

Transductive Zero-Shot Hashing for Multilabel Image Retrieval

Qin Zou^{id}, Senior Member, IEEE, Ling Cao^{id}, Zheng Zhang, Long Chen^{id}, Senior Member, IEEE, and Song Wang^{id}, Senior Member, IEEE

Abstract—Hash coding has been widely used in the approximate nearest neighbor search for large-scale image retrieval. Given semantic annotations such as class labels and pairwise similarities of the training data, hashing methods can learn and generate effective and compact binary codes. While some newly introduced images may contain undefined semantic labels, which we call unseen images, zero-shot hashing (ZSH) techniques have been studied for retrieval. However, existing ZSH methods mainly focus on the retrieval of single-label images and cannot handle multilabel ones. In this article, for the first time, a novel transductive ZSH method is proposed for multilabel unseen image retrieval. In order to predict the labels of the unseen/target data, a visual-semantic bridge is built via instance-concept coherence ranking on the seen/source data. Then, pairwise similarity loss and focal quantization loss are constructed for training a hashing model using both the seen/source and unseen/target data. Extensive evaluations on three popular multilabel data sets demonstrate that the proposed hashing method achieves significantly better results than the comparison methods.

Index Terms—Deep hashing, image retrieval, multilabel image, transductive learning, zero-shot learning.

I. INTRODUCTION

HASHING methods can transform high-dimensional data into compact binary codes while preserving the similarity between them. With high computing efficiency and low storage cost, hashing methods have been widely used for large-scale image retrieval. A number of hashing methods have been proposed in the past decade [1]–[4].

Existing hashing methods can be roughly divided into two categories: supervised [5]–[7] and unsupervised [8]–[12]. The supervised hashing methods incorporate human-annotated information, e.g., semantic labels and pairwise similarities,

into the learning process to find an optimal hash function, whereas the unsupervised methods often learn hash functions by exploiting the intrinsic manifold structure of the unlabeled data. Generally, supervised methods can obtain much higher performance than the unsupervised ones.

In recent years, inspired by the remarkable success of deep neural networks in a broad range of computer-vision applications such as image classification [13]–[15], object detection [16]–[19], semantic segmentation [20], [21], many supervised hashing methods turn to use deep neural networks for hash-code learning [22]–[27]. These deep hashing methods have greatly advanced the retrieval performance on several popular benchmark data sets.

However, with the rapid emerging of new goods and new activities, images may contain concepts (or semantic labels) that are undefined before. For example, various commercial products with different shapes and appearances are released to the market every day, and new sports with novel playing scenes are invented from time to time over the world. The images containing these new products or playing scenes are “unseen” compared with the “seen” images holding predefined labels and are supposed to be annotated with new labels for training the supervised learners. Consequently, the supervised hashing methods may face tremendous challenges due to the lack of timely and reliable annotation of ground-truth labels for the unseen images.

Zero-shot learning (ZSL) [28] is a technique that can potentially solve or alleviate this problem. ZSL bridges the semantic gap between “seen” and “unseen” categories by transferring supervised knowledge from other modalities or domains, e.g., class-attribute descriptors and word vectors. For instance, the word embeddings of similar words that locate closely in the embedding space can capture the distributional similarity in the textual domain [29] based on a large-scale text corpus, such as Wikipedia. Thus, such knowledge transfer can be used to capture the relationship between seen and unseen concepts and can be helpful to handle unseen images in supervised learning.

For image retrieval under the circumstance of unseen images, some zero-shot hashing (ZSH) methods [30]–[34] have also been proposed. Nevertheless, these methods focus on single-label image retrieval, in which a one-to-one visual-semantic representation pair can satisfy the training of a hashing model, whereas in more complicated scenarios, an image often contains multiple object classes and hence more complex semantics and their relationships. How to

Manuscript received November 18, 2019; revised May 31, 2020 and October 4, 2020; accepted November 25, 2020. Date of publication December 25, 2020; date of current version April 5, 2022. This work was supported in part by the National Natural Science Foundation of China under Grant 61872277, Grant 61672376, and Grant U1803264; in part by the Key R&D Program of Hubei Province under Grant 2020BAB018; and in part by the National Key Research Development Program of China under Grant 2019QY(Y)0206. (Corresponding authors: Qin Zou; Long Chen.)

Qin Zou, Ling Cao, and Zheng Zhang are with the School of Computer Science, Wuhan University, Wuhan 430072, China (e-mail: qzou@whu.edu.cn; lingcao@whu.edu.cn; zhengzhang@whu.edu.cn).

Long Chen is with the School of Data and Computer Science, Sun Yat-sen University, Guangzhou 518001, China (e-mail: chenl46@mail.sysu.edu.cn).

Song Wang is with the Department of Computer Science and Engineering, University of South Carolina, Columbia, SC 29201 USA, and also with the School of Computer Science and Technology, Tianjin University, Tianjin 300072, China (e-mail: songwang@cec.sc.edu).

Color versions of one or more figures in this article are available at <https://doi.org/10.1109/TNNLS.2020.3043298>.

Digital Object Identifier 10.1109/TNNLS.2020.3043298

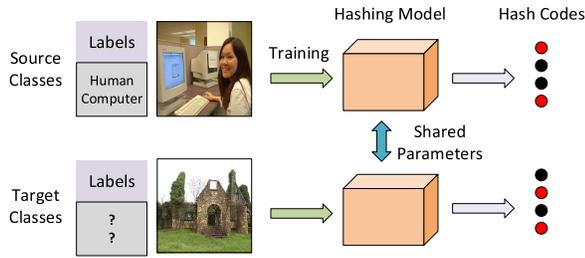


Fig. 1. Illustration of the transductive ZSH. In the learning procedure, both the source and target data are used for training the hashing model. The categories (labels) of the target data are unknown to the learning system.

represent the complex visual-semantic relationships for multilabel images, in a unified framework, is a difficult problem. To the best of our knowledge, there does not exist any work on multilabel ZSH.

Another important but easily ignored problem is that since the underlying distributions of the source data and the target data are different, learning a hash function from a naive knowledge transfer on the source domain without making it adaptive to the target domain may lead to severe domain-shift problems. The knowledge transfer across different domains has become a very important issue in many computer-vision problems [35]–[37], which also needs to be investigated in ZSH.

Considering the problems discussed above, we propose a novel transductive zero-shot hashing method (T-MLZSH) for multilabel image retrieval. Both the labeled source and the unlabeled target data are used in the training phase, as shown in Fig. 1. The labeled source data are used to learn the relationship between visual images and semantic embeddings, and the unlabeled data of target classes are used to alleviate the domain-shift problem. More specifically, we first study a visual-semantic bridge via instance-concept coherence ranking on the source data. In instance-concept coherence ranking, a relatedness score for an image instance and a semantic concept is calculated for each image in the source data, where the score of an instance with a relevant label is larger than that of the same instance with an irrelevant label. Then, we can generate predicted labels for target data and use these predicted labels as supervised information to guide the learning of hashing models. Moreover, we propose a focal quantization loss for fast and efficient hashing learning.

The contributions of this work lie in threefold.

- 1) A T-MLZSH is proposed to solve the domain-shift problem in multilabel image retrieval. To the best of our knowledge, it is the first work studying the ZSH for multilabel image retrieval.
- 2) An instance-concept coherence ranking algorithm is proposed for visual-semantic mapping, which can be used to predict the labels for unseen target data and hence improve the performance of zero-shot deep hashing.
- 3) The proposed method obtains very promising results on three popular multilabel data sets, which constructs the benchmark for zero-shot multilabel image retrieval and paves the road for new research in this field.

The rest of this article is organized as follows. Section II briefly reviews the related work. Section III describes the neural network architecture for zero-shot image retrieval. Section IV demonstrates the effectiveness of the proposed method by experiments. Finally, Section V concludes this article.

II. RELATED WORK

A. Zero-Shot Learning

ZSL refers to training on the seen labeled data with seen labels and testing on the unseen data with unseen labels, where there is no intersection between the seen label set and the unseen label set. The core of ZSL is to obtain the instance labels for the unseen data. According to the ways of obtaining the instances, ZSL can be divided into three categories [38]–[40]: projection-based methods [41]–[43], instance-borrowing methods [44]–[46], and synthesizing methods [47]–[50].

The projection-based methods obtain instances by projecting the elements in the feature space and semantic space into a common space. Generally, the feature space contains the instances of seen classes, and the semantic space contains the prototypes of seen/unseen classes. In [42], a hinge rank loss was constructed to learn a linear transformation from the feature space to the semantic space. In [43], a bilinear mapping function was learned by minimizing the loss of an SVM classifier. In [51], the prediction noise and class bias of label embedding were decreased by deploying multiple classifiers in an ensemble manner. Instance-borrowing methods are based on the similarities of classes, in which the instances belonging to similar classes are taken as positive. In [44], instances in training were borrowed from the seen classes. The borrowed instances have high similarities with the unseen classes. Unlike instance-borrowing methods, the synthesizing methods create the pseudo instances for unseen classes, where the adversarial autoencoder and generative adversarial networks (GANs) are often used. In [52], an optimal latent space was learned to construct a bias-reducing generator network, which can reduce the hubness problem. In [53], an out-of-distribution detector was introduced to reduce the effect of domain shift, and a GAN was employed to synthesize the unseen instances.

To handle multilabel images, the projection is much more complex for the projection-based ZSL. A possible solution is to follow the visual-semantic mapping strategy used in the single-label case. In [75], the meaning of multiple labels for one instance was inferred by summing the word vector representations of individual labels. In [54], an alternative way was proposed to use the direct visual-semantic mapping, which first finds the corresponding area of each semantic label and then extracts the object-level visual presentation for visual-semantic mapping. Some other works try to utilize the co-occurrence among the labels. One typical work is the COSTA [55], which constructs the linear projection matrix between the seen labels and unseen labels by statistic learning on the annotated data sets.

B. Hashing-Based Image Retrieval

Hashing methods for image retrieval can be roughly divided into two categories: the unsupervised and the supervised.

The unsupervised hashing methods generate hash codes without any semantic labels. They use clustering techniques or projection strategies to transfer visual information to feature space and generate an optimized hash function to preserve the similarity in Hamming space [10], [12], [56]–[58]. Some classical algorithms, such as SH [11], formulated hash encoding as a spectral graph-partitioning problem and learned a nonlinear mapping to preserve semantic similarity. Some other methods, e.g., SPQ [59] and muti-k-means [60], decomposed the feature space into a Cartesian product of low-dimensional subspaces and encoded high-dimensional feature vectors into binary codes by clustering-based subspace quantization [61]. ECE [62] treated it as an optimization problem and combined the genetic programming with the boosting-based weight updating. DSTH [63] advocated discrete hash codes and resorted to the semantics augment from auxiliary contextual modalities.

The supervised hashing methods use the annotation information to learn compact hash codes, which usually performs better than the unsupervised methods. Among these supervised methods, CNN-based hashing methods have attracted more and more attention due to the powerful representation ability of deep neural networks [64]. According to the difference of input forms, the existing deep hashing methods can be divided into two kinds. One receives image triplets as the input of the network and generates hash codes by minimizing the triplet ranking loss [5], [6]. These methods consume many computing resources and time to train the hashing model as there are enormous triplet combinations. The other receives the minibatch of images as input and uses the pairwise similarity between images as supervised information to learn the hash codes. Typical works of this form include the DHN [24], DSH [25], HashNet [26], and so on. Considering that existing hashing methods often fall short in concentrating relevant images to be within a small Hamming ball, DCH [65] built a pairwise cross-entropy loss on the Cauchy distribution to improve its capability on this point.

In recent years, hashing-based methods were also developed for multimodal retrieval, which transforms high-dimensional data of different modalities into compact binary codes in a common Hamming space. The multimodal hashing supports image retrieval across other domains, e.g., text, video, audio, and so on. Some representative work can be found in [66]–[71].

