

Generalized Product Quantization Network for Semi-supervised Image Retrieval

Young Kyun Jang

Nam Ik Cho

Department of ECE, INMC, Seoul National University, Seoul Korea

kyun0914@ispl.snu.ac.kr, nicho@snu.ac.kr

Abstract

Image retrieval methods that employ hashing or vector quantization have achieved great success by taking advantage of deep learning. However, these approaches do not meet expectations unless expensive label information is sufficient. To resolve this issue, we propose the first quantization-based semi-supervised image retrieval scheme: **Generalized Product Quantization (GPQ)** network. We design a novel metric learning strategy that preserves semantic similarity between labeled data, and employ entropy regularization term to fully exploit inherent potentials of unlabeled data. Our solution increases the generalization capacity of the quantization network, which allows overcoming previous limitations in the retrieval community. Extensive experimental results demonstrate that GPQ yields state-of-the-art performance on large-scale real image benchmark datasets.

1. Introduction

The amount of multimedia data, including images and videos, increases exponentially on a daily basis. Hence, retrieving relevant content from a large-scale database has become a more complicated problem. There have been many kinds of fast and accurate search algorithms, and the Approximate Nearest Neighbor (ANN) search is known to have high retrieval accuracy and computational efficiency. Recent ANN methods mainly focused on *hashing* scheme [31], because of its low storage cost and fast retrieval speed. To be specific, an image is represented by a binary-valued compact hash code (*binary code*) with only a few tens of bits, and it is utilized to build database and distance computation.

The methods using binary code representation can be categorized as *Binary hashing* (BH) and *Product Quantization* (PQ) [13]. BH-based methods [34, 7, 26] employ a hash function that maps a high-dimensional vector space to a Hamming space, where the distance between two codes can be measured extremely fast via bitwise XOR operation.

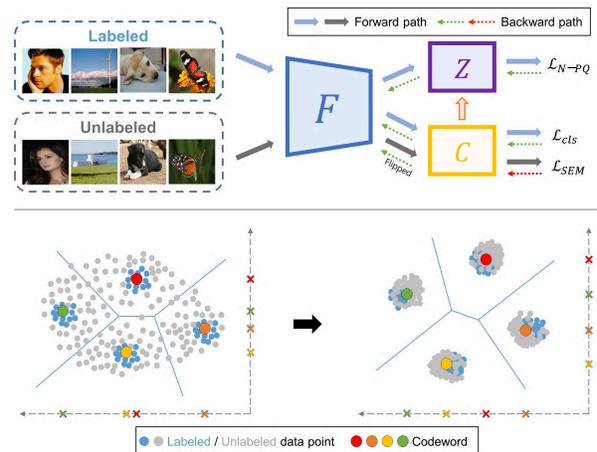


Figure 1. Above: an illustration of the overall framework of GPQ and its three components: feature extractor F , PQ table Z , and classifier C , where C contributes to build Z . Forward path shows how the labeled and unlabeled data pass through the network, and backward path shows the propagation of the gradients originated from the training objectives. \mathcal{L}_{N-PQ} and \mathcal{L}_{cls} train the network to minimize errors using the labeled data, while \mathcal{L}_{SEM} trains the network to simultaneously maximize and minimize entropy using the unlabeled data. Below: a 2D conceptual Voronoi diagram showing one of the codebooks in GPQ. After training, all codewords are evenly distributed, and both labeled and unlabeled data points are clustering around them.

However, BH has a limitation in describing the distance between data points because it can produce only a limited number of distinct values. PQ, which is a kind of vector quantization [8], has been introduced to alleviate this problem in information retrieval [13, 6, 15].

To perform PQ, we first need to decompose the input feature space into a Cartesian product of several disjoint subspaces (*codebooks*) and find the centroid (*codeword*) of each subspace. Then, from the sub-vectors of the input feature vector, sub-binary code is obtained by replacing each sub-vector with the index of the nearest codeword in the codebook. Since codeword consists of real numbers, PQ allows asymmetric distance calculation in real space using

the binary codes, making many PQ-based approaches outperform BH-based ones.

Along with millions of elaborately labeled data, *deep hashing* for both BH [35, 18, 11, 12, 14] and PQ [2, 19, 38, 16] has been introduced to take advantage of deep representations for image retrieval. By employing supervised deep neural networks, deep hashing outperforms conventional ones on many benchmark datasets. Nevertheless, there still is a great deal of potential for improvement, since a significant amount of unlabeled data with abundant knowledge is not utilized. To resolve these issue, some recent methods are considering the *Deep Semi-Supervised Hashing*, based on BH [40, 36, 10]. However, even if PQ generally outperforms BH for both supervised and unsupervised settings, it has not yet been considered for learning in a semi-supervised manner. In this paper, we propose the first PQ-based deep semi-supervised image retrieval approach: Generalized Product Quantization (GPQ) network, which significantly improves the retrieval accuracy with lots of image data and just a few labels per category (class).

Existing deep semi-supervised BH methods construct graphs [40, 36] or apply additional generative models [10] to encode unlabeled data into the binary code. However, due to the fundamental problem in BH; a deviation that occurs when embedding a continuous deep representation into a discrete binary code restricts extensive information of unlabeled data. In our GPQ framework, this problem is solved by including quantization process into the network learning. We adopt intra-normalization [1] and soft assignment [38] to quantize real-valued input sub-vectors, and introduce an efficient metric learning strategy; *N-pair Product Quantization loss* inspired from [28]. By this, we can embed multiple pair-wise semantic similarity between every feature vectors in a training batch into the codewords. It also has an advantage of not requiring any complex batch configuration strategy to learn pairwise relations.

The key point of deep semi-supervised retrieval is to avoid overfitting to labeled data and increase the generalization toward unlabeled one. For this, we suggest a *Subspace Entropy Mini-max Loss* for every codebook in GPQ, which regularizes the network using unlabeled data. Precisely, we first learn a cosine similarity-based classifier, which is commonly used in few-shot learning [27, 32]. The classifier has as many weight matrices as the number of codebooks, and each matrix contains class-specific weight vector, which can be regarded as a *sub-prototype* that indicates class representative centroid of each codebook. Then, we compute the entropy between the distributions of sub-prototypes and unlabeled sub-vectors by measuring their cosine similarity. By maximizing the entropy, the two distributions become similar, allowing the sub-prototypes to move closer to the unlabeled sub-vectors. At the same time, we also minimize the entropy of the distribution of the unlabeled

sub-vectors, making them assemble near the moved sub-prototypes. With the gradient reversal layer generally used for deep domain adaptation [5, 24], we are able to simultaneously minimize and maximize the entropy during the network training.

