

Long-Tail Cross Modal Hashing

Zijun Gao^{1,2}, Jun Wang^{2,*}, Guoxian Yu^{1,2}, Zhongmin Yan^{1,2}, Carlotta Domeniconi³, Jinglin Zhang⁴

¹School of Software, Shandong University, Jinan, China

²SDU-NTU Joint Centre for AI Research, Shandong University, Jinan, China

³Department of Computer Science, George Mason University, Fairfax, VA, USA

⁴School of Control Science and Engineering, Shandong University, Jinan, China

zjgao@mail.sdu.edu.cn, {kingjun, gxyu, yzm}@sdu.edu.cn, carlotta@cs.gmu.edu, jinglin.zhang@sdu.edu.cn

Abstract

Existing Cross Modal Hashing (CMH) methods are mainly designed for balanced data, while imbalanced data with long-tail distribution is more general in real-world. Several long-tail hashing methods have been proposed but they can not adapt for multi-modal data, due to the complex interplay between labels and individuality and commonality information of multi-modal data. Furthermore, CMH methods mostly mine the commonality of multi-modal data to learn hash codes, which may override tail labels encoded by the individuality of respective modalities. In this paper, we propose LtCMH (Long-tail CMH) to handle imbalanced multi-modal data. LtCMH firstly adopts auto-encoders to mine the individuality and commonality of different modalities by minimizing the dependency between the individuality of respective modalities and by enhancing the commonality of these modalities. Then it dynamically combines the individuality and commonality with direct features extracted from respective modalities to create meta features that enrich the representation of tail labels, and binaries meta features to generate hash codes. LtCMH significantly outperforms state-of-the-art baselines on long-tail datasets and holds a better (or comparable) performance on datasets with balanced labels.

Introduction

Hashing aims to map high-dimensional data into a series of low-dimensional binary codes while preserving the data proximity in the original space. The binary codes can be computed in a constant time and economically stored, which meets the need of large-scale data retrieval. In real-world applications, we often want to retrieval data from multiple modalities. For example, when we input key words to information retrieval systems, we expect the systems can efficiently find out related news/images/videos from database. Hence, many cross-modal hashing (CMH) methods have been proposed to deal with such tasks (Wang et al. 2016; Jiang and Li 2017; Yu et al. 2022a; Liu et al. 2019a).

Most CMH methods aim to find a low-dimensional shared subspace to eliminate the modality heterogeneity and to quantify the similarity between samples across modalities. More advanced methods leverage extra knowledge (i.e., labels, manifold structure, neighbors coherence) (Jiang and Li

2017; Yu et al. 2021; Liu et al. 2019b) to more faithfully capture the proximity between multi-modality data to induce hash codes. However, almost all of them are trained and tested on hand-crafted balanced datasets, while *imbalanced* datasets with long tail labels are more prevalence in real-world. Some recent studies (Chen et al. 2021; Cui et al. 2019) have reported that the imbalanced nature of real-world data greatly compromises the retrieval performance.

Real-world samples typically have a skewed distribution with **long-tail labels**, which means that a few labels (a.k.a. *head labels*) annotate to many samples, while the other labels (a.k.a. *tail labels*) have a large quantity but each of them is only annotated to several samples (Liu et al. 2020; Chen et al. 2021). It is a challenging task to train a general model from such distribution, because the head labels gain most of the attention while the tail ones are underestimated. Another inevitable problem of long-tail hashing is the ambiguity of the generated binary codes. Due to the dimension reduction, information loss is unavoidable. In this case, hash codes learned from data-poor tail labels more lack the discrimination, which seriously confuses the retrieval results.

Several efforts have been made toward long-tail single-modal hashing via knowledge transfer (Liu et al. 2020) and information augmentation (Chu et al. 2020; Wang et al. 2020b; Kou et al. 2022). LEAP (Liu et al. 2020) augments each instance of tail labels with certain disturbances in the deep representation space to provide higher intra-class variation to tail labels. OLTR (Liu et al. 2019c) and LTHNet (Chen et al. 2021) propose a meta embedding module to transfer knowledge from head labels to tail ones. In addition, these long-tail hashing methods cannot be directly adapted for multi-modal data, due to the complex interplay between tail labels and heterogeneous data modalities. The head labels can be easily represented by the commonality of multiple modalities, but the tail labels are often represented by individuality of a particular modality. For example, in the left part of Figure 1, tail label ‘Hawaii’ can only be read out from the text-modality. On the other hand, head label ‘Sea’ can be consolidated from the commonality of image and text modality, and head label ‘Island’ comes from the image modality alone. From this example, we can conclude that both the **individuality** and **commonality** should be used for effective CMH on long-tail data, which also account for the complex interplay between head/tail labels and multi-

*Corresponding author.

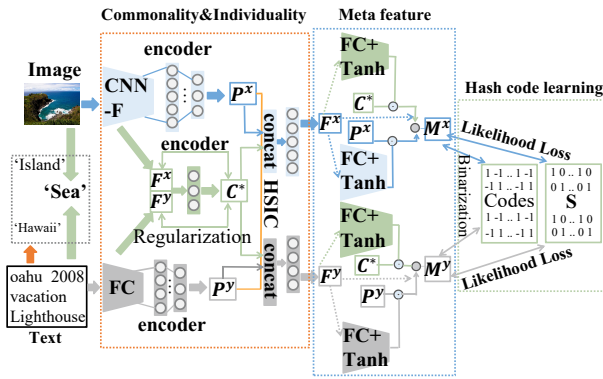


Figure 1: The schematic framework of LtCMH. The direct image/text features F^x/F^y are first extracted by CNN-F and FC (fully-connected) networks with two layers. LtCMH then uses different auto-encoders to mine the commonality C^* and individuality P^x/P^y of multi-modal data via cross-modal regularization and Hilbert-Schmidt independence criterion (HSIC). After that, it creates meta features M^x/M^y by fusing C^* , P^x/P^y and F^x/F^y . Next, LtCMH binarizes meta features into hash codes. The head label ‘Sea’ can be consolidated from the commonality of image and text modality, while tail label ‘Hawaii’ can only be obtained from the text-modality. The complex relations between labels and heterogeneous modalities can be more well explored by the commonality and individuality.

modal data. Unfortunately, most CMH methods solely mine the commonality (shared subspace) to learn hashing functions and assume balanced multi-modal data (Wang et al. 2016; Jiang and Li 2017; Yu et al. 2022a).

