

Singapore Management University

## Institutional Knowledge at Singapore Management University

---

Research Collection School Of Computing and Information Systems

School of Computing and Information Systems

---

10-2022

### Adaptive structural similarity preserving for unsupervised cross modal hashing

Liang LI

Baihua ZHENG

Singapore Management University, bhzheng@smu.edu.sg

Weiwei SUN

Follow this and additional works at: [https://ink.library.smu.edu.sg/sis\\_research](https://ink.library.smu.edu.sg/sis_research)



Part of the [Databases and Information Systems Commons](#)

---

#### Citation

LI, Liang; ZHENG, Baihua; and SUN, Weiwei. Adaptive structural similarity preserving for unsupervised cross modal hashing. (2022). *Proceedings of the 30th ACM International Conference on Multimedia, Lisbon, Portugal, 2022 October 10-14*. 3712-3721.

Available at: [https://ink.library.smu.edu.sg/sis\\_research/7435](https://ink.library.smu.edu.sg/sis_research/7435)

This Conference Proceeding Article is brought to you for free and open access by the School of Computing and Information Systems at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection School Of Computing and Information Systems by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email [cherylds@smu.edu.sg](mailto:cherylds@smu.edu.sg).

# Adaptive Structural Similarity Preserving for Unsupervised Cross Modal Hashing

Liang Li  
lil19@fudan.edu.cn  
School of Computer Science,  
Shanghai Key Laboratory of Data  
Science, Fudan University  
Shanghai Institute of Intelligent  
Electronics & Systems

Baihua Zheng  
bhzheng@smu.edu.sg  
School of Computing and Information  
Systems, Singapore Management  
University

Weiwei Sun\*  
wwsun@fudan.edu.cn  
School of Computer Science,  
Shanghai Key Laboratory of Data  
Science, Fudan University  
Shanghai Institute of Intelligent  
Electronics & Systems

## ABSTRACT

Cross-modal hashing is an important approach for multimodal data management and application. Existing unsupervised cross-modal hashing algorithms mainly rely on data features in pre-trained models to mine their similarity relationships. However, their optimization objectives are based on the static metric between the original uni-modal features, without further exploring data correlations during the training. In addition, most of them mainly focus on association mining and alignment among pairwise instances in continuous space but ignore the latent structural correlations contained in the semantic hashing space. In this paper, we propose an unsupervised hash learning framework *ASSPH* to solve the above problems. Firstly, we propose an adaptive learning scheme, with limited data and training batches, to enrich semantic correlations of unlabeled instances during the training process and meanwhile to ensure a smooth convergence of the training process. Secondly, we present an asymmetric structural semantic representation learning scheme. We introduce structural semantic metrics based on graph adjacency relations and meanwhile align the inter- and intra-modal semantics in the hash space with an asymmetric binary optimization process. Finally, we conduct extensive experiments to validate the enhancements of our work in comparison with existing works.

## CCS CONCEPTS

• **Information systems** → **Multimedia and multimodal retrieval**.

## KEYWORDS

Large-scale Multimodal/Cross-Modal Retrieval; Unsupervised Cross-Modal Hashing; Adaptive Structural Similarity Preserving

### ACM Reference Format:

Liang Li, Baihua Zheng, and Weiwei Sun. 2022. Adaptive Structural Similarity Preserving for Unsupervised Cross Modal Hashing. In *Proceedings*

\*Corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

*MM '22, October 10–14, 2022, Lisboa, Portugal.*

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9203-7/22/10...\$15.00

<https://doi.org/10.1145/3503161.3548431>

of the 30th ACM International Conference on Multimedia (MM '22), October 10–14, 2022, Lisboa, Portugal. ACM, New York, NY, USA, 10 pages.  
<https://doi.org/10.1145/3503161.3548431>

## 1 INTRODUCTION

Hash coding has become an essential method for similarity queries due to its computational efficiency and storage cost advantages. A series of hash representation learning methods [1, 18, 36] developed from local similarity hashing further improves the hashing retrieval efficiency. However, hash learning needs to consider additional factors for multi-modal data. In cross-modal hashing, instances from different modalities need to be compared and ranked together, which means the hamming distance of instances from different modalities can reflect the shared semantic relationship.

Thanks to the rapid development of representation learning, current cross-modal hashing methods often perform fine-tuning on the backbone encoding models with a joint-optimization alignment loss and a binary optimization strategy to obtain the final hashing model. High-quality encoding offered by the pre-trained models predominantly improves the cross-modal retrieval performance, especially the unsupervised performance. In recent years, unsupervised models, such as UnifiedVSE [42], VL-BERT [33], and VLP [50], have made breakthroughs on cross-modal retrieval tasks. In this paper, our hash-learning targets at a more general unsupervised cross-modal scenario [24, 32, 48], where hash models are learned to map the multimodal semantic feature vectors to a binary space such that the hash-learning can improve the performance of multimedia data management and retrieval in a broad manner. As our method is not limited to feature extraction backbones, it can be extended to other cross-modal retrieval between/among other modalities.

According to different optimization objectives, existing unsupervised cross-modal hashing methods can be clustered into alignment methods [47] and reconstruction methods [14, 32, 45]. Both classes of methods lack reliable cross-modal similarity as the supervision during the training [48]. Although some in-depth studies [45, 48] have been conducted in constructing more accurate cross-modal uniform metrics, they ignore the fact that the final goal of our hashing study is to obtain a plausible cross-modal metric, which is difficult to achieve based on a simple combination of raw multimodal similarity distribution.

To exploit the inter- and intra-modal semantic relationships, existing unsupervised hashing methods have widely adopted graph-based paradigms to regularize similarity relationships. However, existing methods often suffer from the “static graph” problem [31].

More specifically, they often employ features from original data or pre-trained models to build the explicit precomputed graphs. However, those pre-defined graphs cannot be adaptively learned during the training to better express the semantic structures. Different from the label-based reliable static association graph in the supervised methods [4, 28], the “static graph” constructed based on the original data measurement in the unsupervised methods [24, 45] actually introduces the bias in the original feature measurement.

As there is no reliable static optimization objective in unsupervised cross-modal hash learning, the difficulty in improving the unsupervised performance lies in how to smoothly refine the optimization objective and reduce the prior noise in the training process, thereby improving the performance. It has been confirmed in existing work [2] that a deep representation learning model with prior knowledge can be trained on middle results and iteratively converge. In recent research, the mainstream unsupervised research [11] has provided excellent examples through the optimization of memory banks and various variants. However, the memory-bank-based methods require massive training data as the basis, which is difficult to obtain for general application scenarios.

We aim at optimizing the optimization goal (“static graph”) in the learning process to mine more information under unsupervised conditions. In order to ensure a smooth convergence of training, we do not directly adopt the continuous similarity between intermediate results, but propose a two-stage correlational relationship mining strategy to get more reliable prior information as supervision while being able to smoothly update the correlational relationship set during the training to enrich the semantic correlations. And then, we introduce a structural semantic representation and an asymmetric binary code learning scheme together to preserve the semantic similarity in the final hash space.

In summary, this paper makes a three-fold contribution. First, we propose a novel adaptive cross-modal correlation mining and structural semantic maintenance strategy. The optimization target is updated indirectly by adaptively expanding the semantic-similar neighborhood to solve the “static graph” problem and maintain a smooth convergence of the optimization process. Though traditional unsupervised learning methods are strictly limited by static constraints, our method does not rely on expensive calculation and data resources in contrastive learning and can be widely adopted in hash retrieval optimization. Second, we introduce a novel cross-modal semantic preservation framework as the backbone of our hash learning to bridge the modality gap. Specifically, we learn common hash representations through multi-level structural semantic consistency constraints, which serve as the basis of adaptive learning. In addition, we design an asymmetric similarity-preserving binary optimization algorithm to reduce the information loss after binarization. Last but not least, we perform extensive experimental studies on two publicly available datasets, NUSWIDE and MIRFlickr-25K, and our approach is able to significantly outperform state-of-the-art methods proposed in the last three years.

