can be *preventive* (they aim to avoid MPCs) or *corrective* (they aim to resolve MPCs when they happen); *individual* (each affected user can do something unilaterally) or *collaborative* (all users potentially affected by a co-owned item work together); and *online* (users make use of social media features) or *offline* (users apply strategies out of social media). An example of a preventive, collaborative and offline strategy is that users are known to sometimes negotiate with whom a photo will be shared before uploading it to social media [56], e.g., friends deciding the photos they take during their holiday trip to only be shared online with common friends. An example of a corrective, individual, and online strategy would be when a user untags herself from a photo uploaded to social media to try to prevent her friends from seeing it [31, 13]. However, these strategies may not always work or provide the desired outputs. For instance, untagging oneself from a photo does not completely remove the possibility of the photo being accessed anyway by the undesired audience, and having to negotiate offline for each and every co-owned item would be unfeasible due to the large number of co-owned items and potential co-owners to deal with.

This has inspired researchers to tackle the problem of how to design support for multi-privacy privacy, particularly by looking to designing interfaces and computational methods for users to manage MPCs in effective and efficient ways, such as [4, 12, 21, 49, 47] as discussed below. However, there is a lack of an empirical base that can help understand when and how MPCs happen, and what is the best way of preventing/resolving them. Even previous works that studied the strategies users follow to cope with MPCs mainly focused on general practices / experiences [31, 56, 13, 24].

In this paper, we present the first large-scale study of actual photo MPCs, from identification of the MPC, to communi-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CHI 2017, May 6–11, 2017, Denver, CO, USA.

Copyright © 2017 ACM ISBN 978-1-4503-4655-9/17/05 ...\$15.00.

<http://dx.doi.org/10.1145/3025453.3025668>

cation and/or resolution. Crucially, we look at photo MPCs through the lens of the uploaders of co-owned items and co-owners, as we are interested to see if there may be any differences between these roles in terms of the information they have about MPCs as well as their point of view and perceptions.

Contribution

The main contribution of this work is the large-scale, empirical study of photo MPCs, including both uploaders and co-owners. We also provide exemplary explorations of the resulting dataset focusing on the research questions below.

Research Questions

Our main research questions are:

1. What are the characteristics of photo MPCs?
2. What is the frequency of communication and resolution strategies, and their effectiveness for resolving MPCs?
3. Are there significant differences between uploaders and co-owners in terms of MPC characteristics and communication/resolution strategies?

RELATED WORK

Previous research on multiparty privacy in social media can be categorised into two main streams.

The first stream uncovered, studied and categorised some of the strategies that users are known to use to deal with multiparty privacy in social media [31, 56, 13]. Lampinen et al. [31] classified the strategies users employ into preventive or corrective, and individual or collaborative. Boyd [6] showed a particular instance of collaborative preventive strategy, discovering and studying how teens share inside jokes and cloak their messages so that any other social media users would not be able to know the actual information teens are sharing with each other. Wisniewski et al. [56] made a further crucial distinction between the strategies which are in some way facilitated by technology, which we refer to as *online* (e.g., untagging), and strategies that are applied outside the confines of social media infrastructures, which we refer to as *offline* (e.g., negotiating the audience before posting). They argued that people are more likely to apply offline mechanisms rather than technological mechanisms, mainly due to the lack of more appropriate built-in technological mechanisms in social media infrastructures. Later on, Cho and Filippova [13] showed that reported strategy use seems to be correlated to perceived efficacy by the individual and the particular group of co-owners in managing privacy. Also, it has been suggested that users enact collaborative strategies based on group perceptions of collective privacy risk, information disclosure, and propensity to value privacy [24]; and that privacy concerns seem to be negatively associated with engaging with friends through tagging activities [57]. Finally, Wisniewski et al. [55] studied the support mainstream social media interfaces provided to users, concluding that these interfaces are ineffective, difficult to use and not easy to be aware of, which leaves users no other option than to apply offline coping strategies. While the studies presented above are

foundational in nature and significantly contributed to the understanding of the problem, they mainly focused on reported general practices. Our study is aimed to complement these studies by providing a large-scale and empirical study of particular cases of MPCs, from the perspective of both uploaders and co-owners.

The second stream focused on designing interfaces and computational mechanisms for supporting multiparty privacy management. Some of these works provide some support for users to resolve MPCs [4, 46]. For example, Besmer and Lipford [4] studied photo tagging and untagging behaviours on Facebook, and proposed a system whereby a user tagged in a photo could send privacy suggestions to the poster. In general, uploaders were very accommodating of the concerns of co-owners, and there was a perception that images posted online were co-owned between those represented within the photograph. Other approaches automatically aggregate individual privacy preferences to come up with a collective sharing policy [12, 21, 49, 47], but the methods used to do so are rarely based on evidence of what users would consider as the best sharing policy for a given situation. There have also been works that propose game-theoretic negotiation mechanisms, in which users [22] or automated software agents [48] negotiate a solution to multiparty privacy conflicts following an established protocol. Both the protocol and negotiation strategies are analysed using game-theoretic solution concepts such as the Nash equilibrium. However, such proposals may not work well in practice since the underlying assumption that users behave in a perfectly rational way may not capture the social idiosyncrasies that users seem to consider when sharing and managing privacy in social media [31, 56]. Finally, Illia et al. [23] presented a mechanism to enforce fine-grained access control in photos by blurring the faces of the users depicted in the photo based on each users' access control list. While multiparty privacy management methods and tools are starting to provide alternatives to the coping techniques adopted by users who are faced with existing system constraints, as a community we still lack more evidence upon which to base the design of these systems. The overarching goal of the present research is to provide such an evidence base.

METHOD

We studied specific multiparty privacy conflicts over co-owned photos in social media using a critical incident method [18] as detailed below, inspired by previous works that had successfully used a similar methodology to collect and study specific experiences of users updating software [51], and security incidents [42].

Survey Instruments

Two questionnaires were designed to be almost identical with small variations/wordings to adapt them to those that had either shared a photograph that led to a MPC (uploaders) or to co-owners who had been involved in MPCs. The two survey questionnaires contained around thirty questions each (see appendix). Then, the guidelines for the critical incident technique were followed to ask for the most recent observation [18] (i.e., the last time participants experienced a MPC),

in order to prevent biasing the study to just the more dramatic or vivid incidents, or some other selected group, such as those which fit the participant's stereotypes. Specifically, a set of closed- and open-ended questions asked, depending on the role, about: a) the last photo they shared that caused an MPC (uploaders); or b) the last photo of themselves shared by someone else that caused a MPC (co-owners). These questions were semi-structured following the known stages of interpersonal conflicts [17]: conflict identification, communication and resolution. In particular, we asked about the context of the photo [37], to understand what was in it, who was in it, relationships between people, the audience with which the photo was shared, the self-assessed severity of the MPC; the nature of any communications between uploaders and co-owners about the MPC; and the approach followed to try to resolve the MPC, and whether the MPC was considered to be resolved or not. Finally, we also asked at the end about the general multiparty sharing practices taken from related literature [31, 56, 13], but also allowing participants to report specific ways in which they prevent/resolve MPCs, together with demographics, and photo taking and social media use.

