

MAD: A Scalable Dataset for Language Grounding in Videos from Movie Audio Descriptions

Mattia Soldan¹, Alejandro Pardo¹, Juan León Alcázar¹, Fabian Caba Heilbron²,
 Chen Zhao¹, Silvio Giancola¹, Bernard Ghanem¹

¹King Abdullah University of Science and Technology (KAUST) ²Adobe Research

{mattia.soldan, alejandro.pardo, juancarlo.alcazar, chen.zhao,
 silvio.giancola, bernard.ghanem}@kaust.edu.sa caba@adobe.com

Abstract

The recent and increasing interest in video-language research has driven the development of large-scale datasets that enable data-intensive machine learning techniques. In comparison, limited effort has been made at assessing the fitness of these datasets for the video-language grounding task. Recent works have begun to discover significant limitations in these datasets, suggesting that state-of-the-art techniques commonly overfit to hidden dataset biases. In this work, we present MAD (Movie Audio Descriptions), a novel benchmark that departs from the paradigm of augmenting existing video datasets with text annotations and focuses on crawling and aligning available audio descriptions of mainstream movies. MAD contains over 384,000 natural language sentences grounded in over 1,200 hours of videos and exhibits a significant reduction in the currently diagnosed biases for video-language grounding datasets. MAD’s collection strategy enables a novel and more challenging version of video-language grounding, where short temporal moments (typically seconds long) must be accurately grounded in diverse long-form videos that can last up to three hours. We have released MAD’s data and baselines code at <https://github.com/Soldelli/MAD>.

1. Introduction

Imagine you want to find the moment in time, in a movie, when your favorite actress is eating a Gumbo dish in a New Orleans restaurant. You could do so by manually scrubbing the film to ground the moment. However, such a process is tedious and labor-intensive. This task is known as natural language grounding [1, 4], and has gained significant momentum in the computer vision community. Beyond smart browsing of movies, the interest in this task stems from multiple real-world applications ranging from smart video search [26, 27], video editing [18, 19, 32], and helping pa-




Figure 1. **Comparison of video-language grounding datasets.** The circle size measures the language vocabulary diversity. The videos in MAD are orders of magnitude longer in duration than previous datasets (~110min), annotated with natural, highly descriptive, language grounding (>60K unique words) with very low coverage in video (~4.1s). Coverage is defined as the average % duration of moments with respect to the total video duration.

tients with memory dysfunction [2, 31]. The importance of solving this task has resulted in novel approaches and large-scale deep-learning architectures that steadily push state-of-the-art performance.

Despite those advances, recent works [17, 34, 37] have diagnosed hidden biases in the most common video-language grounding datasets. Otani *et al.* [17] have highlighted that several grounding methods only learn location priors, to the point of disregarding the visual information and only using language cues for predictions. While recent methods [34, 37] have tried to circumvent these limitations by either proposing new metrics [35] or debiasing strategies [37, 41], it is still unclear if existing grounding datasets [1, 4, 7, 21] provide the right setup to evalu-



Figure 2. **Example from our MAD dataset.** We select the movie “A quiet place” as representative for our dataset. As shown in the figure, the movie contains a large number of densely distributed temporally grounded sentences. The collected annotations can be very descriptive, mentioning people, actions, locations, and other additional information. Note that as per the movies plot, the characters are silent for the vast majority of the movie, rendering audio description essential for visually-impaired audience.

ate progress in this important task. This is partly because datasets used by video-language grounding models were not originally collected to solve this task. These datasets provide valuable video-language pairs for captioning or retrieval, but the grounding task requires high-quality (and dense) temporal localization of the language. Figure 1 shows that the current datasets comprise relatively short videos, containing single structured scenes, and language descriptions that cover most of the video. Furthermore, the temporal anchors for the language are temporally biased, leading to methods not learning from any visual features and eventually overfitting to temporal priors for specific actions, thus limiting their generalization capabilities [8, 17].

In this work, we address these limitations with a novel large-scale dataset called MAD (Movie Audio Descriptions). MAD builds atop (and includes part of) the LSMDC dataset [24], which is a pioneering work in leveraging audio descriptions to enable the investigation of a closely related task: text-to-video retrieval. Similar to LSMDC, we depart from the standard annotation pipelines that rely on crowd-sourced annotation platforms. Instead, we adopt a scalable data collection strategy that leverages professional, grounded audio descriptions of movies for visually impaired audiences. As shown in Figure 2, MAD grounds these audio descriptions on long-form videos, bringing novel challenges to the video-language grounding task.

Our data collection approach consists of transcribing the audio description track of a movie and automatically detecting and removing sentences associated with the actor’s speech, yielding an authentic “untrimmed video” setup where the highly descriptive sentences are grounded in long-form videos. Figures 1 and 3 illustrate the uniqueness of our dataset with respect to the current alternatives.

As showcased in Figure 2, MAD contains videos that, on average, span over 110 minutes, as well as grounded annotations covering short time segments, which are uniformly distributed in the video, and maintain the largest diversity in vocabulary. Video grounding in MAD requires a finer understanding of the video, since the average coverage of the sentences is much smaller than in current datasets.

