

Figure 2: TV-ProxyNet converts the traditional 1-to- N relationship into N distinct 1-to-1 relationships. Each 1-to-1 relationship is unique, compelling the model to recognize the unique data patterns inherent in each pair, excluding irregular factors.

than just powerful representations. There are also inherent data issues (Lin et al. 2022; Wang et al. 2022b, 2024a,b). For instance, some text annotations are overly concise and lack the necessary distinctiveness between different videos. Meanwhile, some videos exhibit complex content changes, corresponding to different texts at various stages, as shown in Fig. 1(a). These irregularities can lead to a paradox where models trained in other dataset spaces with more precise semantic alignment are more prone to errors in TVR matching. Several research efforts have already been dedicated to expanding the matching scope of a single query. By transforming a single text query into a multi-text query, (Wang et al. 2022b) broaden the semantic scope of the query, reframing the task as a K -to- N matching. (Wang et al. 2024a) start from the feature space of text and treat text as a stochastic embedding to enrich the embedding, transforming the task into a space-to- N matching. Although these methods expand the query scope by enriching the text, they lack specific associations with videos during expansion, missing clear directions and quantifiable values for the enrichment.

Our observations highlight that: the matching algorithms are predominantly based on semantics, neglecting other mighty factors that introduce noise and lack consistent patterns across various samples. These factors include, but are not limited to: disparate levels of information richness between videos and texts, and videos that range from having overly smooth frames to those with abrupt changes, which are not consistently reflected in the associated texts. Such discrepancies lead the confusion during the alignment process, as the model struggles to discern patterns amid these irregularities. We term this challenge the “*Modeling Irregularity*” problem, where the model erroneously attempts to extract discernible patterns from semantic inconsistencies.

To alleviate the Modeling Irregularity problem, we propose a framework named **Text-Video-ProxyNet** (TV-ProxyNet). This strategic replaces the single text query with a series of **text proxies** and effectively converts the typical 1-to- N relationship of TVR into N distinct 1-to-1 relationships. Each 1-to-1 relationship is unique, compelling the model to recognize and adapt to the unique data patterns inherent in each pair, not other irregular factors, as shown in Fig. 2. In generating the text proxy, TV-ProxyNet iteratively refines the proxy to balance the textual and video information effectively. To regulate the distribution of proxies,

we control the direction and distance of each proxy through the coordinated efforts of the director and dash mechanisms. The director, in particular, is a directional vector determined by the direction leader and the query. Directed by specific cues from a director, each text proxy leverages a dash to quantitatively adjust the proximity between positive and negative samples, effectively narrowing the gap with positive samples while distancing from negative ones. With each iteration, the direction and distance to adjust the proxy relative to the original text query is assessed and the video information proportion increases. This iterative process is termed as the text proxy path. Fig. 1(b) intuitively illustrates the interactive relationships between the text proxy, the query, the dash, and the director. This approach not only streamlines the retrieval process but also enhances the precision with which semantic relationships are mapped, significantly reducing the propensity for error introduced by irregular factors. Experiments on three representative video-text retrieval benchmarks demonstrate the effectiveness of TV-ProxyNet and provide new insights for building better retrieval models. The contributions of this paper are threefold:

- We disentangle TVR from a 1-to- N relationship into N 1-to-1 relationships to alleviate the impact of irregular factors across numerous sample pairs.
- We propose a novel framework, TV-ProxyNet, to implement this N 1-to-1 relationships, where the direction and distance of text proxy relative to the query is learnable.
- Extensive experiments on MSRVT, DiDeMo and ActivityNet Captions show promising performance on various metrics while outperforms state-of-the-art methods.

Related Work

Text-Video Retrieval

Existing works can be divided into two categories: one is that using pre-trained video and text extractors to obtain offline features. Another is end-to-end training methods mostly benefiting from CLIP (Radford et al. 2021) model.

Many CLIP-based methods perform contrastive learning by learning a single holistic video feature through pooling sampled frames to align with text query embeddings. CLIP4Clip (Luo et al. 2022) obtains a video feature by mean pooling all sampled frames, simply transferring a pre-trained image-language model to the video-text domain. To capture

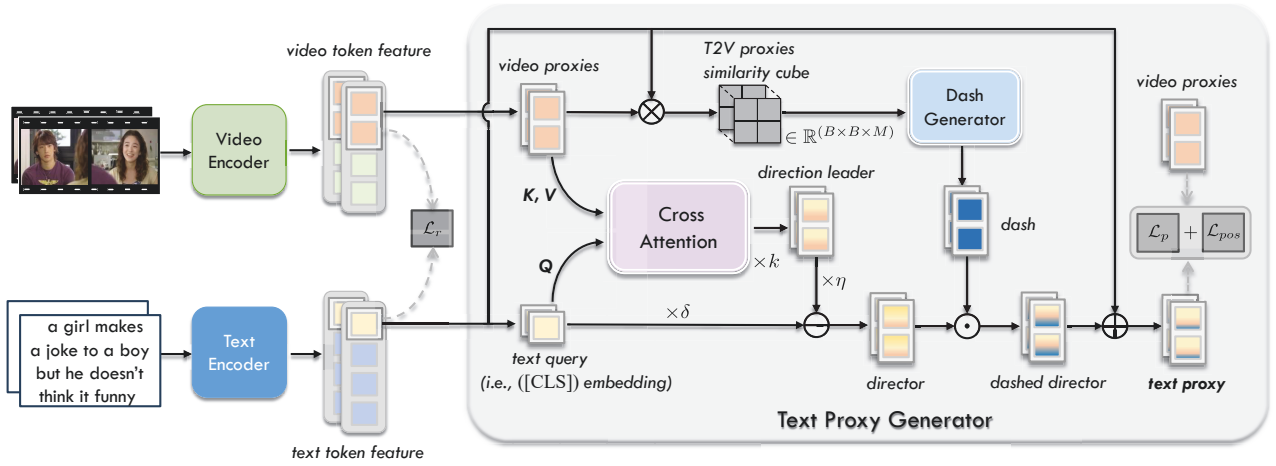


Figure 3: Framework of TV-ProxyNet. After obtaining video proxies and text query from backbone encoders, we generate direction leader to guide the direction of director and dash to specify the distance between text query and text proxy. During the training stage, \mathcal{L}_p aims to optimize the mentioned N 1-to-1 relationships while \mathcal{L}_{pos} aims to regulate the positive proxy.

