

# Privacy-Preserving Visual Learning Using Doubly Permuted Homomorphic Encryption

Ryo Yonetani  
The University of Tokyo  
Tokyo, Japan  
yonetani@iis.u-tokyo.ac.jp

Vishnu Naresh Boddeti  
Michigan State University  
East Lansing, MI, USA  
vishnu@msu.edu

Kris M. Kitani  
Carnegie Mellon University  
Pittsburgh, PA, USA  
kkitani@cs.cmu.edu

Yoichi Sato  
The University of Tokyo  
Tokyo, Japan  
ysato@iis.u-tokyo.ac.jp

## Abstract

We propose a privacy-preserving framework for learning visual classifiers by leveraging distributed private image data. This framework is designed to aggregate multiple classifiers updated locally using private data and to ensure that no private information about the data is exposed during and after its learning procedure. We utilize a homomorphic cryptosystem that can aggregate the local classifiers while they are encrypted and thus kept secret. To overcome the high computational cost of homomorphic encryption of high-dimensional classifiers, we (1) impose sparsity constraints on local classifier updates and (2) propose a novel efficient encryption scheme named doubly-permuted homomorphic encryption (DPHE) which is tailored to sparse high-dimensional data. DPHE (i) decomposes sparse data into its constituent non-zero values and their corresponding support indices, (ii) applies homomorphic encryption only to the non-zero values, and (iii) employs double permutations on the support indices to make them secret. Our experimental evaluation on several public datasets shows that the proposed approach achieves comparable performance against state-of-the-art visual recognition methods while preserving privacy and significantly outperforms other privacy-preserving methods.

## 1. Introduction

An enormous amount of photos and videos are captured everyday thanks to the proliferation of camera technology. Such data has played a pivotal role in the progress of computer vision algorithms and the development of computer vision applications. For instance, a collection of photos up-

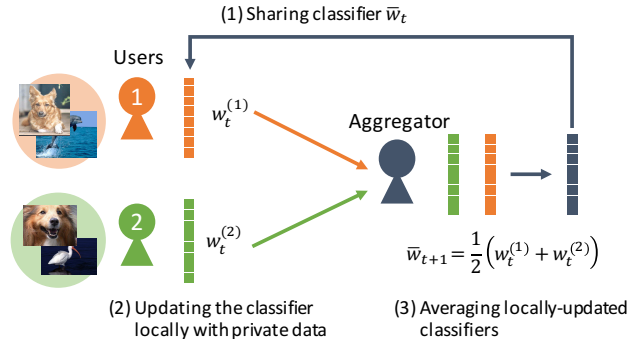


Figure 1. **Learning a Classifier by an Aggregator and Users.** (1) Given classifier  $\bar{w}_t$  shared by an aggregator (e.g., the administrator of a cloud storage); (2) users update it locally using private data and send updated classifier  $w_t^{(1)}, w_t^{(2)}$  to the aggregator; and (3) the aggregator collects and averages the locally-updated classifiers to obtain a better one:  $\bar{w}_{t+1} = \frac{1}{2}(w_t^{(1)} + w_t^{(2)})$ . How can users keep  $w_t^{(1)}, w_t^{(2)}$  secret from the aggregator and other users?

loaded to Flickr has enabled large-scale 3D reconstruction [2] and city attribute analysis [63]. More recently, a large-scale YouTube video dataset has been released to drive the development of the next phase of new computer vision algorithms [1]. Arguably, much of the progress made in computer vision in the last decade has been driven by data that has been shared publicly, typically after careful curation.

Interestingly, we do not always share captured data publicly but just store it on a personal storage [32], and such private data contains various moments of everyday life that are not found in publicly shared data [5, 11, 31, 33]. Accessing extensive amounts of private data could lead to tremendous advances in computer vision and AI technology. However there is a dilemma. On one hand, leveraging this diverse and

abundant source of data could significantly enhance visual recognition capabilities and enable novel applications. On the other hand, using private data comes with its own challenge: *how can we leverage distributed private data sources while preserving the privacy of the data and its owners?*

To address this challenge, we develop a privacy preserving framework that can learn a visual classifier from distributed private data. As illustrated in Figure 1, we consider the following two types of parties: users who own private data and an aggregator who wishes to make use of them. They collaboratively learn a classifier as follows: (1) given a classifier shared by the aggregator; (2) users update it locally with their own data; and (3) the aggregator collects and averages the locally-updated classifiers to obtain a better one. These steps iterate multiple times until the aggregator and the users get a high-performance classifier which has been trained on a large collection of private data. As a practical example, suppose that people capture life-logging videos passively with a wearable camera and store them on a cloud storage. By allowing the cloud administrator to curate these videos, he will be able to learn activity classifiers [19, 20, 37, 45, 48, 52] from realistic and diverse data. However, because such videos inevitably include any private moment that came into the view of cameras, the administrator should ensure that no private information in the videos is leaked during and after the learning procedure.

One possible solution for preserving privacy is to perturb classifier weights by adding noise [43, 46]. Users could also sanitize their data by detecting and removing sensitive images [19, 35, 55] or transforming images into low-resolution ones [49]. We stress here that these approaches however involve strong trade-off between data degradation and data utility – the more information we hide, the less effective the data is for learning classifiers.

An attractive solution for privacy-preserving learning without such trade-off is the use of *homomorphic cryptosystems* — a class of cryptosystems that can allow for basic arithmetic operations over encrypted data [22]. For example, the Paillier cryptosystem [42] calculates the encrypted sum of data as follows:  $\zeta(x^{(1)} + x^{(2)}) = \zeta(x^{(1)}) \zeta(x^{(2)})$  where  $\zeta(x)$  is the ciphertext of value  $x$  and *never reveals the original  $x$*  without a decryption key. This unique property of homomorphic cryptosystems has enabled a variety of privacy-preserving machine learning and data mining techniques (e.g., [7, 14, 25, 43, 54, 61]). By homomorphically encrypting locally-updated classifiers, the aggregator can average them while ensuring that the classifiers never expose private information about the trained data.

However, homomorphic encryption involves prohibitively high computational cost rendering it unsuitable for many tasks that need to learn high-dimensional weights. For example, in our experiments, we aim to learn a classifier for recognizing 101 objects in the Caltech101

dataset [21] using 2048-dimensional deep features. Paillier encryption with a 1024-bit key takes about 3 ms for a single weight using a modern CPU<sup>1</sup>. Therefore, encryption of each classifier will require more than 10 minutes.

To leverage homomorphic encryption for our privacy-preserving framework, we present a novel encryption scheme named *doubly-permuted homomorphic encryption (DPHE)*, which allows high-dimensional classifiers to be updated securely and efficiently. The key observation is that, if we enforce sparsity constraints on the classifier updates, the updated weights can be decomposed into a small number of non-zero values and corresponding support indices. DPHE then applies homomorphic encryption only to non-zero values. Because the support indices could also be private information (e.g., when classifier weights take binary values [10, 13, 27]), we design a shuffling algorithm of the indices based on two different permutation matrices. This shuffling ensures that 1) only the data owners can identify the original indices and 2) an aggregator can still average classifiers while they are encrypted and shuffled. By adopting a sparsity constraint of more than 90%, DPHE reduces the classifier encryption time of the previous example on Caltech101 to about one minute.