The unique configuration of the MAD dataset introduces exciting challenges. First, the video grounding task is now mapped into the domain of long form video, preventing current methods to learn trivial temporal location priors, instead requiring a more nuanced understanding of the video and language modalities. Second, having longer videos means producing a larger number of segment proposals, which will make the localization problem far more challenging. Last, these longer sequences emphasize the necessity for efficient methods in inference and training, mandatory for real-world applications such as long live streaming videos or moment retrieval in large video collections [3].

Contributions. Our contributions are threefold. (1) We propose Movie Audio Description (MAD), a novel large-scale dataset for video-language grounding, containing more than 384K natural language sentences anchored on more than 1.2K hours of video. (2) We design a scalable data collection pipeline that automatically extracts highly valuable video-language grounding annotations, leveraging speech-to-text translation on professionally generated audio descriptions. (3) We provide a comprehensive empirical study that highlights the benefits of our large-scale MAD dataset on video-language grounding as a benchmark, pointing out the difficulties faced by current video-language grounding baselines in long-form videos.

2. Related work

Video Grounding Benchmarks. Most of the current video grounding datasets were previously collected and tailored for other computer vision tasks (*i.e.*, Temporal Activity Localization [5, 21]) and purposes (*i.e.*, Human Action Recognition), then annotated for the video grounding task. This adaptation limits the diversity of the video corpus towards a specific set of actions and objects, and its corresponding natural language sentences to specific sets of verbs and nouns [4, 7, 17, 22]. Currently, ActivityNet-Captions [7] and Charades-STA [4] are the most commonly used benchmarks for the task of video grounding. Both datasets have been collected atop pre-existing video datasets (ActivityNet [5] and Charades [25]) and have been diagnosed with severe biases by Otani *et al.* [17]. These findings show that the annotations contain distinct biases where language tokens are often coupled with specific temporal locations. Moreover, strong priors also affect the temporal endpoints, with a large portion of the annotations spanning the entire video. As a consequence, current methods seem to mainly rely on such biases to make predictions, often disregarding the visual input altogether. In comparison, the unique setup of MAD prevents these drawbacks as keywords are not associated with particular temporal regions, and annotation timestamps are much shorter than the video’s duration. Moreover, different from Charades-STA, MAD defines an official validation set for hyper-parameter tuning.

Unlike Charades-STA and ActivityNet-Captions, TACoS [21] has not been diagnosed with annotation biases. However, its video corpus is small and limited to cooking actions recorded in a static-camera setting. Conversely, spanning over 22 genres across 90 years of cinema history, MAD covers a broad domain of actions, locations, and scenes. Moreover, MAD inherits a diverse set of visual and linguistic content from the broad movie genres, ranging from fiction to everyday life.

Furthermore, DiDeMo [1] was annotated atop Flickr videos with a discrete annotation scheme (*i.e.* in chunks of 5 seconds) for a maximum of 30 seconds, constraining the problem of video grounding to trimmed videos. Given these annotations, the grounding task can be simplified to choosing one out of 21 possible proposals for each video. Conversely, MAD provides a setup to explore solutions for grounding language in long-form videos, whose length can be up to 3 hours. In this scenario, naive sliding window techniques for proposal generation could produce hundreds of thousands of possible candidates. Therefore, developing efficient inference methods becomes a much more urgent requirement compared to previous benchmarks.

State-of-the-art video grounding methods [12, 13, 15, 28, 36, 38, 40] have relied on existing benchmarks to design novel modules (*i.e.*, proposal generation, context modeling,

and multi-modality fusion). However, most of these designs specifically target grounding in short videos and often rely on providing the entire video to the model when making a prediction. As the long-form setup introduced by MAD prohibits this, new methods will have the opportunity to investigate and bridge previous ideas to these new challenging and real-world constraints.

Audio Descriptions. The pioneering works of Rohrbach *et al.* [23] and Torabi *et al.* [30] were the first to exploit audio descriptions to study the text-to-video retrieval task and its counterpart video-to-text. Rohrbach *et al.* [23] introduced the MPII-MD dataset, while Torabi *et al.* [30] presented M-VAD, both collected from audio descriptions in movies. Later, these datasets were fused to create the LSMDC dataset [24], which forms the core of the LSMDC annual challenge. Our annotation pipeline is similar in spirit to these works. However, MAD sizes the potential of movies’ audio descriptions for a new scenario: language grounding in videos.

A concurrent work introduced a new benchmark based on audio descriptions in videos, called QuerYD [16]. It is a dataset for retrieval and event localization in videos crawled from YouTube. This benchmark focuses on the short-form video setup with videos of less than 5 minutes average duration. QuerYD also leverages audio descriptions, which have been outsourced to volunteer narrators. Similar to our takeaways, the authors noticed that audio descriptions are generally more visually grounded and descriptive than previously collected annotations.

3. Collecting the MAD Dataset

In this section, we outline MAD’s data collection pipeline. We follow independent strategies for creating the training and testing set. For the former, we aim at automatically collecting a large set of annotations. For the latter, we re-purpose the manually refined annotations in LSMDC. We provide detailed statistics of MAD’s annotations and compare them to existing datasets.