In summary, the main contributions of our work are as follows:

- To the best of our knowledge, our work is the first deep semi-supervised PQ scheme for image retrieval.
- With the proposed metric learning strategy and entropy regularization term, the semantic similarity of labeled data is well preserved into the codewords, and the underlying structure of unlabeled data can be fully used to generalize the network.
- Extensive experimental results demonstrate that our GPQ can yield the state-of-the-art retrieval results in semi-supervised image retrieval protocols.

2. Related Work

Existing Hashing Methods Referring to the survey [31], early works in Binary Hashing (BH) [34, 7, 26, 30] and Product Quantization (PQ) [13, 6, 15, 9, 21, 37, 29] mainly focused on unsupervised settings. Specifically, Spectral Hashing (SH) [34] considered correlations within hash functions to obtain balanced compact codes. Iterative Quantization (ITQ) [7] addressed the problem of preserving the similarity of original data by minimizing the quantization error in hash functions. There were several studies to improve PQ, for example, Optimized Product Quantization (OPQ) [6] tried to improve the space decomposition and codebook learning procedure to reduce the quantization error. Locally Optimized Product Quantization (LOPQ) [15] employed a coarse quantizer with locally optimized PQ to explore more possible centroids. These methods might reveal some distinguishable results, however, they still have disadvantage of not exploiting expensive label signals.

Deep Hashing Methods After Supervised Discrete Hashing (SDH) [26] has shown the capability to improve using labels, supervised Convolutional Neural Network (CNN)-based BH approaches [35, 18, 11, 12, 14] are leading the mainstream. For examples, CNN Hashing (CNNH) [35] utilized a CNN to simultaneously learn feature representation and hash functions with the given pairwise similarity matrix. Network in Network Hashing (NINH) [18] introduced a sub-network, divide-and-encode module, and a triplet ranking loss for the similarity-preserving hashing. A Supervised, Structured Binary Code (SUBIC) [11], used a block-softmax nonlinear function and computed batch-based entropy error to embed the structure into a binary

form. There also have been researches on supervised learning that uses PQ with CNN [2, 19, 38, 16]. Precisely, Deep Quantization Network (DQN) [2] simultaneously optimizes a pairwise cosine loss on semantic similarity pairs to learn feature representations and a product quantization loss to learn the codebooks. Deep Triplet Quantization (DTQ) [19] designed a group hard triplet selection strategy and trained triplets by triplet quantization loss with weak orthogonality constraint. Product Quantization Network (PQN) [38] applied the asymmetric distance calculation mechanism to the triplets and exploited softmax function to build a differentiable soft product quantization layer to train the network in an end-to-end manner. Our method is also based on PQ, but we try a *semi*-supervised PQ scheme that had not been considered previously.

Deep Semi-supervised Image Retrieval Assigning labels to images is not only expensive but also has the disadvantage of restricting the data structure to the labels. Deep semi-supervised hashing based on BH is being considered in the image retrieval community to alleviate this problem, with the use of a small amount of labeled data and a large amount of unlabeled data. For example, Semi-supervised Deep Hashing (SSDH) [40] employs an online graph construction strategy to train a network using unlabeled data. Deep Hashing with a Bipartite Graph (BGDH) [36] improved SSDH by using the bipartite graph, which is more efficient in building a graph and learning the embeddings. Since Generative Adversarial Network (GAN) had been used for BH and showed good performance as in [14], Semi-supervised Generative Adversarial Hashing (SSGAH) also employed the GAN to fully utilize triplet-wise information of both labeled and unlabeled data. In this paper, we propose GPQ, the first deep semi-supervised image retrieval method applying PQ. In our work, we endeavor to generalize the whole network by preserving the semantic similarity with N-pair Product Quantization loss and extracting the underlying structure of unlabeled data with the Subspace Entropy Mini-max loss.

3. Generalized Product Quantization

Given a dataset \mathcal{X} that is composed of individual images, we split this into two subsets as a labeled dataset $\mathcal{X}^L = \{(I_i^L, y_i) | i = 1, \dots, N^L\}$ and an unlabeled dataset $\mathcal{X}^U = \{I_i^U | i = 1, \dots, N^U\}$ to establish a semi-supervised environment. The goal of our work is learning a quantization function $q : I \rightarrow \mathbf{b} \in \{0, 1\}^B$ which maps a high-dimensional input I to a compact B -bits binary code \mathbf{b} , by utilizing both labeled and unlabeled datasets. We propose a semi-supervised deep hashing framework: GPQ, which integrates q into the deep network as a form of Product Quantization (PQ) [13] to learn the deep representations and the codewords jointly. In the learning process, we aim to pre-

serve the semantic similarity of labeled data and simultaneously explore the structures of unlabeled data to obtain high retrieval accuracy.

GPQ contains three trainable components: 1) a standard deep convolutional neural network-based feature extractor F , e.g. AlexNet [17], CNN-F [3] or modified version of VGG [12] to learn deep representations; 2) a PQ table Z that collects codebooks which are used to map an extracted feature vector to a binary code; 3) a cosine similarity-based classifier C to classify both labeled and unlabeled data. GPQ network is designed to train all these components in an end-to-end manner. In this section, we will describe each component and how GPQ is learned in a semi-supervised way.

3.1. Semi-Supervised Learning

The feature extractor F generates D -dimensional feature vector $\hat{\mathbf{x}} \in R^D$. Under semi-supervised learning condition, we aim to train the F to extract discriminative $\hat{\mathbf{x}}^L$ and $\hat{\mathbf{x}}^U$ from labeled image I^L and unlabeled image I^U , respectively. Also, we leverage the PQ concept to utilize these feature vectors for image retrieval, which requires appropriate codebooks with distinct codewords to replace and store the feature vectors. We introduce three training objectives for our GPQ approach to fully exploit the data structure of labeled and unlabeled images, and we illustrate a conceptual visualization of each loss function in Figure 2 for better understanding.