C. Zero-Shot Hashing

To handle images with unseen categories, some deep learning-based methods formulate the hashing as an unsupervised problem [72], [73]. However, without using reliable supervised information, it is difficult to achieve satisfactory performance. Some other methods [30], [31], [74], from a different perspective, consider it as a ZSH problem. The goal of ZSH [75], [76] is to transfer the model trained on the seen data to unseen data via other available knowledge, such as word vector representations. Since the underlying data distributions of the seen categories and the unseen categories are different, the hashing functions learned by the seen categories

without any adaptation to the unseen categories may cause a domain-shift problem. To narrow the domain gap between seen data and unseen data, ZSH-DA [32] first learns a ZSH model on seen data and then learns the final hashing model with a domain-adaptation algorithm. In [33], a transductive zero-shot hashing (TZSH) network was proposed, which contains a coarse-to-fine similarity mining to find the most presentative target examples of each unseen labels and adds these presentative examples and its corresponding predicted labels to the process of supervised hashing learning. Shen *et al.* [77] used GCN to learn the ZSH model for sketch-image retrieval. Although the GCN method is very promising in exploring the relationship between semantic labels, it is not flexible and requires a predefined adjacency matrix of nodes and additional training costs.

III. TRANSDUCTIVE MULTILABEL ZERO-SHOT HASHING

A. Problem Definition

Suppose that $\mathcal{D}^s = \{\mathcal{I}_i^s, \mathcal{Y}_i^s\}_{i=1}^{N_s}$ is a labeled source data set including N_s images, where \mathcal{I}_i^s is an image and \mathcal{Y}_i^s is the corresponding label annotated with one or more classes and $\mathcal{D}^t = \{\mathcal{I}_i^t\}_{i=1}^{N_t}$ is an unlabeled target data set, which includes N_t images of the unseen target classes and has the labels \mathcal{Y}^t unknown. In the zero-shot setting, the target and source classes are two mutually exclusive label sets, i.e., $\mathcal{Y}^t \cap \mathcal{Y}^s = \phi$. For hash-code learning, we construct the similarity matrix $\mathcal{S} = \{s_{ij} | i, j = 1, 2, \dots, N_s + N_t\}$, where $s_{ij} = 1$ denotes that the pairwise images \mathcal{I}_i and \mathcal{I}_j are similar, and $s_{ij} = 0$ denotes they are dissimilar. The goal of T-MLZSH is to learn a mapping $\mathcal{F} : \mathcal{I} \mapsto \{-1, +1\}^M$ on the labeled source data set \mathcal{D}^s and the unlabeled data set \mathcal{D}^t to encode an input image \mathcal{I}_i into an M -bit binary code $\mathcal{F}(\mathcal{I}_i)$, with the pairwise similarity preserved.

Fig. 2 shows a flowchart of the proposed method. The input images first go through the deep network with the stacked convolutional and fully connected layers and are encoded as a high-dimensional feature representation. Then, the outputs of the last fully connected layer are fed into a hashing layer for compact binary encoding. To transfer the knowledge from seen categories to unseen categories and construct the bridge between visual and semantic modalities, we add a fully connected layer after hashing layer, which maps features from Hamming space to the common embedding space.

B. Instance-Concept Coherence Ranking

Since there is no label information for target images, we should first predict labels for these images by transferring the knowledge from the semantic representations to visual features, before learning a supervised hashing function. If v_i is the visual embedding of the i th image instance \mathcal{I}_i and u_j is the semantic embedding of the j th semantic concept, then we can calculate the relatedness score between \mathcal{I}_i and the j th semantic concept in the embedding space

$$o_{ij} = \langle v_i, u_j \rangle \quad (1)$$

where $\langle a, b \rangle = a^T b$ is the inner product operation. The semantic embeddings can be obtained from the existing word

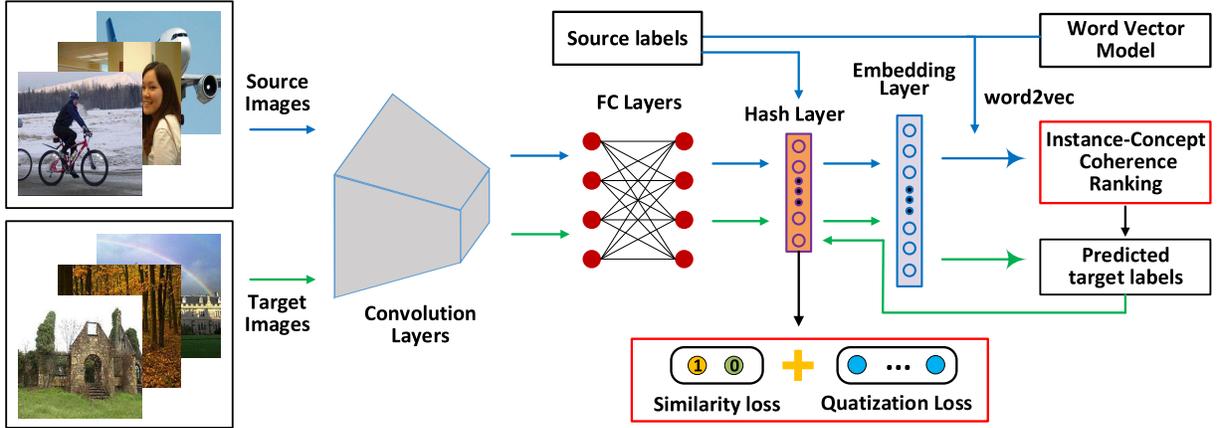


Fig. 2. Overview of the proposed T-MLZSH network. At first, the model constructs a visual-semantic bridge via instance-concept coherence ranking on the source data. It calculates the relatedness scores between visual features and semantic word vectors, under the assumption that related instance-concept pairs should have larger relatedness scores. Given the learned instance-concept coherence relationship, the most relevant concepts for each image instance of the unlabeled target data can be selected to guide the similarity-preserving learning. The whole network can be trained in an end-to-end manner.

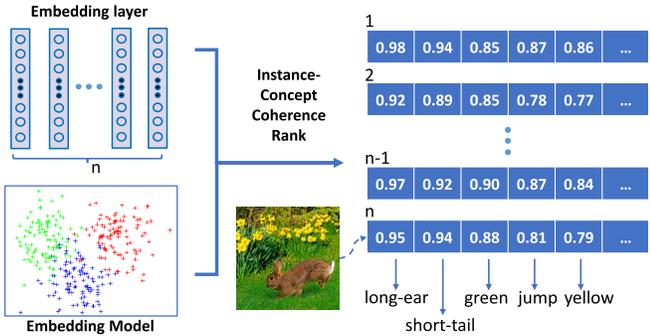


Fig. 3. Illustration of the embedding model. It learns a mapping function where the scores with a relevant label should be higher than that with an irrelevant one.

vector models, and the visual embeddings are variables that should be learned. During the training process, we can get a score list of source labels $\{o_{i1}, o_{i2}, \dots, o_{iL_s}\}$, where L_s is the number of seen labels. The goal of our embedding model is to learn a mapping function that scores with a relevant label should be higher than that with an irrelevant one, as shown in Fig. 3. Inspired by [78], we adopt a RankNet loss function to learn the ranking relationships, for instance, \mathcal{I}_i

$$\mathcal{L}_{\text{rank}} = w_i \cdot \left(\sum_{p \in \mathcal{C}_i^+} \sum_{q \in \mathcal{C}_i^-} \log(1 + \exp(o_{iq} - o_{ip})) + \sum_{j \in \mathcal{C}} \log(1 + \exp(-\psi_{ij} o_{ij})) \right) \quad (2)$$

where \mathcal{C}_i^+ and \mathcal{C}_i^- denote two sets of relevant and irrelevant labels to the i th instance and ψ_{ij} is defined as an indicator function, where $\psi_{ij} = 1$ indicates that the i th instance is related to the j th label and $\psi_{ij} = -1$ indicates that the i th instance is irrelevant to the j th label. $w_i = (|\mathcal{C}_i^+| \cdot |\mathcal{C}_i^-| + |\mathcal{C}|)^{-1}$ plays a regularization role. In the bracket of (2), the first term gives punishment to the situation when the labels irrelevant to \mathcal{I}_i have higher ranking orders than the relevant ones.

The second term is used to enlarge the relatedness scores of the relevant pairs and reduce those of the irrelevant pairs.

Based on the above-trained model, the pairwise relatedness scores for the visual embeddings of target images and the semantic embeddings of target classes can be calculated. We rank the scores $\{o_{i1}, o_{i2}, \dots, o_{iL_t}\}$ (L_t is the number of unseen labels) in descending order and select the classes of top- k highest scores as the predicted target labels.

C. Hash-Code Learning

For efficient nearest neighbor search, the semantic similarity of original images should be preserved in the Hamming space. Generally, the similarity relationships can be defined with image labels. For a multilabel data set, if two images share at least one label, they are considered similar and dissimilar otherwise. If \mathcal{B} is a set of hash codes for all images and $\mathcal{S} = \{s_{ij}\}$ is the pairwise similarity matrix, then the conditional probability of s_{ij} can be defined as

$$p(s_{ij}|\mathcal{B}) = \begin{cases} \sigma(\Omega_{ij}), & s_{ij} = 1 \\ 1 - \sigma(\Omega_{ij}), & s_{ij} = 0 \end{cases} \quad (3)$$

where $\Omega_{ij} = \langle b_i, b_j \rangle$ is the inner product of hash codes b_i and b_j and $\sigma(x) = (1/1 + e^{-x})$ is the sigmoid function, which scales the inner product value within $[0, 1]$.

We adopt negative log-likelihood as the cost function to measure the pairwise similarity loss, as formulated by the following equation:

$$\begin{aligned} \mathcal{L}_p &= - \sum_{s_{ij} \in \mathcal{S}} \log(p(s_{ij}|\mathcal{B})) \\ &= - \sum_{s_{ij} \in \mathcal{S}} (s_{ij} \cdot \log(\sigma(\Omega_{ij})) + (1 - s_{ij}) \cdot \log(1 - \sigma(\Omega_{ij}))) \end{aligned} \quad (4)$$

As $\sigma(\Omega_{ij}) = (1/(1 + e^{-\Omega_{ij}}))$, (4) can be rewritten as

$$\mathcal{L}_p = \sum_{s_{ij} \in \mathcal{S}} (\log(1 + e^{\Omega_{ij}}) - s_{ij} \cdot \Omega_{ij}). \quad (5)$$

It is very challenging to directly optimize this discrete optimization problem, as the binary constraint $b_i \in \{-1, 1\}^M$ requires thresholding on the network outputs, which may result in the vanishing-gradient problem in backpropagation. We adopt the continuous relaxation strategy [24], [25] to solve this problem. The output of deep hashing layer u_i is fed to a tanh function $h_i = \tanh(u_i)$, which is used as a substitute for binary code b_i . Thus, Ω_{ij} is redefined as $h_i^T h_j$.

For more efficient and faster hash learning, we design a focal quantization loss to mitigate the divergence between the discrete binary codes and the continuous output of hashing networks, inspired by [79]. Since the gradient accumulations of a large number of simple samples are not helpful for training, the focal loss attempts to reduce the weights of simple examples to promote the training process. First, we convert the binary code quantization problem into a binary classification problem. We use a sigmoid activation to map the outputs of the hash layer into a probability distribution $\hat{p}_i = \sigma(u_i)$. Note that, tanh and sigmoid are both monotonic increasing functions that hold the same variation trend, i.e., when h_i asymptotically approaches to -1 , p_i also approaches 0, and vice versa (both approach to 1). Thus, the probability of binary classification can reflect the compactness of hash codes effectively.

The focal quantization loss is defined as

$$\mathcal{L}_q = -\frac{1}{N} \sum_{i \in N} \sum_{j \in M} (\hat{y}_{ij} \cdot (1 - \hat{p}_{ij})^\alpha \cdot \log(\hat{p}_{ij}) + (1 - \hat{y}_{ij}) \cdot (\hat{p}_{ij})^\alpha \cdot \log(1 - \hat{p}_{ij})) \quad (6)$$

where \hat{y}_i is a label indicator that indicates which class (0 or 1) the output of hash layer should be classified as. We adopt a weighted sigmoid function to achieve such effect, i.e., $\hat{y}_i = \sigma(\beta \cdot u_i)$, with β a parameter far greater than 1.

By integrating the pairwise similarity loss and focal quantization loss, the overall hashing loss can be defined as

$$\mathcal{L}_{\text{hash}} = \mathcal{L}_p + \mathcal{L}_q. \quad (7)$$

IV. EXPERIMENTS

A. Data Sets

To verify the performance of the proposed method, we compare the proposed method with several baselines on three widely used multilabel image data sets, i.e., NUS-WIDE, VOC2012, and COCO.

NUS-WIDE [80] is a data set containing 269 648 public web images. Each image is annotated with one or more class labels from a total of 81 classes. There exists a widely used subset of images associated with the 21 most common labels and each label associated with at least 5000 images, resulting in a total of 195 834 images.

