

# HiHPQ: Hierarchical Hyperbolic Product Quantization for Unsupervised Image Retrieval

Zexuan Qiu, Jiahong Liu, Yankai Chen, Irwin King

The Chinese University of Hong Kong  
{zxqiu22, jhliu22, ykchen, king}@cse.cuhk.edu.hk

## Abstract

Existing unsupervised deep product quantization methods primarily aim for the increased similarity between different views of the identical image, whereas the delicate multi-level semantic similarities preserved between images are overlooked. Moreover, these methods predominantly focus on the Euclidean space for computational convenience, compromising their ability to map the multi-level semantic relationships between images effectively. To mitigate these shortcomings, we propose a novel unsupervised product quantization method dubbed **Hierarchical Hyperbolic Product Quantization (HiHPQ)**, which learns quantized representations by incorporating hierarchical semantic similarity within hyperbolic geometry. Specifically, we propose a hyperbolic product quantizer, where the hyperbolic codebook attention mechanism and the quantized contrastive learning on the hyperbolic product manifold are introduced to expedite quantization. Furthermore, we propose a hierarchical semantics learning module, designed to enhance the distinction between similar and non-matching images for a query by utilizing the extracted hierarchical semantics as an additional training supervision. Experiments on benchmarks show that our proposed method outperforms state-of-the-art baselines.

## Introduction

Approximate Nearest Neighbor (ANN) search has attracted considerable interest in contemporary image retrieval systems owing to its remarkable search efficiency and exceptional performance. There are two main branches in ANN research: *Binary Hashing* (BH) (Salakhutdinov and Hinton 2009) and *Product Quantization* (PQ) (Jegou, Douze, and Schmid 2010). These two approaches both transform data into binary codes. BH maps high-dimensional data to the Hamming space, allowing fast distance calculations using bit-wise XOR operation. However, it has limitations on image similarity due to integer-only distance values. As for PQ, it decomposes the high-dimensional space and approximates distances between binary codes using pre-computed real-valued inter-codeword distance, reflecting richer similarity information. In industrial practice, PQ has been incorporated into numerous fast retrieval engines (Guo et al. 2020).

The standard PQ (Jegou, Douze, and Schmid 2010) and its variants (e.g., OPQ (Ge et al. 2013) and LOPQ (Kalan-

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

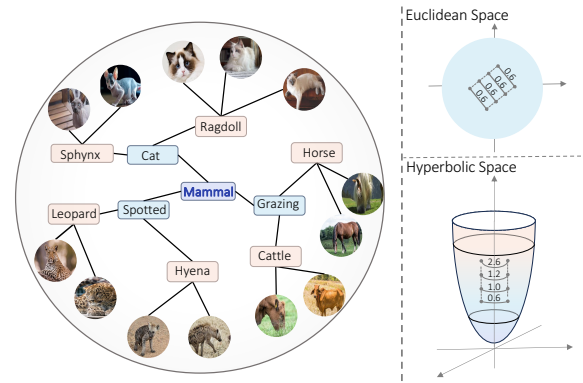


Figure 1: Left: Illustration of hierarchical semantics preserved in images. Right: Distance comparison on Euclidean space (top) and hyperbolic space (bottom). The 2-dimensional hyperbolic space example is depicted using the Lorentz model  $\mathcal{L}^2$  in  $\mathbb{R}^3$ .

tidis and Avrithis 2014)) are widely used in large-scale image retrieval systems. They rely on well-learned representations and heuristic algorithms for training the quantization module. Deep supervised PQ methods (Yue et al. 2016; Yu et al. 2018; Liu et al. 2018; Klein and Wolf 2019) have been developed to learn the neural encoding network and quantization module together using annotated labels. However, their practical application is challenging due to the high cost of annotating a large number of images. To address this, limited unsupervised deep PQ methods have emerged, utilizing reconstruction-based (Morozov and Babenko 2019) and contrastive objectives (Jang and Cho 2021; Wang et al. 2022) to optimize quantized representations.

Current unsupervised product quantization methods based on contrastive learning solely focus on striving for greater similarity between different views of the same image, disregarding the inherent hierarchical semantic structures commonly found in images. For instance, as illustrated in Figure 1, visual characteristics bring together *leopards* and *hyenas* into a *spotted* cluster, allowing them to be categorized more coarsely as *mammals* alongside similar animals like *cats*. Thus, in order to distinguish between the closest neighbors of a query and non-matching items during

the retrieval phase of PQ, it is vital to understand the nuanced levels of similarity among instances, instead of simply treating all non-query images as negatives. As supporting evidence of our motivation, such hierarchical semantic similarity has been shown to be helpful for performance improvement in other fields like visual representations (Guo et al. 2022), metric learning (Yan et al. 2021), classification (Zhang et al. 2022; Liu et al. 2022c,a; Ma et al. 2023) and style transfer (Li et al. 2020a).

On the other hand, *Euclidean space* has become a workhorse for existing PQ methods, where the Euclidean distance between quantized representations serves as a measure of similarity between images. However, *hyperbolic space* (i.e., non-Euclidean space) has shown less distance distortion than Euclidean space when fitting hierarchical data, such as word embeddings (Tifrea, Bécigneul, and Ganea 2018), graphs (Chen et al. 2022b; Liu et al. 2022b) and image-text retrieval (Desai et al. 2023). This can be attributed to the exponentially-evolved distance measurement where samples far away from the origin present a larger distance, as demonstrated on the right side of Figure 1. Consequently, by integrating hyperbolic geometry into the quantization representations, there is great potential to enhance the capturing of the structured semantics of an image.

Motivated by the aforementioned limitations, we introduce a new unsupervised PQ technique, dubbed **H**ierarchical **H**yperbolic **P**roduct **Q**uantization (HiHPQ). Generally, it enhances quantized representation learning by considering the hierarchical semantics within hyperbolic geometry:

- We propose a novel hyperbolic product quantizer. In this framework, each low-dimensional subspace is represented as a Lorentzian manifold with varying curvature, where the data points and codewords are both embedded in the hyperbolic space. Then a novel hyperbolic codebook attention mechanism is proposed to seamlessly extend the differentiable quantization process into the hyperbolic space. Additionally, we design a quantized contrastive learning approach based on the hyperbolic distance for model optimization.
- We design a hierarchical semantics learning module. Specifically, in addition to vanilla view-augmented semantics as the training signal, we utilize hierarchical clustering to extract pseudo hierarchical semantics as extra supervision for product quantization. To ensure retrieval of similar terms for the query, we extract two types of positive samples: the closest prototype in each hierarchy and instances sharing the same closest prototype. Then the prototype-wise and instance-wise contrastive losses that incorporate these two additional positive samples are employed to inject the extracted hierarchical semantics into quantized representations.
- Our HiHPQ surpasses state-of-the-art baselines on benchmark image datasets, with empirical analyses confirming the effectiveness of each proposed component.

