

Text-Video Retrieval with Disentangled Conceptualization and Set-to-Set Alignment

Peng Jin^{1,2}, Hao Li^{1,2}, Zesen Cheng^{1,2}, Jinfa Huang^{1,2}, Zhennan Wang³,
Li Yuan^{1,2,3}, Chang Liu^{4†}, Jie Chen^{1,2,3†}

¹School of Electronic and Computer Engineering, Peking University, Shenzhen, China

²AI for Science (AI4S)-Preferred Program, Peking University Shenzhen Graduate School, Shenzhen, China

³Peng Cheng Laboratory, Shenzhen, China

⁴Department of Automation and BNRist, Tsinghua University, Beijing, China

{jp21, cyanlaser, jinfahuang}@stu.pku.edu.cn, {lihao1984, yuanli-ece}@pku.edu.cn,
wangzhennan2017@email.szu.edu.cn, liuchang2022@tsinghua.edu.cn, chenj@pcl.ac.cn

Abstract

Text-video retrieval is a challenging cross-modal task, which aims to align visual entities with natural language descriptions. Current methods either fail to leverage the local details or are computationally expensive. What’s worse, they fail to leverage the heterogeneous concepts in data. In this paper, we propose the Disentangled Conceptualization and Set-to-set Alignment (DiCoSA) to simulate the conceptualizing and reasoning process of human beings. For disentangled conceptualization, we divide the coarse feature into multiple latent factors related to semantic concepts. For set-to-set alignment, where a set of visual concepts correspond to a set of textual concepts, we propose an adaptive pooling method to aggregate semantic concepts to address partial matching. In particular, since we encode concepts independently in only a few dimensions, DiCoSA is superior at efficiency and granularity, ensuring fine-grained interactions using a similar computational complexity as coarse-grained alignment. Extensive experiments on five datasets, including MSR-VTT, LSMDC, MSVD, ActivityNet, and DiDeMo, demonstrate that our method outperforms the existing state-of-the-art methods.

1 Introduction

Recent years have witnessed encouraging progress in text-video retrieval, which enables humans to associate textual concepts with video entities and vice versa [Wang *et al.*, 2021b; Jin *et al.*, 2023b]. As illustrated in Figure 1, humans perceive the cross-modal matching task by conceptualizing high-dimensional inputs from multiple modalities and reasoning with concepts to achieve partially matched set-to-set alignment. In stark contrast, machine models typically represent each modality as a perceptual whole.

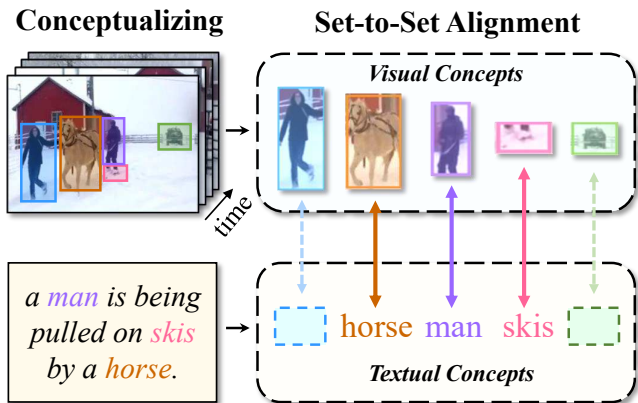


Figure 1: In text-video retrieval, a set of visual concepts corresponds to a set of textual concepts. More importantly, the semantic concepts across text and video are typically partially matched. For example, the description is not exhaustive and does not describe the visual concepts, i.e., “woman” and “car”.

Existing methods for text-video retrieval mainly focus on learning a joint feature representation space for different modalities, where text-video similarities could be measured to enable cross-modal matching. Such cross-modal representation learning methods allow for both global alignment [Liu *et al.*, 2019; Gabeur *et al.*, 2020] and local alignment [Wang *et al.*, 2022; Wray *et al.*, 2019; Chen *et al.*, 2020] between text and video feature representations. The global alignment methods exploit high-level semantics for text-video retrieval. For example, CLIP4Clip [Luo *et al.*, 2022] adopts the text-image pre-training model CLIP [Radford *et al.*, 2021] to transfer the knowledge for enhancing global representation. To leverage the local details, the local alignment methods study fine-grained semantic alignment for text-video retrieval. For instance, T2VLAD [Wang *et al.*, 2021b] shows the potential of local alignment which aligns each word and each frame individually to improve fine-grained retrieval.

However, existing methods are deficient in the following three aspects. (i) The **global alignment** (Figure 2, left) may treat discriminative regions equally and fail to capture local

[†]Corresponding author: Chang Liu, Jie Chen.

*Code is available at <https://github.com/jpthu17/DiCoSA>.

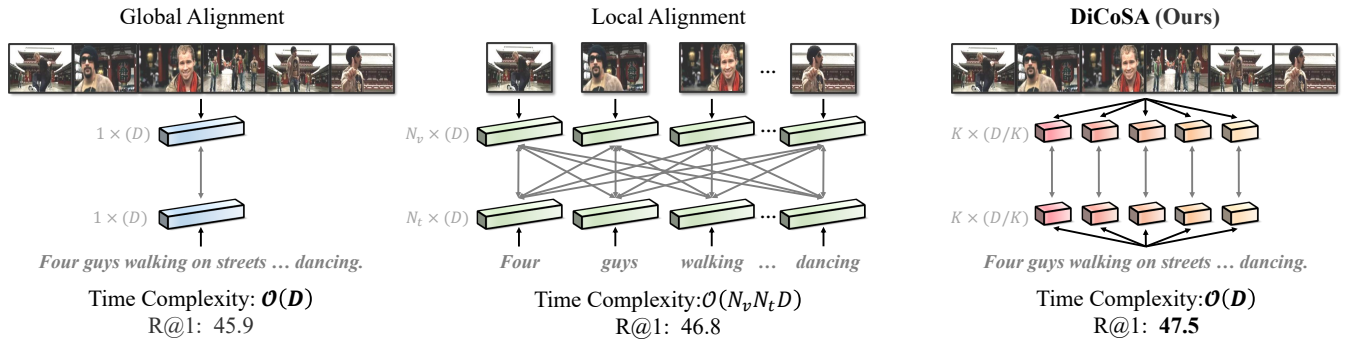


Figure 2: **Global alignment** (left) treats discriminative regions equally and fails to capture local details. **Local alignment** (middle) is computationally expensive due to the exhaustive matching operation. The previous methods typically characterize semantic concepts of texts and videos in a state of entanglement. In stark contrast, our **DiCoSA** (right) aligns texts and videos on heterogeneous semantic concepts to achieve humanlike set-to-set matching. Specifically, we divide the coarse feature into multiple latent factors related to semantic concepts (e.g., “man” and “skis”). Then we propose adaptive pooling to aggregate semantic concepts to address the partial matching in set-to-set alignment. Here, D , N_v , and N_t denote the feature dimension, the frame length, and the text length, respectively.

