
EvoGround: Self-Evolving Video Agents for Video Temporal Grounding

Minjoon Jung^{1,2} Byoung-Tak Zhang¹ Lorenzo Torresani²

¹Seoul National University ²Northeastern University

Abstract

Video temporal grounding (VTG) takes an untrimmed video and a natural-language query as input and localizes the temporal moment that best matches the query. Existing methods rely on large, task-specific datasets requiring costly manual annotation. We introduce EvoGround, a framework of two coupled self-evolving agents, a *proposer* and a *solver*, that learn temporal grounding from raw videos without any human-labeled data. The proposer generates query–moment pairs from raw videos, while the solver learns to ground them and feeds back signals that improve the proposer in return. Through this self-reinforcing reinforcement-learning loop, the two agents are initialized from the *same* backbone and mutually improve across iterations. Trained on 2.5K unlabeled videos, EvoGround matches or surpasses fully supervised models across multiple VTG benchmarks, while emerging as a state-of-the-art fine-grained video captioner without manual labels.¹

1 Introduction

Accurate temporal grounding underpins a wide range of applications, including video retrieval [1], highlight detection [2], and grounded question answering [3], and is increasingly relied upon as a core module in complex video understanding pipelines [4–6]. Given an untrimmed video and a natural language query, video temporal grounding (VTG) [7] aims to localize the temporal moment within the video that corresponds to the query. Despite remarkable progress, existing models typically rely on large-scale human-labeled data for training, fundamentally limiting their scalability and applicability.

Collecting labeled data for VTG is particularly burdensome. Annotators must manually identify temporal segments and provide corresponding natural language descriptions within untrimmed videos, a time-consuming process that inherently limits scalability. This often leads to a sparsity issue [8, 9], where only a fraction of video events are covered, amplifying the narrowness of the training distribution. Beyond scalability, obtaining high-quality annotations remains non-trivial: queries are often ambiguous or misaligned with the underlying visual content, causing models to learn spurious correlations or exploit shortcuts rather than being grounded in true visual evidence [10–12].

Can a model reliably learn temporal information in videos without manual labels? At first glance, this seems unlikely, as temporal grounding requires precise, fine-grained reasoning that one might assume can only emerge from explicit supervision. We challenge this assumption with EvoGround, a framework of two coupled self-evolving agents that learn to ground temporal moments through mutual interaction without manual labels.

As shown in Figure 1, EvoGround consists of a *proposer* and a *solver*, evolved through an iterative two-stage process via reinforcement learning (RL), with both agents initialized from the same base model. Rather than relying on human annotators, the proposer takes on their role: given a raw video, it identifies temporal moments and generates corresponding natural language queries. The solver

¹All resources made publicly available at <https://minjoong507.github.io/projects/EvoGround/>.

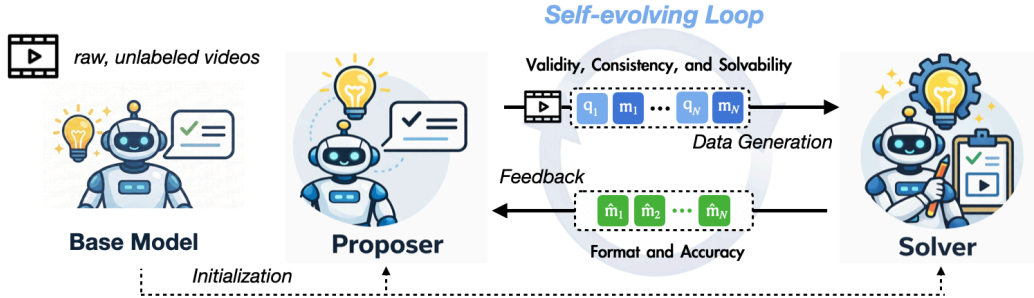


Figure 1: **EvoGround: a self-evolving loop, with unlabeled videos.** A proposer and a solver, both initialized from the same base model, co-evolve through reinforcement learning. The proposer generates query (q)–moment (m) pairs from a raw video; the solver grounds them and produces predictions (\hat{m}) that feed back as a learning signal. Dedicated reward designs guide each agent.

then attempts to ground these queries back to their moments, and its success or failure is fed back to the proposer as a learning signal.

In the first stage, we encourage the proposer to provide valid, semantically consistent, and solvable query-moment pairs; in the second stage, the solver is trained on the data generated by the updated proposer. Through this self-reinforcing loop with dedicated reward designs for each agent, they challenge and sharpen each other: a better proposer generates more meaningful and solvable queries, which in turn train a better solver, whose improved feedback further raises the bar for the proposer. This virtuous cycle, driven entirely by reinforcement learning and requiring no manual labels, yields a system whose temporal understanding improves with each iteration.

Despite never seeing a single manual annotation during training, EvoGround matches or outperforms most fully-supervised models across five VTG benchmarks: Charades-STA [7], ActivityNet-Captions [13], TVGBench [14], ReXTime [15], and E.T.Bench [16]. Moreover, EvoGround demonstrates its strong fine-grained video captioning capability on TemporalBench [17], suggesting that the self-evolution loop fosters a broader and more grounded understanding of video content.

