

Searching Privately by Imperceptible Lying: A Novel Private Hashing Method with Differential Privacy

Yimu Wang, Shiyin Lu, and Lijun Zhang

{wangym,lusy,zhanglj}@lamda.nju.edu.cn

State Key Laboratory for Novel Software Technology, Nanjing University
Nanjing 210023, China

ABSTRACT

In the big data era, with the increasing amount of multi-media data, approximate nearest neighbor (ANN) search has been an important but challenging problem. As a widely applied large-scale ANN search method, hashing has made great progress, and achieved sub-linear search time with low memory space. However, the advances in hashing are based on the availability of large and representative datasets, which often contain sensitive information. Typically, the privacy of this individually sensitive information is compromised. In this paper, we tackle this valuable yet challenging problem and formulate a task termed as *private hashing*, which takes into account both searching performance and privacy protection. Specifically, we propose a novel noise mechanism, *i.e.*, Random Flipping, and two private hashing algorithms, *i.e.*, PHashing and PITQ, with the refined analysis within the framework of differential privacy, since differential privacy is a well-established technique to measure the privacy leakage of an algorithm. Random Flipping targets binary scenarios and leverages the “Imperceptible Lying” idea to guarantee ϵ -differential privacy by flipping each datum of the binary matrix (noise addition). To preserve ϵ -differential privacy, PHashing perturbs and adds noise to the hash codes learned by non-private hashing algorithms using Random Flipping. However, the noise addition for privacy in PHashing will cause severe performance drops. To alleviate this problem, PITQ leverages the power of alternative learning to distribute the noise generated by Random Flipping into each iteration while preserving ϵ -differential privacy. Furthermore, to empirically evaluate our algorithms, we conduct comprehensive experiments on the image search task and demonstrate that proposed algorithms achieve equal performance compared with non-private hashing methods.

CCS CONCEPTS

• Security and privacy; • Information systems → Information retrieval;

KEYWORDS

Differential Privacy, Large-Scale Multimedia Retrieval, Hashing

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MM '20, October 12–16, 2020, Seattle, WA, USA

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-7988-5/20/10...\$15.00

<https://doi.org/10.1145/3394171.3413882>

ACM Reference Format:

Yimu Wang, Shiyin Lu, and Lijun Zhang. 2020. Searching Privately by Imperceptible Lying: A Novel Private Hashing Method with Differential Privacy. In *Proceedings of the 28th ACM International Conference on Multimedia (MM '20)*, October 12–16, 2020, Seattle, WA, USA. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3394171.3413882>

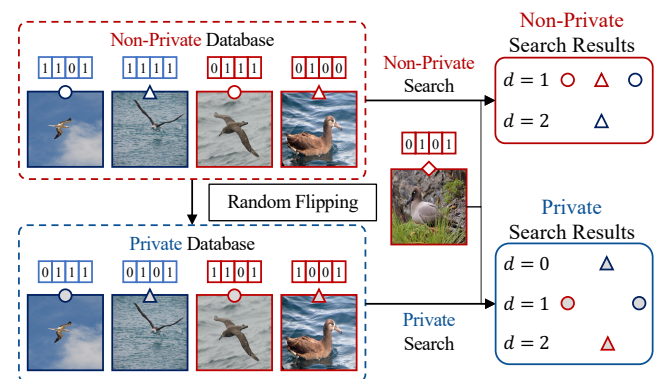


Figure 1: Our proposed *Random Flipping* is the first noise mechanism aimed at hashing scenarios, which changes the *non-private* database into a *private* database, and preserves ϵ -differential privacy. With *non-private* databases, privacy is likely to leak while attaining high search performance. With *private* databases, privacy will be protected but protecting privacy will cause a slight performance drop. Here, different colors of markers and boxes outside images represent different subordinate categories. d is the Hamming distance.

1 INTRODUCTION

To deal with the large-scale data, as a popular technique in multimedia search, approximate nearest neighbor (ANN) search [3, 4] has attracted much attention in recent years. Among many ANN search methods, hashing [17, 33, 35, 43, 51] is an active and representative sub-area. Particularly, hashing maps data points to binary codes with learned hash functions by preserving similarity in the original space of data points. With the power of the binary representation, the storage cost for the large-scale database can be drastically reduced, while the search time complexity is constant or sub-linear. Existing hashing methods can be roughly categorized into unsupervised and supervised hashing. Unsupervised hashing methods, *e.g.*, LSH [17], SH [51], and ITQ [18], learn hash functions from unlabeled data. To improve the quality of hash codes,

supervised hashing methods such as KSH [33], utilize supervised information (e.g., similarity matrix or label information). Inspired by the progress of deep neural networks [21, 50], deep supervised hashing [5, 20, 24, 30, 46, 53, 54] has been proposed, which leverages the power of deep learning to generate high-level semantic features.

With the increasing massive data gathered from customers or individuals in the real world, hashing has been applied in many scenarios, such as person re-identification [56] and face search [31], and achieved promising performance. The availability of large and representative datasets [26, 29], which are often crowd-sourced and contain sensitive information, advances the success of hashing technique. However, that would comprise and leak the privacy of individually sensitive information. The wide use of these sensitive data requires techniques to meet the demands of applications that offer principled and rigorous privacy guarantees. Differential privacy [13] is a well-established metric that mathematically requires the probability of any observable outputs to change very little when any single datum changes. It ensures no attacker can tell whether the algorithm utilizes or not a specific training example, hence addressing the aforementioned problem. Recently, differential privacy has been applied in many active areas, e.g., online learning [2, 19, 22], generic deep learning [1, 44], and natural language processing [37].

In this paper, we formulate this novel problem called *private hashing* as a step towards efficient large-scale multi-media search with privacy. To achieve private hashing, we propose a noise mechanism, Random Flipping, and two algorithms, i.e., PHashing and PITQ, within the framework of differential privacy.

The proposed noise mechanism, Random Flipping, leverages the power of uncertainty (“Imperceptible Lying” proposal) to protect privacy. Consider the following scenario: We are taking a test questioning some private individual information, and we have to answer “yes” or “no”. If we want to keep this private answer uncovered, we can simply protect our privacy by randomly changing our reply. Similarly, the binary hash codes can be regarded as sequential questions of input features, such as black bird head and yellow long tail, which can be used for identifying the input. Random Flipping is based on this proposal and preserves ϵ -differential privacy by randomly flipping each datum (private question) of binary input with a predefined probability as shown in Fig. 1.

Two private hashing algorithms target different scenarios. PHashing utilizes Random Flipping to extend non-private hashing methods into private methods and preserves ϵ -differential privacy. However, the privacy protection causes severe search performance drops, which are inversely proportional to the strength of privacy-preserving. To alleviate this problem, we develop PITQ, which leverages the power of iterative discrete optimization to distribute the noise generated by Random Flipping into several iterations with ϵ -differential privacy.

To empirically evaluate the search performance and cost of privacy protection of our proposed algorithms, i.e., PHashing and PITQ, we conduct comprehensive experiments on five vision datasets, i.e., *ImageNet* [11], *CIFAR-10* [26], *NUS-WIDE* [10], *Aircraft* [36], and *CUB* [48]. Empirical results show that PHashing preserves ϵ -differential privacy at an acceptable price. PITQ achieves equal

performance compared with non-private ITQ, and sometimes outperforms the non-private ITQ as the noise helps to converge to a better local optimum in the discrete binary optimization.

The main contributions of our paper are summarized as follows:

- Our paper is the first to formulate the practical and challenging *private hashing* problem within the framework of differential privacy.
- The proposed *Random Flipping* is the first noise mechanism aimed at hashing scenarios in the differential privacy area, which ensures ϵ -differential privacy by flipping each datum of the binary input with a predefined probability.
- We devise two private hashing algorithms preserving ϵ -differential privacy, i.e., PHashing and PITQ, and present theoretical guarantees. PHashing extends the non-private hashing algorithms into private algorithms by Random Flipping with an acceptable price. To reduce the performance drop, PITQ distributes the noise generated by Random Flipping into several iterations instead of adding the noise only in the last iteration, which leads to better performances than PHashing (ITQ).
- Empirical results show that our algorithms achieve equal performances compared with non-private algorithms on the image search task while preserving ϵ -differential privacy.

