

# Zero-Shot Hashing via Transferring Supervised Knowledge

Yang Yang<sup>†</sup>, Yadan Luo<sup>†</sup>, Weilun Chen<sup>†</sup>, Fumin Shen<sup>†</sup>, Jie Shao<sup>†</sup>, Heng Tao Shen<sup>§,†</sup>

<sup>†</sup>University of Electronic Science and Technology of China, Chengdu, China

<sup>§</sup>The University of Queensland, Brisbane, Australia

{dlyyang, lyadanluol, chenweilunster, fumin.shen}@gmail.com

shaojie@uestc.edu.cn, shenht@itee.uq.edu.au

## ABSTRACT

Hashing has shown its efficiency and effectiveness in facilitating large-scale multimedia applications. Supervised knowledge (e.g., semantic labels or pair-wise relationship) associated to data is capable of significantly improving the quality of hash codes and hash functions. However, confronted with the rapid growth of newly-emerging concepts and multimedia data on the Web, existing supervised hashing approaches may easily suffer from the scarcity and validity of supervised information due to the expensive cost of manual labelling. In this paper, we propose a novel hashing scheme, termed *zero-shot hashing* (ZSH), which compresses images of “unseen” categories to binary codes with hash functions learned from limited training data of “seen” categories. Specifically, we project independent data labels (i.e., 0/1-form label vectors) into semantic embedding space, where semantic relationships among all the labels can be precisely characterized and thus seen supervised knowledge can be transferred to unseen classes. Moreover, in order to cope with the semantic shift problem, we rotate the embedded space to more suitably align the embedded semantics with the low-level visual feature space, thereby alleviating the influence of semantic gap. In the meantime, to exert positive effects on learning high-quality hash functions, we further propose to preserve local structural property and discrete nature in binary codes. Besides, we develop an efficient alternating algorithm to solve the ZSH model. Extensive experiments conducted on various real-life datasets show the superior zero-shot image retrieval performance of ZSH as compared to several state-of-the-art hashing methods.

## Keywords

zero-shot hashing; discrete hashing; supervised knowledge transfer; semantic alignment

## 1. INTRODUCTION

Hashing is a powerful indexing technique for enabling efficient retrieval on large-scale multimedia data, such as image [22, 33] and video [1]. Specifically, in order to achieve shorter response time and less computational cost, hashing encodes high-dimensional data into compact binary codes (i.e., 0 or 1) substantially. In this way,

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 '16, October 15-19, 2016, Amsterdam, Netherlands

© 2016 ACM. ISBN 978-1-4503-3603-1/16/10...\$15.00

DOI: <http://dx.doi.org/10.1145/2964284.2964319>

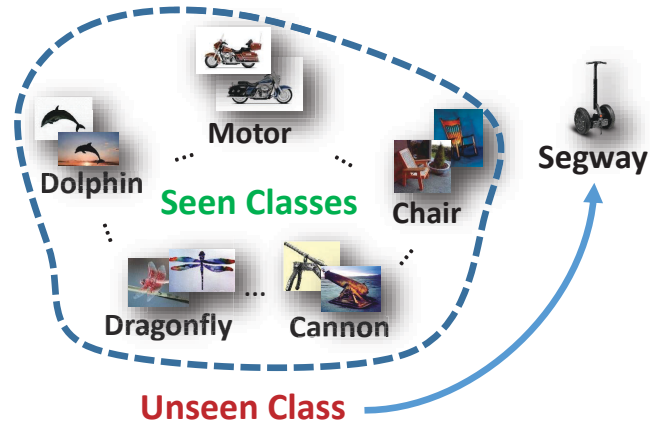


Figure 1: An illustration of newly-emerging concepts and images unseen to the existing learning systems.

data can be compactly stored and Hamming distances can be efficiently calculated with bit-wise XOR operations. Because of its impressive capacity of dealing with “curse of dimensionality” problem, hashing has been extensively employed in various real-world applications, ranging from multimedia indexing [8] to multimedia event detection [20].

There are mainly two branches of hashing, i.e., data-independent hashing and data-dependent hashing. For data-independent hashing, such as Locality Sensitive Hashing [6], no prior knowledge (e.g., supervised information) about data is available, and hash functions are randomly generated. Nonetheless, huge storage and computational overhead might be cost since more than 1,000 bits are usually required to achieve acceptable performance. To address this problem, research directions turn to data-dependent hashing, which leverages information inside data itself. Roughly, data-dependant hashing can be divided into two categories: unsupervised hashing (e.g., Iterative Quantization [7], Order Preserving Hashing [26] and Robust Discrete Spectral Hashing [32]), and supervised hashing (e.g., Supervised Hashing with Kernels [16], Weakly-supervised Hashing [34] and Supervised Discrete Hashing [22]). In general, supervised hashing usually achieves better performance than unsupervised ones because supervised information (e.g., semantic labels and/or pair-wise data relationship) can help to better explore intrinsic data property, thereby generating superior hash codes and hash functions.

Along with the explosive growth of Web data, traditional supervised hashing methods have been facing an enormous challenge, i.e., the generation of reliable supervised knowledge cannot catch up with the rapid increasing speed of newly-emerging semantic

concepts and multimedia data. In other words, due to the expensive cost of manual labelling (time-consuming and labor-intensive), sufficient labelled training data is usually not timely available for learning new hash functions that can accurately encode data of new concepts. As illustrated in Figure 1, within the “seen” zone, where images are attached with known categories, existing supervised hashing algorithms may perform well because they are fed with correct guidance. However, outside the seen area, supervised hashing algorithms may easily fail to generalize to data of new categories that they never observe, *e.g.*, *segway*, a two-wheeled, self-balancing, battery-powered electric vehicle. Moreover, most of current approaches use supervised information in the form of either 0/1 semantic labels or pair-wise data relationship for guiding the learning process, which implies that precious correlation among label semantics are inevitably ignored. One straightforward consequence of the semantic independency is that each category can neither learn from other relevant categories nor distribute its own supervised knowledge to other seen classes and/or even those unseen ones.

The aforementioned disadvantages motivate us to consider if we can encode images of “unseen” categories into binary codes with hash functions learned from limited training samples of “seen” categories? The key challenge of achieving this goal is how to set up a tunnel to transfer supervised knowledge between “seen” and “unseen” categories. In recent years, zero-shot learning (ZSL) [2, 12, 21, 24] has been widely recognized as a way to deal with this problem. The ZSL paradigm aims to learn a general mapping from the feature space to a high-level semantic space, which helps to avoid rebuilding models for unseen categories with extra manually labelled data. ZSL is mostly achieved by using class-attribute descriptors to bridge the semantic gap between low-level features and high-level semantics, where new categories are thus learned using only the relationship between attributes and categories. However, most of existing attribute based ZSL methods still suffer from: (1) erroneous guidance derived from imprecise or incomplete human-labelled attributes [10], which is usually due to the lack of expertise or mislabeling by annotators, etc.; (2) diminishing of discrimination for pre-defined attributes when facing domain shift [14, 19].

