

SoccerNet-Caption: Dense Video Captioning for Soccer Broadcasts Commentaries

Hassan Mkhallati^{*,1} Anthony Cioppa^{*,2,3} Silvio Giancola^{*,3}
Bernard Ghanem³ Marc Van Droogenbroeck²
¹ Université Libre de Bruxelles ² University of Liège ³ KAUST

Abstract

Soccer is more than just a game - it is a passion that transcends borders and unites people worldwide. From the roar of the crowds to the excitement of the commentators, every moment of a soccer match is a thrill. Yet, with so many games happening simultaneously, fans cannot watch them all live. Notifications for main actions can help, but lack the engagement of live commentary, leaving fans feeling disconnected. To fulfill this need, we propose in this paper a novel task of dense video captioning focusing on the generation of textual commentaries anchored with single timestamps. To support this task, we additionally present a challenging dataset consisting of almost 37k timestamped commentaries across 715.9 hours of soccer broadcast videos. Additionally, we propose a first benchmark and baseline for this task, highlighting the difficulty of temporally anchoring commentaries yet showing the capacity to generate meaningful commentaries. By providing broadcasters with a tool to summarize the content of their video with the same level of engagement as a live game, our method could help satisfy the needs of the numerous fans who follow their team but cannot necessarily watch the live game. We believe our method has the potential to enhance the accessibility and understanding of soccer content for a wider audience, bringing the excitement of the game to more people.

1. Introduction

Over the past decade, the quantity and quality of sports data has increased rapidly. This explosion has been driven by the benefits of automated analysis in various applications such as player performance, providing insights into game strategy [68], and audience involvement. Live text commentaries provide a rich summary of the game to increase fan engagement for those who do not have time or the possi-

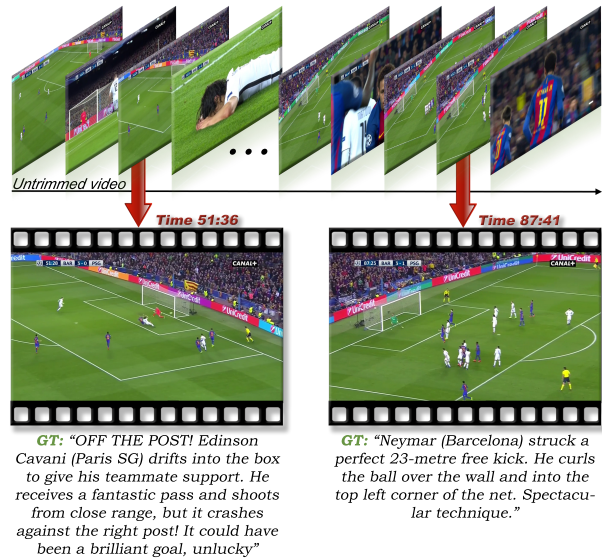


Figure 1. **SoccerNet-Caption.** We provide a large-scale dataset for Single-anchored Dense Video Captioning (SDVC) in untrimmed soccer broadcast videos. Our SoccerNet-Caption dataset is composed of 36,894 textual commentaries, temporally anchored within 715.9 hours of soccer broadcasts. The comments describe the events occurring in the soccer game with rich factual, emotional, and sensational content.

bility to watch the game. However, they are usually exclusive to major professional leagues, while other games are often left out. Moving towards automated solutions relying on already available equipment is therefore essential for lower leagues and amateur soccer. Recently, there has been a growing interest in the research community to automatically generate text based on videos. This task, called Dense Video Captioning (DVC), represents a significant research challenge due to the memory footprint of video data and the complexity of natural language. Besides, most research focuses on generic descriptions of events and activities in open-world scenes. However, in soccer, commentaries need to include rich factual, emotional, and even sensational con-

(*) Equal contributions. Data/code available at www.soccer-net.org.
Contacts: {name.surname}@kaust.edu.sa / soccernet@uliege.be.

tent to engage the fan. Sports is therefore the perfect playground for research in the video and language domain.

In this paper, we publicly release *SoccerNet-Caption*, the first dataset for dense video captioning in soccer broadcast videos. In particular, we provide 36,894 temporally-anchored rich textual commentaries describing 715.9 hours of soccer games. Some examples of comments from our dataset are shown in Figure 1. Along with the data, we introduce the new task of Single-anchored Dense Video Captioning (SDVC), consisting in generating localized captions describing the soccer game. As a first benchmark, we propose a two-stage approach including an action spotting module and a captioning module. More specifically, the spotting module produces temporal proposals for the captions. Then, the videos are trimmed around the proposals and passed to the captioning module to generate captions. We show in Figure 7 that our approach allows to generate relevant captions to describe the game with rich semantics, but that the challenge is still open for major improvements.

Contributions. We summarize our contributions as follows: (i) We publicly release the largest dataset of soccer videos annotated with timestamped textual commentaries describing the game. (ii) We define the novel task of Single-anchored Dense Video Captioning (SDVC), where captions are anchored with a single timestamp and need to be generated in long untrimmed videos. (iii) We propose a first benchmark to tackle this task and provide a thorough ablation study and analysis.

2. Related Works

Sport understanding. The challenging aspect of sports video understanding has contributed to its growing popularity as a research focus [44, 66]. At first, methods focused on video classification [75], including the recognition of specific actions [35, 54], or segmentation of different game phases during the game [11]. More recently, the task of action spotting was introduced by Giancola *et al.* [20], aiming at providing the precise localization of specific actions within an untrimmed soccer broadcast video. Several methods were proposed to automate this process, for instance, using a context-aware loss [9], camera calibration and player localization [10], end-to-end training [29], spatio-temporal encoders [14], graphs-based methods [6], transformer-based methods [83], or anchor-based methods [59, 60]. Other methods focused on other aspects of sports understanding such as player detection [71], player tracking [41] and identification [62, 72], tactics analysis in soccer and fencing [64, 84], pass feasibility [3], 3D ball localization for basketball [69], or 3D shuttle trajectory reconstruction for badminton videos [40].

To support this research, large-scale datasets have been released, including the ones of Pappalardo *et al.* [46], Yu

et al. [79], SoccerTrack [55], SoccerDB [33], and DeepSportRadar [70]. The SoccerNet dataset, introduced by Giancola *et al.* [20], includes benchmarks for 10 different tasks related to soccer understanding, such as action spotting [16], camera calibration [8], and player re-identification [8]. Cioppa *et al.* also introduce the task of player tracking in long sequences, including long-term re-identification [12]. Yearly competitions are organized on this dataset to promote research sports [21]. Our novel SDVC task data is part of the 2023 challenges.

Video-Language Datasets. Initially, research on video combined with language focused on video tagging, including actions and objects [13, 49, 51]. With the successes in image captioning [7, 48] (*i.e.* describing an image with natural language), research has shifted towards deep learning approaches for video captioning. Large-scale multimodal datasets have been introduced thanks to the rise of automatic speech recognition techniques and video-sharing platforms. Youtube has been a major data source for **YouTube-8M** [1], **HowTo100M** [43], and **ViTT** [30]. Other datasets focused on domain-specific videos such as cooking: **Youcook2** [81] and **Tacos** [15], or movies: **MAD** [61]. Further efforts have been proposed for egocentric vision such as **EPIC-KITCHENS-100** [13], and **Ego4D** [23]. On top of those datasets, various tasks have emerged such as video-language grounding [4, 27, 61], video question answering [23, 65], video clip captioning [18, 25, 53, 78] and dense video captioning [37, 81].

Dense Video Captioning. Krishna *et al.* [37] introduced this task, which consists in captioning temporally localized (start and end frame) activities in untrimmed video. This task differs from traditional video captioning [18], where a single caption is generated for short videos, and dense image captioning [34], where captions describe different regions of an image. Currently, **YouCook2** [81], including recipes for non-overlapping sequential events, and **ActivityNet-Captions** [37], including open-domain overlapping activities, are the standard benchmarks for dense video captioning. Our dataset introduces a new task of single-anchored dense video captioning for soccer video comment generation, which requires richer factual, emotional, and sensational comments. Traditionally, the solutions proposed for dense video captioning involve a two-stage “detect-then-describe” framework. The first module produces temporal proposals, and the second module generates captions around these proposals [31, 32, 37, 81]. With the rise of large datasets, many efforts have focused on efficient pre-training of models that combine both language and video [2, 4, 30, 42, 56, 76, 80]. Recent works also tried “YOLO-like” [50] approaches where localization and captioning are generated in one shot [74, 77]. Here, we propose a two-stage approach based on a pre-trained video encoder.