Our primary contributions are summarized as follows:

- We propose EvoGround, a framework of self-evolving agents for VTG that learns temporal information from raw videos without any manual annotation.
- EvoGround adopts a proposer–solver loop in which the two agents, initialized from the same backbone, iteratively challenge and improve each other.
- EvoGround achieves accuracy on par with, and in five VTG benchmarks surpassing, fully supervised models, while demonstrating its fine-grained video captioning on TemporalBench.

2 Related Work

Video Temporal Grounding (VTG). VTG is a task that aims to localize temporal moments from given natural language sentences within untrimmed videos. VTG models typically encode multimodal features using a pre-trained encoder (*e.g.*, CLIP [18]) and adopt proposal-based [19], proposal-free [20, 21], and DETR-based methods [2, 22] for prediction. However, these models are inherently less generalizable and limited by the capacity of pre-trained encoders. Such limitations have motivated a shift toward more general and scalable approaches that employ large language models (LLMs).

Video Large Language Models (Video-LLMs). Video-LLMs leverage the powerful capabilities of LLMs [23, 24] to handle a wide range of downstream tasks, from question answering [25–27] to fine-grained temporal understanding [28–30, 12, 6, 31]. Recently, inspired by the success of RL (*e.g.*, DeepSeek-R1 [32]), RL-based Video-LLMs [33, 14, 34, 35] have emerged and demonstrate superior spatio-temporal reasoning. However, existing models typically rely on labeled data and often require strong teacher models for guidance. For instance, VideoChat-R1.5 [35] leverages DeepSeek to ensure training data quality, and VideoChat-R1 [34] utilizes Qwen2.5-72B [36] to encourage alignment between model outputs and caption supervision. In contrast, EvoGround avoids such reliance and learns temporal information directly from raw videos through iterative self-reinforcing loops.

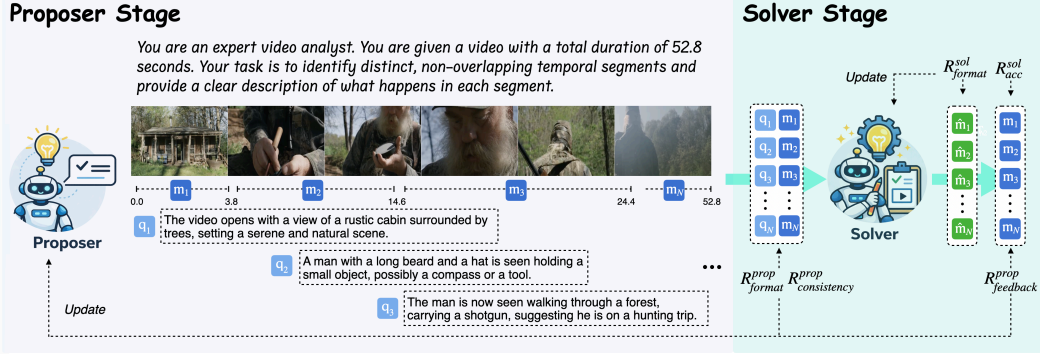


Figure 2: **Overview of EvoGround.** Both agents start from the same backbone. The proposer is updated via three rewards: $R_{\text{format}}^{\text{prop}}$ (validity), $R_{\text{consistency}}^{\text{prop}}$ (consistency, computed with SigLIP-2), and $R_{\text{feedback}}^{\text{prop}}$ (solvability, derived from the solver’s tIoU). The solver is updated via $R_{\text{format}}^{\text{sol}}$ and $R_{\text{acc}}^{\text{sol}}$. Stages alternate: the proposer’s pairs train the solver, the solver’s predictions sharpen the proposer.

Self-Evolving Agents. Self-evolving frameworks [37–39] have improved LLMs with minimal or no training data. Recent works [40–42] have expanded this mechanism from text to vision-language. We adopt this loop, where each agent bootstraps signals for the other without manual annotations, and extend it to video temporal understanding. Both agents share the same backbone, with no stronger video grounder or teacher in the loop. Any gain reflects the procedure itself rather than distillation.

3 Method

As shown in Figure 2, we introduce EvoGround, consisting of two agents: a *proposer* and a *solver*. The proposer identifies candidate temporal events from raw videos and generates corresponding query-moment pairs, while the solver learns to ground temporal moments using the generated data. Both agents are initialized from the same backbone [43], and evolve solely through the self-reinforcing loop without any labeled data. Appendix Figure 6 provides full prompts for each agent.

3.1 Proposer

To guide the proposer toward generating high-quality data, we define three key criteria: *validity* (format reward), *consistency* (consistency reward), and *solvability* (feedback reward). (1) Validity: Temporal moments must be defined by numeric timestamps within a valid range (*i.e.*, non-negative and within the video duration) and avoid extremely short intervals or trivial spans that cover the entire video. (2) Consistency: A query–moment pair should be semantically coherent, *i.e.*, the query should accurately describe the visual content of the moment and be discriminative with respect to other candidates, independent of the solver. (3) Solvability: The generated pairs must be learnable by the solver, *i.e.*, the solver should be able to localize the moment from the query. A pair may be semantically consistent yet unsolvable if its temporal boundaries are too ambiguous or imprecise. We design the corresponding reward functions based on these criteria.