details in texts and videos. (ii) The **local alignment** (Figure 2, middle) aligns each word and per frame individually and is computationally expensive due to the exhaustive matching operation. (iii) **Both global and local alignments** coarsely represent the text (video) as a perceptual whole encoded by a set of concepts. As a result, they may fail to leverage the heterogeneous concepts in data and tend to focus on identifying the invariant features. As an example to illustrate, the video in Figure 1 involves a set of concepts such as “man”, “skis”, “woman”, and “horse”. To align with the text query, the video representations learned by the existing coarse approaches may continuously rely on the salient factors “man” and “horse”, and yet ignore other important factors (e.g., “woman” and “skis”) and the relations among these factors. This flaw might make these approaches stuck in the local invariant matching and misunderstand the cross-modal interaction and context.

A reasonable solution to tackle the cross-modal matching task is to align texts and videos on heterogeneous semantic concepts (Figure 2, right). The core insight is simulating the human process of conceptualizing things and reasoning on the sets of concepts. To learn explanatory and discriminative factors of variations, we adopt disentangled representations learning. A disentangled representation independently encodes information about each latent factor in only a few dimensions.

To this end, we propose the Disentangled Conceptualization and Set-to-set Alignment (DiCoSA), as shown in Figure 3. In detail, we disentangle high-dimensional coarse features into compact latent factors which explicitly encode textual semantics and visual entities. Then, we optimize latent factors from both inter-concept and intra-concept perspectives. In the inter-concept perspective, we minimize the inter-concept mutual information to find representation subspaces with minimal relevance to each other for decoupling representation. In the intra-concept perspective, we maximize the mutual information of each latent factor pair separately to align language and video within each concept. However, due to the information across modalities typically being only partially matched [Liu *et al.*, 2021], we cannot blindly leverage superficial correlations between latent factors for text-video retrieval. To address

this problem, as shown in the bottom panel of Figure 3, we propose an uncertainty-aware module to estimate the confidence of each cross-modal concept matching. Finally, we use the confidence as the weight to aggregate all factor pairs to calculate the similarity of text and video, which is called adaptive pooling.

In particular, since the dimension of the latent factor ($\mathbb{R}^{\frac{D}{K}}$) is lower than the original feature dimension (\mathbb{R}^D), our method ensures fine-grained interactions with local alignment using a similar computational complexity as global alignment. Our contributions are summarized as follows:

- We propose a disentangled conceptualization method for text-video retrieval, which divides the coarse features into multiple latent factors related to semantic concepts and achieves humanlike set-to-set matching.
- To address the partial matching of visual entities and various phrases, we propose adaptive pooling to locate mismatched cross-modal concepts and aggregate all factor pairs to calculate the similarity of text and video.
- We conduct extensive experiments on five datasets, i.e., MSR-VTT, LSMDC, MSVD, ActivityNet and DiDeMo, and achieve new state-of-the-art retrieval performance.

2 Related Work

Text-Video Retrieval. Most works [Jin *et al.*, 2022; Jin *et al.*, 2023b] of text-video retrieval is based on contrastive learning [Zhang *et al.*, 2021; Zhang *et al.*, 2022a; Zhang *et al.*, 2022b; Zhang *et al.*, 2023] to map text and video to the same semantic space. For example, CLIP-ViP [Xue *et al.*, 2022] explores factors that hinder video post-pretraining on pre-trained image-text models and effectively leverage image-text pre-trained model for post-pretraining. HBI [Jin *et al.*, 2023a] designs a new framework of multivariate interaction for cross-modal representation learning [Li *et al.*, 2022a; Li *et al.*, 2022b; Cheng *et al.*, 2023; Ye *et al.*, 2023]. DRL [Wang *et al.*, 2022] aligns each word and each frame individually to achieve token-wise interaction. However, a visual entity may relate to multiple frames, and also a frame may

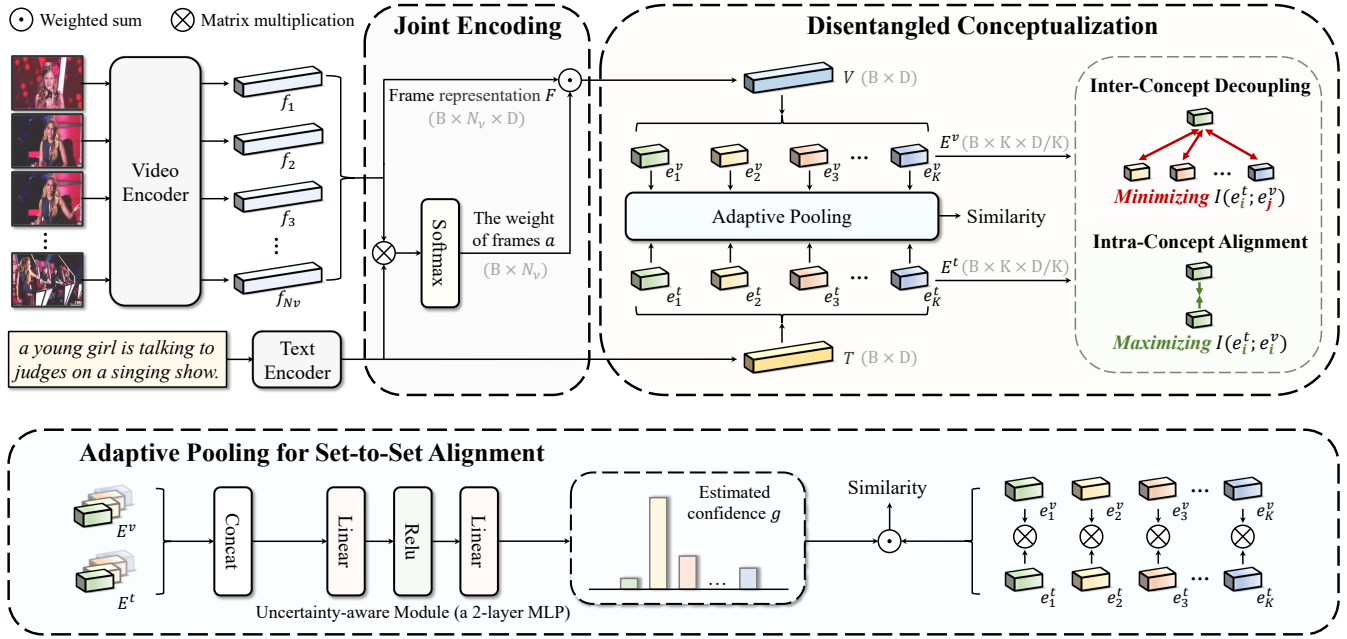


Figure 3: Our DiCoSA framework for text-video retrieval. Here, B , D , K , N_v and N_t denote the batch size, the original feature dimension, the number of concepts, the frame length, and the text length, respectively.

relate to multiple visual entities. Therefore, the previous methods cannot achieve element-level matching between semantic concepts. By contrast, our model benefits from the disentangled latent factors for element-level cross-modal matching.