Following the observation of [27, 32, 38, 19], we normalize the feature vectors and constrain them on a unit hypersphere to focus on the angle rather than the magnitude in measuring the distance between two different vectors. In this way, every data is mapped to the nearest class representative direction, and better performs for the semi-supervised scheme because the distribution divergence between labeled and unlabeled data can be reduced within the constraint. Especially for PQ, we apply intra-normalization [1] for a feature vector $\hat{\mathbf{x}}$ by dividing it into M -sub-vectors $\hat{\mathbf{x}} = [\mathbf{x}_1, \dots, \mathbf{x}_M]$, where $\mathbf{x}_m \in R^d, d = D/M$, and l_2 -normalize each sub-vector as: $\mathbf{x}_m \leftarrow \mathbf{x}_m / \|\mathbf{x}_m\|_2$. In the rest of the paper, $\hat{\mathbf{x}}$ for GPQ denotes the intra-normalized feature vector.

N-pair Product Quantization The PQ table Z collects M -codebooks $Z = [Z_1, \dots, Z_M]$ and each codebook has K -codewords $Z_m = [z_{m1}, \dots, z_{mK}]$ where $z_{mk} \in R^d$, which are used to replace $\hat{\mathbf{x}}$ with the quantized vector $\hat{\mathbf{q}}$. Every codeword is l_2 -normalized to simplify the measurement of cosine similarity as multiplication. We employ the soft assignment $s_m(\cdot)$ [38] to obtain a \mathbf{q}_m from \mathbf{x}_m as:

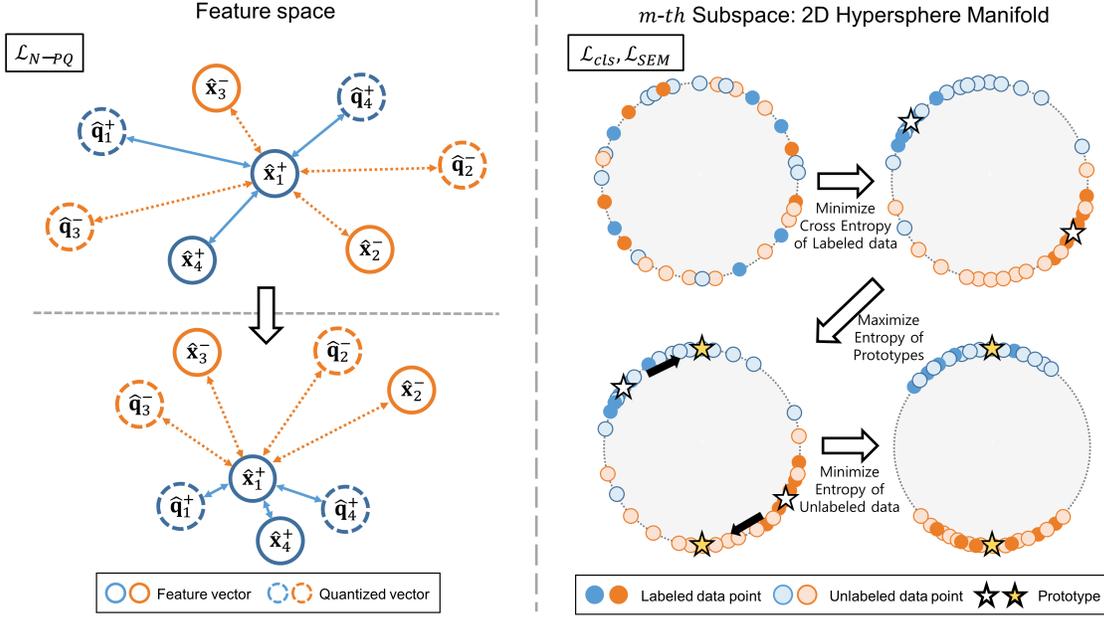


Figure 2. A two class (+: blue, -: orange) visualized examples of our training objectives. 1) The left part shows the learning process of N-pair Product Quantization loss \mathcal{L}_{N-PQ} . When we define an anchor as $\hat{\mathbf{x}}_1^+$, the semantically similar points ($\hat{\mathbf{q}}_1^+$, $\hat{\mathbf{x}}_4^+$, $\hat{\mathbf{q}}_4^+$) are pulled together while the semantically dissimilar points ($\hat{\mathbf{x}}_2^-$, $\hat{\mathbf{q}}_2^-$, $\hat{\mathbf{x}}_3^-$, $\hat{\mathbf{q}}_3^-$) are pushing the anchor. 2) The right part shows the learning process of classification loss \mathcal{L}_{cls} and subspace entropy mini-max loss \mathcal{L}_{SEM} . For the data points constrained on the unit hypersphere, the cross entropy of labeled data points is minimized to find prototypes (white stars). Then, the entropy between the prototypes and the unlabeled data points is maximized to move prototypes toward unlabeled data points and find new prototypes (yellow stars). Finally, the entropy of the unlabeled data points is minimized to cluster them near the new prototypes.

$$\mathbf{q}_m = \sum_k^K \frac{e^{-\alpha(\mathbf{x}_m \cdot \mathbf{z}_{mk})}}{\sum_{k'}^K e^{-\alpha(\mathbf{x}_m \cdot \mathbf{z}_{mk'})}} \mathbf{z}_{mk} \quad (1)$$

where α represents a scaling factor to approximate hard assignment, and $\mathbf{q}_m = s_m(\mathbf{x}_m; \alpha, \mathbf{Z}_m)$ is the sub-quantized vector of $\hat{\mathbf{q}} = [\mathbf{q}_1, \dots, \mathbf{q}_M]$. We multiply \mathbf{x}_m with the every codeword in \mathbf{Z}_m to measure the cosine similarity between them.

Quantization error occurs through the encoding process; therefore, we need to find the codewords that will minimize the error. In addition, the conventional PQ scheme has a limitation that it ignores label information since the sub-vectors are clustered to find codewords without any supervised signals. To fully exploit the semantic labels and reduce the quantization error, we revise the metric learning strategy proposed in [28] from *N-pair Product Quantization loss*: \mathcal{L}_{N-PQ} to learn the F and Z , and set it as one of the training objectives.

