

Theodor Schnitzler^{†*}, Shujaat Mirza[†], Markus Dürmuth, and Christina Pöpper

SoK: Managing Longitudinal Privacy of Publicly Shared Personal Online Data

Abstract: Over the past decade, research has explored managing the availability of shared personal online data, with particular focus on longitudinal aspects of privacy. Yet, there is no taxonomy that takes user perspective and technical approaches into account. In this work, we systematize research on longitudinal privacy management of publicly shared personal online data from these two perspectives: user studies capturing users' interactions related to the availability of their online data and technical proposals limiting the availability of data. Following a systematic approach, we derive conflicts between these two sides that have not yet been addressed appropriately, resulting in a list of challenging open problems to be tackled by future research. While limitations of data availability in proposed approaches and real systems are mostly time-based, users' desired models are rather complex, taking into account content, audience, and the context in which data has been shared. Our systematic evaluation reveals interesting challenges broadly categorized by expiration conditions, data co-ownership, user awareness, and security and trust.

Keywords: Longitudinal Privacy Management, Online Exposure, Personal Data

DOI 10.2478/popets-2021-0013

Received 2020-05-31; revised 2020-09-15; accepted 2020-09-16.

1 Introduction

In their everyday life, users create huge amounts of data, shared online with varying audiences for different pur-


poses. Since data owners mostly do not track the availability of their online data once they initially shared it [22, 41, 57], there is a need for continuous exposure controls and increased attention for the lifetime ending of personal online data. Compared to the non-digital past, in which forgetting information was inherent, today's world with technical capabilities to permanently store information needs actively managed processes to realize forgetting [50].

A high-level overview of users' means to control their online exposure is provided by Bishop et al. [15]. They propose to better control the dissemination of data, e. g., by proactively employing sophisticated access control mechanisms, or by hiding the information within the enormous amount of data available online, such as by the release of large amounts of similar false information to confuse the interpreter. There is evidence that users have detailed perceptions of how to share data, but lack appropriate means to fulfill their goals. It has been shown, for a domesticity context, that users can precisely formulate who may access which of their data [51]. Moreover, users can distinguish different use cases when handling data and, therefore, switch between channels for communication and data sharing, depending on the task and content type [87]. On the downside, it also turned out that users have false perceptions of deleting data shared with others through online services [76] or in instant messengers [83].


While information processing and dissemination are essential aspects of privacy [90], we take a closer look at exposure control in particular. However, there can be a lot of reasons for data revocation or digital forgetting. In many cases, published content is not meant to be available permanently, but is only relevant for a short period of time in a certain context, e. g., when posted impulsively or out of momentum [10, 81]. Reducing exposure due to lack of relevance should not only be attributed to privacy, but can also help keep track of more important content, and fade out the rest. Exposure settings might also not match their data owners' perceptions for cases in which they did not foresee sharing consequences and, therefore, require later adjustment [86, 102].

Specific reasons are not even necessary – in the end, it can be deemed the users' sheer right to determine what is to happen with their data, and how long they


***Corresponding Author: Theodor Schnitzler[†]:** Ruhr-Universität Bochum, Bochum, Germany,

E-mail: theodor.schnitzler@rub.de 

Shujaat Mirza[†]: Courant Institute of Mathematical Sciences, New York University, New York City, NY, USA,

E-mail: shujaat.mirza@nyu.edu 

Markus Dürmuth: Ruhr-Universität Bochum, Bochum, Germany, E-mail: markus.duermuth@rub.de

Christina Pöpper: New York University Abu Dhabi, Abu Dhabi, UAE, E-mail: christina.poepper@nyu.edu 

[†]The first two authors contributed equally.

prefer it to remain available. Data sovereignty has received increased awareness over the past years, also due to the establishment of the *Right to be Forgotten* [28] as part of the European General Data Protection Regulation (GDPR) [27], even though data shared in online spaces is not the focus of this directive.

From a different perspective, research has put great efforts into developing technical approaches to assist users in managing their longitudinal privacy in general, and realizing data revocation in particular. However, such proposals have not found their way to wide-scale adoption, even though there has been a trend towards the use of tools providing better privacy and even some level of ephemerality [84].

In this paper, we take a closer look at this gap between how people use sharing mechanisms and privacy controls for their online data and concepts proposed by academia in order to facilitate online privacy management. To capture how people actually use online sharing mechanisms and privacy, we survey a large body of user studies carried out over the last decade. We categorize these studies along usage patterns, drivers that make users decide to unshare or reduce the exposure of user content, and the desires they have to improve their privacy experience. On the technical side, we survey concepts and proposals that assist users in managing their longitudinal privacy and the availability of their shared online data. We categorize these proposals along the use cases they have been designed for, the adversarial models they take into account, and the underlying protection mechanisms they avail to realize their privacy features.

By evaluating our systematization, we reveal conflicts between these two sides, such as intended use cases that do not appropriately reflect actual usage patterns. Referring to such conflicts, we derive a set of challenging open problems that need to be tackled by future research in order to develop privacy-enhancing technologies that can better assist users in managing their longitudinal online privacy and the availability of their data.

In summary, our paper provides the following contributions:

- We systematize how users interact with online services such as social networking sites in terms of their longitudinal online privacy management.
- We provide a taxonomy for technical systems to realize data revocation or to reduce exposure of publicly shared personal content as proposed in research.
- Based on the systematic analysis of previous work, we derive a set of challenges and open research questions that future research on data revocation

and longitudinal privacy management should aim to tackle.

Our work is first of its kind in combining knowledge from both user studies and technical mechanisms, providing a rich understanding of research efforts on longitudinal privacy management.

2 Systematization Methodology

We start systematizing existing research on longitudinal online privacy management by systematically collecting publications from major academic computer security and privacy venues or broader venues related to and relevant for our topic¹. We focus our targeted paper selection on the last decade. We identified a broad range of papers based on title and abstract and decided upon adding a publication to our final set of literature after having determined its general focus by skim reading its essential sections. We further take into account cross-references starting from the resulting literature set to achieve broad academic coverage of the topic.

Given this body of literature, we study the problem of managing the availability of personal online information from two perspectives: (i) Understanding user habits and desires regarding their longitudinal online privacy and (ii) Collecting technical proposals and concepts that are designed to manage online privacy. We provide an overview of our categorization process in Figure 1 and describe its methodology as follows.

2.1 Categorization Process

The initial systematizations of the two perspectives were drafted by one author each. This included selecting the initial sets of papers, creating a first set of labels as a means to categorize these papers, and assigning each paper such labels. Subsequently, four researchers in our team thoroughly discussed the initial systematizations in several rounds. Any concerns regarding label assignments or the set of papers had to be resolved, and updates required joint agreement of all four researchers.

As we will explain in-depth in Section 3, we systematize research on user attitudes towards privacy management and how users perceive selected aspects of it. For each publication in the list, we provide basic study

¹ We focus on IEEE S&P, USENIX Security, ACM CCS, NDSS, PETS, SOUPS, and CHI.

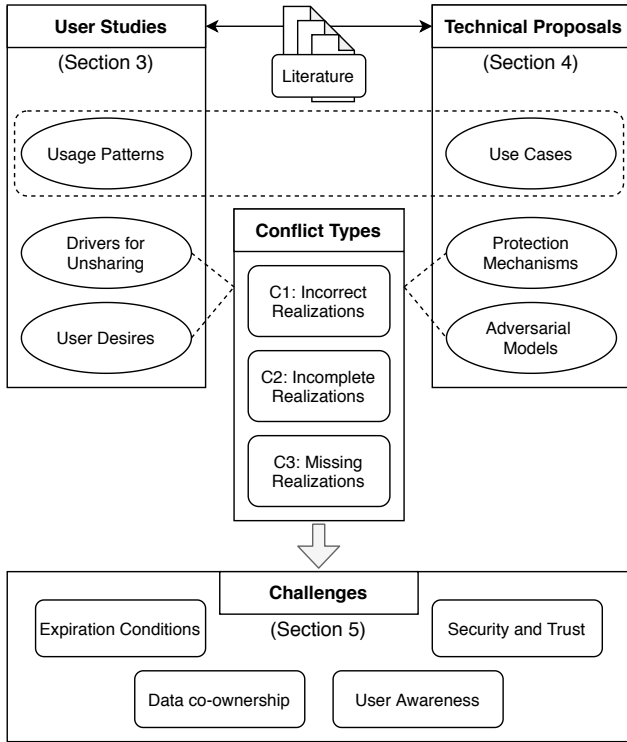


Fig. 1. High-level overview of our systematization methodology. We categorize previous work on User Studies and Technical Proposals along a set of features. Based on the interplay among different features, we derive technical or conceptual challenges worth to be further investigated.

meta-data and extract whether the work explicitly refers to longitudinal aspects of online privacy. We categorize research along the privacy management (*usage patterns*) that is covered, the identified reasons that make users change their initial privacy configuration (*drivers for unsharing*), and what *user desires* lead to a presumably improved privacy management experience.

In Section 4, we examine privacy controls that have been proposed or implemented as proofs-of-concepts. We systematize these controls and mechanisms along the *use cases* they have been designed for. We further categorize the *adversarial models* or adversarial settings that they should protect from, as well as the underlying *protection mechanisms* that they apply.

Discussing the set of papers was particularly necessary in the case of borderline papers, e. g., when it was unclear whether a paper indeed addressed publicly shared online data, which was a requirement for inclusion in the user studies systematization. We agreed that sharing data in cloud storage with an indefinite audience (e. g., co-students) should be sufficient to be considered *publicly* shared (cf. [41]). Similarly, for the systematization of technical proposals, detailed discussions were

held when it was unclear whether a proposal limited the availability of online data. For example, we agreed that adversarial examples helped reduce shared photos’ detection by smart recognition systems and therefore, these perturbations do indeed serve the users’ goal of limiting availability of their online data (cf. [58]).

We further adapted the set of categories using the same process. For example, we initially considered misconceptions expressed in user studies as a separate category; they however turned out to be too diverse to be systematized in detail. We decided to focus on misconceptions that affected users’ decisions about reducing exposure of their data, rendering them a sub-category of *Drivers for Unsharing*. On the technical systematization side, we decided to introduce insider adversary as a separate *adversarial model* after noticing that the existing threat models were not fully capturing the risks covered by this case.

One way to connect the two systematizations is by contrasting *usage patterns*, i. e., how users interact with privacy management options, and the *use cases* technical proposals are intended for, i. e., what they offer users for managing their privacy. Both systematizations capture to what extent content exposure can be limited or entirely ended, and if there is active user interaction involved in this process.

2.2 Deriving Challenges

Starting from the categories identified in either part of the systematization, we identified potential inconsistencies or conflicts between them. Pursuing a user-centric approach, we systematically examined to what extent users’ desires and their drivers for unsharing are reflected in the current state of technical proposals. We identified conflicts, whenever realizations in technical proposals are (i) *incorrect*, i. e., orthogonal to users’ needs, (ii) *incomplete*, i. e., promising but far from satisfying users’ requirements, or (iii) *missing*, i. e., not addressing users’ desires at all. For each conflict, we derived challenges on how such inconsistencies can be addressed.

By combining and contrasting knowledge from both of the obtained systematizations, conflicts were identified and challenges were derived by two researchers individually first and then discussed and iteratively updated. Again, challenges were subject to discussions among four researchers – proposals and concerns brought up by anyone of them had to be resolved and any updates required agreement of all four researchers.

