This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TDSC.2021.3132230, IEEE Transactions on Dependable and Secure Computing

"Do You Know You Are Tracked by Photos That You Didn't Take": Large-Scale Location-Aware Multi-Party Image Privacy Protection

Joshua Morris†, Sara Newman†, Kannappan Palaniappan†, Jianping Fan‡, Dan Lin†
† Department of Electrical Engineering and Computer Science
University of Missouri
jdm6b3@mail.missouri.edu,sn976@mst.edu,pal@missouri.edu,lindan@missouri.edu
‡Department of Computer Science
University of North Carolina - Charlotte
jfan@uncc.edu

Abstract—Most existing image privacy protection works focus mainly on the privacy of photo owners and their friends, but lack the consideration of other people who are in the background of the photos and the related location privacy issues. In fact, when a person is in the background of someone else's photos, he/she may be unintentionally exposed to the public when the photo owner shares the photo online. Not only a single visited place could be exposed, attackers may also be able to piece together a person's travel route from images. In this paper, we propose a novel image privacy protection system, called LAMP, which aims to light up the location awareness for people during online image sharing. The LAMP system is based on a newly designed location-aware multi-party image access control model. Unlike previous works on small scales, the LAMP system is highly efficient and scalable as it can enforce privacy protection for billions of users on social networks in real time. The LAMP system automatically detects the user's occurrences on photos regardless the user is the photo owner or not. Once a user is identified and the location of the photo is deemed sensitive according to the user's privacy policy, the user's face will be replaced with a synthetic face. A prototype of the system was implemented and evaluated to demonstrate its applicability in the real world.

Index terms: image sharing, location privacy, person in the photo background, scalability

1 Introduction

With the growing ubiquity of smartphones and other mobile devices, image sharing is gaining increasing popularity in social networks like Facebook, Instagram and Foursquare. During social image sharing, privacy protection has now become a crucial issue to be addressed since images can intuitively tell *when* and *where* a special moment took place, *who* participated and *what* were their relationships, i.e., sharing images can reveal much of users' personal and social environments and their private lives [3], [32], [34]. News has reported multiple incidents about people being fired due to their private photos being disclosed to undesired audience [5], [23].

Recognizing the importance of image privacy, researchers and social media sites have developed various privacy policies and tools to help users specify the group of people for photo sharing. However, most existing image privacy protection approaches [13]–[16], [19], [21], [29], [33], [37], [38], [40] focus mainly on the privacy of photo owners and at most the photo owners' friends. They lack the consideration of other people who are in the background of the photos and are not related to the photo owners. In fact, when a person is in the background of someone else's photo, he/she may be unintentionally exposed to the public when the photo owner shares the photo online. For example, Alice had a bad day and visited a pub at night. Someone took a photo of the pub with Alice in the background. Alice had no idea about the photo until her supervisor came to show concerns to her because he coincidentally saw her drunk photo online posted by the other person. A recent interview [28] among college students also confirmed such privacy concerns, indicating that more and more undergraduates worry about becoming an Internet meme because their embarrassing moments were photographed by their peers and posted on social media. As an initial effort towards this new privacy problem, Llia et al. [19] suggested the use of face recognition to identify all the people in the photo but their implementation is still limited to identifying photo owner's friends through available image tags and they have not considered the associated location privacy issues as discussed in the following.

With more and more images associated with geo-tags and timestamps, image privacy now comes to the crossroads of the location privacy. Such exposure may cause undesired consequences especially when the person being exposed was visiting sensitive locations. For example, a businessman Bob is meeting an important customer in a restaurant during a business trip while Jack, who usually reviews every restaurant he visits, took a photo of the restaurant with Bob and his customer in the background. Jack published his review along with the photo to a social media site. Bob's company's competitor noticed Bob was in the photo. The photo may have leaked business intelligence since it tells the competitor when and where Bob met the potential business partner whom the two companies are currently competing. There are many other scenarios, such as visiting a specialty

hospital or attending alcoholic counseling, which could cause similar uneasiness to the person if his/her photos at those sensitive locations were posted online by others. Furthermore, attackers may even be able to piece together a person's travel route by analyzing unprotected online photos. Specifically, the photos containing target victim's face may be identified via face recognition; and the photos locations and timestamps may be revealed through various means such as geo-tags, metadata, or landmarks obtained from the advanced image processing tools. In Section 2, we demonstrate an example of such an attack to show its feasibility.

To better understand the aforementioned location related image privacy issues, we have conducted an exploratory user study among more than one hundred people to obtain their privacy opinions over a set of scenarios. The findings from the user study conform with our hypothesis that location sensitive photos could disclose too much of a person's privacy. Unfortunately, there have been very little works on how to help users mitigate such location-dependent image privacy. Thus, we propose a novel image privacy protection system, called LAMP (Location-Aware Multi-party Privacy), which aims to light up the location awareness for people during online image sharing. The LAMP system is based on a newly designed Location-Aware Multi-Party image (LAMPi) access control model that allows individual user to specify sensitive locations and timestamps for any photo in which their faces are identifiable. The proposed access control model goes beyond the traditional owner-centric privacy protection model, and the proposed LAMP system will facilitate social network providers to provide an equal protection for any people in the same photo. Specifically, the LAMP system as an add-on to existing social media sites will automatically detect the user's occurrences on photos to be posted online regardless the user is the photo owner or not. Once a user is identified and the location of the photo is deemed sensitive according to the user's privacy policy, the user's face will be replaced with a virtually generated human face. As we know, face blurring has been commonly used for privacy protection during photo sharing, while face replacement has been provided by existing apps mainly as a fun pastime activity. We hereby argue that the face replacement would be a better way to protect people's privacy as it offers several advantages which cannot be achieved by face blurring. First, it prevents attackers from using the latest image deblurring techniques [22], [26], [42] to uncover the people being protected. Second, the use of face replacement maintains the beauty and intact of the photo and reduces the chance of the photo to become a target of an attack. Considering that a photo with a blurred face and a photo with a swapped face, it is obvious that a blurred face has privacy concerns whereas a nicely swapped face may not even be noticed by the attacker.

The key contribution of the LAMP system is its scalability that it can enforce privacy protection for billions of users on social networks in real time. Such scalability level has never been reached in the past. Specifically, without our system in place, protecting the privacy of every person in a given photo will require the comparison of faces in the photo against the faces of all the users on social network which yields a huge number of comparisons (e.g., 2.4 billion Facebook users). We leverage the LAMPi policies to constraint the facial comparison to only a small group of users who specify privacy concerns on the photo location. This dramatically reduces the number of needed facial comparisons in orders of magnitude. Moreover, we also design face encoding and location policy indexes to further speed up the comparisons. We have implemented a prototype of the proposed system, and conducted a second user study. Our experimental results demonstrate the efficiency and effectiveness of our approach. In a summary, the contributions of our work are the following:

- We define a novel fine-grained location-aware multi-party image access control mechanism which breaks the existing limits of privacy protection only for photo owners and their friends by providing equal privacy protection to every identifiable individual in the photo instead of photo owners and their friends. Moreover, we consider the location-dependent privacy issues that are not studied in the past.
- We build a proof-of-concept application, the LAMP, to automate the location-aware multi-party privacy protection process. We design a graphic-based policy specification tool for users to easily specify sensitive locations at different granularity levels following our proposed LAMPi access control model. The algorithms designed for LAMP are tested to be efficient and scalable to deal with the huge number of photos and users on social media sites.
- We conducted two rounds of user studies involving more than 200 people to obtain valuable user opinions on location-dependent privacy issues and evaluate the effectiveness of privacy protection offered by our approach.

The rest of the paper is organized as follows. Section 2 presents the privacy risk analysis. Section 3 introduces our proposed LAMPi access control model and its implementation. Section 4 describes the LAMP system. Section 5 reviews privacy evaluations. Section 6 reports experimental studies. Section 7 discusses related work. Finally, Section 8 concludes the paper.