We evaluate our visual learning framework on a variety of tasks ranging from object classification on the classic Caltech101/256 datasets [21, 26] to more practical and sensitive tasks with a large-scale dataset including face attribute recognition on the CelebA dataset [40] and sensitive place detection on the Life-logging dataset [19]. Experimental results demonstrate that our framework performs significantly better than several existing privacy-preserving methods [43, 46] in terms of classification accuracy. We also achieve comparable classification performance to some state-of-the-art visual recognition methods [28, 40, 62].

## Related Work

Privacy preservation has been actively studied in several areas including cryptography, statistics, machine learning, and data mining. One long-studied topic in cryptography is secure multiparty computation [59], which aims at computing some statistics (e.g., sum, maximum) securely over private data owned by multiple parties. More practical tasks include privacy-preserving data mining [3, 4], ubiquitous computing [36], and social network analysis [64].

One popular technique for preserving privacy is by perturbing outputs or intermediate results of algorithms based on the theory of differential privacy (DP) [15, 16]. DP was classically introduced “to release statistical information without compromising the privacy of the individual respon-

<sup>1</sup>A single CPU of a MacBook Pro with a 2.9GHz Intel Core i7 was used with `python-paillier` (<https://github.com/nlanalytics/python-paillier>) and `gmpy2` (<https://pypi.python.org/pypi/gmpy2>; a C-coded multiple-precision arithmetic module.)

dents [16],” where the individual respondents are samples in a database and the statistical information is a certain statistic (e.g., average) of those samples. This can be done by adding properly-calibrated random noise to the statistical information such that one cannot distinguish the presence of an arbitrary single sample in the database. DP was then adapted and used also in the context of privacy-preserving machine learning, e.g., [8, 9, 12, 43, 46, 51]. In the classification cases, a training dataset and classifier weights are referred to as a database and its statistical information, respectively. By adding properly-calibrated random noise to the classifier weights (a.k.a. output perturbation [9]) or objective functions (objective perturbation), DP-based classification methods aim to prevent the classifier from leaking the presence of individual samples in the training dataset. While these methods are computationally efficient, the perturbed results are not the exact same as what could be originally learned from the given data. Moreover, the scale of noise typically increases exponentially to the level of privacy preservation (e.g., [17]). Perfect privacy preservation can never be achieved as long as one wishes to get some meaningful results from data.

Another privacy-preserving technique is the use of homomorphic cryptosystems [41], which we will study in this work. The homomorphic cryptosystems also enable a variety of privacy-preserving machine learning and data mining [7, 14, 25, 39, 43, 44, 54, 57, 61] because they can make some weights perfectly secret by encrypting them. Unlike the DP-based approaches, homomorphic encryption does not compromise the accuracy of algorithm outputs at the expense of its high computational cost. Our key contribution is to resolve this computational cost problem by introducing a new efficient encryption scheme.

Finally, studies on privacy preservation in computer vision are still limited. Recent work includes sensitive place detection [19, 35, 55] and privacy-preserving activity recognition [49]. These methods are designed to sanitize private information in a dataset and involve the potential trade-off between data security and utility. Another relevant topic is privacy-preserving face retrieval based on homomorphic encryption [18, 50]. Because these methods encrypt all data samples, they can be applied only to thousands of data, while our approach can accept hundreds of thousands of data as input based on a distributed learning framework.

## 2. Efficient Encryption Scheme for Privacy-Preserving Learning of Classifiers

The goal of this work is privacy-preserving learning of visual classifiers over distributed image data privately owned by people. As described earlier, our framework involves two types of parties: **users**  $\{U^{(n)} \mid n = 1, \dots, N\}$  who individually own labeled image data and an **aggregator** who exploits the data for learning classifiers while en-

suring that no privacy information about the data is leaked during and after the learning procedure.

### 2.1. Privacy-Preserving Learning Framework

We first describe how classifiers can be learned from distributed data. Let us denote classifier weights at step  $t$  by  $D$ -dimensional vector  $\bar{w}_t \in \mathbb{R}^D$ . We assume that the initial weight vector  $\bar{w}_0$  is provided by an aggregator, for example by using some public data already published on the web. As shown in Figure 1, (1) weight vector  $\bar{w}_t$  is first shared from an aggregator to users. (2) User  $U^{(n)}$  updates  $\bar{w}_t$  locally with own labeled data and sends updated classifier  $w_t^{(n)}$  to the aggregator, and (3) the aggregator averages locally-updated classifiers for the initial weights of the next step:  $\bar{w}_{t+1} = \frac{1}{N} \sum_n w_t^{(n)}$ .

This framework, however, is vulnerable to potential privacy leakage at multiple places. An aggregator can have access to  $\{w_t^{(n)} \mid n = 1, \dots, N\}$  each of which reflects user’s private data. In addition, if a communication path between user and aggregator sides is not secure,  $w_t^{(n)}$  could also be intercepted by another user via man-in-the-middle attacks. If  $w_t^{(n)}$  is updated via stochastic gradient descent (SGD), the difference of two classifier weights  $w_t^{(n)} - \bar{w}_t$  could be used to identify a part of the trained private data. Some concrete examples of how private data can be leaked are shown in our supplementary material.

Our new encryption scheme, the doubly-permuted homomorphic encryption (DPHE), is designed to prevent privacy leakage by homomorphically encrypting  $w_t^{(n)}$ . By imposing sparse constraints on local classifier updates (e.g., [24, 56, 65]), the DPHE exploits the sparsity to encrypt high-dimensional classifiers efficiently. In Section 2.2, we first explain briefly an algorithm to average local classifiers securely based on the Paillier cryptosystem [42]. Then Section 2.3 and Section 2.4 present how DPHE ‘doubly-permutes’ sparse data on a step-by-step basis. For simplicity, we will focus on one particular step by omitting time index  $t$ , that is, we consider the problem of computing  $\bar{w} = \frac{1}{N} \sum_n w^{(n)}$  securely.

### 2.2. Secure-Sum with Paillier Encryption

Let us start from a simple secure-sum algorithm where each of  $N$  users  $\{U^{(n)} \mid n = 1, \dots, N\}$  sends a single value  $w^{(n)} \in \mathbb{R}$  to an aggregator, and the aggregator computes their sum  $w_{\text{sum}} = \sum_n w^{(n)}$ , provided that  $w^{(n)}$  is never exposed to any other party but user  $U^{(n)}$ . To accomplish this, we utilize the Paillier cryptosystem [42]. Let us denote by  $\zeta(w)$ , the ciphertext of value  $w \in \mathbb{R}^2$ . As shown in [42],  $\zeta(w)$  cannot be inverted to the original value  $w$  without a

<sup>2</sup>Although the Paillier cryptosystem originally works only on a positive integer, we can deal with real values  $w \in \mathbb{R}$  by scaling them to be a positive integer before encryption and rescaling them after decryption properly.