As we will detail in Section 5, we followed a bottom-up approach: first, we derived fine-grained challenges related to conflicts, and then we put them into a broader context and related them to each other, resulting in a set of four challenge groups. The challenges we identify refer to (i) the *expiration conditions* under which data are supposed to disappear, (ii) *user awareness* of how particular privacy controls actually work, (iii) multi-user conflicts, which originate in the implicit *co-ownership of data*, when data affects the privacy of more than one individual, and (iv) issues regarding *security and trust* w. r. t. specific actors users consider when making changes in their online exposure.

3 Systematizing User Interaction

We first systematize users' preferences and behavior w. r. t. their longitudinal online privacy. We explain the different categories in our taxonomy and summarize our findings in Table 1. We arrange publications in three groups, each of which is ordered chronologically with most recent publications first.

3.1 Study Data

For each piece of research we cover in our systematization, we report the type of the user study that has been conducted: Self-reported data (S), Exploring real-world data with self-reported answers (R), Experiments based on prototype implementations (E), or Analyzing publicly available data sets (P).

Most studies cover scenarios that reflect a situation on a particular online platform, sometimes with a very specific focus, such as fitness social networks. While most studies have covered Facebook (FB) and Twitter (TW), we also find research on Snapchat (SC), Cloud Storage (CL) provided by Dropbox and Google, Fitness (FI) social networking sites, and the subsequently shut down platform Yik Yak (YK).

We further denote the number of participants that have taken part in each study (*Sample Size*), how participants have been recruited (*Participants Sample*), and basic demographics in terms of a gender distribution to provide information about the meaningfulness of results.

Considering the study type and the participants sample can usually hint towards potential study limitations. Qualitative research typically studies significantly smaller sample sizes, thus providing detailed insights

into very specific issues, compared to quantitative studies having larger groups of participants. However, even large samples, e. g., recruited via Amazon Mechanical Turk, do not always generalize for all users of a specific platform under observation, not at all for users of other platforms. Furthermore, it must be considered that self-reported data may not be as meaningful as practical experiments with real user content since alleged privacy attitudes have been shown to differ from actual behavior [23]. On the downside, practical experiments with real user data may deter rather privacy-sensitive users from participating in the study [57].

The focus of our systematization is on studies that explore *Publicly Shared Data* (denoted with ● in the respective column), which applies to all but one study [41] that partially covers public data (◐) since it primarily focuses on data stored in the cloud that *can* be shared with a limited audience. In a similar fashion, we also denote whether a study explicitly refers to longitudinal aspects of data sharing (●) or not (○).

3.2 Usage Patterns

We extract a set of *Usage Patterns* that can be applied to limit the exposure of online content, ranging from explicit deletion operations to exposure reduction, and auto-expiry. We define the patterns we identified within the existing literature as follows:

- *Delete Content* is an explicit action performed by a user to entirely remove content from a platform.
- *Delete Account* is another explicit action performed by users that entirely removes all of their content from the platform and also their account, such that there remains no direct representation of them on that platform.
- *Reduce Exposure (Actively)* covers controls users apply to actively manage the audience for a piece of content, such as, e. g., changing its visibility settings from public to friends only.
- *Reduce Exposure (Passively)* captures features that remove references from exposed content, without actually altering the content availability, such as, e. g., un-tagging a specific person in a shared photo.
- *Auto-expire* covers all mechanisms ensuring that published contents are made unavailable automatically when certain conditions are met. In particular, expiration takes effect without any further action to be taken by the owner or publisher of the content after its initial publication.

Table 1. Systematization of User Studies on Longitudinal Online Privacy. We arrange surveyed publications in three groups, (i) papers explicitly referring to *longitudinal* aspects of privacy, (ii) papers that study publicly shared data without referring to longitudinality, and (iii) papers that are still relevant to the topic but do not cover any of the categories we present in our systematization. Publications within each groups are ranked in chronological order with most recent publications first.

| Publication | | Study Data | | | | | Usage Pattern | | | | | Drivers for Unsharing | | | | User Desires | | | | | | | | |
|-------------|--------------|------------|----------|-------------|---------------------|-----------------|----------------------|-------------------|----------------|----------------|-----------------|----------------------------|-------------|-------------|--------------------|--------------|--------|----------------|-------|--------------------------|------------------------|--------------------------|----------------|-----------|
| Reference | Venue | Study Type | Platform | Sample Size | Participants Sample | Female/Male [%] | Publicly Shared Data | Longitudinal Data | Delete Content | Delete Account | Reduce Exposure | Reduce Exposure (Actively) | Auto-expire | Irrelevance | Change of Opinions | Regrets | Events | Misconceptions | Fears | Reduce Visibility (Time) | Content-based Audience | Control Friends' Content | Confirm Delete | User-view |
| [57] | CCS'19 | R | FB | 78 | AMT | 69/31 | ● | ● | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ |
| [53] | SOUPS'18 | S | - | 30 | UNI | 60/40 | ● | ● | ○ | ○ | ○ | ○ | ● | ● | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ |
| [60] | SOUPS'18 | S | - | 22 | - | 50/50 | ● | ● | ● | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ |
| [41] | CHI'18 | R | CL | 100 | AMT | 41/59 | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ |
| [56] | J-IEEE-IC'17 | P | TW | 100K | [P] | - | ● | ● | ● | ● | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ |
| [68] | J-HCI'17 | S | FB | 272 | AMT | 61/38 | ● | ● | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ |
| [54] | SOUPS'16 | P | TW | 100K | [P] | - | ● | ● | ● | ● | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ |
| [9] | WPES'13 | R | FB | 299 | AMT | 55/44 | ● | ● | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ |
| [6] | SOUPS'13 | S | FB | 193 | AMT | 40/59 | ● | ● | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ |
| [4] | SOUPS'19 | S | FI | 30 | CON | 50/50 | ● | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ |
| [35] | CHI'19 | S | SC | 1515 | Q | 57/43 | ● | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ |
| [77] | SOUPS'18 | S | - | 23 | UNI | 52/48 | ● | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ |
| [81] | CHI'17 | S | YK | 18 | UNI | 56/44 | ● | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ |
| [109] | WWW'16 | P | TW | 30K | [P] | - | ● | ○ | ● | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ |
| [14] | WLSM'16 | P | TW | 203K | [P] | - | ● | ○ | ● | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ |
| [25] | J-CHB'15 | S | FB | 380 | CON | 52/45 | ● | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ |
| [46] | WLSM'14 | P | TW | ALL | [P] | - | ● | ○ | ● | ● | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ |
| [86] | CHI'13 | S | TW | 1221 | AMT | 53/46 | ● | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ |
| [62] | HICCS'13 | R | FB | 68 | UNI | 38/62 | ● | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ |
| [3] | CSCW'13 | P | TW | 292K | [P] | - | ● | ○ | ● | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ |
| [48] | PerCom'12 | S | FB | 65 | UNI | 62/38 | ● | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ |
| [39] | SOUPS'12 | R | FB | 260 | WEB | 75/25 | ● | ○ | ● | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ |
| [102] | SOUPS'11 | S | FB | 569 | AMT | 64/36 | ● | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ |
| [26] | CHI'11 | E | FB | 33 | UNI | 50/50 | ● | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ |
| [80] | IFIP-HCI'11 | P,S | FB | 103 | WEB | 59/41 | ● | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ |
| [13] | CHI'10 | S | FB | 14 | UNI | 57/43 | ● | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ |
| [44] | UPSEC'8 | E | FB | 16 | UNI | 44/56 | ● | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ |
| [23] | PETS'17 | S | - | 60 | AMT | 37/63 | ● | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ |
| [29] | CSCW'17 | R | FB | 1706 | AMT | 58/41 | ● | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ |
| [87] | CHI'16 | S | - | 17 | WEB | 65/35 | ● | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ |
| [55] | SOUPS'14 | R | FB | 1239 | WEB | 24/76 | ● | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ |
| [94] | JPC'13 | P | FB | 5076 | [P] | - | ● | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ |
| [45] | IMC'11 | S | FB | 200 | AMT | 46/54 | ● | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ |

Study Type – **S**: Self-reported data, **P**: Public data analysis, **E**: Experiment based on prototype implementations, **R**: Survey with real user data

Platform – **TW**: Twitter, **FB**: Facebook, **SC**: Snapchat, **CL**: Cloud Storage, **YK**: Yik Yak, **FI**: Fitness Social Networks

Participants Sample – **AMT**: Amazon Mechanical Turk, **Q**: Qualtrics, **WEB**: Other Web Platforms, **UNI**: University Sample (various recruiting methods), **CON**: Convenience Sampling (Offline)

[**P**]: Public data analysis, no participants sample, **-**: No information provided

Previous work studies one or more of these patterns in detail within specific application scenarios. In Table 1, we mark this with a filled circle (●). If the usage pattern is not covered by a paper, we denote this with an empty circle (○).

3.3 Drivers for Unsharing

When it comes to the end of data lifetime, we are interested in users' motivation behind their decision to limit the visibility of data. We identified several drivers that determine users to unshare content on online platforms:

- *Irrelevance* denotes a situation in which content is withdrawn because it has become irrelevant or unimportant for the owner or its audience, and there is no more reason to keep it online.
- *Change of Opinions* indicates that content is withdrawn since the owner changed their opinion about the content exposure, without further specifying reasons.
- *Regrets* captures situations in which users revised their decisions to publish content due to explicitly stated regrets that came up after publication.
- *Events* means that some external event unrelated to the initial publishing has made its owner reason differently about the current level of exposure.
- *Misconceptions* denotes a general term that applies when participants expressed the actual level of exposure does not match what they perceived. In case there is a misconception, other factors (e.g., oversharing) may simultaneously apply.
- *Fears* captures situations in which users stated that they feared that specific groups of people could see their contents.

For all these features, we mark whether they were referred to in the considered publications (●) or they were not covered (○).

3.4 User Desires

In several studies, users have expressed desires for features facilitating their interaction with online services. Whenever such a desire is related to longitudinal online privacy or managing their online exposure, we consider it in our systematization. We identified five related user desires in our literature set:

- *Reduce Visibility (Time)* indicates that users expressed data to become less exposed over time after being published.

- *Content-based Audience* covers cases in which users desired to have the audience composed differently depending on the content of the data being published.
- *Control Friends' Content* means that users desired to control contents owned by their friends (in cases it affected their privacy).
- *Confirm Delete* captures cases in which users expressed that they did not want to have data automatically disappear, but preferred being prompted to confirm its deletion.
- *User-view* denotes a desired feature where users can view their own profile from the perspective of another user to better estimate the specific exposure implications of their privacy configuration.

4 Systematizing Technical Proposals

Technical proposals to tackle longitudinal privacy concerns have been considered and developed for a variety of platforms, such as online social networks (SN) like Facebook (FB) and Twitter (TW), cloud-based applications (CL), and messaging applications (MA); we also consider proposals that are platform-independent (PI). For the systematization of the technical proposals, we consider the *use case* for which they were designed, the *adversarial assumptions* under which they operate, and the *underlying protection mechanisms* they rely upon. We summarize our findings in Table 2 that arranges proposals in a chronological order with most recent publications first.