## Related Work

**Product Quantization** Product quantization was initially proposed in the field of source coding (Gray 1984) and later

applied to ANN search (Jegou, Douze, and Schmid 2010). In large-scale image retrieval systems, PQ is often combined with inverted index (Jégou et al. 2011; Baranchuk, Babenko, and Malkov 2018) or inverted multi-index (Babenko and Lempitsky 2014) to achieve faster similarity search. Typical variants of PQ, such as OPQ (Ge et al. 2013) and Cartesian K-Means (Norouzi and Fleet 2013), focus on computing an optimized rotation matrix to make each subspace orthogonal to each other. Based on OPQ, LOPQ (Kalantidis and Avrithis 2014) employs a locally optimized individual product quantizer per cell of the inverted index. Afterwards, a series of end-to-end supervised deep PQ methods for image retrieval (Yue et al. 2016; Liu et al. 2018; Klein and Wolf 2019; Xiao et al. 2021) are proposed with powerful representation capabilities of neural networks. To address the expensive data annotation problem, limited unsupervised PQ methods have been proposed. Among them, the earlier approach (Chen, Cheung, and Wang 2018; Morozov and Babenko 2019) is based on the auto-encoder, while later methods (Jang and Cho 2021; Wang et al. 2022) are based on the contrastive learning frameworks. Also, Chen, Li, and Sun (2020) and Lu et al. (2023) study improved techniques of differentiable quantization, particularly when the codebook adheres to certain constraints.

**Binary Hashing** In recent years, there has been a surge of research dedicated to unsupervised binary hashing (Chen et al. 2023, 2022a,c; Zhang and Zhu 2020, 2019). An established type of unsupervised hashing method is based on deep generative models. Some of them (Dai et al. 2017; Shen, Liu, and Shao 2019; Shen et al. 2020; Li and van Gemert 2021) utilize the encoder-decoder architecture (Kingma and Welling 2013) which expects the binary code to recover the original input, while others of them (Dizaji et al. 2018; Zieba et al. 2018; Song et al. 2018) utilize generative adversarial networks (Goodfellow et al. 2014) to implicitly maximize the reconstruction likelihood through the discriminator. Another typical category of hashing methods aims to preserve predefined similarity between images using binary codes. In this category, methods like (Yang et al. 2018, 2019; Tu, Mao, and Wei 2020) leverage pre-trained continuous features to define pairwise similarities, while methods like (Lin et al. 2016; Luo et al. 2020; Qiu et al. 2021; Lin et al. 2022) involves utilizing random augmentations to establish pairs of images that share common characteristics. This type of similarity-preserving method is also widely used in hashing-based video search (Li et al. 2019, 2021). Further, Ma et al. (2022) and Wei et al. (2022) both propose to employ fine-grained semantic similarity to improve hashing.

## Preliminaries

### Lorentz Model of Hyperbolic Space

In this part, we briefly review the concepts and basic operations on hyperbolic geometry. A thorough and in-depth explanation can be found in (Nickel and Kiela 2018). Multiple isometric models are commonly used to represent hyperbolic geometry. In this work, we adopt the typical Lorentz model to depict the hyperbolic space.

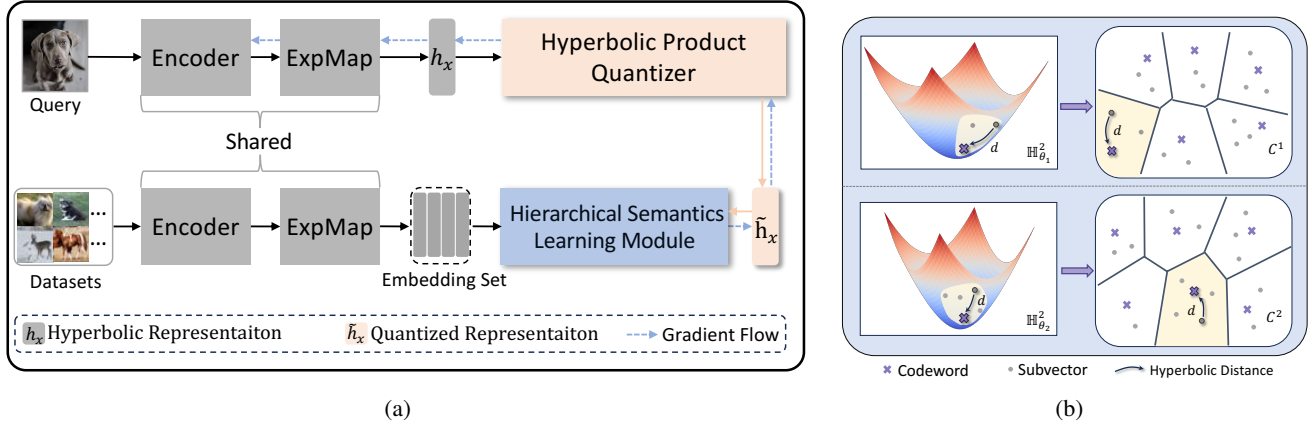


Figure 2: (a) The architecture of HiHPQ. “ExpMap” denotes the exponential map for short. (b) An example of our hyperbolic product quantizer. In this example, there are two hyperbolic spaces  $\mathbb{H}_{\theta_1}^2$  and  $\mathbb{H}_{\theta_2}^2$  of 2-dimension depicted by the Lorentz model in  $\mathbb{R}^3$ , where  $\theta_1$  and  $\theta_2$  denotes the curvature parameter. On the right side, two codebooks  $C^1$  and  $C^2$  are displayed using 2D Voronoi diagrams. Subvectors will be quantized by codewords in corresponding codebooks via the hyperbolic distance metric.

**Lorentz Model** Lorentz model depicts the  $d$ -dimensional hyperbolic space (i.e., Lorentzian manifold) as a submanifold of  $\mathbb{R}^{d+1}$ . Formally, a  $d$ -dimensional Lorentz model with the constant negative curvature  $-\theta$  is defined as the following set of vectors:

$$\mathbb{H}_{\theta}^d = \{\mathbf{x} \in \mathbb{R}^{d+1} : \langle \mathbf{x}, \mathbf{x} \rangle_{\mathcal{L}} = -1/\theta, \theta > 0\}, \quad (1)$$

where  $\langle \cdot, \cdot \rangle_{\mathcal{L}}$  denotes the *Lorentzian inner product* of the hyperbolic space:  $\langle \mathbf{x}, \mathbf{y} \rangle_{\mathcal{L}} = -x_0y_0 + x_1y_1 + x_2y_2 + \dots + x_dy_d$ .

**Lorentz Distance** The Lorentzian distance between  $\mathbf{x}, \mathbf{y} \in \mathbb{H}_{\theta}^d$  is defined as:

$$d_{\mathbb{H}_{\theta}^d}(\mathbf{x}, \mathbf{y}) = \sqrt{1/\theta} \cosh^{-1}(-\theta \langle \mathbf{x}, \mathbf{y} \rangle_{\mathcal{L}}), \quad (2)$$

which depicts the length of the shortest path (i.e., *geodesic*) between two points on the manifold. With  $\theta \rightarrow 0$ , the distance in Eq. (2) reduces the Euclidean distance.