2 Image Sharing Risk Analysis

2.1 Threat Model

As the saying goes, a picture is worth a thousand words. An online photo/image can give out rich information about *who* are doing *what* at *when* and *where*. To better analyze the privacy risks incurred by image sharing, we classify image privacy based on two criteria: human-oriented, and context-oriented.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TDSC.2021.3132230, IEEE Transactions on Dependable and Secure Computing

Human-oriented image privacy can be further classified into three types:

(1) Photo owner's privacy: This type of privacy is currently preserved by allowing the photo owner to specify the groups of people who are permitted to access the shared photo. Most of the research works [13], [15], [29], [40] and commercial social media sites provide policy recommendation and configuration tools to achieve this. For example, Facebook users can choose to share the photos only with their friends but not friends of friends.

(2) Photo owner's friends' privacy: This refers to the privacy of the photo owner's friends who took the photo together with the photo owner. For example, Alice plans to post a party photo that includes her friend Kate. Kate is a shy girl who rarely shares photos online. Considering Kate's privacy, Alice may need to communicate with her before publishing the photo. However, such multi-party privacy issues are mainly discussed in academic world [4], [15], [16]. The current social media sites offer very little functionalities that support the multi-party privacy protection.

(3) Unaware people's privacy: This refers to the privacy of the people who are in the photo but are not aware of their photo being taken by others. For example, when someone took a selfie on the street, other pedestrians may be captured in the photo. These pedestrians will not know when and where their photos would appear on the Internet. Recently, an interview-based study [28] among college students found that undergraduates felt a heightened state of being surveilled by their peers when their photos were taken without their permissions and shared on social media by others. Participants in that study stated that they worried about being judged by others in a negative way based on the images which they were not aware of being taken.

Context-oriented image privacy can be further divided into two categories:

(1) Activity-dependant privacy: There are various scenarios when a person does not feel comfortable of sharing that moment with everyone. For example, a person in a funny costume may just want to share the photo with his/her close friends. In another case, a girl was drunk and someone else took her photo [28]. If the photo was posted online, it could lead to misjudgement of the girl and damage her general reputation. News also reported that some people were fired due to online photos. One case is that a fireman took a sick day off for attending an event, and he was later fired because his supervisor saw the event photo and identified him [27].

(2) Location-dependant privacy: A photo can leak location information of a person in many ways. The photo's embedded EXIF (Exchangeable image file format) [7] is a direct source that tells the date and GPS coordinates a photo was taken. Although some social media sites like Facebook and Instagram stripped the metadata when publishing the photos, they store the metadata in a separate database. If a hacker gains the access to these databases, it is even easier for them to track users since they now just need to look at the collection of metadata from all photos without spending much time on extracting metadata or analyzing photos one by one. Besides metadata, the photo itself may tell where the location is. Advanced image processing algorithms can identify the landmarks and the street signs. Yet another way could be the crowd sourcing. People living in the neighborhood of the place where the photo was taken may easily spot familiar buildings on the photo. With this said, *posting photos without metadata is still not sufficient to guarantee the location-dependant privacy of people in the photo.*

Most existing works on image privacy mainly focus on protecting photo owner and their friends' privacy (detailed review can be found in Section 7). Very limited efforts have been devoted into the other equally important privacy issues. i.e., unaware people's privacy, context-oriented image privacy. Thus, in this work, we aim to design a new access control mechanism to protect people's privacy in the photos which were taken without their knowledge and permissions. Our approach is complementary to existing methods and aims to achieve a full spectrum of image privacy protection. To better motivate our work, we will first present a location tracking attack and an exploratory user study in the following.

2.2 Location Tracking Attack Using Photos

In this experiment, we attempt to track a target person through his/her online photos. The goal is to show that it is not a difficult task for an attacker to cyberstalk a person. To avoid legal issues with a randomly chosen normal person, we decided to select a prominent figure whose images are publicly available in different venues: Joe Biden. Also, we do not hack into any social site to obtain metadata, whereas attackers can certainly gain more information than us by doing so.

We wrote a script to automatically crawl Google images to obtain the target person's photos between 2000 and 2019. We collected around 30,000 photos for the target. From the collected photos, we further analyze the metadata. Although not all the photos contain the

Fig. 1: User Tracking Through Photo Metadata

metadata, it is still amazing that we were able to found 721 days of location and visiting time for the target. Based on the obtained information, we created a tracking map as shown in Figure 1, where each point on the figure shows the location of the target person and the color of the point indicates when the photo was taken.

It is worth noting that Google images may not return photos of a person if he/she is in the background. Even so, the photos obtained from Google images already reveal lots of location information of a person. When an attacker utilizes advanced image processing tools to look for any occurrences of a target (either foreground or background) and combines the knowledge of the metadata stolen from the social media providers, the target movement may be exposed in a much deeper level due to the prevalence of photographing nowadays. Considering that people who took photos of themselves know about the sensitivity of their current locations, whereas people who were in the background of others' photos have no idea their locations have been recorded, one idea in our proposed work is to replace the faces of people who are in the background of the photos to avoid undesired exposure. In this way, even if the attackers run the image processing tool and have all the metadata of photos, the targets' faces have already changed and would not be identifiable.

2.3 A User Study on Unexpected Privacy Disclosure

In order to better understand users' concerns on location-dependent image privacy and gauge their interests in our proposed privacy protection mechanism, we conducted an online user study on Mechanical Turk. The user study is fully anonymous and follows the IRB exempted project guidelines.

We have recruited total 111 participants, including 51 females and 55 males. 15% of them are between 18-25 years old, 41% are between 26-35 years, 23% between 36-45 years, 13% between 46-55, and 8% above 56. The age distribution conforms with the age groups of people who access the social media more often.

At the beginning of the user study, we asked participants if they were aware that online images may tell others where they were and what they were doing, and how much they valued their privacy especially location privacy. From the response, we find that more than 74% of participants were aware of the privacy issues incurred by online images, and more than 76% emphasized that location privacy is important.

From there, we presented 10 different scenarios to the participants and asked them if they would be concerned when their images and their locations are disclosed to unexpected parties. Specifically, each scenario is accompanied with a short paragraph of story and an image. The 10 scenarios were designed with the goal to cover various aspects of our daily life in a nutshell. From the participants' responses, we found that when a photo discloses a critical moment or location of a person without being noticed by the person, more people would hope their identities are protected during the sharing of such kinds of photos. For example, when someone (say Bob) is planning to switch a job and had a job interview at a restaurant, another customer who also dined at the restaurant took a photo of the restaurant with Bob and his interviewers in the background. The customer later shared his photo along with the restaurant review comments in a famous restaurant review website. If Bob's supervisor or colleagues saw the photo and recognized Bob and competitor company's people, that could raise unnecessary tension in Bob's current workplace. Therefore, we see that 92.5% of participants would like their identities be protected in this scenario, which is the scenario with the most concerns. The second mostly concerned privacy breach scenario is when someone's children and home location may be exposed to strangers. About 91.5% of participants desire an identity protection in this case. On the other hand, some scenarios that may not lead to severe consequences, such as having a personal trip and drinking at a bar, have received privacy concerns from a little fewer people, but still close to 70%. Overall, we can observe that majority of people are concerned when their photos being taken and published by others without their knowledge, especially those photos that disclose sensitive locations and reveal their private issues.

At the end of the survey, we also asked the following general question: "*Suppose you are depicted in a photo published on social media by a stranger while you are in a location where you wish not to be seen. If social media websites provided functionality for hiding your identity (e.g., face swapping) when you are in such photos, would you like to use this function?*" More than 93.4% of the participants said that they would like to use such kind of services, which indicates a promisingly high acceptance rate of our proposed system.

