Deconfounded Video Moment Retrieval with Causal Intervention

Xun Yang^{1,2}, Fuli Feng^{1,2}, Wei Ji^{1,2}, Meng Wang³, Tat-Seng Chua^{1,2}

¹Sea-NExT Joint Lab, Singapore

²School of Computing, National University of Singapore, Singapore

³School of Computer Science and Information Engineering, Hefei University of Technology, China
{hfutyangxun,fulifeng93}@gmail.com;wangmeng@hfut.edu.cn;{jiwei,dcscs}@nus.edu.sg

ABSTRACT

We tackle the task of video moment retrieval (VMR), which aims to localize a specific moment in a video according to a textual query. Existing methods primarily model the matching relationship between query and moment by complex cross-modal interactions. Despite their effectiveness, current models mostly exploit dataset biases while ignoring the video content, thus leading to poor generalizability. We argue that the issue is caused by the hidden confounder in VMR, *i.e.*, *temporal location of moments*, that spuriously correlates the model input and prediction. How to design robust matching models against the temporal location biases is crucial but, as far as we know, has not been studied yet for VMR.

To fill the research gap, we propose a causality-inspired VMR framework that builds structural causal model to capture the true effect of query and video content on the prediction. Specifically, we develop a Deconfounded Cross-modal Matching (DCM) method to remove the confounding effects of moment location. It first disentangles moment representation to infer the core feature of visual content, and then applies causal intervention on the disentangled multimodal input based on backdoor adjustment, which forces the model to fairly incorporate each possible location of the target into consideration. Extensive experiments clearly show that our approach can achieve significant improvement over the state-of-the-art methods in terms of both accuracy and generalization.

CCS CONCEPTS

• **Information systems** → **Multimedia and multimodal retrieval**.

KEYWORDS

Multimedia retrieval; cross-media reasoning; query-based moment retrieval; causal intervention; out-of-distribution testing

ACM Reference Format:

Xun Yang, Fuli Feng, Wei Ji, Meng Wang, Tat-Seng Chua. 2021. Deconfounded Video Moment Retrieval with Causal Intervention. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '21)*, July 11–15, 2021, Virtual Event, Canada. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3404835.3462823>

This work is licensed under a Creative Commons Attribution International 4.0 License.

SIGIR '21, July 11–15, 2021, Virtual Event, Canada.

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8037-9/21/07.

<https://doi.org/10.1145/3404835.3462823>

1 INTRODUCTION

Text-based visual retrieval is a fundamental problem in multimedia information retrieval [6, 7, 12, 22, 40] which has achieved significant progress. Recently, Video Moment Retrieval (VMR) has emerged as a new task [1, 9, 10, 17] and received increasing attention [4, 18]. It aims to retrieve a specific moment in a video, according to a textual query. As shown in Figure 1 (a), given a query "Person opens a door", its goal is to predict the temporal location of the target moment.¹ This task is quite challenging, since it requires effective modeling of multimodal content and cross-modal relationship. The past few years have witnessed the notable development of VMR, largely driven by public benchmark datasets, *e.g.*, Charades-STA [9]. Most existing efforts first generate sufficient and diverse moment candidates, and then rank them based on their matching scores with the query. Benefiting from effective modeling of temporal contexts and cross-modal interactions, state-of-the-art (SOTA) models [20, 44, 51] have reported significant performance on public datasets. However, does the reported results of SOTA models truly reflect the progress of VMR? The answer seems to be "No".

Recent studies [23] found empirical evidences that VMR models mostly exploit the temporal location biases in datasets, rather than learning the cross-modal matching. The true effect from the video content may be largely overlooked in the prediction. Such biased models have poor generalizability and are not effective for real-world application. We argue that the critical issue is rooted from two types of temporal location biases in dataset annotations:

- **Long-tailed annotations.** As depicted in Figure 1 (c), most annotations in Charades-STA start at the beginning of videos and last roughly 25% of the video length. This leads to a strong bias toward the frequently annotated locations. If a model is trained to predict the target at the *head* locations (yellow regions in Figure 1 (c)) more frequently than at the *tail* locations (green regions), the former is more likely to prevail over the latter during testing.
- **High correlation between user queries and moment locations.** In Figure 1 (d) and (e), the queries including a verb "open" mostly match target moments at the beginning of videos. While those corresponding to "leave" are often located at the end. That makes user queries and moment locations spuriously correlated.

Although the temporal location bias misleads the model prediction severely, it also has some good, *e.g.*, the temporal location of a moment helps to better model the temporal context for answering the query with temporal languages [10]. Therefore, the key of VMR is how to properly keep the good and remove the bad effect of

¹Through this paper, unless otherwise stated, we use the term "**location**" to depict the "**temporal location**" of a video moment, rather than "**physical location**."

Figure 1: (a) An example of VMR; (b) Causal graph: Q (query), V (video moment), L (moment location), and Y (prediction); The bottom three plots depict the temporal location distributions of moment annotations in Charades-STA based on the 2D structure of moments [51]: (c) Distribution of all moments; (d) and (e) describe the distributions of query-specific moments (i.e., including verbs like "open" or "leave").

moment location for a robust cross-modal matching. As far as we know, no similar effort is devoted to address this critical issue.

In this paper, we present a causal VMR framework that builds structural causal model [24] to capture the causality between different components in VMR. The causal graph is depicted in Figure 1 (b) which consists of four variables: Q (query), V (video moment), Y (prediction), and L (moment location). The prediction of traditional models is inferred by using the probability $P(Y|Q, V)$. In the causal graph, L is a hidden confounder [25] that spuriously associates V and Y , which has been long ignored by traditional VMR models. The effect of L on Y mainly comes from the biases in datasets, while the effect of L on V is mainly due to the entanglement of latent location factor in the representation of V . Here, the confounder L opens a backdoor path: $V \leftarrow L \rightarrow Y$ that hinders us from finding the true effect on Y only caused by (Q, V) . To remove the harmful confounding effects [25] caused by moment location, we design a **Deconfounded Cross-modal Matching** (DCM) method. It first disentangles the moment representation to infer the core feature of visual content, and then applies causal intervention on the multimodal input based on backdoor adjustment [24]. Specifically, we replace $P(Y|Q, V)$ with $P(Y|do(Q, V))$ by *do* calculus which encourages the model to make an unbiased prediction. The query is forced to fairly interact with all possible locations of the target based on the intervention. We expect the deconfounding model to be able to robustly localize the target even when the distribution of moment annotations is significantly changed. This triggers another key question: how to evaluate the generalizability of VMR models?

In general, the split of training and testing sets follows the independent and identically distributed (IID) assumption [32]. Due to the existence of temporal location biases in the datasets, the IID testing sets are insufficient to test the causality between (Q, V) and Y . To address this issue, we introduce the **Out-Of-Distribution** (OOD) testing [30] into evaluation. Our basic idea is to modify the location distribution of moment annotations in the testing sets by inserting a short video at the beginning/end of testing videos. It is only applied on testing sets, which does not introduce any new

biases that affect the model training. Extensive experiments on both IID and OOD testing sets of ActivityNet-Captions [15], Charades-STA [9], and DiDeMo [1], clearly demonstrate that DCM not only achieves higher retrieval performance but also endows the model with strong generalizability against distribution changes.

In summary, this work makes the following main contributions:

- We build a causal model to capture the causalities in VMR and point out that the moment temporal location is a hidden confounding variable that hinders the pursuit of true causal effect.
- We propose a new method, DCM, for improving VMR, that applies causal intervention on the multimodal model input to remove the confounding effects of moment location.
- We conduct extensive experiments on the IID and OOD testing sets of three benchmark datasets, which show the ability of DCM to perform accurate and robust moment retrieval.

2 CAUSAL VIEW OF VMR

In this section, we first formally formulate the task of VMR and then introduce a causal graph to clearly reveal how the confounder L spuriously correlate the multimodal input and prediction.