VOC2012 [81] is a widely used data set for object detection and segmentation, which contains 17 125 images. Each image is associated with at least one of the 20 semantic labels.

COCO [14] is a data set for object detection, semantic scene labeling, and indexing, which contains 123 287 images with semantic labels. Each image is associated with 1-to-16 labels from a total of 90.

B. Implementation Details

To construct a zero-shot scenario, we should further split the data set.¹ Since there are more complex semantic relationships among these multilabel data sets, we use one of these three image data sets as source data and one as target data. For example, we can set NUS-WIDE as source data and VOC2012 as target data and vice versa. Before training models based on these data sets, data preprocessing must be done. We set three experiments, including one between NUS-WIDE and VOC2012, one between NUS-WIDE and COCO, and the last one between COCO and VOC2012.

1) *Experiment Between NUS-WIDE and VOC2012:* In NUS-WIDE, we remove the common concepts (semantic labels) shared by these two data sets and related images because there are much more images in NUS-WIDE than in VOC2012. In VOC2012, we remove several ambiguous concepts and related images. Such data-clean operations result in a subset of NUS-WIDE containing 106 389 images and 18 labels and a subset of VOC2012 containing 16 750 images and 17 labels. For NUS-WIDE, we randomly select 10 000 images as the training set, 2000 images as the test query set, and the rest as the retrieval database. For VOC2012, we randomly select 4000 images as the training set, 1000 images as the test query set, and the rest as the retrieval database.

2) *Experiment Between NUS-WIDE and COCO:* We remove the common concepts and relative images from NUS-WIDE and keep COCO unchanged. Finally, a subset of NUS-WIDE containing 100 303 images and 17 labels and a subset of COCO containing 123 274 images and 80 labels are prepared for the following experiments. For both data sets, we randomly select 10 000 images as the training set, 2000 images as the test query set, and the rest as the retrieval database.

3) *Experiment Between VOC2012 and COCO:* We remove the images related to three ambiguous concepts for VOC2012. After that, all the concepts of VOC2012 are included in COCO. Then, for COCO, we remove the images that have the concepts in VOC2012. Finally, there remains 21 987 images and 60 labels for COCO and 16 750 images and 17 labels for VOC2012. For both data sets, we randomly select 4000 images for training, 1000 images as the test query, and the rest as the retrieval database.

For NUS-WIDE, VOC2012, and COCO, we check the labels and ensure that the training/query set contains all the concepts of the corresponding data set.

We implement the proposed method (T-MLZSH) using the TensorFlow toolkit. In this article, we use AlexNet as the backbone CNN. To validate the versatility of the proposed framework, we will also evaluate by replacing the backbone CNN with VGG16 and ResNet50. We use the pretrained model to initialize the weight parameters and focus on training the hashing layer and embedding layer. The Adam method is adopted for stochastic optimization with a mini-batch size of 128, and all input images are resized to 227×227 .

We compare our method (T-MLZSH) with nine other methods, including the traditional methods KSH [2], IMH [3],

¹<https://github.com/qinnzou/Zero-Shot-Hashing>

SDH [4], ZSH-DA [32], ZSH [30], the deep learning-based methods DHN [24], Hashnet [26], ADSH [82], and TZSH [33]. Among the traditional methods, ZSH-DA and ZSH are two zero-shot hashing methods. Among the deep learning-based methods, TZSH is a T-MLZSH. It is worth noting that all the methods use the same training and test settings on the newly formed data sets.

In the training, for KSH, IMH, SDH, DHN, Hashnet, and ADSH, the images and labels of the training set are used, whereas for ZSH-DA, ZSH, TZSH, and the proposed T-MLZSH, in addition to the images and labels in the training set, the raw images (without labels) in the test query set are also used. Note that, ADSH is a supervised method that treats the query points and database points asymmetrically, and DIHN [83] successfully extends the asymmetric training strategy in an incremental learning framework. As DIHN uses the labels of the unseen test query data for training, which violates the protocols of our experiments, we do not include it in the comparison.

For deep learning-based methods, we use the raw images as input. For the non-CNN approaches, we use the outputs of the fc7 layer in AlexNet as their visual features. All the experiments are conducted on an NVIDIA TITAN Xp GPU and the Ubuntu 16.04 operating system.

C. Metrics

The metrics we used to evaluate the image retrieval quality are four widely used metrics: average cumulative gains (ACGs), normalized discounted cumulative gains (NDCGs), mean average precision (MAP), and weighted mean average precision (WAP) [7], [23], [84], [85].

MAP is the mean of average precision for each query, which can be calculated by

$$\text{MAP} = \frac{1}{Q} \sum_Q \text{AP}(q) \quad (8)$$

where

$$\text{AP}(q) = \frac{1}{N_{tr}(q)@n} \sum_n i \left(\text{Tr}(q, i) \frac{N_{tr}(q)@i}{i} \right) \quad (9)$$

$\text{Tr}(q, i) \in \{0, 1\}$ is an indicator function that if I_q and I_i have the same class labels, $\text{Tr}(q, i) = 1$; otherwise, $\text{Tr}(q, i) = 0$. Q is the number of query sets and $N_{tr}(q)@i$ indicates the number of relevant images with respect to the query image I_q within the top i images. ACG represents the average number of shared labels between the query image and the top n retrieved images. For a given query image I_q , the ACG score of the top n retrieved images is calculated by

$$\text{ACG}@n = \frac{1}{n} \sum_n^i C(q, i) \quad (10)$$

where n denotes the number of top retrieval images and $C(q, i)$ is the number of shared class labels between I_q and I_i . NDCG is a popular evaluation metric in information retrieval. Given a query image I_q , the DCG score of top n retrieved

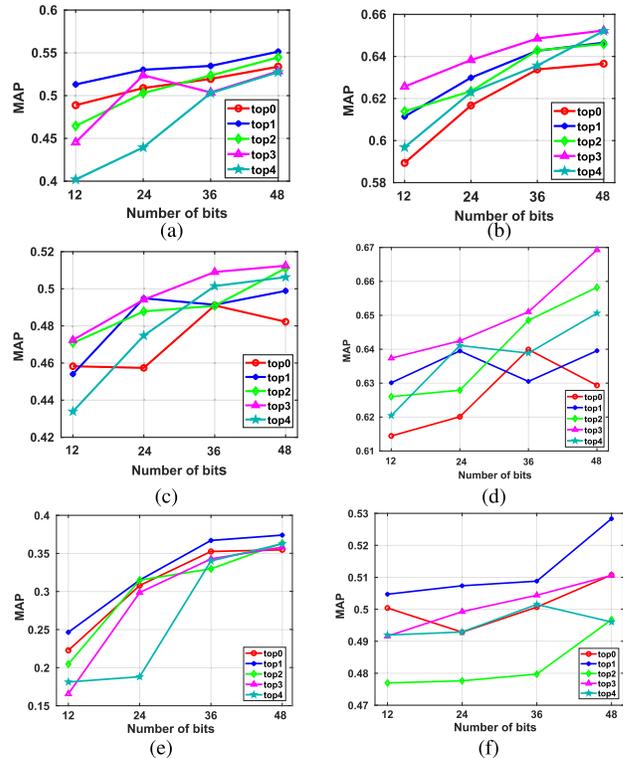


Fig. 4. Performance obtained by using different numbers of predicted labels for target data with hash codes of 12, 24, 36, and 48 bits. The values are computed based on the top-1000 retrieved images. (a) NUS-WIDE \rightarrow VOC2012. (b) VOC2012 \rightarrow NUS-WIDE. (c) NUS-WIDE \rightarrow COCO. (d) COCO \rightarrow NUS-WIDE. (e) VOC \rightarrow COCO. (f) COCO \rightarrow VOC.

images is defined as

$$\text{DCG}@n = \sum_n^i \frac{2^{C(q,i)} - 1}{\log(1+i)}. \quad (11)$$

Then, the normalized DCG(NDCG) score at the position n can be calculated by $\text{NDCG}@n = (\text{DCG}@n / Z_n)$, where Z_n is the maximum value of $\text{DCG}@n$, which constrains the value of NDCG in the range $[0, 1]$.

WAP is similar to MAP, and the only difference is that WAP is the average value of ACG scores at each tip n retrieval image rather than average precision. WAP can be calculated by

$$\text{WAP} = \frac{1}{Q} \sum_Q \left(\frac{1}{N_{Tr}(q)@n} \sum_i^n (\text{Tr}(q, i) \times \text{ACG}@i) \right). \quad (12)$$

D. Overall Performance

In this section, we analyze the retrieval results all evaluated on the unseen target data.

1) *Number of Predicted Unseen Categories*: Fig. 4 shows the results of using different numbers of predicted labels on the target data. Top- k indicates that the first k categories in the correlation-score ranking list are used as predicted labels and top-0 means that the labels of target data are set to a vector of all zeros. From Fig. 4(a) and (b), we can see that when setting VOC2012 as target data, using top-1 predicted labels can achieve the best performance. The possible reason is that the average number of objects in each image on

TABLE I
MAP RESULTS OBTAINED BY USING DIFFERENT NUMBERS OF BITS ON (NUS-WIDE, VOC2012).
THE MAPs ARE COMPUTED BASED ON THE TOP-1000 RETRIEVED IMAGES

Methods	NUS-WIDE \rightarrow VOC2012				VOC2012 \rightarrow NUS-WIDE			
	12-bit	24-bit	36-bit	48-bit	12-bit	24-bit	36-bit	48-bit
KSH [2]	0.4033	0.4079	0.4153	0.4181	0.5476	0.5510	0.5602	0.5670
SDH [4]	0.4097	0.4010	0.4087	0.4095	0.5327	0.5345	0.5368	0.5395
IMH [3]	0.4083	0.4346	0.4320	0.4297	0.5616	0.5723	0.5718	0.5710
DHN [24]	0.4171	0.4282	0.4362	0.4395	0.5664	0.5739	0.5726	0.5688
Hashnet [26]	0.4048	0.4288	0.4349	0.4465	0.5674	0.5940	0.6073	0.6336
ADSH [82]	0.3988	0.4163	0.4060	0.4201	0.5315	0.5473	0.5392	0.5853
ZSH-DA [32]	0.3592	0.3618	0.3770	0.3596	0.5132	0.5166	0.5212	0.5191
ZSH [30]	0.3968	0.4055	0.4111	0.4296	0.5340	0.5566	0.5589	0.5500
TZSH [33]	0.4413	0.4683	0.4644	0.4753	0.5736	0.5805	0.5919	0.5896
T-MLZSH	0.4808	0.4884	0.4894	0.5037	0.6106	0.6131	0.6149	0.6200

TABLE II
MAP RESULTS OBTAINED BY USING DIFFERENT NUMBERS OF BITS ON (NUS-WIDE, COCO).
THE MAPs ARE COMPUTED BASED ON THE TOP-1000 RETRIEVED IMAGES

Methods	NUS-WIDE \rightarrow COCO				COCO \rightarrow NUS-WIDE			
	12-bit	24-bit	36-bit	48-bit	12-bit	24-bit	36-bit	48-bit
KSH [2]	0.3948	0.4069	0.4113	0.4143	0.5948	0.6167	0.6191	0.6224
SDH [4]	0.3782	0.3917	0.3971	0.4050	0.5681	0.5954	0.6102	0.6051
IMH [3]	0.3905	0.4021	0.4114	0.4188	0.5983	0.5961	0.6098	0.6132
DHN [24]	0.4250	0.4325	0.4529	0.4487	0.6177	0.6421	0.6466	0.6559
Hashnet [26]	0.3908	0.4018	0.4129	0.4439	0.5538	0.5674	0.5830	0.5932
ADSH [82]	0.3783	0.3890	0.3968	0.4056	0.5773	0.5906	0.5871	0.6219
ZSH-DA [32]	0.3597	0.3592	0.3772	0.3744	0.5256	0.5220	0.5230	0.5247
ZSH [30]	0.3832	0.4091	0.4109	0.4286	0.5708	0.5727	0.5753	0.5782
TZSH [33]	0.4436	0.4585	0.4660	0.4800	0.5933	0.6368	0.6070	0.6336
T-MLZSH	0.4724	0.4941	0.5090	0.5124	0.6374	0.6425	0.6510	0.6693

VOC2012 is relatively small. With more predicted labels used for supervised hashing learning, it will inevitably incur misleading information and cause performance degradation. When setting NUS-WIDE as target data, the best results are obtained by using top-3 predicted labels. In the following experiments, we use top-1 and top-3 predicted labels as supervised information for the proposed method in default for VOC2012 and NUS-WIDE, respectively. In Fig. 4(c)–(f), we can see that the best performance can be achieved by using top-3 predicted labels for the experiments on COCO and NUS-WIDE and using top-1 for experiments on COCO and NUS-WIDE. It may be because the average number of objects in each image in the (NUS-WIDE, COCO) case is 2.48 and 2.97 and that in the (VOC, COCO) case is 1.30 and 1.48, respectively. Hence, we use top-3 predicted labels as supervised information for the proposed method in default for (COCO, NUS-WIDE) and top-1 for (VOC, COCO) in the following experiments.

