

Dual Encoding for Video Retrieval by Text

Jianfeng Dong, Xirong Li, *Member, IEEE*, Chaoxi Xu, Xun Yang, Gang Yang,
Xun Wang, *Member, IEEE*, Meng Wang, *Fellow, IEEE*,

Abstract—This paper attacks the challenging problem of video retrieval by text. In such a retrieval paradigm, an end user searches for unlabeled videos by ad-hoc queries described exclusively in the form of a natural-language sentence, with no visual example provided. Given videos as sequences of frames and queries as sequences of words, an effective sequence-to-sequence cross-modal matching is crucial. To that end, the two modalities need to be first encoded into real-valued vectors and then projected into a common space. In this paper we achieve this by proposing a *dual deep encoding* network that encodes videos and queries into powerful dense representations of their own. Our novelty is two-fold. First, different from prior art that resorts to a specific single-level encoder, the proposed network performs multi-level encoding that represents the rich content of both modalities in a coarse-to-fine fashion. Second, different from a conventional common space learning algorithm which is either concept based or latent space based, we introduce hybrid space learning which combines the high performance of the latent space and the good interpretability of the concept space. Dual encoding is conceptually simple, practically effective and end-to-end trained with hybrid space learning. Extensive experiments on four challenging video datasets show the viability of the new method.

Index Terms—Video retrieval, cross-modal representation learning, dual encoding, hybrid space learning

1 INTRODUCTION

THIS paper targets at the task of video retrieval by text, where a query is described exclusively in the form of a natural-language sentence, with no visual example attached. The task is scientifically interesting and challenging as it requires establishing proper associations between visual and linguistic information presented in the temporal order.

Retrieving *unlabeled* videos by text attracts initial attention in the form of zero-example multimedia event detection, where the goal is to retrieve video shots showing specific events such as *parking a vehicle*, *dog show* and *birthday party*, but with no training videos provided [1], [2], [3], [4], [5], [6], [7]. All these methods are concept based, representing the video content by automatically detected concepts, which are used to match with textual descriptions of a target event. An attractive property of the concept-based representation is its good interpretability [7], as each of its dimensions has explicit meanings. The concept-based tradition continues in the era of deep learning. For the NIST TRECVID AVS challenge [8], [9], a leading benchmark evaluation for video retrieval by text, we observe that the top performers in the previous years (2016–2018) are mostly concept based [10], [11], [12], [13], [14]. However, the concept-based paradigm faces intrinsic difficulties including



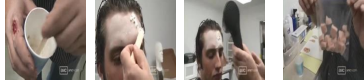

Query: a man in a black shirt is doing an experiment with magnets and batteries			
Multi-level Encoding		Rank of relevant item	Top-1 retrieved video
Query	Video		
✗	✗	113	
✓	✗	55	
✗	✓	23	
✓	✓	1	

Fig. 1. **Showcase of video retrieval by text with and without the proposed encoding.** The ✗ symbol indicates encoding by mean pooling. Numbers in the third column are the rank of the relevant video returned by retrieval models subject to specific query / video encoding strategies. The retrieval model with dual encoding successfully answers this complex query.

- J. Dong and X. Wang are with the College of Computer and Information Engineering, Zhejiang Gongshang University, Hangzhou 310035, China. E-mail: dongjf24@gmail.com
- X. Li, C. Xu and G. Yang are with the Key Lab of Data Engineering and Knowledge Engineering, Renmin University of China, and the AI and Media Computing Lab, School of Information, Renmin University of China, Beijing 100872, China. E-mail: xirong@ruc.edu.cn
- X. Yang is with the School of Computing, National University of Singapore, Singapore 37580, Singapore. E-mail: xunyang@nus.edu.sg
- M. Wang is with the School of Computer Science and Information Engineering, Hefei University of Technology, Hefei 230009, China. E-mail: wangmeng@hfut.edu.cn

Manuscript received Sep 2, 2020; revised Dec 30, 2020 and Feb 9, 2021; accepted Feb 10, 2021. (Corresponding authors: Xirong Li and Xun Wang)

how to specify a set of proper concepts, how to train good classifiers for these concepts, and more crucially how to select relevant and detectable concepts for both video and query representation [6]. These difficulties remain largely unresolved to this day.

Not surprisingly, efforts towards concept-free representation have been made. In the context of cross-modal retrieval between image and text [15], [16], Canonical Correlation Analysis (CCA) has been frequently used to linearly project both visual and textual features into a common space. From the viewpoint of neural networks, CCA es-

essentially applies a fully connected (FC) layers (with no bias term) on the visual side and another FC layer on the textual side. With the quick development of deep neural networks in both computer vision and natural language processing, the simple FC layer has now been replaced by more advanced embedding networks [17], [18], [19], [20], [21], [22]. Still, we see a common pattern among the varied solutions. That is, first encoding videos and textual queries, and then mapping them into a common space where the video-text similarities can be computed directly. Therefore, what matters for this paradigm are forms of video encoding, text encoding and common space learning.

For video encoding, a popular solution is to first extract frame features from videos by pre-trained CNN models, and then aggregate them into a video-level feature by mean pooling [17], [23], max pooling [20], [24], NetVLAD [19], RNN [25], [26] or self-attention mechanisms [21], [22]. As for text encoding, while bag-of-words remains common [7], [27], deep networks are in an increasing use. Given a specific sentence, a typical approach is to first quantize each of its word by word embedding and then aggregate word-level features into a sentence-level feature by max pooling [20], Fisher Vector [28], [29], NetVLAD [19], [24], RNN [17], [30] or Graph Convolutional Network (GCN) [22]. Instead of using one specific encoding strategy, Dong *et al.* [31] and Li *et al.* [32] utilize multiple text encoders including bag-of-words, word2vec and GRU. However, they simply use mean pooling for video encoding. In contrast to the existing works, we propose *dual multi-level encoding* for both videos and text in advance to common space learning. As exemplified in Fig. 1, the new encoding strategy is crucial for describing complex queries and video content.

Our hypothesis is that a given video/query has to be first encoded into a powerful representation of its own. We consider such a decomposition crucial as it allows us to design an encoding network that jointly exploits multiple encoding strategies including mean pooling, recurrent neural networks and convolutional networks. In our design, the output of a specific encoding block is not only used as input of a follow-up encoding block, but also re-used via skip connections to contribute to the final output. It generates new, higher-level features progressively. These features, generated at distinct levels, are powerful and complementary to each other, allowing us to obtain effective video (and text) representations by very simple concatenation.

Philosophically, our dual encoding model is linked to Allan Paivio's dual-coding theory of cognition [33]. Supported by evidence from psychological research, the dual-coding theory postulates that verbal and visual information are processed along distinct channels with separate representations in the human mind. Later, these representations are used for retrieving information previously stored in the mind. In a similar spirit, the dual encoding model stores video and textual information learned from training data in separate representations and recall them in the inference stage.

For common space learning, the state-of-the-art relies on constructing a latent space [19], [20], [21], [22], as such a space can be optimized in an end-to-end manner, and thus permits superior performance against the concept-based alternative. This, however, comes at the cost of losing interpretability. Different from the concept space, each dimension

of the latent space is not directly interpretable. Hence, what a model has truly learned is often agnostic. In order to combine the merits of the latent space and the concept space, we propose to train the dual encoding network with *hybrid space* learning. In particular, a latent space and a concept space are simultaneously learned, for better performance and better interpretability.

In sum, this paper makes the following contributions.

- We propose a novel *dual* network that encodes an input, let it be a query sentence or a video, in a similar manner. By jointly exploiting multi-level encodings, the network explicitly and progressively learns to represent global, local and temporal patterns in videos and sentences. Moreover, dual encoding is orthogonal to common space learning, allowing us to flexibly embrace state-of-the-art common space learning algorithms.
- We propose a new hybrid space learning to learn a hybrid common space for video-text similarity prediction, which inherits the high performance of the latent space and the good interpretability of the concept space.
- We conduct extensive experiments on four challenging video datasets, *i.e.*, MSR-VTT [34], TRECVID AVS 2016-2018 [8], [35], [36], VATEX [37] and MPII-MD [38]. Dual encoding, trained by hybrid space learning, is a new state-of-the-art for video retrieval by text.

The architecture of our model is based on existing components including mean feature pooling, GRU and 1D-CNN. Our novelty is the integration of these components into a dual multi-level encoding architecture for video / text representation learning. We consider the novelty of the system architecture to be more than the sum of its parts. The Dual Encoding network effectively combines these vanilla components into a powerful solution that is competitive with recent Transformer based methods [39], [40]. Also notice that while we use two pre-trained CNNs to extract frame-level features, the two features are concatenated to form a single vector before feeding it into our network. The multi-level features combined in Dual Encoding are generated by the network itself. Hence, our method essentially uses a single visual feature as input, and thus differs fundamentally from previous works that rely on the ensemble of diverse models / features.

A preliminary version of this work was published at CVPR 2019 [41]. The journal extension improves over the conference paper mainly in two aspects. Technically, we introduce hybrid space learning. Compared to the latent space learning used in [41], the new learning algorithm leads to better video retrieval performance. With the learned concept space, the model's interpretability is also improved. Experimentally, our evaluation has been substantially expanded in terms of datasets and competitive baselines.

2 RELATED WORK

2.1 Concept based Methods

Since 2016 the TRECVID starts a new challenge for video retrieval by text, known as Ad-hoc Video Search (AVS) [8]. The majority of the top ranked solutions for this challenge depend on visual concept classifiers to describe video content and linguistic rules to detect concepts in textual queries

[10], [11], [12], [13], [42]. Then the similarity between a textual query and a specific video is typically computed by concept matching. For instance, [11], [42] utilize multiple pre-trained Convolutional Neural Network (CNN) models to detect main objects and scenes in video frames. As for query representation, the authors design relatively complex linguistic rules to extract relevant concepts from a given query. Ueki *et al.* [12] come with a much larger concept bank consisting of more than 50k concepts. In addition to pre-trained CNN models, they train SVM classifiers to automatically annotate the video content. In [43], Snoek *et al.* utilize a model called VideoStory [7] to represent videos and then embed them into a concept space by a linear transformation, while they still represent textual query by selecting concepts based on part-of-speech tagging heuristically. Consequently, the video-text similarity is implemented as the cosine similarity in terms of their concept vectors.

We argue that such a concept-based paradigm has a fundamental disadvantage. That is, it is very difficult to describe the rich sequential information within both video and query using a few selected concepts. Despite the disadvantage, such a paradigm also has its merit where representing video and textual query by concepts make it somewhat interpretable. In this work, we also integrate such interpretation merit into our proposed model, allowing the model to match videos and text in the concept space.

2.2 Latent Space based Methods

Latent space based methods first encode video and textual queries and then map them into a common latent space where the video-text similarity can be measured directly [17], [32], [44], [45]. For these methods, what matters are forms of video encoding, text encoding and similarity learning. So we review recent work in these aspects.

For video encoding, a typical approach is to first extract visual features from video frames by pre-trained CNN models, and subsequently aggregate the frame-level features into a video-level feature. The *de facto* choice is mean pooling [17], [19], [23], [29], [31] or max pooling [20], [24], [28]. To explicitly model the temporal information, Torabi *et al.* [26] use a Long Short-Term Memory (LSTM), where frame-level features are sequentially fed into the LSTM, and the mean pooling of the hidden vectors at each step is used as the video feature. Besides using a Gated Recurrent Unit (GRU) to model the temporal dependency of video frames, Yang *et al.* [21] additionally utilize the multi-head self-attention mechanism [46] to learn the frame-wise correlation thus enhance the video representation. In addition to the frame-level visual features, we also notice efforts on utilizing features extracted from other channels such as audio [17], [19], [28] and motion [17], [19], [20], [28]. However, they still use mean pooling, max pooling or NetVLAD to aggregate different features into video-level features.

For text encoding, while bag-of-words (BoW) remains common [7], [27], [47], deep networks are in increasing use. The common way by deep learning techniques is to first represent each word of the textual query by word2vec models pre-trained on large-scale text corpora, and then aggregate them by max pooling in [20], Fisher Vector in [28], [29] or NetVLAD in [19], [24]. Despite their good performance, the

main drawback of these methods is ignoring the sequential order in text. Recurrent neural networks (RNN), known to be effective for modeling sequence order dependency, are also dominated [17], [23], [30], [44], [48]. Recursive neural networks are investigated in [44] for vectorizing subject-verb-object triplets extracted from a given sentence. Variants of recurrent neural networks are being exploited, see the usage of LSTM, bidirectional LSTM, GRU, bidirectional GRU in [48], [30], [17] and [49], respectively. For instance, Mithun *et al.* [17] utilize the last hidden state of the GRU as the text representation. The work [31] and its extension [32] explore a joint use of multiple text encoding strategies including BoW, word2vec and GRU, and found it is beneficial for video retrieval. However, as aforementioned, those works simply employ mean pooling for video encoding. Recently, Chen *et al.* [22] utilize graph convolutional network to model the connection between words. But it requires text to be well annotated with semantic role relation annotations, which may be not suitable for a new scenario with different linguistic expression patterns. By contrast, our proposed text encoding trained in an end-to-end manner, without using such semantic role annotations. Moreover, different from the above works, this paper aims to explicitly and progressively exploit global, local and temporal patterns in both videos and textual queries.

