

# ProfVLM: A lightweight video-language model for multi-view proficiency estimation<sup>☆</sup>

Edoardo Bianchi<sup>a,\*</sup>, Jacopo Staiano<sup>b</sup>, Antonio Liotta<sup>a</sup>

<sup>a</sup> Free University of Bozen-Bolzano, Via Bruno Buozzi 1, Bozen-Bolzano, 39100, Italy

<sup>b</sup> University of Trento, Via Inama 5, Trento, 38122, Italy

## ARTICLE INFO

Dataset link: <https://ieeexplore.ieee.org/document/10658224>, <https://ego-exo4d-data.org>

### Keywords:

Proficiency estimation  
Action quality assessment  
Skill assessment  
Vision-language modeling  
Video understanding

## ABSTRACT

Most existing approaches formulate action quality assessment and skill proficiency estimation as discriminative prediction tasks, typically producing discrete labels or scores without explicitly modeling the reasoning process underlying the assessment. We instead reformulate the problem as generative vision-language modeling, introducing ProfVLM, a parameter-efficient vision-language model that jointly predicts proficiency levels and generates expert-like natural language feedback from multi-view videos. ProfVLM leverages conditional language generation to provide actionable insights along with quantitative evaluation scores. Central to our method is an AttentiveGatedProjector that dynamically fuses and projects multi-view egocentric and exocentric features from a frozen TimeSformer backbone into a language model fine-tuned for feedback generation. Trained on EgoExo4D with expert commentaries, ProfVLM surpasses state-of-the-art methods while using up to 20x fewer parameters and reducing training time by up to 60% compared to existing classification-based methods. By providing natural language critiques aligned with performance levels, this work shows that generative vision-language modeling offers a powerful and efficient paradigm shift for interpretable action quality assessment.

## 1. Introduction

Understanding proficiency in complex activities from video is crucial for applications such as coaching and rehabilitation. Unlike action recognition, proficiency estimation and Action Quality Assessment (AQA) require a deep and technical understanding of the actions, a non-trivial task even for state-of-the-art architectures.

Moreover, prior approaches present limitations that restrict their practical application. First, proficiency estimation is commonly formulated as a discriminative prediction task (classification or regression), relying on dedicated output heads that produce scores without explicitly modeling the reasoning process experts use when assessing performance. Second, most prior works focus on single-view analysis, despite comprehensive skill assessment often requiring multi-camera perspectives to capture complementary aspects such as hand positioning from egocentric views and overall technique from exocentric angles. Third, although language signals have been incorporated in some approaches, they are typically used as auxiliary supervision or in parallel tasks, rather than as a unified reasoning mechanism that directly drives proficiency prediction. As a result, prediction and explanation are optimized through separate objectives, limiting the model's

ability to align visual evidence with linguistic reasoning in a unified manner.

We address these gaps by reframing proficiency estimation as a conditional generation problem. This formulation enables the model to predict proficiency levels while producing expert-style explanatory commentary, mirroring how human evaluators justify their assessments. Building on this idea, we develop ProfVLM, a parameter- and training-efficient vision-language model that jointly outputs proficiency labels and interpretable, expert-like feedback from multi-view video inputs.

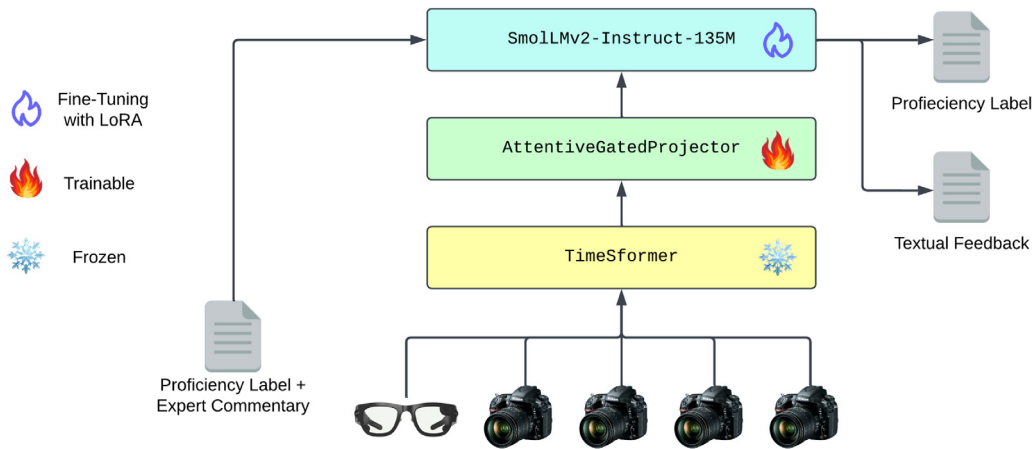
This paper introduces the following key contributions:

- A paradigm shift from classification to a generative formulation of proficiency estimation, enabling models to produce both proficiency labels and natural language commentary;
- ProfVLM, a parameter- and training-efficient vision-language model for generative proficiency estimation that handles multi-view inputs and generalizes across multiple domains;
- An AttentiveGatedProjector module that performs structured multi-view fusion through cross-view attention and learnable

<sup>☆</sup> This article is part of a Special issue entitled: 'CV for Sports' published in Computer Vision and Image Understanding.

\* Corresponding author.

E-mail address: [edbianchi@unibz.it](mailto:edbianchi@unibz.it) (E. Bianchi).



**Fig. 1.** ProfVLM architecture. Multi-view video is encoded by a frozen TimeSformer, fused via a trainable AttentiveGatedProjector, and decoded by a LoRA-tuned SmolLM2 to generate both proficiency label and commentary.

gating, aligning video representations with the LLM embedding space.

To the best of our knowledge, ProfVLM is the first vision-language model to perform multi-view proficiency estimation entirely within an autoregressive generative framework, without relying on a dedicated classification head. The results show how the proposed ProfVLM improves over baselines in multi-view accuracy, while using 20x fewer parameters (5.3M vs. 121M), 2x fewer frames (8 vs. 16), and 60% reduction in training time (6 vs. 15 epochs). Compared to recently published works, our model again achieves better accuracy using 5x fewer parameters (5.3M vs. 27M) and 2-4x fewer frames (8 vs. 16/24/32). This efficiency highlights the benefits of integrating language into skill assessment from multiple visual viewpoints. Additionally, ProfVLM introduces high-quality feedback generation (BERTScore F1 of 85.53), a feature not present in previous approaches, which provides actionable insights for users.

## 2. Background and related works

The intersection of computer vision and sports analysis continues to expand across multiple domains, from team sports officiating systems that leverage multi-view analysis for foul detection (Held et al., 2025) to broadcast enhancement and performance analytics (Gade et al., 2024; Fujii, 2025). These applications demonstrate the growing integration of AI technologies in sports contexts, highlighting the need for robust and interpretable assessment systems that can operate across different skill levels and sporting disciplines.

Action recognition remains a highly active and evolving research area, with comprehensive surveys highlighting its rapid progress across diverse domains and methodologies (Zhao et al., 2024; sedaghati et al., 2025). Recent advances have explored temporal action segmentation in sparse action scenarios (Liu et al., 2024), multi-scale feature fusion with video-text constraints (Tian et al., 2025), graph neural networks for child activity recognition (Mohottala et al., 2025), and skeleton-enhanced spatio-temporal networks for sport action recognition (Bianchi and Lanz, 2025). Domain-specific applications have demonstrated the importance of tailored tracking solutions, as shown in skiing performance analysis where specialized tracking algorithms have been developed to handle the unique challenges of winter sports environments (Dunnhofer et al., 2024).

Moving beyond classical action recognition tasks, Action Quality Assessment (AQA) and Proficiency Estimation target the evaluation of how well an action is performed. Early deep learning approaches relied on video-only inputs and direct score regression using 3D CNNs — see e.g. Tran et al. (2015, C3D), and Carreira and Zisserman (2017,

I3D) — or transformer-based models — see e.g., Liu et al. (2021, VST) — often optimized for continuous or discrete skill labels. More recent approaches employ spatio-temporal feature extraction and deep learning or self-supervised and transfer learning for action assessment in rehabilitation (Kryeem et al., 2025; Kourbane et al., 2025). Although effective, these models offer limited interpretability and require domain-specific supervision (Zhou et al., 2024).

Multi-modal AQA and Proficiency Estimation enhance robustness and expressiveness by leveraging complementary modalities such as RGB, optical flow, depth, pose, and audio. Audio-visual fusion models such as Skating-Mixer (Xia et al., 2022) adopt hierarchical integration strategies to exploit temporal alignment and rhythm. Others use RGB-depth pairs for industrial tasks (Bianchi and Lanz, 2024) or combine skeletal motion with visual input to capture nuanced kinematics in sport performance (Bianchi and Lanz, 2025; Du et al., 2015; Duan et al., 2022). Recent work has also explored domain-specific pose estimation challenges, such as extending human skeleton estimation to include particular sports equipment (Martinelli et al., 2024). Recent benchmarks such as EgoExoLearn (Huang et al., 2024) and EgoExo4D (Grauman et al., 2024) underscore the value of multi-view egocentric and exocentric perspectives. However, the baselines proposed in these works demonstrate that the task is far from trivial, emphasizing the critical need for specialized architectures and tailored training objectives. Parameter-efficient architectures like SkillFormer (Bianchi and Liotta, 2025b) have further advanced multi-view proficiency estimation by using cross-attention mechanisms, while physiological signal extraction approaches such as EgoPulseFormer (Braun et al., 2025) demonstrate how incorporating heart rate data from egocentric cameras can significantly improve skill assessment. Recent advances in temporal sampling strategies have demonstrated that effective skill assessment requires domain-specific adaptations, as different sports and activity categories require distinct temporal parameters to capture fundamental movement patterns while preserving temporal continuity (Bianchi and Liotta, 2025a). While prior works incorporate language components or narrative evaluation (Parmar and Morris, 2019; Zhang et al., 2024), they maintain separate discriminative heads for score or label prediction. In contrast, our approach embeds proficiency prediction directly within the generative language modeling objective, unifying classification and explanation into a single autoregressive process.