Recently, mining other auxiliary datasets has been shown to be helpful to tackle the zero-shot learning problem. For instance, with a huge corpus such as Wikipedia, one can obtain word embeddings that capture distributional similarity in the text corpus [25], such that similar words can be located in similar place. During the learning phase, visual modality can be grounded by the word vectors, and such knowledge can thus be transferred into the learned model. Inspired by this, many approaches choose to utilize auxiliary modalities to address the zero-shot problem. Socher et al. [24] used word embedding as supervision in order to detect novel categories and perform classification accordingly. Frome et al. [5] adopted a similar manner, which connects raw features and word embedding space using the dot-product similarity and hinge rank loss.

As aforementioned, with the explosion of the newly-emerging concepts and multimedia data, we are in urgent demand of a reliable and flexible hash function that can be adopted to hash images of unseen categories. However, in the hashing domain, the zero-shot problem has been rarely studied. In this work, we propose a novel hashing scheme, termed *zero-shot hashing* (ZSH). Inspired by the superior capacity of the word embedding for capturing the semantic correlations among concepts, we map mutually independent labels into a semantic-rich space, where supervised knowledge of both seen and unseen labels can be completely shared. This strategy helps to encode images of unseen categories without any assis-

tance of visual observation in those unknown classes. Besides, even though we cannot retrieve images of exactly the same category, semantically related objects can be returned. Moreover, we recognize the problem of semantic shift caused by off-the-shelf embedding. The embedded space is then rotated to make the hash functions more generalized to images of unseen categories. To further improve the quality of hash functions, we also preserve local structural property and discrete nature in binary codes. We summarize our main contributions as below:

- We address the problem of employing limited training data of seen categories to learn reliable hash functions for transforming images of unseen categories into binary codes. We propose a novel zero-shot hashing scheme, which bridges gaps between originally independent labels through a semantic embedding space. To the best of our knowledge, this is one of the first works that study the problem of hashing data from newly-emerging concepts with limited seen supervised knowledge.
- We devise an effective strategy for transferring available supervised knowledge from seen classes to unseen classes. In particular, we project labels into a word embedding space, where semantic correlations among labels can be quantitatively measured and captured. In this way, unseen labels can leverage the well-established mapping from its semantically close seen categories. For instance, *segway* may learn from *bicycle* and *automobile*.
- Since the initial semantic embedding is from an off-the-shelf word embedding space, which may bring in potential semantic shift between categories and the original visual feature. To alleviate the influence, we propose to further rotate the embedding space to better fit the underlying feature characteristics, thereby narrowing down the semantic gap effectively.
- In order to generate more reliable hash functions, we propose to improve the intermediate binary codes by exploring underlying data properties. Concretely, we impose discrete constraints on binary codes during the learning process as well as preserve data local structure, *i.e.*, if two datums share similar representations in the original space, they are supposed to be close to each other in the learned Hamming space.

The rest of this paper is organized as follows. In Section 2, we briefly review some related work on zero-shot learning and hashing. In Section 3, we elaborate our proposed ZSH, together with an efficient optimization method and the corresponding algorithm analysis. Extensive experimental results on various different datasets are reported and analyzed in Section 4, followed by the conclusion of this work in Section 5.

## 2. RELATED WORK

In this section, we briefly review some related works on zero-shot learning and hashing, and discuss the relationship between our proposed approach and existing solutions.

### 2.1 Zero-Shot Learning

Learning with no data, *a.k.a.*, zero-shot learning, has been proved to be an effective direction to tackle the increasing difficulty posed by insufficient training samples. A promising strategy for solving this problem is to utilize an intermediate semantic layer to represent an image. Specifically, with visual attributes or other semantic abundant descriptors, a novel image can thus be defined as

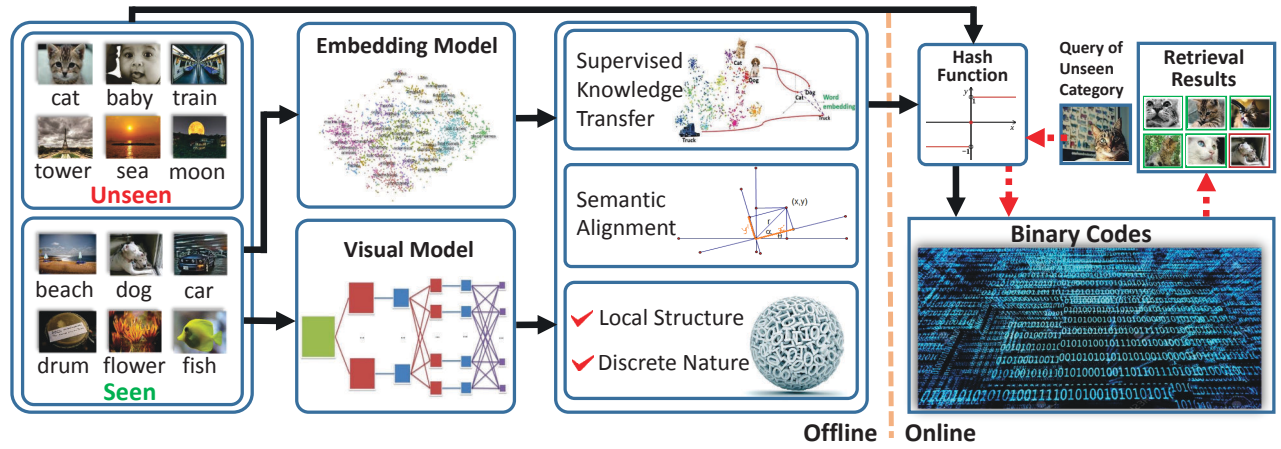


Figure 2: The overall architecture of the proposed zero-shot hashing framework.

the relationship between category and intermediate representation. Farhadi *et al.* [4] proposed to classify unseen objects by describing them with visual attributes. The work in [15] studies the problem of zero-shot learning for facilitating new tasks, which has also proven to be useful when predicting categories not shown in the training dataset. Recently, learning for classifying novel images with auxiliary knowledge (*e.g.*, leveraging textual relationship in a large corpus) has been shown to be powerful for zero-shot tasks. By exploring the correlations among semantic concepts, the label of any unseen data can be reasonably inferred. Socher *et al.* proposed a cross-modal transfer method for zero-shot learning [24], which uses word embedding to detect unseen classes. The work in [5] also adopts the similar scheme as [24]; but the difference is that a new language model and a different loss function is employed to connect two modalities. However, all above methods are only evaluated in classification or recognition scenario. To our best knowledge, our proposed ZSH is one of the first works that focus on handling the large-scale high-dimensional visual indexing problem in zero-shot scenario, *i.e.*, hashing novel images that are not observed by existing learning systems. By adopting a natural language model [9], we can precisely capture the correlations between different concepts, and thus hash unseen images into correct spots in Hamming space.