3.1. MAD Training set

MAD relies on audio descriptions professionally created to make movies accessible to visually-impaired audiences. These descriptions embody a rich narrative describing the most relevant visual information. Thus, they adopt a highly descriptive and diverse language. Audio descriptions are often available as an alternative audio track that can replace the original one. Professional narrators curate them, and devote significant effort to describe a movie. The audio description process demands an average of 30 work hours to narrate a single hour of video [24]. In comparison, previous datasets that have used the Amazon Mechanical Turk service for video-language grounding estimate the annotation effort to be around 3 hours for each video hour [7].

Dataset	Videos			Language Queries							
	Total Duration	Duration / Video	Duration / Moment	Total Queries	# Words / Query	Total Tokens	Vocabulary				
							Adj.	Nouns	Verbs	Total	
TACoS [21]	10.1 h	4.78 min	27.9 s	18.2K	10.5	0.2M	0.2K	0.9K	0.6K	2.3K	
Charades-STA [4]	57.1 h	0.50 min	8.1 s	16.1K	7.2	0.1M	0.1K	0.6K	0.4K	1.3K	
DiDeMo [1]	88.7 h	0.50 min	6.5 s	41.2K	8.0	0.3M	0.6K	4.1K	1.9K	7.5K	
ANet-Captions [7]	487.6 h	1.96 min	37.1 s	72.0K	14.8	1.0M	1.1K	7.4K	3.7K	15.4K	
MAD (Ours)	1207.3 h	110.77 min	4.1 s	384.6K	12.7	4.9M	5.3K	35.5K	13.1K	61.4K	

Table 1. **Statistics of video-language grounding datasets.** We report relevant statistics to compare our MAD dataset against other video grounding benchmarks. MAD provides the largest dataset with 1207hrs of video and 384.6K language queries, the longest form of video (avg. 110.77min), the most diverse language vocabulary with 61.4K unique words, and the shortest moment for grounding (avg. 4.1s).

Data Crawling. Not every commercially available movie is released with audio descriptions. However, we can obtain these audio descriptions from 3rd party creators. In particular, we crawl our audio descriptions from a large open-source and online repository¹. These audio files contain the original movie track mixed with the narrator’s voice, carefully placed when actors are not speaking. One potential problem is that the audio descriptions can be misaligned with the original movie. Such misalignment comes either from a delay in the recording of the audio description (concerning the original movie) or from audio descriptions being created from different versions of the movie (with deleted or trimmed scenes).

Alignment and Cleanup. Since the audio description track also contains the movie’s original audio, we can resolve this misalignment by maximizing the cross-correlation between overlapping segments of the original audio track and the audio description track. We define the original audio signal (f), the audio description signal (g), and the time delay (τ_{delay}) between the two signals. The maximum of the cross-correlation function (denoted with the operator \star) indicates the point in time where the signals exhibit the best alignment. As a result, the time delay τ_{delay} between f and g is defined as follows:

$$\tau_{\text{delay}} = \arg \max_t ((f \star g)(t)) \quad (1)$$

To verify that our single time delay τ_{delay} defines the best possible alignment between the audio descriptions and the original movies, we run our synchronization strategy over several (i.e., 20) uniformly distributed temporal windows. We verify that the delays estimated for all windows are consistent with each other, within a maximum range of ± 0.1 seconds w.r.t the median value of the distribution of the 20 samples. We discard the movies that do not satisfy this criterion to ensure that the whole audio description track correctly aligns with the original movie’s visual content.

Audio Transcriptions and Verification. After the two au-

dio tracks are aligned, we transcribe the audio description file using Microsoft’s Azure Speech-to-Text service². Each recognized word is associated with a temporal timestamp. At this step in our pipeline, we have sentences temporally grounded to the original video stream. As the Speech-to-Text contains both the narrations and the actors’ speech, we set out to remove the latter and only retain the audio description in textual form. To do so, we resort to the movie’s original subtitles and use their timestamps as a surrogate for Voice Activity Detection (VAD). Particularly, we discard closed captions and retain subtitles associated with the actors’ speech or subtitles for songs. Then, we remove from the Speech-to-Text output every sentence overlapping with the VAD temporal locations, obtaining our target audio description sentences. We post-process the text for automatic punctuation refinement using the free tool Punctuator [29].

3.2. From LSMDC to MAD Val/Test

As outlined in Section 3.1, the val/test sets of MAD also relies on audio descriptions for movies. Since the annotations in training are automatically generated, we decided to minimize the noise in the validation and test splits. Hence, we avoid the automatic collection of data for these sets and resort to the data made available by the LSMDC dataset [24]. This dataset collected annotations from audio descriptions in movies targeting the video retrieval task. LSMDC manually refined the grammar and temporal boundaries of sentences. As a consequence, these annotations have clean language and precise temporal boundaries. We reformat a subset of the LSMDC data, adapt it for the video grounding task, and cast it as MAD’s val and test sets.