4.1 Use Cases

For each technical proposal we cover in our systematization, we detail the functionality it is intended to serve:

- *Delete Content* results in removing a piece of content from a platform so that it is no longer publicly accessible. A proposal that provides such guarantees is labeled ●, as opposed to ○.
- *Reduce Exposure* allows users to manage the visibility of a piece of content on a platform such that it is exposed only to a subset of the previous audience. A proposal that allows such functionality is labeled ●, as opposed to ○.
- *User Involvement* captures the nature of the involvement of the data owner while limiting content

Table 2. Systematization of Technical Proposals for Longitudinal Online Privacy. We arrange surveyed mechanisms designed for a variety of platforms, use cases, adversarial assumptions and underlying protection mechanisms. Publications are ranked in a chronological order with most recent publications first.

| Publication | | Use Cases | Adversarial Models | | | | Underlying Protection Mechanisms | | | | | | | | | | |
|-------------|----------------|-----------|--------------------|-----------------|------------------|------------------|----------------------------------|--------------------|-------------|---------|--------------------------|--------------------------|----------------------|----------------------|-------------------------|------------------|-------------------------------------|
| Reference | Venue | Platform | Delete Content | Reduce Exposure | User Involvement | # of Data Owners | Retroactive | Honest-but-curious | Interfering | Insider | Cryptographic/Signatures | Distributed Architecture | Adversarial Examples | Deception & Flooding | Access Control Policies | Game-theoretical | Others/[Specifics] |
| [52] | PETS'19 | TW | ○ | ● | P | 1 | ○ | ○ | ● | ○ | ○ | ○ | ● | ○ | ○ | | Intermittent withdrawal |
| [106] | ForensicSec'19 | CL | ○ | ● | P | n | ○ | ● | ● | ○ | ● | ○ | ○ | ○ | ○ | | [Attribute-based collaboration] |
| [82] | IFIP-SEC'19 | PI | ● | ● | P | 1 | ● | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | | Smart contracts |
| [32] | NeurIPS'19 | PI | ● | ○ | P | 1 | ● | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | | Quantized k-means |
| [66] | NDSS'18 | PI | ○ | ● | A | n | ○ | ● | ● | ○ | ● | ○ | ○ | ○ | ○ | | Identity management system |
| [5] | CODASPY'18 | PI | ● | ○ | P | 1 | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | | [Time-lock puzzles] |
| [38] | CODASPY'17 | SN | ○ | ○ | P | n | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | | [Threshold secret-sharing] |
| [65] | ICCV'17 | SN | ○ | ● | P | 1 | ○ | ● | ● | ○ | ○ | ○ | ○ | ○ | ○ | | [Adversarial Image perturbations] |
| [58] | CVPR'17 | SN | ○ | ● | P | 1 | ○ | ● | ● | ○ | ○ | ○ | ○ | ○ | ○ | | [Adversarial Image perturbations] |
| [74] | GameSec'17 | SN | ○ | ● | P | n | ○ | ● | ● | ○ | ○ | ○ | ○ | ○ | ○ | | [Negotiation] |
| [104] | ETHReport'17 | CL | ● | ○ | A | n | ○ | ● | ● | ● | ○ | ○ | ○ | ○ | ○ | | [Group secret] |
| [7] | CCS'16 | CL | ● | ● | A | 1 | ● | ● | ○ | ○ | ○ | ○ | ○ | ○ | ○ | | Interdependency in encrypted |
| [108] | CODASPY'16 | PI | ● | ○ | P | 1 | ● | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | | [DNS Caching] |
| [95] | TKDE'16 | SN | ○ | ● | P | n | ○ | ● | ○ | ○ | ○ | ○ | ○ | ○ | ○ | | [Computational conflict resolution] |
| [18] | S&P'15 | PI | ● | ○ | P | 1 | ● | ● | ● | ○ | ○ | ○ | ○ | ○ | ○ | | Machine Unlearning |
| [64] | SIGMOD'15 | PI | ● | ● | P | 1 | ○ | ● | ● | ○ | ○ | ○ | ○ | ○ | ○ | | Brain-inspired data retention |
| [1] | ACM-SCC'15 | CL | ● | ○ | P | 1 | ○ | ● | ● | ○ | ○ | ○ | ○ | ○ | ○ | | Forgetful data structures |
| [89] | CCSW'13 | CL | ○ | ○ | P | 1 | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | | Heterogeneous documents |
| [15] | NSPW'13 | PI | ○ | ○ | A | 1 | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | | [False attribution] |
| [93] | IEEE-PST'13 | SN | ○ | ● | A | n | ● | ● | ○ | ○ | ○ | ○ | ○ | ○ | ○ | | User-to-content relations |
| [24] | S&P'12 | TW | ○ | ● | P | 1 | ○ | ● | ● | ○ | ○ | ○ | ○ | ○ | ○ | | [Blind RSA signatures] |
| [79] | WPES'12 | PI | ● | ○ | P | 1 | ● | ● | ○ | ○ | ○ | ○ | ○ | ○ | ○ | | Statistical webpage changes |
| [11] | PETS'11 | SN | ○ | ● | A | 1 | ○ | ● | ● | ○ | ○ | ○ | ○ | ○ | ○ | | [OpenPGP] |
| [20] | ICNP'11 | PI | ● | ○ | P | 1 | ● | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | | [DNS Caching] |
| [31] | UW-CSE'11 | PI | ● | ○ | P | 1 | ● | ● | ○ | ○ | ○ | ○ | ○ | ○ | ○ | | Integrating diverse mechanisms |
| [19] | CollbCom'11 | SN | ○ | ● | P | n | ○ | ● | ○ | ○ | ○ | ○ | ○ | ○ | ○ | | [Aggregation of policies] |
| [97] | PETS'10 | SN | ○ | ● | P | n | ○ | ● | ○ | ○ | ○ | ○ | ○ | ○ | ○ | | [Aggregation of policies] |
| [13] | CHI'10 | FB | ○ | ● | A | n | ○ | ○ | ● | ○ | ○ | ○ | ○ | ○ | ○ | | [Manual conflict resolution] |
| [105] | POLICY'10 | SN | ○ | ● | A | n | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | | [Manual conflict resolution] |
| [72] | ACSAC'10 | MA | ● | ○ | P | 1 | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | | Porter storage |
| [30] | USENIX'09 | PI | ● | ○ | P | 1 | ● | ○ | ○ | ○ | ○ | ○ | ○ | ○ | ○ | | [DHTs of P2P networks] |
| [91] | WWW'09 | SN | ○ | ● | P | n | ○ | ● | ○ | ○ | ○ | ○ | ○ | ○ | ○ | | Auction-based inference |
| [47] | CSE'09 | SN | ○ | ● | P | 1 | ○ | ● | ● | ○ | ○ | ○ | ○ | ○ | ○ | | Third party storage server |
| [16] | SecureCom'09 | PI | ○ | ● | A | 1 | ○ | ○ | ● | ● | ○ | ○ | ○ | ○ | ○ | | Bait information |
| [70] | SMLI'05 | MA | ● | ○ | P | 1 | ● | ○ | ● | ○ | ○ | ○ | ○ | ○ | ○ | | [Centralized server storing keys] |

Platform – TW: Twitter, FB: Facebook, SN: (general) Social Networks, CL: Cloud Storage, MA: Messaging Applications, PI: Platform Independent
User Involvement – A: Active, P: Passive ; **# of data owners** – 1: Single user scenario, n: Multi-user scenario

availability. If the process requires the data owner to actively change the content availability, it is labeled active (A). Otherwise, if the process relies on a mechanism that ensures automatic change in the availability of published content, then we denote it as passive (P). The passive case turns out to be more common.

- *# of Data Owners* captures the number of users making the decision to change the availability of content. In most cases, the data is owned and uploaded by a single user, denoted by 1. Multi-user scenarios that involve content co-owned by more than one user are denoted by n and are also common, but apply to slightly fewer proposals.

4.2 Adversarial Models

The Dolev-Yao (DY) adversary model is widely used to analyze system and network protocols [21]. For many settings, this model is, however, too strong: many legitimate participants of the protocol, such as service providers or fellow users with varying degrees of association, do not qualify to be DY adversaries. This does not imply that these parties cannot be malicious, though, so it is important to consider the relevant threat vectors. We, therefore, analyze the privacy guarantees of existing proposals against the following threat models:

- *Retroactive adversaries* learn which data they are interested in only after the data has been revoked/expired. This threat model makes an assumption that the attacker has no interest in accessing the published data prior to its expiration. Since the data was publicly available during its lifetime, it is not assumed to be private and accessible by everyone. However, past its expiration time, the privacy of deleted data is ensured.
- *Honest-but-curious adversaries* act as a legitimate party in a protocol that will not deviate from the definition but will attempt to learn as much information as possible. The majority of these adversaries are service providers who are handling users' data and running analyses on top of it. These adversaries are also referred to as 'curious-but-non-interfering' or 'passive' mainly due to their tendency to indiscriminately collect data once available in the hope that it may be of interest to them in the future.
- *Interfering adversaries* actively interfere with the private information of the user, either preponing or postponing the event limiting the availability of the content. This threat model treats clients in the sys-

tem as untrusted: they may bypass the system to publish sensitive content without obtaining consent from the target users through means such as colluding with other malicious clients and deviating from the protocol description.

- *Insider adversaries* control user devices, including porter devices, and can compromise users' passwords and passphrases. An insider attack may be intentional or accidental. Insider attackers range from poorly trained administrators who make mistakes, to malicious individuals who intentionally compromise the security of systems.

We rate the adversarial model of each technical proposal w. r. t. these attacker types. If a proposal considers a specific adversary in their threat model, we label it with ●. Otherwise, if it provides no guarantees against a specific adversary, then it is labeled with ○. The honest-but-curious adversary is the most commonly considered threat model, but the other adversaries are also being considered when technical solutions are proposed.

4.3 Underlying Protection Mechanisms

To realize use cases and fulfill adversarial guarantees, each proposal relies on different technical mechanisms. A number of protection mechanism principles have been proposed multiple times in varying realizations; others have occurred less frequently.

- *Cryptographic* mechanisms embed encryption keys into stored data within centralized or distributed storage systems. They may control the extent of the keys' replication to prevent the key from being recovered from the underlying storage after a configurable amount of time. Most of the time-based data revocation proposals rely on encryption by uploading the data in encrypted form along with information on where and how to gather the decryption key during content's lifetime. This category also covers *digital signatures* that allow users to embed signatures to the content.
- *Distributed Architectures* allow members to collectively generate and distribute group secrets among themselves. In order to avoid single-point failures, cryptography-based forgetting schemes avoid putting trust in a central authority for the storage of keys [20, 30]. Instead, they rely on key-sharing and distributing parts of the decryption key on distributed storage. Some approaches have yielded support for an 'expiration date' of a few days by spreading bits of the key among random indices in the

DHT [30] whereas others demonstrated expiration times of up to months by exploiting the evolving nature of webpages and using threshold secret sharing scheme to reconstruct the key [79].

- *Adversarial Examples* confuse AI/recognition systems effectively by generating additive perturbations that are invisible to the human eye, thus without introducing unpleasant artifacts. Given the prevalence of AI systems, such as facial recognition, adversarial examples could allow users to limit their content’s exposure to these algorithms (i. e., go undetected.)
- *Deception & Flooding* approaches require the subject to release large amounts of similar synthetic, but convincing, information that is not correct. The viewer is thus challenged to pick the correct confidential information from the mass of incorrect information.
- *Access Control Policies* are the classical approach to specify how access is managed and who may access information under what circumstances. These policies can be set manually, computed through aggregation, or learned over time using ML algorithms.
- *Game-theoretical* frameworks aim to achieve optimal decision making of independent and competing actors in a strategic setting. It can be used to understand and predict the effect of multi-party involvement in access control decisions on individual behaviors of social network users.
- *Others/[Specifics]*: In addition to the above categories, the existing literature relied on less-frequent protection mechanisms, such as approaches that mimic the human brain, smart contracts, porter storage devices, etc. We list them individually by name. In some cases, we also list specifics of mechanisms covered in one of the above categories. In such a case, we list them in brackets, for it is an explanation instead of a new category.