higher semantic relevance, X-Pool (Gorti et al. 2022) uses cross-attention to generate text-guided video features with varying frame weights. TS2-Net (Liu et al. 2022) captures fine-grained temporal cues by shifting tokens across frames, achieving strong performance. Due to redundant information in videos, textual data often corresponds to specific segments, leading to correspondence ambiguity. TMVM (Lin et al. 2022) and ProST (Li et al. 2023) address this by adaptively aligning fine-grained prototypes. HBI (Jin et al. 2023a) proposes hierarchical Banzhaf interaction for interpretable cross-modal comparison of video frames and text words. Unlike these methods focused on multi-grained representation learning, we aim for the model to learn a text-guided proxy for the basic text query embedding, using it to assist in text-video retrieval without multi-grained alignment.

Discriminative Text Representation Learning

Current TVR methods often rely on simplistic video captioning datasets, limiting the discriminative power of text for video retrieval. Cap4Video (Wu et al. 2023) generates captions for data augmentation, enhancing text-query discrimination. (Yang et al. 2024b) improves text features in image captioning using similarity-based image retrieval (SIIR) for better caption accuracy. (Li et al. 2024) leverages CLIP to select similar image triplets, improving visual task recognition. (Yang et al. 2024a) enhances text features by configuring ICD sequences, boosting CIDEr scores in image captioning and VQA accuracy. CLIP-ViP (Xue et al. 2023) enriches CLIP by incorporating diverse language sources through large-scale video-text datasets. T-MASS (Wang et al. 2024a) extends the semantic range of a text query from a point to a mass. Unlike T-MASS, which uses stochastic text embeddings, the proposed method learns a text proxy path and utilizes the final proxy to improve text discriminability by converting the 1-to- N relationship into N 1-to-1 relationships.

Method

The training objective of TV-ProxyNet is to learn a joint space that minimizes distances between positive pairs and maximizes those between negative pairs. To alleviate the Modeling Irregularity problem, we adopt proxy-to-video retrieval to assist text-to-video retrieval. Our goal is for the proxies to have clearer distance distinctions compared to those between the original text queries and videos, thereby increasing the logit score gap between positive and negative pairs. Effective regulation of proxy distribution requires precise control of each proxy’s direction and distance, arranged by the director and dash mechanisms. The director is a directional vector determined by the direction leader and the query. The illustration of our framework is shown in Fig. 2

Framework

TV-ProxyNet consists of a backbone network and a Text Proxy Generator module. Given a raw video, we uniformly sample T frames, resulting in the video input $X = \{x_1, x_2, \dots, x_T\}$. Each frame includes $H \times W$ patch tokens, where H and W represent the number of patch token in height and width direction respectively. All patch tokens are flattened to obtain a token sequence $X = \{p_i\}_{i=1}^{T \times H \times W}$, where p_i is a patch token. After that, we adopt the same approach as CLIP-ViP (Xue et al. 2023) to obtain M learnable video proxy tokens and concatenate them to the beginning of X . The concatenated token sequence $X_{concat} = \{p_i\}_{i=1}^{M+T \times H \times W}$ is the vision input of backbone network to attend attention score calculation within a video. Also, we have text input $Y = \{t_q, t_1, t_2, \dots, t_w\}$, where t_q is text ([CLS]) token and w is the number of words. We denote video encoder and text encoder as ϕ and ψ respectively. The encoder outputs are:

$$\mathbf{X}_{concat} = \phi(X_{concat}); \mathbf{Y} = \psi(Y). \quad (1)$$

We choose the first video proxy embedding $\mathbf{p}_1 \in \mathbb{R}^{1 \times d}$ from the \mathbf{X}_{concat} as the video feature, and choose the text ([CLS])

embedding $\mathbf{t}_q \in \mathbb{R}^{1 \times d}$ from \mathbf{Y} as the text feature, where d is the dimension of joint space. After calculating the regular text-video pair contrastive loss, we use \mathbf{t}_q to generate its proxy embedding with the guidance of all M video proxy embeddings. For simplicity, we refer to the vision encoder’s output \mathbf{X}_{concat} as \mathbf{X} in the subsequent illustration.

Text Proxy Generation

The text proxy generation is associated with the direction and distance of text proxy relative to the text query.

Given a text query \mathbf{t}_q , we generate a direction leader \mathbf{d}_l with arbitrary video in a batch including N text-video pairs, preliminarily establishing N 1-to-1 relationships. Each \mathbf{d}_l specify a direction by the director vector \mathbf{d} with \mathbf{t}_q . The dash \mathcal{D}_s is generated by calculating similarities between \mathbf{t}_q and M video proxies $\tilde{\mathbf{X}} = \{\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_M\}$ and projecting them into a scalar, representing the distance between the proxy and the text query. The effect of this process is elaborated in detail in Fig. 1(b). The generation of the direction leader and dash forms the core component of TV-ProxyNet. They determine the relative position of proxy to the text query in the joint space. We formulate of generation of proxy \mathbf{t}_p as:

$$\mathbf{t}_p = \mathbf{t}_q + \mathcal{D}_s \cdot \frac{\mathbf{d}}{|\mathbf{d}|}, \quad (2)$$

where $\mathbf{t}_p, \mathbf{d} \in \mathbb{R}^{1 \times d}$, $\mathcal{D}_s \in \mathbb{R}^+$.

