

Context-Dependent Privacy-Aware Photo Sharing based on Machine Learning

Lin Yuan, Joël Theytaz, and Touradj Ebrahimi

Multimedia Signal Processing Group, EPFL
Station 11, 1015 Lausanne, Switzerland
{lin.yuan, joel.theytaz, touradj.ebrahimi}@epfl.ch

Abstract. Photo privacy has raised a growing concern with the advancements of image analytics, face recognition, and deep learning techniques widely applied on social media. If properly deployed, these powerful techniques can in turn assist people in enhancing their online privacy. One possible approach is to build a strong, automatic and dynamic access control mechanism based on analyzing the image content and learning users sharing behavior. This paper presents a model for context-dependent and privacy-aware photo sharing based on machine learning. The proposed model utilizes image semantics and requester contextual information to decide whether or not to share a particular picture with a specific requester at certain context, and if yes, at which granularity. To evaluate the proposed model, we conducted a user study on 23 subjects and collected a dataset containing 1'018 manually annotated images with 12'216 personalized contextual sharing decisions. Evaluation experiments were performed and the results show a promising performance of the proposed model for photo sharing decision making. Furthermore, the influences of different types of features on decision making have been investigated, the results of which validate the usefulness of pre-defined features and imply a significant variance between users sharing behaviors and privacy attitudes.

Keywords: Privacy protection, online social network, photo sharing, access control, decision making, context, machine learning

1 Introduction

Wide spread of smart mobile devices and online social networks (OSNs) make photo sharing an easy and popular activity. However, it has also raised concerns on privacy since the shared content reveals substantial sensitive information about people. Most social networking or photo sharing services provide access control for users to manage their privacy. However, users need to manually set their sharing policies in only a static manner, without the possibility to share their photos to different groups of people dependent to contexts, e.g. the location, time or even nearby people of potential viewer. Most access control mechanisms enforce only binary sharing options, namely “Yes” or “No”, which may not provide the best experience when a user just wants to disable partial information in

photo sharing. With the latest progress in image analytics, pattern recognition, and deep learning techniques, large scale information is mined from the shared multimedia content. Although seemingly compromising privacy, those techniques can in turn be used to enhance privacy, in such a way of helping people estimate the privacy value of their content or control the access of their content automatically and dynamically.

In this paper, we present a machine learning based model that can accurately predict users photo sharing decisions based on their past decisions. To make photo sharing decisions, the proposed model takes into account not only the content of an image, but also the context information about the image capture and potential requester. To validate the proposed model, we conducted a user study on 23 subjects and three sets of evaluation experiments.

The rest of the paper is structured as follows. Section 2 introduces related works. Section 3 describes in detail the proposed model. Then Section 4 and Section 5 present the user study and performance evaluation. Finally, Section 6 outlines some discussions and Section 7 summaries the paper.

2 Related Work

A number of studies have been focused on understanding users privacy concern on photo sharing, as well as the potential privacy implications via both subjective [1,2] and objective [8,13] studies. A number of approaches to privacy protection in photo sharing have been proposed, including usage control scheme in distributed OSNs [6], Secure JPEG scrambling image visual information [22,20,23], separate coding and sharing of JPEG image by P3 [14] and tag-based access control [12]. In addition, a substantial research effort has been made on estimating the privacy value or detecting privacy-sensitive objects in images. These works include private/public image classifications and privacy-sensitive visual information detection, based on not only learning low-level image features (color, edge, faces and SIFT) [24], but also deep learning approaches such as Convolutional Neural Network (CNN) [17,16].

Another branch of research has been focused on context-aware information sharing in the scenario of social networks or cloud services. Smith et al. [15] provided an early investigation on solutions to enable people to share contextual information in mobile social networks. Wiese et al. [19] investigated the impact of various factors on people’s willingness to share information. Harkous et al. [10] present a conceptual framework named C3P for automatic estimation of privacy risk of data based on the sharing context. Bilogrevic et al. [4] present SPISM, an information-sharing system that predicts (semi-)automatically sharing decision, based on personal and contextual features. Despite the substantial works on contextual information sharing, very few have considered context information for privacy protection in online photo sharing. To the best of our knowledge, this paper is the first attempt to investigate the feasibility of deploying both content-related and contextual features of images, to automatically make or “recommend” photo sharing decisions.




Fig. 1. Framework of a photo sharing system based on the proposed model.