## 2.2 Hashing

In this part, we briefly review fast search with binary codes using hashing techniques. Similarity search is a challenge of pursuing data points of smallest distance in a large scale database. One of the easiest hashing schemes is dubbed Local Sensitive Hashing [6], which designs hashing function with no prior knowledge of the data distribution. However, such hashing methods require significantly large code length to achieve acceptable performance, which may cause unnecessary overheads in a database. To address this problem, learning-based hashing comes as a trend. Unsupervised hashing methods mine the statistic distributional information in the data, generating an optimized hashing function to preserve certain data properties in the original space. Classical algorithms (*e.g.*, Spectral Hashing (SH) [27]) learns binary codes which preserve local structural information in the data. Shen *et al.* proposed Inductive Manifold Hashing (IMH) [23], which adopts manifold learning techniques to better model the intrinsic structure embedded in the feature space. In [7], Gong *et al.* devised an effective approach, termed Iterative Quantization, focusing on minimizing quantization error during unsupervised training. Note that a

real-world database is commonly described by multiple modalities/resources, such as visual features or textual information. In [36, 29], they utilized information of at least two different resources to achieve promising performance in cross-media search. Since the unsupervised way is guided with little human-level knowledge, supervised hashing has been proposed to use guiding information to learn binary codes. Hashing techniques in this category have been emerging continuously in recent years, and representative methods include Kernel Supervise Hashing (KSH) [16], Minimal Loss Hashing (MLH) [18], Supervise Discrete Hashing (SDH) [22], Latent Factor Hashing (LFH) [35] and Column Sampling Based Discrete Supervised Hashing (COSDISH) [11], etc. Recently, with the success of deep learning, hashing using CNN has also shown the promising performance [31].

Admittedly, hashing algorithms have provided a possible solution for tackling the “curse of dimensionality” problem. When no training samples of certain “unseen” categories are provided, all above hashing methods may easily fail to generalize to the new data, limiting the hashing capability in the “seen” area where each category possesses some training images. Besides, as the database evolves all the time, re-training hashing function frequently is expensive, thereby leading to impractical usage in large dynamic real-world databases. Based on the above analysis, a hashing method that can perform well on unseen data draws a strong need.

## 3. ZERO-SHOT HASHING

In this section, we elaborate our zero-shot hashing (ZSH). We firstly present a formal definition of hashing in zero-shot scenario, and then depict the details of ZSH, including a brief introduction of overall framework, transferring supervised knowledge, semantic alignment as well as hashing model. Finally, we introduce the optimization process and algorithm analysis.

### 3.1 Problem Definition

Suppose we are given  $n$  training images  $X = [x_1, x_2, \dots, x_n] \in \mathbb{R}^{d \times n}$  labeled with a seen concept set  $\mathcal{C}$ , where  $x_i \in \mathbb{R}^{d \times 1}$ ,  $i = 1, 2, \dots, n$  and  $d$  is the dimensionality of visual feature space. Denote  $Y = [y_1, y_2, \dots, y_n] \in \{0, 1\}^{c \times n}$  is the binary label matrix, where  $y_i \in \{0, 1\}^{c \times 1}$  is the label vector of the  $i$ -th sample  $x_i$  and  $c$  is the number of seen classes in  $\mathcal{C}$ . Different from conventional supervised hashing scenario, where both testing data and training data are associated with the same concept set, *i.e.*,  $\mathcal{C}$ , we intend to cope with the situation where testing data and training data share no



common concepts. In other words, testing data (denoted as  $X^{(u)}$ ) belongs to an “unseen” category set  $\mathcal{C}^{(u)}$ , *i.e.*,  $\mathcal{C}^{(u)} \cap \mathcal{C} = \emptyset$ . Using only the training images  $X$  where no training samples of the “unseen” categories in  $\mathcal{C}^{(u)}$  are available, we aim to learn a hash function  $f: \mathbb{R}^{d \times 1} \mapsto \{-1, 1\}^{l \times 1}$ , which can map images belonging to both  $\mathcal{C}^{(u)}$  and  $\mathcal{C}$  from original visual feature space to  $l$ -bit binary codes. The learned hash function  $f$  not only guarantees that the binary codes of semantically relevant objects have short Hamming distances, but also generalizes well to the testing data belonging to the unseen categories.

The flowchart of our overall framework is illustrated in Figure 2. As we can see, there are two stages: the offline phase and the online phase. In the offline phase, suppose only images of a limited number of categories are visible to our system. We firstly extract visual features of the training images through a convolutional neural network. At the same time, we use an off-the-shelf NLP model to transform seen labels into a semantic-rich embedding space, where each label is represented by a real-valued vector. With the embedded semantics, the relationships among both seen and unseen categories can be well captured and characterized. Instead of 0/1-form label vector, ZSH supervises the learning of hash functions with the embedded semantic vectors to transfer supervised knowledge. We further rotate the off-the-shelf embedding space to better align with the low-level visual feature space. Meanwhile, ZSH preserves local structural information and discrete nature of the intermediate binary codes to improve hash functions. Finally, we use the learned hash functions to transform all the images in the database into binary queries for subsequent retrieval. In the online phase, when a new query image of any unseen category arrives, we encode the new image into binary code following the same mapping and retrieve images that are close to this query in the Hamming space.

In general, most of existing supervised hashing algorithms may retrieve relevant results of queries in the seen categories since there are supervised information for understanding the queries. Nevertheless, when the hashing systems have no knowledge of certain unseen classes, query images from these classes will be probably be misunderstood, thereby leading to inaccurate search. One of the main causes is that the supervised information is in the form of 0/1-form label vectors or pair-wise data relationship, which implicitly makes labels independent to each other and omits the inherent correlation among their high-level semantics (*e.g.*, *cat* is as different from *truck* as from *dog*). As illustrated in Figure 3, using independent labels, each object will be mapped to an independent vertex of a hypercube, and the distance between any two categories will be the same. In order to address such disadvantage, we propose to connect label semantics by taking advantage of the superior ability endowed by neural language processing techniques. Specifically, as illustrated in Figure 3, we map independent labels into a word embedding space, where semantic correlations among labels can be quantitatively measured and captured. Therefore, unseen labels can leverage the well-established mapping from its semantically close seen categories. For example, in the embedding space, *cat* and *dog* will be close to each other. Hence even the hashing systems may never observe any *cat* images, they can still gain some useful clues from the supervised knowledge of *dog*. We adopt the language model [9] pre-trained using free Wikipedia text. This model leverages not only local information but also global document context, thereby achieving superior performance over other competitive ap-

Figure 3: An illustration of independent 0/1-form labels v.s. word embedding.

