

A General Two-Step Approach to Learning-Based Hashing

Guosheng Lin, Chunhua Shen*, David Suter, Anton van den Hengel
University of Adelaide, SA 5005, Australia

Abstract

Most existing approaches to hashing apply a single form of hash function, and an optimization process which is typically deeply coupled to this specific form. This tight coupling restricts the flexibility of the method to respond to the data, and can result in complex optimization problems that are difficult to solve. Here we propose a flexible yet simple framework that is able to accommodate different types of loss functions and hash functions. This framework allows a number of existing approaches to hashing to be placed in context, and simplifies the development of new problem-specific hashing methods. Our framework decomposes the hashing learning problem into two steps: hash bit learning and hash function learning based on the learned bits. The first step can typically be formulated as binary quadratic problems, and the second step can be accomplished by training standard binary classifiers. Both problems have been extensively studied in the literature. Our extensive experiments demonstrate that the proposed framework is effective, flexible and outperforms the state-of-the-art.

1. Introduction

Recently hashing methods have been widely used for a variety of applications, but have been particularly successful when applied to approximate nearest neighbor search. Hashing methods construct a set of hash functions that map the original high-dimensional data into a compact binary space. The resulting binary codes enable fast similarity search on the basis of the Hamming distance between codes. Moreover, compact binary codes are extremely efficient for large-scale data storage. Applications in computer vision include content-based image retrieval, object recognition [12], image matching, etc. In general, hash functions are generated with the aim of preserving some notion of similarity between data points. One of the seminal approaches in this vein is locality-sensitive hashing (LSH) [2], which randomly generates hash functions to approximate cosine similarity. Compared to this *data-independent* method, recent work has focused on *data-dependent* approaches for

generating more effective hash functions. In this category, a number of methods have been proposed, for example: spectral hashing (SPH) [15], multi-dimension spectral hashing (MDSH) [14], iterative quantization (ITQ) [3] and inductive manifold hashing [11]. These methods do not rely on labeled data and are thus categorized as unsupervised hashing methods. Another category is the supervised hashing methods. Recent works include supervised hashing with kernels (KSH) [8], minimal loss hashing (MLH) [10], supervised binary reconstructive embeddings (BRE) [5], semi-supervised sequential projection learning hashing (SPLH) [13] and column generation hashing [7], etc.

Loss functions for hashing are typically defined on the basis of the Hamming distance or Hamming affinity of similar and dissimilar data pairs. Hamming affinity is calculated by the inner product of two binary codes (a binary code takes a value of $\{-1, 1\}$). Existing methods thus tend to optimize a single form of hash functions, the parameters of which are directly optimized against the overall loss function. The common forms of hash functions are linear perceptron functions (MLH, SPLH, LSH), kernel functions (KSH), eigenfunctions (SPH, MDSH). The optimization procedure is then coupled with the selected family of hash function. Different types of hash functions offer a trade-off between testing time and ranking accuracy. For example, compared with kernel functions, the simple linear perceptron function is usually much more efficient for evaluation but can have a relatively low accuracy for nearest neighbor search. Moreover, this coupling often results in a highly non-convex problem which can be very difficult to optimize.

As an example, the loss functions in MDSH, KSH and BRE all take a similar form that aims to minimize the difference between the Hamming affinity (or distance) and the ground truth of data pairs. However, the optimization procedures used in these methods are coupled with the form of hash functions (eigenfunctions, kernel functions) and thus different optimization techniques are needed.

Self-Taught Hashing (STH) [16] is a method which decomposes the learning procedure into two steps: binary code generating and hash function learning. We extend this idea and propose a general two-step approach to hashing of which STH can be seen as a specific example. Note that

*This work was in part supported by ARC grant FT120100969. Correspondence should be addressed to C. Shen (e-mail: chhshen@gmail.com).

STH optimizes the Laplacian affinity loss, which only tries to pull together those similar data pairs but does not push away those dissimilar data pairs and that, as has been shown in manifold learning, this may lead to inferior performance.

Our framework, however, is able to accommodate many different loss functions defined on the Hamming affinity of data pairs, such as the loss function used in KSH, BRE or MLH. This more general family of loss functions may consider both similar and dissimilar data pairs. In order to produce effective binary codes in this first step, we develop a new technique based on coordinate descent. We show that at each iteration of coordinate descent, we can formulate the optimization problem of any Hamming affinity loss as a binary quadratic problem. This formulation unifies different types of objective functions into the same optimization problem, which significantly simplifies the optimization effort. Our main contributions are as follows.

(1) We propose a flexible hashing framework that decomposes the learning procedure into two steps: binary codes inference step and hash function learning step. This decomposition simplifies the problem and enables the use of different types of loss functions and simplifies the hash function learning problem into a standard binary classification problem. An arbitrary classifier, such as linear or kernel support vector machines (SVM), boosting, neural networks, may thus be adopted to train the hash functions.

(2) For binary code inference, we show that optimization using different types of loss functions (e.g., loss functions in KSH, BRE, MLH) can be solved as a series of binary quadratic problems. We show that any type of loss function (e.g., the ℓ_2 loss, exponential loss, hinge loss) defined on Hamming affinity of data pairs can be equivalently converted into a standard quadratic function. Based on this key observation, we propose a general block coordinate decent method that is able to incorporate many different types of loss functions in a unified manner.

(3) The proposed method is simple and easy to implement. We carry out extensive experiments on nearest neighbor search for image retrieval. To show the flexibility, we evaluate our method using different types of loss functions and different formats of hash functions (linear SVM, kernel SVM, Adaboost with decision stumps, etc). Experiments show that our method outperforms the state-of-the-art.