**Tangent Space** For any vector  $\mathbf{p} \in \mathbb{H}_{\theta}^d$ , the tangent space  $\mathcal{T}_{\mathbf{p}}\mathbb{H}_{\theta}^d$  is defined as the first-order Euclidean approximation of  $\mathbb{H}_{\theta}^d$  around  $\mathbf{p}$ :

$$\mathcal{T}_{\mathbf{p}}\mathbb{H}_{\theta}^d = \{\mathbf{v} \in \mathbb{R}^{d+1} : \langle \mathbf{v}, \mathbf{p} \rangle_{\mathcal{L}} = 0\}. \quad (3)$$

For any ambient Euclidean vector  $\mathbf{u} \in \mathbb{R}^{d+1}$ , it can be projected onto the tangent space  $\mathcal{T}_{\mathbf{p}}\mathbb{H}_{\theta}^d$  via the orthogonal projection  $\text{proj}_{\mathbf{p}}(\mathbf{u})$ :

$$\text{proj}_{\mathbf{p}}(\mathbf{u}) = \mathbf{u} + \theta \mathbf{p} \langle \mathbf{p}, \mathbf{u} \rangle_{\mathcal{L}}. \quad (4)$$

**Exponential Map** The exponential map offers a method to project vectors from the tangent space onto the manifold. For a point  $\mathbf{p}$  on the manifold, its exponential map is defined as  $\text{exp}_{\mathbf{p}}^{\theta} : \mathcal{T}_{\mathbf{p}}\mathbb{H}_{\theta}^d \rightarrow \mathbb{H}_{\theta}^d$  is defined as:

$$\text{exp}_{\mathbf{p}}^{\theta}(\mathbf{v}) = \cosh(\sqrt{\theta} \|\mathbf{v}\|_{\mathcal{L}}) \mathbf{p} + \frac{\sinh(\sqrt{\theta} \|\mathbf{v}\|_{\mathcal{L}})}{\sqrt{\theta} \|\mathbf{v}\|_{\mathcal{L}}} \mathbf{v}, \quad (5)$$

where  $\|\mathbf{v}\|_{\mathcal{L}} = \sqrt{\langle \mathbf{v}, \mathbf{v} \rangle_{\mathcal{L}}}$  is the Lorentzian norm of  $\mathbf{v}$ .

## Methodology

### Problem Formulation and Model Overview

The goal of our work is to learn a product quantization mapping:  $x \mapsto \tilde{\mathbf{b}}_x$ , where  $x$  is an image of the unlabeled dataset  $\mathcal{X} = \{x_i\}_{i=1}^N$ , and  $\tilde{\mathbf{b}}_x \in \{0, 1\}^B$  is a  $B$ -bit quantization code. During the retrieval phase,  $\tilde{\mathbf{b}}_x$  will recover the quantized representation  $\tilde{\mathbf{h}}_x \in \mathbb{R}^D$  to enable efficient retrieval. Our proposed model, as illustrated in Figure 2(a), is mainly composed of: (i) an encoder network that utilizes a (CNN) backbone followed by a linear projector, functioning as a feature extractor; (ii) a hyperbolic product quantizer endowed with a proposed soft codebook quantization mechanism for quantizing representations in hyperbolic space; and (iii) a hierarchical semantics learning module designated to enhance the quantized representations based on the extracted hierarchical similarity.

### Hyperbolic Product Quantizer

**Hyperbolic Product Space** Inspired by the recent development of hyperbolic manifold (Gu et al. 2018; Gao et al. 2022), instead of utilizing Euclidean geometry, we propose representing the vector space  $\mathbb{S}$  as the Cartesian product of  $M$   $d$ -dimension hyperbolic subspaces:

$$\mathbb{S} = \mathbb{H}_{\theta_1}^d \times \mathbb{H}_{\theta_2}^d \times \dots \times \mathbb{H}_{\theta_M}^d, \quad (6)$$

where  $-\theta_i$  ( $\theta_i > 0$ ) denotes the negative constant curvature for the  $i$ -th Lorentzian manifold  $\mathbb{H}_{\theta_i}^d$ . Similarly, for the reference point  $\mathbf{p} = (\mathbf{p}^1, \dots, \mathbf{p}^M)$  on the manifold  $\mathbb{S}$ , the approximated tangent space  $\mathcal{T}_{\mathbf{p}}\mathbb{S}$  of  $\mathbf{p}$  can be decomposed by the Cartesian product of  $M$  tangent spaces. Namely, for  $\mathbf{v} \in \mathcal{T}_{\mathbf{p}}\mathbb{S}$ , we have  $\mathbf{v} = (\mathbf{v}^1, \dots, \mathbf{v}^M)$ , where  $\mathbf{v}^m \in \mathcal{T}_{\mathbf{p}^m}\mathbb{H}_{\theta_m}^d$ .

During training, given an image  $x$ , we pass it through the encoder network to get the ambient Euclidean embedding  $\mathbf{z}_x$  and slice it into  $M$  segments  $\{\mathbf{z}_x^m\}_{m=1}^M$ . Then, for each

segment  $\mathbf{z}_x^m \in \mathbb{R}^{d+1}$ , we transform it onto the tangent space  $\mathcal{T}_{\mathbf{p}}$  to get the tangent vector  $\mathbf{v}_x^m$  via the orthogonal projection  $\mathbf{v}_x^m = \text{proj}_{\mathcal{T}_{\mathbf{p}}}(\mathbf{z}_x^m)$ . For the obtained tangent vector  $\mathbf{v}_x = (\mathbf{v}_x^1, \dots, \mathbf{v}_x^M)$ , we transfer it onto the product manifold  $\mathbb{S}$  and get the final hyperbolic embedding  $\mathbf{h}_x \in \mathbb{S}$  by:

$$\mathbf{h}_x = \exp_{\mathbf{p}}(\mathbf{v}_x) = \left( \exp_{\mathbf{p}^1}^{\theta_1}(\mathbf{v}_x^1), \dots, \exp_{\mathbf{p}^M}^{\theta_M}(\mathbf{v}_x^M) \right). \quad (7)$$

In practice, we adopt  $\mathbf{p} = (\mathbf{o}^1, \dots, \mathbf{o}^M)$  as the reference point when applying exponential map operations, where  $\mathbf{o}^m = \left( \sqrt{1/\theta_m}, \mathbf{0} \right)$  ( $\mathbf{0} \in \mathbb{R}^d$ ) is the hyperbolic original point on the Lorentzian Manifold  $\mathbb{H}_{\theta_m}^d$ .

As for the distance measurement on the product manifold  $\mathbb{S}$ , the distance between any pair  $\mathbf{a}, \mathbf{b} \in \mathbb{S}$  simply decomposes into sum of  $M$  Lorentzian distances:

$$d_{\mathbb{S}}(\mathbf{a}, \mathbf{b}) = \sum_{m=1}^M d_{\mathcal{L}_m}(\mathbf{a}, \mathbf{b}), \quad (8)$$

where  $d_{\mathcal{L}_m}(\mathbf{a}, \mathbf{b})$  denotes the Lorentzian distance of the  $m$ -th Lorentzian manifold  $\mathbb{H}_{\theta_m}^d$ .

**Hyperbolic Codebook Quantization Mechanism** The performance of product quantization is greatly influenced by the choice of the quantization technique. In recent methods (Chen, Li, and Sun 2020; Jang and Cho 2021; Wang et al. 2022), the differentiable soft quantization technique has been widely used and proven more effective and feasible than the conventional hard quantization. In order to facilitate end-to-end training with differentiability, we propose a novel soft hyperbolic codebook quantization mechanism. Specifically, we represent the codebook  $C$  as the Cartesian product of  $M$  smaller codebooks, namely  $C = C^1 \times C^2 \times \dots \times C^M$ , where the  $m$ -th codebook  $C^m = \{\mathbf{c}_k^m \in \mathbb{H}_{\theta_m}^d\}_{k=1}^K$  consists of  $K$  codewords distributed on the  $m$ -th Lorentzian manifold. Then, for the representation  $\mathbf{h}_x^m$ , we aim to obtain the soft quantized one  $\tilde{\mathbf{h}}_x^m \in \mathbb{H}_{\theta_m}^d$  as its approximation by calculating the weighted aggregation inside the codebook  $C^m$ . One reasonably desirable property of  $\tilde{\mathbf{h}}_x^m$  is that the expected squared distance from  $\tilde{\mathbf{h}}_x^m$  to the codeword set  $C^m$  should be minimum, *i.e.*,

