

Domain Adaptation of VLM for Soccer Video Understanding

Tiancheng Jiang^{1,2*} Henry Wang² Md Sirajus Salekin²
Parmida Atighehchian² Shinan Zhang²

¹Massachusetts Institute of Technology ²Amazon Web Services

Abstract

Vision Language Models (VLMs) have demonstrated strong performance in multi-modal tasks by effectively aligning visual and textual representations. However, most video understanding VLM research has been domain-agnostic, leaving the understanding of their transfer learning capability to specialized domains underexplored. In this work, we address this by exploring the adaptability of open-source VLMs to specific domains, and focusing on soccer as an initial case study. Our approach uses large-scale soccer datasets and LLM to create instruction-following data, and use them to iteratively fine-tune the general-domain VLM in a curriculum learning fashion (first teaching the model key soccer concepts to then question answering tasks). The final adapted model, trained using a curated dataset of 20k video clips, exhibits significant improvement in soccer-specific tasks compared to the base model, with a 37.5% relative improvement for the visual question-answering task and an accuracy improvement from 11.8% to 63.5% for the downstream soccer action classification task.

1. Introduction

Recent advances in multimodal large language models and vision language models (VLMs) have significantly improved their ability to process and understand visual inputs, including images and videos [3]. However, while general-domain video VLMs have been extensively studied, there remains a gap in understanding how well these models can be adapted to domain-specific applications

[8, 15, 17]. This paper aims to address this gap by exploring an effective methodology and repeatable recipe for adapting a general-purpose video VLM to a specialized domain. Furthermore, since fine-tuning video VLMs tends to be computationally expensive, we wanted to explore an approach that’s not only simple and practical but efficient at the same time.

As a case study, we investigate the adaptation of LLaVA-NeXT-Video [41], a state-of-the-art open-source video VLM, to the soccer domain in sports. Soccer presents unique challenges for video understanding due to its dynamic nature, rapid motion of objects, and fine-grained events [6, 10]. To this end, we employ a structured three-stage fine-tuning approach using synthetically generated data. First, the model undergoes **Concept Alignment**, where it is trained to associate soccer-specific concepts with video clips to improve visual-text alignment. Next, **Instruction Tuning** enhances the model’s ability to answer soccer-specific queries. Finally, **Downstream Task Fine-tuning** aims to boost model performance on specific soccer-related tasks. This methodology ensures a gradual and effective adaptation of the model to the target domain. We evaluate the performance of the adapted model on multiple tasks, including caption generation, visual question answering, and action classification.

Our contributions are summarized below:

- We introduce a simple, yet effective multi-stage fine-tuning strategy for adapting video VLMs to specialized domains. We further demonstrate the success of our proposed adaptation strategy through the lens of a real-world case study in soccer, where a general-purpose video VLM learns to effectively analyze different tasks in this challenging sports domain.

*Work done during an internship at AWS.

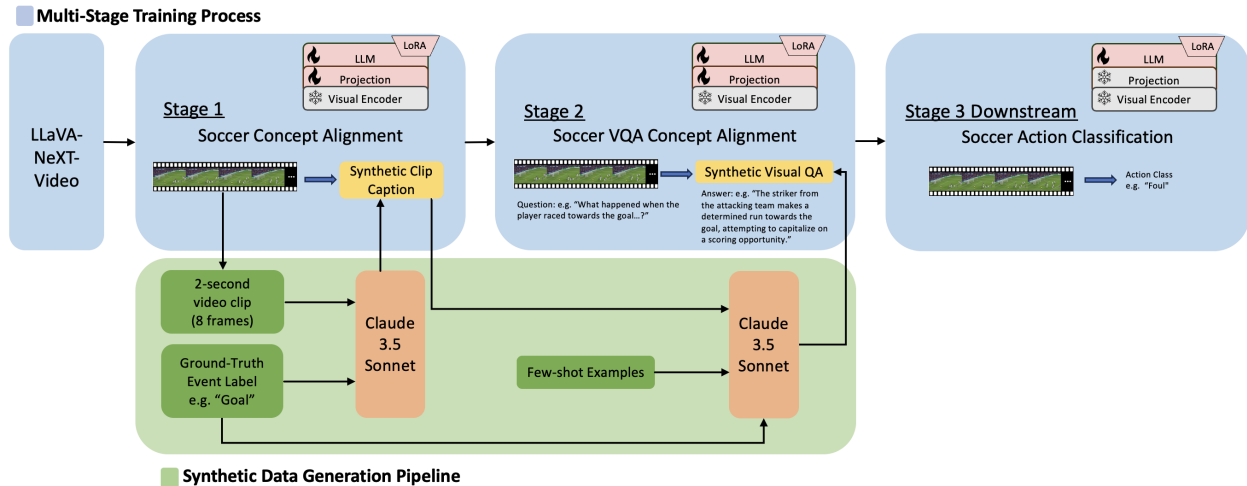


Figure 1. Overall workflow of our proposed multi-stage domain adaptation process. The LLaVA-Next-Video model is iteratively trained in a curriculum learning fashion (first teaching the model key soccer concepts to then question answering tasks). During the first two training stages, visual encoder is frozen while training both the projector and LLM modules using LoRA. For downstream task, only the LLM modules are trained with LoRA. The figure also demonstrates the synthetic data generation pipeline instruction tuning data that is generated for the first two training stages using Claude 3.5 Sonnet V1.