LSMDC data for retrieval is made available only as video chunks, not full movies. To create data suitable for long-form video language grounding, we collect 162 out of the 182 videos in LSMDC and their respective audio descriptions. Again, the full-length video data and the chunked video data provided by LSMDC might not be in sync. To

¹<https://www.audiovault.net/>

²<https://azure.microsoft.com/en-us/services/cognitive-services/speech-to-text/>




Figure 3. **Histograms of moment start/end/duration in video-language grounding datasets.** The plots represent the normalized (by video length) start/end histogram (a-b) and absolute duration distribution (c) for moments belonging to each of the five datasets. We notice severe biases in ActivityNet-Captions and Charades-STA, which show high peaks at the beginning and end of the videos. Conversely MAD does not show any particular preferred start/end temporal location.

align a video chunk from LSMDC with our full-length movies, we follow a similar procedure as the one described in Section 3.1 for the audio alignment, but using visual information. We use CLIP [20] to extract five frame-level features per second for both the video chunks from LSMDC and our full-length movie. Then, we use the maximum cross-correlation score to estimate the delay between the two. We run the alignment on 10 different window lengths and take the median value as the delay of the chunk.

Once the audiovisual data is aligned, we use the text annotations and re-define their original timestamps according to the calculated delays. This process creates full-length movies with curated grounding data for MAD’s validation and test sets. In doing so, we obtain large and clean validation and test sets since LSMDC was curated by humans and has been used for years by the community. Our val/test sets evaluate video-grounding methods with more than 104K grounded phrases coming from more than 160 movies.

3.3. MAD Dataset Analysis

We now present MAD’s statistics and compare them with legacy datasets for the task. Table 1 summarizes the most notable aspects of MAD regarding both the visual and language content of the dataset.

Scale and Scalability. MAD is the largest dataset in video hours and number of sentences. The training, validation, and test sets consist of 488, 50, and 112 movies, respectively. Although other datasets have a larger number of clips, MAD’s videos are full movies which last 2 hours on average. In comparison, the average clip from other datasets spans just a few minutes. Overall, MAD splits contain a combined 50 days of continuous video. We highlight that our test set alone is already larger than any other video grounding dataset. We also emphasize that each movie in MAD is long and composed of several diverse scenes, making it a rich source for long-form video analysis. Also, as the cinema industry is ever-growing, we expect to periodically expand MAD with subsequent releases to fuel further

innovation in this research direction.

Vocabulary Size. Besides having the largest video corpus, MAD also contains the largest and most diverse query set of any dataset. In Table 1, we show that MAD contains the largest set of adjectives, nouns, and verbs among all available benchmarks. In almost every aspect, it is an order of magnitude larger. The number of sentences in MAD training, validation, and test is 280.5K, 32.1K, and 72.0K, respectively, one order of magnitude larger than the equivalent set in any other dataset. Overall, MAD contains 61.4K unique words, almost 4 times more than the 15.4K of ActivityNet-Captions [7] (the highest among the other benchmarks). Finally, the average length per sentence is 12.7 words, which is similar to the other datasets.

Bias Analysis. Figure 3 plots the histograms for start/end timestamps of moments in all grounding datasets. We notice clear biases in current datasets: Charades-STA [4], DiDeMo [1], and ActivityNet-Captions [7]. Charades-STA and ActivityNet-Captions are characterized by tall peaks at the beginning (Figure 3a) and end (Figure 3b) of the video, meaning that most temporal annotations start at the video’s start or finish at the video’s end. This bias is learned quickly by modern machine learning algorithms, resulting in trivial groundings that span a full video. The smaller dataset TACoS also exhibits a similar bias, although it is less prominent. DiDeMo is limited by its annotation strategy, where chunks of 5 seconds are labeled up to a maximum of 30 seconds. This favors structured responses that roughly approximate the start and end points of a moment. In contrast, MAD has an almost uniform histogram. This means that moments of interest can start and end at any point in the video. We only observe a minor imbalance where the end of the movie has slightly more descriptions than the beginning. This is related to the standard structure of a film, where the main plot elements are resolved towards the end, thus creating more situations worth describing. Figure 3c plots the histograms for moment duration. MAD is characterized by

Model	IoU=0.1					IoU=0.3					IoU=0.5				
	R@1	R@5	R@10	R@50	R@100	R@1	R@5	R@10	R@50	R@100	R@1	R@5	R@10	R@50	R@100
Oracle	100.00	—	—	—	—	100.00	—	—	—	—	99.99	—	—	—	—
Random Chance	0.09	0.44	0.88	4.33	8.47	0.04	0.19	0.39	1.92	3.80	0.01	0.07	0.14	0.71	1.40
CLIP [20]	6.57	15.05	20.26	37.92	47.73	3.13	9.85	14.13	28.71	36.98	1.39	5.44	8.38	18.80	24.99
VLG-Net [28]	3.64	11.66	17.89	39.78	51.24	2.76	9.31	14.65	34.27	44.87	1.65	5.99	9.77	24.93	33.95

Table 2. **Benchmarking of grounding baselines on the MAD dataset.** We report the performance of four baselines: *Oracle*, *Random Chance*, *CLIP*, *VLG-Net*, on the test split. The first two validate the choice of proposals by computing the upper bound to the performance and the random performance. CLIP and VLG-Net use visual and language features to score and rank proposals. For all experiments, we adopt the same proposal scheme as in VLG-Net [28], and use CLIP [20] features for video (frames) and language embeddings.

shorter moments on average, having a long tail distribution with moments that last up to one minute.

