

# Ambiguity-Restrained Text-Video Representation Learning for Partially Relevant Video Retrieval

Cheol-Ho Cho, WonJun Moon, Woojin Jun, MinSeok Jung, and Jae-Pil Heo\*

Sungkyunkwan University  
{hoonchcho, wjun0830, junwoojinjin, minseokjung0328, jaepilheo}@gmail.com

## Abstract

Partially Relevant Video Retrieval (PRVR) aims to retrieve a video where a specific segment is relevant to a given text query. Typical training processes of PRVR assume a one-to-one relationship where each text query is relevant to only one video. However, we point out the inherent ambiguity between text and video content based on their conceptual scope and propose a framework that incorporates this ambiguity into the model learning process. Specifically, we propose Ambiguity-Restrained representation Learning (ARL) to address ambiguous text-video pairs. Initially, ARL detects ambiguous pairs based on two criteria: uncertainty and similarity. Uncertainty represents whether instances include commonly shared context across the dataset, while similarity indicates pair-wise semantic overlap. Then, with the detected ambiguous pairs, our ARL hierarchically learns the semantic relationship via multi-positive contrastive learning and dual triplet margin loss. Additionally, we delve into fine-grained relationships within the video instances. Unlike typical training at the text-video level, where pairwise information is provided, we address the inherent ambiguity within frames of the same untrimmed video, which often contains multiple contexts. This allows us to further enhance learning at the text-frame level. Lastly, we propose cross-model ambiguity detection to mitigate the error propagation that occurs when a single model is employed to detect ambiguous pairs for its training. With all components combined, our proposed method demonstrates its effectiveness in PRVR.

## Introduction

With societal advancements, the use of video media for information dissemination has become widespread. Consequently, the field of Text-to-Video Retrieval (T2VR), which allows users to find desired videos using text queries, has also been spotlighted. However, existing T2VR approaches often assume that videos contain only the relevant parts to the text query. This assumption does not align with real-world scenarios, where videos can vary in length and context. To address this issue, partially relevant video retrieval (PRVR) (Dong et al. 2022) is proposed to deal with untrimmed videos where only specific video segments correspond to the text query.

\*Corresponding author

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

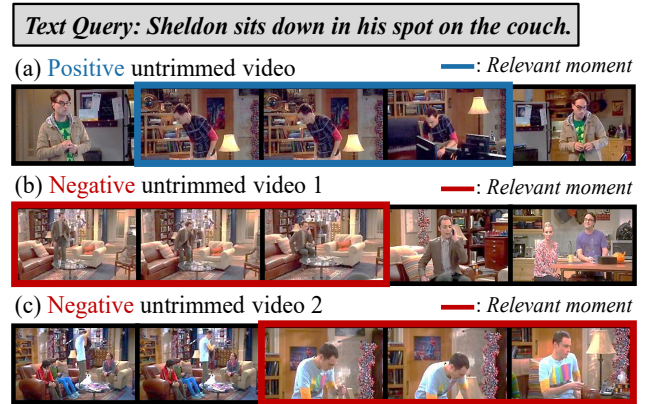


Figure 1: The illustration of ambiguous relationships between text-video pair. For the text query, ‘Sheldon sits down in his spot on the couch.’, all three untrimmed videos contain relevant scenes. However, only video (a) is learned as positive, while videos (b) and (c) are typically treated as negative in prior techniques. Such ambiguous relationships between text and video are more likely to occur with untrimmed videos that often include diverse contexts.

MS-SL (Dong et al. 2022) initially modeled multi-scale video features to prepare for various contexts that a text query might encompass. GMMFormer (Wang et al. 2024), on the other hand, critiqued the exhaustive multi-scale clip modeling of MS-SL for its inefficiency and proposed to use Gaussian attention to encode only the local contexts. Although these methods have made significant progress in PRVR, the obscurity in labeling text-video pairs is a yet-to-be-explored problem.

Typically, text and video instances are labeled pair-wise due to the high costs of exploring relationships between all text and video instances. This has led previous works to learn only the paired instance as positive relations, treating all others as negative, even when similar video-text pairs are present. However, we argue that pair-wisely labeled retrieval datasets often introduce ambiguity between text-video instances. For example, as illustrated in Fig. 1, while the text query is clearly partially relevant to the paired video at the top, it is also relevant to other video instances in the dataset.

In this regard, we propose Ambiguity-Restrained representation Learning (ARL), a framework that detects instances in ambiguous relations with the model’s online (per-epoch) knowledge. These ambiguities are taken into account in the objectives to reduce fallible supervision which can arise when all unpaired pairs are treated as having a negative relationship. To determine ambiguous relationships, we use uncertainty and similarity measures, as shown in Fig.2. Simply put, we define the text-video pairs in ambiguous relations if they exhibit high uncertainty and high pair-wise similarity, indicating that the pair cannot be simply defined as negative since they possess commonly shared semantics throughout the dataset and their resemblance to each other. Our proposed ARL includes these identified ambiguous relationships in training. Particularly, we grant the model the flexibility in treating instances with ambiguous relationships by relaxing the constraints while positive and negative relationships are learned in conventional ways. Beyond the text-video relationships, we further explore text-frame relationships with ambiguity-aware objectives since untrimmed videos often encompass multiple contexts. Lastly, we employ cross-model ambiguity detection to mitigate the risk of error propagation in detecting ambiguous relationships that might occur when a model predicts and uses its prediction for training.

Our main contributions are: (1) We propose ambiguity-restrained representation learning (ARL), the first approach to address label ambiguity in PRVR. By modeling the relationships between instances, ARL mitigates the impact of learning uncertain relationships. (2) We extend ARL to the text-frame level to account for multiple contexts in untrimmed videos, enhancing the effective utilization of all frames during learning. (3) We introduce cross-model ambiguity detection to avoid learning erroneously detected ambiguity repeatedly. (4) We achieve state-of-the-art performance on two datasets, *i.e.*, TVR and ActivityNet.