- We formalize best practices through various ablation studies, providing insight on the specific design choices of the training process for domain adaptation.

2. Related work

The Rise of VLM. Building on the success of large language models (LLMs), vision language models (VLMs), which integrate vision encoders with LLMs for multi-modal reasoning, have seen a tremendous amount of improvements in performance that span across tasks like image captioning, segmentation, classification, etc. over the last few years [20, 21, 25, 26, 31]. Multi-stage training that first aligns vision and language representation followed by instruction tuning appears to emerge as an effective approach [24]. VLMs have also been extended for challenging video understanding tasks, which requires temporal modeling. To this date, however, most of the work has been focused on creating general-purpose video VLMs. [2, 23, 27, 39, 40]

VLM Domain Adaptation. Domain adaptation of VLMs has been an active research area. Med-Flamingo [29] and LLaVA-Med [18], have demonstrated effective adaptation to the medical domain by fine-tuning general-purpose VLMs on biomedical data. Other studies have extended into fields

such as robotics [4], scientific literature [22], remote sensing [33], etc. For example, RT-2 [4] adapted a vision-language model for robotics by training it on web-scale multimodal data alongside robotic demonstrations, allowing it to translate visual understanding into actionable robotic tasks. Despite these advances, video VLM domain adaptation remains under explored as most efforts have been focused on image-based tasks. The scarcity of high-quality labeled datasets complicates domain adaptation for video-based VLMs.

Sports Understanding. Sports understanding [34] has been an evolving field. It includes tasks such as, action recognition [6, 10, 11], commentary generation [30, 32, 36, 38], intelligent refereeing [12, 13], etc. With the rise of multi-modal large language models (MLLMs) and VLMs, there have also been works focused at creating generalized frameworks for multi-sports understanding [19, 37]. Soccer analysis, in particular, traditionally was focused on action spotting [6, 28], replay grounding [12, 42], player tracking [5, 35], and foul recognition [12, 13], and had largely been facilitated by the SoccerNet datasets [5, 7, 9, 10]. As an initial case study of our multi-stage training process, we explore adapting a general-purpose video VLM to the soccer domain.

3. Video Data Source and Preprocessing

We source videos from two large-scale soccer datasets, SoccerNet-V2 [6] and WyScout [16], both of which provide timestamped key soccer event labels.

3.1. Video Sources

SoccerNet-V2 consists of 550 full-length broadcast matches from six major European leagues, with Action Spotting labels marking 17 key soccer events, such as “Goal”, “Throw-in”, etc. WyScout, a subscription-based proprietary soccer dataset, includes over 2,000 games from different leagues worldwide and offer event labels including both primary and secondary actions. It covers a broader range of in-play events like passes, duels, and tactical actions. Given the differences in soccer event labels, the two datasets capture distinct though overlapping distributions of soccer activities, making their combined use beneficial for model training.

For both datasets, we extract 2-second clips based on event timestamps while preserving their respective training-test splits (300/100 games for SoccerNet; a 4:1 ratio for WyScout). These clips are then used to create instruction-following data for model adaptation.

3.2. Preprocessing

Video Clips. To prepare the video clips for fine-tuning, we process full soccer matches into 2-seconds long clips. Each 2-second long game is represented by uniformly selecting 8 frames, a common choice for many open-source video VLMs [23, 27] and allows the context to be continuous. As each clip corresponding to one key soccer event, it starts 0.5 seconds before and ends 1.5 seconds after the event’s timestamp to capture both the lead-up and aftermath of a soccer event. The rationale behind the 2-seconds duration can be found in Sec. 5.3.

Soccer Event Distribution. Once the clips were generated from the event labels, we randomly sample to ensure uniform distribution of soccer events in both the training and test set. In SoccerNet, however, due to the super rare nature of 4 event categories (Red card, Yellow → Red card, Penalty, Kick

off), we decided to remove them so the resulting dataset contains clips representing 13 soccer events instead of the original 17.

4. Methodology

In this section, we introduce LLaVA-NeXT-Video, a general-domain video VLM, as the base VLM used in this research. We then provide an overview of our multi-stage domain adaptation process.

4.1. Base Model and Architecture

LLaVA-Next-Video [41], a state-of-the-art open-source general video VLM, extends the capabilities of LLaVA-Next (LLaVA-1.6) [25] by integrating video-specific processing enhancements while maintaining strong performance in multimodal reasoning tasks.