Disentangled Representation. This concept was first proposed by [Bengio and others, 2009], which aims to separate the latent factors of variations behind the data. Recently, disentangled representation learning has received lots of attention [Sreekar *et al.*, 2021; Wang *et al.*, 2021a; Tran *et al.*, 2017; Suter *et al.*, 2019; Van Steenkiste *et al.*, 2019; Locatello *et al.*, 2019]. However, how to use the disentangled latent factors for fine-grained retrieval remains largely unexplored for text-video retrieval. Based on the information-theoretic view [Chen *et al.*, 2016; Do and Tran, 2019], several works provide a more detailed description [Do and Tran, 2019; Eastwood and Williams, 2018; Ridgeway and Mozer, 2018] by explicitly measuring the properties of a disentangled representation [Eastwood and Williams, 2018; Ridgeway and Mozer, 2018]. We adopt the information-theoretic definition and show that the proposed DSCA converges with semantic disentanglement.

3 Methodology

We focus on the tasks of text-to-video and video-to-text retrieval. In the text-to-video retrieval task, given a query text \mathbf{t} and candidate videos \mathbf{v} , our goal is to rank all videos according to semantic similarity. Similarly, the goal of video-to-text retrieval is to rank all candidate text based on the query video. The problem is formulated as a modality similarity measurement, where the similarity of matched text-video pairs is as high as possible and the similarity of unmatched pairs is as low as possible. Figure 3 illustrates the overall framework.

3.1 Text-Video Joint Encoding

For text representation, we adopt the text encoder of CLIP (ViT-B/32) [Radford *et al.*, 2021]. The output from the [CLS] token is taken as the text representation. For the input text \mathbf{t} , we denote the generated representation as $T \in \mathbb{R}^D$.

For video representation, we first evenly extract the frames from the video clip as the input sequence of video $\mathbf{v} = \{v_1, v_2, \dots, v_{|N_v|}\}$, where N_v denotes the frame length. Subsequently, we use ViT [Dosovitskiy *et al.*, 2021] to encode the frame sequence. Following CLIP, we adopt the output from the [CLS] token as the frame embedding. After that, we use a temporal transformer (a 4-layer transformer) to aggregate the embedding of all frames and obtain the frame representation F . Inspired by [Gorti *et al.*, 2022; Bain *et al.*, 2022], we aggregate frame representation with text as the condition. In detail, we calculate the inner product between the text representation T and frame representation $F = \{f_1, f_2, \dots, f_{N_v}\}$. We get the weight of the frames by:

$$a_i = \frac{\exp((T)^\top f_i / \tau)}{\sum_{i=1}^{N_v} \exp((T)^\top f_k / \tau)}, \quad (1)$$

where τ is the trade-off hyper-parameter. The smaller τ allows visual features to take more textual conditions into account during aggregation. The final video representation $V \in \mathbb{R}^D$ is defined as $V = \sum_{i=1}^{N_v} a_i f_i$.

3.2 Disentangled Conceptualization from Both Inter-Concept and Intra-Concept Perspectives

The above encoding methods only generate the holistic text representations T and video representations V . These representations characterize semantic concepts of the input texts

Method	Text->Video					Video->Text				
	R@1↑	R@5↑	R@10↑	MdR↓	MnR↓	R@1↑	R@5↑	R@10↑	MdR↓	MnR↓
CE [Liu <i>et al.</i> , 2019]	20.9	48.8	62.4	6.0	28.2	20.6	50.3	64.0	5.3	25.1
MMT [Gabeur <i>et al.</i> , 2020]	26.6	57.1	69.6	4.0	24.0	27.0	57.5	69.7	3.7	21.3
Support-Set [Patrick <i>et al.</i> , 2021]	30.1	58.5	69.3	3.0	-	28.5	58.6	71.6	3.0	-
T2VLAD [Wang <i>et al.</i> , 2021b]	29.5	59.0	70.1	4.0	-	31.8	60.0	71.1	3.0	-
TT-CE [Croitoru <i>et al.</i> , 2021]	29.6	61.6	74.2	3.0	-	32.1	62.7	75.0	3.0	-
FROZEN [Bain <i>et al.</i> , 2021]	31.0	59.5	70.5	3.0	-	-	-	-	-	-
CLIP4Clip [Luo <i>et al.</i> , 2022]	44.5	71.4	81.6	2.0	15.3	42.7	70.9	80.6	2.0	11.6
CLIP2Video [Fang <i>et al.</i> , 2021]	45.6	72.6	81.7	2.0	14.6	43.5	72.3	82.1	2.0	10.2
EMCL-Net [Jin <i>et al.</i> , 2022]	46.8	73.1	83.1	2.0	-	46.5	73.5	83.5	2.0	-
X-Pool [Gorti <i>et al.</i> , 2022]	46.9	72.8	82.2	2.0	14.3	44.4	73.3	84.0	2.0	9.0
TS2-Net [Liu <i>et al.</i> , 2022]	47.0	74.5	83.8	2.0	13.0	45.3	74.1	83.7	2.0	9.2
DiCoSA (Ours)	47.5	74.7	83.8	2.0	13.2	46.7	75.2	84.3	2.0	8.9

Table 1: Retrieval performance on the MSR-VTT dataset. “↑” denotes that higher is better. “↓” denotes that lower is better.

and videos in a state of entanglement. Therefore, a direct similarity matching of these representations cannot ensure adequate set-to-set cross-modal matching between a set of various phrases contained in the text description and a set of varying visual entities in the video sequence. In the following, we learn explanatory disentangled factors of variations in text and video representations to explicitly measure and understand the cross-modal relevance.