## Related Works

### Text-to-Video Retrieval

Text-to-video retrieval (T2VR) aims to realize the metadata-free searching paradigm by aligning video contexts and text descriptions. To benefit from the vision-language aligned model, the CLIP model pretrained on large-scale text-image pairs is popularly employed as the initial learning point. Subsequently, tackling the mismatch between the amount of information in videos and texts, several works have focused on designing base units for feature matching (Gorti et al. 2022; Lin et al. 2022), *e.g.*, frame-word (Wang et al. 2022) and frame-sentence (Jin et al. 2023; Wu et al. 2023).

Uncertainty in the scope of contexts has been discussed in the literature of T2VR. Particularly, there have been approaches (Fang et al. 2023; Li et al. 2024; Song and Soleymani 2019) to tackle the uncertainty. PVSE (Song and Soleymani 2019) and UATVR (Fang et al. 2023) extracted multi-faceted representations of text-video pairs. PAU (Li et al. 2024) addressed the aleatoric uncertainty inherent in text-video data by ensuring the consistency between different uncertainty measures. Our study shares a common high-

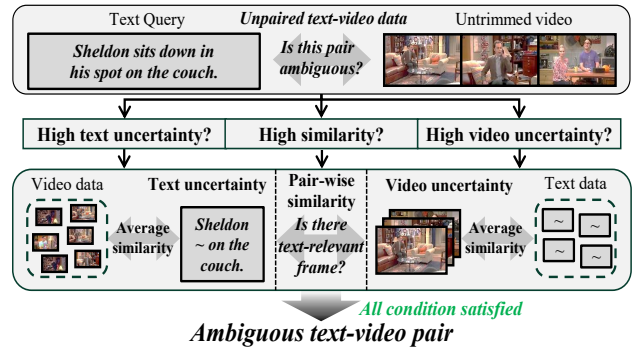


Figure 2: Illustration for ambiguous pair detection. To identify ambiguous text-video pairs, we use two key metrics: uncertainty and similarity. Uncertainty for each text and video instance is calculated by measuring the average similarity across different modalities. This reflects the degree of contextual overlap within the dataset. Similarity is determined by the maximum similarity between the text and frames from the untrimmed video. We define a pair as having an ambiguous relationship when both uncertainty and similarity are high, exhibiting the commonly shared context in the dataset and between the pair.

level concept of tackling uncertainty in text-video data. Yet, our work differs in that the key focus is to explore ambiguous relationships among all text-video pairs.

### Partially Relevant Video Retrieval

Beyond the scenario of T2VR, PRVR further targets the fine-grained capability of search engines to retrieve videos even when only the partial context corresponds to the given text queries. To address PRVR, typical approaches are to clip the videos into multiple segments. MS-SL (Dong et al. 2022) exhaustively constructed clips of varying lengths and performed similarity matching to the text query. GMM-Former (Wang et al. 2024) implemented the feature of locality in forming clip representations by applying normally distributed weights in the attention layers.

On the other hand, our work has a key focus on detecting the ambiguous relationships between text-video instances caused by one-to-one labeling in video retrieval datasets.

### Noisy Label Learning

Learning with noisy labels has been spotlighted due to its practical applications (Han et al. 2018; Li, Socher, and Hoi 2020; Azadi et al. 2015; Wang et al. 2019). Uncertainty estimation and the co-training frameworks are the other popular streams. Uncertainty is often utilized to detect noisy labels (Neverova, Novotny, and Vedaldi 2019; Ju et al. 2022; Northcutt, Jiang, and Chuang 2021; Zheng and Yang 2021) where co-training frameworks are shown to be effective in refining the noisy labels (Han et al. 2018; Wei et al. 2020; Tan et al. 2021; Li, Socher, and Hoi 2020). This issue has also been addressed recently in the context of video-text learning (Lin et al. 2024). To combat possible ambiguity