LLaVA-Next-Video [41] incorporates a *CLIP-ViT* visual encoder, which maps image and video frames into latent feature representations. A lightweight MLP projection layer connects the visual encoder to the language model, ensuring efficient multi-modal integration. The model also applies two architectural changes to enhance video comprehension, which are AnyRes and Length Generalization. The model is pretrained on a mix of multi-modal datasets, including large-scale image-text and video-text data, making it a strong base model to build upon. The same network architecture is used in this research. To fine-tune the model, we create instruction following dataset and train the model using its original auto-regressive training objective. We calculate the probability of the target response X_r as the following.

$$p(X_r|X_v, X_q) = \prod_{i=1}^L p_{\theta}(x_i | X_v, X_{q,<i}, X_{r,<i}) \quad (1)$$

where X_v and X_q indicate the video clip and the instruction or question respectively. θ is the trainable parameters and $X_{q,<i}$ and $X_{r,<i}$ are the instruction and answer tokens before the current prediction token X_i .

4.2. Multi-Stage Fine-tuning Process

Our proposed fine-tuning process as shown in Figure 1 consists of three training stages: 1) Soccer Concept Alignment, 2) Soccer Visual QA In-

struction Tuning, and 3) Downstream Task Fine-tuning. Each stage progressively improves the model’s ability to align and understand the soccer domain. For efficient fine-tuning, we employ the Low Rank Adaptation (LoRA) method [14] in all three training stages. The full training was completed using four 24G A10 GPUs.

Stage 1 - Soccer Concept Alignment. The first training stage, Soccer Concept Alignment, teaches the base model fundamental soccer concepts by associating video clips with their corresponding event labels. We generate a synthetic clip-caption dataset by prompting Claude 3.5 Sonnet V1 [1] to describe each 2-second clip based on its ground-truth event label and the 8 extracted frames. This approach ensures accurate and contextually grounded captions, reducing hallucinations that arise when generating descriptions purely from visual inputs.

From the clip-caption pairs, we then apply a structured expansion method to convert them into instruction-following samples: each sample consists of a video clip and an instruction prompting the model to describe the action, with the corresponding caption as the target response. We manually curate a set of synonymous instructions (e.g. “Please provide a description of what happened in the soccer match video?”) and randomly assign one per sample to create some variability in the instruction.

During training, we only freeze the visual encoder and apply LoRA adapters to the language model, and train using the model’s original task on next-token prediction on the ground-truth captions.

Stage 2 - Soccer VQA Instruction Tuning. The second training stage, Soccer VQA Instruction Tuning, aims to train the model to respond to diverse soccer-related queries. We construct this stage’s instruction-following dataset using the synthetically generated captions from Stage 1, prompting Claude 3.5 Sonnet V1 to generate five distinct question-answer pairs per caption:

- **Description Questions:** Asking for a textual description of the video.
- **Temporal Questions:** Querying the sequence of events to test time-based reasoning.
- **Causal Questions:** Understanding cause-effect relationships between actions.
- **Prediction Questions:** Asking what might hap-

pen next based on visual context.

- **Action Spotting Questions:** Identifying key soccer actions in the clip.

To maintain dataset balance, we randomly sample from each question type, ensuring equal distribution while keeping dataset size consistent with Stage 1. We also enforce a 4:1 ratio of free-form to multiple-choice formats. During training, we also only freeze the visual encoder and apply LoRA adapters to the language model and train using the model’s original next-token prediction task on the ground-truth query answers.

Stage 3 - Downstream Task Fine-tuning. The final stage fine-tunes the model on any downstream task in the domain. The purpose of this stage’s training is mainly to teach the model the specific formatting needed for the downstream task.

Instruction Following Dataset Details. All instruction-following datasets consist of a 2-second long soccer video clip X_v , an instruction or question X_q , and an associated response X_r , all following the same format:

```
Human:  $X_v$   $X_q$ <STOP>\n
Assistant:  $X_r$ <STOP>\n
```

For the first two training stages, we create three dataset sizes, 3k, 10k, and 20k samples, with the smaller datasets as subsets of the larger ones, to understand the effect of dataset sizes on model performance. Also, to ensure privacy and generalization, we prompt Claude 3.5 Sonnet V1 to anonymize player and team names when generating synthetic data.

5. Experiments

5.1. Evaluation Tasks and Metrics

We designed three evaluation tasks.

Task 1: Caption Generation. This task evaluates the model’s ability to generate accurate and descriptive captions for 2-sec soccer clips. Two test sets were curated: a 1,000-sample set and a 200-sample subset for ablation studies. Unlike the training data in Stages 1 and 2, where the 3k and 10k datasets only contain SoccerNet clips, both test sets here contain an equal split between SoccerNet and WyScout samples. Captions were generated using the same approach as in Stage 1.