2.1 Problem Formulation

Given a query and a video with temporal length τ , our task is to identify a target moment $\hat{v}$ in the given video, starting at timestamp τ_s and ending at timestamp τ_e ($\tau_s < \tau_e < \tau$), which semantically matches the query q . The given video is first processed into a set of video moment candidates $\{v_i\}_{i=1}^N$, with diverse temporal durations, either by multi-scale sliding window [9, 17, 18] or hierarchical pooling [51], where N is the number of moment candidates. Each moment v is indexed by a temporal location ℓ , i.e., a pair of timestamps (t_s, t_e) . We compute the relevance score y (e.g., IoU score) between each moment candidate v and target $\hat{v}$ as the supervision for training, where y can be either binary $\{0, 1\}$ or dense $(0, 1)$. The candidates with $y \geq 0.5$ are usually treated as positive. As such, the task of video moment retrieval can be formally defined as:

- **Input:** A corpus of query-moment pairs and their relevance scores: $\{Q, \mathcal{V}, \mathcal{Y}\}$, where Q , $\mathcal{V}$, and $\mathcal{Y}$ denote the sets of queries, moments, and relevance scores, respectively.
- **Output:** A cross-modal matching function $f: Q \times \mathcal{V} \rightarrow \mathbb{R}$ that maps each query-moment pair to a real value by effectively modeling intra-modality and inter-modality feature interactions. During inference, we rank the moment candidates based on the estimated matching scores, and return the best candidate to user.

The solution usually involves: 1) learning the representation of query $\mathbf{q} \in \mathbb{R}^d$ and moment $\mathbf{v} \in \mathbb{R}^d$, and 2) modeling the cross-modal relationship. We mainly focus on the second point. In particular, we will investigate how the temporal locations of video moments spuriously affect the prediction $f(\mathbf{q}, \mathbf{v})$, and then aim to remove the spurious correlation caused by the moment location.

2.2 Causal Graph

We describe the causalities among four variables in VMR: query Q , video moment V , moment location L , and model prediction Y with a simple causal graph in Figure 1 (b). The direct link denotes the causality between two nodes: cause $\rightarrow$ effect. The causal graph is a directed acyclic graph that indicates how the variables $\{Q, V, L, Y\}$

interact with each other through the causal links. Traditional VMR models, e.g., [9, 51], have only two links: $Q \rightarrow Y$ and $V \rightarrow Y$, and predict the target by joint probability $P(Y|Q, V)$. In our causal graph, the moment location L is a confounder [24, 25] that influences both the correlated variable Y and independent variable V , leading to spurious correlation between V and Y :

- $L \rightarrow Y$ is rooted from the frequency prior $P(Y|L)$. Dataset annotators prefer to select short moments at specific locations of videos as the target [1], e.g., the beginning or end of videos, as depicted in Figure 1 (c), which leads to a *long-tailed* location distribution of moment annotations. It introduces a harmful causal effect that misleads the prediction of targets at the *tail* (green regions) location biased towards the data-rich *head* (yellow regions) locations. The location L leads to not only the direct effect $L \rightarrow Y$ but also the joint effect $(Q, L) \rightarrow Y$, which will be investigated in section 4.
- $L \rightarrow V$ denotes the effect of L on moment representation. In VMR, the temporal location of moments leads to a temporal context prior that affects the representation learning of moments in videos. The extracted moment representation from a backbone network is usually entangled with a latent location factor. Although it enriches the representation with location-aware features, it introduces the harmful spurious correlation between (Q, V) and Y via the backdoor path: $V \leftarrow L \rightarrow Y$. Similar feature entanglements are common in video domain [5].

So far, we have clearly revealed how the moment location L confound V and Y via the backdoor path that makes the effect of visual content largely overlooked, as reported by [23]. Next, we propose a deconfounding method to remove the confounding effects.

3 METHODOLOGY

Given the representation of query q and a moment candidate v , our goal is to predict the matching score s between the candidate and the query by a deconfounding method, DCM. As depicted in Figure 2, it is composed of two key steps to achieve the deconfounding: 1) feature disentangling on V that infers the core feature to represent the visual content of moment candidate, reducing the effect of L on v ; 2) causal intervention on the multimodal input to remove the harmful confounding effects of L from Y .

3.1 Moment Representation Disentangling

As discussed before, the moment representation is usually entangled with a location factor. To achieve the deconfounding, the first step is to ensure that the representations of V and L are independent with each other. In this section, we propose to disentangle v into two independent latent vectors: content $c^v \in \mathbb{R}^d$ and location $\ell^v \in \mathbb{R}^d$:

$$c^v = g_c(v), \ell^v = g_\ell(v), \quad (1)$$

where $g_c(\cdot)$ and $g_\ell(\cdot)$ can be implemented by two fully connected layers. To achieve the disentangling, we introduce two learning constraints to optimize the parameters of $g_c(\cdot)$ and $g_\ell(\cdot)$.

- **Reconstruction.** Since we can access the original location (t_s, t_e) of each moment v during training, we use it to supervise $g_\ell(\cdot)$ by a reconstruction loss $\mathcal{L}_{recon}(\ell^v, p)$, where $p \in \mathbb{R}^d$ is a non-learnable positional embedding vector [33] of (t_s, t_e) . $\mathcal{L}_{recon}(\cdot, \cdot)$ can be any reconstruction losses that force ℓ^v to approximate the real location feature. Different from existing efforts [1, 20, 47]

Figure 2: A brief overview of our deconfounded cross-modal matching for the task of VMR.

that exploit (t_s, t_e) as the side information to augment the moment representation, we focus on disentangling v to capture the core effect of visual content and remove the confounding effects of moment location. Injecting extra location feature into v will obviously exacerbate the correlation between L and V .

- **Independence Modeling.** Different from ℓ^v , there is no supervision on the content c^v . Besides, c^v and ℓ^v are both estimated from v , they do not immediately satisfy the *independence* requirement. We introduce an independence learning regularizer $\mathcal{L}_{indep}(\ell^v, c^v)$ that forces ℓ^v to be independent with c^v in the latent space. $\mathcal{L}_{indep}(\cdot, \cdot)$ can be any statistical dependence measures, e.g., distance correlation and mutual information.

We will describe the details of $\mathcal{L}_{indep}$ and $\mathcal{L}_{recon}$ in section 3.3.4. The representation disentangling on V reduces the correlation between V and L , which cuts off the link $L \rightarrow V$ in the feature level. The next question is how to remove the harmful confounding effects of L from Y for making a robust prediction on each candidate.

3.2 Causal Intervention

As described before, L is a confounder that leads to spurious correlation between V and Y . For example, in an unbalanced VMR dataset, the high-frequency location ℓ_h , has higher chance to be selected as the target location than the low-frequency location ℓ_l : $P(Y|\ell_h) > P(Y|\ell_l)$. That will mislead the model towards predicting a higher score on the candidate v_{ℓ_h} at ℓ_h than the candidate v_{ℓ_l} at ℓ_l , without truly looking into the details of multimodal input. Even if the query is considered, if not removing the confounding effects of moment location, the model will still be easily misled by the query-specific location prior (Figure 1 (d) and (e)). In this section, as shown in Figure 3, we propose to use *do*-calculus [25] to intervene the model input based on back-door adjustment. By applying the Bayes rule on the new casual graph, we have

$$P(Y|do(Q, V)) = P(Y|Q, do(V)) = \sum_{\ell \in L} P(Y|Q, V, \ell) P(\ell), \quad (2)$$

where $P(\ell)$ is the prior of location. We see from Eq. (2) that the aim of the *intervention* is to force the query to fairly interact with all the possible locations of the target moment for making an unbiased prediction, subject to the prior $P(\ell)$. In this way, the model is prevented from memorizing the corresponding locations of target moments during training. Note that we force the prior $P(\ell)$ to be a constant $\frac{1}{N}$, following the assumption that each location has equal opportunity to be the target location. For the location confounder set L , since we can hardly enumerate all the moment locations in the real world, we approximate L as the location set of all sampled

Figure 3: A brief illustration of our causal intervention.

moment candidates, i.e., $L = \{\ell_k\}_{k=1}^N$, where ℓ_k is represented by $\ell_k \in \mathbb{R}^d$, i.e., the disentangled location feature from v_k .