2 RELATED WORK

In this section, we briefly review related work on differential privacy and hashing.

2.1 Differential Privacy

Differential privacy, firstly introduced by Dwork et al. [13], is a mathematical notion of privacy that requires the probability of any observable outputs to change very little when any single datum changes.

DEFINITION 1 (DIFFERENTIAL PRIVACY [14]). A randomized algorithm \mathcal{M} with domain $\mathbb{N}^{\|X\|}$ is (ϵ, δ) -differentially private if $\forall S \subseteq \text{Range}(\mathcal{M})$ and $\forall x, y \in \mathbb{N}^{\|X\|}$ with $\|x - y\|_1 \leq 1$,

$$\mathbb{P}[\mathcal{M}(x) \in S] \leq e^\epsilon \mathbb{P}[\mathcal{M}(y) \in S] + \delta. \quad (1)$$

Here, $\epsilon, \delta \geq 0$ are privacy loss parameters, which quantify the strength of privacy protection. If S is a countable set, then we can modify the inequality (1) as

$$\mathbb{P}[\mathcal{M}(x) = s] \leq e^\epsilon \mathbb{P}[\mathcal{M}(y) = s] + \delta, \quad (2)$$

where $s \in S$. If $\delta = 0$, we say that \mathcal{M} is ϵ -differentially private. Furthermore, the strength of privacy-preserving is inversely proportional to privacy parameters, i.e., ϵ and δ .

The above definition provides rigorous, information-theoretic guarantees against an adversary’s ability to infer whether one data is being used in a randomized mechanism \mathcal{M} . In practice, we consider a privacy guarantee meaningful if $\epsilon \approx 1$ and $\delta = O(1/n)$, where n is the size of a private dataset. One important and useful property of differential privacy is that differential privacy is closed under post-processing [14] as shown below.

THEOREM 2 (POST-PROCESSING [14]). Let $\mathcal{M} : \mathbb{N}^{\|X\|} \rightarrow \mathbb{R}$ be a (ϵ, δ) -differentially private algorithm. Let $f : \mathbb{R} \rightarrow \mathbb{R}'$ be an arbitrary

randomized mapping. Then $f \circ M : \mathbb{N}^{\|X\|} \rightarrow \mathbb{R}'$ is (ϵ, δ) -differentially private.

In another word, if any part of an algorithm, including outputs and any components, is differentially private, the whole algorithm is differentially private. Another important property related to our work is the basic composition theorem, which indicates that the composition of k (ϵ_k, δ_k) -differential private algorithms is at worst $(\sum_k \epsilon_k, \sum_k \delta_k)$ -differential privacy.

THEOREM 3 (BASIC COMPOSITION [14]). Let $M_i : \mathbb{N}^{\|X\|} \rightarrow \mathbb{R}_i$ be (ϵ_i, δ_i) -differential private for $i = 1, \dots, k$. Then the composition $M_{[k]} : \mathbb{N}^{\|X\|} \rightarrow \mathbb{R}_1 \times \dots \times \mathbb{R}_k$ defined as: $M_{[k]}(x) = (M_1(x), M_2(x), \dots, M_k(x))$ is $(\sum_{i=1}^k \epsilon_i, \sum_{i=1}^k \delta_i)$ -differential private.

A typical paradigm for achieving differentially private real-valued algorithms M_p is adding noise into the non-private algorithm M . The noise is calibrated to the sensitivity of M , which is the maximum of the absolute distance $\|M(x) - M(x')\|$, where x and x' are adjacent inputs, i.e., $\|x - x'\|_1 \leq 1$. For instance, the Gaussian noise mechanism is defined by

$$M_p(x) = M(x) + \mathcal{N} \left(0, \max_{x, x' \in X, \|x - x'\|_1 \leq 1} \left(\|M(x) - M(x')\| \right)^2 \sigma^2 \right), \quad (3)$$

where $\mathcal{N} \left(0, \max_{x, x' \in X, \|x - x'\|_1 \leq 1} \left(\|M(x) - M(x')\| \right)^2 \sigma^2 \right)$ is the Gaussian distribution with mean 0 and standard deviation $\max_{x, x' \in X, \|x - x'\|_1 \leq 1} \left(\|M(x) - M(x')\| \right) \sigma$. If $\delta \geq \frac{4}{5} \exp^{-(\delta\epsilon)^2/2}$ and $\epsilon < 1$, the private real-valued algorithm M_p satisfies (ϵ, δ) -differential privacy [14]. This paradigm utilizes the post-processing theorem by implying a real-valued random Gaussian noise.

However, the existing literature on differential privacy mostly targets real-valued algorithms, e.g., boosting [15], principal component analysis [9], linear and logistic regression [7, 55], support vector machines [38], risk minimization [8, 12], continuous data processing [39], online learning [2, 19, 22], generic deep learning [1, 44] and natural language processing [37], as opposed to *binary algorithms*, such as hashing [20]. As the privacy of training data will be leaked in the hashing scenarios, hashing is facing more critical situations, and there does not exist suitable solutions for hashing. We firstly formulate this practical and challenging *private hashing* problem, and develop Random Flipping for the binary situation.

2.2 Hashing

Hashing is a widely used method for ANN search in large scale multi-media search with encouraging efficiency in both speed and storage. Existing hashing methods can be roughly categorized into unsupervised and supervised hashing. Unsupervised hashing methods [25] learn hash functions from unlabeled data, e.g., LSH [17], SH [51] and ITQ [18]. Locality Sensitive Hashing (LSH) [17] uses random projections as the hash function. Iterative quantization (ITQ) alternatively rotates the projected data and updates the codes to find an optimal rotation for minimizing the quantization loss. Supervised hashing methods [23, 28, 34, 47, 49, 52] leverage supervised information (e.g., similarity matrix or label information) to improve the quality of hash codes, e.g., KSH [33] and SDH [42]. KSH [33]

utilizes the property of kernels to simplify the original problem while achieving better search accuracy.

Inspired by powerful feature representations in learning with deep neural networks [27], deep supervised hashing [6, 16] that adopts deep learning to generate high-level semantic features has been proposed. Deep Pairwise-Supervised Hashing (DPSH) [30] preserves relative similarities between image triplets straightly by integrating feature learning and hash functions in an end-to-end manner. Deep Supervised Hashing (DSH) [32] splits training data into similar pairs and dissimilar pairs to generate similarity correlations and controls the quantization error to further improve the performance. Moreover, HashNet [5] tackles the data imbalance problem between similar and dissimilar pairs and alleviates this problem by adjusting the weights of similar pairs. To additionally accelerate the training procedure, several asymmetric deep hashing methods [41] are proposed, e.g., Asymmetric Deep Supervised Hashing (ADSH) [24], which only learns the hash function for query points to reduce the computational complexity. For deriving compact hash codes, the objective functions in hashing learning are always designed by the following three principles: i) preserving the similarity of data points in the original space, ii) distributing the codes to uniformly fulfill the code space, and iii) generating compact binary codes.

While the aforementioned hashing methods have shown great success in multi-media searching, they ignore the privacy of data. In this paper, to protect privacy which will be potentially revealed, we are the first to investigate the *private hashing* problem and devise two private hashing algorithms with a private noise mechanism. Besides, to theoretically analyze the proposed algorithms, we derive the privacy and stability guarantees for measuring the influence of hash codes by noise addition.

3 METHOD

In this section, we elaborately introduce our noise mechanism for privacy protection and our proposed private hashing methods. Besides, we present theoretical guarantees of our algorithms.

