

# Dual-Stream Knowledge-Preserving Hashing for Unsupervised Video Retrieval

Pandeng Li<sup>1</sup>, Hongtao Xie<sup>1</sup>\*, Jiannan Ge<sup>1</sup>, Lei Zhang<sup>2</sup>, Shaobo Min<sup>3</sup>,  
and Yongdong Zhang<sup>1</sup>

<sup>1</sup> University of Science and Technology of China

<sup>2</sup> Kuaishou Technology

<sup>3</sup> Tencent Data Platform

{lpd,gejn}@mail.ustc.edu.cn, {htxie,zhyd73}@ustc.edu.cn,  
zhanglei06@kuaishou.com, bobmin@tencent.com

**Abstract.** Unsupervised video hashing usually optimizes binary codes by learning to reconstruct input videos. Such reconstruction constraint spends much effort on frame-level temporal context changes without focusing on video-level global semantics that are more useful for retrieval. Hence, we address this problem by decomposing video information into reconstruction-dependent and semantic-dependent information, which disentangles the semantic extraction from reconstruction constraint. Specifically, we first design a simple dual-stream structure, including a temporal layer and a hash layer. Then, with the help of semantic similarity knowledge obtained from self-supervision, the hash layer learns to capture information for semantic retrieval, while the temporal layer learns to capture the information for reconstruction. In this way, the model naturally preserves the disentangled semantics into binary codes. Validated by comprehensive experiments, our method consistently outperforms the state-of-the-arts on three video benchmarks.


**Keywords:** Unsupervised Video Retrieval; Dual-Stream Hashing;

## 1 Introduction

In view of the explosive growth of informative media (*i.e.*, videos) [57,24,44,46,29], the efficient large-scale retrieval system [3,19,52,11,53] has become an urgent requirement in the real world. Video retrieval system needs to understand the semantic similarity information implicit in videos [55], which can be found by comparing the real-valued features in the last layer of deep networks. Unfortunately, these massive amounts of features take up large storage space [6] and seriously affect the retrieval speed. As a key building block of search algorithms, hashing [12], can alleviate the above issue by compressing high dimensional features into compact binary codes. However, the abundant content and temporal dynamics of videos make it difficult for binary codes to preserve the similarity structure of the real-valued feature space [23,45]. Besides, compared to image

---

\* H. Xie is the corresponding author.



**Fig. 1.** (a) Existing methods usually optimize binary codes by using an encoder-decoder architecture to reconstruct the visual information of input frames. (b) Our method utilizes latent codes in the student model to model temporal changes required for the reconstruction task, thereby allowing binary codes to focus on global semantics.



**Fig. 2.** Essential and superfluous information for the semantic video retrieval task.

datasets, the manual annotation and pre-training costs of standard large-scale video data are very high [10]. Therefore, unsupervised video hashing has intrigued many researchers in practice [37,21,22,56].

As shown in Fig. 1 (a), existing unsupervised video hashing methods usually optimize binary codes by using an encoder-decoder architecture to reconstruct the visual information of input frames. For example, Zhang *et al.* [56] employ an encoder-decoder Recurrent Neural Networks (RNNs) [33] to capture the temporal nature of videos for binary codes. Later, Li *et al.* [20] introduce Variational Auto-Encoders (VAE) [17] to learn a probabilistic latent code of video variations. However, these binary codes are forced to independently satisfy the goal of video reconstruction, which may be sub-optimal for semantic retrieval due to the heterogeneity of two tasks [13] (*i.e.*, the retrieval and reconstruction tasks).

Specifically, given raw videos as input, existing models tend to compress the information that is essential to reconstruction but may be superfluous for similarity search. This argument can be proved in Information Bottleneck (IB) [42] from an information-theoretic perspective. IB models the information flow [50] from input  $x$  to the target  $\hat{x}$  through latent variable  $b$  (*e.g.*, binary codes), where the optimal  $b$  should contain the minimal sufficient information to predict  $\hat{x}$  but discards all superfluous information in  $x$  that is irrelevant for  $\hat{x}$ . This provides an optimization principle that maximizes the mutual information  $I(b; \hat{x})$  between the latent variable and the target, and simultaneously constrains  $I(x; b)$  small. In the above existing hashing methods, maximizing  $I(b; \hat{x})$  corresponds to minimizing the reconstruction error. The reconstruction-essential information [31] may be the sequence of actions, constant changes, etc., which involves more fine-grained temporal understanding. However, as shown in Fig. 2, binary codes require more attention to global semantic concepts like “biking” or “cat” for

ranking videos. Because, the retrieval goal is not to retain all information of the original video data, but to preserve the discriminative similarity information.

Based on the above discussion, we propose a novel Dual-stream Knowledge-Preserving Hashing (DKPH) framework to obtain semantic binary codes by decomposing video information into semantic-dependent and reconstruction-dependent information. As shown in Fig. 3, DKPH fully releases the potential of semantic learning via teacher-student optimization: (1) the student model designs a simple but effective dual-stream structure to disentangle the semantic extraction from reconstruction constraint on a single binary code; (2) the teacher model refines the semantic similarity knowledge to further guide the meaningful information decomposition in the student model.

More concretely, the dual-stream structure contains a parallel temporal layer and hash layer. The temporal layer tries to capture reconstruction-dependent information by learning dynamic frame-level features, while the hash layer focuses on the semantic-dependent part from a global video-level perspective. To achieve the above goal, a teacher model is trained in a self-supervised manner to construct a Gaussian-adaptive similarity graph, which captures the inherent similarity relations between samples. This relation knowledge is preserved into the student hash layer to generate semantic-dependent discriminative binary codes.

**Contributions.** (1) We propose a novel framework, DKPH, to fully release the potential of semantic learning on binary codes and may shed critical insights for the retrieval community. To our best knowledge, our method is the first work that explores the task heterogeneity in video hashing. (2) A simple but effective dual-stream structure is developed to decompose video information, which can generate semantic-dependent discriminative binary codes by preserving the semantic similarity knowledge from the proposed Gaussian-adaptive similarity graph. (3) Extensive experiments demonstrate that DKPH outperforms state-of-the-art video hashing models on FCVID, ActivityNet and YFCC datasets.

## 2 Related Work

**Unsupervised hashing.** Unsupervised hashing aims to learn hash functions that compress data points into binary codes, which are built on training data without manual annotations. Iterative quantization (ITQ) [12] is a traditional representative method that directly explores the minimum quantization error by learning an optimal rotation of principal component directions. However, non-deep image hashing methods only seek a single linear projection, resulting in poor generalization. Then, Deep Hashing (DH) [9] uses a deep neural network to learn binary codes via multiple hierarchical non-linear transformations.