We employed two sets of evaluation metrics. The first was LLM-as-a-Judge metrics, where we asked Claude 3.5 Sonnet V1 to rate the predicted captions against the label captions, and two scores from 1 to 5 (5 being the best) were provided for Correctness and Detailness. The average for either criteria was computed. The second set of evaluation metrics included traditional captioning metrics such as BLEU and ROUGE scores.

Task 2: Visual Question Answering (VQA). This task assessed the model’s ability to answer soccer-related queries. Two test sets (1,000 and 200 samples) were also constructed for the same reason, and both maintain a balanced distribution across the five question types and a 4:1 ratio of free-form to multiple-choice formats. Each test set included an equal split of clips from SoccerNet and WyScout.

For evaluation, we used Claude 3.5 Sonnet V1 to score both the predicted response and a reference upper bound response on helpfulness, relevance, accuracy, and level of details from 1 to 10 (10 being the highest score). We compute a relative score by normalizing the predicted response score with the upper bound response score, and an average relative score for the whole test set is calculated. A higher score indicates better alignment with the reference and a score exceeding 100 suggests the model outperformed the reference in certain cases.

Task 3: SoccerNet Action Classification. This task was chosen as an example downstream task, and it evaluated the model’s ability to classify 2-sec soccer clip into one of 13 predefined soccer action classes in an instruction following format (no classifier head was added). The training set consists of 3.3k samples, which is small as the model is expected to have already acquired soccer domain knowledge from the earlier training stages. The test set consisted of two sizes, a larger set of 1,300 samples (containing 100 samples per event class) and a smaller test set of 100 samples (containing approximately even distribution across classes). The performance was measured using standard classification metrics including accuracy, precision, recall, and F1-score.

5.2. Experimental Results

For the main experiments, we adapted the base Llava-NeXT-Video model through the first two stages of the fine-tuning process using different sizes of training sets, and used the 3.3k training set for the third fine-tuning stage on Action Classification as the downstream task. We report results of the three evaluation tasks below.

Caption Generation Task. We used the model after Stage 2 for this task. The results from Tab. 1 indicate that the 20k model (model trained using the 20k samples dataset) consistently outperforms the other models across both correctness and detailness scores. However, interesting trends can be observed when analyzing the performance of the 3k and 10k models, which were trained only on SoccerNet clips. For example, the 3k model’s average detailness or correctness scores decreased in comparison to that of the Base model, even though SoccerNet-specific scores improved. This decline was attributed to lower WyScout-specific scores, highlighting limitations in the 3k model’s ability to generalize to unseen distributions in WyScout. The 10k model also displayed lower WyScout-specific scores than the Base model, but it overall performed better than the 3k model, indicating perhaps some generalization effects. The inclusion of WyScout data in the training set for the 20k model resulted in significant performance gains across both subsets of the test set, demonstrating the value of a mixed training set. Traditional metrics, such as BLEU-4 and ROUGE in Tab. 6 of the Appendix, showed a large improvement from the Base model to the 3k model, but subsequent increases with larger datasets were less pronounced.

Visual Question Answering Task. For VQA, we also used the model after Stage 2 for this task. From Fig. 2 and Tab. 2, we can see that the 20k model achieved the highest relative scores, reflecting the strongest performance. There was also a steady improvement from models trained using smaller to larger datasets, with the largest gains between the Base and 3k models and between the 10k and 20k models. The sharp improvement between the 10k and 20k models not only signals the importance of having more data but it also emphasize the importance of incorporating WyScout data in the training

(Question Count)	Correctness Scores			Detailness Scores		
	Total Avg. (1000)	SoccerNet (500)	WyScout (500)	Total Avg. (1000)	SoccerNet (500)	WyScout (500)
Claude 3.5 Sonnet V1	2.779	2.594	2.964	2.402	2.420	2.384
LLaMA 3.2 ¹	1.875	1.748	2.002	2.159	2.166	2.159
Base Model	2.130	2.158	2.102	2.226	2.276	2.176
Adapted Models						
3K (SoccerNet clips only)	1.782	2.214	1.350	2.048	2.436	1.660
10K (SoccerNet clips only)	1.905	2.314	1.496	2.2	2.566	1.834
20K	2.528	2.446	2.610	2.675	2.616	2.734

Table 1. Performance comparison of Caption Generation task measured by LLM as Judge scored on Correctness and Detailness.




Figure 2. Comparison of the VQA task abilities measured by the relative score.

set as it broadens the data distribution. Looking at specific question types, the largest gains were observed in Action Recognition, Description, and Temporal question types.

Soccer Action Classification Task. For Soccer Action Classification, since it’s a downstream task, we used the model at the end of stage 3 to run the test set. From Fig. 4, one can see that the accuracy improves significantly from the Base model to that of the 20k model, increasing from 11.8% to 57.8%. However, the confusion matrix of the

¹Only a single frame selected at timestamp of the event was used as input for the model due to API limitation. We tested using the 90B version of LLaMA 3.2.

