

DiffusionRet: Generative Text-Video Retrieval with Diffusion Model

Peng Jin^{1,3*} Hao Li^{1,3*} Zesen Cheng^{1,3} Kehan Li^{1,3} Xiangyang Ji⁴
Chang Liu⁴ Li Yuan^{1,2,3†} Jie Chen^{1,2,3†}

¹School of Electronic and Computer Engineering, Peking University, Shenzhen, China ²Peng Cheng Laboratory, Shenzhen, China

³AI for Science (AI4S)-Preferred Program, Peking University Shenzhen Graduate School, Shenzhen, China

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

{jp21, cyanlaser, kehanli}@stu.pku.edu.cn {lihao1984, yuanli-ec}@pku.edu.cn {liuchang2022, xyji}@tsinghua.edu.cn chenj@pcl.ac.cn

Abstract

Existing text-video retrieval solutions are, in essence, discriminant models focused on maximizing the conditional likelihood, i.e., $p(\text{candidates}|\text{query})$. While straightforward, this de facto paradigm overlooks the underlying data distribution $p(\text{query})$, which makes it challenging to identify out-of-distribution data. To address this limitation, we creatively tackle this task from a generative viewpoint and model the correlation between the text and the video as their joint probability $p(\text{candidates}, \text{query})$. This is accomplished through a diffusion-based text-video retrieval framework (DiffusionRet), which models the retrieval task as a process of gradually generating joint distribution from noise. During training, DiffusionRet is optimized from both the generation and discrimination perspectives, with the generator being optimized by generation loss and the feature extractor trained with contrastive loss. In this way, DiffusionRet cleverly leverages the strengths of both generative and discriminative methods. Extensive experiments on five commonly used text-video retrieval benchmarks, including MSRVT, LSMDC, MSVD, ActivityNet Captions, and DiDeMo, with superior performances, justify the efficacy of our method. More encouragingly, without any modification, DiffusionRet even performs well in out-domain retrieval settings. We believe this work brings fundamental insights into the related fields. Code is available at <https://github.com/jpthu17/DiffusionRet>.


1. Introduction

In recent years, text-video retrieval has made significant progress, allowing humans to associate textual concepts with video entities and vice versa [67, 66]. Existing methods for video-text retrieval typically model the cross-modal interac-


*Equal contribution.

†Corresponding author: Li Yuan, Jie Chen.

Query: *A man speaks to a camera for an interview.*



(a) The joint distribution of query and candidates



(b) Diffusion model for generating joint distribution

Figure 1: **Diffusion model for text-video retrieval.** (a) We propose to model the correlation between the query and the candidates as their joint probability. (b) The diffusion model has demonstrated remarkable generative power in various fields, and due to its coarse-to-fine nature, we utilize the diffusion model for joint probability generation.

tion as discriminant models [24, 51]. Under the discriminant paradigm based on contrastive learning [53], the primary focus of mainstream methods is to improve the dense feature extractor to learn better representation. This has led to the emergence of a large number of discriminative solutions [31, 4, 50, 22], and recent advances in large-scale vision-language pre-training models [56, 41] have pushed their state-of-the-art performance even further.

However, from a probabilistic perspective, discriminant models only learn the conditional probability distribution, i.e., $p(\text{candidates}|\text{query})$. This leads to a limitation of dis-

criminator models that they fail to model the underlying data distribution [6, 48], and their latent space contains fewer intrinsic data characteristics $p(query)$, making it difficult to achieve good generalization on unseen data [48]. In contrast to discriminator models, generative models capture the joint probability distribution of the query and candidates, *i.e.*, $p(candidates, query)$, which allows them to project data into the correct latent space based on the semantic information of the data. As a typical consequence, generative models are often more generalizable and transferrable than discriminator models. Recently, generative models have made significant progress in various fields, such as generating high-quality synthetic images [19, 9, 59], natural language [10], speech [52], and music [18]. One of the most popular generative paradigms is the diffusion model [61, 29, 19], which is a type of likelihood-based model that gradually removes noise via a learned denoising model to generate the signal. In particular, the coarse-to-fine nature of the diffusion model enables it to progressively uncover the correlation between text and video, making it a promising solution for cross-modal retrieval. Therefore, we argue that it is time to rethink the current discriminator retrieval regime from a generative perspective, utilizing the diffusion model.

To this end, we propose a novel diffusion-based text-video retrieval framework, called DiffusionRet, which addresses the limitations of current discriminative solutions from a generative perspective. As shown in Fig. 1, we model the retrieval task as a process of gradually generating joint distribution from noise. Given a query and a gallery of candidates, we adopt the diffusion model to generate the joint probability distribution $p(candidates, query)$. To improve the performance of the generative model, we optimize the proposed method from both generation and discrimination perspectives. During training, the generator is optimized by common generation loss, *i.e.*, Kullback-Leibler divergence [37]. Simultaneously, the feature extractor is trained with the contrastive loss, *i.e.*, InfoNCE loss [53], to enable discriminative representation learning. In this way, DiffusionRet has both the high performance of discriminator methods and the generalization ability of generative methods.

The proposed DiffusionRet has two compelling advantages: **First**, the generative paradigm of our method makes it inherently generalizable and transferrable, enabling DiffusionRet to adapt to out-of-domain samples without requiring additional design. **Second**, the iterative refinement property of the diffusion model allows DiffusionRet to progressively enhance the retrieval results from coarse to fine. Experimental results on five benchmark datasets for text-video retrieval, including MSRVT [70], LSMDC [60], MSVD [11], ActivityNet Captions [36], and DiDeMo [2], demonstrate the advantages of DiffusionRet. To further evaluate the generalization of our method to unseen data, we propose a new out-domain retrieval task. In the out-domain retrieval task [13],

labeled visual data with paired text descriptions are available in one domain (the “source”), but no data are available in the domain of interest (the “target”). Our method not only represents a novel effort to promote generative methods for in-domain retrieval, but also provides evidence of the merits of generative approaches in challenging out-domain retrieval settings. The main contributions are as follows:

- To the best of our knowledge, we are the first to tackle the text-video retrieval from a generative viewpoint. Moreover, we are the first to adapt the diffusion model for cross-modal retrieval.
- Our method achieves new state-of-the-art performance on text-video retrieval benchmarks of MSRVT, LSMDC, MSVD, ActivityNet Captions and DiDeMo.
- More impressively, our method performs well on out-domain retrieval without any modification, which may have a significant impact on the community.