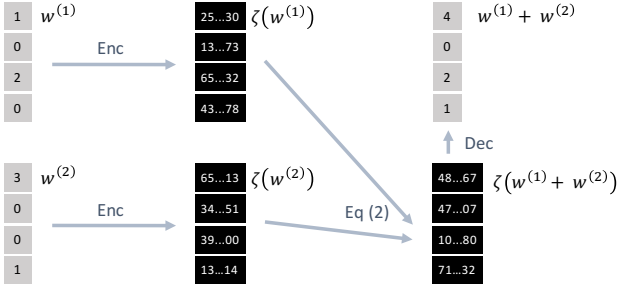


Figure 2. **Secure-Sum with Paillier Encryption.** Vectors  $w^{(1)}, w^{(2)}$  are first encrypted with the Paillier cryptosystem. Ciphertexts  $\zeta(w^{(1)}), \zeta(w^{(2)})$  are then used to compute encrypted sum  $\zeta(w^{(1)} + w^{(2)})$ . After the decryption we obtain  $w^{(1)} + w^{(2)}$ .

decryption key. Also, because the encryption always involves generation of random numbers, multiple encryptions of the same values will result in different ciphertexts.

To enable this secure-sum algorithm, we introduce a **key generator** who issues encryption and decryption keys and distributes the only encryption key to all of the parties. User  $U^{(n)}$  then sends to the aggregator  $\zeta(w^{(n)})$  that is encrypted with the given encryption key. Using the Paillier encryption, the product of two ciphertexts results in a ciphertext of the two original plaintexts, *i.e.*,  $\zeta(w^{(1)})\zeta(w^{(2)}) = \zeta(w^{(1)} + w^{(2)})^3$ . Therefore, the aggregator can compute the encrypted sum of values as follows:

$$\zeta(w_{\text{sum}}) = \zeta\left(\sum_n w^{(n)}\right) = \prod_n \zeta(w^{(n)}). \quad (1)$$

Finally, the aggregator asks the key generator for decrypting  $\zeta(w_{\text{sum}})$  and receives  $w_{\text{sum}}$ . As long as all the parties strictly follow this algorithm, privacy of  $U^{(n)}$  can be preserved by restricting the aggregator's access to  $w^{(n)}$ .

The straightforward extension of this secure-sum algorithm for averaging locally-updated classifiers,  $w^{(n)} \in \mathbb{R}^D$ , is to apply the Paillier encryption element-wise. In what follows, let  $w_{\text{sum}} = \bar{w}N \in \mathbb{R}^D$  and let  $\zeta(w)$  be a  $D$ -dimensional vector which elements are individually Paillier encrypted. Similar to Equation (1),  $w_{\text{sum}}$  can be computed securely as follows (see Figure 2 for a two-vector case):

$$\zeta(w_{\text{sum}}) = \zeta\left(\sum_n w^{(n)}\right) = \odot_n \zeta(w^{(n)}), \quad (2)$$

where  $\odot_n$  is the element-wise product of multiple vectors

<sup>3</sup>Specifically, Paillier encryption is defined as  $\zeta(x) = g^x r^T \bmod T^2$  where the pair of values  $(g, T)$  is an encryption key and  $r$  is a random number generated per encryption. Given  $\zeta(x^{(1)}) = g^{x^{(1)}}(r^{(1)})^T \bmod T^2$  and  $\zeta(x^{(2)}) = g^{x^{(2)}}(r^{(2)})^T \bmod T^2$ , the encrypted sum is derived as follows:  $\zeta(x^{(1)})\zeta(x^{(2)}) = g^{x^{(1)}+x^{(2)}}(r^{(1)}r^{(2)})^T \bmod T^2 = \zeta(x^{(1)} + x^{(2)})$ . Further details can be found in [22, 42].

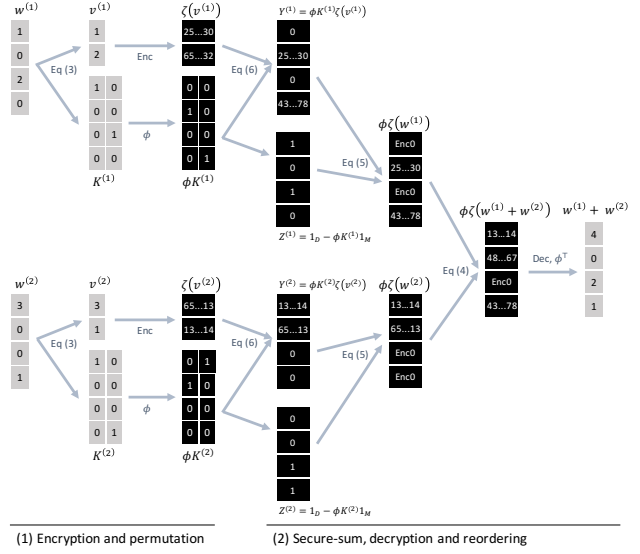


Figure 3. **Secure-sum with Efficient Encryption with Permutations.** (1) While applying the Paillier encryption only to  $v^{(n)}$ , we shuffle index matrix  $K^{(n)}$  with permutation matrix  $\phi$  to keep it secure. (2) An aggregator can compute encrypted and shuffled sum  $\phi\zeta(w^{(1)} + w^{(2)})$  and can ask a key generator for decryption and reordering to obtain  $w^{(1)} + w^{(2)}$ .

over index  $n$ . Then, the aggregator receives  $w_{\text{sum}}$  with help from the key generator and computes  $\bar{w} = \frac{1}{N}w_{\text{sum}}$ .

### 2.3. Encryption with Single Permutation

The extension based on Equation (2) would, however, become infeasible for learning high-dimensional classifiers due to the high computational cost of homomorphic encryption. DPHE can overcome this problem by applying the homomorphic encryption only to a limited part of classifier weights that are updated with sparse constraints.

Let us first introduce *encryption capacity*  $M$  that limits the number of values that are Paillier encrypted. We also denote by  $nz(w^{(n)})$ , the actual number of non-zeros in  $w^{(n)}$ . By choosing  $M$  such that  $nz(w^{(n)}) \leq M \leq D$ ,  $w^{(n)}$  can be decomposed into a pair of real-value vector  $v^{(n)} \in \mathbb{R}^M$  and binary matrix  $K^{(n)} \in \{0, 1\}^{D \times M}$ :

$$w^{(n)} = K^{(n)}v^{(n)}, \quad (3)$$

where  $v^{(n)}$  includes all the constituent non-zero values in  $w^{(n)}$ .  $K^{(n)}$  is an index matrix where the number of ones for each column is exact one and that for each row is at most one, and the  $m$ -th column of  $K^{(n)}$  indicates the index of the  $m$ -th value of  $v^{(n)}$  in vector  $w^{(n)}$ . Users encrypt the only  $v^{(n)}$  with the Paillier cryptosystem to obtain  $\zeta(v^{(n)})$ . If  $nz(w^{(n)}) \ll D$ , *i.e.*,  $w^{(n)}$  is sparse, we can make this encryption more efficient by choosing smaller  $M$ . Note that there are multiple decompositions that satisfy Equation (3), and users can arbitrarily choose one of them.