We start with the text representation $T \in \mathbb{R}^D$. Following the setting of disentangled representation learning [Ma *et al.*, 2019], we assume that each text representation is disentangled into K independent latent factors, i.e., $E^t = [e_1^t, e_2^t, \dots, e_K^t]$. Each latent factor $e_k^t \in \mathbb{R}^{\frac{D}{K}}$ represents a specific semantic concept in the text, and the independence of the latent factors ensures that those semantic concepts are not related to each other. Specifically, we independently project the text representation T into K components, and obtain the k th latent factor e_k^t as follows,

$$e_k^t = W_k^t T, \tag{2}$$

where $W_k^t \in \mathbb{R}^{\frac{D}{K} \times D}$ is trainable parameter. The latent factor e_n^v of video representation can be calculated in the same way, i.e., $e_n^v = W_n^v V$. In this way, we project features explicitly into representation subspaces corresponding to different concepts. The model is then able to optimize and reason information separately from different representation subspaces.

Inter-Concept Decoupling

In order to find representation subspaces with minimal relevance to each other and hence improve the respective discriminative power for semantic matching task, we propose to minimize the inter-concept mutual information. Given two latent factors e_i^t and e_j^v , their mutual information is defined in terms of their probabilistic density functions:

$$I(e_i^t; e_j^v) = \mathbb{E}_{t,v} \left[p(e_i^t, e_j^v) \log \frac{p(e_i^t, e_j^v)}{p(e_i^t)p(e_j^v)} \right]. \tag{3}$$

Obviously, the mutual information is hard to measure directly. To this end, we implicitly measure the mutual information via an encoder discriminator architecture. Concretely, given latent

factors $e^t \in \mathbb{R}^{\frac{D}{K}}$ and $e^v \in \mathbb{R}^{\frac{D}{K}}$, we first normalize them by the following formula:

$$z^t = \frac{e^t - \mathbb{E}[e^t]}{\sqrt{\text{Var}[e^t]}}, \quad z^v = \frac{e^v - \mathbb{E}[e^v]}{\sqrt{\text{Var}[e^v]}}, \tag{4}$$

where z^t, z^v have the same mean and standard deviation. In this manner, we scale the latent factors to the standard scale. Then, we calculate the covariance of z^t and z^v as follows,

$$C_{i,j} = \mathbb{E}_{t,v} [z_i^t z_j^v], \tag{5}$$

where z_i^t and z_j^v are the normalized features of e_i^t and e_j^v , respectively.

Lemma 1. Maximizing (minimizing) $I(e_i^t; e_j^v)$ is equivalent to maximizing (minimizing) $C_{i,j}$, i.e., $C_{i,j} \propto I(e_i^t; e_j^v)$.

Based on Lemma 1, the final inter-concept decoupling loss \mathcal{L}_D is calculated as follows,

$$\mathcal{L}_D = \sum_i \sum_{j \neq i} (C_{i,j})^2. \tag{6}$$

This loss minimizes the mutual information of negative pairs (e_i^t, e_j^v) , thereby decoupling the latent factors.

Intra-Concept Alignment

To comprehensively capture K latent factors from both text and video representations, we are required to mine semantic concepts from the text-video pairs. To this end, we optimize the representation subspace corresponding to each latent factor separately. The key insight here is that we consider each representation subspace independently, instead of text-video pairs, to comprehensively describe the latent factors and capture their relevance. Specifically, we maximize the mutual information between the text latent factor and the corresponding video latent factor within the same subspaces. The intra-concept alignment loss \mathcal{L}_A is formulated as:

$$\mathcal{L}_A = \sum_i (1 - C_{i,i})^2. \tag{7}$$

This loss maximizes the mutual information of each positive pair (e_i^t, e_i^v) separately.

Method	LSMDC					MSVD				
	R@1↑	R@5↑	R@10↑	MdR↓	MnR↓	R@1↑	R@5↑	R@10↑	MdR↓	MnR↓
FROZEN [Bain <i>et al.</i> , 2021]	15.0	30.8	39.8	20.0	-	33.7	64.7	76.3	3.0	-
CLIP4Clip [Luo <i>et al.</i> , 2022]	22.6	41.0	49.1	11.0	61.0	45.2	75.5	84.3	2.0	10.3
EMCL-Net [Jin <i>et al.</i> , 2022]	23.9	42.4	50.9	10.0	-	-	-	-	-	-
TS2-Net [Liu <i>et al.</i> , 2022]	23.4	42.3	50.9	9.0	56.9	-	-	-	-	-
X-Pool [Gorti <i>et al.</i> , 2022]	25.2	43.7	53.5	8.0	53.2	47.2	77.4	86.0	2.0	9.3
DiCoSA (Ours)	25.4	43.6	54.0	8.0	41.9	47.4	76.8	86.0	2.0	9.1

Method	ActivityNet					DiDeMo				
	R@1↑	R@5↑	R@10↑	MdR↓	MnR↓	R@1↑	R@5↑	R@10↑	MdR↓	MnR↓
CE [Liu <i>et al.</i> , 2019]	18.2	47.7	63.9	6.0	23.1	16.1	41.1	-	8.3	43.7
ClipBERT [Lei <i>et al.</i> , 2021]	21.3	49.0	63.5	6.0	-	20.4	48.0	60.8	6.0	-
TT-CE [Croitoru <i>et al.</i> , 2021]	23.5	57.2	-	4.0	-	21.6	48.6	62.9	6.0	-
CLIP4Clip [Luo <i>et al.</i> , 2022]	40.5	72.4	83.6	2.0	7.5	42.8	68.5	79.2	2.0	18.9
TS2-Net [Liu <i>et al.</i> , 2022]	41.0	73.6	84.5	2.0	8.4	41.8	71.6	82.0	2.0	14.8
DiCoSA (Ours)	42.1	73.6	84.6	2.0	6.8	45.7	74.6	83.5	2.0	11.7

Table 2: Text-to-video retrieval performance on other datasets. “↑” denotes that higher is better. “↓” denotes that lower is better.

3.3 Set-to-Set Alignment via Adaptive Pooling

However, the information across text and video is typically partially matched [Liu *et al.*, 2021]. We hence cannot directly and blindly leverage superficial correlations between latent factors for set-to-set alignment. To this end, we propose adaptive pooling to locate those mismatched cross-modal concepts and reduce their impact on the final similarity calculation.