For similarity learning, lots of works [19], [20], [23], [31], [32], [44] typically project encoded videos and text into a common latent space and triplet ranking loss variants are dominantly used for model training. After being projected in a latent space, the video-text similarity can be measured by a standard similarity metric, *e.g.*, cosine similarity. Besides the triplet ranking loss, Zhang *et al.* [18] additionally employ contrastive loss and reconstruction loss to further constrain the latent space. Recently, we notice an increasing used of learning multiple common latent spaces [17], [22], [24], [28] instead of only one latent space. For instance, Miech *et al.* [28] utilize four different features, *i.e.*, appearance, motion, face and audio features, to represent videos, and learn one latent space for each video feature. The final video-text similarity is the fusion of their similarities in the four latent spaces. Wray *et al.* [24] decompose text into nouns and non-noun words, and respectively project them into two different latent spaces. With the similar idea of [24], Chen *et al.* [22] decompose text into events, actions and entities, and three latent spaces are learned for video-text matching.

Different from the existing works where learned common spaces are latent spaces without no explicit interpretability, we propose to learn a hybrid space consisting of a latent space and a concept space, which inherits the merits of high performance of latent space based methods and interpretability of concept based methods.

2.3 Cross-modal Fusion Methods

In contrast to latent space based methods, cross-modal fusion methods do not construct an explicit latent space. Instead, they use a cross-modal fusion subnetwork that takes text and videos as input and directly produces similarity scores [30], [39], [48]. For instance, Yu *et al.* [30] fuse the video and text representation into a 3D tensor using a soft attention, and a convolutional network is further employed

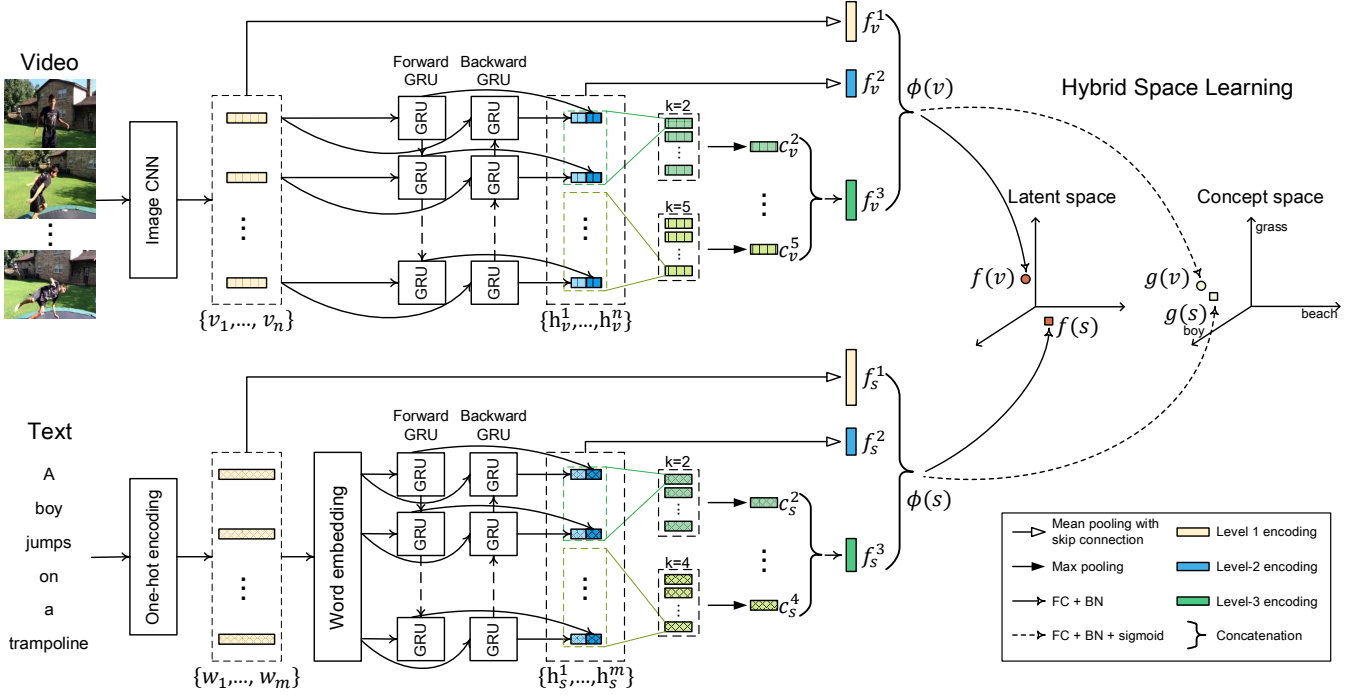


Fig. 2. **A conceptual diagram of the proposed dual encoding network for video retrieval by text.** Given a video v and a sentence s , the network performs in parallel multi-level encodings, *i.e.*, mean pooling, biGRU and biGRU-CNN, eventually representing the two input by two combined vectors $\phi(v)$ and $\phi(s)$, respectively. The vectors are later projected into a hybrid common space which consists of a latent space and a concept space. Once the network is trained, encoding at each side is performed independently, meaning we can process large-scale videos offline and answer ad-hoc queries on the fly.

to directly predict the similarity based on the fused feature. Although these methods [30], [48] are effective, their retrieval efficiency are somewhat low as video and text are coupled with each other. By contrast, our proposed method maps videos and text into a common space by a separated branch respectively, which makes videos and text decoupled. So, all candidate videos can be mapped in the common space off-line, which is efficient for large-scale video retrieval.

2.4 Transformer-based Methods

Transformers [46] are a new genre of deep neural networks that stack multiple self-attended layers. By self-supervised learning with masked language modeling on very large-scale corpora, BERT-like Transformer models have shown superior performance over RNNs on multiple NLP tasks [50]. Unsurprisingly, we see an increasing effort on re-purposing such types of models for video retrieval by text [39], [40], [51], [52]. In Lokoć *et al.* [51] and Zhao *et al.* [52], a pre-trained BERT is adopted to encode the textual query. Gabeur *et al.* [40] also use BERT for query representation, and go one step further by proposing a multi-modal Transformer (MMT) with four stacked Transformer layers to jointly encode diverse video features for video representation. Zhu and Yang propose ActBERT to encode global actions, local regional objects, and text descriptions in a unified framework [39]. We empirically show in Section 5.2.3 that our proposed method, although utilizing relatively simple encoders such as Bag-of-Words, GRU and 1D-CNN, compares favorably to the Transformer based methods. Moreover, our method can be extended with ease to incorporate the Transformers for better performance.

3 THE DUAL ENCODING NETWORK

Both video and sentence are essentially a sequence of items, let it be frames or words. Such a property motivates us to design a dual encoding network to handle the two distinct modalities. Specifically, given a video v and a sentence s , the proposed network encodes them in parallel, in advance to common space learning. As illustrated in Fig. 2, multi-level encodings are performed for each modality. The encoding results are then combined, denoted as $\phi(v)$ and $\phi(s)$, to describe the two modalities in a coarse-to-fine fashion. In our design, $\phi(v)$ and $\phi(s)$ are not directly used for cross-modal matching. So the two vectors do not have to reside in the same feature space and can have distinct dimensions, giving them sufficient freedom to become powerful representations of the corresponding modalities.

In what follows we first depict the network at the video side. We then specify choices that are unique at the text side.

3.1 Video-side Multi-level Encoding

For a given video, we extract uniformly a sequence of n frames with a pre-specified interval of 0.5 seconds. Per frame we extract deep features using a pre-trained ImageNet CNN, as commonly used for video content analysis [7], [42], [44]. Consequently, the video is described by a sequence of feature vectors $\{v_1, v_2, \dots, v_n\}$, where v_t indicates the deep feature vector of the t -th frame. Notice that 3D CNNs [53], [54] can also be used for feature extraction when treating segments of frames as individual items.

3.1.1 Level 1. Global Encoding by Mean Pooling

According to our literature review, mean pooling, which represents a video by simply averaging the features of

its frames, is arguably the most popular choice for text-video retrieval. By definition, mean pooling captures visual patterns that repeatedly present in the video content. These patterns tend to be global. We use f_v^1 to indicate the encoding result at this level, that is:

$$f_v^1 = \frac{1}{n} \sum_{t=1}^n v_t. \quad (1)$$

3.1.2 Level 2. Temporal-Aware Encoding by biGRU

Bi-directional recurrent neural network [55] is known to be effective for making use of both past and future contextual information of a given sequence. We hypothesize that such a network is also effective for modeling the video temporal information. We adopt a bidirectional GRU (biGRU) [56], which has less parameters than the bidirectional LSTM and thus requires less amounts of training data. A biGRU consists of two separated GRU layers, *i.e.*, a forward GRU and a backward GRU. The forward GRU is used to encode frame features in normal order, while the backward GRU encodes frame features in reverse order. Let \vec{h}_t and \overleftarrow{h}_t be their corresponding hidden states at a specific time step $t = 1, \dots, n$. The hidden states are generated as

$$\begin{aligned} \vec{h}_t &= \overrightarrow{GRU}(v_t, \vec{h}_{t-1}), \\ \overleftarrow{h}_t &= \overleftarrow{GRU}(v_{n+1-t}, \overleftarrow{h}_{t-1}), \end{aligned} \quad (2)$$

where \overrightarrow{GRU} and \overleftarrow{GRU} indicate the forward and backward GRUs, with past information carried by \vec{h}_{t-1} and \overleftarrow{h}_{t-1} , respectively. Concatenating \vec{h}_t and \overleftarrow{h}_t , we obtain the biGRU output $h_v^t = [\vec{h}_t, \overleftarrow{h}_t]$. The size of the hidden vectors in the forward and backward GRUs is empirically set to 512. Accordingly, the size of h_v^t is 1,024. Putting all the output together, we obtain a feature map $H_v = \{h_v^1, h_v^2, \dots, h_v^n\}$, with a size of $1,024 \times n$. The biGRU based encoding, denoted $f_v^{(2)}$, is obtained by applying mean pooling on H_v along the row dimension, that is

$$f_v^2 = \frac{1}{n} \sum_{t=1}^n h_v^t. \quad (3)$$

3.1.3 Level 3. Local-Enhanced Encoding by biGRU-CNN

The previous layer treats the output of biGRU at each step equally. To enhance local patterns that help discriminate between videos of subtle difference, we build convolutional networks on top of biGRU. In particular, we adapt 1D CNN originally developed for sentence classification [57].

The input of our CNN is the feature map H_v generated by the previous biGRU module. Let $Conv1d_{k,r}$ be a 1D convolutional block that contains $r = 512$ filters of size k , with $k \geq 2$. Feeding H_v , after zero padding, into $Conv1d_{k,r}$ produces a $n \times r$ feature map. Non-linearity is introduced by applying the ReLU activation function on the feature map. As n varies for videos, we further apply max pooling to compress the feature map to a vector c_k of fixed length r . More formally we express the above process as

$$c_k^r = \text{max-pooling}(\text{ReLU}(\text{Conv1d}_{k,r}(H_v))). \quad (4)$$

A filter with $k = 2$ allows two adjacent rows in H_v to interact with each other, while a filter of larger k means

more adjacent rows are exploited simultaneously. In order to generate a multi-scale representation, we deploy multiple 1D convolutional blocks with $k = 2, 3, 4, 5$. Their output is concatenated to form the biGRU-CNN based encoding, *i.e.*,

$$f_v^3 = [c_v^2, c_v^3, c_v^4, c_v^5]. \quad (5)$$

As f_v^1, f_v^2, f_v^3 are obtained sequentially at different levels by specific encoding strategies, we consider it reasonable to presume that the three encoding results are complementary to each other, with some redundancy. Hence, we obtain multi-level encoding of the input video by concatenating the output from all the three levels, namely

$$\phi(v) = [f_v^1, f_v^2, f_v^3]. \quad (6)$$

In fact, this concatenation operation, while being simple, is a common practice for feature combination [58], [59].

It is worth noting that we have described how to combine three encoders, *i.e.*, mean pooling, biGRU and biGRU-CNN, which naturally capture global, temporal and local patterns. In principle, other encoders, *e.g.*, Transformers, that are complementary to the existing encoders can be combined.