Given the representations of the query and a moment candidate, Eq. (2) is implemented as $\sum_{\ell} p(y|q, v, \ell)P(\ell)$,² where $p(y|q, v, \ell)$ is the output of a cross-modal matching network $f(\cdot)$:

$$p(y|q, v, \ell) = \sigma(f(q, v, \ell)), \quad (3)$$

where $\sigma(\cdot)$ is the sigmoid function $\sigma(x) = \frac{1}{1+\exp(-x)}$ that transforms the output of $f(\cdot)$ into (0, 1). In summary, the implementation of Eq. (2) is formally defined as

$$P(Y|do(Q, V)) := \mathbb{E}_{\ell} [\sigma(f(q, v, \ell))]. \quad (4)$$

To calculate Eq. (4), we found that $\mathbb{E}_{\ell}$ needs expensive sampling. Fortunately, following recent efforts [46, 49], we can first adopt the Normalized Weighted Geometric Mean (NWGM) [39] to approximately move outer expectation $\mathbb{E}_{\ell}$ into the sigmoid function:

$$P(Y|do(Q, V)) := \mathbb{E}_{\ell} [\sigma(f(q, v, \ell))] \approx \sigma(\mathbb{E}_{\ell} [f(q, v, \ell)]). \quad (5)$$

The further calculating of Eq. (5) depends on the implementation of the matching network $f(\cdot)$. If $f(\cdot)$ is a linear model, we can further move the expectation into $f(\cdot)$, i.e., $\mathbb{E}(f(x)) = f(\mathbb{E}(x))$. Then we just need one forward pass to obtain the prediction. Besides, to seek the true effect of visual content of V on Y , the representation v is implemented as the disentangled moment representation. We will describe the detailed implementation of $f(\cdot)$ in section 3.3.3. So far, we have finished the causal intervention via backdoor adjustment [24, 25], thus closing the backdoor path $V \leftarrow L \rightarrow Y$.

Summary: From Eq. (5), we can see that the effect on Y comes from Q , V , and all the possible locations of the target moment, i.e., $\{\ell_k\}_{k=1}^N$, which prevents the model from exploiting the location priors $P(Y|L)$ and $P(Y|Q, L)$. The disentangling module described in section 3.1 supports the deconfounding in the *feature* level while the intervention module contributes more in the *prediction* level. Both of them are indispensable for achieving the deconfounding.

3.3 Model Implementation

This section presents the model implementation. It mainly consists of the moment candidate sampling and encoding, query encoding, cross-modal matching networks $f(\cdot)$, and the learning objective.

3.3.1 Moment Sampling and Encoding. For a given video, we first segment it into a sequence of T non-overlapping video clips by fixed-interval sampling. The clip feature can be obtained by applying pretrained CNNs, such as C3D [31] and I3D [3], on video frames. Then we enumerate all moment candidates from the given video by composing any ordered lists of video clips. Each clip can be also treated as a short moment. The representation $v \in \mathbb{R}^d$ of each candidate is obtained by applying a stacked pooling network over the set of video clips, following [51]. All the feature vectors of

²To keep section 3.2 self-contained, we still use v to denote the moment representation for simplicity. It can either be the original or disentangled representation of V .

sampled moment candidates can be encoded into a moment feature tensor $V \in \mathbb{R}^{d \times T \times T}$. The last two dimensions of V index the start and end coordinates of all sampled candidates.

3.3.2 Query Encoding. For the encoding of the query sentence q , we first extract the word embeddings by the pretrained Glove word2vec model [13]. We then use a three-layer unidirectional LSTM to capture the sequential information in q . The query representation $q \in \mathbb{R}^d$ is the last hidden state of the LSTM output.

3.3.3 Deconfounded Cross-modal Matching (DCM). Our proposed DCM method can be implemented with most existing matching architectures in VMR (proposals-based). In this work, we implement it with two kinds of cross-modal matching networks $f(\cdot)$:

- Context-aware Multi-modal Interaction (CMI). It is modified from the fusion module in [9]. We implement $f(\cdot)$ with CMI as:

$$f(q, v, \ell) = w^T \left(w_1 \left((\bar{q} \odot \phi_v(\bar{v} + h_{qv}(\ell))) \oplus (\bar{q} + \phi_v(\bar{v} + h_{qv}(\ell))) \right) \right), \quad (6)$$

where $\bar{v}$ and $\bar{q}$ denote the transformed moment and query representations, respectively: $\bar{v} = c^v + W_2 \ell^v$ and $\bar{q} = W_3 q$. The $w \in \mathbb{R}^d$, $w_1 \in \mathbb{R}^{d \times 2d}$, $w_2 \in \mathbb{R}^{d \times d}$, and $w_3 \in \mathbb{R}^{d \times d}$ denote the learnable parameters of four fully-connected layers, respectively. The $\odot$ denotes the element-wise multiplication and $\oplus$ denotes the channel-wise vector concatenation. $h_{qv}(\ell)$ is a feature transformation of ℓ , parameterized by the features of q and v . $\phi_v(\bar{v} + h_{qv}(\ell))$ denotes a convolutional operation on the intervened moment tensor $V' \in \mathbb{R}^{d \times T \times T}$, where each element has been changed from v to $\bar{v} + h_{qv}(\ell)$, to aggregate the pre-context and post-context moment features [9] into current moment v , with the weight parameter $W_4 \in \mathbb{R}^{d \times d \times K \times K}$ (K denotes the kernel size). We remove the non-linear activation function behind the convolution operation. Then, we replace $h_{qv}(\ell)$ in Eq. (6) with $\mathbb{E}_{\ell} [h_{qv}(\ell)]$ to approximately compute $\mathbb{E}_{\ell} [f(q, v, \ell)]$.

- Temporal Convolutional Network (TCN) [51]. It is a SOTA matching network in VMR. We implement TCN-based $f(\cdot)$ as:

$$f(q, v, \ell) = w^T \left(\phi_v^* \left(\bar{q} \odot (\bar{v} + h_{qv}(\ell)) \right) \right), \quad (7)$$

where $\phi_v^*(\cdot)$ denotes a multi-layer 2D CNN over the intervened moment tensor $V' \in \mathbb{R}^{d \times T \times T}$ that captures the temporal dependencies between adjacent moments into current moment v . Eq. (7) is non-linear due to the rectified linear unit $Relu(\cdot)$ behind each convolutional layer. Fortunately, based on the theoretical analysis of [2] on the rectified linear unit, the approximation $\mathbb{E}(Relu(x)) \approx Relu(\mathbb{E}(x))$ still holds. Then we can obtain

$$\mathbb{E}_{\ell} [f(q, v, \ell)] \approx w^T \left(\phi_v^* \left(\bar{q} \odot (\bar{v} + \mathbb{E}_{\ell} [h_{qv}(\ell)]) \right) \right). \quad (8)$$

The key of the two above-mentioned cross-modal matching networks $f(\cdot)$ is to compute $\mathbb{E}_{\ell} [h_{qv}(\ell)]$. We implement $h_{qv}(\ell)$ as the scaled Dot-Product attention [33] to adaptively assign weights on different location confounders in $L = [\ell_1, \dots, \ell_N] \in \mathbb{R}^{N \times d}$ with specific input of query q and moment v . Then we have

$$\mathbb{E}_{\ell} [h_{qv}(\ell)] = \sum_{\ell} [\text{softmax}(K^T m / \sqrt{d}) \odot L'] P(\ell), \quad (9)$$

where $\odot$ denotes the element-wise product that supports broadcast, and $m = W_5 q + W_6 c^v$, $K = W_7 L^T$, and $L' = L W_2^T$ with learnable parameters W_5, W_6, W_7 , and $W_2 \in \mathbb{R}^{d \times d}$.

So far, we have introduced the implementation of the proposed DCM. Note that, we do not use the feature of original moment location (t_s, t_e) in the inference stage. Given the representation of query $\mathbf{q}$ and moment candidate $\mathbf{v}$, we can predict the deconfounded cross-modal matching score as $s = \sigma(\mathbb{E}_{\ell} [f(\mathbf{q}, \mathbf{v}, \ell)])$, where the confounder vector ℓ is disentangled from moment representation.

3.3.4 Learning Objective. To train the DCM-based VMR methods, we use the scaled Intersection over Union (IoU) score [51] between the locations of moment candidate v and groundtruth v_{gt} as the supervision y . The basic loss is the binary cross entropy loss:

$$\mathcal{L}_{bce}(\mathbf{q}, \mathbf{v}, y) = -(y \log(s) + (1 - y) \log(1 - s)). \quad (10)$$

In this work, we not only train our method with positive query-video training pairs but also exploit the negative query-video pairs as **counterfactual** training data to optimize the network parameters. Existing efforts all assume that the existence of target moment in the given video is true, forgoing teaching the model the ability of *counterfactual thinking*, i.e., *What if the target moment does not exist in the given video?* By introducing the *counterfactual thinking*, we expect the learned model to focus more on the content of video and query, rather than the location of moment candidates. In summary, our model is trained with a linear combination of four losses:

$$\mathcal{L} = \mathcal{L}_{bce}^+ + \mathcal{L}_{bce}^- + \lambda_1 \mathcal{L}_{recon} + \lambda_2 \mathcal{L}_{indep}, \quad (11)$$

where the **counterfactual loss** $\mathcal{L}_{bce}^-$ is computed over negative query-video pairs with the supervision $y = 0$. $\mathcal{L}_{recon}$ and $\mathcal{L}_{indep}$ are two learning terms for feature disentangling in section 3.1 with two hyperparameters λ_1 and λ_2 . We implement the reconstruction term by $\mathcal{L}_{recon} = \|\ell^v - \mathbf{p}\|_2$ for its simplicity. The independence term is implemented based on distance correlation [29]: $\mathcal{L}_{indep} = \text{dCov}(\ell^v, \mathbf{c}^v) / \sqrt{\text{dVar}(\ell^v) \cdot \text{dVar}(\mathbf{c}^v)}$, where $\text{dCov}(\cdot)$ and $\text{dVar}(\cdot)$ denote the distance covariance and variance, respectively. The overall loss is computed as the average over a training batch.