In parallel, Vision-Language Models (VLMs) have been extended from image-text to video-text alignment. VideoBERT (Sun et al., 2019) pioneered joint video-text modeling by learning bidirectional joint distributions over sequences of visual and linguistic tokens derived from vector quantization. Image-focused VLMs like LLaVA (Liu et al., 2023) established strong baselines for visual instruction tuning and MiniGPT-4 (Zhu et al., 2023) demonstrated effective vision-language alignment

with frozen image encoders. VideoLLaMA (Zhang et al., 2023) introduced instruction-tuned audio-visual language modeling with dual-branch architecture for comprehensive video understanding, while LLaMA-VID (Li et al., 2024) focused on efficient video comprehension through temporal token compression. Video-ChatGPT (Maaz et al., 2024) and Video-LLaVA (Lin et al., 2024) further unified image and video modalities but initially lacked robust temporal grounding. Recent approaches include VALOR (Liu et al., 2025), which improved video-text alignment through contrastive learning, and SmolVLM (Marafioti et al., 2025), a compact model that achieves state-of-the-art performance for its memory footprint while maintaining competitive multi-modal capabilities. Advanced models such as Video-STaR (Zohar et al., 2024a) introduced instruction-tuning with labeled video corpora, while VideoAgent (Wang et al., 2024) and Apollo (Zohar et al., 2024b) framed video understanding as decision-making, improving long-form comprehension with efficient token compression, curriculum learning, or summarization tokens.

Computational efficiency remains a key consideration in both image- and video-based analysis, with recent advances combining knowledge distillation with tensor decomposition (Meneghetti et al., 2025) and applying pruning to transformer architectures to maintain accuracy while reducing computational costs (Li et al., 2025). Such efficiency considerations are particularly important for real-time applications in sports analysis and skill assessment, where rapid feedback can be crucial for training and performance improvement.

Our proposed ProfVLM bridges generative VLMs and skill assessment by treating proficiency estimation as a conditional language generation task. Built on frozen video transformers and LoRA-tuned language models, it jointly predicts structured proficiency levels and generates expert-like feedback.

### 3. Proposed methodology

We introduce ProfVLM, a multimodal vision-language model for skill assessment from videos captured from multiple perspectives. The architecture comprises a pre-trained TimeSformer (Bertasius et al., 2021) for visual encoding, an AttentiveGatedProjector which fuses multi-view features and projects them into the language embedding space, and a parameter-efficient adaptation of the open-source SmoLM2-135M-Instruct (Allal et al., 2025), fine-tuned using LoRA (Hu et al., 2021). Fig. 1 depicts the complete framework.

#### 3.1. Problem formulation

The task of demonstrator proficiency estimation involves analyzing multi-view video recordings of a subject performing a task, with the goal of assessing their skill level and providing qualitative feedback. Each sample consists of time-synchronized video streams from multiple viewpoints, specifically including one egocentric (first-person) and up to four exocentric (third-person) perspectives. Given this input, the objective is twofold: (1) predicting the demonstrator’s proficiency level as one of four predefined categories (Novice, Early Expert, Intermediate Expert, or Late Expert), and (2) generating a natural language comment describing their performance.

Formally, for a given set of  $n$  synchronized video streams  $V_1, V_2, \dots, V_n$ , the model must produce a textual feedback  $c \in \mathcal{T}$ , where  $\mathcal{T}$  denotes the space of valid commentary. The discrete proficiency label  $y \in \{0, 1, 2, 3\}$  is implicitly generated as part of this feedback rather than being predicted through a dedicated classification head.

#### 3.2. Video feature extraction

We employ a TimeSformer backbone (Bertasius et al., 2021), pre-trained on Kinetics-600 (Carreira and Zisserman, 2017), to extract spatio-temporal representations from video. Our choice of TimeSformer is motivated by two key factors: first, it serves as the established baseline architecture for evaluation in the EgoExo4D dataset (Grauman et al., 2024), enabling direct comparison with existing benchmarks; second, its divided space-time attention mechanism is specifically designed for video understanding, making it inherently superior to image-focused architectures that lack temporal modeling capabilities for proficiency estimation tasks. We retain the model frozen throughout training, without any fine-tuning or adaptation.

For multi-view scenarios, we employ parallel processing by flattening the batch and view axes, enabling simultaneous feature extraction across all synchronized video streams through our shared frozen backbone. The extracted features are subsequently reshaped to maintain distinct view-specific representations before entering the fusion and projection module, preserving the view separation critical for effective cross-view integration.

Inputs are fixed to 8-frame clips, matching the TimeSformer pre-training configuration. Preserving this format ensures compatibility with pretrained temporal attention patterns and avoids introducing additional trainable parameters, aligning with our efficiency objectives. Moreover, increasing frames yields diminishing returns: the official TimeSformer reports only a 3.9% Acc@1 gain when moving from 8 to 96 frames on Kinetics-600 despite a 12× increase in input size.<sup>1</sup> Despite using fewer frames, our architecture achieves state-of-the-art performance compared to baselines with 16–32 frames, leveraging more effective feature utilization and synchronized multi-view sampling to maintain consistent semantics across perspectives.

#### 3.3. AttentiveGatedProjector: Multi-view vision-language alignment

We design the AttentiveGatedProjector, an architecture tailored for integrating multi-view video representations and aligning them with the input space of large language models. Though bearing resemblance to a concurrent work (Bianchi and Liotta, 2025b) which targets classification, our module is instead optimized for joint vision-language modeling. While both approaches employ attention-based fusion mechanisms, they operate in fundamentally different learning paradigms: SkillFormer implements a discriminative classification framework that maps multi-view videos through feature fusion and classification heads to discrete labels, whereas our AttentiveGatedProjector enables a generative vision-language framework that projects multi-view videos through vision-language alignment into language models for natural language explanations with embedded labels.

This difference necessitates distinct architectural components and optimization objectives. Unlike discriminative fusion modules that optimize for classification, our projector is designed for video-language integration and includes statistical distribution matching between visual and textual embeddings through learnable normalization parameters, cross-modal dimensional projection to match the target language model’s embedding space, and preservation of semantic coherence across autoregressive generation sequences—requirements that are absent in classification-only architectures. The full architecture is illustrated in Fig. 2 and detailed below.

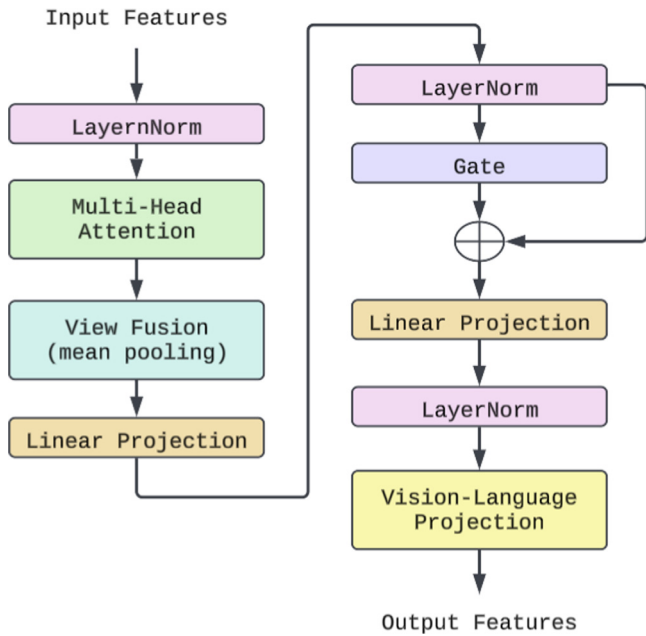
**Independent View Normalization:** Layer normalization is applied to each view (e.g., ego- or exo-centric) independently. This ensures that statistical disparities across views are normalized early, reducing modality-specific biases before fusion.

<sup>1</sup> Results from official implementation: <https://github.com/facebookresearch/TimeSformer>

**Table 1**

Prompt format used for training ProfVLM. During inference, only the System and User messages are provided, and the model generates both the proficiency level classification and explanatory commentary in a unified text response.

Role	Content
System	You are a visual agent for human performance analysis.
User	Here are 8 frames sampled from a video: < video_start >< video >< video_end >. Given this video, analyze the proficiency level of the subject.
Assistant	Proficiency Level: Intermediate Expert. Proficiency Commentary: The subject demonstrates smooth and controlled movements throughout the sequence, with clear intent and minimal hesitation.



**Fig. 2.** AttentiveGatedProjector architecture. Multi-view input features are initially normalized and combined using multi-head attention. The aggregated representation is then refined through a feed-forward block with residual paths and a trainable gating unit. Finally, the output undergoes linear projection, normalization, and alignment to the language model space via learned scaling and shifting parameters.

**Cross-View Multi-Head Attention:** To learn inter-view dependencies, we employ a multi-head attention mechanism where each view can contextualize its representation based on others through self-attention operations. This step flexibly captures complementary cues among perspectives, with attention weights dynamically determining the relevance of each view for the current sample.