Deep metric learning [33, 25] aims to learn an embedding representation of the data with the semantic labels. From a labeled image I^L , we can generate a unique feature vector $\hat{\mathbf{x}}^L$ and its nearest quantized vector $\hat{\mathbf{q}}^L$, which can be regarded as sharing the same semantic information. Thus, for randomly sampled B training examples

$\{(I_1^L, y_1), \dots, (I_B^L, y_B)\}$, the objective function which is based on a standard cross entropy loss \mathcal{L}_{CE} , can be formulated as:

$$\mathcal{L}_{N-PQ} = \frac{1}{B} \sum_{b=1}^B \mathcal{L}_{CE}(\mathcal{S}_b, \mathbf{Y}_b) \quad (2)$$

where $\mathcal{S}_b = [(\hat{\mathbf{x}}_b^L)^T \hat{\mathbf{q}}_1^L, \dots, (\hat{\mathbf{x}}_b^L)^T \hat{\mathbf{q}}_M^L]$ denotes a cosine similarity between b -th feature vector and every quantized vector, and $\mathbf{Y}_b = [(y_b^T \mathbf{y}_1, \dots, (y_b^T \mathbf{y}_B)]$ denotes a similarity between b -th label and every label in a batch. In this case, \mathbf{y} represents one-hot-encoded semantic label, and column-wise normalization is applied to \mathbf{Y}_B . \mathcal{L}_{N-PQ} has an advantage in that no complicated batch construction method is required, and it also allows us to jointly learn the deep feature representation for both feature vectors and the codewords on the same embedding space.

Cosine Similarity-based Classification To embed semantic information into the codewords while reducing correlation between each codebook, we learn a cosine similarity-based classifier C containing M -weight matrices $[\mathbf{W}_1, \dots, \mathbf{W}_M]$, where each matrix includes sub-prototypes as $\mathbf{W}_m = [\mathbf{c}_{m1}, \dots, \mathbf{c}_{mN^c}]$, $\mathbf{W}_m \in R^{d \times N^c}$ and N^c is the number of class. Every sub-prototype is l_2 -normalized to

hold a class-specific angular information. With the m -th sub-vector \mathbf{x}_m^L of \mathbf{x}^L and m -th weight matrix \mathbf{W}_m , we can obtain the labeled class prediction as: $\mathbf{p}_m^L = \mathbf{W}_m^T \mathbf{x}_m^L$. We use \mathcal{L}_{CE} again to train the F and C for classification using $\hat{\mathbf{p}}^L = [\mathbf{p}_1^L, \dots, \mathbf{p}_M^L]$ computed from \mathbf{x}^L with the corresponding semantic label y as:

$$\mathcal{L}_{cls} = \frac{1}{M} \sum_{m=1}^M \mathcal{L}_{CE}(\beta \cdot \mathbf{p}_m^L, y) \quad (3)$$

where β is a scaling factor and y is a label corresponding to \mathbf{x}^L . This classification loss ensures the feature extractor to generate the discriminative features with respect to the labeled examples. In addition, each weight matrix in the classifier is derived to include the class-specific representative sub-prototypes of related subspace.

Subspace Entropy Mini-max On the assumption that the distribution is not severely different between the labeled data and the unlabeled data, we aim to propagate the gradients derived from the divergence between them. To calculate the error arising from the distribution difference, we adopt an entropy in information theory. In particular for PQ setting, we compute the entropy for each subspace to balance the amount of gradients propagating into each subspace. With the m -th sub-vector \mathbf{x}_m^U of the unlabeled feature vector \mathbf{x}^U , and the m -th weight matrix of C that embraces sub-prototypes, we can obtain a class prediction using cosine similarity as: $\mathbf{p}_m^U = \mathbf{W}_m^T \mathbf{x}_m^U$. By using it, the subspace entropy mini-max loss is calculated as:

$$\mathcal{L}_{SEM} = -\frac{1}{M} \sum_{m=1}^M \sum_{l=1}^{N^c} (\beta \cdot p_{ml}^U) \log(\beta \cdot p_{ml}^U) \quad (4)$$

where β is the same as that in Equation 3 and p_{ml}^U denotes the probability of prediction to k' -th class; l -th element of \mathbf{p}_m^U . The generalization capacity of the network can increase by maximizing the \mathcal{L}_{SEM} because high entropy ensures that the sub-prototypes are regularized toward unlabeled data. Explicitly, entropy maximization makes sub-prototypes have a similar distribution with the unlabeled sub-vectors, moving sub-prototypes near the unlabeled sub-vectors. To further improve, we aim to cluster unlabeled sub-vectors near the moved sub-prototype, by applying gradient reversal layer [5, 24] before intra-normalization. Flipped gradients induce F to be learned in the direction of minimizing the entropy, resulting in a skewed distribution of unlabeled data.

From the Equations 2 to 4, total objective function \mathcal{L}_T for B randomly sampled training pairs of I^L and I^U , can be formulated as:

$$\mathcal{L}_T(\mathbf{B}) = \mathcal{L}_{N-PQ} + \frac{1}{B} \sum_{b=1}^B (\lambda_1 \mathcal{L}_{cls} - \lambda_2 \mathcal{L}_{SEM}) \quad (5)$$

where $\mathbf{B} = \{(I_1^L, y_1, I_1^U), \dots, (I_B^L, y_B, I_B^U)\}$ and λ_1 and λ_2 are the hyper-parameters that balance the contribution of each loss function. We force training optimizer to minimize the \mathcal{L}_T , so that the \mathcal{L}_{N-PQ} and \mathcal{L}_{cls} is minimized while the \mathcal{L}_{SEM} is maximized, simultaneously. In this way, F can learn the deep representation of both labeled and unlabeled data. However, to make the codewords robust against unlabeled data, it is necessary to reflect the unlabeled data signals directly into Z . Accordingly, we apply another soft assignment to embed sub-prototype intelligence of \mathbf{W}_m into the of m -th codebook by updating codewords as $\mathbf{z}'_{mk} = s_m(\mathbf{z}_{mk}; \alpha, \mathbf{W}_m)$. As a result, high retrieval performance can be expected by exploiting the potential of unlabeled data for quantization.