For each video, we prompt the proposer to generate a set of temporal moments and corresponding queries. Let the proposer generate N query-moment pairs $\{(q_n, m_n)\}_{n=1}^N$, where q_n denotes the n -th query and $m_n = (s_n, e_n)$ represents the corresponding temporal moment with start time s_n and end time e_n in seconds. For each query q_n , the solver produces a moment prediction $\hat{m}_n = (\hat{s}_n, \hat{e}_n)$. Based on these generated pairs and predictions, we denote the i -th generated sample as $\mathbf{s}_i = \{q_n, m_n, \hat{m}_n\}_{n=1}^N$, which serves as the input to all reward functions defined below.

Format Reward. The format reward verifies whether the proposer generates valid query–moment pairs. To facilitate reliable parsing and evaluation, we prompt the model to provide multiple moments and queries, each enclosed within predefined tags, `<time>...</time>` and `<description>...</description>`, respectively. The format reward is defined as:

$$R_{\text{format}}^{\text{prop}}(\mathbf{s}_i) = \begin{cases} \frac{1}{N} \sum_{n=1}^N \mathbb{I}((s_n, e_n) \in \mathcal{V}), & \text{if the proposer follows the template,} \\ 0, & \text{otherwise,} \end{cases} \quad (1)$$

where $\mathbb{I}(\cdot)$ is the indicator function and \mathcal{V} denotes the set of valid temporal moments satisfying: (1) $s_n \geq 0$ and $e_n \leq T_{\text{video}}$ and (2) $e_n - s_n > \tau_{\text{length}}$, where T_{video} is the video duration and τ_{length} specifies the minimum segment length. In later iterations, we further condition the format reward on the solver’s feedback to ensure that generated query–moment pairs are not only structurally valid but also solvable. We measure the feedback using timestamp-aware IoU (tIoU) [14]:

$$\text{tIoU}(m_n, \hat{m}_n) = \text{IoU}(m_n, \hat{m}_n) \cdot \left(1 - \frac{|s_n - \hat{s}_n|}{T_{\text{video}}}\right) \cdot \left(1 - \frac{|e_n - \hat{e}_n|}{T_{\text{video}}}\right), \quad (2)$$

which extends standard IoU [10] by incorporating penalties for deviations in start and end timestamps relative to the video duration, leading to improved boundary sensitivity. Here, IoU is defined as:

$$\text{IoU}(m_n, \hat{m}_n) = \frac{[s_n, e_n] \cap [\hat{s}_n, \hat{e}_n]}{[s_n, e_n] \cup [\hat{s}_n, \hat{e}_n]}. \quad (3)$$

Based on this, we define the *conditioned format reward* as:

$$\mathbb{I}((s_n, e_n) \in \mathcal{V} \wedge \text{tIoU}(m_n, \hat{m}_n) \geq \delta), \quad (4)$$

where $\delta \in [0, 1]$ is a threshold that conditions the validity of query–moment pairs on the current solver’s accuracy. This design ensures that the proposer generates pairs that are both structurally valid and learnable at each training step. The format reward in Equation 1 corresponds to the case of $\delta = 0$.

Consistency Reward. The consistency reward measures two complementary properties of each query–moment pair: *intra-consistency*, which captures how coherently the query aligns with the frames within its corresponding moment, and *inter-consistency*, which captures how discriminatively the query matches its own moment relative to other candidates. We use SigLIP-2 [44] to compute frame-level similarities $S_{n,t} = \cos(q_n, f_t)$ by uniformly sampling $|m_n|$ frames from each moment at a rate proportional to its duration, and obtain their average μ_n . The intra-consistency score penalizes high standard deviation in $S_{n,t}$, rewarding pairs where the query aligns uniformly with frames in m_n :

$$\text{Intra}_n = \exp(-\gamma \cdot \sigma_n), \quad \sigma_n = \sqrt{\frac{1}{|m_n|} \sum_{t=1}^{|m_n|} (S_{n,t} - \mu_n)^2}, \quad (5)$$

where γ controls sensitivity to alignment fluctuations. However, a pair with uniformly low similarity would still score highly under Intra_n alone. The inter-consistency score addresses this by measuring how discriminatively the query matches its own moment relative to others, computed as a softmax over cross-moment similarities with temperature τ :

$$\text{Inter}_n = \frac{\exp(\mu_n/\tau)}{\sum_{j=1}^N \exp(\mu_{n,j}/\tau)}, \quad (6)$$

where $\mu_{n,j}$ denotes the mean similarity between query q_n and moment m_j . The final consistency reward is their product, requiring both temporal coherence within the moment and semantic discriminability across moments to be simultaneously satisfied:

$$R_{\text{consistency}}^{\text{prop}}(\mathbf{s}_i) = \frac{1}{N} \sum_{n=1}^N \text{Intra}_n \cdot \text{Inter}_n. \quad (7)$$

In Appendix A.3, we provide more details on the intra-consistency formulation and the computation of $|m_n|$. We further explore variants of inter-consistency under different formulations (*e.g.*, hard ranking-based) and find that the softmax formulation consistently outperforms the alternatives.

Feedback Reward. The feedback reward confirms whether the pair is solvable and is measured by the solver’s accuracy:

$$R_{\text{feedback}}^{\text{prop}}(\mathbf{s}_i) = \frac{1}{N} \sum_{n=1}^N \text{tIoU}(m_n, \hat{m}_n), \quad (8)$$

where \hat{m}_n denotes the moment predictions from query q_n .




Figure 3: **Reward dynamics across iterations.** We visualize the evolution of the proposer and solver in (a) and (b), respectively. As the proposer evolves over iterations, the solver correspondingly demonstrates progressively higher accuracy.