2) *Results Under Different Scenarios:* The MAP results of the proposed method and the comparison methods are shown in Tables I–III. It can be seen that the nonzero-shot deep hashing methods, DHN, Hashnet, and ASDH, outperform the traditional hashing methods, KSH, IMH, SDH, ZSH-DA, and ZSH. Among the nonzero-shot deep hashing methods, ADSH has lower performance than the other two. The possible reason

is that it treats the query points and database points in an asymmetric way, which drives the network to pay more attention to the seen data. Thus, when the data set is replaced by the unseen data set, ADSH’s performance will be affected. The two traditional ZSH methods ZSH and ZSH-DA achieve low performance on the multilabel data sets, which indicates that the complex semantics of multilabel images are too hard to be modeled by learning a one-to-one semantic representation. The zero-shot deep hashing methods, T-MLZSH and TZSH, obtain significantly higher performance than the other methods.

It can also be seen that the proposed T-MLZSH outperforms almost all the comparison methods significantly on all target data sets under different scenarios. On NUS-WIDE and VOC2012, the transductive zero-shot hashing methods, i.e., TZSH and T-MLZSH, achieve higher MAP values than other methods, as shown in Table I. Compared with TZSH, T-MLZSH achieves the increments of about 3.1% and 2.8% in the average MAP for different bits on NUS-WIDE and VOC2012, respectively. The possible reason is that TZSH adopts a strategy only utilizing partially selected target data for hash learning, which limits its performance.

From Table II, we can see that compared with TZSH, T-MLZSH achieves increments of about 3.4% or 3.2% in the average MAP for different bits on COCO \rightarrow NUS-WIDE and NUS-WIDE \rightarrow COCO, respectively. In this experiment,

TABLE III
MAP RESULTS OBTAINED BY USING DIFFERENT NUMBERS OF BITS ON (VOC2012, COCO).
THE MAPs ARE COMPUTED BASED ON THE TOP-1000 RETRIEVED IMAGES

Methods	VOC2012 \rightarrow COCO				COCO \rightarrow VOC2012			
	12-bit	24-bit	36-bit	48-bit	12-bit	24-bit	36-bit	48-bit
KSH [2]	0.1763	0.2076	0.2346	0.2490	0.4498	0.4781	0.4912	0.5005
SDH [4]	0.1444	0.1762	0.1871	0.1991	0.4656	0.4869	0.5045	0.5060
IMH [3]	0.2005	0.2151	0.2279	0.2317	0.4083	0.4291	0.4347	0.4380
DHN [24]	0.1635	0.1826	0.1954	0.2351	0.4311	0.4370	0.4573	0.4597
Hashnet [26]	0.1521	0.2029	0.2470	0.2585	0.4331	0.4605	0.4766	0.4924
ADSH [82]	0.1368	0.1472	0.1771	0.1884	0.4310	0.4794	0.5018	0.5061
ZSH-DA [32]	0.0853	0.0746	0.0709	0.1104	0.3924	0.4189	0.4247	0.4393
ZSH [30]	0.1602	0.1935	0.2249	0.2227	0.4658	0.4775	0.5003	0.4981
TZSH [33]	0.2336	0.2632	0.2690	0.2631	0.4967	0.5053	0.5128	0.5079
T-MLZSH	0.2465	0.3149	0.3670	0.3740	0.5047	0.5276	0.5296	0.5345

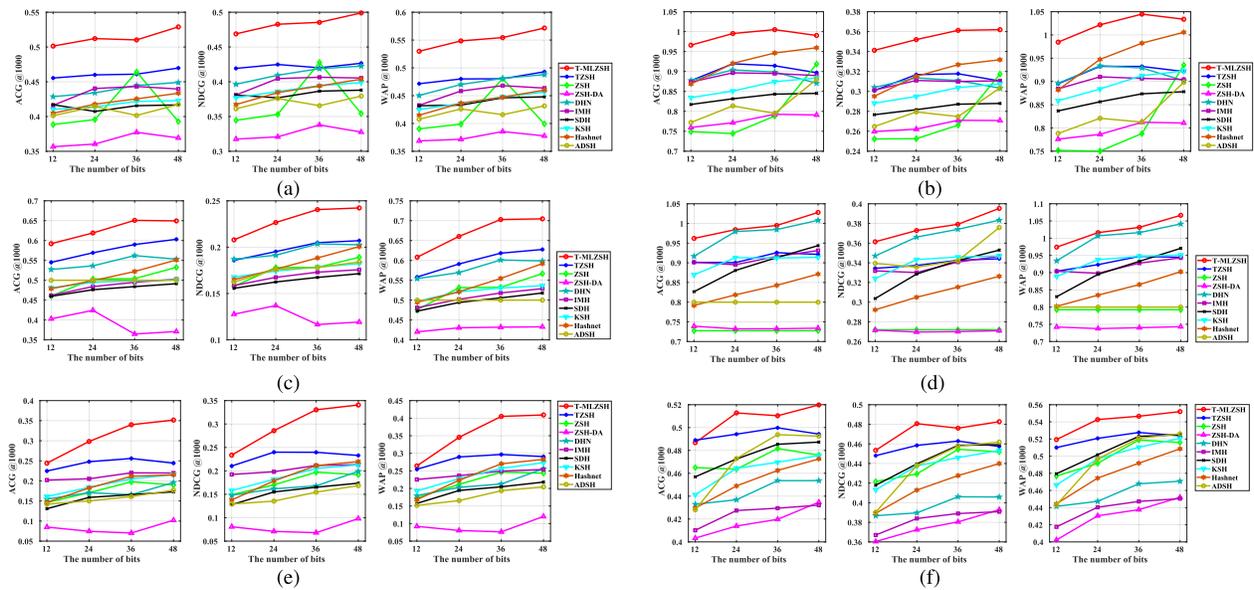


Fig. 5. Comparison of results obtained by the proposed method on ACG, NDCG, and WAP with 12-, 24-, 36-, and 48-bits hash codes, with regarding the top-1000 retrieved images. (a) NUS-WIDE \rightarrow VOC2012. (b) VOC2012 \rightarrow NUS-WIDE. (c) NUS-WIDE \rightarrow COCO. (d) COCO \rightarrow NUS-WIDE. (e) VOC2012 \rightarrow COCO. (f) COCO \rightarrow VOC2012.

the deep supervised hashing method DHN is found to perform better than TZSH on COCO \rightarrow NUS-WIDE. The possible reason is that COCO is categorized into more categories and DHN can get more detailed supervised information when setting COCO as the training set. Nevertheless, T-MLZSH still outperforms DHN by about 0.95% in the average MAP.

Table III shows the results on (COCO, VOC2012). Compared with TZSH, T-MLZSH achieves increments of about 6.8% or 1.8% on average MAP for different bits. The proposed method is found to obtain a significantly better performance on VOC2012 \rightarrow COCO, where COCO is unseen. The possible reason is that although all the instances of VOC2012 are included in COCO, some unlabeled targets in COCO images are similar to that in VOC2012. Features extracted from VOC2012 images may possess high similarity with some COCO images. Since the scale of COCO is much larger than VOC2012, the overall performance is lower when COCO is unseen.

More results in the other three metrics, i.e., ACG, NDCG, and WAP, are presented in Fig. 5. We can see that the overall trends of the performance on the three metrics are consistent. When the code length increases, the performance improves. According to the definition, these three metrics can make a more fair evaluation on multilabel images, as the numbers of shared labels between images are considered. For MAP, the pairwise images that share at least one common object label will be considered as relevant images, and no more comparisons of fine-grained semantic relation between these images are included, which may not stay in step with user demand in multilabel image retrieval. For WAP, the average number of shared class labels among these retrieved similar images is considered. Higher WAP means more high-quality retrieval results having more shared class labels in the nearest search. Although the range of WAP on different data sets would be different, the WAPs of T-MLZSH are stably higher than those of the comparison methods.

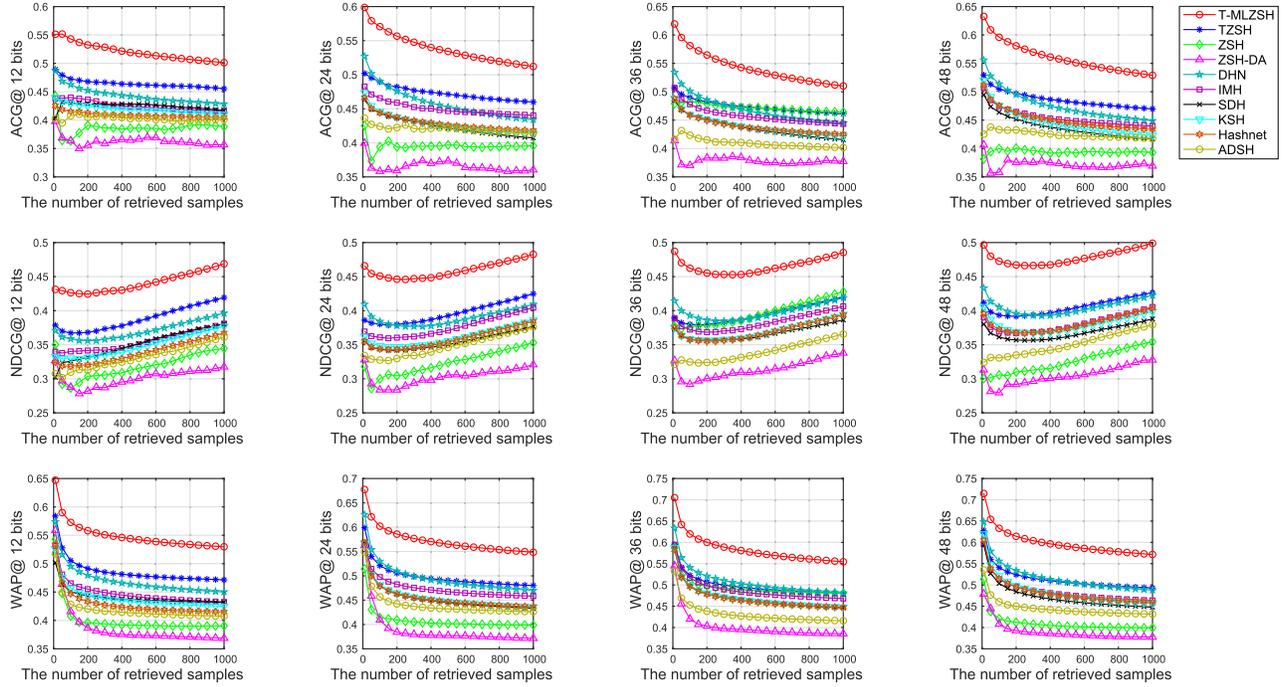


Fig. 6. Performance comparison on NUS-WIDE→VOC2012. The VOC2012 data set is unseen. From top to bottom, there are ACG, NDCG, and WAP with respect to different top returned samples with hash codes of 12, 24, 36, and 48 bits.