To reveal the mismatched cross-modal latent factors, we design an uncertainty-aware module to estimate the confidence of each cross-modal concept matching. Specifically, we concatenate latent factor $e_i^t \in \mathbb{R}^{\frac{D}{K}}$ with the latent factor $e_i^v \in \mathbb{R}^{\frac{D}{K}}$ in the i_{th} subspace, generating the input data of the uncertainty-aware module $\hat{e}_i = [e_i^t, e_i^v] \in \mathbb{R}^{\frac{2D}{K}}$. Then, the confidence of the i_{th} subspace is obtained by:

$$g_i = \text{MLP}(\hat{e}_i), \quad (8)$$

where “MLP” consists of two linear layers and an activation function. Usually, small g_i indicates that the concept corresponding to the i_{th} subspace is matched with a low probability. Then, we use the confidence as the weight to aggregate all factor pairs to calculate the similarity of text and video, which is called adaptive pooling. Finally, the similarity is defined as:

$$S = \sum_{i=1}^K g_i \frac{(e_i^t)^\top e_i^v}{\|e_i^t\| \|e_i^v\|}. \quad (9)$$

3.4 Training Objective

Following common practice, we leverage InfoNCE loss [Oord *et al.*, 2018] to optimize cross-modal similarity:

$$\mathcal{L}_S = -\frac{1}{2} \left(\frac{1}{B} \sum_{l=1}^B \log \frac{\exp(S_{l,l}/\tau')}{\sum_{k=1}^B \exp(S_{l,k}/\tau')} + \frac{1}{B} \sum_{k=1}^B \log \frac{\exp(S_{k,k}/\tau')}{\sum_{l=1}^B \exp(S_{l,k}/\tau')} \right), \quad (10)$$

where B is the batch size and τ' is a pre-defined temperature prior. $S_{l,k}$ is the similarity between the l_{th} text and the k_{th} video. Combining the objective functions for cross-modal similarity \mathcal{L}_S , inter-concept decoupling \mathcal{L}_D and intra-concept alignment \mathcal{L}_A mentioned above, we get the total training loss $\mathcal{L} = \mathcal{L}_S + \alpha \mathcal{L}_D + \beta \mathcal{L}_A$, where α and β are the trade-off hyper-parameters.

3.5 Intuitive Analysis

Current methods mainly perform global alignment or local alignment. Now we explain other advantages of the proposed DiCoSA besides disentangled representation, mainly in three aspects. **(i) Humanlike set-to-set matching.** Humans perceive the world by conceptualizing high-dimensional inputs from multiple modalities such as vision and language. Through conceptualization, humans integrate things into conceptual networks and make inferences based on them. Our method simulates the human process of conceptualizing things and reasoning text and video with concepts to achieve set-to-set cross-modal matching. **(ii) Avoiding curse of dimension.** As the dimension of representation space increases, the model becomes more and more difficult to optimize, which is known as the curse of dimension [Kuo and Sloan, 2005]. Therefore, we choose to optimize the decoupled subspace ($\mathbb{R}^{\frac{D}{K}}$) at lower dimensions rather than directly learn the representation space (\mathbb{R}^D), which we will discuss in experiments (see Figure 4 (c) and (d)). **(iii) Efficiency.** Our method combines efficiency and granularity well. We calculate inference time on the MSRVT dataset, which we will discuss in experiments (see Table 4).

4 Experiments

4.1 Experiment Setup

Datasets. MSR-VTT [Xu *et al.*, 2016] contains 10,000 YouTube videos, each with 20 text descriptions. We follow the 1k-A split [Liu *et al.*, 2019] with 9,000 videos for training and 1,000 for testing. LSMDC [Rohrbach *et al.*, 2015]

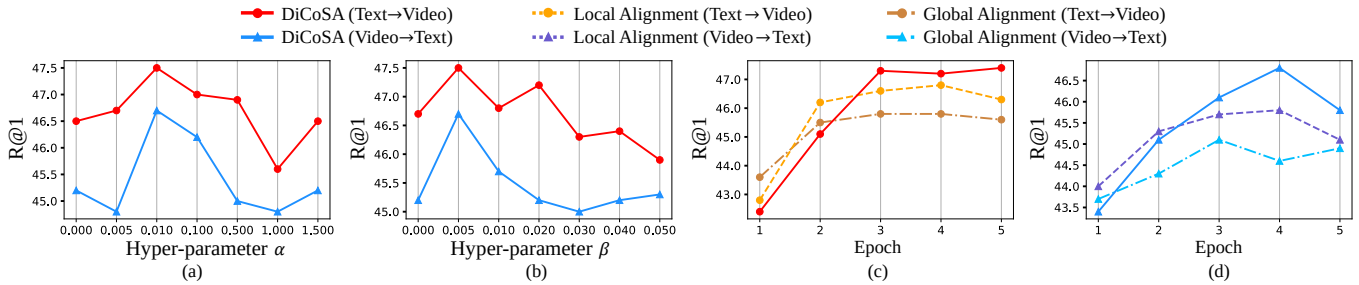


Figure 4: Ablation study on the MSRVTT dataset. Effect of (a) the trade-off hyper-parameter α ; (b) the trade-off hyper-parameter β . The comparison with the global alignment and the local alignment (c) on the text-to-video task; (d) on the video-to-text task.

Method	Text->Video			Video->Text			Method	Text->Video		
	R@1↑	R@5↑	MnR↓	R@1↑	R@5↑	MnR↓		R@1↑	R@5↑	MnR↓
Baseline	45.9	72.0	14.1	45.2	73.5	9.8	K=2	46.2	72.3	13.2
+ Latent Factors	46.2	73.9	14.0	46.0	73.6	10.2	K=4	46.7	73.0	13.0
+ Inter-Concept Decoupling \mathcal{L}_D	46.5	73.4	13.2	46.1	74.1	9.9	K=8	47.5	74.7	13.2
+ Intra-Concept Alignment \mathcal{L}_A	47.1	73.6	13.3	46.1	74.0	10.0	K=16	44.6	72.0	12.9
+ Adaptive Pooling	47.5	74.7	13.2	46.7	75.2	8.9	K=32	43.2	71.5	14.5

Table 3: (I). Ablation study for the architecture design on the MSR-VTT dataset. “Baseline” denotes the global alignment. “↑” denotes that higher is better. “↓” denotes that lower is better. (II). Ablation study for the number of concepts K on the MSR-VTT dataset.

contains 118,081 video clips from 202 movies. We follow the split of [Gabeur *et al.*, 2020] with 1,000 videos for testing. MSVD [Chen and Dolan, 2011] contains 1,970 videos. We follow the official split of 1,200 and 670 as the train and test set, respectively. ActivityNet Caption [Krishna *et al.*, 2017] contains 20K YouTube videos. We report results on the “val1” split of 10,009 and 4,917 as the train and test set. DiDeMo [Anne Hendricks *et al.*, 2017] contains 10k videos annotated 40k text descriptions. We follow the training and evaluation protocol in [Luo *et al.*, 2022].