3 LAMPi Access Control Mechanism

In the previous section, we have discussed both human-oriented and context-oriented image privacy, among which context-oriented image privacy of unaware people is least protected in the literature. To fill the gap, we define the LAMPi (Location-Aware Multi-Party image) access control mechanism, complementing the traditional image access control. The LAMPi policies will allow users to specify location sensitivity at different scenarios and at different granularity levels. Its formal description is given below.

Definition 3.1. A LAMPi policy P of a user u consists of the following components:

- Location range (Loc): The range of locations protected by policy P.
- Location type (Typ): This indicates whether the location is given as a semantic location (denoted as 'S') or an exact address (denoted as 'E').

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TDSC.2021.3132230, IEEE Transactions on Dependable and Secure Computing

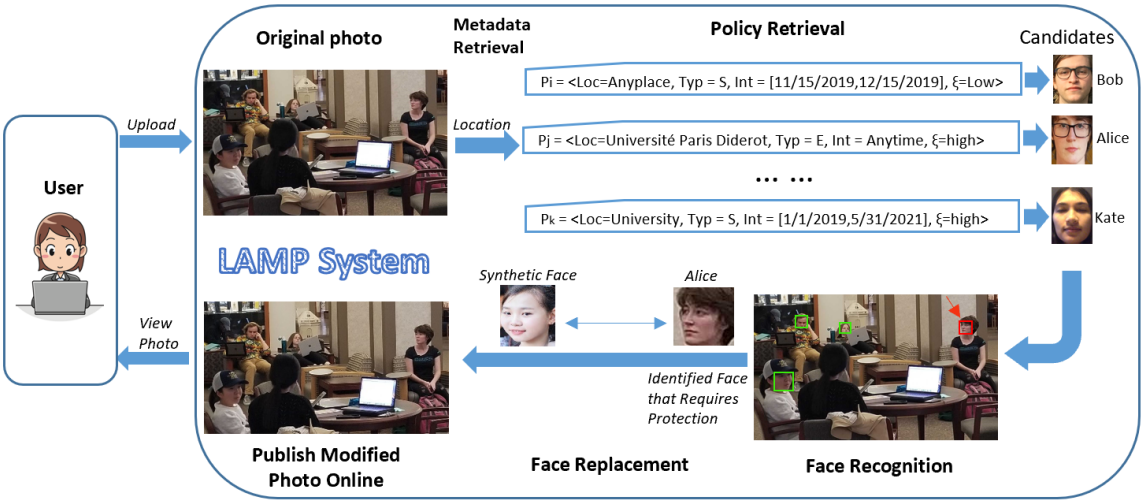

Fig. 2: The Privacy Protection Procedure in the LAMP System

• Time and date interval (Int): The time and date intervals during which the locations within Loc should be protected.
• Sensitiveness (ξ): The sensitiveness of the location Loc, which has two levels: "High" or "Low".

Since the LAMPi policies are mainly used to express the users' privacy concerns when they are unintentionally captured in others' photos, there is no need for the LAMPi policy to specify the sharing group like traditional image privacy policies. Instead, the user can specify how much they care about themselves being exposed in the particular location using the sensitiveness level. When the sensitiveness level is set to high, the user's face at that location will be replaced for protection even if the user's face on the photo is less identifiable, i.e., the face matching score is lower than a threshold (say 50%). When the sensitiveness level is low, our system will only replace the user's face when the face matching score is above the threshold. In this way, we minimize unnecessary image modifications. For locations which are not specified in the user's policies, the locations are simply considered not sensitive for that user.

The following are some example LAMPi policies which demonstrate the usage of different policy settings.

Example 3.1. *Kate does not want others to post photos which show her doing workout and sweating in a gym near her house. She can set her LAMPi policy as P_{Kate}=⟨Loc={gym_address}, Typ=E, Int=anytime, ξ=Low⟩, where 'Loc' is set to the gym's address, 'Typ=E' indicates that this is an exact location, 'Int=anytime' means Kate wants the privacy protection whenever she is in this gym, 'ξ=low' means that as long as she is not easily recognizable in the photo, Kate does not care about the photo being posted.*

Example 3.2. *Alice does not wish her face being recognized on any online photo that shows she is visiting a pub. She can thus set her LAMPi policy as P_{Alice} = ⟨Loc={pub}, Typ=S, Int={8pm-5pm on any day}, ξ=High⟩. That means if anyone attempts to post a pub photo with Alice in the background to a social network site, the social network site where Alice has registered and set the LAMPi policy will automatically replace Alice's face with a synthetic face to preserve Alice's privacy without affecting the photo owner's sharing experience.*

4 The LAMP System

In this section, we discuss how our proposed LAMP system helps preserve privacy of users who have no knowledge of their photos being posted by others. Figure 2 illustrates the data flow in the LAMP system. The LAMPi policy configuration function facilitates the users to specify the LAMPi policy through a graphic-based interface developed using Google Maps API. Users' policies will be indexed and stored in a policy database, and users' face features will be encoded to speed up the future face recognition. When someone wants to upload a photo to share, our LAMP system will first retrieve policies which mark the photo location as sensitive. Among the owners of the retrieved policies, we will further check if their faces depicted on the photo. If so, their faces will be replaced with synthetic faces to avoid undesired disclosure while maintaining the photo quality. The integration of LAMP to the existing social media platforms will be easy since they already have all the needed data to launch our system. If a social media platform is interested in adopting our system, it just needs to extend their current policy settings to include the location components for users to specify their location protection preferences. Then, they can adopt our system to manage LAMPi policies, perform the policy retrieval, facial recognition and privacy enforcement. In the following subsection, we elaborate the user identification and protection algorithms.

4.1 Identify Users with Location Privacy Concerns

In order to provide equal privacy protection to every person on the photo, an inevitable step is to know who

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TDSC.2021.3132230, IEEE Transactions on Dependable and Secure Computing

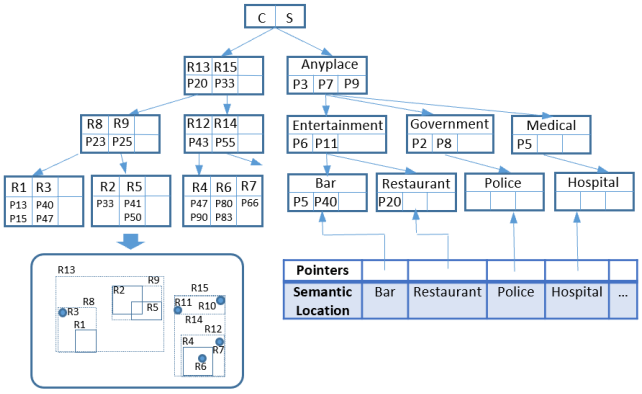

Fig. 3: An Overview of the DLP-tree

(especially those in the background) are in the photo. A brute force method to identifying people in the background of the photo will need to compare the person's face on the photo against all the other users' faces in the social network site (e.g., 2.4 billion users in Facebook), which will be extremely computationally expensive and hardly possible to maintain real-time response for the photo uploading request. This has been a very challenging problem in the image privacy protection as also pointed out by Ilia et al. in [19].

To overcome this challenge, we aim to reduce the total number of faces that need to be compared for each uploaded photo. We achieve this by introducing the LAMPi policies and leveraging the location constraints to reduce the search space to a tractable scope. Unless all the social network users have the privacy concern on the same location, the needed face comparison for a photo will be a much smaller set of the billion number of users.

4.1.1 Indexing LAMPi Policies

Given a photo and its location, we aim to quickly locate users who specify this location (based on the address) or this type of location (based on the semantic keywords) as sensitive so that later on we only need to compare these users' faces with those in the photo. To achieve this, we propose a hybrid data structure called DLP (Dual-Location-Policy)-tree that indexes policies stored in PostgreSQL. PostgreSQL was chosen for a variety of reasons. Most importantly is that PostgreSQL has great read and write performance compared to MySQL. Additionally, the PostGIS extension allows for geospacial data support that is incredibly useful when checking user and image locations. Finally, PostgreSQL easily supports concurrency in reading and writing that is crucial to speed up the face recognition process in our system as discussed in the later part of our approach.