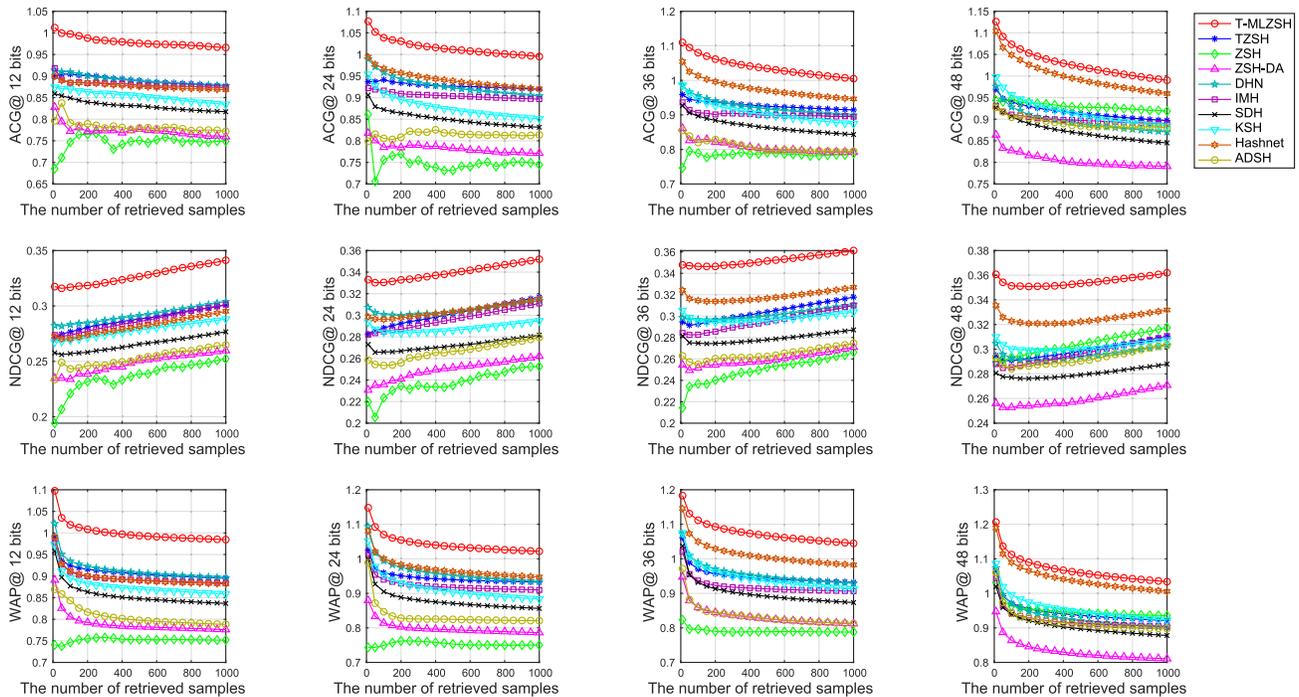


Fig. 7. Performance comparison on VOC2012→NUS-WIDE. The NUS-WIDE data set is unseen. From top to bottom, there are ACG, NDCG, and WAP with respect to different top returned samples with hash codes of 12, 24, 36, and 48 bits.

Figs. 6–11 show more detailed results on ACG, NDCG, and WAP of different numbers of top returned images under different scenarios. Generally, the NDCG curves show a trend of first decreasing and then increasing, while ACG and WAP decrease with the increase of retrieved samples. The more samples retrieved, the more low-quality retrieval results having fewer shared class labels will appear in the rear. This will lead

to the trends that tend to be gentle. Since most of the compared methods are not specific for ZSL, the results of unseen data may be uncertain. The curve of ZSH-DA is unsmooth. The possible reason is that, while using ZSH-DA for multilabel unseen images, the complex semantics are too hard to be modeled by learning a one-to-one semantic representation, which leads to unstable performances.

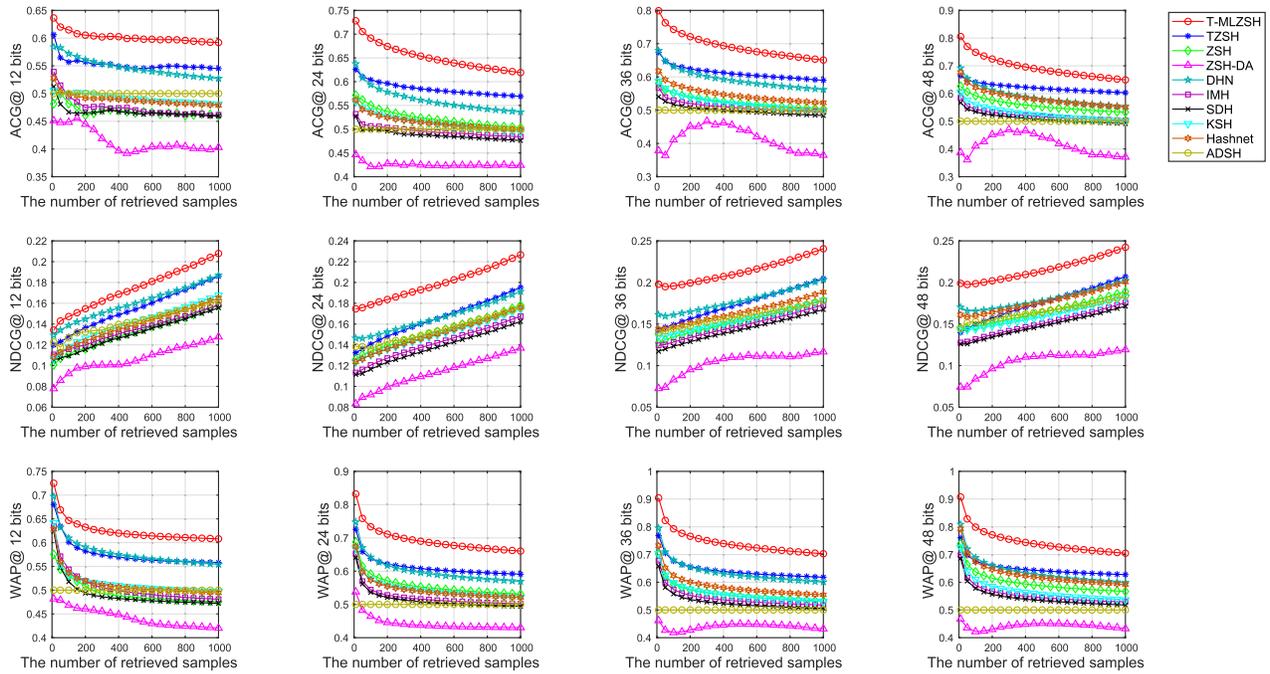


Fig. 8. Performance comparison on NUS-WIDE→COCO. The COCO data set is unseen. From top to bottom, there are ACG, NDCG, and WAP with respect to different top returned samples with hash codes of 12, 24, 36, and 48 bits.

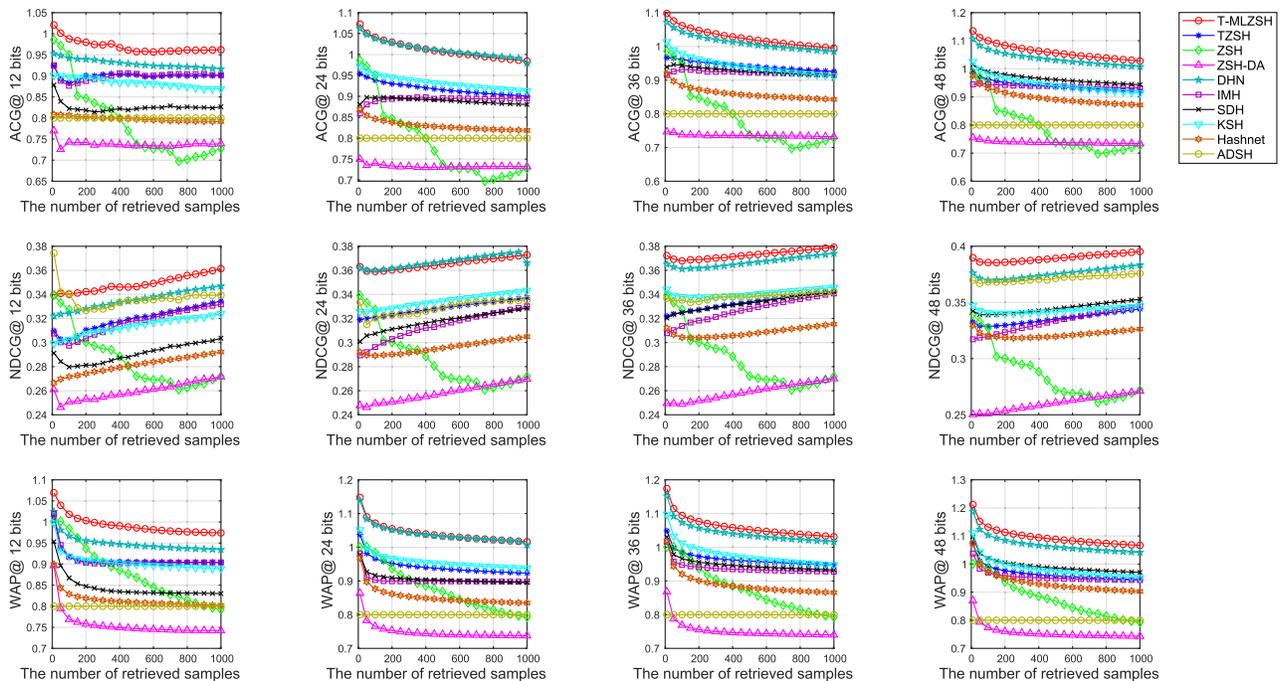


Fig. 9. Performance comparison on COCO→NUS-WIDE. The NUS-WIDE data set is unseen. From top to bottom, there are ACG, NDCG, and WAP with respect to different top returned samples with hash codes of 12, 24, 36, and 48 bits.

3) *Comparison With Different Backbone CNN Blocks:* To justify the versatility of the proposed deep hashing framework, we replace the backbone CNN with VGG16 and ResNet50, both of which achieve more accurate results than AlexNet on the ImageNet competition. We denote these two modifications as “T-MLZSH-VGG16” and “T-MLZSH-ResNet50,” respectively. The results are shown in Table IV. We can see from the results that, with more powerful backbone CNN blocks, T-MLZSH generally achieves higher performance on all these

metrics. It simply indicates a good transfer capability and versatility of the proposed deep hashing framework.

E. Functional Analysis

1) *Influence of the Categories of Data Sets:* In the above two groups of experiments, part of them used NUS-WIDE as the unseen data set and the results are shown in Tables I and II. From the left of the tables, we note that using COCO as seen data set can achieve a better MAP result, which has an

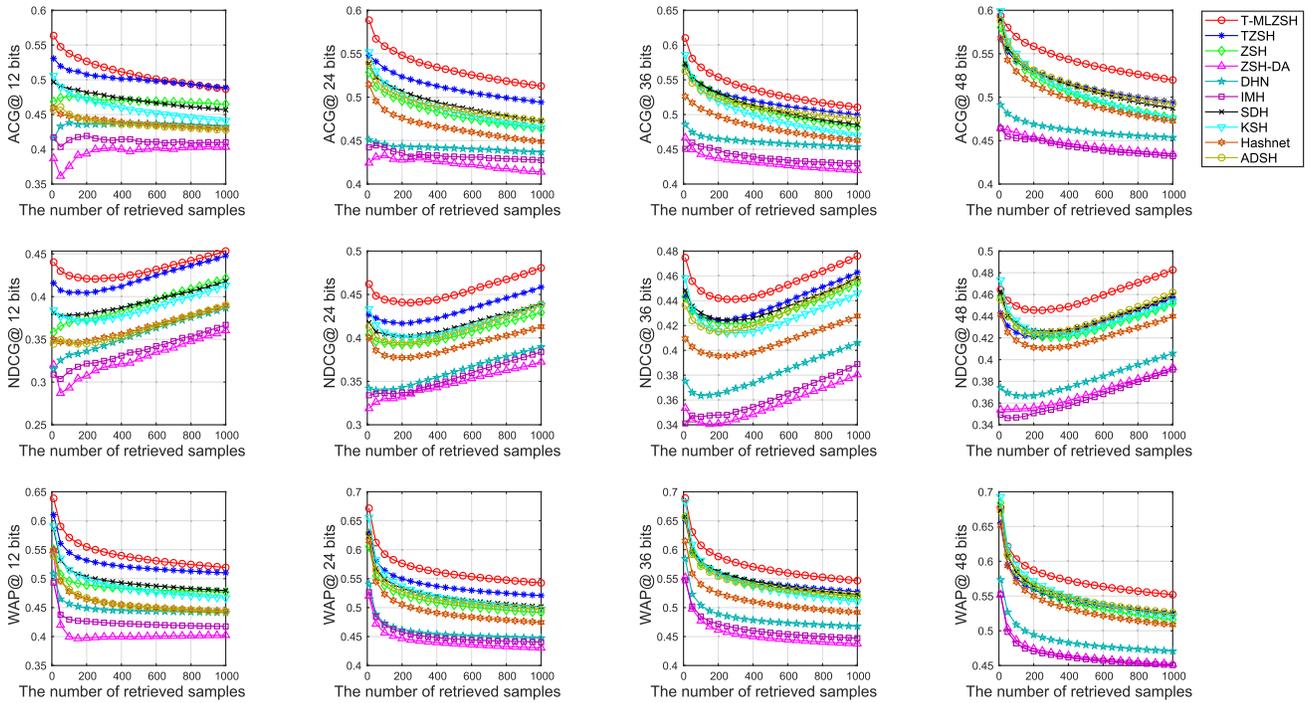


Fig. 10. Performance comparison on COCO→VOC2012. The VOC2012 data set is unseen. From top to bottom, there are ACG, NDCG, and WAP with respect to different top returned samples with hash codes of 12, 24, 36, and 48 bits.