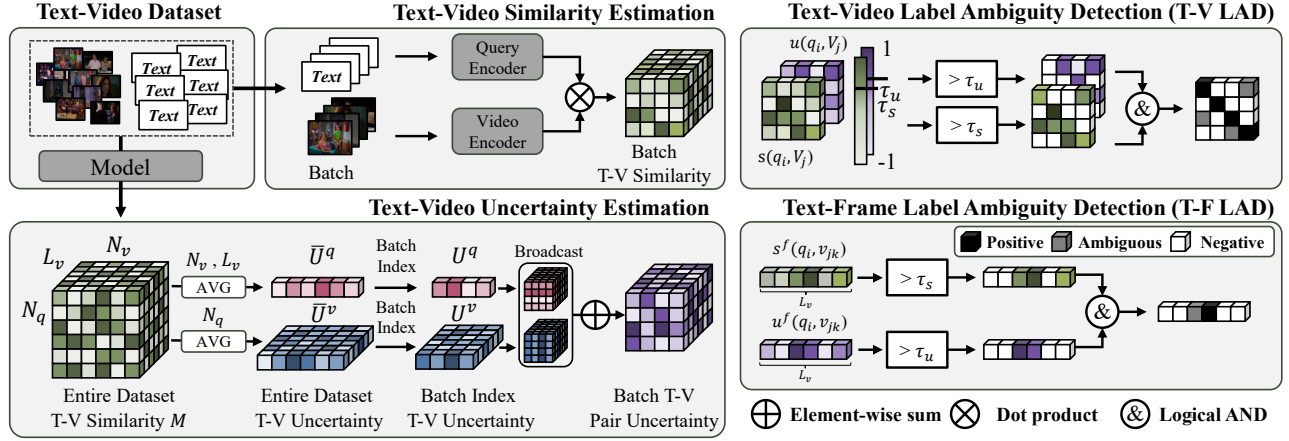


Figure 3: Overview of ARL. (*Left*) Given the text-video train set, we initially calculate the text-video similarity to compute the uncertainty over the dataset level. Batch-wisely indexed uncertainty is used along the batch-wise similarity between texts and videos to explore ambiguous text-video relationships in the mini-batch. (*Right*) Two levels of Label Ambiguity Detection (LAD) modules detect ambiguous relationships. Text-Video LAD employs  $s(q_i, V_j)$  and  $u(q_i, V_j)$ , the similarity and uncertainty maps between the text queries and videos in each mini-batch. Text-Frame LAD utilizes  $s^f(q_i, v_{jk})$  and  $u^f(q_i, v_{jk})$ , the similarity and uncertainty maps between each query and video frames for each text-video pair.

in pair-wisely labeled relationships, our study employs the concept of uncertainty and co-training.

## Method

### Overview

The overview of Ambiguity-Restrained representation Learning (ARL) is shown in Fig. 3. To illustrate, we compute the similarity between all texts and videos in the training set to define text-video uncertainty in each epoch. Subsequently, the batch-wisely indexed uncertainty is processed to the Label Ambiguity Detection (LAD) module along with the similarity between the texts and videos in the mini-batch. LAD identifies ambiguous relationships between text and video modalities at two levels: the text-video and text-frame. Lastly, we perform cross-model ambiguity detection with the dual branch structure which is commonly employed in the literature of PRVR. Employing the same structure and inputs for each branch for cross-model ambiguity detection, cross-model ambiguity detection allows each model to learn with the detected ambiguous sets from another model. Note that the cross-model ambiguity detection is not depicted in Fig.3. For the rest of the paper, we explain the training scenario with an assumption that all text queries and video instances are composed of the same number of elements, i.e.,  $L_q$  words and  $L_v$  frames, for better clarity.

**Text Query Representation.** Given  $i$ -th text query out of  $N_q$  queries in the training dataset, we use the pretrained text encoder to extract features for each word. Subsequently, we embed the word features into a low-dimensional space using a fully connected (FC) layer followed by ReLU activation. After that, we incorporate positional encoding into these features and employ a transformer layer to obtain  $d$ -dimensional word feature vectors  $Q_i \in \mathbb{R}^{L_q \times d}$ . Finally, we

apply an attention pooling module to the word feature vectors to obtain the query text embeddings  $q_i \in \mathbb{R}^d$ .

**Video Representation.** Given  $j$ -th untrimmed video out of  $N_v$  train set videos, we extract frame features  $V_j' \in \mathbb{R}^{L_v \times d_v}$  using a pre-trained 2D or 3D CNN. Symmetric to the text branch, video features are also passed through a fully connected layer with ReLU activation to reduce the dimensions. After that, we incorporate positional encoding  $P$  into the extracted features, before forwarding the features through a transformer layer to obtain  $V_j \in \mathbb{R}^{L_v \times d}$ :

$$V_j = [v_{j1}, v_{j2}, \dots, v_{jL_v}] = \text{Transformer}(\text{FC}(V_j') + P), \quad (1)$$

where  $v_{jk}$  refers to the  $k$ -th frame features in the  $j$ -th video.

**Similarity Measure.** Given text and video representations, the frame-wise similarity score  $s^f$  between a text query features  $q_i$  and video frame features  $v_{jk}$  is derived as:

$$s^f(q_i, v_{jk}) = \cos(q_i, v_{jk}), \quad (2)$$

where  $\cos(\cdot, \cdot)$  is the cosine similarity between text query and video frames. Then, the maximum similarity value between a text query and video frames is used for the retrieval score since only partial video frames are relevant to the text. The retrieval (similarity) score is obtained as follows:

$$s(q_i, V_j) = s^f(q_i, v_{j\hat{k}}), \quad (3)$$

where  $v_{j\hat{k}} = \arg\max_{v_{jk}} \cos(q_i, v_{jk})$  and  $\hat{k}$  denote the video frame and its index with the maximum similarity, respectively.

### Ambiguity-Restrained Representation Learning

**Definition of Ambiguity.** Typical datasets for Partially Relevant Video Retrieval (PRVR) consist of thousands of

matched text-video pairs. It is a popular practice to treat paired text-video data as positive pairs, while all unpaired data are regarded as negative pairs. However, we question whether it is correct to treat all unpaired data as negative.

We address this issue by defining an ambiguous relationship for unpaired text-video pairs that are difficult to classify as simply negative. Specifically, we use uncertainty and similarity measures to identify these ambiguous relationships and grant the model flexibility in treating these instances. First, the uncertainty indicates whether each instance contains commonly shared semantics. It is calculated as the average similarity across the dataset, and thus, the uncertainty is measured high if numerous instances possess high similarity to a specific instance. Second, similarity refers to the degree of pair-wise text-video resemblance, defined as the maximum similarity between a single text query and video frame representations. Thus, the pairs with high similarity indicate that a video instance includes a frame that shares similar attributes with the given query. Although both metrics originate from the same similar operation, we note that these two metrics serve as different criteria with different objectives. We also note that our study in Fig. 4 further validates the claim that these two metrics exhibit different distributions.

**Warmup Phase.** Our method utilizes online knowledge (per epoch) of the model to identify the ambiguous relation between text and video. Therefore, we need to warm up the model for a few epochs to initially train the model to learn general text-video relationships. For the warmup phase, we adopt the commonly used triplet ranking loss (Dong et al. 2023; Faghri et al. 2018) and infoNCE loss (Ma et al. 2022; Zhang et al. 2021) in retrieval tasks.