4 EXPERIMENTS

In this section, we evaluate the effectiveness of our DCM methods by extensive comparison with baseline methods (CMI and TCN) and SOTA VMR methods in both IID and OOD settings.

4.1 Datasets and Experimental Setting

4.1.1 Datasets. We validate the performance of DCM methods on three public datasets: **ActivityNet-Captions** (ANet-Cap) [15], **Charades-STA** [9], and **DiDeMo** [1]: (1) *ANet-Cap* depicts diverse human activities in daily life, consisting of about 20k videos taken from the ActivityNet caption dataset [15]. We follow the setup in [44, 51, 52] to split the dataset; (2) *Charades-STA* is constructed for VMR by [9] based on the Charades video dataset [28]. The videos are about daily indoor activities. Currently, it is the most popular VMR dataset; (3) *DiDeMo* is prepared by [1] with more than 10k videos from YFCC100M. It depicts diverse human activities in real world. The detailed dataset statistics are summarized in Table 1.

4.1.2 Metrics. We adopt the rank-1 accuracy (**R@1**) with different IoU thresholds (**IoU>m**), and **mIoU** as the evaluation metrics, following [9, 50, 51]. The "R@1, IoU>m" denotes the percentage of queries having top-1 retrieval result whose IoU score is larger than m . "mIoU" is the average IoU score of the top-1 results over all

Table 1: The statistics of three public VMR datasets.

Dataset	#Videos	#Anno. (train/val/test)	$\bar{L}_{vid}$	$\bar{L}_{mom} \pm \Delta_{mom}$
ANet-Cap	19,207	37,417/17,505/17,031	117.6s	36.2±40.2s
Charades-STA	6,672	12,408/-/3,720	30.6s	8.2±3.6s
DiDeMo	10,464	132,233/16,765/16,118	30s	6.85±3.5s

¹ #Anno. denotes the number of query-moment annotation pairs in different sets (train/val/test).

² $\bar{L}_{vid}$ and $\bar{L}_{mom}$ denote the average lengths of videos and moments, respectively.

³ Δ_{mom} denotes the standard deviation of moment length.

⁴ In DiDeMo [1], each query may have multiple moment annotations.

testing queries. In our experiments, we use $m \in \{0.5, 0.7\}$ for ANet-Cap and Charades-STA, and $m \in \{0.7, 1.0\}$ for DiDeMo. Different from existing efforts which also report "R@5" accuracy, we only report "R@1, IoU>m" ($m \geq 0.5$) and "mIoU" for stricter evaluation.

4.1.3 Out-of-Distribution Testing. The split of datasets in Table 1 follows the Independent and Identically Distributed (IID) assumption, which is insufficient to evaluate the generalization. Instead, we evaluate the models not only using IID testing, but also adopting Out-of-Distribution (OOD) testing [30]. The idea is to insert a randomly-generated video clip with ρ seconds at the beginning of videos. Then the temporal length of each video is changed to $\tau + \rho$ and the timestamps of target moment are changed accordingly, i.e., $(\tau_s + \rho, \tau_e + \rho)$. For each dataset, we select the best-performing model on original testing set to conduct **two rounds of OOD testing** with different ρ . We use $\rho \in \{30, 60\}$ for ANet-Cap, $\rho \in \{10, 15\}$ for Charades-STA, and $\rho \in \{15, 30\}$ for DiDeMo. The frame feature of the inserted video clip is randomly generated from a normal distribution. By the OOD testing, we keep the matching relationship unchanged while significantly changing the location distribution of moments to test the generalizability of models.

4.1.4 Comparison Methods. (1) *Baselines.* As mentioned in section 3.3.3, we implement our method with two matching networks: **CMI** and **TCN**. We use them as two baseline methods by just replacing $\tilde{\mathbf{v}} + \mathbb{E}_{\ell} [h_{qv}(\ell)]$ with original representation $\mathbf{v}$. Our methods are termed by adding a suffix: **+DCM**. The kernel size of the convolutional operation in CMI is 3×3 and TCN is implemented by a 3-layer CNN with kernel size 5×5 . We also include two **biases-based methods** that explicitly exploit the temporal location biases: a) **Freq-Prior**: it models the moment frequency prior $P(Y|L)$ where the prediction on each moment candidate only depends on the frequency of location annotations in training set; b) **Blind-TAN** [23]: it models $P(Y|Q, L)$ that excludes moment representation from the model input, implemented by replacing the moment feature in TCN with the positional embedding [33] of moment location (t_s, t_e) . (2) *State-of-the-art methods (SOTAs).* We also compare our methods with the reported results of SOTAs on IID testing sets, such as SCDM [44], 2D-TAN [51], DRN [47], LGI [20], and VSLNeT [50].

4.1.5 Implementation Details. For the language query, we use 300d glove vectors as the word embeddings. The dimension of query and moment representations is set to $d = 512$. The batchsize is set to 64 for ANet-Cap and DiDeMo, and 32 for Charades-STA, respectively. We adopt Adam as the optimizer with learning rate $1e - 4$. The hyperparameters λ_1 and λ_2 are fixed as 1 and 0.001, respectively. We use the PCA-reduced C3D [31] features as frame-level representation of ANet-Cap. We use three widely-used backbone networks (C3D, I3D [3], and VGG [14]) to extract video representation of Charades-STA, respectively. We use the flow features in

Table 2: Performance comparison (R@1,%) with state-of-the-arts (SOTAs) and baselines (Bases) on three datasets (The higher the better). The $\dagger$ and $\ddagger$ denote relative improvements larger than 5% and 10% w.r.t. R@1, respectively. The * and $\star$ denote the statistical significance w.r.t. mIoU for $p < 0.05$ and $p < 0.01$, respectively, compared with the baseline counterparts.