After the study was approved by the IRB board at Lancaster University, Amazon Mechanical Turk was used to administer the questionnaires to participants. Two pre-tests were conducted first using a small sample of 50 turkers each time to get feedback and refine the two questionnaires. The final versions of the survey were administered to 1080 turkers. Participants were initially shown with screening questions with eligibility criteria (being user of at least one social sharing platform, having taken photos over the last 30 months, and having shared photos through sharing platforms), and an informed consent statement they needed to accept. After this, participants were randomly assigned to the uploaders or co-owners version of the questionnaire. It took an average of 12 minutes to complete the survey and participants were paid \$3.50, which means an hourly pay of \$17.82, above the minimum average wage of \$7.25 mandated by the US federal government and the maximum per-state minimum average wage of \$10.50.

Data Reliability and Coding Process

Well-known quality-assurance processes were followed when administering the survey instrument through Amazon Mechanical Turk [9, 39, 35, 20], including attention check questions, and selection of US-based participants with at least 50 tasks completed on MTurk and a success rate of at least 95%. For open-ended questions, the data analysis used open coding on the qualitative responses. This was closely followed by the labelling of data with active, specific codes [1]. In particular, a collaborative coding process was followed for this initial stage [43]. One *lead* coder developed an initial version of the code book, which then was handed to a second coder who independently and blindly (without knowing how the other coder had coded the data) coded the data again. Cohen's kappa was used to measure intercoder agreement between the codes that each coder had assigned independently and blindly for each response to each open-ended question, with values of at least 0.7 for all questions, and with some questions that were over 0.8, both values showing a high to

very high degree of agreement between the coders. After this double-blind step, a second step was conducted in which, for every disagreement found between the two coders, both of them worked collaboratively to either refine the code book if the disagreement was attributed to an inconsistency in it, or to reach an agreement on the most appropriate code for the disagreement found [33]. When open coding was finished, the second stage was selective coding, when codes that represented similar concepts were collated into overarching categories, then closely examined for inter-significance to produce a set of themes grounded on the data [8], some of which are preliminarily explored within this study according to the research questions stated above.

RESULTS

Participants

1080 participants submitted the questionnaire, 9 failed at least one of the attention checks, 32 had straight lining responses, and 6 reported a privacy violation but nothing to do with multiparty privacy. These cases were removed from the dataset, and the remaining 1033 (96%) were assessed, including 496 uploaders and 537 co-owners. Table 1 shows demographics and other descriptive statistics.

	Range	Value	%
Age	18-19	19	2%
	20-29	454	44%
	30-39	335	32%
	40-49	139	13%
	50-59	69	7%
	60-69	17	2%
Gender	Male	516	50.0%
	Female	512	49.5%
	Other/N/A	5	0.5%
	Mean	Median	Std.Dev.
Close friends	8.46	5.00	45.67
Online contacts	443.89	218.00	1013.85
Hours spent online / day	6.36	5.00	4.60
Photos taken last 30 days with 1+ people shared with others	65.36	27.00	124.04
	33.09	10.00	77.94
	16.52	5.00	43.86

Table 1. Descriptive statistics of participants.

We also asked frequencies over the last 2 months in which either co-owners were happy/unhappy photos of them were shared by others, or uploaders perceived co-owners being happy/unhappy about photos they shared. The results in Table 2 show both uploaders and co-owners were in general happier than unhappier about multiparty sharing, so participants seem not to be negatively biased against it, which is in line with previous research on the benefits users experience when sharing in social media [25, 29, 10].

# Times	Co-owner		Uploader	
	Happy	Unhappy	Happy	Unhappy
0	31%	47%	27%	61%
1	10%	32%	7%	32%
2	10%	12%	8%	4%
3	5%	5%	3%	1%
4	2%	1%	3%	0%
5+	42%	3%	52%	2%

Table 2. Number of times over the last 2 months in which: i) co-owners were happy/unhappy a photo of them was shared, and ii) uploaders perceived co-owners to be happy/unhappy of the photos their shared.

We also wanted to understand the general practises in multiparty sharing scenarios of our participants beyond the specific

	Co-owner	Uploader
General photo taking/sharing practices		
Multiple accounts to share separately	N/A	83 (17%)
Same account to share personal and work	N/A	243 (49%)
Negotiate rules of thumb for sharing	272 (51%)	288 (58%)
Restrict general privacy settings	445 (83%)	422 (85%)
Problematic photos not shared	417 (78%)	440 (89%)
Share via private message / e-mail	410 (76%)	416 (84%)
Share via private cloud storage	259 (48%)	202 (41%)
Share via Snapchat/photo-vanishing app	163 (30%)	142 (29%)
Change offline behaviour to avoid photo	286 (53%)	N/A
No photo if someone avoiding camera	N/A	412 (82%)
General practices before posting a photo		
Consider others	N/A	127 (25.6%)
Request permission	N/A	127 (26.5%)
Selective/restrictive sharing	N/A	132 (26.6%)
Tag approval	31 (5.8%)	N/A
Decide not to share (or ask not to)	123 (22.9%)	59 (12.9%)
General practices after a conflicting photo was posted		
Ignore conflicts	86 (17%)	25 (6%)
Converse with the other party	156 (30%)	44 (10%)
Remove the photo (or ask to)	231 (43%)	303 (68%)
Restrict access to the photo (or ask to)	15 (2.8%)	15 (3%)
Untag from the photo	87 (17%)	15 (3%)
Unfriend/block the other party	4 (0.7%)	4 (1%)

Table 3. General reported practices for multiparty sharing and to prevent/resolve multiparty privacy conflicts.

instance of MPC they reported. To this end, we presented participants with a list of general practises based on previous research [31, 56, 13, 24], as well as allowing them to report any other general practices they usually follow before and after posting a photo to prevent/resolve MPCs. The results from this part of the survey are shown in Table 3. Note, that although we present these results first to characterise our participants, the questions were actually shown at the end of the survey to avoid biasing what participants would report for the particular MPC — see appendix.

Characteristics of Photo MPCs

In this section we focus on answering RQ1. Also, note correlation coefficients were calculated for characteristics from closed questions to all other closed questions and they were not statistically significant unless explicitly stated otherwise below. A Bonferroni correction for the significances was used to adjust for multiple tests (a total of 61, including correlations and the other tests described later on) to minimise Type 1 errors. The resulting cut-off significance level for correlations as well as for the other tests shown in the paper was $\alpha = 0.05/61 = 0.00082$.

Prevalence

The vast majority of participants (99.3%) reported a multiparty privacy conflict. In particular, 30% of participants experienced their last conflict within 1 month of the survey date, 44% within 6 months, and 66% within 12 months. There was no significant difference between uploaders and co-owners in this regard. The results indicate conflicts are likely to occur more often around key holiday dates such as Halloween, Thanksgiving, Christmas, and New Year, which can also be attributed to an increase in photo taking/uploading during the days around these events. For instance, Figure 1 represents the 4 month period leading up to the close of the study, where we can observe clear spikes around those dates. However, we can also see spikes very close to the date of the study, which,

after analysis, were clearly not related at all to any of the holiday dates, suggesting relatively high frequency of conflicts in the day-to-day sharing as well.