**Uncertainty Estimation.** Uncertainty for each text/video instance is measured by considering similarities across the entire dataset before training progress. Specifically, we calculate the feature similarity map  $M \in \mathbb{R}^{N_q \times N_v \times L_v}$  between all text queries and video frames in the training dataset using the online model. Here,  $M_{xyz}$  denotes the similarity between the  $x$ -th text query and the  $z$ -th frame of the  $y$ -th video in the entire dataset. With the similarity map  $M$ , we define the dataset-wise uncertainty of each text query  $\bar{U}^q \in \mathbb{R}^{N_q}$  as the average similarity between the text and all video frames in the dataset and the uncertainty of video frame  $\bar{U}^v \in \mathbb{R}^{N_v \times L_v}$  as the average similarity to all text queries:

$$\bar{U}_x^q = \frac{1}{N_v L_v} \sum_{y=1}^{N_v} \sum_{z=1}^{L_v} M_{xyz}, \quad \bar{U}_{yz}^v = \frac{1}{N_q} \sum_{x=1}^{N_q} M_{xyz}, \quad (4)$$

where  $\bar{U}_x^q$  is uncertainty value of  $x$ -th text query and  $\bar{U}_{yz}^v$  is uncertainty value of  $z$ -th frame of  $y$ -th video instance. Texts with higher average similarity to video frames and frames with higher average similarity to queries both exhibit greater uncertainty.

Note that an instance with higher uncertainty implies that its context is likely to be commonly shared with other instances. These similarity maps and uncertainties are updated with Eq. 4 every epoch.

For better clarity for indexing uncertainties  $U^q$  and  $U^v$  within each mini-batch, we define the batch-level subsets of uncertainties as  $U^q$  and  $U^v$ . Note that  $U^q$  and  $U^v$  are subsets of  $\bar{U}^q$  and  $\bar{U}^v$ , which include the uncertainties of all queries and videos in the mini-batch.

Below, we define the uncertainty value between the text query and each video frame as  $u^f$  and the uncertainty between the query and the whole video as  $u$ :

$$u(q_i, V_j) = \frac{1}{2}(U_i^q + U_{j\hat{k}}^v); u^f(q_i, v_{jk}) = \frac{1}{2}(U_i^q + U_{jk}^v), \quad (5)$$

where  $\hat{k}$  indicates the frame index with the highest similarity to  $i$ -th text query, as mentioned in Eq. 3.

**Text-Video Label Ambiguity Detection.** To discover pairs in ambiguous relation per text / video instance, we utilize computed uncertainty  $u$  and similarity score  $s$ . For  $i$ -th text query  $q_i$ , a set of ambiguous video pairs are gathered as:

$$\mathcal{A}_i^q = \{V_a \mid s(q_i, V_a) > \tau_s \text{ and } u(q_i, V_a) > \tau_u\}, \quad (6)$$

where  $\tau_s$  and  $\tau_u$  are thresholding hyperparameters.

On the other hand, an ambiguous query set for a video  $V_j$  is defined as follows:

$$\mathcal{A}_j^v = \{q_a \mid s(q_a, V_j) > \tau_s \text{ and } u(q_a, V_j) > \tau_u\}. \quad (7)$$

**Ambiguous Aware Representation Learning.** For training, we utilize margin triplet ranking loss and contrastive learning (Chen et al. 2020), following previous works (Dong et al. 2022; Wang et al. 2024). Below, we enumerate the modifications in objectives to implement with ambiguous relationships. In the case of contrastive learning, we add modifications to supervised contrastive learning (Khosla et al. 2020) to equip with the multi-positive contrastive objective:

$$\mathcal{L}_{ij}^{lv} = -\log \left( \frac{e^{s(q_i, V_j)} + \sum_{V_a \in \mathcal{A}_i^q} e^{s(q_i, V_a)}}{e^{s(q_i, V_j)} + \sum_{V \in \mathcal{A}_i^q \vee \mathcal{N}_i^q} e^{s(q_i, V)}} \right) \quad (8)$$

$$\mathcal{L}_{ij}^{vt} = -\log \left( \frac{e^{s(q_i, V_j)} + \sum_{q_a \in \mathcal{A}_j^v} e^{s(q_a, V_j)}}{e^{s(q_i, V_j)} + \sum_{q \in \mathcal{A}_j^v \vee \mathcal{N}_j^v} e^{s(q, V_j)}} \right) \quad (9)$$

$$\mathcal{L}^{\text{ncc}} = \frac{1}{n} \sum_{(q_i, V_j) \in \mathcal{B}} \mathcal{L}_{ij}^{lv} + \mathcal{L}_{ij}^{vt}, \quad (10)$$

where  $\mathcal{B}$  denotes a mini-batch, and  $(q_i, V_j)$  denotes a positive pairs within this batch.  $\mathcal{N}_j^v$  and  $\mathcal{N}_i^q$  are sets of negative samples for each video and query that  $\mathcal{N}_j^v$  contains samples that are neither in positive nor ambiguous relations to  $j$ -th video. In short, our multi-positive contrastive objective enables more flexible model learning by accommodating ambiguous relationships. While instances within ambiguous sets in the numerator are not trained as negatives, not all instances are necessarily trained to have positive relations with the anchor. This is because of the possibility that maximizing a single similarity value can still facilitate loss convergence.