3.2 Text-side Multi-level Encoding

The above encoding network, after minor modification, is also applicable for the text modality.

Given a sentence s of length m , we represent each of its words by a one-hot vector. Accordingly, a sequence of one-hot vectors $\{w_1, w_2, \dots, w_m\}$ is generated, where w_t indicates the vector of the t -th word. Global encoding f_s^1 is obtained by averaging all the individual vectors in the sequence. This amounts to the classical bag-of-words representation.

For biGRU based encoding, each word is first converted to a dense vector by multiplying its one-hot vector with a word embedding matrix. We initialize the matrix using a word2vec [60] model provided by [31], which trained word2vec on English tags of 30 million Flickr images. The rest is mostly identical to the video counterpart. We denote the biGRU based encoding of the sentence as f_s^2 . Similarly, we have the biGRU-CNN based encoding of the sentence as f_s^3 . Here, we utilize three 1D convolutional blocks with $k = 2, 3, 4$. Multi-level encoding of the sentence is obtained by concatenating the encoding results from all the three levels in the dual network, *i.e.*,

$$\phi(s) = [f_s^1, f_s^2, f_s^3]. \quad (7)$$

As $\phi(v)$ and $\phi(s)$ have not been correlated, they are not directly comparable. For video-text similarity computation, the vectors need to be projected into a common space, the learning algorithm for which will be presented next.

4 HYBRID SPACE LEARNING

We propose to train our dual encoding network with a hybrid space learning algorithm. The hybrid space consists of a latent space which aims for good performance and a concept space which is meant for good interpretability.

4.1 Learning a Latent Space

Network. Given the encoded video vector $\phi(v)$ and the sentence vector $\phi(s)$, we project them into a latent space by affine transformations. From the neural network viewpoint, an affine transformation is essentially a Fully Connected (FC) layer. We additionally use a Batch Normalization (BN) layer after the FC layer, as we find this trick beneficial. Putting everything together, we obtain the video feature vector $f(v)$ and sentence feature vector $f(s)$ in the latent space as:

$$\begin{aligned} f(v) &= \text{BN}(W_1\phi(v) + b_1), \\ f(s) &= \text{BN}(W_2\phi(s) + b_2), \end{aligned} \quad (8)$$

where W_1 and W_2 parameterize the FC layers on each side, with b_1 and b_2 as bias terms.

To measure the video-text similarity $\text{sim}_{\text{lat}}(v, s)$ in the latent space, we use popular cosine similarity between $f(v)$ and $f(s)$:

$$\text{sim}_{\text{lat}}(v, s) = \frac{f(v) \cdot f(s)}{\|f(v)\| \|f(s)\|}. \quad (9)$$

In our preliminary experiment, we also tried the Manhattan and Euclidean distance, but found them less effective than the cosine similarity.

Loss. A desirable similarity function shall make relevant video-sentence pairs near and irrelevant pairs far away in the latent space. Therefore, we use the improved triplet ranking loss [61], which penalizes the model according to the hardest negative examples in the mini-batch. Concretely, given a relevant video-sentence pair (v, s) in a mini-batch, its loss $\mathcal{L}_{\text{lat}}(v, s)$ is:

$$\begin{aligned} \mathcal{L}_{\text{lat}}(v, s) &= \max(0, m + \text{sim}_{\text{lat}}(v, s^-) - \text{sim}_{\text{lat}}(v, s)) \\ &\quad + \max(0, m + \text{sim}_{\text{lat}}(v^-, s) - \text{sim}_{\text{lat}}(v, s)), \end{aligned} \quad (10)$$

where m is the margin constant, while s^- and v^- respectively indicate a negative sentence sample for v and a negative video sample for s . The two negatives are not randomly sampled. Instead, the most similar yet negative sentence and video in the current mini-batch are chosen.

4.2 Learning a Concept Space

As multiple concepts can be used simultaneously to describe a specific video or sentence, learning a concept space can be naturally formulated as a multi-label classification problem.

Network. Suppose the size of the concept vocabulary is K . In order to project $\phi(v)$ and $\phi(s)$ into a K -dimensional concept space, we adopt a network similar to the network used for latent space learning. That is,

$$\begin{aligned} g(v) &= \sigma(\text{BN}(W_3\phi(v) + b_3)), \\ g(s) &= \sigma(\text{BN}(W_4\phi(s) + b_4)), \end{aligned} \quad (11)$$

Note that different from Eq. 8, we additionally use a sigmoid activation σ to produce probabilistic output. For a specific concept indexed by $i = 1, \dots, K$, we use $g(v)_i$ to denote the probability of the concept being relevant with respect to the video v . In a similar vein we define $g(s)_i$.

Consider the absolute scale of elements in the concept vectors matters, cosine similarity which mainly considers the direction of feature vectors is suboptimal to measure similarities between concept vectors. Viewing the two concept vectors $g(v)$ and $g(s)$ as (unnormalized) histograms, we

Training video:

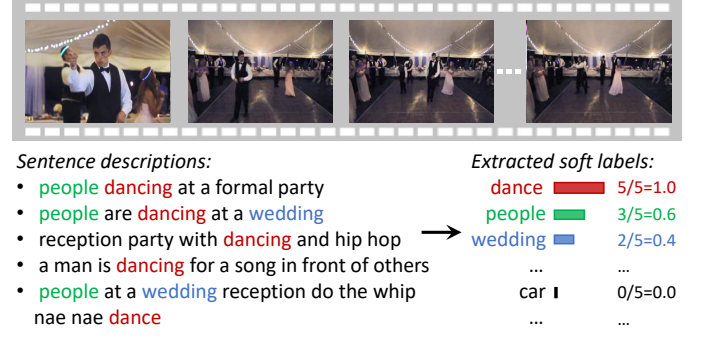


Fig. 3. **An illustration of extracting concept-level annotations from sentence descriptions of training videos.** Instead of binary labels, we extract frequency-based soft labels to better reflect the importance of a specific concept for a given video. E.g., the concept *dance* appears in all the five sentences of the example video, so its y value is 1, whilst the concept *wedding* occurring twice has a y value of 0.4

use generalized Jaccard similarity to compute the video-text similarity $\text{sim}_{\text{con}}(v, s)$ in the concept space, i.e.,

$$\text{sim}_{\text{con}}(v, s) = \frac{\sum_{i=1}^K \min(g(v)_i, g(s)_i)}{\sum_{i=1}^K \max(g(v)_i, g(s)_i)}. \quad (12)$$

Concept-level annotations. Per training video, we extract automatically its concept-level annotations from the associated sentence descriptions as follows. Assume that for a specific training video v , we have access to p sentences, $\{s_1, \dots, s_p\}$, that describe the video content. The relevance of a specific concept w.r.t v is determined by its occurrence in the p sentences. Li *et al.* [62] suggest that a concept appearing in multiple sentences is usually more important than those presented once. Hence, instead of binary labels, we obtain soft labels based on concept frequency. Specifically, let y be a K -dimensional ground-truth vector for v . The value of its i -th dimension, i.e., y_i , is defined as the frequency of the i -th concept divided by the maximum frequency of all concepts within the p sentences, see Fig. 3. Accordingly, we extend a relevant video-sentence pair (v, s) to a triplet training instance (v, s, y) to supervise concept space learning.

Loss. For multi-label classification, the binary cross-entropy (BCE) loss is common. In our context, the loss for a given video-sentence pair (v, s) with respect to their shared ground-truth y is computed as

$$\begin{aligned} \mathcal{L}_{\text{bce}}(v, s, y) &= -\left(\frac{1}{K} \sum_{i=1}^K [y_i \log(g(v)_i) + (1 - y_i) \log(1 - g(v)_i)]\right) \\ &\quad + \frac{1}{K} \sum_{i=1}^K [y_i \log(g(s)_i) + (1 - y_i) \log(1 - g(s)_i)]. \end{aligned} \quad (13)$$

We expect the concept space to be used not only for interpretability but also for improving video-text matching. So in addition to the BCE loss, we also minimize the improved triplet ranking loss in the concept space. That is,

$$\begin{aligned} \mathcal{L}_{\text{con,rank}}(v, s) &= \max(0, m + \text{sim}_{\text{con}}(v, s^-) - \text{sim}_{\text{con}}(v, s)) \\ &\quad + \max(0, m + \text{sim}_{\text{con}}(v^-, s) - \text{sim}_{\text{con}}(v, s)). \end{aligned} \quad (14)$$

The concept space is learned by minimizing the following combined loss:

$$\mathcal{L}_{con}(v, s, y) = \mathcal{L}_{bce}(v, s, y) + \mathcal{L}_{con,rank}(v, s). \quad (15)$$

Both \mathcal{L}_{bce} and $\mathcal{L}_{con,rank}$ matter in Eq. 15. Without \mathcal{L}_{bce} , Eq. 15 is boiled down to learning another latent space that not only lacks interpretability but also makes $sim_{con}(v, c)$ less complementary to $sim_{lat}(v, c)$ for video-text similarity computation. Meanwhile, using \mathcal{L}_{bce} alone makes the concept space suboptimal for video-text matching and sensitive to the concept annotation quality, c.f. the appendix.

4.3 Joint Learning of the Two Spaces

The dual encoding network is trained by minimizing the combination of the latent-space loss \mathcal{L}_{lat} and the concept-based loss \mathcal{L}_{con} . In particular, given a training set $\mathcal{D} = \{(v, s, y)\}$, we have

$$\underset{\theta}{\operatorname{argmin}} \sum_{(v,s,y) \in \mathcal{D}} \mathcal{L}_{lat}(v, s) + \mathcal{L}_{con}(v, s, y), \quad (16)$$

where θ denotes all the trainable parameters in the whole model. Except for image CNNs used for video feature extraction, the dual encoding network is trained in an end-to-end manner.

In order to conceptually explain why hybrid space learning helps, we provide a toy example in Fig. 4 that shows potential weakness when using a latent space or a concept space alone. As the latent space considers only relative positions per triplet, triplets that are semantically close, i.e., red and light-red triplets in Fig. 4(a), could be distant from each other. By contrast, each dimension of the concept space, which corresponds to a unique concept by definition, acts as an anchor to pinpoint the absolute position of a sample w.r.t. the concept. Fig. 4(b) shows the embedding space w.r.t. concept i and concept j , where the concept i alone is insufficient to discern samples near its origin. Other dimensions have to be taken into account. Nonetheless, as the concept set is pre-specified, the other dimensions are not necessarily discriminative, see the vertical axis corresponding to concept j . The hybrid space learning gives us the possibility of combining the best of the two spaces, while overcoming their weaknesses. In addition, the combination of video-text similarities in the two distinct spaces can be viewed as a lightweight instantiation of model ensemble [63], which typically improves over the base models.

4.4 Video-Text Similarity Computation

Once the model is trained, the final similarity between a video v and a sentence s is computed as the sum of their latent-space similarity and concept-space similarity, namely

$$sim(v, s) = \alpha \cdot sim_{lat}(v, c) + (1 - \alpha) \cdot sim_{con}(v, c), \quad (17)$$

where α is a hyper-parameter to balance the importance of two spaces, ranging within $[0, 1]$. Note that raw values of $sim_{lat}(v, s)$ and $sim_{con}(v, s)$ reside in distinct scales. Hence, they are rescaled separately by min-max normalization before being combined. Also note that in the inference stage, the multi-level encoding at the video side can be performed independently. Hence, for a large-scale video collection, their hybrid-space features can be pre-computed, allowing us to answer ad-hoc queries on the fly.

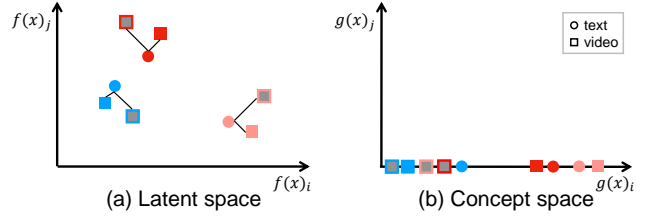


Fig. 4. **A toy example showing potential weakness of a latent space and a concept space in two dimensions.** Markers with the same color are relevant, while negative samples are with gray background. Red and light-red are semantically closer than the blue. Given a specific triplet, only their relative positions matter in the latent space. So even with zero triplet ranking loss, the red triplet can be projected into a region that is closer to the blue triplet rather than the light-red triplet. In the concept space, wherein each dimension corresponds to a distinct concept, the absolute positions of the triplets now matter. Samples positive w.r.t. a specific concept are to be projected to one end, while negatives are to be placed near the origin. As the region around the origin naturally lacks discriminability, other dimensions have to be considered. However, as the concepts are pre-specified, they are not necessarily optimal to describe the video content. Best viewed in color.