2. Two-Step Hashing

Given a set of training points $\mathcal{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n\} \subset \mathbb{R}^d$, the goal of hashing is to learn a set of hash functions that are able to preserve some notion of similarity between data points. A ground truth affinity (or distance) matrix, \mathbf{Y} , is provided (or calculated by a pre-defined rule) for training, which defines the (dis-)similarity relations between data pairs. In this case y_{ij} is the (i, j) -th element of the matrix \mathbf{Y} , which is an affinity value of the data pair $(\mathbf{x}_i, \mathbf{x}_j)$.

As a simple example, if the data labels are available, y_{ij} can be defined as 1 for similar data pairs to t and -1 for dissimilar data pairs. In the case of unsupervised learning, y_{ij} can be defined as the Euclidean distance or Gaussian affinity on data points. $\Phi(\cdot)$ is a set of m hash functions: $\Phi(\cdot) = [h_1(\cdot), h_2(\cdot), \dots, h_m(\cdot)]$. The output of the hash functions are m -bit binary codes: $\Phi(\mathbf{x}) \in \{-1, 1\}^m$. In general, the optimization can be written as:

$$\min_{\Phi(\cdot)} \sum_{i=1}^n \sum_{j=1}^n \delta_{ij} L(\Phi(\mathbf{x}_i), \Phi(\mathbf{x}_j); y_{ij}). \quad (1)$$

Here $\delta_{ij} \in \{0, 1\}$ indicates whether the relation between two data points is defined, and $L(\Phi(\mathbf{x}_i), \Phi(\mathbf{x}_j); y_{ij})$ is a loss function that measures the how well the binary codes match the expected affinity (or distance) y_{ij} . Many different types of loss functions $L(\cdot)$ have been devised, and will be discussed in detail in the next section.

Most existing methods try to directly optimize objective (1) in order to learn the parameters of hash functions [5, 8, 10, 14]. This inevitably means that the optimization process is tightly coupled to the form of hash functions used, which makes it non-trivial to extend a method to use another different format of hash functions. Moreover, this coupling usually results in highly non-convex problems. Following the idea of STH [16], we decompose the learning procedure into two steps: the first step for binary code inference and the second step for hash function learning. The first step is to solve the optimization:

$$\min_{\mathbf{Z}} \sum_{i=1}^n \sum_{j=1}^n \delta_{ij} L(\mathbf{z}_i, \mathbf{z}_j; y_{ij}), \text{ s.t.: } \mathbf{Z} \in \{-1, 1\}^{m \times n}, \quad (2)$$

where \mathbf{Z} is the matrix of m -bit binary codes for all data points, and \mathbf{z}_i is the binary code vector corresponding to i -th data point.

The second step is to learn hash functions based on the binary codes obtained in the first step, which is achieved by solving the optimization problem:

$$\min_{\Phi(\cdot)} \sum_{i=1}^n F(\mathbf{z}_i, \Phi(\mathbf{x}_i)). \quad (3)$$

Here $F(\cdot, \cdot)$ is a loss function. We solve the above optimization independently for each of the m bits. To learn the k -th hash function $h_k(\cdot)$, the optimization can be written:

$$\min_{h_k(\cdot)} \sum_{i=1}^n F'(z_{i,k}, h_k(\mathbf{x}_i)). \quad (4)$$

Here $F'(\cdot, \cdot)$ is an loss function defined on two codes; $z_{i,k}$ is the binary code corresponding to the i -th data point and the k -th bit. Clearly, the above optimization is a binary classification problem which is to minimize a kind of loss given the binary labels. For example, the loss function $F'(\cdot)$ can be an zero-one loss function returning 0 if two inputs have the same value, and 1 otherwise. As in classification, one

Algorithm 1 Block coordinate decent for learning binary codes (Step-1)

1: **Input:** affinity matrix \mathbf{Y} , bit length m , number of cyclic iteration r .
2: **Initialize** the binary code matrix \mathbf{Z} .
3: **Repeat**
4: **For** $t = 1, 2, \dots, m$
5: Solve the binary quadratic problem (BQP) in (13)
 to obtain the binary codes of t -th bit.
6: Update the codes of the t -th bit in the code matrix \mathbf{Z} .
7: **End For**
8: **Until** the maximum cyclic iteration r is reached.
9: **Output:** the matrix of binary codes \mathbf{Z}

can also use a convex surrogate to replace the zero-one loss. Typical surrogate loss functions are hinge loss, logistic loss, etc. The resulting classifier is the hash function that we aim to learn. Therefore, we are able to use any form of classifier. For example, we can learn perceptron hash functions by training a linear SVM. The linear perceptron hash function has the form:

$$h(\mathbf{x}) = \text{sign}(\mathbf{w}^\top \mathbf{x} + b). \quad (5)$$

We could also train, for example, an RBF-kernel SVM, or Adaboost as hash functions. Here we describe a kernel hash function that is learned using a linear SVM on kernel-transferred features (referred to as SVM-KF). The hash function learned by SVM-KF has a form as follows:

$$h(\mathbf{x}) = \text{sign} \left[\sum_{j=q}^Q w_q \kappa(\mathbf{x}'_q, \mathbf{x}) + b \right], \quad (6)$$

in which $\mathcal{X}' = \{\mathbf{x}'_1, \dots, \mathbf{x}'_Q\}$ are Q data points generated from the training set by random or uniform sampling.

We evaluate variety of different kinds of hash function in the Experiments Section below. These tests show that Kernel hash functions often offer better ranking precision but require much more evaluation time than linear perceptron hash functions. The hash functions learned by SVM-KF represents a trade-off between kernel SVM and linear SVM.

The method proposed here is labeled Two-Step Hashing (TSH), the steps are as follows:

- Step-1: Solving the optimization problem in (2) using block coordinate decent (Algorithm 1) to obtain binary codes for each training data point.
- Step-2: Solving the binary classification problem in (4) for each bit based on the binary codes obtained at Step-1.