To address these problems, we propose LtCMH (Long-tail CMH) to achieve CMH on imbalanced multi-modal data, as outlined in Figure 1. Specifically, we adopt different auto-encoders to mine individuality and commonality information from multi-modal data in a collaborative way. We further use the Hilbert-Schmidt independence criterion (HSIC) (Gretton et al. 2005) to extract and enhance the individuality of each modality, and cross-modal regularization to boost the commonality. We model the interplay between head/tail labels and multi-modal data by meta features dynamically fused from the commonality, individuality and direct features extracted from respective modalities. The meta features can enrich tail labels and preserve the correlations between different modalities for more discriminate hash codes. We finally binarize the meta features to generate hash codes. The contributions of this work are summarized as follows:

- (i) We study CMH on long-tail multi-modal data, which is a novel, practical and difficult but understudied topic. We propose LtCMH to achieve effective hashing on both long-tail and balanced datasets.
- (ii) LtCMH can mine the individuality and commonality information of multi-modal data to more comprehensively model the complex interplay between head/tail labels and heterogeneous data modalities. It further defines a dynamic

meta feature learning module to enrich labels and to induce discriminate hash codes.

(iii) Experimental results show the superior robustness and performance of LtCMH to competitive CMH methods (Wang et al. 2020a; Li et al. 2018; Yu et al. 2021) on benchmark datasets, especially for long-tail ones.

Related Work

Cross Modal Hashing

Based on using the semantic labels or not, existing CMH can be divided into two types: unsupervised and supervised. Unsupervised CMH methods usually learn hash codes from data distributions without referring to supervised information. Early methods, such as cross view hashing (CVH) (Kumar and Udupa 2011) and collective matrix factorization hashing (CMFH) (Ding, Guo, and Zhou 2014), typically build on canonical correlation analysis (CCA). UD-CMH (Wu et al. 2018) keeps the maximum structural similarity between the deep hash codes and original data. DJSRH (Su, Zhong, and Zhang 2019) learns a joint semantic affinity matrix to reconstruct the semantic affinity between deep features and hash codes. Supervised CMH methods additionally leverage supervised information such as semantic labels and often gain a better performance than unsupervised ones. For example, DCMH (Jiang and Li 2017) optimizes the joint loss function to maintain the label similarity between unified hash codes and feature representations of each modality. BiNCMH (Sun et al. 2022) uses bi-direction relation reasoning to mine unrelated semantics to enhance the similarity between features and hash codes.

Although these CMH methods perform well on balanced datasets, they are quite fragile on long-tail datasets, in which samples annotated with head labels gain more attention than those with tail labels when jointly optimizing correlation between modalities (Zhou et al. 2020). Given that, we try to enrich the representation of tail labels by mining the individuality and commonality of multi-modal data. This enriched representation can more well model the relation between head/tail labels and multi-modal data.

Long-Tail Data Learning

Long-tail problem is prevalence in real-world data mining tasks and thus has been extensively studied (Chawla 2009). There are three typical solutions: re-sampling, re-weighting and transfer-learning. Re-sampling based solutions use different sampling strategies to re-balance the distribution of different labels (Buda, Maki, and Mazurowski 2018; Byrd and Lipton 2019). Re-weighting solutions give different weights to the head and tail labels of the loss function (Huang et al. 2016). Transfer-learning based techniques learn general knowledge from the head labels, and then use these knowledge to enrich tail labels. Liu et al. (2019c) proposed a meta feature embedding module to combine direct features with memory features to enrich the representation of tail labels. Liu et al. (2020) proposed the concept of ‘feature cloud’ to enhance the diversity of tail ones.

Most of these long-tail solutions manually divide samples into head labels and tail ones, which compromises the gen-

eralization ability. In contrast, our LtCMH does not need to divide head and tail labels beforehand, it leverages both individuality and commonality as memory features, rather than directly stacking direct features (Liu et al. 2019c; Chen et al. 2021), to reduce information overlap. As a result, LtCMH has a stronger generalization and feature representation capability.

The Proposed Methodology

Problem Overview

Without loss of generality, we firstly present LtCMH based on two modalities (image and text). LtCMH can also be applied to other data modality or extended to ≥ 3 modalities (Liu et al. 2022). Suppose $\mathcal{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n\}$ and $\mathcal{Y} = \{\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_n\}$ are the respective image and text modality with n samples, and $\mathbf{L} = \{\mathbf{l}_1, \mathbf{l}_2, \dots, \mathbf{l}_n\}$ is the label matrix of these samples from c distinct labels. A dataset is called long-tail if the numbers of samples annotated with these c labels conform to Zipf’s law distribution (Reed 2001): $z_a = z_1 \times a^{-\mu}$, where z_a is the number of samples with label a in decreasing order ($z_1 > z_2 > \dots \gg z_c$) and μ is the control parameter of imbalance factor (**IF** for short, $\text{IF} = z_1/z_c$). The goal of CMH is to learn hash functions (h^x and h^y) from \mathcal{X} and \mathcal{Y} , and to generate hash codes via $\mathbf{b}^x = h^x(\mathbf{x})$ and $\mathbf{b}^y = h^y(\mathbf{y})$, where $\mathbf{b}^x/\mathbf{b}^y \in \{0, 1\}^k$ is the length- k binary hash code vector. For ease presentation, we denote v as any modality, x/y as the image/text modality indicator.

The whole framework of LtCMH is illustrated in Figure 1. LtCMH firstly captures the individuality and commonality of multi-modal data, and takes the individuality and commonality as the memory features. Then it introduces individuality-commonality selectors on memory features, along with the direct features extracted from respective modality, to create meta features, which not only enrich the representation of tail labels, but also automatically balance the head and tail labels. Finally, it binaries meta features into binary codes through the hash learning module. The following subsections elaborate on these steps.