Now, the remaining problem is how to keep  $K^{(n)}$  secret. This matrix indicates non-zero indices of  $\mathbf{w}^{(n)}$  and could also be private information for users. For instance, when classifier weights take binary values (e.g., [10, 13, 27]), all non-zero values in  $\mathbf{w}^{(n)}$  become one and the private information is found in  $K^{(n)}$ . To solve this problem, we propose shuffling  $K^{(n)}$  with permutation matrix  $\phi \in \{0, 1\}^{D \times D}$ , where the number of ones for each row and for each column is exact one. This permutation matrix is 1) generated randomly by a key generator and 2) shared only among users, which we therefore refer to as *user-shared (U-S) permutation matrix*. As illustrated in Figure 3 (1), each user sends  $\phi K^{(n)}$  to an aggregator. The aggregator cannot identify  $K^{(n)}$  due to the absence of  $\phi$  but is still able to compute the encrypted and shuffled sum of  $\mathbf{w}^{(n)}$ , namely  $\phi \zeta(\mathbf{w}_{\text{sum}})$ , as follows (see also Figure 3 (2)):

$$\phi \zeta(\mathbf{w}_{\text{sum}}) = \phi \zeta\left(\sum_n \mathbf{w}^{(n)}\right) = \odot_n \phi \zeta(\mathbf{w}^{(n)}). \quad (4)$$

Here,  $\phi \zeta(\mathbf{w}^{(n)})$  is constructed as:

$$\phi \zeta(\mathbf{w}^{(n)}) = Y^{(n)} + \text{Enc0} \cdot Z^{(n)}, \quad (5)$$

$$Y^{(n)} = \phi K^{(n)} \zeta(\mathbf{v}^{(n)}), \quad Z^{(n)} = \mathbf{1}_D - \phi K^{(n)} \mathbf{1}_M, \quad (6)$$

where  $\mathbf{1}_D, \mathbf{1}_M$  are all-one vectors of length  $D$  and  $M$ . Enc0 is a ciphertext of a zero computed *only once* by the aggregator. Vector  $Z^{(n)}$  adds Enc0 to the locations of zeros in  $Y^{(n)}$ , which is necessary to compute  $\phi \zeta(\mathbf{w}_{\text{sum}})$  based on Equation (2). Finally, the aggregator asks the key generator for reordering and decrypting  $\phi \zeta(\mathbf{w}_{\text{sum}})$  to obtain  $\mathbf{w}_{\text{sum}}$ .

**Choosing encryption capacity  $M$**  The most efficient approach to choose  $M$  is by letting each user adjust  $M$  independently such that  $M^{(n)} = nz(\mathbf{w}^{(n)})$ . However, this approach will reveal  $nz(\mathbf{w}^{(n)})$  to an aggregator, which is less critical but still private information. Alternatively, one can keep  $nz(\mathbf{w}^{(n)})$  secret as follows: First, the key generator determines  $M$  such that  $M \ll D$ . If  $nz(\mathbf{w}^{(n)}) \leq M$  holds, a user encrypts  $\mathbf{w}^{(n)}$  as presented above. Otherwise, the user arbitrarily splits  $\mathbf{w}^{(n)}$  into several shards  $\mathbf{w}^{(n,1)}, \dots, \mathbf{w}^{(n,F)}$  such that  $\sum_f \mathbf{w}^{(n,f)} = \mathbf{w}^{(n)}$  and  $\max_f \{nz(\mathbf{w}^{(n,f)}) \mid f = 1, \dots, F\} \leq M$ . The user then encrypts and sends  $\mathbf{w}^{(n,f)}$  one-by-one, and the aggregator receives them as if they are individual data. This modification does not affect output  $\mathbf{w}_{\text{sum}}$  and is still efficient than naive encryption of all values as long as  $nz(\mathbf{w}^{(n)}) \ll D$ .

## 2.4. Encryption with Double Permutations

The U-S permutation matrix  $\phi$  ensures that  $K^{(n)}$  is secure against an aggregator. However, because  $\phi$  is shared

---

### Algorithm 1 Averaging Classifiers Securely with DPHE

---

**Require:**  $N$  users  $\{U^{(n)} \mid n = 1, \dots, N\}$  who privately own  $D$ -dimensional classifier weights  $\mathbf{w}^{(n)}$ , aggregator  $A$ , key generator  $G$ , and encryption capacity  $M$ .

**Ensure:** Averaged classifier  $\bar{\mathbf{w}} = \frac{1}{N} \sum_n \mathbf{w}^{(n)}$

- 1:  $G$  generates encryption key  $\zeta(\cdot)$  and the corresponding decryption key, U-S permutation matrix  $\phi$  and U-A permutation matrices  $\{\phi^{(n)} \mid n = 1, \dots, N\}$ .
  - 2:  $G$  distributes  $\zeta(\cdot), \phi, \phi^{(n)}, M$  to  $U^{(n)}$  and  $\zeta(\cdot), \{\phi^{(n)} \mid n = 1, \dots, N\}$  to  $C$ .
  - 3:  $U^{(n)}$  decomposes  $\mathbf{w}^{(n)}$  into  $\mathbf{v}^{(n)}$  and  $K^{(n)}$ .
  - 4:  $U^{(n)}$  sends  $\zeta(\mathbf{v}^{(n)}), \Phi(K^{(n)})$  to  $A$ .
  - 5:  $A$  computes  $\phi \zeta(\mathbf{w}^{(n)})$  for each  $n$  to get  $\phi \zeta(\mathbf{w}_{\text{sum}})$ .
  - 6:  $A$  asks  $G$  for decrypting and reordering  $\phi \zeta(\mathbf{w}_{\text{sum}})$  to receive  $\mathbf{w}_{\text{sum}}$ , and computes  $\bar{\mathbf{w}} = \frac{1}{N} \mathbf{w}_{\text{sum}}$ .
- 

among all users,  $K^{(n)}$  could be identified by a malicious user who intercepts  $\phi K^{(n)}$  via man-in-the-middle attacks.

To address this issue, DPHE involves another set of permutation matrices to shuffle  $K^{(n)}$ : *user-aggregator (U-A) permutation matrices*  $\{\phi^{(n)} \in \{0, 1\}^{D \times D} \mid n = 1, \dots, N\}$  where  $\phi^{(n)}$  is a permutation matrix defined in the same way as  $\phi$  but *is shared only between the user  $U^{(n)}$  and the aggregator*. Both of the U-S and U-A permutation matrices are generated by the key generator and distributed properly. Then, each user *doubly-permutes*  $K^{(n)}$  as follows:

$$\Phi(K^{(n)}) = \phi^{(n)} \phi K^{(n)}. \quad (7)$$

Because the reordering of  $\Phi(K^{(n)})$  requires both  $\phi$  and  $\phi^{(n)}$ , DPHE can now prevent an aggregator as well as any other user than  $U^{(n)}$  from identifying  $K^{(n)}$ .

