Offlinetags – A Novel Privacy Approach to Online Photo Sharing

Frank Pallas
Technical University of Berlin,
Dep. Computers & Society
Marchstr. 23
10587 Berlin, Germany
frank.pallas@tu-berlin.de

Max-Robert Ulbricht
Technical University of Berlin,
Dep. Computers & Society
Marchstr. 23
10587 Berlin, Germany
max-robert.ulbricht@tu-berlin.de

Lorena Jaume-Palasí
Ludwig-Maximilians-University
Munich
Chair for Philosophy IV
Geschwister-Scholl-Platz 1
80539 Munich, Germany
lorena.jaume-palasi@gsi.uni-muenchen.de

Ulrike Höppner
Internet & Society Collaboratory
Sophienstr. 24
10178 Berlin, Germany
ulrike@collaboratory.de

CHI 2014, Apr 26 - May 01 2014, Toronto, ON, Canada
ACM 978-1-4503-2474-8/14/04.
http://dx.doi.org/10.1145/2559206.2581195

Abstract

In this paper, we describe a novel approach to the privacy problem that photos showing persons are often “meddle-shared” by others online. We introduce a set of four elementary privacy preferences a photo subject can have. These preferences are represented by corresponding symbols – “Offlinetags” – which can be worn in the form of stickers or badges and which are designed to be easily recognizable by humans and algorithms. Especially for the context of public events, these Offlinetags can serve as a basis for novel practices of photo sharing that respect the photo subjects’ privacy preferences.

Author Keywords

Privacy; photo sharing; social networks; offlinetags

ACM Classification Keywords

K.4.1. Computers and Society, Public Policy Issues: Privacy
H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

Introduction

Online photo sharing has been a major source of social conflict since the broad adoption of online social networks. Whether in the context of private activities, work life or, in particular, in the context of public events: Nowadays, it is hardly possible to avoid being photographed by others and these photos being shared online. Such sharing of photos without the consent or even the knowledge of the shown persons will herein be called "meddle-sharing", whereas the shown person is referred to as the "photo subject".

Everyday examples of such meddle-sharing include photo subjects being shown as participants of demonstrations (e.g. at a gay parade), as attendees of specific events (e.g. a convention of a political party or a conference), or simply as having been at a certain place at a given time. Depending on the situation shown on the photo and the party taking notice of it, such photos can disclose information about the photo subject that she would otherwise not have revealed to the noticing person. Following a common understanding of privacy as primarily being about "rights to control your public image" [5], meddle-sharing can thus constitute serious privacy infringements for the photo subject. In the following, we therefore present a novel approach for influencing the taking, sharing and further handling of photos of oneself.

Generally speaking, our approach[1] is based on a well-defined set of four symbols that, in the form of stickers, buttons, badges, etc., can be attached to the clothes and represent the wearer's preferences on the desired handling of photos taken of her. The symbols – which we call "Offlinetags" – are designed to be easily understandable to humans and recognition-friendly for computer-vision algorithms, thereby enabling social consideration as well as technological analysis and processing of these preferences.

Related Work

Many approaches for controlling the visibility and handling of personal content like photos have been suggested in the past [3] and are now available in most online social networks and other content sharing platforms. Current scientific discussions go even beyond this platform-focused perspective and suggest rather generic and comprehensive mechanisms for distributed usage control [4, 2]. These mechanisms are based on the assumption that the uploading party is the one that should be provided with possibilities for specifying visibility and usage policies. Meddle-sharing, however, is a concern for the *depicted* party and therefore the mentioned mechanisms do not help.

Photo tagging plays an important role for privacy infringements related to meddle-sharing. Therefore, advanced mechanisms for semi-automated untagging [1] seem highly promising. Nonetheless, such mechanisms are restricted to the platform they are employed in and do not prevent privacy infringements upon non-members. Furthermore, they only come into effect *after* a photo has been uploaded and tagged.

Our mechanism, in contrast, is explicitly designed with the problem of meddle-sharing in mind. Furthermore, it is designed to come into effect much earlier than established models of untagging etc., thereby working against unwanted revelations through photo sharing in general and across the boundaries of single platforms.

[1] The concept presented herein was developed by a multitude of people under the umbrella of the Berlin-based Internet & Society Collaboratory (http://en.collaboratory.de). Besides the authors, significant contributions were made by (in alphabetical order): Thomas Heilmann, Jan Schallaböck, Max Senges and Gordon Süß. The proof-of-concept software was written by Markus Köbele. See also http://offlinetags.net.

