

Deep Semantic Hashing with Generative Adversarial Networks*

Zhaofan Qiu

University of Science and Technology of China
Hefei, China
zhaofanqiu@gmail.com

Ting Yao

Microsoft Research Asia
Beijing, China
tiyao@microsoft.com

Yingwei Pan

University of Science and Technology of China
Hefei, China
panyw.ustc@gmail.com

Tao Mei

Microsoft Research Asia
Beijing, China
tmei@microsoft.com

ABSTRACT

Hashing has been a widely-adopted technique for nearest neighbor search in large-scale image retrieval tasks. Recent research has shown that leveraging supervised information can lead to high quality hashing. However, the cost of annotating data is often an obstacle when applying supervised hashing to a new domain. Moreover, the results can suffer from the robustness problem as the data at training and test stage may come from different distributions. This paper studies the exploration of generating synthetic data through semi-supervised generative adversarial networks (GANs), which leverages largely unlabeled and limited labeled training data to produce highly compelling data with intrinsic invariance and global coherence, for better understanding statistical structures of natural data. We demonstrate that the above two limitations can be well mitigated by applying the synthetic data for hashing. Specifically, a novel deep semantic hashing with GANs (DSH-GANs) is presented, which mainly consists of four components: a deep convolution neural networks (CNN) for learning image representations, an adversary stream to distinguish synthetic images from real ones, a hash stream for encoding image representations to hash codes and a classification stream. The whole architecture is trained end-to-end by jointly optimizing three losses, i.e., adversarial loss to correct label of synthetic or real for each sample, triplet ranking loss to preserve the relative similarity ordering in the input real-synthetic triplets and classification loss to classify each sample accurately. Extensive experiments conducted on both CIFAR-10 and NUS-WIDE image benchmarks validate

*This work was performed at Microsoft Research Asia. The first two authors made equal contributions to the work.

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.

SIGIR '17, August 07-11, 2017, Shinjuku, Tokyo, Japan

© 2017 Association for Computing Machinery.

ACM ISBN 978-1-4503-5022-8/17/08...\$15.00

<https://doi.org/10.1145/3077136.3080842>

the capability of exploiting synthetic images for hashing. Our framework also achieves superior results when compared to state-of-the-art deep hash models.

CCS CONCEPTS

• Information systems → Similarity measures; Learning to rank;

KEYWORDS

Hashing; Similarity Learning; GANs; CNN

1 INTRODUCTION

Accelerated by tremendous increase in Internet bandwidth and storage space, multimedia data have been generated, published and spread explosively. This has led to the surge of research activities in large scale visual search. One fundamental research problem is similarity search, i.e., nearest neighbor search, which attempts to identify similar instances according to a query example. The need to search for millions of visual examples in a high-dimensional feature space, however, makes the task computationally expensive and thus very challenging.

Hashing techniques, one direction of the most well-known Approximate Nearest Neighbor (ANN) search methods, have been studied extensively due to its great efficiency in gigantic data. The basic idea of hashing is to construct a series of hash functions to map each example into compact binary codes, making the Hamming distances on similar examples minimized and simultaneously maximized on dissimilar examples. In the literature, there have been several techniques, including traditional hashing models based on hand-crafted features [3, 4, 14, 25] and deep models [11, 13], being proposed for addressing the problem of hashing. The former seek hashing function on hand-crafted features, which separate the encoding of feature representations and their quantization to hash codes, resulting in sub-optimal solution. The latter jointly learn feature representations and projections from them to hash codes in a deep architecture. While encouraging performances are reported in the aforementioned approaches especially when supervised information is available, we are often facing the problems of applying these methods to new

applications where there is only few labeled training data, not to mention that the distribution of training data may be even different with that in test stage.

We demonstrate in this paper that the above limitations can be mitigated by generating synthetic data for training through Generative Adversarial Networks (GANs). GANs is a new recently proposed framework for estimating generative models via an adversarial process. The spirit behind is a minimax two-player game, in which a generative model is to capture the data distribution and a discriminative model aims to estimate the probability that a sample is from the real training data rather than the generative model. The generative model and discriminative model are trained simultaneously and the learning of the generative model is to fool the discriminative model into making mistakes. Once the training is complete, GANs is capable of generating both diverse and discriminable training examples, which have a great potential to characterize the statistical structures of natural data.

By consolidating the idea of generating training data for boosting hashing, we present a novel Deep Semantic Hashing with GANs (DSH-GANs) architecture, as shown in Figure 1. Specifically, a semi-supervised GANs is first pre-trained on both labeled and unlabeled training data to produce synthetic examples conditioning on class labels. Then, we form a set of real-synthetic triplets and each tuple contains one real image as query image, one synthetic and semantically similar image and another synthetic but dissimilar image. A shared CNN is exploited to capture image representations, followed by importing into an adversary stream for differentiating the synthetic images from real ones, a hash stream to encode hash codes and a classification stream for measuring semantics. An adversarial loss is computed to correct the predicted labels (i.e., synthetic or real) of the images in adversary stream and a triplet ranking loss is devised to preserve relative similarities at the top of hash stream. Meanwhile, a classification error is formulated in classification stream. By jointly learning the three streams, our DSH-GANs is expected to offer a hashing model with high generalization ability and the generated hash codes could better reflect semantic relations between images. It is also worth noting that the whole architecture is trainable in an end-to-end fashion.

In summary, this paper makes the following contributions:

- (1) We explore the problem of supervised hashing by exploiting the synthetic training data from GANs. To the best of our knowledge, this paper represents the first effort towards this target in the information retrieval research community.
- (2) A novel hashing architecture, which combines adversary process, hash coding and classification, is proposed to enhance the generalization ability of hashing model and produce hash codes, which preserve not only relative similarity between images but also semantics of images.
- (3) Extensive experiments on two widely used datasets demonstrate the advantages of our proposal over several state-of-the-art hashing techniques.

2 RELATED WORK

We briefly group the related works into two directions: hashing for image search and image synthesis with Generative Adversarial Networks (GANs). The former draws upon research in encoding visual images into compact binary codes for efficient image search, while the latter investigates synthesizing realistic images by utilizing GANs.

Hashing for Image Search. The research in this direction has proceeded along two dimensions: hand-crafted features based hashing and deep architectures for hashing.