4. Experiments

We proceed with the experimental assessment of video-language grounding on the MAD dataset. We first describe the video grounding task in the MAD dataset along with its evaluation metrics and then report the performance of four selected baselines.

Task. Given an untrimmed video and a language query, the video-language grounding task aims to localize a temporal moment (τ_s, τ_e) in the video that matches the query [1, 4].

Metric. Following the grounding literature [1, 4], we adopt Recall@ K for IoU= θ ($R@K$ -IoU= θ). Given a ranked set of video proposals, this metric measures if any of the top K ranked moments have an IoU larger than θ with the ground truth temporal endpoints. Results are averaged across all test samples. Given the long-form nature of our videos and the large amount of possible proposals, we investigate Recall@ K for IoU= θ with $K \in \{1, 5, 10, 50, 100\}$ and $\theta \in \{0.1, 0.3, 0.5\}$. This allows us to evaluate for loose alignment (*i.e.* IoU=0.1) and approximate ranking (*i.e.*, $K=100$), as well as tight predictions (*i.e.* IoU=0.5) and accurate retrieval (*i.e.* $K=1$).

Baselines. We benchmark MAD using four different grounding strategies, namely: *Oracle*, *Random Chance*, *CLIP* [20], and *VLG-Net* [28]. The first two provide upper bounds and random performance for the recall metric given a predefined set of proposals. *Oracle* chooses the proposal with the highest IoU with the ground-truth annotation, while *Random Chance* chooses a random proposal with uniform probability. We also use *CLIP*, the pre-trained image-text architecture from [20], to extract frame-level and sentence-level features. The frame-level features for each proposal are combined using mean pooling, then we score each proposal using cosine similarity between the visual and the text features. Finally, we adopt VLG-Net [28] as a representative, state-of-the-art method for the grounding task. VLG-Net leverages recent advances in Graph Convolution Networks (GCNs) [9–11, 33, 39] to model individual

modalities, while also enabling the aggregation of non-local and cross-modal context through graph convolutions. We use VLG-Net and adapt it to work with MAD’s long-form videos. See the **supplementary material** for details.

Implementation Details. For the lower bound estimation using *Random Chance*, we average the performance over 100 independent runs. To favor a fair comparison against the CLIP baseline, we train VLG-Net using CLIP features for both modalities. We use the official VLG-Net’s implementation with a clip size of 128 input frames spanning 25.6 seconds (frames are extracted at 5 fps). We train using the Adam optimizer [6] with learning rate of 10^{-4} . For inference over an entire movie, we adopt a sliding window approach, where we stride the input window by 64 frames and discard highly redundant proposals through Non Maximum Suppression (NMS) with a threshold of 0.3.

4.1. Grounding Performance on MAD

Table 2 summarizes the baseline performance on MAD. The *Oracle* evaluation achieves a perfect score across all metrics except for IoU=0.5. Only a negligible portion of the annotated moments cannot be correctly retrieved at a high IoU (0.5), this result showcases the suitability of the proposal scheme. The low performance of the *Random Chance* baseline reflects the difficulty of the task, given the vast pool of proposals extracted over a single video. For the least strict metric ($R@100$ -IoU=0.1), this baseline only achieves 8.47%, while CLIP and VLG-Net baselines are close to 50%, a $\sim 6\times$ relative improvement. An even larger gap is present for the most strict metric, $R@1$ -IoU=0.5, with a relative improvement of two orders of magnitude.

The CLIP [20] baseline is pre-trained for the task of text-to-image retrieval, and we do not fine-tune this model on the MAD dataset. Nevertheless, when evaluated in a zero-shot fashion, it results in a strong baseline for long-form grounding, achieving the best $R@K$ for the least strict IoU=0.1 at $K=\{1, 5, 10\}$. Although IoU=0.1 corresponds to very loose grounding, this result is nonetheless valuable given a large number of negatives in the long-form setup, and the fact that MAD is characterized by containing short moments (4.1s on average). Although VLG-Net is

Model	IoU=0.1		IoU=0.3		IoU=0.5	
	R@1	R@5	R@1	R@5	R@1	R@5
Oracle	100.00	—	99.88	—	99.42	—
Random Chance	3.40	15.69	1.47	7.09	0.52	2.61
CLIP [20]	20.98	45.49	9.74	29.63	4.03	15.90
VLG-Net [28]	23.94	51.46	17.51	43.18	10.17	30.35

Table 3. **Short video setup.** The table showcases the performance of the selected baselines in a short-video setup, where movies are chunked into three minutes (non-overlapping windows). VLG-Net, which falls behind CLIP in the long-form setup, achieves the best grounding performance in most metrics. We can conclude that a new generation of deep learning architectures will have to be investigated to tackle the specific properties of the MAD dataset.

trained for the task at hand, it achieves comparable or better performance with respect to CLIP only when a strict IoU (IoU=0.5) is considered. However, it lags behind CLIP for most other metrics. We believe the shortcomings of VLG-Net are due to two factors. (i) This architecture was developed to ground sentences in short videos, where the entire frame-set can be compared against a sentence in a single forward pass. Thus, it struggles in the long-form setup where we must compare the sentence against all segments of the movie and then aggregate the predictions in a post-processing step. (ii) VLG-Net training procedure defines low IoU moments as negatives, thus favoring high performance only for higher IoUs.