Problem Confinement

If we include the privacy risks that were also present for “traditional” media, we can, from a high-level perspective, now distinguish at least three generic classes of unwanted information revelations:

Unintentional discovery

The most obvious case of information about a person being revealed through photos has been present since the existence of photography in general: A person looking at a photo unintentionally recognizes a known photo subject in a specific context. Such random discoveries constitute a privacy infringement in the above-mentioned sense, at least in the case where the depicted context conflicts with the subject’s intended public image (think, again, of a political party’s convention).

Directed searchability

The possibility to search for photos showing specific persons just by their name, Twitter ID etc. clearly distinguishes current online social networks and other photo sharing platforms from traditional settings. As soon as the respective search results contain “incriminating” information – information that does not match the shown person’s intended public image – this unquestionably heightens the risk of unwanted information revelations and thereby limits the photo subject’s ability to control her public image.

Reverse searchability

Finally, novel technologies from the field of face recognition introduce another risk: Instead of searching for “images attributed to a given identifier”, it is now also possible to revert this search and start with a photo of an unknown person in order to identify her and obtain further information. This leads to even absolute strangers being able to search for information about a given person that this person would by no means have revealed to such strangers. In particular, this also applies to reverted searches being made on the basis of photos meddle-shared by others and to such meddle-shared photos appearing in the results of a reverted search. Again, this heightens the risk of unwanted revelations significantly and thus reduces a person’s ability to control her public image.

As we can see, novel technologies and practices in the field of online photo sharing introduce new risks to an individual’s abilitity to exert control over her public image and reinforce existing ones. This loss of control should be countervailed by means of our Offlinetags.

The Four “Offlinetags”

As already laid out, our Offlinetags shall represent individual preferences of would-be photo subjects on the taking and sharing of photos showing them and counteract the above-mentioned privacy risks. After long-lasting discussions on possibilities for addressing these risks separately, we came to the conclusion that at least the risks of directed and reverse searchability are highly interrelated and can hardly be isolated from each other. For example, countering the risk of directed searches while at the same time ignoring the risk of reverse searchability would hardly make sense: Once a searching party gets hold of just one photo showing the subject she is searching for, directed searches can to a certain extent be substituted through reverse searches. Instead of addressing the different risks separately, we therefore decided to follow a graduated “risk-minimization” approach, leading us to the elementary

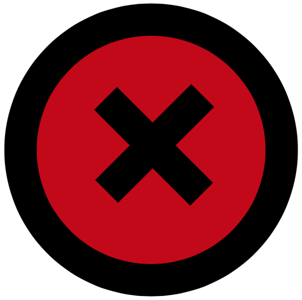

No photos!
Please refrain from taking pictures with me being depicted, no matter if I appear to be not re-cognizable for the person taking the picture.

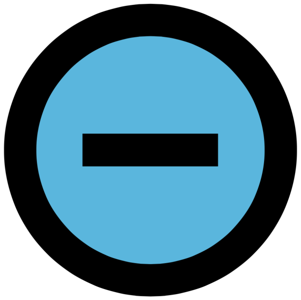

Blur me!
Please ensure before uploading or any application of a picture of me, that I cannot be recognized – especially by means of facial recognition algorithms.

pleas presented below: "No photos", "Blur me", "Upload me" and "Tag me".

Each of these pleas is represented by a respective symbol, the Offlinetag, that shall ensure easy human recognition as well as good machine-readability. We therefore decided for a fundamental design of a bold black circle as an anchor for image recognition algorithms and simple, algorithmically well-distinguishable black symbols inside the circle representing the plea. The symbols are also designed for being intuitively associated with the represented plea by humans. To reinforce the intuitive association with the respective plea, the free space is colored in a corresponding hue. The colors are, however, not intended to be evaluated by image recognition mechanisms to prevent analytical failures for gray photos, for example. Following these fundamental concepts, the meanings and graphical representations of our Offlinetags are as follows:

No photos
The first preference a would-be photo subject can have with regard to the risk of meddle-sharing is not to be seen on any photo in a certain situation – no matter whether this photo is intended to be uploaded somewhere or not. Following the above-mentioned approach of graduated risk-minimization, the "no photos" Offlinetag represents the most rigid plea to take no photos of the person currently wearing the button/badge. Graphically, this rigidness is represented by a cross-symbol in the middle of the circle and a red color hue, transporting a clear "stop" message. This Offlinetag is intended to address all three classes of unwanted information revelation identified above to the strongest possible extent. In particular, this also minimizes the risk of unintentional discovery through recognition of typical clothes or accessories being worn.