There are three main categories on hand-crafted features based hashing: unsupervised hashing, semi-supervised hashing and supervised hashing. Unsupervised hashing [3, 4] refers to the setting when the label information is not available. Locality Sensitive Hashing (LSH) [3] is one of the most popular unsupervised hashing methods, which simply uses random linear projections to construct hash functions. This method is subsequently expanded to Kernelized and Multi-Kernel Locality Sensitive Hashing [10, 27]. Another effective method named Iterative Quantization (ITQ) [4] is proposed for better quantization rather than random projections. Semi-supervised hashing approaches attempt to improve the quality of hash codes by leveraging supervised information into learning procedure. For example, Wang *et al.* develop a Semi-Supervised Hashing (SSH) [25] which utilizes pairwise information on labeled samples to preserve semantic similarity. In another work [8], Semi-Supervised Discriminant Hashing (SSDH) learns hash codes based on Fisher’s discriminant analysis to maximize separability between labeled data from different classes while the unlabeled data are exploited for regularization. When the label information is all available, we refer to the problem as supervised hashing. The representative in this category is Kernel-based Supervised Hashing (KSH) [14] which utilizes pairwise relationship between examples to achieve high quality hashing.

Inspired by recent advances in visual representation learning [9, 17, 19] by using deep convolutional neural networks, several deep architecture based hashing methods have been proposed. Semantic Hashing [23] is one of the early works to exploit deep learning techniques for hashing. It applies the stacked Restricted Boltzman Machine (RBM) [6] to learn binary hash codes for visual search. Recently, Xia *et al.* propose Convolutional Neural Networks Hashing (CNNH) [28] to decompose the hash learning process into a stage of learning approximate hash codes with the pairwise relationship and a following stage of simultaneously learning image feature and hash function. Later in [12], such a two-stage method with pairwise labels is further developed into an end-to-end system, Deep Pairwise-Supervised Hashing (DPSH), which performs simultaneous feature learning and hash encoding. Similar in spirit, Network In Network Hashing (NINH) [11] incorporates the supervised information among triplet labels into the feature learning based deep hashing architecture. More recently, Zhu *et al.* devise Deep Hashing Network (DHN) to simultaneously optimize the pairwise cross-entropy loss on

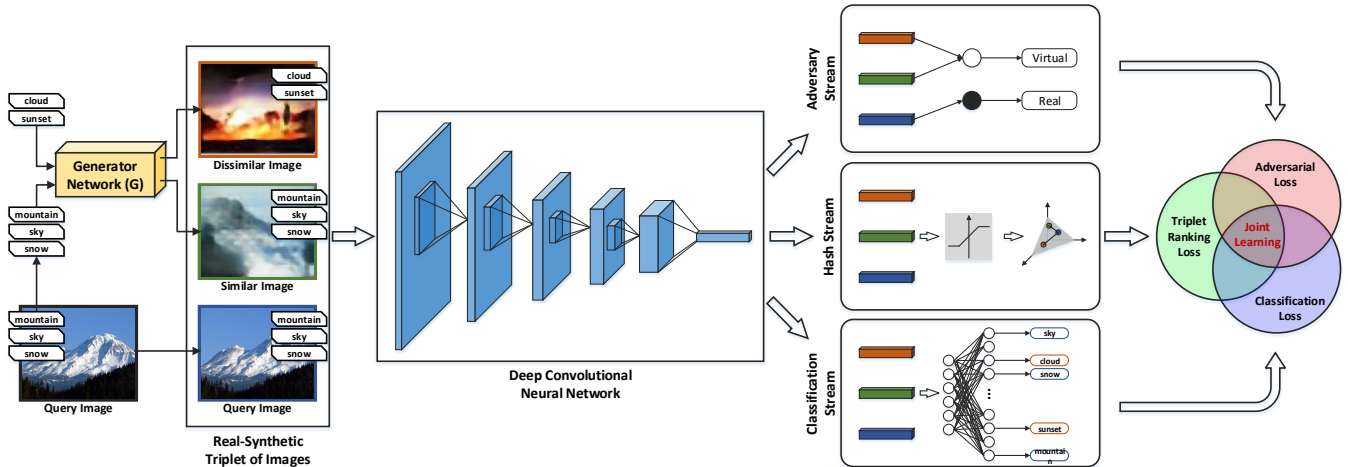


Figure 1: Deep Semantic Hashing with GANs (DSH-GANs) framework (better viewed in color). The input to DSH-GANs architecture is in the form of real-synthetic image triplets and each tuple consists of one real image as query image, one synthetic and similar image produced with same labels of query image through generator network G , and another synthetic but dissimilar image synthesized by G conditioning on different labels. A shared deep convolutional neural networks is exploited for learning image representations, followed by three streams, i.e., hash stream, adversary stream and classification stream. Hash stream is to encode each image into hash codes with relative similarity preservation measured by a triplet ranking loss. Adversary stream is to distinguish synthetic images from real ones trained with an adversarial loss. Classification stream is to characterize the semantic structures on image and softmax loss or cross entropy loss is computed for single label and multi-label classification, respectively. The whole architecture is jointly optimized in an end-to-end fashion.

semantically similar pairs and the pairwise quantization loss on compact hash codes for hashing in [30].

In summary, our work belongs to deep architecture based hashing. The aforementioned deep approaches often focus on leveraging supervised information for training CNNs. Our work in this paper contributes by not only exploring image semantic supervision for hash learning, but also preserving relative similarity between real and synthetic images which are generated through a semi-supervised GANs with intrinsic invariance and global coherence.

Image Synthesis with GANs. Synthesizing realistic images has been studied and analyzed widely in AI systems for characterizing the pixel level structure of natural images. Thanks to the recent development of Generative Adversarial Networks (GANs), researchers have strived to automatically synthesize image with GANs, which could be regarded as the generator network modules learnt with a two-player minimax game mechanism. Goodfellow *et al.* propose a theoretical framework of GANs and utilize GANs to generate images without any supervised information in [5]. Although the early GANs offer a distinct and promising direction for image synthesis, the results are somewhat noisy and blurry. Hence, Laplacian pyramid is further incorporated into GANs in [2] to produce high quality images. Later in [20], Radford *et al.* devise deep convolutional generative adversarial networks (DCGANs) for unsupervised representation learning.

The aforementioned three works mainly explore image synthesis task in an unconditioned manner that generates

synthetic images without any supervised information. Another direction of image synthesis with GANs is to synthesize images by conditioning on supervised information (e.g., class labels or text descriptions). [15] is one of the works that develop a conditional version of GANs by additionally feeding class labels into both discriminator and generator of GANs. Later in [16], this model is further expended with a specialized cost function for classification, named auxiliary classifier GANs (AC-GANs), for generating synthetic images with global coherence and high diversity. Recently, Reed *et al.* utilize GANs for image synthesis based on given text descriptions in [21], enabling translation from character level to pixel level.