Individuality and Commonality Learning

Before learning the individuality and commonality of multi-modal data, we adopt Convolutional Neural Networks (CNN) to extract direct visual features $\mathbf{F}^x = \text{CNN}(\mathcal{X}) \in \mathbb{R}^{n \times d_x}$ and fully connected networks with two layers (FC) to extract textual features $\mathbf{F}^y = \text{FC}(\mathcal{Y}) \in \mathbb{R}^{n \times d_y}$. Other feature extractors can also be adopted for modality adaption. Due to the prevalence of samples of head labels, these direct features are more biased toward head labels and under-represent tail ones. To address this issue, long-tail single modal hashing solutions learn label prototypes to summarize visual features and utilize them to transfer knowledge from head to tail labels (Wei et al. 2022; Tang, Huang, and Zhang 2020; Chen et al. 2021). But they can not account for the complex interplay between head/tail labels and individuality and commonality of multi-modal data, as we exemplified in Figure 1 and discussed in the Introduction. Some

attempts have already explored the individuality and commonality of multi-view data to improve the performance of multi-view learning (Wu et al. 2019; Yu et al. 2022b), and Tan et al. (2021) empirically found the classification of less frequent labels can be improved by the individuality. Inspired by these works, we advocate to mine the individuality and commonality of multi-modal data with long-tail distribution. We then leverage the individuality and commonality to create meta features, which can enrich the representation of tail labels and model the interplay between labels and multi-modal data .

Because of the information overlap of multi-modal high-dimensional data, we want to mine the shared and specific essential information of different modalities in the low-dimensional embedding space. Auto-encoders (AE) (Goodfellow, Bengio, and Courville 2016) is an effective technique that can map high-dimensional data into an informative low-dimensional representation space. AE can encode unlabeled/incomplete data and disregard the non-significant and noisy part, so it has been widely used to extract disentangled latent representation. Here, we propose a new structure for learning the individuality and commonality of multi-modal data. The proposed individuality-commonality AE (as shown in Figure 1) is similar to a traditional AE in both the encoding end of the input and the decoding end of the output. The new ingredients of our AE are the regularization of learning individuality and commonality, and the decoding of each modality via combining the individuality of this modality and commonality of multiple modalities.

As for the encoding part, we use view-specific encoders $\{f_{enI}^v\}$ to initialize the individuality information matrix of each modality as follows:

$$\mathbf{P}^x = f_{enI}^x(\mathbf{F}^x), \mathbf{P}^y = f_{enI}^y(\mathbf{F}^y) \quad (1)$$

where \mathbf{P}^x and \mathbf{P}^y encode the individuality of image and text modality, respectively. To learn the commonality of multi-modal data, we concatenate \mathbf{F}^x and \mathbf{F}^y and then input them into commonality encoder f_{enC} as:

$$\mathbf{C}^* = f_{enC}([\mathbf{F}^x, \mathbf{F}^y]) \quad (2)$$

Although \mathbf{C}^* can encode the commonality of multi-modal data by fusing \mathbf{F}^x and \mathbf{F}^y , it does not concretely consider the intrinsic distribution of respective modalities. Hence, we further define a cross-modal regularization to optimize \mathbf{C}^* and to enhance the commonality using the shared labels of samples and data distribution within each modality as:

$$\begin{aligned} \mathcal{J}_1(\mathbf{C}^*) &= \frac{1}{2} \left(\sum_{a,b=1}^c \|\mathbf{C}_{\cdot a}^* - \mathbf{C}_{\cdot b}^*\|^2 (\mathbf{R}_{ab}^x + \mathbf{R}_{ab}^y) \right) \\ &= \text{tr}(\mathbf{C}^* ((\mathbf{D}^x - \mathbf{R}^x) + (\mathbf{D}^y - \mathbf{R}^y)) (\mathbf{C}^*)^T) \end{aligned} \quad (3)$$

Here $\mathbf{C}_{\cdot a}^*$ is the a -th column of \mathbf{C}^* , $\text{tr}(\cdot)$ denotes the matrix trace operator, \mathbf{D}^v is a diagonal matrix with $\mathbf{D}_{aa}^v = \sum_{b=1}^c \mathbf{R}_{ab}^v$. \mathbf{R}^v quantifies the similarity between different

labels, it is defined as follows:

$$\begin{aligned} \mathbf{R}_{ab}^v &= e^{\frac{-\mathbf{H}_{ab}^v}{(\sigma^v)^2}} \\ \mathbf{H}_{ab}^v &= \frac{\sum_{\mathbf{f}_i^v \in \mathcal{X}_a^v} \min_{\mathbf{f}_j^v \in \mathcal{X}_b^v} d(\mathbf{f}_i^v, \mathbf{f}_j^v)}{|\mathcal{X}_a^v| + |\mathcal{X}_b^v|} \\ &+ \frac{\sum_{\mathbf{f}_j^v \in \mathcal{X}_b^v} \min_{\mathbf{f}_i^v \in \mathcal{X}_a^v} d(\mathbf{f}_i^v, \mathbf{f}_j^v)}{|\mathcal{X}_a^v| + |\mathcal{X}_b^v|} \end{aligned} \quad (4)$$

where \mathbf{H}_{ab}^v is the average Hausdorff distance between two sets of samples separately annotated with label a and b , $d(\mathbf{f}_i^v, \mathbf{f}_j^v)$ is the Euclidean distance between \mathbf{f}_i^v and \mathbf{f}_j^v . $|\mathcal{X}_a^v|$ counts the number of samples annotated with a . σ^v is set to the average of \mathbf{H}^v .

Eq. (3) aims to learn the commonality of each label across modalities, it jointly considers the intrinsic distribution of samples annotated with a particular label within each modality, thus \mathbf{C}^* can bridge \mathcal{X} and \mathcal{Y} and enable CMH. We want to remark that other distance metrics can also be used to setup \mathbf{H}^v . Our choice of Hausdorff distance is for its intuitiveness and effectiveness on qualifying two sets of samples (Hausdorff 2005). For the variants of Hausdorff distance, we choose the average Hausdorff distance because it considers more geometric relations between samples of two sets than the maximum and minimum Hausdorff distances (Zhou et al. 2012). By optimizing Eq. (3) across modalities, we can enhance the quality of extracted commonality of multi-modal data.