3 A Model for Context-/Privacy-Aware Photo Sharing

3.1 Security Assumption and Operating Principle

First of all, we assume the photo sharing service providers are trustworthy. Users allow the service to conduct necessary analysis on their photos, and the system is granted the right to enforce access control of users photos.

Fig. 1 illustrates a photo sharing architecture of the proposed model. The operating procedures of the model can be described by the following story: Alice (the *sender*, who wants to upload and share photos with online friends) uploads a set of pictures on the photo sharing service, and the service system analyzes each picture and extracts a set content and contextual features about those pictures. Meanwhile, the system asks Alice a set of questions on her willingness to share each picture to specified individuals in various scenarios. These individuals can be selected from those who visited Alice’s profile recently or frequently. Each scenario describes a certain context of a possible *requester*, who attempts to visualize a picture shared by the sender. The context includes the identity (either real name or social group), location, nearby people and the time when the requester tries to visualize the image. The system then trains a classifier based on Alice’s answers for different photos in different scenarios. On the other side, Bob (the requester) visits the profile page of Alice. With the help of the classifier, the system analyzes Bob’s context and Alice’s photo information, to decide whether or not to show certain photos to Bob, and if yes, at which granularity.

3.2 Feature Definition

To train such a classifier, we considered two groups of features: **Image Semantic Features (I)** and **Requester Contextual Features (R)**. Instead of using low-level image features such as color, texture, composition and SIFT, we believe higher-level semantic features have more immediate relations with privacy.

Table 1. Feature notations and definitions.

	ID	Feature	Description
What	I_C	Image: Category	Major category of the picture, selected from the eight categories identified in Instagram pictures [11]: <i>Friends, Activity, Selfie, Food, Pets, Gadget, Fashion</i> and <i>Captioned photo</i> .
	I_A	Image: Activities	Activities involved in the picture, selected from 26 keywords partially defined by [5]: <i>working, meeting, reading, presentation, resting, chatting, socializing, family, friends, vacation, TV, cooking, eating, drinking, cleaning, shopping, exercising, traveling, walking, landscape, city, concert, sporting, gaming, gadget</i> and <i>pets</i> .
Who	I_P	Image: # of People	The number of people in the picture.
		Image: Identities	The existence of different identities in the picture. Eight types of identities were defined: <i>Sender him/herself, Family, Close friend, Schoolmate or Colleague, Girl or Boyfriend, Acquaintance, Celebrity</i> and <i>Stranger</i> .
	R_I	Requester: Identity	The relationship between the requester and the sender, categorized in six types: <i>Family, Close friend, Schoolmate or Colleague, Girl or Boyfriend, Acquaintance</i> and <i>Stranger</i> .
	R_G	Requester: Gender	Gender of the requester: <i>Female</i> or <i>Male</i> .
	R_N	Requester: Nearby	Whether or not the requester has other people nearby at requesting time.
Where	I_L	Image: Location	The semantic location where the image was captured, selected from 12 major location categories adopted from Foursquare Location Categories.
		Image: Loc. Coordinates	Latitude and longitude of the image capture location.
		Image: Loc. Frequency	The frequency of the sender being present in such place, selected from <i>Rarely, Sometimes, Often</i> and <i>Almost everyday</i> .
	R_L	Requester: Location	Semantic location of the requester, categorized in <i>Unknown, Friends home, His/her own home, Work place</i> and <i>Public place</i> .
When	I_T	Image: Time	The time of photo capture in a float value, e.g. 14.5 denotes 2:30 PM.
		Image: Day	The day (in a week) of photo capture, selected from <i>Monday</i> to <i>Sunday</i> .

These features include the image category, number/identities of people in image, activities or objects in image and the location and time of image capture. The contextual features of the requester include the requester’s identity, location, nearby people and time.

A detailed description of all the features used in our experiments, grouped in different aspects of context, is shown in Table 1. Note that the time of requester is not used in the current experiment because it would be too cumbersome for subjects to read and analyze the complete information containing all contexts.

3.3 Photo Sharing Decisions

We defined three photo sharing decisions, corresponding to different levels of photo information disclosure. The three decisions and corresponding descriptions presented in the user study are listed in the following:

Decision 1 - Do NOT Share: No, I don’t want to share the picture.

Decision 2 - Partially Share: Yes, but with some image region protected or/and metadata (GPS, time, etc.) removed.