3. Solving binary quadratic problems

Optimizing (2) in Step-1 for the entire binary code matrix can be difficult. Instead, we develop a bit-wise block coordinate descent method so that the problem at each iteration can be solved easily. Moreover, we show that at each iteration, any pairwise Hamming affinity (or distance) based loss can be equivalently formulated as a binary quadratic

problem. Thus we are able to easily work with different loss functions.

Block coordinate decent (BCD) is a technique that iteratively optimizes a subset of variables at a time. For each iteration, we pick one bit for optimization in a cyclic fashion. The optimization for the k -th bit can be written as:

$$\min_{\mathbf{z}^{(k)}} \sum_{i=1}^n \sum_{j=1}^n \delta_{ij} l_k(z_{i,k}, z_{j,k}), \quad \text{s.t. } \mathbf{z}^{(k)} \in \{-1, 1\}^n, \quad (7)$$

where l_k is the loss function defined on the k -th bit: $l_k(z_{i,k}, z_{j,k}) = L(z_{i,k}, z_{j,k}, \bar{\mathbf{z}}_i, \bar{\mathbf{z}}_j; y_{ij})$. Here $\mathbf{z}^{(k)}$ contains the binary codes of the k -th bit. $z_{i,k}$ is the binary code of the i -th data point and the k -th bit. $\bar{\mathbf{z}}_i$ is the binary codes of the i -th data point excluding the k -th bit.

Thus far, we have not described the form of the loss function $L(\cdot)$. Our optimization method is not restricted to optimizing a specified form of the loss function. Based on the following proposition, we are able to rewrite any Hamming affinity (or distance) based loss function $L(\cdot)$ into a standard quadratic problem.

Proposition 1. For any loss function $l(z_1, z_2)$ that is defined on a pair of binary input variables $z_1, z_2 \in \{-1, 1\}$ and $l(1, 1) = l(-1, -1)$, $l(1, -1) = l(-1, 1)$, we can define a quadratic function $g(z_1, z_2)$ that is equal to $l(z_1, z_2)$. We have following equation:

$$l(z_1, z_2) = \frac{1}{2} \left[z_1 z_2 (l^{(11)} - l^{(-11)}) + l^{(11)} + l^{(-11)} \right], \quad (8)$$

$$= \frac{1}{2} z_1 z_2 (l^{(11)} - l^{(-11)}) + \text{const.} \quad (9)$$

$$= g(z_1, z_2). \quad (10)$$

Here $l^{(11)}, l^{(-11)}$ are constants, $l^{(11)}$ is the loss output on identical input pair: $l^{(11)} = l(1, 1)$, and $l^{(-11)}$ is the loss output on distinct input pair: $l^{(-11)} = l(-1, 1)$.

Proof. This proposition can be easily proved by exhaustively checking all possible inputs of the loss function. Notice that there are only two possible output values of the loss function. For the input ($z_1 = 1, z_2 = 1$):

$$g(1, 1) = \frac{1}{2} \left[1 \times 1 \times (l^{(11)} - l^{(-11)}) + l^{(11)} + l^{(-11)} \right] \\ = l(1, 1),$$

For the input ($z_1 = -1, z_2 = 1$):

$$g(-1, 1) = \frac{1}{2} \left[-1 \times 1 \times (l^{(11)} - l^{(-11)}) + l^{(11)} + l^{(-11)} \right] \\ = l(-1, 1),$$

The input ($z_1 = -1, z_2 = -1$) is the same as ($z_1 = 1, z_2 = 1$) and the input ($z_1 = 1, z_2 = -1$) is the same as ($z_1 = -1, z_2 = 1$). In conclusion, the function l and g have the same output for any possible inputs. \square

Any hash loss function $l(\cdot, \cdot)$ which is defined on the Hamming affinity between, or Hamming distance of, data

pairs is able to meet the requirement that: $l(1, 1) = l(-1, -1)$, $l(1, -1) = l(-1, 1)$. Applying this proposition, the optimization of (7) can be equivalently reformulated as:

$$\min_{\mathbf{z}_{(k)} \in \{-1, 1\}^n} \sum_{i=1}^n \sum_{j=1}^n \delta_{ij} (l_{k,i,j}^{(11)} - l_{k,i,j}^{(-11)}) z_{i,k} z_{j,k}, \quad (11)$$

The above optimization is an unconstrained binary quadratic problem. Let $a_{i,j}$ denote the (i, j) -th element of matrix \mathbf{A} , which we define as:

$$a_{i,j} = \delta_{ij} (l_{k,i,j}^{(11)} - l_{k,i,j}^{(-11)}). \quad (12)$$

The above optimization (11) can be written in matrix form:

$$\min_{\mathbf{z}_{(k)}} \mathbf{z}_{(k)}^\top \mathbf{A} \mathbf{z}_{(k)}, \quad \text{s.t. } \mathbf{z}_{(k)} \in \{-1, 1\}^n. \quad (13)$$

We have shown that at each iteration, the original optimization in (7) can be equivalently reformulated as a binary quadratic problem (BQP) in (13). BQP has been extensively studied. To solve (13), we first apply the spectral relaxation to get an initial solution. Spectral relaxation drops the binary constraints. The optimization becomes

$$\min_{\mathbf{z}_{(k)}} \mathbf{z}_{(k)}^\top \mathbf{A} \mathbf{z}_{(k)}, \quad \text{s.t. } \|\mathbf{z}_{(k)}\|_2^2 = n. \quad (14)$$

The solution (denoted $\mathbf{z}_{(k)}^0$) of the above optimization is simply the eigenvector that corresponds to the minimum eigenvalue of the matrix \mathbf{A} . To achieve a better solution, here we take a step further. We solve the following relaxed problem of (13) as follows

$$\min_{\mathbf{z}_{(k)}} \mathbf{z}_{(k)}^\top \mathbf{A} \mathbf{z}_{(k)}, \quad \text{s.t. } \mathbf{z}_{(k)} \in [-1, 1]^n. \quad (15)$$

This relaxation is tighter than the spectral relaxation and provides a solution of better quality. To solve the above problem, we use the solution $\mathbf{z}_{(k)}^0$ of spectral relaxation in (14) as initialization and solve it using the efficient LBFGS-B solver [17]. The algorithm for binary code inference in Step-1 is summarized in Algorithm 1.