Figure 1. Number of MPCs during the 4 months before administering the survey.

The majority of photos that led to conflicts were shared on Facebook (80.8%), Instagram (5.4%), Snapchat (1.4%), and Twitter (1.3%). Other photo sharing platforms such as imgur, Flickr, MySpace, and Tumblr were used but individually represented less than 1% of the study.

Content

Regarding the content of the conflicting photos, 4.3% contained no person (e.g., landscape or object), 13.6% had 1 person, 32.7% had 2 people, 18.2% had 3 people, and 31.3% had 4 or more.

Regarding the relationships between the depicted people in the photo (when more than one person), family made up 41.36%, friends 40.20%, relationships such as boyfriend/girlfriend 9.29%, acquaintances 4.06%, ex-partners 2.47%, co-workers 2.32%, and others 0.29%.

Regarding specific topics specified by respondents (60% of MPCs), they included people drinking or being drunk (23.57%), old photos from the past (10.27%), children (7.41%), people eating or food-related photos (6.23%), people playing around such as making faces (6.06%), intimate photos or photos with some degree of nudity (4.04%), people sleeping (2.86%), selfies (2.56%), photos taken off the subjects' guard (2.36%), people having just got up or with pajamas on (2.36%), people in costumes (1.52%), people celebrating (1.52%), photos taken with poor-quality (1.52%), etc.

Regarding specific events specified by respondents (54% of MPCs), it related to parties (31.52%), including birthdays (11%), new year (7.22%), Thanksgiving and 4th July celebrations (3.6%), and non date-specific parties (13.37%). Christmas (15.50%), vacations (10.09%), meals and gatherings (8.65%), weddings (3.6%), sport events (3.24%), Halloween events (3.24%), graduations (2.6%), and concerts (1.62%) also contributed to occasions that were explicitly cited as times in which conflicting photos had been taken at.

Severity

Participants were asked to identify on a Likert scale how severe the conflict was, with 5 being very severe and 1 not severe. Table 4 represents the number of conflicts and severity, showing a very similar pattern for both uploaders and co-owners of fewer cases as the severity increases.

In terms of contexts and content for which conflicts are perceived to be more severe than others (levels 4 and 5), there

Severity	Not Severe				Very Severe	Total
	1	2	3	4	5	
Uploader	184 (37.1%)	165 (33.2%)	99 (20.0%)	44 (8.9%)	4 (0.8%)	496 (100%)
Co-owner	189 (35.2%)	190 (35.4%)	107 (19.9%)	44 (8.2%)	8 (1.3%)	537 (100%)

Table 4. Respondents' perceived severity of conflicts.

was a predominance of drink-, party-, and friend-related contents, with some familiar and ex-* relationships also appearing. In terms of relationship closeness, we could find examples of varying relationship closeness at all levels of severity, e.g., “*he is my husband we are/were very close*” (P407u); “*he is my cousin*” (P502c); “*friend, from high school*” (P338u). Finally, the gender of the respondent ($r = 0.12$, $p < 0.00005$) was significantly correlated with severity.

Process of Photo MPCs

Now, we focus on RQ2, i.e., what is the frequency of communication and resolution strategies, and their effectiveness for resolving MPCs.

Communication before sharing

Regarding communication *before* posting, 57.3% of uploaders stated that they had not asked for any kind of feedback or permission from co-owners, 32.7% had not asked for feedback or permission but later regretted not doing so, 6.5% had asked everyone before sharing the photo, 2.6% asked some of the people represented, and 1.0% asked someone that was not in the photo, such as “*I asked his mother if she would mind me sharing it*” (P88u). When co-owners were asked whether they had been approached for feedback or permission before a photograph was shared, 96% stated no, and only 4% reported that they had been asked for feedback and/or permission. Thus, asking for feedback/permission before posting was very rare regardless of the group being asked.

Communication after sharing

We focus now on what happened *after* a photo was shared. In general, complaints related to individual privacy issues and confirmed previous observations in multiparty scenarios [4, 31], with three main overarching themes: self-presentation issues [45], context collapse [5], and controlling who has access and when [56, 24]. Regarding how the complaints were made, we found: in person (38.8%), texting (20.4%), online (17.6%), telephoning (9.6%), messaging through apps such as whatsapp (7.3%), email (3.8%), and other (2.3%).

For uploaders, all the reported MPCs described a situation in which co-owners did express their discontent. The majority, 74%, of complaints were made by a person depicted in the photo. Complaints reported by someone not in the photo amounted 23%, and they included people who complained about someone they felt responsible for or about something of them that was shown in the photo. One example was parents complaining about a photo of their children: “*[It was] a picture of my child holding a piece of paper ... my husband complained that i should have asked before submitting the photo*” (P50u). Another example was people complaining about their house, room, car, pet and so on: “*It was a picture out the front windshield of a friend's vehicle from*

when we went storm chasing, and the windshield is cracked. My friend felt embarrassed about the aforementioned cracks in the windshield, he thought it made him look bad” (P3u). There were also a small number of MPCs in which someone complained after being tagged but without being depicted in the photo: “*It was a photo of a crab on a beach. I [...] tagged a friend in it. It's a longstanding joke that I call her crablady and started a rumor that she has crabs*” (P277u). Finally, very few uploaders reported the photo being flagged to the social media provider (1.6%), and 3% of uploaders claimed not to have received a complaint, as they got to know about the conflict when they contacted co-owners before sharing the photo.

For co-owners, there were two main cases, either they did or did not communicate their unhappiness with the photo to uploaders. In particular, 46.8% of co-owners complained about the photo they were unhappy about while the remaining 53.2% did not. When co-owners did complain, 95% addressed directly the person who uploaded the photo after it was shared, and 5% of complaints were directed towards the provider of the website or application, though no association with severity was noted. There was a significant correlation between the severity of the conflict and whether co-owners complained or not ($r = 0.3$, $p < 0.00001$).

For the 53.2% co-owners who did not complain about a photo they were unhappy about, 33.9% stated that they did not wish to upset both the person who published the photo, nor those represented in the photo, e.g., “*too awkward, didn't want to get on people's bad side*” (P698c). Of the remaining co-owners who did not complain, 25.0% did not consider it worth complaining; 12.1% thought complaining was petty, would cause additional attention, or would be considered vain; 7.3% thought it was too late once the photo had been published; 3.2% suggested complaining would complicate other people's ability to share the photo; and 5.6% had already been approached for feedback.

Conflict Resolution

Table 5 shows all the conflict resolution strategies and whether the MPC was considered to be resolved or not broken down by the reporter, i.e., uploader or co-owner. Also, we split co-owners into co-owners who did complain and co-owners who did not complain. We can clearly see that there is a relationship between the conflict resolution strategy and whether or not the MPC was perceived to be resolved regardless of the reporter, which is significant ($\chi^2 = 667$, $p < 0.0001$).