3. Dataset

Data collection. Our SoccerNet-Caption dataset comprises 471 SoccerNet untrimmed broadcast games, including the top five European leagues (EPL, La Liga, Ligue 1, Bundesliga, and Serie A) as well as the Champions League from 2014 to 2017. All videos are available at 25fps in two resolutions: 224p and 720p, alongside frame features at 2fps pre-extracted using *e.g.* ResNET or the Baidu feature encoder [82] at 1fps, following SoccerNet’s original data format. Since the role of the producer is to select the right camera to convey the story of the game to the viewer in the best possible way, these broadcasted games are perfectly suited for a commentary generation task.

In this work, we provide novel textual comments embedded in time describing the game. We collect those comments by scrapping the flashscore website for 471 out of the 500 games included in the SoccerNet dataset. The commentaries for the remaining games were not unavailable. The comments typically describe in a few sentences the main events occurring at specific times in the game, by giving insights about the involved players or teams, and the sequence of actions that led to this situation in a rich factual, emotional, and sensational way. Our data collection efforts resulted in 36,894 timestamped comments across 715.9 hours of video footage, including some metadata such as the type of event the comment relates to (*e.g.* an action, a fun fact, *etc.*).

Alongside these textual comments, we also collected metadata about the game, including the list of all players with their jerseys number, the referees’ name, and the teams’ name. These metadata also includes the starting lineups for each team, with their tactics, the 11 starting players, the substitutes, and any events associated with a player, such as goals, assists, substitutions, and yellow/red cards.

Data anonymization. Following the traditional captioning dataset, we provide an anonymized version of the captions, where each player, referee, team, and coach names are replaced with generic tokens. In fact, most captioning methods are not suited to recognize the exact identity of the people shown in the videos. Hence, without including specific modules for identity classification, player tracking, and re-identification, it would be almost impossible to generate the correct names. However, we still provide the original captions for future research.

To anonymize the dataset, we leverage the game metadata to retrieve the team, coach, referee, and player names. Then, we automatically search through the comments for these specific names and replace them with a generic token ([TEAM], [COACH], [REFEREE], or [PLAYER]). Sometimes, the player and coach names available in the lineups may differ from those mentioned in the comments as some particles or longer names may be truncated. To re-

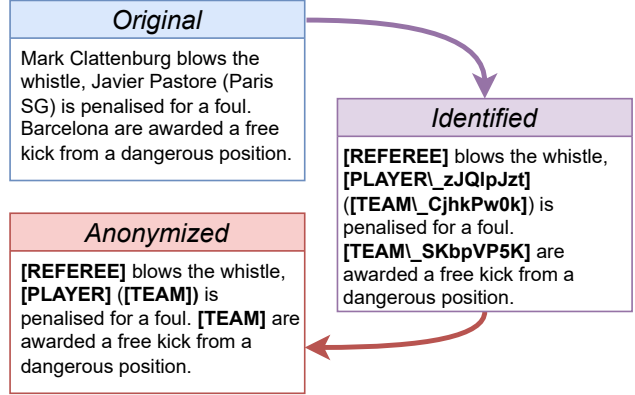


Figure 2. **Comment anonymization.** We provide three versions for each comment. The original commentary, an identified version where each player is associated with a unique id token, and an anonymized version where each entity is replaced by a specific token: [TEAM], [COACH], [REFEREE], and [PLAYER].

trieve those names, we used advanced string-matching techniques to identify and reconcile any discrepancies between the names in the lineups and those mentioned in the comments. Finally, for the remaining names (*e.g.* compound surnames are formulated differently between the lineups and the comments), we manually refined the annotation. This approach ensures that we accurately link each name to the correct token.

Eventually, we provide an intermediate anonymization, where each player is identified with a unique id (uid) inside the token ([Player_uid]). This alleviates the issue of the different names for the same player in the original captions and maintains his identity. Examples of such comments are provided in Figure 2.

Data format. Following the SoccerNet format, we organize our textual annotations into individual JSON files for each game. Each file contains a dictionary that includes all metadata of the game (*e.g.* the starting lineups for each team, with their tactics, the 11 starting players, the substitutes, their jersey numbers, *etc.*) and the list of annotated comments. Each annotation is associated with a timestamp, the three versions of the comment (original description, identified, and anonymized), a boolean value indicating whether the comment is related to a key moment of the game, and a contextual label (*e.g.* corner, substitution, yellow card, whistle, soccer ball, time, injury, fun fact, attendance, penalty, red card, own goal, or missed penalty).

Data statistics. SoccerNet-Caption dataset contains an average of 78.33 temporally localized comments per game, resulting in a total of 36,894 captions for the entire dataset. This is equivalent to almost one comment every minute. As can be seen in Figure 3, the distribution of the comments within a single game over time shows a peak at the start of




Figure 3. **Distribution of the comments.** Most comments are uniformly scattered in each half-time, except at the start of the game where a peak is followed by fewer comments for 10 minutes.




Figure 4. **Distribution of the number of words per comment** This plot shows that the number of words per comment follows a long tail distribution with 21.38 words on average.

the game that usually corresponds to the comment related to the first whistle of the referee. Then, there is a period of fewer comments in the first 10 minutes compared to the rest of the half-time that follows a uniform distribution. This shows that no bias can be used to find a good location for the comments, except at the very start of the game.

Finally, we analyze the content of each comment on a textual and semantic level. Figure 4 shows that the number of words per comment ranges from 4 to 93 words following a long tail distribution with 21.38 words on average. As mentioned, soccer has its own specific terminology for describing events. We can observe in Figure 5 that, apart from the generic tokens, the most commonly used words are soccer action-related verbs (*e.g.* kick, pass, cross) or soccer-related nouns (*e.g.* corner, box, goal). Additionally, it is important to note the over-representation of player and team names in the comments which motivates the use of anonymized comments.




Figure 5. **Distribution of the most common words.** The most frequent words are the names of the teams and the players, followed by words semantically related to soccer verbs and soccer elements. There is a high imbalance in the distribution.

Novelty. We compare our dataset with other recent captioning and dense video captioning datasets in Table 1. Our dataset provides the longest videos on average by a large margin. Processing such long videos in an end-to-end approach is still an open challenge in most video understanding tasks, which makes our dataset the perfect playground to innovate. SoccerNet-Caption also ranks third in terms of total video length, making it a large-scale dataset for video and language training. The other soccer-related dataset either focus on captioning short clips or highlights. SoccerNet-Caption is the first dataset to anchor the comments with a single timestamp instead of a bounding box with a start and end timestamp for each comment. Following the work of SoccerNet on action spotting, we believe that it is arduous to annotate when a specific action starts and ends in soccer. Finally, compared to more generic datasets, soccer commentaries require richer content, including emotional or sensational sentences. By providing the original comments with the game metadata, we aim to push research toward identifying the players in the video to accurately describe what is happening in the game.

4. Single-anchored Dense Video Captioning

Task. We define the novel task of Single-anchored Dense Video Captioning (SDVC) as follows: Given a video, spot all instants where a comment should be anchored and generate sentences describing the events occurring around that time using natural language. This is different from the previously defined Dense Video Captioning (DVC) task [37], where the captions have temporal boundaries (start and end timestamps). For soccer games, this task is particularly challenging, as it requires describing complex sequences of actions involving subtle player movements, rather than well-separated activities with clearly defined boundaries.