	Method	ANet-Cap (C3D)			Charades-STA(C3D)			Charades-STA(I3D)			Charades-STA(VGG)			DiDeMo (Flow)		
		IoU>m		mIoU	IoU>m		mIoU	IoU>m		mIoU	IoU>m		mIoU	IoU>m		mIoU
		0.5	0.7		0.5	0.7		0.5	0.7		0.5	0.7		0.7	1.0	
SOTAs (IID)	MAN[48]	-	-	-	-	-	-	-	-	-	46.5	22.7	-	-	27.0	41.2
	CBP [35]	35.8	17.8	36.9	36.8	18.9	35.7	-	-	-	-	-	-	-	-	-
	SCDM [44]	36.8	19.9	-	-	-	-	54.4	33.4	-	-	-	-	-	-	-
	2D-TAN[51]	44.5	26.5	43.2	47.0	27.2	42.0	53.7	31.2	47.0	42.8	23.3	-	35.3	25.6	47.8
	DRN [47]	45.5	24.4	-	45.4	26.4	-	53.1	31.8	-	42.9	23.7	-	-	-	-
	LGI [20]	43.0	25.1	42.6	50.5	27.7	44.9	59.5	35.5	51.4	44.7	24.0	40.8	-	-	-
VSLNeT [50]	43.2	26.2	43.2	47.3	30.2	45.2	53.3	33.7	49.9	39.2	20.8	40.3	-	-	-	
Bases (IID)	Freq-Prior	29.7	13.9	32.0	29.7	16.3	30.1	29.7	16.3	30.1	29.7	16.3	30.1	23.3	19.4	31.9
	Blind-TAN	45.3	28.6	43.9±0.4	40.2	23.5	36.7±0.2	40.2	23.5	36.7±0.2	40.2	23.5	36.7±0.2	24.4	19.4	34.5±1.4
	CMI	40.6	23.2	41.1±0.3	43.3	24.6	40.2±0.4	48.5	28.1	43.6±0.5	40.1	21.9	37.7±0.4	32.9	25.1	44.1±0.5
	CMI+DCM	41.4	23.9	41.6±0.3*	52.8‡	32.1‡	47.0±0.4*	57.5‡	34.9‡	50.4±0.3*	45.2‡	26.0‡	41.5±0.4*	37.6‡	27.8‡	49.4±0.5*
	TCN	43.3	26.1	42.3±0.5	48.0	28.9	43.1±0.5	52.6	32.3	46.3±0.4	43.1	25.0	39.3±0.8	35.1	26.0	47.4±0.7
TCN+DCM	44.9	27.7†	43.3±0.2*	55.8‡	34.4‡	48.7±0.5*	59.7‡	37.8‡	51.5±0.4*	47.8‡	28.0‡	43.1±0.4*	37.5†	27.6†	49.9±0.5*	
OOD-1	Freq-Prior	3.9	0	20.6	0.1	0	7.7	0.1	0	7.7	0.1	0	7.7	0	0	0
	Blind-TAN	16.2	6.4	22.0±0.2	15.0	6.1	19.2±0.4	15.0	6.1	19.2±0.4	15.0	6.1	19.2±0.4	3.9	2.1	6.8±2.8
	LGI [20]	16.3	6.2	22.2	26.2	11.4	26.8	42.1	18.6	41.2	24.1	8.2	27.8	-	-	-
	VSLNeT [50]	-	-	-	21.9	12.0	30.0	17.5	8.8	26.5	-	-	-	-	-	-
	CMI	12.3	5.2	19.1±0.5	32.1	15.3	36.1±2.5	30.4	16.4	30.3±4.2	20.5	8.0	25.6±7.4	21.4	18.1	32.7±0.6
	CMI+DCM	16.8‡	6.4‡	23.9±0.3*	33.7†	17.2‡	37.4±0.6	37.8‡	19.7‡	39.6±0.8*	24.5‡	11.5‡	31.3±0.8*	29.9‡	23.7‡	40.0±2.0*
	TCN	16.4	6.6	23.2±0.9	30.6	13.2	30.8±2.3	27.1	13.1	25.7±3.2	21.8	9.1	25.9±1.2	29.9	21.9	42.1±1.6
	TCN+DCM	18.2‡	7.9‡	24.4±0.5*	39.7‡	16.3‡	39.6±0.9*	44.4‡	19.7‡	42.3±1.1*	31.6‡	12.2‡	34.3±0.9*	31.6†	24.4‡	42.0±1.6
OOD-2	Freq-Prior	0.5	0	18.4	0	0	2.5	0	0	2.5	0	0	2.5	0	0	0
	Blind-TAN	11.1	3.7	17.8±0.4	13.4	4.4	16.1±0.5	13.4	4.4	16.1±0.5	13.4	4.4	16.1±0.5	3.9	2.1	6.7±2.5
	LGI [20]	11.0	3.9	17.3	20.3	7.4	22.1	35.8	13.5	37.1	18.8	5.3	24.3	-	-	-
	VSLNeT [50]	-	-	-	17.4	9.2	20.7	10.2	4.7	18.4	-	-	-	-	-	-
	CMI	10.0	4.2	16.8±0.6	28.5	12.3	35.3±2.4	28.1	13.6	29.0±4.8	16.0	5.31	24.5±6.9	18.3	15.7	25.9±3.4
	CMI+DCM	13.1	5.2	21.6±0.4*	30.5†	15.2‡	35.4±0.8	33.2‡	17.1‡	36.7±1.1*	21.1‡	9.6‡	29.1±0.5*	29.5‡	23.2‡	39.3±2.6*
	TCN	11.5	3.9	19.4±0.9	25.7	9.3	28.4±2.5	21.1	8.8	22.5±2.9	17.6	6.2	22.2±1.5	28.0	21.2	40.1±1.6
	TCN+DCM	12.9‡	4.8‡	20.7±0.5*	33.8‡	12.4‡	36.3±1.1*	38.5‡	15.4‡	39.0±1.2*	27.8‡	9.3‡	32.1±1.0*	30.0†	23.4‡	39.8±1.7

¹ We report the average performance of **10 runs on Charades-STA and DiDeMo**, and **5 runs on ANet-Cap**, with different random seeds for network initialization. For each run, we select the best-performing model on original testing set to conduct two rounds of OOD testing, i.e., OOD-1 and OOD-2, with different values of ρ mentioned in section 4.1.3. "-" means that the result on the dataset is not reported by the paper or its model is unavailable.

[1] as video representation of DiDeMo. We follow [51] to sample moment candidates from video and set the number of clips $T=16$ for ANet-Cap and Charades-STA and $T=6$ for DiDeMo. Besides, we exclude the long moments ($L_{mom}/L_{vid} \geq 0.5$, nearly 4k in testing set) for OOD testing on ANet-Cap, since we found the OOD performance of Freq-Prior on these long moments is quite high.

4.2 Overall Performance Comparison

4.2.1 Comparison with Baselines. We report the empirical results of all baseline methods in Table 2 on both IID (i.e., original) and OOD testing sets. The relative improvement and statistical significance test are performed between DCM methods and their baseline counterparts. We have the following observations:

- All three datasets have strong temporal location biases, revealed by the two biased methods. By using the prior $P(Y|L)$ without training, Freq-Prior reports nearly 30% mIoU on all IID testings. By using $P(Y|Q, L)$, Blind-TAN reports much higher R@1(IoU>0.5) accuracy. In particular, a SOTA result (45.3%) is reported in ANet-Cap. When evaluated by OOD testing, their performances drop significantly. The results support the empirical analysis in [23] and indicate the necessity of OOD testing.
- The DCM methods consistently improve TCN and CMI in all settings, which suggests that DCM is agnostic to the methods,

datasets, backbones, and distributions of testing sets. In particular, the improvements on OOD testing are typically larger than those on IID testing. For example, the relative improvement w.r.t. R@1(IoU>0.5) is 2.82% on IID testing and 23.68% on OOD testing on ANet-Cap. It reflects that traditional VMR models are vulnerable to the biases in datasets and sensitive to the distribution changes of moment annotations, and clearly demonstrates the high effectiveness of our DCM. We attribute such improvement to the following aspects: 1) the disentangling module reduces the effect of location on moment representation in the feature level, endowing the models with a stronger ability of precisely exploiting the video content; and 2) by the causal intervention, DCM is forced to fairly consider all the possible locations of target moment, via backdoor adjustment, to make a robust prediction which avoids exploiting the temporal location biases in datasets and therefore improves the generalizability of models.

- We find that the improvements of DCM on Charades-STA are much larger than those on the other datasets. It might be due to its relatively small training set, as shown in Table 1. Models can be more easily misled by the temporal location biases when trained with less data. In particular, the TCN method, with more training parameters, reports much worse generalization performance than the lightweight method, CMI, on the small dataset.
- Even on the two OOD testing sets, Blind-TAN still reports more than 10% R@1(IoU>0.5) score on ANet-Cap and Charades-STA,

Table 3: Ablation studies. The * and ★ denote the statistical significance for $p < 0.05$ and $p < 0.01$, respectively.

	TCN+DCM	ANet-Cap		Charades.		DiDeMo	
		IoU>m		IoU>m		IoU>m	
IID Test		0.5	0.7	0.5	0.7	0.7	1.0
	Full Model	44.9	27.7	59.7	37.8	37.5	27.6
	w/o feat. disent.	43.7*	26.7*	55.0*	33.7*	36.2*	26.7*
	w/o indep. loss	44.7	27.7	59.8	37.4	37.1*	27.3
	w/o recon. loss	42.7*	25.7*	55.8*	34.5*	35.7*	26.6*
	w/o loc. feat.	44.4*	27.1*	59.8	37.3*	37.1*	27.4*
	w/o counterf. loss	45.3	28.9*	58.2*	37.3*	37.0*	26.7*
	w/o causal interv.	45.6	29.5	57.7*	37.1*	36.8*	26.9*
OOD-1 Test	Full Model	18.2	7.9	44.4	19.7	31.6	24.4
	w/o feat. disent.	17.7	7.3*	34.3*	17.0*	29.3*	22.4*
	w/o indep. loss	18.1	7.7	43.8	19.1	31.0	23.9
	w/o recon. loss	17.8	7.2	38.2*	18.6	30.7	22.7*
	w/o loc. feat.	17.7	7.1*	43.5	19.2	28.4*	22.3*
	w/o counterf. loss	17.0*	7.3	40.2*	18.0*	31.1	23.5
	w/o causal interv.	16.4	6.7*	34.1*	15.9*	31.1	23.3*

which indicates the high correlation between query and location. In fact, the prior $P(Y|Q, L)$ can be mitigated by re-split the dataset based on moment locations, i.e., forcing the locations of target moments in testing set to be unseen in training set. For easy comparison with existing results, we do not re-split the datasets.