$$\tilde{\mathbf{h}}_x^m \triangleq \arg \min_{\mu \in \mathbb{H}_{\theta_m}^d} w_k^m d_{\mathcal{L}_m}^2(\mathbf{c}_k^m, \mu), \quad (9)$$

where  $w_k^m$  is the computed weight of the  $k$ -th hyperbolic codeword  $\mathbf{c}_k^m$  in the  $m$ -th Lorentzian manifold. Intuitively, it ensures that the quantized embedding is close to those codewords with larger weights. It was proved in (Law et al. 2019) that with the squared Lorentzian distance defined as:

$$d_{\mathcal{L}_m}^2(\mathbf{a}, \mathbf{b}) = \|\mathbf{a} - \mathbf{b}\|_{\mathcal{L}}^2 = -2\theta_m - 2\langle \mathbf{a}, \mathbf{b} \rangle_{\mathcal{L}}, \quad (10)$$

then Eq. (9) has the closed-form solution:

$$\tilde{\mathbf{h}}_x^m = \frac{\sum_{k=1}^K w_k^m \mathbf{c}_k^m}{\sqrt{-1/\theta_m} \left\| \sum_{k=1}^K w_k^m \mathbf{c}_k^m \right\|_{\mathcal{L}}}, \quad (11)$$

where  $\|\mathbf{a}\|_{\mathcal{L}} = \sqrt{\|\mathbf{a}\|_{\mathcal{L}}^2}$  is the modulus (Ratcliffe 2006) of the imaginary Lorentzian norm of the vector  $\mathbf{a}$ . Here, the

attention weight  $w_k^m$  is computed using softmax relaxation:

$$w_k^m = \frac{\exp(-d_{\mathcal{L}_m}^2(\mathbf{c}_k^m, \mathbf{h}_x^m)/\tau)}{\sum_j^K \exp(-d_{\mathcal{L}_m}^2(\mathbf{c}_j^m, \mathbf{h}_x^m)/\tau)}, \quad (12)$$

where  $\tau$  is a hyper-parameter. After quantizing each subspace separately, we can finally get the quantized hyperbolic embedding  $\tilde{\mathbf{h}}_x = (\tilde{\mathbf{h}}_x^1, \dots, \tilde{\mathbf{h}}_x^M)$  as the approximation of original continuous embedding  $\mathbf{h}_x$  on the product manifold  $\mathbb{S}$ . In inference, we simply use the hard quantization operation to obtain the quantized representation. For more inference details, please refer to the supplementary material.

**Hyperbolic Quantized Contrastive Learning** Inspired by the recent success (Yan et al. 2021; Ge et al. 2023) of distance learning on hyperbolic space, we introduce a hyperbolic quantized contrastive learning scheme aimed at jointly both the encoder network and the quantization module. Specifically, given a mini-batch of training images  $\mathcal{B}$  with the batch size  $N_B$ , we apply random augmentations to each image twice, resulting in total  $2N_B$  views  $\{(x_i^{(1)}, x_i^{(2)})\}_{i=1}^{N_B}$ . After passing them through our encoder and the quantization module, we can obtain the quantized hyperbolic representations  $\{(\tilde{\mathbf{h}}_{x_i}^{(1)}, \tilde{\mathbf{h}}_{x_i}^{(2)})\}_{i=1}^{N_B}$ . Then, the hyperbolic quantized contrastive loss is defined as:

$$\mathcal{L}_{aug} = \frac{1}{N_B} \sum_{x \in \mathcal{B}} (\ell^1(x) + \ell^2(x)), \quad (13)$$

where  $aug$  denotes ‘‘augmented’’ for short, and  $\ell^j(x)$  for  $j = \{1, 2\}$  is defined as:

$$\ell^j(x) = -\log \frac{\mathcal{S}(\tilde{\mathbf{h}}_x^{(1)}, \tilde{\mathbf{h}}_x^{(2)})}{\mathcal{S}(\tilde{\mathbf{h}}_x^{(1)}, \tilde{\mathbf{h}}_x^{(2)}) + \sum_{\substack{t \in \mathcal{B} \setminus x \\ n=1,2}} \mathcal{S}(\tilde{\mathbf{h}}_x^{(1)}, \tilde{\mathbf{h}}_t^{(n)})}. \quad (14)$$

Intuitively, for the hyperbolic quantized embedding  $\tilde{\mathbf{h}}_x^{(1)}$  as the query,  $\tilde{\mathbf{h}}_x^{(2)}$  is defined as the positive key, while other  $2N_B - 2$  quantized embeddings from other images within mini-batch are considered negative keys. Different from previous works, our distance metric works on the hyperbolic product manifold  $\mathbb{S}$ , thus  $\mathcal{S}(\tilde{\mathbf{h}}_1, \tilde{\mathbf{h}}_2)$  is defined as:

$$\mathcal{S}(\tilde{\mathbf{h}}_1, \tilde{\mathbf{h}}_2) \triangleq \exp(-d_{\mathbb{S}}(\tilde{\mathbf{h}}_1, \tilde{\mathbf{h}}_2)/\tau_{qc}). \quad (15)$$

In Eq. (15), the negative of the distance defined on the product manifold  $\mathbb{S}$  is taken as the similarity score on  $\mathcal{S}$ , which implies that the similarity between two images is determined by hyperbolic distances on  $M$  decomposed Lorentzian manifold jointly; and  $\tau_{qc}$  is a non-negative temperature.

## Hierarchical Semantics Learning Module

**Hierarchical Semantics Extraction** Inspired by (Yan et al. 2021; Guo et al. 2022), we propose a hierarchical semantics learning module to enhance the semantics of hyperbolic quantized representation with explicit hierarchical semantic similarity as the extra training signal. To this end, we resort to bottom-up hierarchical clustering to extract pseudo

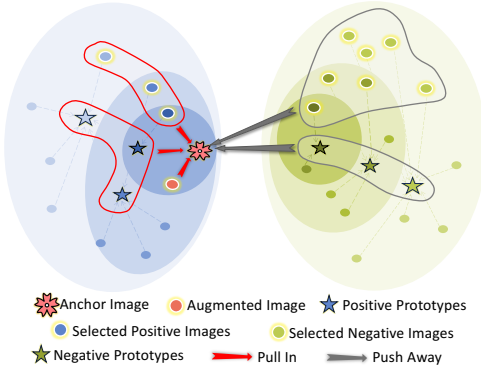


Figure 3: Illustration of both instance-wise and prototype-wise contrastive learning based on our extracted hierarchy.

hierarchical semantic similarity. Notably, since direct clustering in the hyperbolic space is challenging, our pseudo hierarchical semantics between images is extracted based on the Euclidean tangent space  $\mathcal{T}_p\mathbb{S}$ . Our subsequent experiments showed that the semantic hierarchy refined with the Euclidean tangent space has sufficiently benefited the quality of our quantized hyperbolic representations.

Specifically, in the bottom-up merging step, the two closest sub-clusters will be merged to form a new cluster. The distance between two sub-clusters is approximated by the distance between their corresponding prototypes as:

$$d_{ab} = \|\mathbf{e}_a - \mathbf{e}_b\|, \quad \mathbf{e}_a = \frac{1}{n_a} \sum_{\mathbf{v}_j \in G_a} \mathbf{v}_j, \quad (16)$$

where  $\|\cdot\|$  denotes the Euclidean distance,  $n_a$  represents the number of samples belonging to the sub-cluster  $G_a$ , and the prototype  $\mathbf{e}_a$  is the average of all tangent vectors in  $G_a$ .