Suppose that we have m query data points and n database points denoted as $\mathbf{X} = \{x_i\}_{i=1}^m$ and $\mathbf{Y} = \{y_j\}_{j=1}^n$. The pairwise information is denoted by $\mathbf{S} \in \{-1, +1\}^{m \times n}$, which is available in the supervised setting. If point x_i and point y_j are similar, $S_{ij} = +1$, otherwise $S_{ij} = -1$. Under this definition, the goal of hashing is to learn binary hash codes $\mathbf{b} \in \{-1, +1\}^c$ for each point, and the corresponding hash function $\mathbf{h}(\cdot)$, where c is the target length of binary codes. We denote $\mathbf{V} \in \{-1, +1\}^{m \times c}$ and $\mathbf{U} \in \{-1, +1\}^{n \times c}$ as the hash codes of the query points and the database points, respectively.

The mean average precision (mAP) is one of the most widely used criteria to evaluate search accuracy. The core idea of mAP is to evaluate ranked lists by averaging the precision at each position. Given m query $\mathbf{V} = [\mathbf{v}_1^T, \dots, \mathbf{v}_m^T] \in \{-1, +1\}^{m \times c}$ and lists of n ranked retrieval database results, mAP is defined as:

$$mAP(\mathbf{V}) = \frac{1}{m} \sum_{i=1}^m \frac{1}{N_i^+} \sum_{j=1}^n \text{precision}(\mathbf{v}_i, \mathbf{U}, S, j) \mathbb{I}(S_{ij} = 1)$$

where N_i^+ represents the count of ground-truth relevant instances in the database for the query \mathbf{v}_i , $\mathbb{I}(\cdot)$ is the indicator function, and $\text{precision}(\mathbf{v}_i, \mathbf{U}, S, j)$ stands for the precision at cut-off j in the

Algorithm 1 RandomFlipping(\mathbf{U}, β)

Input: $\mathbf{U} \in \{-1, +1\}^{n \times c}$ represents the input binary matrix. β is the hyper-parameter for controlling the strength of privacy-protecting.

Output: The perturbed binary matrix \mathbf{U} .

```

1: for  $i \rightarrow 1, \dots, n$  do
2:   for  $j \rightarrow 1, \dots, c$  do
3:     Get a random number  $0 \leq k < 1$ ;
4:     if  $k < \beta$  then
5:        $U_{ij} \leftarrow -U_{ij}$ ;
6:     end if
7:   end for
8: end for
9: return  $\mathbf{U}$ .
```

ranked list. mAP is related to the Hamming distance between query and database which is further influenced by hash codes of database.

3.1 Random Flipping

Here we present our novel noise mechanism, which is the key novelty of this paper and also an important component of the proposed private hashing algorithms.

The intuition behind our novel noise mechanism algorithm is “Imperceptible Lying”. If we are taking a test questioning some private individual information and have to answer either “yes” or “no”, and we want to keep this private answer uncovered, we can simply change our reply with a fixed probability. Similarly, “yes” and “no” can be seen as “+1” and “-1”, while the c -length binary hash codes of an input data can be regarded as c sequential questions, such as “does the input data have black bird head?”, which can be used for identifying the input. Hence, by changing each datum of c -length binary hash codes with a predefined probability β , the individually sensitive information of input data can be protected. This mechanism is suitable for the hashing scenario, in another word, the binary algorithms.

As shown in Algorithm 1, our proposed noise mechanism named as Random Flipping, takes a binary matrix, e.g., hash codes \mathbf{U} , as input, and randomly flips each datum of the input binary matrix with a predefined probability β , i.e., flipping “+1” to “-1” or flipping “-1” to “+1”, and then outputs the private binary matrix \mathbf{U}' . Random Flipping is straightforward and has the following privacy guarantee:

THEOREM 4 (PRIVACY GUARANTEE OF RANDOM FLIPPING). *Algorithm 1 is ϵ -differentially private, given $\beta = e^{-\epsilon}$, the parameter of Random Flipping.*

The above theorem shows that the size of the input matrix does not affect the privacy guarantee, while only the probability β will influence privacy protection. But, perturbing the input matrix will cause unpredictable consequences. To estimate the difference between the input matrix \mathbf{U} and the perturbed matrix \mathbf{U}' which preserves ϵ -differential privacy, we present the following lemma.

LEMMA 5 (STABILITY OF INPUT BINARY MATRIX). *Let $\mathbf{U} \in \{-1, +1\}^{n \times c}$ be the input binary matrix and $\mathbf{U}' \in \{-1, +1\}^{n \times c}$ be the perturbed binary matrix derived by $\mathbf{U}' = \text{RandomFlipping}(\mathbf{U}, \beta)$. The expected*

Algorithm 2 PHashing($\mathbf{Y}, \mathcal{S}, \mathcal{H}(\cdot, \cdot), \beta$)

Input: $\mathbf{Y} = \{y_j\}_{j=1}^n$ represents n data points. $\mathcal{S} \in \{-1, +1\}^{n \times n}$ is the similarity matrix between data points \mathbf{Y} . $\mathcal{H}(\cdot, \cdot)$ is a hashing learning algorithm. β is the hyper-parameter of the Random Flipping shown in Algorithm 1.

Output: $\mathbf{h}(\cdot)$ is the learned hash function derived by $\mathcal{H}(\cdot, \cdot)$. \mathbf{U}' is the private binary hash codes for data points.

```

1: Train: Learn the hash function  $\mathbf{h}(\cdot)$  and binary hash codes  $\mathbf{U}$  for data points by  $\mathbf{h}(\cdot)$ ,  $\mathbf{U} \leftarrow \mathcal{H}(\mathbf{Y}, \mathcal{S})$ .
2: Perturb: Obtain the private binary hash codes  $\mathbf{U}'$  for data points by  $\mathbf{U}' \leftarrow \text{RandomFlipping}(\mathbf{U}, \beta)$ .
3: return  $\mathbf{h}(\cdot)$  and  $\mathbf{U}'$ .
```

F -norm $\|\mathbf{U}' - \mathbf{U}\|_F$ satisfies

$$\mathbb{E} \left[\|\mathbf{U}' - \mathbf{U}\|_F \right] = 2\sqrt{\beta nc}.$$

The above lemma can be further expanded as

$$\mathbb{E} \left[\|\mathbf{U}' - \mathbf{U}\|_F \right] = 2\sqrt{\beta nc} = 2\sqrt{\frac{nc}{e^\epsilon}}.$$

It indicates that the stability of the input matrix is proportional to the size of the input matrix and the privacy parameter ϵ , and inversely proportional to the strength of privacy-protecting, which is obvious in the sense if we want to hide more individual information, more stability will lose.

3.2 Private Hashing (PHashing)

In this part, we present our PHashing (private hashing) algorithm which can equip all the non-private hashing algorithms with the ability of protecting privacy.

Inspired by the post-processing theorem [14], PHashing combines our proposed novel *Random Flipping* noise mechanism with the non-private hashing algorithms to achieve private hashing. Specifically, after the training of a non-private hashing algorithm $\mathcal{H}(\cdot, \cdot)$ and obtaining learned hash codes \mathbf{U} and the hash function $\mathbf{h}(\cdot)$, we perturb the learned hash codes \mathbf{U} for privacy by $\text{RandomFlipping}(\cdot, \cdot)$. In this way, the perturbed hash codes \mathbf{U}' is private. PHashing is widely applicable as we can employ all the existing hashing methods in the training procedure, and PHashing involves the following steps:

1. *Train:* Given any non-private hashing algorithm $\mathcal{H}(\cdot, \cdot)$, we train it until convergence, and obtain the learned hash function $\mathbf{h}(\cdot)$ and the binary hash codes \mathbf{U} for the database;
2. *Perturb:* Given the privacy parameter β , we employ the Random Flipping mechanism shown in Algorithm 1 to derive the private hash codes \mathbf{U}' .

The above procedure is summarized in Algorithm 2.