3.2. Retrieval

Building Retrieval Database After learning the entire GPQ framework, we can build a retrieval database using images in \mathcal{X}^U . Given an input image $I^R \in \mathcal{X}^U$, we first extract $\hat{\mathbf{x}}^R$ from F . Then, we find the nearest codeword \mathbf{z}_{mk^*} of each sub-vector \mathbf{x}_m^R from the corresponding codebook \mathbf{Z}_m , by computing the cosine similarity. After that, formatting a index k^* of the nearest codeword as binary to generate a sub-binary code \mathbf{b}^R . Finally, concatenate all the sub-binary codes to obtain a $M \cdot \log_2(K)$ -bits binary code $\hat{\mathbf{b}}^R$, where $\hat{\mathbf{b}}^R = [\mathbf{b}_1^R, \dots, \mathbf{b}_M^R]$. This procedure is repeated for all images to store them as binary, and Z is also stored for distance calculation.

Asymmetric Search For a given query image I^Q , $\hat{\mathbf{x}}^Q$ is extracted from F . To conduct the image retrieval, we take the m -th sub-vector \mathbf{x}_m^Q as an example, compute the cosine similarity between \mathbf{x}_m^Q and every codeword belonging to the m -th codebook, and store measured similarities on the look-up-table (LUT). Similarly, the same operation is done for the other sub-vectors, and the results are also stored on the LUT. The distance between the query image and a binary code in the database can be calculated asymmetrically, by loading the pre-computed distance from the LUT using a sub-binary codes, and aggregating all the loaded distances.

4. Experiments

We evaluate GPQ for two semi-supervised image retrieval protocols against several hashing approaches. We conduct experiments on the two most popular image retrieval benchmark datasets. Extensive experimental results show that GPQ achieves superior performance to existing methods.

4.1. Setup

Evaluation Protocols and Metrics Following the semi-supervised retrieval experiments in [40, 36, 10], we adopt two protocols as follows.

- **Protocol 1: Single Category Image Retrieval** Assuming all categories (classes) used for image retrieval are known and only a small number of label data is provided for each class. The labeled data is used for training, and the unlabeled data is used for building the retrieval database and the query dataset. In this case, labeled data and the unlabeled data belonging to the retrieval database are used for semi-supervised learning.
- **Protocol 2: Unseen Category Image Retrieval** In line with semi-supervised learning, suppose that the information of the categories in the query dataset is unknown, and consider building a retrieval database using both known and unknown categories. For this situation, we divide image dataset into four parts: train75, test75, train25, and test25, where train75 and test75 are the data of 75% categories, while train25 and test25 are the data of the remaining 25% categories. We use train75 for training, train25 and test75 for the retrieval database, and test25 for the query dataset. In this case, train75 with labels, and train25, test75 without labels are used for semi-supervised learning.

The retrieval performance of hashing method is measured by mAP (mean Average Precision) with bit lengths of 12, 24, 32, and 48 for all images in the query dataset. In particular, we set Protocol 1 as the primary experiment and observe the contribution of each training objective.

Datasets We set up two benchmark datasets, different for each protocol as shown in the Table 1, and each dataset is configured as follows.

- **CIFAR-10** is a dataset containing 60,000 color images with the size of 32×32 . Each image belongs to one of 10 categories, and each category includes 6,000 images.
- **NUS-WIDE** [4] is a dataset consisting nearly 270,000 color images with various resolutions. Images in the dataset associate with one or more class labels of 81 semantic concepts. We select the 21 most frequent concepts for experiments, where each concept has more than 5,000 images, with a total of 169,643.

	CIFAR-10		NUS-WIDE	
	Protocol 1	Protocol 2	Protocol 1	Protocol 2
Query	1,000	9,000	2,100	35,272
Training	5,000	21,000	10,500	48,956
Retrieval Database	54,000	30,000	157,043	85,415

Table 1. Detailed composition of two benchmark datasets.

Implementation Details We implement our GPQ based on the Tensorflow framework and perform it with a NVIDIA Titan XP GPU. When it comes to conducting experiments on non-deep learning-based hashing methods [34, 7, 26, 13, 6, 15], we utilize hand-crafted features as inputs following [18], extracting 512-dimensional GIST [22] features from CIFAR-10 images, and 500-dimensional bag-of-words features from NUS-WIDE images. For deep hashing methods [35, 18, 11, 2, 19, 38, 40, 36, 10], we use raw images as inputs and adopt ImageNet pretrained AlexNet [17] and CNN-F [3] as the backbone architectures to extract the deep representations. Since AlexNet and CNN-F significantly downsample the input image at the first convolution layer, they are inadequate especially for small images. Therefore, we employ modified VGG architecture proposed in [12] as the baseline architecture for feature extractor of GPQ. For a fair comparison, we also employ CNN-F to evaluate our scheme (GPQ-F), and details will be discussed in the section 4.2.

With regard to network training, we adopt ADAM algorithm to optimize the network, and apply exponential learning rate decay with the initial value 0.0002 and the $\beta_1 = 0.5$. We configure the number of labeled and unlabeled images equally in the learning batch. To simplify the experiment, we set several hyper-parameters as follows: the scaling factors α and β are fixed for 20 and 4 respectively, the codeword number K is fixed at 2^4 while M is adjusted to handle multiple bit lengths, and the dimension d of the sub-vector \mathbf{x}_m is fixed at 12. Detail analysis of these hyper-parameters is shown in the supplementary. We will make our code publicly open for further research and comparison.

4.2. Results and Analysis

Overview Experimental results for protocol 1 and protocol 2 are shown in Tables 2 and 3, respectively. In each Table, the methods are divided into several basic concepts and listed by group. We investigate the variants of GPQ for ablation study, and the results can be seen in Figures 3 to 5. From the results, we can observe that our GPQ scheme outperforms other hashing methods, demonstrating that the proposed loss functions effectively improve the GPQ network by training it in a semi-supervised fashion.