**Learnable Gating Mechanism:** The fused features are passed through a gated transformation layer comprising a feed-forward network followed by a learnable sigmoid gate. This gate modulates each hidden unit independently through element-wise multiplication, enhancing discriminative signals while suppressing irrelevant or redundant information. The gating mechanism is particularly crucial for generative tasks where noise suppression affects text coherence.

**Vision-Language Projection:** Finally, the fused video features undergo a two-stage projection process specifically designed for vision-language integration. The features are first mapped to the target language model dimensionality, then undergo statistical normalization that centers them to zero mean, scales to unit variance, and applies learnable parameters to match the statistical distribution of language model token embeddings. This statistical alignment is essential for maintaining semantic coherence during text generation.

### 3.4. Tokenization and LLM prompt structure

To enable conditional generation, we extend the tokenizer vocabulary with three special tokens: <|video\_start|>, <|video\_end|>, and <|video|>. These delimiters explicitly mark the location of the visual input and allow the language model to attend to the fused video features in a controlled way. Importantly, our model performs both tasks — proficiency classification and commentary generation — entirely through language modeling. That is, the discrete proficiency label is not produced via a separate classification head but is instead embedded as part of the autoregressively generated textual response.

We adopt a structured prompt composed of system, user, and assistant messages in a chat-style format, as shown in Table 1. Our approach follows the chat template used during the pretraining and instruction-tuning of the open-source SmoLM2-135M-Instruct (Allal et al., 2025) to maintain consistency with the model’s expected input structure. The system message defines the model’s role; the user message includes the visual input placeholder bracketed by the special tokens; and the assistant message contains the target label and commentary. During training, all three messages are used to condition and supervise generation. During inference, only the system and user messages are provided as input, and the model is required to generate the assistant message, producing both the predicted proficiency label and the natural language feedback as a textual output.

The structured output format is provided by the training instance targets, allowing the model to progressively learn the required output schema during supervised fine-tuning. Since the inference prompt does not explicitly enforce the template, schema compliance emerges during fine-tuning: the ablations reported in Section 5.4.4 show that before fine-tuning the model does not follow the required output format, but after only a single epoch it achieves full compliance for the label field and near-complete compliance for the commentary field. At convergence, both fields reach 100% compliance, indicating that the autoregressive decoder reliably learns to generate template-aligned outputs under structured supervision. For deployment scenarios requiring strict guarantees, constrained decoding strategies<sup>2</sup> can be applied to enforce schema compliance.

## 4. Experimental setup

### 4.1. Dataset

We conduct experiments on the EgoExo4D dataset (Grauman et al., 2024), which contains over 1200 h of aligned egocentric and exocentric video from 740 participants across 123 locations. The dataset spans eight activities — cooking, health, bike maintenance, music, basketball, bouldering, soccer, and dance — captured in diverse real-world settings.

For the demonstrator proficiency estimation benchmark, six domains are used: cooking, music, basketball, bouldering, soccer, and

<sup>2</sup> Outlines: <https://github.com/dottxt-ai/outlines>

**Table 2**

Comparison of original and experimental dataset statistics for train and validation splits. Our train set is larger because each sample includes multiple expert commentary annotations. The validation set is kept unchanged to ensure a fair comparison.

Scenario	Official		Experiment	
	Train	Val	Train	Val
Basketball	576	105	1104	105
Cooking	83	20	381	20
Dancing	408	124	1220	124
Music	149	39	431	39
Bouldering	620	162	923	162
Soccer	68	16	179	16
Total	1904	466	4238	466

dance. Each sample consists of one first-person video (Project Aria glasses (Engel et al., 2023)) and up to four third-person views (GoPro cameras), alongside synchronized audio, eye gaze, 3D pose, and language annotations. Notably, the proficiency label distribution is skewed toward higher skill levels (Intermediate and Late Experts), reflecting the dataset’s intentional focus on recruiting participants with advanced abilities.

The dataset provides two key supervision signals (Fig. 3): a discrete proficiency level (*Novice*, *Early Expert*, *Intermediate Expert*, or *Late Expert*) and a free-form expert commentary from domain specialists. These commentaries were originally provided as verbal feedback by domain experts and subsequently transcribed from raw audio without manual curation or editing. This approach results in commentary text that preserves the evaluation style of domain experts, including spontaneous observations, colloquial language, and the natural flow of expert reasoning. However, it also introduces potential challenges such as transcription errors, incomplete sentences, repetitions, and varying levels of detail across commentaries.

We use all videos with proficiency estimation labels from the dataset, adhering to the official benchmark’s training and validation splits. As in prior work (Braun et al., 2025), we train on the official training set and evaluate on the full, held-out validation set to ensure comparability.

Our training set contains more samples than the official benchmark because each unique video-commentary pair is treated as a separate training instance. Classification labels are explicitly integrated into the input text as prompts (see Section 3.4 and Table 1), allowing the model to leverage textual and label information jointly. This provides richer supervision compared to classification baselines that rely solely on discrete labels, but also introduces a more challenging generative objective: the model must produce consistent proficiency predictions across varied linguistic contexts rather than mapping to fixed label embeddings. Table 2 summarizes the dataset statistics, highlighting the expanded training size while preserving the original validation set for fair evaluation.

#### 4.2. Implementation details

We fine-tune ProfVLM in three view configurations: Ego (1 ego-centric view), Exos (4 exocentric views), and Ego+Exos (all 5 views combined). Our training protocol consists of 6 epochs with a batch size of 32, using a cosine annealing learning rate schedule starting at  $3e-4$  with 1 warmup epoch. All experiments run on a single NVIDIA A100 GPU (80 GB).

For video processing, we employ a frozen TimeSformer (Bertasius et al., 2021) pre-trained on Kinetics-600 (Carreira and Zisserman, 2017). From each video, we extract 8 uniformly distributed frames from a randomly positioned temporal window, applying a sampling rate of 8 to ensure temporal coverage across variable-length clips. Each

frame is pre-processed by resizing its shortest edge to 224 pixels, center-cropping to  $224 \times 224$ , rescaling to  $[0, 1]$ , and normalizing with mean  $[0.45, 0.45, 0.45]$  and standard deviation  $[0.225, 0.225, 0.225]$ .

For language modeling, we use the open-source SmoLLM2-135M-Instruct (Allal et al., 2025), fine-tuned via Low-Rank Adaptation (LoRA) (Hu et al., 2021) with rank  $r = 8$  and scaling factor  $\alpha = 32$ . The AttentiveGatedProjector variant implements a hidden size of 1024 with 4 attention heads, while the simpler MLP variant uses a hidden size of 576 to match the LLM’s input dimension.

To further validate the contributions of our approach, we conduct ablation studies (Section 5.4) examining (1) architectural components, (2) supervision modality and volume, and (3) prompt design choices.

#### 4.3. Evaluation metrics

We employ a comprehensive evaluation framework that addresses both the classification and natural language generation aspects of our model. For classification, we use accuracy as the primary metric since it is the standard measure in the official EgoExo4D benchmark (Grauman et al., 2024), ensuring direct comparability with existing methods. Given the class imbalance present in the dataset, which can affect accuracy scores, we additionally report F1-score to provide a more robust evaluation.

### 5. Results


We evaluate ProfVLM on the task of multimodal proficiency estimation using the EgoExo4D (Grauman et al., 2024) benchmark. The following sections provide a detailed analysis of our model’s performance: the rationale behind our baseline selection (Section 5.1), overall accuracy and efficiency (Section 5.2), per-scenario accuracy (Section 5.3), and qualitative evaluation of the generated textual feedback (Section 5.5).

#### 5.1. Baseline selection rationale

Before presenting our results, we clarify our choice of baselines. We focus our comparison on established multi-view proficiency estimation methods rather than general-purpose vision-language models (VLMs) for several key reasons. First, our approach differs fundamentally from general-purpose VLMs by reformulating proficiency estimation as a unified generation task with specialized multi-view fusion mechanisms—capabilities that general-purpose VLMs inherently lack. Second, adapting general-purpose VLMs to our task would require substantial architectural modifications (e.g., multi-view attention mechanisms, task-specific training protocols), making direct comparison unfair.

To validate this approach, we tested general-purpose vision-language models on the EgoExo4D (Grauman et al., 2024) benchmark using zero-shot evaluation with basic frame concatenation (no fusion mechanisms). We chose three representative models with different scales and architectures. First, PandaGPT (Su et al., 2023) serves as our larger-scale baseline—a multimodal model that combines an ImageBind encoder (Girdhar et al., 2023) with Vicuna-13B (Chiang et al., 2023) and natively handles video input. Second, we evaluated SmoVLM and SmoVLM2 (Marafioti et al., 2025), which use a SigLip-93M encoder (Zhai et al., 2023) paired with SmoLLM2-135M (Allal et al., 2025). This second baseline is particularly important because it shares the same language model backbone as our approach, allowing us to assess whether using SmoLLM2-135M alone could achieve competitive results without our specialized architectural components.

Results are reported in Table 3 and confirm that general-purpose VLMs perform significantly worse than our approach due to their lack of proper multi-view fusion mechanisms and task-specific training. Therefore, we focus our evaluation against established multi-view methods like SkillFormer (Bianchi and Liotta, 2025b) and EgoPulseFormer (Braun et al., 2025), which represent the current state-of-the-art for this specific task on the EgoExo4D dataset.