5 EVALUATION

Our evaluation is organized as follows. Firstly in Section 5.1, we compare the proposed Dual Encoding model (with its best setup) against the state-of-the-art on four datasets, i.e., MSR-VTT [34], TRECVID AVS 2016-2018 [8], [35], [36], VATEX [37] and MPII-MD [38]. Constructed independently by the dataset developers, the first three datasets consist of short web videos with very diverse content, while the last dataset contains video clips from 72 movies. Second, in order to verify the influence of major components in the proposed model, Section 5.2 presents an ablation study on the MSR-VTT dataset. Additional experiments concerning the effect of the hyper-parameter α and the concept annotation quality on the performance and the efficiency of our model are provided in the appendix.

Before proceeding to the experiments, we detail common implementations regarding text preprocessing, video features, concept vocabulary extraction, and model training. For sentence preprocessing, we first convert all words to the lowercase and then replace words that occurring less than five times in the training set with a special token. For video features, on VATEX we adopt 1,024-d I3D [54] video features provided by the dataset developers [37]. As for the other datasets, we extract frame-level ResNeXt-101 [64], [65] and ResNet-152 [66] using an open-source toolbox¹. The two feature vectors are concatenated to obtain a combined 4,096-d CNN feature, which we refer to as ResNeXt-ResNet.

To obtain the concept vocabulary, we conduct part-of-speech tagging by NLTK toolkit on all training sentences, and only keep the nouns, verbs and adjectives. All the English stopwords also removed. Besides, we also lemmatize the words, making *dog* and *dogs* to be a same concept. Finally, the top $K = 512$ frequent words are selected as the final concept vocabulary.

The proposed model is implemented using PyTorch (<http://pytorch.org>). Following [61], the parameter m for the improved triplet ranking loss is set to 0.2. The weight α in the combined similarity is empirically set to 0.6. We

1. <https://github.com/xuchaoxi/video-cnn-feat>

learn from our earlier studies [32], [41] that setting the dimensionality of the common space to 2,048 is a good practice. Hence, we let the overall dimensionality of the hybrid space be 2,048. Recall that the concept space is 512-dimensional. Accordingly, the dimensionality of the latent space is $2,048 - 512 = 1,536$. We use stochastic gradient descent with Adam [67]. The mini-batch size is 128. With an initial learning rate of 0.0001, we take an adjustment schedule similar to [31]. That is, once the validation loss does not decrease in three consecutive epochs, we divide the learning rate by 2. Early stop occurs if the validation performance does not improve in ten consecutive epochs. The maximal number of epochs is 50. In practice, early stop occurs typically after 15 epochs.

5.1 Comparison with the State-of-the-art

5.1.1 Experiments on MSR-VTT

Data. The MSR-VTT dataset [34], originally developed for video captioning, consists of 10k web video clips and 200k natural sentences describing the visual content of the clips. The number of sentences per clip is 20. For this dataset, we notice there are three distinct editions of data partition in the literature [28], [30], [34]. The official partition [34] uses 6,513 clips for training, 497 clips for validation, and the remaining 2,990 clips for testing. For the partition by [28], there are 6,656 clips for training and 1,000 clips for testing. The partition of [30] uses 7,010 and 1,000 clips for training and testing, respectively. As the last two data partitions provide no validation set, we build a validation set by randomly sample 1,000 clips from MSR-VTT with [28], [30] excluded, respectively. For a comprehensive evaluation, our experiments are performed on all the three data partitions.

Performance Metrics. We use rank-based metrics, namely $R@K$ ($K = 1, 5, 10$), Median rank (Med r) and mean Average Precision (mAP) to evaluate the performance. $R@K$ is the percentage of test queries for which at least one relevant item is found among the top- K retrieved results. Med r is the median rank of the first relevant item in the search results. Higher $R@K$, mAP and lower Med r mean better performance. For overall comparison, we report the Sum of all Recalls (SumR).

Baselines. The following thirteen state-of-the-art models are compared:

- VSE++ [61]: A state-of-the-art text-image retrieval model, which is commonly used as the strong baseline model for text-video retrieval. We replace its image-side branch with mean pooling on frame-level feature followed by a FC layer.
- W2VV [31]: Learn to project text into a visual feature space by minimizing the distance of relevant video-text pairs in the visual space. Multiple text encoding strategies including BoW, word2vec and GRU are jointly used for text encoding.
- MEE [28]: Use four different features to represent videos, and learn one latent space for each video feature. The weighted sum of similarities in four latent spaces is regarded as the final video-text similarity.
- W2VV++ [32]: An improved version of W2VV, it employs a better sentence encoding strategy and an improved triplet ranking loss.
- CE [19]: Use a collaborative gating to fuse multiple features to obtain a strong video representation.

- TCE [21]: Utilize a latent semantic tree [69] augmented encoder to represent text, and a GRU with multi-head self-attention mechanism [46] to encode videos.
- HGR [22]: Utilize graph convolutional network to model the connection between words, and project text and videos into three latent spaces.
- Mithun *et al.* [17]: Project videos and text into two latent spaces, and a weighted triplet ranking loss is used for training.
- Francis *et al.* [68]: Fuse multimodal features, *i.e.*, counting, activity and concept features to obtain stronger video and text representations.
- JPoSE [24]: Decompose text into nouns and non-noun words, and respectively project them into two different latent spaces.
- CT-SAN [48]: Learn to directly predict the similarity based on the fused video-text features without learning a common space, and a concept word detector is used for enhancing the video representation.
- JSFusion [30]: With the same idea of [48] that directly predicts the video-text similarity, a stronger joint sequence fusion is employed to fuse video and text features.
- Miech *et al.* [20]: Project videos and text into a common space by a gated embedding module respectively. The model is pre-trained on large-scale video-text dataset HowTo100M [20] and fine-tuned on MSR-VTT.

For a direct comparison, we cite numbers from the original papers whenever applicable. Meanwhile, we notice that video features used by specific papers vary. So, to make the comparison fairer, we have re-trained the following seven models which have been open-sourced, *i.e.*, W2VV, MEE, VSE++, W2VV++, CE, TCE and HGR, using the same ResNeXt-ResNet feature².

Results. Table 1 summarizes the performance comparison on three different data partitions of MSR-VTT. Though our goal is video retrieval by text, which corresponds to text-to-video retrieval in the table, video-to-text retrieval is also included for completeness. Note that the number of the candidate videos/sentences to be retrieved in the official partition is larger than that in the other two partitions. Hence, the official partition is more challenging. As a consequence, for all the models, their performance scores on the official partition are lower than their counterparts on the other partitions. Consider the results using the same video features, our dual encoding model achieves the best overall performance.

Among the results marked with asterisks (*), CE* which utilizes seven video features performs the best. Still, the proposed Dual Encoding model using only two visual features outperforms CE* on the first two data partitions, while slightly worse on Test1k-Yu. Note that on this test set, Dual Encoding was trained on the corresponding training set of 7,010 videos as specified by [30], whilst CE* was trained on an enlarged set of 9,000 videos. The results demonstrate the effectiveness of our proposed model.

2. MEE and CE employ a separated branch to handle the two video features respectively, while others utilize the concatenated feature as the whole input.

TABLE 1

State-of-the-art on MSR-VTT. Larger R@{1,5,10}, mAP and smaller Med r indicate better performance. Symbol asterisk (*) indicates numbers directly cited from the original papers, and the others are obtained by our re-training given the same ResNeXt-ResNet feature. On all the three distinct editions of data partition, the proposed Dual Encoding model obtains the best overall performance.

Method	Text-to-Video Retrieval					Video-to-Text Retrieval					SumR
	R@1	R@5	R@10	Med r	mAP	R@1	R@5	R@10	Med r	mAP	
Official full-size test set [34]											
Francis <i>et al.</i> * [68]	6.5	19.3	28.0	42	-	-	-	-	-	-	-
Mithun <i>et al.</i> * [17]	7.0	20.9	29.7	38	-	12.5	32.1	42.4	16	-	144.6
TCE* [21]	7.7	22.5	32.1	30	-	-	-	-	-	-	-
HGR* [22]	9.2	26.2	36.5	24	-	15.0	36.7	48.8	11	-	172.4
CE* [19]	10.0	29.0	41.2	16	-	15.6	40.9	55.2	8.3	-	191.9
W2VV [31]	1.1	4.7	8.1	236	3.7	17.0	37.9	49.1	11	7.6	117.9
MEE [28]	6.8	20.7	31.1	28	14.7	13.4	32.0	44.0	14	6.6	148.0
CE [19]	7.9	23.6	34.6	23	16.5	11.0	31.9	46.1	13	6.8	155.1
VSE++ [61]	8.7	24.3	34.1	28	16.9	15.6	36.6	48.6	11	7.4	167.9
TCE [21]	9.3	27.3	38.6	19	18.7	15.1	36.8	50.2	10	8.0	177.3
W2VV++ [32]	11.1	29.6	40.5	18	20.6	17.5	40.2	52.5	9	8.5	191.4
HGR [22]	11.1	30.5	42.1	16	20.8	18.7	44.3	57.6	7	9.9	204.4
Dual Encoding	11.6	30.3	41.3	17	21.2	22.5	47.1	58.9	7	10.5	211.7
Test1k-Miech [28]											
JPoSE* [24]	14.3	38.1	53.0	9	-	16.4	41.3	54.4	8.7	-	217.5
MEE* [28]	16.8	41.0	54.4	9	-	-	-	-	-	-	-
TCE* [21]	17.1	39.9	53.7	9	-	-	-	-	-	-	-
CE* [19]	18.2	46.0	60.7	7	-	18.0	46.0	60.3	6.5	-	249.2
W2VV [31]	2.7	12.5	17.3	83	7.9	17.3	42.0	53.5	9	29.3	145.3
MEE [28]	15.7	39.0	52.3	9	27.1	15.3	41.9	54.5	8	28.1	218.7
VSE++ [61]	17.0	40.9	52.0	10	16.9	18.1	40.4	52.1	9	29.2	220.5
CE [19]	17.8	42.8	56.1	8	30.3	17.4	42.9	56.1	8	29.8	233.1
TCE [21]	17.0	44.7	58.3	7	30.0	15.1	43.3	58.2	7	28.3	236.6
W2VV++ [32]	21.7	48.6	60.9	6	34.4	18.6	46.4	59.1	6	31.7	255.3
HGR [22]	22.9	50.2	63.6	5	35.9	20.0	48.3	60.9	6	33.2	265.9
Dual Encoding	23.0	50.6	62.5	5	36.1	25.1	52.1	64.6	5	37.7	277.9
Test1k-Yu [30]											
CT-SAN* [48]	4.4	16.6	22.3	35	-	-	-	-	-	-	-
JSFusion* [30]	10.2	31.2	43.2	13	-	-	-	-	-	-	-
TCE* [21]	16.1	38.0	51.5	10	-	-	-	-	-	-	-
Miech <i>et al.</i> * [20]	14.9	40.2	52.8	9	-	-	-	-	-	-	-
CE* [19]	20.9	48.8	62.4	6	-	20.6	50.3	64.0	5.3	-	267.0
W2VV [31]	1.9	9.9	15.2	79	6.8	17.3	39.3	50.2	10	27.8	133.8
VSE++ [61]	16.0	38.5	50.9	10	27.4	16.2	39.3	51.2	10	27.4	212.1
MEE [28]	14.6	38.4	52.4	9	26.1	15.2	40.9	53.8	9	27.9	215.3
W2VV++ [32]	19.0	45.0	58.7	7	31.8	16.9	42.7	54.6	8	29.0	236.9
CE [19]	17.2	46.2	58.5	7	30.3	15.8	44.9	59.2	7	30.4	241.8
TCE [21]	17.8	46.0	58.3	7	31.1	18.9	43.5	58.8	7	31.4	243.3
HGR [22]	21.7	47.4	61.1	6	34.0	20.4	47.9	60.6	6	33.4	259.1
Dual Encoding	21.1	48.7	60.2	6	33.6	21.7	49.4	61.6	6	34.7	262.7

5.1.2 Experiments on TRECVID AVS 2016-2018

Data. IACC.3 dataset is the largest test bed for video retrieval by text to this date, which developed for TRECVID (Ad-hoc Video Search) AVS 2016, 2017 and 2018 task [8], [35], [36]. The dataset contains 4,593 Internet Archive videos with duration ranging from 6.5 minutes to 9.5 minutes and a mean duration of almost 7.8 minutes. Shot boundary detection results in 335,944 shots in total. Given an ad-hoc query, *e.g.*, *Find shots of military personnel interacting with protesters*, the task is to return for the query a list of 1,000 shots from the test collection ranked according to their likelihood of containing the given query. Per year TRECVID specifies 30 distinct queries of varied complexity.