20k model prediction revealed that the model struggles to distinguish between visually similar actions, such as “Shots on Target” and “Shots off Target,” suggesting that additional training strategies may be required to refine fine-grained classification. As benchmarks, LLaMA 3.2 and Claude 3.5 Sonnet V1 scored an accuracy of 24.2% and 26.7% respectively.

5.3. Ablation Studies and Best Practices

Training the Projector or Not. We compared training only the LLM versus training both the LLM and the projector through the first two stages of the training process. Based on results shown in Tab. 8 and 9 in the appendix, we suggest training on both components as it provides incremental benefits without substantial computational overhead.

Rationale for 2-Second Clip Duration. To adapt video VLMs for soccer analysis, we segmented full matches into 2-sec clips. While longer clips are typically used as pre-training data, we found that for soccer domain such duration can introduce excessive noise as it captures overlapping soccer events and increases hallucination risks in synthetic data generation. In contrast, 2-sec clips provide a focused context, ensuring a strong association with the soccer event label. This effectively reduces ambiguity and enhances data alignment. Additionally, with a fixed eight-frame constraint, shorter clips minimize information loss in this fast-paced domain. Fig. 5 illustrates how shorter clips offer clearer event representation, making them more interpretable even for humans.

Multi-Stage Training Process. Our results highlighted the importance of the multi-stage train-

(Question Count)	Question Sources		Domains					Overall (1000)
	SoccerNet (500)	WyScout (500)	Desc. (250)	Temp. (250)	Causal (250)	Pred. (250)	Action Rec. (250)	
Claude 3.5 Sonnet V1	75.79	74.85	42.72	91.74	83.57	95.50	63.07	75.32
LLaMA 3.2 ¹	69.97	85.66	40.60	67.42	97.38	126.77	56.92	77.82
Base Model	57.44	62.97	36.21	61.49	83.49	79.26	40.57	60.20
Adapted Models								
3K (SoccerNet clips only)	78.96	63.67	47.62	69.91	83.84	89.95	65.25	71.31
10K (SoccerNet clips only)	82.77	64.81	51.31	72.49	82.41	93.58	69.07	73.77
20K	83.76	81.83	59.14	81.91	91.07	97.85	84.02	82.79

Table 2. Performance comparison of VQA task abilities, measured by the relative score via Claude 3.5 Sonnet V1 evaluation, normalized against a reference response. A higher score indicates better alignment with the reference and a score exceeding 100 suggests the model outperformed the reference in certain cases.

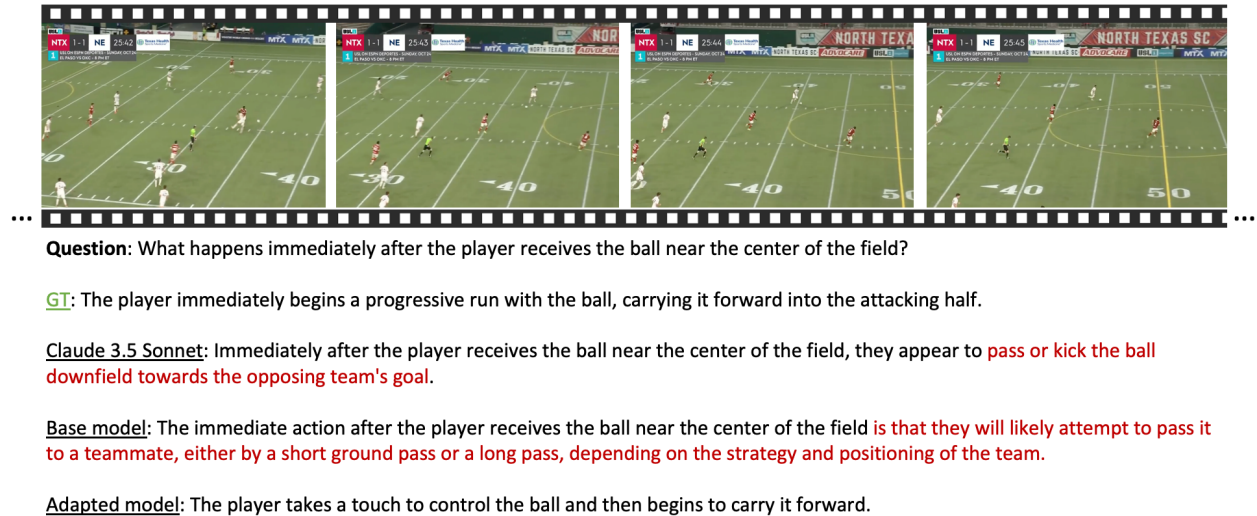


Figure 3. A qualitative example from the VQA task test set showing the ground-truth answer and responses from the different models. The red indicates hallucination