Importantly, because the aggregator knows all U-A permutation matrices  $\{\phi^{(n)} \in \{0, 1\}^{D \times D} \mid n = 1, \dots, N\}$ ,  $\Phi(K^{(n)})$  can be reordered partially to obtain  $\phi K^{(n)}$ :

$$\phi K^{(n)} = (\phi^{(n)})^\top \Phi(K^{(n)}). \quad (8)$$

Note that  $(\phi^{(n)})^\top = (\phi^{(n)})^{-1}$ . This allows the aggregator to compute  $\phi \zeta(\mathbf{w}_{\text{sum}})$  based on Equation (4) by replacing  $\phi K^{(n)}$  in  $Y^{(n)}$  and  $Z^{(n)}$  in Equation (6) with  $(\phi^{(n)})^\top \Phi(K^{(n)})$ . Algorithm 1 summarizes how DPHE can average local classifiers securely and efficiently.

## 2.5. Security Evaluation

This section briefly describes that DPHE is guaranteed to be secure under certain assumptions. A more formal evaluation is present in our supplementary material.

**Assumptions** Recall that our framework involves users, an aggregator, and a key generator. Our security evaluation

is built upon one of the classical assumptions in cryptography that they are all *semi-honest* — each party “follows the protocol properly with the exception that it keeps a record of all its intermediate computations” [41]. For instance, the aggregator is not allowed to ask the key generator to decrypt individual classifier  $\phi\zeta(\mathbf{w}^{(n)})$  while he may use  $\bar{\mathbf{w}}$  to identify  $\mathbf{w}^{(n)}$ . We also assume that there is no collusion among the parties. For example, we will not consider cases where the aggregator and the key generator collude to share a Paillier decryption key and where the aggregator and all but one users collude to collect  $\{\mathbf{w}^{(j)} \mid j \neq n\}$  to reveal  $\mathbf{w}^{(n)}$ . These assumptions are justified in some practical applications such as crowdsourcing [34] and multiparty machine learning [43]. Finally, as described in Section 2.1, we consider a case where malicious users may intercept data sent from other users to the aggregator by slightly abusing the semi-honesty assumption.

**Security on Algorithm 1** With the assumptions above, we can guarantee that *no one but user  $U^{(n)}$  can identify  $\mathbf{w}^{(n)}$  and its non-zero indices  $K^{(n)}$  during and after running Algorithm 1 if the number of users satisfies  $N \geq 3$* . Specifically, non-zero weights of  $\mathbf{w}^{(n)}$ ,  $\mathbf{v}^{(n)}$ , are encrypted with the Paillier cryptosystem, which prevents an aggregator and all users from identifying  $\mathbf{v}^{(n)}$  from its ciphertext  $\zeta(\mathbf{v}^{(n)})$  as they do not have a decryption key. Also, they cannot identify  $K^{(n)}$  from its doubly-permuted form  $\Phi(K^{(n)})$  due to the lack of either U-S permutation matrix  $\phi$  or U-A matrix  $\phi^{(n)}$ . Although the key generator owns the decryption key and all of the permutation matrices, he does not access  $\mathbf{v}^{(n)}$  and  $K^{(n)}$  in the algorithm. Finally,  $\bar{\mathbf{w}}$  cannot be used to identify  $\mathbf{w}^{(n)}$  when  $N \geq 3$ . For example, when  $N = 3$ , user  $U^{(1)}$  gets  $\bar{\mathbf{w}} = \frac{1}{3}(\mathbf{w}^{(1)} + \mathbf{w}^{(2)} + \mathbf{w}^{(3)})$ . However,  $U^{(1)}$  can only compute  $3\bar{\mathbf{w}} - \mathbf{w}^{(1)} = \mathbf{w}^{(2)} + \mathbf{w}^{(3)}$ , which still remains ambiguous to identify  $\mathbf{w}^{(2)}$  and  $\mathbf{w}^{(3)}$ . Similarly, non-zero indices of  $\bar{\mathbf{w}}$  is just the union of those of  $\mathbf{w}^{(n)}$  and not useful to reveal  $K^{(n)}$ . This security on  $\bar{\mathbf{w}}$  holds also for the aggregator and the key generator.

**Limitations** The requirement of  $N \geq 3$  implies that, if only one or two users are available at one time, the aggregator will never publish classifiers without revealing each user’s updates. Moreover, we cannot prevent certain attacks using published classifiers to infer potentially-private data, *e.g.*, using a face recognition model and its output to reconstruct face images specific to the output [23], although such attacks are not currently able to identify which users privately owned the reconstructed data in a distributed setting.

### 3. Experiments

In this section, we address several visual recognition tasks with our privacy-preserving learning framework based

on DPHE. Specifically, we first evaluate DPHE empirically under various conditions systematically with object classification tasks on Caltech101 [21] and Caltech256 [26] datasets. Then we tackle more practical and sensitive tasks: face attribute recognition on the CelebA dataset [40] and sensitive place detection on the Life-logging dataset [19].

#### 3.1. Settings of Experiments

Throughout our experiments, we learned a linear SVM via SGD. We employed the elastic net [65], *i.e.*, the combination of L1 and L2 regularizations, to enforce sparsity on locally-updated classifiers. For a simulation purpose, multiple users, an aggregator, and a key generator were all implemented in a single computer and thus no transmissions among them were considered in the experiments.

As a baseline, we adopted the following two off-the-shelf privacy-preserving machine learning methods. Similar to our approach, these methods were designed to learn classifiers over distributed private data while ensuring that no private information about distributed data was leaked during and after the learning procedure.

**PRR10 [43]** Users first train  $D$ -dimensional linear classifiers locally using their own data. They then average the locally-trained classifiers followed by adding a  $D$ -dimensional Laplace noise vector to the output classifier based on differential privacy. Because the scale of noise depends on the size of local user data and thus involves private information, the noise vector is encrypted with the Paillier cryptosystem (*i.e.*, this method needs  $D$ -times encryptions). An L2-regularized logistic regression classifier was trained by each user via SGD.

**RA12 [46]** Similar to our approach, RA12 iteratively averages locally-updated classifiers to get a better classifier. Unlike our method and PRR10, each local classifier is updated in a gradient descent fashion while adding the combination of Laplace and Gaussian noises to its objective function to prevent learned classifiers from leaking private information based on differential privacy. We employed an L2-regularized linear SVM for each of local classifiers.

These methods involve hyper-parameters  $\epsilon \in \mathbb{R}_+$ ,  $\delta \in \mathbb{R}_+$  to control the strength of privacy preservation. We followed the original papers [43, 46] and set to  $\epsilon = 0.2$ ,  $\delta = 0.05$ , which were adjusted and justified in the papers to preserve privacy while keeping high recognition performances.

We set  $N = 5$  for these two baselines and our method, that is, five users were assumed to participate in a task. Encryption capacity  $M$  was set to  $M = \lceil 0.1D \cdot NC \rceil$ , where  $D$  was the dimension of learned classifiers and  $NC$  was the number of classes. If each of locally-updated classifiers has

Table 1. **Object Recognition on Caltech101/256**: mean and standard deviation of classification accuracies (%) averaged over object categories. Note that the scores of HZRS14 and ZF13 were cited from their original papers [28, 62].

