

G2L: Semantically Aligned and Uniform Video Grounding via Geodesic and Game Theory

Hongxiang Li¹, Meng Cao^{2,1}, Xuxin Cheng¹, Yaowei Li¹, Zhihong Zhu¹, Yuexian Zou^{1†}

¹School of Electronic and Computer Engineering, Peking University

²International Digital Economy Academy (IDEA)

{lihongxiang, chengxx, zhihongzhu, ywl}@stu.pku.edu.cn; {mengcao, zouyx}@pku.edu.cn

Abstract

The recent video grounding works attempt to introduce vanilla contrastive learning into video grounding. However, we claim that this naive solution is suboptimal. Contrastive learning requires two key properties: (1) alignment of features of similar samples, and (2) uniformity of the induced distribution of the normalized features on the hypersphere. Due to two annoying issues in video grounding: (1) the co-existence of some visual entities in both ground truth and other moments, i.e. semantic overlapping; (2) only a few moments in the video are annotated, i.e. sparse annotation dilemma, vanilla contrastive learning is unable to model the correlations between temporally distant moments and learned inconsistent video representations. Both characteristics lead to vanilla contrastive learning being unsuitable for video grounding. In this paper, we introduce Geodesic and Game Localization (G2L), a semantically aligned and uniform video grounding framework via geodesic and game theory. We quantify the correlations among moments leveraging the geodesic distance that guides the model to learn the correct cross-modal representations. Furthermore, from the novel perspective of game theory, we propose semantic Shapley interaction based on geodesic distance sampling to learn fine-grained semantic alignment in similar moments. Experiments on three benchmarks demonstrate the effectiveness of our method.

1. Introduction

Video grounding [18, 59, 36, 45, 17, 75, 8, 69, 72, 67, 6, 25] aims to identify the timestamps semantically corresponding to a given query within an untrimmed video, which is a challenging multimedia retrieval task due to flexibility and complexity of the query and video content. The video grounding model needs well to model the complex cross-modal correlations and semantic information.

[†] Corresponding author.

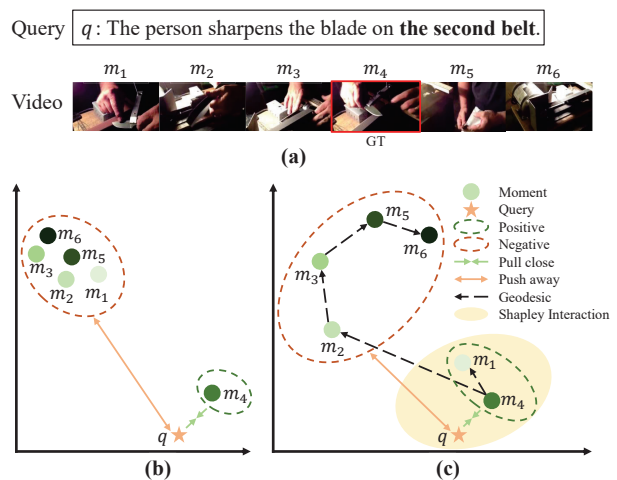


Figure 1. (a) Illustration of video grounding. ‘GT’ indicates the ground truth. A comparison of (b) existing contrastive learning-based methods and (c) our proposed G2L method. G2L makes semantically similar video moments closer in representation space while exploring nuances among similar moments.

Contrastive learning [10, 23, 43] is proposed to learn representations by contrasting positive pairs against negative pairs. With the popularity of contrastive learning in vision-language tasks [34, 13, 46, 7, 5, 64, 62], several works also apply it to video grounding. Nan *et al.* [45] propose a dual contrast learning to learn more informative feature representations by maximizing the mutual information between the query and the corresponding video clips. This naive solution, however, achieves sub-optimum performance.

Generally, contrastive learning requires two key properties [58]: *alignment* and *uniformity*. *Alignment* favors encoders that assign similar features to similar samples. *Uniformity* prefers a feature distribution that preserves maximal information, i.e., the uniform distribution on the unit hypersphere. We argue that these two issues are not satisfied in current video grounding works with contrastive learning.

Firstly, the semantic overlapping issue is widespread,

i.e., the co-existence of some visual entities in both ground truth and other moments. As shown in Figure 1(a), the entity of ‘person’, ‘blade’ and ‘belt’ appear in both ground-truth moment m_4 and others. Since there exist no classification labels in video grounding, previous methods distinguish positive and negative samples only based on the annotated moments. This strict scheme, however, ignores the semantic overlapping among video moments, which leads to the contradiction in feature representations. As shown in Figure 1(a)(b), the moment m_1 of “sharpen the blade on the first belt” only differs from target moment m_4 in the ‘second’ order. They share similar semantic meanings but are forced to be pushed away in feature space. It is not consistent with the *alignment* principle of the ideal contrastive learning.

Another issue lies in the sparse annotation dilemma [67, 31]. Due to the costly labeling process, only a few moments are annotated regardless of the thousands of frames contained. Such severe data imbalance leads to significant learning bias for vanilla contrastive learning, *i.e.*, unannotated moments are pushed away by different queries, regardless of the semantic relationships. Consequently, their representations are close, although they don’t necessarily have strong semantic similarities. This undermines the *uniformity* requirement of contrastive learning. For example in Figure 1(b), wrong results always exist when encountering these unannotated moments (*e.g.*, m_1).

To address the issues mentioned above, we propose a novel Geodesic and Game Localization (G2L), a semantically aligned and uniform video grounding framework as shown in Figure 1(c). We propose to measure the similarity according to the geodesic distance [29] between two video moments along the manifold. In Figure 1(c), the geodesic distance between m_4 to m_6 is the length of the shortest path as the moment graph, *i.e.* $m_4 \rightarrow m_2 \rightarrow m_3 \rightarrow m_5 \rightarrow m_6$. In contrast to previous methods, we construct positive and negative pairs based on the geodesic distance rather than the temporal moment to relax the strict positional principle. The geodesics from the target moment to other moments are used to guide the maximizing mutual information. In this manner, the distance between video moments correctly reflects semantic relevance.

Unfortunately, the relaxed contrastive objective with geodesic leads to one side-effect, *i.e.*, the model may confuse similar video moments. As shown in Figure 1(c), the model may falsely map m_1 to be as close to the query q since m_1 shares a similar appearance with the ground truth m_4 . To further prevent the model from confusing similar video moments, we formulate video moments and queries as multiple players into a cooperative game and quantify their game-theoretic interactions (*i.e.*, Shapley interactions [49, 20]). Through this, we evaluate the marginal contributions of each fine-grained component, which leads

to a more accurate division. However, computing the exact Shapley interaction for all players is an NP-hard problem [41] and is difficult to achieve solutions in the video grounding setting. Therefore, we further propose a semantic Shapley interaction module, which samples similar intra-video moments by geodesic distance with a focus on their nuances.