In practice, the distance threshold for merging can be set as sufficiently high, allowing the merging process to continue regardless of the distance of the closest sub-cluster pair. Then, we can pre-define  $L$  levels, with each level containing  $N_l$  clusters. When the bottom-up merging reaches a cluster count of  $N_l$ , we save the corresponding clustering result. Ultimately, a hierarchical similarity structure with  $L$  levels is obtained. After completing the hierarchical clustering, as the obtained prototypes are in the tangent space, we need to map them to the product manifold  $\mathbb{S}$  so that they can interact with our quantized hyperbolic representations. For this purpose, the prototype  $\mathbf{e}_j = (\mathbf{e}_j^1, \dots, \mathbf{e}_j^M)$  is transformed to the product manifold via the exponential map:

$$\tilde{\mathbf{e}}_j = \exp_{\mathbf{p}}(\mathbf{e}_j) = \left( \exp_{\mathbf{p}^1}^{\theta_1}(\mathbf{e}_j^1), \dots, \exp_{\mathbf{p}^M}^{\theta_M}(\mathbf{e}_j^M) \right). \quad (17)$$

Practically, directly performing hierarchical clustering on all training images is time-consuming. Therefore, during the hierarchical clustering phase, we first apply K-means clustering to the training data to obtain a sufficient number of sub-clusters, and then perform iterative merging to get the desired  $L$  hierarchies. Moreover, we perform hierarchical clustering at the beginning of every epoch for training efficiency.

## Prototype-Wise & Instance-Wise Semantics Learning

The multi-level hierarchical semantics relationship we uncover consists of two components: the affiliation relationship between images and their corresponding prototypes, and the assignment relationship between images within the semantics structure. Therefore, we resort to both prototype-wise (Li et al. 2020b; Guo et al. 2022) and instance-wise contrastive loss to jointly learn these two aspects. Specifically, given a query image  $x_i$ , we denote the hyperbolic prototype  $\tilde{\mathbf{e}}_j^l(i)$  that  $x_i$  is assigned to in  $l$ -th hierarchy and regard  $(\tilde{\mathbf{h}}_{x_i}, \tilde{\mathbf{e}}_j^l(i))$  as a positive pair, while other prototypes at the same level in the hierarchy play the role of negative prototypes. Then the prototype-wise contrastive loss is defined as:

$$\mathcal{L}_{prot} = - \sum_{i=1}^{N_B} \frac{1}{L} \sum_{l=1}^L \log \frac{\mathcal{S}(\tilde{\mathbf{h}}_{x_i}, \tilde{\mathbf{e}}_j^l(i))}{\sum_{n=1}^{N_l} \mathcal{S}(\tilde{\mathbf{h}}_{x_i}, \tilde{\mathbf{e}}_n^l)}, \quad (18)$$

where the similarity metric  $\mathcal{S}(\cdot, \cdot)$  is the negative hyperbolic distance defined in Eq. (15).

While the prototypical contrastive loss in Eq. (18) can preserve similarity between the image and their closest prototype in a multi-granularity hierarchy, it neglects the proximity between images belonging to the same prototype in each hierarchy. To alleviate this issue, we fall back on instance-wise contrastive learning. Specifically, given an image  $x_i$ , we randomly select an image  $x_{il}$  from the pool of images assigned to the prototypes that  $x_i$  belongs to at each  $l$ -th hierarchy. Then the hierarchical instance-wise contrastive loss is defined as:

$$\mathcal{L}_{ins} = - \sum_{i=1}^{N_B} \frac{1}{L} \sum_{l=1}^L \log \frac{\mathcal{S}(\tilde{\mathbf{h}}_{x_i}, \tilde{\mathbf{h}}_{x_{il}})}{\sum_{t \in \mathcal{B} \setminus x_i} \mathcal{S}(\tilde{\mathbf{h}}_{x_i}, \tilde{\mathbf{h}}_t)}. \quad (19)$$

Then, our training objective for the pseudo hierarchical similarity is defined as:

$$\mathcal{L}_{hs} = \lambda_1 \mathcal{L}_{prot} + \lambda_2 \mathcal{L}_{ins}, \quad (20)$$

where  $\lambda_1$  and  $\lambda_2$  are the hyper-parameters used to control the relative importance of the two losses. Figure 3 visualizes our hierarchical semantics learning framework. Overall, the ultimate training objective for HiHPQ is:

$$\mathcal{L} = \mathcal{L}_{aug} + \mathcal{L}_{hs}. \quad (21)$$

## Experiments

### Experimental Setup

**Datasets** The proposed method is evaluated using **Flickr25K** (Huiskes and Lew 2008), **NUS-WIDE** (Chua et al. 2009), as well as two experimental protocols CIFAR-10 (I) and CIFAR-10 (II) both of which are based on **CIFAR-10** (Krizhevsky, Hinton et al. 2009). Further details can be found in the supplementary material.

**Evaluation Metrics** We utilize the Mean Average Precision (MAP) at top N as the evaluation metric to assess the proposed method. We adopt the MAP@1000 for CIFAR-10(I) and CIFAR-10 (I) and (II), and MAP@5000 for both NUS-WIDE and Flickr25K.

Method	Type	Flickr25K			CIFAR-10 (I)			CIFAR-10 (II)			NUS-WIDE		
		16bits	32bits	64bits	16bits	32bits	64bits	16bits	32bits	64bits	16bits	32bits	64bits
DeepBit	BH	62.04	66.54	68.34	19.43	24.86	27.73	20.60	28.23	31.30	39.20	40.30	42.90
SGH	BH	72.10	72.84	72.83	34.51	37.04	38.93	43.50	43.70	43.30	59.30	59.00	60.70
HashGAN	BH	72.11	73.25	75.46	44.70	46.30	48.10	42.81	47.54	47.29	68.44	70.56	71.71
GreedyHash	BH	69.97	70.85	73.03	44.80	47.20	50.10	45.76	48.26	53.34	63.30	69.10	73.10
BinGAN	BH	-	-	-	-	-	-	47.60	51.20	52.00	65.40	70.90	71.30
BGAN	BH	-	-	-	-	-	-	52.50	53.10	56.20	68.40	71.40	73.00
SSDH	BH	75.65	77.10	76.68	36.16	40.17	44.00	33.30	38.29	40.81	58.00	59.30	61.00
DVB	BH	-	-	-	-	-	-	40.30	42.20	44.60	60.40	63.20	66.50
TBH	BH	74.38	76.14	77.87	54.68	58.63	62.47	53.20	57.30	57.80	71.70	72.50	73.50
Bi-half Net	BH	76.07	77.93	78.62	56.10	57.60	59.50	49.97	52.04	55.35	76.86	78.31	79.94
CIBHash	BH	77.21	78.43	79.59	59.39	63.67	65.16	59.00	62.20	64.10	79.00	80.70	81.50
PQ+VGG	PQ	62.75	66.63	69.40	27.14	33.30	37.67	28.16	30.24	30.61	65.39	67.41	68.56
OPQ+VGG	PQ	63.27	68.01	69.86	27.29	35.17	38.48	32.17	33.50	34.46	65.74	68.38	69.12
DeepQuan	PQ	-	-	-	39.95	41.25	43.26	-	-	-	-	-	-
SPQ	PQ	77.35	78.74	79.98	63.17	66.88	68.02	56.56	61.45	63.30	78.51	80.41	81.70
MeCoQ	PQ	<b>81.31</b>	81.71	82.68	68.20	69.74	71.06	62.88	64.09	65.07	78.31	81.59	81.80
HiHPQ	PQ	80.74	<b>82.55</b>	<b>82.95</b>	<b>70.56</b>	<b>73.22</b>	<b>73.71</b>	<b>63.31</b>	<b>65.81</b>	<b>67.11</b>	<b>79.86</b>	<b>82.05</b>	<b>82.63</b>

Table 1: MAP (%) comparison between different unsupervised state-of-the-art efficient retrieval methods on benchmark datasets. The types of methods used are denoted as ‘‘PQ’’ for ‘‘Product Quantization’’ and ‘‘BH’’ for ‘‘Binary Hashing’’.