5 Technical Key Challenges

Based on our systematizations in Sections 3 and 4, we determine a set of technically challenging problems that have not been solved to date. We explore to what extent users’ desires and their drivers for unsharing, as expressed in user studies, have been realized as part of technical proposals. Whenever we identify factors that have not been appropriately addressed on the technical side, i. e., when realizations are incorrect, incomplete,

or missing, we identify this as a conflict to be resolved, each resulting in one or more challenges.

We determine these challenges first and then group similar ones and consider them also in context with each other. Our systematization results in challenges that are broadly categorized regarding (i) the expiration conditions under which data are supposed to be rendered unavailable (Section 5.1), (ii) the co-ownership of data resulting in potential conflicts among multiple users (Section 5.2), (iii) user awareness regarding the functionality of privacy controls (Section 5.3), and (iv) security and trust relations among the parties involved in data publishing (Section 5.4). The overall list of challenges per group is illustrated in Figure 2.

5.1 Expiration Conditions

Multiple studies reported in Section 3 have found that participants did not want contents to fade away wholesale with age [9, 41, 68]. Whereas participants of these studies have shown a preference for a handful of posts to become more private over time, they demonstrated their desire to make some posts *more visible* over time. Thus, the decision on content’s exposure control is a complicated one, hardly captured in the true sense by focusing alone on the age of posting.

Studies have identified other contextual factors such as inactivity of the post (e. g., lack of viewing/sharing) [54, 56] and major life events (e. g., moving to a new city or graduation) [6] that could impact users’ desire to keep the data publicly available. Users’ preference to limit exposure also largely depends on the content of their data, and effective audience control mechanisms can facilitate their openness to share [51, 60, 68]. In this regard, private-by-default interfaces, such as Snapchat, that allow audience-related considerations to be made on a per-post basis, result in users being much more audience-aware [2, 35]. In contrast, content sharing interfaces that are not as intuitive to per-post based audience decisions result in content being overexposed w. r. t. the uploaders’ intentions [12].

The overview of technical proposals in Section 4 shows that the most commonly considered condition for data revocation in previous academic proposals is the time passed since publication [30, 70, 79]. Solutions for end-users also use time as an expiration condition [67, 73, 88, 96]. Time-based mechanisms for data revocation are easily comprehensible and provide transparently decidable expiration conditions. However, each

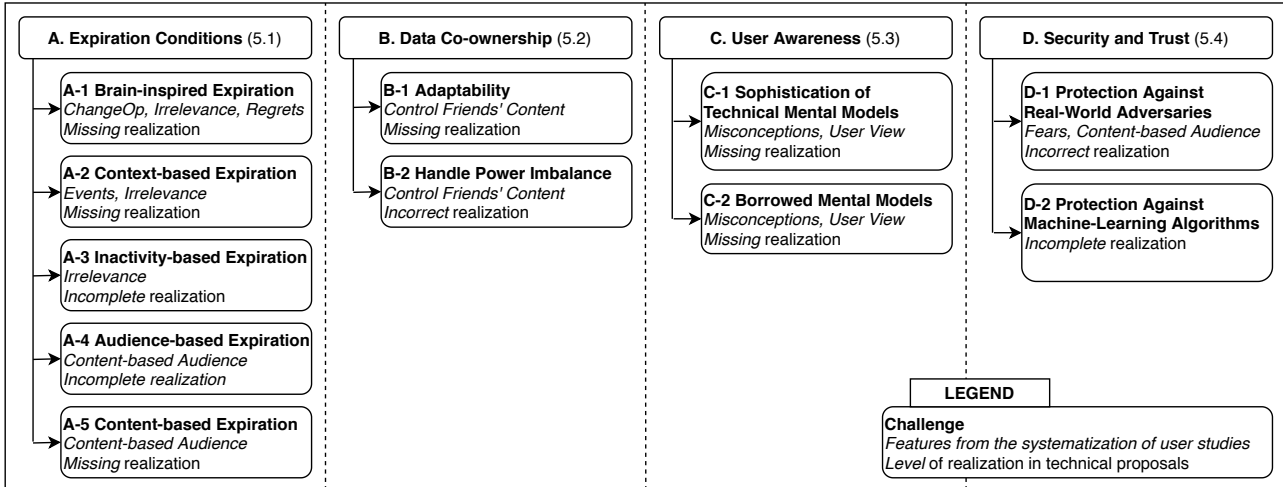


Fig. 2. Overview of the challenges we derived from conflicts identified in the systematizations of user studies and technical proposals, grouped by four topic areas: Expiration Conditions, Data Co-ownership, User Awareness, and Security and Trust. We denote to which feature(s) of the user studies systematization each challenge refers (bottom line) and to what extent they are currently addressed in technical proposals (in terms of realization level).

expiration time is determined and set at the time of publishing of data, which leads to a three-fold conflict:

- (i) the appropriate time for data revocation is often difficult to determine in advance,
- (ii) the context in which data is published (and in which the expiration condition is set) can change, which may require to adapt the expiration condition, and
- (iii) no context information or other potentially relevant aspects for deciding whether data should remain online or not are taken into consideration when the expiration condition is determined.

Improving revocation mechanisms is a complex problem, as it must take into account multiple contradictory factors, such as the desire to retain some old content while allowing other content to be completely removed. Based on our systematization of user studies and technical proposals, we identify and present challenging research dimensions that are desired by the users but have not yet been effectively realized in the technical implementations.

The first two challenges, A-1 and A-2, tackle missing realizations, taking into account multiple drivers for unsharing as expressed by users. Challenge A-3 takes up on work that already considers relevance as a factor to determine expiration, focusing on how to overcome its yet incomplete realization. We emphasize that there is an overlap between A-1 and the two subsequent challenges. Whereas A-1 provides a more holistic viewpoint, the other two can be considered specific cases of it. However, A-2 and A-3 can also be tackled independently and

do not require A-1 to be resolved. Finally, challenges A-4 and A-5 deal with incomplete and missing realizations in the interplay between published contents and audiences.

Challenge A-1: Brain-inspired Expiration

All existing mechanisms proposed have in common that the data revocation mechanism is implemented as a feature in terms of an explicit process. In contrast, Müller and Pilzecker's classical work [59] on retroactive inhibition in human memory found that forgetting is not a process that is actively triggered, but an implicit result of multiple information interfering with each other with more relevant information suppressing other information. What gets preserved in long-term memory may depend on multiple factors, including the 'meaningfulness' of the memory [17]. This can be transferred to our observations in the user studies systematization, where also multiple different factors implicitly contribute to the appropriateness of expiration conditions.

The technical challenge here is to imitate this behavior within a file storage system, i. e., to make access to information more difficult, the more new information is added, thus, waiving the need for explicitly revoking such information. In recent years, some research efforts have provided a promising start towards formalizing models imitating workings of human memory for their information management processes [1, 63, 64]. That being said, we are far from letting go of hard demarcation

of data availability and realizing mechanisms that have contents fade away over time, which is why we keep labelling this challenge as missing (cf. Figure 2).

Challenge A-2: Context-based Expiration

External factors, such as changes in life circumstances, can impact users' privacy preferences for online content, possibly due to changes in social circles or individual preferences. Since users do not explicitly formulate contextual factors, such as major life events, reflecting them in the deletion mechanism is still a major technical challenge. Service providers who aggregate a lot of information about individual users would possibly be able to design mechanisms that incorporate information about users and their social circles to change the visibility of published data. However, this is rather difficult for cryptographic erasure mechanisms applied to standalone information that is published anonymously and/or not related to any other source of information. Besides its limited technical feasibility, additional information aggregation also raises questions about privacy implications.

Challenge A-3: Inactivity-based Expiration

Some mechanisms [56, 108] have attempted, with varying levels of success, to realize expiration based on the amount of attention/interactions attracted by the data object. However, sole reliance on this model does not fully capture all practical aspects: some users choose to keep/archive some content even after it becomes inactive. Thus, it is technically challenging to realize an inactivity-based expiration solution that is equipped to identify user-specific content features which contribute to their willingness to keep the content alive despite its inactive status. Another challenging aspect under such implementations is that posts containing controversial content will elicit considerable attention and thus will continue to remain in the public domain for longer.

Challenge A-4: Audience-based Expiration

People do share not only different types of data but also have multiple heterogeneous groups of audiences accessing their contents. While cryptographic erasure mechanisms assume that everyone can read published data under the same conditions, there is a variety of access con-

trol settings available in social networks or cloud storage systems to satisfy the need to manage data for different target audiences. Adoption of audience-specific privacy controls suggests that not all readers of ephemeral data should be affected by exposure control decisions in the same way, but that there should be different conditions for individual users or groups of users. This leads to the technical challenge of realizing mechanisms that implement audience-dependent expiration conditions.

Challenge A-5: Content-based Expiration

Studies on changes in users' preferences about data availability have also captured the contents of data [51, 60, 68]. The challenge to realize more sophisticated expiration conditions is not limited to incorporating appropriate external factors. The data items themselves should also be taken into account, both in terms of their file formats and their contents or structural parameters. This requires to determine appropriate conditions for each type of data and to analyze data upon publishing in order to map them according to the categorization.

5.2 Data Co-ownership

A significant number of items uploaded to Online Social Networks (OSNs) involve multiple parties who are supposed to be interested in controlling its exposure to the public. Such items range from photos that depict multiple users to comments that mention multiple users to events in which multiple users are invited. Existing implementations of OSNs have not successfully tackled the problem of conflicting privacy preferences among users that co-own a piece of data.

In real-world applications such as Instagram, users uploading a photo can tag other users who are also present in or related to that photo. The tagged user can then control the visibility of the photo on their profile by hiding the tagged photo or deleting the tag itself. Neither of these options affects the visibility of the tagged photo on the whole platform since followers of the uploader are guaranteed access regardless of other tagged users' visibility preferences. When we recall that even preferences of individual users do not remain constant, it appears reasonable that merging the privacy preferences of multiple users is likely to end in conflict. The lack of appropriate conflict resolution mechanisms in the current implementations of OSNs can lead to pri-

vacy violations with serious outcomes for the parties involved.

User studies on online privacy management often refer to multi-user scenarios as a use case, for example, for photos being taken at parties or social events. However, the set of research that actually covers multi-user scenarios and their implications is rather small, even though users have expressed a desire to control their friends' content when it affects them already ten years ago [13]. The only privacy management measure suitable in multi-user scenarios that is covered by several studies is untagging but from different perspectives such as its overall prevalence [25], or revisiting initially set and possibly erroneous privacy settings [39, 48]. Eventually, users' strategy to overcome the risk of being unintentionally exposed publicly is preventing photos from being taken at all [77].

Research proposals that require users to collectively solve their privacy conflicts [91, 105] comprise promising concepts but lack practical evaluations of their acceptance in real-world applications. Other proposed mechanisms that automate this process rely heavily on fixed rules (majority voting, veto voting, etc.) [19, 97], thus, resulting in oversimplification of the conflict resolution process and mismatch between actual user behavior and the suggested method for resolving privacy conflicts. Such and Criado [95] proposed a promising computational model that adapts conflict resolution strategy based on the sensitivity of the item being shared and relative importance of the conflict (estimated through the strength of the relationship between owners and the target audiences). However, their mechanism does not take into account the strength of the relationship between negotiators and the role of history of previous negotiations on concessions in the current conflict. Furthermore, the approach does not take into account the effect of types of data items under consideration. In a rather restrictive proposal by Olteanu et al. [66], photos can only be uploaded to a social network site with all faces detected in it being removed, only allowing to display them after the corresponding person has explicitly agreed.