In sum, our contributions are summarized as follows:

- We present G2L which introduces geodesic and game theory to learn the semantic alignment and uniformity between video and query for video grounding.
- We propose a novel geodesic-guided contrastive learning scheme that considers the correct semantics of all moments in the video.
- We introduce an effective semantic Shapley interaction strategy based on geodesic distance.
- Extensive experiments on three public datasets demonstrate the effectiveness of our G2L.

2. Related Work

Video Grounding. Video grounding proposed by [18, 1], which aims to predict the start and end boundaries of the activity described by a given language query within a video. Early approaches focus on carefully designed complex video-text interaction modules. Yuan *et al.* [66] propose an approach which can directly predict the coordinates of the queried video clip using attention mechanism. Zeng *et al.* [67] propose a pyramid neural network to consider multi-scale information. Liu *et al.* [39] advise a memory attention to emphasize the visual features and simultaneously utilize the context information. Xu *et al.* [61] introduce a multi-level model to integrate visual and textual features earlier and further re-generate queries as an auxiliary task. To improve the representation of the model, several methods introduce contrastive learning or cross-modal discrimination. Ding *et al.* [17] propose to combine discriminative contrastive objective and generative caption objective to optimize dual-encoder. Nan *et al.* [45] introduce introducing causal intervention and dual contrastive learning to improve representation. In this paper, we propose to model semantic alignment and uniformity by approximate geodesics and game-theoretic interactions.

Contrastive Learning. Contrastive learning (CL) often serves as an unsupervised objective to learn representations by contrasting positive pairs against negative pairs [10, 23, 43, 35, 12, 11]. Some prior works consider maximizing the mutual information (MI) between latent representations [24]. MI quantifies the “amount of information” gained about one random variable by observing another random variable [2]. Contrastive learning has been applied to vision-language tasks to learn the joint representations of visual and textual modalities [42, 51]. Sun *et al.* [71] pro-

pose a retrieval and localization network with contrastive learning for video corpus moment retrieval.

Shapley Value. The Shapley value [49, 26] originates from cooperative game theory. It has been theoretically proven to be the unique metric to fairly estimate the contribution of each player in a cooperative game such that satisfying certain desirable axioms [60] and is widely used in deep learning. Li *et al.* [32] propose a semantically aligned vision-language pre-training based on Shapley value to model fine-grained semantics. Ren *et al.* [48] propose to explain adversarial attacks by Shapley value. Li *et al.* [33] propose to explicit credit assignment for multi-agent reinforcement learning using Shapley value.

3. Geodesic and Game Localization (G2L)

3.1. Problem Formulation and Model Overview

Let Q_i and V_i be a given textual query and untrimmed video respectively. The purpose of video grounding is to locate the most relevant video interval $A_i = (t_i^s, t_i^e)$, where t_i^s and t_i^e are starting and ending times respectively. The key to video grounding is to learn semantics between video and query. To this end, previous methods [45, 17] incorporate vanilla contrastive learning into existing cross-modal interaction architectures. Typically, at the training stage, they employ two loss functions: video grounding loss \mathcal{L}_{VG} and vanilla contrastive loss \mathcal{L}_{VCL} . \mathcal{L}_{VG} computes the cross entry between the target timestamp and the prediction timestamp to optimize the model. \mathcal{L}_{VCL} discriminates between positive samples and negative samples based on the temporal moment and adopts noise-contrastive estimation (NCE) [21] to obtain MI of videos and queries.

However, vanilla contrastive learning is not appropriate for video grounding due to the semantic overlapping and sparse annotation dilemma. To learn correct semantics and improve representations, as illustrated in Figure 2, we propose Geodesic and Game Localization (G2L), a semantically aligned and uniform video grounding framework that germinates from geodesic and cooperative game theory. Our G2L learns semantic alignment and uniformity from two components. With geodesic-guided contrastive learning (GCL), complete semantic alignment and uniformly distributed video features are learned. By semantic Shapley interaction (SSI), we learn the fine-grained semantic alignment between similar moments and target queries. Combined with the two novel proposed training objectives, the full training objective of semantically aligned and uniform video grounding can be formulated as:

$$\mathcal{L} = \mathcal{L}_{VG} + \mathcal{L}_{GCL} + \mathcal{L}_{SSI} \quad (1)$$

During inference, it can be directly removed, rendering a semantics-sensitive dual encoder.

3.2. Feature Encoder

Video Encoder. For an input video V_i , we first segment it into small video clips and perform a fixed-interval sampling over these clips. Then the video clips are fed into a pre-trained 3D CNN model (*e.g.* C3D) to extract the video features F_i^V . Referring to the previous works [75], we employ sparse sampling and a proposal network to construct the feature map $F_i^M = \{m_i\}_{i=1}^{N_m}$ of moment candidates based on F_i^V , where N_m is the number of proposals in the video.

Query Encoder. For a textual query Q_i , we generate the tokens of words by the tokenizer and add a class embedding token ‘[CLS]’ in the beginning. Then the tokens are fed into pre-trained BERT [27], and we perform average pooling of its last two hidden states to obtain the sentence feature q_i .

3.3. Geodesic Distance Computation

Video content is intricate and flexible, and for video clips, temporal proximity does not always equate to semantic similarity. Conversely, two video clips can have strong correlations even though they are temporally distant. We propose to measure correlations of video representations using the geodesic distance. Formally, we define a mini-batch of video-query pairs as $\{V_i, Q_i\}_{i=1}^B$, where B is the size of the mini-batch. After fed into the video and query encoders, we obtain textual representation $L = \{q_i\}_{i=1}^B$ and moment representations $M = \{m_i\}_{i=1}^{B \times N_m}$.

The video moments representations may lie in a high-dimensional manifold, and our purpose is to measure the geodesic distance between two points along the manifold. However, computing the exact geodesic distance [29] is difficult without explicit knowledge of the manifold structure. To obtain the geodesic distance, we first employ the K-NN graph [15] to approximate the manifold structure [52, 14]. In this graph, each moment m_i forms a node, and each node connects to at most n other nodes. A directed edge exists from node m_i to node m_j if node m_j is among the n nearest neighbors of m_i . The edge weight $d(m_i, m_j)$ is defined using cosine similarity: $d(m_i, m_j) = 1 - m_i m_j^T$. Finally, we utilize the shortest path algorithm, *i.e.*, Dijkstra [16], to compute the length of the shortest path between two moments along the obtained weighted directed graph as the geodesic distance $\mathcal{G}(m_i, m_j)$.