Director Generation. Director \mathbf{d} is the direction vector pointing from text query \mathbf{t}_q to direction leader \mathbf{d}_l or vice versa. To generate director, we need to generate direction leader first. As we iteratively generate k rounds direction leader to form a path, we denote the final direction leader on the path as \mathbf{d}_{l_k} .

Given a text query \mathbf{t}_q and N videos $\tilde{\mathbf{X}}_B = \{\tilde{\mathbf{X}}_j\}_{j=1}^N \in \mathbb{R}^{N \times (M \times d)}$, we duplicate \mathbf{t}_q into $\mathbf{D}_{l_0} = \{\mathbf{t}_q\}_{j=1}^N$ and consider \mathbf{D}_{l_0} as \mathbf{Q} , and $\tilde{\mathbf{X}}_B$ as \mathbf{K} , \mathbf{V} of cross attention for generating direction leaders $\mathbf{D}_{l_k} = \{\mathbf{d}_{l_{k,j}}\}_{j=1}^N \in \mathbb{R}^{N \times (1 \times d)}$. Next we have the i -th iteration of direction leader \mathbf{D}_{l_i} :

$$\mathbf{D}_{l_i} = \text{softmax}(\mathbf{Q}_{i-1} \cdot \mathbf{K}_{i-1}^\top) \cdot \mathbf{V}_{i-1} + \mathbf{Q}_{i-1}, \quad (3)$$

where $i \in [1, k]$, \mathbf{Q}_{i-1} and \mathbf{K}_{i-1} , \mathbf{V}_{i-1} are the $(i-1)$ -th linear transformation of $\mathbf{D}_{l_{i-1}}$ and $\tilde{\mathbf{X}}_B$ respectively, and $\mathbf{K}_{i-1}^\top \in \mathbb{R}^{N \times (d \times M)}$. This operation generates N direction leaders \mathbf{d}_{l_k} for each of text query \mathbf{t}_q in this batch, preliminarily establishing N 1-to-1 relationships.

Generating a direction leader should consider the spatial distribution of M video proxy embeddings containing various semantic information of the video. However, the position of direction leader shouldn’t be merely geometric center of M video proxy embeddings but additionally considering semantic of text query. That’s why we denote text query as \mathbf{Q} of cross attention to calculate attention score. For $\mathbf{D}_{l_{i-1}}$, it needs to be projected into \mathbf{Q}_{i-1} , which means $\mathbf{D}_{l_{i-1}}$ acts as a text guide of the generation of \mathbf{D}_{l_i} . Then the last generated direction leader \mathbf{D}_{l_k} takes responsibility to generate the director \mathbf{d} with \mathbf{t}_q .

During the training stage, direction leader \mathbf{d}_{l_k} and text query \mathbf{t}_q jointly indicate in which direction the text proxy \mathbf{t}_p should be generated relative to the text query \mathbf{t}_q . Thus we have director \mathbf{d} :

$$\mathbf{d} = \delta \mathbf{t}_q - \eta \mathbf{d}_{l_k}, \quad (4)$$

where δ and η are hyperparameters to determine the pointing direction and scale the range of absolute position of \mathbf{t}_q and \mathbf{d}_{l_k} . We scale the range to control the dominant of either one. Since the number of patch tokens under ViT-B/16 is much larger than that under ViT-B/32, directly weighting same importance on both of \mathbf{t}_q and \mathbf{d}_{l_k} lacks of specificity under ViT-B/16. Specifically, when $|\delta|$ is lower and $|\eta|$ is higher, the absolute position of \mathbf{t}_q in joint space will be proportionally scaled relative to the origin, resulting in a director \mathbf{d} relatively dominated by \mathbf{d}_{l_k} . The lower $|\delta|$, the greater the dominant role of \mathbf{d}_{l_k} , or vice versa. In a word, \mathbf{d} indicates that $\eta \mathbf{d}_{l_k}$ points at $\delta \mathbf{t}_q$ when $\delta > 0$ and $\eta > 0$. We will further discuss the influence of both of these hyperparameters in ablation study.

Dash Generation. Dash \mathcal{D}_s is the relative distance between proxy and text query. Since the substantial disparities in data distribution characteristics among diverse datasets, we adopt different dash generation options for different datasets.

Specifically, given text query \mathbf{t}_q and M video proxy embeddings $\tilde{\mathbf{X}} = \{\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_M\}$, \mathcal{D}_s is generated with the guidance of M similarities between \mathbf{t}_q and $\{\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_M\}$. Denote the similarity vector as $\mathbf{S} = [s_1, s_2, \dots, s_M]$, $s_i = s(\mathbf{t}_q, \mathbf{p}_i)$, where $s(\cdot)$ is cosine similarity function. Then we have dash formulation:

$$\mathcal{D}_s = \exp\left(\frac{1}{M} \sum_i \theta \cdot s_i\right), \quad (5)$$

where θ is a learnable scaling factor to scale the mean similarity. And this operation considers all M similarities to be of equal importance. The exponential operation projects the scalar into \mathbb{R}^+ as we only need director to specify the direction and dash to specify the distance, meanwhile, it provides higher flexibility. Also, for further enhancing the flexibility, we introduce a dash variant:

$$\mathbf{d}_s = \exp(\mathbf{S}\mathbf{W}), \quad (6)$$

where $\mathbf{d}_s \in \mathbb{R}_+^{1 \times d}$ and $\mathbf{W} \in \mathbb{R}^{M \times d}$. \mathbf{W} is a learnable linear projection matrix, aiming to project M similarities into a d -dimension vector \mathbf{d}_s , which indicates that we allow the proxy to undergo minor disturbances along a specific direction, while each similarity has different weight.

