2022 IEEE/ACM International Conference on Big Data Computing, Applications and Technologies (BDCAT) | 978-1-6654-6090-3/22/$31.00 ©2022 IEEE | DOI: 10.1109/BDCAT56447.2022.00029

Learning and Preserving Relationship Privacy in Photo Sharing

Jialin Liu
Department of Computer Science
Prairie View A&M University
Prairie View, USA
jialinliu528@gmail.com

Lin Li
Department of Computer Science
Prairie View A&M University
Prairie View, USA
lilin@pvamu.edu

Na Li*
Department of Computer Science
Prairie View A&M University
Prairie View, USA
nali@pvamu.edu

***Abstract*—In recent years, Online Social Networks (OSN) have become popular content-sharing environments. With the emergence of smartphones with high-quality cameras, people like to share photos of their life moments on OSNs. The photos, however, often contain private information that people do not intend to share with others (e.g., their sensitive relationship). Solely relying on OSN users to manually process photos to protect their relationship can be tedious and error-prone. Therefore, we designed a system to automatically discover sensitive relations in a photo to be shared online and preserve the relations by face blocking techniques. We first used the Decision Tree model to learn sensitive relations from the photos labeled private or public by OSN users. Then we defined a face blocking problem and developed a linear programming model to optimize the trade-off between preserving relationship privacy and maintaining the photo utility. In this paper, we generated synthetic data and used it to evaluate our system performance in terms of privacy protection and photo utility loss.**

***Index Terms*—relationship privacy, utility, decision tree, face blocking, linear programming**

I. Introduction

Photo-sharing in social media has become prevalent at an unprecedented scale. Unfortunately, semantically rich photos not only contain what users want to share but may also reveal sensitive information they do not intent to expose, such as users' location, personal habits, or their relations [1], [2]. Today, most photo-sharing platforms allow users to configure their privacy preference to a certain degree. For instance, Facebook allows users to determine who can view the photos (e.g., friends or family). However, studies have shown that users have difficulty configuring and maintaining those privacy settings [1]. Given the amount and depth of shared information, such case-by-case decision-making configuration can be tiresome and error-prone.

This paper presents the design of an automatic system to discover and preserve sensitive relations in photo sharing. We first simulate that users label their photos private or public based on their own privacy preferences but without sharing the preferences with the system. Then our system automatically learns the sensitive relations which dominate users' photo labeling. With the learned relations, the system discovers and hides sensitive relations in the photos through face blocking. Particularly, the system handles the trade-off between privacy protection and photo utility loss by minimizing the number of blocked faces. The contribution of this paper can be summarized as follows:

- Propose a new photo privacy problem and investigate how to hide sensitive relations in a single photo (Note that research on the exposure of sensitive relations across multiple photos is out of the scope of this work)
- Design a decision tree based algorithm to learn sensitive relations from the labeled photos
- Define an optimization problem of face blocking to handle the trade-off between privacy protection and utility loss and solve the problem using linear programming

The roadmap of this paper is outlined as follows: Section II gives a literature review of photo sharing privacy. Section III describes the two major components of our system, *Relationship Learning* and *Face Blocking*. Section IV presents the experimental results to evaluate the system in terms of privacy preservation and utility loss. Section V concludes the paper.

II. Related Work

To address the concerns, research has been conducted on privacy policy recommendation and access control mechanisms to assist users in decision-making before sharing photos online. Some researchers [3]–[5] focused on developing access control systems to prevent unwanted individuals from recognizing users in a photo. They used descriptive photo tags labeled by people or rendered by the tagging service to create access-control rules. In [3] and [4], the granularity of access control is changed from the level of the photo to that of a user's personally identifiable information (e.g., face).

A number of technologies focused on defining privacy policies to decide whether a photo should be considered public or private. Photo-related information, such as tags, comments, and content, is an indicator for creating users' privacy policies. Some work trained learning models using the information for private and public image classification [5]–[7]. [6] proposed seven types of privacy concerns and 268 privacy-sensitive object classes for image categorization. Additionally, they integrated the user trustworthiness into the classifier to recommend fine-grained privacy settings for image sharing. [7] designed a framework, called HideMe, to preserve the associated users' privacy. It allows users to build a scenario-based access control model by combining the factors, such as

978-1-6654-6090-3/22/$31.00 ©2022 IEEE
DOI 10.1109/BDCAT56447.2022.00029

Authorized licensed use limited to: Prairie View A M University. Downloaded on August 05,2023 at 02:11:57 UTC from IEEE Xplore. Restrictions apply.

temporal, spatial, and interpersonal attributes, and then decide to blur or show their faces for each scenario. The framework also protects the privacy of the bystanders in the photos.