On the other hand, we organize dual triplets for margin triplet ranking loss; one with ambiguous sets for  $\mathcal{L}_a^{\text{trip}}$  and the other with negative pairs for  $\mathcal{L}_n^{\text{trip}}$  as follows:

$$\mathcal{L}_a^{\text{trip}} = \frac{1}{n} \sum_{(q_i, V_j) \in \mathcal{B}} \{ \max(0, m_a + s(q_a, V_j) - s(q_i, V_j)) + \max(0, m_a + s(q_i, V_a) - s(q_i, V_j)) \} \quad (11)$$

$$\mathcal{L}_n^{\text{trip}} = \frac{1}{n} \sum_{(q_i, V_j) \in \mathcal{B}} \{ \max(0, m + s(q_n, V_j) - s(q_i, V_j)) + \max(0, m + s(q_i, V_n) - s(q_i, V_j)) \}, \quad (12)$$

where  $m_a$  and  $m$  denote the respective margins, and  $q_a \in \mathcal{A}_j^v$  and  $V_a \in \mathcal{A}_i^q$  represent the ambiguous samples for each video and query, respectively.  $q_n$  and  $V_n$  are the negative samples for each query and video. We set the margin for  $\mathcal{L}_a^{\text{trip}}$  smaller, ensuring that ambiguous instances remain more similar to the anchor than contextually irrelevant instances ( $m_a < m$ ). By alleviating distance constraints on ambiguous instances, we allow ambiguous sets with potential positive relationships to the anchor to be excluded from negative training. The margin  $m_a$  is used to maintain the hierarchy, ensuring that paired instances are trained as positives in Eq. 8- 9. In conclusion, the ambiguity-restrained objective for text-video pairs is formulated as follows:  $\mathcal{L}^{\text{video}} = \lambda_{\text{ncc}} \mathcal{L}^{\text{ncc}} + \mathcal{L}_a^{\text{trip}} + \mathcal{L}_n^{\text{trip}}$ , where  $\lambda_{\text{ncc}}$  is a hyperparameter to balance the losses.

**Text-Frame Label Ambiguity Detection Within Untrimmed Video.** For untrimmed videos, diverse contexts may be present within a single instance. Yet, exploring the relationships with the same video is a yet-to-be-explored issue. Therefore, we delve deep into the relation between the text query and frame-wise representations. Symmetrical to the process of discovering the ambiguity relations between text and video instances in the mini-batch, we apply Eq. 6 and Eq. 7 between the query feature  $q$  and video frame feature  $v$ . Note that the similarity and uncertainty measures, *i.e.*,  $s$  and  $u$ , are also substituted with frame-wise similarity  $s^f$  (Eq. 2) and uncertainty  $u^f$  (Eq. 5), respectively. Consequently, the same objectives for learning text-video relationships in Eq. 8- 12 are applied to learn text-frame relationships for each video. The objective for text-frame pairs  $\mathcal{L}^{\text{frame}}$  is also the same with  $\mathcal{L}^{\text{video}}$ . We note that  $\lambda_{\text{ncc}}$  is shared.

### Cross-Model Ambiguity Detection

Ambiguous pairs can be detected using the model’s own predictions, similar to self-training (Balcan, Beygelzimer, and Langford 2006; Freund, Schapire, and Abe 1999). However, we point out the vulnerability of error propagation of ambiguity detection when the model relies on its knowledge, progressively reinforcing its initial imperfect predictions.

To address this challenge, we utilize two identical encoders that reciprocally transfer one’s detected ambiguous sets to the other to mitigate the impact of noisy labels (Han et al. 2018). Given two models denoted as  $\theta$  and

$\Phi$ , each model computes ambiguous text-video pairs following Eq. 4-Eq. 7 with its online knowledge and provides it as the training guidance to the other model.

Finally, the retrieval score for the prediction is yielded as the average of Eq. 3 for each model:

$$s(q_i, V_j) = \frac{1}{2} (s_\theta(q_i, V_j) + s_\Phi(q_i, V_j)), \quad (13)$$

where  $s_\theta$  and  $s_\Phi$  denote the retrieval scores from model  $\theta$  and  $\Phi$ , respectively.

## Experiments

### Datasets and Metrics

We evaluate our method on two large-scale video datasets, *i.e.* TVR (Lei et al. 2020) and ActivityNet Captions (Krishna et al. 2017). We adopt the data split provided by previous works (Zhang et al. 2020, 2021). Following (Dong et al. 2022), we use rank-based recall as our evaluation metric, *i.e.*, R@K (K=1, 5, 10, 100), where R@K denotes the fraction of queries that successfully retrieve the desired items within the top K of the ranking list. Moreover, we report the Sum of all Recalls (SumR) for comprehensive comparisons.

### Implementation Details

We use ResNet (He et al. 2016) and I3D (Carreira and Zisserman 2017) for TVR and only the I3D for the ActivityNet-Captions to extract visual features. For text query representation, we use the RoBERTa (Liu et al. 2019) feature for both datasets. Also, to demonstrate the effectiveness of our method with the large model, we conducted experiments using CLIP-L/14 (Radford et al. 2021). While the typical parameters are set the same as in (Wang et al. 2024), the thresholds  $\tau_s$  and  $\tau_u$  were defined at each epoch using the distribution values of similarity and uncertainty from the training dataset. Particularly,  $\tau_s$  is set to the mean value of the similarity distribution of the positive pairs, and  $\tau_u$  is set to the value corresponding to the mean of uncertainty distribution of train dataset. More details are provided in the appendix.

### Comparison With the State-of-the-Arts

**Retrieval Performance.** We report the results comparing our approach with state-of-the-art methods for video corpus moment retrieval, and PRVR. Note that all performances are yielded without using the moment supervision. In Tab. 1 and Tab. 2, we show the results on the TVR and ActivityNet Captions datasets. As observed, our proposed method outperforms previous works in all recall metrics, achieving a margin of 9.3%, 2.3% in SumR compared to GMMFormer when using ResNet, I3D, and Roberta. Moreover, we highlight the consistent gaps between ours and DL-DKD (Dong et al. 2023) while DL-DKD employed an additional CLIP-B/32 model for knowledge distillation (marked with †). We owe this superior performance to the study of ambiguity in one-to-one relationship learning with untrimmed videos for PRVR. This trend was similarly observed with CLIP-L/14, showing improvements of 9.1% and 3.3% compared to MS-SL and GMMFormer respectively.