2. Related Work

Text-Video Retrieval. Text-video retrieval is one of the most popular cross-modal tasks [43, 46, 39, 15, 40]. Most existing text-video retrieval works [21, 65, 44, 45] employ a mapping technique that places both text and video inputs in the same latent space for direct similarity calculation. For instance, CLIP4Clip [51] transfers knowledge from the text-image pre-training model, *i.e.*, CLIP [56], to enhance video representation. EMCL-Net [31] improves representation capabilities by bridging the gap between text and video. HBI [32] models video-text as game players with multivariate cooperative game theory to handle the uncertainty during fine-grained semantic interaction. These methods are discriminator models and focus primarily on maximizing the conditional likelihood $p(candidates|query)$, without considering the underlying data distribution $p(query)$. As a result, identifying out-of-distribution data becomes a challenging task. In contrast, our DiffusionRet adopts a generative viewpoint to tackle the task and provides evidence of the merits of generative approaches in a challenging, out-domain retrieval. To the best of our knowledge, we are the first to tackle text-video retrieval from a generative viewpoint.

Diffusion models. Diffusion models [61, 29, 19] are a type of neural generative model that uses the stochastic diffusion process, which is based on thermodynamics. The process involves gradually adding noise to a sample from the data distribution, and then training a neural network to reverse this process by gradually removing the noise. Recent developments in diffusion models have focused on generative tasks, *e.g.*, image generation [29, 62, 19, 30, 68], natural language generation [3, 47, 26], and audio generation [55]. Some other works have attempted to adapt the




Figure 2: **Our DiffusionRet framework for generative text-video retrieval.** We model the retrieval task as a process of gradually generating joint distribution from Gaussian noise. In contrast to the prior works, which typically optimize the posterior probabilities $p(v|t) + p(t|v)$, our method builds the joint probabilities $p(v, t)$.

diffusion model for discriminant tasks, *e.g.*, image segmentation [1, 5, 8], visual grounding [16], and detection [12]. However, there are no previous works that adapt generative diffusion models for cross-modal retrieval. This work addresses this gap by modeling the correlation between text and video as their joint probability and utilizing the diffusion model to gradually generate the joint probability distribution from noise. To the best of our knowledge, we are the first to adapt the diffusion model for cross-modal retrieval.

3. Method

In this paper, we tackle the tasks of text-to-video and video-to-text retrieval. In the task of text-to-video retrieval, we are given a text query t and a gallery of videos \mathbf{V} . The goal is to rank all videos $v \in \mathbf{V}$ so that the video corresponding to the text query t is ranked as high as possible. Similarly, the goal of video-to-text retrieval is to rank all text candidates $t \in \mathbf{T}$ based on the video query v .

3.1. Existing Solutions: Discriminant Modeling

A common method for the retrieval problem is similarity learning [69]. Specifically, the caption as well as the video are represented in a common multi-dimensional embedding space, where the similarity can be calculated as the dot product of their corresponding representations. In the task of text-to-video retrieval, such methods [51, 28, 31] use the logits to calculate the posterior probability:

$$p(v|t; \theta^t, \theta^v) = \frac{\exp(f_{\theta^t}(t)^\top f_{\theta^v}(v)/\tau)}{\sum_{v' \in \mathbf{V}} \exp(f_{\theta^t}(t)^\top f_{\theta^v}(v')/\tau)}, \quad (1)$$

where τ is the temperature hyper-parameter. θ^t and θ^v are parameters of the text feature extractor and video feature extractor, respectively. Finally, existing methods rank

all video candidates based on the posterior probability $p(v|t; \theta^t, \theta^v)$. Similarly, in the video-to-text retrieval, they rank all text candidates based on $p(t|v; \theta^t, \theta^v)$. The parameters $\{\theta^t, \theta^v\}$ of text feature extractor and video feature extractor are optimized by minimizing the contrastive learning loss [53, 72, 71, 73]:

$$\theta_t^*, \theta_v^* = \arg \min_{\theta_t, \theta_v} -\frac{1}{2} \mathbb{E}_{(t,v) \in \mathcal{D}} [\log p(v|t; \theta_t, \theta_v) + \log p(t|v; \theta_t, \theta_v)], \quad (2)$$

where \mathcal{D} is a corpus of text-video pairs (t, v) . Such learning strategy is equivalent to maximizing conditional likelihood, *i.e.*, $\prod_{(t,v) \in \mathcal{D}} p(v|t) + \prod_{(t,v) \in \mathcal{D}} p(t|v)$, which is called *discriminative training* [6].

Since existing methods directly model the conditional probability distribution $p(v|t) + p(t|v)$, without considering the input distribution $p(t)$ and $p(v)$, they fail to achieve good generalization on unseen data.

3.2. DiffusionRet: Generation Modeling

Our method reformulates the retrieval task from a generative modeling perspective. As Feynman’s mantra “what I cannot create, I do not understand” shows, we argue that it is time to rethink the current discriminative retrieval regime.

Inspired by the great success of diffusion models [61, 29, 19], we adopt the diffusion model as the generator. Concretely, given a query and N candidates, our goal is to synthesize the distribution $x^{1:N} = \{x^i\}_{i=1}^N$ from Gaussian noise $\mathcal{N}(0, \mathbf{I})$. In contrast to the prior works, which typically optimize the posterior probabilities $p(v|t; \theta_t, \theta_v) + p(t|v; \theta_t, \theta_v)$, our method builds the joint probabilities:

$$x^{1:N} = p(v, t|\phi) = f_\phi(v, t, \mathcal{N}(0, \mathbf{I})), \quad (3)$$

where ϕ is the parameter of the generator. It is worth noting that the learning objective of the generator is equivalent to approximating the data distribution, *i.e.*, $\prod_{(t,v) \in \mathcal{D}} p(v,t)$, which is called *generative training* [6]. The overview of our DiffusionRet is shown in Fig. 2.

3.2.1 Text-Frame Attention Encoder

To extract the joint encoding of text and video, we propose the text-frame attention encoder, which takes text representation as query and frame representation as key and value.

Specifically, for text representation, we adopt the text encoder of CLIP (ViT-B/32) [56] and take the output from the [CLS] token as the text representation $C_t \in \mathbb{R}^D$, where D is the size of the dimension. For video frame representation, we extract the frames evenly from the video clip as the input frame sequence. Then we use ViT [20] to encode the frame sequence and adapt the output from the [CLS] token as the frame embedding. After extracting the frame embedding, we use a 4-layer transformer to aggregate the embedding of all frames and obtain the frame representation F .