**Proficiency Level:** Early Expert. **Proficiency Commentary:** This body position is really good. This is kind of like a twist lock technique. She's got left hand on, right foot on, so she's opposite posed. She's got that knee pretty high, so she really is kind of sitting up over that foot. She's able to apply that pressure down straight through that foothold. And now the left foot is kind of smearing into the wall, staying tensioned. And that's going to let her kind of push herself up to the hold. Again, it does a good job walking the feet through, but again, this non-tensioned foot is just loose and flopping around. It's actually kind of pulling her away from the wall.

**Fig. 3.** Synchronized frames from the five views (1 egocentric and 4 exocentric) of a single EgoExo4D test sample for the bouldering domain. The proficiency level and commentary apply to the entire multi-view sequence, demonstrating the multi-perspective skill assessment setup.

**Table 3**

Comparison with EgoExo4D baselines, with 0-shot general-purpose VLMs (Su et al., 2023; Marafioti et al., 2025), and with concurrent approaches which adopt the official training-validation splits: EgoPulseFormer (Braun et al., 2025), and SkillFormer (Bianchi and Liotta, 2025b). Results are reported for Ego, Exos, and Ego+Exos views. Metric is top-1 accuracy (%) as in the official benchmark (Grauman et al., 2024) for direct comparison. ProfVLM+AGP surpasses state-of-the-art accuracy in the Ego+Exos setting using significantly less parameters and input frames. Best in bold, second-best underlined.

Method	Pretrain	Ego	Exos	Ego+Exos	Frames	Train Params	Epochs
Random	–	24.9	24.9	24.9	–	–	–
Majority-class	–	31.1	31.1	31.1	–	–	–
TimeSFormer	–	42.3	40.1	40.8	16	121M	15
TimeSFormer	K400	42.9	39.1	38.6	16	121M	15
TimeSFormer	HowTo100M	<b>46.8</b>	38.2	39.7	16	121M	15
TimeSFormer	EgoVLP	44.4	40.6	39.5	16	121M	15
TimeSFormer	EgoVLPv2	<u>45.9</u>	38.0	37.8	16	121M	15
PandaGPT (0-shot)	ImageBind + Vicuna	23.5	25.6	25.8	–	52M	–
SmolVLM (0-shot)	–	34.2	30.7	31.1	–	256M	–
SmolVLM2 (0-shot)	–	26.8	30.5	32.2	–	256M	–
EgoPulseFormer	EgoPPG-DB	45.3	35.9	42.4	16	121M	15
SkillFormer	K600	<u>45.9</u>	<b>46.3</b>	<u>47.5</u>	16/24/32	14/20/27M	4
<b>ProfVLM (MLP)</b>	K600 + SmolLM2	38.9	44.2	42.9	8	2.2M	6
<b>ProfVLM (AGP)</b>	K600 + SmolLM2	44.2	<u>45.1</u>	<b>48.2</b>	8	5.3M	6

**Table 4**

Accuracy and F1 comparison between ProfVLM variants (MLP vs. AGP) across different view settings. All models share identical hyperparameters: learning rate =  $3e-4$ , LoRA rank = 8, scaling = 32, 8 frames per video, 6 training epochs, batch size = 32. Bold: best results; underlined: second-best.

Method	View	Hid	Heads	Train Params	Acc (%)	F1 (%)
ProfVLM (MLP)	Ego	576	–	2.2M	38.9	35.4
	Exos	576	–	2.2M	44.2	36.9
	Ego+Exos	576	–	2.2M	42.9	36.0
ProfVLM (AGP)	Ego	1024	4	5.3M	44.2	38.5
	Exos	1024	4	5.3M	<u>45.1</u>	<u>38.9</u>
	Ego+Exos	1024	4	5.3M	<b>48.2</b>	<b>44.4</b>

## 5.2. Overall classification accuracy and efficiency

Table 3 presents classification accuracy (%) across egocentric (1 Ego), exocentric (4 Exos), and combined (1 Ego + 4 Exos) view configurations. Baseline results are taken directly from the EgoExo4D benchmark (Grauman et al., 2024), where models are trained exclusively on video inputs and perform classification using a dedicated output head. In contrast, ProfVLM adopts a vision-language formulation: it receives both visual input and textual supervision, and infers

the proficiency level directly within the generated response, without relying on an explicit classification head.

Our best-performing variant, ProfVLM with AttentiveGateProjector (AGP), surpasses the state-of-the-art with 48.2% accuracy in the combined Ego+Exos setting, outperforming the established EgoExo4D baselines as well as two contemporary works, SkillFormer (Bianchi and Liotta, 2025b) and EgoPulseFormer (Braun et al., 2025), that adopt the same standardized dataset splits. This alignment ensures a fair comparison under identical evaluation protocols. The consistent gains

**Table 5**

Accuracy (%) comparison across six scenarios using egocentric (Ego), exocentric (Exos), and combined (Ego+Exos) views. Baseline from [Grauman et al. \(2024\)](#). *Maj.* refers to the majority-class baseline. ProfVLM is evaluated with both a simple MLP projector and our AttentiveGateProjector (AGP). Best results are in bold; second-best are underlined.

Scenario	Maj.	Baseline			ProfVLM (MLP)			ProfVLM (AGP)		
		Ego	Exos	Ego+Exos	Ego	Exos	Ego+Exos	Ego	Exos	Ego+Exos
Basketball	36.19	<u>51.43</u>	52.30	<b>55.24</b>	34.00	38.00	40.00	36.00	33.00	41.00
Cooking	50.00	45.00	35.00	35.00	41.00	<u>51.00</u>	36.00	31.0	<b>56.00</b>	<u>51.00</u>
Dancing	51.61	<u>55.65</u>	42.74	42.74	40.02	<u>48.97</u>	53.03	51.41	53.85	<b>60.35</b>
Music	58.97	46.15	<u>69.23</u>	56.41	58.89	56.26	55.26	<b>72.05</b>	61.53	56.26
Bouldering	0.00	25.31	<u>17.28</u>	17.28	33.08	<u>38.11</u>	34.33	37.48	37.48	<b>38.74</b>
Soccer	62.50	56.25	<u>75.00</u>	<u>75.00</u>	69.75	69.75	44.75	57.25	<b>76.00</b>	69.75

across these benchmarks underscore the robustness of our approach. [Table 4](#) shows that ProfVLM (AGP) consistently outperforms its MLP counterpart across all configurations. The gain is most notable in the Ego+Exos setting, where AGP achieves 48.2% accuracy and 44.4% F1, compared to 42.9% and 36.0% for the MLP—a relative improvement of 12% in accuracy and 23% in F1. This highlights the effectiveness of our structured, attention-based cross-view fusion.

ProfVLM is also considerably more efficient: it requires just 8 frames per video compared to 16–32 in prior approaches, and utilizes only 5.3M trainable parameters versus 14–27M in SkillFormer (5x reduction) and 121M in TimeSFormer-based models (20x reduction). Furthermore, our model converges in just 6 training epochs compared to 15 for TimeSFormer and EgoPulseFormer variants, representing a 60% reduction in training time. These findings validate our vision-language formulation, which unifies classification and explanation within a compact, interpretable framework.

### 5.3. Cross-domain accuracy analysis

[Table 5](#) breaks down performance across six diverse scenarios in the EgoExo4D dataset. Our analysis reveals that ProfVLM outperforms the baselines in five out of six scenarios, with Basketball being the notable exception. Specifically, ProfVLM (AGP) achieves the highest accuracy in Dancing (60.35%), Bouldering (38.74%), Cooking (56.00%), Soccer (76.00%), and Music (72.05%), demonstrating remarkable cross-domain effectiveness.

Several interesting patterns emerge from this analysis. First, the benefit of combining views varies by scenario: in Dancing, the Ego+Exos configuration significantly outperforms individual approaches (60.35% vs. 51.41%/53.8%), while in Soccer the exocentric view provides the strongest signal (76.00%). Second, the structured fusion mechanism in AGP consistently enhances performance over the simpler MLP projector, particularly in visually complex scenarios like Dancing and Bouldering. These domain-specific results highlight ProfVLM’s ability to leverage complementary information from multiple viewpoints, where activities requiring spatial relationships and body positioning analysis (Dancing, Bouldering) benefit from integrated egocentric and exocentric information, while activities with standardized visual cues visible from specific angles (Music) perform equally well with single-view configurations.

The Basketball scenario warrants deeper investigation as it represents the only domain where our approach underperforms the baseline (41.00% vs. 55.24%). This challenge likely stems from basketball’s unique proficiency assessment requirements that differ fundamentally from other activities ([Pan et al., 2025](#)). Recent research has established basketball as one of the most analytically complex sports, with the BASKET dataset ([Pan et al., 2025](#)) demonstrating that state-of-the-art video models achieve accuracy barely above random chance despite advances in computer vision. Unlike domains where proficiency manifests through broader movement patterns, basketball skills depend on highly precise technical details: wrist angle during shooting, fingertip control in dribbling, and subtle postural adjustments ([Sarlis and Tjortjis, 2020](#)). Studies have shown that basketball analysis requires

**Table 6**

Progressive ablation of AttentiveGatedProjector components using Ego+Exos views. Hyperparameters follow [Section 4.2](#). Percentage increase (%  $\Delta$ ) is computed relative to the previous row.