Name	Domain	# Video	Avg Duration (s)	Total Duration (hr)	# Sentences	Anchors	Task
ActivityNet-Caption [37]	Open	20k	180	849	100k	$[t_s, t_e]$	DVC
Youcook2 [81]	Cooking	2k	30	176	15.4k	t_s, t_e	DVC
TACoS [15] / TACoS-Multilevel [52]	Cooking	127/185	360	15.9/27.1	18.2k/52.5k	t_s, t_e	Retrieval
Charades-STA [58]	Human	9.8k	30	82.01	27.7k	t_s, t_e	Retrieval
VideoStory [19]	Social Media	20k	70	396	123k	t_s, t_e	Storytelling
VTT [30]	Cooking + open	8.2k	–	–	– (Tag)	t_s, t_e	DVC
Epic Kitchen-100 [13]	Cooking-Ego	700	514	100	–	t_s, t_e	Action Recognition
Ego4D [23]	Ego	9.6k	1369.8	3,670	3.85M	t_s, t_e	Moment Queries
DiDeMo [27]	Open-human	10K	30	88.7	40.5k	t_s, t_e	Retrieval
MAD [61]	Movie	650	6646.2	1207.3	384.6k	t_s, t_e	Grounding
Fine-grained Sports Narrative [78]	Basketball	2K	–	–	6520	–	Video Captioning
Soccer captioning [25]	Soccer action clips	22k clip-action	–	–	22k	–	Video Captioning
GOAL (Qi <i>et al.</i> [47])	Soccer action clips	8.9k	10.31	25.5	22k	–	KGVC
GOAL (Suglia <i>et al.</i> [63])	Soccer highlights	1.1k	238	73.1	53k	–	Video Captioning
SoccerNet-caption (ours)	Soccer games	942	2735.9	715.9	36,894	t	SDVC

Table 1. **Comparison of SoccerNet-Caption with other captioning datasets.** Our dataset contains the second longest video sequences as well as the third longest total video length. This shows that SoccerNet-Caption is a great dataset for research in dense video captioning. Also, it is the first dataset on untrimmed soccer broadcast games, unlike GOAL (Suglia *et al.* [63]) which only focuses on soccer highlights and GOAL (Qi *et al.* [47]) which only focuses on soccer actions for Knowledge-Grounded Video Captioning (KGVC) [47].

Metric. Defining a metric for this task requires evaluating both the temporal accuracy of the detected anchors and the quality of the generated commentaries. Finding the exact timestamp for a commentary is challenging as it depends on the game evolution or certain actions that need to be highlighted, rather than being associated with a specific action like the events defined in the action spotting task. Therefore, we need to include some tolerance around the ground-truth action spot. Additionally, evaluating the quality of generated commentary is not trivial as the expressions used are semantically more related compared to an open vocabulary. Hence, subtle variations in the chosen words need to be accurately evaluated when describing the game.

In the literature, several metrics have been proposed to evaluate dense video captioning methods. The SODA metric [17] evaluates the video narrative by finding the temporally optimal matching between generated and reference captions. Hammoudeh *et al.* [25] proposed another approach that focuses on the precision and recall of generated words with specific expressions defined beforehand and matched with the ground truth. This approach is useful in terms of semantic accuracy in captions but strongly depends on the chosen words dictionary. EMScore [57] focuses on the consistency between the video and candidate captions, relying on an Embedding Matching-based score. However, the EMScore depends on the performance of the used vision-language pre-trained (VLP) model. Since no VLP models have been trained on sports videos yet, we cannot use one of these metrics for our dataset. In order to quantify the spotting quality of the generated captions, we use the $\text{mAP}@ \delta$ for action spotting introduced in Giancola *et al.* [22], where δ is the tolerance in seconds.

For the SDVC task, we first choose the metric proposed in ActivityNet-captions (that still has a consensus within the community) and adapt it for our single-anchored task as follows: for each ground-truth caption in a video, we

build a time window with a chosen tolerance centered on its timestamp. We then use established captioning evaluation metrics (METEOR [38], BLEU [45], ROUGE [39], and CIDEr [73]) to estimate the language similarity between all generated captions with any ground-truth caption for which its timestamps fall within a δ tolerance. The performances are finally averaged over the video and the dataset. We call our metric $\text{METEOR}@ \delta$ (resp. $\text{BLEU}@ \delta$, $\text{ROUGE}@ \delta$, and $\text{CIDEr}@ \delta$). As a second metric, we similarly adapt the SODA_c [17] metric by adding a time window around the ground-truth and generated captions.

5. Benchmarks

In this section, we benchmark a first baseline model on our task of Single-anchored Dense Video Captioning (SDVC). Particularly, we study the performance of several architectures and hyperparameters of our model on the spotting, captioning, and global SDVC task. We finally provide qualitative results on real sequences.

5.1. SDVC baseline

Following the literature on dense video captioning, we propose a two-stage approach [31, 32, 37, 81] as an initial baseline model for our SDVC task. Our model, denoted by M , consists of a spotting model and a captioning model, which are cascaded together. The sub-models are trained independently and both consist of a frozen feature encoder E , followed by an aggregator module A , and either a spotting head S or a captioning head C . The feature encoder generates a compressed per-frame feature representation of the video clip, which is then temporally pooled by the aggregator. The resulting clip feature representation is subsequently passed either to the spotting head (locating where to generate the comments) or to the captioning head (generating a comment). During inference, the spotting model generates

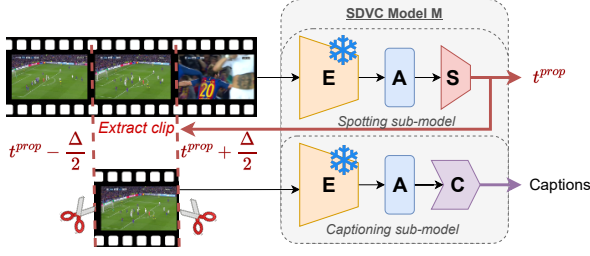


Figure 6. **Pipeline of our Single-anchored Dense Video Captioning (SDVC) baseline model M.** To generate dense comments with a single timestamp, we propose a two-stage approach. **M** consists of a spotting model followed by a captioning model. Both models use a shared frozen feature extractor **E** to generate a compact per-frame representation of the video. The spotting model uses an aggregator module **A** to combine the frame features into a single clip feature representation that is then passed to the spotting head **S** to generate proposal timestamps t^{prop} . The timestamps are then used to trim clips of size Δ that are subsequently processed by the captioning model through **E**, **A**, and a captioning head **C** to generate the anchored comment.

a series of temporal proposals for a video:

$$\{t_0^{prop}, \dots, t_N^{prop}\} = \mathbf{S}(\mathbf{A}(\mathbf{E}(\text{video}))),$$

where N is the total number of proposals. The original video is then trimmed around each proposal t_n^{prop} into a clip that is passed to the captioning model to generate a caption, following:

$$\text{caption}_n = \mathbf{C}(\mathbf{A}(\mathbf{E}(\text{video}[t_n^{prop} - \frac{\Delta}{2}, t_n^{prop} + \frac{\Delta}{2}]))) ,$$

where Δ is the window size of the captioning model. Our baseline model is illustrated in Figure 6.

Feature encoder E. We use the feature encoders provided in SoccerNet-v2, *i.e.* ResNet-152 [26], I3D [5], C3D [67], and Baidu [82]. The features are extracted at 1 or 2 fps from the original soccer broadcast videos. To speed up the training, we reduce the feature dimensionality to 512 for each image using PCA for the first three encoders, and a linear transformation for the Baidu features as suggested in [22].

Aggregator A. The aggregator module pools the frame feature vectors into one single compact feature representation of the clip. For our baseline, we use four trainable pooling modules proposed by Giancola *et al.* [22]: NetVLAD, NetRVLAD, and the temporally aware pooling modules NetVLAD++ and NetRVLAD++.

Spotting head S. We build our spotting head as a dense layer with sigmoid activation that outputs 2 classes: the presence of comment (foreground), and absence of comment (background). During training, the video is randomly cropped into video chunks, and a binary cross-entropy loss function is applied to the output. During inference, we split

the whole half-time video into overlapping clips and concatenate the predictions over time. We then use a Non-Maximum Suppression (NMS) algorithm to reduce redundant spots in a specific time window.

