

One Loss for Quantization: Deep Hashing with Discrete Wasserstein Distributional Matching

Khoa D. Doan, Peng Yang, Ping Li
Cognitive Computing Lab
Baidu Research

10900 NE 8th St. Bellevue, WA 98004, USA

{khoodoan106, pengyang01, pingli98}@gmail.com

Abstract

Image hashing is a principled approximate nearest neighbor approach to find similar items to a query in a large collection of images. Hashing aims to learn a binary-output function that maps an image to a binary vector. For optimal retrieval performance, producing balanced hash codes with low-quantization error to bridge the gap between the learning stage’s continuous relaxation and the inference stage’s discrete quantization is important. However, in the existing deep supervised hashing methods, coding balance and low-quantization error are difficult to achieve and involve several losses. We argue that this is because the existing quantization approaches in these methods are heuristically constructed and not effective to achieve these objectives. This paper considers an alternative approach to learning the quantization constraints. The task of learning balanced codes with low quantization error is re-formulated as matching the learned distribution of the continuous codes to a pre-defined discrete, uniform distribution. This is equivalent to minimizing the distance between two distributions. We then propose a computationally efficient distributional distance by leveraging the discrete property of the hash functions. This distributional distance is a valid distance and enjoys lower time and sample complexities. The proposed single-loss quantization objective can be integrated into any existing supervised hashing method to improve code balance and quantization error. Experiments confirm that the proposed approach substantially improves the performance of several representative hashing methods.

1. Introduction

An important challenge associated with massive image datasets is to efficiently and effectively search for images containing *semantically similar* content in these datasets. Hashing is a principled approximate nearest neighbor search approach with applications in many domains, ranging from text or image retrieval [22, 47] to spam

or duplicate-scene detection [8, 38]. Hashing approaches learn binary encoding of the original images so that the “candidate” subset of images can be efficiently discovered from the binary-coding space. Binary codes are efficient to store and the high cost of pairwise distance calculations in the high-dimensional space is reduced to the significantly lower cost of discrete Hamming distance calculations. A Hamming-distance calculation only requires a bit-wise XOR and a bit-count operation, which can be efficiently computed in most of conventional systems.

To ensure that the retrieved images are relevant, hashing methods learn hash functions that preserve the pairwise similarity of the images in the discrete space. Supervised hashing methods additionally leverage the annotated similarity to learn the hash functions and achieve superior retrieval performance compared to unsupervised hashing methods [4, 5, 20, 40, 48, 50, 54, 57, 60]. Since discrete optimization is intractable, these methods solve a relaxed problem that replaces the discrete constraint with a continuous output. The continuous output is “quantized” to obtain the binary during inference. Such a relaxation results in a discrepancy between the discrete and continuous optimizations that must be compensated for in the learning process. Two important criteria to consider are quantization error and coding balance [22, 52, 53]. Quantization error is the information loss when the discrete function is represented by a continuous function. Quantization error penalizes the cases of assigning “very” similar data points to binary codes with large distances [53]. Coding balance, on the other hand, encourages a uniform distribution of the images into the binary codes, which helps reduce the time complexity of the retrieval operations in the worst and average scenarios [24].

Existing supervised hashing methods, especially those that are based on neural networks, include one or more penalty terms, besides the similarity-preserving loss, to force the continuous output as discrete as possible. However, these relaxation schemes still introduce non-trivial

quantization error and coding unbalanced, which eventually leads to sub-optimal hash codes [52, 53]. Multiple penalty terms also lead to longer model trainings due to the time-consuming hyperparameter-tuning step.

This paper proposes a faster and more performant quantization approach for the deep supervised hashing methods. First, we empirically show that low-quantization error and balance coding induce a uniform discrete distribution. The ultimate goal of hash-function learning is to project data into this uniform discrete distribution. Thus, we formulate the quantization objective as minimizing the distributional distance between the relaxed, continuous hash distribution and this uniform discrete distribution. The proposed formulation has two advantages: (i) achieving low-quantization error and coding balance is easier with this formulation and (ii) low-quantization error and coding balance are simultaneously optimized in a unified formulation, thus reducing the number of hyperparameters to tune. Our **main contributions** are as follows:

- We achieve a low quantization error and coding balance by minimizing the single-loss distributional distance between the learned hash distribution and the uniform discrete distribution. This new quantization objective can be used in conjunction with any existing deep supervised hashing methods to further improve their retrieval performances and reduce their time-consuming hyperparameter tuning processes.
- We propose a low computation- and sample-complexity Sliced-Wasserstein-based distributional distance. The proposed distance, called HSWD, is theoretically a valid distance with better computational efficiency than other Wasserstein-distance types, including the original Sliced Wasserstein Distance.
- We demonstrate the efficiency and effectiveness of the proposed quantization technique in several well-known deep supervised hashing methods on various widely used real-world datasets using both quantitative and qualitative performance analysis.

Paper Organization: The rest of the paper is organized as follows. We review the related work in Section 2. In Section 3, we present the empirical analysis of the quantization error and coding balance, and the details of the proposed methodology. We evaluate the effectiveness of the proposed quantization approach in Section 4. Finally, Section 5 presents remarks and concludes this paper. We present more details about experimental settings and results as well as supporting proofs in the supplementary material.

2. Related Work

In this section, we first discuss the prior works in deep supervised hashing and the quantization constraints. Then

we discuss the Wasserstein distances, which are related to the computational approach proposed in the paper.

2.1. Image Hashing

Hashing has been intensively investigated in both theory and practice, with applications ranging from retrieval [47] to compressed sensing [34] and feature learning [41]. In the image retrieval domain, existing data-dependent hashing methods can be organized into two categories: shallow hashing and deep hashing. Shallow hashing methods rely on hand-crafted features and learn linear hash functions and feature extraction techniques [22, 33, 44, 53]. Along with the rapid development of deep neural networks, deep hashing methods [6, 43, 55, 56] combine representation learning and hash-function learning into an end-to-end model and have demonstrated significant performance improvements over the hand-crafted shallow hashing approaches. Hashing methods can also be broadly classified into (data dependent) unsupervised [12, 13, 15, 22, 25, 26, 28, 29, 42, 47, 53, 56] and supervised methods [4–6, 20, 40, 48, 54, 57, 58]. Supervised hashing methods demonstrate superior performance over the unsupervised ones by utilizing the labeled similarity information in the training data to learn the hash functions. Note that there is also a popular line of hashing works on “data independent” (unsupervised) hashing with quantizations, including (i) quantized random projections [7, 9, 21, 37, 39, 51]; (ii) quantized stable random projections [34]; (iii) quantized random Fourier features [41, 59]; (iv) b-bit minwise hashing [3, 36] and b-bit consistent weighted sampling [30, 35, 45], etc.