Due to the explosive growth of short videos, some works [36,21] also focus on video hashing. Multiple Feature Hashing (MFH) [36] mines local structural information while ignoring inter-frame temporal consistency [51]. Later, a series of methods based on encoder-decoder structure have become mainstream methods for video hashing. For example, Self-Supervised Temporal Hashing (SSTH) [56] employs an encoder-decoder RNNs to capture the temporal nature of videos. Li *et*

*al.* [18] jointly model static visual appearance and temporal pattern into binary codes via two special reconstruction losses. Unsupervised Deep Video Hashing (UDVH) [45] emphasizes balancing dimensional variation for each binary representation. Self Supervised Video Hashing (SSVH) [37] attempts more powerful Bi-LSTM to model more granular inter-frame dependencies. Despite a similar network architecture to SSVH, Neighborhood Preserving Hashing (NPH) [21] encodes the neighborhood-dependent video content as a binary code. Bidirectional Transformer Hashing (BTH) [22] introduces the BERT architecture [7] in NLP to explore inter-frame correlations, and achieves excellent results. However, these video hashing methods fail to consider the heterogeneity between reconstruction and retrieval tasks for optimizing binary codes. Recently, Shen *et al.* [34] propose twin bottlenecks to extract continuous features, but the similarity optimization process for binary codes is still implicit and heavily depends on the reconstruction effects. Besides, more efficient sample relations have not been explored, which affects the semantic discriminative of binary codes.

**Knowledge distillation.** [14,32] first propose to transfer knowledge from teacher models to student models through the soft outputs or intermediate layer features. Recently, Knowledge Distillation (KD) is extended to training deep networks in generations and [1,25] find that KD can refine ground truth labels. In unsupervised video hashing, to preserve and distill the semantic knowledge, we refine pre-trained CNN features to visual embeddings in the teacher model, which can further construct an efficient similarity graph for training student model.

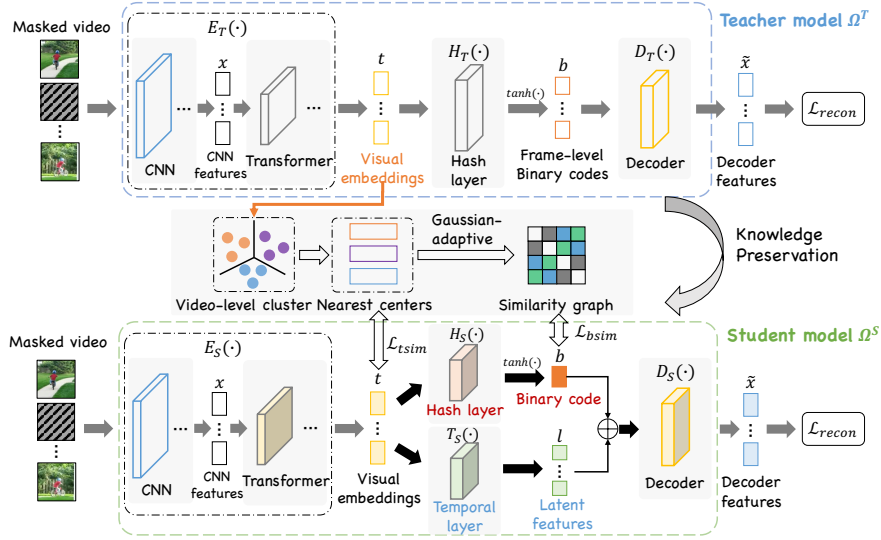
### 3 Method

#### 3.1 Problem Definition

We introduce some notations and the problem definition of unsupervised video hashing. Generally, learning hash functions is considered in an unsupervised manner from a training set of  $N$  video data points  $\mathcal{V} = \{\mathbf{v}_i\}_{i=1}^N \in \mathbb{R}^{N \times M \times D}$ , where each  $\mathbf{v}_i = [\mathbf{x}_1, \dots, \mathbf{x}_M] \in \mathbb{R}^{M \times D}$  is a CNN feature set,  $M$  is the number of frames, and  $D$  is the feature dimension of each frame. DKPH aims to learn nonlinear hash functions based on transformer blocks that map each video data point  $\mathbf{v}_i$  into a  $K$ -dimensional Hamming space  $\mathbf{b}_i \in \{-1, 1\}^K$ , which needs to keep relative semantic similarity between videos.

#### 3.2 Network Overview

DKPH consists of a teacher model  $\Omega^T$  and a student model  $\Omega^S$ . As shown in Fig. 3,  $\Omega^T$  is a common encoder-decoder architecture that can exchange inter-frame information through transformers to obtain long-term semantic knowledge.  $\Omega^S$  is a dual-stream encoder-decoder architecture that can disentangle the semantic extraction and reconstruction constraint on a single binary code to better capture the semantic information transmitted by  $\Omega^T$ . In this section, we introduce three key sub-networks: transformer encoder, hash layer and temporal layer, where the structure of transformer encoder is the same in  $\Omega^T$  and  $\Omega^S$ .



**Fig. 3.** The proposed DKPH framework which involves (1) training teacher model  $\Omega^T$  and student model  $\Omega^S$  in an unsupervised manner, (2) distilling semantic knowledge from  $\Omega^T$  to guide the information decomposition in  $\Omega^S$ .

**Transformer encoder.** To model long-term semantic correlation in videos, we first employ transformer blocks to handle the pre-processing CNN frame features. Each transformer encoder block has a multi-head self-attention and a feed-forward layer. Different from splitting images into several tokens in ViT [8], we treat frame features  $[x_i^1, \dots, x_i^M]$  as token units, which contain rich visual content information. Besides, to learn the ordering information of each frame inside the original video, we follow the standard procedure in ViT by adding trainable positional encoding embeddings  $\mathbf{E}_{pos}$ . Thus, the input video matrix  $\mathbf{X}_i$  is defined as follows:

$$\mathbf{X}_i = [x_i^1, \dots, x_i^M] + \mathbf{E}_{pos}. \quad (1)$$

Given the input matrix  $\mathbf{X}_i$ , we calculate queries  $\mathbf{Q}_i$ , keys  $\mathbf{K}_i$  and values  $\mathbf{V}_i$  as follows:  $\mathbf{Q}_i = \mathbf{X}_i \mathbf{W}_i^Q$ ,  $\mathbf{K}_i = \mathbf{X}_i \mathbf{W}_i^K$ ,  $\mathbf{V}_i = \mathbf{X}_i \mathbf{W}_i^V$ , where  $\mathbf{W}_i^Q$ ,  $\mathbf{W}_i^K$  and  $\mathbf{W}_i^V$  are linear projections with an output of  $d$  dimensions. Then the self-attention outputs can be calculated by

$$\text{Att}(\mathbf{Q}_i, \mathbf{K}_i, \mathbf{V}_i) = \text{softmax}\left(\frac{\mathbf{Q}_i \mathbf{K}_i^T}{\sqrt{d}}\right) \mathbf{V}_i^T. \quad (2)$$

Finally, these frame token units undergo multiple informative interactions, which are transformed into a sequence of visual embeddings  $[t_i^1, \dots, t_i^M]$ .