Methods	Caltech101	Caltech256	Privacy
HZRS14 [28]	93.4 ± 0.5	N/A	✗
ZF13 [62]	85.4 ± 0.4	72.6±0.1	✗
PRR10 [43]	41.6 ± 1.2	55.9 ± 0.5	✓
RA12 [46]	83.8 ± 1.1	68.0 ± 0.3	✓
<b>DPHE</b>	89.3 ± 0.8	74.7 ± 0.4	✓

more than 90% sparsity, its weights can be encrypted and transmitted at one time (see Section 2.3).

### 3.2. Performance Analysis on Caltech101/256

We compared DPHE against the privacy-preserving baselines [43, 46] as well as some of the state-of-the-art object classification methods [28, 62] on Caltech101/256 object classification datasets. Our evaluation was based on the protocol shown in [28, 62] as follows. For Caltech101 dataset [21], we generated training and testing data by randomly picking 30 and no more than 50 images respectively for each category. On the other hand, for the Caltech256 dataset [26], we chose random 60 images per category for training data and the rest of images for testing data.

Recall that our learning framework requires some initialization data for obtaining the initial weights  $\bar{w}_0$ . Therefore, we first 1) left 10% of the training data for the initialization data and 2) split 90% of them into five subsets of the same size, which we regarded as private image data of five users. For the two baselines [43, 46], we added 1) to each of 2) to serve as private data of each user. This ensured that in both of our method and the baselines each user had access to the same image set. We ran the evaluation ten times to get a mean and a standard deviation of classification accuracies averaged over all the categories.

For our method and the baselines, we extracted 2048-dimensional outputs of the global-average pooling layer of the deep residual network [29] with 152 layers trained on ImageNet [47] as a feature, which was then normalized to have zero-mean and unit-variance for each dimension. In our method, the relative strength of L1 regularization in the elastic net was fixed to 0.5, and the overall strength of elastic-net regularization was adjusted adaptively so that local classifiers had more than 90% sparsity on average.

**Comparisons to other methods** Table 1 presents classification accuracies of all the methods. Our method performed almost comparably well to state-of-the-art methods [28, 62] while preserving privacy. We also found that DPHE outperformed the privacy-preserving baselines [43, 46]. We here emphasize that DPHE was designed to preserve privacy

Table 2. **Sparsity (%), Classification Accuracy (%), and Encryption Time (sec) of DPHE on Caltech101.**

Sparsity	Accuracy	Time
0.01	89.7	620
79.1	89.6	186
95.6	88.2	62




Figure 4. **Classification Accuracy (%) with Different Numbers of Users on Caltech101.** Left: users own the fixed number of partially-duplicated images. Right: users own the decreasing number (inversely proportional to  $N$ ) of non-duplicated images.

without adding any noise on locally-updated or averaged classifiers, which would be one main reason of our performance improvements over the two baselines. RA12 [46] inevitably needs to add noise on local classifiers, and still does not preserve privacy perfectly due to the trade-off between data utility and security. On the other hand, PRR10 [43] achieves better privacy preservation thanks to the combination of homomorphic encryption and perturbed classifiers. However, this baseline did not employ iterative updates of classifiers unlike other methods and thus could not effectively leverage distributed training data.

**Encryption time** Table 2 depicts the relationship between the sparsity of locally-updated classifiers (averaged over all users and all iterations), classification accuracies, and encryption times of DPHE on Caltech101. We adjusted the sparsity by changing the relative strength of L1 regularization in the elastic net. Overall, the increase of sparsity little affected the classification performance. The comparison between the top and the bottom rows of the table indicates that we achieved ten times faster encryption at the cost of the only 1.5 percent points decrease of classification accuracies. Note that PRR10 [43] requires to encrypt a 2048-dimensional dense noise vector for privacy preservation, which took about 620 sec in this condition.

**Number of users** We also investigate how the number of users  $N$  affected classification performances on Caltech101 under two different conditions. In the left of Figure 4, we repeated multiple (1, 2, 5, 10, and 20) times the random split of training data into five subsets. This yielded  $N = 5, 10, 25, 50, 100$  subsets each of which has the fixed number of partially duplicated images. We found

Table 3. **Face Attribute Recognition on CelebA Dataset** : mean accuracies (%) over 40 face attribute labels. The score of LLWT15 was cited from their original paper [40].

Methods	Accuracy	Privacy
LLWT15 [40]	87	✗
PRR10 [43]	78	✓
RA12 [46]	64	✓
<b>DPHE</b>	84	✓

that the increase of user sizes hardly affected classification performances but slightly made it worse when  $N = 100$ . On the right of the figure, we split the training data into  $N = 5, 10, 25, 50, 100$  subsets. While this split ensured that there were no duplicated images across multiple subsets, the number of images in each subset decreased as  $N$  became larger. The result with this condition implies that DPHE requires each user to own a sufficient number of images to make each local update stable.

### 3.3. Practical Examples

As proof-of-concept applications, this section addresses two practical visual learning tasks where training images could be highly private and need privacy preservation: face attribute recognition and sensitive place detection. Note that the following experiments are practical also in terms of data size: the two datasets used in the experiments contained more than one hundred thousand images in total.

**Face attribute recognition** We first adopted a face attribute recognition task on the CelebFaces Attributes (CelebA) Dataset [40] comprising 202,599 face images with 40 attribute labels (*e.g.*, big lips, oval faces, young). As a feature, we used 512-dimensional outputs of the FC5 layer of the face recognition network [58] trained on the CASIA WebFace [60]. All face images were cropped and aligned using provided facial landmarks. We used validation data (19,868 images) for initialization data and split training data (162,770 images) into five subsets to run our method with  $N = 5$ . Similar to the previous experiments, we added the initialization data to each of the five subsets to serve as one of the five private data for privacy-preserving baselines [43, 46]. The regularization strength of our method was adjusted so that the sparsity of local classifiers was about 75% on average, which required about 14 seconds for encrypting each classifier. Table 3 reports a recognition accuracy averaged over all of the 40 face attributes. DPHE worked comparably well to the state-of-the-art method [40] and outperformed privacy-preserving baselines [43, 46].

**Sensitive place detection** Another interesting application where privacy preservation plays an important role is sensitive place detection presented in [19, 35, 55]. This task

Table 4. **Sensitive Place Detection on Life-logging Dataset**: average precision for detecting images of sensitive places.

Methods	PRR10 [43]	RA12 [46]	DPHE
Average Precision	0.546	0.704	0.729

aims at detecting images of a place where sensitive information could be involved, such as those of bathrooms and laptop screens. We used the Life-logging dataset [19] that contained 131,287 training images selected from the Microsoft COCO [38] and Flickr8k [30] datasets, and 7,210 testing images recorded by a wearable life-logging camera. Following [19], images with specific annotations: ‘toilet,’ ‘bathroom,’ ‘locker,’ ‘lavatory,’ ‘washroom,’ ‘computer,’ ‘laptop,’ ‘iphone,’ ‘smartphone,’ and ‘screen’ were regarded as sensitive ones. We used a 4096-dimensional deep feature provided in the work [19]. 132 images containing both sensitive and non-sensitive samples were chosen randomly for initialization data and the rest was split into five subsets to serve as private data. The regularization strength was chosen such that sparsity of local updates was more than 95% on average, which required only 1.2 seconds for encryption. Table 4 presents average precision scores for detecting sensitive places with our method and the two privacy-preserving baselines. Note that the method in the original paper [19] could not be compared directly due to the absence of its average precision score and different experimental conditions (*i.e.*, smaller numbers of training and testing images were used). We confirmed that DPHE performed fairly well compared to the two baselines.