Since learning the discrete hash function is computationally intractable, existing deep hashing methods approximate the discrete hash function with a continuous-output one, by replacing the discrete-output constraint with a *tanh* activation [6, 40, 58, 60]. To learn the hash functions, these methods optimize the similarity-preserving loss and quantization loss, simultaneously. The similarity-preserving loss tries to preserve the input-space similarity structure in the discrete space while the quantization loss minimizes the discrepancy between the discrete and continuous optimizations. It has been shown that an effective quantization approach is crucial in improving the retrieval performance [18, 37, 53]. However, such quantization objective is often ignored or not effectively learned via multiple losses (> 3) in the existing deep supervised hashing methods. This paper studies the effectiveness of quantization objectives in these approaches and proposes an alternative, more effective single-loss quantization objective. Note that, we consider our work to be orthogonal to OrthoHash [27] although OrthoHash also aims to reduce the number of losses (but only in the point-wise supervised image hashing). Our approach does not require Batch Normalization for coding balance as in OrthoHash, can be used in the point-wise or pair-wise

setting, and can be incorporated into any existing hashing methods to improve its quantization effectiveness.

2.2. Wasserstein Distances

Wasserstein distances have been extensively studied in the literature and especially in recent applications such as GANs [2, 17, 23]. One advantage of the Wasserstein distances over other types of distributional distances, such as Kullback–Leibler (KL) or Jensen-Shannon (JS) divergence, is that they consider the geometry of the data and are well-defined even for distributions without overlapping supports [1]. However, estimating the Wasserstein distance is a challenging task. Computing the Wasserstein-2 distance from its primal domain is computationally intractable [2]. Computing the Wasserstein-2 distance via the Optimal-Transport (OT) formulation has also been explored [16, 18], but the OT’s high computational cost is prohibitive. Employing the Kantorovich-Rubinstein dual requires estimating a Lipschitz continuous function to calculate the distance. Parameterizing the Lipschitz function with a neural network is challenging [10]. Furthermore, the dual approach results in a minimax optimization, which requires non-trivial modifications to learning algorithms such as those in hashing. Note that, estimating the KL or JS divergence also requires minimax optimization [19].

Wasserstein-2 distance (also KL and JS) requires an exponential number of samples to reliably estimate the distance [11]. In contrast, a variant of the Wasserstein distance, called Sliced Wasserstein Distance (SWD), approximates the Wasserstein distance by averaging the one-dimensional Wasserstein distances of the data points when they are projected onto many random, one-dimensional directions [31]. SWD has a polynomial sample complexity [10]. The SWD estimation has a computational cost of $O(LN \log(Nd))$, where L is the number of random directions, N is the number of samples, and d is the dimension of the data. Nevertheless, in the high dimensional space, it becomes very likely that many random, one-dimensional directions do not lie on the manifold of the data. In other words, in several of these directions, the projected distances are close to zero. Consequently, in practice, the number of random directions L is often large. For example, in [11], for a mini-batch size of 64, SWD needs $L = 10,000$ projections. To address this problem, some works find the best directions and estimate the Wasserstein distance using these directions [10, 14, 15].

3. Distributional Matching for Hash Function Learning

3.1. Hash Function Learning

We focus on the general supervised setting of learning a discrete hash function with a neural network. In this setting, a training dataset consists of N images, $X = \{x\}_N$,

and a corresponding labeling function Y . Y can be point-wise (i.e., assigning labels to each image) or pair-wise (i.e., assigning a similarity score to a pair of images) and is pre-defined. The labeling function essentially describes the semantic similarity of the input. The goal of deep supervised hashing is to learn a discrete hash function $\mathbf{H} : x \rightarrow c$ that preserves the semantic similarity of the input. Without loss of generality, we consider the point-wise labeling case and the θ is the parameter of the hash function. The optimization objective can be defined as:

$$\min_{\theta} L_s(\mathbf{H}(x), Y(x)), \quad s.t. \mathbf{H}(x) \in \{-1, 1\}^m \quad (1)$$

where m is the number of bits. In this objective function, L_s preserves the semantic similarity of the input, defined via Y . In practice, such a discrete optimization, especially when \mathbf{H} is a neural network, is intractable. To this end, the discrete \mathbf{H} is relaxed with a continuous function $h : x \rightarrow [-1, 1]^m$ with a regularization. The general, continuously-relaxed objective function can be written as,

$$\min_{\theta} L_s(h(x), Y(x)) + \lambda L_q(h(x)) \quad (2)$$

where the $h(x)$ is usually modeled with a \tanh activation function at the output layer. The term L_q denotes the quantization objective, which minimizes the gap between the discrete and continuous solutions.

3.2. Quantization and Code-balance

The quantization objective is very important for learning a hash function with a good retrieval performance [52, 53]. This objective typically includes three constraints: quantization error, bit balance, and bit uncorrelation. Low-quantization error alleviates the loss that occurs in assigning “very” similar data points to the codes with large Hamming distances. Bit balance ensures that each bit has the same chance of being -1 or 1 , while bit uncorrelation encourages each bit to represent the orthogonal feature of the data. Together, bit-balance and bit-uncorrelation essentially characterize “code-balance”, i.e., the condition where training data points are uniformly assigned to each hash code. Code balance is important because it helps minimize the time complexity of both the worst and average cases in retrieval and improve the retrieval quality [24].

Consider Figure 1, where the data, from four classes, is projected on a two-dimensional space (Figure 1(a)). The goal is to learn a two-bit hash function. In Figure 1(c), since the bits are correlated, it is not possible to avoid the collision of data points from different classes in the discrete, hashing space. When the hash function has a high quantization error (Figure 1(d)), some similar data points near the boundary are likely to be assigned to two different hash codes, which results in a higher false-negative rate in retrieval. The optimal hash function, as seen in Figure 1(b)




Figure 1. A visual illustration of the optimal function (Figure (b)) and learned hash function with poor the code balance (Figure (c)) and quantization constraint (Figure (d)). Figure (a) shows the original data clusters from four classes.

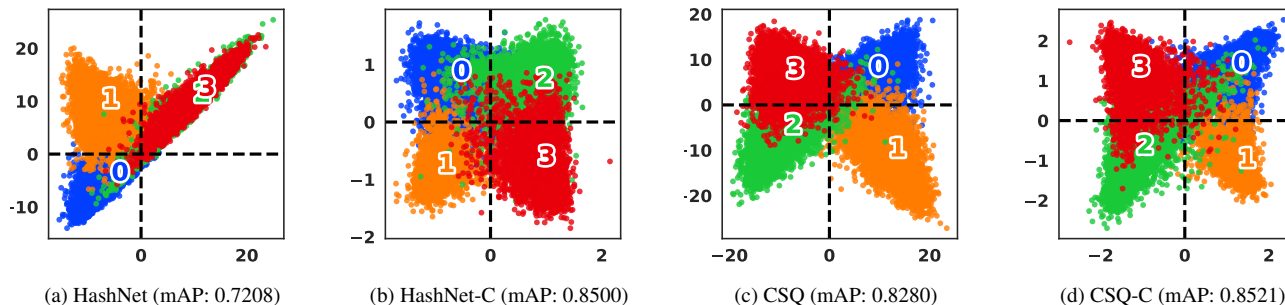


Figure 2. Visualization of the learned 2-bit hashing space (i.e. learn a 2-bit hash function) using the HashNet and DSDH methods. HashNet-C and DSDH-C replace the quantization in HashNet and DSDH, respectively, with the proposed HSWD quantization approach.

pulls the data points toward the corners (low quantization error) and equally divided the data into four quadrants (code balance) while preserving the semantic similarity of the data (data from different classes occupy different quadrants).