Most of the above approaches focus on leveraging GANs for image synthesis. Our work is different that we apply the synthetic images generated from GANs learnt on both largely unlabeled and limited labeled images for hash learning, leading to more effective and robust binary image representation for image retrieval task.

3 DEEP SEMANTIC HASHING WITH GANs (DSH-GANs)

In this section, we will present the proposed Deep Semantic Hashing with GANs (DSH-GANs) in detail. Figure 1 illustrates an overview of our architecture for hash learning, which consists of four components: a shared CNN for learning image representations, an adversary stream for distinguishing synthetic images from real ones, a hash stream for encoding each

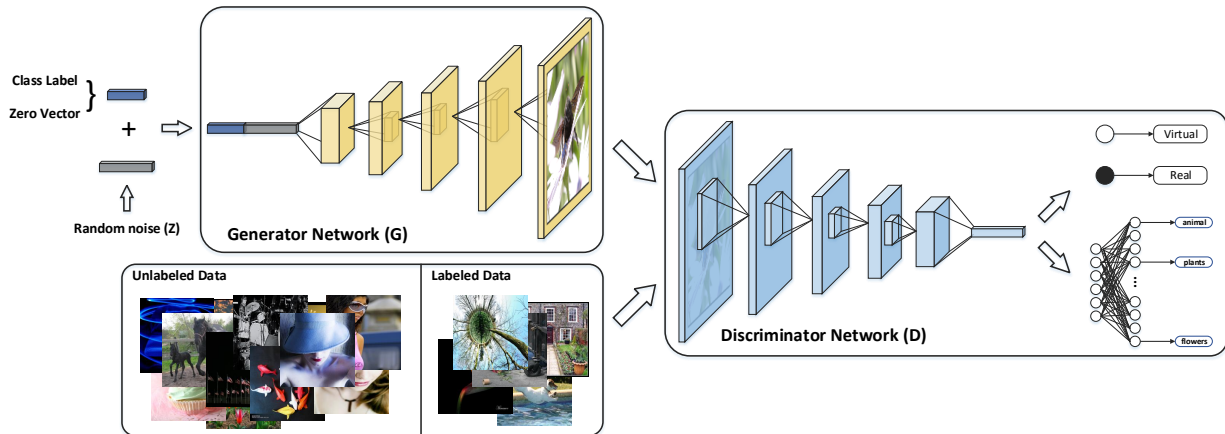


Figure 2: Our semi-supervised GANs framework mainly consists of a generator network G and a discriminator network D (better viewed in color). For the generator network G , it tries to synthesize realistic images with the concatenation input of the class label vector C and random noise vector z . For the discriminator network D , it tries to simultaneously distinguish real images from synthetic ones and classify input images with correct class labels. The whole architecture is trained with the adversarial loss for assigning correct source and the classification loss for assigning correct class label in a two-player minimax game mechanism.

image into hash codes and a classification stream for leveraging semantic supervision. Specifically, a semi-supervised GANs is first devised to leverage both unlabeled and labeled images for producing synthetic images conditioning on class labels, followed by the three streams in our proposed DSH-GANs framework. In particular, hash stream is trained with the input real-synthetic triplets in a triplet-wise manner, adversary stream recognizes the label of synthetic or real for each image example while classification stream reinforces the hash learning to preserve semantic structures on both real and synthetic images. Finally, the whole optimization of DSH-GANs and hash codes generation for image retrieval are elaborated.

3.1 Notation

Suppose there are n images in the whole set, represented as: $\mathcal{X} = \{x_i | i = 1, \dots, n\}$ and each image can be presented as x . Similarly, assume there are L ($L < n$) labeled images and the set of the labeled images are denoted as $\mathcal{X}_l = \{x_j | j = 1, \dots, L\}$. The goal of image hashing is to learn a mapping $\mathcal{H} : x \rightarrow \{0, 1\}^K$, such that an input image x will be encoded into a K -bit binary code $\mathcal{H}(x)$.

3.2 Semi-supervised GANs

An unconditional generative adversarial networks (GANs) consists of two networks: a generator network G that captures the data distribution for synthesizing image and a discriminator network D that distinguishes real images from synthetic ones. In particular, the generator network G takes a random noise vector \mathbf{z} as input and produces a synthetic image $x_{syn} = G(\mathbf{z})$. For the discriminator network D , it takes an image x as input stochastically chosen (with equal probability) from training real images or synthetic images through G and produces a probability distribution $P(S|x) = D(x)$ over

the two image sources. As proposed in [5], the whole GANs can be trained in a two-player minimax game. Concretely, given an image sample x , the discriminator network D is trained to minimize the adversarial loss, i.e., maximizing the log-likelihood of assigning correct source to this sample:

$$l_a(x) = \begin{cases} -\log P(S = real|x), x \in \mathcal{X} \\ -\log P(S = synthetic|x), x \in \mathcal{X}_{syn} \end{cases}, \quad (1)$$

where \mathcal{X} and \mathcal{X}_{syn} denote the collections of real images in training and synthetic images produced by G , respectively. Meanwhile, the generator network G is trained to maximize the adversarial loss in Eq.(1), targeting for maximally fooling the discriminator network D with its generated synthetic images \mathcal{X}_{syn} .

To characterize the pixel level structure of both unlabeled and labeled natural images in one architecture elegantly, we take the inspiration from conditional GANs [15, 16] purely trained with supervised samples and devise a novel semi-supervised GANs architecture as shown in Figure 2. Similar to aforementioned architectures of unconditional GANs, our semi-supervised GANs consists of a generator network G for synthesizing images conditioning on class labels, and a discriminator network D that simultaneously distinguishes real images from synthetic ones and classify the input images with correct semantics. Specifically, given the whole image set \mathcal{X} including L labeled images in c classes, the class label information of each labeled image is first encoded into a c -dimensional vector $\mathbf{C} \in \{0, 1\}^c$, whose element is a class label indicator. The indicator is 1 if the image contains this label otherwise the indicator is 0. As such, the class label vector \mathbf{C} of each unlabeled image is set as zero vector $\mathbf{0}$. Then the generator network G takes the concatenation of the class label vector \mathbf{C} and random noise vector $\mathbf{z} \in \mathcal{N}(0, 1)$ as the input for producing a synthetic image $x_{syn} = G(\mathbf{C}, \mathbf{z})$. The discriminator network D generates both a probability

distribution over two sources and a probability distribution over all the c class labels, i.e., $\{P(S|x), P(C|x)\} = D(x)$, for each image example x from either real images or synthetic images through G . It is worth noting that both the unlabeled and labeled images are included in the real image selection pool of D for better understanding the statistical structures of natural data.