### 3.4 Semantic Alignment

$$\left\| R^T Y - W^T B \right\|_F^2, \quad (1)$$

### 3.5 Hashing Model

<sup>1</sup>In practice, we find that by setting word vector to unit length, retrieval performance can be augmented with no distortion of the cosine similarities. Thus, we empirically normalize word vector to be unit length.

$$\begin{aligned}
\min_{f, W, B, R} & \left\| R^T Y - W^T B \right\|_F^2 + \lambda \|W\|_F^2 + \alpha \|f(X) - B\|_F^2 \\
& + \beta \Omega(f) + \gamma \sum_{i=1}^n \sum_{j=1}^n S_{ij} \|f(x_i) - f(x_j)\|_F^2 \\
\text{s.t. } & B \in \{-1, 1\}^{l \times n} \wedge R^T R = I_c,
\end{aligned} \quad (2)$$

where  $R \in \mathbb{R}^{c \times c}$  is the semantic rotation matrix.  $W \in \mathbb{R}^{l \times c}$  is the mapping matrix.  $B \in \{-1, 1\}^{l \times n}$  denotes the binary codes.  $\|\cdot\|_F$  denotes the Frobenius norm of a matrix.  $\lambda > 0, \alpha > 0, \beta > 0$  and  $\gamma > 0$  are balancing parameters.  $f: \mathbb{R}^{m \times 1} \rightarrow \mathbb{R}^{l \times 1}$  define a hash function from a non-linear embedded feature space to the desired Hamming space:

$$f(x) = P^T \phi(x), \quad (3)$$

where  $f(X) = [f(x_1), f(x_2), \dots, f(x_n)]$ .  $P \in \mathbb{R}^{m \times l}$  is the transformation matrix.  $\Omega(f) = \|P\|_F^2$  is a regularizer  $P$ . Following the successful practice for learning hash functions in [16], we employ kernel mapping to handle the potential problem of linear inseparability:

$$\phi(x) = \left[ \exp\left(-\frac{\|x - a_1\|^2}{\delta}\right), \dots, \exp\left(-\frac{\|x - a_m\|^2}{\delta}\right) \right]^T, \quad (4)$$

where  $\{a_i\}_{i=1}^m$  are  $m$  anchors randomly sampled from  $X$  and  $\delta$  is the bandwidth parameter.

Note that we keep the discrete constraint on the variable  $B$  to prevent information loss of binary codes to the greatest extent. The term  $\gamma \sum_{i=1}^n \sum_{j=1}^n S_{ij} \|f(x_i) - f(x_j)\|_F^2$  in Eq. (2) preserves local structural information of training data, *i.e.*, if two samples are similar in the original feature space, then they are enforced to share similar binary codes in the Hamming space. Here,  $S_{ij}$  is the similarity of  $x_i$  and  $x_j$  in the original visual feature space.

In the next part, we introduce an efficient algorithm to optimize our zero-shot hashing model.

### 3.6 Optimization

We first rewrite the model in matrix form as follows:

$$\begin{aligned}
\min_{P, W, B, R} & \left\| R^T Y - W^T B \right\|_F^2 + \lambda \|W\|_F^2 + \alpha \left\| P^T \phi(X) - B \right\|_F^2 \\
& + \beta \|P\|_F^2 + \gamma \text{Tr}(P^T \phi(X) L \phi(X)^T P) \\
\text{s.t. } & B \in \{-1, 1\}^{l \times n} \wedge R^T R = I_c,
\end{aligned} \quad (5)$$

where  $\phi(X) = [\phi(x_1), \phi(x_2), \dots, \phi(x_n)]$ .  $\text{Tr}(\cdot)$  is the trace of a matrix. The Laplacian matrix  $L$  is computed as:

$$L = D - S, \quad (6)$$

where  $D$  is a diagonal matrix with its  $i$ -th diagonal element computed as  $D_{ii} = \sum_{j=1}^n S_{ij}$  and  $S = \{S_{ij}\}_{i,j=1,2,\dots,n}$  is the similarity matrix of  $X$ .

Next, we present an algorithm to optimize the model in Eq. (5) by alternately updating  $P, W, B, R$ .

#### 3.6.1 Update $P$

Fixing all variables except for  $P$ , we get the quadratic problem as:

$$\min_P \alpha \|P^T \phi(X) - B\|_F^2 + \beta \|P\|_F^2 + \gamma \text{Tr}(P^T \phi(X) L \phi(X)^T P). \quad (7)$$

By setting its derivative with respect to  $P$  to 0, we have the following solution:

$$P = \left( \phi(X) \phi(X)^T + \frac{\beta}{\alpha} I_m + \frac{\gamma}{\alpha} \phi(X) L \phi(X)^T \right)^{-1} \phi(X) B^T, \quad (8)$$

where  $I_m$  is an identity matrix of size  $m \times m$ .

#### 3.6.2 Update $B$

In this step, we fix all other variables and learn binary codes  $B$  with discrete constraint. The objective function can be reduced to

$$\begin{aligned}
\min_B & \left\| R^T Y - W^T B \right\|_F^2 + \alpha \left\| P^T \phi(X) - B \right\|_F^2 \\
\text{s.t. } & B \in \{-1, 1\}^{l \times n}.
\end{aligned} \quad (9)$$

The above formulation can be further written as

$$\begin{aligned}
\min_B & \left\| W^T B \right\|_F^2 - 2 \text{Tr}(B^T H) \\
\text{s.t. } & B \in \{-1, 1\}^{l \times n},
\end{aligned} \quad (10)$$

where  $H = W R^T Y + \alpha P^T \phi(X)$ .

Inspired by [30], we apply the *discrete coordinate descent (DCC)* algorithm to solve the problem (10). Let  $B = [q_1^T, q_2^T, \dots, q_l^T]$ ,  $H = [h_1^T, h_2^T, \dots, h_l^T]$  and  $W = [u_1^T, u_2^T, \dots, u_l^T]$ , where  $q_i^T$ ,  $h_i^T$  and  $u_i^T$  are the  $i$ -th row of  $B$ ,  $H$  and  $W$ , respectively. Furthermore, for convenience, we denote

$$\begin{cases} B_{-i} = [q_1^T, \dots, q_{i-1}^T, q_{i+1}^T, \dots, q_l^T], \\ H_{-i} = [h_1^T, \dots, h_{i-1}^T, h_{i+1}^T, \dots, h_l^T], \\ W_{-i} = [u_1^T, \dots, u_{i-1}^T, u_{i+1}^T, \dots, u_l^T]. \end{cases} \quad (11)$$

Then, we can have

$$\begin{aligned}
\left\| W^T B \right\|_F^2 &= \text{Tr}(B^T W W^T B) \\
&= \text{const} + \|q_i u_i^T\|_F^2 + 2 u_i^T W_{-i}^T B_{-i} q_i \\
&= \text{const} + 2 u_i^T W_{-i}^T B_{-i} q_i.
\end{aligned} \quad (12)$$

Here,  $\|q_i u_i^T\|_F^2 = \text{Tr}(u_i q_i^T q_i u_i^T) = \text{const}$ . Following the same rule, we also have the following conclusion

$$\text{Tr}(B^T H) = \text{const} + h_i^T q_i. \quad (13)$$

The sub-problem can be transformed to

$$\begin{aligned}
\min_{q_i} & (u_i^T W_{-i}^T B_{-i} - h_i) q_i \\
\text{s.t. } & h_i \in \{-1, 1\}^{n \times 1}.
\end{aligned} \quad (14)$$

The optimal solution of above equation is

$$q_i = \text{sgn}(h_i - B_{-i}^T W_{-i} u_i), \quad (15)$$

where  $\text{sgn}(\cdot)$  is the sign function. We can see that each bit of the desired binary code  $B$  can be learned based on other  $l - 1$  bits. Thus, we can use cyclic coordinate descent approach to generate the optimal codes until the entire procedure converges.