As TRECVID does not specify training data for the AVS task, we train the dual encoding network using the joint collection of MSR-VTT and the TGIF [70] which contains 100K animated GIFs and 120K sentences describing visual content of the GIFs. Although animated GIFs are a very different domain, TGIF was constructed in a way to resemble user-generated video clips, *e.g.*, with cartoon, static, and textual

content removed. For IACC.3, MSR-VTT and TGIF, we use the ResNeXt-ResNet video feature.

Performance Metrics. We utilize inferred Average Precision (infAP), the official performance metric used by the TRECVID AVS task. The overall performance is measured by averaging infAP scores over the queries. Their values are reported in percentage (%). Note that the TRECVID ground truth is partially available at the shot-level. The task organizers employ a pooling strategy to collect the ground truth, *i.e.*, a pool of candidate shots are formed by collecting the top-1000 shots from each submission and a random subset is selected for manual verification. The ground truth thus favors official participants. As the top ranked items found by our method can be outside of the subset, infAP scores of our method are likely to be underestimated.

Baselines. We include the top 3 entries of each year, *i.e.*, [10], [11], [71] for 2016, [12], [13], [43] for 2017 and [72], [73], [74] for 2018. Besides we include publications on the tasks, *i.e.*, [7], [42]. The most of above methods are concept based, except [72], [73], [74]. Among them, [72]

TABLE 2

State-of-the-art on the TRECVID AVS 2016 / 2017 / 2018. Symbol asterisk (*) indicates numbers directly cited from the original papers. The proposed dual encoding model consistently perform the best.

	TRECVID edition			OVERALL
	2016	2017	2018	
Top-3 TRECVID finalists:				
Rank 1*	5.4 [10]	20.6 [43]	12.1 [72]	—
Rank 2*	5.1 [11]	15.9 [12]	8.7 [73]	—
Rank 3*	4.0 [71]	12.0 [13]	8.2 [74]	—
Literature methods:				
VideoStory* [7], [75]	8.7	15.0	—	—
Markatopoulou <i>et al.</i> * [42]	6.4	—	—	—
CE [19]	7.4	14.5	8.6	10.2
VSE++ [61]	13.5	16.3	10.6	13.5
W2VV [31]	14.9	19.8	10.3	15.0
W2VV++ [32]	15.1	21.3	10.6	15.7
Dual Encoding	15.2	23.1	12.1	16.8

TABLE 3

State-of-the-art on VATEX. Our proposed model performs the best.

Method	Text-to-Video			Video-to-Text			SumR
	R@1	R@5	R@10	R@1	R@5	R@10	
W2VV [31]	14.6	36.3	46.1	39.6	69.5	79.4	285.5
VSE++ [61]	31.3	65.8	76.4	42.9	73.9	83.6	373.9
CE [19]	31.1	68.7	80.2	41.3	71.0	82.3	374.6
W2VV++ [32]	32.0	68.2	78.8	41.8	75.1	84.3	380.2
HGR [22]	35.1	73.5	83.5	—	—	—	—
<i>Dual Encoding</i>	36.8	73.6	83.7	46.8	75.7	85.1	401.7

fuses three W2VV++ variants with different settings, [73] uses two attention networks, besides the classical concept-based representation, while [74] is based on VSE++. Notice that visual features and training data used by these methods vary, meaning the comparison and consequently conclusions drawn from this comparison is at a system level. So for a more conclusive comparison, we re-train VSE++ [61], W2VV [31], W2VV++ [32] and CE [19] using the same training data and the ResNeXt-ResNet feature.

Results. Table 2 shows the performance of different methods on the TRECVID AVS 2016, 2017 and 2018 tasks, and the overall performance is the mean score of the three years. The proposed method again performs the best, with infAP of 15.2, 23.1 and 12.1 respectively. While [72] has a same infAP of 12.1 on the TRECVID AVS 2018 task, their solution ensembles three models. Their best single model, *i.e.*, W2VV++ [32] which uses the same training data and the same ResNeXt-ResNet feature, has a lower infAP of 10.6. Given the same training data and feature, the proposed method outperforms VSE++, W2VV, CE with a clear margin. These results confirm the effectiveness of our dual encoding for large-scale video retrieval by text.

5.1.3 Experiments on VATEX

Data. VATEX [37] a large-scale multilingual video description dataset. Each video, collected for YouTube, has a duration of 10 seconds. Per video there are 10 English sentences and 10 Chinese sentences to describe the corresponding

TABLE 4


State-of-the-art on MPII-MD. Our proposed model performs the best.

Method	Text-to-Video			Video-to-Text			SumR
	R@1	R@5	R@10	R@1	R@5	R@10	
W2VV++ [32]	0.3	1.0	1.7	0.1	0.7	1.3	5.1
VSE++ [61]	0.2	0.9	1.6	0.8	2.2	3.6	9.3
W2VV [31]	0.1	0.3	0.5	1.3	4.0	6.1	12.3
CE [19]	0.9	3.1	5.7	1.1	3.6	5.8	20.2
<i>Dual Encoding</i>	1.7	4.8	7.0	1.4	4.7	7.0	26.6

Top-5 shots retrieved from the MPII-MD test set by our model

Ground truth


Query: Someone smiles



Query: The car stops before a chain link fence



Query: They wrap their arms around each other



Query: In a restaurant, Someone sits at a table with the guy



Fig. 5. **Selected examples of movie retrieval by text on MPII-MD.** The top retrieved shots, though not being ground truth, appear to be correct.

video content. Here, we only use the English sentences. We adopt the dataset partition provided by [22], *i.e.*, 25,991 video clips for training, 1,500 clips for validation and 1,500 clips for testing, where validation and test set are obtained by randomly split the official validation set of 3,000 clips into two equal parts.

Performance Metrics. $R@K$ ($K = 1, 5, 10$) and SumR are used as the performance metrics.

Baselines. For method comparison, we consider HGR [22], the first work reporting video retrieval performance on VATEX. We also compare the VSE++ [61], W2VV [31], W2VV++ [32] and CE [19].

Results. Table 3 summarizes the performance, where all the models use the same I3D [54] video features. Among them, VSE, VSE++ and W2VV++ are all use mean pooling over the frame-level feature to encode videos, while our dual encoding model explores multi-level features to represent videos, consistently achieving better performance. The result shows the benefit of using the multi-level feature for video representation. Although HGR utilizes the extra semantic role annotation of sentences for text representation, our dual encoding model still slightly outperforms HGR.

5.1.4 Experiments on MPII-MD

Data. To evaluate the effectiveness of our proposed methods for a specific video domain, we conduct the experiment on MPII-MD [38] a movie description dataset. We use the official data partition, that is, 56,828, 4,929 and 6,580 movie clips

TABLE 5

Effectiveness of Dual Encoding. The overall performance, as indicated by Sum of Recalls, goes up as more encoding layers are added. Dual encoding exploiting all the three levels is the best.

Encoding strategy	Text-to-Video Retrieval					Video-to-Text Retrieval					SumR
	R@1	R@5	R@10	Med r	mAP	R@1	R@5	R@10	Med r	mAP	
Level 1 (Mean pooling)	9.7	26.8	37.0	23	18.5	17.7	40.0	51.7	10	8.3	182.9
Level 2 (biGRU)	10.3	28.5	39.4	19	19.7	18.4	41.9	54.3	8	9.3	192.7
Level 3 (biGRU-CNN)	11.1	30.1	41.6	17	20.9	18.1	41.6	55.3	8	9.6	197.9
Level 1 + 2	10.6	28.8	39.2	20	19.9	19.1	43.1	54.5	8	9.2	195.3
Level 1 + 3	11.5	30.0	40.8	18	20.9	19.8	42.7	55.2	8	9.5	200.1
Level 2 + 3	11.4	30.6	41.7	17	21.2	19.9	44.3	55.8	8	10.1	203.8
Level 1 + 2 + 3	11.6	30.3	41.3	17	21.2	22.5	47.1	58.9	7	10.5	211.7

TABLE 6

Performance of Dual encoding with distinct common spaces. Dataset: MSR-VTT.

Common Space	Text-to-Video Retrieval					Video-to-Text Retrieval					SumR
	R@1	R@5	R@10	Med r	mAP	R@1	R@5	R@10	Med r	mAP	
2,048-d Latent Space (Conference version [41])	11.0	29.2	39.8	19	20.2	18.8	42.7	56.2	8	9.3	197.7
1,536-d Latent Space	11.0	29.3	39.9	19	20.3	19.7	43.6	55.6	8	9.3	199.0
512-d Concept Space	9.9	26.8	37.4	23	18.7	17.9	41.5	53.9	8	9.0	187.4
2,048-d Hybrid, without $\mathcal{L}_{con,rank}$	9.8	26.3	36.0	25	18.2	17.2	40.7	53.3	9	8.7	183.3
2,048-d Hybrid: 1,536-d Latent Space + 512-d Concept Space	11.6	30.3	41.3	17	21.2	22.5	47.1	58.9	7	10.5	211.7

for training, evaluation and testing, respectively. Each movie clip is associated with one or two textual descriptions.

Performance Metrics. Performance of $R@1$, $R@5$, $R@10$ and SumR are reported.

Baselines. In this experiments, we compare the VSE++ [61], W2VV [31], W2VV++ [32] and CE [19]. All the models are trained using the ResNeXt-ResNet feature.

Results. Table 4 summarizes the performance on the MPII-MD dataset. Our proposed dual encoding model outperforms the other counterparts. It is worth noting that the performance of all models are lower on MPII-MD than that on MSR-VTT, we attribute it to the challenging nature of movie retrieval by vague descriptions on MPII-MD. See Fig. 5 for some qualitative results.

5.2 Ablation Study

In this section, we evaluate the viability of each component on MSR-VTT, with performance reported on the full-size test set unless otherwise stated. A comparison with recent multi-modal Transformer methods is also conducted. In addition, we investigate the effectiveness of our multi-level text encoding for image-text retrieval.

5.2.1 Multi-level Encoding versus Single-Level Encoding

To exam the usefulness of each encoding component in the dual encoding network, we conduct an ablation study as follows. Given varied combinations of the encoding components, seven models are trained. Table 5 summarizes the choices of video and text encodings and the corresponding performance.

Among the individual encoders, biGRU-CNN is found to be the most effective. As more encoding layers are included, the overall performance goes up. For the last four models which combine the output from previous layers, they all outperform the corresponding counterpart using

the output of a specific layer. E.g., the model with Level 1 + 2 encoding strategy outperforms the ones with Level 1 or Level 2. The results suggest that features of different levels are complementary to each other. The full multi-level encoding setup, *i.e.*, Level 1 + 2 + 3, performs the best.

We also investigate single-side encoding, that is, video-side multi-level encoding with mean pooling on the text side, and text-side multi-level encoding with mean pooling on the video side. These two strategies obtain SumR of 194.5 and 191, respectively. The lower scores justify the necessity of dual encoding. The result also suggests that video-side multi-level encoding is more beneficial.

5.2.2 Hybrid Space versus Single Space

In order to verify the effectiveness of the hybrid space, we have re-trained Dual Encoding with two alternative spaces, *i.e.*, fully latent space (the CVPR version [41]) and fully concept space, respectively. Table 6 summarizes their performance on MSR-VTT.

Our model with the hybrid space consistently outperforms the other two counterparts with a clear margin, which shows the effectiveness of hybrid space for video-text retrieval. Among them, although the model with the concept space is able to give some interpretation of the retrieval model, its performance is the worst. The latent space counterpart gives better performance than the concept space, but lacks interpretability. Moreover, simply increasing the dimensionality of the latent space, from 1,536 to 2,048, does not improve the performance. By contrast, the hybrid space strikes a proper balance between the retrieval performance and the interpretability.