3.4. Geodesic-guided Contrastive Learning

Previous Contrastive Scheme. In video grounding, a query usually corresponds to multiple clips. An intuitive method [17] to learn representations is to set the clips in ground truth intervals as positive samples, while others are negatives. Another scheme [59] is to calculate the intersection over union (IoU) between the ground truth and other moments, where those with higher IoU are considered as positive samples and the lower ones are regarded as negative samples. Then they learn the joint representation by

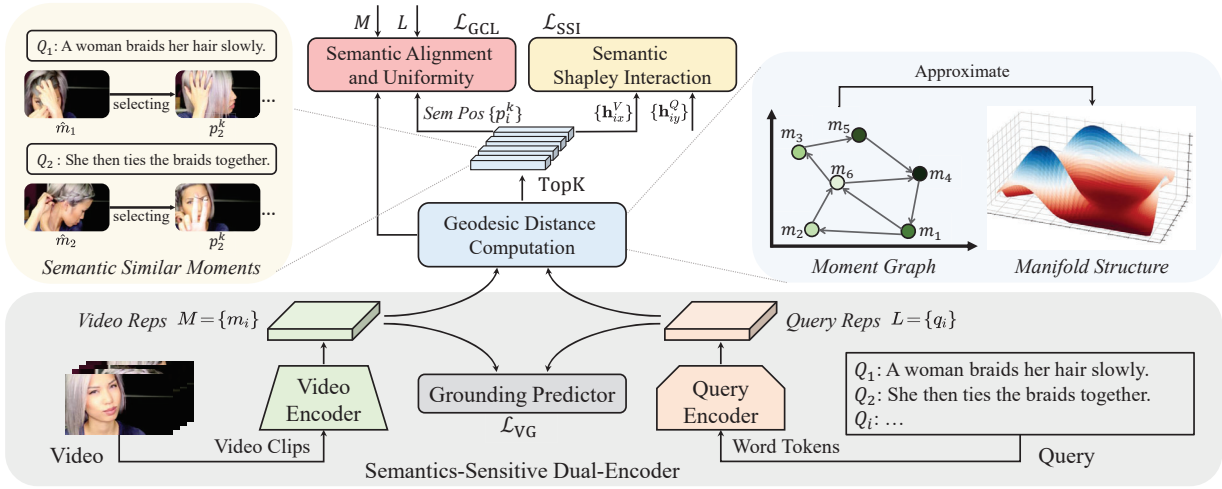


Figure 2. Overview of Geodesic and Game Localization (G2L). Our framework encourages the model to learn semantically aligned and uniform joint representations. In the inference stage, we directly fuse the video features and query features to compute the predicted moments. In the training stage, the grounding loss \mathcal{L}_{VG} is obtained by calculating the cross-entropy between the predicted moment and the target moment. Then, we approximate the high-dimensional manifold structure of the video representations through a moment graph and calculate the geodesic distance from the target moment to other moments. Finally, we leverage geodesic distance for cross-modal discrimination and semantic Sharpley interaction modeling.

pulling the query features and ground truth moment features together, and pushing the query features and non-ground truth moment features apart. The previous contrastive loss \mathcal{L}_{VCL} can be formulated as:

$$\mathcal{L}_{VCL} = - \sum_{i=1}^B \left(\log \frac{\exp(q_i m_i^\top / \tau)}{\sum_{m_j \in M} \exp(q_i m_j^\top / \tau)} \right) \quad (2)$$

where τ is the temperature hyper-parameter.

While intuitive, such a manner fails to capture semantic *alignment* and *uniformity* between queries and video content. To learn complete semantic *alignment* and *uniformity* for video grounding, we propose geodesic-guided contrastive learning.

Semantic Alignment. In vanilla supervised contrastive learning [56, 28], all classes play the same role: pulling is done within every class and pushing between every pair of different classes. Such uniform contrastive learning works well for symmetric and equal cross-modal learning. In fact, due to semantic overlapping, the semantics of the video and text are asymmetric and unequal in video grounding. Temporally distant moments may have similar semantics. To model semantic *alignment*, we discriminate positive pairs based on the geodesic distance instead of the temporal moment. For a query feature q_i , we define the k moments with the closest geodesic distance from the target moment $\hat{m}_i \in M$ as its semantic positive samples P_i :

$$P_i = \{p_i^k\} = \arg \text{topk} \mathcal{G}(\hat{m}_i, m_k) \quad (3)$$

P_i contains m_i and moments adjacent to m_i geodesic which is used to construct positive pairs with q_i to relax the previous strict positional principle.

Semantic Uniformity. Due to the sparse annotation dilemma [67, 31], most of the moments are only marked as negative while pushed away by different queries. Notably, this push operation is undifferentiated, *i.e.*, the model is encouraged to learn to push all unannotated moments from the annotated ones. This clearly undermines the model to learn correct semantics. To learn semantic *uniformity*, we introduce the geodesic distance to differentially push negative samples away from the query based on semantic relationships. The similarity between query and moment $s(q_i, m_j)$ is defined as:

$$s(q_i, m_j) = \exp \left(q_i m_j^\top \left(\hat{m}_i m_j^\top \log \frac{1}{\exp(\mathcal{G}(\hat{m}_i, m_j) + 1)} \right) \right) \quad (4)$$

We assign corresponding weights according to the geodesic distance $\mathcal{G}(\hat{m}_i, m_j)$ between the target moment \hat{m}_i and all moments m_j . $s(q_i, m_j)$ considers the relationships between negative samples while maximizing mutual information.

Finally, our geodesic-guided contrastive loss can be formulated as:

$$\mathcal{L}_{GCL} = - \sum_{i=1}^B \left(\log \frac{\sum_{p_i^k \in P_i} \exp(q_i p_i^k / \tau)}{\sum_{m_j \in M} s(q_i, m_j) / \tau} \right) \quad (5)$$

In contrast to Equation 2, \mathcal{L}_{GCL} is designed for representa-

tion learning rather than directly learning localization.

3.5. Semantic Shapley Interaction Modeling

To prevent the model from confusing similar video moments due to the relaxed contrastive objective in Equation 3, we propose semantic Shapley interaction to model fine-grained semantic alignment.

3.5.1 Preliminaries

Shapley Value. The Shapley value [49] is a classical game theory solution for the unbiased estimation of the contribution of each player in a cooperation game. Assume a game consists of \mathcal{N} players, $\mathcal{U} \subseteq \mathcal{N}$ represents a potential subset of players. A game function $f(\cdot)$ maps each subset \mathcal{U} of players to a score, estimating the cooperated contribution of a set of players. For a player i , its Shapley value $\phi(i|\mathcal{N})$ is computed as the average marginal contribution of player i to all possible coalitions \mathcal{U} without i :

$$\phi(i|\mathcal{N}) = \sum_{\mathcal{U} \subseteq \mathcal{N} \setminus \{i\}} p(\mathcal{U}) [f(\mathcal{U} \cup \{i\}) - f(\mathcal{U})] \quad (6)$$

$$p(\mathcal{U}) = \frac{|\mathcal{U}|!(|\mathcal{N}| - |\mathcal{U}| - 1)!}{|\mathcal{N}|!} \quad (7)$$

where $p(\mathcal{U})$ is the likelihood of \mathcal{U} being sampled. The Shapley value has been proved to be the unique metric that satisfies the following axioms: *Linearity*, *Symmetry*, *Dummy*, and *Efficiency*. We summarize these axioms in the supplementary material.