Then, we aggregate frame representation with the text-frame attention encoder. Concretely, we feed the text representation $C_t \in \mathbb{R}^D$ as query and the frame representation F as key and value into an attention module. The final video representation is defined as:

$$C_v = \text{Softmax}(C_t F^\top / \tau') F, \quad (4)$$

where τ' is the trade-off hyper-parameter. The smaller τ' allows visual features to take more textual information into account when aggregated.

3.2.2 Query-Candidate Attention Denoising Network

Different from other generation tasks which only focus on the authenticity and diversity of generation, the key to the retrieval task is to mine the correspondence between the query and candidates. To this end, we propose the query-candidate attention denoising network to capture the correspondence between query and candidates in the generation process. The overview of the denoising network is shown in Fig. 3.

We start with the task of text-to-video retrieval. To elaborate, we first project the text representation $C_t \in \mathbb{R}^D$ into the query and video representation $C_v \in \mathbb{R}^{N \times D}$ into key and value, where N is the number of video candidates. The projections are formulated as:

$$\begin{aligned} Q_t &= W_{Q_t}(C_t + \text{Proj}(k)), \\ K_v &= W_{K_v}(C_v + \text{Proj}(k)), \\ V_v &= W_{V_v}(C_v + \text{Proj}(k)), \end{aligned} \quad (5)$$

where W_{Q_t} , W_{K_v} and W_{V_v} are projection matrices. ‘‘Proj’’ projects the noise level k into D dimensional embedding. To




Figure 3: **The model architecture of the denoising network.** We first leverage the text-frame attention encoder to extract the joint encoding of text and video. Then, we feed a distribution x_k of length N , as well as noise level k , and the text and video representations into the query-candidate attention network. In each sampling step, the denoising network predicts the final clean distribution \hat{x}_0 .

give more weight to the video candidates with higher joint probabilities of the previous noise level, we add the distribution x_k to the attention weight. The attention mechanism can be defined as:

$$E_t = \text{Softmax}(Q_t K_v^\top + x_k) V_v. \quad (6)$$

We treat the output E_t of the attention module as high semantic level embedding which contains textual query information. Then, we concatenate video representation C_v and embedding E_t , generating the input data of the denoising decoder $[C_v, E_t] \in \mathbb{R}^{N \times 2D}$. The denoising decoder is a multi-layer perceptron (MLP) containing a linear layer with a Relu activation function [25] for encoding features and a linear layer for calculating output distribution.

Similarly, in the video-to-text retrieval, we feed the projection of video representation as query and the projection of text representation as key and value into an attention module. The output distribution is calculated in the same way.

3.2.3 Optimization from both Generation Perspective and Discrimination Perspective

Compared to the generation methods, discriminant methods usually give good predictive performance [6]. To leverage




Figure 4: **Similarity distribution for positive and negative pairs in in-domain retrieval and out-domain retrieval.** “MSRVT T” represents the similarity distribution on the MSRVT T test setting using pre-trained models on the MSRVT T dataset. “A->B” indicates that “A” is the source domain and “B” is the target domain. For example, “MSRVT T->DiDeMo” denotes that the generalization results on unseen DiDeMo test setting using pre-trained models on the MSRVT T dataset.

the benefit of both generative and discriminative methods, we optimize the proposed generation model from both generation and discrimination perspectives.

Probabilistic Diffusion (Generation Perspective). In the generation perspective, we model the distribution $x = p(v, t|\phi)$ as the reversed diffusion process of the K -length Markov Chain. Specifically, our method learns the joint distribution of text and video by gradually denoising a variable sampled from Gaussian distribution. In a forward diffusion process $q(x_k|x_{k-1})$, noised sampled from Gaussian distribution is added to a ground truth data distribution x_0 at every noise level k :

$$q(x_k|x_{k-1}) = \mathcal{N}(x_k; \sqrt{1 - \beta_k}x_{k-1}, \beta_k\mathbf{I}),$$

$$q(x_{1:K}|x_0) = \prod_{k=1}^K q(x_k|x_{k-1}), \quad (7)$$

where β_k decides the noise schedule which gradually increases. We can sample x_k by the following formula:

$$x_k = \sqrt{\bar{\alpha}_k}x_0 + \sqrt{1 - \bar{\alpha}_k}\epsilon, \quad (8)$$

where $\bar{\alpha}_k = \prod_{i=1}^k \alpha_i$ and $\alpha_k = 1 - \beta_k$. ϵ is a noise sampled from $\mathcal{N}(0, 1)$. Instead of predicting the noise component ϵ , we follow [58] and predict the data distribution itself, i.e., $\hat{x}_0 = f_\phi(v, t, x_k)$. The training objective of the diffusion model can be defined as:

$$\mathcal{L}_G = \mathbb{E}_{(t,v) \in \mathcal{D}} [\text{KL}(x_0 \| f_\phi(v, t, x_k))]. \quad (9)$$

This loss maximizes the likelihood of $p(v, t)$ by bringing $f_\phi(v, t, x_k)$ and x_0 closer together.

Contrastive Learning (Discrimination Perspective).

In the discrimination perspective, we optimize the features which are input into the generator so that these features contain discriminant semantic information. Inspired by [65], we align representations of text and video at the token level. Specifically, we take all tokens output by the text encoder as word-level features $\{w^i\}_{i=1}^{N_t}$, where N_t is the length of the text. Frame-level features $\{f^j\}_{j=1}^{N_v}$ are all tokens output by the video encoder, where N_v is the length of the video. Then, we calculate the alignment matrix $A = [a_{ij}]^{N_t \times N_v}$, where $a_{ij} = \frac{(w^i)^\top f^j}{\|w^i\| \|f^j\|}$ is the cosine similarity between the i th word and the j th frame. The total similarity score consists of two parts: text-to-video similarity and video-to-text similarity. For the text-to-video similarity, we first calculate the maximum alignment score of the i th word as $\max_j a_{ij}$. We then take the weighted average maximum alignment score over all words. For the video-to-text similarity, we take the weighted average maximum alignment score over all frames. The total similarity score can be defined as:

$$s = \frac{1}{2} \left(\underbrace{\sum_{i=1}^{N_t} g_t^i \max_j a_{ij}}_{\text{text-to-video similarity}} + \underbrace{\sum_{j=1}^{N_v} g_v^j \max_i a_{ij}}_{\text{video-to-text similarity}} \right), \quad (10)$$

where $\{g_t^i\}_{i=1}^{N_t} = \text{Softmax}(\text{MLP}_t(\{w^i\}_{i=1}^{N_t}))$ and $\{g_v^j\}_{j=1}^{N_v} = \text{Softmax}(\text{MLP}_v(\{f^j\}_{j=1}^{N_v}))$ are the weights of the text words and video frames, respectively. Then, the contrastive