With the fast development of big data and deep learning technologies, they were used to extract features for photo classification [8]–[10]. [8] trained a deep learning model to identify sensitive objects and provide recommendation of privacy settings. [9] categorized personal information in images into 68 attributes and trained models that predict such attributes. It also proposed models that predict user specific privacy score from images in order to enforce the users' privacy preferences. [10] proposed a system, AutoPri, to automatically detect private photos in a user-specific manner through a model based on a multimodal variational autoencoder and pinpoint sensitive regions in private photos. Nevertheless, these work did not consider the sensitive relationship between users. In the literature, several image processing techniques, such as blurring, warping, pixelation, and cropping, have been applied to processing sensitive objects detected to preserve photo privacy [11], [12].

III. System Overview

A framework of our system is illustrated as in Figure 1. The system consists of two major components, *Relationship Learning* and *Face Blocking*, which are introduced below:

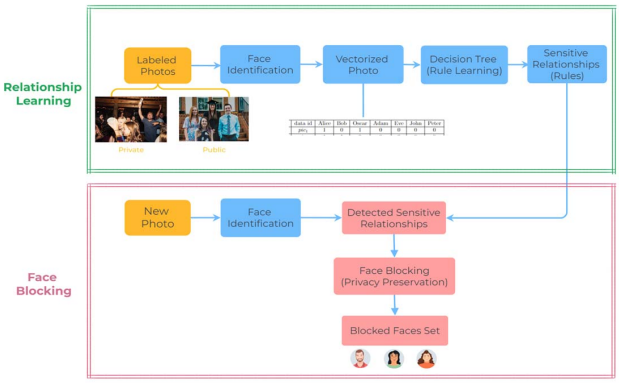

Fig. 1: System flowchart

A. Relationship Learning

This component applies supervised learning to learn the sensitive relations from a training set of labeled photos. Sensitive relationship is defined as a relationship in which the involved people are not willing to expose to others. A sensitive relationship can involve two or more people. Different users have different relations they consider private, which essentially guide how they label their photos as *private* or *public*. In this research, we made the following assumptions to simplify practical photo sharing scenarios: (1) labeled photos contributed by users must include themselves, and (2) users are concerned about the disclosure of their sensitive relations with others. For example, if a user u_1 considers his relationship with u_2 and u_3 sensitive, then we claim that u_1's privacy is compromised whenever a photo involving the three users is published.

Given a labeled photo, we vectorize it and make each user in the OSN a binary feature of the vector. If a user exists in the photo, the corresponding feature is one, otherwise zero. We use the Decision Tree algorithm to learn the sensitive rules (i.e., relations) from the photos. We first build the tree to partition the photos according to their labels. Then, for each path from the root to a private leaf node, we consider it as a learned rule. The learned rule from each path of the decision tree contains feature nodes with different conditions (i.e., "absent" or "present"). Since we only focus on whether a user's presence contributes to the decision-making of the photo class (i.e., "Private" or "Public"), we remove the "absent" features from each rule and then generate the whole rule set.

It should be noted that our system only needs users to label limited photos at the beginning of their use of the OSN. Afterwards, the system will automatically use the learned rules to detect sensitive relations in posted photos Additionally, our system will re-learn the rules whenever a new user joins the OSN and contributes his photos to the training dataset or an existing user requests to update his rules by labeling more data. The updated rules will be applied to all photos uploaded thereafter. This design ensures the system scalability.

B. Face Blocking

To preserve the sensitive relations of a photo which contains the rule(s) learned previously, we need to block one or more people's faces in the photo. In an extreme case, we may block all faces to fully protect people's sensitive relations as no one can tell who are concurrently in the photo. However, the aesthetic value of the photo will be reduced to zero. Thus, we need to handle the trade-off between privacy protection and utility loss. Accordingly, we define the problem as:

Optimizing Face Blocking: Given a set of rules (i.e., sensitive relations) discovered in a photo, minimize the number of faces to be blocked to hide all the rules.