Blur me
To allow photos of multiple persons to be taken without infringing upon the privacy of individual photo subjects, we introduced the "blur me" Offlinetag. It represents the preference of the wearer to be made unrecognizable in case of the photo being shared. A common way for this anonymization is to blur out single faces, but other mechanisms can also be thought of. Following the known practice of blurring, we decided on a light blue color for this Offlinetag. As for the symbol, anything "blurry" would have been hard for algorithms to automatically recognize. We therefore chose a single black horizontal bar, referencing the one usually put over a photo subject's eyes for anonymization purposes in traditional media. If the represented plea is followed, this Offlinetag primarily minimizes the risks arising from reverse searches and from directed searches on the basis of taggings made by automated face recognition mechanisms. The risk of unintentional discovery is also limited to a certain extent, even if such discoveries can still happen on the basis of other recognized features than the photo subject's face like clothes, accessories, etc.

Upload me
Besides not wanting to be seen or recognized online, a photo subject can also accept or even desire being seen online in a specific context while still feeling uncomfortable with being subject to excessive directed and reverted searches. This less restrictive preference is reflected by the "upload me" Offlinetag, meaning that uploading and sharing the photo is accepted while rejecting tagging or face recognition mechanisms. The

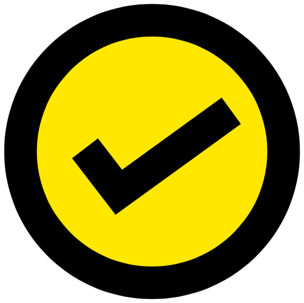

Upload me!
Feel free to upload and share pictures of me, but please refrain from tagging or facial recognition.

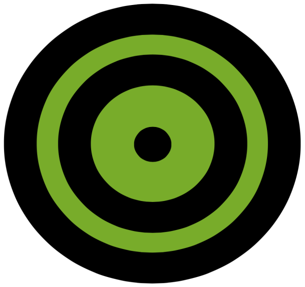

Tag Me!
Feel free to take pictures of me, upload them, tag them, and make them available for facial recognition or any other means.

general acceptance of uploading is represented by a checkmark symbol and the yellow color hue transports the intuitive message that at least some attention is necessary during the handling of the photo. Regarding the categories of unwanted revelations, this Offlinetag does not prevent unintentional discoveries but does, at least to a certain extent, work against the risks of directed and reverted searches.

Tag me
Finally, a photo subject can also have no objections against being tagged on photos showing her in certain contexts or being subject to face recognition mechanisms etc. Moreover, a photo subject can even have a vital interest in being tagged and algorithmically identified. The "tag me" Offlinetag represents this preference. It has a green color hue, signaling an "anything goes" attitude and carries another circle and a dot at the center as symbolic references to an abstract target. Different from the other Offlinetags, this one is not intended to counteract the identified privacy risks. Instead, it shall give the wearer a possibility to signal that she is aware of any potential risks and has consciously decided that they don't matter to her when using this badge/button.

Intended Use and Enforcement

As laid out above, Offlinetags are intended to be worn in the form of stickers, buttons, badges, etc. that represent the wearers privacy preferences in a given situation. The question is, then, how the so-formulated preferences are to be enforced. First of all, Offlinetags are *not* meant to "replace" or "overwrite" existing legal rules already regulating the handling of photos showing individuals. Instead, Offlinetags shall *complement* legal rules by providing a simple and intuitive way for communicating individual preferences within specific situations. This being said, we envisage different modes of enforcement:

First, the intuitive design shall make those persons taking and handling photos aware of the photo subjects' preferences. Of course, it is always possible to act against these preferences, but this would necessarily require a conscious decision against the subject's explicitly stated preference and therefore break a moral convention. In this vein, Offlinetags function as a means of communication among humans and allow for a more consensual practice in the field of photo sharing.

Second, Offlinetags are explicitly designed to be easily recognizable by algorithms. This enables automatic enforcement of the preferences at any point of the photo-sharing chain from the camera to the final recipient. In particular, one can think of cameras that don't take photos or automatically blur them as soon as the respective Offlinetag is recognized. Based on OpenCV and Qt, we implemented a proof-of concept desktop application realizing exactly this functionality for photos being taken by a webcam. Figure 1 shows this application in operation for a "blur me" Offlinetag.