An interesting observation from Table 5 (columns for *All*) is that there is a large majority of cases (around 75%) in which the conflict resolution was either removing the photo or doing nothing about it, which shows a very *all-or-nothing* approach for conflict resolution. However, the difference between these two main strategies is their effectiveness. In particular, removing the photo was perceived to *always* resolve the conflict in all 338 cases that were reported by both uploaders and co-owners. In contrast, doing nothing was perceived not to resolve a conflict around 70% of the time, though there are

Strategy	Co-owner (complain:no)			Co-owner (complain:yes)			Uploader			All		
	Resolved Yes	No	Total	Resolved Yes	No	Total	Resolved Yes	No	Total	Resolved Yes	No	Total
Apologised				6	1	7	29	2	31	35	3	38
Converse				6		6	24		24	30		30
Converse (flattery)					1	1	17		17	17	1	18
Agreed before sharing	3		3	1		1	4		4	8		8
Forfeit							1		1	1		1
Seek permission				10		10	11		11	21		21
Stopped sharing							1		1	1		1
Remove				129		129	209		209	338		338
Remove from timeline	5		5							5		5
Replace (another)				6		6	14		14	20		20
Replace (blur)				1		1				1		1
Replace (crop)				1		1	9		9	10		10
Photo vanished							1		1	1		1
Restrict access	3		3	1		1		8	8	4	8	12
Untag (uploader)	2		2	3		3		16	16	5	16	21
Untag (co-owner)	36	4	40	7	1	8	3		3	46	5	51
Did nothing (uploader)				18	63	81	77	62	139	95	125	220
Did nothing (co-owner)	47	174	221							47	174	221
Blocked	2		2	2		2	2		2	6		6
Unfriend (uploader)							4		4	4		4
Unfriend (co-owner)	1		1					2	2	1	2	3
Not speaking	2		2				1		1	3		3
Changed offline		1	1							1		1
Grand Total	101	179	280	191	66	257	430	66	496	722	311	1033

Table 5. Contingency table with # times each conflict resolution strategy was observed and whether or not it resolved the conflict.

slight differences here between uploaders and co-owners that we discuss more in-depth in the next section.

For the remaining approaches to conflict resolution (25% all together), we found a lot of variety, with a significant amount of people apologising, and others just conversing with the co-owners. These conversations had either an element of alleviating co-owners concern, e.g., “*I had a good talk and laugh with the person*” (P166u), very often including elements of persuasion to make think co-owners it was all right and no action was required. One particular type of persuasive strategy very common within the converse category (and hence we show the results for this separated) was flattery. For instance “*I reminded her how beautiful she is and pointed out the tons of comments confirming this*” (P14u). In some other cases, there was an explicit promise made from uploaders to co-owners that they would seek their permission before posting in the future, that they would not share again a photo of the co-owner, or there was one case of compensation in the form of a forfeit.

There were also some resolutions that were alternatives to removing. In particular, there were cases in which uploaders replaced the photo with another one of the same people (even in the same event sometimes), they cropped the co-owners out, or blurred faces. These strategies are interesting, as they are alternative ways of resolving an MPC that seek a compromise between what the uploader and co-owners want — as opposed to the more prevalent *all-or-nothing* approach described above. Another interesting resolution approach was that a photo shared using snapchat vanished after a while, resolving the conflict but also letting co-owners to talk about the photo and avoiding it from uploading to other social media platforms.

Other approaches to find a compromise included restricting access to the photo, or the uploader or co-owner removing the tag in an attempt to restrict access to the photo. More radical but hardly applied approaches were to block, unfriend, or stop speaking with the co-owners who complained. Importantly, these type of strategies were applied most of the time when the relationship between uploaders and co-owners was distant, which might be related to distant relationships being perceived as less important to maintain. Other than this, there were no other apparent correlations between relationships and strategies used.

Finally, there was a significant association between the strategy being used and the severity ($\chi^2 = 146.395$, $p = 0.00081$). In particular, Figure 2 shows that the higher the severity the more times a photo was removed. Also, there seems to be an interesting trend, which is that doing nothing happens more often than removing for the least severe items but then the inverse happens as severity increases. However, we can see cases of doing nothing at all levels, and even at the highest severity level, over 40% of the cases are solved this way.




Figure 2. Percentage of strategy use (X axis) per Severity level (Y axis).

Differences between uploaders and co-owners

We finally show the results obtained regarding RQ3, focusing on the differences between uploaders and co-owners that were found to be significant.

Divergence in perceived audience

Uploaders reported that 72% of the photos were shared with friends, 15% with public, and 1% did not know. In contrast, co-owners reported 47% of the photos were shared with friends, 34% with public, and 10% did not know. The disparity between groups was also apparent regarding the number of people perceived to have access to the photo, with a median of 200 for uploaders and 250 for co-owners; as well as people perceived to have looked at the photo, with median of 30 for uploaders and 50 for co-owners. Means and SDs were orders of magnitude higher than medians, which pointed to many outliers and answers not being normally distributed, but values were always much higher for co-owners than for uploaders. A Kruskal-Wallis test [30], which is a non-parametric extension of one-way ANOVA for data that may not be normally-distributed, was used to compare the values of uploaders and co-owners for how many people had access to the photo ($\chi^2 = 17.56$, $p < 0.00005$) and how many people looked at it ($\chi^2 = 30.17$, $p < 0.00001$), finding significant differences between groups in both cases.

Differences in perceptions of conflict resolution

Figure 3 shows that, overall, 70% of conflicts were resolved and 30% remained unresolved even after a conversation between the affected parties. However, we could again observe clear discrepancies between uploaders and co-owners, which a chi-squared test confirmed to be statistically significant ($\chi^2 = 129.75$, $p < 0.00001$). Uploaders considered 86% of issues resolved compared to only 54% of co-owners. However, recall from above that many co-owners did not complain (53% of them). When co-owners did not complain, 36% of the conflicts were considered resolved and 64% unresolved. In contrast, when co-owners did complain, 74% of the conflicts were considered resolved and 26% unresolved. Therefore, conflict communication plays and important role for a conflict to be resolved or not, which is consistent with current conflict resolution theories [17]. We can see that when co-owners complain the perceptions of the conflict being resolved or not are closer to those of uploaders. However, even if we just take the co-owners who complained, the differences are still statistically significant between uploaders and co-owners as confirmed by a chi-squared test ($\chi^2 = 16.34$, $p < 0.00005$), pointing to other factors that influence those differences beyond complaining/not-complaining. This divergence can be observed in Table 5 regarding uploaders doing nothing about a conflict. We can see that co-owners consider the conflict not to be resolved when uploaders did nothing about it in the majority of cases (78%), while uploaders only consider the conflict not to be resolved when they do nothing about it in less than 44% of the cases.