To justify the necessity of the combined loss for concept space learning, we also report the performance of the hybrid space that excludes the triplet ranking loss from Eq. 15. This variant suffers a noticeable performance decrease in terms of

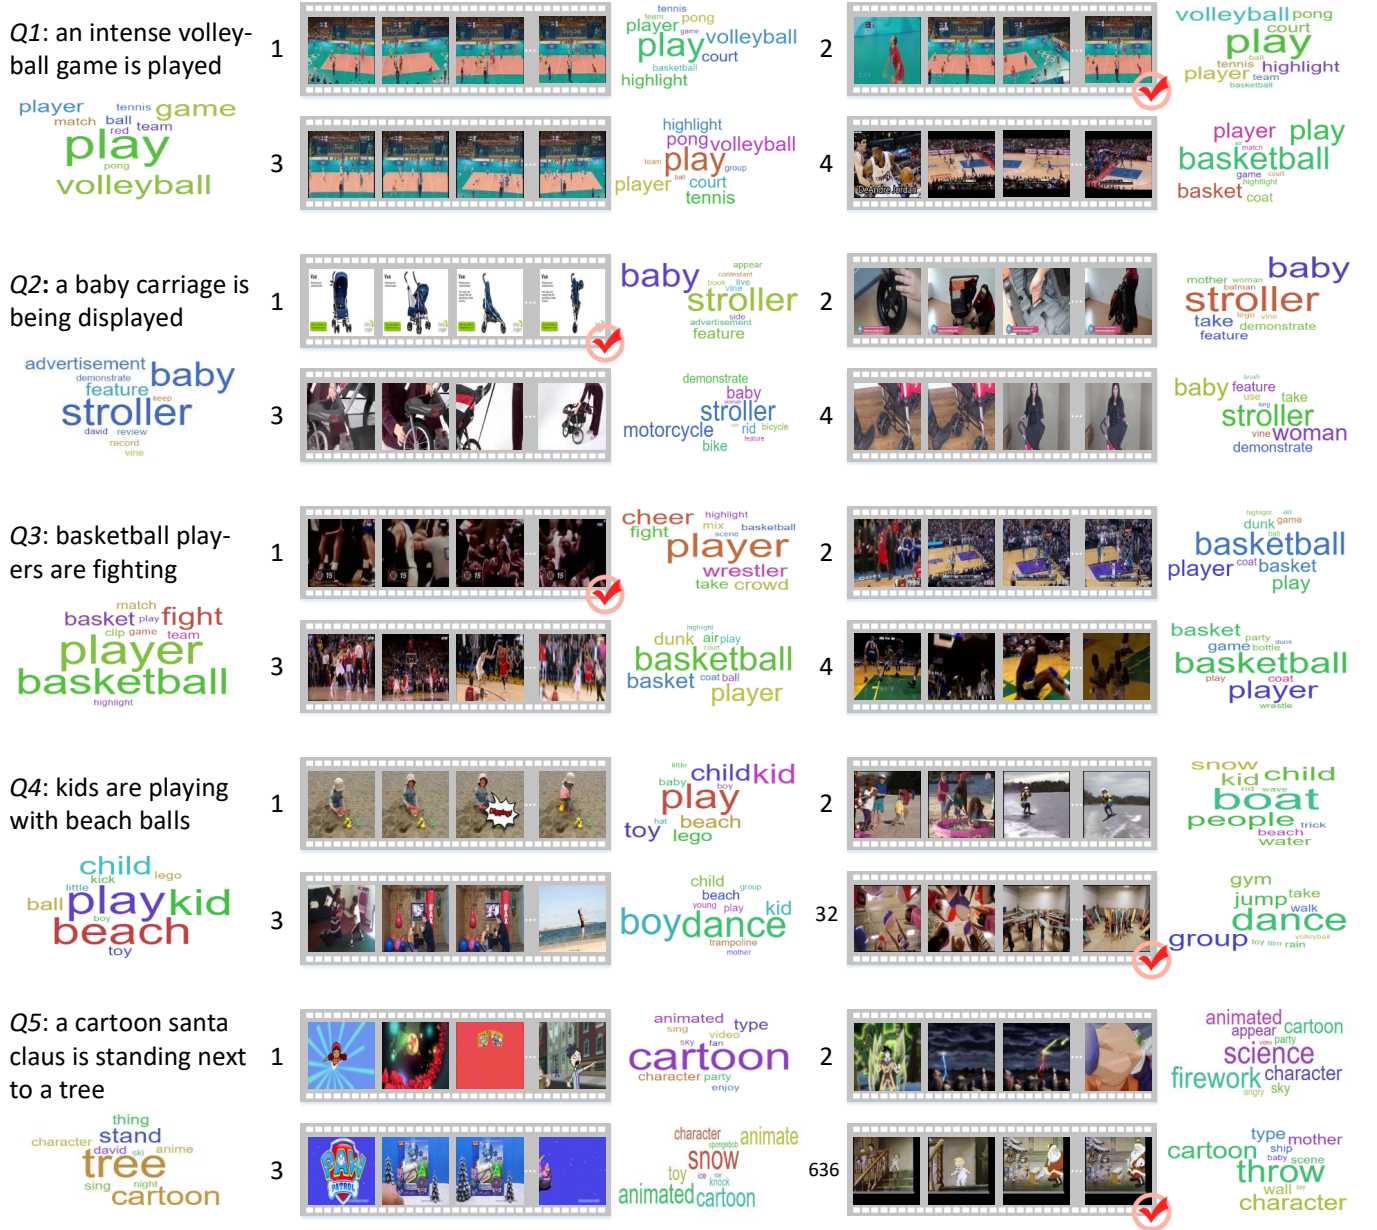


Fig. 6. **Selected examples of text-to-video retrieval by our model on MSR-VTT.** For each query, the top 3 ranked videos and the ground-truth video (marked with red ticks) are shown. In case the ground-truth video is among the top three, the fourth video will be included as well. By definition, each query has only one ground-truth video. Number on the left hand side of each video indicates the video's rank in the retrieval result. Below a specific query are its predicted concepts, visualized in the form of a tag cloud, bigger font meaning larger predicted scores. Next to the videos are their predicted concepts. Putting these tag clouds together helps us better understand the video retrieval results.

SumR, from 211.7 to 183.3. The result shows the importance of considering the triplet ranking loss for learning a concept space that is beneficial for video-text matching.

Interpreting retrieval results with predicted concepts. Fig. 6 shows some examples returned by our proposed dual encoding. Although only one correct video is annotated for each query, the top retrieved videos for Q1 (Query 1), Q2 and Q3 are typically relevant to the given query to some extent. In Q4, as the word *beach* is used to describe the object *ball*, the former is less important than the latter in this query. However, the predicted concept vector of Q4 shows that the model over emphasizes *beach*, as visualized in the tag cloud. This explains that the top 2 retrieved videos are all about

activities on beach. Meanwhile, for the truly relevant video, which is ranked at the position of 32, the predicted concepts are *dance* and *group*. Although these concepts are semantically relevant to the video content, they are irrelevant for the query. For Q5, concepts predicted our model, e.g., *cartoon* and *tree*, are not precise enough to capture *santa claus* the key role in the query. So our model also fails to answer this query. In general, we find concepts predicted by our dual encoding model reasonable, and useful for understanding the retrieval model.

5.2.3 Dual Encoding versus-and-with Transformers

Dual Encoding versus multi-modal Transformers. We compare with two recent multi-modal Transformers, i.e., Act-

TABLE 7

Comparison with Transformer-based multi-modal methods for video retrieval by text. Dataset: MSR-VTT. Even though Dual Encoding uses relatively simple 2D CNN video features, it outperforms MMT on Test1k-Miech [28], while less effective on Test1k-Yu [30]. The inclusion of Transformer features as shown in Fig. 7 makes Dual Encoding top the performance table consistently on all the three test sets.

Method	Official full-size test set [34]					Test1k-Miech [28]					Test1k-Yu [30]				
	R@1	R@5	R@10	Med r	mAP	R@1	R@5	R@10	Med r	mAP	R@1	R@5	R@10	Med r	mAP
ActBERT [39], HowTo100M pre-training	-	-	-	-	-	-	-	-	-	-	8.6	23.4	33.1	36	-
MMT [40]	-	-	-	-	-	20.3	49.1	63.9	6	-	24.6	54.0	67.1	4	-
MMT [40], HowTo100M pre-training	-	-	-	-	-	-	-	-	-	-	26.6	57.1	69.6	4	-
Dual Encoding	11.6	30.3	41.3	17	21.2	23.0	50.6	62.5	5	36.1	21.1	48.7	60.2	6	33.6
Dual Encoding, with Transformer features only	8.5	25.1	36.1	22	17.4	16.8	41.6	55.8	8	28.7	15.0	40.7	54.2	8	27.6
Dual Encoding, with Transformer features included	13.3	34.0	45.7	13	23.8	25.7	52.9	66.4	5	38.9	32.4	62.3	72.8	3	46.0

TABLE 8

Comparison with MMT in terms of model size and computation overhead at the inference stage. For each model, we measure the amount of FLOPs it takes to encode a given video-text pair. The computational cost of video feature extraction is excluded as that step is typically performed once in an offline mode. Numbers in parentheses indicate relative changes against MMT.

Model	Model Complexity	
	Parameters (M)	FLOPs (G)
MMT	133.4	12.64
Dual Encoding	95.9 (↓ 28.1%)	3.64 (↓ 71.2%)

BERT [39] and MMT [40], that have been evaluated on specific test sets of MSR-VTT. As Table 7 shows, Dual Encoding clearly outperforms ActBERT and is comparable to MMT in terms of the overall performance (better on Test1k-Miech and worse on Test1k-Yu). It is worth pointing out that MMT uses a diverse set of seven video features that describe varied aspects of the video content including motion, audio, scene, OCR, face, speech, and visual appearance features. By contrast, Dual Encoding uses only two regular 2D-CNN features. Moreover, on Test1k-Yu [30], Dual Encoding was trained on the corresponding training set of 7,010 videos, whilst MMT was trained on an enlarged set of 9,000 videos. We consider in this context that Dual Encoding compares favorably to MMT. In addition, Table 8 shows a comparison concerning model complexity. Note that ActBERT is not included, as that method remains closed-source, making a precise estimation of its model complexity impossible. Dual Encoding is smaller (28.1% less parameters) and much faster (71.2% less FLOPs) than MMT.

Dual Encoding with Transformers. Dual Encoding is a generic framework that allows pre-trained Transformers to be included with ease. As illustrated in Fig. 7, we extend the Dual Encoding network by adopting MMT trained on the HowTo100M dataset [20] as the visual Transformer and BERT [76] trained on Wikipedia and book corpora [77] as the textual Transformer. The two Transformers are integrated as is without fine tuning. Their outputs, *i.e.*, a 3,584-d video feature vector and a 1,024-d sentence vector, are concatenated with the video-side and text-side multi-level encodings, respectively, in advance to hybrid space learning. As shown in Table 7, the inclusion of the Transformer features clearly and consistently boosts the retrieval performance on all the three test sets. Meanwhile, using the Transformer features alone is less effective than the previous Dual Encoding model. Hence, our multi-level encoding modules, while

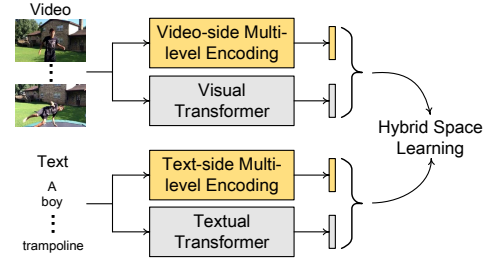


Fig. 7. Integrating pre-trained visual / textual Transformers into the Dual Encoding network. The outputs of the two transformers are respectively concatenated with the video-side and text-side multi-level encodings in advance to hybrid space learning.

TABLE 9

Performance of image-text retrieval on Flickr30k and MSCOCO. The proposed text-side multi-level encoding (MLE) is beneficial for VSE++ [61].

Method	Text-to-Image			Image-to-Text			SumR
	R@1	R@5	R@10	R@1	R@5	R@10	
<i>On Flickr30k</i>							
VSE++	23.1	49.2	60.7	31.9	58.4	68.0	291.3
VSE++, MLE	24.7	52.3	65.1	35.1	62.2	71.3	310.7
<i>On MSCOCO</i>							
VSE++	33.7	68.8	81.0	43.6	74.8	84.6	389.6
VSE++, MLE	34.8	69.6	82.6	46.7	76.2	85.8	395.7

built upon the relatively simple mean feature pooling, GRU and 1D-CNN encoders, remain competitive and can be used together with the Transformers to maximize the retrieval performance.

5.2.4 Multi-level Encoding for Image-Text Retrieval

Setup. We investigate if the VSE++ model [61] can be improved in its original context of image-text retrieval, when replacing its textual encoding module, which is a GRU, by the proposed multi-level encoding module. To that end, we fix all other choices, adopting the exact evaluation protocol of [61]. That is, we use the same data split, where the training / validation / test test has 30,000 / 1,000 / 1,000 images for Flickr30K, and 82,783 / 5,000 / 5,000 images for MSCOCO. We also use the same VGGNet feature provided by [61]. Performance of $R@1$, $R@5$ and $R@10$ are reported. On MSCOCO, the results are reported by averaging over 5 folds of 1,000 test images.

Results. Table 9 shows the performance of image-text retrieval on Flickr30k and MSCOCO. Integrating text-side

multi-level encoding into VSE++ brings improvements on both datasets. The results suggest that the proposed text-side multi-level encoding is also beneficial for VSE++ in its original context.