The approach proposed above is applicable to many different types of loss functions, which are defined on Hamming distance or Hamming affinity, such as the ℓ_2 loss, exponential loss, hinge loss. Here we describe a selection of such loss functions, most of which arise from recently proposed hashing methods. We evaluate these loss functions in the Experiments Section below. Note that m is the number of bits, and $d_h(\cdot, \cdot)$ is the Hamming distance on data pairs. If not specified, $y_{ij} = 1$ if the data pair is similar, and $y_{ij} = -1$ if the data pair is dissimilar. $\delta(\cdot) \in \{0, 1\}$ is an indicator function.

TSH-KSH The KSH loss function is based on Hamming affinity using ℓ_2 loss function. MDSH also uses a similar form of loss function (weighted Hamming affinity instead):

$$L_{\text{KSH}}(\mathbf{z}_i, \mathbf{z}_j) = (\mathbf{z}_i^\top \mathbf{z}_j - m y_{ij})^2. \quad (16)$$

TSH-BRE The BRE loss function is based on Hamming distance using the ℓ_2 loss function:

$$L_{\text{BRE}}(\mathbf{z}_i, \mathbf{z}_j) = (d_h(\mathbf{z}_i, \mathbf{z}_j)/m - \delta(y_{ij} < 0))^2. \quad (17)$$

Table 2: Training time (in seconds) for TSH using different loss functions, and several other supervised methods on 3 datasets. The value inside a brackets is the time used in the first step for inferring the binary codes. The results show that our method is efficient. Note that the second step of learning the hash functions can be easily parallelised.

	LABELME	MNIST	CIFAR10
TSH-KSH	198 (107)	341 (294)	326 (262)
TSH-BRE	133 (33)	309 (264)	234 (175)
TSH-EE	124 (29)	302 (249)	287 (225)
TSH-ExpH	128 (43)	334 (281)	344 (256)
STHs-RBF	133	99	95
KSH	326	355	379
BREs	216	615	231
MLH	670	805	658

TSH-SPLH It uses an exponential loss outside the loss function proposed in SPLH which is based on the Hamming affinity of data pairs:

$$L_{\text{SPLH}}(\mathbf{z}_i, \mathbf{z}_j) = \exp \left[\frac{-y_{ij} \mathbf{z}_i^\top \mathbf{z}_j}{m} \right]. \quad (18)$$

TSH-EE Elastic Embedding (EE) is a dimension reduction method proposed in [1]. Here we use their loss function with some modifications, which is a exponential based on distance. Here λ is a trade-off parameter:

$$L_{\text{EE}}(\mathbf{z}_i, \mathbf{z}_j) = \delta(y_{ij} > 0) d_h(\mathbf{z}_i, \mathbf{z}_j) + \lambda \delta(y_{ij} < 0) \exp[-d_h(\mathbf{z}_i, \mathbf{z}_j)/m]. \quad (19)$$

TSH-ExpH Here ExpH is an exponential loss function using the Hamming distance:

$$L_{\text{ExpH}}(\mathbf{z}_i, \mathbf{z}_j) = \exp \left[\frac{y_{ij} d_h(\mathbf{z}_i, \mathbf{z}_j) + m \delta(y_{ij} < 0)}{m} \right]. \quad (20)$$

4. Experiments

We compare with a few state-of-the-art hashing methods, including 6 (semi-)supervised methods: Supervised Hashing with Kernels (KSH) [8], Iterative Quantization with supervised embedding (ITQ-CCA) [3], Minimal Loss Hashing (MLH) [10], Supervised Binary Reconstructive Embeddings (BREs) [5] and its unsupervised version BRE, Supervised Self-Taught Hashing (STHs) [16] and its unsupervised version STH, Semi-supervised sequential Projection Learning Hashing(SPLH) [13], and 7 unsupervised methods: Locality-Sensitive Hashing (LSH) [2], Iterative Quantization (ITQ) [3], Anchor Graph Hashing (AGH) [9], Spectral Hashing (SPH [15]), Spherical Hashing (SPHER) [4], Multi-dimension Spectral Hashing (MDSH) [14], Kernelized Locality-Sensitive Hashing KLSH [6]. For comparison methods, we follow the original papers for parameter setting. For SPLH, the regularization trade-off parameter is picked from 0.01 to 1. We use the hierarchical variant of AGH. For each dataset, the bandwidth parameters of Gaussian affinity in MDSH and RBF kernel in KLSH, KSH and our method TSH is set as $\sigma = t\bar{d}$. Here \bar{d} is the average Euclidean distance of top 100 nearest neighbours and t is

Table 1: Results (using hash codes of 32 bits) of TSH using different loss functions, and a selection of other supervised and unsupervised methods on 3 datasets. The upper part relates the results on training data and the lower on testing data. The results show that Step-1 of our method is able to generate effective binary codes that outperforms those of competing methods on the training data. On the testing data our method also outperforms others by a large margin in most cases.