Comparison with Others As shown in Table 2, the proposed GPQ performs substantially better over all bit lengths than compared methods. Specifically, when we averaged mAP scores for all bit lengths, GPQ are 4.8%p and 4.6%p higher than the previous semi-supervised retrieval methods on CIFAR-10 and NUS-WIDE respectively. In particular, the performance gap is more pronounced as the number of bits decreases. This tendency is intimately related to the baseline hashing concepts. Comparing the results of the PQ-based and BH-based methods, we can identify that PQ-

Concept	Method	CIFAR-10				NUS-WIDE			
		12-bits	24-bits	32-bits	48-bits	12-bits	24-bits	32-bits	48-bits
Deep Semi-supervised	GPQ (Ours)	0.858	0.869	0.878	0.883	0.852	0.865	0.876	0.878
	SSGAH [10]	0.819	0.837	0.847	0.855	0.838	0.849	0.863	0.867
	BGDH [36]	0.805	0.824	0.826	0.833	0.810	0.821	0.825	0.829
	SSDH [40]	0.801	0.813	0.812	0.814	0.783	0.788	0.791	0.794
Deep Quantization	PQN [38]	0.795	0.819	0.823	0.830	0.803	0.818	0.822	0.824
	DTQ [19]	0.785	0.789	0.790	0.792	0.791	0.798	0.808	0.811
	DQN [2]	0.527	0.551	0.558	0.564	0.764	0.778	0.785	0.793
Deep Binary Hashing	SUBIC [11]	0.635	0.689	0.713	0.721	0.652	0.783	0.792	0.796
	NINH [18]	0.600	0.667	0.689	0.702	0.597	0.627	0.647	0.651
	CNNH [35]	0.496	0.580	0.582	0.583	0.536	0.522	0.533	0.531
Product Quantization	LOQP [15]	0.279	0.324	0.366	0.370	0.436	0.452	0.463	0.468
	OPQ [6]	0.265	0.315	0.323	0.345	0.429	0.433	0.450	0.458
	PQ [13]	0.237	0.265	0.268	0.266	0.398	0.406	0.413	0.422
Binary Hashing	SDH [26]	0.255	0.330	0.344	0.360	0.414	0.465	0.451	0.454
	ITQ [7]	0.158	0.163	0.168	0.169	0.428	0.430	0.432	0.435
	SH [34]	0.124	0.125	0.125	0.126	0.390	0.394	0.393	0.396

Table 2. The mean Average Precision (mAP) scores of different hashing algorithms on experimental protocol 1.

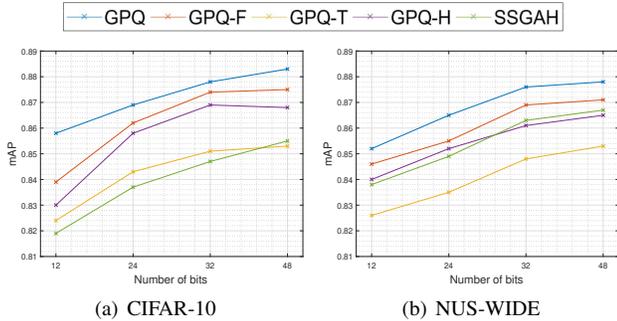


Figure 3. The comparison results of GPQ and its variants.

based ones are generally superior for both deep and non-deep cases especially for smaller bits. This is because unlike BH, PQ-based methods have the codewords of real values which enable mitigating the deviations generated during the encoding time, and they also allow more diverse distances through asymmetric calculation between database and query inputs. For the same reason, PQ-based GPQ with these advantages is able to achieve the state-of-the-art results in semi-supervised image retrieval.

Ablation Study To evaluate the contribution and importance of each component and training objectives in GPQ, we build three variants: 1) replace the feature extractor F with CNN-F [3] for $GPQ-F$; 2) remove the classifier C and learn the network only with the PQ table Z for $GPQ-H$; 3) exchange the N-pair Product Quantization loss \mathcal{L}_{N-PQ} with a standard Triplet loss [25] for $GPQ-T$. We perform retrieval experiments on protocol 1 for these variants, and empirically determine the sensitivity of GPQ.

The mAP scores of these variants and the original GPQ are reported in Figure 3. In this experiment, the hyper-parameters are all set identically as defaults in section 4.1, and the balancing parameter λ_1 and λ_2 are all set to 0.1. GPQ-F employs all the training objectives and shows the best mAP scores. It even outperforms other semi-

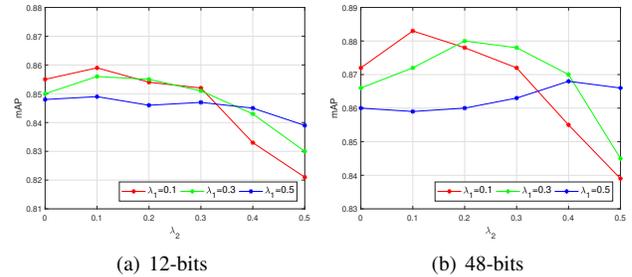


Figure 4. The sensitivity investigation of two balancing parameters: λ_1 and λ_2 .

supervised retrieval methods including SSGAH [10] in all experimental settings of protocol 1.

However, the modified VGG-based [12] original GPQ is superior to GPQ-F on both datasets. It reverses the general idea that CNN-F with more trainable parameters would have higher generalization capacity. This happens because high complexity does not always guarantee performance gains, same as the observations in [39]. Therefore, instead of increasing the network complexity, we focus on network structures that can extract more general features, and it results in high retrieval accuracy especially for semi-supervised situation.

To determine the contributions of Classification loss \mathcal{L}_{cls} and Subspace Entropy Mini-max loss \mathcal{L}_{SEM} , we carry out experiments with varying the balancing parameters of each loss function λ_1 and λ_2 , respectively. GPQ-H is equivalent to both λ_1 and λ_2 being zero, and the result is reported in Figure 3. Experimental results for the different options for λ_1 and λ_2 are detailed in Figure 4 for the bit lengths of 12 and 48. In general, high accuracy is achieved when the influences of \mathcal{L}_{cls} and \mathcal{L}_{SEM} are similar, and optimal results are obtained when both of two balancing parameters are set to 0.1 for 12 and 48 bit lengths.