\mathbf{P}^x and \mathbf{P}^y may be highly correlated, because they are simply obtained from the individuality auto-encoders without any contrast and collaboration, and samples from different modalities share the same labels. As such, the individuality of each modality cannot be well preserved and tail labels encoded by the individuality are under-represented. Here, we minimize the correlation between \mathbf{P}^x and \mathbf{P}^y to capture the intrinsic individuality of each modality. Meanwhile, \mathbf{C}^* can include more shared information of multi-modal data.

For this purpose, we use HSIC (Gretton et al. 2005) to approximately quantify the correlation between \mathbf{P}^x and \mathbf{P}^y , for its simplicity and effectiveness on measuring linear and nonlinear interdependence/correlation between two sets of variables. The correlation between \mathbf{P}^x and \mathbf{P}^y is approximated as:

$$\begin{aligned} \mathcal{J}_2(\mathbf{P}^v) &= HSIC(\mathbf{P}^x, \mathbf{P}^y) = (n-1)^{-2} tr(\mathbf{K}^x \mathbf{A} \mathbf{K}^y \mathbf{A}) \\ s.t. \mathbf{K}_{ab}^v &= \kappa(\mathbf{P}_{a.}^v, \mathbf{P}_{b.}^v) = e^{\left(\frac{-\|\mathbf{P}_{a.}^v - \mathbf{P}_{b.}^v\|^2}{\sigma^v}\right)} \end{aligned} \quad (5)$$

where \mathbf{K}^x and \mathbf{K}^y are the kernel-induced similarity matrix from \mathbf{P}^x and \mathbf{P}^y , respectively. \mathbf{A} is a centering matrix: $\mathbf{A} = \mathbf{I} - \mathbf{e}\mathbf{e}^T/n$, where $\mathbf{e} = (1, \dots, 1)^T \in \mathbb{R}^n$ and \mathbf{I} is the identity matrix.

For the decoding part, we use the extracted commonality \mathbf{C}^* shared across modalities and the individuality of each modality to reconstruct this original modality as follows:

$$\mathcal{J}_3(\mathbf{C}^*, \mathbf{P}^v) = \sum_{v=x,y} \frac{\|\mathbf{F}^v - f_{de}^v([\mathbf{C}^*, \mathbf{P}^v])\|_F^2}{nd_v} \quad (6)$$

where $\{f_{de}^v\}$ is the corresponding decoder of each modality.

Then we can define the loss function of the individuality-commonality AE as:

$$\min_{\theta_z} Loss1 = \alpha \mathcal{J}_1(\mathbf{C}^*) + \beta \mathcal{J}_2(\mathbf{P}^v) + \mathcal{J}_3(\mathbf{C}^*, \mathbf{P}^v) \quad (7)$$

θ_z are the parameters of the individuality-commonality AE. α and β ($\in (0, 1]$) are the parameters to control the weight of individuality and commonality. To this end, LtCMH can capture the commonality and individuality of different modalities, which will be used to create dynamic meta features in the next subsection.

Dynamic Meta Features Learning

For long-tail datasets, the head labels have abundant representations while the tail labels don't (Cui et al. 2019). As a result, head labels can be easily distinguished but tail labels can not. To enrich label representation and transfer knowledge of head labels to tail ones, we propose the dynamic meta memory embedding module using direct features of \mathbf{F}^x and \mathbf{F}^y , and the commonality \mathbf{C}^* and individuality \mathbf{P}^v .

The modality heterogeneity is a typical issue in cross modal information fusion, which means that different modalities may have completely different feature representations and statistical characteristics for the same semantic labels. This makes it difficult to directly measure the similarity between different modalities. However, multi-modal data that describe the same instance usually have close semantic meanings. In the previous subsection, we have learnt a commonality information matrix \mathbf{C}^* by mining the geometric relations between samples of two labels in the embedding space. Therefore, \mathbf{C}^* can be seen as a bridge between different modalities. Meanwhile, different modalities often have their own individual information, so we learn \mathbf{P}^v to capture the individuality of each modality. We take \mathbf{C}^* and \mathbf{P}^v as the memory features, and obtain the meta features of a modality as follows:

$$\mathbf{M}^v = \mathbf{F}^v + \mathbf{E}_1 \odot \mathbf{C}^* + \mathbf{E}_2 \odot \mathbf{P}^v \quad (8)$$

Data-poor tail labels need more memory features to enrich their representations than data-rich head labels, we design two adaptive selection factors \mathbf{E}_1 and \mathbf{E}_2 on \mathbf{C}^* and \mathbf{P}^v for this purpose, which can be adaptive computed from \mathbf{F}^v to derive selection weights as $\mathbf{E}_1 = \text{Tanh}(FC(\mathbf{F}^v))$ and $\mathbf{E}_2 = \text{Tanh}(FC(\mathbf{F}^v))$. \odot is the Hadamard product. To match the general multi-modal data in a soft manner, we adopt the lightweight 'Tanh+FC' to avoid complex and inefficient parameters adjusting. Different from long-tail hashing methods (Kou et al. 2022) that use prototype networks (Snell, Swersky, and Zemel 2017) or simply stack the direct features to learn meta features for each label, our \mathbf{M}^v is created from both individuality and commonality, and direct features of multi-modal data. In addition, the dynamic factors can integrate both the enhanced commonality across modalities and individuality of each modality for cross-modal hashing, thus \mathbf{M}^v has a good interpretability and enables a better performance without the clear division of head and tail labels.