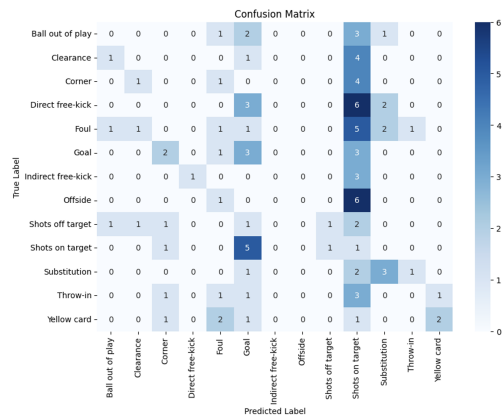
ing approach in both VQA performance and action classification accuracy. As shown in Tab. 3, while Instruction Tuning (IT) provided the largest boost in VQA performance, the best results were achieved when the model underwent both Concept Alignment (CA) and IT, indicating that CA establishes essential soccer knowledge before instruction tuning. Similarly, Tab.14 in Appendix demonstrates that directly fine-tuning the base model on action classification training set yields minimal improvement, whereas models trained through the full CA → IT → Action Classification (AC) sequence achieved the highest accuracy.

Extension to longer duration clips The fine-tuned model, trained on 2-sec clips, performed similarly when classifying 5-sec clips, achieving a macro F1


Visual QA Task	RS
Base Model	58.57
20k - Base→Concept Alignment(CA)	54.85
20k - Base→Instruction Tuning(IT)	79.09
20k - Base→CA→IT	85.80

Table 3. Comparison of performance for the VQA task across different training sequence variants, measured by the average relative score.

score of 0.61 compared to 0.63 on 2-sec clips. This suggests early evidence of the model’s generalizability to varying clip lengths, which could improve efficiency in downstream tasks.




(a) Base Model Pred.



(b) 20k Model Pred.

Figure 4. Confusion Matrices from the Action Classification task for the Base model (LLaVA-NeXT-Video), and our adapted 20k Model. Their testset prediction accuracies are 0.118, and 0.635 respectively. Full results can be found in Tab. 7 of the Appendix

2-second Clip by Frames



14-second Clip by Frames




Figure 5. Demonstration of short clip vs. long clip on information continuity. The above 2s clip is much more continuous hence easier to comprehend while the bottom 14s clip has more information loss due to the wide time gap between the frames.

5.4. Limitations

While our training relies on synthetic data, which may introduce stylistic biases, it provides a strong foundation for developing and validating our methodology. Our data pipeline, though effective, can be prone to hallucination for longer clips. We evaluated our approach using LLaVA-Next-Video but have not yet tested other video VLMs, particularly those with longer context windows, which presents opportunities for future explorations. We also did not fine-tune the visual encoder, which could further enhance the model’s ability to capture fine-grained, domain-specific visual cues.

6. Conclusion

This work demonstrates the feasibility of adapting a general-purpose video VLM to a specialized domain through a structured multi-stage fine-tuning approach using synthetically generated instruction-following data. Looking ahead, we aim to extend our evaluation to additional downstream tasks beyond soccer action classification to more real-world soccer applications to further validate the generalizability of our approach. Additionally, we plan to apply this approach in other specialized domains to assess and refine the broader generalizability of our multi-stage fine-tuning procedure.