Training Objective

We adopt the baseline’s approach: select the first video proxy embedding \mathbf{p}_1 to represent the video \mathbf{X} . During the training stage, we expect the distance between the i -th video \mathbf{p}_{1_i} and the text proxy $\mathbf{t}_{p_{i,i}}$ generated by the i -th text query and the i -th video should be closer than that between the j -th video \mathbf{p}_{1_j} and the proxy $\mathbf{t}_{p_{i,j}}$, which is similar to the regular text-video pair training objective. The distinction lies in our utilization of proxy-video pair as a substitute for text-video

pair to execute contrastive learning. Next we have the proxy-video pair infoNCE loss (Oord, Li, and Vinyals 2018):

$$\mathcal{L}_p^{p2v} = -\frac{1}{B} \sum_{i=1}^B \log \frac{\exp(s(\mathbf{t}_{p_{i,i}}, \mathbf{p}_{1_i})/\sigma)}{\sum_{j=1}^B \exp(s(\mathbf{t}_{p_{i,j}}, \mathbf{p}_{1_j})/\sigma)}, \quad (7)$$

$$\mathcal{L}_p^{v2p} = -\frac{1}{B} \sum_{i=1}^B \log \frac{\exp(s(\mathbf{p}_{1_i}, \mathbf{t}_{p_{i,i}})/\sigma)}{\sum_{j=1}^B \exp(s(\mathbf{p}_{1_j}, \mathbf{t}_{p_{i,j}})/\sigma)}, \quad (8)$$

$$\mathcal{L}_p = \frac{1}{2}(\mathcal{L}_p^{p2v} + \mathcal{L}_p^{v2p}), \quad (9)$$

where σ is learnable temperature and B is batch size.

For the positive text proxy $\mathbf{t}_{p_{i,i}}$, not only the distance between $\mathbf{t}_{p_{i,i}}$ and \mathbf{p}_{1_i} should be closer than that between the negative proxy $\mathbf{t}_{p_{i,j}}$ and \mathbf{p}_{1_j} , but also the distance between $\mathbf{t}_{p_{i,i}}$ and \mathbf{p}_{1_i} should be closer than that between $\mathbf{t}_{p_{i,i}}$ and \mathbf{p}_{1_j} . Therefore, we need to establish a hard negative constraint for this type of proxy to further regulate the direction and distance of positive proxy. In this regard, we have positive proxy infoNCE loss:

$$\mathcal{L}_{pos}^{p2v} = -\frac{1}{B} \sum_{i=1}^B \log \frac{\exp(s(\mathbf{t}_{p_{i,i}}, \mathbf{p}_{1_i})/\sigma)}{\sum_{j=1}^B \exp(s(\mathbf{t}_{p_{i,i}}, \mathbf{p}_{1_j})/\sigma)}, \quad (10)$$

$$\mathcal{L}_{pos} = \frac{1}{2}(\mathcal{L}_{pos}^{p2v} + \mathcal{L}_{pos}^{v2p}). \quad (11)$$

Since we require a text query to locate the proxy, we first ensure that the text query is correctly distributed. If the text query is not properly distributed, the proxy path cannot be normalized. Therefore, we employ the infoNCE loss of regular text-video pairs:

$$\mathcal{L}_r^{t2v} = -\frac{1}{B} \sum_{i=1}^B \log \frac{\exp(s(\mathbf{t}_{q_i}, \mathbf{p}_{1_i})/\sigma)}{\sum_{j=1}^B \exp(s(\mathbf{t}_{q_j}, \mathbf{p}_{1_j})/\sigma)}, \quad (12)$$

$$\mathcal{L}_r = \frac{1}{2}(\mathcal{L}_r^{t2v} + \mathcal{L}_r^{v2t}). \quad (13)$$

Finally, the total loss is

$$\mathcal{L} = \mathcal{L}_r + \alpha \mathcal{L}_p + \beta \mathcal{L}_{pos}. \quad (14)$$

where α and β are hyperparameters that weigh the distribution characteristics of the text query and text proxy within the joint space.

Inference Pipeline

During the inference stage, we employ proxy-video logit score $s(\mathbf{t}_{p_{i,j}}, \mathbf{p}_{1_j})$ to remedy text-video logit score $s(\mathbf{t}_{q_i}, \mathbf{p}_{1_j})$. This alters the logit scores for both positive and negative text-video sample pairs, aiming to enhance the likelihood of retrieving the ground-truth videos. Thus, the final retrieval logit score is $s(\mathbf{t}_{q_i}, \mathbf{p}_{1_j}) + \gamma s(\mathbf{t}_{p_{i,j}}, \mathbf{p}_{1_j})$, where $\gamma \in [0.1, 0.8]$ is a weight coefficient. Notice that this logit score is no longer a simple cosine similarity but rather a linear combination of different cosine similarities. An ablation study of γ , a more detailed analysis of the formula derivation and visualization are provided in the appendix.

It's crucial to note that we avoid computing the logit score $s(\mathbf{t}_{p_{i,i}}, \mathbf{p}_{1_j})$ for remedying the logit score of (i, j) -th text-video pair, as this would constitute cheating. This is because the positive proxy $\mathbf{t}_{p_{i,i}}$ is derived from the i -th video, inevitably leading to a higher similarity.