Designing a model that is complex enough to emulate user behavior most of the time, and that requires minimum intervention from the user's side is indeed challenging. From a legal perspective, proposals that use, e. g., *majority voting* do not seem to uphold users' right to be forgotten as prescribed in the recent regulations – as soon as one of the involved users wants an item to be deleted, it has to be removed if we strictly interpret the European GDPR [71].

While multiple or evolving drivers for unsharing already apply to single-user scenarios [6, 68] (cf. Section 5.1), expanding their concepts to multi-user settings raises additional challenges. The challenges listed here are related to realizations of users' desires to control their friends' contents in case it also affects themselves.

Challenge B-1: Adaptability

It is technically challenging to devise a model that takes into account the past history of negotiations between co-owners when deciding on the privacy preferences for new items. Since major OSNs keep a record of all postings on one's profile, it is likely that exposure settings for the past co-owned postings may no longer serve users' privacy requirements in the present context. Individual preferences for existing items may equally evolve and need to be adapted. Allowing users the option to renegotiate the privacy settings for co-owned items might be necessary for these models to be widely adopted. However, realizations of adaptable exposure controls for co-owned data items are missing in current realizations.

Challenge B-2: Handling Power Imbalance

Another challenge involving co-ownership of data on OSNs is that users' attitudes towards each other do not remain constant. On most of the platforms, users have the option to unfriend or even 'block' other users, rendering their profiles inaccessible. In the aftermath of such an event, users are denied the power to access the co-owned data items on the other user's profile. It is challenging to come up with a solution that honors users' unfriending decision while still ensuring their right to manage the co-owned data items.

5.3 User Awareness

Kang et al. identified that people with more articulated technical models on average expressed higher awareness of who could access their data [40]. A Better understanding of the number of privacy threats was found to be correlated with the protective actions taken by the individuals [69]. Internet users have been found to struggle to update their existing models at a rate comparable to the change in the internet and online platforms. In fact, prior privacy studies have identified that only a

few participants expressed awareness that their models might be outdated [40]. Prior work has also called for serious attention towards the presence of age gap in information behavior. Yong found out that older people are less skillful in privacy control and, therefore, are more susceptible to become the victims of privacy-related breaches [69]. The situation is further complicated by a lack of enthusiasm on older users' part in seeking help with privacy-related technology to avoid social embarrassment. To put the demographics into perspective, Facebook alone has at least 20% of its user base aged above 45 [92]. The matters are worsened as technical mechanisms operate under various levels of adversarial assumptions and rely on a variety of different protection mechanisms; the average user is usually not technically proficient or aware to update their mental model about different security functionalities. It is, therefore, not surprising that multiple studies reported misconceptions as one of the major drivers behind users' unsharing of data [9, 14, 86]

There also exist vast differences in the implementation of security-related features across different services (e.g. social networks vs. messaging applications) and different platforms within a service (e.g. Facebook vs. Twitter). Talking specifically about implementations of content deletion, there exist inconsistencies across:

- (i) services – the way Facebook (SN) implements deletion for shared postings within a group is different from the way Facebook Messenger (MA) tackles deletion of messages in a group. Similarly, users lack information on how deletion would work for cloud storage. Findings of Ramokapane et al. study attribute users' failure to delete from cloud storage to the lack of information about how cloud and deletion within the cloud functions [76].
- (ii) platforms – whereas deletion of a post on Facebook (SN) makes the related comments and re-shares on the post unavailable, it is not the same for Twitter (SN), where residual tweets (interactions associated with the withdrawn post) continue to leak information about the withdrawn tweet [54]. Similarly, disparities in the implementation of deletion functionality exist for messaging platforms. Skype (MA) allows the message sender to delete messages from the logs of all participants in the conversation, whereas Facebook messenger (MA) allows the sender to delete messages from their own conversation history only [83].

The challenges C-1 and C-2 below relate to the missing realizations taking into account drivers for unshar-

ing (misconceptions) and desires (user view) reported by users, and inconsistencies in implementations.

Challenge C-1: Sophistication of Technical Mental Models

Users are known to formulate their own incorrect mental models when they are faced with a task to complete with their limited knowledge [103]. Given extreme fluctuation among users' technical understanding and variation among mechanisms' promised adversarial guarantees, the technical challenge here is to work within existing mental models to make actual functions clearer and communicate complex privacy issues to regular users in an intuitive and correct way. Since service providers make regular changes to their interfaces and features, it is important and challenging to simultaneously update the knowledge of the end-users, to minimize the risks associated with outdated mental models.

Challenge C-2: Borrowed Mental Models

Any given internet user is likely to be a member of multiple online services as well as platforms within those services. Some users naively transfer their mental models from one platform to another. These borrowed mental models considerably hinder the correct understanding of features and can expose users' data to unintended audiences. The technical challenge is the design of user interfaces, tutorials, and control setting pages that effectively convey the consequences of different actions taken by users on a specific platform.

5.4 Security and Trust

The process of making data available online typically involves multiple parties interacting with the data, such as friends or contacts in social networks, service providers, advertising companies aggregating individual user profiles for marketing purposes, or other third parties proactively crawling all available web contents. Such activities are usually carried out as soon as pieces of data appear online. In contrast, the common security model used in research proposals on automated data revocation is security against a retrospective adversary [20, 30, 70, 79, 93, 108]. Basically, this type of attacker is not interested in tampering with published data during its lifetime, but only after its expiration.

In the same way, a large body of proposals rely on distributed architectures to realize expiration since centralized service providers are considered untrusted [5, 20, 30, 38, 93, 108]. As a particular flaw, all types of entities are considered equally, and there are no differences between types of audiences. This is not in line with users publishing photos on platforms of large companies such as Facebook, who rather express fears such as specific groups of people (e. g. their parents or other family members) seeing their content and considering it inappropriate [77, 102].

Data deletion in artificial intelligence environments is a complicated task and poses a serious threat to longitudinal aspects of users' privacy. Legal scholars have questioned the legality of using of AI systems trained on deleted data in the context of the Right to be Forgotten [101]. In fact, model inversion and membership inference attacks have already demonstrated that the information used in training a model could be reconstructed afterwards by an adversary [100]. Our systematization of technical proposals identified that few of them enable control over the availability of data that is fed into machine learning models.

In light of the failure of the existing (theoretical) adversary models to capture the actual security requirements reported by users through drivers for unsharing (Fears) and desires (Content-based Audience), challenge D-1 brings attention to incorrect realizations of real-world threats. Challenge D-2 focuses on incomplete realizations of threat models that could provide guarantees against the emergent threat posed by machine learning algorithms.

Challenge D-1: Protection against real-world adversaries and threat scenarios

There is currently a gap between security under a given (theoretical) adversary model and actual security requirements in a real-world scenario. Instead of trying to provide security guarantees under unrealistic assumptions such as the presence of a solely retrospective adversary, solutions should incorporate effective mechanisms to reduce the unauthorized use of published data during all stages of their life-cycle (such as preventing screen-capturing in Snapchat [88]).

The key challenge here is to develop adversarial models that represent real-world threats, that incorporate users' fears regarding their privacy and unintended exposure in real data publishing scenarios and to secure data sharing mechanisms under these models.

Challenge D-2: Protection against machine learning algorithms

Prevalent use of artificially intelligent systems by service providers adds a new threat dimension to the exposure of users' data. When the data is used to aggregate statistics or to train machine-learning models, e. g., for image classification or recommender systems, the information that data carries will implicitly remain in the model, even when the original data and everything explicitly linked to it is deleted. This limits users' control over the availability of information encoded in their previously shared data. Similarly, AI-based recognition algorithms also hinder users' capacity to effectively manage the visibility of their data from service providers. Despite some promising initial work, such as the use of adversarial examples [58, 65], it remains a challenge to counter the capabilities of AI systems and provide security guarantees against their use.

6 Further Issues

In Section 5, we presented a set of succinct, yet unresolved challenges regarding longitudinal online privacy management. Inherently, not all challenges can be approached from a purely technical perspective, e. g., challenges relating to flawed mental models require more holistic approaches, centered around end-users' issues. Our systematization is supposed to trigger activities in both the technical and the human-factor research communities, as a number of identified issues can only be resolved conjointly, taking into account both technical and user perspectives. One key takeaway is that technical solutions point towards promising directions, such as proposals targeting to overcome purely time-based exposure control mechanisms. However, it is critical to match users' actual needs in order to find adoption and to serve users by providing tools that they need to appropriately control the exposure of their personal online data.

We finally discuss five open issues that did not make it to our list of challenges because these were not directly derived out of the systematizations or were not specifically limited to publicly shared data. However, these aspects still provide further insights to the community about the landscape of longitudinal privacy of publicly shared data.

Control over Inversely Private Information

Gurevich et al. term an item of personal information about you *inversely private* if some party has access to it, but you do not [34]. The situation described here elicits similar challenges as Data Co-ownership (cf. B-2) but is different in that users may not be aware of this particular information to exist. Daily interactions with various institutions ranging from toll roads operators to social networks generate vast amount of data about users. Processing users' private data and their pattern of interactions with the platform yields more inversely private data. In some cases, this private information held by companies can even contradict users' preferences in the current context. For example, a social network user can continue to receive ads related to a preference derived from one of their old posts despite choosing to limit its lifetime. It is not straightforward to realize technical proposals that can allow users to manage and erase vast amounts of inversely private data about them held by different entities. The information is typically used for gaining a competitive edge, which is one of the reasons why corporations have been denying the inverse privacy entitlement to their users [34]. Regulations on service providers' processing of data could prove helpful, but it is unclear if existing laws, such as GDPR, provide users the right to erasure of inversely private information.

Content Obfuscation versus Usability

While transformations targeting automated classifiers as means to solving the Security and Trust challenge (cf. D-2) may have only little impact on an image's appearance to humans, it also needs to be further investigated to what degree visible image perturbation is acceptable for users as a trade-off between privacy and vision comfort. There has been research on viewer satisfaction for blurring and pixelating photo scene elements that need to be protected [36, 42, 43], as well as on how the overall photo can be modified equally using aesthetic transforms to increase satisfaction [37].

Response to Privacy Paradox

While users claim to be very concerned about their privacy, they nevertheless undertake very little to protect their personal data. Recent research on the privacy paradox has revealed discrepancies between users' preferences and their actual behavior [8, 23, 98]. Various studies have reported instances of users not taking the logical step of limiting the disclosure in their so-

cial networks despite being aware of privacy concerns [49, 61, 107]. These results hint that User Awareness (cf. C-1) alone is not going to lead to widespread adoption of longitudinal privacy technologies. To bridge the gap between users' desires and mechanisms' functionalities, it is equally important to investigate and understand the causes and implications of the privacy paradox. Such an understanding will allow for design decisions that will increase the adoption of privacy-enhancing technologies.

Complications with Metadata Obfuscation

Correlation and analysis of individual metadata can allow to draw conclusions about a person. Information deduced from communication flows can create privacy concerns in the same way as sensitive information obtained from posted contents [33]. Depending on the extent of metadata generation, sensitive information may still be preserved even if there is a technically perfect revocation mechanism for the actual data. For example, Facebook includes a feature that automatically adds descriptive keywords to photos to assist visually impaired users in comprehending its contents. In the case of photos of human subjects, their faces are detected, and users are suggested to enter the name of the person. While such features can be easily observed in the application interface, it remains unclear what types of additional data collection invisibly run in the background. One approach to counteract potential privacy threats by metadata aggregation and its residuals can be achieved by preventing metadata from being generated in the first place. This could be realized by applying image perturbation techniques to hamper metadata generation. While this strategy renders targeted classifiers unable to correctly assess image content, users would still be able to see the content. Related approaches have been developed with a different mindset, i. e., adversarial perturbations, e. g., used to interfere with traffic sign recognition used by self-driving cars [85]. More universal approaches to falsely classify images have also been demonstrated [58]. However, such protective mechanisms come along with new potential conflicts. Whenever the use of such a perturbation mechanism is transparent, or its presence becomes apparent, service providers (if considered in an adversarial setting) can adapt their classification techniques to circumvent the protection. This game-theoretic consideration, already laid out by Oh et al. [65], is yet interesting to be investigated when developing even more sophisticated protection mechanisms.