**Hash layer.** As shown in the teacher model  $\Omega^T$  of Fig. 3, the intuitive approach [22] is to directly reduce the visual embedding dimension of each frame through linear mapping  $[\hat{t}_i^1, \dots, \hat{t}_i^M] = H_T([t_i^1, \dots, t_i^M]) \in \mathbb{R}^{M \times K}$ , and then

binarize them to obtain frame-level binary codes  $[\mathbf{b}_i^1, \dots, \mathbf{b}_i^M] \in \{-1, 1\}^{M \times K}$ . However, according to the settings of existing methods [22] in the testing phase,  $\Omega^T$  needs to average  $[\mathbf{b}_i^1, \dots, \mathbf{b}_i^M]$  to obtain a real-valued code, which is binarized to video-level binary code for retrieval. This leads to two issues: (1) there is a quantization error between real-valued codes and binary codes, resulting in a sub-optimal solution; (2) when  $+1$  and  $-1$  numbers of frame binary codes are the same,  $\{-1, 0, 1\}^K$  may be generated, which violates the principle of hashing.

Therefore, we directly concatenate the frame visual features  $[\mathbf{t}_i^1, \dots, \mathbf{t}_i^M]$  of the video from a global perspective in  $\Omega^S$ , and then extract a real-valued code through the Fully Connected (FC) layer:

$$\hat{\mathbf{t}}_i = H_S(\text{Concat}([\mathbf{t}_i^1, \dots, \mathbf{t}_i^M])) \in \mathbb{R}^K. \quad (3)$$

Finally, we can obtain a K-bit binary code:

$$\mathbf{b}_i = \text{sgn}(\tanh(\hat{\mathbf{t}}_i)) \in \{-1, 1\}^K. \quad (4)$$

Besides, to avoid the discrete optimization problem [5], we follow [15] for back-propagating gradients. In this way, the encoder-decoder methods [56, 37] can compress the visual information as much as possible. However, to meet the goal of video reconstruction, the compression process may contain lots of retrieval-supervising information, which affects the discriminativeness of binary codes.

**Temporal layer.** To alleviate the task heterogeneity problem, a simple but effective dual-stream structure is introduced to decompose video information in  $\Omega^S$ . Specifically, we design a temporal layer  $T_S$  parallel to the hash layer  $H_S$  in the dual-stream structure.  $T_S$  directly reduces the dimension of frame visual features  $[\mathbf{t}_i^1, \dots, \mathbf{t}_i^M]$  to obtain frame-level latent features via FC:

$$[\mathbf{l}_i^1, \dots, \mathbf{l}_i^M] = T_S([\mathbf{t}_i^1, \dots, \mathbf{t}_i^M]) \in \mathbb{R}^{M \times K}. \quad (5)$$

The temporal layer attempts to model complex information such as dynamic temporal changes via the reconstruction constraint, while for the hash layer, we will design similarity constraints to guide the flow of semantic information. Next, we will introduce how to perform dual-stream reconstruction and similarity knowledge preservation respectively.

### 3.3 Dual-stream Reconstruction Learning

Existing video hashing works usually design the reconstruction task to compress visual information into binary codes. Inspired by masked language modeling in BERT [7, 43],  $\Omega^T$  [22] exploits the visual cloze task to optimize transformer blocks and capture inter-frame correlations, which randomly masks the input frame features as tokens and reconstructs the masked tokens in the decoder. In this way, frame-level binary codes in  $\Omega^T$  can retain all the essential information for reconstruction, rather than retrieval.

To avoid this issue in  $\Omega^S$ , we first mix  $[\mathbf{l}_i^1, \dots, \mathbf{l}_i^M]$  and  $\mathbf{b}_i$  derived in Eq. 4 and Eq. 5, and then leverage the FC layer to reconstruct:

$$[\tilde{\mathbf{x}}_i^1, \dots, \tilde{\mathbf{x}}_i^M] = D_S([\mathbf{l}_i^1 + \mathbf{b}_i, \dots, \mathbf{l}_i^M + \mathbf{b}_i]) \in \mathbb{R}^{M \times D}. \quad (6)$$

Then, we can use the mean square error loss to measure the difference between CNN features  $[\mathbf{x}_i^1, \dots, \mathbf{x}_i^M]$  and decoder features  $[\tilde{\mathbf{x}}_i^1, \dots, \tilde{\mathbf{x}}_i^M]$  in  $\Omega^S$ :

$$\mathcal{L}_{recon} = \frac{1}{DNM} \sum_{i=1}^N \sum_{m=1}^M \|\mathbf{x}_i^m - \tilde{\mathbf{x}}_i^m\|_2^2. \quad (7)$$

Please note that, in  $\Omega^T$ , the definition of  $\mathcal{L}_{recon}$  is the same as Eq. 7, but teacher decoder features can only be generated from frame-level binary codes  $[\mathbf{b}_i^1, \dots, \mathbf{b}_i^M]$ , where the code length is fixed to 128 in experiments.


### 3.4 Semantic Knowledge Preservation

Using the frame-level reconstruction task alone does not make the two layers perform the desired role, so we further guide  $H_S$  to learn video-level semantic similarity information. For unsupervised learning, some image hashing works [49] prove that neighborhood structures learned from original CNN features can capture the similarity relations between samples. However, this strategy is time-consuming due to building a similarity graph for all samples directly, and has lots of noisy predictions, which confuses the learning of hash functions. Benefiting from the teacher-student distillation framework in Fig. 3, we construct a Gaussian-adaptive similarity graph from  $\Omega^T$  that captures the inherent semantic relations by estimating positives and hard negatives of training videos. These relations can guide  $H_S$  to generate discriminative binary codes and maintain the neighborhood structure in Hamming space.