3.2.1 Privacy and Stability Analysis. We now present the privacy and stability guarantees. With Theorem 2, we can easily obtain the following privacy guarantee.

THEOREM 6 (PRIVACY GUARANTEE). *With $\beta = e^{-\epsilon}$, Algorithm 2 is ϵ -differentially private.*

Algorithm 3 PITQ(Y, β, E)

Input: $Y = \{y_j\}_{j=1}^n$ represents n centralized data points. β is the hyper-parameter of Random Flipping. E is the iterative epoch.

Output: U and R are the private binary hash codes for data points and the rotation matrix.

- 1: Employ proper embedding methods, *e.g.*, PCA, to project the database data Y and return the optimal projection matrix W ;
- 2: **for** $i \leftarrow 1, \dots, E$ **do**
- 3: *Fix R and update U:* $U \leftarrow \text{sign}(YWR)$;
- 4: *Perturb U:* $U \leftarrow \text{RandomFlipping}(U, \beta)$;
- 5: *Fix U and update R:* first compute the SVD of $U^T(YW)$ as $U^T(YW) = N\Omega N^T$, and then $R \leftarrow NN^T$;
- 6: **end for**
- 7: **return** R and $U = \text{sign}(YWR)$.

Now, we focus on the stability of PHashing. Due to sparse and random properties of hash codes and the mechanism of calculating mAP, it is hard to measure the performance drop of an unconstrained algorithm, while if we add some assumptions about hashing algorithms and hash codes, the guarantee will not suit most of scenarios. Thus, for deriving general theorems estimating the privacy impact, we characterize the influence on the hash codes and Hamming distance between query and database, which directly influence the search accuracy, *i.e.*, mAP, and present the following lemma.

LEMMA 7 (STABILITY OF HASH CODES AND HAMMING DISTANCE). *Let $v \in \{-1, +1\}^c$ be a query code containing only one point. Let $d' \in \{0, \dots, c\}^n$ be the Hamming distance between v and the private hash code $U' \in \{-1, +1\}^{n \times c}$ for the database, and $d \in \{0, \dots, c\}^{n \times 1}$ be the Hamming distance between v and the non-private hash code $U \in \{-1, +1\}^{n \times c}$. The expected F -norm $\|U - U'\|_F$ satisfies Lemma 5, and expected 1-norm $\|d - d'\|_1$ satisfies*

$$\mathbb{E} \left[\|d - d'\|_1 \right] \leq \frac{n^{3/2}c}{\sqrt{e^\epsilon}} = n^{3/2}c\sqrt{\beta}.$$

3.3 Private Iterative Quantization (PITQ)

In this section, to reduce performance drop, we devise a private unsupervised iterative hashing algorithm. The proposed algorithm is named as PITQ (private iterative quantization), whose performance is better than PHashing with $\mathcal{H}(\cdot, \cdot) = \text{ITQ}$. Different from PHashing, PITQ distributes the noise to each iteration and alternately updates the interdependent parameters for improving the search performance, while PHashing adds the noise directly on the output.

As one of the most popular unsupervised conventional hashing methods, ITQ [18] first obtains the binary hash codes $U = \text{sign}(YWR)$ by rotating and quantizing the projected centralized data YW , where the projection matrix W is learned by embedding methods, *e.g.*, PCA, and then updates the rotation matrix R to minimize the quantization loss (Eq. (4)) iteratively. This procedure totally utilizes the structure of data points in the original space.

Inspired by the power of iterative discrete learning in ITQ, PITQ is motivated as follows: in each iteration $i \in [E]$, PITQ first obtains

the binary hash codes $U = \text{sign}(YWR)$ by rotating (rotation matrix R) and quantizing the projected centralized data YW , where the projection matrix W is learned by embedding methods, *e.g.*, PCA. Now, for the sake of the privacy in the database, PITQ perturbs the learned hash codes U utilizing Random Flipping (Algorithm 1). After that, PITQ updates the rotation matrix R to minimize the quantization loss (Eq. (4)). In this way, the rotation matrix R and hash codes U are alternatively updated. In the last iteration, PITQ returns the learned rotation matrix R and private hash codes $U = \text{sign}(YWR)$, where the $\text{sign}(\cdot)$ is the signum function. PITQ minimizes the following loss function:

$$Q(U, R) = \|U - YWR\|_F^2. \quad (4)$$

Specifically, PITQ involves the following steps:

1. *Fix R and update U:* Expanding Eq. (4), we have

$$\begin{aligned} Q(U, R) &= \|U\|_F^2 + \|YW\|_F^2 - 2\text{tr}(UR^T(YW)^T) \\ &= nc + \|YW\|_F^2 - 2\text{tr}(UR^T(YW)^T). \end{aligned} \quad (5)$$

Minimizing (5) is equivalent to maximizing $\text{tr}(UR^T(YW)^T)$. It is obvious that the trace is maximized by letting $U = \text{sign}(YWR)$;

2. *Perturb U:* Given the parameter β , we employ *Random Flipping* mechanism (Algorithm 1) to obtain private hash codes U ;
3. *Fix U and update R:* Now the objective function (5) corresponds to the classic Orthogonal Procrustes problem [40], which tries to find a rotation for aligning two sets. Thus, we first compute the SVD of $U^T(YW)$ as $U^T(YW) = N\Omega N^T$, and then set $R \leftarrow NN^T$;
4. Loop the step 1-3 for predefined iterations E , then return R and $U = \text{sign}(YWR)$.

As summarized in Algorithm 3, our PITQ algorithm alternately updates the parameters R and U for several predefined iterations E to find a locally optimal solution. The noise added when perturbing learned binary codes U in each iteration may help the discrete optimization to find a better local optimum in some cases practically, as shown in Tab. 2.

3.3.1 Privacy Analysis. Similarly, based on Theorem 2 and Theorem 3 of differential privacy, we have the following privacy guarantee.

THEOREM 8 (PRIVACY GUARANTEE). *Algorithm 3 is ϵ -differentially private, given the parameter $\beta = e^{-\epsilon/E}$ of Random Flipping and the iterations E .*

Nevertheless, due to the limitation of analysis on discrete optimization, it is still hard to derive a general guarantee. But it is easy to show that the stability of learned hash codes U in each iteration satisfies Lemma 5. Also, empirical results show that PITQ achieves equal performances with non-private algorithms with barely performance drop.

4 EXPERIMENTS

In this section, we analyze the search performance and privacy cost by conducting comprehensive experiments of several non-private hashing algorithms and our proposed private hashing algorithms.

Table 1: Comparisons of mAP w.r.t. the different lengths of bits on *ImageNet*, *CIFAR-10*, and *NUS-WIDE* of private versions (PHashing) and non-private versions (original algorithms) of several hashing algorithms. The rows ticked with private represents the private versions (PHashing) with ϵ -differential privacy.