3.2 Solver

From the query-moment pairs generated by the proposer, each pair is processed independently. We consider a single query-moment pair (q, m) . Given the query q , the solver generates multiple rollout samples $\{o_k\}_{k=1}^K$, from which we extract the corresponding predicted moments $\{\hat{m}_k\}_{k=1}^K$.

Format Reward. To facilitate easy answer extraction, we encourage the solver to follow the pre-defined template: `<think>...</think>` and `<answer>...</answer>`, where the `<think>` block contains the model’s intermediate reasoning steps and the `<answer>` block contains the predicted temporal moment as start and end timestamps. Based on this, the format reward is defined as:

$$R_{\text{format}}^{\text{sol}}(o_k) = \begin{cases} 1, & \text{if the solver follows the template,} \\ 0, & \text{otherwise.} \end{cases} \quad (9)$$

Accuracy Reward. To improve accuracy, we extract the prediction \hat{m}_k from the content within the answer tags and measure accuracy by comparing it with the ground-truth moment m :

$$R_{\text{acc}}^{\text{sol}}(o_k) = \text{tIoU}(m, \hat{m}_k). \quad (10)$$

Note that while $R_{\text{acc}}^{\text{sol}}$ and $R_{\text{feedback}}^{\text{prop}}$ share the same formulation, they serve distinct roles: the former optimizes the solver toward accurate grounding, while the latter signals to the proposer whether its generated pairs are learnable.

3.3 Optimization

In this section, we describe how EvoGround is optimized. EvoGround alternates between two stages: in the *proposer stage*, the proposer is optimized while the solver is frozen; in the *solver stage*, the solver is updated using query-moment pairs generated by the frozen proposer.

Objective. We adopt GDPO [45] as our RL optimizer, which improves upon GRPO [32] by weighting and normalizing each reward component individually before aggregation, rather than summing first and normalizing after. This improves the resolution of the training signals and prevents easier rewards from dominating the optimization, which is critical in our multi-reward setting. We provide the detailed formulation of each objective in Appendix A.1.

Curriculum Design. In the first iteration, we warm up the proposer with the format reward only and update the solver, since neither agent is yet tailored to the task. In subsequent iterations, we apply consistency and feedback rewards to the proposer and adopt GDPO for this multi-reward setting, assigning higher weights to the feedback reward than to the format reward to emphasize more challenging objectives. We also progressively increase δ (in Equation 4) across iterations, shifting the focus from coarse matches toward more precise temporal alignments. We validate the effectiveness of these design choices in Section 5. For the solver, we keep the design simple, applying naive GRPO without additional reward designs or the multi-reward weighting scheme (*i.e.*, GDPO).

Reward Dynamics. Figure 3 visualizes the reward dynamics of EvoGround across iterations. Both agents quickly saturate the format reward in the first iteration, suggesting that satisfying the format constraint is easy. In subsequent iterations, the conditioned format reward decreases as δ increases, exhibiting a significantly lower scale compared to the naive format reward due to the stricter constraint. Meanwhile, the feedback reward steadily improves, tracking the solver’s increasing accuracy. This confirms that our curriculum design effectively guides the optimization process across iterations.

Table 1: **Results on Charades-STA and ActivityNet-Captions.** EvoGround achieves competitive or superior performance to existing models without using any manual labels, outperforming all SFT-based models and matching the best RL-based models. * denotes our reproduced results.

Method	Charades-STA				ActivityNet-Captions			
	R1@0.3	R1@0.5	R1@0.7	mIoU	R1@0.3	R1@0.5	R1@0.7	mIoU
<i>SFT-based Models</i>								
TimeChat [30]	-	32.2	13.4	32.2	36.2	20.2	9.5	21.8
HawkEye [46]	50.6	31.4	14.5	-	49.1	29.3	10.7	-
VTimeLLM [28]	51.0	27.5	11.4	31.2	44.0	27.8	14.3	30.4
Momentor [47]	42.6	26.6	11.6	28.5	42.9	23.0	12.4	29.3
Qwen2.5-VL* [43]	68.5	48.8	22.5	45.0	38.7	25.6	14.9	28.6
Grounded-VideoLLM [48]	51.8	34.3	18.3	34.7	43.9	29.1	18.3	34.5
TimeSuite [31]	69.9	48.7	24.0	-	-	-	-	-
TRACE [49]	58.6	40.3	19.4	38.7	-	-	-	-
ED-VTG [50]	59.5	39.3	19.8	40.2	52.1	33.1	16.0	35.2
<i>RL-based Models</i>								
VideoChat-R1 [34]	-	-	-	-	50.4	32.2	16.2	34.3
VideoChat-R1.5 [35]	-	-	-	-	52.4	32.3	16.8	35.5
Time-R1 [14]	78.1	60.8	<u>35.3</u>	58.1	58.6	<u>39.0</u>	<u>21.4</u>	<u>40.5</u>
EvoGround (Ours)	<u>77.2</u> ± 1.3	<u>60.5</u> ± 1.6	<u>35.5</u> ± 1.5	<u>53.1</u> ± 0.9	62.9 ± 0.8	43.6 ± 0.7	25.0 ± 0.6	43.9 ± 0.5