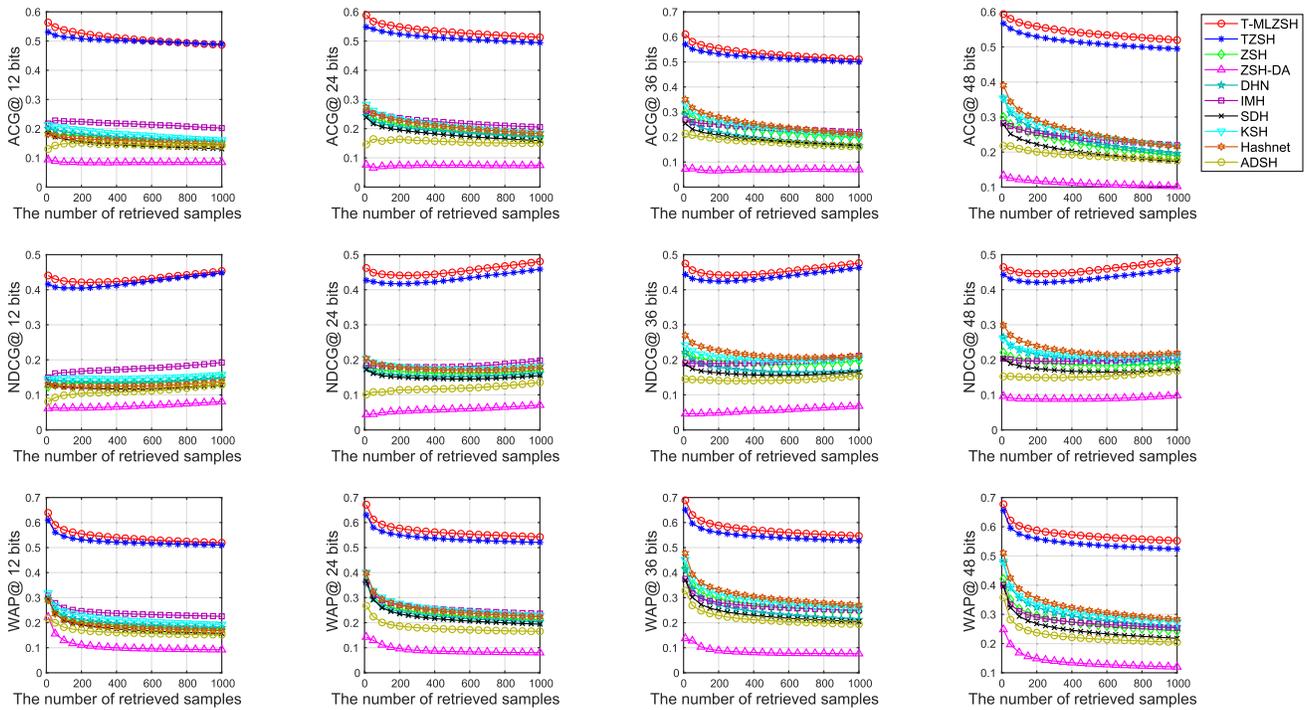


Fig. 11. Performance comparison on VOC2012→COCO. The COCO data set is unseen. From top to bottom, there are ACG, NDCG, and WAP with respect to different top returned samples with hash codes of 12, 24, 36, and 48 bits.

improvement of about 2.6%, 2.9%, 3.6%, and 4.9% in average MAP with different hash bits. These two groups of experiments have the same target domain and the only difference is the source domain. We guess that the possible reason leading to different MAPs is the difference of categories. COCO data set is more finely divided and more semantic information can be used, which makes the network much stronger.

2) *Influence of the Size of Data Sets:* Moreover, we explore the influence of the sizes of the source and target data sets by using different amounts of “seen” and “unseen” images to train the model. Three orders of magnitude are considered. Exactly, the number of images from the source data set and target data set is 10000 and 10000, 10000 and 4000, and 4000 and 10000, respectively. The results are shown in Fig. 12.

TABLE IV
PERFORMANCE UNDER DIFFERENT BACKBONE CNN BLOCKS. THE VALUES ARE COMPUTED ON THE TOP-1000 RETRIEVED IMAGES

Metrics	MAP				NDCG				ACG				WAP			
	12b	24b	36b	48b												
NUS-WIDE->VOC2012																
T-MLZSH-AlexNet	0.4808	0.4884	0.4894	0.5037	0.4689	0.4827	0.4855	0.4990	0.5298	0.5484	0.5545	0.5716	0.5298	0.5484	0.5545	0.5716
T-MLZSH-VGG16	0.5340	0.5346	0.5843	0.6149	0.6363	0.6270	0.6520	0.6782	0.5299	0.5124	0.5546	0.5798	0.5504	0.5521	0.6065	0.6379
T-MLZSH-ResNet50	0.5711	0.6054	0.6518	0.6541	0.6471	0.6820	0.6915	0.7377	0.5505	0.5761	0.5850	0.6171	0.5902	0.6274	0.6739	0.6791
VOC2012->NUS-WIDE																
T-MLZSH-AlexNet	0.6106	0.6131	0.6149	0.6200	0.3412	0.3520	0.3613	0.3619	0.9660	1.0207	1.0448	1.0338	0.9845	1.0217	1.0448	1.0338
T-MLZSH-VGG16	0.6049	0.6503	0.6635	0.6799	0.3280	0.3670	0.3657	0.3885	0.9175	1.0104	1.0238	1.0587	0.9315	1.0361	1.0597	1.1040
T-MLZSH-ResNet50	0.6542	0.6715	0.6957	0.6977	0.3677	0.3836	0.4056	0.4037	1.0267	1.0599	1.1146	1.0882	1.0485	1.0858	1.1373	1.1282
NUS-WIDE->COCO																
T-MLZSH-AlexNet	0.4724	0.4941	0.5291	0.5124	0.2097	0.2296	0.2405	0.2423	0.5921	0.6191	0.6508	0.7046	0.6080	0.6604	0.6508	0.7046
T-MLZSH-VGG16	0.4970	0.5705	0.6031	0.6075	0.3014	0.3683	0.3978	0.3935	0.6504	0.6724	0.6889	0.7093	0.6374	0.6609	0.6657	0.6700
T-MLZSH-ResNet50	0.4994	0.5519	0.6030	0.6209	0.3022	0.3639	0.3983	0.4112	0.6543	0.6698	0.6934	0.7102	0.6309	0.6532	0.6694	0.6741
COCO->NUS-WIDE																
T-MLZSH-AlexNet	0.6374	0.6425	0.6510	0.6693	0.3614	0.3728	0.3791	0.3953	0.9621	0.9014	0.9948	1.0280	0.9741	1.0162	1.0314	1.0665
T-MLZSH-VGG16	0.6333	0.6686	0.6709	0.6858	0.4380	0.4808	0.4798	0.4985	0.9718	0.9796	1.0297	1.1145	0.9932	1.0345	1.1136	1.1203
T-MLZSH-ResNet50	0.6588	0.6843	0.6883	0.6914	0.4713	0.4925	0.5044	0.5020	0.9821	0.9902	1.1342	1.1201	1.0023	1.0962	1.1361	1.1903
VOC2012->COCO																
T-MLZSH-AlexNet	0.2465	0.3149	0.3670	0.3740	0.2336	0.2857	0.3304	0.3406	0.2443	0.2983	0.3399	0.3515	0.2642	0.3455	0.4050	0.4097
T-MLZSH-VGG16	0.3707	0.4183	0.4757	0.4901	0.3631	0.3890	0.4430	0.4651	0.4033	0.4504	0.5151	0.4295	0.3968	0.4234	0.4596	0.4775
T-MLZSH-ResNet50	0.4161	0.4695	0.5184	0.5320	0.3992	0.4372	0.4766	0.4921	0.4558	0.5052	0.5673	0.5804	0.4033	0.4504	0.5151	0.5295
COCO->VOC1012																
T-MLZSH-AlexNet	0.5047	0.5074	0.5088	0.5283	0.4502	0.4567	0.4665	0.4860	0.4868	0.4791	0.4930	0.5134	0.5193	0.5244	0.5338	0.5555
T-MLZSH-VGG16	0.5583	0.6538	0.6652	0.6844	0.5131	0.6084	0.6100	0.6314	0.5761	0.6804	0.6905	0.7146	0.5506	0.6149	0.6224	0.6298
T-MLZSH-ResNet50	0.6402	0.6672	0.6796	0.6761	0.6056	0.6148	0.6330	0.6272	0.6210	0.6223	0.6327	0.6274	0.6614	0.6939	0.7071	0.7027

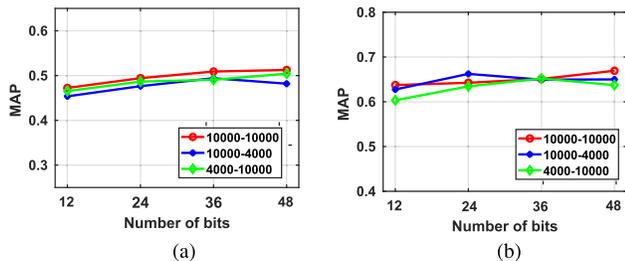


Fig. 12. Performance obtained by using different amounts of training data. The numbers indicate the training data in the form of “source-target.” (a) NUS-WIDE \rightarrow COCO. (b) COCO \rightarrow NUS-WIDE.

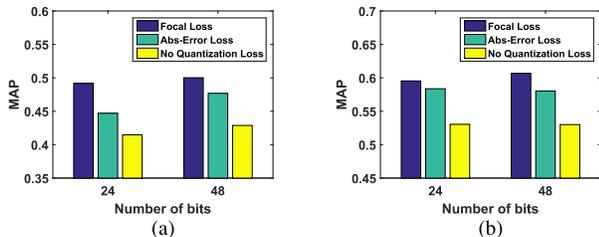


Fig. 13. Performance obtained by different quantization losses. (a) NUS-WIDE \rightarrow VOC2012. (b) VOC2012 \rightarrow NUS-WIDE.

It can be seen that there is a slight difference in the use of different orders of magnitude, but the performance is stable on the whole. It manifests that the proposed model has certain stability even if the number of images from two domains used for training varies from each other.

3) *Necessary of Quantization Loss:* We also explore the effectiveness of the proposed quantization loss. We compare the proposed method with its variant versions: one adopts the widely used absolute error loss that measures the Euclidean distance between continuous outputs and discrete codes directly, and the other does not use quantization loss. The results are presented in Fig. 13. It can be seen that, without quantization loss, there is a rapid degradation of

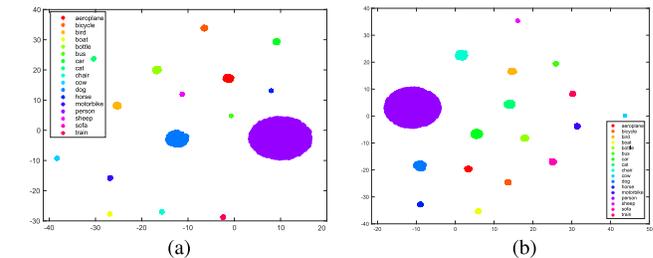


Fig. 14. Visualization of the predicted labels and the true labels using t-SNE. (a) Predicted label. (b) True label.

the performance. The difference in the evaluation index of the MAP is about 0.5%, which illustrates the importance of using quantization loss in deep hashing learning. We can also see that applying quantified losses has greatly improved the results, and the proposed focal quantization loss has a much more advanced performance among all the proposed architecture.