Projector	Gating	Adapt. Norm	Params	Acc (%)	% $\Delta$
MLP	–	–	2.2M	42.9	–
AGP (base)	✗	✗	4.2M	44.2	+3.0%
+ Gating	✓	✗	5.3M	45.5	+2.9%
+ Gating + Norm	✓	✓	5.3M	<b>48.2</b>	+5.8%

sophisticated spatio-temporal modeling to capture the complex interactions between players, ball dynamics, and tactical formations ([Zhang et al., 2025](#); [Lingrui et al., 2025](#)). Our current temporal sampling strategy and spatial resolution may be insufficient to capture these fine-grained discriminative features that distinguish skill levels, particularly given basketball’s requirements for spatial accuracy and high temporal resolution ([Zhang et al., 2025](#)).

Additionally, the attentive fusion mechanism, while effective for integrating complementary information across views in other domains, may struggle with basketball’s localized critical regions (hands, ball trajectory) where relevant signals could be suppressed by global attention patterns. The baseline’s superior performance suggests that simpler aggregation strategies may be more robust when discriminative features are highly localized and require precise temporal alignment, pointing toward the need for domain-adaptive attention mechanisms and higher-resolution temporal modeling for technical sports analysis in future work.

### 5.4. Ablation studies

We conduct systematic ablations to validate our AGP design choices ([Section 5.4.1](#)), evaluate the impact of supervision modality and training data volume ([Section 5.4.2](#)), and examine an alternative prompt formulation ([Section 5.4.3](#)).

#### 5.4.1. Projector architecture

[Table 6](#) reports the progressive ablation of the AttentiveGatedProjector. Starting from the MLP baseline (42.9%), we introduce: (1) multi-head attention for cross-view fusion, which corresponds to switching from the MLP to the base AGP; (2) learnable gating for feature selection; and (3) adaptive normalization for vision-language alignment. Each component contributes to the final performance (48.2%), with the full AGP achieving a 12.3% percentage increase over the MLP baseline. This validates that all three mechanisms are essential for effective multi-view proficiency estimation.

#### 5.4.2. Supervision modality and volume

We investigate how different supervision strategies affect proficiency estimation performance ([Table 7](#)).

**Table 7**

Supervision modality and volume ablation using Ego+Exos views with the AGP projector. Discriminative baselines are reported for reference. Hyperparameters follow Section 4.2.

Training condition	Samples	Acc (%)	F1 (%)
Label-only generation	1904	45.0	36.7
Single random commentary	1904	40.7	35.9
Full (all commentaries)	4238	<b>48.2</b>	<b>44.4</b>
<i>Discriminative baselines:</i>			
SkillFormer (Bianchi and Liotta, 2025b)	1904	47.5	–
EgoPulseFormer (Braun et al., 2025)	1904	42.4	–
TimeSformer (best) (Grauman et al., 2024)	1904	40.8	–

**Table 8**

Impact of prompt structure using Ego+Exos views. Both variants use the full AGP architecture and multiple commentaries. Hyperparameters follow Section 4.2.

Prompt structure	Acc (%)	F1 (%)
Proficiency Commentary → Proficiency Level	45.3	41.5
Proficiency Level → Proficiency Commentary	<b>48.2</b>	<b>44.4</b>

*Label-only generation.* We train ProfVLM to generate only the discrete proficiency label (e.g., “Intermediate Expert”) without commentary, using the official train set (1904 samples). This configuration achieves 45.0% accuracy, demonstrating that the generative vision-language formulation remains competitive with baselines even when limited to classification-only supervision.

*Single commentary per video.* Training with exactly one randomly sampled commentary per video yields 40.7% accuracy, which is below the label-only configuration. This counter-intuitive result reveals an important training dynamic: the commentary generation task introduces loss imbalance. Since loss is computed over all generated tokens, the commentary dominates the gradient compared to the single label token. With only 1904 training samples, the model overfits to generating plausible commentary at the expense of label prediction accuracy. This phenomenon does not occur in the label-only setting (single prediction token) or in the multiple commentary setting where increased data volume and label-commentary diversity provide implicit regularization.

*Multiple commentaries per video.* Our full training regime leverages all available expert commentaries per video (4238 samples total, 2.2x data augmentation). This achieves 48.2% accuracy through: (1) increased training data volume, and (2) implicit regularization: observing the same video with multiple valid commentaries teaches the model that proficiency labels remain invariant to commentary style, strengthening label prediction robustness. The gap between single commentary (40.7%) and full supervision (48.2%) demonstrates that effective commentary integration requires both sufficient data volume to overcome loss imbalance and diversity to provide regularization benefits.

#### 5.4.3. Prompt structure

Our default prompt format generates the proficiency label before the commentary (Table 1). To investigate whether this ordering affects performance, we evaluate an alternative format where the commentary is generated first, followed by the label.

Table 8 shows that generating the label first yields significantly better accuracy (48.2% vs. 45.3%). This 2.9% gap reveals important training dynamics in autoregressive vision-language models. When generating commentary-first, the label prediction at position  $T$  depends causally on all preceding tokens at position  $\{1, \dots, T-1\}$  through the attention mechanism. During training, the model conditions on ground-truth commentaries, but at inference must condition on its own generated text. Even when the generated commentary is semantically correct, lexical divergence from training distributions introduces spurious dependencies that harm label prediction.

**Table 9**

Format compliance results.

Model	Epochs	Label field (%)	Commentary field (%)
Base (no fine-tuning)	0	0.0	0.0
ProfVLM (early)	1	100.0	96.7
ProfVLM (final)	6	100.0	100.0

Additionally, the commentary-first ordering introduces gradient dilution: the label loss must backpropagate through 50–100 commentary tokens, weakening the classification signal with limited training data. In contrast, the label-first ordering establishes the discrete prediction directly from visual features before elaborating with natural language, avoiding spurious causal dependencies in the autoregressive chain. This design choice proves essential for maintaining strong classification performance while enabling interpretable feedback generation.

#### 5.4.4. Schema compliance

As shown in Table 1, our inference prompt does not provide an output template; therefore, adherence to the expected output structure must be learned during supervised fine-tuning. In Table 9 we show how schema compliance emerges during fine-tuning; compliance is measured by verifying both output JSON validity and the presence of the two schema-required keys (“Proficiency Level” and “Proficiency Commentary”).

It can be seen how the 0-epoch model does not follow the required template. After a single epoch, the model already achieves 100% compliance for the label field and 96.7% for the commentary field. At convergence, both fields reach 100% compliance. These results indicate that structured output behavior is progressively induced by the supervised training signal and stabilizes with continued optimization.

#### 5.5. Language generation quality

A key capability of ProfVLM is its ability to generate detailed, expert-like feedback in natural language that explains the reasoning behind proficiency assessments. This dual functionality — providing both classification scores and interpretable commentary — addresses a critical gap in automated skill evaluation where practitioners need actionable insights beyond numerical ratings. We evaluate this language generation capability using established text quality metrics as shown in Table 10. Since, to the best of our knowledge, no prior work addresses this task, we do not report comparisons with other models in this evaluation.

##### 5.5.1. Semantic alignment and BERTScore validity

Our approach achieves consistently high BERTScore (Zhang et al., 2020) F1 values (85.41–85.53) across all configurations, indicating strong semantic alignment with expert commentaries. For BERTScore, we use RoBERTa-large (Liu et al., 2019) as the underlying model, which captures semantic similarity beyond surface-level token matching through contextual embeddings. Importantly, the high BERTScore values reflect genuine semantic coherence rather than mere topical relevance—the metric effectively captures whether the generated feedback addresses the same technical aspects and maintains the domain-appropriate language patterns found in expert annotations.

These observations also clarify why the AGP and MLP variants yield similar text-generation quality. The language-modeling objective:

$$\mathcal{L} = - \sum_t \log P(\text{token}_t | \text{tokens}_{<t}, \text{video\_features})$$

creates a strong gradient imbalance (approximately 35:1 between commentary tokens and label tokens), naturally emphasizing fluent and semantically coherent commentary over fine-grained label discrimination. Furthermore, proficiency levels form a continuous spectrum,

**Table 10**

Evaluation of generated feedback quality for ProfVLM across different input views and projector architectures. BERTScore (F1) captures semantic similarity to expert commentaries, while METEOR and ROUGE assess lexical overlap. The variations in metric values are minimal across all configurations, suggesting consistent language generation quality regardless of view type or projector design.

Method	View	BERTScore (F1)	METEOR	ROUGE-L
ProfVLM (MLP)	Ego	85.41	18.07	15.55
	Exos	85.41	17.98	15.63
	Ego+Exos	85.51	18.22	15.49
ProfVLM (AGP)	Ego	85.41	18.06	14.47
	Exos	85.51	17.33	15.67
	Ego+Exos	85.53	18.23	15.65

**Table 11**

ProfVLM commentary outputs with corresponding BERTScore (F1), METEOR, and ROUGE-L evaluation metrics. We present three samples representing high, medium, and low quality predicted proficiency feedback.

Prediction	BERTScore (F1)	METEOR	ROUGE-L
<p><i>High-quality example (bouldering):</i> So here we can see the climber is going up to their right hand hold. And they have this really nice, pretty strong hold on this left side. But that’s not a good choice for this move because it would make them lose power in the same way as they are using it. So the climber was able to kind of get through with this one by just pushing off from the left foot and then putting down into the net so he could come out of his feet. The climbers got out of the grip and now they’re hitting this wall.</p>	87.14	36.22	20.22
<p><i>Medium-quality example (dancing):</i> So here we have the dancer performing a step up. This is really great because it’s so simple and quick, but also very effective. So she does this step up with her arms out to the sides. She’s turning around at that point of time. So it feels really nice to see her do this movement in a different direction. Here she steps through and into the space. This is what I like about this move as well. It’s kind of like a dance-step where you’re stepping forward.</p>	85.10	24.04	14.39
<p><i>Low-quality example (basketball):</i> Again, good job on that. She’s playing with her hips and just getting a little bit of power here. But again, you can see she’s got a lot more energy in her hands than she should be having. And she has to really focus on the ball right now, which is why we want to know that you’re throwing it straight down to your body. Right now, you have to hit this ball up high above your head so that when you release it, it will go off at that point. So I would definitely give this one another shot.</p>	82.87	5.25	7.59

where expert descriptors often apply to adjacent classes (e.g., “*controlled movement*” may characterize both Intermediate and Late-Expert performances). This structural property of the task enables the model to maintain high semantic alignment even when classification boundaries are mispredicted.