Captioning head C. Our captioning head is composed of two fully connected layers with ReLU activation and dropout, and a vanilla LSTM [28] module with softmax activation. The fully connected layers project the output of the aggregator from the video hidden space to the language hidden space. The projected features are then used to initialize the hidden state of the LSTM module that outputs the list of word confidence scores. During training, we use the cross-entropy loss between the predicted and ground-truth words. To stabilize the learning process, we use the teacher forcing method presented by Graves *et al.* [24]. As the soccer vocabulary is much more specific than generic language, we do not use pre-trained word embeddings, but learn the whole language during the training. During inference, we sample words with a greedy approach, *i.e.* the next word is the one with the highest confidence.

5.2. Training parameters

To train and evaluate our model, we use SoccerNet-Caption without the “fun fact” and “attendance” commentaries as they describe out-of-the-game content. During training, we use the Adam optimizer [36] with PyTorch’s default β parameters. We reduce the learning rate by a factor of 10 when the validation loss plateaus during 10 consecutive epochs, with an initial value of 10^{-3} , and a stopping criterion of 10^{-6} . The temporal pooling module is initialized with 64 clusters and the dimension of the hidden vector of the LSTM is 512, the dropout value after the fully connected layers is set to 0.4, and the word embedding dimension is set to 256. Finally, the size of the vocabulary in SoccerNet-Caption is 1,769 words.

5.3. Results

We now first separately study the performance of our spotting model and the captioning module, then provide the performance of the whole pipeline for our new SDVC task.

Commentary spotting results. We first study the use of different combinations of feature encoders and aggregators, by fixing both the size of the input video chunk and the NMS window to 15 seconds. The results are presented in Table 2. As can be seen, the best results are obtained using the Baidu feature encoder and NetVLAD or NetVLAD++ pooling. This suggests that using a feature extractor fine-tuned on soccer data leads to better performance even for commentary spotting. However, there is not such a clear tendency regarding the aggregator module layer, with a slight advantage to NetVLAD or NetVLAD++. Hence, for the following spotting model, we use the Baidu feature encoder with NetVLAD as it provides the best mAP@5.

Encoder	Aggregator	mAP@ (%)		
		5	30	60
RN_PCA	NetVLAD	6.4	39.1	38.0
RN_PCA	NetVLAD++	5.3	41.3	39.4
RN_PCA	NetRVLAD	7.0	39.5	37.8
RN_PCA	NetRVLAD++	5.4	41.5	39.4
I3D_PCA	NetVLAD	4.5	33.0	33.3
I3D_PCA	NetVLAD++	6.7	34.8	34.8
I3D_PCA	NetRVLAD	4.1	32.8	32.3
I3D_PCA	NetRVLAD++	5.9	34.5	34.1
C3D_PCA	NetVLAD	4.8	38.1	37.1
C3D_PCA	NetVLAD++	4.5	40.7	39.1
C3D_PCA	NetRVLAD	6.9	39.3	38.1
C3D_PCA	NetRVLAD++	3.9	39.2	38.2
Baidu	NetVLAD	10.5	42.1	40.7
Baidu	NetVLAD++	6.3	44.5	41.8
Baidu	NetRVLAD	6.7	39.5	38.7
Baidu	NetRVLAD++	3.6	44.0	41.2

Table 2. **Spotting results.** We train our spotting model to detect and localize comments and compare different combinations of encoder and pooling modules.

WS (s)	NMS (s)	mAP@ (%)		
		5	30	60
15	10	12.5	35.3	34.9
15	30	9.3	49.4	47.0
15	60	8.5	46.2	41.2
30	10	5.3	29.2	27.7
30	30	2.5	45.7	43.0
30	60	1.7	46.5	40.2
45	10	6.8	27.0	26.1
45	30	5.2	47.8	41.7
45	60	2.6	45.4	39.8
60	10	4.9	19.0	19.5
60	30	1.9	38.0	35.2
60	60	1.1	39.5	35.6

Table 3. **Spotting window size and NMS.** We compare several window and NMS sizes. Small windows achieve better results.

Next, we study the influence of the window and NMS size on the quality of the localization. As shown in Table 3, the best performance is obtained with a window of 15 seconds and a NMS of 10 or 30 seconds. As the performance (mAP@30 and mAP@60) largely decreases for a NMS window of 10 seconds, we choose 15 and 30 seconds for the window and NMS size of the following experiments. Our analysis shows that the commentary spotting performances are significantly lower than action spotting on SoccerNet-v2 [21], as a comment may describe several actions, unlike spotting that focuses on a single action.

Captioning results. We conduct similar ablations on the captioning task. As shown in Table 4, the best results are also achieved using the Baidu feature encoder and NetVLAD pooling. This suggests again that pre-training

Encoder	Aggregator	B@4	M	R	C
RN_PCA	NetVLAD	4.0	21.8	21.5	10.8
RN_PCA	NetVLAD++	4.4	22.3	21.4	10.9
RN_PCA	NetRVLAD	4.2	21.9	21.4	10.6
RN_PCA	NetRVLAD++	4.2	21.8	21.2	11.0
I3D_PCA	NetVLAD	3.2	21.9	20.4	7.5
I3D_PCA	NetVLAD++	2.9	20.0	18.3	6.6
I3D_PCA	NetRVLAD	3.1	21.1	20.7	8.0
I3D_PCA	NetRVLAD++	3.7	21.7	20.8	9.2
C3D_PCA	NetVLAD	3.9	21.9	21.4	10.3
C3D_PCA	NetVLAD++	3.9	21.2	20.4	10.6
C3D_PCA	NetRVLAD	4.0	21.6	21.1	11.0
C3D_PCA	NetRVLAD++	3.9	21.3	20.7	11.0
Baidu	NetVLAD	5.7	23.6	23.5	15.7
Baidu	NetVLAD++	5.4	23.2	23.4	17.0
Baidu	NetRVLAD	5.5	23.5	23.2	15.4
Baidu	NetRVLAD++	5.7	23.5	23.3	16.0

Table 4. **Captioning results.** We train our captioning model to generate comments and compare different combinations of encoder and pooling modules with the Bleu (B), METEOR(M), ROUGE-L (R), and CIDEr (C) metrics.

Window (s)	B@4	M	R	C
15	5.4	23.5	23.2	15.2
30	5.9	23.7	23.7	17.8
45	5.9	23.6	23.9	18.2
60	5.6	23.0	23.6	17.5

Table 5. **Captioning window size.** We compare several window sizes with the Bleu (B), METEOR (M), ROUGE-L (R), and CIDEr (C) metrics.

L	TF ratio	B@4	M	R	C
2	0.5	2.1	21.5	17.5	5.5
2	1	6.0	23.7	23.8	17.8
4	0.5	4.1	22.4	21.7	8.8
4	1	6.0	23.7	24.1	18.5
8	0.5	1.3	18.1	15.8	3.4
8	1	2.5	20.7	21.6	7.3

Table 6. **Captioning ablation.** We compare several numbers of LSTM layers (L) and teacher forcing (TF) ratio with the Bleu (B), METEOR (M), ROUGE-L (R), and CIDEr (C) metrics.

a video encoder on the spotting task helps generating better commentaries. In Table 5, we also study the effect of the window size on the captioning model. Unlike the commentary spotting task, the best results are obtained with a window of 45 seconds, showing that captioning requires more contextual information. For the following experiments, we use 45-seconds video clips as input to the captioning model.