4 Experiments

4.1 Experiment Setup

Baselines. We consider Video-LLMs as baselines, most of which focus on VTG. We categorize them into two folds: SFT-based models, including TimeChat [30], HawkEye [46], VTimeLLM [28], Momentor [47], Qwen2.5-VL [43], TimeSuite [31], Grounded-VideoLLM [48], VideoChat-Flash [51], LITA [52], TRACE [49], and ED-VTG [50], and RL-based models, including VideoChat-R1 [34], VideoChat-R1.5 [35], and Time-R1 [14]. For captioning, we include various Video-LLMs, including Qwen2-VL [36], LLaVA-OneVision [53], LLaVA-NeXT-Video [54], InternLM-XC2.5 [55], VideoLLaVA [56], MiniCPM-V2.6 [57], Phi-3.5-Vision [58], MA-LLM [59], and Matryoshka Multimodal Models (M^3) [60]. Models are 7B-scale unless otherwise specified. Among all baselines, EvoGround is the only model explicitly trained for VTG without manual annotations.

Benchmarks. We conduct experiments on five VTG benchmarks: Charades-STA [7], ActivityNet-Captions [13], TVGBench [14], ReXTime [15], and E.T.Bench [16]. Charades-STA contains short indoor activity videos, while ActivityNet-Captions consists of long, untrimmed videos. TVGBench integrates multiple benchmarks, providing a balanced range of video durations and high diversity. ReXTime and E.T.Bench involve test samples rigorously curated by human annotators. Additionally, we use TemporalBench [17] to evaluate the model’s captioning capability.

Evaluation Metrics. For grounding, we use $R@n$, $IoU=m$ metric, which measures the percentage of top- n predicted moments whose IoU with the ground-truth moment exceeds a threshold $m \in \{0.3, 0.5, 0.7\}$, and also report the mean IoU (mIoU) and F1 score. For further reliable comparison, we report 95% confidence intervals estimated via nonparametric bootstrap (*i.e.*, 1,000 resamples with replacement of the test set)². Point estimates are computed on the full test set, and intervals correspond to the 2.5th–97.5th percentiles of the bootstrap distribution. For captioning, we adopt classical captioning metrics: CIDEr [61], BLEU [62] at various n-gram levels, ROUGE [63], and sentence similarity between generated and ground-truth captions with a pre-trained Sentence Transformer [64].

Implementation Details. We use Qwen2.5-VL-7B [43] as a backbone for a fair comparison with prior RL-based models [34, 35, 14] and utilize 2.5K raw videos from TimeRFT [14]. We perform three iterations to train both agents. At each iteration, the proposer generates approximately 9K query–moment pairs and updates the solver, completing within 24 hours using $4 \times A100$ GPUs. For the proposer, we progressively increase δ ($0 \rightarrow 0.3 \rightarrow 0.5$) and assign reward weights of $w_{\text{format}} = 0.5$, $w_{\text{consistency}} = 0.5$, and $w_{\text{feedback}} = 1.0$. Following [14], video frames are sampled at 2 FPS and adaptively resized to 2.8M pixels to control efficiency during both training and testing; results marked with * in the following tables are obtained from the official checkpoint under the same setup. Please refer to Appendix A.2 for more implementation details.

²We exclude ReXTime for this since its ground-truth data is released in the official test server

Table 2: **Results on TVGBench, ReXTime, and E.T.Bench.** Across diverse benchmarks spanning varying video durations and domains, EvoGround surpasses SFT-based models and remains competitive with RL-based models trained on human-annotated data. * denotes our reproduced results.

Method	TVGBench				ReXTime			E.T.Bench
	R1@0.3	R1@0.5	R1@0.7	mIoU	R1@0.3	R1@0.5	mIoU	TVG _{F1}
<i>SFT-based Models</i>								
TimeChat [30]	22.4	11.9	5.3	-	14.4	7.6	11.6	26.2
VTimeLLM [28]	-	-	-	-	28.8	17.4	20.1	7.6
Qwen2.5-VL* [43]	28.1	19.5	10.5	20.4	16.3	11.2	13.3	46.6
LITA-13B [52]	-	-	-	-	29.4	16.2	21.4	22.2
TimeSuite [31]	31.1	18.0	8.9	-	-	-	-	-
VideoChat-Flash [51]	32.8	19.8	10.4	-	-	-	-	-
TRACE [49]	37.0	25.5	14.6	-	-	-	-	-
<i>RL-based Models</i>								
VideoChat-R1.5* [35]	31.5	20.8	11.3	22.0	28.3	18.0	20.5	50.3
Time-R1* [14]	<u>39.3</u>	<u>28.0</u>	16.0	<u>27.8</u>	<u>32.2</u>	<u>22.1</u>	<u>24.1</u>	69.4
EvoGround (Ours)	42.1 ± 3.3	28.5 ± 3.1	<u>15.3</u> ± 2.6	29.4 ± 2.1	33.5	22.4	25.5	<u>69.0</u> ± 2.2

Table 3: **Results on TemporalBench.** EvoGround outperforms all prior models across all captioning metrics, demonstrating broader video understanding beyond temporal grounding.