Metrics. We choose Recall at rank L (R@L, higher is better), Median Rank (MdR, lower is better) and mean rank (MnR, lower is better) to evaluate the performance.

Implementation Details. Following previous works [Luo *et al.*, 2022; Liu *et al.*, 2022; Jin *et al.*, 2022], we utilize the CLIP (ViT-B/32) [Radford *et al.*, 2021] as the pre-trained model. The dimension of the feature is 512. The temporal transformer [Vaswani *et al.*, 2017; Li *et al.*, 2022c] is composed of 4-layer blocks, each including 8 heads and 512 hidden channels. The temporal position embedding and parameters are initialized from the CLIP’s text encoder. We use the Adam optimizer [Kingma and Ba, 2014] with a linear warmup. The initial learning rate is $1e-7$ for the text encoder and video encoder and $1e-3$ for other modules. If not otherwise specified, we set $\tau' = 0.01$, $K = 8$, $\alpha = 0.01$, $\beta = 0.005$. The network is optimized with the batch size of 128 in 5 epochs. During the inferring phase, we assume that only the candidate set is known in advance. We follow inferring schedules from [Bogolin *et al.*, 2022].

4.2 Comparisons with State-of-the-art Methods

We compare the proposed DiCoSA with other methods on five benchmarks. In Table 1, we show the results of our

method on the MSR-VTT dataset. Our model outperforms the recently proposed SOTA methods on both text-to-video retrieval and video-to-text retrieval tasks. Table 2 shows text-to-video retrieval results on the LSMDC, MSVD, ActivityNet and DiDeMo datasets. The results on all five datasets demonstrate that our method is capable of dealing with both short and long videos. DiCoSA achieves consistent improvements across different datasets, which demonstrates the effectiveness and generalization ability of our method.

4.3 Ablation Study

Architecture Design. To illustrate the importance of each part of our method, we conduct ablation experiments on the MSR-VTT dataset. From Table 3 (I), we can draw the following observations: (i) The model using latent factors achieves comparable or better performance than the baseline on the two retrieval tasks. We consider that it is because the latent factor dimension is small, thus alleviating the curse of dimension. (ii) Compared with the optimization of latent factors only from the inter-concept perspective, the improvement of the optimization from both inter-concept and intra-concept perspectives is more significant. We consider that it is because intra-concept alignment avoids the interference of mismatched cross-modal concepts in the learning of other concepts. (iii) “Adaptive pooling” can locate the mismatched cross-modal concepts and improve the respective discriminative power for set-to-set matching. Our full model achieves the best performance. This demonstrates that the four parts are beneficial for aligning visual content and textual semantics.

Effect of the Number of Concepts. The concept size K controls the number of latent factors $E = [e_1, e_2, \dots, e_K]$, where $e_k \in \mathbb{R}^D$. We start with a small size and increase it to large ones. In Table 3 (II), overall performance improves



Figure 5: Attention visualization of latent factors. We take Video0070 in the MSR-VTT dataset as an example.

Method	Complexity	Time (ms)	Text->Video					Video->Text				
			R@1↑	R@5↑	R@10↑	MdR↓	MnR↓	R@1↑	R@5↑	R@10↑	MdR↓	MnR↓
Global Alignment	$\mathcal{O}(D)$	806	45.9	72.0	81.7	2.0	14.1	45.2	73.5	82.6	2.0	9.8
Local Alignment	$\mathcal{O}(N_t N_v D)$	1158	46.8	72.6	82.6	2.0	13.4	46.4	72.2	82.3	2.0	13.4
DiCoSA (Ours)	$\mathcal{O}(D)$	978	47.5	74.7	83.8	2.0	13.2	46.7	75.2	84.3	2.0	8.9

Table 4: The comparison with the global alignment and the local alignment on the MSR-VTT dataset. “↑” denotes that higher is better. “↓” denotes that lower is better. We report the average inference time for processing the test set (1k videos and 1k text queries) using two Tesla V100 GPUs. Here, D , N_v and N_t denote the feature dimension, the frame length and the text length, respectively.

and then decreases. On the one hand, we find that fewer concepts limit the ability to leverage fine-grained information. On the other hand, a larger number of concepts reduces the dimension of each latent factor $e_k \in \mathbb{R}^{\frac{D}{K}}$, which limits the discriminability of the factors. We set the concept size $K = 8$ to achieve the best performance in practice.

Parameter Sensitivity. The parameter α indicates the importance of \mathcal{L}_D . We evaluate the scale range setting $\alpha \in [0.0, 1.5]$ as shown in Figure 4 (a). We find that R@1 is improved from 46.7% to 47.5% when $\alpha = 0.005$ and saturated with $\alpha = 0.01$ for text-to-video retrieval. As a result, we adopt $\alpha = 0.01$ to achieve the best performance. In Figure 4 (b), we show the influence of the hyper-parameter β . We evaluate the scale range setting $\beta \in [0.0, 0.05]$. We find that the model achieves the best performance at $\beta = 0.005$, so we set it as the default.

Comparisons to Other Baseline Methods. We further compare our method with other baseline methods in Table 4. Since the dimension of latent factor ($\mathbb{R}^{\frac{D}{K}}$) is lower than the original feature dimension (\mathbb{R}^D), our method introduces negligible computational overhead. Moreover, our method brings remarkable improvements by disentangled conceptualization and set-to-set alignment. We also show the training process of our method and baseline methods in Figure 4 (c) and (d). Because we optimize the decoupled subspace at a lower dimension ($\mathbb{R}^{\frac{D}{K}}$), our method has the most efficient performance improvement. We observe that the performance of our method in early epochs is lower than other methods. We consider that it is because our disentangled alignment is far removed from the CLIP pre-training objective and therefore cannot directly transfer the knowledge of CLIP in early epochs.

4.4 Qualitative Analysis

In Figure 5, we show the visualization of latent factors. In our method, each latent factor individually focuses on a specific detail. Specifically, Factor 6 focuses on “talking”; Factor 7 notices “woman”; and Factor 8 corresponds to “man”. Our latent factors are explainable to some extent. Therefore, the proposed method can be used as a tool for visualizing the cross-modal interaction and help us understand the existing retrieval model. Interestingly, we find that the factor pair corresponding to action (“talking”) has higher confidence than those corresponding to entities (“woman” and “man”). This result illustrates that the model tends to judge cross-modal similarity by actions rather than entities.