As a final ablation study, we compare the performance of our captioning model when changing the number of stacked LSTM layers or the teacher forcing ratio. The teacher-forcing ratio is the probability of the LSTM module receiv-

Aggregator	Spotting (mAP@ (%))			Captioning				Single-anchored dense video captioning				
	5	30	60	B@4	M	R	C	B@4@30	M@30	R@30	C@30	SODA_c
$A_{scratch}$	9.24	49.40	46.97	6.37	23.85	24.04	18.70	21.07	27.94	22.06	27.02	7.72
$A_{frozen}^{spot. \rightarrow capt.}$	2.99	49.06	46.32	6.04	23.68	24.05	18.53	21.10	27.24	21.99	26.96	7.71
$A_{fine-tuned}^{spot. \rightarrow capt.}$	4.67	49.61	46.97	6.18	23.71	24.03	19.00	21.03	28.14	22.03	27.01	7.76
$A_{frozen}^{capt. \rightarrow spot.}$	9.23	49.49	46.97	5.50	23.03	23.11	17.20	21.52	24.99	21.46	26.63	7.44
$A_{fine-tuned}^{capt. \rightarrow spot.}$	9.25	49.40	46.97	6.62	23.84	24.53	21.45	21.63	29.44	22.09	27.20	7.79

Table 7. **SDVC Results.** We train SDVC models using Baidu encoder, NetVLAD as pooling layer, a window size of 15 seconds for the spotting model, a NMS of 30 seconds, a window size of 45 seconds for the captioning module, four LSTM layers and with teacher forcing. We study different pre-trainings for the feature aggregators where the spotting and captioning models are trained separately from scratch, or the model is trained first on spotting (resp. captioning), and then weights are transferred to captioning (resp. spotting). The captioning (resp. spotting) aggregator is then either fine-tuned or frozen. We report the spotting mAP @5/30/60, the captioning Bleu (B), METEOR (M), ROUGE-L (R), and CIDEr (C) with the corresponding single-anchored dense video captioning metrics @30, and the SODA metric.



Figure 7. **Qualitative results.** (a) Our baseline is able to provide accurate captions for some actions with rich vocabulary, but (b) sometimes generates captions on unlabeled events. (c) The scores are almost never correctly generated since the model generates captions based on clips without long temporal context and (d) sometimes generates completely wrong captions.

ing the true output sequence rather than its own prediction. For example, a teacher-forcing ratio of 0.5 means that half of the time, the LSTM receives the true output sequence, and half of the time, it receives its own prediction. As shown in Table 6, the best performance is achieved with 4 LSTM layers and using the teacher forcing technique.

Single-anchored dense video captioning results. Finally, we evaluate the performance of our whole pipeline. Particularly, we study the influence of pre-training the aggregator in three ways: (1) no-pre-training for the spotting and captioning aggregator, (2) training the spotting model and transferring the weights to the captioning model, and (3) training the captioning model and transferring the weight to the spotting model. We also either freeze or fine-tune the

transferred weights on the second task. As can be seen from Table 7, the best performance for our SDVC task is obtained when training the captioning model from scratch, transferring the aggregator weights to the spotting aggregator, and fine-tuning those weights on the spotting task.

Qualitative Results. We provide four predictions of our SDVC baseline in Figure 7 and compare them with the ground truth: (a) Our method is able to generate good commentaries for some actions. (b) Our spotting model shows a tendency to generate proposals not close to any commentary, yet our captioning model still describes the ongoing action perfectly. This is also shown in the difference between the METEOR@30 and SODA_c performances in Table 7. However, in a real-world application, such captions would add value rather than be considered a mistake. (c) When generating a caption, our pipeline considers only the short time window around the proposal, hence the generated scores after a goal are almost never accurate. (d) We show a case of hard failure, where our model confuses a serious injury situation while the team was actually celebrating a goal. These results show that our baseline is already able to generate accurate captions, but that there is room for improvement, especially in gathering temporal context.

6. Conclusion

This paper proposes the novel task of single-anchored dense video captioning focusing on generating textual commentaries anchored with single timestamps. To support this task, we present SoccerNet-Caption, a challenging dataset consisting of 37k timestamped commentaries across 715.9 hours of soccer broadcast videos. We benchmarked a first baseline algorithm on this dataset, highlighting the difficulty of temporally anchoring commentaries yet showing the capacity to generate meaningful commentaries.

Acknowledgement. This work was partly supported by KAUST OSR through the VCC funding and the SDAIA-KAUST Center of Excellence in Data Science and Artificial Intelligence. A. Cioppa is funded by the F.R.S.-FNRS.

References

- [1] Sami Abu-El-Haija, Nisarg Kothari, Joonseok Lee, Paul Natsev, George Toderici, Balakrishnan Varadarajan, and Sudheendra Vijayanarasimhan. YouTube-8M: A large-scale video classification benchmark. *CoRR*, abs/1609.08675, 2016. 2
- [2] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katie Millican, Malcolm Reynolds, Roman Ring, Eliza Rutherford, Serkan Cabi, Tengda Han, Zhitao Gong, Sina Samangooei, Marianne Monteiro, Jacob Menick, Sebastian Borgeaud, Andrew Brock, Aida Nematzadeh, Sahand Sharifzadeh, Mikolaj Binkowski, Ricardo Barreira, Oriol Vinyals, Andrew Zisserman, and Karen Simonyan. Flamingo: a visual language model for few-shot learning. *CoRR*, abs/2204.14198, 2022. 2
- [3] Adrià Arbués Sangüesa, Adrià Martín, Javier Fernández, Coloma Ballester, and Gloria Haro. Using player’s body-orientation to model pass feasibility in soccer. In *IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Work. (CVPRW)*, pages 3875–3884, Seattle, WA, USA, Jun. 2020. Inst. Electr. Electron. Eng. (IEEE). 2
- [4] Max Bain, Arsha Nagrani, Gul Varol, and Andrew Zisserman. Frozen in time: A joint video and image encoder for end-to-end retrieval. In *IEEE Int. Conf. Comput. Vis. (ICCV)*, pages 1708–1718, Montreal, QC, Canada, Oct. 2021. Inst. Electr. Electron. Eng. (IEEE). 2
- [5] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. *CoRR*, abs/1705.07750, 2017. 6
- [6] Alejandro Cartas, Coloma Ballester, and Gloria Haro. A graph-based method for soccer action spotting using unsupervised player classification. In *Int. ACM Work. Multimedia Content Anal. Sports (MMSports)*, pages 93–102, Lisbon, Port., Oct. 2022. ACM. 2
- [7] Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. UNITER: UNiversal image-TExt representation learning. In *Eur. Conf. Comput. Vis. (ECCV)*, volume 12375 of *Lect. Notes Comput. Sci.*, pages 104–120. Springer Int. Publ., 2020. 2
- [8] Anthony Cioppa, Adrien Delière, Silvio Giancola, Bernard Ghanem, and Marc Van Droogenbroeck. Scaling up SoccerNet with multi-view spatial localization and re-identification. *Sci. Data*, 9(1):1–9, Jun. 2022. 2
- [9] Anthony Cioppa, Adrien Delière, Silvio Giancola, Bernard Ghanem, Marc Van Droogenbroeck, Rikke Gade, and Thomas B. Moeslund. A context-aware loss function for action spotting in soccer videos. In *IEEE Int. Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pages 13123–13133, Seattle, WA, USA, Jun. 2020. 2
- [10] Anthony Cioppa, Adrien Delière, Silvio Giancola, Florian Magera, Olivier Barnich, Bernard Ghanem, and Marc Van Droogenbroeck. Camera calibration and player localization in SoccerNet-v2 and investigation of their representations for action spotting. In *IEEE Int. Conf. Comput. Vis. Pattern Recognit. Work. (CVPRW)*, *CVSports*, pages 4532–4541, Nashville, TN, USA, Jun. 2021. 2
- [11] Anthony Cioppa, Adrien Delière, and Marc Van Droogenbroeck. A bottom-up approach based on semantics for the interpretation of the main camera stream in soccer games. In *IEEE Int. Conf. Comput. Vis. Pattern Recognit. Work. (CVPRW)*, *CVSports*, pages 1846–1855, Salt Lake City, UT, USA, Jun. 2018. 2
- [12] Anthony Cioppa, Silvio Giancola, Adrien Deliege, Le Kang, Xin Zhou, Zhiyu Cheng, Bernard Ghanem, and Marc Van Droogenbroeck. SoccerNet-tracking: Multiple object tracking dataset and benchmark in soccer videos. In *IEEE Int. Conf. Comput. Vis. Pattern Recognit. Work. (CVPRW)*, *CVSports*, pages 3490–3501, New Orleans, LA, USA, Jun. 2022. Inst. Electr. Electron. Eng. (IEEE). 2
- [13] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Antonino Furnari, Evangelos Kazakos, Jian Ma, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, and Michael Wray. Rescaling egocentric vision: Collection, pipeline and challenges for EPIC-KITCHENS-100. *Int. J. Comput. Vis.*, 130(1):33–55, Oct. 2021. 2, 5
- [14] Abdulrahman Darwish and Tallel El-Shabrway. STE: Spatio-temporal encoder for action spotting in soccer videos. In *Int. ACM Work. Multimedia Content Anal. Sports (MMSports)*, pages 87–92, Lisbon, Port., Oct. 2022. ACM. 2
- [15] Pradipto Das, Chenliang Xu, Richard F. Doell, and Jason J. Corso. A thousand frames in just a few words: Lingual description of videos through latent topics and sparse object stitching. In *IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*. Inst. Electr. Electron. Eng. (IEEE), Jun. 2013. 2, 5
- [16] Adrien Delière, Anthony Cioppa, Silvio Giancola, Meisam J. Seikavandi, Jacob V. Dueholm, Kamal Nasrollahi, Bernard Ghanem, Thomas B. Moeslund, and Marc Van Droogenbroeck. SoccerNet-v2: A dataset and benchmarks for holistic understanding of broadcast soccer videos. In *IEEE Int. Conf. Comput. Vis. Pattern Recognit. Work. (CVPRW)*, *CVSports*, pages 4508–4519, Nashville, TN, USA, Jun. 2021. Best CVSports paper award. 2
- [17] Soichiro Fujita, Tsutomu Hirao, Hidetaka Kamigaito, Manabu Okumura, and Masaaki Nagata. SODA: Story oriented dense video captioning evaluation framework. In *Eur. Conf. Comput. Vis. (ECCV)*, volume 12351 of *Lect. Notes Comput. Sci.*, pages 517–531. Springer Int. Publ., 2020. 5
- [18] Lianli Gao, Zhao Guo, Hanwang Zhang, Xing Xu, and Heng Tao Shen. Video captioning with attention-based LSTM and semantic consistency. *IEEE Trans. Multimedia*, 19(9):2045–2055, Sept. 2017. 2
- [19] Spandana Gella, Mike Lewis, and Marcus Rohrbach. A dataset for telling the stories of social media videos. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 2018. 5
- [20] Silvio Giancola, Mohieddine Amine, Tarek Dghaily, and Bernard Ghanem. SoccerNet: A scalable dataset for action spotting in soccer videos. In *IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Work. (CVPRW)*, pages 1792–179210, Salt