- We find the OOD evaluation results on ANet-Cap are much lower than those of other datasets, since we exclude the long moments from its OOD testing. The reason is that moment length also leads to a prior in $P(Y|L)$. For the long target moments with length $L_{\text{mom}}/L_{\text{vid}} \geq 0.5$, we can obtain 100% R@1(IoU>0.5) by just returning the whole video as the prediction. The number of long moments in the testing set of ANet-Cap is nearly 4K. We remove them to keep the effectiveness of OOD testing.

4.2.2 Comparison with State-of-the-art Results. We include the reported results of recent SOTA methods, such as LGI [20] and VSLNet [50], in Table 2. We observe that TCN+DCM and CMI+DCM achieve new SOTA results w.r.t. R@1(IoU>0.7) on IID testing: 27.7% on ANet-Cap, 27.8% on DiDeMo, and 34.4%, 37.8%, 28% on Charades-STA with different backbones. We also report the OOD testing results of LGI and VSLNet using their released models or codes. We find that our DCM methods significantly outperform SOTA methods w.r.t. R@1(IoU>0.7) on OOD testing, e.g., 27.4% relative improvement on OOD-1 on ANet-Cap and 26.7% improvement on OOD-2 on Charades-STA (I3D) over LGI. The result is consistent with the comparison with baseline counterparts. It demonstrates that DCM can not only retain the good location bias to improve the accuracy in original testing set, but also remove the bad confounding effects of location to improve the generalization.

4.3 Study of DCM

For a better understanding of DCM, we use TCN+DCM as an example to analyze the contribution of individual components. Empirical results on the three datasets³ are shown in Table 3 and Figure 4.

4.3.1 Effect of Feature Disentangling. As a core of DCM, the disentangling module decomposes moment representation into two latent factors: content and location using two specific losses, which

³If not specifically stated, we use the I3D backbone for Charades-STA by default.

Figure 4: (a) Distance correlation of TCN+DCM on 3 datasets; (b) Performance comparison (R@1, IID) on 4 query subsets of ANet-Cap with specific temporal languages.

ensures the visual content of moments be fairly treated and perceived. Here, we investigate, in Table 3, how this module affect the performance by alternatively removing the whole module (w/o feat. disent.), the independence loss (w/o indep. loss), the reconstruction loss (w/o recon. loss), and the disentangled location feature ℓ^v in $\bar{v}$ (w/o loc. feat.) from the model. We have several findings:

- We observe clear performance drop in all settings when removing the disentangling module. The reasons are two folds: 1) The entangled moment feature is weak in capturing the core feature of moment content. The latent location factor may dominate the moment representation, leading to poor generalization; and 2) without the disentangling, the confounder set has to be composed by the positional embedding of original locations (τ_s, τ_e). As mentioned in section 3.1, injecting the extra location feature into $\bar{v}$ may further exacerbate the correlation between V and L , which hurts the model expressiveness. The results support the analysis in [5] on the importance of video representation disentangling and clearly indicate the necessity of the module in DCM.
- We use distance correlation to model the independence between the moment content and location vectors. The comparisons, w.r.t. distance correlation and R@1 accuracy, are reported in Figure 4 (a) and Table 3, respectively. We observe that a better OOD performance is substantially coupled with a higher independence across the board. This also supports the correlation between the performance and feature disentangling mentioned above. w/o indep. loss hurts less the IID performance, since we use a small $\lambda_2 = 0.001$ to modulate its effect for stable training.
- w/o recon. loss leads to consistently worse IID and OOD results. The reason is that the latent location factor is not disentangled accurately, and then causal intervention fails to cut off the backdoor path. The result is consistent with the above analyses.
- In our implementation of DCM, we additionally keep the disentangled location feature ℓ^v in the moment representation $\bar{v}$ to slightly improve the weight of current location. We observe that w/o loc. feat. leads to clear performance drop in OOD testing. The results suggest the necessity of keeping the good location bias. The location feature facilitates the temporal context modeling, which is especially crucial to handle the queries including temporal languages, such before or after. That can be supported by the performance comparison across 4 query subsets with different temporal languages in Figure 4 (b). We find that, compared with the overall comparison (44.9% vs. 43.3% w.r.t. R@1(IoU>0.5)) between TCN+DCM and TCN on ANet-Cap, the improvements in Figure 4 (b) are more significant. It further verifies the advantages of DCM in exploiting the good location bias.

Figure 5: Performance comparison (R@1, IID) w.r.t. the temporal lengths of annotated moments in three datasets. The ratio of moment annotations in each group to the total number in training set is also plot.

4.3.2 Effect of Causal Intervention. Causal Intervention is the key module in DCM. It intervenes the multimodal input based on backdoor adjustment to cut off the backdoor path. It consists of two components: 1) representation intervention, *i.e.*, the addition of $\mathbb{E}_{\ell} [h_{qv}(\ell)]$ to $\bar{v}$; and 2) counterfactual loss $\mathcal{L}_{bce}^-$ that endows DCM with the ability of counterfactual thinking. We investigate its effect in Table 3 by alternatively removing the whole module (*w/o causal interv.*) and the counterfactual loss (*w/o counterf. loss*) from the model. We have the following observations:

- *w/o causal interv.* leads to significant performance drops in OOD testing results. Compared with the full model, the relative performance drops w.r.t. $R@1(IoU>0.7)$ are 15.2%, 19.3%, and 4.5% on three datasets, respectively. In particular, we surprisingly observe a large IID improvement (27.7→29.5) but a significant OOD drop (7.9→6.7) on ANet-Cap (the largest dataset) when removing causal intervention. This suggests that, without causal intervention, the model is at higher risk of exploiting dataset biases while making less use of video content. It is consistent with the remarkable performance of Blind-TAN on ANet-Cap in Table 2, which demonstrates that the issue of dataset biases is more serious in ANet-Cap. Overall, the results and analysis clearly indicate the effectiveness of causal intervention and further highlight the importance of introducing OOD testing into evaluation.
- If only removing the counterfactual loss $\mathcal{L}_{bce}^-$ from our learning objective in Eq. (11), we observe substantial performance drop of OOD testing results. This is consistent with the observation of *w/o causal interv.*, which suggests that counterfactual loss is a perfect complement to the representation intervention in the loss level. It is able to prevent DCM from exploiting the location biases by penalizing high prediction scores on the biased locations in a negative video, like the prediction of Freq-Prior and Blind-TAN.

4.3.3 On the Moment Length. We investigate how DCM affect the IID performance w.r.t. the moment length (*the shorter the more challenging*). The empirical results are reported in Figure 5. We have the following findings:

Figure 6: The sampled moment retrieval results, w.r.t. R@1, on both IID and OOD testing sets of Charades-STA.

- We find the temporal lengths of annotated moments in ANet-Cap are quite diverse and there are nearly 25% annotations with duration $L_{mom}/L_{vid} \geq 0.5$. The baseline TCN achieves higher performance w.r.t. $R@1(IoU>0.5)$ on the long moment groups. The improvement might be attributed to the usage of dataset bias on moment lengths, as analyzed at the end of section 4.2.1. The finding supports the observation in [23]. With causal intervention and counterfactual training, DCM-based method is prevented from exploiting such kind of dataset bias, thus reporting lower performance on these moment groups. While on the short moment groups in Figure 5 (a), TCN+DCM reports significantly better performance over TCN. These results are consistent with the observation on ANet-Cap in Table 3.
- We observe substantial improvement of TCN+DCM over its counterpart, w.r.t. $R@1(IoU>0.5)$, across all groups in Charades-STA. In particular, on the first group in Figure 5 (b) where the average duration of moments is just about 3s, we observe 168% relative improvement of TCN+DCM over TCN. The finding clearly demonstrates the advantages of DCM-based methods in retrieving very short video moments. We also observe a bad performance of TCN on the long moment group (0.6). It is not contrast to the observation from Figure 5 (a), since there are just 0.2% long moments with normalized length 0.6 in the training set of Charades-STA and thus the model can not learn the length bias of long moments with such limited data.
- For the comparison on DiDeMo, we observe from Figure 5 (c) that most of its annotations fall into the short moment group (0.2). Our DCM method achieves significant improvement (13.4%, w.r.t. $R@1(IoU>0.7)$) over its counterpart on this group.