5 Conclusion

In this paper, we propose the Disentangled Conceptualization and Set-to-set Alignment (DiCoSA) for text-video retrieval. DiCoSA is superior at computation efficiency and granularity, ensuring fine-grained interactions with local alignment using a similar computational complexity as global alignment. Experimental results on five text-video retrieval benchmark datasets show the advantages of the proposed method. In the future, we hope that the disentangled features learned by our method could also be applied to other cross-modal tasks.

Acknowledgements

This work was supported in part by the National Key R&D Program of China (No. 2022ZD0118201), Natural Science Foundation of China (No. 61972217, 32071459, 62176249, 62006133, 62271465), and the Natural Science Foundation of Guangdong Province in China (No. 2019B1515120049).

References

- [Anne Hendricks *et al.*, 2017] Lisa Anne Hendricks, Oliver Wang, Eli Shechtman, Josef Sivic, Trevor Darrell, and Bryan Russell. Localizing moments in video with natural language. In *ICCV*, pages 5803–5812, 2017.
- [Bain *et al.*, 2021] Max Bain, Arsha Nagrani, Gül Varol, and Andrew Zisserman. Frozen in time: A joint video and image encoder for end-to-end retrieval. *arXiv preprint arXiv:2104.00650*, 2021.
- [Bain *et al.*, 2022] Max Bain, Arsha Nagrani, Gül Varol, and Andrew Zisserman. A clip-hitchhiker’s guide to long video retrieval. *arXiv preprint arXiv:2205.08508*, 2022.
- [Bengio and others, 2009] Yoshua Bengio et al. Learning deep architectures for ai. *Foundations and trends® in Machine Learning*, 2(1):1–127, 2009.
- [Bogolin *et al.*, 2022] Simion-Vlad Bogolin, Ioana Croitoru, Hailin Jin, Yang Liu, and Samuel Albanie. Cross modal retrieval with querybank normalisation. In *CVPR*, pages 5194–5205, 2022.
- [Chen and Dolan, 2011] David Chen and William B Dolan. Collecting highly parallel data for paraphrase evaluation. In *ACL*, pages 190–200, 2011.
- [Chen *et al.*, 2016] Xi Chen, Yan Duan, Rein Houthoofd, John Schulman, Ilya Sutskever, and Pieter Abbeel. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. *NeurIPS*, 29, 2016.
- [Chen *et al.*, 2020] Shizhe Chen, Yida Zhao, Qin Jin, and Qi Wu. Fine-grained video-text retrieval with hierarchical graph reasoning. In *CVPR*, pages 10638–10647, 2020.
- [Cheng *et al.*, 2023] Zesen Cheng, Kehan Li, Peng Jin, Xiangyang Ji, Li Yuan, Chang Liu, and Jie Chen. Parallel vertex diffusion for unified visual grounding. *arXiv preprint arXiv:2303.07216*, 2023.
- [Croitoru *et al.*, 2021] Ioana Croitoru, Simion-Vlad Bogolin, Marius Lordeanu, Hailin Jin, Andrew Zisserman, Samuel Albanie, and Yang Liu. Teachtext: Crossmodal generalized distillation for text-video retrieval. In *ICCV*, pages 11583–11593, 2021.
- [Do and Tran, 2019] Kien Do and Truyen Tran. Theory and evaluation metrics for learning disentangled representations. *arXiv preprint arXiv:1908.09961*, 2019.
- [Dosovitskiy *et al.*, 2021] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021.
- [Eastwood and Williams, 2018] Cian Eastwood and Christopher KI Williams. A framework for the quantitative evaluation of disentangled representations. In *ICLR*, 2018.
- [Fang *et al.*, 2021] Han Fang, Pengfei Xiong, Luhui Xu, and Yu Chen. Clip2video: Mastering video-text retrieval via image clip. *arXiv preprint arXiv:2106.11097*, 2021.
- [Gabeur *et al.*, 2020] Valentin Gabeur, Chen Sun, Karteek Alahari, and Cordelia Schmid. Multi-modal transformer for video retrieval. In *ECCV*, volume 12349, pages 214–229, 2020.
- [Gorti *et al.*, 2022] Satya Krishna Gorti, Noël Vouitsis, Junwei Ma, Keyvan Golestan, Maksims Volkovs, Animesh Garg, and Guangwei Yu. X-pool: Cross-modal language-video attention for text-video retrieval. In *CVPR*, pages 5006–5015, 2022.
- [Jin *et al.*, 2022] Peng Jin, Jinfa Huang, Fenglin Liu, Xian Wu, Shen Ge, Guoli Song, David Clifton, and Jie Chen. Expectation-maximization contrastive learning for compact video-and-language representations. *NeurIPS*, 35:30291–30306, 2022.
- [Jin *et al.*, 2023a] Peng Jin, Jinfa Huang, Pengfei Xiong, Shangxuan Tian, Chang Liu, Xiangyang Ji, Li Yuan, and Jie Chen. Video-text as game players: Hierarchical banzhaf interaction for cross-modal representation learning. In *CVPR*, 2023.
- [Jin *et al.*, 2023b] Peng Jin, Hao Li, Zesen Cheng, Kehan Li, Xiangyang Ji, Chang Liu, Li Yuan, and Jie Chen. Diffusion-ret: Generative text-video retrieval with diffusion model. *arXiv preprint arXiv:2303.09867*, 2023.
- [Kingma and Ba, 2014] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [Krishna *et al.*, 2017] Ranjay Krishna, Kenji Hata, Frederic Ren, Li Fei-Fei, and Juan Carlos Niebles. Dense-captioning events in videos. In *ICCV*, pages 706–715, 2017.
- [Kuo and Sloan, 2005] Frances Y Kuo and Ian H Sloan. Lifting the curse of dimensionality. *Notices of the AMS*, 52(11):1320–1328, 2005.
- [Lei *et al.*, 2021] Jie Lei, Linjie Li, Luwei Zhou, Zhe Gan, Tamara L. Berg, Mohit Bansal, and Jingjing Liu. Less is more: Clipbert for video-and-language learning via sparse sampling. In *CVPR*, pages 7331–7341, 2021.
- [Li *et al.*, 2022a] Hao Li, Jinfa Huang, Peng Jin, Guoli Song, Qi Wu, and Jie Chen. Toward 3d spatial reasoning for human-like text-based visual question answering. *arXiv preprint arXiv:2209.10326*, 2022.
- [Li *et al.*, 2022b] Hao Li, Xu Li, Belhal Karimi, Jie Chen, and Mingming Sun. Joint learning of object graph and relation graph for visual question answering. In *ICME*, pages 01–06. IEEE, 2022.
- [Li *et al.*, 2022c] Kehan Li, Runyi Yu, Zhennan Wang, Li Yuan, Guoli Song, and Jie Chen. Locality guidance for improving vision transformers on tiny datasets. In *ECCV*, pages 110–127. Springer, 2022.
- [Liu *et al.*, 2019] Yang Liu, Samuel Albanie, Arsha Nagrani, and Andrew Zisserman. Use what you have: Video retrieval using representations from collaborative experts. In *BMVC*, page 279, 2019.
- [Liu *et al.*, 2021] Yang Liu, Qingchao Chen, and Samuel Albanie. Adaptive cross-modal prototypes for cross-domain