The DLP-tree is an integration of the coordinate-based index, a keyword-based index, and a policy hash table, with the goal to speed up policy queries that contain either regions or semantic keywords in a coherent way. For the coordinate-based search, most of the existing spatial indexes will work. We select an R-tree like structure since R-trees [2] have been implemented in many commercial database systems. As for the semantic-based policy queries, we employ the hierarchical keyword organization. Therefore, the overall performance of this hybrid index structure is similar to any tree like data structure, which reduces the complexity of the exhaustive face comparison O(n) to $O(log(n))$.

The structure of the DLP-tree is illustrated in Figure 3. The DLP-tree consists of two main parts to index exact locations and semantic locations in the LAMPi policies, respectively. The left side of the DLP-tree organizes exact locations in a hierarchical way from nation (N), state (S), city (C) to address (A). Each entry in the leaf node is in the form of $\langle street, city, state, nation, PIDs \rangle$, where the first four attributes are the address used in user specified policies, and $PIDs$ is the list of IDs of policies that specify this address as sensitive. Nearby locations are grouped together in the same leaf node or sibling leaf nodes. An entry in an internal node is in the form of $\langle region, CPT, PIDs \rangle$, where $region$ described the region that covers all its child node pointed by CPT while $PIDs$ stores a list of IDs of policies that specify this region as sensitive. In this way, the internal nodes can efficiently facilitate the LAMPi policy search in a top-down manner. Note that the DLP-tree is different from a map since the DLP does not need to store all the places (e.g., all the addresses, all the cities) if no policies have been specified there yet.

The right side of the DLP-tree indexes semantic locations based on the hierarchical relationship among their semantic meanings. In particular, semantic locations are first classified into the basic categories, such as "bar", "hospital", "shopping mall" and "company". Basic categories are further classified into more generic categories which are the upper level of the DLP-tree. For example, basic categories like "bar" and "shopping mall" can be classified as a more generic category: "entertainment", and basic categories like "hospital", "clinic", and "urgent care" can be classified as "medical". Unlike the top-down search in the left side of the DLP-tree, the search in this part of the DLP-tree is from the leaf nodes to the root node. This is because users are allowed to specify their sensitive locations using semantic words at different granularity. Some users may specify "entertainment" in their policies while some users may specify only "bar" in their policies. Thus, the user policies are attached to different levels of the DLP-tree correspondingly. An entry in a node of this side of the DLP-tree is in the form of $\langle \varpi, PIDs, PPT \rangle$, where ϖ is the semantic keyword specified in the list of LAMPi policies (denoted as $PIDs$), and PPT is the pointer to the parent node in the DLP-tree. The policy hash table maps the semantic locations to the lowest level of the right part of the tree.

Here is the complete process for retrieving LAMPi policies of a given photo. When a user uploads a photo, we first extract its exact location from the photo's metadata and obtain its semantic meaning from map apps such as Google Map API. Then, we search the exact

location in the left part of the DLP-tree and the semantic location in the right part of the DLP-tree, respectively. Specifically, for the exact location, the search starts from the root of the DLP-tree and traverses to the left side to find the node whose region encloses the photo's exact location. We follow its child pointer to conduct the same boundary check in its child node until we reach the leaf level. If we find a matching location in the leaf node, we will collect the policies associated with the nodes along this search path. For example, if an exact location R_3 (as shown in Figure 3) is located, we will retrieve the policies associated with R_3 and its ancestor nodes R_8 and R_{13} since the two nodes enclose this location R_3 as shown in the bottom left part of Figure 3. The set of retrieved policies contains policies P_{40}, P_{47}, P_{23}, and P_{20}. As for the semantic meaning of the photo's location, the search starts from the policy hash table and finds the entries that contain the matching semantic location which points to the corresponding leaf node in the right part of the DLP-tree. For example, if the semantic location is "bar", we will follow the pointer in the hash table to locate the leaf node in the right part of the DLP-tree that contains "bar", and move up to its ancestor nodes "Entertainment" and "Anyplace" to retrieve the policies P_5. P_{40}, P_6, P_{11}, P_3, P_7, P_9. These policies all specify semantic locations that enclose "bar".

The owners of the policies retrieved from the DLP-tree will be compared against the people on the photo using face recognition as discussed in the following subsection.

4.1.2 Speed Up Face Recognition

To speed up the individual face comparison, we adopt two strategies. One is to pre-compute the user's face features when the user set up his/her LAMPi policies, which helps save the face recognition time during the photo uploading phase. The other is to employ multi-thread programming to conduct individual pairs of face recognition simultaneously.

Specifically, we calculate the face feature using the $load\_photo(image\_path)$ and $encode(image)$ functions in an open source python face recognition tool by Geitgey [6]. As reported, this face recognition tool has achieve 99.38% accuracy. The $load\_photo(image\_path)$ function loads an image from an image path using PIL. This image is then converted to a numpy array and returned. Then, the $encode(image)$ function takes the image numpy array as the input, and utilizes a pre-trained algorithm from dlib's facial recognition library to convert the image numpy array to a 128 dimension facial description. We then store the 128 dimension face features along with the user ID for the future face recognition.

Given a new photo, we first detect faces on the photo using the $locations(image)$ in the Geitgey's face recognition tool. For the detected faces, we calculate their face features in the similar way as aforementioned. Then, we compare the face features of those in the photo with those associated with the retrieved LAMPi policies that specify the photo location as sensitive. The face comparison is conducted using the $compare(source, destination, tolerance)$ function in the face recognition tool, which takes a source facial feature, a destination facial feature, and a tolerance. The tolerance value is set based on the sensitiveness value in the corresponding LAMPi policy. The function calculates the Euclidean distance between the facial feature vectors and checks if the distance is below the tolerance value. If the distance is smaller than the tolerance value, the two faces are considered match.

It is worth noting that social network sites can also use their existing face recognition tools when adopting our proposed privacy preservation function.

4.2 Protect Users with Location Privacy Concerns

After the face recognition, we will obtain a set of users who are in the photo and concerned about their privacy at the photo location. For these users, we propose to replace their faces so that they will not be recognizable even if the photo is shared publicly. There have been several face replacement algorithms and software [9], [36]. We revised an open source software for the face replacement [36] and integrated it into the LAMP system.

Figure 2 illustrates a running example of how the LAMP system protects privacy. Assume that a student reporter took some photos on campus and plans to post them online to show the student life at a university in Paris. When she uploaded the photos to the social media site that deployed the LAMP system, the LAMP system will check each photo and conduct the following privacy preservation procedure.

First, the LAMP extracts the photo's metadata to obtain the location information. In the example, the location information includes both the university address and the semantic keyword "university". Given the university address, the LAMP system will search the left part of the DLP-tree to locate the internal nodes and leaf node that contain this university address. Following the pointer from the located leaf entry to the policy hash table, we further retrieve the policies associated with this address. Assume that Alice's policy is associated with this address. Alice is a celebrity who was taking a year off to study abroad. Alice does not want to be followed or disturbed by her fans whenever she was at the university, and thus she has set her LAMPi policy as follows: P_j=⟨Loc=Universite Paris Diderot, Typ=E, Int=Anytime, ξ=High⟩. Note that Alice sets the time interval of protection to "Anytime" for convenience instead of using the exact time and date duration. She can simply remove this policy to release the protection after she returns home. Next, given the "university" as the semantic keyword, we search the hash table and find the pointer to the leaf node in the right part of the DLP-tree. Moving up in the DLP-tree, we will collect the policies associated with all the ancestor nodes of this semantic word "university", such as "Education" and "Anyplace". Assume

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TDSC.2021.3132230, IEEE Transactions on Dependable and Secure Computing