The overall objective function of our semi-supervised GANs is composed of two parts: the adversarial loss $l_a(x)$ in Eq.(1) for assigning correct source to the image example x , and the classification loss $l_c(x)$ for assigning correct class label to this image. The details of how to measure the classification loss for images with single label or multiple labels will be presented in Section 3.5. Accordingly, the discriminator network D is learnt to minimize $l_c(x) + l_a(x)$ for recognizing both correct source and class label, while the generator network G is trained to minimize $l_c(x) - l_a(x)$ for fooling D on source prediction and meanwhile preserving the correct class label. After training the whole semi-supervised GANs with unlabeled and labeled natural images, the learnt generator network G is directly utilized as the pre-trained generator network in our DSH-GANs architecture for synthesizing realistic images conditioning on class labels.

3.3 Hash Stream

In the traditional binary representation learning, the hash encoding of each image is always treated independently in point-wise hashing learning methods, regardless of the relationships of similar or dissimilar between images. More importantly, the relative similarity relations like “for query image x , it should be more similar to image x^+ than to image x^- ,” are reflected in the image class labels in view that image x and x^+ belong to the same class while image x^- comes from other categories. The utilization of these relative similarity relations has also been proved to be effective in hash coding [1, 11, 18, 29]. Inspired by the idea of preserving relative similarity in deep architecture [11], we propose a hash stream for encoding hash codes learnt in a triplet-wise manner, which aims to preserve the relative similarity ordering in the input real-synthetic triplets.

Specifically, we can easily obtain a set of real-synthetic triplets \mathcal{T} based on image labels, where each tuple $(x, x_{syn}^+, x_{syn}^-)$ consists of one real image x as query image, one synthetic and semantically similar image x_{syn}^+ , and another synthetic but dissimilar image x_{syn}^- . Note that x_{syn}^+ is synthesized by generator network G conditioning on the same class labels of query image x , while x_{syn}^- is produced through G conditioning on different labels of x . To preserve the similarity relations in the real-synthetic triplets, we aim to learn a hash mapping $\mathcal{H}(\cdot)$ which makes the compact code $\mathcal{H}(x)$ more similar to $\mathcal{H}(x_{syn}^+)$ than to $\mathcal{H}(x_{syn}^-)$. Hence, the triplet ranking loss is employed and defined as

$$\begin{aligned} \hat{l}_{triplet}(x, x_{syn}^+, x_{syn}^-) \\ = \max(0, 1 - \|\mathcal{H}(x) - \mathcal{H}(x_{syn}^-)\|_H + \|\mathcal{H}(x) - \mathcal{H}(x_{syn}^+)\|_H), \\ \text{s.t. } \mathcal{H}(x), \mathcal{H}(x_{syn}^+), \mathcal{H}(x_{syn}^-) \in \{0, 1\}^K \end{aligned} \quad (2)$$

where $\|\cdot\|_H$ represents Hamming distance. For ease of optimization, natural relaxation tricks are utilized on Eq.(2) to change integer constraint to the range constraint and replace Hamming norm with l_2 norm. Then, the triplet ranking loss function is reformulated as

$$\begin{aligned} \hat{l}_{triplet}(x, x_{syn}^+, x_{syn}^-) \\ = \max(0, 1 - \|\mathcal{H}(x) - \mathcal{H}(x_{syn}^-)\|_2^2 + \|\mathcal{H}(x) - \mathcal{H}(x_{syn}^+)\|_2^2). \end{aligned} \quad (3)$$

s.t. $\mathcal{H}(x), \mathcal{H}(x_{syn}^+), \mathcal{H}(x_{syn}^-) \in [0, 1]^K$

3.4 Adversary Stream

Noticing that the input real-synthetic triplets of aforementioned hash stream contain not only different semantics, but also are from distinctly different sources. As a result, we additionally devise an adversary stream to distinguish synthetic images from real ones within each real-synthetic triplet, targeting for exploiting the mutual but also fuzzy relationship between the hash codes learning and source discrimination in GANs. In particular, for the adversary stream, the shared CNN for learning image representation can be treated as the discriminator network D in GANs, followed by a cross entropy loss layer for source prediction. Thus, given the real-synthetic triplet $(x, x_{syn}^+, x_{syn}^-)$, an adversarial loss is used to measure the correctness of the predicted source (i.e., real or synthetic) of all the three images:

$$\hat{l}_a(x, x_{syn}^+, x_{syn}^-) = \frac{1}{3} (l_a(x) + l_a(x_{syn}^+) + l_a(x_{syn}^-)), \quad (4)$$

where $l_a(\cdot)$ denotes the log-likelihood adversarial loss for each image as in Eq.(1).

3.5 Classification Stream

Image labels not only provide knowledge in classification but also are useful supervised information for mining semantic structures in images. A valid question is how to leverage the semantic supervision into both hashing and GANs, and make the generated hash codes better reflecting semantic similarities between images. Hence, we propose a joint learning mechanism by combining hash stream, adversary stream and classification stream. In the classification stream, a classification error is measured based on the input real-synthetic triplets. Specifically, for the single label classification, we use softmax optimization method. Given an input image x , the softmax loss is then formulated as

$$l_c(x) = - \sum_{j=1}^c I_{(y=j)} \log \frac{e^{\theta_j^\top \mathbf{x}}}{\sum_{l=1}^c e^{\theta_l^\top \mathbf{x}}}, \quad (5)$$

where \mathbf{x} is the output image representation of the shared CNN for image x , θ_j denotes the parameter matrix in a softmax layer and $y \in \{1, 2, \dots, c\}$ represents image class label. The indicator function $I_{condition} = 1$ if *condition* is true; otherwise $I_{condition} = 0$.

If an image contains multiple class labels, we refer to this problem as multi-label classification. Cross entropy loss is then employed in this case. Similar to softmax loss, cross

entropy loss is computed by

$$l_c(x) = - \sum_{j=1}^c [I_{(\mathbf{C}_j=1)} \log(P(\mathbf{C}_j = 1|\mathbf{x})) + (1 - I_{(\mathbf{C}_j=1)}) \log(1 - P(\mathbf{C}_j = 1|\mathbf{x}))], \quad (6)$$

$$P(\mathbf{C}_j = 1|\mathbf{x}) = \frac{1}{1 + e^{-\delta_j^T \mathbf{x}}}$$

where \mathbf{C}_j denotes the j -th element in class label vector \mathbf{C} and δ_j denotes the parameter matrix in a sigmoid layer.