Experiment

Datasets and Implementation

We employ three TVR benchmarks for evaluation: (1) **MSRVTT** (Xu et al. 2016) consists of 10k videos and each video has 20 captions. We train 9k train+val videos and test on a test set of 1k text-video pairs. (2) **DiDeMo** (Anne Hendricks et al. 2017) contains 10K Flickr videos and 40K captions. All caption descriptions of a video are concatenated as a query to perform paragraph-video retrieval. (3) **ActivityNet Captions** (Krishna et al. 2017) contains 20K YouTube videos annotated with 100K sentences. We follow the paragraph-video retrieval setting to train models on 10K videos and report results on the val set with 4.9K videos.

We employ CLIP-ViP (Xue et al. 2023) (both ViT-B/32 and ViT-B/16) as our backbone. All videos are resized to 224×224 . We set frame number to 12 for a fair comparison except for ActivityNet (set to 32) as its videos are much longer. The text proxy generator module consists of k Transformer layers with a single attention head and a linear dash generator. We finally set k to 2 to perform two rounds direction leader generation. For learning rate, we set backbone to $1e-6$ and text proxy generator module to $2e-6$ for MSRVTT, $2e-6$ and $4e-6$ for DiDeMo and ActivityNet. We optimize the model for 5 epochs for MSRVTT on 2 RTX 4090 GPUs, 20 epochs for DiDeMo on 1 RTX 4090 GPU and ActivityNet on an A800 GPU, employ AdamW optimizer with weight decay 0.2. The training batchsize is 128 for MSRVTT, 64 for DiDeMo and ActivityNet. The α and β hyperparameters are (0.5, 0.25) for MSRVTT, (1, 0.25) for DiDeMo and (0.75, 0.25) for ActivityNet Captions. The dimension of proxy is 512 same as the backbone's output.

Comparison with State-of-the-art Methods

We compare the performance of TV-ProxyNet with recent state-of-the-art methods and list the results in Tab. 1 and Tab. 2. We have achieved the state-of-the-art performance on certain metrics of two datasets. On MSRVTT, TV-ProxyNet based on ViT-B/32 improves the R@1 metric by 2.2% (52.3% v.s. 50.1%) compared to CLIP-ViP, while based on ViT-B/16, it improves the R@5 metric by 3.2% (80.4% v.s. 77.2%). Meanwhile, under DiDeMo, although we did not exceed T-MASS, we gain the improvement of 2.0% (50.6% v.s. 48.6%) under ViT-B/32 compared to CLIP-ViP. As shown in Tab. 2, under ActivityNet Captions, TV-ProxyNet has improved by 1.9% (53.0% v.s. 51.1%) under ViT-B/32 on R@1 metric compared to the baseline while under ViT-B/16 it has improved by 3.3% (56.7% v.s. 53.4%). It can be seen that the proposed method has a certain degree of continuous performance improvement on different datasets. It should be explained that due to the limitation of experimental resources, we cannot use the learning rate ($1e-6$) provided in the baseline configuration file to reproduce the DiDeMo results provided in the baseline article. This may be due to the large batch size of a single GPU and the small learning rate. Therefore, we changed the learning rate of the backbone to $2e-6$ (batch size 64 on ViT-B/32) on DiDeMo to obtain results close to the baseline. The reproduction result on DiDeMo is R@1=48.3, which indicates that we can achieve

Method	MSRVTT Retrieval					DiDeMo Retrieval				
	R@1	R@5	R@10	MdR↓	MnR↓	R@1	R@5	R@10	MdR↓	MnR↓
<i>CLIP-ViT-B/32</i>										
EMCL (Jin et al. 2022)	46.8	73.1	83.1	2.0	-	-	-	-	-	-
X-Pool (Gorti et al. 2022)	46.9	72.8	82.2	2.0	14.3	44.6	73.2	82.0	2.0	15.4
TS2-Net (Liu et al. 2022)	47.0	74.2	83.3	2.0	13.6	41.8	71.6	82.0	2.0	14.8
UATVR (Fang et al. 2023)	47.5	73.9	83.5	2.0	12.3	43.1	71.8	82.3	2.0	15.1
ProST (Li et al. 2023)	48.2	74.6	83.4	2.0	12.4	44.9	72.7	82.7	2.0	13.7
HBI (Jin et al. 2023a)	48.6	74.6	83.4	2.0	12.0	46.9	74.9	82.7	2.0	12.1
Cap4Video‡ (Wu et al. 2023)	49.3	74.3	83.8	2.0	12.0	52.0	79.4	87.5	1.0	10.5
CLIP-ViP‡ (Xue et al. 2023)	50.1	74.8	84.6	1.0	-	48.6	77.1	84.4	2.0	-
T-MASS (Wang et al. 2024a)	<u>50.2</u>	<u>75.3</u>	<u>85.1</u>	1.0	<u>11.9</u>	<u>50.9</u>	<u>77.2</u>	85.3	1.0	12.1
TV-ProxyNet (Ours)	52.3	77.8	85.8	1.0	11.1	50.6	76.9	<u>86.0</u>	1.0	<u>11.5</u>
<i>CLIP-ViT-B/16</i>										
X-Pool (Gorti et al. 2022)	48.2	73.7	82.6	2.0	12.7	47.3	74.8	82.8	2.0	14.2
ProST (Li et al. 2023)	49.5	75.0	84.0	2.0	11.7	47.5	75.5	84.4	2.0	12.3
UATVR (Fang et al. 2023)	50.8	76.3	85.5	1.0	12.4	45.8	73.7	83.3	2.0	13.5
CLIP-ViP‡ (Xue et al. 2023)	<u>54.2</u>	<u>77.2</u>	84.8	1.0	-	50.5	78.4	87.1	1.0	-
Cap4Video‡ (Wu et al. 2023)	51.4	<u>75.7</u>	83.9	1.0	12.4	-	-	-	-	-
T-MASS (Wang et al. 2024a)	52.7	77.1	<u>85.6</u>	1.0	<u>10.5</u>	53.3	<u>80.1</u>	<u>87.7</u>	1.0	<u>9.8</u>
TV-ProxyNet (Ours)	55.2	80.4	86.8	1.0	9.3	<u>52.1</u>	80.4	88.7	1.0	8.7

Table 1: Text-to-video Retrieval comparisons on MSRVTT and DiDeMo. Bold denotes the best performance. Underline denotes the second performance. “-” result is unavailable. ‡ denotes the method using extra data.