Model	TVR					ActivityNet Captions				
	R@1	R@5	R@10	R@100	SumR	R@1	R@5	R@10	R@100	SumR
<i>VCMR methods without moment localization</i>										
XML (Lei et al. 2020)	10.0	26.5	37.3	81.3	155.1	5.3	19.4	30.6	73.1	128.4
ReLoCLNet (Zhang et al. 2021)	10.7	28.1	38.1	80.3	157.1	5.7	18.9	30.0	72.0	126.6
CONQUER (Hou, Ngo, and Chan 2021)	11.0	28.9	39.6	81.3	160.8	6.5	20.4	31.8	74.3	133.1
<i>PRVR models</i>										
MS-SL (Dong et al. 2022)	13.5	32.1	43.4	83.4	172.4	7.1	22.5	34.7	75.8	140.1
DL-DKD (Dong et al. 2023)†	14.4	34.9	45.8	84.9	179.9	8.0	25.0	37.4	77.1	147.6
GMMFormer (Wang et al. 2024)	13.9	33.3	44.5	84.9	176.6	<b>8.3</b>	24.9	36.7	76.1	146.0
<b>Ours</b>	<b>15.6</b>	<b>36.3</b>	<b>47.7</b>	<b>86.3</b>	<b>185.9</b>	<b>8.3</b>	24.6	<b>37.4</b>	<b>78.0</b>	<b>148.3</b>

Table 1: Performance comparison for Resnet, I3D and Roberta features. † indicates the usage of additional CLIP-B/32 model.

Model	TVR					ActivityNet Captions				
	R@1	R@5	R@10	R@100	SumR	R@1	R@5	R@10	R@100	SumR
MS-SL (Dong et al. 2022)	31.9	57.6	67.7	93.8	251.0	14.7	37.1	50.4	84.6	186.7
GMMFormer (Wang et al. 2024)	29.8	54.2	64.6	92.5	241.1	15.2	37.7	50.5	83.7	187.1
<b>Ours</b>	<b>34.6</b>	<b>60.4</b>	<b>70.7</b>	<b>94.4</b>	<b>260.1</b>	<b>15.3</b>	<b>38.4</b>	<b>51.5</b>	<b>85.2</b>	<b>190.4</b>

Table 2: Performance comparison on TVR and ActivityNet Captions for CLIP-L/14 features.

	MS-SL	GMMFormer	Ours
FLOPs (G)	1.29	1.95	1.23
Params (M)	4.85	12.85	5.34

Table 3: Model complexity comparison.

<i>inference runtime (ms)</i>				
Video size	1000	1500	2000	2500
MS-SL	0.366	0.606	0.759	0.893
GMMFormer	0.264	0.267	0.270	0.293
Ours	0.294	0.391	0.427	0.612

Table 4: Inference time comparison.

**Complexity Analysis.** In this section, we present the model complexity analysis, as shown in Tab 3 and 4. We analyzed three aspects: FLOPs, the number of parameters, and runtime required to process a single text query on Nvidia RTX 3090 GPU. While our method demonstrates high efficiency in terms of FLOPs and parameters among PRVR methods, its runtime is relatively slower than that of GMMFormer. The efficiency in FLOPs and the number of parameters is attributed to our streamlined transformer architecture, whereas GMMFormer employs multiple attention blocks in parallel with different Gaussian kernels. The relatively slower runtime of our method is due to the absence of aggregated frame-level feature used in GMMFormer. However, our method still achieves a real-time runtime of under 1 ms, which represents a well-balanced tradeoff between performance and runtime speed.

## Ablation and Further Studies

**Component Analysis.** To understand the effectiveness of each component, we conduct component ablation in Tab. 5. Comparing rows (a) and (b), our base model and the model trained with text-video ambiguity in representation learning, we observe that it is effective in that it results in a 3.6-

	T-V	T-F	C.L	R@1	R@5	R@10	R@100	SumR
(a)	-	-	-	32.8	58.1	68.2	93.7	252.8
(b)	✓	-	-	33.6	58.9	69.4	94.5	256.4
(c)	✓	✓	-	34.3	59.9	70.1	94.4	258.7
(d)	✓	-	✓	34.3	59.9	69.9	94.3	258.4
(e)	✓	✓	✓	34.6	60.4	70.7	94.4	260.1

Table 5: Ablation study to investigate the effectiveness of different components on TVR. T-V Ambiguity and T-F Ambiguity indicate the use of ambiguity-aware representation learning within text-video representation learning and the use of text-frame representation learning.

$\mathcal{A}$	R@1	R@5	R@10	R@100	SumR
Positive	34.0	59.7	70.1	94.5	258.3
Ignore	34.5	60.1	70.1	94.5	259.2
Ours	34.6	60.4	70.7	94.4	260.1

Table 6: Performances comparison between different usages of ambiguous sets.

point increase in SumR. It demonstrates that the flexibility in learning text-video relationships is important for PRVR because text queries and videos often exhibit ambiguity in contextual similarity. Moreover, the increase in rows (c) and (e) spotlight the potential of exploring the text-frame relationships within each video. Subsequently, cross-model learning is verified to be beneficial for PRVR as shown in rows (d) and (e). To conclude, the benefits of our components imply that considering only the paired text-video instances as contextually similar to each other in pair-wisely labeled datasets is vulnerable to forming ambiguous relationships between text-video instances (especially for PRVR).