4.2. The Challenges of Long-form Video Grounding

This section presents an in-depth analysis of the performance of the selected baselines in the long-form setup. We first investigate how the performance changes when the evaluation is constrained over segments of the movie, whose length is comparable to current datasets. Then explore how methods behave as the size of the movie chunks changes. To this end, we split each video into non-overlapping windows (short videos), and assign the annotations to the short-video with the highest temporal overlap.

Short-video Setup. In Table 3, we set the short-video window length to three minutes. This duration is a candidate representative for short videos. The upper bound performance *Oracle* slightly decreases to 99.42% for IoU=0.5. This is a consequence of the division into short videos, which occasionally breaks down a few ground truth moments. The *Random Chance* baseline reports increased performance as the number of proposals generated is reduced. In particular for R@5-IoU=0.1, the performance increases from 0.44% (Table 2) to 15.69% (Table 3), demonstrating that the short-video setup is less challenging compared to MAD’s original long-form configuration. In a similar trend, performance substantially increases for both CLIP and VLG-Net baselines, with the latter now obtaining the best performances, in all cases.




Figure 4. **Performance trend across different windows lengths.** We observe from the graph the decrease in performance for both CLIP and VLG-Net, as the evaluation window length increases. This demonstrates that current grounding methods cannot tackle the task in the long-form video setting.

From Short- to Long-form Results. Figure 4 showcases the performance trend for the metrics $R@{1, 5}$ -IoU=0.5, when the window length is changed from a small value (30 seconds) to the entire movie duration (average duration is 2hrs). The graph displays how the performance steadily drops as the window length increases, showing the challenging setup of long-form grounding enabled by MAD.

Takeaway. This set of experiments verifies that VLG-Net could successfully outperform the zero-shot CLIP baseline when evaluated in a short-video setup. We can conclude that current state-of-the-art grounding methods are not ready to tackle the long-form setting proposed by MAD. This opens the door to opportunities for the community to leverage previously developed techniques in a more challenging setting and potentially incorporate new constraints when designing deep learning architectures for this task.

5. Ablation Study

This section focuses on the empirical assessment of the quality of MAD training data. To help the reader navigate the following experiments, Table 4 defines a naming convention associated with different splits of the data. LSMDC16 refers to the original data collected by [23, 24, 30] for the task of text-to-video retrieval. LSMDC-G is our adaptation of this dataset for the grounding task, as described in Section 3.2. We remind the reader that we could only retrieve 162 out of 182 movies included in LSMDC16. For these two datasets, we follow the original train/val/test partitions. Finally, we include MAD’s split details.

Improving Grounding Performance with MAD Data.

We are interested in evaluating the contribution our data can bring to the grounding task. We investigate the performance of VLG-Net in the long-form grounding setup when the

Dataset Name	Task	Videos			Annotations		
		Train / Val / Test	Train / Val / Test				
LSMDC16 [24]	Retrieval	155 / 12 / 17	101.1 K / 7.4 K / 10.1 K				
LSMDC-G	Grounding	138 / 11 / 13	89.7 K / 6.7 K / 7.6 K				
MAD	Grounding	488 / 50 / 112	280.5 K / 32.1 K / 72.0 K				

Table 4. **Data split cheat-sheet.** This table clarifies the data splits used in the following experiments (Table 5 and Table 6). LSMDC16 [24] is the original data collected for retrieval. LSMDC-G is our adaptation to the grounding task. MAD is our proposed dataset.

training data change. All trained models are evaluated on the same test split: LSMDC-G test.

The first row of Table 5 shows the performance when VLG-Net is exclusively trained on the LSMDC-G training split. This set only contains data that was manually curated in [24]. The second row is trained with 32% of MAD-training data, which is equivalent to the size of LSMDC-G training split. We observe a drop in performance, which can be associated with the presence of noise introduced by the automatic annotation process in MAD. In the third row, we use the complete MAD training set. Here, MAD’s scale allows us to overcome the performance issues associated with noisy data and improve performance to be comparable to only using the clean LSMDC-G for training. Using 100% of MAD data yields a relative improvement of 20% for R@5-IoU=0.5. Then, we investigate whether the performance of VLG-Net saturates given the amount of data available. To this end, we use 100% of LSMDC-G training and gradually augment it by adding MAD training samples. In these three experiments (rows 4-6), the performance steadily increases. These results suggest that current models for video grounding do benefit from larger-scale datasets, even though the automatically collected training data may be noisy. Moreover, designing scalable strategies for automatic dataset collection is crucial, as the drawbacks of noisy data can be off-

Training Set		Testing Set	IoU=0.5		
% LSMDC-G	% MAD	LSMDC-G	R@1	R@5	R@10
100%	0%	Test	1.36	5.18	8.82
0%	32%	Test	0.60	2.60	5.11
0%	100%	Test	1.61	6.23	10.18
100%	32%	Test	2.18	6.63	10.73
100%	64%	Test	2.23	7.79	11.74
100%	100%	Test	2.82	8.74	13.36

Table 5. **Grounding performance with varying training data.** We investigate VLG-Net [28] grounding performance on LSMDC-G test, when different data regimens are used for training. This compares our automatically collected data (MAD training) against the manually curated one (LSMDC-G). We conclude that expensive manual curation can be avoided if large scale data is available.