Differences in conflict resolution strategies

Beyond the similarities of uploaders and co-owners when it comes to frequencies of use of conflict resolution strategies




Figure 3. Whether or not MPCs were resolved by group and whether co-owners complained.

already stated above, there were also some statistically significant differences on the strategies used by uploaders and co-owners ($\chi^2 = 482$, $p < 0.00001$). Looking back at Table 5, it is already apparent the differences between removing the photo and doing nothing (which together account for 75% of cases) depending on the role. Interestingly enough, these differences are more apparent if we further split co-owners into those who complained and those who did not. In particular, Table 5 (columns for co-owners) show that most co-owners who did not complain ended up doing nothing about the conflict, which actually shows as leaving the conflict unresolved in the vast majority of cases (almost 80%) — note that if co-owners did not express their discontent to or ask uploaders about what to do directly, there are strategies that were not available to them, e.g., the photo would not be removed as the only one able to take a photo down is the one who uploaded it. In contrast, co-owners who did complain, they did something about the MPC, and it was the uploader who may or may not have done anything about the MPC.

Another strategy where differences were found was untagging. In particular, co-owners seemed to use it more often than uploaders ($\chi^2 = 95.89$, $p < 0.00001$). On a closer look, however, there seems to be again differences within co-owners depending on whether they did complain or not ($\chi^2 = 14.67$, $p < 0.0005$). If co-owners did not complain, there was little left for them to do unilaterally but untagging from the photo, or trying other means of restricting access, such as hiding the photo from the timeline without removing the tag. Once a co-owner complained, untagging was used much less than when a co-owner did not complain. Still, we could observe that untagging is sometimes used as a way of trying to hide the photo from potential undesired viewers. However, some of the users understand this does not completely resolve a conflict, e.g., “The issue was not entirely resolved. I untagged myself from the photo, but the photo remained online” (P845c).

DISCUSSION

We now discuss the main implications for design and for future research that stemmed from the results, particularly discussing and comparing to existing work in the area of multi-party privacy.

When is a photo co-owned?

In the majority of cases, co-ownership was somewhat clear, e.g., the people depicted in a photo could be considered the co-owners, and they felt entitled to upload/share the photo or to complain/contact the uploader as co-owners if they were

unhappy about the photo. However, there were many other *types* of co-ownership. One type was uploaders being co-owners despite not being themselves in the photo, as the authors and posters of the photo online. There were also co-owners who were not in the photo themselves, yet something *of them* was in the photo (car, house, room, pet) or someone they felt responsible for (e.g. their children), even when they may not be tagged in the photo either. There were also cases in which, despite being tagged in a photo, co-owners were not depicted in the photo but a reference was made to them (e.g. the crabs lady case discussed before). In other cases, uploaders intentionally avoided tagging someone *in* the photo as an attempt to hide the photo from them.

Understanding co-ownership is crucial to design tools for managing co-owned items as suggested earlier too [24]. However, our study uncovered co-ownership is a complex phenomenon with different manifestations, and engineering systems that are capable of being aware of co-owned items need to consider this complexity. Further research is needed on this topic, as most of the existing frameworks for multiparty access control [4, 46, 12, 21, 48, 49, 47, 22] start from the underlying assumption that the co-owners of an item are known in advance.

Friends and Family are danger

MPCs were reported much more often related to close relationships between friends (including best and close friends) and family (including siblings, parents, etc.), rather than more distant relationships such as acquaintances and co-workers, which together represented only 5% of the cases (see Characteristics of Photo MPCs). This was already suggested in [4] and our work quantitatively confirmed it. This, however, contrasts with previous research that suggested that trusted relationships may be an antidote for MPCs [13, 31]. What we found, however, is that even in very trusted relationships (wife/husband) privacy conflicts do exist and are not always resolved. One of the reasons that could explain the mismatch with previous results is that, as already anticipated in [31] and confirmed by our participants (53% of co-owners stated they would change their offline behavior to avoid a photo), users do consider the consequences in the online world and behave differently in the offline world, particularly in situations with less trusted people, e.g., “*I don’t allow untrustworthy people to take photos of me*” (P828c). This is also known as the extended *chilling* effect of online social networks in the off-line world [34]. This suggests there may be less opportunities for conflict with untrusted relationships, as offline interactions with untrusted people are controlled more tightly and less photos may be taken with them.

Audience Awareness

We observed a significant difference in the perceived audience of conflicting photos depending on the role. In particular, uploaders claimed to know who the audience was, stated the audience not to be the public, and reported less people looking at the conflicting photo. Although it is very well-known in the literature that users’ perceptions of the audience are never accurate [3], it suggests that co-owners perceive the

audience to be larger than uploaders. Importantly, the anticipated audience is known to influence very much disclosure decisions [26]. Audience visualisation tools could be used so that not only uploaders but also co-owners could have access to who the audience is for a particular photo. This would have the potential to reduce audience perception mismatch between co-owners and uploaders, and hence, help reduce the perceived privacy threat a photo may pose to them. Examples of such tools include AudienceView [32] and PViz [36].

Designing for minimally-intrusive pre-sharing interfaces

In the vast majority of cases (90%), MPCs were dealt with *after* a photo had been shared (see Process of Photo MPCs above). This has the drawback that conflict resolution happens after a privacy violation may have already happened. Indeed, a prevalent emerging theme from the coded free-text data in our study was a sense that the damage surrounding conflicts had already been done, and could not be rectified. There were preventative strategies uploaders reported to avoid conflicts *in general*, such as negotiating general rules of thumb for sharing (around 58%) or considering others before sharing (30%). However, even when participants explicitly reported having considered others before the particular MPC they reported, they were not able to predict whether/what co-owners would complain about. The need for support *before* sharing was acknowledged in previous research [31], but no specific recommendations were given for design. Based on our study, we suggest below three main avenues for future research in this area.

The first avenue could be context-aware nudges that help users manage multiparty sharing in particular scenarios, as our study revealed that there are dates, events, places, topics and relationships that seem to favour the emergence of MPCs (see Characteristics of MPCs above). In this sense, context-aware mechanisms [44] could be used to focus nudges similar to those described in [54, 53] on particular contexts. Another approach could be to learn users’ behaviour in particular contexts and social groups and make recommendations for other but similar contexts and social groups [11, 15].

The second avenue would be to explore the potential of restricted audience previews among co-owners before a photo is shared. This stems from: i) uploaders and co-owners perceived audience significantly different and would benefit from audience awareness mechanisms as already mentioned, and ii) some participants reported MPCs within a restricted audience in one social media infrastructure, which acted as a kind of firewall, preventing uploaders from posting the conflicting photo to other social media — e.g., a participant reported an MPC using Snapchat with a reduced audience, which resulted in the photo not being posted to Facebook. The audience visualisation tools already discussed before could also be helpful here to preview the prospective audience before sharing a photo.

The third avenue is support for conflict resolution, detailed in the next section, which could equally be used *before* or *after* sharing, depending on *when* the MPC is discovered.