## 2 RELATED WORKS

Hashing learning [35, 36] in uni-modal and cross-modal datasets aims to learn a hashing function to map the original data to a binary hashing space while the hashing metric is expected to preserve the

semantic similarities between original multimedia data. For cross-modal hashing, the additional modality gives new opportunities to describe and dig up the correlation between/among modalities, and thus, the contrastive-based methods can perform well. However, due to the distribution difference in various data modalities, the structural-semantic difference, called the *semantic gap*, impedes the construction of semantic mapping. As labels play an important role in existing cross-modal learning to determine whether two instances are similar, related works in the past decade can be categorized into the following two classes, i.e., *supervised hashing* and *unsupervised hashing*.

**Supervised Cross-Modal Hashing.** Early supervised cross-modal hashing methods, e.g., Semantic Correlation Maximization (SCM) [46] and Semantics-Preserving Hashing (SePH) [23], perform the reconstruction via different approaches, including semantic similarity matrix preserving and KL-divergence based correlation measurement. Recent works [13, 18, 20, 34] start using the deep neural network and semantic embeddings provided by pre-trained models. A subsequent series of works [5, 25, 26, 29, 39, 44] explores the relationships of semantic tags and latent semantics with different conditions and improves the hashing performance. Supervised methods achieve good performance by using labels as supervision information to align different modalities; however, due to the high cost of labeling, they cannot be widely used. On the contrary, unsupervised methods, to be reviewed next, are not limited by labels and thus can be applied to a wider range of applications.

**Unsupervised Cross-Modal Hashing.** Unsupervised cross-modal hashing methods [7, 10, 27, 40] aim to maximize the correlation between different modals or to preserve the inter-modal and intra-modal similarity, like Canonical Correlation Analysis (CCA) [10] and Composite Correlation Quantization (CCQ) [27]. With the help of deep representation learning, some recent methods [2, 19] can describe inter-modal correlations more precisely and perform well on semantic alignment.

According to their optimization objectives, recent unsupervised hashing methods can be broadly divided into alignment methods and reconstruction methods. To overcome the lack of semantic supervision under unsupervised conditions, the first category usually utilizes some data-mining methods to describe the correlational relationship (pairwise or triplet relationship), a measurement of the similarity or semantic difference between two cross-modal instances [12, 15, 37, 38, 41, 47, 51]. Some contrastive-learning-based frameworks are introduced to learn the hash representation. Some works [9, 21] even try to use the generative models [30] to describe the distribution of each modality and then make an alignment while introducing additional overhead during the hash coding.

The second category adopts the semantic reconstruction [24, 32, 43, 45, 48], which considers the similarity as the exact distance value in the hash space instead of relevance similarity or probability. While keeping the semantic metric in the hash space directly, they are all constrained by the noise arising in the feature extraction. Some works [38, 48] use a two-stage optimization and independent codebook to update the similarity guideline simultaneously, e.g., UKD [14] proposes a two-stage data distillation framework to the cross-modal hashing learning and demonstrates the effectiveness of the cross-modal hidden embeddings to optimize the following cross-modal hashing network further.

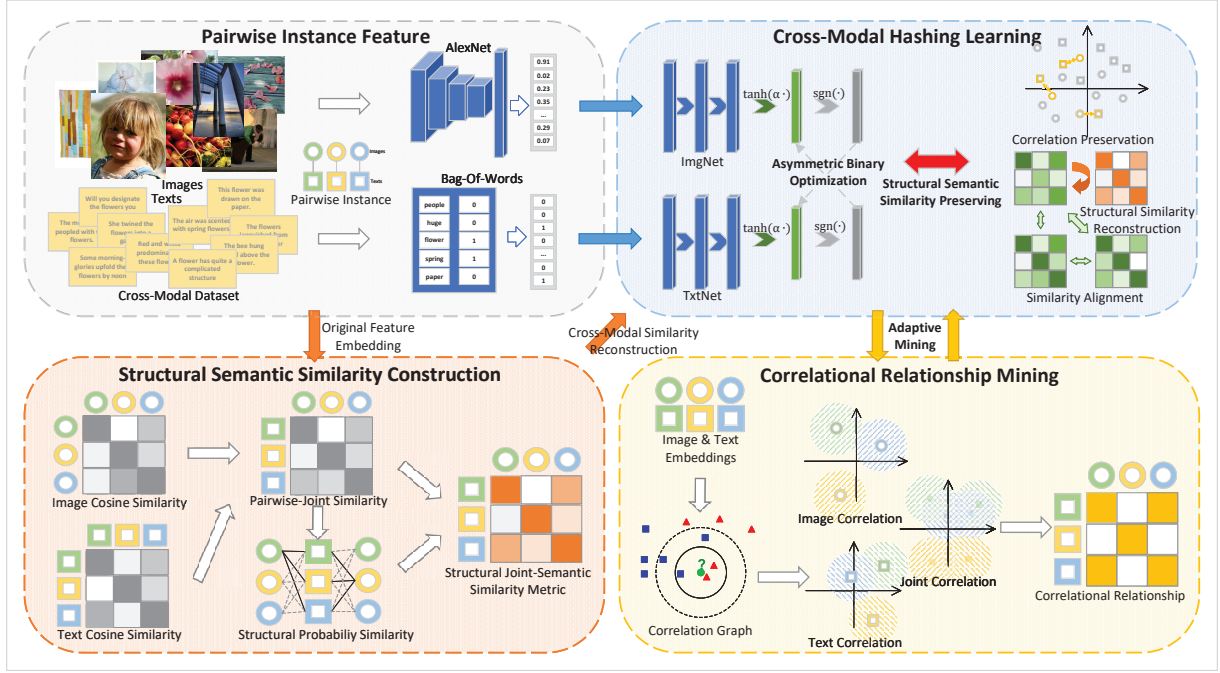


Figure 1: An overview of our Adaptive Structural Similarity Preserving Hash (ASSPH) model.

### 3 PROPOSED APPROACH

This paper focuses on cross-modal retrieval between image and text modalities. Let  $O_i = \{v_i, t_i\}$ ,  $i \in \{1, 2, \dots, M\}$  indicate an  $i^{\text{th}}$  cross-modal instance,  $M$  represent the total number of instances,  $v_i \in \mathbb{R}^{d_o} / t_i \in \mathbb{R}^{d_t}$  refer to the visual/text feature in the pairwise instance, and  $d_o/d_t$  represents the dimension of the image/text features.

**Problem Formulation.** Given the cross-modal set  $O$  with the corresponding image feature set  $I = \{v_i\}_{i=1}^M$  and text feature set  $T = \{t_i\}_{i=1}^M$ , we want to learn hash functions  $f_I(v, \theta^I)$  and  $f_T(t, \theta^T)$  to generate hash code  $B^I, B^T \in \{-1, +1\}^{M \times K}$  for image and text modalities respectively.  $K$  is the length of the hash code,  $\theta^I$  and  $\theta^T$  are the parameters of ImgNet and TxtNet for hash code to be optimized. In addition, the similarity between hash codes  $b_i^I$  and  $b_j^T$ , i.e., their Hamming distance  $D(b_i^I, b_j^T) = \frac{1}{2}(K - \langle b_i^I, b_j^T \rangle)$ , is able to reflect the semantic similarity between original instances  $v_i$  and  $t_j$ .