- Lake City, UT, USA, Jun. 2018. Inst. Electr. Electron. Eng. (IEEE). 2
- [21] Silvio Giancola, Anthony Cioppa, Adrien Delière, Floriane Magera, Vladimir Somers, Le Kang, Xin Zhou, Olivier Barnich, Christophe De Vleeschouwer, Alexandre Alahi, Bernard Ghanem, Marc Van Droogenbroeck, Abdulrahman Darwish, Adrien Maglo, Albert Clapés, Andreas Luyts, Andrei Boiarov, Artur Xarles, Astrid Orcesi, Avijit Shah, Baoyu Fan, Bharath Comandur, Chen Chen, Chen Zhang, Chen Zhao, Chengzhi Lin, Cheuk-Yiu Chan, Chun Chuen Hui, Dengjie Li, Fan Yang, Fan Liang, Fang Da, Feng Yan, Fufu Yu, Guanshuo Wang, H. Anthony Chan, He Zhu, Hongwei Kan, Jiaming Chu, Jianming Hu, Jianyang Gu, Jin Chen, João V. B. Soares, Jonas Theiner, Jorge De Corte, José Henrique Brito, Jun Zhang, Junjie Li, Junwei Liang, Leqi Shen, Lin Ma, Lingchi Chen, Miguel Santos Marques, Mike Azatov, Nikita Kasatkin, Ning Wang, Qiong Jia, Quoc Cuong Pham, Ralph Ewerth, Ran Song, Rengang Li, Rikke Gade, Ruben Debieen, Runze Zhang, Sangrok Lee, Sergio Escalera, Shan Jiang, Shigeyuki Odashima, Shimin Chen, Shoichi Masui, Shouhong Ding, Sin-wai Chan, Siyu Chen, Tallal El-Shabrawy, Tao He, Thomas B. Moeslund, Wan-Chi Siu, Wei Zhang, Wei Li, Xiangwei Wang, Xiao Tan, Xiaochuan Li, Xiaolin Wei, Xiaoqing Ye, Xing Liu, Xinying Wang, Yandong Guo, Yaqian Zhao, Yi Yu, Yingying Li, Yue He, Yujie Zhong, Zhenhua Guo, and Zhiheng Li. SoccerNet 2022 challenges results. In *Int. ACM Work. Multimedia Content Anal. Sports (MMSports)*, pages 75–86, Lisbon, Port., Oct. 2022. ACM. 2, 7
- [22] Silvio Giancola and Bernard Ghanem. Temporally-aware feature pooling for action spotting in soccer broadcasts. In *IEEE Int. Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pages 4490–4499, Nashville, TN, USA, Jun. 2021. 5, 6
- [23] Kristen Grauman, Andrew Westbury, Eugene Byrne, Zachary Chavis, Antonino Furnari, Rohit Girdhar, Jackson Hamburger, Hao Jiang, Miao Liu, Xingyu Liu, Miguel Martin, Tushar Nagarajan, Ilija Radosavovic, Santhosh Kumar Ramakrishnan, Fiona Ryan, Jayant Sharma, Michael Wray, Mengmeng Xu, Eric Zhongcong Xu, Chen Zhao, Siddhant Bansal, Dhruv Batra, Vincent Cartillier, Sean Crane, Tien Do, Morrie Doulaty, Akshay Erapalli, Christoph Feichtenhofer, Adriano Fragomeni, Qichen Fu, Abraham Gebreelasie, Cristina Gonzalez, James Hillis, Xuhua Huang, Yifei Huang, Wenqi Jia, Weslie Khoo, Jachym Kolar, Satwik Kotur, Anurag Kumar, Federico Landini, Chao Li, Yanghao Li, Zhenqiang Li, Karttikeya Mangalam, Raghava Modhugu, Jonathan Munro, Tullie Murrell, Takumi Nishiyasu, Will Price, Paola Ruiz Puentes, Merey Ramazanova, Leda Sari, Kiran Somasundaram, Audrey Southerland, Yusuke Sugano, Ruijie Tao, Minh Vo, Yuchen Wang, Xindi Wu, Takuma Yagi, Ziwei Zhao, Yunyi Zhu, Pablo Arbelaez, David Crandall, Dima Damen, Giovanni Maria Farinella, Christian Fuegen, Bernard Ghanem, Vamsi Krishna Ithapu, C. V. Jawahar, Hanbyul Joo, Kris Kitani, Haizhou Li, Richard Newcombe, Aude Oliva, Hyun Soo Park, James M. Rehg, Yoichi Sato, Jianbo Shi, Mike Zheng Shou, Antonio Torralba, Lorenzo Torresani, Mingfei Yan, and Jitendra Malik. Ego4d: Around the world in 3,000 hours of egocentric video. In *IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*. Inst. Electr. Electron. Eng. (IEEE), Jun. 2022. 2, 5
- [24] Alex Graves. Supervised sequence labelling with recurrent neural networks. *Stud. Comput. Intell.*, 2012. 6
- [25] Ahmad Hammoudeh, Bastien Vanderplaetse, and Stéphane Dupont. Soccer captioning: dataset, transformer-based model, and triple-level evaluation. *Procedia Computer Science*, 210:104–111, 2022. 2, 5
- [26] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*. Inst. Electr. Electron. Eng. (IEEE), Jun. 2016. 6
- [27] Lisa Anne Hendricks, Oliver Wang, Eli Shechtman, Josef Sivic, Trevor Darrell, and Bryan Russell. Localizing moments in video with natural language. In *IEEE Int. Conf. Comput. Vis. (ICCV)*. Inst. Electr. Electron. Eng. (IEEE), Oct. 2017. 2, 5
- [28] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, Nov. 1997. 6
- [29] James Hong, Haotian Zhang, Michaël Gharbi, Matthew Fisher, and Kayvon Fatahalian. Spotting temporally precise, fine-grained events in video. *CoRR*, abs/2207.10213, 2022. 2
- [30] Gabriel Huang, Bo Pang, Zhenhai Zhu, Clara Rivera, and Radu Soricut. Multimodal pretraining for dense video captioning. *CoRR*, abs/2011.11760, 2020. 2, 5
- [31] Vladimir Iashin and Esa Rahtu. A better use of audio-visual cues: Dense video captioning with bi-modal transformer. *CoRR*, abs/2005.08271, 2020. 2, 5
- [32] Vladimir Iashin and Esa Rahtu. Multi-modal dense video captioning. *CoRR*, abs/2003.07758, 2020. 2, 5
- [33] Yudong Jiang, Kaixu Cui, Leilei Chen, Canjin Wang, and Changliang Xu. SoccerDB: A large-scale database for comprehensive video understanding. In *Int. ACM Work. Multimedia Content Anal. Sports (MMSports)*, page 1–8. ACM, Oct. 2020. 2
- [34] Justin Johnson, Andrej Karpathy, and Li Fei-Fei. DenseCap: Fully convolutional localization networks for dense captioning. In *IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*. Inst. Electr. Electron. Eng. (IEEE), Jun. 2016. 2
- [35] Muhammad Zeeshan Khan, Summra Saleem, Muhammad A. Hassan, and Muhammad Usman Ghanni Khan. Learning deep C3D features for soccer video event detection. In *Int. Conf. Emerg. Technol. (ICET)*, pages 1–6, Islamabad, Pakistan, Nov. 2018. 2
- [36] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2014. 6
- [37] Ranjay Krishna, Kenji Hata, Frederic Ren, Li Fei-Fei, and Juan Carlos Nieves. Dense-captioning events in videos. *CoRR*, abs/1705.00754, 2017. 2, 4, 5
- [38] Alon Lavie and Abhaya Agarwal. Meteor. In *Proceedings of the Second Workshop on Statistical Machine Translation - StatMT '07*. Association for Computational Linguistics, 2007. 5
- [39] Chin-Yew Lin. ROUGE: A package for automatic evaluation of summaries. In *Annual Meeting of the Association for Computational Linguistics*, pages 74–81, Jul. 2004. 5