Hence, given the real-synthetic triplet $(x, x_{syn}^+, x_{syn}^-)$, the classification error is calculated on all the three examples by

$$\hat{l}_c(x, x_{syn}^+, x_{syn}^-) = \frac{1}{3} (l_c(x) + l_c(x_{syn}^+) + l_c(x_{syn}^-)). \quad (7)$$

3.6 Optimization

The overall training objective of DSH-GANs integrates the triplet ranking loss in Eq.(3), adversarial loss in Eq.(4) and classification error in Eq.(7). As our DSH-GANs is a variant of GANs architecture which mainly consists of generator network G for image synthesis with labels and the shared CNN for image representation learning, we train the whole architecture in a two-player minimax game mechanism. In particular, for the shared CNN, we update its parameters according to the following overall loss:

$$\hat{l}_{CNN} = \sum_{\mathcal{T}} [\hat{l}_{triplet}(x, x_{syn}^+, x_{syn}^-) + \hat{l}_a(x, x_{syn}^+, x_{syn}^-) + \hat{l}_c(x, x_{syn}^+, x_{syn}^-)], \quad (8)$$

where \mathcal{T} is the set of real-synthetic triplets. By minimizing this term, the shared CNN is trained to preserve the relative similarity ordering in the real-synthetic triplets and simultaneously recognize both correct sources and class labels of images in the triplets.

For the generator network G , its parameters are adjusted with the following loss:

$$\hat{l}_G = \sum_{\mathcal{T}} [\hat{l}_{triplet}(x, x_{syn}^+, x_{syn}^-) - \hat{l}_a(x, x_{syn}^+, x_{syn}^-) + \hat{l}_c(x, x_{syn}^+, x_{syn}^-)]. \quad (9)$$

Thus, the generator network G is trained to fool the shared CNN on source prediction and meanwhile preserve the relative similarity ordering and correct class labels of the real-synthetic triplets.

3.7 Image Retrieval

After the optimization of DSH-GANs, we can employ hash stream in the architecture followed by a sigmoid layer to generate K -bit hash codes for each input image. In this procedure, an image x is first encoded into a K -dimension feature vector $\mathbf{h} = \mathcal{H}(x)$. Then, a quantization operation $\mathbf{b} = \mathcal{Q}(\mathbf{h})$ is exploited to generate hash codes \mathbf{b} , where $\mathcal{Q}(\mathbf{h})$ is a sign function on vector \mathbf{h} with $\mathcal{Q}(h_i) = 1$ if $h_i > 0.5$ and otherwise $\mathcal{Q}(h_i) = 0$. Given a query image, the retrieval list of images is produced by sorting the Hamming distances of hash codes between the query image and images in search pool.

4 EXPERIMENTS

We conducted extensive evaluations of our proposed architecture on two image datasets, i.e., CIFAR-10¹ which is a collection of tiny images and NUS-WIDE² of a large-scale Web image dataset.

4.1 Datasets

The **CIFAR-10** dataset consists of 60,000 real world tiny images (32×32 pixels), which can be divided into 10 categories and 6,000 images for each category. We randomly select 1,000 images (100 images per class) as the test query set. For the unsupervised setting, all the rest images are used as training samples. For the supervised setting, we additionally sample 500 images from each class in the training samples and constitute a subset of 5,000 labeled images for training. The rest training images are treated as the unlabeled data.

The **NUS-WIDE** dataset contains 269,648 images collected from Flickr. Each of these images is associated with one or multiple labels in 81 semantic concepts. For a fair comparison, we follow the settings in [11] to employ the subset of images associated with 21 most frequent labels, where each label associates with at least 5,000 images. Similar to the split in CIFAR-10, we randomly select 2,100 images (100 images per class) as the test query set. For the unsupervised setting, all the rest images are used as the training set. For the supervised setting, we uniformly sample 500 images from each class to construct the labeled subset for training and the rest training images are all treated as unlabeled data.

4.2 Experimental Settings

On both datasets, we utilize AlexNet [9] as our basic CNN architecture and take the outputs of fc6 layer from AlexNet as the image representation. The shared CNN architecture is pre-trained on ImageNet dataset [22] and the generator network G is pre-trained with our proposed semi-supervised GANs on each dataset.

We mainly implement our proposed method based on Caffe [7], which is one of the widely adopted deep learning frameworks. For the semi-supervised GANs, we follow the standard settings in [20] and train our GANs models on both datasets by utilizing Adam optimizer with a mini-batch size of 128. All weights are initialized from a zero-centered Normal distribution with standard deviation 0.02 and the slope of the leak is set to 0.2 in the LeakyReLU. We fix the learning rate and momentum to 0.0002 and 0.9, respectively. For our DSH-GANs architecture, it is trained by stochastic gradient descent with 0.9 momentum. The start learning rate is set to 0.0001, and we decrease it to 10% after 10,000 iterations on CIFAR-10 and after 40,000 iterations on NUS-WIDE, respectively. The mini-batch size of images is 64 and the weight decay parameter is 0.0005.

¹<http://www.cs.toronto.edu/~kriz/cifar.html>

²<http://lms.comp.nus.edu.sg/research/NUS-WIDE.htm>

Table 1: Accuracy in terms of MAP. The best MAPs on each number of hash bits are shown in boldface. Note that the MAP performance is calculated on the top 5,000 returned images for NUS-WIDE dataset.

Method	CIFAR-10 (MAP)				NUS-WIDE (MAP)			
	12-bits	24-bits	32-bits	48-bits	12-bits	24-bits	32-bits	48-bits
DSH-GANs	0.735	0.781	0.787	0.802	0.838	0.856	0.861	0.863
DSH-GANs ⁻	0.726	0.769	0.772	0.783	0.823	0.847	0.845	0.854
DPSH	0.713	0.727	0.744	0.757	0.794	0.822	0.838	0.851
NINH	0.552	0.566	0.558	0.581	0.674	0.697	0.713	0.715
CNNH	0.439	0.476	0.472	0.489	0.611	0.618	0.625	0.608
KSH+CNN	0.446	0.502	0.518	0.516	0.746	0.774	0.765	0.749
ITQ+CNN	0.212	0.230	0.234	0.240	0.728	0.707	0.689	0.661
SH+CNN	0.158	0.157	0.154	0.151	0.620	0.611	0.620	0.591
LSH+CNN	0.134	0.157	0.173	0.185	0.438	0.586	0.571	0.507
KSH	0.303	0.337	0.346	0.356	0.556	0.572	0.581	0.588
ITQ	0.162	0.169	0.172	0.175	0.452	0.468	0.472	0.477
SH	0.127	0.128	0.126	0.129	0.454	0.406	0.405	0.400
LSH	0.121	0.126	0.120	0.120	0.403	0.421	0.426	0.441

4.3 Protocols and Baseline Methods

We follow four evaluation protocols, i.e., mean average precision (MAP), hash lookup, precision-recall curve, and precision curves w.r.t. different numbers of top returned samples, which are widely used in [4, 11, 14]. We compare the following approaches for performance evaluation:

(1) Locality Sensitive Hashing [3] (LSH) aims to map similar examples to the same bucket with high probability by using a Gaussian random projection matrix. The property of locality in the original space will be largely preserved in the Hamming space.