Datasets	Precision-Recall			MAP			Precision at K (K=300)		
	LABELME	MNIST	CIFAR10	LABELME	MNIST	CIFAR10	LABELME	MNIST	CIFAR10
Results on training data									
TSH-KSH	0.501	1.000	1.000	0.570	1.000	1.000	0.229	0.667	0.667
TSH-BRE	0.527	1.000	1.000	0.600	1.000	1.000	0.230	0.667	0.667
TSH-SPLH	0.504	1.000	1.000	0.524	1.000	1.000	0.230	0.667	0.667
TSH-EE	0.485	1.000	1.000	0.524	1.000	1.000	0.224	0.667	0.667
TSH-ExpH	0.475	1.000	1.000	0.541	1.000	1.000	0.225	0.667	0.667
STHs	0.335	0.800	0.629	0.387	0.882	0.774	0.176	0.575	0.433
KSH	0.283	0.892	0.585	0.316	0.967	0.652	0.168	0.647	0.481
BREs	0.161	0.445	0.220	0.153	0.504	0.190	0.097	0.376	0.171
SPLH	0.166	0.500	0.292	0.153	0.588	0.302	0.092	0.422	0.260
MLH	0.120	0.547	0.190	0.142	0.685	0.235	0.100	0.478	0.200
Results on testing data									
TSH-KSH	0.175	0.843	0.282	0.296	0.893	0.440	0.293	0.889	0.410
TSH-BRE	0.169	0.844	0.283	0.293	0.896	0.439	0.293	0.890	0.409
TSH-SPLH	0.174	0.840	0.284	0.291	0.895	0.444	0.288	0.891	0.416
TSH-EE	0.169	0.843	0.280	0.288	0.896	0.438	0.286	0.892	0.410
TSH-ExpH	0.172	0.844	0.282	0.287	0.892	0.441	0.286	0.887	0.410
STHs	0.094	0.385	0.144	0.162	0.639	0.229	0.156	0.634	0.218
STHs-RBF	0.151	0.674	0.178	0.274	0.897	0.354	0.271	0.893	0.352
KSH	0.165	0.781	0.249	0.279	0.884	0.407	0.158	0.881	0.398
BREs	0.106	0.409	0.151	0.178	0.703	0.226	0.171	0.702	0.210
MLH	0.100	0.470	0.150	0.181	0.648	0.264	0.174	0.623	0.215
SPLH	0.093	0.452	0.191	0.168	0.714	0.321	0.158	0.708	0.315
ITQ-CCA	0.077	0.619	0.206	0.143	0.792	0.333	0.133	0.784	0.325
MDSH	0.100	0.298	0.150	0.178	0.691	0.288	0.155	0.685	0.228
SHPER	0.102	0.296	0.152	0.185	0.624	0.244	0.176	0.623	0.233
ITQ	0.116	0.386	0.161	0.206	0.750	0.264	0.197	0.751	0.252
AGH	0.096	0.404	0.144	0.194	0.743	0.252	0.187	0.744	0.244
STH	0.077	0.361	0.135	0.135	0.593	0.216	0.125	0.644	0.204
BRE	0.091	0.323	0.137	0.160	0.651	0.238	0.147	0.582	0.185
LSH	0.069	0.211	0.123	0.116	0.459	0.188	0.103	0.448	0.162

picked from 0.01 to 50. For STHs and our method TSH, the trade-off parameter in SVM is picked from $10/n$ to $10^5/n$, n is the number of data points. For our TSH-EE using EE lost function, we simply set the trade-off parameter λ to 100. If not specified, our method TSH use SVM with RBF kernel as hash functions. The cyclic iteration number r in Algorithm 1 is simply set to 1.

We use 2 large scale image datasets and another 3 datasets for evaluation. 2 large image datasets are 580,000 tiny image dataset (Tiny-580K) [3], and Flickr 1 Million image dataset. Another 3 datasets include CIFAR10, MNIST and LabelMe [10].

For the LabelMe dataset, the ground truth pairwise affinity matrix is provided. For other datasets, we use the multi-class labels to define the ground truth affinity by label agreement. Tiny-580K is used in [3]. Flickr-1M dataset consists of 1 million thumbnail images of the MIRFlickr-1M We generate 320-dimension GIST features. For these 2 large datasets, there is no semantic ground truth affinity provided. Following the same setting as other hash methods [8, 13], we generate pseudo-labels for supervised methods according to the ℓ_2 distance. In detail, a data point is labelled as a relevant neighbour to the query if it lies in the top 2 per-

centile points of the whole database. For all datasets, following a common setting in many supervised hashing methods [5, 8, 10], we randomly select 2000 examples as testing queries, and the rest is served as database. We train methods using a subset of the database: 5000 examples for large datasets (Tiny-580K and Flickr-1M) and 2000 examples for the rest. We use 4 types of evaluation measures: precision of top-K retrieved examples (Precision-at-K), Mean Average Precision (MAP), the area under the Precision-Recall curve, precision of retrieved examples within the Hamming distance 2.

4.1. Using different loss functions

We evaluate the performance of our method TSH using different loss functions on 3 datasets: LabelMe, MNIST, CIFAR10. 3 types of evaluation measures are used here: Precision-at-K, Mean Average Precision (MAP) and the area under the Precision-Recall curve. The loss function is defined in Section 3. In particular, our method TSH-KSH uses the KSH [8] loss function, TSH-BRE uses the BRE [5] function. STHs-RBF is the STHs method using RBF kernel hash functions. Our method also uses SVM with RBF kernel as hash functions.



Figure 1: Some retrieval examples of our method TSH on CIFAR10. The first column shows query images, and the rest are top 30 retrieved images in the database. False predictions are marked by red boxes.

First, we evaluate the effectiveness of the Step-1 in our method. We compare the quality of the generated binary codes on the *training data points*. The results are shown in the upper part of the table in Table. 1. The results show that our methods generate high-quality binary codes and outperform others by a large margin. In CIFAR10 and MNIST, we are able to generate perfect codes that match the ground truth similarity. This demonstrates the effectiveness of coordinate descent based hashing codes learning procedure (Step 1 of our framework).