- [40] Paul Liu and Jui-Hsien Wang. MonoTrack: Shuttle trajectory reconstruction from monocular badminton video. In *IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Work. (CVPRW)*, pages 3512–3521, New Orleans, LA, USA, Jun. 2022. Inst. Electr. Electron. Eng. (IEEE). 2
- [41] Adrien Maglo, Astrid Orcesi, and Quoc-Cuong Pham. Efficient tracking of team sport players with few game-specific annotations. In *IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Work. (CVPRW)*, pages 3460–3470, New Orleans, LA, USA, Jun. 2022. Inst. Electr. Electron. Eng. (IEEE). 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 *IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*. Inst. Electr. Electron. Eng. (IEEE), Jun. 2020. 2
- [43] Antoine Miech, Dimitri Zhukov, Jean-Baptiste Alayrac, Makarand Tapaswi, Ivan Laptev, and Josef Sivic. HowTo100M: Learning a text-video embedding by watching hundred million narrated video clips. In *IEEE Int. Conf. Comput. Vis. (ICCV)*. Inst. Electr. Electron. Eng. (IEEE), Oct. 2019. 2
- [44] Thomas B. Moeslund, Graham Thomas, and Adrian Hilton. *Computer vision in sports*. Springer, 2014. 2
- [45] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. BLEU. *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics - ACL '02*, 2001. 5
- [46] Luca Pappalardo, Paolo Cintia, Alessio Rossi, Emanuele Massucco, Paolo Ferragina, Dino Pedreschi, and Fosca Giannotti. A public data set of spatio-temporal match events in soccer competitions. *Scientific Data*, 6(1):236, Oct 2019. 2
- [47] Ji Qi, Jifan Yu, Teng Tu, Kunyu Gao, Yifan Xu, Xinyu Guan, Xiaozhi Wang, Yuxiao Dong, Bin Xu, Lei Hou, and et al. GOAL: A challenging knowledge-grounded video captioning benchmark for real-time soccer commentary generation. *CoRR*, abs/2303.14655, 2023. 5
- [48] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In Marina Meila and Tong Zhang, editors, *Int. Conf. Mach. Learn. (ICML)*, volume 139, pages 8748–8763. PMLR, Jul. 2021. 2
- [49] Muhammad Rafiq, Ghazala Rafiq, and Gyu Sang Choi. Video description: Datasets & evaluation metrics. *IEEE Access*, 9:121665–121685, 2021. 2
- [50] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*. Inst. Electr. Electron. Eng. (IEEE), Jun. 2016. 2
- [51] Michaela Regneri, Marcus Rohrbach, Dominikus Wetzel, Stefan Thater, Bernt Schiele, and Manfred Pinkal. Grounding action descriptions in videos. *Trans. Assoc. For Comput. Linguistics*, 1:25–36, Dec. 2013. 2
- [52] Anna Rohrbach, Marcus Rohrbach, Wei Qiu, Annemarie Friedrich, Manfred Pinkal, and Bernt Schiele. Coherent multi-sentence video description with variable level of detail. In *Pattern Recognit.*, volume 8753 of *Lect. Notes Comput. Sci.*, pages 184–195. Springer Int. Publ., 2014. 5
- [53] Anna Rohrbach, Marcus Rohrbach, Niket Tandon, and Bernt Schiele. A dataset for movie description. In *IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*. Inst. Electr. Electron. Eng. (IEEE), Jun. 2015. 2
- [54] Himangi Saraogi, Rahul Anand Sharma, and Vijay Kumar. Event recognition in broadcast soccer videos. In *Indian Conference on Computer Vision, Graphics and Image Processing*, pages 1–7, Dec. 2016. 2
- [55] Atom Scott, Ikuma Uchida, Masaki Onishi, Yoshinari Kameda, Kazuhiro Fukui, and Keisuke Fujii. SoccerTrack: A dataset and tracking algorithm for soccer with fish-eye and drone videos. In *IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Work. (CVPRW)*, pages 3568–3578, New Orleans, LA, USA, Jun. 2022. Inst. Electr. Electron. Eng. (IEEE). 2
- [56] Paul Hongsuck Seo, Arsha Nagrai, Anurag Arnab, and Cordelia Schmid. End-to-end generative pretraining for multimodal video captioning. In *IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*. Inst. Electr. Electron. Eng. (IEEE), Jun. 2022. 2
- [57] Yaya Shi, Xu Yang, Haiyang Xu, Chunfeng Yuan, Bing Li, Weiming Hu, and Zheng-Jun Zha. EMScore: Evaluating video captioning via coarse-grained and fine-grained embedding matching. In *IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*. Inst. Electr. Electron. Eng. (IEEE), Jun. 2022. 5
- [58] Gunnar A. Sigurdsson, Gül Varol, Xiaolong Wang, Ali Farhadi, Ivan Laptev, and Abhinav Gupta. Hollywood in homes: Crowdsourcing data collection for activity understanding. In *Eur. Conf. Comput. Vis. (ECCV)*, volume 9905 of *Lect. Notes Comput. Sci.*, pages 510–526. Springer Int. Publ., 2016. 5
- [59] João V. B. Soares and Avijit Shah. Action spotting using dense detection anchors revisited: Submission to the SoccerNet challenge 2022. *CoRR*, abs/2206.07846, 2022. 2
- [60] Joao V. B. Soares, Avijit Shah, and Topojoy Biswas. Temporally precise action spotting in soccer videos using dense detection anchors. In *IEEE Int. Conf. Image Process. (ICIP)*, pages 2796–2800, Bordeaux, France, Oct. 2022. Inst. Electr. Electron. Eng. (IEEE). 2
- [61] Mattia Soldan, Alejandro Pardo, Juan Leon Alcazar, Fabian Caba Heilbron, Chen Zhao, Silvio Giancola, and Bernard Ghanem. MAD: A scalable dataset for language grounding in videos from movie audio descriptions. In *IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*. Inst. Electr. Electron. Eng. (IEEE), Jun. 2022. 2, 5
- [62] Vladimir Somers, Christophe De Vleeschouwer, and Alexandre Alahi. Body part-based representation learning for occluded person Re-Identification. In *IEEE/CVF Winter Conf. Appl. Comput. Vis. (WACV)*, pages 1613–1623, Waikoloa, HI, USA, Jan. 2023. Inst. Electr. Electron. Eng. (IEEE). 2
- [63] Alessandro Suglia, José Lopes, Emanuele Bastianelli, Andrea Vanzo, Shubham Agarwal, Malvina Nikandrou, Lu Yu, Ioannis Konstas, and Verena Rieser. Going for GOAL:

- A resource for grounded football commentaries. *CoRR*, abs/2211.04534, 2022. 5
- [64] Genki Suzuki, Sho Takahashi, Takahiro Ogawa, and Miki Haseyama. Team tactics estimation in soccer videos based on a deep extreme learning machine and characteristics of the tactics. *IEEE Access*, 7:153238–153248, 2019. 2
- [65] Makarand Tapaswi, Yukun Zhu, Rainer Stiefelhamen, Antonio Torralba, Raquel Urtasun, and Sanja Fidler. MovieQA: Understanding stories in movies through question-answering. In *IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*. Inst. Electr. Electron. Eng. (IEEE), Jun. 2016. 2
- [66] Graham Thomas, Rikke Gade, Thomas B. Moeslund, Peter Carr, and Adrian Hilton. Computer vision for sports: current applications and research topics. *Comput. Vis. Image Underst.*, 159:3–18, Jun. 2017. 2
- [67] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3D convolutional networks. In *IEEE Int. Conf. Comput. Vis. (ICCV)*. Inst. Electr. Electron. Eng. (IEEE), Dec. 2015. 6
- [68] Karl Tuyls, Shayegan Omidshafiei, Paul Muller, Zhe Wang, Jerome Connor, Daniel Hennes, Ian Graham, William Spearman, Tim Waskett, Dafydd Steel, Pauline Luc, Adria Recasens, Alexandre Galashov, Gregory Thornton, Romuald Elie, Pablo Sprechmann, Pol Moreno, Kris Cao, Marta Garnelo, Praneet Dutta, Michal Valko, Nicolas Heess, Alex Bridgland, Julien Pérolat, Bart De Vylder, S. M. Ali Eslami, Mark Rowland, Andrew Jaegle, Remi Munos, Trevor Back, Razia Ahamed, Simon Bouton, Nathalie Beauguerlange, Jackson Broshear, Thore Graepel, and Demis Hassabis. Game plan: What AI can do for football, and what football can do for AI. *Journal of Artificial Intelligence Research*, 71:41–88, May 2021. 1
- [69] Gabriel Van Zandycke and Christophe De Vleeschouwer. 3D ball localization from a single calibrated image. In *IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Work. (CVPRW)*, pages 3471–3479, New Orleans, LA, USA, Jun. 2022. Inst. Electr. Electron. Eng. (IEEE). 2
- [70] Gabriel Van Zandycke, Vladimir Somers, Maxime Istasse, Carlo Del Don, and Davide Zambrano. DeepSportradar-v1: Computer vision dataset for sports understanding with high quality annotations. In *Int. ACM Work. Multimedia Content Anal. Sports (MMSports)*, pages 1–8, Lisbon, Port., Oct. 2022. ACM. 2
- [71] Renaud Vandeghen, Anthony Cioppa, and Marc Van Droogenbroeck. Semi-supervised training to improve player and ball detection in soccer. In *IEEE Int. Conf. Comput. Vis. Pattern Recognit. Work. (CVPRW)*, *CVSports*, pages 3480–3489, New Orleans, LA, USA, Jun. 2022. Inst. Electr. Electron. Eng. (IEEE). 2
- [72] Kanav Vats, William McNally, Pascale Walters, David A. Clausi, and John S. Zelek. Ice hockey player identification via transformers and weakly supervised learning. In *IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Work. (CVPRW)*, pages 3450–3459, New Orleans, LA, USA, Jun. 2022. Inst. Electr. Electron. Eng. (IEEE). 2
- [73] Ramakrishna Vedantam, C. Lawrence Zitnick, and Devi Parikh. CIDEr: Consensus-based image description evaluation. In *IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*. Inst. Electr. Electron. Eng. (IEEE), Jun. 2015. 5
- [74] Teng Wang, Ruimao Zhang, Zhichao Lu, Feng Zheng, Ran Cheng, and Ping Luo. End-to-end dense video captioning with parallel decoding. In *IEEE Int. Conf. Comput. Vis. (ICCV)*. Inst. Electr. Electron. Eng. (IEEE), Oct. 2021. 2
- [75] Fei Wu, Qingzhong Wang, Jian Bian, Haoyi Xiong, Ning Ding, Feixiang Lu, Jun Cheng, and Dejing Dou. A survey on video action recognition in sports: Datasets, methods and applications. *CoRR*, abs/2206.01038, 2022. 2
- [76] Antoine Yang, Antoine Miech, Josef Sivic, Ivan Laptev, and Cordelia Schmid. Just ask: Learning to answer questions from millions of narrated videos. In *IEEE Int. Conf. Comput. Vis. (ICCV)*. Inst. Electr. Electron. Eng. (IEEE), Oct. 2021. 2
- [77] Antoine Yang, Arsha Nagrani, Paul Hongsuck Seo, Antoine Miech, Jordi Pont-Tuset, Ivan Laptev, Josef Sivic, and Cordelia Schmid. Vid2Seq: Large-scale pretraining of a visual language model for dense video captioning. *CoRR*, abs/2302.14115, 2023. 2
- [78] Huanyu Yu, Shuo Cheng, Bingbing Ni, Minsi Wang, Jian Zhang, and Xiaokang Yang. Fine-grained video captioning for sports narrative. In *IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*. Inst. Electr. Electron. Eng. (IEEE), Jun. 2018. 2, 5
- [79] Junqing Yu, Aiping Lei, Zikai Song, Tingting Wang, Hengyou Cai, and Na Feng. Comprehensive dataset of broadcast soccer videos. In *IEEE Conf. Multimedia Inf. Process. Retr. (MIPR)*, pages 418–423, Miami, FL, USA, Apr. 2018. Inst. Electr. Electron. Eng. (IEEE). 2
- [80] Qi Zhang, Yuqing Song, and Qin Jin. Unifying event detection and captioning as sequence generation via pre-training. *CoRR*, abs/2207.08625, 2022. 2
- [81] Luowei Zhou, Chenliang Xu, and Jason Corso. Towards automatic learning of procedures from Web instructional videos. *AAAI Conf. Artif. Intell.*, 32(1), Apr. 2018. 2, 5
- [82] X. Zhou, L. Kang, J. Xin, B. He, and J. Xin. Feature combination meets attention: Baidu soccer embeddings and transformer based temporal detection. *CoRR*, abs/2106.14447, Jun. 2021. 3, 6
- [83] He Zhu, Junwei Liang, Chengzhi Lin, Jun Zhang, and Jianming Hu. A transformer-based system for action spotting in soccer videos. In *Int. ACM Work. Multimedia Content Anal. Sports (MMSports)*, pages 103–109, Lisbon, Port., Oct. 2022. ACM. 2
- [84] Kevin Zhu, Alexander Wong, and John McPhee. FenceNet: Fine-grained footwork recognition in fencing. In *IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Work. (CVPRW)*, pages 3588–3597, New Orleans, LA, USA, Jun. 2022. Inst. Electr. Electron. Eng. (IEEE). 2