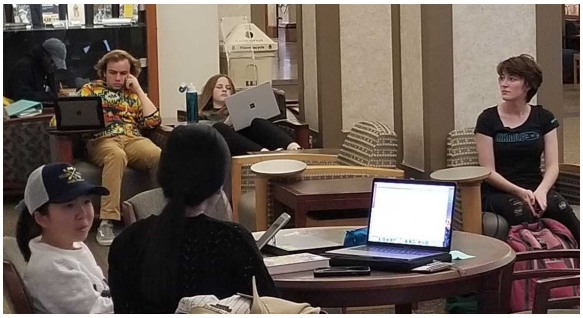

(a) Original Photo

(b) A Synthetic Face

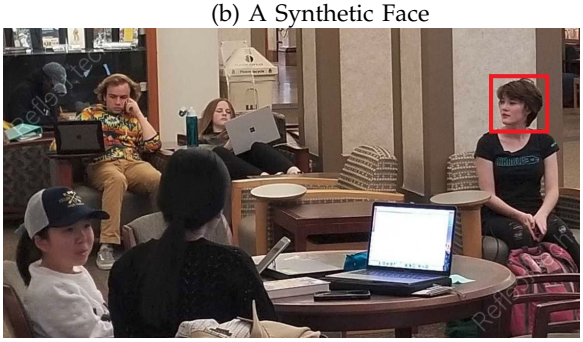

(c) Face Replaced

Fig. 4: Face Replacement

that Bob's policy aims to keep his photos at anyplace private during the period of 11/15/2019 to 12/15/2019. His policy may look like P_i=⟨Loc=Anyplace, Typ=S, Int=[11/15/2019,12/15/2019], ξ=Low⟩. As Bob's policy contains the keyword "Anyplace" which is in the ancestor node of "university", Bob's policy will also be retrieved for further examination.

From the set of retrieved policies, the LAMP system next loads the face feature vectors of these policies' owners. These candidate face features will be used for face recognition, i.e., compared with faces on the uploaded photos which are highlighted in boxes in Figure 2. In the example, Alice's face was identified (pointed by the red arrow in the figure). Since Alice has wished to remain private in this location, the LAMP system will then help privatize this user through face replacement. Specifically, a synthetic face will be automatically generated to replace Alice' face in the photo. Figure 4(a) shows the original photo uploaded by the user and Figure 4(c) shows the photo after the face in the red box is replaced using the synthetic face shown in Figure 4(b). Observe that face replacement takes care of the skin tones and facial expression. The modified photo looks very natural. Therefore, we expect that the modified face will not raise special attention from users who are viewing the photo.

Future work in this would be to utilize an artificial face generation method so that the face replacement can be conducted automatically without selecting a candidate face that does not have privacy concerns. Moreover, besides face replacement, it may be interesting to incorporate latest AI techniques such as CycleGAN [41] that replace other portions of a user such as a user's attire or even entire body to prevent someone who is familiar with the user from identifying the person. However, we also argue that since current face replacement results in a natural look, people viewing the photo do not know the photo has been modified or not, and may not try hard to match each person in a photo with someone they know.

5 Privacy Evaluation

In order to evaluate the effectiveness of our proposed privacy protection, we conducted another round of user study to see if participants are still able to identify the person who requires privacy protection and has been processed by our system. The idea is to present a set of testing photos and two reference photos to the participants, and ask them to try to identify the female and male references from the photos. Among the testing photos, some contain the referenced female and male without modification which represent the scenarios when people did not mark that location as sensitive; some contain the reference female and male with replaced faces which represent the scenarios when the people require privacy protection at that location; some contain blurred faces which represent the traditional privacy protection approach; and some do not contain any of the referenced people which are used as comparisons. The details of the user study are described as follows.

We have recruited total 102 participants on Mechanical Turk. There are 51 females and 51 males. Among them, 22% are 18 to 25 years old, 43% are between 26 and 35, 19% are between 36 and 45, 12% are 46 to 55, and 4% are above 56 years old. The user study is fully anonymous and follows the IRB exempted project guidelines.

Our study starts by telling participants that they will review images with numbered people within them. They will also have a reference photo for a person's face. They are told that if they can identify the reference person in the photo with a large degree of certainty, they just mark down the number of the person on the photo. They are also told that some photos may not include the reference person, and if they can not identify the reference person with a high degree of certainty, they need to input 0.

Each participant was asked to view 10 images. Each image contains 4 to 10 faces in the foreground and background. Half of these photos were asked about a male reference, and the remaining half were asked about a female reference. Photos for both male and female references included 1 photo with the reference not in the photo, 1 photo with the references whose face have been replaced, 1 photo with the references whose faces have been blurred, and 2 photos of the references in clear view. The photos in clear view give us an understanding of if the majority of participants can correctly identify

the person within the image without any changes, and then we can compare that with how they react when they see images that have been modified. Lastly, we ask if they noticed any image that has been altered in some way by presenting them unaltered photos and photos with replaced faces. We summarize participants' responses using the misidentification ratio which is the percentage of participants who did not correctly identify the reference person in the photo. From the study, we have the following two major findings.

Finding 1: *Photos with replaced faces have on average the highest misidentification ratio.* Specifically, as shown in Table 1, 84% of participants did not recognize the male reference in the photos where his face is replaced. Similarly for the female reference, 77% of participants did not recognize her in the photo with her face swapped. Both ratios are higher than those for the photos with blurred faces. This is possibly because when a face is blurred in the photo, the participants of the study clearly know which face to examine, and they can pay closer attention to the person of the blurred face including checking the hair style and other features. Thus, more participants were able to guess that the blurred faces were the references with high confidence. When it comes to the photos with swapped faces, the participants do not know which face is swapped. There is more work for them to examine all the faces in the photo in great details. Thus, fewer people were able to correctly identify the references. In addition, the misidentification ratios for photos with full face shown are much lower than the misidentification ratio of modified faces.

Finding 2: *Unaltered and face swapped photos are hard to distinguish*. At the end of the user study, we present an unaltered photo and a photo that contains a swapped face to the participants. For each photo, the participants were asked to check if the photo has been altered. The results are very interesting. Given an unaltered photo, 45% of participants said it had been altered. However, given the photo with a swapped face, only 32% of participants said it was altered. Such results could be because participants believed that some photos must be altered since the survey asked the question, so they were taking a guess whether a photo had been altered even though they were not 100% sure. This interesting result indicates that it would be really hard for human eyes to distinguish a face swapped photo from unaltered photos.

The result from our study demonstrates the effectiveness of privacy protection of our proposed use of face replacement. Moreover, in the real world scenario, when an attacker suspects a blurred face, he can utilize deblurring technique to further verify. In contrast, the replaced faces may not arouse much attention from viewers, and hence we expect higher misidentification ratio in the real social media sites. However, we would also like to point out the limitation of the face replacement strategy when the observer is one of the co-owners of a group photo. In that case, the observe may notice that one of his/her friends' faces have been modified as he/she knows who were taking the photo together.

TABLE 1: Privacy Evaluation Results

Image Type	Misidentification Ratio
Male reference with full face shown.	30%
Female reference with full face shown.	26%
Male reference with 40% of face showing.	76%
Female reference with 50% of face showing.	42%
Male reference not within photo.	44%
Female reference not within photo.	11%
Male reference with face blurred.	79%
Female reference with face blurred.	68%
Male reference with face swapped.	84%
Female reference with face swapped.	77%

6 Efficiency Evaluation

In this section, we aim to examine the efficiency and scalability of our proposed LAMP system. All of our experiments were conducted on custom Intel based computer with an i5-6600k at 3.9 GHz turbo clock, and 24 GB of 2133 Mhz RAM.

Users' LAMPi policies are synthetically generated. We randomly select exact locations and semantic locations along with time and date constraints for each user to create the LAMPi policies. In the experiments, we vary the number of sensitive locations (i.e., the number of policies) specified per user to test the efficiency of our algorithms.