Compared to STHs-RBF, even though we are using the same format of hash function, our overall objective function and the bit-wise binary code inference algorithm may be more effective. Thus our method achieves better performance than STH.

The second part of the result in Table. 1 shows the testing performance. Our method also outperforms others in most cases. Note that MNIST is an ‘easy’ dataset and not as challenging as CIFAR10 and LabelMe. Thus many methods manage to achieve good performance. In the challenging dataset CIFAR10 and LabelMe, our method significantly outperforms others by a large margin.

Overall, for preserving the semantic similarity, supervised methods usually perform much better than those unsupervised methods, which is expected. Our method performs the best, and the running-up methods are STHs-RBF, KSH, and ITQ-CCA.

We show further results of using different numbers of bits in Fig. 2 for supervised methods and Fig. 3 for unsupervised methods on the dataset CIFAR10 and LabelMe. In the figures, TSH denotes our method using BRE loss function. Our method still performs the best in most cases. Some search examples are shown in Fig. 1.

Training time In Table 2, we compare the training time of different methods. It shows that our method is fast compared to the state-of-the-art. We also present the binary code learning time in the table. Notice that in the second step, learning hash functions by binary classification can be easily paralleled which would make our method even more efficient.

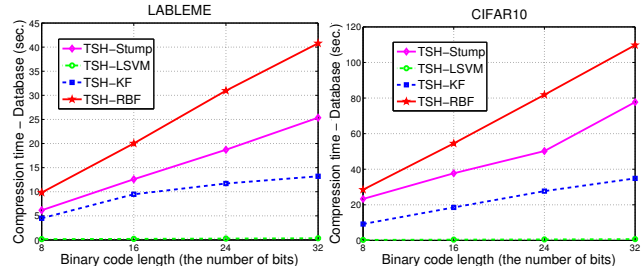


Figure 5: Code compression time using different hash functions. Results show that using kernel transferred feature (TSH-KF) is much faster than SVM with RBF kernel (TSH-RBF). Linear SVM is the fastest one.

4.2. Using different hash functions

We evaluate our method using different hash functions. The hash functions are SVM with RBF kernel (TSH-RBF), linear SVM with kernel transferred feature (TSH-KF), linear SVM (TSH-LSVM), Adaboost with decision-stump (TSH-Stump, 2000 iterations). Results on 3 datasets are shown in Fig. 4. The testing time for different hash functions are shown in Fig. 5.

It shows that the kernel hash functions (TSH-RBF and TSH-KF) achieve best performance in similarity search. However, the testing of linear hash functions is much faster than kernel hash functions. We also find that the testing time of TSH-KF is much faster than TSH-RBF. The TSH-KF is a trade-off between testing time and search performance.

4.3. Large datasets

We carry out experiments on 2 large scale datasets: Flickr 1 million image dataset (Flickr1M) and 580,000 Tiny image dataset (Tiny580k). Results are shown in Fig. 6. Our method TSH achieve on par results with KSH. KSH and our TSH significantly outperform other supervised or unsupervised methods. Notice that there is no semantic similarity ground truth provided on these two datasets. We generate the similarity ground truth using the Euclidean distance. Some unsupervised methods are also able to perform well in this setting (e.g., MDSH, SPHER and ITQ).

5. Conclusion

We have shown that it is possible to place a wide variety of learning-based hashing methods into a unified framework, and that doing so provides insights into the strengths, weaknesses, and commonality between various competing methods. One of the key insights is the fact that the code generation and hash function learning processes may be seen as separate steps, and that the latter may accurately be formulated as a classification problem. This insight enables the development of new approaches to hashing, one of which is detailed above. Experimental testing has validated this approach, and shown that this new approach outperforms the state-of-the-art.

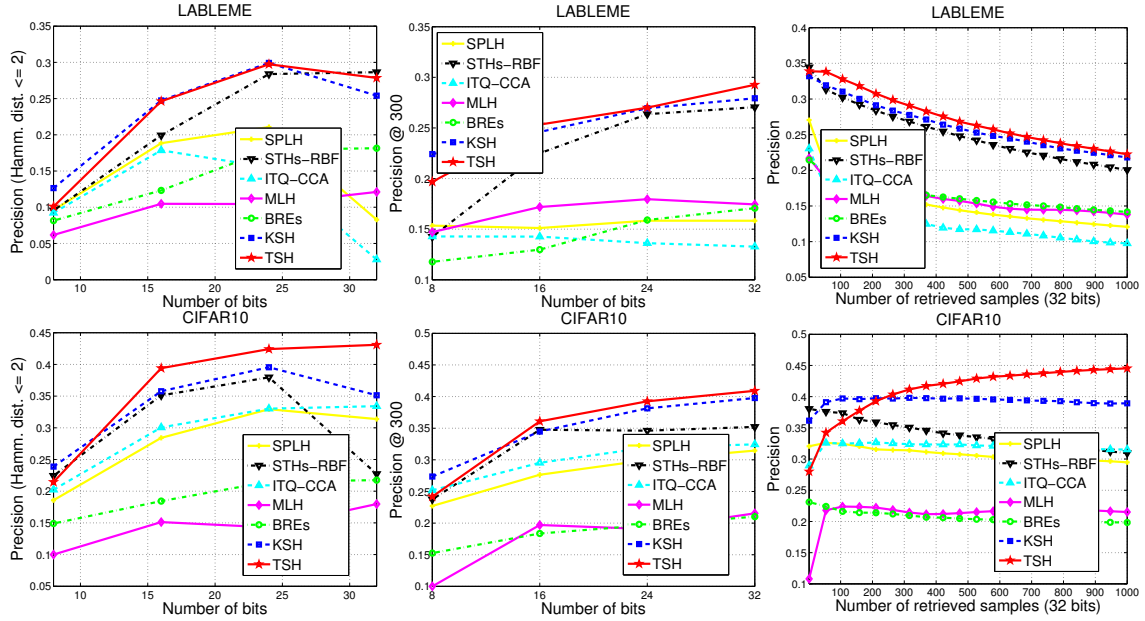


Figure 2: Results on 2 datasets of supervised methods. Results show that TSH outperforms others usually by a large margin. The running up methods are STHs-RBF and KSH.