Hash Code Learning

After obtaining the dynamic meta features, the information across modalities is reserved and each sample’s representation is enriched. Alike DCMH (Jiang and Li 2017), we use $\mathbf{S} \in \mathbb{R}^{n \times n}$ to store the similarity of n training samples. $S_{ij} = 1$ means that \mathbf{x}_i and \mathbf{y}_j are with the same label, and $S_{ij} = 0$ otherwise. Based on \mathbf{M}^x , \mathbf{M}^y and \mathbf{S} , we can define the likelihood function as follows:

$$p(\mathbf{S}|\mathbf{M}^x, \mathbf{M}^y) = \begin{cases} \sigma(\phi^{xy}) & S_{ij} = 1 \\ 1 - \sigma(\phi^{xy}) & S_{ij} = 0 \end{cases} \quad (9)$$

where $\phi^{xy} = 1/2(\mathbf{M}^x)^T \mathbf{M}^y$ and $\sigma(\phi^{xy}) = \frac{1}{1+e^{-\phi^{xy}}}$. The smaller the angle between \mathbf{M}^x and \mathbf{M}^y is, the larger the ϕ^{xy} is, which makes it a higher probability that $S_{ij} = 1$ and vice versa. Then we can define the loss function for learning hash codes as follows:

$$\begin{aligned} \min_{\mathbf{B}, \theta_x, \theta_y} Loss2 = & - \sum_{i,j=1}^n (S_{ij} \phi_{ij}^{xy} - \log(1 + e^{\phi_{ij}^{xy}})) \\ & + \gamma(\|\mathbf{B} - \mathbf{M}^x\|_F^2 + \|\mathbf{B} - \mathbf{M}^y\|_F^2) + \eta(\|\mathbf{M}^x \mathbf{1}\|_F^2 + \|\mathbf{M}^y \mathbf{1}\|_F^2) \\ & s.t. \quad \mathbf{B} \in \{+1, -1\}^{c \times n} \end{aligned} \quad (10)$$

ϕ_{ij}^{xy} means the i -th row and j -th column of ϕ^{xy} . \mathbf{B} is the unified hash codes and $\mathbf{1}$ is a vector with all elements being 1. θ_x and θ_y are the parameters of image and text meta feature learning module. The first term is the negative log likelihood function, it aims at minimizing the inner product between \mathbf{M}^x and \mathbf{M}^y and preserving the semantic similarities of samples from different modalities. The Hamming distance between similar samples will be reduced and between dissimilar ones will be enlarged. The second term aims to minimize the difference between the unified hash codes and meta features of each modality. Due to the similarity preservation of \mathbf{M}^x and \mathbf{M}^y in \mathbf{S} , enforcing the unified hash codes closer to \mathbf{M}^x and \mathbf{M}^y is expected to preserve the cross modal similarity to match the goal of cross modal hashing. The last term pursues balanced binary codes with fixed length for a larger coding capacity.

Optimization

There are three parameters θ_x , θ_y , \mathbf{B} in our hash learning loss function in Eq. (10), it is difficult to simultaneously optimize them all and find the global optimum. Considering that, we adopt a canonically-used alternative strategy that optimizes one of them with others fixed, and give the optimization as below.

Optimize θ_x with fixed θ_y and \mathbf{B} : We first calculate the derivative of the loss function $Loss2$ with respect to \mathbf{M}^x for image modality and then we take the back-propagation (BP) and stochastic gradient decent (SGD) to update θ_x until convergence or the preset maximum epoch. $\frac{\partial Loss2}{\partial \mathbf{M}^x}$ can be calculated as:

$$\begin{aligned} \frac{\partial Loss2}{\partial \mathbf{M}_{*i}^x} = & \frac{1}{2} \sum_{j=1}^n (\sigma(\phi_{ij}) \mathbf{M}_{*j}^y - S_{ij} \mathbf{M}_{*j}^y) \\ & + 2\gamma(\mathbf{M}_{*i}^x - \mathbf{B}_{*i}) + 2\eta \mathbf{M}^x \mathbf{1} \end{aligned} \quad (11)$$

Dataset	N_{base}	N_{query}	$z_1(z_c)$	c
Flicker25K	18015	2000	3000(60)	24
NUS-WIDE	195834	2000	5000(100)	21

Table 1: Statistics of long-tail datasets. z_a is the number of samples annotated with label a , which conforms to Zipf’s law distribution; c is the number of distinct labels.

where \mathbf{M}_{*i}^x represents the i -th column of \mathbf{M}^x . By the chain rule, we update θ_x based on BP algorithm.

The way to optimize θ_y is similar as that to optimize θ_x .

Optimize \mathbf{B} with fixed θ_x and θ_y : When θ_x and θ_y are fixed, the optimization problem can be reformulated as:

$$\begin{aligned} \max_{\mathbf{B}} \quad & tr(\mathbf{B}(\mathbf{M}^x + \mathbf{M}^y)^T) \\ s.t. \quad & \mathbf{B} \in \{+1, -1\}^{c \times n} \end{aligned} \quad (12)$$

To maximize Eq. (12), \mathbf{B} should have the same sign as $(\mathbf{M}^x + \mathbf{M}^y)$ and \mathbf{B} can only take values 1 or -1. Opposite sign leads to the decreased value.

$$\mathbf{B} = sign(\mathbf{M}^x + \mathbf{M}^y) \quad (13)$$

We illustrate the whole framework of LtCMH in Figure 1, and defer its algorithmic procedure into the Supplementary file.

Experiments

Experimental Setup

There is no off-the-shelf benchmark long-tail multi-modal dataset for experiments, so we pre-process two hand-crafted multi-modal datasets (Flicker25K (Huiskes and Lew 2008) and NUS-WIDE (Chua et al. 2009)), to make them fit long-tail settings. We do not directly use the original NUS-WIDE dataset since it contains many meaningless samples, which does not well match the cross modal hashing task and long-tail setting. The statistics of pre-processed datasets are reported in Table 1, and more information of the pre-processings are given in the Supplementary file. We also take the public Flicker25K and NUS-WIDE as the balanced datasets for experiments.

Alike DCMH (Jiang and Li 2017), we use a pre-trained CNN named CNN-F with 8 layers, to extract the direct image features \mathbf{F}^x , and another network with two fully connected layers to extract direct text features \mathbf{F}^y . These two networks are with the same hyper-parameters as DCMH. Other networks can also be used here, which is not the main focus of this work. SGD is used to optimize model parameters. Learning rate of image and text feature extraction is set as $1e-1.5$, the learning rate of individuality-commonality AE is set as $1e-2$. Other hyper-parameters are set as: batch size=128, $\alpha=0.05$, $\beta=0.05$, $\gamma=1$, $\eta=1$, d_x and d_y is equal to hash code length k , the max epoch is 500. Parameter sensitivity is studied in the Supplementary file.