Specifically, we first warm up  $\Omega^T$  with the reconstruction task  $\mathcal{L}_{recon}$ , and exploit the transformer visual embeddings  $[\mathbf{t}_i^1, \dots, \mathbf{t}_i^M]$  instead of the high dimensional CNN frame features to mine similarity relations. Although visual embeddings may contain redundant information due to  $\mathcal{L}_{recon}$ , they model inter-frame correlations compared to CNN features, which are helpful for mining long-term semantic concepts. To obtain video-level graph, we average  $[\mathbf{t}_i^1, \dots, \mathbf{t}_i^M]$  to video embedding  $\bar{\mathbf{t}}_i$ . Then, to solve the time-consuming problem, we follow [26] to use the neighbor graph between each video point  $\bar{\mathbf{t}}_i$  and the cluster center  $\{\mathbf{c}_i\}_{i=1}^{N_c}$  of video points to approximate similarity relations between  $\bar{\mathbf{t}}_i$  and  $\{\bar{\mathbf{t}}_i\}_{i=1}^N$ , where  $N_c$  is the number of K-means clustering center. For each  $\bar{\mathbf{t}}_i$ , we calculate  $p$  nearest centers  $\{\mathbf{c}_{ij}\}_{j=1}^p$ , and the similarity matrix  $\mathbf{Z} \in \mathbb{R}^{N \times N_c}$  is expressed as:

$$\mathbf{Z}_{ij} = \frac{\exp(-\|\bar{\mathbf{t}}_i, \mathbf{c}_{ij}\|_2/\alpha)}{\sum_{j'=1}^p \exp(-\|\bar{\mathbf{t}}_i, \mathbf{c}_{ij'}\|_2/\alpha)}, \quad (8)$$

where  $\alpha$  is a bandwidth parameter. Note that the similarity values between  $\bar{\mathbf{t}}_i$  and corresponding  $N_c - p$  non-nearest centers in  $\mathbf{Z}$  are set to 0, for simplicity. Finally, an approximate graph adjacency  $\mathbf{A} \in \mathbb{R}^{N \times N}$  is calculated as:  $\mathbf{A} = \mathbf{Z} \Lambda^{-1} \mathbf{Z}^T$ , where  $\Lambda = \text{diag}(\mathbf{Z}^T \mathbf{1}) \in \mathbb{R}^{N_c \times N_c}$ . However,  $\mathbf{A}$  may still be a noisy similarity signal, where the nearest center number  $p$  greatly affects the prediction quality. To avoid this dilemma, the existing work [22] builds multiple large adjacency



**Fig. 4.** The cumulative distribution and corresponding histogram distribution of similarity values for 5 video points in  $\mathbf{A}$  on the FCVID dataset [16], where  $p = 10$ .

matrices to vote for credible sample relations, but it requires careful parameter tuning and takes up huge storage space on large video datasets.

Different from [22], we develop the Gaussian-adaptive similarity graph inspired by SSDH [48], which requires only one matrix to estimate more efficient sample relations. Compared to SSDH, our novelty lies in building a graph based on each sample and mining hard negative samples. Specifically, we first investigate the cumulative distribution and corresponding histogram of similarity values for each video point in  $\mathbf{A}$ . For better visualization, we randomly select the similarity values corresponding to 5 video points on the FCVID dataset [16], and use kernel density estimation [2] to simulate the real distribution curve in Fig. 4. Observing the cumulative distribution shows that the similarity values between most graph nodes are relatively small, while the histogram of the similarity value corresponding to each video point tends to a Gaussian distribution. This shows from the real data that it is very noisy to directly treat all the similarity signals in  $\mathbf{A}$  as positive samples. To ensure high confidence in the supervision signal, we adaptively obtain positive samples for each video point. For the video point  $v_i$ , the mean and standard deviation of similarity values between the nodes can be expressed as  $\mu_i$  and  $\epsilon_i$ , then we take the positive sample estimator as  $PT_i = \mu_i + \lambda_1 * \epsilon_i$ . Some metric learning works [58,28] show that hard negative samples are beneficial to model, so we add negative sample estimator  $NT_i = \mu_i - \lambda_2 * \epsilon_i$  to mine hard negative samples for training. In this way, the Gaussian-adaptive graph adjacency matrix can be expressed as:

$$\hat{\mathbf{A}}_{ij} = \begin{cases} 1, & \text{if } \mathbf{A}_{ij} \geq PT_i \\ -1, & \text{if } NT_i < \mathbf{A}_{ij} < \mu_i \\ 0, & \text{otherwise} \end{cases} \quad (9)$$

To preserve the similarity graph relations mined in  $\mathbf{\Omega}^T$  for binary codes, we design a binary structure similarity loss:

$$\mathcal{L}_{bsim} = \frac{1}{N} \sum_{\{i,j\} \in \mathcal{S}} |\hat{\mathbf{A}}_{ij}| (\hat{\mathbf{A}}_{ij} - \frac{1}{K} \mathbf{b}_i \mathbf{b}_j^T)^2, \quad (10)$$

where  $\mathcal{S}$  is the equal sampling strategy that samples positive or negative pairs with probability 0.5 based on  $\hat{\mathbf{A}}$ . Finally, we can obtain discriminative codes.



Furthermore, some works [32] argue that the middle layer of the teacher network can serve as a hint to the corresponding layer of the student network, thereby improving the effect of semantic knowledge transfer. Therefore, we consider aligning the visual embeddings between  $\Omega^S$  and  $\Omega^T$ . Inspired by [41,39], we design a visual embedding similarity loss:

$$\mathcal{L}_{tsim} = \frac{1}{N} \sum_{\{i,j\} \in \mathcal{S}} \|\bar{\mathbf{t}}_i - \mathbf{c}_{i1}\|_2^2 + \eta |\hat{\mathbf{A}}_{ij}| (1 - \hat{\mathbf{A}}_{ij}) [\|\bar{\mathbf{t}}_i - \mathbf{c}_{i1}\|_2^2 - \|\bar{\mathbf{t}}_i - \mathbf{c}_{j1}\|_2^2 + \beta]_+, \quad (11)$$

where  $\bar{\mathbf{t}}_i$  is the mean visual embedding of  $i$ -th video in  $\Omega^S$ ,  $\mathbf{c}_{i1}$  or  $\mathbf{c}_{j1}$  is 1-NN nearest center of corresponding teacher visual embedding,  $\eta$  controls the balance and  $[x]_+$  means the hinge function  $\max(0, x)$ , which makes  $\bar{\mathbf{t}}_i$  closer to  $\mathbf{c}_{i1}$  than negative pair  $\mathbf{c}_{j1}$  by a fixed margin  $\beta = 1.0$ .

### 3.5 Overall Learning

The overall training objectives of  $\Omega^T$  and  $\Omega^S$  are as follows:

$$\begin{aligned} \mathcal{L}_{\text{teacher}} &= \mathcal{L}_{\text{recon}}, \\ \mathcal{L}_{\text{student}} &= \mathcal{L}_{\text{recon}} + \gamma_1 \mathcal{L}_{\text{bsim}} + \gamma_2 \mathcal{L}_{\text{tsim}}, \end{aligned} \quad (12)$$

where  $\gamma_1$  and  $\gamma_2$  relatively weight the losses.