Shapley Interaction. In game theory, some players tend to form a small cooperative coalition and always participate in the game together. This cooperation provides additional contributions to the game. The Shapley interaction [20] measures the additional contributions made by the coalition compared to players working individually. We define $[\mathcal{U}]$ as a single hypothetical player, which is the union of the players in \mathcal{U} . A reduced game is formed by removing the individual players in \mathcal{U} from the game and adding $[\mathcal{U}]$ to the game. Finally, according to Equation 6, the Shapley interaction for coalition \mathcal{U} is formulated as:

$$\mathfrak{J}([\mathcal{U}]) = \phi([\mathcal{U}] | \mathcal{N} \setminus \mathcal{U} \cup \{[\mathcal{U}]\}) - \sum_{i \in \mathcal{U}} \phi(i | \mathcal{N} \setminus \mathcal{U} \cup \{i\}) \quad (8)$$

where $\mathfrak{J}([\mathcal{U}])$ reflects the interactions inside \mathcal{U} . The higher value of $\mathfrak{J}([\mathcal{U}])$ indicates that players in \mathcal{U} cooperate closely with each other.

3.5.2 Fine-Grained Semantic Alignment via Shapley

According to Equation 6 and Equation 8, it can be known that the computational complexity of the Shapley interaction grows factorially as the number of players increases. To reduce the computational cost, we propose semantic

Shapley interaction based on geodesic distance sampling. Specifically, as in Equation 3, we sample semantic similar moments based on geodesic distances to investigate their nuances. Then, we consider queries and semantic similar moments from the same video as players in the same cooperative game.

To avoid confusion with the previous symbols, we define $\mathcal{H}_i^V = \{\mathbf{h}_{ix}^V\}_{x=1}^{N_i^v}$ and $\mathcal{H}_i^Q = \{\mathbf{h}_{iy}^Q\}_{y=1}^{N_i^q}$ are a set of semantic positive samples and queries from the video V_i respectively, where $N_i^v = KN_i^q$ and N_i^q indicate the number of similar moments and queries in video V_i respectively. K is the number of moments sampled for each query \mathbf{h}_{iy}^Q . We investigate the effect of K on computational cost and performance in the supplementary material.

If a query and a video moment have strong semantic correspondence, then they tend to cooperate with each other and contribute to the coalition. Thus, we can take $\mathcal{H}^i = \mathcal{H}_i^V \cup \mathcal{H}_i^Q$ as the players in the same game inspired by [32]. We define the alignment matrix as: $\mathcal{A}_i = [a_{xy}^i]$, where $a_{xy}^i = h_{ix}^V \top h_{iy}^Q$ represents the alignment score between x -th moment and y -th query in video V_i . Next, $\tilde{\mathcal{A}}_i$ is obtained by applying softmax-normalization over each row of \mathcal{A}_i . Then we average the maximum alignment score $\max_y \tilde{a}_{xy}^i$ as the fine-grained moment-to-query similarity ψ_1 . Similarly, we can obtain the fine-grained query-to-moment similarity ψ_2 . The total fine-grained similarity score can be defined: $\psi = (\psi_1 + \psi_2)/2$, which be considered as the game score $f(\cdot)$ in our game.

According to Equation 8, the semantic Shapley interaction of them can be formulated as:

$$\begin{aligned} \mathfrak{J}([\mathcal{H}_{xy}^i]) &= \phi([\mathcal{H}_{xy}^i] | \mathcal{H}^i \setminus \mathcal{H}_{xy}^i \cup \{[\mathcal{H}_{xy}^i]\}) \\ &\quad - \phi(\mathbf{h}_{ix}^V | \mathcal{H}^i \setminus \mathcal{H}_{xy}^i \cup \{\mathbf{h}_{ix}^V\}) \\ &\quad - \phi(\mathbf{h}_{iy}^Q | \mathcal{H}^i \setminus \mathcal{H}_{xy}^i \cup \{\mathbf{h}_{iy}^Q\}) \quad (9) \\ &= \mathbb{E}_{\mathcal{C}} \left\{ \sum_{\substack{\mathcal{U} \subseteq \mathcal{H}^i \setminus \mathcal{H}_{xy}^i \\ |\mathcal{U}| = \mathcal{C}}} [f(\mathcal{U} \cup \mathcal{H}_{xy}^i) - f(\mathcal{U} \cup \{\mathbf{h}_{ix}^V\}) \right. \\ &\quad \left. - f(\mathcal{U} \cup \{\mathbf{h}_{iy}^Q\}) + f(\mathcal{U})] \right\} \quad (10) \end{aligned}$$

where $\mathfrak{J}([\mathcal{H}_{xy}^i])$ represents the single player formed by the coalition of x -th moment and y -th query in video V_i , \mathcal{C} represents the coalition size. We take normalized $\mathfrak{J}'([\mathcal{H}_{xy}^i])$ as soft labels, the fine-grained semantic alignment loss can be defined as:

$$\mathcal{L}_{\text{SSI}} = - \sum_{i=1}^T \frac{1}{N_i^v N_i^q} \sum_{x=1}^{N_i^v} \sum_{y=1}^{N_i^q} \mathfrak{J}'([\mathcal{H}_{xy}^i]) \log(\tilde{a}_{xy}^i) \quad (11)$$

where T is the total number of unique videos in a mini-batch, *i.e.* $T \leq B$.

Model	R@1	R@1	R@5	R@5
	IoU0.5	IoU0.7	IoU0.5	IoU0.7
CTRL [18]	29.01	10.34	59.17	37.54
ACRN [39]	31.67	11.25	60.34	38.57
TripNet [22]	32.19	13.93	-	-
SCDM [65]	36.75	19.86	64.99	41.53
LGI [44]	41.51	23.07	-	-
VSLNet [73]	43.22	26.16	-	-
2D-TAN [75]	44.51	26.54	77.13	61.96
DRN [67]	45.45	24.39	77.97	50.30
DPIN [54]	47.27	28.31	77.45	60.03
CBLN [37]	48.12	27.60	79.32	63.41
SMIN [55]	48.46	30.34	<u>81.16</u>	62.11
MATN [74]	48.02	<u>31.78</u>	78.02	63.18
GTR [4]	<u>50.57</u>	29.11	80.43	<u>65.14</u>
CMAS [63]	46.23	29.48	77.04	60.25
MMN [59]	48.59	29.26	79.50	64.76
CL-based:				
IVG-DCL [45]	43.84	27.10	-	-
SSCS [17]	46.67	27.56	78.37	63.78
G2L (ours)	51.68	33.35	81.32	67.60

Table 1. Performance comparisons on ActivityNet-Captions using C3D features.

4. Experiments

4.1. Datasets and Evaluation

ActivityNet-Captions. ActivityNet-Captions [30] contains 20,000 untrimmed videos and 100,000 descriptions [3], covering a wide range of complex human behavior. The video clips with annotations have much larger variations. Following the public split [75], we use 37417, 17505 and 17031 sentence-video pairs for training, validation and testing, respectively.