Method	R@1	R@5	R@10	MdR
<i>CLIP-ViT-B/32</i>				
CLIP4Clip (Luo et al. 2022)	40.5	72.4	-	2.0
EMCL (Jin et al. 2022)	41.2	72.7	-	2.0
HBI (Jin et al. 2023a)	42.2	73.0	84.6	2.0
CenterCLIP (Zhao et al. 2022)	43.9	74.6	85.8	2.0
CLIP2TV (Gao et al. 2021)	44.1	75.2	-	2.0
DRL (Wang et al. 2022a)	44.2	74.5	86.1	2.0
CAMoE* (Cheng et al. 2021)	51.0	77.7	-	-
CLIP-ViP‡ (Xue et al. 2023)	<u>51.1</u>	<u>78.4</u>	<u>88.3</u>	1.0
TV-ProxyNet (Ours)	53.0	80.9	89.6	1.0
<i>CLIP-ViT-B/16</i>				
CenterCLIP (Zhao et al. 2022)	46.2	77.0	87.6	2.0
DRL (Wang et al. 2022a)	46.2	77.3	88.2	2.0
CLIP-ViP‡ (Xue et al. 2023)	<u>53.4</u>	<u>81.4</u>	<u>90.0</u>	1.0
TV-ProxyNet (Ours)	56.7	83.1	91.1	1.0

Table 2: Text-to-video Retrieval comparisons on ActivityNet Captions. Bold denotes the best performance. Underline denotes the second performance. “-” result is unavailable. * denotes that the use of DSL as post-processing operation. ‡ denotes the method using extra data.

a performance improvement of 2.3% (50.6% vs.48.3%) on this dataset using text proxy.

Compared to T-MASS, which randomly generated text embedding with a mass range, TV-ProxyNet improves the R@1 metrics on ViT-B/32 and ViT-B/16 by 2.1% and 2.0% on MSRVTT respectively. It is worthy noting that, T-MASS still establishes a 1-to- N relationship for a given text query in the training stage. Instead, TV-ProxyNet aims to exaggerate the relative distances between a given text query and all

\mathcal{L}_r	\mathcal{L}_p	\mathcal{L}_{pos}	R@1	R@5	R@10	MnR
✓			50.1	74.8	84.6	-
✓	✓		52.0	77.7	85.6	11.1
✓		✓	50.6	76.5	85.2	11.3
✓	✓	✓	52.3	77.8	85.8	11.1

Table 3: The ablation study of different losses on MSRVTT-9k. Row 1 only using \mathcal{L}_r refers to baseline.

videos through the distribution of the proxies, thereby improving the likelihood of retrieving the target video, which is implemented by establishing N 1-to-1 relationships.

Ablation Studies

We conducted the ablation experiments on different variables of the Texy Proxy Generator module on MSRVTT-1k. The first two ablation experiments are conducted under ViT-B/32, and the last is under ViT-B/16.

The Effect of Losses. We evaluated the impact of adding two proposed training objectives to the baseline, as shown in Tab. 3. The baseline (Row 1) uses only \mathcal{L}_r . Adding \mathcal{L}_p (Row 2) significantly improves performance by aligning positive proxy-video pairs more closely while separating negative pairs. This highlights how \mathcal{L}_p effectively decomposes the 1-to- N relationship into N 1-to-1 relationships, reducing interference and increasing flexibility in managing relative distribution.

The best results (Row 4) are achieved by combining \mathcal{L}_p and \mathcal{L}_{pos} , confirming that these objectives together enhance the N 1-to-1 framework’s effectiveness. In contrast, adding only \mathcal{L}_{pos} (Row 3) yields minimal improvement, as it reg-

Round(s)	$\exp(\sum \theta \cdot s_i / M)$			
	R@1	R@5	R@10	MnR
1	51.9	77.5	85.3	11.5
2	52.3	77.8	85.8	11.1
3	52.5	77.8	85.4	11.1
4	52.4	77.6	85.8	11.0
5	52.3	77.4	85.6	11.1
Round(s)	$\exp(\mathbf{SW})$			
	R@1	R@5	R@10	MnR
1	52.2	77.9	85.3	11.2
2	51.9	77.5	85.3	11.2
3	52.2	77.8	85.1	11.3
4	52.0	77.6	85.2	11.0
Round(s)	w/o dash			
	R@1	R@5	R@10	MnR
1	51.1	76.5	84.6	11.3
2	51.9	76.4	84.5	10.8
3	50.3	76.3	84.5	11.8

Table 4: Evaluation of direction leader iteration on MSRVT-9k. “w/o dash” means we directly adopt direction leader \mathbf{d}_l as text proxy without specifying the distance between proxy and query.

ulates positive proxies without transitioning to the 1-to- N decomposition. Notably, in Row 3, we still use proxy-video logits to complement text-video logits during inference, assigning varying proportions by γ . However, performance peaks when $\gamma = 0$, suggesting that \mathcal{L}_{pos} works best when paired with \mathcal{L}_p .

The Effect of Dash and Iteration of Direction Leader.