Methods	Text-to-Video						Video-to-Text					
	R@1↑	R@5↑	R@10↑	Rsum↑	MdR↓	MnR↓	R@1↑	R@5↑	R@10↑	Rsum↑	MdR↓	MnR↓
MMT [24] <small>ECCV20</small>	26.6	57.1	69.6	153.3	4.0	24.0	27.0	57.5	69.7	154.2	3.7	21.3
Support-Set [54] <small>ICLR21</small>	30.1	58.5	69.3	157.9	3.0	-	28.5	58.6	71.6	158.7	3.0	-
T2VLAD [67] <small>CVPR21</small>	29.5	59.0	70.1	158.6	4.0	-	31.8	60.0	71.1	162.9	3.0	-
TT-CE [17] <small>ICCV21</small>	29.6	61.6	74.2	165.4	3.0	-	32.1	62.7	75.0	169.8	3.0	-
FROZEN [4] <small>ICCV21</small>	31.0	59.5	70.5	161.0	3.0	-	-	-	-	-	-	-
CLIP4Clip [51] <small>Neurocomputing22</small>	44.5	71.4	81.6	197.5	2.0	15.3	42.7	70.9	80.6	194.2	2.0	11.6
X-Pool [28] <small>CVPR22</small>	46.9	72.8	82.2	201.9	2.0	14.3	44.4	73.3	84.0	201.7	2.0	9.0
TS2-Net [50] <small>ECCV22</small>	47.0	74.5	83.8	205.3	2.0	13.0	45.3	74.1	83.7	203.1	2.0	9.2
EMCL-Net [31] <small>NeurIPS22</small>	46.8	73.1	83.1	203.0	2.0	12.8	46.5	73.5	83.5	203.5	2.0	8.8
TABLE [14] [†] <small>AAAI23</small>	47.1	74.3	82.9	204.3	2.0	13.4	47.2	74.2	84.2	205.6	2.0	11.0
DiCoSA [33] <small>IJCAI23</small>	47.5	74.7	83.8	206.0	2.0	13.2	46.7	75.2	84.3	206.2	2.0	8.9
HBI [32] <small>CVPR23</small>	48.6	74.6	83.4	206.6	2.0	12.0	46.8	74.3	84.3	205.4	2.0	8.9
DiffusionRet (Ours)	49.0	75.2	82.7	206.9	2.0	12.1	47.7	73.8	84.5	206.0	2.0	8.8
+ QB-Norm [7]	48.9	75.2	83.1	207.2	2.0	12.1	49.3	74.3	83.8	207.4	2.0	8.5

Table 1: **Comparisons to current state-of-the-art methods on the MSRVTT dataset.** “↑” denotes that higher is better. “↓” denotes that lower is better. “†” denotes that the model uses additional expert features, *e.g.*, TABLE uses the object, person, scene, motion, and audio features from multiple pretrained experts. We gray out TABLE for a fair comparison.

Method	LSMDC						MSVD					
	R@1↑	R@5↑	R@10↑	Rsum↑	MdR↓	MnR↓	R@1↑	R@5↑	R@10↑	Rsum↑	MdR↓	MnR↓
FROZEN [4] <small>ICCV21</small>	15.0	30.8	39.8	85.6	20.0	-	33.7	64.7	76.3	174.7	3.0	-
TT-CE [17] <small>ICCV21</small>	17.2	36.5	46.3	100.0	13.7	-	25.4	56.9	71.3	153.6	4.0	-
CLIP4Clip [51] <small>Neurocomputing22</small>	22.6	41.0	49.1	112.7	11.0	61.0	45.2	75.5	84.3	205.0	2.0	10.3
TS2-Net [50] <small>ECCV22</small>	23.4	42.3	50.9	116.6	9.0	56.9	-	-	-	-	-	-
EMCL-Net [31] <small>NeurIPS22</small>	23.9	42.4	50.9	117.2	10.0	-	42.1	71.3	81.1	194.5	2.0	17.6
DiffusionRet (Ours)	24.4	43.1	54.3	121.8	8.0	40.7	46.6	75.9	84.1	206.6	2.0	15.7
+ QB-Norm [7]	23.9	42.7	53.6	120.2	8.0	40.7	47.9	77.2	84.8	209.9	2.0	15.6

Method	ActivityNet Captions						DiDeMo					
	R@1↑	R@5↑	R@10↑	Rsum↑	MdR↓	MnR↓	R@1↑	R@5↑	R@10↑	Rsum↑	MdR↓	MnR↓
ClipBERT [38] <small>CVPR21</small>	21.3	49.0	63.5	133.8	6.0	-	20.4	48.0	60.8	129.2	6.0	-
CLIP4Clip [51] <small>Neurocomputing22</small>	40.5	72.4	83.6	196.5	2.0	7.5	42.8	68.5	79.2	190.5	2.0	18.9
TS2-Net [50] <small>ECCV22</small>	41.0	73.6	84.5	199.1	2.0	8.4	41.8	71.6	82.0	195.4	2.0	14.8
EMCL-Net [31] <small>NeurIPS22</small>	41.2	72.7	83.6	197.5	2.0	8.6	45.3	74.2	82.3	201.8	2.0	12.3
DiCoSA [33] <small>IJCAI23</small>	42.1	73.6	84.6	200.3	2.0	6.8	45.7	74.6	83.5	203.8	2.0	11.7
HBI [32] <small>CVPR23</small>	42.2	73.0	84.6	199.8	2.0	6.6	46.9	74.9	82.7	204.5	2.0	12.1
DiffusionRet (Ours)	45.8	75.6	86.3	207.7	2.0	6.5	46.7	74.7	82.7	204.1	2.0	14.3
+ QB-Norm [7]	48.1	75.6	85.7	209.4	2.0	6.8	48.9	75.5	83.3	207.7	2.0	14.1

Table 2: **Text-to-video retrieval performance on other datasets.** “↑” denotes that higher is better. “↓” denotes that lower is better. Video-to-text retrieval performance is provided in the supplementary material.

loss [53] can be formulated as:

$$\mathcal{L}_D = -\frac{1}{2} \mathbb{E}_{(t,v) \in \mathcal{D}} \left[\log \frac{\exp(s_{t,v}/\hat{\tau})}{\sum_{v' \in \mathcal{V}} \exp(s_{t,v'}/\hat{\tau})} + \log \frac{\exp(s_{t,v}/\hat{\tau})}{\sum_{t' \in \mathcal{T}} \exp(s_{t',v}/\hat{\tau})} \right], \quad (11)$$

where $s_{t,v}$ is the similarity score between text t and video v . $\hat{\tau}$ is the temperature hyper-parameter. This loss brings semantically similar texts and videos closer together in the representation space, thus helping the diffusion model to generate the joint distribution of text and video.