Method	TemporalBench						
	Similarity	CIDEr	ROUGE	BLEU_1	BLEU_2	BLEU_3	BLEU_4
Qwen2-VL [36]	51.9	<u>6.9</u>	<u>18.0</u>	12.5	6.1	3.0	<u>1.6</u>
LLaVA-OneVision [53]	50.1	0.3	14.5	11.1	5.1	2.2	1.1
LLaVA-NeXT-Video [54]	50.1	2.3	15.8	18.1	7.0	2.6	1.1
InternLM-XC2.5 [55]	<u>52.4</u>	2.3	15.9	17.8	7.1	2.8	1.2
VideoLLaVA [56]	46.0	4.5	16.9	12.6	5.4	2.3	1.0
MiniCPM-V2.6 [57]	47.2	1.5	14.2	15.5	5.4	1.9	0.8
Phi-3.5-Vision [58]	42.9	3.7	16.5	<u>20.4</u>	<u>8.4</u>	<u>3.4</u>	<u>1.6</u>
MA-LMM [59]	38.7	3.1	15.0	10.1	4.8	2.2	1.1
M ³ [60]	47.8	3.0	16.4	16.7	6.9	2.8	1.2
EvoGround (Ours)	53.8	11.4	20.5	26.9	12.5	5.6	2.6

4.2 Main Results

Despite never seeing any manual labels, EvoGround demonstrates strong performance across all benchmarks. In Tables 1 and 2, EvoGround surpasses all SFT-based models and achieves first- or second-best performance among RL-based models, which, however, rely on manually labeled data. In Table 3, we evaluate on TemporalBench, which reflects the model’s fine-grained captioning capability to describe a sequence of events throughout the video. To this end, we aggregate the event descriptions generated by the proposer and compare it against the ground-truth captions. EvoGround outperforms all prior models across every metric, with particularly large margins on CIDEr and BLEU, despite never being explicitly trained for captioning. Beyond performance, EvoGround is significantly more data-efficient than existing models, which typically involve data generation and curation pipelines. For instance, VideoChat-R1.5 [35] constructs VTTS-80K from 11 public datasets, and ED-VTG [50] collects 136K samples from 8 public datasets, whereas we use only 2.5K raw videos. Overall, these results demonstrate that EvoGround effectively learns temporal information without manual labels.

5 Analysis

In this section, we provide in-depth analyses of EvoGround.

GDPO outperforms GRPO. Figure 4 (a) reports the performance using different RL learning objectives across iterations. GDPO consistently outperforms GRPO across all iterations and already surpasses GRPO’s best performance at the second iteration, achieving faster convergence. These results support our claim that GDPO is well-suited for our multi-reward settings.




Figure 4: **Improvements across iterations on TVGBench.** (a) shows performance using different learning objectives. (b) and (c) show performance across different video and moment lengths.

Table 4: **Ablation on data scaling.** FT indicates fine-tuning on Charades-STA training data. In contrast, EvoGround (7.5K) additionally uses 5K raw videos from Charades-STA without manual labels.

Method	FT	Charades-STA			
		R1@0.3	R1@0.5	R1@0.7	mIoU
Hawkeye [46]	✓	72.5	58.3	28.8	49.3
TimeChat [30]	✓	-	46.7	23.7	-
TimeChat-VT [12]	✓	-	58.4	34.7	-
ED-VTG [50]	✓	78.2	62.1	35.0	52.6
VideoChat-R1 [34]	✓	-	71.7	50.2	60.8
VideoChat-R1.5 [35]	✓	82.8	71.6	48.3	60.6
Time-R1 [14]	✓	82.8	72.2	50.1	60.9
EvoGround (2.5K)	✗	77.2	60.5	35.5	53.1
EvoGround (7.5K)	✗	81.7	67.7	45.6	60.3

Table 5: **Ablation on rewards.** Each reward contributes incrementally, with all rewards combined achieving the best performance.

Iter	R_{fmt}	R_{con}	R_{feed}	TVGBench		
				R1@0.3	R1@0.5	R1@0.7
1	✓			34.0	21.2	12.0
2	✓			37.0	23.8	12.5
2	✓	✓		38.7	26.2	14.6
2	✓		✓	38.1	26.0	14.4
2	✓	✓	✓	40.1	27.4	14.9
3	✓			37.2	23.6	13.5
3	✓	✓		41.7	27.3	15.0
3	✓		✓	41.8	27.4	14.8
3	✓	✓	✓	42.1	28.5	15.3

Increasing δ shifts the model’s focus toward more accurate grounding. We compare the cases where δ is increased to 0.5 versus kept at 0.3. In Figure 4 (b) and (c), $\delta = 0.5$ shows a more balanced improvement than $\delta = 0.3$ across different video durations and moment lengths. Specifically, $\delta = 0.3$ yields only marginal improvement over the previous iteration, and its effectiveness diminishes particularly for longer videos and moments. This suggests that increasing δ encourages the proposer to foster a more robust solver. This does not bias toward trivially easy samples, as the improvement is most pronounced for longer videos and moments where grounding is inherently harder.

Adding more videos restarts the loop. Figure 4 (a) shows that gains plateau by iteration 3. With only 2.5K raw videos, the proposer and solver have already learned what they can from this small pool, and later iterations mostly recycle similar pairs. In Table 4, training EvoGround on 5K additional raw videos from Charades-STA (7.5K total, no overlap) for one extra iteration yields further gains, remaining competitive with models fine-tuned on human annotations. As in-domain gains are expected, the Appendix Table 10 shows smaller but positive transfer to other benchmarks.

Gains transfer across backbone scales and families. Across four backbones spanning 3B–8B parameters and two families (*i.e.*, QwenVL and InternVL), EvoGround delivers nearly identical gains and benefits from larger capacity, with Qwen3-VL-8B [65] reaching the highest absolute scores. We will explore more aggressive scaling in future work.