Decision 3 - Entirely Share: Yes, I want to share the picture completely.




Fig. 2. Workflow of user study using ProShare S.

The reasons of using the specific three sharing decisions instead of conventional binary decisions (“Yes” or “No”) are twofold: First, in many scenarios of online photo sharing, people may want to simply remove partial privacy-sensitive visual information in an image, such as ID card, license plate or their children faces. Second, most images shared from smart mobile devices contain metadata such as geotags, camera model and time, which could also compromise privacy. Therefore, an option should be provided for users to partially protect and share their image content.

4 User Study and Data Collection

We conducted a study that put participants in personalized photo sharing scenarios, and collected an image dataset containing manual-annotated image semantic features and contextual sharing decisions.

To conduct the user study, we developed an Android app¹, named *ProShare S*. The application allows a user to create an account, take pictures, conduct a set of surveys for each, protect privacy-sensitive image regions, and finally upload them to a dedicated server. The workflow of a user study using ProShare S is illustrated in Fig. 2. Particularly, the survey part is structured in two sets of questionnaires:

Q1 - Image Semantic Information The first questionnaire (Q1) requires the user to add necessary image semantic tags. This questionnaire appears once a picture has been taken from either gallery or camera. The questions in Q1 cover all the semantic features defined in Section 3.2. A build-in face detector offered by Android API is applied to count the number of people in image, which can be manually modified if not correct. Location coordinates and capture time are automatically extracted from image metadata.

Q2 - Contextual Photo Sharing Decisions Once Q1 is finished, the user is directed to the second questionnaire (Q2), where he/she is presented with 12 sharing contexts/questions. For each context, the user needs to decide how he/she would like to share the picture with the specific requester, by selecting one of the three decisions defined in Section 3.3. An example context is “Would

¹ The application is publicly available at <http://grebvm2.epfl.ch/proshare-s/proShare-rd2.1.apk>.




Fig. 3. Distribution of (a) images in each category and (b) subjects sharing decisions.

you share this picture with a *close friend*, when *he* is at a *public place* with *other people*?” The 12 contexts/questions are selected in a special way such that each of the six requester identities appears twice in a random order, with the other contextual features (gender, location, nearby people) sampled at random. In the study, basic user profile is also collected through the App. We therefore present the sharing contexts adaptively based on user’s profile. For instance, for a female user we present the requester as “your boyfriend” instead of “girl or boyfriend”.

We recruited 23 volunteers to participate in our user study, and assigned each of them a task of uploading at least 50 daily pictures of their own and completing corresponding surveys using ProShare S. Each subject was required to complete the task within a week and to try to cover a wide range of image content². Finally, 20 out of the 23 subjects successfully finished the required task. We therefore kept only the data of the 20 effective subjects for the later evaluation. A total of 1’018 images including 12216 sharing decisions were contributed by the 20 subjects, each providing 50.9 images on average. Fig. 3 shows the histogram of images in each category and the contextual sharing decisions made on all the images.

5 Evaluation and Analysis

5.1 Methodology

To evaluate the performance of the proposed model for decision making, we conducted three sets of experiments based on the data collected from our user study. We take the working hypothesis that users photo sharing behaviors and privacy attitudes are highly subjective and the difference in users behaviors may cause the proposed model to perform differently between subjects.

The first experiment focused on the performance of the proposed model with respect to each user, namely, within-subject analysis. In the second experiment,

² The instruction and agreement sheet for the user study including several screenshots of the ProShare S App is available at http://grebvm2.epfl.ch/proshare-s/instruction_sheet_rd2.1.pdf

we explored a universal one-size-fits-all classifier trained on all users data for predicting a new user’s decisions. In the third experiment, we investigated the influences of different image and requester features on the decision making performance of the proposed model.

The WEKA machine learning library [9] was used in experiments and three representative classification methods were considered: logistic regression, support vector machine (SVM) and random forest. We started with a preliminary test by running a 10-fold cross validation on each user’s data using the three methods and random forest always outperformed the other two. We therefore kept using random forest for the rest of the experiment.

To evaluate the decision making performance, the following metrics are used:

- **Correct Decision rate:** The proportion of correctly predicted decisions.
- **Over-Sharing rate:** The proportion of cases where image information is shared more than what user expect to share, which compromises privacy.
- **Under-Sharing rate:** The proportion of cases where image information is shared less than what user expects to share, which may compromise usability.
- **Kappa statistic:** Cohen’s kappa score [18] that measures the chance-corrected agreement between predicted and ground truth decisions.