Charades-STA. The Charades dataset [50] is collected for video action recognition and video captioning. Gao *et al.* [18] adapt the Charades dataset to the video grounding task by collecting the query annotations. The Charades-STA dataset contains 6672 videos and involves 16128 video-query pairs, where 12408 pairs are used for training and 3720 for testing. We follow the same split of the dataset as in Gao *et al.* [18] for fair comparisons.

TACoS. TACoS [47] contains 127 videos from the cooking scenarios. We follow the standard split [18], which has 10146, 4589 and 4083 video query pairs for training, validation and testing, respectively.

Evaluation. Following previous work [18, 75], we adopt “R@n, IoU=m” as the evaluation metric. It calculates the percentage of IoU greater than “m” between at least one of the top “n” video moments retrieved and the ground truth.

Model	R@1	R@1	R@5	R@5
	IoU0.5	IoU0.7	IoU0.5	IoU0.7
MCN [1]	17.46	8.01	48.22	26.73
SAP [9]	27.42	13.36	66.37	38.15
TripNet [22]	36.61	14.50 -	-	-
MAN [70]	41.21	20.54	83.21	51.85
2D-TAN [75]	39.70	23.31	80.32	51.26
DRN [67]	42.90	23.68	87.80	54.87
FVMR [19]	42.36	24.14	83.97	50.15
CBLN [37]	43.67	24.44	88.39	56.49
MMRG [68]	44.25	-	60.22	-
MMN [59]	<u>47.31</u>	<u>27.28</u>	<u>83.74</u>	<u>58.41</u>
CL-based:				
SSCS [17]	43.15	25.54	84.26	54.17
G2L(Ours)	47.91	28.42	84.80	59.33

Table 2. Performance comparisons on Charades-STA using VGG features.

4.2. Implementation Details

For a fair comparison, we extracted video features from a pre-trained 3D CNN (C3D for ActivityNet-Captions and TACoS, VGG for Charades-STA) following previous works [59, 17, 45]. We uniformly sampled (256, 32, 256) clips as the input video sequence and set the length of 2D feature map [75] (64, 16, 128) for ActivityNet-Captions, Charades-STA, and TACoS, respectively. For the language query, the pre-trained BERT[27] was employed for each word of the query. The average pooling outputs of the first and last layers were used to obtain the embedding of the whole sentence. We employed the AdamW optimizer [40] to train our model and set the temperature weight τ to 0.1. The learning rates were set to (8×10^{-4} , 8×10^{-4} , 1×10^{-4}) for ActivityNet-Captions, Charades-STA, and TACoS, respectively. We conducted experiments on 8 A100 GPUs with batch size 48 for ActivityNet-Captions and Charades-STA, and on 4 A100 GPUs with batch size 8 for TACoS.

4.3. Comparisons with State-of-the-art Methods

Comparison on ActivityNet-Captions. In Table 1, we report our performance in comparison with other state-of-the-art methods on ActivityNet-Captions. Compared with IVG-DCL [45] and SSCS [17], which are also based on contrastive learning, our method achieves significant improvements. Notably, ActivityNet-Captions is currently the dataset with the most severe semantic overlapping and sparse annotation dilemma. We observe that our model showed even more significant improvements on ActivityNet-Captions, achieving absolute improvements of up to 7.8% and 5.7% compared to IVG-DCL [45] and SSCS [17], respectively. They encourage the model to focus on cross-modal alignment in favor of grounding while ig-

Model	R@1 IoU0.3	R@1 IoU0.5	R@5 IoU0.3	R@5 IoU0.5
CTRL [18]	18.32	13.30	36.69	25.42
CBP [57]	27.31	24.79	43.64	37.40
SCDM [65]	26.11	21.17	40.16	32.18
2D-TAN [75]	37.29	25.32	57.81	45.04
DRN [67]	-	23.17	-	33.36
CMIN [76]	24.64	18.05	-	-
CSMGAN [38]	33.90	27.09	53.98	41.22
CBLN [37]	38.98	27.65	59.96	46.24
MMN [59]	39.24	26.17	62.03	47.39
GTR [4]	40.39	<u>30.22</u>	61.94	47.73
FVMR [19]	<u>41.48</u>	29.12	<u>64.53</u>	50.00
CL-based:				
IVG-DCL [45]	38.84	29.07	-	-
SSCS [17]	41.33	29.56	60.65	48.01
G2L(Ours)	42.74	30.95	65.83	49.86

Table 3. Performance comparisons on TACoS using C3D features.

noring the semantics among full moments, especially unannotated moments. In contrast, our method conducts cross-modal discrimination guided by the geodesic distance such that overcoming data bias.

Comparison on Charades-STA. Table 2 reports the result comparison between state-of-the-art methods on Charades-STA. Compared with SSCS [17], our method achieves a performance improvement of up to 5.1%. On more stringent evaluation metrics, such as “R@1 IoU=0.7”, our method achieves a performance improvement of 1.1% compared to the cutting-edge method MMN [59], which indicates that exploring fine-grained semantic alignment information between similar video moments can improve grounding quality. Notably, MMN [59] and SSCS [17] have similar loss functions, and MMN [59] from a perspective on temporal grounding as a metric-learning problem proposes a mutual matching network to enhance joint representation learning by mining more negative samples achieving better performances. However, it also constructs cross-modal pairs based on the temporal moment, which results in a large number of constructed negative pairs containing potential weak semantic positive pairs, thus hindering learning.

Comparison on TACoS. Table 3 summarizes the comparisons on the TACoS. We observe that our model achieves state-of-the-art results in most settings. However, the performance gain on this dataset is smaller than the previous three datasets. The reason is that the sparse annotation dilemma and semantic overlapping are insignificant on TACoS. It contains only 127 videos but has about 20,000 queries and focuses on cooking activities with more uniform objects, roles, and actions. Nevertheless, our method still outperforms previous contrastive learning-based methods in various metrics.

Model	R@1 IoU0.5	R@1 IoU0.7	R@5 IoU0.5	R@5 IoU0.7
Full Model	51.68	33.35	81.32	67.60
w/o SA	50.85	30.73	79.22	65.47
w/o SU	49.10	31.14	79.48	65.46
w/o GCL	48.74	28.62	78.89	64.44
w/o SSI	50.65	30.01	79.54	66.76
w/o GCL+SSI	45.74	26.74	77.26	61.64

Table 4. Ablation studies of main components on ActivityNet-Captions. “SA” and “SU” denote the Semantic Alignment and the Semantic Uniformity in GCL, respectively.

Metric	R@1 IoU0.5	R@1 IoU0.7	R@5 IoU0.5	R@5 IoU0.7
Euclidean	47.83	28.44	79.42	65.46
Timestamp	49.85	29.73	78.15	65.47
Cosine	50.94	30.09	79.57	65.84
Geodesic	51.68	33.35	81.32	67.60

Table 5. Comparison of different distance metrics on ActivityNet-Captions.