References

- [1] Anthropic. Claude 3.5 sonnet: Our best model yet, 2024. Accessed: 2025-03-11. [4](#)
- [2] Kirolos Ataallah, Xiaoqian Shen, Eslam Abdelrahman, Essam Sleiman, Deyao Zhu, Jian Ding, and Mohamed Elhoseiny. Minigt4-video: Advancing multimodal llms for video understanding with interleaved visual-textual tokens. *arXiv preprint arXiv:2404.03413*, 2024. [2](#)
- [3] Florian Bordes, Richard Yuanzhe Pang, Anurag Ajay, Alexander C. Li, Adrien Bardes, Suzanne Petryk, Oscar Mañas, Zhiqiu Lin, Anas Mahmoud, Bargav Jayaraman, Mark Ibrahim, Melissa Hall, Yunyang Xiong, Jonathan Lebensold, Candace Ross, Srihari Jayakumar, Chuan Guo, Diane Bouchacourt, Haider Al-Tahan, Karthik Padthe, Vasu Sharma, Huijuan Xu, Xiaoqing Ellen Tan, Megan Richards, Samuel Lavoie, Pietro Astolfi, Reyhane Askari Hemmat, Jun Chen, Kushal Tirumala, Rim Assouel, Mazda Moayeri, Arjang Talattof, Kamalika Chaudhuri, Zechun Liu, Xilun Chen, Quentin Garrido, Karen Ullrich, Aishwarya Agrawal, Kate Saenko, Asli Celikyilmaz, and Vikas Chandra. An introduction to vision-language modeling. *ArXiv*, abs/2405.17247, 2024. [1](#)
- [4] Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Xi Chen, Krzysztof Choromanski, Tianli Ding, Danny Driess, Avinava Dubey, Chelsea Finn, Pete Florence, Chuyuan Fu, Montse Gonzalez Arenas, Keerthana Gopalakrishnan, Kehang Han, Karol Hausman, Alex Herzog, Jasmine Hsu, Brian Ichter, Alex Irpan, Nikhil Joshi, Ryan Julian, Dmitry Kalashnikov, Yuheng Kuang, Isabel Leal, Lisa Lee, Tsang-Wei Edward Lee, Sergey Levine, Yao Lu, Henryk Michalewski, Igor Mordatch, Karl Pertsch, Kanishka Rao, Krista Reymann, Michael Ryoo, Grecia Salazar, Pannag Sanketi, Pierre Sermanet, Jaspier Singh, Anikait Singh, Radu Soricut, Huong Tran, Vincent Vanhoucke, Quan Vuong, Ayzaan Wahid, Stefan Welker, Paul Wohlhart, Jialin Wu, Fei Xia, Ted Xiao, Peng Xu, Sichun Xu, Tianhe Yu, and Brianna Zitkovich. Rt-2: Vision-language-action models transfer web knowledge to robotic control. In *arXiv preprint arXiv:2307.15818*, 2023. [2](#)
- [5] 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 3491–3502, 2022. [2](#)
- [6] Adrien Deliege, 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 4508–4519, 2021. [1](#), [2](#), [3](#)
- [7] Adrien Deliege, 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 4508–4519, 2021. [2](#)
- [8] Xinyu Fang, Kangrui Mao, Haodong Duan, Xiangyu Zhao, Yining Li, Dahua Lin, and Kai Chen. Mmbench-video: A long-form multi-shot benchmark for holistic video understanding. *Advances in Neural Information Processing Systems*, 37: 89098–89124, 2024. [1](#)
- [9] Sushant Gautam, Mehdi Houshmand Sarkhoosh, Jan Held, Cise Midoglu, Anthony Cioppa, Silvio Giancola, Vajira Thambawita, Michael A Riegler, Pal Halvorsen, and Mubarak Shah. Soccernet-echoes: A soccer game audio commentary dataset. In *2024 International Symposium on Multimedia (ISM)*, pages 71–78. IEEE, 2024. [2](#)
- [10] Silvio Giancola, Mohieddine Amine, Tarek Dghaily, and Bernard Ghanem. Soccernet: A scalable dataset for action spotting in soccer videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 1711–1721, 2018. [1](#), [2](#)
- [11] Xiaofan Gu, Xinwei Xue, and Feng Wang. Fine-grained action recognition on a novel basketball dataset. In *International Conference on Acoustics, Speech, and Signal Processing*, pages 2563–2567. IEEE, 2020. [2](#)
- [12] Jan Held, Anthony Cioppa, Silvio Giancola, Abdullah Hamdi, Bernard Ghanem, and Marc Van Droogenbroeck. Vars: Video assistant referee system for automated soccer decision making from multiple views. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 5085–5096, 2023. [2](#)
- [13] Jan Held, Hani Itani, Anthony Cioppa, Silvio Giancola, Bernard Ghanem, and Marc Van Droogenbroeck. X-vars: Introducing explainability in football refereeing with multi-modal large language