Method	Private?	ϵ	ImageNet				CIFAR-10				NUS-WIDE			
			12 bits	24 bits	32 bits	48 bits	12 bits	24 bits	32 bits	48 bits	12 bits	24 bits	32 bits	48 bits
LSH		0	0.0520	0.1134	0.1625	0.2268	0.1162	0.1215	0.1264	0.1244	0.1850	0.2222	0.2088	0.2116
PHashing (LSH)	✓	1	0.0130	0.0150	0.0162	0.0186	0.1036	0.1047	0.1057	0.1054	0.1704	0.1783	0.1752	0.1755
	✓	2	0.0280	0.0532	0.0761	0.1140	0.1111	0.1147	0.1183	0.1170	0.1789	0.2052	0.1954	0.1968
	✓	4	0.0475	0.1033	0.1492	0.2110	0.1154	0.1205	0.1252	0.1234	0.1841	0.2199	0.2069	0.2095
ITQ		0	0.2917	0.4929	0.5586	0.6337	0.1544	0.1607	0.1630	0.1656	0.3189	0.3264	0.3277	0.3322
PHashing (ITQ)	✓	1	0.0213	0.0346	0.0451	0.0715	0.1110	0.1134	0.1148	0.1167	0.1960	0.1987	0.1999	0.2029
	✓	2	0.1429	0.3192	0.3982	0.5116	0.1374	0.1442	0.1471	0.1509	0.2718	0.2805	0.2835	0.2907
	✓	4	0.2684	0.4710	0.5389	0.6208	0.1522	0.1585	0.1612	0.1638	0.3128	0.3207	0.3223	0.3271
SH		0	0.1468	0.2545	0.3039	0.3739	0.1316	0.1289	0.1287	0.1274	0.2270	0.2134	0.2092	0.2066
PHashing (SH)	✓	1	0.0142	0.0159	0.0168	0.0187	0.1061	0.1055	0.1054	0.1048	0.1724	0.1716	0.1716	0.1711
	✓	2	0.0535	0.0917	0.1133	0.1591	0.1204	0.1185	0.1184	0.1171	0.1936	0.1890	0.1875	0.1862
	✓	4	0.1275	0.2251	0.2712	0.3415	0.1299	0.1273	0.1272	0.1259	0.2205	0.2090	0.2050	0.2028
KSH		0	0.1458	0.1913	0.2150	0.2544	0.2339	0.2586	0.2669	0.2751	0.2891	0.2813	0.2773	0.3055
PHashing (KSH)	✓	1	0.0148	0.0182	0.0206	0.0257	0.1228	0.1320	0.1387	0.1487	0.1886	0.1893	0.1889	0.2022
	✓	2	0.0591	0.0868	0.0982	0.1345	0.1911	0.2187	0.2306	0.2444	0.2470	0.2451	0.2433	0.2737
	✓	4	0.1285	0.1718	0.1937	0.2345	0.2287	0.2537	0.2626	0.2717	0.2833	0.2763	0.2728	0.3017
SDH		0	0.1296	0.1732	0.2245	0.3058	0.1746	0.2115	0.2033	0.2153	0.1659	0.1659	0.2519	0.2656
PHashing (SDH)	✓	1	0.0092	0.0198	0.0219	0.0301	0.1181	0.1205	0.1210	0.1286	0.1665	0.1662	0.1690	0.1691
	✓	2	0.0493	0.0769	0.1012	0.1435	0.1525	0.1775	0.1740	0.1912	0.1662	0.1656	0.1798	0.1797
	✓	4	0.1242	0.1688	0.2117	0.2981	0.1710	0.2069	0.1994	0.2126	0.1659	0.1657	0.2266	0.2240
DPSH		0	0.2591	0.5038	0.5980	0.6970	0.6872	0.7024	0.7281	0.7437	0.4169	0.4911	0.5103	0.5438
PHashing (DPSH)	✓	1	0.0204	0.0326	0.0460	0.0751	0.1920	0.2418	0.2816	0.3616	0.2358	0.2724	0.2999	0.3269
	✓	2	0.1196	0.3108	0.4307	0.5819	0.5720	0.6350	0.6700	0.7059	0.3704	0.4503	0.4994	0.5011
	✓	4	0.2358	0.4797	0.5782	0.6851	0.6752	0.6991	0.7196	0.7386	0.4119	0.4968	0.5085	0.5123
DSH		0	0.6773	0.7164	0.7053	0.7202	0.7230	0.7644	0.7746	0.7920	0.5127	0.5093	0.5184	0.5168
PHashing (DSH)	✓	1	0.0245	0.0515	0.0671	0.1264	0.1749	0.2903	0.3435	0.4302	0.2687	0.3225	0.3492	0.3864
	✓	2	0.3700	0.6262	0.6379	0.6950	0.5264	0.7109	0.7242	0.7411	0.4734	0.4979	0.5039	0.5050
	✓	4	0.6555	0.7122	0.6997	0.7173	0.6958	0.7548	0.7579	0.7653	0.5100	0.5078	0.5158	0.5142
HashNet		0	0.3005	0.4903	0.5503	0.6313	0.7261	0.7614	0.7858	0.7950	0.5516	0.5663	0.5723	0.5782
PHashing (HashNet)	✓	1	0.0231	0.0320	0.0366	0.0638	0.2206	0.3139	0.3917	0.4925	0.2844	0.3325	0.3570	0.3963
	✓	2	0.1536	0.3113	0.3752	0.5054	0.6413	0.7139	0.7489	0.7634	0.5083	0.5337	0.5399	0.5512
	✓	4	0.2743	0.4684	0.5308	0.6181	0.7162	0.7491	0.7710	0.7784	0.5469	0.5623	0.5681	0.5745
ADSH		0	0.0630	0.1696	0.3029	0.5720	0.6599	0.7413	0.7590	0.7672	0.5691	0.5938	0.6022	0.6137
PHashing (ADSH)	✓	1	0.0111	0.0129	0.0249	0.0618	0.1745	0.2802	0.3503	0.3980	0.2738	0.3196	0.3460	0.3750
	✓	2	0.3594	0.0745	0.0648	0.2149	0.5119	0.6969	0.7260	0.7432	0.5031	0.5372	0.5475	0.5249
	✓	4	0.5834	0.1090	0.2139	0.4864	0.6379	0.7360	0.7519	0.7645	0.5596	0.5797	0.5877	0.5581

Table 2: Comparisons of mAP w.r.t. the different lengths of bits on *ImageNet*, *CIFAR-10*, *NUS-WIDE* of our proposed PITQ and non-private ITQ method.

Method	Private?	ϵ	ImageNet				CIFAR-10				NUS-WIDE			
			12 bits	24 bits	32 bits	48 bits	12 bits	24 bits	32 bits	48 bits	112 bits	24 bits	32 bits	48 bits
ITQ		0	0.2917	0.4929	0.5586	0.6337	0.1544	0.1607	0.1630	0.1656	0.3189	0.3264	0.3277	0.3322
PITQ	✓	0.25	0.2921	0.4887	0.5628	0.6398	0.1543	0.1608	0.1625	0.1655	0.3221	0.3257	0.3275	0.3311
	✓	0.5	0.2918	0.4880	0.5633	0.6406	0.1544	0.1608	0.1626	0.1657	0.3223	0.3257	0.3277	0.3313
	✓	1	0.2922	0.4883	0.5598	0.6389	0.1543	0.1607	0.1626	0.1659	0.3217	0.3256	0.3276	0.3310
	✓	4	0.2928	0.4881	0.5663	0.6374	0.1538	0.1610	0.1624	0.1659	0.3224	0.3257	0.3273	0.3306

4.1 Datasets and Evaluation Details

We evaluate the performance of our proposed methods on five vision datasets, three generic datasets and two fine-grained datasets.

- **ImageNet [11] (Generic Dataset):** ImageNet is a classical benchmark image dataset and contains over 1.2M images in the training set and 50K images in the validation set. Each image has a single label of one of the 1,000 categories. We