**Baselines** We consider following baselines for comparison: (i) binary hashing methods: ; DeepBit (Lin et al. 2016), SGH (Dai et al. 2017), HashGAN (Dizaji et al. 2018), GreedyHash (Su et al. 2018), BinGAN (Zieba et al. 2018), BGAN (Song et al. 2018), SSDH (Yang et al. 2018), DVB (Shen, Liu, and Shao 2019), TBH (Shen et al. 2020), Bi-half Net (Li and van Gemert 2021), CIBHash (Qiu et al. 2021); and (ii) product quantization methods: PQ (Jegou, Douze, and Schmid 2010), OPQ (Ge et al. 2013); DeepQuan (Chen, Cheung, and Wang 2018), SPQ<sup>1</sup> (Jang and Cho 2021), MeCoQ (Wang et al. 2022).

### Implementation Details

The pre-defined hierarchies for different datasets are: [50, 20] on Flickr25K; [200, 100, 50] on CIFAR-10 (I); [100, 50, 25] on CIFAR-10 (II); [200, 100, 75] on NUS-WIDE. Please refer to the supplementary material for more implementation details.

### Results and Analyses

**Overall Performance** Table 1 presents the performances of our proposed model and existing baselines on public benchmark image datasets. It can be seen that our proposed HiHPQ outperforms previous shallow and deep quantization methods by a significant margin across almost all experimental settings. Among all the baselines compared, our method closely resembles SPQ and MeCoQ, which propose cross-quantized contrastive learning and debiased contrastive learning with code memory, respectively. However, these two methods significantly underperform HiHPQ, which can be attributed to HiHPQ’s hierarchical semantics

<sup>1</sup>As SPQ uses a totally different backbone and training strategy, we reproduce SPQ in our setting using their codebase.

learning in the hyperbolic product manifold. In Figure 4(a), we further demonstrate the retrieval performance of HiHPQ with varying numbers of codebooks. It can be observed that regardless of the number of codebooks, our model outperforms the state-of-the-art SPQ and MeCoQ by a significant margin, showing the superiority of the HiHPQ.

**Ablation Study** We configure six variants of HiHPQ for investigation. (i) *Vanilla-Hyperbolic*: by setting both  $\lambda_1 = 0$  and  $\lambda_2 = 0$ , it discards the extracted hierarchy as the extra training signal. (ii) *Instance-Hyperbolic*: by setting  $\lambda_1 = 0$  and  $\lambda_2 = 1$ , it learns the hierarchical similarity by only using the instance-wise contrastive loss in Eq. (19). (iii) *Prototype-Hyperbolic*: by setting  $\lambda_1 = 1$  and  $\lambda_2 = 0$ , it learns the hierarchical similarity by only using the prototype-wise contrastive loss in Eq. (18). To compare the representative power of hyperbolic space with Euclidean space, we implement the Euclidean counterparts of above three mentioned variants: (iv) *vanilla-Euclidean*, (v) *Instance-Euclidean*, and (vi) *Prototype-Euclidean*. For the Euclidean scenario, the quantization process and the distance metric refer to MeCoQ (Wang et al. 2022), and the continuous features ahead of the product quantization module are utilized to construct the hierarchical similarity.

As seen from Table 2, when employing the vanilla view-augmented contrastive objective, quantized representations in hyperbolic space may not have necessarily better quality than the counterpart in Euclidean space. Moreover, by introducing the extracted hierarchical similarity as the extra training signal, the retrieval performance of both embedding spaces can be significantly improved. Nevertheless, if only using the instance-wise contrastive loss to learn the hierarchical similarity between images, quantized representations in Euclidean space achieve greater performance gain

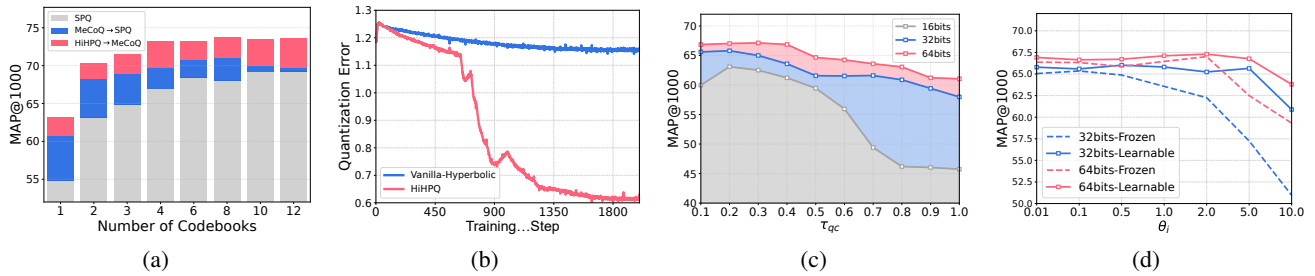


Figure 4: (a) MAP@1000 with varying numbers of codebooks on CIFAR10 (I); “A → B” denotes the performance gain achieved by Model A compared to Model B. (b) 32-bit quantization errors of HiHPQ and its variant on CIFAR10 (II). (c) Effects of the temperature  $\tau_{qc}$  on CIFAR-10 (II). (d) Effect of the initial curvature parameters on CIFAR10 (II).

Variants	16bits	32bits	64bits
Vanilla-Euclidean	59.26	62.52	62.72
Vanilla-Hyperbolic	58.40	61.86	62.17
Instance-Euclidean	61.56	64.70	65.40
Instance-Hyperbolic	61.80	63.84	63.98
Prototype-Euclidean	60.84	64.36	65.10
Prototype-Hyperbolic	62.65	65.54	66.85
HiHPQ (Full Model)	<b>63.31</b>	<b>65.81</b>	<b>67.11</b>

Table 2: MAP (%) comparison between HiHPQ and its variants on CIFAR10 (II).

than the hyperbolic scenario. It is worth noting that the enhancement in performance is particularly more significant when employing the prototype-wise contrastive loss on hyperbolic space to acquire hierarchical similarity. We conjecture this is because the prototype-wise objective directly incorporates prototypes as internodes in a tree-like hierarchy, thereby leveraging the advantages of hyperbolic geometry.

In Figure 4(b), we illustrate the effect of hierarchical semantics on the quantization error. Here, the quantization error refers to the hyperbolic distance between the continuous and quantized representation of an image. It is shown that including hierarchical information reduces quantization error significantly, while its absence slows error reduction.

**Impacts of Pre-defined Clustering Sizes** We investigate the impact of different pre-defined hierarchies on retrieval performance. As seen from Table 3, setting a cluster number for each hierarchy level that is significantly larger than the number of ground truth labels in the dataset (*e.g.*, 200 for CIFAR10 (I)) often has a more positive impact on performance. Moreover, setting the hierarchy level to 2 or 3 generally yields great performance improvements. In particular, we observe that prototype-wise contrastive loss in Eq. (18) always outperforms the instance-wise one in Eq. (19) in terms of performance gain.