The difference between the two variants manifests instead in visual reasoning: the AGP module integrates multi-view information through cross-view attention and learnable gating, producing more discriminative features and achieving a +12.3% improvement in prediction accuracy over the MLP. Since both variants share the same language backbone (SmolLM2-135M with LoRA), their commentary fluency remains comparable, while AGP provides stronger visual grounding for proficiency estimation.

### 5.5.2. Lexical vs. Semantic metrics

The relatively modest METEOR (Lavie and Agarwal, 2007) (17.33–18.23) and ROUGE-L (Lin, 2004) (14.47–15.67) scores reflect the open-ended nature of the feedback generation task. Unlike classification or short-answer generation, our model produces multi-sentence technical commentary where semantic appropriateness is more critical than exact lexical overlap with reference text. This is particularly evident in Table 11, which shows that while BERTScore remains relatively

stable across quality levels (87.14–82.87), METEOR and ROUGE-L metrics show substantial degradation in lower-quality examples (from 36.22 to 5.25 for METEOR and 20.22 to 7.59 for ROUGE-L). This pattern suggests that ProfVLM preserves domain-appropriate language and conceptual alignment even when diverging from the specific phrasing of expert commentaries.

### 5.5.3. Dataset-driven feedback characteristics

The nature of our generated feedback reflects the characteristics of the EgoExo4D dataset’s expert commentaries, which primarily focus on technical observations and skill assessment rather than prescriptive coaching or instructional guidance. Since the original expert commentaries were transcribed from raw audio without manual curation, our model learns to emulate both the expert evaluation style and the naturalistic speech patterns present in the training data, including occasional incomplete sentences and informal language. Our model generates feedback that identifies technical aspects of performance (e.g., “grip positioning could be improved”, “timing of the movement shows good coordination”) rather than providing detailed corrective instructions. This alignment with the dataset’s annotation style explains both the semantic consistency captured by BERTScore and the technical, assessment-oriented nature of the generated text.

#### 5.5.4. Model efficiency considerations

Our choice of SmolLM2-135M as the language generation backbone represents an explicit trade-off between quality and deployment efficiency. With only 5.3M trainable parameters in the language component, ProfVLM maintains semantic coherence while enabling practical deployment in resource-constrained environments. The consistently high BERTScore values across configurations demonstrate that this efficiency-focused design does not compromise semantic alignment with expert-level technical commentary.

Overall, both quantitative metrics and qualitative analysis confirm that ProfVLM produces coherent and contextually appropriate commentary that maintains semantic alignment with expert annotations while reflecting the technical, observational style characteristic of the underlying dataset.

## 6. Limitations

Despite its promising performance, ProfVLM has some limitations that suggest directions for future work.

First, the model relies on language annotations during training, assuming that both proficiency labels and expert commentaries are available. In real-world scenarios, such supervision may be difficult to obtain at scale or may reflect subjective judgments that vary across annotators.

Second, while ProfVLM demonstrates strong performance in generating coherent and semantically aligned feedback, as confirmed by our automatic evaluation metrics, exploring additional dimensions of quality represents an exciting direction for future work. Although BERTScore effectively captures semantic similarity, aspects such as factual accuracy, pedagogical effectiveness, and actionability could be further assessed only through expert evaluation.

Third, the model generates the class label as part of free-form text. Although this enables seamless integration of classification and explanation, it introduces potential downsides—e.g., misformatted or ambiguous outputs may complicate downstream usage, especially in settings that require structured prediction or reliability guarantees.

Fourth, the current design assumes a fixed number of input views (e.g., Ego and Exos). Adapting the architecture to dynamically handle a variable number of camera perspectives or synchronize inputs with inconsistent frame counts is an open challenge that would improve generalizability.

Lastly, our implementation keeps the video encoder (TimeSFormer) frozen, which simplifies training and reduces computational cost but may limit domain adaptation. Fine-tuning the visual backbone or incorporating domain-specific motion priors could further enhance performance, particularly for tasks requiring fine-grained temporal understanding.

In addition, although ProfVLM leverages multi-view inputs, it primarily relies on high-level visual features, without explicitly modeling pose, body kinematics, or action-specific cues. Explicitly incorporating structured pose and motion representations can potentially provide a more fine-grained understanding of proficiency across views. Investigating how to integrate such structured signals within the vision-language framework represents an important direction for future work.

Addressing these limitations through improved supervision strategies, human-centered evaluation protocols, and more flexible fusion mechanisms represents a promising direction for advancing vision-language models for skill assessment and related multimodal evaluation tasks.

## 7. Discussion and conclusions

Our evaluation demonstrates that ProfVLM advances proficiency estimation through a unified vision-language framework, surpassing state-of-the-art classification accuracy (48.2% in Ego+Exos) on the EgoExo4D benchmark while generating semantically aligned textual

feedback. The model outperforms video-only baselines and contemporary methods (SkillFormer, EgoPulseFormer) by integrating three key innovations: (1) proficiency inference via language generation instead of explicit classification heads, (2) the AttentiveGateProjector (AGP) for structured cross-view fusion, and (3) a parameter-efficient design requiring only 5.3M trainable parameters and 6 training epochs—yielding up to 20×computational efficiency gains versus transformer-based alternatives.

The AGP mechanism proves critical for multi-view integration, particularly in spatially demanding scenarios (e.g., Dancing, Bouldering), where it outperforms MLP-based fusion by 12%–23% in accuracy and F1. Cross-domain analysis reveals scenario-dependent view utility: exocentric inputs dominate in standardized tasks (e.g., Soccer), while combined views excel in activities requiring spatial reasoning. This adaptability underscores ProfVLM’s capacity to balance complementary perspectives without architectural overcomplication.

Performance in the Basketball domain remains challenging across all approaches, likely reflecting basketball’s unique analytical complexity where proficiency depends on highly precise technical details that may exceed our current temporal and spatial resolution capabilities. Addressing these domain-specific challenges represents an important direction for future work.

Beyond classification, ProfVLM generates feedback with strong semantic alignment to expert annotations (BERTScore > 85), even when lexical overlap metrics (METEOR, ROUGE-L) are modest—a reflection of its focus on conceptual correctness over rigid template matching. Qualitative analysis confirms coherent explanations across proficiency levels, validating the framework’s dual strength in perception and interpretability.

In summary, ProfVLM establishes a new paradigm unifying classification with natural language reasoning. Its efficient design bridges automated evaluation with expert-like analysis, offering a transparent solution for multimodal proficiency estimation.

### CRedit authorship contribution statement

**Edoardo Bianchi:** Writing – review & editing, Writing – original draft, Validation, Software, Methodology, Conceptualization. **Jacopo Staiano:** Writing – review & editing, Writing – original draft, Validation, Supervision, Methodology, Conceptualization. **Antonio Liotta:** Writing – review & editing, Supervision, Methodology, Conceptualization.

### Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

### Acknowledgments

We acknowledge ISCR for awarding this project access to the LEONARDO supercomputer, owned by the EuroHPC Joint Undertaking, hosted by CINECA (Italy).

### Data availability

The dataset used in the research is named EgoExo4D (Grauman et al.). The official dataset paper can be found at <https://ieeexplore.ieee.org/document/10658224>. The official website is <https://ego-exo4d-data.org>.