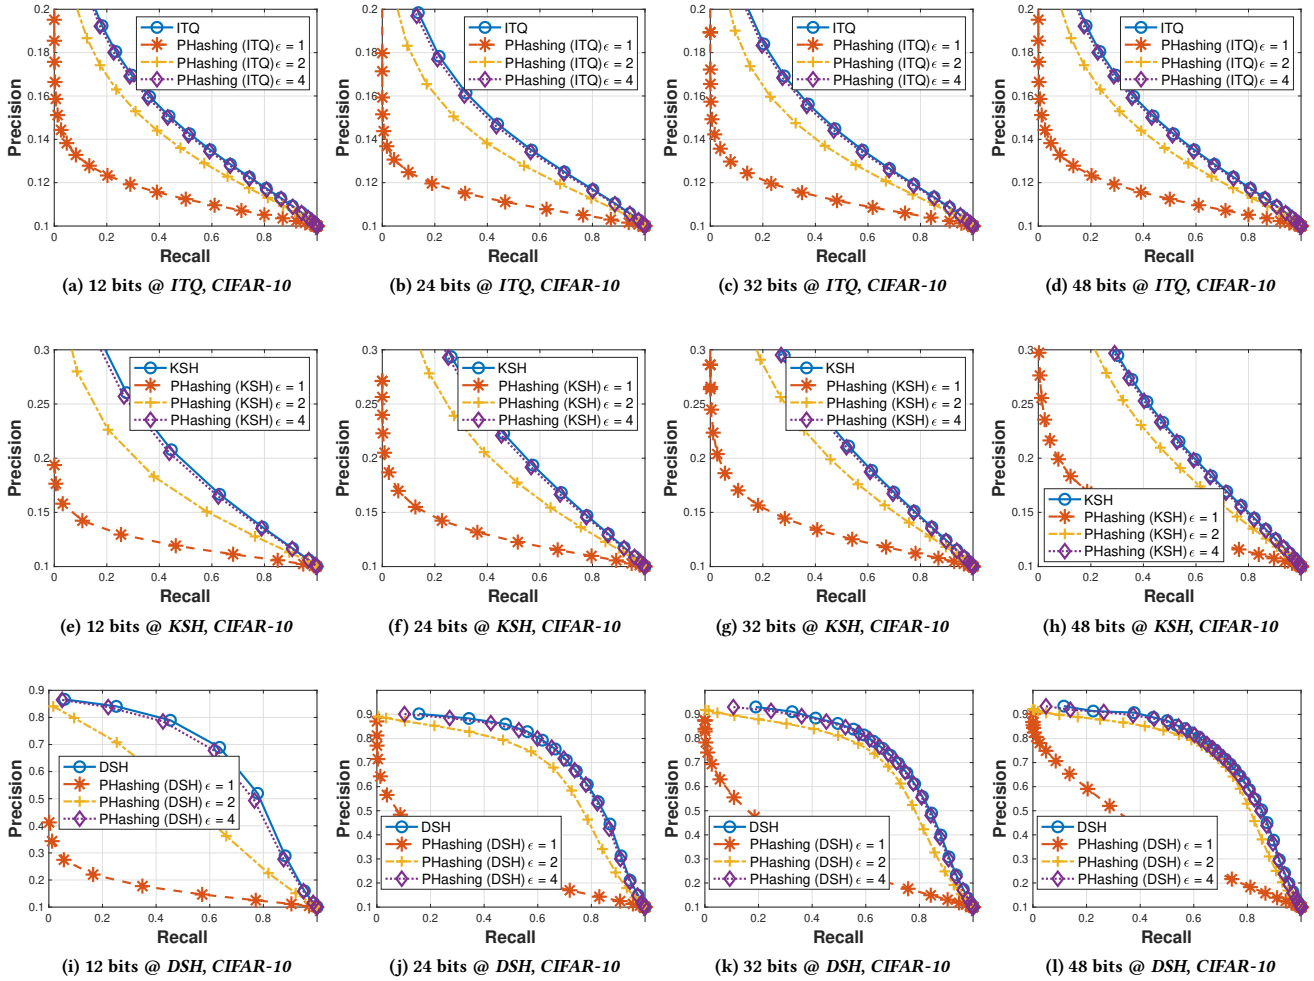


Figure 2: Performance of precision-recall curves on *CIFAR-10* comparing several non-private and private methods, i.e., ITQ and PHashing (ITQ), KSH and PHashing (KSH), DSH and PHashing (DSH). The four sub-figures in each row are the precision-recall curves for 12 bits, 24 bits, 32 bits, and 48 bits, respectively. Best view in color.

randomly select 100 categories, use all the images of these categories in the training set as the database, and use all the images in the validation set as the query.

- **CIFAR-10** [26] (*Generic Dataset*): This dataset contains 60,000 32×32 color images. Each image belongs to one of the ten classes. We follow the protocol proposed in [26], where 10,000 images (1000 images per class) in the test set is used for evaluation, with the remaining 50,000 images as database points.
- **NUS-WIDE** [10] (*Generic Dataset*): This is a public Web image dataset which contains 269,648 images downloaded from Flickr.com. Each image is manually annotated by multi-classes of 81 categories for evaluating retrieval models. We randomly sample 5,000 images as queries, with the remaining images as the database.

- **CUB** [48] (*Fine-grained Dataset*): This is a fine-grained dataset focusing on birds with 11,788 images for 200 various birds species. Database and test sets have 5,994 and 5,794 images.
- **Aircraft** [36] (*Fine-grained Dataset*): Aircraft contains 10,000 images of aircraft, with 100 images for each of 100 different aircraft model variants. We also follow the standard split in [36].

For the above datasets, two images are treated as a ground-truth similar pair if they share at least one label. We evaluate the search performance by adopting two evaluation metrics: mean Average Precision (mAP) and Precision-recall Curves (PR Curves) based on lookup. PR Curves are obtained by varying the Hamming radius from 0 to c with the step-size of 1 to demonstrate the hash lookup procedure. All the data are reported with average values running five times.




Figure 3: mAP comparison on *ImageNet*, *CIFAR-10*, and *NUS-WIDE* with $\epsilon = 1$ for some private methods. For a clear presentation, we omit the legend of Fig. 3(a) and Fig. 3(c) whose legends are the same as Fig. 3(b). Best view in color.

We use several non-private state-of-the-art hashing methods to compare the performance of PHashing, including shallow methods, *i.e.*, LSH [17], ITQ [18], SH [51], SDH [42] and KSH [33], and deep supervised methods, *i.e.*, DSH [32], DPSH [30], HashNet [5], and ADSH [24]. To simplify the notation, we denote the hashing algorithm extended by PHashing as PHashing (\cdot). For example, PHashing (ITQ) represents the private version of ITQ extended by PHashing, while $\mathcal{H}(\cdot, \cdot)$ is ITQ. For deep hashing methods, we use raw images resized to 224×224 as inputs. For traditional shallow methods, we extract 4096-dimensional deep features by the VGG-D [45] model pre-trained with ImageNet to conduct fair comparisons. Besides, for all the non-private hashing methods, we employ hyper-parameters introduced in their papers.

4.2 Performance Drop Caused by Privacy Protection – Parameter Sensitivity

To validate the cost of privacy protection, we choose ϵ from the set $\{1, 2, 4\}$ of PHashing and $\{0.25, 0.5, 1, 4\}$ of PITQ. The mAP of different methods on three generic datasets is presented in Tab. 1 and Tab. 2 with code-lengths of 12 bits, 24 bits, 32 bits, and 48 bits, respectively. PR Curves of some methods are shown in Fig. 2. The left mAP results and PR Curves are presented in Appendix due to the limitation of space.

Search accuracy: In Tab. 1, the results demonstrate that the search accuracy is inversely proportional to the privacy parameter ϵ , which is consistent with the theoretical analysis. And as the parameter ϵ increases, accuracies of PHashing are approaching non-private accuracies. The search accuracy in Tab. 2 shows that distributing the noise to several iterations significantly improves the search performance, as our PITQ achieves equal performances with the non-private ITQ and outperforms PHashing (ITQ).

PR Curves: Fig. 2 presents several precision-recall curves for PHashing using ITQ, KSH, and DSH on *CIFAR-10*. As the curves show, the private curves are approaching the non-private curves with ϵ increasing. The private curves of $\epsilon = 4$ are almost the same as non-private curves.

4.3 Search Performance

To validate the search performance, we set $\epsilon = 1$ for PHashing and PITQ. The mAP of different methods on three generic datasets is presented in Fig. 3 with the code-lengths of 12 bits, 24 bits, 32 bits, and 48 bits, respectively. Due to the limitation of space, the mAP comparison of fine-grained datasets and PR Curves of all five datasets are presented in Appendix.

Search accuracy: In Fig. 3, PITQ outperforms PHashing (ITQ). As PITQ distributes the noise to each iteration, the performance drop is reduced and PITQ achieves satisfying performance. Due to the power of supervised learning and deep learning, PITQ is worse than some deep supervised private hashing methods, which is reasonable.