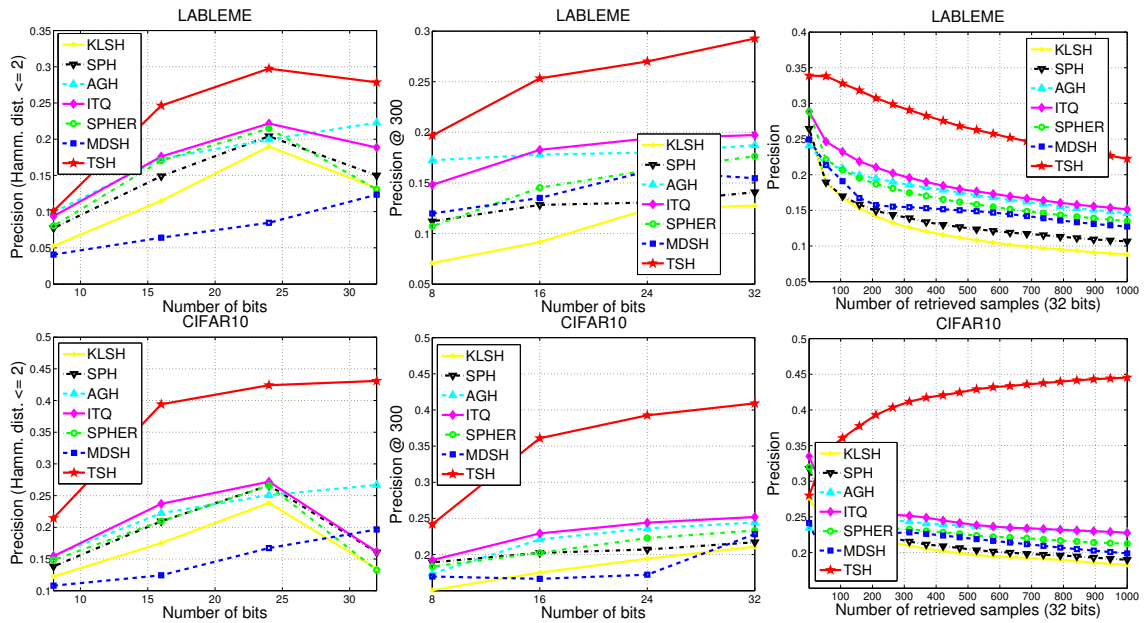


Figure 3: Results on 2 datasets for comparing unsupervised methods. Results show that TSH outperforms others usually by a large margin.

References

- [1] M. Á. Carreira-Perpiñán. The elastic embedding algorithm for dimensionality reduction. In *Proc. Int. Conf. Mach. Learn.*, 2010. 4
- [2] A. Gionis, P. Indyk, and R. Motwani. Similarity search in high dimensions via hashing. In *Proc. Int. Conf. Very Large Data Bases*, 1999. 1, 4
- [3] Y. Gong, S. Lazebnik, A. Gordo, and F. Perronnin. Iterative quantization: a procrustean approach to learning binary codes for large-scale image retrieval. *IEEE Trans. Pattern Analysis Mach. Intelli.*, 2012. 1, 4, 5
- [4] J.-P. Heo, Y. Lee, J. He, S.-F. Chang, and S.-E. Yoon. Spherical hashing. In *Proc. IEEE Conf. Comp. Vis. Pattern Recogn.*, 2012. 4
- [5] B. Kulis and T. Darrell. Learning to hash with binary reconstructive embeddings. In *Proc. Adv. Neural Info. Process. Syst.*, 2009. 1, 2, 4, 5
- [6] B. Kulis and K. Grauman. Kernelized locality-sensitive hashing. *IEEE Trans. Pattern Analysis Mach. Intelli.*, 2012. 4
- [7] X. Li, G. Lin, C. Shen, A. van den Hengel, and A. R. Dick. Learning hash functions using column generation. In *Proc. Int. Conf. Mach. Learn.*, 2013. 1
- [8] W. Liu, J. Wang, R. Ji, Y. Jiang, and S. Chang. Supervised hashing with kernels. In *Proc. IEEE Conf. Comp. Vis. Pattern Recogn.*, 2012. 1, 2, 4, 5
- [9] W. Liu, J. Wang, S. Kumar, and S. F. Chang. Hashing with graphs. In *Proc. Int. Conf. Mach. Learn.*, 2011. 4
- [10] M. Norouzi and D. Fleet. Minimal loss hashing for compact binary codes. In *Proc. Int. Conf. Mach. Learn.*, 2011. 1, 2, 4, 5
- [11] F. Shen, C. Shen, Q. Shi, A. van den Hengel, and Z. Tang. Inductive hashing on manifolds. In *Proc. IEEE Conf. Comp. Vis. Pattern*