4. Experiments

4.1. Experimental Settings

Datasets. MSRVTT [70] contains 10,000 YouTube videos, each with 20 text descriptions. We follow the 1k-A split [49] with 9,000 videos for training and 1,000 for testing. LSMDC [60] contains 118,081 video clips from 202 movies. We follow the split of [24] with 1,000 videos for testing. MSVD [11] contains 1,970 videos. We follow the official split of 1,200 and 670 as the train and test set, respectively. ActivityNet Captions [36] contains 20,000 YouTube videos. We report results on the “val1” split of 10,009 and 4,917 as

Loss type	R@1 \uparrow	R@5 \uparrow	R@10 \uparrow	MnR \downarrow	Sampling	R@1 \uparrow	R@5 \uparrow	R@10 \uparrow	MnR \downarrow	Schedule of β	R@1 \uparrow	R@5 \uparrow	R@10 \uparrow	MnR \downarrow
MSE	46.9	75.0	82.5	12.2	DDPM	48.9	75.2	82.6	12.1	Linear	48.7	75.1	82.8	12.3
KL	49.0	75.2	82.7	12.1	DDIM	49.0	75.2	82.7	12.1	Cosine	49.0	75.2	82.7	12.1

(a) Effect of the generation loss type.

Training strategy	R@1 \uparrow	R@5 \uparrow	R@10 \uparrow	MnR \downarrow
CLIP zero-shot	25.3	46.5	56.0	75.8
Generation (Gen)	40.2	68.6	78.8	16.1
Discrimination (Dis)	46.8	73.5	82.9	13.0
Both Gen and Dis	49.0	75.2	82.7	12.1

(d) Effect of the training strategy.

	eval	10	50	100	1000
train					
10	48.7	\times	\times	\times	\times
50	48.6	49.0	\times	\times	\times
100	48.1	48.5	48.6	\times	\times
1000	48.2	48.4	48.5	48.6	

(e) Effect of the diffusion steps on R@1.

Scale of β	R@1 \uparrow	R@5 \uparrow	R@10 \uparrow	MnR \downarrow
0.1	48.4	75.3	82.7	12.1
0.5	48.5	75.1	82.6	12.0
1.0	49.0	75.2	82.7	12.1
1.5	48.3	75.3	82.7	12.2
2.0	48.5	75.2	82.5	12.1

(f) Effect of the scale of β .

Table 3: **Ablation study on the MSRVTT dataset.** “KL” denotes Kullback-Leibler divergence. “Gen” denotes generation. “Dis” denotes discrimination. “ \times ” denotes that the evaluation process fails. Default settings are marked in blue .

the train and test set. **DiDeMo** [2] contains 10,464 videos annotated 40,543 text descriptions. We follow the training and evaluation protocol in [51].

Metrics. We choose Recall at rank K (R@K), the Sum of Recall at rank $\{1, 5, 10\}$ (Rsum), Median Rank (MnR), and mean rank (MnR) to evaluate the retrieval performance.

Implementation Details. Following previous works [51, 50, 31, 32], we utilize the CLIP (ViT-B/32) [56] as the pre-trained model. The dimension of the feature is 512. The temporal transformer [64, 42] is composed of 4-layer blocks, each including 8 heads and 512 hidden channels. The temporal position embedding and parameters are initialized from the text encoder of the CLIP. We use the Adam optimizer [34] and set the batch size to 128. The initial learning rate is $1e-7$ for the text encoder and video encoder and $1e-3$ for other modules. We set the temperature $\hat{\tau}$ to 0.01 and τ' to 1. For short video datasets, *i.e.*, MSRVTT, LSMDC, and MSVD, the word length is 32 and the frame length is 12. For long video datasets, *i.e.*, ActivityNet Captions and DiDeMo, the word length is 64 and the frame length is 64. The training is divided into two stages. In the first stage, we train the feature extractor from the discrimination perspective. In the second stage, we optimize the generator from the generation perspective. For the MSRVTT and LSMDC datasets, the experiments are carried out on 2 NVIDIA Tesla V100 GPUs. For the MSVD, ActivityNet Captions, and DiDeMo datasets, the experiments are carried out on 8 NVIDIA Tesla V100 GPUs. In both of the tasks of text-to-video and video-to-text retrieval, we assume that only the candidate sets are known in advance. In the inference phase, we consider both the distance of video and text representations in the representation space and the joint probability of video and text.

4.2. Comparison with State-of-the-art

We compare the proposed DiffusionRet with other methods on five benchmarks. In Tab. 1, we show the retrieval

results on the MSRVTT dataset. Our model outperforms the recently proposed state-of-the-art methods on both text-to-video retrieval and video-to-text retrieval tasks. Tab. 2 shows text-to-video retrieval results on the LSMDC, MSVD, ActivityNet Captions, and DiDeMo datasets. DiffusionRet achieves consistent improvements across different datasets, which demonstrates the effectiveness of our method.

4.3. Ablation Study

Generation loss type. In Tab. 3a, we compare two common generation losses, *i.e.*, mean-squared loss (MSE) and Kullback-Leibler (KL) divergence. Results demonstrate that the KL divergence achieves optimal retrieval performance. We explain that it is because KL divergence can better measure the distance between probabilities than MSE, so it is more suitable for probability generation.

Sampling strategy. Denoising diffusion probabilistic models [29] (DDPM) learn the underlying data distribution by a Markov chain as shown in Eq. 7. To accelerate the sampling process of diffusion models, denoising diffusion implicit models [62] (DDIM) formulate a Markov chain that reverses a non-Markovian perturbation process. As shown in Tab. 3b, we find that the two sampling strategies have similar performance. To speed up the sampling process, we use DDIM by default in practice.

Schedule of β . The schedule of β controls how the step size increases. We compare the linear schedule and cosine schedule in Tab. 3c. As shown in Tab. 3c, we find that the cosine schedule performs well, so we adopt the cosine schedule by default, which is the same as the default setting in the motion generation task [63].