Practicality of Referencing Data

The current way to distribute data is to upload it to online platforms and copy-share it through various channels in order to make it available for different types of audiences [87]. In an entirely different approach, users could have only one instance of all their data hosted in a single location of their choice, providing them the individual level of privacy they desire. Instead of creating multiple copies of data and uploading them to different platforms, those services would be allowed or licensed to reference the data, without actually obtaining a copy or possessing them. Such a solution will enable tracking of all interactions with data objects and could facilitate the realization of challenging Expiration Conditions (cf. A-3). Bishop et al. [15] came up with ideas in a similar direction when discussing dissemination control as a means to manage online privacy.

The approach is not without challenges since interactions with the data entail modifications of the data itself. For example, multiple instant messaging platforms provide popular features enabling users to add text and drawings to the images sent in the chats. In such settings, each transformed output of the original data needs to be tracked in order to uphold the integrity of data provenance and ensure effective control over dissemination of the data.

In the light of applying such a scenario equally to end-users' personal data, one must also discuss if large companies such as Google or Facebook would already consider themselves such hosting platforms, providing almost every kind of service for one's online actions from a single source. It is unclear how the data object's single source of origin might impact its availability since providers would need to be willing to adapt their practices, and interfaces, to facilitate sharing of data hosted on their competitors' platforms.

7 Related Surveys

Earlier surveys focused on deletion from the standpoint of a specific platform. In [78], authors classify secure deletion approaches into layers through which they access the physical medium. The focus of their work is on systematizing approaches that remove data objects from physical media. They emphasize understanding the properties of the environment before choosing a suitable approach for data deletion. Unger et al. systematized secure messaging solutions and proposed an evaluation framework for their security, usability, and

ease-of-adoption properties [99]. Their approach yielded three sets of challenges in the form of trust establishment, conversation security, and transport privacy.

In [75], authors survey requirements and challenges of assured deletion in the cloud considering two different adversarial models in terms of cloud provider honesty. It is stated that future work should focus on finding better ways of tracking the location of data to help providers make better decisions on which methods of assured deletion to apply. Geambasu et al. focus on identifying approaches to practical self-destructing data systems that secure sensitive data from disclosure in the highly mobile, social-networked world [31]. They make a case for a framework that combines multiple key-storage mechanisms into a single self-destructing data system.

In light of the effect of the GDPR [27] across the European Union, Politou et al. reviewed controversies regarding definitions of consent revocation and the right to be forgotten [71]. They evaluated the preparedness of existing technical infrastructure to incorporate these new requirements and report the feasibility of doing so provided clear implementation guidelines are in place.

8 Conclusion

We provided the first systematization to capture users' interactions related to longitudinal privacy management on existing platforms, as well as the landscape of diverse technical proposals dealing with the availability of online data. Our broad approach afforded us the ability to contrast end-users' desires and mental models against the technical proposals' use cases and adversarial assumptions. This enabled us to uncover open challenges and identify interesting problems where effective solutions have not yet been realized. By pointing the research community's direction towards these challenges, we hope this paper serves as an inspiration and a basis for the development of longitudinal privacy-enhancing solutions that will assist millions of end-users with managing the availability of their publicly-shared data.

Acknowledgments

We would like to thank our shepherd Michelle L. Mazurek and the anonymous reviewers for their helpful comments. This research was supported by the BMBF project 'InStruct' under grant 16KIS0581 (Germany) and by the Center for Cyber Security at New York University Abu Dhabi (NYUAD).

References

- [1] A. Abouzied and J. Chen, “Harnessing Data Loss With Forgetful Data Structures,” in *ACM Symposium on Cloud Computing*, ser. SoCC '15. Kohala Coast, HI, USA: ACM, Aug. 2015, pp. 168–173.
- [2] S. Ahern, D. Eckles, N. S. Good, S. King, M. Naaman, and R. Nair, “Over-exposed? privacy patterns and considerations in online and mobile photo sharing,” in *CHI Conference on Human Factors in Computing Systems*, ser. CHI '07. San Jose, CA, USA: ACM, Apr. 2007, pp. 357–366.
- [3] H. Almuhammedi, S. Wilson, B. Liu, N. Sadeh, and A. Acquisti, “Tweets are Forever: A Large-scale Quantitative Analysis of Deleted Tweets,” in *ACM Conference on Computer supported cooperative work*, ser. CSCW '13. San Antonio, TX, USA: ACM, Feb. 2013, pp. 897–908.
- [4] A. Alqhatani and H. R. Lipford, ““There is nothing that I need to keep secret”: Sharing Practices and Concerns of Wearable Fitness Data,” in *USENIX Symposium on Usable Privacy and Security*, ser. SOUPS '19. Santa Clara, CA, USA: USENIX Association, Aug. 2019.
- [5] G. Amjad, M. S. Mirza, and C. Pöpper, “Forgetting with Puzzles: Using Cryptographic Puzzles to support Digital Forgetting,” in *ACM Conference on Data and Application Security and Privacy*, ser. CODASPY '18. Tempe, AZ, USA: ACM, Mar. 2018, pp. 342–353.
- [6] O. Ayalon and E. Toch, “Retrospective privacy: managing longitudinal privacy in online social networks,” in *Symposium on Usable Privacy and Security*, ser. SOUPS '13. Newcastle, UK: USENIX Association, Jul. 2013.
- [7] E. Bacis, S. De Capitani di Vimercati, S. Foresti, S. Paraboschi, M. Rosa, and P. Samarati, “Mix&Slice: Efficient Access Revocation in the Cloud,” in *Conference on Computer and Communications Security*, ser. CCS '16. Vienna, Austria: ACM, Oct. 2016, p. 217–228.
- [8] S. Barth and M. D. de Jong, “The Privacy Paradox – Investigating Discrepancies Between Expressed Privacy Concerns and Actual Online Behavior – A Systematic Literature Review,” *Telematics and Informatics*, vol. 34, no. 7, pp. 1038–1058, 2017.
- [9] L. Bauer, L. F. Cranor, S. Komanduri, M. L. Mazurek, M. K. Reiter, M. Sleeper, and B. Ur, “The Post Anachronism: The Temporal Dimension of Facebook Privacy,” in *Workshop on Privacy in the Electronic Society*, ser. WPES '13. Berlin, Germany: ACM, Nov. 2013, pp. 1–12.
- [10] J. B. Bayer, N. B. Ellison, S. Y. Schoenebeck, and E. B. Falk, “Sharing the Small Moments: Ephemeral Social Interaction on Snapchat,” *Information, Communication & Society*, vol. 19, no. 7, pp. 956–977, Apr. 2016.
- [11] F. Beato, M. Kohlweiss, and K. Wouters, “Scramble! Your Social Network Data,” in *Privacy Enhancing Technologies Symposium*, ser. PETS '11. Waterloo, ON, Canada: Springer, Jul. 2011, pp. 211–225.
- [12] M. S. Bernstein, E. Bakshy, M. Burke, and B. Karrer, “Quantifying the Invisible Audience in Social Networks,” in *CHI Conference on Human Factors in Computing Systems*, ser. CHI '13. Paris, France: ACM, Apr. 2013, pp. 21–30.
- [13] A. Besmer and H. Richter Lipford, “Moving Beyond Untagging: Photo Privacy in a Tagged World,” in *CHI Conference on Human Factors in Computing Systems*, ser. CHI '10. Atlanta, GA, USA: ACM, Apr. 2010, pp. 1563–1572.
- [14] P. Bhattacharya and N. Ganguly, “Characterizing Deleted Tweets and Their Authors,” in *AAAI Conference on Weblogs and Social Media*, ser. ICWSM '16. Cologne, Germany: AAAI, May 2016.
- [15] M. Bishop, E. R. Butler, K. Butler, C. Gates, and S. Greenspan, “Forgive and Forget: Return to Obscurity,” in *New Security Paradigms Workshop*, ser. NSPW '13. Banff, Alberta, Canada: ACM, Sep. 2013, pp. 1–10.
- [16] B. M. Bowen, S. Hershkop, A. D. Keromytis, and S. J. Stolfo, “Baiting Inside Attackers Using Decoy Documents,” in *International Conference on Security and Privacy in Communication Systems*, ser. SecureComm '09. Athens, Greece: Springer, Sep. 2009, pp. 51–70.
- [17] T. F. Brady, T. Konkle, and G. A. Alvarez, “A Review of Visual Memory Capacity: Beyond Individual Items and Toward Structured Representations,” *Journal of Vision*, vol. 11, no. 5, pp. 1–34, 05 2011.
- [18] Y. Cao and J. Yang, “Towards Making Systems Forget With Machine Unlearning,” in *Symposium on Security and Privacy*, ser. S&P '15. San Jose, CA, USA: IEEE, May 2015, pp. 463–480.
- [19] B. Carminati and E. Ferrari, “Collaborative access control in online social networks,” in *Conference on Collaborative Computing: Networking, Applications and Worksharing*, ser. CollaborateCom '11. Orlando, FL, USA: IEEE, Oct. 2011, pp. 231–240.
- [20] C. Castelluccia, E. De Cristofaro, A. Francillon, and M.-A. Kaafar, “EphPub: Toward Robust Ephemeral Publishing,” in *IEEE Conference on Network Protocols*, ser. ICNP '11. Vancouver, BC, Canada: IEEE, Oct. 2011, pp. 165–175.
- [21] I. Cervesato, “The Dolev-Yao Intruder is the Most Powerful Attacker,” in *Annual Symposium on Logic in Computer Science*, ser. LICS '01. Boston, MA, USA: IEEE, Jun. 2001.
- [22] J. W. Clark, P. Snyder, D. McCoy, and C. Kanich, ““I Saw Images I Didn’t Even Know I Had”: Understanding User Perceptions of Cloud Storage Privacy,” in *CHI Conference on Human Factors in Computing Systems*, ser. CHI '15. Seoul, Republic of Korea: ACM, Apr. 2015, pp. 1641–1644.
- [23] K. P. Coopamootoo and T. Groß, “Why Privacy is All But Forgotten: An Empirical Study of Privacy & Sharing Attitude,” in *Privacy Enhancing Technologies Symposium*, ser. PETS '17. Minneapolis, MN, USA: Sciencdo, Jul. 2017, pp. 97–118.
- [24] E. De Cristofaro, C. Soriente, G. Tsudik, and A. Williams, “Hummingbird: Privacy at the time of twitter,” in *Symposium on Security and Privacy*, ser. S&P '12. San Francisco, CA, USA: IEEE, May 2012, pp. 285–299.
- [25] A. Dhir, P. Kaur, K. Lonka, and M. Nieminen, “Why Do Adolescents Untag Photos on Facebook?” *Computers in Human Behavior*, vol. 55, no. PB, pp. 1106–1115, Feb. 2016.
- [26] S. Egelman, A. Oates, and S. Krishnamurthi, “Oops, I Did it Again: Mitigating Repeated Access Control Errors on