## References

- Allal, L.B., Lozhkov, A., Bakouch, E., Blázquez, G.M., Penedo, G., Tunstall, L., Marafioti, A., Kydlíček, H., Lajarín, A.P., Srivastav, V., Lochner, J., Fahlgren, C., Nguyen, X.S., Fourrier, C., Burtenshaw, B., Larcher, H., Zhao, H., Zakka, C., Morlon, M., Raffel, C., von Werra, L., Wolf, T., 2025. SmolLM2: When smol goes big – data-centric training of a small language model. *arXiv:2502.02737*. URL: <https://arxiv.org/abs/2502.02737>.
- Bertasius, G., Wang, H., Torresani, L., 2021. Is space-time attention all you need for video understanding?. *arXiv:2102.05095*. URL: <https://arxiv.org/abs/2102.05095>.
- Bianchi, E., Lanz, O., 2024. Egocentric video-based human action recognition in industrial environments. In: Concli, F., Maccioni, L., Vidoni, R., Matt, D.T. (Eds.), *Latest Advancements in Mechanical Engineering*. Springer Nature Switzerland, Cham, pp. 257–267.
- Bianchi, E., Lanz, O., 2025. Gate-shift-pose: Enhancing action recognition in sports with skeleton information. In: *Proceedings of the Winter Conference on Applications of Computer Vision (WACV) Workshops*. pp. 1257–1264.
- Bianchi, E., Liotta, A., 2025a. PATS: Proficiency-aware temporal sampling for multi-view sports skill assessment. In: 2025 IEEE International Workshop on Sport, Technology and Research. STAR, pp. 1–6. <http://dx.doi.org/10.1109/STAR66750.2025.11264769>.
- Bianchi, E., Liotta, A., 2025b. SkillFormer: Unified multi-view video understanding for proficiency estimation. *arXiv:2505.08665*. URL: <https://arxiv.org/abs/2505.08665>.
- Braun, B., Armani, R., Meier, M., Moebus, M., Holz, C., 2025. egoPPG: Heart rate estimation from eye-tracking cameras in egocentric systems to benefit downstream vision tasks. *arXiv:2502.20879*. URL: <https://arxiv.org/abs/2502.20879>.
- Carreira, J., Zisserman, A., 2017. Quo vadis, action recognition? A new model and the kinetics dataset. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition. CVPR, pp. 4724–4733. <http://dx.doi.org/10.1109/CVPR.2017.502>.
- Chiang, W.-L., Li, Z., Lin, Z., Sheng, Y., Wu, Z., Zhang, H., Zheng, L., Zhuang, S., Zhuang, Y., Gonzalez, J.E., Stoica, I., Xing, E.P., 2023. Vicuna: an open-source chatbot impressing GPT-4 with 90%\* ChatGPT quality. <https://lmsys.org/blog/2023-03-30-vicuna/>.
- Du, Y., Wang, W., Wang, L., 2015. Hierarchical recurrent neural network for skeleton based action recognition. In: 2015 IEEE Conference on Computer Vision and Pattern Recognition. CVPR, pp. 1110–1118. <http://dx.doi.org/10.1109/CVPR.2015.7298714>.
- Duan, H., Zhao, Y., Chen, K., Lin, D., Dai, B., 2022. Revisiting skeleton-based action recognition. In: 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition. CVPR, pp. 2959–2968. <http://dx.doi.org/10.1109/CVPR52688.2022.00298>.
- Dunnhofer, M., Sordi, L., Martinel, N., Micheloni, C., 2024. Tracking skiers from the top to the bottom. In: 2024 IEEE/CVF Winter Conference on Applications of Computer Vision. WACV, pp. 8496–8506. <http://dx.doi.org/10.1109/WACV57701.2024.00832>.
- Engel, J., Somasundaram, K., Goesele, M., Sun, A., Gamino, A., Turner, A., Talatof, A., Yuan, A., Souti, B., Meredith, B., Peng, C., Sweeney, C., Wilson, C., Barnes, D., DeTone, D., Caruso, D., Valleroy, D., Ginpupalli, D., Frost, D., Miller, E., Mueggler, E., Oleinik, E., Zhang, F., Somasundaram, G., Solaira, G., Lanaras, H., Howard-Jenkins, H., Tang, H., Kim, H.J., Rivera, J., Luo, J., Dong, J., Straub, J., Bailey, K., Eckenhoff, K., Ma, L., Pesqueira, L., Schwesinger, M., Monge, M., Yang, N., Charron, N., Raina, N., Parkhi, O., Borschowa, P., Moulon, P., Gupta, P., Mur-Artal, R., Pennington, R., Kulkarni, S., Miglani, S., Gondi, S., Solanki, S., Diener, S., Cheng, S., Green, S., Saarinen, S., Patra, S., Mourikis, T., Whelan, T., Singh, T., Balntas, V., Baiyya, V., Dreewes, W., Pan, X., Lou, Y., Zhao, Y., Mansour, Y., Zou, Y., Lv, Z., Wang, Z., Yan, M., Ren, C., Nardi, R.D., Newcombe, R., 2023. Project aria: A new tool for egocentric multi-modal AI research. *arXiv:2308.13561*. URL: <https://arxiv.org/abs/2308.13561>.
- Fujii, K., 2025. Computer vision for sports analytics. In: *Machine Learning in Sports: Open Approach for Next Play Analytics*. Springer Nature Singapore, Singapore, pp. 21–57. [http://dx.doi.org/10.1007/978-981-96-1445-5\\_2](http://dx.doi.org/10.1007/978-981-96-1445-5_2).
- Gade, R., Merler, M., Thomas, G., Moeslund, T., 2024. The (Computer) vision of sports: Recent trends in research and commercial systems for sport analytics. In: Ahad, M., Mahbub, U., Turk, M., Hartley, R. (Eds.), *Computer Vision*. CRC Press, Netherlands, pp. 296–311. <http://dx.doi.org/10.1201/9781003328957-14>.
- Girdhar, R., El-Nouby, A., Liu, Z., Singh, M., Alwala, K.V., Joulin, A., Misra, I., 2023. ImageBind one embedding space to bind them all. In: 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition. CVPR, pp. 15180–15190. URL: <https://api.semanticscholar.org/CorpusID:258564264>.
- Grauman, K., Westbury, A., Torresani, L., Kitani, K., Malik, J., Afouras, T., Ashutosh, K., Baiyya, V., Bansal, S., Boote, B., Byrne, E., Chavis, Z., Chen, J., Cheng, F., Chu, F.J., Crane, S., Dasgupta, A., Dong, J., Escobar, M., Forigua, C., Gebreselasie, A., Hareesh, S., Huang, J., Islam, M.M., Jain, S., Khirodkar, R., Kukeja, D., Liang, K.J., Liu, J.W., Majumder, S., Mao, Y., Martin, M., Mavroudi, E., Nagarajan, T., Ragusa, F., Ramakrishnan, S.K., Seminar, L., Somayazulu, A., Song, Y., Su, S., Xue, Z., Zhang, E., Zhang, J., Castillo, A., Chen, C., Fu, X., Furuta, R., Gonzalez, C., Gupta, P., Hu, J., Huang, Y., Huang, Y., Khoo, W., Kumar, A., Kuo, R., Lakhavani, S., Liu, M., Luo, M., Luo, Z., Meredith, B., Miller, A., Oguntola, O., Pan, X., Peng, P., Pramanick, S., Ramazanov, M., Ryan, F., Shan, W., Somasundaram, K., Song, C., Southerland, A., Tateno, M., Wang, H., Wang, Y., Yagi, T., Yan, M., Yang, X., Yu, Z., Zha, S.C., Zhao, C., Zhao, Z., Zhu, Z., Zhuo, J., Arbelaez, P., Bertasius, G., Damen, D., Engel, J., Farinella, G.M., Furnari, A., Ghanem, B., Hoffman, J., Jawahar, C., Newcombe, R., Park, H.S., Rehg, J.M., Sato, Y., Savva, M., Shi, J., Shou, M.Z., Wray, M., 2024. Ego-Exo4D: Understanding skilled human activity from first- and third-person perspectives. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. CVPR*, pp. 19383–19400.
- Held, J., Cioppa, A., Giancola, S., Hamdi, A., Devue, C., Ghanem, B., Van Droogenbroeck, M., 2025. Towards an AI-powered video assistant referee system (VARs) for association football. In: Petersen, J., Dahl, V.A. (Eds.), *Image Analysis*. Springer Nature Switzerland, Cham, pp. 295–309.
- Hu, E.J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W., 2021. LoRA: Low-rank adaptation of large language models. *arXiv:2106.09685*. URL: <https://arxiv.org/abs/2106.09685>.
- Huang, Y., Chen, G., Xu, J., Zhang, M., Yang, L., Pei, B., Zhang, H., Lu, D., Wang, Y., Wang, L., Qiao, Y., 2024. EgoExoLearn: A dataset for bridging asynchronous ego and exo-centric view of procedural activities in real world. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Kourbane, I., Papadakis, P., Andries, M., 2025. SSL-Rehab: Assessment of physical rehabilitation exercises through self-supervised learning of 3D skeleton representations. *Comput. Vis. Image Underst.* 251, 104275. <http://dx.doi.org/10.1016/j.cviu.2024.104275>, URL: <https://www.sciencedirect.com/science/article/pii/S1077314224003564>.
- Kryeem, A., Boutboul, N., Bear, I., Raz, S., Eluz, D., Itah, D., Hel-Or, H., Shimshoni, I., 2025. Action assessment in rehabilitation: Leveraging machine learning and vision-based analysis. *Comput. Vis. Image Underst.* 251, 104228. <http://dx.doi.org/10.1016/j.cviu.2024.104228>, URL: <https://www.sciencedirect.com/science/article/pii/S1077314224003096>.
- Lavie, A., Agarwal, A., 2007. Meteor: an automatic metric for MT evaluation with high levels of correlation with human judgments. In: *Proceedings of the Second Workshop on Statistical Machine Translation. StatMT '07*, Association for Computational Linguistics, USA, pp. 228–231.
- Li, W., Liu, M., Liu, H., Wang, P., Lu, S., Sebe, N., 2025. H2OT: Hierarchical hourglass tokenizer for efficient video pose transformers. *IEEE Trans. Pattern Anal. Mach. Intell.* 1–15. <http://dx.doi.org/10.1109/TPAMI.2025.3608284>.
- Li, Y., Wang, C., Jia, J., 2024. LLaMA-VID: An image is worth 2 tokens in large language models. In: *Computer Vision – ECCV 2024: 18th European Conference, Milan, Italy, September 29–October 4, 2024, Proceedings, Part XLVI*. Springer-Verlag, Berlin, Heidelberg, pp. 323–340. [http://dx.doi.org/10.1007/978-3-031-72952-2\\_19](http://dx.doi.org/10.1007/978-3-031-72952-2_19).
- Lin, C.Y., 2004. ROUGE: A package for automatic evaluation of summaries. In: *Text Summarization Branches Out*. Association for Computational Linguistics, Barcelona, Spain, pp. 74–81. URL: <https://aclanthology.org/W04-1013/>.
- Lin, B., Ye, Y., Zhu, B., Cui, J., Ning, M., Jin, P., Yuan, L., 2024. Video-LLaVA: Learning unified visual representation by alignment before projection. *arXiv:2311.10122*. URL: <https://arxiv.org/abs/2311.10122>.
- Lingrui, X., Mandi, L., Lei, Z., 2025. TacticExpert: Spatial-temporal graph language model for basketball tactics. *arXiv:2503.10722*. URL: <https://arxiv.org/abs/2503.10722>.
- Liu, J., Chen, S., He, X., Guo, L., Zhu, X., Wang, W., Tang, J., 2025. VALOR: Vision-audio-language omni-perception pretraining model and dataset. *IEEE Trans. Pattern Anal. Mach. Intell.* 47 (2), 708–724. <http://dx.doi.org/10.1109/TPAMI.2024.3479776>.
- Liu, Y., Cheng, X., Li, Y., Ikenaga, T., 2024. Bidirectional temporal and frame-segment attention for sparse action segmentation of figure skating. *Comput. Vis. Image Underst.* 249, 104186. <http://dx.doi.org/10.1016/j.cviu.2024.104186>, URL: <https://www.sciencedirect.com/science/article/pii/S1077314224002674>.
- Liu, H., Li, C., Wu, Q., Lee, Y.J., 2023. Visual instruction tuning. In: *NeurIPS*.
- Liu, Z., Ning, J., Cao, Y., Wei, Y., Zhang, Z., Lin, S., Hu, H., 2021. Video swin transformer. *arXiv:2106.13230*. URL: <https://arxiv.org/abs/2106.13230>.
- Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., Stoyanov, V., 2019. RoBERTa: A robustly optimized BERT pretraining approach. *arXiv:1907.11692*. URL: <https://arxiv.org/abs/1907.11692>.
- Maaz, M., Rasheed, H., Khan, S., Khan, F.S., 2024. Video-ChatGPT: Towards detailed video understanding via large vision and language models. *arXiv:2306.05424*. URL: <https://arxiv.org/abs/2306.05424>.
- Marafioti, A., Zohar, O., Farré, M., Noyan, M., Bakouch, E., Cuenca, P., Zakka, C., Allal, L.B., Lozhkov, A., Tazi, N., Srivastav, V., Lochner, J., Larcher, H., Morlon, M., Tunstall, L., von Werra, L., Wolf, T., 2025. SmolVLM: Redefining small and efficient multimodal models. *arXiv preprint arXiv:2504.05299*.
- Martinelli, G., Diprima, F., Bisagno, N., Conci, N., 2024. Ski pose estimation. In: 2024 IEEE International Workshop on Sport, Technology and Research. STAR, pp. 120–125. <http://dx.doi.org/10.1109/STAR62027.2024.10635966>.
- Meneghetti, L., Bianchi, E., Demo, N., Rozza, G., 2025. KD-AHOSVD: Neural network compression via knowledge distillation and tensor decomposition. In: *Lorandei, J., Kamaleldin, A. (Eds.), Design and Architecture for Signal and Image Processing*. Springer Nature Switzerland, Cham, pp. 81–92.
- Mohottala, S., Gawesha, A., Kasthurirathna, D., Samarasinghe, P., Abhayaratne, C., 2025. Spatio-temporal graph neural network based child action recognition using data-efficient methods: A systematic analysis. *Comput. Vis. Image Underst.* 259, 104410. <http://dx.doi.org/10.1016/j.cviu.2025.104410>, URL: <https://www.sciencedirect.com/science/article/pii/S107731422500133X>.