5 CONCLUSION

In this paper, we presented a novel noise mechanism and two private algorithms for the *private hashing* task. One of the key contributions was the novel *Random Flipping* algorithm designed for the binary situation, in other words, hashing scenario, which ensured ϵ -differential privacy. Besides, our first algorithm, PHashing, showed its ability to extend any non-private hashing methods to preserve ϵ -differential privacy with acceptable performance drops. Also, our second algorithm, PITQ, improved the search accuracy by distributing the noise generated by Random Flipping to each iteration while achieving equal performances comparing with the non-private ITQ, as the noise could practically help to find a better local optimum with ϵ -differential privacy. Experimental results on diverse vision datasets showed the effectiveness of our methods. In the future, we would like to explore novel private hashing methods under the supervised setting and reduce performance drop.

ACKNOWLEDGMENTS

This work was partially supported by NSFC-NRF Joint Research Project (61861146001), the Fundamental Research Funds for the Central Universities (14380074), and the Collaborative Innovation Center of Novel Software Technology and Industrialization.

REFERENCES

- [1] Martin Abadi, Andy Chu, Ian Goodfellow, H. Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. 2016. Deep Learning with Differential Privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. Association for Computing Machinery, New York, NY, USA, 308–318. <https://doi.org/10.1145/2976749.2978318>
- [2] Naman Agarwal and Karan Singh. 2017. The Price of Differential Privacy for Online Learning. In *Proceedings of the 34th International Conference on Machine Learning*. PMLR, International Convention Centre, Sydney, Australia, 32–40.
- [3] Alexandr Andoni and Piotr Indyk. 2008. Near-Optimal Hashing Algorithms for Approximate Nearest Neighbor in High Dimensions. *Commun. ACM* 51, 1 (Jan. 2008), 117–122. <https://doi.org/10.1145/1327452.1327494>
- [4] Alexandr Andoni and Ilya Razenshteyn. 2015. Optimal Data-Dependent Hashing for Approximate Near Neighbors. In *Proceedings of the 47th Annual ACM Symposium on Theory of Computing*. Association for Computing Machinery, New York, NY, USA, 793–801. <https://doi.org/10.1145/2746539.2746553>
- [5] Zhangjie Cao, Mingsheng Long, Jianmin Wang, and Philip S. Yu. 2017. HashNet: Deep Learning to Hash by Continuation. In *The IEEE International Conference on Computer Vision*. 5608–5617.
- [6] Zhangjie Cao, Ziping Sun, Mingsheng Long, Jianmin Wang, and Philip S. Yu. 2018. Deep Priority Hashing. In *Proceedings of the 26th ACM International Conference on Multimedia*. Association for Computing Machinery, New York, NY, USA, 1653–1661. <https://doi.org/10.1145/3240508.3240543>
- [7] Kamalika Chaudhuri and Claire Monteleoni. 2009. Privacy-preserving logistic regression. In *Advances in Neural Information Processing Systems 21*, D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou (Eds.). Curran Associates, Inc., 289–296.
- [8] Kamalika Chaudhuri, Claire Monteleoni, and Anand D. Sarwate. 2011. Differentially Private Empirical Risk Minimization. *Journal of Machine Learning Research* 12 (July 2011), 1069–1109.
- [9] Kamalika Chaudhuri, Anand D. Sarwate, and Kaushik Sinha. 2013. A Near-Optimal Algorithm for Differentially-Private Principal Components. *Journal of Machine Learning Research* 14, 1 (Jan. 2013), 2905–2943.
- [10] Tat-Seng Chua, Jinhui Tang, Richang Hong, Haojie Li, Zhiping Luo, and Yantao Zheng. 2009. NUS-WIDE: A Real-World Web Image Database from National University of Singapore. In *Proceedings of the ACM International Conference on Image and Video Retrieval*. Association for Computing Machinery, New York, NY, USA, Article 48, 9 pages. <https://doi.org/10.1145/1646396.1646452>
- [11] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. ImageNet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*. 248–255.
- [12] John C. Duchi, Michael I. Jordan, and Martin J. Wainwright. 2014. Privacy Aware Learning. *J. ACM* 61, 6, Article 38 (Dec. 2014), 57 pages. <https://doi.org/10.1145/2666468>
- [13] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating Noise to Sensitivity in Private Data Analysis. In *Proceedings of the Third Conference on Theory of Cryptography*. Springer-Verlag, Berlin, Heidelberg, 265–284. https://doi.org/10.1007/11681878_14
- [14] Cynthia Dwork and Aaron Roth. 2014. The Algorithmic Foundations of Differential Privacy. *Foundations and Trends in Theoretical Computer Science* 9, 3–4 (Aug. 2014), 211–407. <https://doi.org/10.1561/04000000042>
- [15] Cynthia Dwork, Guy N. Rothblum, and Salil Vadhan. 2010. Boosting and Differential Privacy. In *Proceedings of the 2010 IEEE 51st Annual Symposium on Foundations of Computer Science*. IEEE, 51–60.
- [16] Lianli Gao, Jingkuan Song, Fuhao Zou, Dongxiang Zhang, and Jie Shao. 2015. Scalable Multimedia Retrieval by Deep Learning Hashing with Relative Similarity Learning. In *Proceedings of the 23rd ACM International Conference on Multimedia*. Association for Computing Machinery, New York, NY, USA, 903–906. <https://doi.org/10.1145/2733373.2806360>
- [17] Aristides Gionis, Piotr Indyk, and Rajeev Motwani. 1999. Similarity Search in High Dimensions via Hashing. In *Proceedings of the 25th International Conference on Very Large Data Bases*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 518–529.
- [18] Yunchao Gong, Svetlana Lazebnik, Albert Gordo, and Florent Perronnin. 2013. Iterative Quantization: A Procrustean Approach to Learning Binary Codes for Large-Scale Image Retrieval. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 35, 12 (2013), 2916–2929.
- [19] Abhradeep Guha Thakurta and Adam Smith. 2013. (Nearly) Optimal Algorithms for Private Online Learning in Full-information and Bandit Settings. In *Advances in Neural Information Processing Systems 26*, C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger (Eds.). Curran Associates, Inc., 2733–2741.
- [20] Kun He, Fatih Cakir, Sarah Adel Bargal, and Stan Sclaroff. 2018. Hashing as Tie-Aware Learning to Rank. In *2018 IEEE Conference on Computer Vision and Pattern Recognition*. 4023–4032.
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition*. 770–778.
- [22] Prateek Jain, Pravesh Kothari, and Abhradeep Thakurta. 2012. Differentially Private Online Learning. In *Proceedings of the 25th Annual Conference on Learning Theory*. Shie Mannor, Nathan Srebro, and Robert C. Williamson (Eds.). PMLR, Edinburgh, Scotland, 24.1–24.34.
- [23] I-Hong Jhuo. 2019. Supervised Set-to-Set Hashing in Visual Recognition. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence*. International Joint Conferences on Artificial Intelligence Organization, 803–810. <https://doi.org/10.24963/ijcai.2019/113>
- [24] Qing-Yuan Jiang and Wu-Jun Li. 2018. Asymmetric Deep Supervised Hashing. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, New Orleans, Louisiana, USA*, Sheila A. McIlraith and Kilian Q. Weinberger (Eds.). AAAI Press, 3342–3349. <https://www.aaai.org/ocs/index.php/AAAI/AAAI18/paper/view/17296>
- [25] Qing-Yuan Jiang and Wu-Jun Li. 2015. Scalable Graph Hashing with Feature Transformation. In *Proceedings of the 24th International Conference on Artificial Intelligence*. AAAI Press, 2248–2254.
- [26] Alex Krizhevsky and Geoffrey Hinton. 2009. *Learning multiple layers of features from tiny images*. Technical Report. Citeseer.
- [27] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. 2017. ImageNet Classification with Deep Convolutional Neural Networks. *Commun. ACM* 60, 6 (May 2017), 84–90. <https://doi.org/10.1145/3065386>
- [28] Brian Kulis and Trevor Darrell. 2009. Learning to Hash with Binary Reconstructive Embeddings. In *Advances in Neural Information Processing Systems 22*, Y. Bengio, D. Schuurmans, J. D. Lafferty, C. K. I. Williams, and A. Culotta (Eds.). Curran Associates, Inc., 1042–1050.
- [29] Yann Lecun, Leon Bottou, Yoshua Bengio, and Patrick Haffner. 1998. Gradient-based learning applied to document recognition. *Proc. IEEE* 86, 11 (1998), 2278–2324.
- [30] Wu-Jun Li, Sheng Wang, and Wang-Cheng Kang. 2016. Feature Learning Based Deep Supervised Hashing with Pairwise Labels. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*. AAAI Press, 1711–1717.
- [31] Jie Lin, Zechao Li, and Jinhui Tang. 2017. Discriminative Deep Hashing for Scalable Face Image Retrieval. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence*. 2266–2272. <https://doi.org/10.24963/ijcai.2017/315>
- [32] Haomiao Liu, Ruiping Wang, Shiguang Shan, and Xilin Chen. 2016. Deep Supervised Hashing for Fast Image Retrieval. In *2016 IEEE Conference on Computer Vision and Pattern Recognition*. 2064–2072.
- [33] Wei Liu, Jun Wang, Rongrong Ji, Yu-Gang Jiang, and Shih-Fu Chang. 2012. Supervised hashing with kernels. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*. 2074–2081.
- [34] Xingbo Liu, Xiushan Nie, Quan Zhou, Xiaoming Xi, Lei Zhu, and Yilong Yin. 2019. Supervised Short-Length Hashing. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence*. International Joint Conferences on Artificial Intelligence Organization, 3031–3037. <https://doi.org/10.24963/ijcai.2019/420>
- [35] Xu Lu, Lei Zhu, Zhiyong Cheng, Jingjing Li, Xiushan Nie, and Huaxiang Zhang. 2019. Flexible Online Multi-Modal Hashing for Large-Scale Multimedia Retrieval. In *Proceedings of the 27th ACM International Conference on Multimedia*. Association for Computing Machinery, New York, NY, USA, 1129–1137. <https://doi.org/10.1145/3343031.3350999>
- [36] Subhransu Maji, Esa Rahtu, Juho Kannala, Matthew Blaschko, and Andrea Vedaldi. 2013. *Fine-Grained Visual Classification of Aircraft*. Technical Report. arXiv:cs-cv/1306.5151
- [37] H. Brendan McMahan, Daniel Ramage, Kunal Talwar, and Li Zhang. 2018. Learning Differentially Private Recurrent Language Models. In *International Conference on Learning Representations*.
- [38] Benjamin I. P. Rubinstein, Peter L. Bartlett, Ling Huang, and Nina Taft. 2012. Learning in a Large Function Space: Privacy-Preserving Mechanisms for SVM Learning. *Journal of Privacy and Confidentiality* 4, 1 (Jul. 2012). <https://doi.org/10.29012/jpc.v4i1.612>
- [39] Anand D. Sarwate and Kamalika Chaudhuri. 2013. Signal Processing and Machine Learning with Differential Privacy: Algorithms and Challenges for Continuous Data. *IEEE Signal Processing Magazine* 30, 5 (2013), 86–94.
- [40] Peter H Schönemann. 1966. A generalized solution of the orthogonal procrustes problem. *Psychometrika* 31, 1 (1966), 1–10.
- [41] Fumin Shen, Xin Gao, Li Liu, Yang Yang, and Heng Tao Shen. 2017. Deep Asymmetric Pairwise Hashing. In *Proceedings of the 25th ACM International Conference on Multimedia*. Association for Computing Machinery, New York, NY, USA, 1522–1530. <https://doi.org/10.1145/3123266.3123345>
- [42] Fumin Shen, Chunhua Shen, Wei Liu, and Heng Tao Shen. 2015. Supervised Discrete Hashing. In *2015 IEEE Conference on Computer Vision and Pattern Recognition*. 37–45.
- [43] Xiaobo Shen, Fumin Shen, Quan-Sen Sun, and Yun-Hao Yuan. 2015. Multi-View Latent Hashing for Efficient Multimedia Search. In *Proceedings of the 23rd ACM International Conference on Multimedia*. Association for Computing Machinery, New York, NY, USA, 831–834. <https://doi.org/10.1145/2733373.2806342>