While the quantization objective is important, existing deep supervised approaches do not effectively learn this objective. Consider the case of learning a two-bit hash function on CIFAR-10 using data from four classes. Figures 2(a) and 2(c) show the quantization results of two representative supervised hashing methods, HashNet [6] and CSQ [58], respectively. HashNet suffers from high quantization error (unbounded continuous output) and bit correlation (several samples from class 1 and class 3 occupy the hash code (1, 1)). CSQ is slightly better than HashNet but cannot balance well between preserving semantic similarity and the quantization objective. Our proposed quantization constraint in HashNet-C and CSQ-C, which revises the corresponding quantization objective in HashNet and CSQ, respectively, learns hash codes with better balance and lower quantization error. This leads to improvements in the retrieval performance (from 0.7208 and 0.8280 to 0.8500 and 0.8521 for HashNet and CSQ, respectively).

In the next section, we will describe the proposed quantization objective, i.e., a single distributional-distance minimization between the learned continuous hash codes and a fixed, uniform discrete distribution. This approach can

be easily adapted to existing methods and facilitates these methods to learn better hash functions. Our approach reduces the number of quantization hyperparameters to one, significantly reducing the training process.

3.3. Distributional Quantization Distance

As discussed, following the uniform discrete distribution (Figure 1(b)) results in better coding balance and low quantization error. Under this distribution (denoted as B), a sample $b \in \{-1, 1\}^m$ can be drawn as follows: for each dimension, independently and randomly samples a value of -1 or 1 with equal probability. We can see that the samples from B exhibit coding balance: each bit is balanced because its value has an equal chance of being -1 or 1 , and the bits are uncorrelated since each bit is sampled independently. We propose to minimize the discrepancy between the learned hash distribution and B to improve the coding balance in the learned distribution. This approach also reduces the quantization error since the learned continuous codes will gradually become binary. Formally, we minimize the following quantization objective:

$$L_q(h(X)) = \mathcal{D}(h(X)||B) \quad (3)$$

where \mathcal{D} denotes the distributional distance function. Note that the constraint is applied on the output space of h (thus, $h(X)$), instead on an individual sample ($h(x)$). Nonetheless,

less, minimizing \mathcal{D} is a non-trivial task, especially when the hashing space is high dimensional and the distribution density of the output space that we are learning cannot be estimated. Furthermore, for a choice of a distribution distance, its estimation should be efficient to compute to be used in practical applications.

Some well-known divergences, such as Kullback–Leibler (KL) or Jensen-Shannon (JS) divergence, generally require the ability to estimate the learned output distribution of $h(x)$. While we can use an additional network (i.e., similar to the discriminators in GANs [23]) to estimate KL (or JS), learning the neural network with the minimax optimization is computationally expensive. KL and JS are also not well-defined for distributions with non-overlapping supports [1]. To remedy this issue, we first consider the Wasserstein-2 distance. To improve the computation efficiency of the Wasserstein-2 distance, we propose a Sliced-based Wasserstein-2 distance that is more computationally efficient and effective to estimate. The quantization objective \mathcal{D} using the Wasserstein-2 distance can be formulated as follows:

$$\mathcal{D}(\mu, \nu) = \left(\inf_{\gamma \in \Pi(\mu, \nu)} \int_{(z, b) \sim \gamma} p(z, b) \|z - b\|_2 dz db \right)^{1/2} \quad (4)$$

where $\Pi(\mu, \nu)$ is the set of all transportation plans γ such that their marginal distributions are μ and ν . μ and ν define the output distribution of the continuous hash function h and the distribution B , respectively. Computing the infimum in Equation (4) is difficult since the distribution induced by $z = h(x)$ is not fixed or unknown while employing the Kantorovich-Rubinstein duality to optimize Equation (4) requires modeling a Lipschitz function using an additional neural network [2]. Fortunately, Wasserstein distance has an elegant yet closed-form solution for one-dimensional continuous measures. Denoting q_μ and q_ν as the density functions of μ and ν , respectively, the Wasserstein-2 distance between one-dimensional measures μ and ν is given by:

$$\mathcal{W}(\mu, \nu) = \left(\int_0^1 \|F_\mu^{-1}(w) - F_\nu^{-1}(w)\|_2 dw \right)^{1/2} \quad (5)$$

where $F_\mu(w) = \int_{-\infty}^w q_\mu(\rho) d\rho$ and $F_\nu(w) = \int_{-\infty}^w q_\nu(\rho) d\rho$ are the cumulative distribution functions. Inspired by the efficiency of estimating the one-dimensional Wasserstein-2 distance, we propose to find a family of one-dimensional representations, e.g., through the linear projections, and approximate the Wasserstein-2 distance as a function of these one-dimensional marginals:

$$\mathcal{D}(h(X), B) \approx \left(\frac{1}{L} \sum_{l=1}^L \mathcal{W}(\omega_l^T h(X), \omega_l^T B) \right)^{1/2} \quad (6)$$

where $\omega_l^T h(X)$ and $\omega_l^T B$ are the projections of the output samples of h and samples from B onto a one-dimensional direction defined by ω_l (a slice). Typically, ω_l is drawn from a uniform distribution on the unit sphere. This is also known as the Sliced-Wasserstein distance (SWD) [11, 31]. While this approach has successful applications in a variety of tasks [11, 31], the random nature of the slices could lead to several non-informative directions where the sliced distances along these directions are close to 0. Consequently, a large number L of random directions is needed to approximate the sliced-Wasserstein distance, which increases the computational complexity of the estimation.