Training strategy. In Tab. 3d, we compare different training strategies. Similar to the findings of the pioneer work [6], we find that pure discriminant training has better performance on limited data than pure generative training. The hybrid training method we proposed achieves the best

Methods	Steps	Candidate gallery size is 1,000			Candidate gallery size is 3,000			Candidate gallery size is 5,000		
		Inference Time (s)↓	GPU Memory (M)↓	R@1↑	Inference Time (s)↓	GPU Memory (M)↓	R@1↑	Inference Time (s)↓	GPU Memory (M)↓	R@1↑
CLIP4Clip [51]	×	53.20	3043	43.3	137.24	3095	37.5	244.63	3147	32.6
X-Pool [28]	×	63.02	5585	45.9	330.02	5587	31.3	833.11	5590	27.0
TS2-Net [50]	×	117.71	2803	46.5	702.65	3053	37.8	1866.18	3303	32.8
DiffusionRet (Ours)	1	59.02	3253	23.4	143.04	3747	16.6	260.67	4297	13.1
DiffusionRet (Ours)	3	59.32	3253	47.8	144.42	3747	37.8	260.96	4297	30.5
DiffusionRet (Ours)	5	59.55	3253	47.9	144.63	3747	37.8	261.11	4297	32.9
DiffusionRet (Ours)	10	59.70	3253	48.6	145.97	3747	37.9	261.69	4297	33.0
DiffusionRet (Ours)	30	60.66	3253	48.9	148.11	3747	38.1	274.49	4297	33.1
DiffusionRet (Ours)	50	62.04	3253	49.0	151.31	3747	38.2	281.27	4297	33.2

Table 4: **Evaluation of the Inference costs with a single RTX3090 GPU.** We report the R@1 metric for the text-to-video task and repeat the experiment three times to take the average value. “Steps” denotes the diffusion steps. “×” denotes that this method does not apply this parameter. “↑” denotes that higher is better. “↓” denotes that lower is better.

Method	MSRVTT				MSRVTT->DiDeMo				MSRVTT->LSMDC			
	R@1↑	R@5↑	R@10↑	MdR↓	R@1↑	R@5↑	R@10↑	MdR↓	R@1↑	R@5↑	R@10↑	MdR↓
CLIP4Clip [51] [‡] <small>Neurocomputing22</small>	43.8	70.6	81.4	2.0	31.8	57.0	66.1	4.0	15.3	31.3	40.5	21.0
EMCL-Net [31] [‡] <small>NeurIPS22</small>	47.0	72.3	82.6	2.0	30.0	56.1	65.8	4.0	16.6	29.3	36.5	24.0
DiffusionRet (Ours)	49.0	75.2	82.7	2.0	33.2	59.3	68.4	3.0	17.1	32.4	41.0	21.0

Table 5: **Text-to-video retrieval performance in out-domain retrieval settings.** “A->B” indicates that “A” is the source domain and “B” is the target domain. “[‡]” denotes our own re-implementation of baselines.

results, which indicates that our hybrid training strategy can combine the advantages of the two training methods.

The number of steps. In Tab. 3e, we explore the influence of the number of diffusion steps. Results demonstrate that the number of steps of 50 achieves optimal performance in the retrieval task, outperforming the standard value of 1000 of the image generation task [29, 19]. We consider that this is due to the simpler probability distribution in retrieval compared to the pixel distribution in natural images. Therefore, compared with the image generation task, the cross-modal retrieval task only needs fewer diffusion steps.

Scale of β . The scale of β indicates the signal-to-noise ratio of the diffusion process [12]. We evaluate the scale range [0.1, 2.0] as shown in Tab. 3f. We find that the model achieves the best performance at the scaling of 1.0, so we set the scale of β to 1.0 in practice, which is the same as the default setting in the image generation task [19].

4.4. The Efficiency of DiffusionRet


We provide comparisons of the inference time and memory consumption of our method with other methods under different conditions (the number of diffusion steps and the size of the candidate gallery) in Tab. 4. As shown in Tab. 4, our method is as efficient as the existing state-of-the-art methods during the inference stage, which we explain in the following three aspects. **(1) Lightweight denoising network.** Our denoising network is lightweight (2.50 M parameter),

while other methods use complex matching networks, *e.g.*, TS2-Net which uses the token shift transformer and token selection transformer for fine-grained alignment. **(2) Efficient feature extractor.** About 80% of the inference time is spent on feature extraction such as extracting query text features. Compared with other methods with bulky feature extractors, *e.g.*, X-Pool which uses an additional transformer-based pooling module to aggregate features, our method only uses vanilla transformer to extract features and is therefore more efficient. **(3) Scalability.** We can increase the number of diffusion steps to boost performance at a negligible time cost, indicating the scalability of our method. Compared with other methods, our method is more flexible and better suited to different retrieval scenarios where accuracy and speed are required differently.

4.5. Out-domain Retrieval

Current text-video retrieval methods are mainly evaluated on the same dataset. To further evaluate the generalization to unseen data, we conduct out-domain retrieval [13]: we first pre-train a model on one dataset (the “source”) and then measure its performance on another dataset (the “target”) that is unseen in the training. As shown in Tab. 5, we compare the proposed DiffusionRet with other baselines (*i.e.*, CLIP4Clip [51] and EMCL-Net [31]) in out-domain retrieval settings. We find that the discriminant approaches fail to migrate the performance of in-domain retrieval to

Query: A man is playing guitar and singing.



Query: A woman in black dress and a man in a black suit sits together.




Figure 5: **The visualization of the diffusion process of the probability distribution.** We highlight the ground truth in green, and show the process from randomly initialized noise input (x_{50}) to the final predicted distribution (x_0).

out-of-domain retrieval well. For example, the performance of EMCL-Net is significantly higher than that of CLIP4Clip in in-domain retrieval. However, in out-domain retrieval, the performance of EMCL-Net is slightly lower than that of CLIP4Clip. In contrast, DiffusionRet performs well for both in-domain and out-of-domain retrieval.

To further illustrate the benefits of our generative modeling, we provide the visualization of the similarity distribution in in-domain retrieval and out-domain retrieval. As shown in Fig. 4, compared to the baseline models, our DiffusionRet maximally keeps the distributions of positive and negative pairs separate from each other in both in-domain and out-domain settings. These results confirm that our method is more generalizable and transferrable for the unseen data than typical discriminant models.