## 4. Conclusion

We developed a privacy-preserving framework with a new encryption scheme DPHE for learning visual classifiers securely over distributed private data. Our experiments show that the proposed framework outperformed other privacy-preserving baselines [43, 46] in terms of accuracy and worked comparably well to several state-of-the-art visual recognition methods.

Although we focused exclusively on the learning of linear classifiers, DPHE can encrypt any type of sparse high-dimensional data efficiently and thus could be used for other tasks. One interesting task is distributed unsupervised or semi-supervised learning which will not require users to annotate all private data. Another promising direction for the future work is learning much higher-dimensional models like sparse convolutional neural networks [6, 53]. Our privacy-preserving framework will make it easy to provide diverse data sources for learning such complex models.

**Acknowledgements** This work was supported by JST ACT-I Grant Number JPMJPR16UT and JST CREST Grant Number JPMJCR14E1, Japan.



## References

- [1] S. Abu-El-Haija, N. Kothari, J. Lee, P. Natsev, G. Toderici, B. Varadarajan, and S. Vijayanarasimhan. YouTube-8M: A Large-Scale Video Classification Benchmark. *arXiv preprint arXiv:1609.08675*, 2016.
- [2] S. Agarwal, Y. Furukawa, and N. Snavely. Building Rome in a Day. *Communications of the ACM*, 54(10):105–112, 2011.
- [3] C. C. Aggarwal and P. S. Yu. A General Survey of Privacy-Preserving Data Mining Models and Algorithms. In *Privacy-Preserving Data Mining*, volume 34, pages 11–52. Springer US, 2008.
- [4] R. Agrawal and R. Srikant. Privacy-Preserving Data Mining. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*, volume 29, pages 439–450, 2000.
- [5] S. Ahern, D. Eckles, N. Good, S. King, M. Naaman, and R. Nair. Over-Exposed? Privacy Patterns and Considerations in Online and Mobile Photo Sharing. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, pages 357–366, 2007.
- [6] Baoyuan Liu, Min Wang, H. Foroosh, M. Tappen, and M. Pensky. Sparse Convolutional Neural Networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 806–814, 2015.
- [7] R. Bost, R. Popa, S. Tu, and S. Goldwasser. Machine Learning Classification over Encrypted Data. In *Proceedings of the Annual Network and Distributed System Security Symposium*, pages 1–14, 2015.
- [8] K. Chaudhuri and C. Monteleoni. Privacy-Preserving Logistic Regression. In *Proceedings of the Advances in Neural Information Processing Systems*, pages 289–296, 2008.
- [9] K. Chaudhuri, C. Monteleoni, and A. D. Sarwate. Differentially Private Empirical Risk Minimization. *Journal of Machine Learning Research*, 12:1069–1109, 2011.
- [10] M. M. Cheng, Z. Zhang, W. Y. Lin, and P. Torr. BING: Binarized Normed Gradients for Objectness Estimation at 300fps. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3286–3293, 2014.
- [11] S. Chowdhury. Exploring Lifelog Sharing and Privacy. In *Proceedings of the ACM International Joint Conference on Pervasive and Ubiquitous Computing*, pages 553–558, 2016.
- [12] A. Chu, M. Abdi, and I. Goodfellow. Deep Learning with Differential Privacy. In *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security*, pages 308–318, 2016.
- [13] M. Courbariaux, Y. Bengio, and J.-P. David. BinaryConnect: Training Deep Neural Networks with Binary Weights during Propagations. In *Proceedings of the Advances in Neural Information Processing Systems*, pages 3123–3131, 2015.
- [14] N. Dowlin, R. Gilad-Bachrach, K. Laine, K. Lauter, M. Naehrig, and J. Wernsing. CryptoNets: Applying Neural Networks to Encrypted Data with High Throughput and Accuracy. Technical report, Microsoft Research, 2016.
- [15] C. Dwork. Differential Privacy. In *Proceedings of the International Colloquium on Automata, Languages and Programming*, pages 1–12, 2006.
- [16] C. Dwork. Differential Privacy: A Survey of Results. In *Proceedings of the Theory and Applications of Models of Computation*, volume 4978, pages 1–19, 2008.
- [17] C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating Noise to Sensitivity in Private Data Analysis. In *Proceedings of the Third conference on Theory of Cryptography*, pages 265–284, 2006.
- [18] Z. Erkin, M. Franz, J. Guajardo, S. Katzenbeisser, I. Lagendijk, and T. Toft. Privacy-Preserving Face Recognition. In *Proceedings of the International Symposium on Privacy Enhancing Technologies Symposium*, pages 235–253, 2009.
- [19] C. Fan and D. J. Crandall. Deepdiary: Automatically Captioning Lifelogging Image Streams. In *Proceedings of the Workshop on Egocentric Perception, Interaction, and Computing*, volume 9913, pages 459–473, 2016.
- [20] A. Fathi, J. K. Hodgins, and J. M. Rehg. Social Interactions: A First-Person Perspective. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1226–1233, 2012.
- [21] L. Fei-Fei, R. Fergus, and P. Perona. Learning Generative Visual Models from Few Training Examples: An Incremental Bayesian Approach Tested on 101 Object Categories. *Computer Vision and Image Understanding*, 106(1):59–70, 2007.
- [22] C. Fontaine and F. Galand. A Survey of Homomorphic Encryption for Nonspecialists. *Eurasip Journal on Information Security*, 2007(15):1–10, 2007.
- [23] M. Fredrikson, S. Jha, and T. Ristenpart. Model Inversion Attacks That Exploit Confidence Information and Basic Countermeasures. In *Proceedings of the SIGSAC Conference on Computer and Communications Security*, pages 1322–1333, 2015.
- [24] J. H. Friedman. Fast Sparse Regression and Classification. *International Journal of Forecasting*, 28(3):722–738, 2012.
- [25] T. Graepel, K. Lauter, and M. Naehrig. ML Confidential: Machine Learning on Encrypted Data. In *Proceedings of the International Conference on Information Security and Cryptology*, pages 1–21, 2012.
- [26] G. Griffin, A. Holub, and P. Perona. Caltech-256 Object Category Dataset. Technical report, Caltech Technical Report 7694, 2007.
- [27] S. Hare, A. Saffari, and P. H. S. Torr. Efficient Online Structured Output Learning for Keypoint-Based Object Tracking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1894–1901, 2012.
- [28] K. He, X. Zhang, S. Ren, and J. Sun. Spatial Pyramid Pooling in Deep Convolutional Networks for Visual Recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(9):1904–1916, 2015.
- [29] K. He, X. Zhang, S. Ren, and J. Sun. Deep Residual Learning for Image Recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 171–180, 2016.
- [30] M. Hodosh, P. Young, and J. Hockenmaier. Framing Image Description as a Ranking Task: Data, Models and Evaluation Metrics. *Journal of Artificial Intelligence Research*, 47(1):853–899, 2013.