To evaluate the impact of the dash mechanism, we conducted ablation experiments under three settings: $\exp(\frac{1}{M} \sum \theta \cdot s_i)$, $\exp(\mathbf{SW})$, and without dash (i.e., using the direction leader as the text proxy).

As shown in Tab. 4, generating proxies consistently outperformed the baseline, especially under $\exp(\frac{1}{M} \sum \theta \cdot s_i)$. Optimal performance was observed after three iterations with this setting, where improvements stabilized as iterations increased. However, $\exp(\mathbf{SW})$ showed no consistent pattern across iterations. This suggests that while the dash generator has fewer parameters than a single Transformer layer, its design significantly affects proxy distribution. Simpler dash generation methods, such as $\exp(\frac{1}{M} \sum \theta \cdot s_i)$, introduce fewer uncontrollable factors, improving both the controllability of direction leader generation and proxy distribution.

Since each iteration adds computational cost, we opted for two iterations for efficiency, as this configuration also achieved optimal performance under $\exp(\mathbf{SW})$ on two other datasets. To illustrate the effect of iteration count, we focus on $\exp(\frac{1}{M} \sum \theta \cdot s_i)$. A single iteration exaggerates the text query’s relative distance to videos, boosting retrieval accuracy but overfitting the text query’s semantics, reducing generalization. Conversely, five iterations overly incorporate video information, weakening the text query’s semantic guidance. Three iterations strike a balance, ensuring both

$\{\delta, \eta\}$	R@1	R@5	R@10	MnR
$\{1.5, 1\}$	54.3	79.4	87.0	8.9
$\{1, 1\}$	54.2	78.9	87.2	8.8
$\{0.5, 1\}$	54.1	79.0	86.9	9.2
$\{-1.5, -1\}$	55.1	80.2	86.9	9.2
$\{-1, -1\}$	55.2	80.4	86.8	9.3
$\{-0.5, -1\}$	54.9	79.9	86.7	8.7

Table 5: Evaluation to investigate the effect of δ and η under ViT-B/16 on MSRVT-9k. Row 1-3 imply the direction leader points at the text proxy, Row 4-6 imply the opposite.

guidance and generalization.

To assess the role of dash, we also evaluated proxies using only the direction leader across iteration schemes. As shown in Tab. 4, this setup performed well after two iterations but lacked scalability with more iterations. Solely relying on the direction leader as the proxy imposes a burden on its generation, as it cannot effectively determine the proxy’s relative distance to the query.

Finally, we adopt $\exp(\frac{1}{M} \sum \theta \cdot s_i)$ for MSRVT and $\exp(\mathbf{SW})$ for DiDeMo and ActivityNet Captions. We attribute this difference to the concise text information in MSRVT, which aligns better with the former, while the latter provides more flexibility for the diverse semantics in the other datasets.

The Effect of δ and η . We control the dominance and pointing direction between the text query and the direction leader by adjusting δ and η , respectively. While performance fluctuations under ViT-B/32 are minor (within 0.3%), both δ and η significantly influence performance under ViT-B/16. Thus, this ablation experiment focuses on ViT-B/16. Due to resource constraints, our reproduced result on MSRVT achieves R@1=53.5 (baseline: R@1=54.2) using a batch size of 56 on an A800 GPU.

As shown in Tab. 5, when $\delta > 0$ and $\eta > 0$, the R@1 metric shows limited improvement, as the direction leader primarily points towards the text query. Conversely, with $\delta < 0$ and $\eta < 0$, R@1 increases more significantly. This improvement is likely because ViT-B/16 has a larger number of patch tokens than ViT-B/32, meaning proxies under $\delta > 0$ and $\eta > 0$ are insufficiently exaggerated to enhance text-video logits. For ViT-B/32, we set $\delta = 1$ and $\eta = 1$.

Conclusion

This work introduces TV-ProxyNet, a novel framework for addressing challenges in text-video retrieval (TVR). By reformulating the task from a traditional 1-to- N matching problem into N 1-to-1 relationships, our method effectively mitigates modeling irregularities. Leveraging text proxies guided by the director and dash mechanisms, TV-ProxyNet achieves precise alignment, enhancing retrieval accuracy and robustness. Extensive evaluations on multiple benchmarks demonstrate its superiority over state-of-the-art methods, offering a robust solution for complex multimodal datasets and paving the way for improved semantic understanding and retrieval in real-world applications.

Acknowledgments

This work was supported by the NSFC NO. 62172138 and No. 62202139. This work was also partially supported by the Fundamental Research Funds for the Central Universities NO. JZ2024HGTG0310