To remedy this issue, we avoid the random projections and select directions that contain discriminant information of the two data sources. We utilize the observation that the objective L_s aims to learn discriminative projections, via h , of the data points onto the discrete space. Intuitively, since each output dimension of the hash function describes a discriminative feature of the space, projecting the data into this dimension can capture meaningful separation of the two distributions. Empirically, **how to define HSWD?** we can show that the averaged one-dimensional Wasserstein distance along these dimensions, denoted as HSWD, captures better separations of the data than SWD. Specifically, we provide the estimated HSWD and SWD (using 10,000 projections in CIFAR-10) when the model is trained to minimize SWD (Figures 3(a) and 3(c)), or HSWD (Figures 3(b) and 3(d)) as the quantization objective. As we can observe, HSWD describes better separations of the two distributions (i.e., higher distances). Minimizing SWD does not effectively bring the two distributions closer in some cases (e.g., HSWD increases in Figure 3(a), indicating the

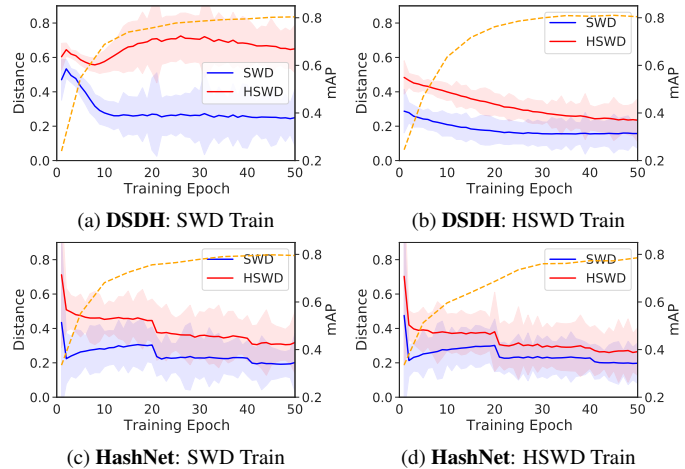


Figure 3. Estimated SWD and HSWD during training the 32-bit hash functions on CIFAR-10 dataset. For each method (DSDH [40] or HashNet [6]), SWD or HSWD is used as the quantization objective. The dashed (orange) line denotes the retrieval performance (mAP).

separations even increase in some directions). However, minimizing HSWD effectively forces $h(X)$ to be closer to B , as both HSWD and SWD decrease in Figures 3(b) and 3(d). More importantly, the number of projections, which is also the number of dimensions of $h(X)$, is fixed and generally small ($m \leq 128$ in most applications). The one-dimensional distances along these dimensions are exactly the distances along the coordinate axis of the discrete space. Thus, HSWD is defined, as follows:

$$\mathcal{D}(h(X), B) \approx \left(\frac{1}{m} \sum_{l=1}^m [\mathcal{W}(h(X)_{l,:}, B_{l,:})]^2 \right)^{1/2} \quad (7)$$

where $h(X)_{l,:}$ and $B_{l,:}$ are the one-dimensional samples along the dimension l of $h(X)$ and B , respectively. Note that, using the discrete dimensions as projected directions is similar to the works in [15] for the unsupervised study. We can show that HSWD is a valid distance.

Theorem 1. *The proposed distance Sliced Wasserstein calculation in Theorem 1, denoted as HSWD, is a valid distance function of probability measures in this space.*

Remark 1. *The proposed HSWD is similar to the Max Sliced Wasserstein Distance (M-SWD), which has successful applications [10, 31]. However, different from M-SWD, HSWD does not require a separate discriminative network to estimate the distance.*

HSWD enjoys a better computational efficiency than SWD ($O(mN \log(Nd))$ compared to $O(LN \log(Nd))$). In most problems, SWD requires a large number ($L \gg N$) of random directions, typically between 1000 to 10,000, to provide a reliable estimate of the distance [10, 46]. In HSWD, the number of directions is fixed to the dimension of the hashing space m , which is typically between 16 to 128 for many image hashing applications.

4. Experiments

We present the experimental results to demonstrate the effectiveness of the proposed quantization method over the existing quantization approaches using extensive combinations of representative datasets, deep supervised hashing methods, and hash-function architectures.

4.1. Datasets Used

We utilize the following datasets:

- **CIFAR10:** This dataset consists of 60,000 natural images categorized uniformly into 10 labels. We randomly select 100 and 500 images from each label for the query set and the training set, respectively. The remaining images form the retrieval set. Hence, there are 1000, 5,000, and 54,000 images in the query, training, and retrieval sets, respectively.

- **NUS-WIDE:** This dataset contains 269,648 images, each of which belongs to at least one of 21 concepts. We randomly select 5,000 images for the query set, with the remaining images used as the retrieval set. 10,000 images randomly selected in the retrieval set are used for training.
- **COCO:** This dataset contains 123,287 images, labeled with at least 1 out of 80 semantic concepts. Similarly, the query set is randomly constructed with 5,000 images, with the remaining images used as the retrieval set (10,000 randomly selected images in this set are used for training).

4.2. Evaluation Metrics

For evaluating the performance of the proposed model, we follow the standard evaluation mechanism that is widely accepted for the problem of image hashing - the **precision@R (P@R)** and **mean average precision (mAP)**.

4.3. Comparison Methods

We compare the performance of the proposed quantization approach in various representative deep supervised image hashing methods: **DSDH** [40], **HashNet** [6], **Greedy-Hash** [50], **DCH** [5], **CSQ** [58], and **DBDH** [60]. Note that the performance comparisons between the methods are beyond the scope of this paper and we only focus on the performance improvement when the proposed quantization approach is incorporated in these methods.

For each method, report the retrieval performance using the original quantization proposal and our quantization approach using SWD (using the suffix -S) and HSWD (using the suffix -C). For example, for the CSQ method, we report CSQ (original algorithm), CSQ-S (CSQ with SWD quantization), and CSQ-C (CSQ with HSWD quantization).

4.4. Implementation Details

For a fair comparison, we use the same underlying deep hash function (e.g., VGG11) for the original algorithm and those with the proposed quantization approaches. We perform a hyperparameter selection step for each method and report the average performances for the best configuration across multiple runs with different random initializations.

4.5. Retrieval Performance

In this section, we measure the performance of the proposed quantization approach and the original quantization proposal for each selected deep supervised hashing method. Table 1 shows the mAP results in our experiments for learning the 16-bit, 32-bit, and 64-bit hash functions with the VGG11 backbone [49]. We also report the P@1000 results for learning the 16-bit and 32-bit hash functions with the AlexNet backbone [32] in Table 2. For each of the proposed

Table 1. mAP for different numbers of bits on the three image datasets. The blue value (in bold) along the mAP value of each proposed approach shows the relative improvement over the original algorithm, while the italicized value indicates no improvement.