We collected 5000 facial images from "Labeled Faces in the Wild" database [17]. These images are used to simulate uploaded images that we need to check the privacy compliance. Since not all the images that we collected associate with location information and our goal is to evaluate only the efficiency of our algorithms, we randomly generate location tags for each image. Each image is associated with both a coordinate location and 1 to 5 keywords indicating the semantic meanings of the location. The keywords for semantic locations are generated based on a four-level hierarchy similar to the one shown in Figure 3. In the experiments, we vary the total number of distinct locations to test their impact on our system performance.

(1) Varying the Number of Distinct Locations: In the first round of experiments, we evaluate the effect of the total number of distinct locations on policy retrieval. We vary the number of distinct locations from 100,000 to 1 million. We set the total number of users to 1 million and the number of policies per location to 1000. Under this setting, the total number of policies ranges from 100 million to 1 billion, and each user has 10 to 100 policies, i.e., each user specifies 10 to 100 locations as sensitive. We tested two different location distributions. One is

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TDSC.2021.3132230, IEEE Transactions on Dependable and Secure Computing

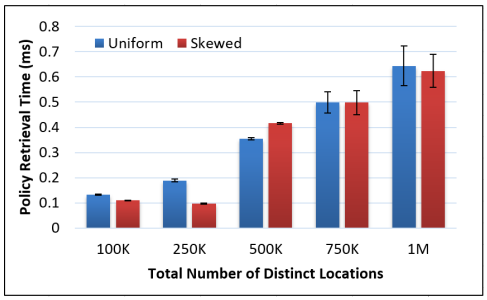

Fig. 5: Policy Retrieval Time When Varying the Total Number of Distinct Locations

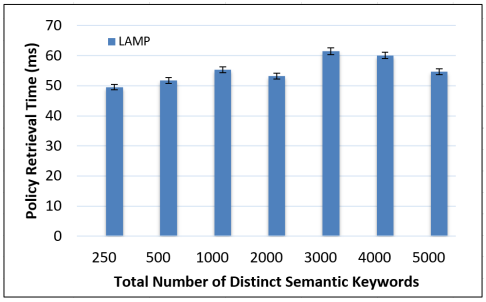

Fig. 6: Policy Retrieval Time When Varying the Total Number of Distinct Semantic Keywords

the uniform distribution whereby locations are evenly spread in the space. The other is a skewed distribution whereby locations are clustered around several hot spots. We randomly generated 10,000 location queries and recorded the total policy retrieval time.

Figure 5 shows the average of the total time needed to retrieve all the policies for a given photo location. Although the policy retrieval time increases with the number of locations, the LAMP system requires less than a millisecond to retrieve policies for a single photo location even when there are 1 million locations and 1 billion policies. This indicates the scalability of the underlying data structure, the DLP-tree. Specifically, considering 100 entries per node in the DLP-tree, 5 levels of the tree will be able to index about 1 billion policies. In other words, given a location either an address or a semantic keyword, we only need to check a few nodes (a few hundred entries out of 1 billion) in the DLP-tree to locate the group of policies that specify this location as sensitive. In addition, we can also observe that the performance on skewed data is sometime faster, which is likely because some of the randomly generated queries fall in the areas of fewer locations and retrieve fewer policies.

(2) Varying the Number of Distinct Semantic Keywords: We now proceed to test the performance of retrieving semantic-based policies. We also set the total number of users to 1M, and the total number of exact locations to 100K. The spatial distribution of the locations does not matter in this case as we are searching policies based on semantic keywords. The number of policies per location is 1000, which results in 100M total policies. Each location is associated with 5 semantic keywords. We vary the total number of distinct keywords from 250 to 5000. 5000 categories of places are considered an extreme case in the real world scenario, and hence we think it is sufficient to be used to test the scalability of our approach.

Figure 6 shows the total time taken to retrieve all the policies containing the semantic keywords associated with a given photo location. As we can see, the policy retrieval time stays relatively constant around 55ms with the increases of the keywords. This is due to two interacting factors. On one hand, the more the distinct keywords, the fewer number of policies contain the same keywords, and hence fewer policies to be retrieved. On the other hand, the more the distinct keywords, the more nodes in the semantic-based index to look up to locate the querying keywords, which increases the search time. The overall policy retrieval time shows the combined effects of these two factors. Moreover, we also observe that the semantic-based policy retrieval requires more time compared to the coordinate-based policy. This is mainly because the number of policies associate with the querying semantic keywords is hundreds of times more than that associated with a physical location. For example, when there are 5000 distinct semantic keywords and 100M policies each of which has 5 keywords, the average number of policies per semantic keyword reaches 100K which is much larger than 1000 policies per exact location.

(3) Performance of Face Recognition: After locating candidate users who specify the photo's location as sensitive, the next step is to check if the user actually appears on the photo through face recognition. Note that in our system, we only need to compare faces in the photo with candidate users who specify the photo's location as sensitive, rather than the users in the whole social network. Therefore, we vary the number of candidate users from 100 to 100K whereby the value 100K represents the relatively extreme case where 100K policies that may be associated with one semantic location as discussed in the previous experiment. Then, we select a group photo that contains 10 faces as the photo to be uploaded. That means, there will be up to 100K*10=1M face comparisons. We tested both linear and parallel face recognition performance. In Figure 7, it is not surprising to see that the face recognition time increases with the number of face comparisons. In all cases, the multi-thread face comparison yields reasonable time as it takes only 1 second for 1M face comparisons.

From the above results, we would like to highlight one of our key contributions – scalability. As for a system without any LAMPi policies in place, it will need to scan all the users' faces to identify the people in a given photo. Consider a photo that contains 10 faces as in the above experiment. When there are 1 billions of users on social networks, the time for facial comparison in a

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TDSC.2021.3132230, IEEE Transactions on Dependable and Secure Computing

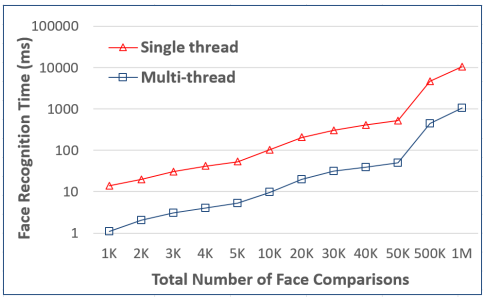

Fig. 7: Face Recognition Time

system without LAMPi policies could go up to 2.7 hours. Spending 2.7 hours for checking each photo is obvious not practical and will not be able to meet any user's image sharing needs. With the use of our approach, the number of faces to be compared is bounded by the number of policies per location which is not linear to the total number of users in the entire social networks, but determined by a much smaller population who are concerned about a specific location.

(4) Performance of Face Replacement: Finally, we look into the last step of the system which is the face replacement. Since face replacement is not the research focus of our work, we adopt an existing open source software [36] to gain the idea of the time needed for face swapping. Synthetic faces are generated beforehand. When there is a need for face replacement, a synthetic face of the same gender will be randomly selected for the replacement. Figure 8 shows the time needed to replace

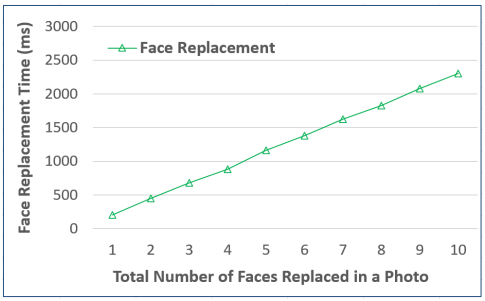

Fig. 8: Face Replacement Time

1 to 10 faces in a single photo. Replacing two faces need less than 500 millisecond while replacing 10 faces take just around 2 seconds.

Taking into account all the three steps in our system, i.e., policy retrieval, face recognition, and face replacement, the overall time to protect a photo with 10 faces will be around 3 to 4 seconds. We expect the whole process to be faster at the server side, and a few seconds of delay before the photo go live online would not be very noticeable by users considering there are also network and webpage refresh delays.