5.2 Within-Subject Analysis

In the first experiment, we used different proportions (from 10% to 90%) of each subject’s data to train a classifier, and evaluated the classifier on the rest of the data (evaluation set). This is to examine the trade-off between user-burden and prediction accuracy of the proposed model. The evaluation results measured by different metrics across the 20 subjects are shown as box plots in Fig. 4. In this figure, one observes that the median correct decision rate has already reached 0.75 at a training set of only 10%, which corresponds to only 5 images in average. This means we could already build an acceptable model for half of the users using a very small number of images and their decisions. Above the training set of 50%, most users obtained the correct decision rate higher than 0.8. The median Kappa score at the training set of 10% is below 0.5 and rapidly reaches 0.6 at the training set of 20%. Above the training size of 60%, an almost perfect prediction is observed for half of the users with a median Kappa statistic greater than 0.8. On the other hand, both the over-sharing and under-sharing rates of most users are very low, even at the training set of 10%. However, we observe the over-sharing rate is always higher than the under-sharing rate. A possible explanation is that most users tend to share images and the numbers of different decisions in the dataset are imbalanced. From the results, one also observes a significant variance between users. At the training size of 10%, the maximum difference in correct decision rate between users is up to 0.44. At the training size of 80%, where the optimal performance is obtained for most of the users, such difference still remains around 0.2. Such results agree with our hypothesis made in the beginning of this section that users subjective behaviors may influence the performance of the proposed model.




Fig. 4. Performance of sharing decision prediction at different sizes of training sets.

Table 2. The cost matrix of the applied cost-sensitive learning.

↓ classified as →	Decision 1	Decision 2	Decision 3
Decision 1	0	$C_{1 \rightarrow 2} = c$	$C_{1 \rightarrow 3} = 2c$
Decision 2	1	0	$C_{2 \rightarrow 3} = c$
Decision 3	1	1	0

Cost-Sensitive Decision Making To address the issue of over-sharing, we introduced the cost-sensitive learning [7] in our decision making core. The aim is to evaluate the extent to which incorrect decisions can be biased towards the under-sharing cases instead of over-sharing, when users concern their privacy more than usability. We specified different error-penalties $C_{i \rightarrow j}$ (> 1) for over-sharing cases and the penalty of 1 for under-sharing. Therefore, the training process tries to minimize the following cost: $\sum_{1 \leq i < j \leq 3} (C_{i \rightarrow j} \times N_{i \rightarrow j} + 1 \times N_{j \rightarrow i})$, where $N_{i \rightarrow j}$ denotes the number of cases where Decision i is misclassified classified as Decision j . Specially, we assigned a double error-penalty $2c$ for the over-sharing cases $C_{1 \rightarrow 3}$ compared to the other two over-sharing cases. This is because a mistake by classifying “Do NOT Share” to “Entirely Share” may severely compromise privacy. The cost matrix for the cost-sensitive learning is shown in Table 2.

We experimented with a set of values for c (from 1.5 to 5), on each user’s data using the same random forest classification. The results at $c = 2$ and $c = 4$ are shown in Fig. 5. With an error-penalty $c = 2$, the over-sharing rate is greatly reduced to a level lower than the under-sharing rate. When increasing c to 4, the over-sharing rate is further reduced, in sacrifice of a significant increase on the under-sharing rate. This indicates a significant trade-off between the capability of privacy protection and system usability. In any cases of cost-sensitive learning, the overall correct decision rate and Kappa statistic do not change much, as the introduced error-penalty mainly acts as a parameter to tune the weights of different incorrect decisions.




Fig. 5. Performance of cost-sensitive decision making with two different values of c .




Fig. 6. Performance of a One-Size-Fits-All classifier.