We compare LtCMH against with six representative and related CMH methods, which include CMFH (Ding, Guo, and Zhou 2014), JIMFH (Wang et al. 2020a), DCMH (Jiang and Li 2017), DGCPN (Yu et al. 2021), SSAH (Li et al.

2018), and MetaCMH (Wang et al. 2021). The first two are shallow solutions, and latter four are deep ones. They all focus on the commonality of multi-modal data to learn hash codes. CMFH, JIMFH and DGCPN are unsupervised solutions, while the others are supervised ones. We also take the recent long-tail single-modal hashing method LTHNet (Chen et al. 2021) as another baseline. For fair evaluation with LTHNet, we solely train and test LtCMH and LTHNet on the image-modality. The parameters of compared methods are fixed as reported in original papers or selected by a validation process. All experiments are independently repeated for ten times. The code of LtCMH is shared at www.sdu-idea.cn/codes.php?name=LtCMH. We implement LtCMH in Python 3.7 with the MindSpore deep learning framework.

Result Analysis

We adopt the typical mean average precision (MAP) as the evaluation metric, and report the average results and standard deviation in Table 2 (long-tailed), Table 3 (balanced) and Table 4 (single-modality). From these results, we have several important observations:

(i) LtCMH can effectively handle long-tail multi/single-modal data, this is supported by the clear performance gap between LtCMH and other compared methods. DCMH and SSAH have greatly compromised performance on long-tailed datasets, because they use semantic labels to guide hash codes learning, while tail labels do not have enough training samples to preserve the modality relationships in both the common semantic space and Hamming space. Although unsupervised CMH methods give hash codes without referring to the skewed label distributions, they are also misled by the overwhelmed samples of head labels. Another cause is that they all focus the commonality of multi-modal data, while LtCMH considers both the commonality and individuality. We note each compared method has a greatly improved performance on the balanced datasets, since they all target at balanced multi-modal data. Compared with results on long-tail datasets, the performance drop of LtCMH is the smallest, this proves its generality. LtCMH sometimes slightly loses to SSAH on the balanced datasets, that is because the adversarial network of SSAH can utilize more semantic information in some cases.

(ii) LtCMH can learn more effective meta features and achieve better knowledge transfer from head labels to tail labels than LTHNet and MetaCMH. The latter two stack direct features as memory features to transfer knowledge, they perform better than other compared methods (except LtCMH), but the stacked features may cause information overlap and they only enrich tail labels within each modality, and suffer the inconsistency caused by modality heterogeneity. LtCMH captures both individuality and commonality of multi-modal data, and the enhanced commonality helps to keep cross-modal consistency. We further separately measure the performance of LtCMH, MetaCMH and LTHNet on the head labels and tail ones, and report the results in Figure 2 and Supplementary file. We find that LtCMH gives better results than MetaCMH and LTHNet on both head and tail labels, which further suggest the effectiveness of our meta features.




Figure 2: Performance comparison of LtCMH and LTHNet on *head(-h)* and *tail(-t)* labels. For Flicker, the first 14 labels are head; and for NUS-WIDE, the first 15 labels are head.

(iii) The consideration of label information and data distribution improve the performance of CMH. We notice that most supervised methods perform better than unsupervised ones. Deep neural networks based solutions also perform better than shallow ones. Besides deep feature learning, SSAH and MetaCMH leverage both the label information and the data distribution, they obtain a better performance than other compared methods. LtCMH leverages these information sources in a more sensible way, and thus achieves a better performance on both long-tail and balanced datasets.

Ablation Experiments

To gain an in-depth understanding of LtCMH, we introduce five variants: LtCMH-w/oC, LtCMH-w/oI, LtCMH-w/oIC, LtCMH-w/oMI, LtCMH-w/oMT, which separately disregards the individuality, commonality and the both, dynamic meta features from the image modality and text modality. Figure 3 shows the results of these variants on Flicker25K and NUS-WIDE. We have several important observations. (i) Both the commonality and individuality are important for LtCMH on long-tail datasets. This is confirmed by clearly reduced performance of LtCMH-w/oC and LtCMH-w/oI, and by the lowest performance of LtCMH-w/oIC. We find LtCMH-w/oI performs better than LtCMH-w/oC, since commonality captures the shared and complementary information of multi-modal data, and bridges two modalities for cross-modal retrieval. (ii) Both the meta features from the text- and image-modality are helpful for hash code learning. We note a significant performance drop when meta features of target query modality are unused.

Besides, we also study the impact of input parameters α , β , γ and η , which control the learning loss of commonality, individuality, consistency of hash codes across modalities, balance of hash codes. We find that a too small α or β cannot ensure to learn commonality and individuality across modalities well, but a too large of them brings down the validity of AE. A too small γ and η cannot keep consistency of hash codes across modalities and generate balanced hash codes. Given these results, we set $\alpha=0.05$, $\beta=0.05$, $\gamma=1$, $\eta=1$.

Conclusion

In this paper, we study how to achieve CMH on the prevalence long-tail multi-modal data, which is a practical and important, but largely unexplored problem in CMH. We pro-

		Flicker25K			NUS-WIDE		
		16-bits	32-bits	64-bits	16-bits	32-bits	64bits
I→T	CMFH	.354±.030	.377±.005	.382±.003	.254±.010	.256±.007	.261±.020
	JIMFH	.400±.009	.414±.010	.433±.026	.308±.032	.323±.027	.347±.016
	DGCPN	.538±.008	.551±.013	.579±.005	.379±.006	.381±.007	.409±.004
	DCMH	.477±.020	.492±.003	.506±.017	.347±.014	.367±.003	.376±.010
	SSAH	.571±.004	.588±.005	.603±.009	.371±.007	.383±.006	.418±.004
	MetaCMH	.608±.016	.621±.009	.624±.025	.409±.016	.421±.004	.430±.007
	LtCMH	.687±.015	.732±.004	.718±.014	.433±.004	.475±.008	.532±.017
T→I	CMFH	.366±.015	.382±.005	.396±.010	.269±.025	.279±.004	.287±.001
	JIMFH	.433±.008	.449±.007	.448±.014	.368±.009	.372±.018	.379±.020
	DGCPN	.529±.014	.541±.007	.577±.016	.377±.003	.388±.008	.420±.012
	DCMH	.500±.010	.510±.007	.514±.005	.348±.020	.380±.008	.401±.009
	SSAH	.566±.012	.579±.006	.630±.008	.382±.005	.415±.013	.426±.007
	MetaCMH	.624±.002	.640±.006	.643±.032	.416±.004	.433±.003	.438±.009
	LtCMH	.729±.008	.738±.015	.750±.006	.441±.008	.458±.012	.463±.006

Table 2: Results (MAP) of each method on *long-tail* Flickr25K and NUS-WIDE. The best results are in boldface.