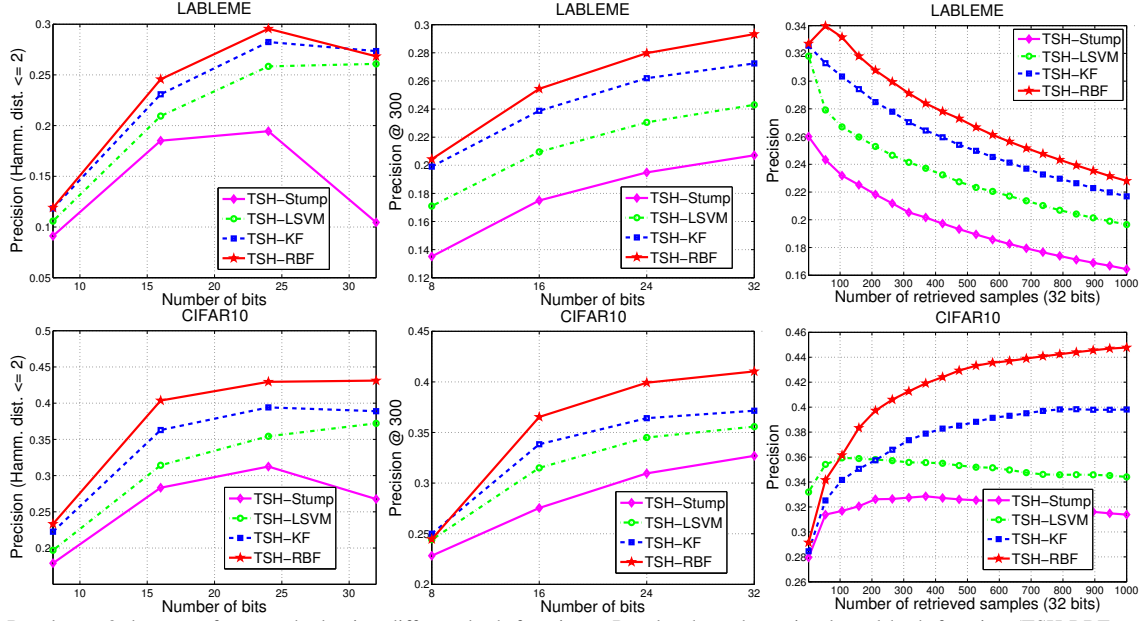


Figure 4: Results on 2 datasets of our method using different hash functions. Results show that using kernel hash function (TSH-RBF and TSH-KF) achieves best performances.

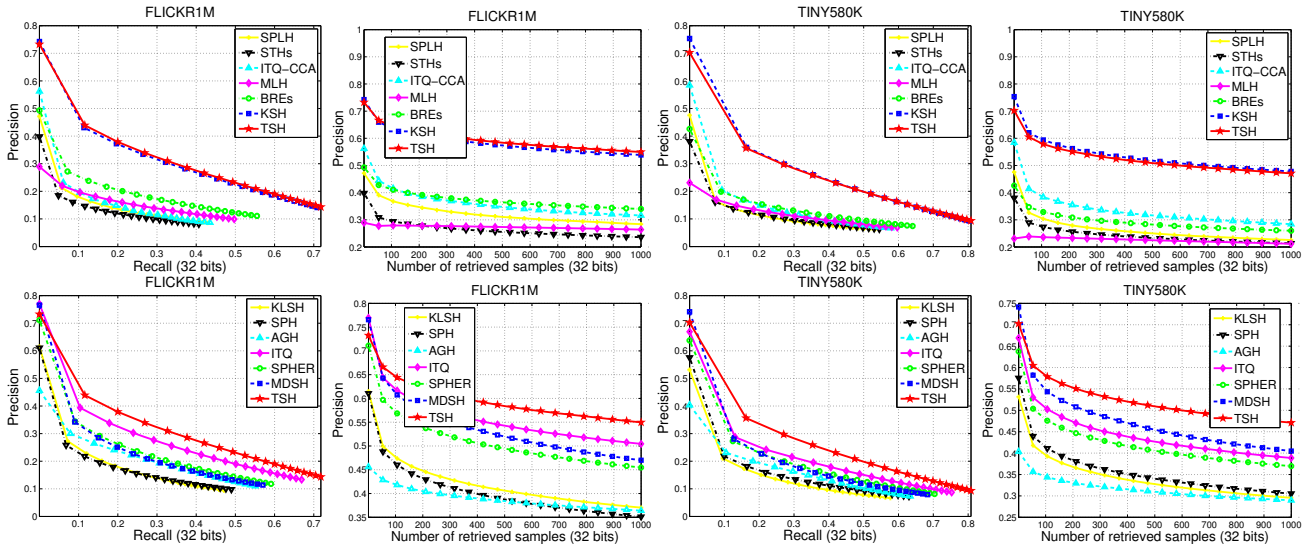


Figure 6: Results on 2 large scale datasets: Flickr1M and Tiny580k. The first row shows the results of supervised methods and the second row for unsupervised methods. Our method TSH achieves on par result with KSH. TSH and KSH significantly outperform other methods.

Recogn., 2013. 1

- [12] A. Torralba, R. Fergus, and W. Freeman. 80 million tiny images: A large data set for nonparametric object and scene recognition. *IEEE Trans. Pattern Analysis Mach. Intell.*, 2008. 1
- [13] J. Wang, S. Kumar, and S. Chang. Semi-supervised hashing for large scale search. *IEEE Trans. Pattern Analysis Mach. Intell.*, 2012. 1, 4, 5
- [14] Y. Weiss, R. Fergus, and A. Torralba. Multidimensional spectral hashing. In *Proc. Eur. Conf. Comp. Vis.*, 2012. 1, 2, 4
- [15] Y. Weiss, A. Torralba, and R. Fergus. Spectral hashing. In *Proc. Adv. Neural Info. Process. Syst.*, 2008. 1, 4
- [16] D. Zhang, J. Wang, D. Cai, and J. Lu. Self-taught hashing for fast similarity search. In *Proc. Annual ACM SIGIR Conf.*, 2010. 1, 2, 4
- [17] C. Zhu, R. H. Byrd, P. Lu, and J. Nocedal. Algorithm 778: L-BFGS-B: Fortran subroutines for large-scale bound-constrained optimization. *ACM T. Math. Softw.*, 1997. 4