Training Set		Testing Set	R@1	R@5	R@10
% LSMDC16	% MAD	LSMDC16			
100%	0%	Test	20.9	39.4	48.5
0%	36%	Test	19.2	35.5	44.8
0%	100%	Test	20.5	38.8	48.7
100%	36%	Test	23.3	40.3	48.8
100%	72%	Test	23.6	41.4	49.3
100%	100%	Test	24.8	40.5	50.0

Table 6. **Retrieval performance on LSMDC16 with model CLIP4Clip [14].** This experiment showcases how MAD data can be valuable for a related task, beyond grounding.

set and even overcome with more training data.

Improving Retrieval Performance with MAD Data.

Adopting the same ablation procedure, we evaluate the possible contribution of our data in a related task, namely text-to-video retrieval. Here, we format our MAD data similarly to LSMDC16, where short videos are trimmed around the annotated timestamps. For this experiment, we use CLIP4Clip [14], a state-of-the-art architecture for the retrieval task, as the baseline. Table 6 reports the performance when different amounts of data is used for training. Again, we see that training with the whole LSMDC16 or MAD leads to a very similar performance. Moreover, our previous argument also holds true in this task. Pouring more data into the task boosts the performance, motivating the benefit of having a scalable dataset like MAD.

Takeaway. The MAD dataset is able to boost performance in two closely related tasks, video grounding and text-to-video retrieval, where we show that scale can compensate for potential noise present due to automatic annotation.

6. Conclusion

The paper presents a new video grounding benchmark called MAD, which builds on high-quality audio descriptions in movies. MAD alleviates the shortcomings of previous grounding datasets. Our automatic annotation pipeline allowed us to collect the largest grounding dataset to date. The experimental section provides baselines for the task solution and highlights the challenging nature of the long-form grounding task introduced by MAD. Our methodology comes with two main hypotheses and limitations: (i) Noise cannot be avoided but can be dealt with through scale. (ii) Due to copyright constraints, MAD’s videos will not be publicly released. However, we will provide all necessary features for our experiments’ reproducibility and promote future research in this direction.

Acknowledgments. This work was supported by the King Abdullah University of Science and Technology (KAUST) Office of Sponsored Research through the Visual Computing Center (VCC) funding.

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 (ICCV)*, 2017.
- [2] Andrew E Budson and Bruce H Price. Memory dysfunction. *New England Journal of Medicine*, 352(7):692–699, 2005.
- [3] Victor Escorcia, Mattia Soldan, Josef Sivic, Bernard Ghanem, and Bryan Russell. Temporal localization of moments in video collections with natural language. *arXiv preprint arXiv:1907.12763*, 2019.
- [4] Gao Jiyang, Sun Chen, Yang Zhenheng, Nevatia, Ram. TALL: Temporal Activity Localization via Language Query. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [5] Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 961–970, 2015.
- [6] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [7] 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 (ICCV)*, 2017.
- [8] Xiaohan Lan, Yitian Yuan, Xin Wang, Zhi Wang, and Wenwu Zhu. A survey on temporal sentence grounding in videos. *arXiv preprint arXiv:2109.08039*, 2021.
- [9] Guohao Li, Matthias Müller, Guocheng Qian, Itzel Carolina Delgadillo Perez, Abdulallah Abualshour, Ali Kassem Thabet, and Bernard Ghanem. Deepgcn: Making gcn as deep as cnns. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.
- [10] Guohao Li, Matthias Müller, Ali Thabet, and Bernard Ghanem. Deepgcn: Can gcn as deep as cnns? In *The IEEE International Conference on Computer Vision (ICCV)*, 2019.
- [11] Guohao Li, Chenxin Xiong, Ali Thabet, and Bernard Ghanem. Deepgcn: All you need to train deeper gcn, 2020.
- [12] 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 (CVPR)*, pages 11235–11244, June 2021.
- [13] 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.
- [14] Huaishao Luo, Lei Ji, Ming Zhong, Yang Chen, Wen Lei, Nan Duan, and Tianrui Li. Clip4clip: An empirical study of clip for end to end video clip retrieval. *arXiv preprint arXiv:2104.08860*, 2021.
- [15] 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 (CVPR)*, 2020.
- [16] Andreea-Maria Oncescu, João F. Henriques, Yang Liu, Andrew Zisserman, and Samuel Albanie. Queryd: A video dataset with high-quality text and audio narrations. In *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2265–2269, 2021.
- [17] Mayu Otani, Yuta Nakashima, Esa Rahtu, and Janne Heikkilä. Uncovering hidden challenges in query-based video moment retrieval. In *The British Machine Vision Conference (BMVC)*, 2020.
- [18] Alejandro Pardo, Fabian Caba, Juan León Alcázar, Ali K Thabet, and Bernard Ghanem. Learning to cut by watching movies. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6858–6868, 2021.
- [19] Alejandro Pardo, Fabian Caba Heilbron, Juan León Alcázar, Ali Thabet, and Bernard Ghanem. Moviecuts: A new dataset and benchmark for cut type recognition. *arXiv preprint arXiv:2109.05569*, 2021.
- [20] 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. *arXiv preprint arXiv:2103.00020*, 2021.
- [21] 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 (ACL)*, 2013.
- [22] Anna Rohrbach, Marcus Rohrbach, Wei Qiu, Annemarie Friedrich, Manfred Pinkal, and Bernt Schiele. Coherent multi-sentence video description with variable level of detail. In *German conference on pattern recognition*, pages 184–195. Springer, 2014.
- [23] Anna Rohrbach, Marcus Rohrbach, Niket Tandon, and Bernt Schiele. A dataset for movie description. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3202–3212, 2015.
- [24] Anna Rohrbach, Atousa Torabi, Marcus Rohrbach, Niket Tandon, Christopher Pal, Hugo Larochelle, Aaron Courville, and Bernt Schiele. Movie description. *International Journal of Computer Vision*, 123(1):94–120, 2017.
- [25] Gunnar A. Sigurdsson, Gül Varol, Xiaolong Wang, Ali Farhadi, Ivan Laptev, and Abhinav Gupta. Hollywood in homes: Crowdsourcing data collection for activity understanding. In Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling, editors, *Computer Vision – ECCV 2016*, pages 510–526, Cham, 2016. Springer International Publishing.
- [26] Josef Sivic and Andrew Zisserman. Video google: A text retrieval approach to object matching in videos. In *Computer Vision, IEEE International Conference on*, volume 3, pages 1470–1470. IEEE Computer Society, 2003.
- [27] Cees Snoek, Kvd Sande, OD Rooij, Bouke Huurnink, J Uijlings, M van Liempt, M Bugalhoj, I Trancosoy, F Yan, M