- [44] Reza Shokri and Vitaly Shmatikov. 2015. Privacy-Preserving Deep Learning. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*. Association for Computing Machinery, New York, NY, USA, 1310–1321. <https://doi.org/10.1145/2810103.2813687>
- [45] Karen Simonyan and Andrew Zisserman. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv:1409.1556* (2014).
- [46] Ge Song and Xiaoyang Tan. 2017. Hierarchical deep hashing for image retrieval. *Frontiers of Computer Science* 11, 2 (2017), 253–265. <https://doi.org/10.1007/s11704-017-6537-3>
- [47] Jingkuan Song, Lianli Gao, Yan Yan, Dongxiang Zhang, and Nicu Sebe. 2015. Supervised Hashing with Pseudo Labels for Scalable Multimedia Retrieval. In *Proceedings of the 23rd ACM International Conference on Multimedia*. Association for Computing Machinery, New York, NY, USA, 827–830. <https://doi.org/10.1145/2733373.2806341>
- [48] Catherine Wah, Steve Branson, Peter Welinder, and Serge Belongie Pietro Perona. 2011. *The Caltech-UCSD Birds-200-2011 Dataset*. Technical Report CNS-TR-2011-001. California Institute of Technology.
- [49] Jun Wang, Sanjiv Kumar, and Shih-Fu Chang. 2010. Semi-supervised hashing for scalable image retrieval. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. 3424–3431.
- [50] Yimu Wang, Renjie Song, Xiu-Shen Wei, and Lijun Zhang. 2020. An Adversarial Domain Adaptation Network For Cross-Domain Fine-Grained Recognition. In *2020 IEEE Winter Conference on Applications of Computer Vision*. 1217–1225.
- [51] Yair Weiss, Antonio Torralba, and Rob Fergus. 2009. Spectral Hashing. In *Advances in Neural Information Processing Systems 21*, D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou (Eds.). Curran Associates, Inc., 1753–1760.
- [52] Liang Xie, Jialie Shen, Jungong Han, Lei Zhu, and Ling Shao. 2017. Dynamic Multi-View Hashing for Online Image Retrieval. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*. AAAI Press, 3133–3139.
- [53] Xinyu Yan, Lijun Zhang, and Wu-Jun Li. 2017. Semi-Supervised Deep Hashing with a Bipartite Graph. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*. 3238–3244.
- [54] Chengyuan Zhang, Lei Zhu, Shichao Zhang, and Weiren Yu. 2020. TDHPPIR: An Efficient Deep Hashing Based Privacy-Preserving Image Retrieval Method. *Neurocomputing* 406 (2020), 386 – 398. <https://doi.org/10.1016/j.neucom.2019.11.119>
- [55] Jun Zhang, Zhenjie Zhang, Xiaokui Xiao, Yin Yang, and Marianne Winslett. 2012. Functional Mechanism: Regression Analysis under Differential Privacy. *Proceedings of the VLDB Endowment* 5, 11 (July 2012), 1364–1375. <https://doi.org/10.14778/2350229.2350253>
- [56] Ruimao Zhang, Liang Lin, Rui Zhang, Wangmeng Zuo, and Lei Zhang. 2015. Bit-Scalable Deep Hashing With Regularized Similarity Learning for Image Retrieval and Person Re-Identification. *IEEE Transactions on Image Processing* 24, 12 (2015), 4766–4779.