(2) Spectral Hashing [26] (SH) is based on quantizing the values of analytical eigenfunctions computed along PCA directions of the data.

(3) Iterative Quantization [4] (ITQ) learns similarity-preserving binary codes by directly minimizing the quantization error of mapping data to vertices of the binary hypercube.

(4) Kernel-based Supervised Hashing [14] (KSH) employs a kernel formulation for learning the hash functions to handle linearly inseparable data.

(5) Convolutional Neural Networks Hashing [28] (CNNH) firstly learns approximate hash codes with the supervised pairwise relationship and then trains CNN architecture with approximate hash codes and image tags.

(6) Network In Network Hashing [11] (NINH) utilizes a triplet ranking loss to preserve relative similarity and divide-and-encode modules to encode hash bits.

(7) Deep Pairwise-Supervised Hashing [12] (DPSH) performs simultaneous feature learning and hash learning by leveraging pairwise labels in an end-to-end system.

(8) Deep Semantic Hashing with Generative Adversarial Networks (DSH-GANs) is our proposal in this paper. A slightly different of this run is named as DSH-GANs⁻, which is trained without classification stream.

Note that for the four hashing methods using hand-crafted features (i.e., LSH, SH, ITQ and KSH), each image in CIFAR-10 and NUS-WIDE is represented by a 512-dimensional GIST

vector and an officially available 500-dimensional bag-of-words vector, respectively. For the deep hashing methods, we resize all images to be 224×224 pixels and then directly exploit the raw image pixels as input. Moreover, we also conduct the experiments by using the outputs of fc6 layer in AlexNet as image representation in the four traditional hashing approaches and name them as LSH+CNN, SH+CNN, ITQ+CNN and KSH+CNN, respectively.

4.4 Results on CIFAR-10 Dataset

The left half of Table 1 shows the MAP performance comparisons on CIFAR-10 dataset. Overall, the results across different number of hash bits indicate that our DSH-GANs consistently outperforms others. In particular, the MAP of DSH-GANs with 48-bits makes the relative improvement over the best traditional competitor KSH with GIST features or the outputs of fc6 layer in AlexNet, and deep model DPSH by 125.3%, 55.4% and 5.9%, respectively. Furthermore, traditional approaches with image representations extracted from CNN architecture lead to a large performance boost against these methods with GIST features, which is expected as deep CNN has demonstrated its high capability in generating image representations. Compared to the traditional models with deep image representations, deep hash models which benefit from the joint learning of image representations and hash coding exhibit better performances. DSH-GANs⁻ outperforms DPSH and NINH. The result basically indicates the advantage of exploring synthetic images in hashing. DSH-GANs further improves DSH-GANs⁻ with a relative increase of 1.2%~2.4%, demonstrating the strength of boosting hashing by additionally preserving semantics of images through classification. In addition, when utilizing a deeper CNN architecture VGG-19 [24] networks as our basic CNN, the MAP performance of our DSH-GANs with 12-bits, 24-bits, 32-bits and 48bits will be boosted up to 86.1%, 88.1%, 87.9% and 88.4%, respectively.

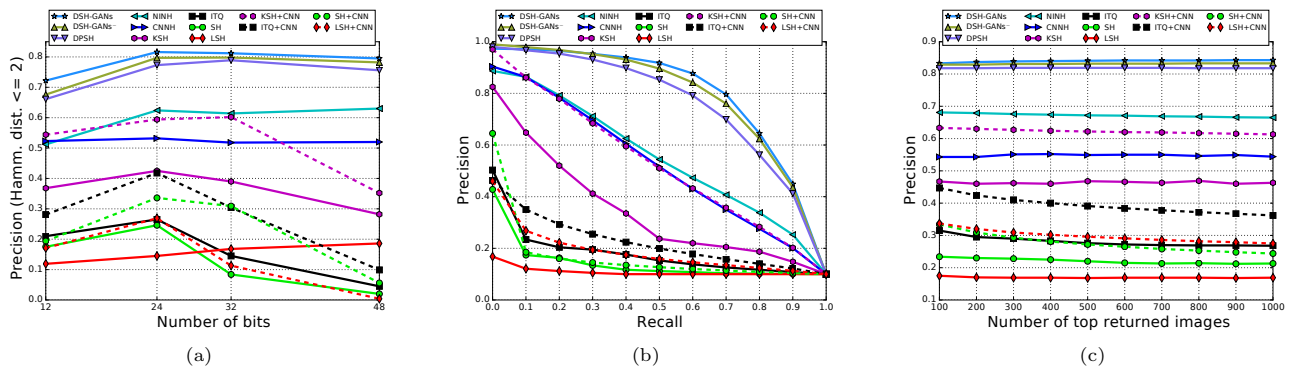


Figure 3: Comparisons with state-of-the-art approaches on CIFAR-10 dataset. (a) Precision within Hamming radius 2 using hash lookup. (b) Precision-Recall curves with 48-bits. (c) precision curves with 48-bits w.r.t. different number of top returned samples. Better viewed in original color pdf file.

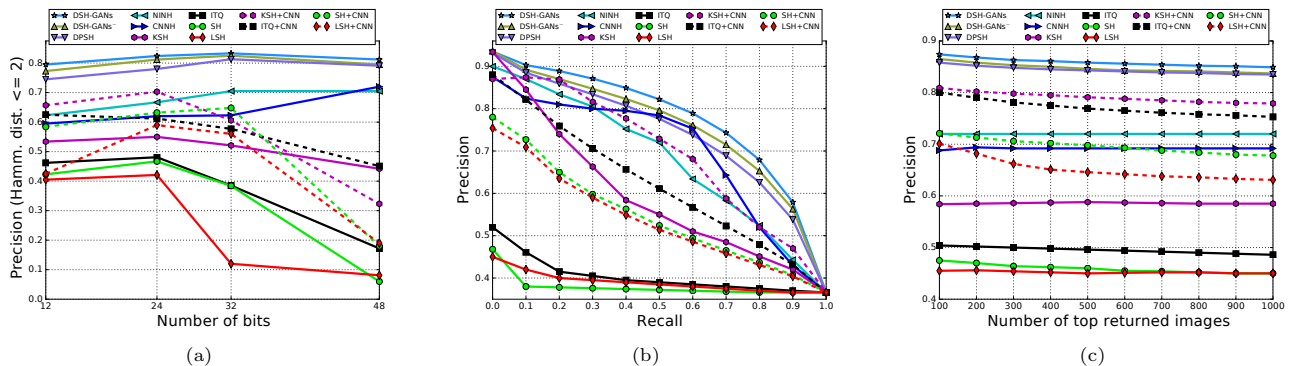


Figure 4: Comparisons with state-of-the-art approaches on NUS-WISE dataset. (a) Precision within Hamming radius 2 using hash lookup. (b) Precision-Recall curves with 48-bits. (c) precision curves with 48-bits w.r.t. different number of top returned samples. Better viewed in original color pdf file.