Overall, the observations from Figure 5 clearly indicate the effectiveness of our DCM on moment retrieval in videos.

4.3.4 Qualitative Results. In this section, we present some qualitative results, as shown in Figure 6, to give an intuitive impression of the effectiveness of our DCM for VMR. Figure 6 shows a real example from Charades-STA: retrieving the specific moment of "Person opens a cabinet door" in the given video. We illustrate the R@1 retrieval results of TCN+DCM and its counterpart by both IID and OOD testing. We have the following findings:

- For the query in Figure 6, TCN+DCM performs stably against the distribution changes of moment location. In particular, in Figure 6 (a), its top-1 retrieval result is the moment at timestamps (0.0s, 3.7s), which has more than 70% overlap with the groundtruth annotation. When moved to the OOD testing by inserting a randomly-generated video segment at the beginning of the video, TCN+DCM can still find the best candidate with high IoU score. It reflects that DCM is not only robust to the *distribution changes* but also robust to the noisy temporal context.
- TCN is sensitive to the distribution changes. The IoU score of its top-1 retrieval result significantly drops from 58% in Figure 6 (a) to zero in (b). The result is consistent with the previous analysis. It indicates that the location biases have severely affected the training of TCN, leading to low robustness.

5 RELATED WORK

5.0.1 Video Moment Retrieval. Video moment retrieval (VMR) [1, 9], also termed temporal grounding/localization, attracts increasing attention in recent years [4, 10, 44]. It aims to retrieve a short video segments in videos using language queries. The key to solving VMR is how to effectively learn the cross-modal matching relationship between the language and video segments. Existing works mainly follow a cross-modal matching framework [1, 9, 10, 17, 18] that first generates a set of candidate moments, and then ranks them based on the matching score between the moment candidates and the language query. Gao *et al.* [9] and Liu *et al.* [17, 18] used the traditional multi-scale sliding window to sample the candidates, which is simple but has low recall of short candidates. Zhang *et al.* [51] proposed a stacked pooling network to densely enumerate all the valid candidates and store them in a three-dimensional tensor matrix to preserve the temporal structure, which allows CNNs to be applied in VMR for learning moment representation. The modeling of temporal structure of video moments, introduced in [51], becomes increasingly popular [16, 34]. The difference among existing efforts mainly lies in the design of multimodal fusion module, *e.g.*, cross-modal feature fusion [17, 20, 50] or pairwise similarity learning [1] following metric learning framework [42]. The former is more popular. Chen *et al.* [4], Yuan *et al.* [45], and Mun *et al.* [20] designed a complex visual-textual interaction/attention modules for cross-modal interaction. In particular, Mun *et al.* [20] presented a video-text interaction algorithm in capturing relationships of semantic phrases and video segments by modeling local and global contexts, reporting SOTA performance. Yuan *et al.* [44] and Zhang *et al.* [48, 51] applied layer-wise temporal CNNs to model the temporal relations between moments, receiving more attention recently.

Difference between DCM and existing works: 1) We focus on improving the generalizability of models, since recent findings [23] reported that SOTA models tend to exploit dataset biases while agnostic to video content. Specifically, our DCM applies causal intervention to remove the confounding effects of moment location, which encourages the model to fairly consider all the possible locations of target for making a robust prediction; 2) We introduce counterfactual training that endows the model with the ability of counterfactual thinking to answer the question "What if the target moment does not exist in the given video?". Different from our work, existing methods all assume that the target moment must be

in the give video; 3) We demonstrate the importance and necessity of applying OOD testing to evaluate the generalizability of VMR models, beyond the widely-used IID testing; 4) We empirically indicate that existing methods are vulnerable to the temporal location biases in datasets. Our DCM has the potential to be coupled with exiting methods to improve their generalizability.

5.0.2 Causal Inference. Recently, causal inference [24, 25] has attracted increasing attention in information retrieval and multimedia for removing dataset biases in domain-specific applications, such as recommendation [8, 27, 37, 38, 43], visual dialog [26], segmentation [49], unsupervised feature learning [36], video action localization [19], and scene graph [49], *etc.* The general purpose of causal inference is to empower the models the ability of pursuing the causal effect, thus leading to more robust decision. In particular, Qi *et al.* [26] presented causal principles to improve visual dialog, where causal intervention is used to remove the effect of an unobserved confounder. Zhang *et al.* [49] introduced causal inference into weakly-supervised semantic segmentation by removing confounding bias of context prior. Wang *et al.* [36] proposed to use causal intervention to learn "sense-making" visual knowledge beyond traditional visual co-occurrences. **Different from existing work,** we make a new attempt of causal inference [24] to solve the task of VMR. Specifically, we reveal the causalities between different components in VMR and point out that the temporal location of a video moment is a confounder that spuriously correlates the model prediction and multimodal input. To remove the harmful confounding effects, we develop a deconfounding method, DCM, that first disentangles moment representation to learn the core feature of visual content and then intervenes the multimodal input based on backdoor adjustment. This is the first causality-based work that addresses the temporal location biases of VMR, which is significantly different from a new VMR work [21].

6 CONCLUSION

This paper proposed a deconfounded cross-modal matching method for the VMR task. By introducing out-of-distribution testing into the model evaluation, we empirically demonstrated that existing VMR models are sensitive to the distribution changes of moment annotations, and found that the temporal location biases, hidden in current datasets, may severely affect the model training and prediction. Our proposed deconfounding algorithm is effective in improving the generalizability of VMR models. It can not only eliminate the bad but also retain the good bias of moment temporal location in videos. Our method has the potential to be effectively integrated with most exiting methods to improve their generalizability. This work showed initial attempts toward robust video moment retrieval against the dataset biases. In the future, we will further explore the intervention strategy on the user query to ensure a comprehensive understanding of complex user intents. Besides, we will also apply our deconfounding approach on other location-sensitive tasks, *e.g.*, visual grounding [11, 41] and temporal activity localization [19], to mitigate their location biases in datasets.

7 ACKNOWLEDGMENTS

This research/project is supported by the Sea-NExT Joint Lab.