4) *Effectiveness of the Visual-Semantic Bridge:* The visual-semantic bridge is built to predict labels for unseen data. It helps the images belonging to the same category to be with higher similarity by transferring the knowledge from the semantic representations to visual features. Since the labels of unseen images are unknown, the predicted labels will be represented as meaningless digital codes for the unseen images.

To validate the effectiveness of the visual-semantic bridge in linking the semantic representations and visual features, we use the proposed model to predict labels for the seen images and compare them to true labels. In the experiment, we use the model trained on the training set of VOC2012 to predict the labels for 10000 images from the database set. Since the average number of objects in each image from VOC2012 is 1.301, we predict one label for each image. The distribution of labels is visualized by t-SNE, as shown in Fig. 14. Comparing Fig. 14(a) and (b), we can find that

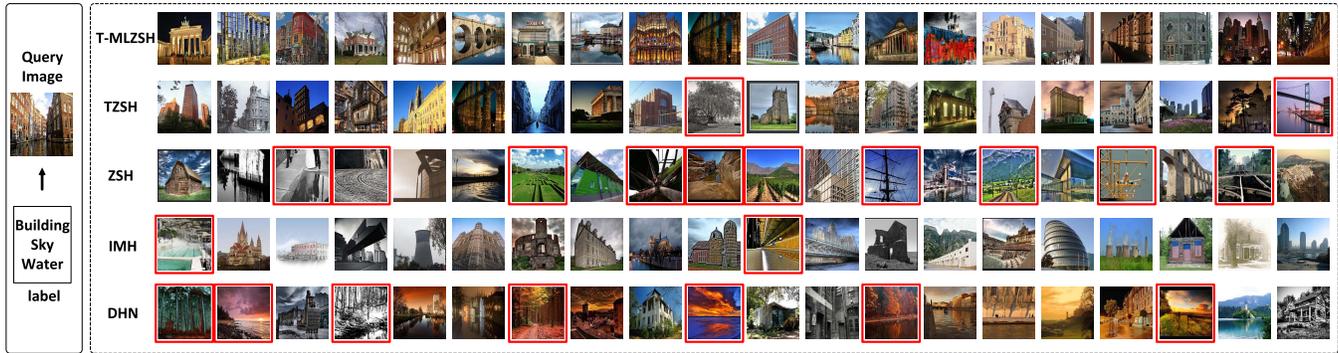


Fig. 15. Top-20 images retrieved by the proposed method and the comparison methods using the Hamming ranking on 48-bit hash codes.

the overall distributions are similar for the true label and the predicted label. However, because one image may contain multiple labels but is predicted with only one, the area of each predicted label category is observed to be a little smaller than that of the true label category. Meanwhile, we count the correctness of the predicted labels. The predicted labels for 8283 images fall within their true labels. These experimental results show that the proposed visual-semantic bridge has a high performance.

5) *Top Retrieval Results*: Fig. 15 shows some retrieval samples of some typical hashing methods according to the ascending Hamming ranking. The query image contains three semantic labels, i.e., building, sky, and water, with the main content of a building. We mark the mismatched image with the red box from the perspective of human vision. The retrieval results of T-MLZSH are better visual plausible while focusing on the main object of the query image, while other compared methods may return some mismatched results like the forest or return some images related to the less important part of the query image with higher ranking orders.

6) *Running Efficiency*: We do the image retrieval by returning the top 1000 similar images from 16900 images based on the trained models. The running time includes calculating the hash codes and calculating the hamming distances. In our experiments, when calculating the 36-bit hash codes of 16901 images, it takes 17.4616, 48.2418, and 34.6509 s for T-MLZSH-AlexNet, T-MLZSH-VGG16, and T-MLZSH-ResNet50, respectively, that is, it will take about 1.03, 2.85, and 2.05 ms for the three methods to calculate the hash codes of a new image. The hash codes are computed on an NVIDIA TITAN Xp GPU. The GPU memory usage is about 2.5, 4.3, and 1.2 GB for the three models in the retrieval test. For calculating the hamming distance of 36-bit hash codes, the running time on the 16900 images is 0.006742 s, which is about 4×10^{-7} s for one calculation. The hamming distance is computed by a 2.0 GHz core of an Intel Xeon E5-2620 CPU.

V. CONCLUSION

In this article, a novel T-MLZSH was proposed for multi-label image retrieval. It utilized the instance-concept coherence to construct a bridge for connecting the seen and unseen labels. Based on these connections, several categories

with the highest relatedness scores were selected as the predicted labels for target data. Then, hashing learning was performed on both the source data and target data in a supervised manner. Experimental results on three widely used multilabel data sets showed that the proposed method outperformed state-of-the-art methods with a significant margin. Moreover, the superiority of the proposed focal loss was verified by ablation studies, the effectiveness of the visual-semantic bridge was demonstrated through feature visualization, and the high performance on image retrieval was illustrated with visual comparisons.

ACKNOWLEDGMENT

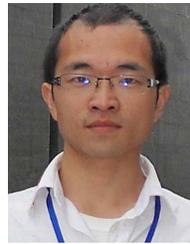
The authors would like to thank Dr. Song Bai from the University of Oxford for helpful suggestions.

REFERENCES

- [1] B. Kulis and K. Grauman, "Kernelized locality-sensitive hashing," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 34, no. 6, pp. 1092–1104, Jun. 2012.
- [2] W. Liu, J. Wang, R. Ji, Y.-G. Jiang, and S.-F. Chang, "Supervised hashing with kernels," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2012, pp. 2074–2081.
- [3] F. Shen, C. Shen, Q. Shi, A. van den Hengel, and Z. Tang, "Inductive hashing on manifolds," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2013, pp. 1562–1569.
- [4] F. Shen, C. Shen, W. Liu, and H. T. Shen, "Supervised discrete hashing," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2015, pp. 37–45.
- [5] H. Lai, Y. Pan, Y. Liu, and S. Yan, "Simultaneous feature learning and hash coding with deep neural networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2015, pp. 3270–3278.
- [6] R. Zhang, L. Lin, R. Zhang, W. Zuo, and L. Zhang, "Bit-scalable deep hashing with regularized similarity learning for image retrieval and person re-identification," *IEEE Trans. Image Process.*, vol. 24, no. 12, pp. 4766–4779, Dec. 2015.
- [7] Z. Zhang, Q. Zou, Y. Lin, L. Chen, and S. Wang, "Improved deep hashing with soft pairwise similarity for multi-label image retrieval," *IEEE Trans. Multimedia*, vol. 22, no. 2, pp. 540–553, Feb. 2020.
- [8] W. Liu, J. Wang, S. Kumar, and S.-F. Chang, "Hashing with graphs," in *Proc. ICML*, 2011, pp. 1–8.
- [9] M. Datar, N. Immorlica, P. Indyk, and V. S. Mirrokni, "Locality-sensitive hashing scheme based on p-stable distributions," in *Proc. 20th Annu. Symp. Comput. Geometry (SCG)*, 2004, pp. 253–262.
- [10] J.-P. Heo, Y. Lee, J. He, S.-F. Chang, and S.-E. Yoon, "Spherical hashing: Binary code embedding with hyperspheres," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 37, no. 11, pp. 2304–2316, Nov. 2015.
- [11] Y. Weiss, A. Torralba, and R. Fergus, "Spectral hashing," in *Proc. NIPS*, 2009, pp. 1753–1760.