In the evaluation of hash lookup within Hamming radius 2 as shown in Figure 3(a), the precisions for most of the traditional methods drop when a longer size of hash codes is used (48 bits in our case). This is because the number of samples falling into a bucket decreases exponentially for longer sizes of hash codes. Therefore, for some query images, there are not even any neighbor in a Hamming ball of radius 2. Even in this case, the precision of our proposed DSH-GANs only has a slight decrease from 80.6% of 32 bits to 79.7% of 48 bits, indicating fewer failed queries for DSH-GANs. We further detail the precision-recall curves and precision curves with 48-bits w.r.t. different number of top returned samples in Figure 3(b) and 3(c). The results confirm the trends observed in Figure 3(a) and demonstrate performance improvements by our proposed DSH-GANs approach over other methods.

4.5 Results on NUS-WIDE Dataset

The right half of Table 1 lists the MAP performance comparisons on NUS-WIDE dataset. Precision with Hamming radius 2 using hash lookup, precision-recall curves with 48-bits and precision curves with 48-bits w.r.t. different number of top

returned samples is given in Figure 4(a), 4(b) and 4(c), respectively. DSH-GANs constantly exhibits better performance than other baselines across different performance metrics. Specifically, the MAP performance and precision with Hamming radius 2 using hash lookup of DSH-GANs achieve 86.3% and 81.2% with 48-bits, which make the improvements over the best competitor DPSH by 1.4% and 2.7%, respectively. This again verifies the effectiveness of generating synthetic and discriminable training data through GANs for hashing. Furthermore, DSH-GANs is benefited from utilizing semantic supervision and thus shows a relative increase of 1.1%~1.9% over DSH-GANs⁻ in terms of MAP.

Figure 5 further showcases the top ten image search results by different methods in response to two query images. We can see that the proposed DSH-GANs method achieves the most satisfying results and retrieves eight “excellent images” in the returned top ten images to each query image. It is worth noticing that “excellent images” here refer to images whose annotations completely contain all the labels of the query image (e.g., “water,” “clouds,” “ocean” and “beach” of the first image example). As a result, the images retrieved

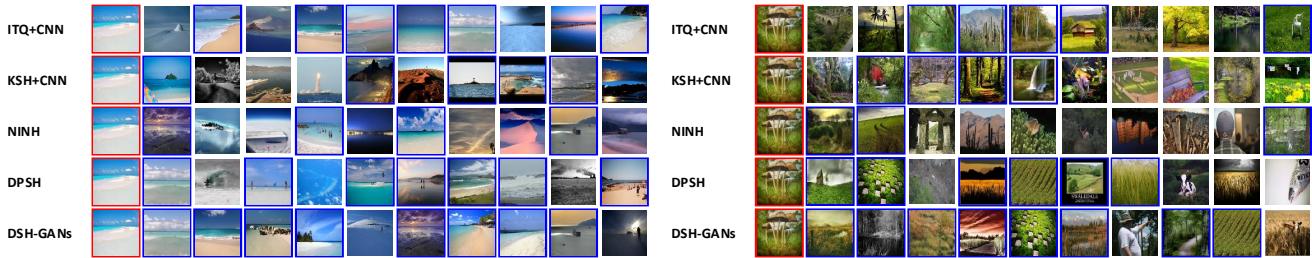


Figure 5: Examples showing the top 10 image retrieval results by different methods in response to two query images on NUS-WIDE dataset (better viewed in color). In each row, the first image with a red bounding box is the query image and the images whose annotations completely contain all the labels of the query image are regarded as excellent ones, which are enclosed in a blue bounding box.

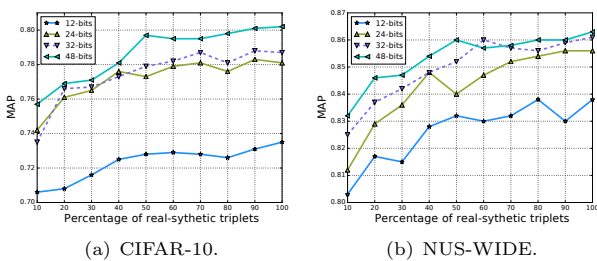


Figure 6: MAP performance comparison with different percentage of synthetic data in training triplets.

by our DSH-GANs approach are more similar in semantics with the query image.

4.6 Comparison between Synthetic and Real Examples for Hashing

In order to examine how performance is affected when exploiting synthetic examples in training triplets of different degree by DSH-GANs, we compare the MAP performances of using synthetic data with percentage ranging from 10% to 100%. In the previous experiments, the similar and dissimilar images in the training triplets are all synthetic images, which refers to 100% in this analysis. We control the ratio between real and synthetic data in training by replacing part of synthetic images with real ones. Figure 6 shows the results on both CIFAR-10 and NUS-WIDE datasets across different hash bits. The results are encouraging in the way that involving more synthetic data tends to achieve better performance. This empirically validates our proposal of generating synthetic data through semi-supervised GANs which additionally leverages largely unlabeled data, making the generated examples more discriminable to characterize the structure of the data.

4.7 Visualization of Synthetic Images

Figure 7 illustrates image examples on CIFAR-10 dataset, which are both randomly selected from each class in the dataset (left half) and generated for each class through our semi-supervised GANs (right half). In general, the generated images are plausible and semantically relevant to each class.



Figure 7: Visualization of image examples on CIFAR-10 dataset. Left half: images randomly selected from each class in the dataset; Right half: synthetic image examples for each class through our semi-supervised GANs.

Figure 8 further visualizes the synthetic image examples on NUS-WIDE dataset. The images in the right half of each row are semantically related to the images in the left half. Take the first row as an example, the images in the left half are generated with label “clouds,” while the images in the right half are synthesized with labels “clouds” and “sunset.” All the images look real and the generated images in the right part could clearly manifest the semantics of “sunset” and differentiate them from the images in the left part with only semantics of “clouds.”