- Pan, Y., Zhang, C., Bertasius, G., 2025. BASKET: A large-scale video dataset for fine-grained skill estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. CVPR, pp. 28952–28962.
- Parmar, P., Morris, B.T., 2019. What and how well you performed? A multitask learning approach to action quality assessment. In: 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition. CVPR, pp. 304–313. <http://dx.doi.org/10.1109/CVPR.2019.00039>.
- Sarlis, V., Tjortjis, C., 2020. Sports analytics — Evaluation of basketball players and team performance. *Inf. Syst.* 93, 101562. <http://dx.doi.org/10.1016/j.is.2020.101562>, URL: <https://www.sciencedirect.com/science/article/pii/S0306437920300557>.
- sedaghati, N., ardebili, S., Ghaffari, A., 2025. Application of human activity/action recognition: a review. *Multimedia Tools Appl.* 84 (28), 33475–33504. <http://dx.doi.org/10.1007/s11042-024-20576-2>.
- Su, Y., Lan, T., Li, H., Xu, J., Wang, Y., Cai, D., 2023. PandaGPT: One model to instruction-follow them all. *arXiv preprint arXiv:2305.16355*.
- Sun, C., Myers, A., Vondrick, C., Murphy, K., Schmid, C., 2019. VideoBERT: A Joint Model for Video and Language Representation Learning. In: 2019 IEEE/CVF International Conference on Computer Vision. ICCV, IEEE Computer Society, Los Alamitos, CA, USA, pp. 7463–7472. <http://dx.doi.org/10.1109/ICCV.2019.00756>, URL: <https://doi.ieeecomputersociety.org/10.1109/ICCV.2019.00756>.
- Tian, Q., Zeng, F., Ning, J., Zhang, L., 2025. UniMultNet: Action recognition method based on multi-scale feature fusion and video-text constraint guidance. *Comput. Vis. Image Underst.* 260, 104456. <http://dx.doi.org/10.1016/j.cviu.2025.104456>, URL: <https://www.sciencedirect.com/science/article/pii/S1077314225001791>.
- Tran, D., Bourdev, L., Fergus, R., Torresani, L., Paluri, M., 2015. Learning spatiotemporal features with 3D convolutional networks. *arXiv:1412.0767*. URL: <https://arxiv.org/abs/1412.0767>.
- Wang, X., Zhang, Y., Zohar, O., Yeung-Levy, S., 2024. VideoAgent: Long-form video understanding with large language model as agent. *arXiv:2403.10517*. URL: <https://arxiv.org/abs/2403.10517>.
- Xia, J., Zhuge, M., Geng, T., Fan, S., Wei, Y., He, Z., Zheng, F., 2022. Skating-mixer: Long-term sport audio-visual modeling with MLPs. *arXiv:2203.03990*. URL: <https://arxiv.org/abs/2203.03990>.
- Zhai, X., Mustafa, B., Kolesnikov, A., Beyer, L., 2023. Sigmoid loss for language image pre-training. In: 2023 IEEE/CVF International Conference on Computer Vision. ICCV, pp. 11941–11952. <http://dx.doi.org/10.1109/ICCV51070.2023.01100>.
- Zhang, S., Bai, S., Chen, G., Chen, L., Lu, J., Wang, J., Tang, Y., 2024. Narrative action evaluation with prompt-guided multimodal interaction. In: 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition. CVPR, pp. 18430–18439. <http://dx.doi.org/10.1109/CVPR52733.2024.01744>.
- Zhang, T., Kishore, V., Wu, F., Weinberger, K.Q., Artzi, Y., 2020. BERTScore: Evaluating text generation with BERT. *arXiv:1904.09675*. URL: <https://arxiv.org/abs/1904.09675>.
- Zhang, H., Li, X., Bing, L., 2023. Video-LLaMA: An instruction-tuned audio-visual language model for video understanding. In: Feng, Y., Lefever, E. (Eds.), Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: System Demonstrations. Association for Computational Linguistics, Singapore, pp. 543–553. <http://dx.doi.org/10.18653/v1/2023.emnlp-demo.49>, URL: <https://aclanthology.org/2023.emnlp-demo.49/>.
- Zhang, Z., Liu, W., Zheng, Y., Du, L., Sun, L., 2025. Learning spatio-temporal context for basketball action pose estimation with a multi-stream network. *Sci. Rep.* 15 (1), 29173. <http://dx.doi.org/10.1038/s41598-025-14985-y>.
- Zhao, L., Lin, Z., Sun, R., Wang, A., 2024. A review of state-of-the-art methodologies and applications in action recognition. *Electronics* 13 (23), <http://dx.doi.org/10.3390/electronics13234733>, URL: <https://www.mdpi.com/2079-9292/13/23/4733>.
- Zhou, K., Cai, R., Wang, L., Shum, H.P.H., Liang, X., 2024. A comprehensive survey of action quality assessment: Method and benchmark. *arXiv:2412.11149*. URL: <https://arxiv.org/abs/2412.11149>.
- Zhu, D., Chen, J., Shen, X., Li, X., Elhoseiny, M., 2023. MiniGPT-4: Enhancing vision-language understanding with advanced large language models. *arXiv:2304.10592*. URL: <https://arxiv.org/abs/2304.10592>.
- Zohar, O., Wang, X., Bitton, Y., Szpektor, I., Yeung-Levy, S., 2024a. Video-STaR: Self-training enables video instruction tuning with any supervision. *arXiv:2407.06189*. URL: <https://arxiv.org/abs/2407.06189>.
- Zohar, O., Wang, X., Dubois, Y., Mehta, N., Xiao, T., Hansen-Estruch, P., Yu, L., Wang, X., Juefei-Xu, F., Zhang, N., Yeung-Levy, S., Xia, X., 2024b. Apollo: An exploration of video understanding in large multimodal models. *arXiv:2412.10360*. URL: <https://arxiv.org/abs/2412.10360>.