## 4 Experimental Results

### 4.1 Datasets, Metrics and Implementation Details

**Datasets.** We run experiments on three popular video datasets. **FCVID** [16] is a web video dataset consisting of 91,223 YouTube videos annotated into 239 categories. It covers a wide range of topics, with the majority of them being real-world events such as “biking”, “making coffee” and “yoga”. The dataset is evenly split into training and testing partitions with 45,585 and 45,600 videos. We use the testing partition as the query set and retrieval database. **ActivityNet** [4] consists of 20K YouTube videos annotated with 200 class descriptions. As the testing set labels are not publicly available, the evaluation is performed on the validation set. Following [21], we pick 9,722, 1,000 and 3,760 videos as training set, query set and retrieval database, respectively. **YFCC** [40] is a massive dataset from the Yahoo Webscope program containing 0.8M videos. We randomly select 409,788 unlabeled videos for training and 101,256 labeled videos with 80 semantic concepts [47] for testing. In these labeled videos, we sample 1000 videos as the query set and the remaining ones as retrieval database.

**Metrics.** We measure the retrieval performance with standard metrics in information retrieval, including Mean Average Precision at top-k retrieved results (MAP@k) and Precision-Recall (PR) curves.

**Implementation Details.** Our experiments are based on the Pytorch framework [30]. In the video encoding process, we uniformly sample 25 frames from

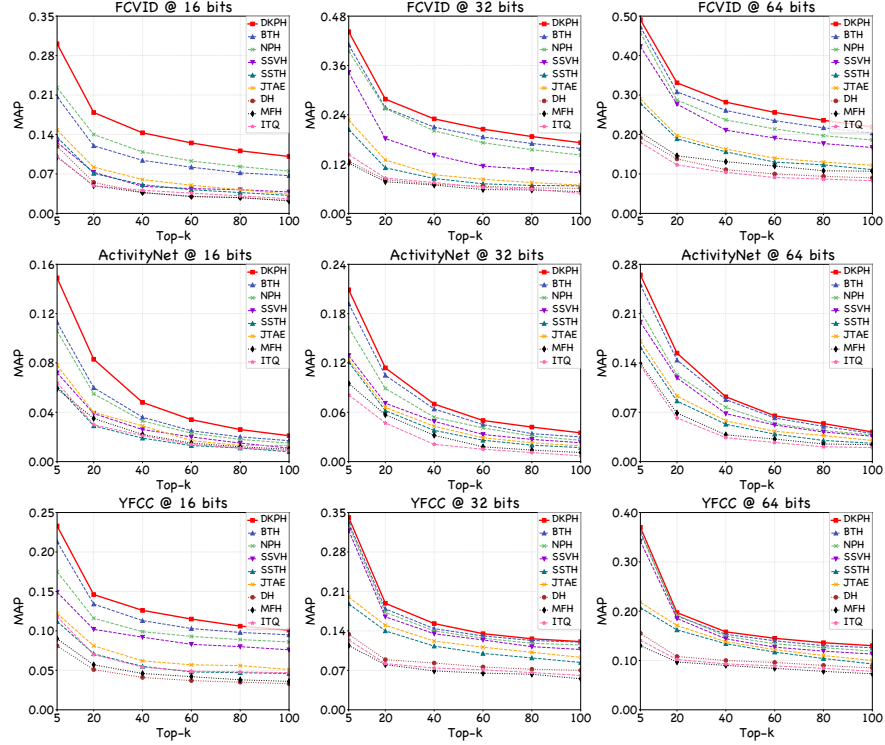


Fig. 5. Performance comparison (*w.r.t.* MAP@k) of DKPH and SOTA methods.

each video and use VGG-16 pretrained on Imagenet [35] to extract frame-wise features. To ensure fair comparison [22], we use a single transformer block with a single attention head as the transformer encoder. For the teacher model  $\Omega^T$ , we warm up 60, 300 and 200 epochs on FCVID, ActivityNet and YFCC, respectively. Considering the trade-off of visual information compression loss and inter-frame correlations, we set the dimension of the visual embeddings and binary codes to 256 and 128 in  $\Omega^T$ . In the graph construction process, the number of clustering center  $N_c$  is set as 2,000, 1,000 and 2,000 on FCVID, ActivityNet and YFCC respectively. We employ Adam optimizer to train the model with a mini-batch size of 256 and train the student model  $\Omega^S$  for 48 epochs, where the initial learning rate is  $5 \times 10^{-4}$ . The default hyper-parameters setting is:  $\lambda_1 = 2$ ,  $\lambda_2 = 1$ ,  $\eta = 0.1$ ,  $\gamma_1 = 0.11$ ,  $\gamma_2 = 0.9$ . In the testing phase, we only use the student model  $\Omega^S$ , where the lengths of binary codes are 16, 32 and 64.

## 4.2 Comparisons with State-of-the-art (SOTA) Methods

To prove the effectiveness of DKPH, we compare the retrieval performance with two image hashing methods: ITQ [12], DH [9], and six SOTA video hashing methods: MFH [36], SSTH [56], JTAE [18], SSVH [37], NPH [21], and BTH [22]. Fig. 5




Fig. 6. Performance comparison (*w.r.t.* PR curves) of DKPH and SOTA methods.

Table 1. The impact of different frame feature encoders with or without the dual-stream structure on FCVID. Some abbreviations: TF-Transformer, D-Dual-stream.

Method	16 bits				32 bits				64 bits			
	k=5	k=20	k=60	k=100	k=5	k=20	k=60	k=100	k=5	k=20	k=60	k=100
CNN [54]	0.229	0.116	0.080	0.065	0.395	0.242	0.172	0.140	0.460	0.294	0.207	0.171
CNN [54]+D	<b>0.273</b>	<b>0.152</b>	<b>0.106</b>	<b>0.084</b>	<b>0.407</b>	<b>0.252</b>	<b>0.175</b>	<b>0.147</b>	<b>0.464</b>	<b>0.308</b>	<b>0.223</b>	<b>0.189</b>
LSTM [38]	0.227	0.114	0.077	0.062	0.393	0.240	0.168	0.139	0.457	0.291	0.210	0.174
LSTM [38]+D	<b>0.272</b>	<b>0.150</b>	<b>0.104</b>	<b>0.083</b>	<b>0.404</b>	<b>0.248</b>	<b>0.172</b>	<b>0.146</b>	<b>0.462</b>	<b>0.301</b>	<b>0.224</b>	<b>0.192</b>
TF [7]	0.235	0.122	0.083	0.069	0.421	0.252	0.172	0.143	0.477	0.313	0.238	0.202
TF [7]+D	<b>0.297</b>	<b>0.174</b>	<b>0.120</b>	<b>0.097</b>	<b>0.441</b>	<b>0.275</b>	<b>0.203</b>	<b>0.171</b>	<b>0.494</b>	<b>0.331</b>	<b>0.255</b>	<b>0.228</b>

shows the MAP@K results on three datasets. Compared with SOTA methods, DKPH achieves the best results on three video datasets. Specifically, we obtain 0.7%–8.6% MAP@5 gains for various bits, which demonstrates the efficiency of DKPH. We owe the great advantage of DKPH over these two methods [22,21] to the full use of dual-stream structure and Gaussian-adaptive similarity graph. Note that the model performance gaps are larger at 16 bits, as we expected. Because the amount of information carried by the binary code is limited by the length. Therefore, in low-bit scenarios, the impact of task heterogeneity will be more serious, leading to inferior results from existing methods [22].