The face blocking problem can be represented by a bipartite graph to depict the people and their sensitive relations in a photo. An example is shown in Figure 2, where all people detected in the photo are on the left, and all rules detected are on the right. There is a connection between a user and a rule if the user is involved in the rule. If a user's face is blocked, then the rule(s) he is associated with are preserved. Hence, the research problem is whose faces should be blocked so that the utility loss of the photo is minimized whiling preserving all the rules detected in the photo. For instance, in Figure 2, two blocking options: $\{u_0, u_1, u_5\}$ and $\{u_1, u_6\}$ can both preserve the four rules, but the latter is better as it blocks fewer faces.

The face blocking problem can be converted into an equivalent Set-Cover problem if we take the detected rule set as the universe and each user as a subset covering part of the universe (i.e., the rules the user involves in). Since Set-Cover is a classical NP-complete problem, the face blocking problem is also NP-hard. Accordingly, we developed a linear programming model to find the optimal solution, which can be formulated as in Algorithm 1.

Authorized licensed use limited to: Prairie View A M University. Downloaded on August 05,2023 at 02:11:57 UTC from IEEE Xplore. Restrictions apply.

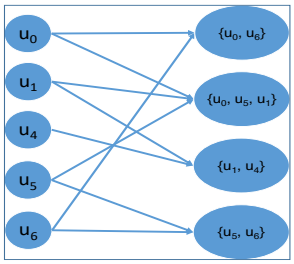

Fig. 2: A bipartite illustration of users and sensitive relations

Algorithm 1: Face blocking with Linear Programming

Input: Detected rule set $\mathcal{R}$ and user set $\mathcal{U}$

$$\text{Minimize:} \quad \sum_{i=1}^{|\mathcal{U}|} S_i \tag{1}$$

$$\text{Subject To:} \quad \sum_{i=1}^{|\mathcal{U}|} S_i \cdot X_{ij} \geq 1 \tag{2}$$

$$S_i \in \{0, 1\} \tag{3}$$

where S_i indicates whether a person u_i is selected to be blocked, and X_{ij} represents whether a person u_i is involved in a sensitive rule r_j.

Output: The selected users to block

IV. Experimental Study

A. Datasets

In this research, we used a synthetic dataset to evaluate the system on learning sensitive relations and handling the trade-off between privacy preservation and utility loss, due to the inaccessibility of a real-world OSN dataset that fits our needs. Additionally, even with the latest face recognition technologies, there are false alarms which will lead to errors in rule learning and face blocking, thereby affecting the accuracy of what we intend to evaluate in the system.

1) Synthetic Dataset Generation: We created a dummy OSN with 30 users and generated a friend circle for each user. Assuming users only share photos of them with their friends and the friend circle size s_{fc_i} follows a normal distribution with $\mu = 5$ and $\sigma = 1$, for each u_i, we randomly selected s_{fc_i} users to form his friend circle. We took three steps to generate the datasets. *First*, we generated the vectorized photos that a user may contribute to the OSN. Each photo is represented by a 31-dimension vector with the features corresponding to the 30 users in the OSN and a class label indicating private or public. Since the labeled photos contributed by a user i to the training dataset may not cover all $2^{(s_{fc_i}-1)}$ combinations of his circle friends, we set a parameter, p_t, to define the average percentage of the photos a user contributes. We varied the value of p_t to evaluate the learning model performance. *Second*, we generated the sensitive relations (i.e., the ground truth rules). For each user, a random number K is selected in $[1, 5]$ as his rule number. We created all of his possible rules by selecting and combining his circle friends and randomly selected K rules. Then we merged the rules of all users to generate the ground truth rule set. *Last*, we labeled the photos "private" or "public" based on the ground truth rule set.