7 Related Work

(1) Image Privacy Protection for the Photo Owner

There have been a large body of image privacy protection works which focus on protecting the privacy of the photo owners [13], [29], [33], [37], [40]. These types of solutions include frameworks meant for suggesting privacy policies (i.e., which groups of people to share the image) for the photo owners at the time of the image being uploaded. For example, an early work [29] is by Squicciarini et al. who propose a privacy policy prediction system called A3P which considers image content, image metadata as well as the photo owners' historic privacy preferences when generating the policies. In [13], Hu et al. propose an interesting idea of calculating a level of sensitivity for each photo based on both user-defined rules and general rules discovered by machine learning. Users can then use the sensitivity levels as guidance for their privacy settings. In [40], Yuan et al. employ machine learning algorithms to analyze a social media user's photo sharing behavior, taking into account both the content of the image and the social context of the users who may see the photo. From that information, the system then determines whether or not to share the photo, entirely or partially, with a certain user.

(2) Multi-Party Image Privacy Protection

The protection for a single photo owner has later been extended to co-owners of the photo, i.e., people who took a group photo together. This type of multi-party privacy protection [18], [31] is typically achieved by considering privacy preferences of each party, solve policy conflicts among multiple parties, and then blur the faces with privacy restriction [14]–[16], [19], [21], [38]. For example, Hu et al. [15] define an access control model to capture the multiparty authorization requirements, based on which they develop multiparty policy specification scheme and algorithms to solve policy conflicts among multiple parties. To make the privacy preference aggregation more dynamic, Such et al. [30] propose a computational mechanism to resolve conflicts for co-owners of photos. Ilia et al. [19] propose a new way for multi-party privacy protection. They employ face recognition to automatically detect faces on the photo, and present the photo with the restricted faces blurred out. Vishwamitra et al. [35] also developed a multiparty access control model for collaborative access control for friends in the same photo. They implemented their system as Facebook application and adopted face blurring for privacy enforcement. Mosca et al. [24] propose an interesting idea that adds the consideration of moral values of users during the negotiation process of the multiuser privacy agreement. Most recently, deep learning techniques have been leveraged to directly learn the privacy settings for newly uploaded images [38], [39]. Then the identified privacy-sensitive objects will be blurred for privacy protection. The aforementioned works focus only on co-owners of a photo. None of them addresses the privacy concerns of people in the background of the photo who are not aware of their photos being taken by strangers. They do not specifically

consider the location privacy incurred by photo sharing either. Moreover, these existing works utilize obscured images for privacy protection, which may affect the utility and aesthetic of images as pointed out by recent studies [10], [11]. Our work adopts face replacement which preserves the quality of the images.

Very limited efforts have been devoted into privacy protection of people who occur in the background of others' photos like what we discuss in our work. As acknowledged in [19], identifying a person who is not related to the photo owner, i.e., not in the photo owner's contact list, would require a huge amount of computing resources due to the need to scan the whole enormous social network user set. A related work by Henne et al. [12] needed up to 1 hour to check an image and notify the bystanders. It detects only 50%-70% privacy violations in many cases and did not enforce the protection. In [25], Olteanu et al. propose a system that can ask the photo owner to label the people in the photo to seek their sharing consents. The process takes about 5s to identify a face among 10K registered users. Their processing time is linear to the number of users, which is not as scalable as our system whereby the number of faces to be compared is much fewer and is not directly determined by the total number of users. Moreover, they do not consider location related privacy issues.

(3) Privacy Issues Regarding Photo Metadata

Besides achieving protection through proper policy configurations, recent studies also look into potential privacy breach caused by metadata associated with photos [1]. Metadata like geotags and timestamps can easily disclose a person's location information, and multiple photos with geo-tags and timestamps may be used to track a person. To prevent undesired exposure, researchers [8] have proposed to remove metadata. However, such strategy may not be sufficient since the context of the photo may still reveal the location with the advance of the image processing technology. Our approach takes another route by hiding the person's face so as to avoid any location privacy breach. In another work [20], Chandra et al. developed a mobile app which can detect human subjects and issue a privacy alert if the location is sensitive. Their location privacy protection component is relatively preliminary which simply stores users' sensitive locations as link lists.

The Uniqueness of Our Work

To sum up, our work is different from existing works as we integrate the following new aspects into our system. First, we aim to protect location privacy of every human subject depicted on a photo regardless he/she is the owner of the photo or happens to appear in the background of the photo. Second, we achieve scalability of privacy protection without affecting the photo sharing experience. We design a highly efficient and scalable approach that minimizes the needed facial comparisons and is able to enforce privacy protection for billions of users on social networks in real time.

8 Conclusion

In this paper, we propose a novel idea to address the increasing concerns of location tracking of an individual through online images posted by others. Specifically, we define a new access control model, namely Location-Aware Multi-Party image (LAMPi) access control, which goes beyond the traditional access control that offers protection to only the owners of the image. Our proposed LAMPi access control mechanism provides equal privacy protection to every human subject on an online photo, no matter the human subject is the owner of the image or not, is at the foreground or background of the image. We also design an efficient policy management system that leverages policy indexing techniques and uses face replacement as policy enforcement, and we achieve the privacy protection in real time of photo uploading process. Our user studies and experimental results on the system prototype demonstrate both effectiveness and efficiency of our approach.

Acknowledgment

This work is supported by NSF projects DGE-1946619 and CNS-2027398.

References

[1] K. Albrecht and L. McIntyre. Psst...your location is showing!: Metadata in digital photos and posts could be revealing more than you realize. *IEEE Consumer Electronics Magazine*, 4(1):94–96, 2015.

[2] Norbert Beckmann, Hans-Peter Kriegel, Ralf Schneider, and Bernhard Seeger. The r*-tree: An efficient and robust access method for points and rectangles. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*, pages 322–331, 1990.

[3] Andrew Besmer and Heather Richter Lipford. Privacy perceptions of photo sharing in facebook. 01 2008.

[4] Andrew Besmer and Heather Richter Lipford. Moving beyond untagging: Photo privacy in a tagged world. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '10, pages 1563–1572, New York, NY, USA, 2010. ACM.

[5] John Boone. Florida teacher olivia sprauer fired for bikini modeling pics. Website, 10 2013.

[6] Adam Geitgey. Facial recognition. https://github.com/ageitgey/face_recognition/blob/master/LICENSE, 2017. MIT License.

[7] Thomas Germain. How a photo's hidden 'exif' data exposes your personal information. https://www.consumerreports.org/privacy/what-can-you-tell-from-photo-exif-data/, 2019.

[8] Kambiz Ghazinour and John Ponchak. Hidden privacy risks in sharing pictures on social media. *Procedia Computer Science*, 113:267–272, 12 2017.

[9] Han Guo, Dongmei Niu, Xiangyu Kong, and Xiuyang Zhao. Face replacement based on 2d dense mapping. In *Proceedings of the 2Nd International Conference on Image and Graphics Processing*, ICIGP '19, pages 23–28, New York, NY, USA, 2019. ACM.

[10] Rakibul Hasan, Eman Hassan, Yifang Li, Kelly Caine, David J Crandall, Roberto Hoyle, and Apu Kapadia. Viewer experience of obscuring scene elements in photos to enhance privacy. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, pages 1–13, 2018.

[11] Rakibul Hasan, Yifang Li, Eman Hassan, Kelly Caine, David J. Crandall, Roberto Hoyle, and Apu Kapadia. *Can Privacy Be Satisfying? On Improving Viewer Satisfaction for Privacy-Enhanced Photos Using Aesthetic Transforms*, pages 1–13. 2019.

[12] Benjamin Henne, Christian Szongott, and Matthew Smith. Snapme if you can: Privacy threats of other peoples' geo-tagged media and what we can do about it. pages 95–106, 04 2013.