References

- Anne Hendricks, L.; Wang, O.; Shechtman, E.; Sivic, J.; Darrell, T.; and Russell, B. 2017. Localizing moments in video with natural language. In *Proceedings of the IEEE international conference on computer vision*, 5803–5812.
- Cheng, F.; Wang, X.; Lei, J.; Crandall, D.; Bansal, M.; and Bertasius, G. 2023. Vindlu: A recipe for effective video-and-language pretraining. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10739–10750.
- Cheng, X.; Lin, H.; Wu, X.; Yang, F.; and Shen, D. 2021. Improving video-text retrieval by multi-stream corpus alignment and dual softmax loss. *arXiv preprint arXiv:2109.04290*.
- Fang, B.; Wu, W.; Liu, C.; Zhou, Y.; Song, Y.; Wang, W.; Shu, X.; Ji, X.; and Wang, J. 2023. Uatvr: Uncertainty-adaptive text-video retrieval. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 13723–13733.
- Fu, T.-J.; Li, L.; Gan, Z.; Lin, K.; Wang, W. Y.; Wang, L.; and Liu, Z. 2021. Violet: End-to-end video-language transformers with masked visual-token modeling. *arXiv preprint arXiv:2111.12681*.
- Gao, Z.; Liu, J.; Sun, W.; Chen, S.; Chang, D.; and Zhao, L. 2021. Clip2tv: Align, match and distill for video-text retrieval. *arXiv preprint arXiv:2111.05610*.
- Gorti, S. K.; Vouitsis, N.; Ma, J.; Golestan, K.; Volkovs, M.; Garg, A.; and Yu, G. 2022. X-pool: Cross-modal language-video attention for text-video retrieval. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5006–5015.
- Jin, P.; Huang, J.; Liu, F.; Wu, X.; Ge, S.; Song, G.; Clifton, D.; and Chen, J. 2022. Expectation-maximization contrastive learning for compact video-and-language representations. *Advances in neural information processing systems*, 35: 30291–30306.
- Jin, P.; Huang, J.; Xiong, P.; Tian, S.; Liu, C.; Ji, X.; Yuan, L.; and Chen, J. 2023a. Video-text as game players: Hierarchical banzhaf interaction for cross-modal representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2472–2482.
- Jin, P.; Li, H.; Cheng, Z.; Huang, J.; Wang, Z.; Yuan, L.; Liu, C.; and Chen, J. 2023b. Text-video retrieval with disentangled conceptualization and set-to-set alignment. *arXiv preprint arXiv:2305.12218*.
- Krishna, R.; Hata, K.; Ren, F.; Fei-Fei, L.; and Carlos Niebles, J. 2017. Dense-captioning events in videos. In *Proceedings of the IEEE international conference on computer vision*, 706–715.
- Li, L.; Peng, J.; Chen, H.; Gao, C.; and Yang, X. 2024. How to configure good in-context sequence for visual question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 26710–26720.
- Li, P.; Xie, C.-W.; Zhao, L.; Xie, H.; Ge, J.; Zheng, Y.; Zhao, D.; and Zhang, Y. 2023. Progressive spatio-temporal prototype matching for text-video retrieval. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4100–4110.
- Lin, C.; Wu, A.; Liang, J.; Zhang, J.; Ge, W.; Zheng, W.-S.; and Shen, C. 2022. Text-adaptive multiple visual prototype matching for video-text retrieval. *Advances in neural information processing systems*, 35: 38655–38666.
- Liu, S.; Fan, H.; Qian, S.; Chen, Y.; Ding, W.; and Wang, Z. 2021a. Hit: Hierarchical transformer with momentum contrast for video-text retrieval. In *Proceedings of the IEEE/CVF international conference on computer vision*, 11915–11925.
- Liu, Y.; Xiong, P.; Xu, L.; Cao, S.; and Jin, Q. 2022. Ts2-net: Token shift and selection transformer for text-video retrieval. In *European conference on computer vision*, 319–335. Springer.
- Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021b. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, 10012–10022.
- Luo, H.; Ji, L.; Zhong, M.; Chen, Y.; Lei, W.; Duan, N.; and Li, T. 2022. Clip4clip: An empirical study of clip for end to end video clip retrieval and captioning. *Neurocomputing*, 508: 293–304.
- Oord, A. v. d.; Li, Y.; and Vinyals, O. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.
- Wang, J.; Ge, Y.; Yan, R.; Ge, Y.; Lin, K. Q.; Tsutsui, S.; Lin, X.; Cai, G.; Wu, J.; Shan, Y.; et al. 2023. All in one: Exploring unified video-language pre-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6598–6608.
- Wang, J.; Sun, G.; Wang, P.; Liu, D.; Dianat, S.; Rabbani, M.; Rao, R.; and Tao, Z. 2024a. Text Is MASS: Modeling as Stochastic Embedding for Text-Video Retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16551–16560.
- Wang, Q.; Zhang, Y.; Zheng, Y.; Pan, P.; and Hua, X.-S. 2022a. Disentangled Representation Learning for Text-Video Retrieval. *arXiv:2203.07111*.
- Wang, Y.; Yuan, S.; Jian, X.; Pang, W.; Wang, M.; and Yu, N. 2024b. HaVTR: Improving Video-Text Retrieval Through Augmentation Using Large Foundation Models. *arXiv preprint arXiv:2404.05083*.

- Wang, Z.; Wu, Y.; Narasimhan, K.; and Russakovsky, O. 2022b. Multi-query video retrieval. In *European Conference on Computer Vision*, 233–249. Springer.
- Wu, W.; Luo, H.; Fang, B.; Wang, J.; and Ouyang, W. 2023. Cap4video: What can auxiliary captions do for text-video retrieval? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10704–10713.
- Xu, J.; Mei, T.; Yao, T.; and Rui, Y. 2016. Msr-vtt: A large video description dataset for bridging video and language. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5288–5296.
- Xue, H.; Sun, Y.; Liu, B.; Fu, J.; Song, R.; Li, H.; and Luo, J. 2023. Clip-vip: Adapting pre-trained image-text model to video-language alignment. In *The Eleventh International Conference on Learning Representations*.
- Yang, X.; Peng, Y.; Ma, H.; Xu, S.; Zhang, C.; Han, Y.; and Zhang, H. 2024a. Lever LM: Configuring In-Context Sequence to Lever Large Vision Language Models. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Yang, X.; Wu, Y.; Yang, M.; Chen, H.; and Geng, X. 2024b. Exploring diverse in-context configurations for image captioning. *Advances in Neural Information Processing Systems*, 36.
- Ye, Q.; Xu, G.; Yan, M.; Xu, H.; Qian, Q.; Zhang, J.; and Huang, F. 2023. Hitea: Hierarchical temporal-aware video-language pre-training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 15405–15416.
- Zhao, S.; Zhu, L.; Wang, X.; and Yang, Y. 2022. Centerclip: Token clustering for efficient text-video retrieval. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 970–981.