5. Ablation Study

Effectiveness of Individual Components. In Table 4, we conduct a thorough ablation study on the proposed components to verify their effectiveness. As shown in Table 4, removing the entire GCL will result in up to 4.5% of performance degradation, demonstrating the contribution of GCL by learning semantic alignment and unification. We also observe that removing either SA or SU results in about 2 points of performance drop on average, indicating that our method enables the vanilla contrastive learning to be more suitable for video grounding setting. Removing SSI leads to a 3.3% drop in performance on the more stringent metric (*i.e.*, “R@1 IoU=0.7”), which highlights the importance of fine-grained semantic alignment for high-quality moment retrieval. The last row shows our baseline, and our method achieves up to 6 points of performance improvement without modifying the architecture compared to the baseline, confirming the superiority of our model.

Effectiveness of Geodesic Distance. To demonstrate the effectiveness of the geodesic, we substitute it with different distance metrics in the model, including euclidean distance, timestamp distance, and cosine distance, and compare their performances as shown in Table 5. It can be observed that when using euclidean distance, the performances drop by up to 5%. This is due to the fact that video representations often lie in a high dimensional manifold and euclidean distance cannot accurately quantify the correlations among video moments. When using timestamp distance, the performances decrease by 2% on average, indicating that the temporal adjacency does not necessarily

Model	Inference Time	Inference Time
2D-TAN [75]	0.13s	32s
MMN [59]	0.32s	37s
Base Model	0.39s	40s
G2L(Ours)	0.84s	43s

Table 6. Time consumption on ActivityNet-Captions.




Figure 3. Projected video moment features (a): learned representations of the previous method with vanilla contrastive learning; (b): learned representations of our method.

correspond to semantic similarity due to the flexibility and complexity of video content. Using cosine distance results in performances dropping by up to 3%. Due to the insufficient representation capacity of the model in the early training stage to accurately calculate similarity. We approximate the manifold structure of video features using the moment graph, where the reachability and the shortest path in the graph can facilitate similarity measurement.

The efficiency of our method. In Table 6, we compute the average training time per iteration and total inference time. Due to the K-NN graph and Shapley interaction, G2L requires more training costs. During inference, GCL and SSI can be removed thus our G2L only needs additional 3s compared to the base model.

6. Qualitative Analysis

We visualize a sample video from the ActivityNet-Captions into the 2D images using t-SNE [53]. In Figure 3(a), we observe that there is a clear “island” in the learned representations of the previous method. We argue that it is formed by a few moments which are temporally adjacent to the ground truth. The reason is that the previous method divides positive and negative samples based on the strict temporal moment while pushing negative samples away by different queries, regardless of the semantic relationships. In contrast, our method relaxes the previous strict contrastive objective using geodesic distance, mitigating semantic overlapping and sparse annotation dilemma. Furthermore, the semantic Shapley interaction enables our model to capture discriminative features between similar moments. Our method learns aligned and uniform representations thus eliminating the “island” in Figure 4(b),

We present two qualitative results from ActivityNet-

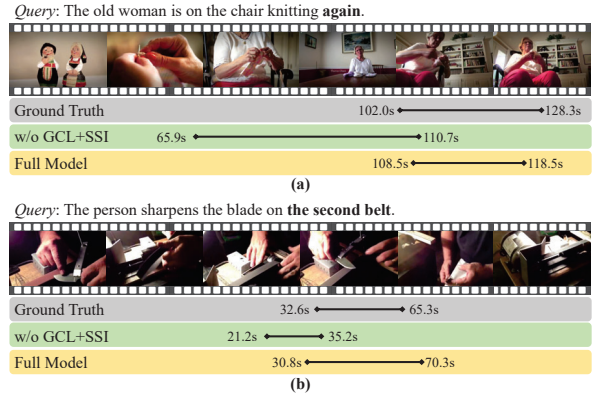


Figure 4. Qualitative results of our method on the ActivityNet-Captions.

Captions in Figure 4. We observe that our approach can successfully model the relationships of temporally distant moments. For instance, “knitting again” in Figure 4(a) and “the second belt” in Figure 4(b) are far away from their first appearances. Our method can capture the information of “first knitting” and “the first belt” and infer precisely the location of the target moment. Previous contrastive learning-based methods fail to capture this semantic association due to tending to perform cross-modal contrast by the strict positional principle. These models only focus on the contextual content related to the target moment. Our method demonstrates that these large volumes of unannotated moments contain rich information and modeling their relationships can improve representation learning.

7. Conclusion and Future Work

This paper introduces a novel semantically aligned and uniform video grounding method, Geodesic and Game Localization (G2L), which explores more semantic information by measuring the correlation between video moments through the geodesic and models the nuances of similar moments using game-theoretic interactions. By contrasting the video features and query features in the shared space, the learned bi-modal features become similar when their semantics match, and the similarity between video moments is modeled by the geodesic-guided pushing operation. Extensive experiments demonstrate that our novel training objective significantly improves the performance of existing contrastive learning-based methods in video grounding. This paper provides a novel perspective for cross-modal contrastive learning. In the future, we intend to apply this idea to multi-modal pre-training.

Acknowledgment. This paper was partially supported by NSFC (No: 62176008) and Shenzhen Science & Technology Research Program (No: GXWD20201231165807007-20200814115301001).