[13] Donghui Hu, Fan Chen, Xintao Wu, and Zhongqiu Zhao. A framework of privacy decision recommendation for image sharing in online social networks. pages 243–251, 06 2016.
[14] H. Hu, G. Ahn, and J. Jorgensen. Enabling collaborative data sharing in google+. In *2012 IEEE Global Communications Conference (GLOBECOM)*, pages 720–725, 2012.
[15] H. Hu, G. Ahn, and J. Jorgensen. Multiparty access control for online social networks: Model and mechanisms. *IEEE Transactions on Knowledge and Data Engineering*, 25(7):1614–1627, 2013.
[16] Hongxin Hu and Gail-Joon Ahn. Multiparty authorization framework for data sharing in online social networks. In *Proceedings of the 25th Annual IFIP WG 11.3 Conference on Data and Applications Security and Privacy*, DBSec'11, pages 29–43, 2011.
[17] Gary B. Huang, Manu Ramesh, Tamara Berg, and Erik Learned-Miller. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. Technical Report 07-49, University of Massachusetts, Amherst, October 2007.
[18] Mathias Humbert, Benjamin Trubert, and Kévin Huguenin. A survey on interdependent privacy. *ACM Computing Survey*, 52(6), 2019.
[19] Panagiotis Ilia, Iasonas Polakis, Elias Athanasopoulos, Federico Maggi, and Sotiris Ioannidis. Face/off: Preventing privacy leakage from photos in social networks. In *Proceedings of the 22Nd ACM SIGSAC Conference on Computer and Communications Security*, pages 781–792, 2015.
[20] D. Keerthi Chandra, W. Chowgule, Y. Fu, and D. Lin. Ripa: Real-time image privacy alert system. In *2018 IEEE 4th International Conference on Collaboration and Internet Computing (CIC)*, pages 136–145, 2018.
[21] 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*, pages 154–162, 2019.
[22] Lerenhan Li, Jinshan Pan, Wei-Sheng Lai, Changxin Gao, Nong Sang, and Ming-Hsuan Yang. Blind image deblurring via deep discriminative priors. *International Journal of Computer Vision*, 127(8):1025–1043, Aug 2019.
[23] Pauline Morrissey. 6 people who were fired for social media posts. Website, 11 2019.
[24] Francesca Mosca and Jose M. Such. Elvira: An explainable agent for value and utility-driven multiuser privacy. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems*, pages 916–924, 2021.
[25] Alexandra-Mihaela Olteanu, Kévin Huguenin, Italo Dacosta, and J-P Hubaux. Consensual and privacy-preserving sharing of multi-subject and interdependent data. In *Proceedings of the 25th Network and Distributed System Security Symposium (NDSS)*, pages 1–16, 2018.
[26] Liyuan Pan, Richard Hartley, Miaomiao Liu, and Yuchao Dai. Phase-only image based kernel estimation for single image blind deblurring. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
[27] Lydia Price. 20 tales of employees who were fired because of social media posts. https://people.com/celebrity/employees-who-were-fired-because-of-social-media-posts/, 2016.
[28] Yasmeen Rashidi, Tousif Ahmed, Felicia Patel, Emily Fath, Apu Kapadia, Christena Nippert-Eng, and Norman Makoto Su. "you don't want to be the next meme": College students' workarounds to manage privacy in the era of pervasive photography. In *Fourteenth Symposium on Usable Privacy and Security (SOUPS 2018)*, pages 143–157, Baltimore, MD, 2018.
[29] Anna Cinzia Squicciarini, Smitha Sundareswaran, Dan Lin, and Josh Wede. A3p: Adaptive policy prediction for shared images over popular content sharing sites. In *Proceedings of the 22Nd ACM Conference on Hypertext and Hypermedia*, HT '11, pages 261–270, 2011.
[30] Jose M Such and Natalia Criado. Resolving multi-party privacy conflicts in social media. *IEEE Transactions on Knowledge and Data Engineering*, 28(7):1851–1863, 2016.
[31] Jose M Such and Natalia Criado. Multiparty privacy in social media. *Communications of the ACM*, 61(8):74–81, 2018.
[32] Jose M. Such, Joel Porter, Sören Preibusch, and Adam Joinson. Photo privacy conflicts in social media: A large-scale empirical study. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pages 3821–3832, 2017.
[33] Gang Sun, Yuxia Xie, Dan Liao, Yu Hongfang, and Victor Chang. User-defined privacy location-sharing system in mobile online social networks. *Journal of Network and Computer Applications*, 86, 11 2016.
[34] Kurt Thomas, Chris Grier, and David Nicol. unfriendly: Multi-party privacy risks in social networks. pages 236–252, 07 2010.
[35] Nishant Vishwamitra, Yifang Li, Kevin Wang, Hongxin Hu, Kelly Caine, and Gail-Joon 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 (SACMAT)*, pages 155–166, 2017.
[36] wuhuikai, XiangFugui, and niczem. Faceswap. https://github.com/wuhuikai/FaceSwap, 2018.
[37] Xi Xiao, Chunhui Chen, Arun Sangaiah, Guangwu Hu, Runguo Ye, and Yong Jiang. Cenlocshare: A centralized privacy-preserving location-sharing system for mobile online social networks. *Future Generation Computer Systems*, 86, 02 2017.
[38] 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*, 12(5):1005–1016, 2017.
[39] Jun Yu, Zhenzhong Kuang, Baopeng Zhang, Wei Zhang, Dan Lin, and Jianping Fan. Leveraging content sensitiveness and user trustworthiness to recommend fine-grained privacy settings for social image sharing. *IEEE transactions on information forensics and security*, 13(5):1317–1332, 2018.
[40] Lin Yuan, Joël Theytaz, and Touradj Ebrahimi. Context-dependent privacy-aware photo sharing based on machine learning. In *IFIP Advances in Information and Communication Technology*, pages 93–107, 2017.
[41] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *IEEE International Conference on Computer Vision (ICCV)*, pages 2242–2251, 2017.
[42] Li Zhu, Long Jin, Jihua Zhu, Zhongyu Li, Zhiqiang Tian, and Huimin Lu. Blind image deblurring based on local rank. *Mobile Networks and Applications*, Sep 2019.

Joshua Morris is currently pursuing his Ph.D. in Computer Science at the University of Missouri - Columbia. He received his Bachelor's degrees in computer science at Missouri University of Science and Technology. His research interests focus on privacy protection for users on social networks.

Sara Newman received her Bachelor's degree in Computer Science at Missouri University of Science and Technology and is currently pursuing her PhD in Computer Science at University of Missouri - Columbia. Her research interests include the robustness of, and various adversarial attacks against, deep neural networks, particularly in the context of image classification.

Kannappan Palaniappan is a professor at University of Missouri - Columbia. He received his PhD from the University of Illinois - Urbana Champaign. His current, multidisciplinary interests in computer vision, high performance computing, data science and biomedical image analysis range across orders of scale from subcellular microscopy at the molecular level to aerial and satellite remote sensing imaging at the macro level.

Jianping Fan is Lenovo Vice President and Head of Artificial Intelligence Lab of Lenovo Research. Before joining Lenovo, Dr. Fan was a tenured full professor at the University of North Carolina at Charlotte. He received his PhD degree in optical storage and computer science from Shanghai Institute of Optics and Fine Mechanics of Chinese Academy of Sciences in 1997. He was a researcher at Fudan University (Shanghai, China) form 1997 to 1998. His research interests include image/video privacy protection, computer vision, statistical machine learning and deep learning.

Dan Lin is an associate professor and Director of I-Privacy Lab at University of Missouri - Columbia. She received her PhD degree in Computer Science from National University of Singapore in 2007, and was a postdoctoral research associate at Purdue University for two years. Her research interests cover many areas in the fields of information security and artificial intelligence. She is an IEEE senior member.