- Facebook,” in *CHI Conference on Human Factors in Computing Systems*, ser. CHI '11. Vancouver, BC, Canada: ACM, May 2011, pp. 2295–2304.
- [27] European Parliament, “Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation),” 2016.
- [28] European Union, “Factsheet on the “Right to be Forgotten” Ruling (C-131/12),” May 2014, http://ec.europa.eu/justice/data-protection/files/factsheets/factsheet_data_protection_en.pdf, as of October 20, 2020.
- [29] C. Fiesler, M. Dye, J. L. Feuston, C. Hiruncharoenvate, C. J. Hutto, S. Morrison, P. Khanipour Roshan, U. Pavalanathan, A. S. Bruckman, M. De Choudhury, and E. Gilbert, “What (or Who) is Public? Privacy Settings and Social Media Content Sharing,” in *Conference on Computer Supported Cooperative Work*, ser. CSCW '17. Portland, OR, USA: ACM, Feb. 2017, pp. 567–580.
- [30] R. Geambasu, T. Kohno, A. A. Levy, and H. M. Levy, “Vanish: Increasing Data Privacy with Self-Destructing Data,” in *USENIX Security Symposium*, ser. USENIX '09. Montreal, QC, Canada: USENIX Association, Aug. 2009, pp. 299–316.
- [31] R. Geambasu, T. Kohno, A. Krishnamurthy, A. Levy, H. M. Levy, P. Gardner, and V. Moscaritolo, “New Directions for Self-Destructing Data,” University of Washington, Tech. Rep. Tech. Rep. UW-CSE-11-08-01, 2011.
- [32] A. Ginart, M. Guan, G. Valiant, and J. Y. Zou, “Making ai forget you: Data deletion in machine learning,” in *Advances in Neural Information Processing Systems*, 2019, pp. 3513–3526.
- [33] B. Greschbach, G. Kreitz, and S. Buchegger, “The Devil is in the Metadata — New Privacy Challenges in Decentralised Online Social Networks,” in *Conference on Pervasive Computing and Communications Workshops*, ser. PerCOM '12. Lugano, Switzerland: IEEE, Mar. 2012, pp. 333–339.
- [34] Y. Gurevich, E. Hudis, and J. M. Wing, “Inverse Privacy,” *Communications of the ACM*, vol. 59, no. 7, pp. 38–42, Jun. 2016.
- [35] H. Habib, N. Shah, and R. Vaish, “Impact of Contextual Factors on Snapchat Public Sharing,” in *CHI Conference on Human Factors in Computing Systems*, ser. CHI '19. Glasgow, UK: ACM, May 2019.
- [36] R. Hasan, E. Hassan, Y. Li, K. Caine, D. J. Crandall, R. Hoyle, and A. Kapadia, “Viewer Experience of Obscuring Scene Elements in Photos to Enhance Privacy,” in *CHI Conference on Human Factors in Computing Systems*, ser. CHI '18. Montreal, QC, Canada: ACM, Apr. 2018.
- [37] R. Hasan, Y. Li, E. Hassan, K. Caine, D. J. Crandall, R. Hoyle, and A. Kapadia, “Can Privacy Be Satisfying? On Improving Viewer Satisfaction for Privacy-Enhanced Photos Using Aesthetic Transforms,” in *CHI Conference on Human Factors in Computing Systems*, ser. CHI '19. Glasgow, UK: ACM, May 2019.
- [38] P. Ilia, B. Carminati, E. Ferrari, P. Fragopoulou, and S. Ioannidis, “SAMPAC: Socially-Aware Collaborative Multi-Party Access Control,” in *Conference on Data and Application Security and Privacy*, ser. CODASPY '17. Scottsdale, AZ, USA: ACM, Mar. 2017, p. 71–82.
- [39] M. Johnson, S. Egelman, and S. M. Bellovin, “Facebook and Privacy: It’s Complicated,” in *Symposium on Usable Privacy and Security*, ser. SOUPS '12. Washington, D.C., USA: ACM, Jul. 2012.
- [40] R. Kang, L. Dabbish, N. Fruchter, and S. Kiesler, ““My Data Just Goes Everywhere.” User Mental Models of the Internet and Implications for Privacy and Security,” in *Symposium On Usable Privacy and Security*, ser. SOUPS '15. Ottawa, ON, Canada: USENIX Association, Jul. 2015, pp. 39–52.
- [41] M. T. Khan, M. Hyun, C. Kanich, and B. Ur, “Forgotten But Not Gone: Identifying the Need for Longitudinal Data Management in Cloud Storage,” in *CHI Conference on Human Factors in Computing Systems*, ser. CHI '18. Montreal, QC, Canada: ACM, Apr. 2018.
- [42] Y. Li, N. Vishwamitra, B. P. Knijnenburg, H. Hu, and K. Caine, “Blur vs. Block: Investigating the Effectiveness of Privacy-Enhancing Obfuscation for Images,” in *IEEE Conference on Computer Vision and Pattern Recognition Workshops*, ser. CVPRW '17. Honolulu, HI, USA: IEEE, Jul. 2017, pp. 1343–1351.
- [43] —, “Effectiveness and Users’ Experience of Obfuscation as a Privacy-Enhancing Technology for Sharing Photos,” *Proceedings of the ACM on Human-Computer Interaction*, vol. 1, no. 67, pp. 1–24, 2017.
- [44] H. R. Lipford, A. Besmer, and J. Watson, “Understanding Privacy Settings in Facebook with an Audience View,” in *USENIX Workshop on Usability, Psychology, and Security*, ser. UPSEC '08. San Francisco, CA, USA: USENIX Association, Apr. 2008, pp. 1–8.
- [45] Y. Liu, K. P. Gummadi, B. Krishnamurthy, and A. Mislove, “Analyzing Facebook Privacy Settings: User Expectations vs. Reality,” in *Internet Measurement Conference*, ser. IMC '11. Berlin, Germany: ACM, Nov. 2011, pp. 61–70.
- [46] Y. Liu, C. Kliman-Silver, and A. Mislove, “The Tweets They Are a-Changin’: Evolution of Twitter Users and Behavior,” in *AAAI Conference on Weblogs and Social Media*, ser. ICWSM '14. Ann Arbor, MI, USA: AAAI, Jun. 2014.
- [47] W. Luo, Q. Xie, and U. Hengartner, “FaceCloak: An Architecture for User Privacy on Social Networking Sites,” in *Conference on Computational Science and Engineering*, ser. CSE '09. Vancouver, BC, Canada: IEEE, Aug. 2009, pp. 26–33.
- [48] M. Madejski, M. Johnson, and S. M. Bellovin, “A Study of Privacy Settings Errors in an Online Social Network,” in *IEEE International Conference on Pervasive Computing and Communications Workshops*, ser. PerCOM '12. Lugano, Switzerland: IEEE, Mar. 2012, pp. 340–345.
- [49] M. J. Mainier, R. Morris, and M. O. Louch, “Social Networks and the Privacy Paradox: A Research Framework,” *Issues in Information Systems*, vol. 11, no. 1, pp. 513–517, 2010.
- [50] V. Mayer-Schönberger, *Delete: The Virtue of Forgetting in the Digital Age*. Princeton, New Jersey, USA: Princeton University Press, Jul. 2011.
- [51] M. L. Mazurek, J. Arseneault, J. Breece et al., “Access Control for Home Data Sharing: Attitudes, Needs and Practices,” in *CHI Conference on Human Factors in Computing*

- Systems*, ser. CHI '10. Atlanta, GA, USA: ACM, Apr. 2010, pp. 645–654.
- [52] M. Minaei, M. Mondal, P. Loiseau, K. Gummadi, and A. Kate, “Lethe: Conceal Content Deletion from Persistent Observers,” in *Privacy Enhancing Technologies Symposium*, ser. PETS '19. Stockholm, Sweden: Sciendo, Jul. 2019, pp. 206–226.
- [53] R. E. Mohamed and S. Chiasson, “Online Privacy and Aging of Digital Artifacts,” in *USENIX Symposium on Usable Privacy and Security*, ser. SOUPS '18. Baltimore, MD, USA: USENIX Association, Aug. 2018.
- [54] M. Mondal, M. Johnnatan, S. Ghosh, K. P. Gummadi, and A. Kate, “Forgetting in Social Media: Understanding and Controlling Longitudinal Exposure of Socially Shared Data,” in *Symposium on Usable Privacy and Security*, ser. SOUPS '16. Denver, CO, USA: USENIX Association, Jun. 2016, pp. 287–299.
- [55] M. Mondal, Y. Liu, B. Viswanath, K. P. Gummadi, and A. Mislove, “Understanding and Specifying Social Access Control Lists,” in *Symposium on Usable Privacy and Security*, ser. SOUPS '14. Menlo Park, California, USA: USENIX Association, Jul. 2014.
- [56] M. Mondal, J. Messias, S. Ghosh, K. P. Gummadi, and A. Kate, “Longitudinal Privacy Management in Social Media: The Need for Better Controls,” *IEEE Internet Computing*, vol. 21, no. 3, pp. 48–55, May 2017.
- [57] M. Mondal, G. S. Yilmaz, N. Hirsch, M. T. Khan, M. Tang, C. Tran, C. Kanich, B. Ur, and E. Zheleva, “Moving Beyond Set-It-And-Forget-It Privacy Settings on Social Media,” in *ACM Conference on Computer and Communications Security*, ser. CCS '19. London, UK: ACM, Nov. 2019, pp. 991–1008.
- [58] S.-M. Moosavi-Dezfooli, A. Fawzi, O. Fawzi, and P. Frossard, “Universal Adversarial Perturbations,” in *IEEE Conference on Computer Vision and Pattern Recognition*, ser. CVPR '17. Honolulu, HI, USA: IEEE, Jul. 2017, pp. 86–94.
- [59] G. E. Müller and A. Pilzecker, *Experimentelle Beiträge zur Lehre vom Gedächtniss*. JA Barth, 1900, vol. 1.
- [60] A. Murillo, A. Kramm, S. Schnorf, and A. De Luca, ““If I press delete, it's gone” - User Understanding of Online Data Deletion and Expiration,” in *USENIX Symposium on Usable Privacy and Security*, ser. SOUPS '18. Baltimore, MD, USA: USENIX Association, Aug. 2018, pp. 329–339.
- [61] J. Nagy and P. Pecho, “Social Networks Security,” in *Conference on Emerging Security Information, Systems and Technologies*, ser. SECUREWARE '09. Athens, Greece: IARIA, Jun. 2009, pp. 321–325.
- [62] M. Netter, M. Riesner, M. Weber, and G. Pernul, “Privacy Settings in Online Social Networks—Preferences, Perception, and Reality,” in *Hawaii International Conference on System Sciences*, ser. HICSS '13. Wailea, HI, USA: IEEE, Jan. 2013, pp. 3219–3228.
- [63] C. Niederée, “Learning from Human Memory: Managed Forgetting and Contextualized Remembering for Digital Memories,” in *Conference on Theory and Practice of Digital Libraries*, ser. TPD L '15. Poznań, Poland: Springer, Sep. 2015, pp. 1–6.
- [64] C. Niederée, N. Kanhabua, F. Gallo, and R. H. Logie, “Forgetful Digital Memory: Towards Brain-inspired Long-term Data and Information Management,” *ACM SIGMOD Record*, vol. 44, no. 2, pp. 41–46, 2015.
- [65] S. J. Oh, M. Fritz, and B. Schiele, “Adversarial Image Perturbation for Privacy Protection: A Game Theory Perspective,” in *IEEE International Conference on Computer Vision*, ser. ICCV '17. Venice, Italy: IEEE, Oct. 2017, pp. 1482–1491.
- [66] A.-M. Olteanu, K. Huguenin, I. Dacosta, and J.-P. Hubaux, “Consensual and Privacy-preserving Sharing of Multi-subject and Interdependent Data,” in *Symposium on Network and Distributed System Security*, ser. NDSS '18. San Diego, CA, USA: Internet Society, Feb. 2018, pp. 1–16.
- [67] Open Whisper Systems, “Signal,” May 2010, <https://signal.org/>, as of October 20, 2020.
- [68] A. Oshrat and T. Eran, “Not Even Past: Information Aging and Temporal Privacy in Online Social Networks,” vol. 32, no. 2. Taylor & Francis, 2017, pp. 73–102.
- [69] Y. J. Park, “Digital Literacy and Privacy Behavior Online,” *Communication Research*, vol. 40, no. 2, pp. 215–236, 2013.
- [70] R. Perlman, “The Ephemerizer: Making Data Disappear,” Sun Microsystems Laboratories, Inc., Mountain View, CA, USA, Tech. Rep. SMLI TR-2005-140, Feb. 2005.
- [71] E. Politou, E. Alepis, and C. Patsakis, “Forgetting Personal Data and Revoking Consent Under the GDPR: Challenges and Proposed Solutions,” *Journal of Cybersecurity*, vol. 4, no. 1, pp. 1–20, Mar. 2018.
- [72] C. Pöpper, D. Basin, S. Capkun, and C. Cremers, “Keeping Data Secret under Full Compromise using Porter Devices,” in *Annual Computer Security Applications Conference*, ser. ACSAC '10. Orlando, FL, USA: ACM, Dec. 2010, pp. 241–250.
- [73] Proton Technologies AG, “ProtonMail,” May 2014, <https://protonmail.com/>, as of October 20, 2020.
- [74] S. Rajtmajer, A. Squicciarini, J. M. Such, J. Semonsen, and A. Belmonte, “An Ultimatum Game Model for the Evolution of Privacy in Jointly Managed Content,” in *Conference on Decision and Game Theory for Security*, ser. GameSec '07. Vienna, Austria: Springer, Oct. 2017, pp. 112–130.
- [75] K. M. Ramokapane, A. Rashid, and J. M. Such, “Assured Deletion in the Cloud: Requirements, Challenges and Future Directions,” in *ACM Cloud Computing Security Workshop*, ser. CCSW '16. Vienna, Austria: ACM, Oct. 2016, pp. 97–108.
- [76] —, ““I Feel Stupid I Can't Delete...”: A Study of Users' Cloud Deletion Practices and Coping Strategies,” in *Symposium on Usable Privacy and Security*, ser. SOUPS '17. Santa Clara, CA: USENIX Association, 2017, pp. 241–256.
- [77] Y. Rashidi, T. Ahmed, F. Patel, E. Fath, A. Kapadia, C. Nippert-Eng, and N. M. Su, ““You Don't Want to be the Next Meme”: College Students' Workarounds to Manage Privacy in the Era of Pervasive Photography,” in *Symposium on Usable Privacy and Security*, ser. SOUPS '18. Baltimore, MD, USA: USENIX Association, Aug. 2018, pp. 143–157.
- [78] J. Reardon, D. Basin, and S. Capkun, “SoK: Secure Data Deletion,” in *IEEE Symposium on Security and Privacy*, ser. SP '13. San Francisco, California, USA: IEEE, May 2013, pp. 301–315.