| Method | CIFAR-10 | | | NUS-WIDE | | | COCO | | |
|-----------------|---------------------|---------------------|---------------------|---------------------|---------------------|---------------------|---------------------|---------------------|---------------------|
| | 16 bits | 32 bits | 64 bits | 16 bits | 32 bits | 64 bits | 16 bits | 32 bits | 64 bits |
| DSDH [40] | 0.7909 | 0.8072 | 0.8278 | 0.8270 | 0.8455 | 0.8640 | 0.7331 | 0.7853 | 0.8074 |
| DSDH-S | 0.8187/ 3.5% | 0.8439/ 4.6% | 0.8517/ 2.9% | 0.8282/ 0.1% | 0.8461/ 0.1% | 0.8712/ 0.8% | 0.7330/ <i>0.0%</i> | 0.8030/ 2.3% | 0.8404/ 4.1% |
| DSDH-C | 0.8531/ 7.9% | 0.8620/ 6.8% | 0.8658/ 4.6% | 0.8433/ 2.0% | 0.8631/ 2.1% | 0.8749/ 1.3% | 0.7424/ 1.3% | 0.8032/ 2.3% | 0.8408/ 4.1% |
| HashNet [6] | 0.6922 | 0.8311 | 0.8566 | 0.7728 | 0.8336 | 0.8654 | 0.6899 | 0.7666 | 0.8098 |
| HashNet-S | 0.8131/ 17% | 0.8573/ 3.2% | 0.8749/ 2.1% | 0.8062/ 4.3% | 0.8438/ 1.2% | 0.8713/ 0.7% | 0.7215/ 4.6% | 0.7764/ 1.3% | 0.8189/ 1.1% |
| HashNet-C | 0.7939/ 14% | 0.8467/ 1.9% | 0.8691/ 1.5% | 0.8002/ 3.5% | 0.8437/ 1.2% | 0.8791/ 1.6% | 0.7202/ 4.4% | 0.7789/ 1.6% | 0.8202/ 1.3% |
| GreedyHash [50] | 0.8223 | 0.8474 | 0.8646 | 0.7802 | 0.8081 | 0.8328 | 0.6533 | 0.7219 | 0.7561 |
| GreedyHash-S | 0.8280/ 0.7% | 0.8497/ 0.3% | 0.8653/ 0.1% | 0.7815/ 0.1% | 0.8083/ <i>0.0%</i> | 0.8390/ 0.7% | 0.6668/ 2.1% | 0.7291/ 1.0% | 0.7618/ 0.8% |
| GreedyHash-C | 0.8375/ 1.9% | 0.8536/ 0.7% | 0.8722/ 0.9% | 0.7890/ 1.1% | 0.8179/ 1.2% | 0.8477/ 1.8% | 0.6637/ 1.6% | 0.7299/ 1.1% | 0.7712/ 2.0% |
| DCH [5] | 0.8302 | 0.8432 | 0.8558 | 0.8015 | 0.8061 | 0.8040 | 0.7578 | 0.7792 | 0.7723 |
| DCH-S | 0.8372/ 0.8% | 0.8515/ 1.0% | 0.8602/ 0.5% | 0.8058/ 0.5% | 0.8079/ 0.2% | 0.8067/ 0.3% | 0.7657/ 1.1% | 0.7831/ 0.5% | 0.7803/ 1.0% |
| DCH-C | 0.8446/ 1.7% | 0.8596/ 1.9% | 0.8711/ 1.8% | 0.8159/ 1.8% | 0.8145/ 1.0% | 0.8155/ 1.4% | 0.7702/ 1.6% | 0.7892/ 1.3% | 0.7807/ 1.1% |
| CSQ [58] | 0.8069 | 0.8291 | 0.8366 | 0.7992 | 0.8384 | 0.8596 | 0.6783 | 0.7550 | 0.8146 |
| CSQ-S | 0.8401/ 4.1% | 0.8555/ 3.2% | 0.8554/ 2.3% | 0.8044/ 0.7% | 0.8495/ 1.3% | 0.8626/ 0.4% | 0.7036/ 3.7% | 0.7765/ 2.8% | 0.8234/ 1.0% |
| CSQ-C | 0.8457/ 4.8% | 0.8558/ 3.2% | 0.8652/ 3.4% | 0.8054/ 0.8% | 0.8511/ 1.5% | 0.8701/ 1.2% | 0.6989/ 3.0% | 0.7752/ 2.7% | 0.8255/ 1.3% |
| DBDH [60] | 0.7660 | 0.8223 | 0.8492 | 0.8305 | 0.8552 | 0.8666 | 0.7202 | 0.7826 | 0.8042 |
| DBDH-S | 0.8458/ 10% | 0.8587/ 4.4% | 0.8603/ 1.3% | 0.8387/ 1.0% | 0.8577/ 0.3% | 0.8680/ 1.8% | 0.7461/ 2.2% | 0.7996/ 3.7% | 0.8336/ 4.3% |
| DBDH-C | 0.8466/ 10% | 0.8593/ 4.5% | 0.8668/ 2.1% | 0.8395/ 1.1% | 0.8633/ 0.9% | 0.8760/ 1.1% | 0.7389/ 2.6% | 0.7889/ 0.8% | 0.8308/ 3.9% |

Table 2. P@1000 for 16 and 32 bits on CIFAR-10 and NUS-WIDE. The blue value (in bold) along the P@1000 value of each proposed approach shows the relative improvement over the original algorithm, while the italicized value indicates no improvement.

| Method | CIFAR-10 | | NUS-WIDE | |
|--------------|----------------------|---------------------|----------------------|----------------------|
| | 16 bits | 32 bits | 16 bits | 32 bits |
| DSDH | 0.8252 | 0.8406 | 0.8117 | 0.8294 |
| DSDH-S | 0.8526/ 3.3% | 0.8543/ 1.6% | 0.8162/ 0.6% | 0.8312/ 0.2% |
| DSDH-C | 0.8645/ 4.8% | 0.8739/ 4.0% | 0.8195/ 1.0% | 0.8391/ 1.2% |
| HashNet | 0.6193 | 0.8613 | 0.7581 | 0.8158 |
| HashNet-S | 0.8470/ 36.8% | 0.8755/ 1.7% | 0.7743/ 2.1% | 0.8199/ 0.5% |
| HashNet-C | 0.7698/ 24.3% | 0.8715/ 1.2% | 0.7456/ <i>-1.7%</i> | 0.8078/ <i>-1.0%</i> |
| GreedyHash | 0.8561 | 0.8616 | 0.7601 | 0.8009 |
| GreedyHash-S | 0.8583/ 0.3% | 0.8656/ 0.5% | 0.7657/ 0.7% | 0.7973/ <i>-0.5%</i> |
| GreedyHash-C | 0.8517/ <i>-0.5%</i> | 0.8700/ 1.0% | 0.7630/ 0.4% | 0.7931/ <i>-1.0%</i> |
| DCH | 0.8621 | 0.8568 | 0.7843 | 0.7898 |
| DCH-S | 0.8622/ <i>0.0%</i> | 0.8761/ 2.3% | 0.7846/ <i>0.0%</i> | 0.7923/ 0.3% |
| DCH-C | 0.8654/ 0.4% | 0.8635/ 0.8% | 0.7893/ 0.6% | 0.7914/ 0.2% |
| CSQ | 0.8510 | 0.8571 | 0.7903 | 0.8285 |
| CSQ-S | 0.8661/ 1.8% | 0.8732/ 1.9% | 0.8034/ 1.7% | 0.8318/ 0.4% |
| CSQ-C | 0.8670/ 1.9% | 0.8688/ 1.4% | 0.8007/ 1.3% | 0.8353/ 0.8% |
| DBDH | 0.8440 | 0.8421 | 0.8122 | 0.8323 |
| DBDH-S | 0.8626/ 2.2% | 0.8675/ 3.0% | 0.8177/ 0.7% | 0.8388/ 0.8% |
| DBDH-C | 0.8658/ 2.6% | 0.8731/ 3.7% | 0.8135/ 0.1% | 0.8380/ 0.7% |

quantization approaches, we report the relative improvement besides the mAP or P@1000 values. We highlight the statistically-significant relative improvements in blue (and in bold values), while the values are italic when there are not statistically-significant improvements.