As, however, such mechanisms could easily be circumvented, combined modes of enforcement seem most promising to us. For example, we envisage cameras and upload routines which automatically analyze photos, raise indicative warnings or even *offer* to automatically blur certain faces once the respective tag has been found. This would allow to take the best of both modes without patronizing individual users and introduce a whole new method of self-regulation since other users of a certain platform would be aware that a meddle-sharer must have *consciously* infringed upon the well-stated preferences of the photo subject.

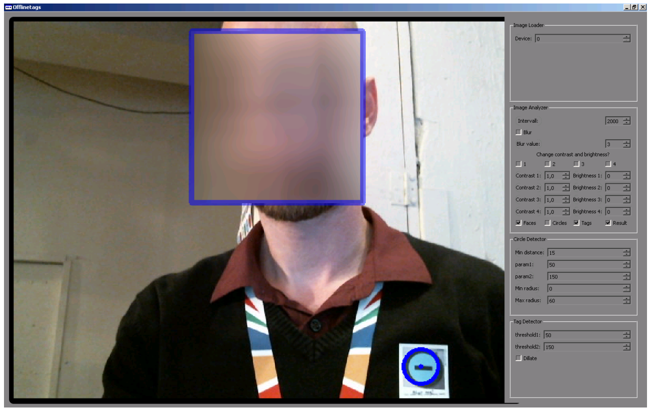
Offlinetags

Image Loader
Device: 0

Image Analyzer
Intervall: 2000
□ Blur
Blur value: 3
Change contrast and brightness?
□ 1 □ 2 □ 3 □ 4
Contrast 1: 1,0 Brightness 1: 0
Contrast 2: 1,0 Brightness 2: 0
Contrast 3: 1,0 Brightness 3: 0
Contrast 4: 1,0 Brightness 4: 0
☑ Faces □ Circles ☑ Tags ☑ Result

Circle Detector
Min distance: 15
param1: 50
param2: 150
Min radius: 0
Max radius: 60

Tag Detector
threshold1: 50
threshold2: 150
□ Dillate

Figure 1: Proof-of-concept implementation of the software showing the ability to recognize faces and tags appropriately and to automatically perform the desired behavior.

Conclusion and Next Steps

As we have seen, the concept of Offlinetags introduces novel ways of influencing the way photos are shared online. In particular, they address the problem of photos being "meddle-shared" without the consent or even knowledge of the photo subject. The design is optimized for being recognized by humans as well as by computer vision algorithms and thereby supports different modalities of enforcement.

For the future, we are planning to focus on more detailed analyses of the Offlinetags' interplay with existing legal regulations regarding the handling of photos. For example, it is unclear whether wearing an "upload me" tag would suffice as an explicit statement of consent to publicly share a photo as it is formally required under several legislations. Furthermore, we are also planning for structured user-studies on the appropriateness and acceptance of the overall concept as well as of the chosen elementary pleas.

From a more general perspective, the effectiveness of social enforcement mechanisms based on norm compliance also needs to be further explored. We will therefore relate the concept of Offlinetags to broader debates on anonymity, privacy and public space as well as considering its relationship to current debates on non-state governance. Finally, we will try to find partners who will implement and test Offlinetags-based mechanisms in camera-apps, upload routines and other contexts in order to evaluate the concept's practical applicability.

References

[1] Besmer, A., Lipford, H.R. Moving beyond untagging: photo privacy in a tagged world. Proc. CHI '10, pp. 1563-1572. Doi: 10.1145/1753326.1753560

[2] Kumari, P., Pretschner, A., Peschla, J., Kuhn, J.-M. Distributed data usage control for web applications: a social network implementation. Proc. CODASPY '11, pp. 85-96. Doi: 10.1145/1943513.1943526

[3] Mannan, M., Oorschot, P.C. Privacy-enhanced sharing of personal content on the web. Proc. WWW '08, pp. 487-496. Doi: 10.1145/1367497.1367564

[4] Pretschner, A., Hilty, M., Basin, D. Distributed usage control. Comm. ACM 49(9), 2006, pp. 39-44. Doi: 10.1145/1151030.1151053

[5] Whitman, J.Q. The Two Western Cultures of Privacy: Dignity versus Liberty. *Yale Law Journal 113(6),* 2004, pp. 1151-1221