In this paper, we propose *Adaptive Structural Similarity Preserving Hash (ASSPH)* model, as a solution to the problem formulated above. ASSPH designs a comprehensive structured semantic retention strategy with a new positive sample expansion technique to stabilize the training process while smoothening the updates of the optimization target to enrich the cross-modal semantic [2, 3]. In addition, ASSPH introduces an asymmetric deep hash learning method to further reduce the quantization loss and to preserve the structural semantic in the hashing space to improve the retrieval performance. ASSPH consists of four major components, as shown in Figure 1, and Algorithm 1 provides the pseudocode of the proposed hash learning method. In the rest of this section, we will detail the four components.

#### 3.1 Feature Extraction

Feature extraction is an independent component in ASSPH. We follow the setup of existing work and use AlexNet and Bag-of-Words for extracting image feature  $F^I$  and text feature  $F^T$  respectively. Switching to a more efficient model like VGG or ResNet does improve the final result, but it is not the main focus of this work. ASSPH supports different feature extraction models in order to support a wider range of application scenarios.

#### 3.2 Structural Semantic Similarity Construction

Once the original semantic features  $F^I$  and  $F^T$  of the cross-modal training set  $O = \{(v_i, t_i)\}_1^M$  are obtained, ASSPH constructs their structural similarities. It first transforms the original semantic metric into a similarity probability based on the metric distribution and practical application implications; it then defines the joint-modal similarity and structural semantic similarity between pairs of data respectively; it finally combines these two similarities together to form the structural semantic similarity metric.

To be more specific, we firstly follow the previous works [13] to calculate the pair-wise modality-specific cosine metrics within different modals,  $S^* = \{s_{ij}^* = \cos(f_i^*, f_j^*)\}_{i,j=1}^M$  to capture pair-wise semantic similarities between instances within a single modality. Note, in this paper,  $*$   $\in \{I, T\}$  represents either image or text modality. As we find that the distribution shape of  $S^*$  is approximately a skewed normal distribution, we map  $S^*$  from  $[-1, 1]$  to the interval  $[0, 1]$ , denoting it as  $S_p^* = 2S^* - 1$ .  $S_p^*$  is regarded as the probability that can describe whether two instances in a single modality (i.e.,  $v_i$  and  $v_j$  in image modality or  $t_i$  and  $t_j$  in text modality) are semantically similar. We then calculate its cross-modal similarity  $S_{fusion}$  under the independence condition to subtract the common part and map to  $[0, 1]$ , i.e.,  $S_{fusion} = S_p^I + S_p^T - S_p^I \cdot S_p^T$ .

To define the structural semantic, existing works [24, 32] directly compare the fusion similarity of neighbor distance distributions between instances as the structural similarity. However, we observe that when the fusion similarity value is small, its accuracy as a metric is correspondingly reduced. Take the fusion similarity constructed on MIRFlickr as an example. MAP@50 is around 0.9 when MAP@all decreases to about 0.75. To tackle this issue, we incorporate two optimizations. Firstly, we keep only the top  $K_S$  nearest distances based on fusion similarity  $S_{fusion}$ , but remove the rest of distances that are much noisier and less useful. In our implementation,  $NN_{K_S}(i)$  w.r.t. an instance  $i$  captures its top- $K_S$  nearest neighbors. Secondly, we reweight the distances between the remaining instances based on their original cosine distance to reduce the impact of uneven density. That is to say, if  $j \in NN_{K_S}(i)$ ,  $\hat{S}(i, j) = \frac{S_{fusion}(i, j)}{\sum_{k \in NN_{K_S}(i)} S_{fusion}(i, k)}$ ; otherwise,  $\hat{S}(i, j) = 0$ . The structural similarity is defined as  $S_{structure} = K_S \hat{S} \times \hat{S}^T$ . Here,  $K_S$  is used again to rebalance data range;  $\hat{S}$  refers to the matrix transpose of  $\hat{S}$ , and we will follow this notation later.

Based on the above definition, we define the structural semantic similarity  $S$  by considering both cross-modal similarity and structural similarity as a multi-level semantic metric to fully represent the structural semantic information. We finally map  $S$  back to the range of the cosine metric, i.e.  $S = 2\hat{S} - 1$ .

$$\hat{S}' = (1 - \gamma)S_{fusion} + \gamma S_{structure}, \gamma \in [0, 1] \quad (1)$$

### 3.3 Correlational Relationship Mining

Our ASSPH model proposes to construct a set of correlational relationships as the positive samples to further guide the pair-wise cross-modal hashing alignment. In addition to correcting bias in static structural semantic similarity, we introduce the correlational relationships because of the following considerations. On the one hand, limited by the similarity distribution between the original data, the existing static similarity measures [24, 32, 45] tend to have poor local consistency. That is, most of the measurement results are in the insensitive range, which makes it challenging to effectively map correlated instances into adjacent hash codes. On the other hand, adjusting the whole continuous similarity metric space directly can have uncontrollable effects on the training process and lead to training collapse, especially when the training data is not large enough. Inspired by DeepCluster [2], we try to solve this problem by revising the discrete constraints. We introduce a new discrete training target here, which can use the knowledge learned by the model during each training step to further strengthen the local consistency of the hash representation and ensure that the training process can eventually converge smoothly.

To be more specific, we define a strict correlation discrimination scheme based on the common adjacency relations in inter- and intra-modals. We initialize the correlated instance set by considering their original semantic features, and then adaptively expand the set based on the currently learned cross-modal hidden-layer representations during the model training process.

**Initialization of the Correlational Relationship.** We introduce a two-stage positive sample expansion strategy to initialize the the binary-valued correlational relationship set defined as  $R =$

$\{r_{ij} \in \{0, 1\}\}_{i, j=1}^M$ . In the first stage, we perform proximity instance mining based on  $S^I$  and  $S^T$  (the pair-wise modality-specific cosine metrics introduced previously) respectively. To be more specific, for  $j \in NN_{K_R}(i)$ ,  $R_1^*[i, j] = 1$ ; otherwise,  $R_1^*[i, j] = 0$ . Note  $K_R$  controls the range of neighborhood. To preserve the semantic consistency of proximity instances,  $K_R$  will be much smaller than  $K_S$  and it will be dynamically expanded in adaptive similarity learning to be detailed next.

However, KNN-based results generated in the above stage still consist of many errors, e.g., the neighboring instances of some marginal data or outliers often have different semantics. Therefore, we further define the second-order proximity relation  $R_2^*$  to measure the similarity between instances by their shared neighbors. If  $R_1^* \times R_1^{*'} \geq \tau$ ,  $R_2^* = 1$ ; otherwise,  $R_2^* = 0$ . Here,  $\tau$  is a similarity threshold, and it is set to 1 in practice to ensure better results. Similarly, we can define the cross-modal correlation  $R_2^{cross}$ . If  $\max(R_1^I \times R_1^{T'}, R_1^T \times R_1^{I'}) \geq \tau$ ,  $R_2^{cross} = 1$ ; otherwise,  $R_2^{cross} = 0$ . Finally, Eq. (2) defines the binary-valued correlational relationship  $R$ .

$$R = R_2^I \cup R_2^T \cup R_2^{cross} \quad (2)$$

**Adaptive Similarity Learning.** Adaptive learning aims to keep the optimization target consistent [11] to avoid training collapse due to convergence to a trivial solution [8]. To maintain stable iterations during the training process under the conditions of limited training batches and small datasets, we adopt a clear and effective expansion strategy based on the correlational relationship defined above.