As we can observe in Tables 1 and 2, the proposed quantization approaches achieve significant improvement in the retrieval results across different supervised hashing methods. Specifically, the SWD variants result in up to over 10% improvement in CIFAR-10, 4% in NUS-WIDE, and almost 5% in COCO. Similarly, the HSWD variants also

consistently improve the retrieval results in these supervised hashing methods. As mentioned previously, the primary advantage of HSWD over SWD is its computational efficiency. We conjecture that the superior performance of the proposed approach is due to the following reasons:

- The quantization constraints in the existing deep supervised hashing methods are less effective to minimize the gaps between the relaxed hash function $h(x)$ and the optimal, discrete hash function $H(x)$. The observation in Figure 2 also confirms our discussion.
- The proposed quantization approaches can help learn better-quantized hash functions. Specifically, it is more effective to optimize the combination of the similarity preserving loss L_s and L_q when L_q is one of our proposed approaches.
- Finally, low-quantization error and code balance are crucial in improving the retrieval performance in the existing deep supervised hashing methods. Additional empirical results for this relationship are presented in Section 4.7.2.

4.6. Computational Efficiency

We compare the training time per epoch of the proposed quantization approaches when learning 64-bit hash functions. For each method, we capture the average running time (in seconds) of the original algorithm and the algorithm with SWD- and HSWD-quantizations. For SWD, we report the corresponding running time of the number of projections that results in the optimal performance. We observe that the optimal number of projections ranges from 1000 to 10,000. Note that for HSWD, the number of projections is fixed and equals the size of the hash codes.

Table 3. Averaged running time per epoch across different supervised hashing methods (in seconds). The blue values are the relative running-time decreases.

| Dataset | Original | SWD | HSWD |
|----------|----------|------|------------------|
| CIFAR-10 | 19.4 | 24.2 | 17.1/ 40% |
| NUS-WIDE | 58.3 | 71.2 | 50.1/ 41% |
| COCO | 55.6 | 68.1 | 49.5/ 37% |

Table 3 reports the average running times across different deep hashing methods. We can observe that HSWD is more computationally efficient than SWD because of its fewer projections and the omission of the matrix-multiplication operation that projects the data points into random directions. Compared to the original quantizations, HSWD quantization is also faster since each original quantization (in L_q) usually comprises of several losses.

4.7. Qualitative Analysis

4.7.1 Hash Code Visualization

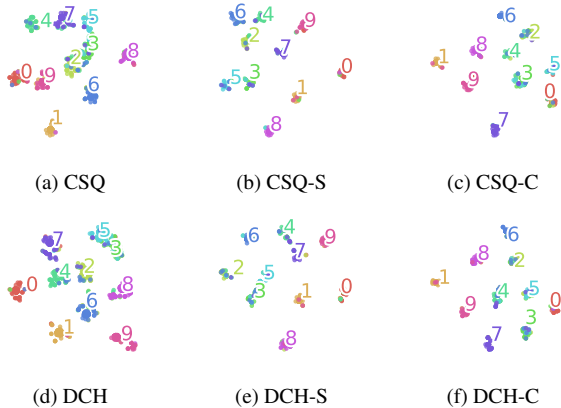


Figure 4. Two-dimensional t-SNE visualizations of the quantized 16-bit hash codes learned by different quantization approaches in CSQ and DCH on CIFAR-10.

We present the visualization of the hash codes generated by the learned hash functions for the different quantization approaches, using the CSQ and DCH methods, in Figure 4. We project the hash codes into 2-dimensional embeddings using the t-SNE method. As Figure 4 displays, the learned hash codes of the methods with SWD and HSWD quantizations exhibit better inter-cluster separation and intra-cluster closeness, which makes their retrieval performance superior. For example, for the CSQ’s result in Figure 4(a), we can observe that the samples from classes 2 and 3 are closer to each other. For CSQ-S and CSQ-C, the samples from these two classes are more separated.

4.7.2 Analysis of Effective Quantization and Performance Improvement

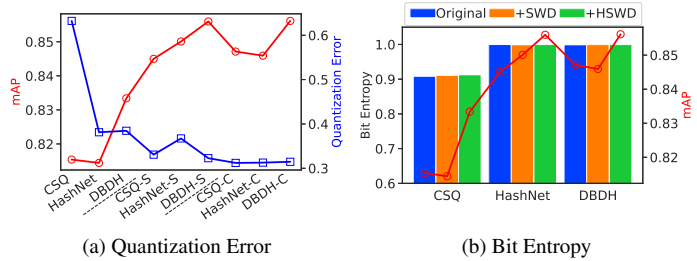


Figure 5. Quantization Error (angle, in radian, between continuous codes and corresponding hash codes) and Bit Entropy (CIFAR10).

The proposed method replaces several quantization losses with a single loss (SWD/HSWD), allowing efficient and effective quantization in deep supervised hashing methods. In Figure 5, we show that the proposed quantization approaches simultaneously achieve low-quantization and balanced bits. As we can observe in this figure, when the quantization error is reduced (as in the proposed approaches), the retrieval performance increases. Furthermore, the proposed approaches also consistently achieve highly balanced bits (i.e., high bit entropies). Higher balanced bits also lead to better performance.

5. Conclusions

In this paper, we proposed a novel single-loss quantization objective for the deep supervised image hashing problem. Different from the existing quantization approaches, our model selects an optimal uniform discrete distribution and directly minimizes the distribution distance between the output distribution of the hash function and this uniform distribution. We considered the Sliced Wasserstein Distance as the choice of the distributional distance for its low-sample complexity and easier estimation. However, the Sliced Wasserstein Distance can require projections into several random directions that do not contain useful information about the separation of the data. By studying the properties of the distributions under consideration, we proposed a variant of SWD, called HSWD, and showed that it only requires a small number of informative directions. HSWD can achieve significant computational gains than SWD. Our experiments validate that the proposed quantization approach can improve the performance of several representative deep supervised-hashing methods on several datasets. Our work makes one leap towards leveraging an efficient, robust quantization approach for deep supervised hashing and we envision that our model will serve as a motivation for improving other related hashing applications.