References

- [1] Lisa Anne Hendricks, Oliver Wang, Eli Shechtman, Josef Sivic, Trevor Darrell, and Bryan Russell. Localizing moments in video with natural language. In *Proceedings of the IEEE international conference on computer vision*, pages 5803–5812, 2017. 2, 6
- [2] Anthony J Bell and Terrence J Sejnowski. An information-maximization approach to blind separation and blind deconvolution. *Neural computation*, 7(6):1129–1159, 1995. 2
- [3] Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 961–970, 2015. 6
- [4] Meng Cao, Long Chen, Mike Zheng Shou, Can Zhang, and Yuexian Zou. On pursuit of designing multi-modal transformer for video grounding. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 9810–9823, 2021. 6, 7
- [5] Meng Cao, Ji Jiang, Long Chen, and Yuexian Zou. Correspondence matters for video referring expression comprehension. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 4967–4976, 2022. 1
- [6] Meng Cao, Fangyun Wei, Can Xu, Xiubo Geng, Long Chen, Can Zhang, Yuexian Zou, Tao Shen, and Daxin Jiang. Iterative proposal refinement for weakly-supervised video grounding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6524–6534, 2023. 1
- [7] Meng Cao, Tianyu Yang, Junwu Weng, Can Zhang, Jue Wang, and Yuexian Zou. Locvtp: Video-text pre-training for temporal localization. In *European Conference on Computer Vision*, pages 38–56. Springer, 2022. 1
- [8] Meng Cao, Can Zhang, Long Chen, Mike Zheng Shou, and Yuexian Zou. Deep motion prior for weakly-supervised temporal action localization. *IEEE Transactions on Image Processing*, 31:5203–5213, 2022. 1
- [9] Shaoxiang Chen and Yu-Gang Jiang. Semantic proposal for activity localization in videos via sentence query. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 8199–8206, 2019. 6
- [10] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020. 1, 2
- [11] Xuxin Cheng, Bowen Cao, Qichen Ye, Zhihong Zhu, Hongxiang Li, and Yuexian Zou. MI-Imcl: Mutual learning and large-margin contrastive learning for improving asr robustness in spoken language understanding. In *Proc. of ACL Findings*, 2023. 2
- [12] Xuxin Cheng, Wanshi Xu, Ziyu Yao, Zhihong Zhu, Yaowei Li, Hongxiang Li, and Yuexian Zou. FC-MTLF: A Fine- and Coarse-grained Multi-Task Learning Framework for Cross-Lingual Spoken Language Understanding. In *Proc. of InterSpeech*, 2023. 2
- [13] Xuxin Cheng, Zhihong Zhu, Hongxiang Li, Yaowei Li, and Yuexian Zou. Ssvm: Saliency-based self-training for video-music retrieval. *arXiv preprint arXiv:2302.09328*, 2023. 1
- [14] Somnath Basu Roy Chowdhury, Nicholas Monath, Avinava Dubey, Amr Ahmed, and Snigdha Chaturvedi. Unsupervised opinion summarization using approximate geodesics. *arXiv preprint arXiv:2209.07496*, 2022. 3
- [15] Thomas Cover and Peter Hart. Nearest neighbor pattern classification. *IEEE transactions on information theory*, 13(1):21–27, 1967. 3
- [16] Edsger W Dijkstra. A note on two problems in connexion with graphs. In *Edsger Wybe Dijkstra: His Life, Work, and Legacy*, pages 287–290, 2022. 3
- [17] Xinpeng Ding, Nannan Wang, Shiwei Zhang, De Cheng, Xiaomeng Li, Ziyuan Huang, Mingqian Tang, and Xinbo Gao. Support-set based cross-supervision for video grounding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11573–11582, 2021. 1, 2, 3, 6, 7
- [18] Jiyang Gao, Chen Sun, Zhenheng Yang, and Ram Nevatia. Tall: Temporal activity localization via language query. In *Proceedings of the IEEE international conference on computer vision*, pages 5267–5275, 2017. 1, 2, 6, 7
- [19] Junyu Gao and Changsheng Xu. Fast video moment retrieval. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1523–1532, 2021. 6, 7
- [20] Michel Grabisch and Marc Roubens. An axiomatic approach to the concept of interaction among players in cooperative games. *International Journal of game theory*, 28:547–565, 1999. 2, 5
- [21] Michael Gutmann and Aapo Hyvärinen. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 297–304. JMLR Workshop and Conference Proceedings, 2010. 3
- [22] Meera Hahn, Asim Kadav, James M Rehg, and Hans Peter Graf. Tripping through time: Efficient localization of activities in videos. *arXiv preprint arXiv:1904.09936*, 2019. 6
- [23] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9729–9738, 2020. 1, 2
- [24] R Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Phil Bachman, Adam Trischler, and Yoshua Bengio. Learning deep representations by mutual information estimation and maximization. *arXiv preprint arXiv:1808.06670*, 2018. 2
- [25] Ji Jiang, Meng Cao, Tengtao Song, and Yuexian Zou. Video referring expression comprehension via transformer with content-aware query. *arXiv preprint arXiv:2210.02953*, 2022. 1
- [26] Peng Jin, Jinfa Huang, Pengfei Xiong, Shangxuan Tian, Chang Liu, Xiangyang Ji, Li Yuan, and Jie Chen. Video-text as game players: Hierarchical banzhaf interaction for cross-modal representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2472–2482, 2023. 3

- [27] Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT*, pages 4171–4186, 2019. 3, 6
- [28] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. *Advances in Neural Information Processing Systems*, 33:18661–18673, 2020. 4
- [29] Ron Kimmel and James A Sethian. Computing geodesic paths on manifolds. *Proceedings of the national academy of Sciences*, 95(15):8431–8435, 1998. 2, 3
- [30] Ranjay Krishna, Kenji Hata, Frederic Ren, Li Fei-Fei, and Juan Carlos Niebles. Dense-captioning events in videos. In *Proceedings of the IEEE international conference on computer vision*, pages 706–715, 2017. 6
- [31] Hongxiang Li, Meng Cao, Xuxin Cheng, Zhihong Zhu, Yaowei Li, and Yuexian Zou. Generating templated caption for video grounding. *arXiv preprint arXiv:2301.05997*, 2023. 2, 4
- [32] Juncheng Li, Xin He, Longhui Wei, Long Qian, Linchao Zhu, Lingxi Xie, Yueting Zhuang, Qi Tian, and Siliang Tang. Fine-grained semantically aligned vision-language pre-training. *arXiv preprint arXiv:2208.02515*, 2022. 3, 5
- [33] Jiahui Li, Kun Kuang, Baoxiang Wang, Furui Liu, Long Chen, Fei Wu, and Jun Xiao. Shapley counterfactual credits for multi-agent reinforcement learning. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pages 934–942, 2021. 3
- [34] Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. Align before fuse: Vision and language representation learning with momentum distillation. *Advances in neural information processing systems*, 34:9694–9705, 2021. 1
- [35] Yaowei Li, Bang Yang, Xuxin Cheng, Zhihong Zhu, Hongxiang Li, and Yuexian Zou. Unify, align and refine: Multi-level semantic alignment for radiology report generation. In *Proc. of ICCV*, 2023. 2
- [36] Daizong Liu and Wei Hu. Skimming, locating, then perusing: A human-like framework for natural language video localization. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 4536–4545, 2022. 1
- [37] Daizong Liu, Xiaoye Qu, Jianfeng Dong, Pan Zhou, Yu Cheng, Wei Wei, Zichuan Xu, and Yulai Xie. Context-aware biaffine localizing network for temporal sentence grounding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11235–11244, 2021. 6, 7
- [38] Daizong Liu, Xiaoye Qu, Xiao-Yang Liu, Jianfeng Dong, Pan Zhou, and Zichuan Xu. Jointly cross-and self-modal graph attention network for query-based moment localization. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 4070–4078, 2020. 7
- [39] Meng Liu, Xiang Wang, Liqiang Nie, Xiangnan He, Baoquan Chen, and Tat-Seng Chua. Attentive moment retrieval in videos. In *The 41st international ACM SIGIR conference on research & development in information retrieval*, pages 15–24, 2018. 2, 6
- [40] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2018. 6
- [41] Yasuko Matsui and Tomomi Matsui. Np-completeness for calculating power indices of weighted majority games. *Theoretical Computer Science*, 263(1-2):305–310, 2001. 2
- [42] Antoine Miech, Jean-Baptiste Alayrac, Lucas Smaira, Ivan Laptev, Josef Sivic, and Andrew Zisserman. End-to-end learning of visual representations from uncurated instructional videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9879–9889, 2020. 2
- [43] Ishan Misra and Laurens van der Maaten. Self-supervised learning of pretext-invariant representations. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6707–6717, 2020. 1, 2
- [44] Jonghwan Mun, Minsu Cho, and Bohyung Han. Local-global video-text interactions for temporal grounding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10810–10819, 2020. 6
- [45] Guoshun Nan, Rui Qiao, Yao Xiao, Jun Liu, Sicong Leng, Hao Zhang, and Wei Lu. Interventional video grounding with dual contrastive learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2765–2775, 2021. 1, 2, 3, 6, 7
- [46] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 1
- [47] Michaela Regneri, Marcus Rohrbach, Dominikus Wetzel, Stefan Thater, Bernt Schiele, and Manfred Pinkal. Grounding action descriptions in videos. *Transactions of the Association for Computational Linguistics*, 1:25–36, 2013. 6
- [48] Jie Ren, Die Zhang, Yisen Wang, Lu Chen, Zhanpeng Zhou, Yiting Chen, Xu Cheng, Xin Wang, Meng Zhou, Jie Shi, et al. A unified game-theoretic interpretation of adversarial robustness. *arXiv preprint arXiv:2103.07364*, 2021. 3
- [49] Lloyd S Shapley. A value for n-person games. *Classics in game theory*, 69, 1997. 2, 3, 5
- [50] Gunnar A Sigurdsson, Gül Varol, Xiaolong Wang, Ali Farhadi, Ivan Laptev, and Abhinav Gupta. Hollywood in homes: Crowdsourcing data collection for activity understanding. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part I 14*, pages 510–526. Springer, 2016. 6
- [51] Chen Sun, Fabien Baradel, Kevin Murphy, and Cordelia Schmid. Learning video representations using contrastive bidirectional transformer. *arXiv preprint arXiv:1906.05743*, 2019. 2
- [52] Vitaly Surazhsky, Tatiana Surazhsky, Danil Kirsanov, Steven J Gortler, and Hugues Hoppe. Fast exact and approximate geodesics on meshes. *ACM transactions on graphics (TOG)*, 24(3):553–560, 2005. 3