For each epoch online, we update the correlational relationship. After the  $e^{th}$  training round, we can obtain hidden embeddings  $H$  containing cross-modal semantics from the latest hash functions  $ImgNet(F_I, \theta_e^I)$  and  $TxtNet(F_T, \theta_e^T)$ , i.e.,  $H_I^I = ImgNet_e(F^I, \theta_e^I)$  and  $H_I^T = TxtNet_e(F^T, \theta_e^T)$ . Then, we replace the original features with the hidden features of the current hash layer and use the two-stage correlation mining strategy presented above to obtain new cross-modal relationships  $R_e$ . Finally, we union it with the current  $R_{e-1}$ , as stated in Eq. (3).

$$R_e = R_{e-1} \cup R(H^I, H^T) \quad (3)$$

To ensure convergence, we adopt a lower learning rate while strictly limiting the size of  $K_R$  to control the reliability of the positive correlations. Subsequent experiments demonstrate that our correlation mining process further enriches the semantic correlation distribution while maintaining convergence. Due to structural semantic construction, our method requires  $O(M^2)$  complexity for training, where  $M$  represents the training set size. We argue that it is acceptable in most cases because of limited training set size.

### 3.4 Cross-Modal Hashing Learning

The key point of hash learning is to ensure that the distance relationship between hash codes can reflect the semantic similarity of original data instances. However, it is challenging to optimize binary representations in gradient backpropagation. Therefore, in our cross-modal hash learning module, we first design a set of cross-modal alignment loss functions to optimize the continuous hash hidden layer activated by the  $\tanh()$  function, and then introduce an asymmetric binary optimization approach to improve the semantic expression ability of the binary hashing representation.

**Structural Semantic Similarity Preservation.** We adopt two independent networks  $ImgNet(F^I, \theta^I)$  and  $TxtNet(F^T, \theta^T)$  to get the continuous cross-modal representations. Take the hidden features  $H^I$  from  $ImgNet()$  and  $H^T$  from  $TxtNet()$  as input, we design a multi-task learning paradigm to jointly learn the hash code to maintain the semantics within the original cross-modal data. The optimization task consists of three parts: structural semantic reconstruction, correlational relationship preservation, and semantic alignment.

First, we follow the similarity preservation method in existing unsupervised learning approaches. Existing cross-modal unsupervised methods can be classified into contrastive learning, autoencoders, and similarity preservation, according to the category of the loss function used. Among them, similarity preservation has the lowest dependence on training batch size, which benefits scenarios with limited computational resources. Therefore, we first define a similarity preserving loss in Eq. (4). That is, both intra-modal and inter-modal metric results should be consistent with a pre-defined structural similarity metric.

$$\mathcal{L}_{sr}(H^I, H^T) = \left\| S - \cos(H^I, H^T) \right\|_F^2 + \left\| S - \cos(H^I, H^I) \right\|_F^2 + \left\| S - \cos(H^T, H^T) \right\|_F^2 \quad (4)$$

Second, we want to achieve symmetric consistency, i.e., both intra and inter-modal similarities are consistent. Therefore, we construct a semantic alignment loss in Eq. (5) to ensure that both the intra-modality and inter-modality metrics are consistent.

$$\mathcal{L}_{sa}(H^I, H^T) = \left\| \cos(H^I, H^I) - \cos(H^T, H^T) \right\|_F^2 + \left\| \cos(H^I, H^T) - \cos(H^I, H^I) \right\|_F^2 + \left\| \cos(H^I, H^T) - \cos(H^T, H^T) \right\|_F^2 \quad (5)$$

As the above-mentioned two optimization goals are both based on static graph and hence will be affected by the prior deviation of the static graph, we introduce the set of correlational relationships  $R$  to rectify the training process. That is, the distance of related instances shall be as close as possible. Meanwhile, considering the massive difference in the number of positive and negative instance pairs, we introduce a re-balance coefficient  $\beta$  to optimize the performance. The semantic consistency loss is defined as follows:

$$\mathcal{L}_{cp}(H^I, H^T) = \left\| \cos(H^I, H^T) \cdot R - \beta R \right\|_F^2 \quad (6)$$

Eq. (7) defines the final training objective of our ASSPH, where  $\mu_1$  and  $\mu_2$  are the trade-off parameters to balance different optimization goals.

$$\min_{\theta^I, \theta^T} \mathcal{L} = \mathcal{L}_{sr} + \mu_1 \mathcal{L}_{cp} + \mu_2 \mathcal{L}_{sa} \quad (7)$$

**Asymmetric Similarity-Preserving Binary Optimization.** As our goal is to optimize the semantic loss  $\mathcal{L}(B^I, B^T)$  between binary hash codes, we design the asymmetric binary optimization to constrain the hidden semantic space and the binary hashing space together. Traditional work mainly optimizes hash coding by reducing the quantization loss from real value to binary value. Still, gradient descent based algorithms in symmetric hashing learning will force the pair-wise instance to be updated towards the

---

**Algorithm 1:** Adaptive Structural Sim. Preserving Hashing
 

---

**Input:** Training set  $\{I_k, T_k\}_{k=1}^M$  and their corresponding features  $F^I$  and  $F^T$ ; batch size  $m$ ; train epoch  $E$ ,  $NN_{K_R}$  scale  $K_R$  and  $NN_{K_S}$  scale  $K_S$ .  
**Output:** Hashing function  $ImgNet(F^I, \theta^I)$  and  $TxtNet(F^T, \theta^T)$

- 1 Initialize  $\theta^I, \theta^T$ , epoch  $e \leftarrow 0$ ,  $S \leftarrow$  Eq. (1),  $R \leftarrow$  Eq. (2);
- 2 **for** each  $e \in [1, E]$  **do**
- 3     **for**  $\lfloor \frac{M}{m} \rfloor$  iterations **do**
- 4         Sample  $m$  instances  $\{f_k^I, f_k^T\}_{k=1}^m$  from training set;
- 5         Calculate the representations  $H_k^I \leftarrow ImgNet(f_k^I, \theta^I)$ ,  $H_k^T \leftarrow TxtNet(f_k^T, \theta^T)$  and get the corresponding  $S(H^I, H^T)$ ;
- 6         Update  $\theta^I$  and  $\theta^T$  using  $\mathcal{L}(H^I, H^T)$ ;
- 7         Calculate  $B^I \leftarrow sgn(H^I)$  and  $B^T \leftarrow sgn(H^T)$ ;
- 8         Update  $\theta^I/\theta^T$  using  $\mathcal{L}(H^I, B^T)/\mathcal{L}(B^I, H^T)$ ;
- 9         Update  $R_e$  with Eq. (3);
- 10     Return  $ImgNet(F^I, \theta^I)$  and  $TxtNet(F^T, \theta^T)$ ;

---

same directions and switch the sign of the hash code in the worst case [16].

Inspired by the asymmetric hashing methods [29, 49], we propose a two-stage asymmetric binary similarity optimization strategy and stabilize the hash learning by optimizing one side of the sub-network at each stage. Specifically, we optimize the continuous coding of one modality with the help of the binary coding of the other modality using  $\mathcal{L}(H^I, B^T)$  and  $\mathcal{L}(B^I, H^T)$ , which also ensures the consistency of both the structure and the semantics of the hash space.

Besides, we also follow the previous approximation strategy [24] to reduce the quantization loss by taking advantage of the convergence of the  $\tanh()$  function to  $sgn()$ . We replace the activation function of the final layer with  $\tanh()$  and adopt the following strategy  $h_i = \tanh(\eta x_i)$ ,  $\eta \rightarrow +\infty$ . During the training stage,  $\eta$  will gradually increase with the increase of epoch.