References

- [1] Martín Arjovsky and Léon Bottou. Towards principled methods for training generative adversarial networks. In *Proceedings of the 5th International Conference on Learning Representations (ICLR)*, Toulon, France, 2017. 3, 5
- [2] Martín Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, pages 214–223, Sydney, Australia, 2017. 3, 5
- [3] Andrei Z. Broder, Steven C. Glassman, Mark S. Manasse, and Geoffrey Zweig. Syntactic clustering of the web. In *Proceedings of the World Wide Web (WWW)*, pages 1157 – 1166, Santa Clara, CA, 1997. 2
- [4] Yue Cao, Bin Liu, Mingsheng Long, and Jianmin Wang. Hashgan: Deep learning to hash with pair conditional wasserstein gan. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1287–1296, Salt Lake City, UT, 2018. 1, 2
- [5] Yue Cao, Mingsheng Long, Bin Liu, and Jianmin Wang. Deep cauchy hashing for hamming space retrieval. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1229–1237, Salt Lake City, UT, 2018. 1, 2, 6, 7
- [6] Zhangjie Cao, Mingsheng Long, Jianmin Wang, and Philip S Yu. Hashnet: Deep learning to hash by continuation. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 5608–5617, Venice, Italy, 2017. 2, 4, 5, 6, 7
- [7] Moses S. Charikar. Similarity estimation techniques from rounding algorithms. In *Proceedings on 34th Annual ACM Symposium on Theory of Computing (STOC)*, pages 380–388, Montreal, Canada, 2002. 2
- [8] Ondrej Chum, James Philbin, and Andrew Zisserman. Near duplicate image detection: min-hash and tf-idf weighting. In *Proceedings of the British Machine Vision Conference (BMVC)*, pages 1–10, 2008. 1
- [9] Mayur Datar, Nicole Immorlica, Piotr Indyk, and Vahab S. Mirrokni. Locality-sensitive hashing scheme based on p-stable distributions. In *Proceedings of the 20th ACM Symposium on Computational Geometr (SCG)*, pages 253 – 262, Brooklyn, NY, 2004. 2
- [10] Ishan Deshpande, Yuan-Ting Hu, Ruoyu Sun, Ayis Pyrros, Nasir Siddiqui, Sanmi Koyejo, Zhizhen Zhao, David A. Forsyth, and Alexander G. Schwing. Max-sliced wasserstein distance and its use for gans. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10648–10656, Long Beach, CA, 2019. 3, 6
- [11] Ishan Deshpande, Ziyu Zhang, and Alexander G. Schwing. Generative modeling using the sliced wasserstein distance. In *Proceedings of the 2018 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3483–3491, Salt Lake City, UT, 2018. 3, 5
- [12] Kamran Ghasedi Dizaji, Feng Zheng, Najmeh Sadoughi, Yanhua Yang, Cheng Deng, and Heng Huang. Unsupervised deep generative adversarial hashing network. In *Proceedings of the 2018 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3664–3673, Salt Lake City, UT, 2018. 2
- [13] Thanh-Toan Do, Anh-Dzung Doan, and Ngai-Man Cheung. Learning to hash with binary deep neural network. In *Proceedings of the 14th European Conference on Computer Visino (ECCV), Part V*, pages 219–234, 2016. 2
- [14] Khoa D Doan, Yingjie Lao, and Ping Li. Backdoor attack with imperceptible input and latent modification. *Advances in Neural Information Processing Systems (NeurIPS)*, 2021. 3
- [15] Khoa D Doan, Saurav Manchanda, Sarkhan Badirli, and Chandan K Reddy. Image hashing by minimizing discrete component-wise wasserstein distance. *arXiv preprint arXiv:2003.00134*, 2020. 2, 3, 6
- [16] Khoa D. Doan, Saurav Manchanda, Suchismit Mahapatra, and Chandan K. Reddy. Interpretable graph similarity computation via differentiable optimal alignment of node embeddings. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*, pages 665–674, Virtual Event, Canada, 2021. 3
- [17] Khoa D Doan, Saurav Manchanda, Fengjiao Wang, Sathiya Keerthi, Avradeep Bhowmik, and Chandan K Reddy. Image generation via minimizing fréchet distance in discriminator feature space. *arXiv preprint arXiv:2003.11774*, 2020. 3
- [18] Khoa D Doan and Chandan K Reddy. Efficient implicit unsupervised text hashing using adversarial autoencoder. In *Proceedings of The Web Conference (WWW)*, pages 684–694, Taipei, 2020. 2, 3
- [19] Khoa D. Doan, Pranjul Yadav, and Chandan K. Reddy. Adversarial factorization autoencoder for look-alike modeling. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management (CIKM)*, pages 2803–2812, Beijing, China, 2019. 3
- [20] Tiezheng Ge, Kaiming He, and Jian Sun. Graph cuts for supervised binary coding. In *Proceedings of the 13th European Conference on Computer Vision (ECCV), Part VII*, pages 250–264, Zurich, Switzerland, 2014. 1, 2
- [21] Michel X. Goemans and David P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. ACM*, 42(6):1115–1145, 1995. 2
- [22] Yunchao Gong, Svetlana Lazebnik, Albert Gordo, and Florent Perronnin. Iterative quantization: A procrustean approach to learning binary codes for large-scale image retrieval. *IEEE Trans. Pattern Anal. Mach. Intell.*, 35(12):2916–2929, 2013. 1, 2
- [23] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron C. Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems (NIPS)*, pages 2672–2680, Montreal, Canada, 2014. 3, 5
- [24] Junfeng He, Shih-Fu Chang, Regunathan Radhakrishnan, and Claus Bauer. Compact hashing with joint optimization of search accuracy and time. In *Proceedings of the 24th IEEE Conference on Computer Vision and Pattern Recogni-*