In order to investigate the effect of the proposed metric learning strategy, we substitute N-pair Product Quantiza-

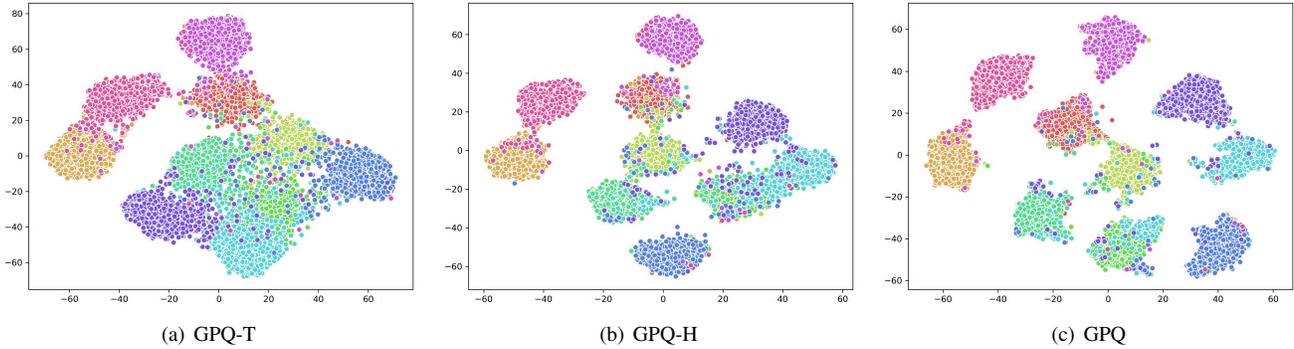


Figure 5. The t-SNE visualization of deep representations learned by GPQ-T, GPQ-H and GPQ on CIFAR-10 dataset respectively.

Concept	Method	CIFAR-10				NUS-WIDE			
		12-bits	24-bits	32-bits	48-bits	12-bits	24-bits	32-bits	48-bits
Deep Hashing	GPQ (Ours)	0.321	0.333	0.350	0.358	0.554	0.565	0.578	0.586
	SSGAH [10]	0.309	0.323	0.341	0.339	0.539	0.553	0.565	0.579
	SSDH [40]	0.285	0.291	0.311	0.325	0.510	0.533	0.549	0.551
	NINH [18]	0.241	0.249	0.253	0.272	0.484	0.483	0.485	0.487
	CNNH [35]	0.210	0.225	0.227	0.231	0.445	0.463	0.471	0.477
CNN features + non-Deep Hashing	SDH [26]	0.185	0.193	0.199	0.213	0.471	0.490	0.489	0.507
	ITQ [7]	0.157	0.165	0.189	0.201	0.488	0.493	0.508	0.503
	LOPQ [15]	0.134	0.127	0.126	0.124	0.416	0.386	0.380	0.379
	OPQ [6]	0.107	0.119	0.125	0.138	0.341	0.358	0.371	0.373

Table 3. The mean Average Precision scores (mAP) of different hashing algorithms on experimental protocol 2.

tion loss with triplet loss and visualize each result in Figure 5. While learning GPQ-T, the margin value is fixed at 0.1, and anchor, positive and negative pair are constructed using quantized vectors of different images. We utilize t-SNE algorithm [20] to examine the distribution of deep representations extracted from 1,000 images for each category in CIFAR-10. From the Figure where each color represents a different category, we can observe that GPQ better separates the data points.

Transfer Type Retrieval Following the transfer type investigation proposed in [23], we conduct experiments under protocol 2. Datasets are divided into two partitions of equal size, each of which is assigned to train-set and test-set, respectively. To be specific with each dataset, CIFAR-10 takes 7 categories and NUS-WIDE takes 15 categories to construct train75 and test75, and the rest categories for train25 and test25 respectively. For further comparison with hashing methods of non-deep learning concept [26, 7, 15, 6], we employ 7-th fully-connected features of pre-trained AlexNet as inputs and conduct image retrieval.

As we can observe in table 3, the average mAP score decreased compared to the results in protocol 1, because the label information of unseen categories disappears. In consequence, there is a noticeable mAP drop in supervised-based schemes, and the performance gap between the supervised and unsupervised learning concepts decreased. However, our GPQ method still outperforms other hashing methods. This is because, in GPQ, labeled information of known categories is fully exploited to learn the discriminative and ro-

bust codewords via metric learning algorithm, and at the same time, unlabeled data is fully explored to generalize the overall architecture through entropy control of feature distribution.

5. Conclusion

In this paper, we have proposed the first quantization based deep semi-supervised image retrieval technique, named Generalized Product Quantization (GPQ) network. We employed a metric learning strategy that preserves semantic similarity within the labeled data for the discriminative codebook learning. Further, we compute an entropy for each subspace and simultaneously maximize and minimize it to embed underlying information of the unlabeled data for the codebook regularization. Comprehensive experimental results justify that the GPQ yields state-of-the-art performance on large-scale image retrieval benchmark datasets.

Acknowledgement This work was supported in part by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No. 1711075689, Decentralised cloud technologies for edgeIoT integration in support of AI applications), and in part by Samsung Electronics Co., Ltd.