## 4 EXPERIMENTAL EVALUATION

### 4.1 Experiment Setup

**Datasets.** Two commonly used multi-label image-text cross-modal datasets, NUS-WIDE and MIRFlickr-25K, are chosen for our experiments. As a piece of unsupervised work, we do not use the label information in the dataset during the training process but only use it as a basis when verifying the performance. *NUS-WIDE* [6] consists of 269,648 labeled multimodal instances, each containing an image and a text description in pairs. Following previous work [24, 32, 41], we keep the instances corresponding to the top-10 most frequent labeled classes, resulting in 186,577 instances. *MIRFlickr-25K* [17] consists of 25,000 tagged multimodal instances collected from the Flickr website. Each instance contains an image and a corresponding text description, and they are divided into 24 categories in total. After removing the unlabeled instances, there are 20,015 instances left. In both datasets, we follow previous work [24, 32, 41] to select 2,000 random instances as the query set and use the remaining

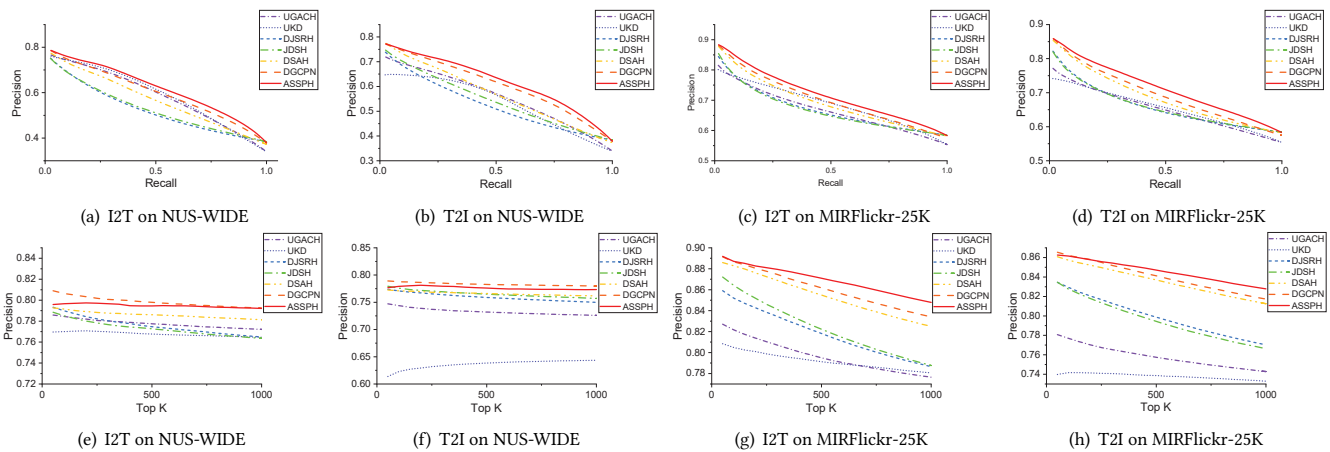


Figure 2: The precision-recall and Top- $k$  precision curves on NUS-WIDE and MIRFlickr-25K datasets with 64 bit hash code.

instances to form the database, which also includes a training set consisting of 5,000 instances. As for the single-label data set, it is a complex scene for our method because the positive correlation between data is much sparser than that of the multi-label scene, resulting in more errors in the unsupervised construction of the correlation. However, our model still achieves some improvements over existing methods. We have evaluated our ASSPH on the single-label Wiki dataset and please refer to [22] for the details.

**Evaluation Protocol.** We verify the accuracy of hash retrieval in the task of mutual image and text retrieval, i.e., using an image as a query and retrieving the text associated with it (I2T) and vice versa (T2I). To evaluate the retrieval quality, we adopt three standard evaluation metrics, including Mean Average Precision (MAP), Precision-Recall Curve (PR-Curve), and Top- $k$ -Precision. At the same time, considering the actual application scenario, we also evaluate our ASSPH under MAP@50 and again please refer to [22] for the detailed MAP@50 results.

**Baselines.** We implement eight state-of-the-art methods based on deep hash learning proposed in the last three years as competitors, including UGACH [47], UKD [14], SRCH [38], DJSRH [32], DSAH [43], JDSH [24], DGCPN [45] and UCH [21].

**Implementation Details.** In our implementation, we take 4,096 dimensional features extracted from AlexNet for images and use the 1,000-dimensional BoW features for texts. For the hashing representation network, we directly use two nonlinear layers with 4,096 dimensions in the hidden layer. For the competitors, we adopt their original implementations released by the authors and follow their best settings to perform the experiments. We also choose the same image and text embeddings (AlexNet and BoW features) for fair comparisons. The training epoch is uniformly set to 50 on both datasets, except UGACH and UKD with their own settings. Our method and competitors are implemented based on python 3.6 and trained on a single 2080ti GPU. Last but not least, we fix the batch size to 32 and use the SGD optimizer with 0.9 momentum and 0.0005 weight decay. The learning rates are set to 0.001. For other parameters in our ASSPH, we perform cross-validation and finally select the following set of parameters with the selection range in brackets:  $K_R = 50, [5, 500]$ ,  $K_S = 2000, [1000, 4000]$ ,  $\mu_1 = 2[0.5, 5]$ ,

$\mu_2 = 1[0.5, 5]$ ,  $\beta = 1.5, [1, 3]$  and  $\gamma = 0.3[0, 1]$ . We use the same parameters on both datasets.

## 4.2 Retrieval Performance

We compare the performance of ASSPH against its competitors. Note that since the implementations of UCH and SRCH have not been released, we directly copy the MAP results from the original papers without PR-Curve and Top- $k$  Precision Curve.

First, we report in the top part of Table 1 the MAP performance (averaged over three times) based on different code lengths for two tasks, I2T and T2I. Our ASSPH consistently performs the best.

Then, we report PR-Curve (without the two data points when recall is 0 and recall is 1) and Top- $k$  in Figure 2 with 64-bit code length. As expected, our method also outperforms all the competitors in the PR curves for both tasks. For the Top- $k$  indicator, when  $k$  is small, our improvement is not obvious, but as  $k$  increases, our ASSPH performs better. This is because these original static similarity metrics in competitors are reliable in top-ranked results. The semantic metric reconstructed in DJSRH, JDSH and DGCPN depends heavily on the accuracy of the original semantic metric, while the current deep learning features have a high confidence level in top-ranked results (e.g., Top-50). In other words, they can only take advantage of the relationship between a few data (such as the relationship between the top-50 related data), and most of the relationships can only be filtered out as noise. However, ASSPH can make full use of the relationship between all data. While the improvement of ASSPH may be relatively limited in the top-50 retrieval performance, ASSPH can instead optimize the overall retrieval results.