Each reward plays a distinct role. We optimize our model under different reward configurations: (1) format only, (2) format + feedback, (3) format + consistency, and (4) all rewards, and visualize their moment distributions and query–moment length correlation (r) in Figure 5, with performance in Table 5. Each reward addresses a specific failure mode of the proposer, jointly preventing collapse to trivial outputs. The format reward enforces structural validity, but formatting alone yields the worst performance, suggesting the proposer collapses into biased, repetitive data without quality constraints. The feedback reward enriches query descriptiveness (r : 0.11 \rightarrow 0.25), though r is a diagnostic signal rather than a quality metric: a longer query does not necessarily describe the moment better. The consistency reward sharpens the moment distribution, producing tighter temporal moments ($27.0 \pm 37.4 \rightarrow 19.8 \pm 24.6$) and reducing extension toward the video end. The tri-modal start time distribution reflects the consistency reward’s tendency to anchor moments at discriminative positions, avoiding overlap with other candidates. Combining all rewards moderates r to 0.14 and yields the balanced distribution and best performance overall.

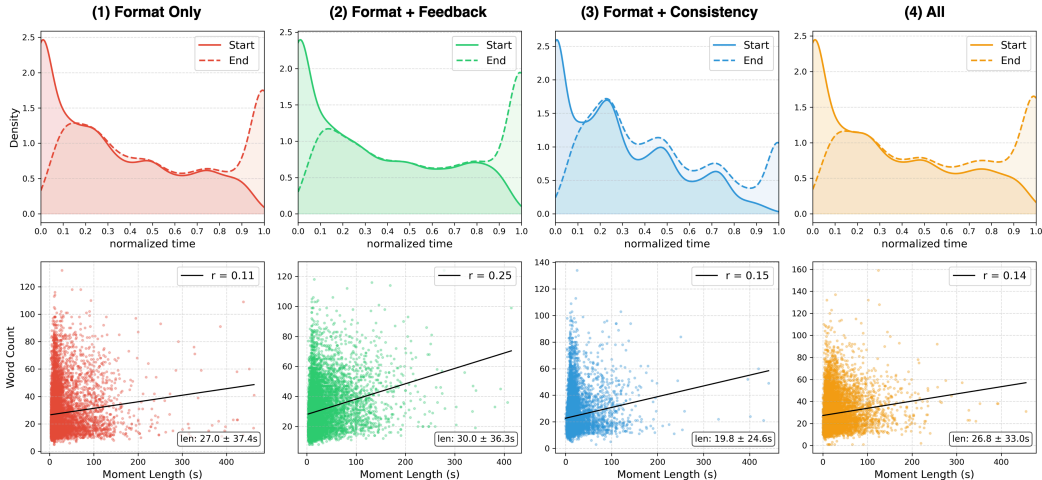


Figure 5: **Generated data distributions across different reward configurations.** *Top:* Kernel density estimates of normalized start and end times of the moment, shown as solid and dashed lines, respectively. All times are normalized by video duration. *Bottom:* We report the correlation (r) between query and moment length, along with the mean \pm standard deviation of moment lengths.

Table 6: **Ablation on backbones.** Gray columns show initial performance, and blue indicates improvements from EvoGround. Gains are stable across backbone size and family, suggesting the framework rather than backbone-specific behavior drives improvement.

Backbone	TVGBench					
	R1@0.3		R1@0.5		R1@0.7	
Qwen2.5-VL-3B [43]	17.5	31.3 (+13.8†)	10.0	19.4 (+9.4†)	5.2	9.9 (+4.7†)
Intern3.5-VL-4B [66]	14.2	30.6 (+16.4†)	8.5	18.1 (+9.6†)	4.0	9.6 (+5.6†)
Qwen2.5-VL-7B [43]	28.1	42.1 (+14.0†)	19.5	28.5 (+9.0†)	10.5	15.3 (+4.8†)
Qwen3-VL-8B [65]	30.2	44.6 (+14.4†)	20.8	30.1 (+9.3†)	10.5	17.2 (+6.7†)

Table 7: **Dense video captioning results.** EvoGround remains competitive with existing models. FT denotes training on dense captions in ActivityNet-Captions.

Method	FT	ActivityNet-Captions	
		SODA_c	METEOR
Momentor [47]	✗	2.3	4.7
VTimeLLM [28]	✓	5.8	6.8
TRACE [49]	✓	6.4	6.0
Grounded-VideoLLM [48]	✓	6.2	6.4
EvoGround (Ours)	✗	4.1	7.0

EvoGround generalizes beyond grounding. We evaluate the proposer on dense video captioning [13]. Table 7 reports results on ActivityNet-Captions, showing that EvoGround achieves the highest METEOR score among all models. Despite a lower SODA_c score, it remains competitive without involving any caption data, unlike prior methods fine-tuned on human caption supervision. We further evaluate the solver on representative VideoQA benchmarks, Video-MME [5] and TempCompass [67], in Appendix Table 9, where EvoGround achieves slight improvements over the backbone without disrupting its original QA capabilities.