		Flicker25K			NUS-WIDE		
		16-bits	32-bits	64-bits	16-bits	32-bits	64bits
I→T	CMFH	.597±.037	.597±.036	.597±.037	.409±.041	.419±.051	.417±.062
	JIMFH	.621±.038	.635±.050	.635±.045	.503±.051	.524±.062	.529±.064
	DGCPN	.732±.010	.742±.004	.751±.008	.625±.003	.635±.007	.654±.020
	DCMH	.710±.022	.721±.017	.735±.014	.573±.027	.603±.003	.609±.002
	SSAH	.738±.018	.750±.010	.779±.009	.630±.001	.636±.010	.659±.004
	MetaCMH	.708±.007	.716±.005	.724±.013	.612±.011	.619±.005	.644±.018
	LtCMH	.745±.018	.753±.012	.781±.011	.635±.010	.654±.020	.678±.003
T→I	CMFH	.598±.048	.602±.055	.601±.065	.412±.072	.417±.061	.418±.068
	JIMFH	.650±.042	.662±.064	.657±.059	.584±.081	.604±.052	.655±.069
	DGCPN	.729±.015	.741±.008	.749±.014	.631±.008	.648±.013	.660±.004
	DCMH	.738±.020	.752±.020	.760±.019	.638±.010	.641±.011	.652±.012
	SSAH	.750±.028	.795±.023	.799±.008	.655±.002	.662±.012	.669±.009
	MetaCMH	.741±.009	.758±.014	.763±.004	.594±.004	.611±.007	.649±.008
	LtCMH	.770±.008	.795±.004	.802±.023	.608±.029	.644±.010	.688±.007

Table 3: Results (MAP) of each method on *balanced* Flickr25K and NUS-WIDE. The best results are in boldface.

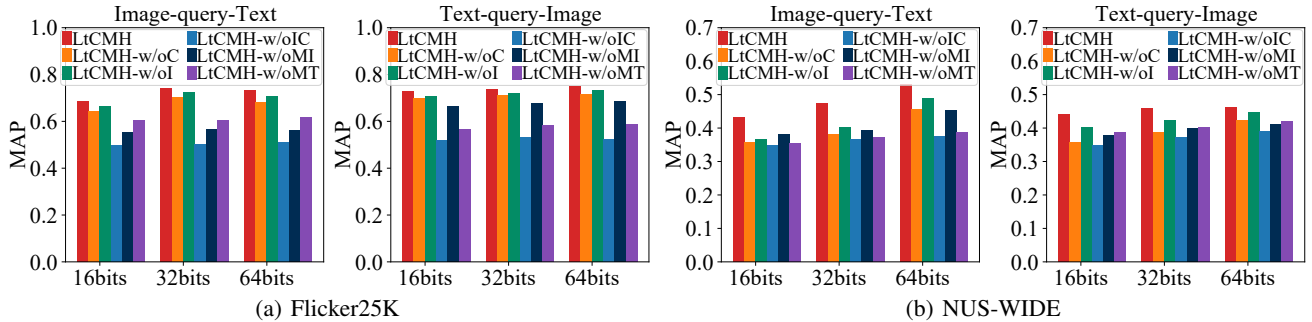


Figure 3: Results of LtCMH and its variants on long-tailed Flickr25K and NUS-WIDE.

		16bits	32bits	64bits
Flickr25K	LTHNet	.602	.573	.610
	LtCMH	.651	.647	.654
NUS-WIDE	LTHNet	.401	.411	.420
	LtCMH	.473	.508	.513

Table 4: Results of LtCMH and LTHNet on *long-tailed* Flickr25K and NUS-WIDE. Better results are in boldface.

pose an effective approach LtCMH that leverages the individuality and commonality of multi-modal data to create dynamic meta features, which enrich the representations of tail labels and give discriminant hash codes. The effectiveness and adaptivity of LtCMH are verified by experiments on long-tail and balanced multi-modal datasets.

Acknowledgements

This work is supported by NSFC (No. 62031003 and 62072380) and CAAI-Huawei MindSpore Open Fund.