Furthermore, we examine DKPH with PR curves on FCVID and ActivityNet in Fig. 6. DKPH delivers higher precision than SOTA methods at the same recall rate, and improves more significantly at low recall requirements. This illustrates that the model is suitable for real-world video retrieval systems, as people tend to focus more on results with high accuracy rather than finding all similar results.


### 4.3 Ablation study

To provide further insight into DKPH, we conduct critical ablation studies.

**Analysis of the dual-stream structure with different encoders.** DKPH employs a transformer encoder and a dual-stream structure to generate binary codes. Thus, we explore the impact of different frame feature encoders (CNN [54], LSTM [38], and transformer [7]) with or without the dual-stream structure in Table 1. Specifically, based on the dual-stream structure, CNN, LSTM and transformer obtain 0.4%–6.2% MAP@5 gains at different bits on FCVID. The following advantages can be clearly observed: (1) the dual-stream structure is a **general-purpose and important design** that consistently improves three en-

**Table 2.** Contributions of different modules on FCVID.

Method	16 bits				32 bits				64 bits			
	k=5	k=20	k=60	k=100	k=5	k=20	k=60	k=100	k=5	k=20	k=60	k=100
DKPH+A	0.202	0.119	0.080	0.062	0.399	0.244	0.170	0.139	0.458	0.294	0.211	0.180
DKPH+M	0.269	0.139	0.091	0.075	0.421	0.258	0.176	0.151	0.484	0.321	0.243	0.204
DKPH+T	0.228	0.135	0.090	0.073	0.419	0.251	0.172	0.147	0.463	0.298	0.218	0.181
DKPH-D	0.235	0.122	0.083	0.069	0.421	0.252	0.172	0.143	0.477	0.313	0.238	0.202
DKPH-DR	0.276	0.159	0.111	0.092	0.435	0.267	0.189	0.158	0.487	0.326	0.245	0.218
DKPH-TS	0.158	0.088	0.067	0.051	0.311	0.123	0.078	0.058	0.345	0.169	0.118	0.103
DKPH-Lb	0.174	0.093	0.072	0.058	0.322	0.142	0.080	0.062	0.366	0.187	0.123	0.107
DKPH-Lt	0.209	0.121	0.098	0.080	0.405	0.239	0.177	0.141	0.432	0.276	0.205	0.162
<b>DKPH</b>	<b>0.297</b>	<b>0.174</b>	<b>0.120</b>	<b>0.097</b>	<b>0.441</b>	<b>0.275</b>	<b>0.203</b>	<b>0.171</b>	<b>0.494</b>	<b>0.331</b>	<b>0.255</b>	<b>0.228</b>

**Fig. 7.** The MAP@k scores with various configurations about the positive estimator factor and the negative estimator factor ( $\lambda_1, \lambda_2$ ) on FCVID.

coders, especially at low bits; (2) the transformer outperforms CNN and LSTM due to its strong ability to model long-term inter-frame correlations.

**Analysis of model components.** We compare DKPH with the following variations: (1) DKPH+A. The Gaussian-adaptive graph adjacency matrix  $\hat{\mathbf{A}}$  is replaced by  $\mathbf{A}$ ; (2) DKPH+M.  $\hat{\mathbf{A}}$  is replaced by multiple matrices [22]; (3) DKPH+T. We replace the dual-stream structure with twin bottlenecks designed for image hashing [34]; (4) DKPH-D. The dual-stream structure is removed; (5) DKPH-DR. Both the dual-stream structure and  $\mathcal{L}_{recon}$  in  $\Omega^S$  are removed; (6) DKPH-TS. We remove the teacher-student distillation strategy and only use  $\mathcal{L}_{recon}$ ; (7) DKPH-Lb. We remove the binary structure similarity loss  $\mathcal{L}_{bsim}$ ; (8) DKPH-Lt. We remove the visual embedding similarity loss  $\mathcal{L}_{tsim}$ . Table 2 shows the performance of DKPH and its variations at different bits on FCVID, and proves that each module significantly contributes to the final result.

We have the following observations. First, reasonable mining of positive and hard negative pairs helps discriminate binary codes. DKPH adaptively explores the similarity relations of each video point through the sample estimators. However, in DKPH+M, the multiple matrices strategy consumes more time and space resources, and requires careful adjustment of matrix parameters, which cannot achieve optimal results. Second, twin bottlenecks (DKPH+T) are still difficult to replace the dual-stream structure designed for video hashing. There are two reasons: (1) the mechanism of twin bottlenecks is to learn better reconstructed images to feedback binary codes, which cannot exhibit the advantages of video information decomposition; (2) twin bottlenecks generate frame-level binary codes, resulting in quantization errors during testing. Third, DKPH-D and

**Table 3.** Cross-dataset MAP@20 results when training on FCVID and test on YFCC at 64 bits. The blue number indicates the performance drop compared with training and testing both on YFCC (black number).

Method	SSTH [56]	SSVH [37]	NPH [21]	BTH [22]	DKPH
MAP@20	0.155 ( -6.3% )	0.173 ( -7.8% )	0.180 ( -6.0% )	0.191 ( -5.7% )	0.199 ( -2.8% )

**Table 4.** The MAP@k results of latent features  $\mathbf{l}$  and binary codes  $\mathbf{b}$  at 16 bits.

Method	k=5	k=20	k=60
Only $\mathbf{l}$	0.098	0.077	0.065
Only $\mathbf{b}$	<b>0.297</b>	<b>0.174</b>	<b>0.120</b>

**Table 5.** The reconstruction errors in a category.

Method	Error
DKPH	<b>0.4849</b>
DKPH <sub>f</sub>	0.9527

**Table 6.** The reconstruction errors in test set.

Method	Error
DKPH	<b>0.5586</b>
Remove $\mathbf{b}$	0.5615
Remove $\mathbf{l}$	0.9769


DKPH-DR explore the effects of task heterogeneity, which suggests a conflict between  $\mathcal{L}_{recon}$  and similarity learning in existing methods [37,22]. We decouple the tasks, which allows binary codes to retain useful information and avoids the conflict. Fourth, results in DKPH-Lb yield an excessive drop due to the lack of similarity guidance, where  $\mathcal{L}_{bsim}$  is the core loss of information decomposition.