- [79] S. Reimann and M. Dürmuth, “Timed Revocation of User Data: Long Expiration Times from Existing Infrastructure,” in *ACM Workshop on Privacy in the Electronic Society*, ser. WPES '12. Raleigh, NC, USA: ACM, Oct. 2012, pp. 65–74.
- [80] B. Reynolds, J. Venkatanathan, J. Gonçalves, and V. Kostakos, “Sharing Ephemeral Information in Online Social Networks: Privacy Perceptions and Behaviours,” in *IFIP Conference on Human-Computer Interaction*, ser. INTERACT '11. Lisbon, Portugal: IFIP, 2011, pp. 204–215.
- [81] A. Schlesinger, E. Chandrasekharan, C. A. Masden, A. S. Bruckman, W. K. Edwards, and R. E. Grinter, “Situated Anonymity: Impacts of Anonymity, Ephemerality, and Hyper-locality on Social Media,” in *CHI Conference on Human Factors in Computing Systems*, ser. CHI '17. Denver, CO, USA: ACM, May 2017, pp. 6912–6924.
- [82] T. Schnitzler, M. Dürmuth, and C. Pöpper, “Towards Contractual Agreements for Revocation of Online Data,” in *ICT Systems Security and Privacy Protection*, ser. IFIP SEC '19. Lisbon, Portugal: Springer, Jun. 2019.
- [83] T. Schnitzler, C. Utz, F. Farke, C. Pöpper, and M. Dürmuth, “User Perception and Expectations on Deleting Instant Messages – or – “What Happens If I Press This Button?”,” in *European Workshop on Usable Security*, ser. EuroUSEC '18. London, UK: Internet Society, Apr. 2018, pp. 1–9.
- [84] E. Shein, “Ephemeral Data,” *Communications of the ACM*, vol. 56, no. 9, pp. 20–22, Sep. 2013.
- [85] C. Sitawarin, A. N. Bhagoji, A. Mosenia, P. Mittal, and M. Chiang, “Rogue Signs: Deceiving Traffic Sign Recognition with Malicious Ads and Logos,” in *IEEE Deep Learning and Security Workshop*, ser. DLS '18. San Francisco, CA, USA: IEEE, May 2018.
- [86] M. Sleeper, J. Cranshaw, P. G. Kelley, B. Ur, A. Acquisti, L. F. Cranor, and N. Sadeh, ““I Read My Twitter the Next Morning and Was Astonished”: A Conversational Perspective on Twitter Regrets,” in *CHI Conference on Human Factors in Computing Systems*, ser. CHI '13. Paris, France: ACM, 2013, pp. 3277–3286.
- [87] M. Sleeper, W. Melicher, H. Habib, L. Bauer, L. F. Cranor, and M. L. Mazurek, “Sharing Personal Content Online: Exploring Channel Choice and Multi-Channel Behaviors,” in *CHI Conference on Human Factors in Computing Systems*, ser. CHI '16. Santa Clara, CA, USA: ACM, May 2016, pp. 101–112.
- [88] Snap Inc., “Snapchat,” Sep. 2011, <https://www.snapchat.com/>, as of October 20, 2020.
- [89] P. Snyder and C. Kanich, “Cloudsweeper: Enabling Data-centric Document Management for Secure Cloud Archives,” in *ACM Cloud Computing Security Workshop*, ser. CCSW '13. Berlin, Germany: ACM, Nov. 2013, pp. 47–54.
- [90] D. J. Solove, “A Taxonomy of Privacy,” *University of Pennsylvania Law Review*, vol. 154, no. 3, pp. 477–560, Januar 2006.
- [91] A. C. Squicciarini, M. Shehab, and F. Paci, “Collective Privacy Management in Social Networks,” in *Conference on World Wide Web*, ser. WWW '09. Madrid, Spain: ACM, Apr. 2009, pp. 521–530.
- [92] Statista Inc., “Distribution of Facebook users worldwide as of January 2020, by age and gender,” Feb. 2020, <https://www.statista.com/statistics/376128/facebook-global-user-age-distribution/>, as of October 20, 2020.
- [93] K. Stokes and N. Carlsson, “A Peer-to-Peer Agent Community for Digital Oblivion in Online Social Networks,” in *IEEE Conference on Privacy, Security and Trust*, ser. PST '13. Tarragona, Spain: IEEE, Jul. 2013, pp. 103–110.
- [94] F. D. Stutzman, R. Gross, and A. Acquisti, “Silent Listeners: The Evolution of Privacy and Disclosure on Facebook,” *IEEE Security and Privacy Magazine*, vol. 4, no. 2, pp. 7–41, 2013.
- [95] J. M. Such and N. Criado, “Resolving Multi-Party Privacy Conflicts in Social Media,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 28, no. 7, pp. 1851–1863, Jul. 2016.
- [96] Telegram Messenger LLP, “Telegram,” Aug. 2013, <https://telegram.org/>, as of October 20, 2020.
- [97] K. Thomas, C. Grier, and D. M. Nicol, “unFriendly: Multi-party Privacy Risks in Social Networks,” in *Privacy Enhancing Technologies Symposium*, ser. PETS '10. Berlin, Germany: Springer, 2010, pp. 236–252.
- [98] J. Y. Tsai, S. Egelman, L. F. Cranor, and A. Acquisti, “The Effect of Online Privacy Information on Purchasing Behavior: An Experimental Study,” *Information Systems Research*, vol. 22, no. 2, pp. 254–268, 2011.
- [99] N. Unger, S. Dechand, J. Bonneau, S. Fahl, H. Perl, I. Goldberg, and M. Smith, “SoK: Secure Messaging,” in *IEEE Symposium on Security and Privacy*, ser. SP '15. San Jose, CA, USA: IEEE, May 2015, pp. 232–249.
- [100] M. Veale, R. Binns, and L. Edwards, “Algorithms that remember: model inversion attacks and data protection law,” *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 376, no. 2133, p. 20180083, 2018.
- [101] E. F. Villaronga, P. Kieseberg, and T. Li, “Humans forget, machines remember: Artificial intelligence and the right to be forgotten,” *Computer Law & Security Review*, vol. 34, no. 2, pp. 304–313, 2018.
- [102] Y. Wang, G. Norcie, S. Komanduri, A. Acquisti, P. G. Leon, and L. F. Cranor, ““I Regretted the Minute I Pressed Share”: A Qualitative Study of Regrets on Facebook,” in *Symposium on Usable Privacy and Security*, ser. SOUPS '11. Pittsburgh, PA, USA: ACM, 2011, pp. 1–16.
- [103] R. Wash, “Folk Models of Home Computer Security,” in *Symposium on Usable Privacy and Security*, ser. SOUPS '10. Redmond, WA, USA: ACM, Jul. 2010.
- [104] G. Wegberg, H. Ritzdorf, and S. Capkun, “Multi-User Secure Deletion on Agnostic Cloud Storage,” ETH Zurich, Zurich, Switzerland, Tech. Rep., Oct. 2017.
- [105] R. Wishart, D. Corapi, S. Marinovic, and M. Sloman, “Collaborative Privacy Policy Authoring in a Social Networking Context,” in *Symposium on Policies for Distributed Systems and Networks*, ser. POLICY '10. Washington D. C., USA: IEEE, Jul. 2010, pp. 1–8.
- [106] Y. Xue, K. Xue, N. Gai, J. Hong, D. S. L. Wei, and P. Hong, “An Attribute-Based Controlled Collaborative Access Control Scheme for Public Cloud Storage,” *IEEE Transactions on Information Forensics and Security*, vol. 14, no. 11, pp. 2927–2942, Nov. 2019.

- [107] C. W. Yoo, H. J. Ahn, and H. R. Rao, "An Exploration of the Impact of Information Privacy Invasion," in *International Conference on Information Systems*, ser. ICIS '12. Orlando, FL, USA: AIS, Dec. 2012.
- [108] A. Zarras, K. Kohls, M. Dürmuth, and C. Pöpper, "Neuralyzer: Flexible Expiration Times for the Revocation of Online Data," in *ACM Conference on Data and Application Security and Privacy*, ser. CODASPY '16. New Orleans, Louisiana, USA: ACM, Mar. 2016, pp. 14–25.
- [109] L. Zhou, W. Wang, and K. Chen, "Tweet Properly: Analyzing Deleted Tweets to Understand and Identify Regrettable Ones," in *The World Wide Web Conference*, ser. WWW '16. Montreal, QC, Canada: ACM, Apr. 2016, pp. 603–612.