- models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3267–3279, 2024. 2
- [14] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2022. 4
- [15] Bin Huang, Xin Wang, Hong Chen, Zihan Song, and Wenwu Zhu. Vtimellm: Empower llm to grasp video moments. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14271–14280, 2024. 1
- [16] Hudl. Wycout: Football data and analytics platform, 2024. Accessed: March 5, 2025. 3
- [17] Yixing Jiang, Jeremy Irvin, Ji Hun Wang, Muhammad Ahmed Chaudhry, Jonathan H. Chen, and Andrew Y. Ng. How well does gpt-4v(ision) adapt to distribution shifts? a preliminary investigation. *arXiv preprint arXiv:2401.06794*, 2024. 1
- [18] Chunyuan Li, Cliff Wong, Sheng Zhang, Naoto Usuyama, Haotian Liu, Jianwei Yang, Tristan Naumann, Hoifung Poon, and Jianfeng Gao. Llava-med: Training a large language-and-vision assistant for biomedicine in one day. *Advances in Neural Information Processing Systems*, 36: 28541–28564, 2023. 2
- [19] Haopeng Li, Andong Deng, Qihong Ke, Jun Liu, Hossein Rahmani, Yulan Guo, Bernt Schiele, and Chen Chen. Sports-qa: A large-scale video question answering benchmark for complex and professional sports. *arXiv preprint arXiv:2401.01505*, 2024. 2
- [20] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. BLIP-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *Proceedings of the 40th International Conference on Machine Learning*, pages 19730–19742. PMLR, 2023. 2
- [21] Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. Visualbert: A simple and performant baseline for vision and language. *arXiv preprint arXiv:1908.03557*, 2019. 2
- [22] Xue Li, Yiyun Sun, Wei Cheng, Yinglun Zhu, and Haifeng Chen. Chain-of-region: Visual language models need details for diagram analysis. In *The Thirteenth International Conference on Learning Representations*, 2025. 2
- [23] Bin Lin, Bin Zhu, Yang Ye, Munan Ning, Peng Jin, and Li Yuan. Video-llava: Learning united visual representation by alignment before projection. *arXiv preprint arXiv:2311.10122*, 2023. 2, 3
- [24] Ji Lin, Hongxu Yin, Wei Ping, Pavlo Molchanov, Mohammad Shoeybi, and Song Han. Vila: On pre-training for visual language models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 26689–26699, 2024. 2
- [25] Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee. Llava-next: Improved reasoning, ocr, and world knowledge, 2024. 2, 3
- [26] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *Advances in Neural Information Processing Systems*. Curran Associates, Inc., 2019. 2
- [27] Muhammad Maaz, Hanoona Rasheed, Salman Khan, and Fahad Shahbaz Khan. Video-chatgpt: Towards detailed video understanding via large vision and language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (ACL 2024)*, 2024. 2, 3
- [28] Hassan Mkhallati, Anthony Cioppa, Silvio Giancola, Bernard Ghanem, and Marc Van Droogenbroeck. Soccernet-caption: Dense video captioning for soccer broadcasts commentaries. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 5074–5085, 2023. 2
- [29] Michael Moor, Qian Huang, Shirley Wu, Michihiro Yasunaga, Yash Dalmia, Jure Leskovec, Cyril Zakka, Eduardo Pontes Reis, and Pranav Rajpurkar. Med-flamingo: a multimodal medical few-shot learner. In *Machine Learning for Health (MLAH)*, pages 353–367. PMLR, 2023. 2
- [30] Mengshi Qi, Yunhong Wang, Annan Li, and Jiebo Luo. Sports video captioning via attentive motion representation and group relationship modeling. *IEEE Transactions on Circuits and Systems for Video Technology*, 30(8):2617–2633, 2019. 2
- [31] 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 *Proceedings of the 38th International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021. 2
- [32] Jiayuan Rao, Haoning Wu, Chang Liu, Yanfeng Wang, and Weidi Xie. Matchtime: Towards automatic soccer game commentary generation. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, 2024. 2
- [33] João Daniel Silva, João Magalhães, Devis Tuia, and Bruno Martins. Multilingual vision-language pre-training for the remote sensing domain. In *Proceedings of the 32nd ACM International Confer-*

- ence on *Advances in Geographic Information Systems*, pages 220–232, 2024. [2](#)
- [34] Graham Thomas, Rikke Gade, Thomas B Moeslund, Peter Carr, and Adrian Hilton. Computer vision for sports: Current applications and research topics. *Computer Vision and Image Understanding*, 159:3–18, 2017. [2](#)
- [35] Renaud Vandeghen, Anthony Cioppa, and Marc Van Droogenbroeck. Semi-supervised training to improve player and ball detection in soccer. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 3481–3490, 2022. [2](#)
- [36] Zeyu Xi, Ge Shi, Lifang Wu, Xuefen Li, Junchi Yan, Liang Wang, and Zilin Liu. Knowledge graph supported benchmark and video captioning for basketball. *arXiv preprint arXiv:2401.13888*, 2024. [2](#)
- [37] Haotian Xia, Zhengbang Yang, Yuqing Wang, Rhys Tracy, Yun Zhao, Dongdong Huang, Zezhi Chen, Yan Zhu, Yuanfang Wang, and Weining Shen. Sportqa: A benchmark for sports understanding in large language models. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics*, 2024. [2](#)
- [38] Huanyu Yu, Shuo Cheng, Bingbing Ni, Minsi Wang, Jian Zhang, and Xiaokang Yang. Fine-grained video captioning for sports narrative. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6006–6015, 2018. [2](#)
- [39] Rowan Zellers, Ximing Lu, Jack Hessel, Youngjae Yu, Jae Sung Park, Jize Cao, Ali Farhadi, and Yejin Choi. Merlot: Multimodal neural script knowledge models. In *Advances in Neural Information Processing Systems*, pages 23634–23651. Curran Associates, Inc., 2021. [2](#)
- [40] Hang Zhang, Xin Li, and Lidong Bing. Videollama: An instruction-tuned audio-visual language model for video understanding. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 543–553, Singapore, 2023. Association for Computational Linguistics. [2](#)
- [41] Yuanhan Zhang, Bo Li, haotian Liu, Yong jae Lee, Liangke Gui, Di Fu, Jiashi Feng, Ziwei Liu, and Chunyuan Li. Llava-next: A strong zero-shot video understanding model, 2024. [1](#), [3](#)
- [42] Xin Zhou, Le Kang, Zhiyu Cheng, Bo He, and Jingyu Xin. Feature combination meets attention: Baidu soccer embeddings and transformer based temporal detection. *arXiv preprint arXiv:2106.14447*, 2021. [2](#)