**Hyperparameter analysis.** We investigate various configurations about the positive and negative estimator factors  $(\lambda_1, \lambda_2)$ , as shown in Fig. 7. From this experiment, we find that as  $\lambda_1$  grows (*i.e.*,  $PT_i$  grows), the performance increases at first and reaches the best results, then decreases as a whole. A small  $PT_i$  may cause the model to be trained on more noisy signals, , while a large  $PT_i$  may not fully exploit the underlying positive similarity relations. Moreover,  $\lambda_2$  has a greater impact on model performance than  $\lambda_1$ , reflecting the vital contribution of hard negative samples to the model.

#### 4.4 Further Analysis

**Cross-dataset evaluation comparisons.** To investigate the generalization of DKPH for cross-dataset retrieval, we train various methods on FCVID and test on YFCC in Table 3, which shows MAP@20 results for cross-dataset retrieval at 64 bits. DKPH can not only achieve SOTA in the single-dataset setting, but also the performance drop (-2.8%) is the lowest in the cross-dataset setting. This may be because binary codes focus more on semantic concepts rather than the underlying reconstruction information, which ensures good transferability and generalization of DKPH when retrieving unknown datasets.

**Information decomposition analysis.** Table 4 shows the MAP@k results of latent features  $\mathbf{l}$  and binary codes  $\mathbf{b}$  at 16 bits on FCVID. Results of Only  $\mathbf{l}$  are much lower than those of Only  $\mathbf{b}$ , which indicate that  $\mathbf{l}$  may not have enough semantics to support the retrieval task. Next, we examine the effect of dual-stream features for reconstruction at 16 bits on FCVID. In Table 5, we randomly input a category of test videos and then calculate the mean square error between  $[\mathbf{x}_i^1, \dots, \mathbf{x}_i^M]$  and  $[\tilde{\mathbf{x}}_i^1, \dots, \tilde{\mathbf{x}}_i^M]$ . When we replace  $\mathbf{l}$  with fixed values (*i.e.*, the mean of latent features), the reconstruction error in DKPH<sub>f</sub>



**Fig. 8.** t-SNE visualizations [27] of BTH and DKPH. Videos are randomly sampled on FCVID database, and samples with different labels are marked with different colors.

increases by 96.5%. In Table 6, we directly remove  $\mathbf{b}$  or  $\mathbf{l}$  for reconstruction and calculate errors in all test videos. Removing  $\mathbf{b}$ , the error increases by 0.52%, while Removing  $\mathbf{l}$  increases the error by 74.9%. Table 5 and 6 prove that  $\mathbf{l}$ , rather than  $\mathbf{b}$ , contains sufficient essential information (dynamic changes) for reconstruction. **Qualitative results.** Fig. 8 shows the t-SNE visualization [27] of binary codes learned by BTH and DKPH. To facilitate the observation, we randomly sample 8 categories of videos twice on 16 bits and 64 bits, respectively, to obtain binary codes. At 16 bits, there is a clear distinction between most categories in our model. In particular, t-SNE embeddings of DKPH in some categories (*e.g.*, 0, 2, 6) can be mapped onto a small circle. This proves that DKPH pays more attention to the learning of global semantics and binary codes of a category are almost very close in Hamming space, so the phenomenon of t-SNE embedding aggregation occurs. At 64 bits, t-SNE embeddings of our model in different categories are well separated, which proves the good discriminativeness.

## 5 Conclusion

We propose a novel unsupervised video hashing framework, DKPH, to tackle the task heterogeneity problem. Firstly, we design the dual-stream structure to decompose video information, which disentangles the semantic extraction from reconstruction constraint. Then, a Gaussian-adaptive similarity graph is developed to explore the semantic similarity knowledge between samples. With the help of this knowledge, the hash layer in the dual-stream structure can further generate discriminative semantic binary codes. In this paper, we hope not only to present insights into the importance of information decomposition but also to facilitate future work that advances video hashing by solving design flaws rather than mostly trial and error.

**Acknowledgements.** This work is supported by the National Nature Science Foundation of China (62121002, 62022076, U1936210), the Fundamental Research Funds for the Central Universities under Grant WK3480000011, the Youth Innovation Promotion Association Chinese Academy of Sciences (Y2021122). We acknowledge the support of GPU cluster built by MCC Lab of Information Science and Technology Institution, USTC.

## References

1. Bagherinezhad, H., Horton, M., Rastegari, M., Farhadi, A.: Label refinery: Improving imagenet classification through label progression. In: AAAI (2021)
2. Botev, Z.I., Grotowski, J.F., Kroese, D.P.: Kernel density estimation via diffusion. *Ann Stat* (2010)
3. Brown, A., Xie, W., Kalogeiton, V., Zisserman, A.: Smooth-ap: Smoothing the path towards large-scale image retrieval. In: ECCV (2020)
4. Caba Heilbron, F., Escorcia, V., Ghanem, B., Carlos Niebles, J.: Activitynet: A large-scale video benchmark for human activity understanding. In: CVPR (2015)
5. Cao, Z., Long, M., Wang, J., Yu, P.S.: Hashnet: Deep learning to hash by continuation. In: ICCV (2017)
6. Cui, Q., Jiang, Q.Y., Wei, X.S., Li, W.J., Yoshie, O.: Exchnet: A unified hashing network for large-scale fine-grained image retrieval. In: ECCV (2020)
7. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. In: NAACL (2019)
8. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. In: ICLR (2021)
9. Erin Liong, V., Lu, J., Wang, G., Moulin, P., Zhou, J.: Deep hashing for compact binary codes learning. In: CVPR (2015)
10. Gabeur, V., Sun, C., Alahari, K., Schmid, C.: Multi-modal transformer for video retrieval. In: ECCV (2020)
11. Ge, J., Xie, H., Min, S., Zhang, Y.: Semantic-guided reinforced region embedding for generalized zero-shot learning. In: AAAI (2021)
12. Gong, Y., Lazebnik, S., Gordo, A., Perronnin, F.: Iterative quantization: A procrustean approach to learning binary codes for large-scale image retrieval. *TPAMI* (2012)
13. Guo, M., Haque, A., Huang, D.A., Yeung, S., Fei-Fei, L.: Dynamic task prioritization for multitask learning. In: ECCV (2018)
14. Hinton, G., Vinyals, O., Dean, J., et al.: Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531* (2015)
15. Hubara, I., Courbariaux, M., Soudry, D., El-Yaniv, R., Bengio, Y.: Binarized neural networks. *NeurIPS* (2016)
16. Jiang, Y.G., Wu, Z., Wang, J., Xue, X., Chang, S.F.: Exploiting feature and class relationships in video categorization with regularized deep neural networks. *TPAMI* (2017)
17. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. In: ICLR (2014)
18. Li, C., Yang, Y., Cao, J., Huang, Z.: Jointly modeling static visual appearance and temporal pattern for unsupervised video hashing. In: CIKM (2017)
19. Li, P., Li, Y., Xie, H., Zhang, L.: Neighborhood-adaptive structure augmented metric learning. In: AAAI (2022)
20. Li, S., Chen, Z., Li, X., Lu, J., Zhou, J.: Unsupervised variational video hashing with 1d-cnn-lstm networks. *TMM* (2019)
21. Li, S., Chen, Z., Lu, J., Li, X., Zhou, J.: Neighborhood preserving hashing for scalable video retrieval. In: ICCV (2019)
22. Li, S., Li, X., Lu, J., Zhou, J.: Self-supervised video hashing via bidirectional transformers. In: CVPR (2021)
23. Liong, V.E., Lu, J., Tan, Y.P., Zhou, J.: Deep video hashing. *TMM* (2016)