#### 3.6.3 Update $R$

With  $B, W, P$  fixed, we then have

$$\begin{aligned}
\min_R & \left\| R^T Y - W^T B \right\|_F^2 \\
\text{s.t. } & R^T R = I_c,
\end{aligned} \quad (16)$$

which can efficiently solved by the algorithm in [28].

### 3.6.4 Update $W$

By fixing  $P, B, R$ , we arrive at a classic ridge regression problem:

$$\min_W \|R^T Y - W^T B\|_F^2 + \lambda \|W\|_F^2. \quad (17)$$

The above equation has a closed-form solution:

$$W = (BB^T + \lambda I_l)^{-1} B Y^T R, \quad (18)$$

where  $I_l$  is a diagonal matrix of size  $l \times l$ .

By iteratively updating  $P, W, B, R$  until convergence, we can arrive at an optima. The overall algorithm is summarised in Algorithm 1.

---

#### Algorithm 1 Algorithm for optimizing Zero-Shot Hashing

---

**Input:** Training data  $X$  and the embedded label matrix  $Y$ ;

**Output:** Binary codes  $B$ , rotation matrix  $R$ , hash function  $P$  and mapping matrix  $W$ ;

- 1: Randomly initialize  $B, P$  and  $W$ ;
  - 2: Randomly initialize  $R$  to be orthogonal;
  - 3: Map  $X$  to  $\phi(X)$  using  $m$  anchors randomly selected from  $X$ ;
  - 4: Construct Laplacian matrix  $L$ ;
  - 5: **repeat**
  - 6:   Update  $P$  according to Eq. (8);
  - 7:   Update  $B$  iteratively by using the solution of (15);
  - 8:   Update  $R$  by solving the problem in Eq. (16);
  - 9:   Update  $W$  according to Eq. (18);
  - 10: **until** there is no change to  $P, W, B, R$
  - 11: **return**  $P, W, B, R$ ;
- 

## 3.7 Algorithm Analysis

In this section, we analyze the convergence and time complexity of our algorithm.

### 3.7.1 Convergence Study

As shown in Algorithm 1, in each iteration, the updates of all variables make the value of the objective function decreased. We also conducted empirical study on the convergence property using ImageNet [3]. Specifically, we trained our zero-shot hashing model with 30,000 seen images randomly sampled from the ImageNet dataset, with label embedding as supervised information. We selected 1,000 anchors and set the code length to 64. As Figure 4 shows, our algorithm starts with cost function value at roughly 30,000, but descends dramatically within only 10 iterations, and reaches a stable local minima at the 20-th iteration. This phenomenon clearly indicates the efficiency of our algorithm.




Figure 4: Convergence study on ImageNet.

### 3.7.2 Computational Complexity

In each iteration (line 6-9), the time cost is analyzed as follows. The computation of  $P$  in Eq. (8) is  $O(m^2 n + nml + m^3)$ . The DCC algorithm for updating  $B$  costs  $O(cl^2 + l^2 n)$ . As to the optimization of the sub-problem in Eq. (16), the time cost is  $O(c^3)$ . Finally, the computational cost of updating  $W$  is  $O(l^2 n + lnc + lc^2 + l^3)$ . Given that  $m \ll n, l \ll n, c \ll n$  and our algorithm converges within a few iterations (less than 10), the overall time cost of our algorithm is  $O(n)$ . It is worth noting that the dominant operation of our algorithm is matrix multiplication, which can be greatly speeded up by using parallel and/or distributed algorithms.

## 4. EXPERIMENT

### 4.1 Experimental Settings

In our experiments, we employ three real-life image datasets, including CIFAR-10<sup>2</sup>, ImageNet<sup>3</sup> and MIRFlickr<sup>4</sup>.

**CIFAR-10** consists of 60,000 images which are manually labelled with 10 classes including *airplane, automobile, bird, cat, deer, dog, frog, horse, ship, truck*, with 6,000 samples in each class. The classes are completely mutually exclusive, *i.e.*, no overlap between classes (*e.g.*, automobiles and trucks).

**ImageNet** is an image dataset organized according to the WordNet [17] hierarchy. The subset of ImageNet for the Large Scale Visual Recognition Challenge 2012 (ILSVRC2012) is used for our experiments, consisting of over 1.2 million Web images, manually labeled with 1,000 object categories.

**MIRFlickr** comprises 25,000 images collected from the social photography site Flickr through its public API. Firstly introduced in 2008, this dataset is widely used in multimedia research. MIRFlickr is a multi-label dataset with every image associating with 24 popular tags such as *sky, river*, etc.

For all image data, we adopted the winning model for the 1000-class ImageNet Large Scale Visual Recognition Challenge 2012 [13] to extract the fully connected layer fc-7 as visual feature.

Various metrics are employed for measuring performance of different evaluation tasks. For image retrieval, we used the two traditional metrics *i.e.*, Precision and Mean Average Precision (MAP). MAP focuses on the ranking of retrieval results and we reported the results over the top 5,000 retrieved samples. Precision mainly concentrates on the retrieval accuracy and we reported the results with Hamming radius  $r \leq 2$ .

We compared our proposed ZSH with four state-of-the-art supervised hashing approaches, including COSDISH [11], SDH [22], KSH [16] and LFH [35]. For all anchor-based algorithms, we randomly sampled 1,000 anchors from the training dataset. Furthermore, we compared to one of the most representative unsupervised hashing method, *i.e.*, Inductive Hashing on Manifolds (IMH) [23].