6 SUMMARY AND CONCLUSIONS

For video retrieval by text, this paper proposes a Dual Encoding network with hybrid space learning. By jointly exploiting multiple encoding strategies at different levels, the proposed network encodes both videos and text into powerful dense representations. Followed by hybrid space learning, these representations can be transformed to perform sequence-to-sequence cross-modal matching effectively. Extensive experiments on four video datasets, *i.e.*, MSR-VTT, TRECVID AVS 2016-2018, VATEX, and MPII-MD, support the following conclusions. Among the three levels of encoding, biGRU-CNN that builds a 1D convolutional network on top of bidirectional GRU is the most effective when used alone. Video-side multi-level encoding is more beneficial when compared with its text-side counterpart. Our multi-level encoding modules also show competitive performance against recent multi-modal Transformers. Compared with the widely used latent space learning, our hybrid space learning not only improves the retrieval performance but also enhances the interpretability of what the dual encoding network has learned. For state-of-the-art performance, we recommend Dual Encoding with hybrid space learning for video-text matching.

APPENDIX

Effect of α on the retrieval performance. The influence of the hyper-parameter α in Eq. 17 is studied as follows. We try α with its value ranging from 0.1 to 0.9 with an interval of 0.1. As shown in Fig. 8, when the α is larger than 0.2, the performance of our model with the hybrid space are all over 200 on MSR-VTT, which consistently outperform the counterparts using the latent space or the concept space alone. The results show that our hybrid space is not very sensitive to this parameter.

Effect of the concept annotation quality on the retrieval performance. The concept annotations used to supervise the hybrid space learning process are automatically extracted from video descriptions. In order to explore how their quality affects the retrieval performance, we re-train the model using concept annotations with varied levels of simulated noise. Concretely, suppose a specific training video is originally associated with q concepts. Given a noise level of $h \in \{0.1, 0.2, \dots, 1\}$, we replace a random subset of $q \times h$ concepts in the original annotations by the same number of concepts selected from the remaining $K - q$ concepts (recall that K is the size of the concept vocabulary). A larger h means lower annotation quality. As shown in Fig. 9(a), the retrieval performance is hardly affected by the concept annotation quality at a large range of noise levels. We attribute this result to two factors. First, the hybrid-space similarity is dominated by the latent-space similarity, which is independent of concept annotations by definition. Hence, while the retrieval performance of the concept-space similarity degenerates as more noise is added, the performance

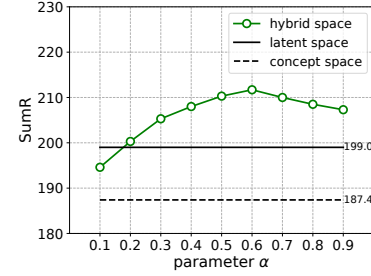


Fig. 8. The influence of the parameter α of Eq. 17 on our proposed model. Performance are evaluated on MSR-VTT.

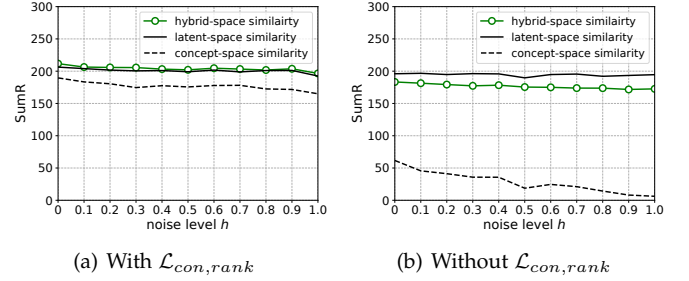


Fig. 9. The influence of the concept annotation quality on the overall performance of text-to-video and video-to-text retrieval. Per noise level, hybrid space learning is performed with $\mathcal{L}_{con,rank}$ (a) and without $\mathcal{L}_{con,rank}$ (b), respectively. The curve of “hybrid-space similarity” shows how the noise affects the model performance. To further reveal the effect of the noise on the individual spaces, we plot performance curves when using the latent-space similarity and concept-space similarity separately. The use of $\mathcal{L}_{con,rank}$ allows our model to be highly insensitive to the concept annotation quality.

of the latent-space similarity is largely stable. Second, the use of $\mathcal{L}_{con,rank}$ maintains to a large extent the effectiveness of the concept space for cross-modal matching. Note the extreme case of $h = 1$ at the right end of the curves, where the concept annotations become fully random and consequently the correspondence between the individual dimensions and concepts is completely lost. By minimizing $\mathcal{L}_{con,rank}$, the “concept” space remains a latent space that can be used for matching. Without the loss, the space will be meaningless, see Fig. 9(b). Therefore, the retrieval performance of our model is highly insensitive to the concept annotation quality. Although \mathcal{L}_{bce} is relatively sensitive to the quality of annotation, we consider the loss necessary as it produces an auxiliary common space not only interpretable but also more complementary to the latent space, see the better performance of the hybrid space against the fully latent space in Table 6.

Efficiency Test. Recall that the dual encoding network is designed to represent both videos and sentences into a common space respectively. Once the network is trained, representing them in the common space can be performed independently. This means we can process large-scale videos offline and answer ad-hoc queries on the fly. Specifically, given a natural-sentence query, it takes approximately 0.2 seconds to retrieve videos from the largest IACC.3 dataset, which consists of 335,944 videos. The performance is tested on a normal computer with 64G RAM and a GTX 1080TI GPU. The retrieval speed is adequate for instant response.

ACKNOWLEDGMENTS

This work was supported by NSFC (No. 61902347, No. 61672523), BJNSF (No. 4202033), ZJNSF (No. LQ19F020002), the Public Welfare Technology Research Project of Zhejiang Province (No. LGF21F020010), the Fundamental Research Funds for the Central Universities and the Research Funds of Renmin University of China (No. 18XNLG19), and the Alibaba-ZJU Joint Research Institute of Frontier Technologies.

REFERENCES

- [1] J. Dalton, J. Allan, and P. Mirajkar, "Zero-shot video retrieval using content and concepts," in *CIKM*, 2013, pp. 1857–1860.
- [2] L. Jiang, D. Meng, T. Mitamura, and A. G. Hauptmann, "Easy samples first: Self-paced reranking for zero-example multimedia search," in *ACM Multimedia*, 2014, pp. 547–556.
- [3] A. Habibiian, T. Mensink, and C. G. M. Snoek, "Composite concept discovery for zero-shot video event detection," in *ICMR*, 2014, pp. 17–24.
- [4] S. Wu, S. Bondugula, F. Luisier, X. Zhuang, and P. Natarajan, "Zero-shot event detection using multi-modal fusion of weakly supervised concepts," in *CVPR*, 2014, pp. 2665–2672.
- [5] X. Chang, Y. Yang, A. Hauptmann, E. P. Xing, and Y.-L. Yu, "Semantic concept discovery for large-scale zero-shot event detection," in *AAAI*, 2015, pp. 2234–2240.
- [6] Y.-J. Lu, H. Zhang, M. de Boer, and C.-W. Ngo, "Event detection with zero example: select the right and suppress the wrong concepts," in *ICMR*, 2016, pp. 127–134.
- [7] A. Habibiian, T. Mensink, and C. G. M. Snoek, "Video2vec embeddings recognize events when examples are scarce," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 39, no. 10, pp. 2089–2103, 2016.
- [8] G. Awad, J. Fiscus, D. Joy, M. Michel, A. Smeaton, W. Kraaij, G. Quenot, M. Eskevich, R. Aly, R. Ordelman, G. Jones, B. Huet, and M. Larson, "TRECVID 2016: Evaluating video search, video event detection, localization, and hyperlinking," in *TRECVID Workshop*, 2016.
- [9] G. Awad, D.-D. Le, C.-W. Ngo, V.-T. Nguyen, G. Quenot, C. Snoek, and S. Satoh, "Video indexing, search, detection, and description with focus on TRECVID," in *ICMR*, 2017, pp. 3–4.
- [10] D. Le, S. Phan, V.-T. Nguyen, B. Renoust, T. A. Nguyen, V.-N. Hoang, T. D. Ngo, M.-T. Tran, Y. Watanabe, M. Klinkigt *et al.*, "NII-HITACHI-UIT at TRECVID 2016," in *TRECVID Workshop*, 2016.
- [11] F. Markatopoulou, A. Moutzidou, D. Galanopoulos, T. Mironidis, V. Kaltsa, A. Ioannidou, S. Symeonidis, K. Avgerinakis, S. Andreadis *et al.*, "ITI-CERTH participation in TRECVID 2016," in *TRECVID Workshop*, 2016.
- [12] K. Ueki, K. Hirakawa, K. Kikuchi, T. Ogawa, and T. Kobayashi, "Waseda Meisei at TRECVID 2017: Ad-hoc video search," in *TRECVID Workshop*, 2017.
- [13] P. A. Nguyen, Q. Li, Z.-Q. Cheng, Y.-J. Lu, H. Zhang, X. Wu, and C.-W. Ngo, "Vireo at TRECVID 2017: Video-to-text, ad-hoc video search and video hyperlinking," in *TRECVID Workshop*, 2017.
- [14] K. Ueki, Y. Nakagome, K. Hirakawa, K. Kikuchi, Y. Hayashi, T. Ogawa, and T. Kobayashi, "Waseda Meisei at TRECVID 2018: Ad-hoc video search," in *TRECVID Workshop*, 2018.
- [15] J. C. Pereira, E. Coviello, G. Doyle, N. Rasiwasia, G. R. Lanckriet, R. Levy, and N. Vasconcelos, "On the role of correlation and abstraction in cross-modal multimedia retrieval," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 36, no. 3, pp. 521–535, 2014.
- [16] F. Wu, X. Lu, Z. Zhang, S. Yan, Y. Rui, and Y. Zhuang, "Cross-media semantic representation via bi-directional learning to rank," in *ACM Multimedia*, 2013, pp. 877–886.
- [17] N. C. Mithun, J. Li, F. Metze, and A. K. Roy-Chowdhury, "Learning joint embedding with multimodal cues for cross-modal video-text retrieval," in *ICMR*, 2018, pp. 19–27.
- [18] B. Zhang, H. Hu, and F. Sha, "Cross-modal and hierarchical modeling of video and text," in *ECCV*, 2018, pp. 374–390.
- [19] Y. Liu, S. Albanie, A. Nagrani, and A. Zisserman, "Use what you have: Video retrieval using representations from collaborative experts," *arXiv preprint arXiv:1907.13487*, 2019.
- [20] A. Miech, D. Zhukov, J.-B. Alayrac, M. Tapaswi, I. Laptev, and J. Sivic, "Howto100m: Learning a text-video embedding by watching hundred million narrated video clips," in *ICCV*, 2019, pp. 2630–2640.
- [21] X. Yang, J. Dong, Y. Cao, X. Wang, M. Wang, and T.-S. Chua, "Tree-augmented cross-modal encoding for complex-query video retrieval," in *SIGIR*, 2020, pp. 1339–1348.
- [22] S. Chen, Y. Zhao, Q. Jin, and Q. Wu, "Fine-grained video-text retrieval with hierarchical graph reasoning," in *CVPR*, 2020, pp. 10638–10647.
- [23] M. Otani, Y. Nakashima, E. Rahtu, J. Heikkilä, and N. Yokoya, "Learning joint representations of videos and sentences with web image search," in *ECCV Workshop*, 2016, pp. 651–667.
- [24] M. Wray, D. Larlus, G. Csurka, and D. Damen, "Fine-grained action retrieval through multiple parts-of-speech embeddings," in *ICCV*, 2019, pp. 450–459.
- [25] J. Dong, S. Huang, D. Xu, and D. Tao, "DI-61-86 at TRECVID 2017: Video-to-text description," in *TRECVID Workshop*, 2017.
- [26] A. Torabi, N. Tandon, and L. Sigal, "Learning language-visual embedding for movie understanding with natural-language," *arXiv preprint arXiv:1609.08124*, 2016.
- [27] M. Kratochvíl, P. Veselý, F. Mejzlík, and J. Lokoč, "Som-hunter: Video browsing with relevance-to-som feedback loop," in *MMM*. Springer, 2020, pp. 790–795.
- [28] A. Miech, I. Laptev, and J. Sivic, "Learning a text-video embedding from incomplete and heterogeneous data," *arXiv preprint arXiv:1804.02516*, 2018.
- [29] D. Shao, Y. Xiong, Y. Zhao, Q. Huang, Y. Qiao, and D. Lin, "Find and focus: Retrieve and localize video events with natural language queries," in *ECCV*, 2018, pp. 200–216.
- [30] Y. Yu, J. Kim, and G. Kim, "A joint sequence fusion model for video question answering and retrieval," in *ECCV*, 2018, pp. 471–487.
- [31] J. Dong, X. Li, and C. G. M. Snoek, "Predicting visual features from text for image and video caption retrieval," *IEEE Transactions on Multimedia*, vol. 20, no. 12, pp. 3377–3388, 2018.
- [32] X. Li, C. Xu, G. Yang, Z. Chen, and J. Dong, "W2VV++: fully deep learning for ad-hoc video search," in *ACM Multimedia*, 2019, pp. 1786–1794.
- [33] A. Paivio, "Dual coding theory: Retrospect and current status," *Canadian Journal of Psychology*, vol. 45, no. 3, p. 255, 1991.
- [34] J. Xu, T. Mei, T. Yao, and Y. Rui, "MSR-VTT: A large video description dataset for bridging video and language," in *CVPR*, 2016, pp. 5288–5296.
- [35] G. Awad, A. Butt, J. Fiscus, D. Joy, A. Delgado, M. Michel, A. F. Smeaton, Y. Graham, W. Kraaij, G. Quenot, M. Eskevich, R. Ordelman, G. J. F. Jones, and B. Huet, "TRECVID 2017: Evaluating ad-hoc and instance video search, events detection, video captioning and hyperlinking," in *TRECVID Workshop*, 2017.
- [36] G. Awad, A. Butt, K. Curtis, Y. Lee, J. Fiscus, A. Godil, D. Joy, A. Delgado, A. Smeaton, Y. Graham *et al.*, "TRECVID 2018: Benchmarking video activity detection, video captioning and matching, video storytelling linking and video search," in *TRECVID Workshop*, 2018.
- [37] X. Wang, J. Wu, J. Chen, L. Li, Y.-F. Wang, and W. Y. Wang, "Vatex: A large-scale, high-quality multilingual dataset for video-and-language research," in *ICCV*, 2019, pp. 4581–4591.
- [38] A. Rohrbach, M. R., N. Tandon, and B. Schiele, "A dataset for movie description," in *CVPR*, 2015, pp. 3202–3212.
- [39] L. Zhu and Y. Yang, "ActBERT: Learning global-local video-text representations," in *CVPR*, 2020, pp. 8746–8755.
- [40] V. Gabeur, C. Sun, K. Alahari, and C. Schmid, "Multi-modal transformer for video retrieval," in *ECCV*, 2020, pp. 1–16.
- [41] J. Dong, X. Li, C. Xu, S. Ji, Y. He, G. Yang, and X. Wang, "Dual encoding for zero-example video retrieval," in *CVPR*, 2019, pp. 9346–9355.
- [42] F. Markatopoulou, D. Galanopoulos, V. Mezaris, and I. Patras, "Query and keyframe representations for ad-hoc video search," in *ICMR*, 2017, pp. 407–411.
- [43] C. G. M. Snoek, X. Li, C. Xu, and D. C. Koelma, "University of Amsterdam and Renmin university at TRECVID 2017: Searching video, detecting events and describing video," in *TRECVID Workshop*, 2017.
- [44] R. Xu, C. Xiong, W. Chen, and J. J. Corso, "Jointly modeling deep video and compositional text to bridge vision and language in a unified framework," in *AAAI*, 2015, pp. 2346–2352.