**Hyper-parameter Sensitivity Analysis** We analyze the effect of two key hyper-parameters: the temperature  $\tau_{qc}$  in contrastive learning, and the curvature parameters  $\theta_i$ . As shown in Figure 4(c), a typically small temperature value

Dataset	Hierarchy Size	Prototype	Instance
CIFAR10 (I)	Vanilla-Hyperbolic → 66.87		
	200	72.12	71.54
	100	71.85	72.01
	200,100	72.97	72.49
	200,100,50	73.58	72.06
	100,50,25	71.72	70.39
Flickr25K	Vanilla-Hyperbolic → 80.63		
	50	81.35	81.21
	100,50	81.79	81.55
	100,30	81.99	82.33
	50,20	82.63	81.93
	100,50,20	81.91	81.71

Table 3: MAP (%) performance on CIFAR10 (I) at 32bits using different preset hierarchical structures. The “Prototype” and “Instance” columns show the performances of previously mentioned variants “Prototype-Hyperbolic” and “Instance-Hyperbolic”, respectively.

(*e.g.*, 0.2) has a positive impact on retrieval performance. Figure 4(d) illustrates the influence of the initial value of  $\theta_i$  on model performance, indicating that our method is robust within the range (0.01, 2.0) of initial values, whereas larger values lead to degradation. Notably, setting  $\theta_i$  as a learnable parameter can alleviate this decline in model performance.

## Conclusion

In this paper, we have proposed a novel unsupervised deep product quantization method, namely HiHPQ. In HiHPQ, we managed to build a hyperbolic product quantizer, in which both the soft hyperbolic codebook quantization mechanism and the quantized contrastive learning based on the hyperbolic product manifold were introduced to facilitate quantization. We further enhanced the semantics of quantized representations by utilizing the extracted hierarchical semantics supervision in our semantics learning module, which helped better distinguish between the query’s similar images and non-matching ones. Extensive experiments have shown the superiority of our model over existing baselines.

## Acknowledgements

The work described in this paper was substantially supported by InnoHK initiative, The Government of the HK-SAR, and Laboratory for AI-Powered Financial Technologies. In addition, the work was partially supported by a grant from the Research Grants Council of the Hong Kong Special Administrative Region, China (CUHK 14222922, RGC GRF No. 2151185).

## References

- Babenko, A.; and Lempitsky, V. 2014. The inverted multi-index. *TPAMI*, 37(6): 1247–1260.
- Baranchuk, D.; Babenko, A.; and Malkov, Y. 2018. Revisiting the inverted indices for billion-scale approximate nearest neighbors. In *ECCV*, 202–216.
- Chen, J.; Cheung, W. K.; and Wang, A. 2018. Learning Deep Unsupervised Binary Codes for Image Retrieval. In *IJCAI*, 613–619.
- Chen, T.; Li, L.; and Sun, Y. 2020. Differentiable product quantization for end-to-end embedding compression. In *ICML*, 1617–1626.
- Chen, Y.; Fang, Y.; Zhang, Y.; and King, I. 2023. Bipartite Graph Convolutional Hashing for Effective and Efficient Top-N Search in Hamming Space. In *WWW*, 3164–3172.
- Chen, Y.; Guo, H.; Zhang, Y.; Ma, C.; Tang, R.; Li, J.; and King, I. 2022a. Learning binarized graph representations with multi-faceted quantization reinforcement for top-k recommendation. In *SIGKDD*, 168–178.
- Chen, Y.; Yang, M.; Zhang, Y.; Zhao, M.; Meng, Z.; Hao, J.; and King, I. 2022b. Modeling scale-free graphs with hyperbolic geometry for knowledge-aware recommendation. In *WSDM*, 94–102.
- Chen, Y.; Zhang, Y.; Guo, H.; Tang, R.; and King, I. 2022c. An Effective Post-training Embedding Binarization Approach for Fast Online Top-K Passage Matching. In *AAACL*, 102–108.
- Chua, T.-S.; Tang, J.; Hong, R.; Li, H.; Luo, Z.; and Zheng, Y. 2009. Nus-wide: a real-world web image database from national university of singapore. In *CIVR*, 1–9.
- Dai, B.; Guo, R.; Kumar, S.; He, N.; and Song, L. 2017. Stochastic generative hashing. In *ICML*, 913–922.
- Desai, K.; Nickel, M.; Rajpurohit, T.; Johnson, J.; and Vedantam, R. 2023. Hyperbolic Image-Text Representations. *arXiv preprint arXiv:2304.09172*.
- Dizaji, K. G.; Zheng, F.; Sadoughi, N.; Yang, Y.; Deng, C.; and Huang, H. 2018. Unsupervised deep generative adversarial hashing network. In *CVPR*, 3664–3673.
- Gao, Z.; Wu, Y.; Harandi, M.; and Jia, Y. 2022. Curvature-adaptive meta-learning for fast adaptation to manifold data. *TPAMI*, 45(2): 1545–1562.
- Ge, S.; Mishra, S.; Kornblith, S.; Li, C.-L.; and Jacobs, D. 2023. Hyperbolic contrastive learning for visual representations beyond objects. In *CVPR*, 6840–6849.
- Ge, T.; He, K.; Ke, Q.; and Sun, J. 2013. Optimized product quantization for approximate nearest neighbor search. In *CVPR*, 2946–2953.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. In *NeurIPS*.
- Gray, R. 1984. Vector quantization. *IEEE Assp Magazine*, 1(2): 4–29.
- Gu, A.; Sala, F.; Gunel, B.; and Ré, C. 2018. Learning mixed-curvature representations in product spaces. In *ICLR*.
- Guo, R.; Sun, P.; Lindgren, E.; Geng, Q.; Simcha, D.; Chern, F.; and Kumar, S. 2020. Accelerating Large-Scale Inference with Anisotropic Vector Quantization. In *ICML*, 3887–3896.
- Guo, Y.; Xu, M.; Li, J.; Ni, B.; Zhu, X.; Sun, Z.; and Xu, Y. 2022. Hcsc: Hierarchical contrastive selective coding. In *CVPR*, 9706–9715.
- Huiskes, M. J.; and Lew, M. S. 2008. The mir flickr retrieval evaluation. In *MIR*, 39–43.
- Jang, Y. K.; and Cho, N. I. 2021. Self-supervised product quantization for deep unsupervised image retrieval. In *ICCV*, 12085–12094.
- Jégou, H.; Douze, M.; and Schmid, C. 2010. Product quantization for nearest neighbor search. *TPAMI*, 33(1): 117–128.
- Jégou, H.; Tavenard, R.; Douze, M.; and Amsaleg, L. 2011. Searching in one billion vectors: re-rank with source coding. In *ICASSP*, 861–864.
- Kalantidis, Y.; and Avrithis, Y. 2014. Locally optimized product quantization for approximate nearest neighbor search. In *CVPR*, 2321–2328.
- Kingma, D. P.; and Welling, M. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Klein, B.; and Wolf, L. 2019. End-to-end supervised product quantization for image search and retrieval. In *CVPR*, 5041–5050.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning multiple layers of features from tiny images.
- Law, M.; Liao, R.; Snell, J.; and Zemel, R. 2019. Lorentzian distance learning for hyperbolic representations. In *ICML*, 3672–3681.
- Li, J.; Li, Z.; Mou, L.; Jiang, X.; Lyu, M.; and King, I. 2020a. Unsupervised text generation by learning from search. In *NeurIPS*, 10820–10831.
- Li, J.; Zhou, P.; Xiong, C.; and Hoi, S. C. 2020b. Prototypical contrastive learning of unsupervised representations. *arXiv preprint arXiv:2005.04966*.
- Li, S.; Chen, Z.; Lu, J.; Li, X.; and Zhou, J. 2019. Neighborhood preserving hashing for scalable video retrieval. In *ICCV*, 8212–8221.
- Li, S.; Li, X.; Lu, J.; and Zhou, J. 2021. Structure-adaptive neighborhood preserving hashing for scalable video search. *IEEE Transactions on Circuits and Systems for Video Technology*, 32(4): 2441–2454.
- Li, Y.; and van Gemert, J. 2021. Deep unsupervised image hashing by maximizing bit entropy. In *AAAI*, 2002–2010.
- Lin, K.; Lu, J.; Chen, C.-S.; and Zhou, J. 2016. Learning compact binary descriptors with unsupervised deep neural networks. In *CVPR*, 1183–1192.