24. Liu, B., Yeung, S., Chou, E., Huang, D.A., Fei-Fei, L., Niebles, J.C.: Temporal modular networks for retrieving complex compositional activities in videos. In: ECCV (2018)
25. Liu, Q., Xie, L., Wang, H., Yuille, A.L.: Semantic-aware knowledge preservation for zero-shot sketch-based image retrieval. In: ICCV (2019)
26. Liu, W., Wang, J., Kumar, S., Chang, S.F.: Hashing with graphs. In: ICML (2011)
27. Van der Maaten, L., Hinton, G.: Visualizing data using t-sne. JMLR (2008)
28. Milbich, T., Roth, K., Bharadhwaj, H., Sinha, S., Bengio, Y., Ommer, B., Cohen, J.P.: Diva: Diverse visual feature aggregation for deep metric learning. In: ECCV (2020)
29. Min, S., Yao, H., Xie, H., Wang, C., Zha, Z.J., Zhang, Y.: Domain-aware visual bias eliminating for generalized zero-shot learning. In: CVPR (2020)
30. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al.: Pytorch: An imperative style, high-performance deep learning library. NeurIPS (2019)
31. Qiu, Z., Su, Q., Ou, Z., Yu, J., Chen, C.: Unsupervised hashing with contrastive information bottleneck. In: IJCAI (2021)
32. Romero, A., Ballas, N., Kahou, S.E., Chassang, A., Gatta, C., Bengio, Y.: Fitnets: Hints for thin deep nets. In: ICLR (2015)
33. Rumelhart, D.E., Hinton, G.E., Williams, R.J.: Learning representations by back-propagating errors. *nature* (1986)
34. Shen, Y., Qin, J., Chen, J., Yu, M., Liu, L., Zhu, F., Shen, F., Shao, L.: Auto-encoding twin-bottleneck hashing. In: CVPR (2020)
35. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. In: ICLR (2015)
36. Song, J., Yang, Y., Huang, Z., Shen, H.T., Hong, R.: Multiple feature hashing for real-time large scale near-duplicate video retrieval. In: ACM MM (2011)
37. Song, J., Zhang, H., Li, X., Gao, L., Wang, M., Hong, R.: Self-supervised video hashing with hierarchical binary auto-encoder. TIP (2018)
38. Srivastava, N., Mansimov, E., Salakhudinov, R.: Unsupervised learning of video representations using lstms. In: ICML (2015)
39. Su, S., Zhang, C., Han, K., Tian, Y.: Greedy hash: Towards fast optimization for accurate hash coding in cnn. In: NeurIPS (2018)
40. Thomee, B., Shamma, D.A., Friedland, G., Elizalde, B., Ni, K., Poland, D., Borth, D., Li, L.J.: The new data and new challenges in multimedia research. arXiv preprint arXiv:1503.01817 (2015)
41. Tian, K., Zhou, S., Guan, J.: Deepcluster: A general clustering framework based on deep learning. In: ECML (2017)
42. Tishby, N., Zaslavsky, N.: Deep learning and the information bottleneck principle. In: ITW (2015)
43. Wang, Y., Xie, H., Fang, S., Wang, J., Zhu, S., Zhang, Y.: From two to one: A new scene text recognizer with visual language modeling network. In: ICCV (2021)
44. Wang, Y., Xie, H., Zha, Z.J., Xing, M., Fu, Z., Zhang, Y.: Contournet: Taking a further step toward accurate arbitrary-shaped scene text detection. In: CVPR (2020)
45. Wu, G., Han, J., Guo, Y., Liu, L., Ding, G., Ni, Q., Shao, L.: Unsupervised deep video hashing via balanced code for large-scale video retrieval. TIP (2018)
46. Wu, W., Zhang, D., Fu, Y., Shen, C., Zhou, H., Cai, Y., Luo, P.: End-to-end video text spotting with transformer. arXiv preprint arXiv:2203.10539 (2022)
47. Xiao, J., Hays, J., Ehinger, K.A., Oliva, A., Torralba, A.: Sun database: Large-scale scene recognition from abbey to zoo. In: CVPR (2010)



48. Yang, E., Deng, C., Liu, T., Liu, W., Tao, D.: Semantic structure-based unsupervised deep hashing. In: IJCAI (2018)
49. Yang, E., Liu, T., Deng, C., Liu, W., Tao, D.: Distillhash: Unsupervised deep hashing by distilling data pairs. In: CVPR (2019)
50. Yang, K., Zhou, T., Tian, X., Tao, D., et al.: Class-disentanglement and applications in adversarial detection and defense. NeurIPS (2021)
51. Ye, G., Liu, D., Wang, J., Chang, S.F.: Large-scale video hashing via structure learning. In: ICCV (2013)
52. Yu, T., Yang, Y., Li, Y., Liu, L., Fei, H., Li, P.: Heterogeneous attention network for effective and efficient cross-modal retrieval. In: SIGIR (2021)
53. Yu, T., Yuan, J., Fang, C., Jin, H.: Product quantization network for fast image retrieval. In: ECCV (2018)
54. Yue-Hei Ng, J., Hausknecht, M., Vijayanarasimhan, S., Vinyals, O., Monga, R., Toderici, G.: Beyond short snippets: Deep networks for video classification. In: CVPR (2015)
55. Zhang, B., Hu, H., Sha, F.: Cross-modal and hierarchical modeling of video and text. In: ECCV (2018)
56. Zhang, H., Wang, M., Hong, R., Chua, T.S.: Play and rewind: Optimizing binary representations of videos by self-supervised temporal hashing. In: ACM MM (2016)
57. Zhang, X., Zhang, T., Hong, X., Cui, Z., Yang, J.: Graph wasserstein correlation analysis for movie retrieval. In: ECCV (2020)
58. Zhao, Y., Jin, Z., Qi, G.j., Lu, H., Hua, X.s.: An adversarial approach to hard triplet generation. In: ECCV (2018)