- Tahir, et al. The mediamill trecvid 2009 semantic video search engine. In *TRECVID workshop*. University of Surrey, 2009.
- [28] Mattia Soldan, Mengmeng Xu, Sisi Qu, Jesper Tegner, and Bernard Ghanem. Vlg-net: Video-language graph matching network for video grounding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3224–3234, 2021.
- [29] Ottokar Tilk and Tanel Alumäe. Bidirectional recurrent neural network with attention mechanism for punctuation restoration. In *Interspeech 2016*, 2016.
- [30] Atousa Torabi, Christopher Pal, Hugo Larochelle, and Aaron Courville. Using descriptive video services to create a large data source for video annotation research. *arXiv preprint arXiv:1503.01070*, 2015.
- [31] Takumi Toyama and Daniel Sonntag. Towards episodic memory support for dementia patients by recognizing objects, faces and text in eye gaze. In *Joint German/Austrian Conference on Artificial Intelligence (Künstliche Intelligenz)*, pages 316–323. Springer, 2015.
- [32] Yu Xiong, Fabian Caba Heilbron, and Dahua Lin. Transcript to video: Efficient clip sequencing from texts. *arXiv preprint arXiv:2107.11851*, 2021.
- [33] Mengmeng Xu, Chen Zhao, David S. Rojas, Ali Thabet, and Bernard Ghanem. G-tad: Sub-graph localization for temporal action detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [34] Yitian Yuan, Xiaohan Lan, Long Chen, Wei Liu, Xin Wang, and Wenwu Zhu. A closer look at temporal sentence grounding in videos: Datasets and metrics. *CoRR*, abs/2101.09028, 2021.
- [35] Yitian Yuan, Xiaohan Lan, Long Chen, Wei Liu, Xin Wang, and Wenwu Zhu. A closer look at temporal sentence grounding in videos: Datasets and metrics. *arXiv preprint arXiv:2101.09028*, 2021.
- [36] 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 (CVPR)*, 2020.
- [37] Hao Zhang, Aixin Sun, Wei Jing, and Joey Tianyi Zhou. Towards debiasing temporal sentence grounding in video. *arXiv preprint arXiv:2111.04321*, 2021.
- [38] Zhang Songyang, Peng Houwen, Fu Jianlong, Luo, Jiebo. Learning 2D Temporal Adjacent Networks for Moment Localization with Natural Language. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2020.
- [39] Chen Zhao, Ali K Thabet, and Bernard Ghanem. Video self-stitching graph network for temporal action localization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 13658–13667, 2021.
- [40] Yang Zhao, Zhou Zhao, Zhu Zhang, and Zhijie Lin. Cascaded prediction network via segment tree for temporal video grounding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4197–4206, June 2021.
- [41] Hao Zhou, Chongyang Zhang, Yan Luo, Yanjun Chen, and Chuanping Hu. Embracing uncertainty: Decoupling and de-bias for robust temporal grounding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8445–8454, 2021.