**Ambiguous Set Learning Strategy.** In our work, we granted flexibility to the model in learning ambiguous relationships. Yet, other options exist to treat all text-video instances in ambiguous sets as positives for a given anchor




Figure 4: Similarity and uncertainty distributions for positively-paired and unpaired text-video pairs on TVR. Similarity distributions are illustrated in (a). Distributions of positive sets are generally higher than that of negative sets. In (b), uncertainty distributions are shown. As uncertainty value is not much affected by the similarity of single pair (for positive pairs), it is shown that distributions of positive and unpaired are formed similarly.

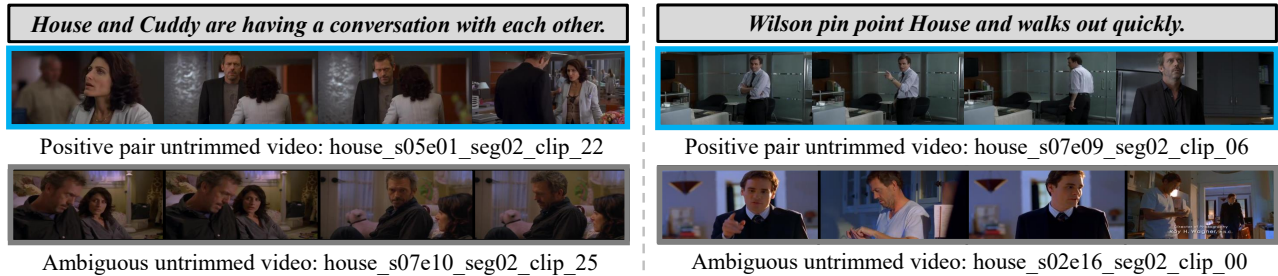


Figure 5: Ambiguity detection results on TVR dataset. For given queries, we visualize untrimmed videos that are detected to be in ambiguous relations after the training process (gray boxes). The videos are shown to be highly relevant to the query although they are not paired as positive instances of the query.

or to exclude them from training. In Tab. 6, we report the performances of other options. Specifically, we employed a supervised contrastive objective (Khosla et al. 2020) to maximize the similarity between every ambiguous pair for positive or utilized masking operation to ignore. As the performances decrease in both scenarios, we believe that ambiguous sets  $\mathcal{A}$  are the mixture of instances in positive and negative relations with an anchor. Consequently, these results indicate that simply defining an ambiguous set as positive or negative could lead to suboptimal results.

**Uncertainty & Similarity Distributions.** we argued that the two metrics used to detect ambiguity serve different roles and exhibit distinct distributions. In Fig. 4, we present the distributions of each metric on the TVR. (ActivityNet Captions is similar to the TVR.) As illustrated, the distributions of similarity (left) and uncertainty (right) display different shapes. Specifically, the positive and negative distributions are distinguishable in terms of similarity, indicating that similarity between positively paired pairs is generally higher than that of unpaired text-video pairs. Conversely, we observe that uncertainty values for positive and unpaired sets are similarly distributed since the degree of semantic overlap throughout the dataset does not depend on a single pair-wise similarity. Note that the text-video averaged uncertainty values are used for plots as pairwise uncertainty is utilized to detect ambiguous pairs (Eq. 5). These results indicate that similarity and uncertainty hold independent meanings that both are preferred to be considered simultaneously.

## Qualitative Results

We analyzed to verify whether detected text-video pairs actually include ambiguous relationships to an anchor. In Fig. 5, we plot two text queries with the paired videos and the videos in ambiguous relationships (with gray boxes). To illustrate, we observe that the characters in the scenes are identical while the overall context is also very similar. Especially, the video on the left example includes the moment that can be expressed with the given query. This confirms that our ambiguity detection effectively captures ambiguous relationships and reduces the impact of erroneous supervision that occurs when treating all unpaired text-video pairs as negative sets.

## Conclusion

**Conclusion.** In this paper, we addressed the issue of ambiguous relationships between the text-video pairs in pairwise labeled text-video data. To tackle this challenge, we proposed Ambiguity-Restrained representation Learning (ARL), designed to mitigate the impact of learning from ambiguous relationships between text and video. ARL utilizes uncertainty to first assess whether each text or video likely includes common contexts within the dataset. Subsequently, the similarity within each mini-batch is computed to identify ambiguous relationships. These relationships are then incorporated into the Ambiguity-Aware Representation Learning framework, allowing the model flexibility in learning these relationships. Our results demonstrate that one-to-one relationship learning is vulnerable to ambiguous relationships between texts and videos.

## Acknowledgments

This work was supported in part by MSIT/IITP (No. 2022-0-00680, 2020-0-01821, 2019-0-00421, RS-2024-00459618, RS-2024-00360227, RS-2024-00437102, RS-2024-00437633), and MSIT/NRF (No. RS-2024-00357729).