For all comparing approaches, we followed their suggested parameter settings. For ZSH, we empirically set  $\alpha$  to  $10^{-5}$  and  $\gamma$  to  $10^{-6}$ . For regularization parameters  $\lambda$  and  $\beta$ , we set them to  $10^{-2}$  and  $10^{-4}$ , respectively. The number of iterations is set to 10. We define the similarity matrix  $S$  as follows:

$$S_{ij} = \begin{cases} \exp(-\frac{\|x_i - x_j\|}{2\sigma^2}), & \text{if } x_i \in \mathcal{N}_k(x_j) \text{ or } x_i \in \mathcal{N}_k(x_j) \\ 0, & \text{otherwise,} \end{cases}$$

<sup>2</sup><https://www.cs.toronto.edu/~kriz/cifar.html>

<sup>3</sup><http://image-net.org/>

<sup>4</sup><http://press.liacs.nl/mirflickr/>

where  $\mathcal{N}_k(\cdot)$  is the function of searching  $k$  nearest neighbors. In our experiment, we set  $\sigma = 1$ .

## 4.2 Results on CIFAR-10

### 4.2.1 Overall Comparison of Zero-Shot Image Retrieval

To evaluate the efficacy of retrieving images in unseen categories, we split CIFAR-10 into a “seen” training set and an “unseen” testing set. In particular, we selected *truck* as unseen testing category and left the rest 9 categories as seen training set. For all comparing algorithms, we randomly sampled 10,000 images for learning hash functions. For testing purpose, we randomly selected 1,000 images from the unseen category as query images, and the remaining 5,000 test images together with the 54,000 images of seen categories are combined to form the retrieval database.

The performance of all comparing approaches *w.r.t.* different codes lengths (*i.e.*, {16, 32, 64, 96, 128}) is illustrated in Figure 5. As we can see, the proposed ZSH outperforms all the other hashing algorithms in terms of MAP at all code lengths. As to Precision, ZSH still shows superior image retrieval performance in most cases. The underlying principle is that our method not only utilizes inherent semantic relationship among labels to transfer supervised knowledge, but also preserves discrete and structural properties of data in the learning of hash codes and hash functions. An interesting observation is that the performance of IMH, which is an unsupervised method, gains competitive even better retrieval results in terms of Precision as compared to some supervised methods such as KSH, SDH. While unsupervised methods encode images solely with the distributional properties in the feature space, the supervised ones may be misled by independent semantic labels in the learning processing.

Besides, MAP increases rapidly for all methods when code length varies from 16 to 64, and then reaches a slow-growth stage from 64 bits to 128 bits. When code length is short, more codes are required to guarantee the descriptive and discriminative power. However, after encoding space is large enough (*e.g.* 64 bits), with the expression ability saturated, providing more bits cannot significantly improve the performance. As to Precision, hashing performance significantly deteriorates as code length is larger than 64. Recall that our searching radius is empirically set to 2, forming a hyper-ball of radius 2 in Hamming space. When the code length increases from 16 to 64, significant improvement in retrieval ability counteracts the searching difficulty. However, as Hamming space becomes larger, searching difficulty grows linearly, thereby degrading the Precision performance. Therefore, as a trade-off between efficiency and effectiveness, an eclectic code length should be chosen.




Figure 5: Performance (MAP and Precision) of different comparing methods on zero-shot image retrieval over CIFAR-10 dataset.

### 4.2.2 Effect of Different Unseen Category

In this experiment, we aim to evaluate the performance of zero-shot image retrieval on different unseen categories. The experimen-

tal settings are the same as that in the previous subsection. Figure 6 illustrates the MAP and Precision performance of ZSH using each individual label as unseen testing data.




Figure 6: Performance (Precision and MAP) of zero-shot image retrieval for each individual unseen category on CIFAR-10 dataset.

We can observe that zero-shot image retrieval performance varies from one class to another, reaching peak at *bird* and bottom at *automobile*. Intuitively, if an unseen class is semantically closer to other seen categories, more relevant supervised knowledge can be transferred from word embedding space for boosting the retrieval performance. To dig deeper about the reason behind the fluctuation of performance on different unseen objects, we computed the average cosine similarity between each unseen category and other seen categories, and list the corresponding MAP in Table 1.

Category	Average Cosine Similarity	MAP
airplane	0.2191	0.2791
automobile	0.1567	0.2603
bird	0.3565	0.3751
cat	0.3661	0.3621
deer	0.2981	0.2912
dog	0.3826	0.3467
frog	0.3485	0.2991
horse	0.3015	0.3125
ship	0.1663	0.2987
truck	0.3358	0.3120

Table 1: Average cosine similarity of each category and all other categories, together the corresponding MAP performance.

We observe that the MAP performance is positively related to the average cosine similarity. For instance, those of larger cosine similarity (*e.g.*, *dog*, *cat*) performs relatively well, while those of smaller similarity (*e.g.*, *airplane*, *automobile*) gain relatively poor performance. This observation implies that in order to achieve satisfactory retrieval results, unseen classes should have sufficient correlation with seen ones.

As we can see in Figure 6, we also compare the effects of embedded labels and binary labels. The performance of embedded labels is obviously better than that of binary labels. The underlying reason is that the embedding space can help to capture the relationship between seen and unseen categories for transferring supervised knowledge. In contrast, binary labels neglect semantic correlations, thereby leading to irregular fluctuations of retrieval performance.

### 4.2.3 Effect of Seen Category Ratio

In this experiment, we evaluate the performance of our proposed ZSH *w.r.t.* different numbers of seen categories. Specifically, we varied the ratio of seen categories in the training set from 0.1 to 0.9. For each ratio, we randomly sampled 10,000 images from the






Figure 7: A demonstration of the zero-shot image retrieval exemplars using different comparing hashing algorithms on CIFAR-10, where top 15 retrieval images are reported. Pictures with green bounding box indicate the correct results while those with red outlines indicate failure results. As can be seen, our proposed ZSH method returns the largest number of correct retrieved results with query from unseen category, followed by COSDISH which returns four correct samples.

seen categories for training. Further, we randomly selected 1,000 images from the unseen set as queries to search in the remaining 59,000 images. Note that when the ratio of seen categories decreases to 0.1, we use all 6,000 datums of that class as training set.




Figure 8: Effects of different ratios of seen categories on CIFAR-10 dataset.

We report the experimental results in Figure 8, from which we have the following observations: (1) The performance of both MAP and Precision boosts as the ratio of the seen categories grows; (2) As the ratio increases from 0.1 to 0.3, we see a dramatic leap of the retrieval performance, followed by a relatively slight performance improvement from 0.3 to 0.9. We analyze that by observing more “seen” categories, we have higher possibility to find relevant supervision for the unseen class, which guides to learn better intermediate hash codes, thereby simultaneously improving the quality of hash functions.

#### 4.2.4 Effect of Training Size

This part of experiment mainly focuses on evaluating the effect of training size on the searching quality of ZSH. For simplicity, we chose Precision as evaluation metric and selected *truck* as the unseen object and varied the size of training data in the range of {1000, 2000, ..., 10000, 20000, ..., 50000}.

The experimental results are demonstrated in Figure 9. As we can see, when the size increases from 1,000 to 10,000, we observe a rapid rise of the Precision performance. Nonetheless, when fed with more training data, ZSH does not gain noticeable performance boost. For the balance of training efficiency and effectiveness, in the rest experiments, we consistently set the training size to 10,000.