Designing to facilitate compromise

Previous works encouraged the development of built-in capabilities for social media to facilitate a discussion between uploaders and co-owners [4, 31, 56]. However, what we found in our study is that a discussion may not be enough, as when actual negotiations happened (when uploader and co-owners communicated), there was little or no room given for compromise, and the photo was either removed or nothing was done about it in 75% of the MPCs reported in a very *all-or-nothing* approach. This could be related to the lack of awareness about other alternatives [55]. Therefore, while we are not suggesting that conflicts are always resolvable [41, 50], interfaces should show and facilitate the application of alternative conflict resolution actions to compromise and, hence, reach agreement beyond just removing or doing nothing.

The other reason for alternative conflict resolution mechanisms that facilitate compromise relate to the issue of whether participants complained or not about a conflict. In particular, participants showed a preference for *direct* methods to complain (e.g., in person, via telephone, etc.), which is in line with earlier evidence [56]. However, having to deal with MPCs directly is actually what put many co-owners off complaining (e.g., not to upset others, to avoid further attention, etc.). That is, if co-owners decided to complain, they were most likely to do it in a direct way, but there were many co-owners who decided not to complain precisely to avoid this direct interaction.

One could be tempted to suggest untagging as one of those alternative resolution approaches. However, when we observed the use of untagging, it was mainly as a last resort not a compromise, which was many times an alternative to interaction, and it was significantly reported more by co-owners who did not complain than co-owners who did complain or than uploaders. Moreover, a prevalent theme was a misunderstanding of what was achieved by untagging, and we only few participants who were knowledgeable about the limitations of such an approach: *“I will untag myself and make it so none of my friends can see it, unless they are friend with the person that posted it”* (P897c). We give two examples of alternative conflict resolution strategies that we found in the study that could help find a compromise.

The first one is audience negotiation. Restricting the audience of a photo was shown as one of the conflict resolution strategies participants enacted that always resolved a conflict. Some interfaces have been proposed to collaboratively select audiences, such as [46, 22]. However, the problem they have relate to the effort needed from users as well as the very direct approach followed. In contrast, other more automated approaches like [12, 21, 47, 49] aggregate the preferred audiences of all co-owners using pre-defined rules, which may not hold or render acceptable results in all situations. Therefore, the challenge for future research is to be able to recommend a subset of the audience that makes all co-owners happy with the minimal intervention possible, which could be then visualised using the audience visualisation tools mentioned earlier. Recommendations could be based on users' preferences and their reasons and/or simulated negotiations [19, 27].

The second refers to quick ways to replace/modify a conflicting photo, e.g., a built-in option/interface in the social media infrastructure to suggest another (non-conflicting) photo in the same event, to blur someone's face [23], to crop some people out (there are apps that do this automatically), etc. These were all conflict resolution strategies found in the study that were always perceived to resolve the conflicts regardless of whether participants were uploaders or co-owners, and participants applied them even if social media infrastructures do not provide built-in support for them yet.

Limitations and Further Work

As one of the usual practices for data reliability in MTurk, only US-based participants were used [35, 20, 24]. Thus, the results obtained might not generalise to other countries, though recent work did not find significant differences between US and Singapore citizens when it came to general multiparty privacy practices [13].

In line with previous research, we focused on photos for this study. However, users also share other co-owned items in social media, like posts, videos, events, etc. Beyond social media, multiparty privacy should also be studied in smart phones, augmented reality glasses [16], lifelogging [14, 28], and digital health devices [7].

CONCLUSION

For the first time in the related literature, we studied particular Multiparty Privacy Conflicts (MPCs) over co-owned photos from identification to resolution, using a survey designed following a critical incident methodology to collect photo MPCs from both uploaders and co-owners. This allowed us to establish the first empirical basis for the prevalence, context and severity of photo MPCs. Also, we uncovered nuances and complexities about MPCs not known before, including: the different notions of what a co-owned photo is and how it manifests in practice; that MPCs seem to be more frequent in relationships closer than previously thought but, at the same time, the most radical conflict resolution approaches (e.g., unfriending) seem to be used in more distant relationships; that uploaders and co-owners have significantly different perceptions of the audience of a co-owned photo; that more than half of the co-owners did not communicate with the uploaders about a MPC; that there is an all-or-nothing approach of either removing the photo or doing nothing about MPCs, which leads to many MPCs not being resolved; and that there are other alternative strategies shown to always resolve MPCs that could act as a compromise if supported appropriately from social media interfaces. Based on all of this, we derived key insights and specific implications to design next-generation multiparty privacy interfaces, pointing to particular areas of interest for further research.

ACKNOWLEDGEMENTS

We would like to thank the EPSRC for supporting this research under grant EP/M027805/1. We also thank Semo Abiola for helping code responses to open-ended questions, and Natalia Criado for her comments on draft versions of this paper. Finally, we thank the anonymous Associate Chairs and reviewers for their very helpful comments.

DATA STATEMENT

The anonymised version of the dataset is available at <http://doi.org/10.18742/RDM01-122>

APPENDIX

SURVEY QUESTIONS (uploaders' version, see dataset for co-owners)

For the following questions, please think back to the last time someone was unhappy about a personal photo of other people that you shared online (e.g., through Facebook or flickr or other online platforms).