## References

- Azadi, S.; Feng, J.; Jegelka, S.; and Darrell, T. 2015. Auxiliary image regularization for deep cnns with noisy labels. *arXiv preprint arXiv:1511.07069*.
- Balcan, M.-F.; Beygelzimer, A.; and Langford, J. 2006. Agnostic active learning. In *Proceedings of the 23rd international conference on Machine learning*, 65–72.
- Carreira, J.; and Zisserman, A. 2017. Quo vadis, action recognition? a new model and the kinetics dataset. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 6299–6308.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, 1597–1607. PMLR.
- Dong, J.; Chen, X.; Zhang, M.; Yang, X.; Chen, S.; Li, X.; and Wang, X. 2022. Partially Relevant Video Retrieval. In *Proceedings of the 30th ACM International Conference on Multimedia*, 246–257.
- Dong, J.; Zhang, M.; Zhang, Z.; Chen, X.; Liu, D.; Qu, X.; Wang, X.; and Liu, B. 2023. Dual Learning with Dynamic Knowledge Distillation for Partially Relevant Video Retrieval. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 11302–11312.
- Faghri, F.; Fleet, D. J.; Kiros, J. R.; and Fidler, S. 2018. VSE++: Improving Visual-Semantic Embeddings with Hard Negatives.
- Fang, B.; Liu, C.; Zhou, Y.; Yang, M.; Song, Y.; Li, F.; Wang, W.; Ji, X.; Ouyang, W.; et al. 2023. Uatvr: Uncertainty-adaptive text-video retrieval. *arXiv preprint arXiv:2301.06309*.
- Freund, Y.; Schapire, R.; and Abe, N. 1999. A short introduction to boosting. *Journal-Japanese Society For Artificial Intelligence*, 14(771-780): 1612.
- 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.
- Han, B.; Yao, Q.; Yu, X.; Niu, G.; Xu, M.; Hu, W.; Tsang, I.; and Sugiyama, M. 2018. Co-teaching: Robust training of deep neural networks with extremely noisy labels. *Advances in neural information processing systems*, 31.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Hou, Z.; Ngo, C.-W.; and Chan, W. K. 2021. CONQUER: Contextual query-aware ranking for video corpus moment retrieval. In *Proceedings of the 29th ACM International Conference on Multimedia*, 3900–3908.
- Jin, P.; Li, H.; Cheng, Z.; Huang, J.; Wang, Z.; Yuan, L.; Liu, C.; and Chen, J. 2023. Text-video retrieval with disentangled conceptualization and set-to-set alignment. *arXiv preprint arXiv:2305.12218*.
- Ju, L.; Wang, X.; Wang, L.; Mahapatra, D.; Zhao, X.; Zhou, Q.; Liu, T.; and Ge, Z. 2022. Improving medical images classification with label noise using dual-uncertainty estimation. *IEEE transactions on medical imaging*, 41(6): 1533–1546.
- Khosla, P.; Teterwak, P.; Wang, C.; Sarna, A.; Tian, Y.; Isola, P.; Maschinot, A.; Liu, C.; and Krishnan, D. 2020. Supervised contrastive learning. *Advances in neural information processing systems*, 33: 18661–18673.
- 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.
- Lei, J.; Yu, L.; Berg, T. L.; and Bansal, M. 2020. Tvr: A large-scale dataset for video-subtitle moment retrieval. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXI 16*, 447–463. Springer.
- Li, H.; Song, J.; Gao, L.; Zhu, X.; and Shen, H. 2024. Prototype-based Aleatoric Uncertainty Quantification for Cross-modal Retrieval. *Advances in Neural Information Processing Systems*, 36.
- Li, J.; Socher, R.; and Hoi, S. C. 2020. Dividemix: Learning with noisy labels as semi-supervised learning. *arXiv preprint arXiv:2002.07394*.
- 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.
- Lin, Y.; Zhang, J.; Huang, Z.; Liu, J.; Wen, Z.; and Peng, X. 2024. Multi-granularity correspondence learning from long-term noisy videos. *arXiv preprint arXiv:2401.16702*.
- Liu, Y.; Ott, M.; Goyal, N.; Du, J.; Joshi, M.; Chen, D.; Levy, O.; Lewis, M.; Zettlemoyer, L.; and Stoyanov, V. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Ma, Y.; Xu, G.; Sun, X.; Yan, M.; Zhang, J.; and Ji, R. 2022. X-clip: End-to-end multi-grained contrastive learning for video-text retrieval. In *Proceedings of the 30th ACM International Conference on Multimedia*, 638–647.
- Neverova, N.; Novotny, D.; and Vedaldi, A. 2019. Correlated uncertainty for learning dense correspondences from noisy labels. *Advances in Neural Information Processing Systems*, 32.
- Northcutt, C.; Jiang, L.; and Chuang, I. 2021. Confident learning: Estimating uncertainty in dataset labels. *Journal of Artificial Intelligence Research*, 70: 1373–1411.

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.

Song, Y.; and Soleymani, M. 2019. Polysemous visual-semantic embedding for cross-modal retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1979–1988.

Tan, C.; Xia, J.; Wu, L.; and Li, S. Z. 2021. Co-learning: Learning from noisy labels with self-supervision. In *Proceedings of the 29th ACM International Conference on Multimedia*, 1405–1413.

Wang, Q.; Zhang, Y.; Zheng, Y.; Pan, P.; and Hua, X.-S. 2022. Disentangled Representation Learning for Text-Video Retrieval. *arXiv:2203.07111*.

Wang, Y.; Ma, X.; Chen, Z.; Luo, Y.; Yi, J.; and Bailey, J. 2019. Symmetric cross entropy for robust learning with noisy labels. In *Proceedings of the IEEE/CVF international conference on computer vision*, 322–330.

Wang, Y.; Wang, J.; Chen, B.; Zeng, Z.; and Xia, S.-T. 2024. GMMFormer: Gaussian-Mixture-Model based Transformer for Efficient Partially Relevant Video Retrieval. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 5767–5775.

Wei, H.; Feng, L.; Chen, X.; and An, B. 2020. Combating noisy labels by agreement: A joint training method with co-regularization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 13726–13735.

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.

Zhang, B.; Hu, H.; Lee, J.; Zhao, M.; Chammas, S.; Jain, V.; Ie, E.; and Sha, F. 2020. A hierarchical multi-modal encoder for moment localization in video corpus. *arXiv preprint arXiv:2011.09046*.

Zhang, H.; Sun, A.; Jing, W.; Nan, G.; Zhen, L.; Zhou, J. T.; and Goh, R. S. M. 2021. Video corpus moment retrieval with contrastive learning. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 685–695.

Zheng, Z.; and Yang, Y. 2021. Rectifying pseudo label learning via uncertainty estimation for domain adaptive semantic segmentation. *International Journal of Computer Vision*, 129(4): 1106–1120.