REFERENCES

- [1] Lisa Anne Hendricks, Oliver Wang, Eli Shechtman, Josef Sivic, Trevor Darrell, and Bryan Russell. 2017. Localizing moments in video with natural language. In *ICCV*. 5803–5812.
- [2] Pierre Baldi and Peter Sadowski. 2014. The dropout learning algorithm. *Artificial intelligence* 210 (2014), 78–122.
- [3] Joao Carreira and Andrew Zisserman. 2017. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*. 6299–6308.
- [4] Jingyuan Chen, Xinpeng Chen, Lin Ma, Zequn Jie, and Tat-Seng Chua. 2018. Temporally grounding natural sentence in video. In *EMNLP*. 162–171.
- [5] Emily L Denton et al. 2017. Unsupervised learning of disentangled representations from video. In *NeurIPS*. 4414–4423.
- [6] Jianfeng Dong, Xirong Li, and Cees GM Snoek. 2018. Predicting visual features from text for image and video caption retrieval. *IEEE Transactions on Multimedia* 20, 12 (2018), 3377–3388.
- [7] Jianfeng Dong, Xirong Li, Chaoxi Xu, Xun Yang, Gang Yang, Xun Wang, and Meng Wang. 2021. Dual encoding for video retrieval by text. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2021).
- [8] Fuli Feng, Weiran Huang, Xiangnan He, Xin Xin, Qifan Wang, and Tat-Seng Chua. 2021. Should Graph Convolution Trust Neighbors? A Simple Causal Inference Method. In *SIGIR*.
- [9] Jiyang Gao, Chen Sun, Zhenheng Yang, and Ram Nevatia. 2017. Tall: Temporal activity localization via language query. In *ICCV*. 5267–5275.
- [10] Lisa Anne Hendricks, Oliver Wang, Eli Shechtman, Josef Sivic, Trevor Darrell, and Bryan Russell. 2018. Localizing Moments in Video with Temporal Language. In *EMNLP*. 1380–1390.
- [11] Richang Hong, Daqing Liu, Xiaoyu Mo, Xiangnan He, and Hanwang Zhang. 2019. Learning to compose and reason with language tree structures for visual grounding. *IEEE transactions on pattern analysis and machine intelligence* (2019).
- [12] Richang Hong, Yang Yang, Meng Wang, and Xian-Sheng Hua. 2015. Learning visual semantic relationships for efficient visual retrieval. *IEEE Transactions on Big Data* 1, 4 (2015), 152–161.
- [13] Richard Socher, Jeffrey Pennington, and Christopher D Manning. 2014. Glove: Global vectors for word representation. In *EMNLP*.
- [14] Andrew Zisserman, Karen Simonyan. 2015. Very Deep Convolutional Networks for Large-Scale Image Recognition. In *ICLR*.
- [15] Ranjay Krishna, Kenji Hata, Frederic Ren, Li Fei-Fei, and Juan Carlos Niebles. 2017. Dense-captioning events in videos. In *ICCV*. 706–715.
- [16] Daizong Liu, Xiaoye Qu, Xiao-Yang Liu, Jianfeng Dong, Pan Zhou, and Zichuan Xu. 2020. Jointly Cross-and Self-Modal Graph Attention Network for Query-Based Moment Localization. In *ACM MM*. 4070–4078.
- [17] Meng Liu, Xiang Wang, Liqiang Nie, Xiangnan He, Baoquan Chen, and Tat-Seng Chua. 2018. Attentive moment retrieval in videos. In *SIGIR*. ACM, 15–24.
- [18] Meng Liu, Xiang Wang, Liqiang Nie, Qi Tian, Baoquan Chen, and Tat-Seng Chua. 2018. Cross-modal moment localization in videos. In *ACM MM*. 843–851.
- [19] Yuan Liu, Jingyuan Chen, Zhenfang Chen, Bing Deng, Jianqiang Huang, and Hanwang Zhang. 2021. The Blessings of Unlabeled Background in Untrimmed Videos. In *CVPR*.
- [20] Jonghwan Mun, Minsu Cho, and Bohyung Han. 2020. Local-Global Video-Text Interactions for Temporal Grounding. In *CVPR*.
- [21] Guoshun Nan, Rui Qiao, Yao Xiao, Jun Liu, Sicong Leng, Hao Zhang, and Wei Lu. 2021. Interventional Video Grounding with Dual Contrastive Learning. In *CVPR*.
- [22] Liqiang Nie, Shuicheng Yan, Meng Wang, Richang Hong, and Tat-Seng Chua. 2012. Harvesting visual concepts for image search with complex queries. In *ACM MM*. 59–68.
- [23] Mayu Otani, Yuta Nakashima, Esa Rahtu, and Janne Heikkilä. 2020. Uncovering Hidden Challenges in Query-Based Video Moment Retrieval. In *BMVC*.
- [24] Judea Pearl, Madelyn Glymour, and Nicholas P Jewell. 2016. *Causal inference in statistics: A primer*. John Wiley & Sons.
- [25] Judea Pearl and Dana Mackenzie. 2018. *The book of why: the new science of cause and effect*. Basic Books.
- [26] Jiaxin Qi, Yulei Niu, Jianqiang Huang, and Hanwang Zhang. 2020. Two causal principles for improving visual dialog. In *CVPR*. 10860–10869.
- [27] Masahiro Sato, Sho Takemori, Janmajay Singh, and Tomoko Ohkuma. 2020. Unbiased Learning for the Causal Effect of Recommendation. In *Fourteenth ACM Conference on Recommender Systems*. 378–387.
- [28] Gunnar A Sigurdsson, Gül Varol, Xiaolong Wang, Ali Farhadi, Ivan Laptev, and Abhinav Gupta. 2016. Hollywood in homes: Crowdsourcing data collection for activity understanding. In *ECCV*. Springer, 510–526.
- [29] Gábor J Székely, Maria L Rizzo, Nail K Bakirov, et al. 2007. Measuring and testing dependence by correlation of distances. *The annals of statistics* 35, 6 (2007), 2769–2794.
- [30] Damien Teney, Kushal Kafle, Robik Shrestha, Ehsan Abbasnejad, Christopher Kanan, and Anton van den Hengel. 2020. On the Value of Out-of-Distribution Testing: An Example of Goodhart’s Law. In *NeurIPS*.
- [31] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. 2015. Learning spatiotemporal features with 3d convolutional networks. In *ICCV*. 4489–4497.
- [32] Vladimir N Vapnik. 1999. An overview of statistical learning theory. *IEEE transactions on neural networks* 10, 5 (1999), 988–999.
- [33] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *NeurIPS*. 5998–6008.
- [34] Hao Wang, Zheng-Jun Zha, Xuejin Chen, Zhiwei Xiong, and Jiebo Luo. 2020. Dual Path Interaction Network for Video Moment Localization. In *ACM MM*. 4116–4124.
- [35] Jingwen Wang, Lin Ma, and Wenhao Jiang. 2020. Temporally Grounding Language Queries in Videos by Contextual Boundary-aware Prediction. In *AAAI*.
- [36] Tan Wang, Jianqiang Huang, Hanwang Zhang, and Qianru Sun. 2020. Visual commonsense r-cnn. In *CVPR*. 10760–10770.
- [37] Wenjie Wang, Fuli Feng, He Xiangnan, Hanwang Zhang, and Tat-Seng Chua. 2021. Clicks can be Cheating: Counterfactual Recommendation for Mitigating Clickbait Issue. In *SIGIR*.
- [38] Yixin Wang, Dawen Liang, Laurent Charlin, and David M Blei. 2020. Causal Inference for Recommender Systems. In *Fourteenth ACM Conference on Recommender Systems*. 426–431.
- [39] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. 2015. Show, attend and tell: Neural image caption generation with visual attention. In *ICML*. 2048–2057.
- [40] Xun Yang, Jianfeng Dong, Yixin Cao, Xun Wang, Meng Wang, and Tat-Seng Chua. 2020. Tree-Augmented Cross-Modal Encoding for Complex-Query Video Retrieval. In *SIGIR*. 1339–1348.
- [41] Xun Yang, Xueliang Liu, Meng Jian, Xinjian Gao, and Meng Wang. 2020. Weakly-Supervised Video Object Grounding by Exploring Spatio-Temporal Contexts. In *ACM MM*. 1939–1947.
- [42] Xun Yang, Meng Wang, and Dacheng Tao. 2018. Person re-identification with metric learning using privileged information. *IEEE Transactions on Image Processing* 27, 2 (2018), 791–805.
- [43] Ghim-Eng Yap, Ah-Hwee Tan, and Hwee-Hwa Pang. 2007. Discovering and exploiting causal dependencies for robust mobile context-aware recommenders. *IEEE Transactions on Knowledge and Data Engineering* 19, 7 (2007), 977–992.
- [44] Yitian Yuan, Lin Ma, Jingwen Wang, Wei Liu, and Wenwu Zhu. 2019. Semantic Conditioned Dynamic Modulation for Temporal Sentence Grounding in Videos. In *NeurIPS*. 536–546.
- [45] Yitian Yuan, Tao Mei, and Wenwu Zhu. 2019. To find where you talk: Temporal sentence localization in video with attention based location regression. In *AAAI*, Vol. 33. 9159–9166.
- [46] Zhongqi Yue, Hanwang Zhang, Qianru Sun, and Xian-Sheng Hua. 2020. Inter-ventional Few-Shot Learning. In *NeurIPS*.
- [47] Runhao Zeng, Haoming Xu, Wenbing Huang, Peihao Chen, Minghui Tan, and Chuang Gan. 2020. Dense regression network for video grounding. In *CVPR*.
- [48] Da Zhang, Xiyang Dai, Xin Wang, Yuan-Fang Wang, and Larry S Davis. 2019. Man: Moment alignment network for natural language moment retrieval via iterative graph adjustment. In *CVPR*. 1247–1257.
- [49] Dong Zhang, Hanwang Zhang, Jinhui Tang, Xiansheng Hua, and Qianru Sun. 2020. Causal intervention for weakly-supervised semantic segmentation. In *NeurIPS*.
- [50] Hao Zhang, Aixun Sun, Wei Jing, and Joey Tianyi Zhou. 2020. Span-based Localizing Network for Natural Language Video Localization. In *ACL*. Online, 6543–6554.
- [51] Songyang Zhang, Houwen Peng, Jianlong Fu, and Jiebo Luo. 2020. Learning 2D Temporal Adjacent Networks for Moment Localization with Natural Language. In *AAAI*.
- [52] Zhu Zhang, Zhijie Lin, Zhou Zhao, and Zhenxin Xiao. 2019. Cross-modal interaction networks for query-based moment retrieval in videos. In *SIGIR*. 655–664.