5.3 One-Size-Fits-All Model

In the second experiment, we evaluated a one-size-fits-all model, to examine the potential of building a global classifier trained on the data of all users, to make or “recommendation” decisions for new users. To be fair, for each subject i , we trained a classifier with random forest on the data of the remaining subjects, which was then evaluated on the data of subject i . Cost-sensitive learning was also included in this experiment for comparison. The results over all the 20 subjects are shown in Fig. 6. The median correct decision, over-/under-sharing rates and the Kappa statistic without cost-sensitive learning are 0.636, 0.218, 0.155 and 0.348 respectively. With cost-sensitive learning, the over-sharing rates are reduced below under-sharing, without greatly degrading the correct decisions and Kappa score. The overall performance of such a one-size-fits-all model is not as good as the personalized classifier built on each user’s own data. This again implies that users may have very different behaviors and privacy attitudes towards photo sharing. However, such a classifier could already provide an acceptable performance better than a random guess. This experiment provides the insight of building a one-size-fits-all classifier to predict or “recommend” photo sharing decisions for a new user, until the user has enough data to build a personalized classifier.




Fig. 7. Distribution of photo sharing decisions distinguished by different features.

5.4 Influences of Features on Decision Making

At the end, we investigated the influences of different types of features on users photo sharing decisions and on the performance of our prediction model.

First, the histograms of three sharing decisions distinguished by different types of features are shown in Fig. 7. The variation in distributions over different feature values indicates the degree of influence of a particular type of features. One observes a significant difference in decision histograms across different requester identities, which implies that the requester identity influences users decision making the most. On the other hand, although the decision histograms do not change much between other contextual features of the requester, there is still a small decrease of the “Entirely Share” decisions at the cases where the requester is at an “Unknown” place or with “Other people” nearby. Image semantic features also influence users decision making significantly. For instance, users prefer sharing photos without people or with a lot of people (≥ 6), to sharing photos with 1 ~ 5 people. Also, users favor sharing those pictures containing strangers or celebrities, over personal photos with intimate connections like family and close friends.

We then evaluated the performance of decision making on different combinations of image and requester features, by conducting a 10-fold cross validation on each user’s data. The correct decision rates of cross validation of all the 20 subjects are shown in Fig. 8. Please refer to the feature notations in Table 1. We gradually remove certain features, and the leftmost and rightmost box plots in Fig. 8 show two extreme cases where all the features ($\mathbf{I}_{All} + \mathbf{R}_{All}$) or only the requester features (\mathbf{R}_{All}) were used. As is shown, when reducing features, the correct decision rate of the majority of subjects decreases, which implies that all those features in general have a positive impact on decision making for most users. When reducing image-related features, a significant variance across different subjects is observed, which indicates that those image features are important for modeling many users sharing decisions. However, for two of those subjects, the prediction model still performs well (correct decision rate higher than 0.9) even using only the requester features (\mathbf{R}_{All}). A possible reason is that




Fig. 8. Correct decision rates of all users obtained on different combinations of features.




Fig. 9. Performance of five example users obtained on combination of all Image Semantic Features (\mathbf{I}_{Au}) and different Requester Contextual Features (\mathbf{R}).

the two users made their sharing decisions mostly dependent on the context of requesters, regardless of the image content.

One also finds that by removing certain requester contextual features, like requester gender (\mathbf{R}_G), location (\mathbf{R}_L), or nearby people (\mathbf{R}_N), the overall accuracy does not significantly change. With merely the requester identity (\mathbf{R}_I) + all image features (\mathbf{I}_{Au}), the overall decision making accuracy still remains high. This implies that the requester contextual information than the requester identity has very weak or even negative influence on decision making. However, this is not always the case for every subject. Fig. 9 illustrates the results of five example subjects obtained on different combinations of requester contextual features (along with all image features \mathbf{I}_{Au}). Here, one observes that the inclusion of requester contextual features other than the requester identity influences decision making quite differently between users. For instance, the correct decision rate of User *C* obtained on all requester features (~ 0.8) is much higher than that on only requester identity \mathbf{R}_I . For User *A* or *D*, combining different requester features ($\mathbf{R}_{I,L,N}$ or $\mathbf{R}_{I,N}$ respectively) generates better accuracy than just using requester identity \mathbf{R}_I . However, for User *B* and *E*, using only the requester identity \mathbf{R}_I provides the best performance, in which case the other contextual features of requester are considered as noise in machine learning. Such a variance between users again proved our hypothesis that users have different personalized behaviors in photo sharing.