- [53] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008. 8
- [54] Hao Wang, Zheng-Jun Zha, Xuejin Chen, Zhiwei Xiong, and Jiebo Luo. Dual path interaction network for video moment localization. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 4116–4124, 2020. 6
- [55] Hao Wang, Zheng-Jun Zha, Liang Li, Dong Liu, and Jiebo Luo. Structured multi-level interaction network for video moment localization via language query. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7026–7035, 2021. 6
- [56] Haotao Wang, Aston Zhang, Yi Zhu, Shuai Zheng, Mu Li, Alex J Smola, and Zhangyang Wang. Partial and asymmetric contrastive learning for out-of-distribution detection in long-tailed recognition. In *International Conference on Machine Learning*, pages 23446–23458. PMLR, 2022. 4
- [57] Jingwen Wang, Lin Ma, and Wenhao Jiang. Temporally grounding language queries in videos by contextual boundary-aware prediction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 12168–12175, 2020. 7
- [58] Tongzhou Wang and Phillip Isola. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *International Conference on Machine Learning*, pages 9929–9939. PMLR, 2020. 1
- [59] Zhenzhi Wang, Limin Wang, Tao Wu, Tianhao Li, and Gangshan Wu. Negative sample matters: A renaissance of metric learning for temporal grounding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 2613–2623, 2022. 1, 3, 6, 7, 8
- [60] Robert J Weber. Probabilistic values for games. *The Shapley Value. Essays in Honor of Lloyd S. Shapley*, pages 101–119, 1988. 3
- [61] Huijuan Xu, Kun He, Bryan A Plummer, Leonid Sigal, Stan Sclaroff, and Kate Saenko. Multilevel language and vision integration for text-to-clip retrieval. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 9062–9069, 2019. 2
- [62] Shuzhou Yang, Moxuan Ding, Yanmin Wu, Zihan Li, and Jian Zhang. Implicit neural representation for cooperative low-light image enhancement, 2023. 1
- [63] Xun Yang, Shanshan Wang, Jian Dong, Jianfeng Dong, Meng Wang, and Tat-Seng Chua. Video moment retrieval with cross-modal neural architecture search. *IEEE Transactions on Image Processing*, 31:1204–1216, 2022. 6
- [64] Yuhang Yang, Wei Zhai, Hongchen Luo, Yang Cao, Jiebo Luo, and Zheng-Jun Zha. Grounding 3d object affordance from 2d interactions in images. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2023. 1
- [65] Yitian Yuan, Lin Ma, Jingwen Wang, Wei Liu, and Wenwu Zhu. Semantic conditioned dynamic modulation for temporal sentence grounding in videos. *Advances in Neural Information Processing Systems*, 32, 2019. 6, 7
- [66] Yitian Yuan, Tao Mei, and Wenwu Zhu. To find where you talk: Temporal sentence localization in video with attention based location regression. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 9159–9166, 2019. 2
- [67] Runhao Zeng, Haoming Xu, Wenbing Huang, Peihao Chen, Mingkui Tan, and Chuang Gan. Dense regression network for video grounding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10287–10296, 2020. 1, 2, 4, 6, 7
- [68] Yawen Zeng, Da Cao, Xiaochi Wei, Meng Liu, Zhou Zhao, and Zheng Qin. Multi-modal relational graph for cross-modal video moment retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2215–2224, 2021. 6
- [69] Can Zhang, Meng Cao, Dongming Yang, Jie Chen, and Yuexian Zou. Cola: Weakly-supervised temporal action localization with snippet contrastive learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16010–16019, 2021. 1
- [70] Da Zhang, Xiyang Dai, Xin Wang, Yuan-Fang Wang, and Larry S Davis. Man: Moment alignment network for natural language moment retrieval via iterative graph adjustment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1247–1257, 2019. 6
- [71] Hao Zhang, Aixin Sun, Wei Jing, Guoshun Nan, Liangli Zhen, Joey Tianyi Zhou, and Rick Siow Mong Goh. Video corpus moment retrieval with contrastive learning. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 685–695, 2021. 2
- [72] Hao Zhang, Aixin Sun, Wei Jing, Liangli Zhen, Joey Tianyi Zhou, and Rick Siow Mong Goh. Natural language video localization: A revisit in span-based question answering framework. *IEEE transactions on pattern analysis and machine intelligence*, 2021. 1
- [73] Hao Zhang, Aixin Sun, Wei Jing, and Joey Tianyi Zhou. Span-based localizing network for natural language video localization. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6543–6554, 2020. 6
- [74] Mingxing Zhang, Yang Yang, Xinghan Chen, Yanli Ji, Xing Xu, Jingjing Li, and Heng Tao Shen. Multi-stage aggregated transformer network for temporal language localization in videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12669–12678, 2021. 6
- [75] Songyang Zhang, Houwen Peng, Jianlong Fu, and Jiebo Luo. Learning 2d temporal adjacent networks for moment localization with natural language. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 12870–12877, 2020. 1, 3, 6, 7, 8
- [76] Zhu Zhang, Zhijie Lin, Zhou Zhao, and Zhenxin Xiao. Cross-modal interaction networks for query-based moment retrieval in videos. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 655–664, 2019. 7