Next, we report some results with a different CNN backbone. Some unsupervised hash works [45, 47] use VGG to get the image features with higher quality to improve the effectiveness of unsupervised works. Accordingly, we report the retrieval performance of different methods based on VGG-16 in the bottom part of Table 1. For the convenience of comparison, we directly list the experimental results reported by each competitor in the original literature. Note that DJSRH, JDSH, and DSAH are missing here because their original papers only report the results of MAP@50, and the metric we compare are based on MAP@all. Using AlexNet will result in a

**Table 1: The MAP performances of MIRFlickr-25K and NUS-WIDE at various hashing code lengths**

	Task	Method	MIRFlickr-25K				NUS-WIDE			
			16-bit	32-bit	64-bit	128-bit	16-bit	32-bit	64-bit	128-bit
CNN backbone: AlexNet	I2T	UGACH	0.6595	0.6674	0.6777	0.6805	0.6107	0.6180	0.6167	0.6227
		UKD-SS	0.6922	0.6903	0.6982	0.6963	0.6026	0.6202	0.6264	0.6112
		<i>UCH*</i>	0.6540	0.6690	0.6790	-	-	-	-	-
		<i>SRCH*</i>	0.6808	0.6916	0.6997	-	0.5441	0.5565	0.5671	-
		DJSRH	0.6576	0.6606	0.6758	0.6856	0.5073	0.5239	0.5422	0.5530
		JDSH	0.6417	0.6567	0.6731	0.6850	0.5128	0.5170	0.5475	0.5563
		DSAH	0.6914	0.6981	0.7035	0.7086	0.5677	0.5839	0.5959	0.6053
		DGCPN	0.6991	0.7116	0.715	0.7243	0.5993	0.6150	0.6262	0.6338
		<b>ASSPH</b>	<b>0.7138</b>	<b>0.7244</b>	<b>0.7301</b>	<b>0.7341</b>	<b>0.6238</b>	<b>0.6340</b>	<b>0.6477</b>	<b>0.6512</b>
	T2I	UGACH	0.6512	0.6587	0.6677	0.6693	0.5993	0.6022	0.6019	0.6006
		UKD-SS	0.6760	0.6799	0.6689	0.6722	0.5830	0.6042	0.5963	0.6075
		<i>UCH*</i>	0.6610	0.6670	0.6680	-	-	-	-	-
		<i>SRCH*</i>	0.6971	0.7081	0.7146	-	0.5533	0.5670	0.5754	-
		DJSRH	0.6594	0.6528	0.6688	0.6801	0.4936	0.5325	0.5508	0.5470
		JDSH	0.6352	0.6608	0.6686	0.6736	0.5153	0.4982	0.5683	0.5704
		DSAH	0.6903	0.6916	0.7013	0.7038	0.5695	0.5874	0.6066	0.6114
		DGCPN	0.6991	0.7074	0.7133	0.7233	0.6141	0.6353	0.6456	0.6522
		<b>ASSPH</b>	<b>0.7170</b>	<b>0.7265</b>	<b>0.7315</b>	<b>0.7380</b>	<b>0.6332</b>	<b>0.6502</b>	<b>0.6646</b>	<b>0.6628</b>
CNN Backbone: VGG-16	I2T	UGACH	0.685	0.693	0.704	0.702	0.613	0.623	0.628	0.631
		UKD-SS	0.714	0.718	0.725	0.720	0.614	0.637	0.638	0.645
		<i>UCH</i>	0.654	0.669	0.679	-	-	-	-	-
		<i>SRCH</i>	0.681	0.692	0.700	-	0.544	0.557	0.567	-
		DGCPN	0.732	0.742	0.751	-	0.625	0.635	0.654	-
		<b>ASSPH</b>	<b>0.739</b>	<b>0.753</b>	<b>0.757</b>	<b>0.762</b>	<b>0.639</b>	<b>0.660</b>	<b>0.667</b>	<b>0.669</b>
	T2I	UGACH	0.673	0.676	0.686	0.690	0.603	0.614	0.640	0.641
		UKD-SS	0.715	0.716	0.721	0.719	0.630	0.656	0.657	0.663
		<i>UCH</i>	0.661	0.667	0.668	-	-	-	-	-
		<i>SRCH</i>	0.6971	0.7081	0.7146	-	0.5533	0.567	0.5754	-
		DGCPN	0.729	0.741	0.749	-	0.631	0.648	0.660	-
		<b>ASSPH</b>	<b>0.746</b>	<b>0.756</b>	<b>0.764</b>	<b>0.767</b>	<b>0.652</b>	<b>0.673</b>	<b>0.676</b>	<b>0.674</b>

2% to 4% reduction in the overall effect of the experiments compared with the performance in the top part of Table 1. This is the main reason why our results reported in the top part of Table 1 deviate from the results in their original literature.

We observe that in the experiment based on VGG, our improvement is smaller than that based on AlexNet. This is because the prior knowledge used as the training targets in existing work, such as the neighboring network in UGACH [47] or the cross-mode similarity measure in DGCPN [45], will be more reliable when the feature quality improves. However, our method still outperforms its competitors regardless of the original feature quality. The difference between the performance with different backbone models can further reflect the versatility of ASSPH, i.e., it remains effective even when the original similarity metric has low quality.

### 4.3 Ablation Studies

We design several variants of ASSPH to show the effectiveness of several key modules. To be more specific, we implement four variants models, including i) **ASSPH\_NoAdapt** that removes the adaptive learning strategy from ASSPH, so  $R$  will remain unchanged during the whole training process; ii) **ASSPH\_PairCorr** that replaces the structural correlation with the pairwise correlation; iii)

**ASSPH\_NoCorr** that removes the correlational relationship and the corresponding correlation preserving loss  $\mathcal{L}_{cp}$ ; and iv) **ASSPH\_NoBinOpt** that removes the similarity-preserving binary optimization strategy from ASSPH.

The performance comparison among different variants is reported in Table 2. In most cases, we can observe that these four key modules effectively improve the performance by different extents and the results fully illustrate their effectiveness in ASSPH. For the BinOpt module, it is introduced to speedup the convergence process to the optimal binary representation. As the training on MIRFlickr tends to converge quickly, our BinOpt module does not improve much in this fast-converging scenario, so ASSPH performs only slightly better than ASSPH\_NoBinOpt.

### 4.4 Parameter analysis

In this set of experiments, we conduct further analysis on the effects of different parameters in our method, including the  $K_S$  and  $\gamma$  in structural similarity construction,  $K_R$  in correlational relationship mining and the training batch size. Parameter  $\beta$ , a balance parameter to adjust the data bias in Eq. (6), has a rather small optional range (e.g., [1,2]). We follow previous works [24, 32, 43, 45] to set  $\beta$  to 1.5, and perform cross-validation to confirm its plausibility.



First, we study the value of  $K_R$ , denoted by  $N$ , that records the neighborhood range used as positive samples in correlation mining. Figure 3 shows the MAP performance achieved by our method under different  $N$  values. As it is a single-peaked function of neighbors, we can see that when the adaptive learning strategy is effective under proper settings and our ASSPH can increase positive relations further while maintaining its reliability.

Next, we study the value of  $K_S$  in Table 3. The table shows the cross-modal retrieval performance under different values of  $K_S$ . We can observe that, on the one hand, the final result is not sensitive to the value of  $K_S$ , and on the other hand, when  $K_S$  increases significantly, the final result decreases. This is because as  $K_S$  is approaching the training set size, we actually consider the correlation between all instances, which undoubtedly introduces additional noise because the similarity relations between distant instances are less credible.

In the same table, we also report the results of our experiments with different batch sizes/ $\gamma$  values. As ASSPH has a low dependence on batch size, it allows the unsupervised learning strategy to learn effectively even in small samples and limited datasets. As the batch size increases from its default value 32 to 256, the performance changes are small. Parameter  $\gamma$  controls the weight of structural similarity when calculating the structural semantic similarity  $S$ . We can observe that our method is still effective even when  $\gamma = 0$  (i.e.,  $S$  is purely based on structural similarity). However, combining structural similarity and static cross-modal similarity as a multi-level cross-mode similarity can further improve the performance.

#### 4.5 Further Analysis of the Adaptive Process

Figure 4 respectively describes the convergence of the adaptive mining process and the cross-modal semantic complement brought by it. Figure 4(a) shows the variation of the number of correlated instances under different scales  $K_R = N$ , where  $Num$  on  $y$ -axis indicates that there are  $2^{Num} \times 10^4$  instances. We can observe that our adaptive mining strategy expands the correlation set 4 times larger than the initial one. At the same time, the growth tends to be zero, which ensures the convergence of the optimization process. Figure 4(b) reports the variation of the similarity distribution among the correlated instances mined during the training process. The number of correlations increases with the training process while their similarity distribution gradually shifts to a minor direction. Meanwhile, we test the accuracy of the correlation relationships, and it rises from the initial 0.861 to 0.885 after 25 training epochs, which indicates the reliability of our adaptive learning. Therefore, we can conclude that ASSPH obtains additional joint-modal correlations independent of the original static metrics, while most initial correlations are of higher similarity. These low-similarity correlations can be seen as hard samples that contribute to the enrichment of the semantic representation of our hashing model.