6 Discussions

Image Visual Information and Metadata Protection Prior to this study, we have proposed and researched on different approaches [20,21,22,23] to protect image privacy (visual information and metadata) such that the protected photos can be publicly shown to any party (service provider and individuals) while original photos being secretly shared to authorized individuals. The principal idea of these approaches is to utilize JPEG application markers to secretly preserve partial image information or metadata, which not only enables the reversibility of the obfuscated image but also minimizes the storage burden. Such approaches are collectively named Secure JPEG. In the proposed photo sharing model, Secure JPEG can be used to create a secure version of a photo. Depending on the predicted decision, the system can release the corresponding version (protected or recovered original form) of the image to a requester.

Security Discussion As mentioned in Section 3.1, we assume the service provider in proposed model is trusted. The reasons are twofold: First, it is still not possible to perform certain pattern recognition tasks on client devices efficiently, e.g. image semantic recognition; Second, the system makes sharing decisions in a dynamic way by analyzing both image content and requester context, which means the decision making core must lie on the service provider. However, as the development of pattern recognition on mobile devices, the security requirement of the proposed model can be relaxed. In another specific case of the proposed model, where only requester’s identity is taken into account (no other context), the security assumption can be discarded. In this case, the photo sharing decisions are made in a static way equivalent to using an access policy. According to a privacy-preserving photo sharing architecture proposed in our previous work [23], the access policy can be integrated in a Ciphertext-Policy Attribute-Based Encryption (CP-ABE) [3] and secure photo sharing can be achieved through an “Honest but Curious” untrusted server.

Feature Extraction In this study, most image semantic features were manually annotated. This is because we lack the access and control to a popular social network, and that the automatic tools for extraction of some semantic features (e.g. activities in image [5]) are not mature enough. In practice, with the advances in deep learning, content understanding and ubiquitous sensors, automatic extraction of different semantic or contextual features is becoming more accurate and fine-grained.

7 Conclusion

This paper presents a conceptual model for context-dependent and privacy-aware photo sharing based on machine learning. The proposed model utilizes the images semantic and requesters contextual information to predict photo sharing

decisions for users, based on their previous shared photos and past decisions. To evaluate the proposed model, we first conducted a user study on 23 subjects and collected a dataset containing 1'018 manually annotated images with 12'216 personalized sharing decisions in different contexts. Evaluation experiments have been performed and show a promising performance of the proposed method. Furthermore, the influence of different content- and context-related features on decision making has been investigated, the results of which validated the importance of pre-defined features and implied a significant variance between users sharing behaviors and privacy attitudes.

As our future work, we intend to conduct larger-scale user study based on realistic social networking environment. This will further help us understand users photo sharing behaviors. In addition, we will investigate more sophisticated machine learning or even deep learning approaches to build more accurate and secure photo sharing systems. We believe machines will become intelligent enough to understand people's privacy concerns towards their photo content and this is how we define "privacy-aware".

Acknowledgement

This research was possible thanks to the Swiss National Science Foundation funded project LEADME (200020-149259). We also acknowledge Thomas Mizraji and Vincent Debieux for contributions in the development of ProShare S app. A special thank goes to all the patient subjects who participated in the user study.

References

1. Ahern, S., Eckles, D., Good, N., King, S., Naaman, M., Nair, R.: Over-exposed?: Privacy patterns and considerations in online and mobile photo sharing. In: CHI. pp. 357–366. ACM (2007)
2. Besmer, A., Richter Lipford, H.: Moving beyond untagging: Photo privacy in a tagged world. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. pp. 1563–1572. CHI '10, ACM, New York, NY, USA (2010)
3. Bethencourt, J., Sahai, A., Waters, B.: Ciphertext-policy attribute-based encryption. In: Proceedings of the 2007 IEEE Symposium on Security and Privacy. pp. 321–334. SP '07, IEEE Computer Society, Washington, DC, USA (2007)
4. Bilogrevic, I., Huguenin, K., Agir, B., Jadliwala, M., Gazaki, M., Hubaux, J.P.: A machine-learning based approach to privacy-aware information-sharing in mobile social networks. *Pervasive and Mobile Computing* 25, 125 – 142 (2016)
5. Castro, D., Hickson, S., Bettadapura, V., Thomaz, E., Abowd, G., Christensen, H., Essa, I.: Predicting daily activities from egocentric images using deep learning. ISWC (2015)
6. Cutillo, L.A., Molva, R., Önen, M.: Privacy preserving picture sharing: Enforcing usage control in distributed on-line social networks. In: 5th ACM Workshop on Social Network Systems. Bern, Switzerland (April 2012)
7. Elkan, C.: The foundations of cost-sensitive learning. In: Proceedings of the 17th International Joint Conference on Artificial Intelligence - Volume 2. pp. 973–978. IJCAI'01, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA (2001)