References

- [1] Relja Arandjelovic and Andrew Zisserman. All about vlad. In *CVPR*, pages 1578–1585, 2013. 2, 3
- [2] Yue Cao, Mingsheng Long, Jianmin Wang, Han Zhu, and Qingfu Wen. Deep quantization network for efficient image retrieval. In *AAAI*, 2016. 2, 3, 6, 7
- [3] Ken Chatfield, Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Return of the devil in the details: Delving deep into convolutional nets. In *BMVC*, 2014. 3, 6, 7
- [4] Tat-Seng Chua, Jinhui Tang, Richang Hong, Haojie Li, Zhiping Luo, and Yantao Zheng. Nus-wide: a real-world web image database from national university of singapore. In *CIVR*, page 48. ACM, 2009. 6
- [5] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. *ICML*, 2015. 2, 5
- [6] Tiezheng Ge, Kaiming He, Qifa Ke, and Jian Sun. Optimized product quantization for approximate nearest neighbor search. In *CVPR*, pages 2946–2953, 2013. 1, 2, 6, 7, 8
- [7] Yunchao Gong, Svetlana Lazebnik, Albert Gordo, and Florent Perronnin. Iterative quantization: A procrustean approach to learning binary codes for large-scale image retrieval. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(12):2916–2929, 2012. 1, 2, 6, 7, 8
- [8] Robert M. Gray and David L. Neuhoff. Quantization. *IEEE Transactions on Information Theory*, 44(6):2325–2383, 1998. 1
- [9] Jae-Pil Heo, Zhe Lin, and Sung-Eui Yoon. Distance encoded product quantization for approximate k-nearest neighbor search in high-dimensional space. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018. 2
- [10] Qinghao Hu, Jian Cheng, Zengguang Hou, et al. Semi-supervised generative adversarial hashing for image retrieval. In *ECCV*, pages 491–507. Springer, 2018. 2, 6, 7, 8
- [11] Himalaya Jain, Joaquin Zepeda, Patrick Pérez, and Rémi Gribonval. Subic: A supervised, structured binary code for image search. In *ICCV*, pages 833–842, 2017. 2, 6, 7
- [12] Young Kyun Jang, Dong-ju Jeong, Seok Hee Lee, and Nam Ik Cho. Deep clustering and block hashing network for face image retrieval. In *ACCV*, pages 325–339. Springer, 2018. 2, 3, 6, 7
- [13] Herve Jegou, Matthijs Douze, and Cordelia Schmid. Product quantization for nearest neighbor search. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(1):117–128, 2010. 1, 2, 3, 6, 7
- [14] Dong-ju Jeong, Sung-Kwon Choo, Wonkyo Seo, and Nam Ik Cho. Classification-based supervised hashing with complementary networks for image search. In *BMVC*, page 74, 2018. 2, 3
- [15] Yannis Kalantidis and Yannis Avrithis. Locally optimized product quantization for approximate nearest neighbor search. In *CVPR*, pages 2321–2328, 2014. 1, 2, 6, 7, 8
- [16] Benjamin Klein and Lior Wolf. End-to-end supervised product quantization for image search and retrieval. In *CVPR*, pages 5041–5050, 2019. 2, 3
- [17] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *NeurIPS*, pages 1097–1105, 2012. 3, 6
- [18] Hanjiang Lai, Yan Pan, Ye Liu, and Shuicheng Yan. Simultaneous feature learning and hash coding with deep neural networks. In *CVPR*, pages 3270–3278, 2015. 2, 6, 7, 8
- [19] Bin Liu, Yue Cao, Mingsheng Long, Jianmin Wang, and Jingdong Wang. Deep triplet quantization. *ACM Multimedia*, 2018. 2, 3, 6, 7
- [20] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(Nov):2579–2605, 2008. 8
- [21] Qingqun Ning, Jianke Zhu, Zhiyuan Zhong, Steven CH Hoi, and Chun Chen. Scalable image retrieval by sparse product quantization. *IEEE Transactions on Multimedia*, 19(3):586–597, 2016. 2
- [22] Aude Oliva and Antonio Torralba. Modeling the shape of the scene: A holistic representation of the spatial envelope. *International Journal of Computer Vision*, 42(3):145–175, 2001. 6
- [23] Alexandre Sablayrolles, Matthijs Douze, Nicolas Usunier, and Hervé Jégou. How should we evaluate supervised hashing? In *ICASSP*, pages 1732–1736. IEEE, 2017. 8
- [24] Kuniaki Saito, Shohei Yamamoto, Yoshitaka Ushiku, and Tatsuya Harada. Open set domain adaptation by backpropagation. In *ECCV*, pages 153–168, 2018. 2, 5
- [25] Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In *CVPR*, pages 815–823, 2015. 4, 7
- [26] Fumin Shen, Chunhua Shen, Wei Liu, and Heng Tao Shen. Supervised discrete hashing. In *CVPR*, pages 37–45, 2015. 1, 2, 6, 7, 8
- [27] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *NeurIPS*, pages 4077–4087, 2017. 2, 3
- [28] Kihyuk Sohn. Improved deep metric learning with multi-class n-pair loss objective. In *NeurIPS*, pages 1857–1865, 2016. 2, 4
- [29] Jingkuan Song, Lianli Gao, Li Liu, Xiaofeng Zhu, and Nicu Sebe. Quantization-based hashing: a general framework for scalable image and video retrieval. *Pattern Recognition*, 75:175–187, 2018. 2
- [30] Jun Wang, Sanjiv Kumar, and Shih-Fu Chang. Sequential projection learning for hashing with compact codes. In *ICML*, pages 1127–1134, 2010. 2
- [31] Jingdong Wang, Ting Zhang, Nicu Sebe, Heng Tao Shen, et al. A survey on learning to hash. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(4):769–790, 2017. 1, 2
- [32] Yen-Cehng Wei-Yu and Jia-Bin Zsolt, Yu-Chiang. A closer look at few-shot classification. In *ICLR*, 2019. 2, 3
- [33] Kilian Q Weinberger and Lawrence K Saul. Distance metric learning for large margin nearest neighbor classification. *Journal of Machine Learning Research*, 10(Feb):207–244, 2009. 4
- [34] Yair Weiss, Antonio Torralba, and Rob Fergus. Spectral hashing. In *NeurIPS*, pages 1753–1760, 2009. 1, 2, 6, 7

- [35] Rongkai Xia, Yan Pan, Hanjiang Lai, Cong Liu, and Shuicheng Yan. Supervised hashing for image retrieval via image representation learning. In *AAAI*, 2014. [2](#), [6](#), [7](#), [8](#)
- [36] Xinyu Yan, Lijun Zhang, and Wu-Jun Li. Semi-supervised deep hashing with a bipartite graph. In *IJCAI*, pages 3238–3244, 2017. [2](#), [3](#), [6](#), [7](#)
- [37] Litao Yu, Zi Huang, Fumin Shen, Jingkuan Song, Heng Tao Shen, and Xiaofang Zhou. Bilinear optimized product quantization for scalable visual content analysis. *IEEE Transactions on Image Processing*, 26(10):5057–5069, 2017. [2](#)
- [38] Tan Yu, Junsong Yuan, Chen Fang, and Hailin Jin. Product quantization network for fast image retrieval. In *ECCV*, pages 186–201, 2018. [2](#), [3](#), [6](#), [7](#)
- [39] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. 2017. [7](#)
- [40] Jian Zhang and Yuxin Peng. Ssdh: semi-supervised deep hashing for large scale image retrieval. *IEEE Transactions on Circuits and Systems for Video Technology*, 29(1):212–225, 2017. [2](#), [3](#), [6](#), [7](#), [8](#)