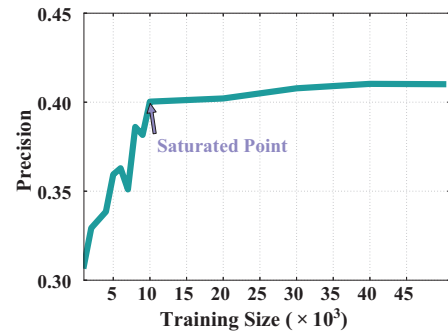


Figure 9: Effects of training size on Precision performance over CIFAR-10 dataset.

## 4.3 Results on ImageNet

### 4.3.1 Overall Comparison of Zero-Shot Image Retrieval

In this part, we evaluate our proposed ZSH on zero-shot image retrieval as compared to other state-of-the-art methods using the Large Scale Visual Recognition Challenge 2012 (ILSVRC2012) dataset. Recall that the ILSVRC2012 dataset contains more than 1.2 million images tagged with 1,000 synsets without any overlap. For evaluation purpose, we randomly chose 100 categories which have corresponding word embedding learned from Wikipedia text corpus. This gives us a set of roughly 130,000 images. We split the data into a training set (90 seen categories) and a testing set (10 unseen categories). For all comparing algorithms, we randomly selected 10,000 images of seen categories for training. As to image queries, we randomly sampled 1,000 images from the unseen categories. We used the learned hash function to encode all the remaining images to form the retrieval database.

The performance of our proposed ZSH and other four state-of-the-art supervised hashing methods with different code lengths are reported in Figure 10. As we can see, ZSH consistently outperforms all other competitors in most cases. As code length varies from 16 to 128, we can observe the similar variation tendency of performance on ImageNet to that on CIFAR-10. This phenomenon again implies that we should choose a trade-off code length to guarantee the retrieval performance.






Figure 10: Performance (MAP and Precision) of different comparing methods on zero-shot image retrieval over ImageNet dataset.

#### 4.3.2 Image Retrieval in Related Categories

In zero-shot image retrieval scenario, we expect that even though we fail to retrieve relevant images of the same category, we can still obtain semantically related images. For instance, if the query image describes a *cat*, we may prefer to retrieve images of *dog* rather than images of *car*. Our proposed ZSH utilizes semantic embedding to set up connections between semantically similar labels in the embedded space. In this way, the supervised knowledge of seen categories can be transferred into the learning of hash functions, which can effectively encode images of unseen categories.




Figure 11: Comparison of ZSH and other hashing approaches on the capability of retrieving semantically similar images from related categories.

Since we need to search more related categories, all remaining images of both seen and unseen categories are used to form retrieval database. All the other settings are the same as that in Section 4.3.1. In order to evaluate the performance of retrieving related categories, we use two modified metrics, named  $\text{MAP}_{\text{related}}$  and  $\text{Precision}_{\text{related}}$ , which are defined as

$$\text{MAP}_{\text{related}} = \sum_{i=1}^K \frac{\text{AP}_{\text{related}@i}}{K}, \quad (19)$$

$$\text{Precision}_{\text{related}} = \frac{n_{\text{related}}}{n_{\text{retrieved}}}, \quad (20)$$

where  $\text{MAP}_{\text{related}}$  is calculated based on the top  $K$  retrieved results,  $\text{AP}_{\text{related}@i}$  is the average precision based on the related results, calculated by

$$\text{AP}_{\text{related}@i} = \frac{n_{\text{related}}^{(i)}}{i}, \quad (21)$$

where  $n_{\text{related}}^{(i)}$  is the number of related images in top  $i$  retrieved results.  $n_{\text{related}}$  and  $n_{\text{retrieved}}$  are the related retrieval under Hamming radius 2 and total examples retrieved under Hamming radius 2, respectively. Using WordNet [17], which is a lexical database for the English language, we define query  $A$  and retrieved object  $B$  are related if: 1)  $A$  and  $B$  are not of the same category; and 2)  $A$  can reach  $B$  on WordNet within 5 hops.

In practice, we set  $K = 5,000$ . Figure 11 shows the experimental results. We can see that in terms of  $\text{MAP}_{\text{related}}$ , our method always outperforms other methods at every code length. When we look at  $\text{Precision}_{\text{related}}$ , our proposed ZSH achieves 0.3262, 0.2636, 0.2129 at 32 bits, 64 bits and 96 bits, respectively, which significantly outperforms the second best method. This observation indicates that ZSH is capable of detecting the semantically similar images from the most related categories.

#### 4.4 Results on MIRFlickr




Figure 12: Performance (MAP and Precision) of different comparing methods on zero-shot image retrieval over MIRFlickr dataset.

In real-world images, especially in user-generated photos, there often exists multiple tags in each individual picture. To further examine the practical efficacy of our proposed ZSH, we conducted an extra experiment on a real-life multi-label dataset, *i.e.*, MIRFlickr, which contains 25,000 images downloaded from the social photography site Flickr. Each image is associated with 24 tags. In multi-label image dataset, different categories share overlapping images, which makes it difficult to divide the dataset into training set and testing set. Therefore, we employed ImageNet as an auxiliary dataset to train our hash functions and evaluated the zero-shot image retrieval performance on MIRFlickr. Specifically, from the ILSVRC2012 dataset we selected 100 categories which does not overlap with the 24 tags in MIRFlickr. For fair comparison, all hashing approaches used 10,000 randomly sampled images for training. After the hash function was learned, we directly applied them to transform the MIRFlickr images into binary codes. We then sampled 1,000 datums as query images and searched relevant results in the remaining 24,000 images. We regarded the retrieval images sharing at least two tags with the query as the true neighbors, and computed MAP on the top 5,000 retrieved results and Precision under Hamming distance 2. Figure 12 illustrates the results of our ZSH and other comparing algorithms on MIRFlickr. In the left sub-figure, we can see that with different code lengths, our ZSH can consistently achieve the best MAP performance among all the comparing algorithms. As the code length increases, the MAP performance of each algorithm keeps increasing, reaching 0.2488 at 128 bits, which outperforms the second best hashing method COSDISH by 19% at the same length. In terms of Precision, ZSH exceeds all other methods in most cases. Similar to that of CIFAR-10 and ImageNet, we can see a variation pattern with an increasing trend from 16 to 64 and a performance drop from 64 to 128. The promising performance on MIRFlickr demonstrates the potential of ZSH in indexing and searching real-life image data.