4.6. Qualitative Analysis

The coarse-to-fine nature of the diffusion model enables it to progressively uncover the correlation between text and video, rendering it an effective approach for cross-modal retrieval. To better understand the diffusion process, we show the additional visualization of the diffusion process in Fig. 5. These results demonstrate that our method can progressively uncover the correlation between text and video.

4.7. Why Diffusion Models

Diffusion models have demonstrated remarkable generative power in various fields. Besides the powerful generative power of diffusion models, we explain other advantages of applying the diffusion model rather than other generative approaches to cross-modal retrieval, mainly in two aspects. **First**, the coarse-to-fine nature of the diffusion model enables it to progressively uncover the correlation between text and video, rendering it a more effective approach for

retrieval tasks than other generation training methods, such as generative adversarial network [27] and variational autoencoder [35]. **Second**, the many-to-many nature of the diffusion model makes it more suitable for generating joint probabilities than the auto-regressive networks [23, 57]. In Fig. 5, we show the visualization of the diffusion process. These results further demonstrate the above two advantages of applying the diffusion model rather than other generative approaches to cross-modal retrieval.

5. Conclusion

In this paper, we propose DiffusionRet, the first generative diffusion-based framework for text-video retrieval. By explicitly modeling the joint probability distribution of text and video, DiffusionRet shows promise to solve the intrinsic limitations of the current discriminant regime. It successfully optimizes the DiffusionRet from both the generation perspective and discrimination perspective. This makes DiffusionRet principled and well-applicable in both in-domain retrieval and out-domain retrieval settings. To the best of our knowledge, we are the first to tackle text-video retrieval from the generative viewpoint. Besides, we show that our generation modeling method is superior to existing discriminant methods in terms of performance and generalization ability. This work is the first effort to promote generative methods for text-video retrieval, which may be meaningful and helpful to the community.

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, 62202014). Li Yuan is also sponsored by CCF Tencent Open Research Fund.