## 5 CONCLUSION


This paper proposes a novel unsupervised cross-modal hashing framework, named *Adaptive Structural Similarity Preserving Hash (ASSPH)*, and conducts extensive experiments to verify its effectiveness. The framework is based on asymmetric structural semantic

**Table 2: Ablations experiments with 64-bit hashcode**


Method	NUS-WIDE		MIRFlickr-25K	
	I2T	T2I	I2T	T2I
ASSPH	0.648	0.665	0.730	0.732
ASSPH_NoAdapt	0.638	0.654	0.723	0.729
ASSPH_PairCorr	0.633	0.654	0.718	0.728
ASSPH_NoCorr	0.604	0.617	0.696	0.691
ASSPH_NoBinOpt	0.639	0.648	0.725	0.732

**Table 3: The MAP performance with different parameters**

		NUS-WIDE		MIRFlickr-25K	
		I2T	T2I	I2T	T2I
$K_S$	1000	0.642	0.656	0.727	0.726
	2000	0.648	0.665	0.730	0.732
	3000	0.639	0.654	0.726	0.728
	4000	0.627	0.655	0.722	0.719
batch size	32	0.648	0.665	0.730	0.732
	64	0.649	0.663	0.731	0.729
	128	0.652	0.667	0.727	0.735
	256	0.651	0.669	0.729	0.733
$\gamma$	0	0.642	0.658	0.722	0.721
	0.2	0.644	0.662	0.730	0.729
	0.4	0.638	0.652	0.727	0.727
	0.6	0.632	0.637	0.719	0.722
	0.8	0.618	0.623	0.705	0.710
	1	0.603	0.619	0.668	0.662



**Figure 3: The MAP performance with different  $K_R = N$  on NUS-WIDE and MIRFlickr-25K datasets (64 bit hash code).**



**Figure 4: Descriptions of the adaptive correlation mining on MIRFlickr-25K with 64 bit hash code.**

preserving with additional adaptive correlation expansion and constraints to learn the joint-semantic relationship adaptively. Combining the two parts is also a good solution for avoiding training collapse caused by unbalanced samples and small datasets, while mitigating the limitation from static metric reconstruction in traditional unsupervised cross-modal hashing.

preserving with additional adaptive correlation expansion and constraints to learn the joint-semantic relationship adaptively. Combining the two parts is also a good solution for avoiding training collapse caused by unbalanced samples and small datasets, while mitigating the limitation from static metric reconstruction in traditional unsupervised cross-modal hashing.

## 6 ACKNOWLEDGEMENTS

This research is supported in part by the National Natural Science Foundation of China under grant 62172107 and the National Key Research and Development Program of China under grant 2018YFB0505000.