2) Testing Dataset: Similarly, we generated 1000 vectorized photos for the testing set. A parameter p is set to determine the probability of a user appearing in a photo. For each user feature in a photo, we randomly set it 1 with p probability, indicating the user's presence. Practically, p decides the average user number appeared in a photo. For example, if $p = 0.5$ and the OSN has 30 users, the average users in a photo equals 15. We conducted experiments with different p values, 0.1, 0.3, and 0.5, to measure its impact on the trade-off between privacy protection and utility loss. After generating the 1000 vectorized photos, we removed duplicates and the vectors which also appear in the training dataset. In the research, we have 1000 photos for p = 0.5, and 998 photos for p = 0.3, and 728 photos for p = 0.1. We only present the results for $p = 0.5$ in this paper due to the page limit.

B. Privacy Preservation

We evaluated the privacy preservation of our system from two aspects: *rule learning* and *the final output of the system*. The more accurate the rules learned from the training data, the better we can identify the rules from testing photos and thereby protect privacy. Additionally, we measure the privacy leakage existing in the photo after face blocking, which results from the rule learning inaccuracy. We evaluated the system performance by generating two sets of ground truth rules, with 68 and 58 rules, respectively.

1) Rule Learning: In this group of experiments, we varied the percentage of photos that individual users may contribute in training, p_t, from 30% to 90% at an interval of 20%. Given a percentage, we ran five trials and averaged the results. For each trial, the ground truth rules remain the same, but the training data is different due to the random selection of photos contributed for training. While evaluating the learned rules, three possible cases exist: C_1 - the ground truth rules that match learned rules; C_2 - the ground truth rules that are subset-covered by learned rules; and C_3 - the ground truth rules not covered, where $|C_3| = T - |C_1| - |C_2|$ and T is the number of ground truth rules.

TABLE I: Rule learning w.r.t p_t, LR: learned rules, $p = 0.5$

p_t	# of LR	# of C_1	# of C_2	# of C_3
30%	26.4	8.6	45.6	13.8
50%	41.4	19.4	37.8	10.8
70%	57	33.4	26.8	7.8
90%	67.6	44.4	20.4	3.2

Table I presents the learning results while $p = 0.5$ and $T = 68$. We can see that the number of ground truth rules learned exactly ($|C_1|$) increases with p_t. With more data contributed in training, more ground truth rules can be learned. Additionally, both the numbers of ground truth rules subset-covered by learned rules ($|C_2|$) and the number of ground truth rules not covered ($|C_3|$) decrease as p_t increases, which indicates that more ground truth rules can be learned accurately and there is less subset-coverage or non-coverage.

Authorized licensed use limited to: Prairie View A M University. Downloaded on August 05,2023 at 02:11:57 UTC from IEEE Xplore. Restrictions apply.

Fig. 3: Average number of rules not covered in privacy leaked photos with varied p_t when $p = 0.5$ (58 rules)

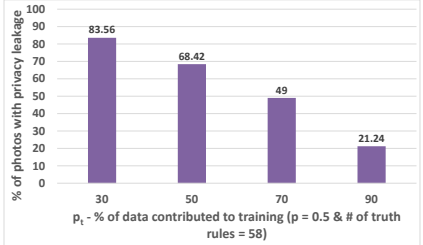

Fig. 4: Percentage of photos with privacy leakage with varied p_t when $p = 0.5$ (58 rules)

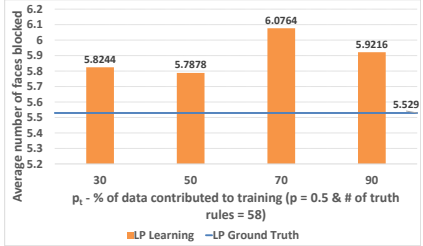

Fig. 5: Average number of faces blocked with varied p_t when $p = 0.5$ (58 rules)

2) Privacy Leak of the System: We analyzed privacy leak by counting the number of rules which are not covered in the photos after anonymization. Similar to previous experiments, we ran each experiment five times and averaged the results. Figure 3 and Figure 4 show that both the average number of rules not covered in the photos with privacy leak and the percentage of photos which still have privacy disclosure after anonymization decrease as p_t increases. This means that with more data contributed to training, the learning is more accurate and fewer rules are exposed in the anonymized photos.

C. Utility Loss

We measured utility loss by counting the number of blocked faces in the anonymized photos from the previous experiments. In practice, blocking different faces may cause different utility losses, but it is beyond the scope of this work. Figure 5 shows the average number of blocked faces with different p_t values. The horizontal line represents the average number of faces blocked based on the ground truth rules. It should be noted that more rules detected in a photo may not always lead to more faces being blocked. It depends on the users detected and their involvement in the private relations (i.e., the structure of the bipartite graph), which can be case-dependent.