- [45] Y. Xiong, Q. Huang, L. Guo, H. Zhou, B. Zhou, and D. Lin, "A graph-based framework to bridge movies and synopses," in *ICCV*, 2019, pp. 4592–4601.
- [46] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *NIPS*, 2017, pp. 5998–6008.
- [47] A. Habibiyan, T. Mensink, and C. G. Snoek, "VideoStory: A new multimedia embedding for few-example recognition and translation of events," in *ACM Multimedia*, 2014, pp. 17–26.
- [48] Y. Yu, H. Ko, J. Choi, and G. Kim, "End-to-end concept word detection for video captioning, retrieval, and question answering," in *CVPR*, 2017, pp. 106–111.
- [49] Y. Song and M. Soleymani, "Polysemous visual-semantic embedding for cross-modal retrieval," in *CVPR*, 2019, pp. 1979–1988.
- [50] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018.
- [51] J. Lokoć, T. Souček, P. Veselý, F. Mejzlík, J. Ji, C. Xu, and X. Li, "A W2VV++ case study with automated and interactive text-to-video retrieval," in *ACM Multimedia*, 2020, pp. 2553–2561.
- [52] R. Zhao, K. Zheng, and Z.-J. Zha, "Stacked convolutional deep encoding network for video-text retrieval," in *ICME*, 2020, pp. 1–6.
- [53] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri, "Learning spatiotemporal features with 3d convolutional networks," in *ICCV*, 2015, pp. 4489–4497.
- [54] J. Carreira and A. Zisserman, "Quo vadis, action recognition? a new model and the kinetics dataset," in *CVPR*, 2017, pp. 6299–6308.
- [55] M. Schuster and K. K. Paliwal, "Bidirectional recurrent neural networks," *IEEE Transactions on Signal Processing*, vol. 45, no. 11, pp. 2673–2681, 1997.
- [56] K. Cho, B. Van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio, "Learning phrase representations using RNN encoder-decoder for statistical machine translation," in *EMNLP*, 2014.
- [57] Y. Kim, "Convolutional neural networks for sentence classification," in *EMNLP*, 2014, pp. 1746–1751.
- [58] B. Zhou, Y. Tian, S. Sukhbaatar, A. Szlam, and R. Fergus, "Simple baseline for visual question answering," *arXiv preprint arXiv:1512.02167*, 2015.
- [59] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in *CVPR*, 2017, pp. 4700–4708.
- [60] T. Mikolov, K. Chen, G. Corrado, and J. Dean, "Efficient estimation of word representations in vector space," in *ICLR*, 2013.
- [61] F. Faghri, D. J. Fleet, J. R. Kiros, and S. Fidler, "VSE++: Improved visual-semantic embeddings," in *BMVC*, 2018, pp. 1–13.
- [62] S. Li, S. Purushotham, C. Chen, Y. Ren, and C.-C. J. Kuo, "Measuring and predicting tag importance for image retrieval," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 39, no. 12, pp. 2423–2436, 2017.
- [63] J. Kittler, "Combining classifiers: A theoretical framework," *Pattern Analysis and Applications*, vol. 1, no. 1, pp. 18–27, 1998.
- [64] S. Xie, R. Girshick, P. Dollár, Z. Tu, and K. He, "Aggregated residual transformations for deep neural networks," in *CVPR*, 2017, pp. 1492–1500.
- [65] P. Mettes, D. C. Koelma, and C. G. M. Snoek, "Shuffled ImageNet banks for video event detection and search," *ACM Transactions on Multimedia Computing, Communications, and Applications*, vol. 16, no. 2, pp. 1–21, 2020.
- [66] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *CVPR*, 2016, pp. 770–778.
- [67] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [68] D. Francis, P. A. Nguyen, B. Huët, and C.-W. Ngo, "Fusion of multimodal embeddings for ad-hoc video search," in *ICCV Workshop*, 2019.
- [69] K. S. Tai, R. Socher, and C. D. Manning, "Improved semantic representations from tree-structured long short-term memory networks," *arXiv preprint arXiv:1503.00075*, 2015.
- [70] Y. Li, Y. Song, L. Cao, J. Tetreault, L. Goldberg, A. Jaimes, and J. Luo, "TGIF: A new dataset and benchmark on animated GIF description," in *CVPR*, 2016, pp. 4641–4650.
- [71] J. Liang, J. Chen, P. Huang, X. Li, L. Jiang, Z. Lan, P. Pan, H. Fan, Q. Jin, J. Sun *et al.*, "Informedia @ Trecvid 2016," in *TRECVID Workshop*, 2016.
- [72] X. Li, J. Dong, C. Xu, J. Cao, X. Wang, and G. Yang, "Renmin University of China and Zhejiang Gongshang University at TRECVID 2018: Deep cross-modal embeddings for videotext retrieval," in *TRECVID Workshop*, 2018.
- [73] P.-Y. Huang, J. Liang, V. Vaibhav, X. Chang, and A. Hauptmann, "Informedia@TRECVID 2018: Ad-hoc video search with discrete and continuous representations," *TRECVID Workshop*, 2018.
- [74] M. Bastan, X. Shi, J. Gu, Z. Heng, C. Zhuo, D. Sng, and A. Kot, "NTU ROSE lab at TRECVID 2018: Ad-hoc video search and video to text," *TRECVID Workshop*, 2018.
- [75] D. C. Koelma and C. G. M. Snoek, "Query understanding is key for zero-example video search," in *TRECVID Workshop*, 2017.
- [76] Google Research, "Tensorflow code and pre-trained models for bert," <https://github.com/google-research/bert>.
- [77] Y. Zhu, R. Kiros, R. Zemel, R. Salakhutdinov, R. Urtasun, A. Torralba, and S. Fidler, "Aligning books and movies: Towards story-like visual explanations by watching movies and reading books," in *ICCV*, 2015, pp. 19–27.



Jianfeng Dong received the B.E. degree in software engineering from Zhejiang University of Technology in 2009, and the Ph.D. degree in computer science from Zhejiang University in 2018, all in Hangzhou, China. He is currently a Research Professor at the College of Computer and Information Engineering, Zhejiang Gongshang University, Hangzhou, China. His research interests include multimedia understanding, retrieval and recommendation. He was awarded the ACM Multimedia Grand Challenge Award in 2016. Has won a number of international competitions including the TRECVID 2016, 2017, 2018 Video-to-Text (VTT) Matching and Ranking task, the MSR Bing Image Retrieval Challenge at ACM Multimedia 2015, and so on.



Xirong Li received the B.S. and M.E. degrees from Tsinghua University, Beijing, China, in 2005 and 2007, respectively, and the Ph.D. degree from the University of Amsterdam, Amsterdam, The Netherlands, in 2012, all in computer science. He is currently an Associate Professor with the Key Lab of Data Engineering and Knowledge Engineering, Renmin University of China, Beijing, China. His research focuses on multimedia intelligence. Dr. Li was recipient of the ACM MM 2016 Grand Challenge Award, the ACM SIGMM Best Ph.D. Thesis Award 2013, the IEEE TRANSACTIONS ON MULTIMEDIA Prize Paper Award 2012, and the Best Paper Award of ACM CIVR 2010. He served as program co-chair of Multimedia Modeling 2021 and is serving as associate editor of ACM TOMM and the Multimedia Systems Journal.



Chaoxi Xu received his B.S. and M.E. degrees in Computer Science from Renmin University of China, Beijing, China, in 2017 and 2020, respectively. He is currently an assistant engineer at the Institute of Computing Technology, Chinese Academy of Sciences, Beijing, China.



Xun Yang received the Ph.D. degree in 2018 from School of Computer and Information at Hefei University of Technology, China. He is currently a postdoctoral research fellow at the NExT++ Research Center, National University of Singapore. His current research interests include information retrieval, multimedia content analysis, and computer vision. He has served as the PC member and the invited reviewer for top-tier conferences and prestigious journals including ACM MM, IJCAI, AAAI, ACM TOMM, IEEE

TNNLS, IEEE TKDE, and IEEE TCSVT.



Gang Yang received his Ph.D. degree in Innovative Life Science from University of Toyama, Toyama, Japan in 2009. He is currently an Associate Professor at School of Information, Renmin University of China, Beijing, China. His current research interests include computational intelligence, multimedia computing and machine learning.



Xun Wang received his B.S. in mechanics, Ph.D. degrees in computer science, all from Zhejiang University, Hangzhou, China, in 1990 and 2006, respectively. He is currently a professor at the School of Computer Science and Information Engineering, Zhejiang Gongshang University, China. His current research interests include mobile graphics computing, image/video processing, pattern recognition, intelligent information processing and visualization. He is a member of the IEEE and ACM, and a senior

member of CCF.



Meng Wang (F'20) received his B.E. degree and Ph.D. degree in the Special Class for the Gifted Young and the Department of Electronic Engineering and Information Science from the University of Science and Technology of China (USTC), Hefei, China, in 2003 and 2008, respectively. He is currently a Professor with the Hefei University of Technology, China. His current research interests include multimedia content analysis, computer vision, and pattern recognition. He has authored more than 200 book chapters, journal and conference papers in these areas. He is the recipient of the ACM SIGMM Rising Star Award 2014. He is an associate editor of IEEE Transactions on Knowledge and Data Engineering (IEEE TKDE), IEEE Transactions on Circuits and Systems for Video Technology (IEEE TCSVT), IEEE Transactions on Multimedia (IEEE TMM), and IEEE Transactions on Neural Networks and Learning Systems (IEEE TNNLS).

He is the recipient of the ACM SIGMM Rising Star Award 2014. He is an associate editor of IEEE Transactions on Knowledge and Data Engineering (IEEE TKDE), IEEE Transactions on Circuits and Systems for Video Technology (IEEE TCSVT), IEEE Transactions on Multimedia (IEEE TMM), and IEEE Transactions on Neural Networks and Learning Systems (IEEE TNNLS).