## REFERENCES

- [1] Zhangjie Cao, Mingsheng Long, Jianmin Wang, and Philip S Yu. 2017. Hashnet: Deep learning to hash by continuation. In *Proceedings of the IEEE international conference on computer vision*. 5608–5617.
- [2] Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. 2018. Deep clustering for unsupervised learning of visual features. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 132–149.
- [3] Jianlong Chang, Lingfeng Wang, Gaofeng Meng, Shiming Xiang, and Chunhong Pan. 2017. Deep adaptive image clustering. In *Proceedings of the IEEE international conference on computer vision*. 5879–5887.
- [4] Yudong Chen, Sen Wang, Jianglin Lu, Zhi Chen, Zheng Zhang, and Zi Huang. 2021. Local Graph Convolutional Networks for Cross-Modal Hashing. In *MM '21: ACM Multimedia Conference*, Heng Tao Shen, Yueting Zhuang, John R. Smith, Yang Yang, Pablo Cesar, Florian Metzke, and Balakrishnan Prabhakaran (Eds.). 1921–1928. <https://doi.org/10.1145/3474085.3475346>
- [5] Zhen-Duo Chen, Yongxin Wang, Hui-Qiong Li, Xin Luo, Liqiang Nie, and Xin-Shun Xu. 2019. A two-step cross-modal hashing by exploiting label correlations and preserving similarity in both steps. In *Proceedings of the 27th ACM International Conference on Multimedia*. 1694–1702.
- [6] Tat-Seng Chua, Jinhui Tang, Richang Hong, Haojie Li, Zhiping Luo, and Yantao Zheng. 2009. Nus-wide: a real-world web image database from national university of singapore. In *Proceedings of the ACM international conference on image and video retrieval*. 1–9.
- [7] Guiguang Ding, Yuchen Guo, Jile Zhou, and Yue Gao. 2016. Large-scale cross-modality search via collective matrix factorization hashing. *IEEE Transactions on Image Processing* 25, 11 (2016), 5427–5440.
- [8] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre H. Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Daniel Guo, Mohammad Gheshlaghi Azar, et al. 2020. Bootstrap your own latent: A new approach to self-supervised learning. *arXiv preprint arXiv:2006.07733* (2020).
- [9] Jiuxiang Gu, Jianfei Cai, Shafiq R Joty, Li Niu, and Gang Wang. 2018. Look, imagine and match: Improving textual-visual cross-modal retrieval with generative models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 7181–7189.
- [10] David R Hardoon, Sandor Szedmak, and John Shawe-Taylor. 2004. Canonical correlation analysis: An overview with application to learning methods. *Neural computation* 16, 12 (2004), 2639–2664.
- [11] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. 2020. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 9729–9738.
- [12] Li He, Xing Xu, Huimin Lu, Yang Yang, Fumin Shen, and Heng Tao Shen. 2017. Unsupervised cross-modal retrieval through adversarial learning. In *2017 IEEE International Conference on Multimedia and Expo (ICME)*. IEEE, 1153–1158.
- [13] Di Hu, Feiping Nie, and Xuelong Li. 2018. Deep binary reconstruction for cross-modal hashing. *IEEE Transactions on Multimedia* 21, 4 (2018), 973–985.
- [14] Hengtong Hu, Lingxi Xie, Richang Hong, and Qi Tian. 2020. Creating something from nothing: Unsupervised knowledge distillation for cross-modal hashing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 3123–3132.
- [15] Qinghao Hu, Jiuxiang Wu, Jian Cheng, Lifang Wu, and Hanqing Lu. 2017. Pseudo label based unsupervised visual discriminative hashing for image retrieval. In *Proceedings of the 25th ACM international conference on Multimedia*. 1584–1590.
- [16] Long-Kai Huang, Jianda Chen, and Sinno Jialin Pan. 2019. Accelerate learning of deep hashing with gradient attention. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 5271–5280.
- [17] Mark J Huiskes and Michael S Lew. 2008. The mir flickr retrieval evaluation. In *Proceedings of the 1st ACM international conference on Multimedia information retrieval*. 39–43.
- [18] Qing-Yuan Jiang and Wu-Jun Li. 2017. Deep cross-modal hashing. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 3232–3240.
- [19] Dong-Hyun Lee et al. 2013. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on challenges in representation learning, ICML*, Vol. 3.
- [20] Chao Li, Cheng Deng, Ning Li, Wei Liu, Xinbo Gao, and Dacheng Tao. 2018. Self-supervised adversarial hashing networks for cross-modal retrieval. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 4242–4251.
- [21] Chao Li, Cheng Deng, Lei Wang, De Xie, and Xianglong Liu. 2019. Coupled cyclegan: Unsupervised hashing network for cross-modal retrieval. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 176–183.
- [22] Liang Li, Baihua Zheng, and Weiwei Sun. 2022. Adaptive Structural Similarity Preserving for Unsupervised Cross Modal Hashing. [arXiv:2207.04214 \[cs.LG\]](https://arxiv.org/abs/2207.04214)
- [23] Zijia Lin, Guiguang Ding, Mingqing Hu, and Jianmin Wang. 2015. Semantics-preserving hashing for cross-view retrieval. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 3864–3872.
- [24] Song Liu, Shengsheng Qian, Yang Guan, Jiawei Zhan, and Long Ying. 2020. Joint-modal Distribution-based Similarity Hashing for Large-scale Unsupervised Deep Cross-modal Retrieval. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1379–1388.
- [25] Xin Liu, Zhikai Hu, Haibin Ling, and Yiu-ming Cheung. 2019. Mtfh: A matrix tri-factorization hashing framework for efficient cross-modal retrieval. *IEEE transactions on pattern analysis and machine intelligence* (2019).
- [26] Xingbo Liu, Xiushan Nie, Wenjun Zeng, Chaoran Cui, Lei Zhu, and Yilong Yin. 2018. Fast discrete cross-modal hashing with regressing from semantic labels. In *Proceedings of the 26th ACM international conference on Multimedia*. 1662–1669.
- [27] Mingsheng Long, Yue Cao, Jianmin Wang, and Philip S Yu. 2016. Composite correlation quantization for efficient multimodal retrieval. In *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval*. 579–588.
- [28] Xu Lu, Lei Zhu, Li Liu, Liqiang Nie, and Huaxiang Zhang. 2021. Graph Convolutional Multi-modal Hashing for Flexible Multimedia Retrieval. In *MM '21: ACM Multimedia Conference*, Heng Tao Shen, Yueting Zhuang, John R. Smith, Yang Yang, Pablo Cesar, Florian Metzke, and Balakrishnan Prabhakaran (Eds.). 1414–1422. <https://doi.org/10.1145/3474085.3475598>
- [29] Min Meng, Haitao Wang, Jun Yu, Hui Chen, and Jigang Wu. 2020. Asymmetric Supervised Consistent and Specific Hashing for Cross-Modal Retrieval. *IEEE Transactions on Image Processing* 30 (2020), 986–1000.
- [30] Mehdi Mirza and Simon Osindero. 2014. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784* (2014).
- [31] Yuming Shen, Jie Qin, Jiabin Chen, Mengyang Yu, Li Liu, Fan Zhu, Fumin Shen, and Ling Shao. 2020. Auto-Encoding Twin-Bottleneck Hashing. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2815–2824. <https://doi.org/10.1109/CVPR42600.2020.00289>
- [32] Shupeng Su, Zhisheng Zhong, and Chao Zhang. 2019. Deep joint-semantics reconstructing hashing for large-scale unsupervised cross-modal retrieval. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 3027–3035.
- [33] Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. 2019. Vi-bert: Pre-training of generic visual-linguistic representations. *arXiv preprint arXiv:1908.08530* (2019).
- [34] Bokun Wang, Yang Yang, Xing Xu, Alan Hanjalic, and Heng Tao Shen. 2017. Adversarial cross-modal retrieval. In *Proceedings of the 25th ACM international conference on Multimedia*. 154–162.
- [35] Jun Wang, Wei Liu, Sanjiv Kumar, and Shih-Fu Chang. 2015. Learning to hash for indexing big data—A survey. *Proc. IEEE* 104, 1 (2015), 34–57.
- [36] Jingdong Wang, Ting Zhang, Nicu Sebe, Heng Tao Shen, et al. 2017. A survey on learning to hash. *IEEE transactions on pattern analysis and machine intelligence* 40, 4 (2017), 769–790.
- [37] Lu Wang, Jie Yang, Masoumeh Zareapoor, and Zhonglong Zheng. 2021. Cluster-wise unsupervised hashing for cross-modal similarity search. *Pattern Recognition* 111 (2021), 107732.
- [38] Weiwei Wang, Yuming Shen, Haofeng Zhang, Yazhou Yao, and Li Liu. 2020. Set and rebase: determining the semantic graph connectivity for unsupervised cross modal hashing. In *International Joint Conference on Artificial Intelligence*. 853–859.
- [39] Yongxin Wang, Xin Luo, and Xin-Shun Xu. 2020. Label Embedding Online Hashing for Cross-Modal Retrieval. In *Proceedings of the 28th ACM International Conference on Multimedia*. 871–879.
- [40] Yair Weiss, Antonio Torralba, Robert Fergus, et al. 2008. Spectral hashing. In *Nips*, Vol. 1. Citeseer, 4.
- [41] Gengshen Wu, Zijia Lin, Jungong Han, Li Liu, Guiguang Ding, Baochang Zhang, and Jialie Shen. 2018. Unsupervised Deep Hashing via Binary Latent Factor Models for Large-scale Cross-modal Retrieval. In *IJCAI*. 2854–2860.
- [42] Hao Wu, Jiayuan Mao, Yufeng Zhang, Yuning Jiang, Lei Li, Weiwei Sun, and Wei-Ying Ma. 2019. Unified visual-semantic embeddings: Bridging vision and language with structured meaning representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 6609–6618.
- [43] Dejie Yang, Dayan Wu, Wanqian Zhang, Haisu Zhang, Bo Li, and Weiping Wang. 2020. Deep Semantic-Alignment Hashing for Unsupervised Cross-Modal Retrieval. In *Proceedings of the 2020 International Conference on Multimedia Retrieval*. 44–52.
- [44] Yang Yang, Chubing Zhang, Yi-Chu Xu, Dianhai Yu, De-Chuan Zhan, and Jian Yang. 2021. Rethinking Label-Wise Cross-Modal Retrieval from A Semantic Sharing Perspective. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence*, Zhi-Hua Zhou (Ed.). 3300–3306. <https://doi.org/10.24963/ijcai.2021/454>
- [45] Jun Yu, Hao Zhou, Yibing Zhan, and Dacheng Tao. 2021. Deep Graph-neighbor Coherence Preserving Network for Unsupervised Cross-modal Hashing. (2021).
- [46] Dongqing Zhang and Wu-Jun Li. 2014. Large-scale supervised multimodal hashing with semantic correlation maximization. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 28.
- [47] Jian Zhang, Yuxin Peng, and Mingkuan Yuan. 2018. Unsupervised generative adversarial cross-modal hashing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 32.

- [48] Peng-Fei Zhang, Yadan Luo, Zi Huang, Xin-Shun Xu, and Jingkuan Song. 2021. High-order nonlocal Hashing for unsupervised cross-modal retrieval. *World Wide Web* 24, 2 (2021), 563–583.
- [49] Zheng Zhang, Zhihui Lai, Zi Huang, Wai Keung Wong, Guo-Sen Xie, Li Liu, and Ling Shao. 2019. Scalable supervised asymmetric hashing with semantic and latent factor embedding. *IEEE Transactions on Image Processing* 28, 10 (2019), 4803–4818.
- [50] Luowei Zhou, Hamid Palangi, Lei Zhang, Houdong Hu, Jason Corso, and Jianfeng Gao. 2020. Unified vision-language pre-training for image captioning and vqa. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 13041–13049.
- [51] Lei Zhu, Jialie Shen, Liang Xie, and Zhiyong Cheng. 2016. Unsupervised visual hashing with semantic assistant for content-based image retrieval. *IEEE Transactions on Knowledge and Data Engineering* 29, 2 (2016), 472–486.