- [31] R. Hoyle, R. Templeman, D. Anthony, D. Crandall, and A. Kapadia. Sensitive Lifelogs: A Privacy Analysis of Photos from Wearable Cameras. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, pages 1645–1648, 2015.
- [32] R. Hoyle, R. Templeman, S. Armes, D. Anthony, and D. Crandall. Privacy Behaviors of Lifeloggers using Wearable Cameras. In *Proceedings of the ACM International Joint Conference on Pervasive and Ubiquitous*, pages 571–582, 2014.
- [33] S. Kairam, J. J. Kaye, J. A. Guerra-Gomez, and D. A. Shamma. Snap Decisions?: How Users, Content, and Aesthetics Interact to Shape Photo Sharing Behaviors. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, pages 113–124, 2016.
- [34] H. Kajino, H. Arai, and H. Kashima. Preserving Worker Privacy in Crowdsourcing. *Data Mining and Knowledge Discovery*, 28(5-6):1314–1335, 2014.
- [35] M. Korayem, R. Templeman, D. Chen, D. Crandall, and A. Kapadia. Enhancing Lifelogging Privacy by Detecting Screens. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, pages 4309–4314, 2016.
- [36] J. Krumm. A Survey of Computational Location Privacy. *Personal and Ubiquitous Computing*, 13(6):391–399, 2009.
- [37] Y. J. Lee and K. Grauman. Predicting Important Objects for Egocentric Video Summarization. *International Journal of Computer Vision*, 114(1):38–55, 2015.
- [38] T. Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft COCO: Common Objects in Context. In *Proceedings of the European Conference on Computer Vision*, pages 740–755, 2014.
- [39] Y. Lindell and B. Pinkas. Secure Multiparty Computation for Privacy-Preserving Data Mining. *Journal of Privacy and Confidentiality*, 1(1):59–98, 2009.
- [40] Z. Liu, P. Luo, X. Wang, and X. Tang. Deep Learning Face Attributes in the Wild. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3730–3738, 2015.
- [41] Oded Goldreich. *Foundations of Cryptography: Basic Applications*. Cambridge University Press, 2004.
- [42] P. Paillier. Public-Key Cryptosystems based on Composite Degree Residuosity Classes. In *Proceedings of the International Conference on the Theory and Applications of Cryptographic Techniques*, volume 1592, pages 223–238, 1999.
- [43] M. Pathak, S. Rane, and B. Raj. Multiparty Differential Privacy via Aggregation of Locally Trained Classifiers. In *Proceedings of the Advances in Neural Information Processing Systems*, pages 1876–1884, 2010.
- [44] B. Pinkas and Benny. Cryptographic Techniques for Privacy-Preserving Data Mining. *ACM SIGKDD Explorations Newsletter*, 4(2):12–19, 2002.
- [45] H. Pirsiavash and D. Ramanan. Detecting Activities of Daily Living in First-Person Camera Views. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2847–2854, 2012.
- [46] A. Rajkumar and S. Agarwal. A Differentially Private Stochastic Gradient Descent Algorithm for Multiparty Classification. In *Proceedings of the International Conference on Artificial Intelligence and Statistics*, pages 933–941, 2012.
- [47] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015.
- [48] M. S. Ryoo and L. Matthies. First-Person Activity Recognition: What Are They Doing to Me? In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2730–2737, 2013.
- [49] M. S. Ryoo, B. Rothrock, and C. Fleming. Privacy-Preserving Egocentric Activity Recognition from Extreme Low Resolution. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 1–8, 2017.
- [50] A. R. Sadeghi, T. Schneider, and I. Wehrenberg. Efficient Privacy-Preserving Face Recognition. In *Proceedings of the International Conference on Information Security and Cryptology*, volume 5984, pages 229–244, 2010.
- [51] R. Shokri and V. Shmatikov. Privacy-Preserving Deep Learning. In *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security*, pages 1310–1321, 2015.
- [52] K. K. Singh, K. Fatahalian, and A. A. Efros. KrishnaCam: Using a Longitudinal, Single-Person, Egocentric Dataset for Scene Understanding Tasks. In *Proceedings of the IEEE Winter Conference on Applications of Computer Vision*, pages 1–9, 2016.
- [53] Y. Sun, X. Wang, and X. Tang. Sparsifying Neural Network Connections for Face Recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4856–4864, 2016.
- [54] H. Takabi, E. Hesamifard, and M. Ghasemi. Privacy Preserving Multi-party Machine Learning with Homomorphic Encryption. In *Proceedings of the Workshop on Private Multi-Party Machine Learning*, pages 1–5, 2016.
- [55] R. Templeman, M. Korayem, D. Crandall, and K. Apu. PlaceAvoider: Steering First-Person Cameras away from Sensitive Spaces. In *Proceedings of the Annual Network and Distributed System Security Symposium*, pages 23–26, 2014.
- [56] R. Tibshirani. Regression Shrinkage and Selection via the Lasso. *Journal of the Royal Statistical Society. Series B: Statistical Methodology*, 58:267–288, 1994.
- [57] J. Vaidya and C. Clifton. Privacy-Preserving K-Means Clustering over Vertically Partitioned Data. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 206–215, 2003.
- [58] Y. Wen, K. Zhang, Z. Li, and Y. Qiao. A Discriminative Feature Learning Approach for Deep Face Recognition. In *Proceedings of the European Conference on Computer Vision*, pages 499–515, 2016.
- [59] A. C.-C. Yao. How to Generate and Exchange Secrets. In *Proceedings of the Annual Symposium on Foundations of Computer Science*, pages 162–167, 1986.
- [60] D. Yi, Z. Lei, S. Liao, and S. Z. Li. Learning Face Representation from Scratch. *arXiv preprint arXiv:1411.7923*, 2014.

- [61] J. Yuan and S. Yu. Privacy Preserving Back-Propagation Neural Network Learning Made Practical with Cloud Computing. *IEEE Transactions on Parallel and Distributed Systems*, 25(1):212–221, 2014.
- [62] M. D. Zeiler and R. Fergus. Visualizing and Understanding Convolutional Networks. In *Proceedings of the European Conference on Computer Vision*, volume 8689, pages 818–833, 2014.
- [63] B. Zhou, L. Liu, A. Oliva, and A. Torralba. Recognizing City Identity via Attribute Analysis of Geo-Tagged Images. In *Proceedings of the European Conference on Computer Vision*, volume 8691, pages 519–534, 2014.
- [64] B. Zhou, J. Pei, and W. Luk. A Brief Survey on Anonymization Techniques for Privacy Preserving Publishing of Social Network Data. *ACM SIGKDD Explorations Newsletter*, 10(2):12–22, 2008.
- [65] H. Zou and T. Hastie. Regularization and Variable Selection via the Elastic Net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2):301–320, 2005.