5 CONCLUSIONS

We have presented a Deep Semantic Hashing with Generative Adversarial Networks (DSH-GANs) architecture which explores semi-supervised GANs to generate synthetic training data for hashing. Particularly, a semi-supervised GANs is trained on both labeled and unlabeled data to produce compelling and discriminable examples conditioning on class




Figure 8: Visualization of synthetic image examples on NUS-WIDE dataset. All the image examples are generated with multiple labels. The images in the right half of each row are semantically related to the images in the left half.

labels. To verify our claim, we optimize the whole architecture of our hashing model by simultaneously distinguishing synthetic images from real ones and preserving not only relative similarity between images but also semantics of images. Experiments conducted on both CIFAR-10 and NUS-WIDE datasets validate our proposal and analysis. Performance improvements are clearly observed when comparing to other hashing techniques.

Our future works are as follows. First, as our architecture is a joint learning procedure, how the architecture performs on classification task will be further evaluated. Next, more in-depth studies of how to fuse the three streams in a principled way could be explored. Finally, more advanced GANs (e.g., Stacked GANs) and CNN architectures (e.g., ResNet) will be investigated in our architecture.

REFERENCES

- [1] Qi Dai, Jianguo Li, Jingdong Wang, and Yu-Gang Jiang. 2016. Binary Optimized Hashing. In *ACM MM*.
- [2] Emily Denton, Soumith Chintala, Arthur Szlam, and Rob Fergus. 2015. Deep Generative Image Models using a Laplacian Pyramid of Adversarial Networks. In *NIPS*.
- [3] Aristides Gionis, Piotr Indyk, and Rajeev Motwani. 1999. Similarity search in high dimensions via hashing. In *VLDB*.
- [4] Yunchao Gong, Svetlana Lazebnik, Albert Gordo, and Florent Perronnin. 2013. Iterative Quantization: A Procrustean Approach to Learning Binary Codes for Large-scale Image Retrieval. *IEEE Trans. Pattern Anal. Mach. Intell.* (2013).
- [5] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. In *NIPS*.
- [6] Geoffrey E Hinton and Ruslan R Salakhutdinov. 2006. Reducing the Dimensionality of Data with Neural Networks. *Science* (2006).
- [7] Yangqing Jia, Evan Shelhamer, Jeff Donahue, Sergey Karayev, Jonathan Long, Ross Girshick, Sergio Guadarrama, and Trevor Darrell. 2014. Caffe: Convolutional architecture for fast feature embedding. In *ACM MM*.
- [8] Saehoon Kim and Seungjin Choi. 2011. Semi-supervised Discriminant Hashing. In *ICDM*.
- [9] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. 2012. ImageNet Classification with Deep Convolutional Neural Networks. In *NIPS*.
- [10] Brian Kulis and Kristen Grauman. 2012. Kernelized Locality-sensitive Hashing. *IEEE Trans. Pattern Anal. Mach. Intell.* (2012).
- [11] Hanjiang Lai, Yan Pan, Ye Liu, and Shuicheng Yan. 2015. Simultaneous Feature Learning and Hash Coding with Deep Neural Networks. In *CVPR*.
- [12] Wu-Jun Li, Sheng Wang, and Wang-Cheng Kang. 2016. Feature learning based deep supervised hashing with pairwise labels. In *IJCAI*.
- [13] Venice Erin Liong, Jiwen Lu, Gang Wang, Pierre Moulin, and Jie Zhou. 2015. Deep Hashing for Compact Binary Codes Learning. In *CVPR*.
- [14] Wei Liu, Jun Wang, Rongrong Ji, Yu-Gang Jiang, and Shih-Fu Chang. 2012. Supervised hashing with kernels. In *CVPR*.
- [15] Mehdi Mirza and Simon Osindero. 2014. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784* (2014).
- [16] Augustus Odena, Christopher Olah, and Jonathon Shlens. 2016. Conditional Image Synthesis With Auxiliary Classifier GANs. *arXiv preprint arXiv:1610.09585* (2016).
- [17] Yingwei Pan, Yehao Li, Ting Yao, Tao Mei, Houqiang Li, and Yong Rui. 2016. Learning deep intrinsic video representation by exploring temporal coherence and graph structure. *IJCAI*.
- [18] Yingwei Pan, Ting Yao, Houqiang Li, Chong-Wah Ngo, and Tao Mei. 2015. Semi-supervised hashing with semantic confidence for large scale visual search. In *SIGIR*.
- [19] Zhaofan Qiu, Ting Yao, and Tao Mei. 2017. Deep Quantization: Encoding Convolutional Activations with Deep Generative Model. In *CVPR*.
- [20] Alec Radford, Luke Metz, and Soumith Chintala. 2016. Unsupervised representation learning with deep convolutional generative adversarial networks. In *ICLR*.
- [21] Scott Reed, Zeynep Akata, Xinchun Yan, Lajanugen Logeswaran, Bernt Schiele, and Honglak Lee. 2016. Generative adversarial text to image synthesis. In *ICML*.
- [22] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. 2015. ImageNet Large Scale Visual Recognition Challenge. *IJCV* (2015).
- [23] Ruslan Salakhutdinov and Geoffrey Hinton. 2009. Semantic Hashing. *IJAR* (2009).
- [24] Karen Simonyan and Andrew Zisserman. 2015. Very Deep Convolutional Networks for Large-scale Image Recognition. In *ICLR*.
- [25] Jun Wang, Sanjiv Kumar, and Shih-Fu Chang. 2012. Semi-Supervised Hashing for Large-Scale Search. *IEEE Trans. Pattern Anal. Mach. Intell.* (2012).
- [26] Y. Weiss, A. Torralba, and R. Fergus. 2008. Spectral hashing. In *NIPS*.
- [27] Hao Xia, Pengcheng Wu, Steven CH Hoi, and Rong Jin. 2012. Boosting Multi-kernel Locality-sensitive Hashing for Scalable Image Retrieval. In *SIGIR*.
- [28] Rongkai Xia, Yan Pan, Hanjiang Lai, Cong Liu, and Shuicheng Yan. 2014. Supervised Hashing for Image Retrieval via Image Representation Learning. In *AAAI*.
- [29] Ting Yao, Fuchen Long, Tao Mei, and Yong Rui. 2016. Deep semantic-preserving and ranking-based hashing for image retrieval. In *IJCAI*.
- [30] Han Zhu, Mingsheng Long, Jianmin Wang, and Yue Cao. 2016. Deep Hashing Network for Efficient Similarity Retrieval. In *AAAI*.