- visual-language retrieval. In *CVPR*, pages 14954–14964, 2021.
- [Liu *et al.*, 2022] Yuqi Liu, Pengfei Xiong, Luhui Xu, Shengming Cao, and Qin Jin. Ts2-net: Token shift and selection transformer for text-video retrieval. In *ECCV*, pages 319–335. Springer, 2022.
- [Locatello *et al.*, 2019] Francesco Locatello, Michael Tschannen, Stefan Bauer, Gunnar Rätsch, Bernhard Schölkopf, and Olivier Bachem. Disentangling factors of variation using few labels. *arXiv preprint arXiv:1905.01258*, 2019.
- [Luo *et al.*, 2022] Huaishao Luo, Lei Ji, Ming Zhong, Yang Chen, Wen Lei, Nan Duan, and Tianrui Li. Clip4clip: An empirical study of clip for end to end video clip retrieval and captioning. *Neurocomputing*, 508:293–304, 2022.
- [Ma *et al.*, 2019] Jianxin Ma, Peng Cui, Kun Kuang, Xin Wang, and Wenwu Zhu. Disentangled graph convolutional networks. In *ICML*, pages 4212–4221, 2019.
- [Oord *et al.*, 2018] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- [Patrick *et al.*, 2021] Mandela Patrick, Po-Yao Huang, Yuki Asano, Florian Metze, Alexander G Hauptmann, Joao F. Henriques, and Andrea Vedaldi. Support-set bottlenecks for video-text representation learning. In *ICLR*, 2021.
- [Radford *et al.*, 2021] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *ICML*, pages 8748–8763, 2021.
- [Ridgeway and Mozer, 2018] Karl Ridgeway and Michael C Mozer. Learning deep disentangled embeddings with the f-statistic loss. *NeurIPS*, 31, 2018.
- [Rohrbach *et al.*, 2015] Anna Rohrbach, Marcus Rohrbach, Niket Tandon, and Bernt Schiele. A dataset for movie description. In *CVPR*, pages 3202–3212, 2015.
- [Sreekar *et al.*, 2021] P Aditya Sreekar, Ujjwal Tiwari, and Anoop Namboodiri. Mutual information based method for unsupervised disentanglement of video representation. In *ICPR*, pages 6396–6403, 2021.
- [Suter *et al.*, 2019] Raphael Suter, Djordje Miladinovic, Bernhard Schölkopf, and Stefan Bauer. Robustly disentangled causal mechanisms: Validating deep representations for interventional robustness. In *ICML*, pages 6056–6065, 2019.
- [Tran *et al.*, 2017] Luan Tran, Xi Yin, and Xiaoming Liu. Disentangled representation learning gan for pose-invariant face recognition. In *CVPR*, pages 1415–1424, 2017.
- [Van Steenkiste *et al.*, 2019] Sjoerd Van Steenkiste, Francesco Locatello, Jürgen Schmidhuber, and Olivier Bachem. Are disentangled representations helpful for abstract visual reasoning? *NeurIPS*, 32, 2019.
- [Vaswani *et al.*, 2017] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, volume 30, pages 5998–6008, 2017.
- [Wang *et al.*, 2021a] Tan Wang, Zhongqi Yue, Jianqiang Huang, Qianru Sun, and Hanwang Zhang. Self-supervised learning disentangled group representation as feature. *NeurIPS*, 34, 2021.
- [Wang *et al.*, 2021b] Xiaohan Wang, Linchao Zhu, and Yi Yang. T2vlad: Global-local sequence alignment for text-video retrieval. In *CVPR*, pages 5079–5088, 2021.
- [Wang *et al.*, 2022] Qiang Wang, Yanhao Zhang, Yun Zheng, Pan Pan, and Xian-Sheng Hua. Disentangled representation learning for text-video retrieval. *arXiv preprint arXiv:2203.07111*, 2022.
- [Wray *et al.*, 2019] Michael Wray, Gabriela Csurka, Diane Larlus, and Dima Damen. Fine-grained action retrieval through multiple parts-of-speech embeddings. In *ICCV*, pages 450–459, 2019.
- [Xu *et al.*, 2016] Jun Xu, Tao Mei, Ting Yao, and Yong Rui. Msr-vtt: A large video description dataset for bridging video and language. In *CVPR*, pages 5288–5296, 2016.
- [Xue *et al.*, 2022] Hongwei Xue, Yuchong Sun, Bei Liu, Jianlong Fu, Ruihua Song, Houqiang Li, and Jiebo Luo. Clip-vip: Adapting pre-trained image-text model to video-language representation alignment. *arXiv preprint arXiv:2209.06430*, 2022.
- [Ye *et al.*, 2023] Qichen Ye, Bowen Cao, Nuo Chen, Weiyuan Xu, and Yuexian Zou. Fits: Fine-grained two-stage training for knowledge-aware question answering. *arXiv preprint arXiv:2302.11799*, 2023.
- [Zhang *et al.*, 2021] Shaofeng Zhang, Feng Zhu, Junchi Yan, Rui Zhao, and Xiaokang Yang. Zero-cl: Instance and feature decorrelation for negative-free symmetric contrastive learning. In *ICLR*, 2021.
- [Zhang *et al.*, 2022a] Shaofeng Zhang, Meng Liu, Junchi Yan, Hengrui Zhang, Lingxiao Huang, Xiaokang Yang, and Pinyan Lu. M-mix: Generating hard negatives via multi-sample mixing for contrastive learning. In *KDD*, pages 2461–2470, 2022.
- [Zhang *et al.*, 2022b] Shaofeng Zhang, Lyn Qiu, Feng Zhu, Junchi Yan, Hengrui Zhang, Rui Zhao, Hongyang Li, and Xiaokang Yang. Align representations with base: A new approach to self-supervised learning. In *CVPR*, pages 16600–16609, 2022.
- [Zhang *et al.*, 2023] Shaofeng Zhang, Feng Zhu, Rui Zhao, and Junchi Yan. Patch-level contrasting without patch correspondence for accurate and dense contrastive representation learning. In *ICLR*, 2023.