References

- [1] Tomer Amit, Eliya Nachmani, Tal Shaharbany, and Lior Wolf. Segdiff: Image segmentation with diffusion probabilistic models. *arXiv preprint arXiv:2112.00390*, 2021.
- [2] 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.
- [3] Jacob Austin, Daniel D Johnson, Jonathan Ho, Daniel Tarlow, and Rianne van den Berg. Structured denoising diffusion models in discrete state-spaces. In *NeurIPS*, pages 17981–17993, 2021.
- [4] Max Bain, Arsha Nagrani, Gül Varol, and Andrew Zisserman. Frozen in time: A joint video and image encoder for end-to-end retrieval. In *ICCV*, pages 1728–1738, 2021.
- [5] Dmitry Baranchuk, Ivan Rubachev, Andrey Voynov, Valentin Khulkov, and Artem Babenko. Label-efficient semantic segmentation with diffusion models. *arXiv preprint arXiv:2112.03126*, 2021.
- [6] JM Bernardo, MJ Bayarri, JO Berger, AP Dawid, D Heckerman, AFM Smith, and M West. Generative or discriminative? getting the best of both worlds. *Bayesian Statistics*, 8(3):3–24, 2007.
- [7] Simion-Vlad Bogolin, Ioana Croitoru, Hailin Jin, Yang Liu, and Samuel Albanie. Cross modal retrieval with querybank normalisation. In *CVPR*, pages 5194–5205, 2022.
- [8] Emmanuel Asiedu Brempong, Simon Kornblith, Ting Chen, Niki Parmar, Matthias Minderer, and Mohammad Norouzi. Denoising pretraining for semantic segmentation. In *CVPR*, pages 4175–4186, 2022.
- [9] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*, 2018.
- [10] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. In *NeurIPS*, pages 1877–1901, 2020.
- [11] David Chen and William Dolan. Collecting highly parallel data for paraphrase evaluation. In *ACL*, pages 190–200, 2011.
- [12] Shoufa Chen, Peize Sun, Yibing Song, and Ping Luo. Diffusiondet: Diffusion model for object detection. *arXiv preprint arXiv:2211.09788*, 2022.
- [13] Shizhe Chen, Yida Zhao, Qin Jin, and Qi Wu. Fine-grained video-text retrieval with hierarchical graph reasoning. In *CVPR*, pages 10638–10647, 2020.
- [14] Yizhen Chen, Jie Wang, Lijian Lin, Zhongang Qi, Jin Ma, and Ying Shan. Tagging before alignment: Integrating multi-modal tags for video-text retrieval. *arXiv preprint arXiv:2301.12644*, 2023.
- [15] Zesen Cheng, Peng Jin, Hao Li, Kehan Li, Siheng Li, Xiangyang Ji, Chang Liu, and Jie Chen. Wico: Win-win cooperation of bottom-up and top-down referring image segmentation. *arXiv preprint arXiv:2306.10750*, 2023.
- [16] 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.
- [17] Ioana Croitoru, Simion-Vlad Bogolin, Marius Leordeanu, Hailin Jin, Andrew Zisserman, Samuel Albanie, and Yang Liu. Teactext: Crossmodal generalized distillation for text-video retrieval. In *ICCV*, pages 11583–11593, 2021.
- [18] Prafulla Dhariwal, Heewoo Jun, Christine Payne, Jong Wook Kim, Alec Radford, and Ilya Sutskever. Jukebox: A generative model for music. *arXiv preprint arXiv:2005.00341*, 2020.
- [19] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. In *NeurIPS*, pages 8780–8794, 2021.
- [20] 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.
- [21] Bo Fang, Chang Liu, Yu Zhou, Min Yang, Yuxin Song, Fu Li, Weiping Wang, Xiangyang Ji, Wanli Ouyang, et al. Uatvr: Uncertainty-adaptive text-video retrieval. *arXiv preprint arXiv:2301.06309*, 2023.
- [22] Han Fang, Pengfei Xiong, Luhui Xu, and Yu Chen. Clip2video: Mastering video-text retrieval via image clip. *arXiv preprint arXiv:2106.11097*, 2021.
- [23] Brendan J Frey, Geoffrey E Hinton, and Peter Dayan. Does the wake-sleep algorithm produce good density estimators? In *NeurIPS*, 1995.
- [24] Valentin Gabeur, Chen Sun, Karteek Alahari, and Cordelia Schmid. Multi-modal transformer for video retrieval. In *ECCV*, pages 214–229, 2020.
- [25] Xavier Glorot, Antoine Bordes, and Yoshua Bengio. Deep sparse rectifier neural networks. In *JMLR*, pages 315–323, 2011.
- [26] Shanshan Gong, Mukai Li, Jiangtao Feng, Zhiyong Wu, and LingPeng Kong. Diffuseq: Sequence to sequence text generation with diffusion models. *arXiv preprint arXiv:2210.08933*, 2022.
- [27] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020.
- [28] 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.
- [29] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*, pages 6840–6851, 2020.
- [30] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022.
- [31] Peng Jin, Jinfa Huang, Fenglin Liu, Xian Wu, Shen Ge, Guoli Song, David A. Clifton, and Jie Chen. Expectation-maximization contrastive learning for compact video-and-language representations. In *NeurIPS*, pages 30291–30306, 2022.
- [32] 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*, pages 2472–2482, 2023.
- [33] Peng Jin, Hao Li, Zesen Cheng, Jinfa Huang, Zhennan Wang, Li Yuan, Chang Liu, and Jie Chen. Text-video retrieval with disentangled conceptualization and set-to-set alignment. *arXiv preprint arXiv:2305.12218*, 2023.
- [34] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [35] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [36] Ranjay Krishna, Kenji Hata, Frederic Ren, Li Fei-Fei, and Juan Carlos Niebles. Dense-captioning events in videos. In *ICCV*, pages 706–715, 2017.
- [37] Solomon Kullback. *Information theory and statistics*. Courier Corporation, 1997.
- [38] 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.
- [39] Hao Li, Jinfa Huang, Peng Jin, Guoli Song, Qi Wu, and Jie Chen. Weakly-supervised 3d spatial reasoning for text-based visual question answering. *TIP*, 32:3367–3382, 2023.
- [40] Hao Li, Peng Jin, Zesen Cheng, Songyang Zhang, Kai Chen, Zhennan Wang, Chang Liu, and Jie Chen. Tg-vqa: Ternary game of video question answering. *arXiv preprint arXiv:2305.10049*, 2023.
- [41] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. *arXiv preprint arXiv:2201.12086*, 2022.
- [42] 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, 2022.
- [43] Pandeng Li, Yan Li, Hongtao Xie, and Lei Zhang. Neighborhood-adaptive structure augmented metric learning. In *AAAI*, pages 1367–1375, 2022.
- [44] Pandeng Li, Chen-Wei Xie, Hongtao Xie, Liming Zhao, Lei Zhang, Yun Zheng, Deli Zhao, and Yongdong Zhang. Momentdiff: Generative video moment retrieval from random to real. *arXiv preprint arXiv:2307.02869*, 2023.
- [45] Pandeng Li, Hongtao Xie, Jiannan Ge, Lei Zhang, Shaobo Min, and Yongdong Zhang. Dual-stream knowledge-preserving hashing for unsupervised video retrieval. In *ECCV*, pages 181–197, 2022.
- [46] Pandeng Li, Hongtao Xie, Shaobo Min, Jiannan Ge, Xun Chen, and Yongdong Zhang. Deep fourier ranking quantization for semi-supervised image retrieval. *TIP*, 31:5909–5922, 2022.
- [47] Xiang Lisa Li, John Thickstun, Ishaan Gulrajani, Percy Liang, and Tatsunori B Hashimoto. Diffusion-lm improves controllable text generation. *arXiv preprint arXiv:2205.14217*, 2022.
- [48] Chen Liang, Wenguan Wang, Jiaxu Miao, and Yi Yang. Gmm-seg: Gaussian mixture based generative semantic segmentation models. *arXiv preprint arXiv:2210.02025*, 2022.
- [49] Yang Liu, Samuel Albanie, Arsha Nagrani, and Andrew Zisserman. Use what you have: Video retrieval using representations from collaborative experts. In *BMVC*, 2019.
- [50] 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, 2022.
- [51] 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.
- [52] Aaron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew Senior, and Koray Kavukcuoglu. Wavenet: A generative model for raw audio. *arXiv preprint arXiv:1609.03499*, 2016.
- [53] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- [54] Mandela Patrick, Po-Yao Huang, Yuki Asano, Florian Metzger, Alexander G Hauptmann, Joao F. Henriques, and Andrea Vedaldi. Support-set bottlenecks for video-text representation learning. In *ICLR*, 2021.
- [55] Vadim Popov, Ivan Vovk, Vladimir Gogoryan, Tasnima Sadekova, and Mikhail Kudinov. Grad-tts: A diffusion probabilistic model for text-to-speech. In *ICML*, pages 8599–8608, 2021.
- [56] 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.
- [57] Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. Improving language understanding by generative pre-training. *OpenAI*, 2018.
- [58] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022.
- [59] Ali Razavi, Aaron Van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with vq-vae-2. In *NeurIPS*, 2019.
- [60] Anna Rohrbach, Marcus Rohrbach, Niket Tandon, and Bernt Schiele. A dataset for movie description. In *CVPR*, pages 3202–3212, 2015.
- [61] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *ICML*, pages 2256–2265, 2015.
- [62] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020.
- [63] Guy Tevet, Sigal Raab, Brian Gordon, Yonatan Shafir, Daniel Cohen-Or, and Amit H Bermano. Human motion diffusion model. *arXiv preprint arXiv:2209.14916*, 2022.
- [64] 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*, pages 5998–6008, 2017.

- [65] 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.
- [66] Wenzhe Wang, Mengdan Zhang, Runnan Chen, Guanyu Cai, Penghao Zhou, Pai Peng, Xiaowei Guo, Jian Wu, and Xing Sun. Dig into multi-modal cues for video retrieval with hierarchical alignment. In *IJCAI*, pages 1113–1121, 2021.
- [67] Xiaohan Wang, Linchao Zhu, and Yi Yang. T2vlad: Global-local sequence alignment for text-video retrieval. In *CVPR*, pages 5079–5088, 2021.
- [68] Yinhuai Wang, Jiwen Yu, and Jian Zhang. Zero-shot image restoration using denoising diffusion null-space model. *arXiv preprint arXiv:2212.00490*, 2022.
- [69] Eric Xing, Michael Jordan, Stuart J Russell, and Andrew Ng. Distance metric learning with application to clustering with side-information. In *NeurIPS*, 2002.
- [70] 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.
- [71] 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.
- [72] 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.
- [73] Shaofeng Zhang, Feng Zhu, Rui Zhao, and Junchi Yan. Patch-level contrasting without patch correspondence for accurate and dense contrastive representation learning. *arXiv preprint arXiv:2306.13337*, 2023.