#### 5. CONCLUSION

With the explosion of newly-emerging concepts and multimedia data on the Web, it is impossible to supply existing supervised hashing methods with sufficient labeled data in time. In this paper, we studied the problem of how to map images of unseen categories using hash functions learned from limited seen classes. We pro-

posed a novel hashing scheme, termed *zero-shot hashing* (ZSH), which is capable of transmitting supervised knowledge from seen categories to unseen categories. Independent 0/1-form labels were projected into an off-the-shelf embedding space with abundant semantics, where label semantic correlations can be fully characterized and quantified. Considering the issues of domain difference and semantic shift, we further narrowed down the gap between binary codes and high-level semantics by a semantic alignment operation. Specifically, we intentionally rotated the embedding space to adjust the supervised knowledge more suitable for learning high-quality hash codes. Besides, we also preserved local structural property and discrete nature of hash codes in the ZSH model. An efficient algorithm was designed to optimize the model in an alternating manner and the empirical study showed the convergency and efficiency. We evaluated our proposed ZSH hashing approach on three real-world image datasets, including CIFAR-10, ImageNet and MIRFlickr. The experimental results demonstrated the superiority of ZSH as compared to several state-of-the-art hashing approaches on zero-shot image retrieval task.

In the future, we plan to enhance the exploration of label semantic correlations by integrating knowledge from multiple sources, including textual corpus and visual clues. We expect this will compensate the incomplete representation of each individual modality, thereby solving the problem of domain difference and semantic shift fundamentally.

## 6. ACKNOWLEDGMENTS

This work was supported in part by the National Natural Science Foundation of China under Project 61572108 and Project 61502081, the National Thousand-Young-Talents Program of China, and the Fundamental Research Funds for the Central Universities under Project ZYGX2014Z007 and Project ZYGX2015J055.

## 7. REFERENCES

- [1] L. Cao, Z. Li, Y. Mu, and S.-F. Chang. Submodular video hashing: a unified framework towards video pooling and indexing. In *ACM Multimedia*, 2012.
- [2] G. Castañón, Y. Chen, Z. Zhang, and V. Saligrama. Efficient activity retrieval through semantic graph queries. In *ACM Multimedia*, 2015.
- [3] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *IEEE CVPR*, 2009.
- [4] A. Farhadi, I. Endres, D. Hoiem, and D. Forsyth. Describing objects by their attributes. In *IEEE CVPR*, 2009.
- [5] A. Frome, G. S. Corrado, J. Shlens, S. Bengio, J. Dean, T. Mikolov, et al. Devise: A deep visual-semantic embedding model. In *NIPS*, 2013.
- [6] A. Gionis, P. Indyk, R. Motwani, et al. Similarity search in high dimensions via hashing. In *VLDB*, 1999.
- [7] Y. Gong and S. Lazebnik. Iterative quantization: A procrustean approach to learning binary codes. In *IEEE CVPR*, 2011.
- [8] Y. Hu, Z. Jin, H. Ren, D. Cai, and X. He. Iterative multi-view hashing for cross media indexing. In *ACM Multimedia*, 2014.
- [9] E. H. Huang, R. Socher, C. D. Manning, and A. Y. Ng. Improving word representations via global context and multiple word prototypes. In *ACL*, 2012.
- [10] D. Jayaraman and K. Grauman. Zero-shot recognition with unreliable attributes. In *NIPS*, 2014.
- [11] W.-C. Kang, W.-J. Li, and Z.-H. Zhou. Column sampling based discrete supervised hashing. In *AAAI*, 2016.
- [12] P. Kankuekul, A. Kawewong, S. Tangruamsub, and O. Hasegawa. Online incremental attribute-based zero-shot learning. In *IEEE CVPR*, 2012.
- [13] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012.
- [14] C. H. Lampert, H. Nickisch, and S. Harmeling. Attribute-based classification for zero-shot visual object categorization. *IEEE TPAMI*, 2014.
- [15] H. Larochelle, D. Erhan, and Y. Bengio. Zero-data learning of new tasks. In *AAAI*, 2008.
- [16] W. Liu, J. Wang, R. Ji, Y.-G. Jiang, and S.-F. Chang. Supervised hashing with kernels. In *IEEE CVPR*, 2012.
- [17] G. A. Miller. Wordnet: a lexical database for english. *Communications of the ACM*, 1995.
- [18] M. Norouzi and D. J. Fleet. Minimal loss hashing for compact binary codes. *ICML*, 2011.
- [19] D. Parikh and K. Grauman. Relative attributes. In *IEEE ICCV*, 2011.
- [20] S. Petrović, M. Osborne, and V. Lavrenko. Streaming first story detection with application to twitter. In *ACL*, 2010.
- [21] B. Romera-Paredes and P. Torr. An embarrassingly simple approach to zero-shot learning. In *ICML*, 2015.
- [22] F. Shen, C. Shen, W. Liu, and H. Tao Shen. Supervised discrete hashing. In *IEEE CVPR*, 2015.
- [23] F. Shen, C. Shen, Q. Shi, A. Hengel, and Z. Tang. Inductive hashing on manifolds. In *IEEE CVPR*, 2013.
- [24] R. Socher, M. Ganjoo, C. D. Manning, and A. Ng. Zero-shot learning through cross-modal transfer. In *NIPS*, 2013.
- [25] J. Turian, L. Ratinov, and Y. Bengio. Word representations: a simple and general method for semi-supervised learning. In *ACL*, 2010.
- [26] J. Wang, J. Wang, N. Yu, and S. Li. Order preserving hashing for approximate nearest neighbor search. In *ACM Multimedia*, 2013.
- [27] Y. Weiss, A. Torralba, and R. Fergus. Spectral hashing. In *NIPS*, 2009.
- [28] Z. Wen and W. Yin. A feasible method for optimization with orthogonality constraints. *Mathematical Programming*, 2013.
- [29] F. Wu, Z. Yu, Y. Yang, S. Tang, Y. Zhang, and Y. Zhuang. Sparse multi-modal hashing. *IEEE TMM*, 2014.
- [30] T. T. Wu and K. Lange. Coordinate descent algorithms for lasso penalized regression. *The Annals of Applied Statistics*, 2008.
- [31] R. Xia, Y. Pan, H. Lai, C. Liu, and S. Yan. Supervised hashing for image retrieval via image representation learning. In *AAAI*, 2014.
- [32] Y. Yang, F. Shen, H. T. Shen, H. Li, and X. Li. Robust discrete spectral hashing for large-scale image semantic indexing. *IEEE TBD*, 2015.
- [33] Y. Yang, H. Zhang, M. Zhang, F. Shen, and X. Li. Visual coding in a semantic hierarchy. In *ACM Multimedia*, 2015.
- [34] H. Zhang, N. Zhao, X. Shang, H. Luan, and T. seng Chua. Discrete image hashing using large weakly annotated photo collections. In *AAAI*, 2016.
- [35] P. Zhang, W. Zhang, W.-J. Li, and M. Guo. Supervised hashing with latent factor models. In *ACM SIGIR*, 2014.
- [36] X. Zhu, Z. Huang, H. T. Shen, and X. Zhao. Linear cross-modal hashing for efficient multimedia search. In *ACM Multimedia*, 2013.