1. When was the last time someone was unhappy about a photo that you shared online?
If you cannot recall the exact date, please use an approximate date as close to the event as you can remember.
2. What was shown in the photo and what was the photo about? (free text)
3. Were you shown in the photo yourself? (Y/N)
4. Were there other people shown in the photo?
 - No, there was no person (photo of a landscape)
 - No, only me
 - Yes, one other person
 - Yes, two other persons
 - Yes, three or more other
5. On which Website or through which app did you share the photo? (free text)
6. Did you ask other people feedback or permission before sharing the photo?
 - No
 - No, but I wish I had
 - Yes, I asked everyone in the photo
 - Yes, I asked some of the people in the photo
 - Yes, I only asked someone not in the photo
 - Yes, I asked some of the people in the photo, and some additional people not in the photo
7. If yes: How did you ask for feedback or permission? (free text)
8. How many times have you been abducted by aliens during the past 30 days? (attention check)
9. Who complained after you had shared the photo?
 - The provider of the Website or app removed or flagged my photo. (Y/N)
 - Someone in the photo complained. (Y/N)
 - Someone not in the photo complained. (Y/N)
10. How did they complain?
 - Online, directly on the photo (e.g., in the comments) (Y/N)
 - By email (Y/N)
 - By telephone (voice) (Y/N)
 - By telephone (SMS, MMS) (Y/N)
 - By messaging app (e.g., Whatsapp) (Y/N)
 - In person (Y/N)
 - Other (Y/N)
11. Why did they complain and what did they complain about? (free text)
12. What was your relationship with the person who complained? *Please provide details on how you knew them and how close you were.* (free text)
13. Did you understand their concerns? (free text)
14. How severe was the issue? (Likert scale with 1-not severe and 5-very severe)
15. Was the issue resolved? (Y/N)
16. Have you seen a five-legged zebra hiding in your wallet? (attention check)
 - Yes, once
 - Yes, two or more times
 - No
17. How was the issue resolved? (free text)
18. With whom did you share this photo initially?
 - With no one
 - With only one person
 - With a few people I manually selected (e.g. by entering their names or email addresses)
 - With my friends or followers only
 - With the public
 - Don't know/unsure
19. How many people ... *If unsure, please make a guess*
 - could have had access to the photo?
 - actually looked at the photo?
20. Have you later deleted the online photo or restricted who can see it?
 - Yes, I deleted / restricted the photo less than a day later
 - Yes, I deleted/restricted the photo less than a week later
 - Yes, I deleted/restricted the photo more than a week later
 - No, not this time, but I have deleted/restricted photos in the past
 - No
 - I don't know
21. If the photo was shared through Facebook or similar, has any of the people in the photo untagged from it?
 - Yes
 - No
 - Not shared through Facebook or equivalent
 - I don't know
22. In general, how often over the last 2 months have people been ...
 - unhappy you shared a certain photo or album?
 - happy you shared a certain photo or album?
 - complaining you shared a certain photo online?
23. In general, how do you prevent other people complaining who are unhappy about the photos you share online? (free text)
24. In general, how do you manage privacy and public access to the photos you share online? (Y/N for each)
 - I create multiple accounts in one or more social networks (Facebook, Instagram, flickr, etc.) to separate different audiences.
 - I apply a rule of thumb in decisions on sharing (e.g., I have an agreement with my friends that I will not photograph them during specific events, etc.)
 - I adjust my privacy settings to only allow access to certain people (e.g., permissions in Facebook only allow certain friends access, etc.)
 - I avoid publicizing content that could be problematic (e.g., if I take photos of friends who are drinking, I make sure that photos are not shared online.)
 - I use Snapchat or a similar app in which shared photos vanish after a while.
 - I use private communication channels (e.g., private messaging, email etc.)
 - I use private cloud storage (e.g., I share my iCloud, Google Drive, OneDrive account, etc.)
 - I make sure I do not take photos of people when I see someone actively avoiding my camera (e.g., I put my camera away if someone indicates to me that they do not wish to be photographed.)
 - I share personal and work photos through the same account (e.g., friends and work colleagues can collectively see my photos.)
25. In general, how do you resolve it when other people complain or are unhappy about the photos you share online? (free text)
26. Your gender
27. Your age
28. To what degree do the following statements apply to you personally? (Likert scale)
 - If someone does me a favour, I am prepared to return it.
 - If I suffer a serious wrong, I will take revenge as soon as possible, no matter what the cost.
 - If somebody puts me in a difficult position, I will do the same to him/her.
 - I go out of my way to help somebody who has been kind to me before.
 - If somebody offends me, I will offend him/her back.
 - I am ready to undergo personal costs to help somebody who helped me before.
 - I get over it relatively quickly when someone hurts my feelings.
 - When somebody has wronged me I often think about it for quite a while.
 - I tend to bear grudges.
 - When other people wrong me I try to just forgive and forget.
 - I have a positive attitude toward myself.
29. How many close friends do you have?
30. Which platforms do you use to upload or share your photos? *Please list all you use.*
31. How many online contacts do you have in total across these photo sharing sites?
32. How many hours per day are you online?
33. If you have any comments on this survey or on online photo sharing in general, please let us know. (optional, free text)

REFERENCES

1. P. Alasuutari, L. Bickman, and J. Brannen. The sage handbook of social research methods. 2008.
2. I. Altman. The environment and social behavior: Privacy, personal space, territory, and crowding. 1975.
3. M. S. Bernstein, E. Bakshy, M. Burke, and B. Karrer. Quantifying the invisible audience in social networks. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 21–30. ACM, 2013.
4. A. Besmer and H. Richter Lipford. Moving beyond untagging: photo privacy in a tagged world. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1563–1572. ACM, 2010.
5. J. Binder, A. Howes, and A. Sutcliffe. The problem of conflicting social spheres: effects of network structure on experienced tension in social network sites. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 965–974. ACM, 2009.

6. D. Boyd and A. Marwick. Social steganography: Privacy in networked publics. *International Communication Association, Boston, MA*, 2011.
7. I. Brown, A. A. Adams, et al. The ethical challenges of ubiquitous healthcare. *International Review of Information Ethics*, 8(12):53–60, 2007.
8. A. Bryman. *Social Research Methods*. Oxford University Press, New York, forth edition, 2012.
9. M. Buhrmester, T. Kwang, and S. D. Gosling. Amazon's mechanical turk a new source of inexpensive, yet high-quality, data? *Perspectives on psychological science*, 6(1):3–5, 2011.
10. M. Burke, C. Marlow, and T. Lento. Social network activity and social well-being. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 1909–1912. ACM, 2010.
11. G. Calikli, M. Law, A. K. Bandara, A. Russo, L. Dickens, B. A. Price, A. Stuart, M. Levine, and B. Nuseibeh. Privacy dynamics: Learning privacy norms for social software. In *Proceedings of the 11th International Workshop on Software Engineering for Adaptive and Self-Managing Systems*, pages 47–56. ACM, 2016.
12. B. Carminati and E. Ferrari. Collaborative access control in on-line social networks. In *IEEE CollaborateCom*, pages 231–240, 2011.
13. H. Cho and A. 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*, pages 503–514. ACM, 2016.
14. S. Clinch, N. Davies, M. Mikusz, P. Metzger, M. Langheinrich, A. Schmidt, and G. Ward. Collecting shared experiences through lifelogging: Lessons learned. *IEEE Pervasive Computing*, 15(1):58–67, 2016.
15. N. Criado and J. M. Such. Implicit contextual integrity in online social networks. *Information Sciences*, 325:48–69, 2015.
16. T. Denning, Z. Dehlawi, and T. Kohno. In situ with bystanders of augmented reality glasses: Perspectives on recording and privacy-mediating technologies. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems*, pages 2377–2386. ACM, 2014.
17. M. Deutsch, P. T. Coleman, and E. C. Marcus. *The handbook of conflict resolution: Theory and practice*. John Wiley & Sons, 2011.
18. J. C. Flanagan. The critical incident technique. *Psychological bulletin*, 51(4):327, 1954.
19. R. L. Fogues, P. Murukanniah, J. M. Such, and M. P. Singh. Sharing policies in multiuser privacy scenarios: Incorporating context, preferences, and arguments in decision making. *ACM Transactions on Computer-Human Interaction*, page In press., 2017.
20. U. Gadiraju, R. Kawase, S. Dietze, and G. Demartini. Understanding malicious behavior in crowdsourcing platforms: The case of online surveys. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, pages 1631–1640. ACM, 2015.
21. H. Hu, G.-J. Ahn, and J. Jorgensen. Multiparty access control for online social networks: model and mechanisms. *Knowledge and Data Engineering, IEEE Transactions on*, 25(7):1614–1627, 2013.
22. H. Hu, G.-J. Ahn, Z. Zhao, and D. Yang. Game theoretic analysis of multiparty access control in online social networks. In *Proceedings of the 19th ACM symposium on Access control models and technologies*, pages 93–102. ACM, 2014.
23. P. Ilia, I. Polakis, E. Athanasopoulos, F. Maggi, and S. Ioannidis. Face/off: Preventing privacy leakage from photos in social networks. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, pages 781–792. ACM, 2015.
24. H. Jia and H. Xu. Autonomous and interdependent: Collaborative privacy management on social networking sites. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pages 4286–4297. ACM, 2016.
25. A. N. Joinson. Looking at, looking up or keeping up with people?: motives and use of facebook. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*, pages 1027–1036. ACM, 2008.
26. S. Kairam, M. Brzozowski, D. Huffaker, and E. Chi. Talking in circles: selective sharing in google+. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 1065–1074. ACM, 2012.
27. D. Keküllüoğlu, N. Kökciyan, and P. Yolum. Strategies for privacy negotiation in online social networks. In *Proceedings of the 1st International Workshop on AI for Privacy and Security*, page 2. ACM, 2016.
28. M. Korayem, R. Templeman, D. Chen, and D. C. A. Kapadia. Enhancing lifelogging privacy by detecting screens. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pages 4309–4314. ACM, 2016.
29. H. Krasnova, S. Spiekermann, K. Koroleva, and T. Hildebrand. Online social networks: why we disclose. *Journal of Information Technology*, 25(2):109–125, 2010.
30. W. H. Kruskal and W. A. Wallis. Use of ranks in one-criterion variance analysis. *Journal of the American statistical Association*, 47(260):583–621, 1952.
31. A. Lampinen, V. Lehtinen, A. Lehmuskallio, and S. Tamminen. We're in it together: interpersonal management of disclosure in social network services. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 3217–3226. ACM, 2011.