- [12] V. E. Liong, J. Lu, G. Wang, P. Moulin, and J. Zhou, "Deep hashing for compact binary codes learning," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2015, pp. 2475–2483.
- [13] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Proc. NIPS*, 2012, pp. 1097–1105.
- [14] T. Y. Lin *et al.*, "Microsoft COCO: Common objects in context," in *Proc. ECCV*, 2014, pp. 740–755.
- [15] L. Chen, W. Zhan, W. Tian, Y. He, and Q. Zou, "Deep integration: A multi-label architecture for road scene recognition," *IEEE Trans. Image Process.*, vol. 28, no. 10, pp. 4883–4898, Oct. 2019.
- [16] Y. Sun, Y. Chen, X. Wang, and X. Tang, "Deep learning face representation by joint identification-verification," in *Proc. NIPS*, 2014, pp. 1988–1996.
- [17] L. Zhou, Z. Wang, Y. Luo, and Z. Xiong, "Separability and compactness network for image recognition and superresolution," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 30, no. 11, pp. 3275–3286, Nov. 2019.
- [18] Z. Shao, L. Wang, Z. Wang, W. Du, and W. Wu, "Saliency-aware convolution neural network for ship detection in surveillance video," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 30, no. 3, pp. 781–794, Mar. 2020.
- [19] L. Chen *et al.*, "Surrounding vehicle detection using an FPGA panoramic camera and deep CNNs," *IEEE Trans. Intell. Transp. Syst.*, vol. 21, no. 12, pp. 5110–5122, Dec. 2020.
- [20] J. Long, E. Shelhamer, and T. Darrell, "Fully convolutional networks for semantic segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2015, pp. 3431–3440.
- [21] Q. Zou, Z. Zhang, Q. Li, X. Qi, Q. Wang, and S. Wang, "DeepCrack: Learning hierarchical convolutional features for crack detection," *IEEE Trans. Image Process.*, vol. 28, no. 3, pp. 1498–1512, Mar. 2019.
- [22] R. Xia, Y. Pan, H. Lai, C. Liu, and S. Yan, "Supervised hashing for image retrieval via image representation learning," in *Proc. AAAI*, 2014, pp. 1–38.
- [23] F. Zhao, Y. Huang, L. Wang, and T. Tan, "Deep semantic ranking based hashing for multi-label image retrieval," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2015, pp. 1556–1564.
- [24] H. Zhu, M. Long, J. Wang, and Y. Cao, "Deep hashing network for efficient similarity retrieval," in *Proc. AAAI*, 2016, pp. 2415–2421.
- [25] H. Liu, R. Wang, S. Shan, and X. Chen, "Deep supervised hashing for fast image retrieval," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 2064–2072.
- [26] Z. Cao, M. Long, J. Wang, and P. S. Yu, "HashNet: Deep learning to hash by continuation," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 5609–5618.
- [27] C. Yan, B. Gong, Y. Wei, and Y. Gao, "Deep multi-view enhancement hashing for image retrieval," *IEEE Trans. Pattern Anal. Mach. Intell.*, early access, Feb. 24, 2020, doi: [10.1109/TPAMI.2020.2975798](https://doi.org/10.1109/TPAMI.2020.2975798).
- [28] C. H. Lampert, H. Nickisch, and S. Harmeling, "Learning to detect unseen object classes by between-class attribute transfer," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2009, pp. 951–958.
- [29] J. Pennington, R. Socher, and C. Manning, "Glove: Global vectors for word representation," in *Proc. Conf. Empirical Methods Natural Lang. Process. (EMNLP)*, 2014, pp. 1532–1543.
- [30] Y. Yang, Y. Luo, W. Chen, F. Shen, J. Shao, and H. T. Shen, "Zero-shot hashing via transferring supervised knowledge," in *Proc. ACM Multimedia Conf. (MM)*, 2016, pp. 1286–1295.
- [31] Y. Guo, G. Ding, J. Han, and Y. Gao, "SitNet: Discrete similarity transfer network for zero-shot hashing," in *Proc. 26th Int. Joint Conf. Artif. Intell.*, Aug. 2017, pp. 1767–1773.
- [32] S. Pachori, A. Deshpande, and S. Raman, "Hashing in the zero shot framework with domain adaptation," *Neurocomputing*, vol. 275, pp. 2137–2149, Jan. 2018.
- [33] H. Lai, "Transductive zero-shot hashing via coarse-to-fine similarity mining," in *Proc. ACM Int. Conf. Multimedia Retr.*, Jun. 2018, pp. 196–203.
- [34] H. Zhang, Y. Long, and L. Shao, "Zero-shot hashing with orthogonal projection for image retrieval," *Pattern Recognit. Lett.*, vol. 117, pp. 201–209, Jan. 2019.
- [35] C. Yan, L. Li, C. Zhang, B. Liu, Y. Zhang, and Q. Dai, "Cross-modality bridging and knowledge transferring for image understanding," *IEEE Trans. Multimedia*, vol. 21, no. 10, pp. 2675–2685, Oct. 2019.
- [36] C. Yan *et al.*, "A fast uyghur text detector for complex background images," *IEEE Trans. Multimedia*, vol. 20, no. 12, pp. 3389–3398, Dec. 2018.
- [37] C. Yan, B. Shao, H. Zhao, R. Ning, Y. Zhang, and F. Xu, "3D room layout estimation from a single RGB image," *IEEE Trans. Multimedia*, vol. 22, no. 11, pp. 3014–3024, Nov. 2020.
- [38] W. Wang, V. W. Zheng, H. Yu, and C. Miao, "A survey of zero-shot learning: Settings, methods, and applications," *ACM Trans. Intell. Syst. Technol.*, vol. 10, no. 2, pp. 1–37, 2019.
- [39] X. Xu, T. Hospedales, and S. Gong, "Transductive zero-shot action recognition by word-vector embedding," *Int. J. Comput. Vis.*, vol. 123, no. 3, pp. 309–333, Jul. 2017.
- [40] A. Lazaridou, E. Bruni, and M. Baroni, "Is this a wampimuk? cross-modal mapping between distributional semantics and the visual world," in *Proc. 52nd Annu. Meeting Assoc. Comput. Linguistics*, vol. 1, 2014, pp. 1403–1414.
- [41] T. Mikolov, Q. V. Le, and I. Sutskever, "Exploiting similarities among languages for machine translation," 2013, *arXiv:1309.4168*. [Online]. Available: <http://arxiv.org/abs/1309.4168>
- [42] A. Frome *et al.*, "Devise: A deep visual-semantic embedding model," in *Proc. NIPS*, 2013, pp. 2121–2129.
- [43] Z. Akata, S. Reed, D. Walter, H. Lee, and B. Schiele, "Evaluation of output embeddings for fine-grained image classification," in *Proc. CVPR*, Jun. 2015, pp. 2927–2936.
- [44] C. Gan, Y. Yang, L. Zhu, D. Zhao, and Y. Zhuang, "Recognizing an action using its name: A knowledge-based approach," *Int. J. Comput. Vis.*, vol. 120, no. 1, pp. 61–77, Oct. 2016.
- [45] Y. Guo, G. Ding, J. Han, and Y. Gao, "Zero-shot learning with transferred samples," *IEEE Trans. Image Process.*, vol. 26, no. 7, pp. 3277–3290, Jul. 2017.
- [46] Y. Guo, G. Ding, J. Han, and Y. Gao, "Zero-shot recognition via direct classifier learning with transferred samples and pseudo labels," in *Proc. AAAI*, 2017, pp. 1–7.
- [47] Y. Xian, T. Lorenz, B. Schiele, and Z. Akata, "Feature generating networks for zero-shot learning," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 5542–5551.
- [48] Y. Zhu, M. Elhoseiny, B. Liu, X. Peng, and A. Elgammal, "A generative adversarial approach for zero-shot learning from noisy texts," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 1004–1013.
- [49] Y. Guo, G. Ding, J. Han, and Y. Gao, "Synthesizing samples for zero-shot learning," in *Proc. 26th Int. Joint Conf. Artif. Intell.*, Aug. 2017, pp. 1774–1780.
- [50] Y. Zhu, J. Xie, B. Liu, and A. Elgammal, "Learning feature-to-feature translator by alternating back-propagation for generative zero-shot learning," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 9844–9854.
- [51] M. Ye and Y. Guo, "Progressive ensemble networks for zero-shot recognition," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 11728–11736.
- [52] A. Paul, N. C. Krishnan, and P. Munjal, "Semantically aligned bias reducing zero shot learning," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 7056–7065.
- [53] D. Mandal *et al.*, "Out-of-distribution detection for generalized zero-shot action recognition," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 9985–9993.
- [54] Z. Ren, H. Jin, Z. Lin, C. Fang, and A. Yuille, "Multi-instance visual-semantic embedding," 2015, *arXiv:1512.06963*. [Online]. Available: <http://arxiv.org/abs/1512.06963>
- [55] T. Mensink, E. Gavves, and C. G. M. Snoek, "COSTA: Co-occurrence statistics for zero-shot classification," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2014, pp. 2441–2448.
- [56] G. Irie, Z. Li, X.-M. Wu, and S.-F. Chang, "Locally linear hashing for extracting non-linear manifolds," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2014, pp. 2115–2122.
- [57] A. Torralba, R. Fergus, and Y. Weiss, "Small codes and large image databases for recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2008, pp. 1–8.
- [58] Y. Weiss, R. Fergus, and A. Torralba, "Multidimensional spectral hashing," in *Proc. ECCV*, 2012, pp. 340–353.
- [59] Q. Ning, J. Zhu, Z. Zhong, S. C. H. Hoi, and C. Chen, "Scalable image retrieval by sparse product quantization," *IEEE Trans. Multimedia*, vol. 19, no. 3, pp. 586–597, Mar. 2017.
- [60] S. Ercoli, M. Bertini, and A. D. Bimbo, "Compact hash codes for efficient visual descriptors retrieval in large scale databases," *IEEE Trans. Multimedia*, vol. 19, no. 11, pp. 2521–2532, Nov. 2017.
- [61] S. Xia *et al.*, "A fast adaptive k-means with no bounds," *IEEE Trans. Pattern Anal. Mach. Intell.*, early access, Jul. 13, 2020, doi: [10.1109/TPAMI.2020.3008694](https://doi.org/10.1109/TPAMI.2020.3008694).

- [62] L. Liu and L. Shao, "Sequential compact code learning for unsupervised image hashing," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 27, no. 12, pp. 2526–2536, Dec. 2016.
- [63] L. Zhu, Z. Huang, Z. Li, L. Xie, and H. T. Shen, "Exploring auxiliary context: Discrete semantic transfer hashing for scalable image retrieval," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 29, no. 11, pp. 5264–5276, Nov. 2018.
- [64] J. Gui, T. Liu, Z. Sun, D. Tao, and T. Tan, "Supervised discrete hashing with relaxation," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 29, no. 3, pp. 608–617, Mar. 2018.
- [65] Y. Cao, M. Long, B. Liu, and J. Wang, "Deep cauchy hashing for Hamming space retrieval," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 1229–1237.
- [66] D. Wang, Q. Wang, and X. Gao, "Robust and flexible discrete hashing for cross-modal similarity search," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 28, no. 10, pp. 2703–2715, Oct. 2018.
- [67] D. Wang, X. Gao, X. Wang, and L. He, "Label consistent matrix factorization hashing for large-scale cross-modal similarity search," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 41, no. 10, pp. 2466–2479, Oct. 2019.
- [68] D. Tian, Y. Wei, and D. Zhou, "Learning decorrelated hashing codes with label relaxation for multimodal retrieval," *IEEE Access*, vol. 8, pp. 79260–79272, 2020.
- [69] S. Su, Z. Zhong, and C. Zhang, "Deep joint-semantics reconstructing hashing for large-scale unsupervised cross-modal retrieval," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 3027–3035.
- [70] D. Wang, X. Gao, X. Wang, and L. He, "Semantic topic multimodal hashing for cross-media retrieval," in *Proc. IJCAI*, 2015, pp. 3890–3896.
- [71] D. Wang, X. Gao, X. Wang, L. He, and B. Yuan, "Multimodal discriminative binary embedding for large-scale cross-modal retrieval," *IEEE Trans. Image Process.*, vol. 25, no. 10, pp. 4540–4554, Oct. 2016.
- [72] H. Zhang, L. Liu, Y. Long, and L. Shao, "Unsupervised deep hashing with pseudo labels for scalable image retrieval," *IEEE Trans. Image Process.*, vol. 27, no. 4, pp. 1626–1638, Apr. 2018.
- [73] F. Shen, Y. Xu, L. Liu, Y. Yang, Z. Huang, and H. T. Shen, "Unsupervised deep hashing with similarity-adaptive and discrete optimization," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 40, no. 12, pp. 3034–3044, Dec. 2018.
- [74] S. Jialin Pan and Q. Yang, "A survey on transfer learning," *IEEE Trans. Knowl. Data Eng.*, vol. 22, no. 10, pp. 1345–1359, Oct. 2010.
- [75] Y. Fu, T. M. Hospedales, T. Xiang, and S. Gong, "Transductive multi-view zero-shot learning," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 37, no. 11, pp. 2332–2345, 2015.
- [76] H. Venkateswara, J. Eusebio, S. Chakraborty, and S. Panchanathan, "Deep hashing network for unsupervised domain adaptation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 5018–5027.
- [77] Y. Shen, L. Liu, F. Shen, and L. Shao, "Zero-shot sketch-image hashing," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 3598–3607.
- [78] Q. Wang and K. Chen, "Multi-label zero-shot human action recognition via joint latent ranking embedding," 2017, *arXiv:1709.05107*. [Online]. Available: <http://arxiv.org/abs/1709.05107>
- [79] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollar, "Focal loss for dense object detection," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 42, no. 2, pp. 318–327, Feb. 2020.
- [80] T.-S. Chua, J. Tang, R. Hong, H. Li, Z. Luo, and Y. Zheng, "NUS-WIDE: A real-world Web image database from National University of Singapore," in *Proc. ICIVR*, 2009, p. 48.
- [81] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman, "The Pascal visual object classes (VOC) challenge," *Int. J. Comput. Vis.*, vol. 88, no. 2, pp. 303–338, Jun. 2010.
- [82] Q. Jiang and W. Li, "Asymmetric deep supervised hashing," in *Proc. AAAI*, 2018, pp. 3342–3349.
- [83] D. Wu, Q. Dai, J. Liu, B. Li, and W. Wang, "Deep incremental hashing network for efficient image retrieval," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 9069–9077.
- [84] Y. Cao, L. Ju, Q. Zou, C. Qu, and S. Wang, "A multichannel edge-weighted centroidal Voronoi tessellation algorithm for 3D super-alloy image segmentation," in *Proc. CVPR*, Jun. 2011, pp. 17–24.
- [85] K. Järvelin and J. Kekäläinen, "Cumulated gain-based evaluation of IR techniques," *ACM Trans. Inf. Syst.*, vol. 20, no. 4, pp. 422–446, Oct. 2002.



Qin Zou (Senior Member, IEEE) received the B.E. degree in information engineering and the Ph.D. degree in computer vision from Wuhan University, Wuhan, China, in 2004 and 2012, respectively.

From 2010 to 2011, he was a Visiting Ph.D. Student with the Computer Vision Lab, University of South Carolina, Columbia, SC, USA. He is currently an Associate Professor with the School of Computer Science, Wuhan University. His research activities involve computer vision, pattern recognition, and machine learning.

Dr. Zou is a member of the Association for Computing Machinery (ACM). He was a co-recipient of the National Technology Invention Award of China in 2015.



Ling Cao received the B.S. degree in computer science from Yunnan University, Kunming, China, in 2018. She is currently pursuing M.S. degree in computer application with the School of Computer Science, Wuhan University, Wuhan, China.

Her research interests include deep learning and image/video retrieval.



Zheng Zhang received the B.S. degree in computer science and the M.S. degree in computer application from the School of Computer Science, Wuhan University, Wuhan, China, in 2016 and 2019, respectively.

His research interests include deep learning and its application in image classification and retrieval.

Dr. Zhang received the first prize in the China Undergraduate Contest in Internet of Things in 2015.



Long Chen (Senior Member, IEEE) received the B.Sc. degree in communication engineering and the Ph.D. degree in signal and information processing from Wuhan University, Wuhan, China, in 2007 and 2013, respectively.

From October 2010 to November 2012, he was co-trained Ph.D. Student at the National University of Singapore, Singapore. He is currently an Associate Professor with the School of Data and Computer Science, Sun Yat-sen University, Guangzhou, China. His areas of interest include deep learning and cognitive perception.



Song Wang (Senior Member, IEEE) received the Ph.D. degree in electrical and computer engineering from the University of Illinois at Urbana-Champaign (UIUC), Champaign, IL, USA, in 2002.

From 1998 to 2002, he also worked as a Research Assistant with the Image Formation and Processing Group, Beckman Institute, UIUC. In 2002, he joined the Department of Computer Science and Engineering, University of South Carolina, Columbia, SC, USA, where he is currently a Professor. His research interests include computer vision, medical image processing, and machine learning.

Dr. Wang is a Senior Member of the IEEE Computer Society. He is also serving as an Associate Editor of IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE and an Associate Editor of *Pattern Recognition Letters*.