References

- Buda, M.; Maki, A.; and Mazurowski, M. A. 2018. A systematic study of the class imbalance problem in convolutional neural networks. *Neural Netw.*, 106: 249–259.
- Byrd, J.; and Lipton, Z. 2019. What is the effect of importance weighting in deep learning? In *ICML*, 872–881.
- Chawla, N. V. 2009. Data mining for imbalanced datasets: An overview. *Data Mining and Knowledge Discovery Handbook*, 875–886.
- Chen, Y.; Hou, Y.; Leng, S.; Zhang, Q.; Lin, Z.; and Zhang, D. 2021. Long-tail hashing. In *SIGIR*, 1328–1338.
- Chu, P.; Bian, X.; Liu, S.; and Ling, H. 2020. Feature space augmentation for long-tailed data. In *ECCV*, 694–710.
- 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.
- Cui, Y.; Jia, M.; Lin, T.-Y.; Song, Y.; and Belongie, S. 2019. Class-balanced loss based on effective number of samples. In *CVPR*, 9268–9277.
- Ding, G.; Guo, Y.; and Zhou, J. 2014. Collective matrix factorization hashing for multimodal data. In *CVPR*, 2075–2082.
- Goodfellow, I.; Bengio, Y.; and Courville, A. 2016. *Deep learning*. MIT press.
- Gretton, A.; Bousquet, O.; Smola, A.; and Schölkopf, B. 2005. Measuring statistical dependence with Hilbert-Schmidt norms. In *ALT*, 63–77.
- Hausdorff, F. 2005. *Set theory*, volume 119. American Mathematical Soc.
- Huang, C.; Li, Y.; Loy, C. C.; and Tang, X. 2016. Learning deep representation for imbalanced classification. In *CVPR*, 5375–5384.
- Huiskes, M. J.; and Lew, M. S. 2008. The mir flickr retrieval evaluation. In *MIR*, 39–43.
- Jiang, Q.-Y.; and Li, W.-J. 2017. Deep cross-modal hashing. In *CVPR*, 3232–3240.
- Kou, X.; Xu, C.; Yang, X.; and Deng, C. 2022. Attention-guided contrastive hashing for long-tailed image retrieval. In *IJCAI*, 1017–1023.
- Kumar, S.; and Udupa, R. 2011. Learning hash functions for cross-view similarity search. In *IJCAI*, 1360–1365.
- Li, C.; Deng, C.; Li, N.; Liu, W.; Gao, X.; and Tao, D. 2018. Self-supervised adversarial hashing networks for cross-modal retrieval. In *CVPR*, 4242–4251.
- Liu, J.; Sun, Y.; Han, C.; Dou, Z.; and Li, W. 2020. Deep representation learning on long-tailed data: A learnable embedding augmentation perspective. In *CVPR*, 2970–2979.
- Liu, X.; Yu, G.; Domeniconi, C.; Wang, J.; Ren, Y.; and Guo, M. 2019a. Ranking-based deep cross-modal hashing. In *AAAI*, 4400–4407.
- Liu, X.; Yu, G.; Domeniconi, C.; Wang, J.; Xiao, G.; and Guo, M. 2019b. Weakly Supervised Cross-Modal Hashing. *TBD*, 552–563.
- Liu, X.; Yu, G.; Domeniconi, C.; Wang, J.; Xiao, G.; and Guo, M. 2022. Weakly supervised cross-modal hashing. *TBD*, 8(2): 552–563.
- Liu, Z.; Miao, Z.; Zhan, X.; Wang, J.; Gong, B.; and Yu, S. X. 2019c. Large-scale long-tailed recognition in an open world. In *CVPR*, 2537–2546.
- Reed, W. J. 2001. The Pareto, Zipf and other power laws. *Economics Letters*, 74(1): 15–19.
- Snell, J.; Swersky, K.; and Zemel, R. 2017. Prototypical networks for few-shot learning. In *NeurIPS*, 4080–4090.
- Su, S.; Zhong, Z.; and Zhang, C. 2019. Deep joint-semantics reconstructing hashing for large-scale unsupervised cross-modal retrieval. In *ICCV*, 3027–3035.
- Sun, C.; Latapie, H.; Liu, G.; and Yan, Y. 2022. Deep normalized cross-Modal hashing with bi-direction relation reasoning. In *CVPR*, 4941–4949.
- Tan, Q.; Yu, G.; Wang, J.; Domeniconi, C.; and Zhang, X. 2021. Individuality-and commonality-based multiview multilabel learning. *TCYB*, 51(3): 1716–1727.
- Tang, K.; Huang, J.; and Zhang, H. 2020. Long-tailed classification by keeping the good and removing the bad momentum causal effect. *NeurIPS*, 33: 1513–1524.
- Wang, D.; Wang, Q.; He, L.; Gao, X.; and Tian, Y. 2020a. Joint and individual matrix factorization hashing for large-scale cross-modal retrieval. *Pat. Recog.*, 107: 107479.
- Wang, K.; Yin, Q.; Wang, W.; Wu, S.; and Wang, L. 2016. A comprehensive survey on cross-modal retrieval. *arXiv preprint arXiv:1607.06215*.
- Wang, R.; Yu, G.; Domeniconi, C.; and Zhang, X. 2021. Meta cross-modal hashing on long-tailed data. *arXiv preprint arXiv:2111.04086*.
- Wang, T.; Li, Y.; Kang, B.; Li, J.; Liew, J.; Tang, S.; Hoi, S.; and Feng, J. 2020b. The devil is in classification: A simple framework for long-tail instance segmentation. In *ECCV*, 728–744.
- Wei, T.; Shi, J.-X.; Li, Y.-F.; and Zhang, M.-L. 2022. Prototypical classifier for robust class-imbalanced learning. In *PAKDD*, 44–57.
- Wu, G.; Lin, Z.; Han, J.; Liu, L.; Ding, G.; Zhang, B.; and Shen, J. 2018. Unsupervised deep hashing via binary latent factor models for large-scale cross-modal retrieval. In *IJCAI*, 2854–2860.
- Wu, X.; Chen, Q.-G.; Hu, Y.; Wang, D.; Chang, X.; Wang, X.; and Zhang, M.-L. 2019. Multi-view multi-label learning with view-specific information extraction. In *IJCAI*, 3884–3890.
- Yu, G.; Liu, X.; Wang, J.; Domeniconi, C.; and Zhang, X. 2022a. Flexible cross-modal hashing. *TNNLS*, 33(1): 304–314.
- Yu, G.; Xing, Y.; Wang, J.; Domeniconi, C.; and Zhang, X. 2022b. Multiview multi-instance multilabel active learning. *TNNLS*, 33(8): 1–12.

Yu, J.; Zhou, H.; Zhan, Y.; and Tao, D. 2021. Deep graph-neighbor coherence preserving network for unsupervised cross-modal hashing. In *AAAI*, 4626–4634.

Zhou, B.; Cui, Q.; Wei, X.-S.; and Chen, Z.-M. 2020. Bbn: Bilateral-branch network with cumulative learning for long-tailed visual recognition. In *CVPR*, 9719–9728.

Zhou, Z.-H.; Zhang, M.-L.; Huang, S.-J.; and Li, Y.-F. 2012. Multi-instance multi-label learning. *Artif. Intell.*, 176(1): 2291–2320.