- Lin, Q.; Chen, X.; Zhang, Q.; Cai, S.; Zhao, W.; and Wang, H. 2022. Deep unsupervised hashing with latent semantic components. In *AAAI*, volume 36, 7488–7496.
- Liu, B.; Cao, Y.; Long, M.; Wang, J.; and Wang, J. 2018. Deep triplet quantization. In *MM*, 755–763.
- Liu, J.; Fournier-Viger, P.; Zhou, M.; He, G.; and Nouioua, M. 2022a. CSPM: Discovering compressing stars in attributed graphs. *Information Sciences*, 611: 126–158.
- Liu, J.; Yang, M.; Zhou, M.; Feng, S.; and Fournier-Viger, P. 2022b. Enhancing hyperbolic graph embeddings via contrastive learning. *arXiv preprint arXiv:2201.08554*.
- Liu, J.; Zhou, M.; Fournier-Viger, P.; Yang, M.; Pan, L.; and Nouioua, M. 2022c. Discovering Representative Attribute-stars via Minimum Description Length. In *ICDE*, 68–80.
- Lu, Z.; Lian, D.; Zhang, J.; Zhang, Z.; Feng, C.; Wang, H.; and Chen, E. 2023. Differentiable Optimized Product Quantization and Beyond. In *WWW*, 3353–3363.
- Luo, X.; Wu, D.; Ma, Z.; Chen, C.; Deng, M.; Ma, J.; Jin, Z.; Huang, J.; and Hua, X.-S. 2020. Cimon: Towards high-quality hash codes. *arXiv preprint arXiv:2010.07804*.
- Ma, Y.; Song, Z.; Hu, X.; Li, J.; Zhang, Y.; and King, I. 2023. Graph component contrastive learning for concept relatedness estimation. In *AAAI*, volume 37, 13362–13370.
- Ma, Z.; Luo, X.; Chen, Y.; Hou, M.; Li, J.; Deng, M.; and Lu, G. 2022. Improved deep unsupervised hashing with fine-grained semantic similarity mining for multi-label image retrieval. In *IJCAI*, 1254–1260.
- Morozov, S.; and Babenko, A. 2019. Unsupervised neural quantization for compressed-domain similarity search. In *ICCV*, 3036–3045.
- Nickel, M.; and Kiela, D. 2018. Learning continuous hierarchies in the lorentz model of hyperbolic geometry. In *ICML*, 3779–3788.
- Norouzi, M.; and Fleet, D. J. 2013. Cartesian k-means. In *CVPR*, 3017–3024.
- Qiu, Z.; Su, Q.; Ou, Z.; Yu, J.; and Chen, C. 2021. Unsupervised Hashing with Contrastive Information Bottleneck. In *IJCAI*, 959–965.
- Ratcliffe, J. G. 2006. Hyperbolic Geometry. *Foundations of Hyperbolic Manifolds*, 54–99.
- Salakhutdinov, R.; and Hinton, G. 2009. Semantic hashing. *IJAR*, 50(7): 969–978.
- Shen, Y.; Liu, L.; and Shao, L. 2019. Unsupervised binary representation learning with deep variational networks. *IJCV*, 127(11-12): 1614–1628.
- Shen, Y.; Qin, J.; Chen, J.; Yu, M.; Liu, L.; Zhu, F.; Shen, F.; and Shao, L. 2020. Auto-encoding twin-bottleneck hashing. In *CVPR*, 2818–2827.
- Song, J.; He, T.; Gao, L.; Xu, X.; Hanjalic, A.; and Shen, H. T. 2018. Binary generative adversarial networks for image retrieval. In *AAAI*.
- Su, S.; Zhang, C.; Han, K.; and Tian, Y. 2018. Greedy hash: Towards fast optimization for accurate hash coding in cnn. In *NeurIPS*.
- Tifrea, A.; Bécigneul, G.; and Ganea, O.-E. 2018. Poincaré glove: Hyperbolic word embeddings. *arXiv preprint arXiv:1810.06546*.
- Tu, R.-C.; Mao, X.; and Wei, W. 2020. MLS3RDUH: Deep Unsupervised Hashing via Manifold based Local Semantic Similarity Structure Reconstructing. In *IJCAI*, 3466–3472.
- Wang, J.; Zeng, Z.; Chen, B.; Dai, T.; and Xia, S.-T. 2022. Contrastive quantization with code memory for unsupervised image retrieval. In *AAAI*, volume 36, 2468–2476.
- Wei, R.; Liu, Y.; Song, J.; Xie, Y.; and Zhou, K. 2022. Hyperbolic Hierarchical Contrastive Hashing. *arXiv preprint arXiv:2212.08904*.
- Xiao, S.; Liu, Z.; Shao, Y.; Lian, D.; and Xie, X. 2021. Matching-oriented product quantization for ad-hoc retrieval. *arXiv preprint arXiv:2104.07858*.
- Yan, J.; Luo, L.; Deng, C.; and Huang, H. 2021. Unsupervised hyperbolic metric learning. In *CVPR*, 12465–12474.
- Yang, E.; Deng, C.; Liu, T.; Liu, W.; and Tao, D. 2018. Semantic structure-based unsupervised deep hashing. In *IJCAI*, 1064–1070.
- Yang, E.; Liu, T.; Deng, C.; Liu, W.; and Tao, D. 2019. Distillhash: Unsupervised deep hashing by distilling data pairs. In *CVPR*, 2946–2955.
- Yu, T.; Yuan, J.; Fang, C.; and Jin, H. 2018. Product quantization network for fast image retrieval. In *ECCV*, 186–201.
- Yue, C.; Long, M.; Wang, J.; Han, Z.; and Wen, Q. 2016. Deep quantization network for efficient image retrieval. In *AAAI*, 3457–3463.
- Zhang, S.; Xu, R.; Xiong, C.; and Ramaiah, C. 2022. Use all the labels: A hierarchical multi-label contrastive learning framework. In *CVPR*, 16660–16669.
- Zhang, Y.; and Zhu, H. 2019. Doc2hash: Learning discrete latent variables for documents retrieval. In *NAACL*.
- Zhang, Y.; and Zhu, H. 2020. Discrete wasserstein autoencoders for document retrieval. In *ICASSP*, 8159–8163.
- Zieba, M.; Semberecki, P.; El-Gaaly, T.; and Trzcinski, T. 2018. Bingan: Learning compact binary descriptors with a regularized gan. In *NeurIPS*.