8. Friedland, G., Sommer, R.: Cybercasing the joint: On the privacy implications of geo-tagging. In: Proceedings of the 5th USENIX Conference on Hot Topics in Security. pp. 1–8. HotSec'10, USENIX Association, Berkeley, CA, USA (2010)
9. Hall, M., Frank, E., Holmes, G., Pfahringer, B., Reutemann, P., Witten, I.H.: The WEKA data mining software: An update. *SIGKDD Explor. Newsl.* 11(1), 10–18 (Nov 2009)
10. Harkous, H., Rahman, R., Aberer, K.: C3P: Context-Aware Crowdsourced Cloud Privacy. In: Privacy Enhancing Technologies, Pets 2014. vol. 8555, pp. 102–122. Springer-Verlag Berlin (2014)
11. Hu, Y., Kambhampati, L.M.Y.S.: What we Instagram: A first analysis of instagram photo content and user types. In: Proceedings of the 8th International Conference on Weblogs and Social Media, ICWSM 2014. pp. 595–598. The AAAI Press (2014)
12. Klemperer, P., Liang, Y., Mazurek, M., Sleeper, M., Ur, B., Bauer, L., Cranor, L.F., Gupta, N., Reiter, M.: Tag, you can see it!: Using tags for access control in photo sharing. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. pp. 377–386. CHI '12, ACM, New York, NY, USA (2012)
13. Pesce, J.a.P., Casas, D.L., Rauber, G., Almeida, V.: Privacy attacks in social media using photo tagging networks: A case study with facebook. In: Proceedings of the 1st Workshop on Privacy and Security in Online Social Media. pp. 4:1–4:8. PSOSM '12, New York, NY, USA (2012)
14. Ra, M.R., Govindan, R., Ortega, A.: P3: Toward privacy-preserving photo sharing. In: Presented as part of the 10th USENIX Symposium on Networked Systems Design and Implementation. pp. 515–528. USENIX, Berkeley, CA (2013)
15. Smith, I., Consolvo, S., Lamarca, A., Hightower, J., Scott, J., Sohn, T., Hughes, J., Iachello, G., Abowd, G.D.: Social Disclosure of Place: From Location Technology to Communication Practices, pp. 134–151. Berlin, Heidelberg (2005)
16. Tonge, A., Caragea, C.: Image privacy prediction using deep features. In: Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence (2016)
17. Tran, L., Kong, D., Jin, H., Liu, J.: Privacy-cnh: A framework to detect photo privacy with convolutional neural network using hierarchical features. In: Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence (2016)
18. Viera, A.J., Garrett, J.M., et al.: Understanding interobserver agreement: the kappa statistic. *Fam Med* 37(5), 360–363 (2005)
19. Wiese, J., Kelley, P.G., Cranor, L.F., Dabbish, L., Hong, J.I., Zimmerman, J.: Are you close with me? are you nearby?: investigating social groups, closeness, and willingness to share. In: Proceedings of the 13th international conference on Ubiquitous computing. pp. 197–206. ACM (2011)
20. Yuan, L., Korshunov, P., Ebrahimi, T.: Secure jpeg scrambling enabling privacy in photo sharing. In: 11th IEEE International Conference and Workshops on Automatic Face and Gesture Recognition (FG). vol. 04, pp. 1–6 (May 2015)
21. Yuan, L., Ebrahimi, T.: Image transmorphing with JPEG. In: Image Processing (ICIP), 2015 IEEE International Conference on. pp. 3956–3960 (Sept 2015)
22. Yuan, L., Korshunov, P., Ebrahimi, T.: Privacy-preserving photo sharing based on a secure JPEG. In: 2015 IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS). pp. 185–190. IEEE (2015)
23. Yuan, L., McNally, D., Kupcu, A., Ebrahimi, T.: Privacy-preserving photo sharing based on a public key infrastructure. In: Proc. SPIE. vol. 9599 (2015)
24. Zerr, S., Siersdorfer, S., Hare, J., Demidova, E.: Privacy-aware image classification and search. In: Proceedings of the 35th International ACM SIGIR Conference on Research and Development in Information Retrieval. pp. 35–44. SIGIR '12, ACM, New York, NY, USA (2012)