V. Conclusion

In this paper, we designed a system to protect the sensitive relationship in photo sharing on OSN. We first applied Decision Tree to learn sensitive relationships from the labeled photos and then protected photo privacy by blocking faces. We defined a new face blocking problem to minimize the utility loss and handle the trade-off between privacy protection and utility retention. A linear programming model was developed to optimize the result. Experimental results show that in general, the more the photos contribute to the training process, the fewer the rules will be unlearned. Additionally, our model performs effectively in handling the trade-off aforementioned. Future work includes creating a real-world dataset and applying face recognition technology to identify users for relationship detection. Additionally, we will investigate photo privacy protection across multiple photos.

Acknowledgment

This project is supported in part by the National Science Foundation (NSF) under grant DUE-1712496.

References

[1] A. Acquisti and R. Gross, "Imagined communities: Awareness, information sharing, and privacy on the facebook," *Privacy Enhancing Technologies*, p. 36–58, 2006.

[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," *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 357–366, 2007.

[3] P. Ilia, I. Polakis, E. Athanasopoulos, F. Maggi, and S. Ioannidis, "Face/off: Preventing privacy leakage from photos in social networks," in *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, (New York, NY, USA), p. 781–792, Association for Computing Machinery, 2015.

[4] N. Vishwamitra, Y. Li, K. Wang, H. Hu, K. Caine, and G.-J. Ahn, "Towards pii-based multiparty access control for photo sharing in online social networks," in *Proceedings of the 22nd ACM on Symposium on Access Control Models and Technologies*, (New York, NY, USA), p. 155–166, Association for Computing Machinery, 2017.

[5] A. Squicciarini, D. Lin, S. Sundareswaran, and J. Wede, "Privacy policy inference of user-uploaded images on content sharing sites," *IEEE Transactions on Knowledge & Data Engineering*, vol. 27, no. 01, pp. 193–206, 2015.

[6] J. Yu, Z. Kuang, B. Zhang, W. Zhang, D. Lin, and J. Fan, "Leveraging content sensitiveness and user trustworthiness to recommend fine-grained privacy settings for social image sharing," *IEEE Transactions on Information Forensics and Security*, vol. 13, no. 5, pp. 1317–1332, 2018.

[7] F. Li, Z. Sun, A. Li, B. Niu, H. Li, and G. Cao, "Hideme: Privacy-preserving photo sharing on social networks," in *IEEE INFOCOM 2019 - IEEE Conference on Computer Communications*, pp. 154–162, 2019.

[8] J. Yu, B. Zhang, Z. Kuang, D. Lin, and J. Fan, "iprivacy: Image privacy protection by identifying sensitive objects via deep multi-task learning," *IEEE Transactions on Information Forensics and Security*, vol. 12, no. 5, pp. 1005–1016, 2017.

[9] T. Orekondy, B. Schiele, and M. Fritz, "Towards a visual privacy advisor: Understanding and predicting privacy risks in images," *CoRR*, vol. abs/1703.10660, 2017.

[10] N. Vishwamitra, Y. L. Li, H. Hu, K. Caine, L. Cheng, Z. Zhao, and G.-J. Ahn, "Towards automated content-based photo privacy control in user-centered social networks," in *2022 the Twelth ACM Conference on Data and Application Security and Privacy (CODASPY)*, 2022.

[11] J. He, B. Liu, D. Kong, X. Bao, N. Wang, H. Jin, and G. Kesidis, "Puppies: Transformation-supported personalized privacy preserving partial image sharing," in *2016 46th Annual IEEE/IFIP International Conference on Dependable Systems and Networks*, pp. 359–370, 2016.

[12] L. Yuan, P. Korshunov, and T. Ebrahimi, "Secure jpeg scrambling enabling privacy in photo sharing," in *2015 11th IEEE International Conference and Workshops on Automatic Face and Gesture Recognition (FG)*, vol. 04, pp. 1–6, 2015.

Authorized licensed use limited to: Prairie View A M University. Downloaded on August 05,2023 at 02:11:57 UTC from IEEE Xplore. Restrictions apply.