32. H. Lipford, A. Besmer, and J. Watson. Understanding privacy settings in facebook with an audience view. In *Proceedings of the 1st Conference on Usability, Psychology, and Security*, pages 1–8. USENIX Association Berkeley, CA, USA, 2008.
33. K. M. MacQueen, E. McLellan, K. Kay, and B. Milstein. Codebook development for team-based qualitative analysis. *Cultural anthropology methods*, 10(2):31–36, 1998.
34. B. Marder, A. Joinson, A. Shankar, and D. Houghton. The extended chilling effect of facebook: The cold reality of ubiquitous social networking. *Computers in Human Behavior*, 60:582–592, 2016.
35. W. Mason and S. Suri. Conducting behavioral research on amazons mechanical turk. *Behavior research methods*, 44(1):1–23, 2012.
36. A. Mazzia, K. LeFevre, and E. Adar. The pviz comprehension tool for social network privacy settings. In *Proceedings of the Eighth Symposium on Usable Privacy and Security*, page 13. ACM, 2012.
37. H. Nissenbaum. Privacy as contextual integrity. *Wash. L. Rev.*, 79:119, 2004.
38. L. Palen and P. Dourish. Unpacking privacy for a networked world. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 129–136. ACM, 2003.
39. E. Peer, J. Vosgerau, and A. Acquisti. Reputation as a sufficient condition for data quality on amazon mechanical turk. *Behav Res methods*, pages 1–9, 2013.
40. S. Petronio. *Boundaries of privacy: Dialectics of disclosure*. Suny Press, 2012.
41. S. Preibusch, B. Hoser, S. Gürses, and B. Berendt. Ubiquitous social networks—opportunities and challenges for privacy-aware user modelling. In *Workshop on Data Mining for User Modelling at UM*, pages 50–62, 2007.
42. E. Rader, R. Wash, and B. Brooks. Stories as informal lessons about security. In *Proceedings of the Eighth Symposium on Usable Privacy and Security*, page 6. ACM, 2012.
43. J. Saldaña. *The coding manual for qualitative researchers*. Sage, 2015.
44. F. Schaub, B. Könings, and M. Weber. Context-adaptive privacy: Leveraging context awareness to support privacy decision making. *IEEE Pervasive Computing*, 14(1):34–43, 2015.
45. M. Sleeper, R. Balebako, S. Das, A. L. McConahy, J. Wiese, and L. F. Cranor. The post that wasn't: exploring self-censorship on facebook. In *Proceedings of the 2013 conference on Computer supported cooperative work*, pages 793–802. ACM, 2013.
46. A. C. Squicciarini, M. Shehab, and F. Paci. Collective privacy management in social networks. In *Proceedings of the 18th international conference on World wide web*, pages 521–530. ACM, 2009.
47. J. M. Such and N. Criado. Resolving multi-party privacy conflicts in social media. *IEEE Transactions on Knowledge and Data Engineering*, 28(7):1851–1863, 2016.
48. J. M. Such and M. Rovatsos. Privacy policy negotiation in social media. *ACM Transactions on Autonomous and Adaptive Systems*, 11(1):4:1–4:29, 2016.
49. K. Thomas, C. Grier, and D. M. Nicol. unfriendly: Multi-party privacy risks in social networks. In *Privacy Enhancing Technologies*, pages 236–252. Springer, 2010.
50. K. W. Thomas. Conflict and conflict management: Reflections and update. *Journal of organizational behavior*, 13(3):265–274, 1992.
51. K. Vaniea and Y. Rashidi. Tales of software updates: The process of updating software. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pages 3215–3226. ACM, 2016.
52. J. Vitak, P. Wisniewski, X. Page, A. Lampinen, E. Litt, R. De Wolf, P. G. Kelley, and M. Sleeper. The future of networked privacy: Challenges and opportunities. In *Proceedings of the 18th ACM Conference Companion on Computer Supported Cooperative Work & Social Computing*, pages 267–272. ACM, 2015.
53. Y. Wang, P. G. Leon, A. Acquisti, L. F. Cranor, A. Forget, and N. Sadeh. A field trial of privacy nudges for facebook. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 2367–2376. ACM, 2014.
54. Y. Wang, P. G. Leon, K. Scott, X. Chen, A. Acquisti, and L. F. Cranor. Privacy nudges for social media: an exploratory facebook study. In *Proceedings of the 22nd International Conference on World Wide Web*, pages 763–770. ACM, 2013.
55. P. Wisniewski, A. Islam, H. Richter Lipford, and D. C. Wilson. Framing and measuring multi-dimensional interpersonal privacy preferences of social networking site users. *Communications of the Association for Information Systems*, 38(1):10, 2016.
56. P. Wisniewski, H. Lipford, and D. Wilson. Fighting for my space: Coping mechanisms for sns boundary regulation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 609–618. ACM, 2012.
57. P. Wisniewski, H. Xu, H. Lipford, and E. Bello-Ogunu. Facebook apps and tagging: The trade-off between personal privacy and engaging with friends. *Journal of the Association for Information Science and Technology*, 66(9):1883–1896, 2015.
58. H. Xu. Reframing privacy 2.0 in online social network. *U. Pa. J. Const. L.*, 14:1077, 2011.