- tion (CVPR), pages 753–760, Colorado Springs, CO, 2011. 1, 3
- [25] Kaiming He, Fang Wen, and Jian Sun. K-means hashing: An affinity-preserving quantization method for learning binary compact codes. In *Proceedings of the 2013 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2938–2945, Portland, OR, 2013. 2
- [26] Jae-Pil Heo, Youngwoon Lee, Junfeng He, Shih-Fu Chang, and Sung-Eui Yoon. Spherical hashing. In *Proceedings of the 2012 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2957–2964, Providence, RI, 2012. 2
- [27] Jiun Tian Hoe, Kam Woh Ng, Tianyu Zhang, Chee Seng Chan, Yi-Zhe Song, and Tao Xiang. One loss for all: Deep hashing with a single cosine similarity based learning objective. *Advances in Neural Information Processing Systems (NeurIPS)*, 2021. 2
- [28] Chen Huang, Chen Change Loy, and Xiaoou Tang. Unsupervised learning of discriminative attributes and visual representations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5175–5184, Las Vegas, NV, 2016. 2
- [29] Shanshan Huang, Yichao Xiong, Ya Zhang, and Jia Wang. Unsupervised triplet hashing for fast image retrieval. In *Proceedings of the Thematic Workshops of ACM Multimedia*, pages 84–92, 2017. 2
- [30] Sergey Ioffe. Improved consistent sampling, weighted min-hash and L1 sketching. In *The 10th IEEE International Conference on Data Mining (ICDM)*, pages 246–255, Sydney, Australia, 2010. 2
- [31] Soheil Kolouri, Kimia Nadjahi, Umut Simsekli, Roland Badeau, and Gustavo Rohde. Generalized sliced wasserstein distances. pages 261–272, Vancouver, Canada, 2019. 3, 5, 6
- [32] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems (NIPS)*, pages 1106–1114, Lake Tahoe, NV, 2012. 6
- [33] Brian Kulis and Trevor Darrell. Learning to hash with binary reconstructive embeddings. In *Advances in Neural Information Processing Systems (NIPS)*, pages 1042–1050, Vancouver, Canada, 2009. 2
- [34] Ping Li. Binary and multi-bit coding for stable random projections. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 1430–1438, Fort Lauderdale, FL, 2017. 2
- [35] Ping Li. Linearized GMM kernels and normalized random Fourier features. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, pages 315–324, 2017. 2
- [36] Ping Li and Arnd Christian König. b-bit minwise hashing. In *Proceedings of the 19th International Conference on World Wide Web (WWW)*, pages 671–680, Raleigh, NC, 2010. 2
- [37] Ping Li, Michael Mitzenmacher, and Anshumali Shrivastava. Coding for random projections. In *Proceedings of the 31th International Conference on Machine Learning (ICML)*, pages 676–684, Beijing, China, 2014. 2
- [38] Ping Li, Anshumali Shrivastava, Joshua Moore, and Arnd Christian König. Hashing algorithms for large-scale learning. In *Advances in Neural Information Processing Systems (NIPS)*, pages 2672–2680, Granada, Spain, 2011. 1
- [39] Ping Li and Martin Slawski. Simple strategies for recovering inner products from coarsely quantized random projections. In *Advances in Neural Information Processing Systems (NIPS)*, pages 4567–4576, Long Beach, CA, 2017. 2
- [40] Qi Li, Zhenan Sun, Ran He, and Tieniu Tan. Deep supervised discrete hashing. In *Neural Information Processing Systems (NIPS)*, pages 2479–2488, Long Beach, CA, 2017. 1, 2, 5, 6, 7
- [41] Xiaoyun Li and Ping Li. One-sketch-for-all: Non-linear random features from compressed linear measurements. In *Proceedings of the 24th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 2647–2655, Virtual Event, 2021. 2
- [42] Kevin Lin, Jiwen Lu, Chu-Song Chen, and Jie Zhou. Learning compact binary descriptors with unsupervised deep neural networks. In *Proceedings of the 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1183–1192, Las Vegas, NV, 2016. 2
- [43] Kevin Lin, Huei-Fang Yang, Jen-Hao Hsiao, and Chu-Song Chen. Deep learning of binary hash codes for fast image retrieval. In *Proceedings of the 2015 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 27–35, Boston, MA, 2015. 2
- [44] Wei Liu, Jun Wang, Rongrong Ji, Yu-Gang Jiang, and Shih-Fu Chang. Supervised hashing with kernels. In *Proceedings of the 2012 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2074–2081, Providence, RI, 2012. 2
- [45] Mark Manasse, Frank McSherry, and Kunal Talwar. Consistent weighted sampling. Technical Report MSR-TR-2010-73, Microsoft Research, 2010. 2
- [46] Khai Nguyen, Nhat Ho, Tung Pham, and Hung Bui. Distributional sliced-wasserstein and applications to generative modeling. In *Proceedings of the 9th International Conference on Learning Representations (ICLR)*, Virtual Event, Austria, 2021. 6
- [47] Ruslan Salakhutdinov and Geoffrey Hinton. Semantic hashing. *International Journal of Approximate Reasoning*, 50(7):969–978, 2009. 1, 2
- [48] Fumin Shen, Chunhua Shen, Wei Liu, and Heng Tao Shen. Supervised discrete hashing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 37–45, Boston, MA, 2015. 1, 2
- [49] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *Proceedings of the 3rd International Conference on Learning Representations (ICLR)*, San Diego, CA, 2015. 6
- [50] Shupeng Su, Chao Zhang, Kai Han, and Yonghong Tian. Greedy hash: Towards fast optimization for accurate hash coding in CNN. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 806–815, Montréal, Canada, 2018. 1, 6, 7
- [51] Santosh S. Vempala. *The Random Projection Method*. American Mathematical Society, 2004. 2

- [52] Jingdong Wang, Ting Zhang, Jingkuan Song, Nicu Sebe, and Heng Tao Shen. A survey on learning to hash. *IEEE Trans. Pattern Anal. Mach. Intell.*, 40(4):769–790, 2018. [1](#), [2](#), [3](#)
- [53] Yair Weiss, Antonio Torralba, and Robert Fergus. Spectral hashing. In *Advances in neural information processing systems (NIPS)*, pages 1753–1760, Vancouver, Canada, 2008. [1](#), [2](#), [3](#)
- [54] Rongkai Xia, Yan Pan, Hanjiang Lai, Cong Liu, and Shuicheng Yan. Supervised hashing for image retrieval via image representation learning. In Carla E. Brodley and Peter Stone, editors, *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence (AAAI)*, pages 2156–2162, Québec City, Canada, 2014. [1](#), [2](#)
- [55] Erkun Yang, Cheng Deng, Tongliang Liu, Wei Liu, and Dacheng Tao. Semantic structure-based unsupervised deep hashing. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1064–1070, Stockholm, Sweden, 2018. [2](#)
- [56] Erkun Yang, Tongliang Liu, Cheng Deng, Wei Liu, and Dacheng Tao. Distillhash: Unsupervised deep hashing by distilling data pairs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2946–2955, Long Beach, CA, 2019. [2](#)
- [57] Huei-Fang Yang, Kevin Lin, and Chu-Song Chen. Supervised learning of semantics-preserving hash via deep convolutional neural networks. *IEEE Trans. Pattern Anal. Mach. Intell.*, 40(2):437–451, 2018. [1](#), [2](#)
- [58] Li Yuan, Tao Wang, Xiaopeng Zhang, Francis EH Tay, Zequn Jie, Wei Liu, and Jiashi Feng. Central similarity quantization for efficient image and video retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3083–3092, Seattle, WA, 2020. [2](#), [4](#), [6](#), [7](#)
- [59] Jian Zhang, Avner May, Tri Dao, and Christopher Ré. Low-precision random fourier features for memory-constrained kernel approximation. In *Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 1264–1274, Naha, Okinawa, Japan, 2019. [2](#)
- [60] Xiangtao Zheng, Yichao Zhang, and Xiaoqiang Lu. Deep balanced discrete hashing for image retrieval. *Neurocomputing*, 403:224–236, 2020. [1](#), [2](#), [6](#), [7](#)