6 Conclusion

We presented EvoGround, a framework of coupled self-evolving agents that learns temporal grounding from raw videos without any manual labels. Through a proposer-solver loop driven entirely by reinforcement learning, the two agents iteratively challenge and improve each other, yielding a system that matches or outperforms most fully supervised models across multiple benchmarks for both grounding and captioning. This suggests that self-evolution fosters broader video understanding than the training objective alone would imply. We hope this work inspires a broader shift toward annotation-free video understanding systems that learn directly from the raw visual world.

Limitations. Despite the effectiveness of EvoGround, several challenges remain. First, it becomes costly when handling ultra-long videos, as it generates a large number of query–moment pairs and needs the solver’s feedback for each pair. Second, as with other approaches, it can be influenced by the underlying video corpus. While this plays a crucial role, we discuss how the proposed self-evolution framework contributes more to performance gains than dataset-specific advantages (see Appendix A.4). Additionally, the scaling experiments suggest room for further improvement. Addressing these aspects could strengthen EvoGround, and we will explore them in future work.

References

- [1] M. Jung, S. Choi, J. Kim, J.-H. Kim, and B.-T. Zhang, “Modal-specific pseudo query generation for video corpus moment retrieval,” in *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pp. 7769–7781, Dec. 2022.
- [2] J. Lei, T. L. Berg, and M. Bansal, “Detecting moments and highlights in videos via natural language queries,” *Advances in Neural Information Processing Systems*, vol. 34, pp. 11846–11858, 2021.
- [3] J. Xiao, A. Yao, Y. Li, and T.-S. Chua, “Can i trust your answer? visually grounded video question answering,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 13204–13214, 2024.
- [4] Z. Shangguan, C. Li, Y. Ding, Y. Zheng, Y. Zhao, T. Fitzgerald, and A. Cohan, “Tomato: Assessing visual temporal reasoning capabilities in multimodal foundation models,” *arXiv preprint arXiv:2410.23266*, 2024.
- [5] C. Fu, Y. Dai, Y. Luo, L. Li, S. Ren, R. Zhang, Z. Wang, C. Zhou, Y. Shen, M. Zhang, *et al.*, “Video-mme: The first-ever comprehensive evaluation benchmark of multi-modal llms in video analysis,” *arXiv preprint arXiv:2405.21075*, 2024.
- [6] M. Jung, J. Xiao, J. Kim, B.-T. Zhang, and A. Yao, “Egoexo-con: Exploring view-invariant video temporal understanding,” *arXiv preprint arXiv:2510.26113*, 2025.
- [7] J. Gao, C. Sun, Z. Yang, and R. Nevatia, “Tall: Temporal activity localization via language query,” in *Proceedings of the IEEE international conference on computer vision*, pp. 5267–5275, 2017.
- [8] R. Zeng, H. Xu, W. Huang, P. Chen, M. Tan, and C. Gan, “Dense regression network for video grounding,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10287–10296, 2020.
- [9] H. Li, M. Cao, X. Cheng, Y. Li, Z. Zhu, and Y. Zou, “Exploiting auxiliary caption for video grounding,” in *Proceedings of the AAAI conference on artificial intelligence*, 2024.
- [10] Y. Yuan, X. Lan, X. Wang, L. Chen, Z. Wang, and W. Zhu, “A closer look at temporal sentence grounding in videos: Dataset and metric,” in *Proceedings of the 2nd international workshop on human-centric multimedia analysis*, pp. 13–21, 2021.
- [11] M. Otani, Y. Nakashima, E. Rahtu, and J. Heikkilä, “Uncovering hidden challenges in query-based video moment retrieval,” *arXiv preprint arXiv:2009.00325*, 2020.
- [12] M. Jung, J. Xiao, B.-T. Zhang, and A. Yao, “On the consistency of video large language models in temporal comprehension,” in *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 13713–13722, 2025.
- [13] R. Krishna, K. Hata, F. Ren, L. Fei-Fei, and J. C. Niebles, “Dense-captioning events in videos,” in *International Conference on Computer Vision (ICCV)*, 2017.
- [14] Y. Wang, Z. Wang, B. Xu, Y. Du, K. Lin, Z. Xiao, Z. Yue, J. Ju, L. Zhang, D. Yang, *et al.*, “Time-rl: Post-training large vision language model for temporal video grounding,” *arXiv preprint arXiv:2503.13377*, 2025.
- [15] J.-J. Chen, Y.-C. Liao, H.-C. Lin, Y.-C. Yu, Y.-C. Chen, and F. Wang, “Rextime: A benchmark suite for reasoning-across-time in videos,” *Advances in Neural Information Processing Systems*, vol. 37, pp. 28662–28673, 2024.
- [16] Y. Liu, Z. Ma, Z. Qi, Y. Wu, Y. Shan, and C. W. Chen, “Et bench: Towards open-ended event-level video-language understanding,” *Advances in Neural Information Processing Systems*, vol. 37, pp. 32076–32110, 2024.
- [17] M. Cai, R. Tan, J. Zhang, B. Zou, K. Zhang, F. Yao, F. Zhu, J. Gu, Y. Zhong, Y. Shang, *et al.*, “Temporalbench: Benchmarking fine-grained temporal understanding for multimodal video models,” *arXiv preprint arXiv:2410.10818*, 2024.

- [18] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, *et al.*, “Learning transferable visual models from natural language supervision,” in *International conference on machine learning*, pp. 8748–8763, PmLR, 2021.
- [19] S. Zhang, H. Peng, J. Fu, and J. Luo, “Learning 2d temporal adjacent networks for moment localization with natural language,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 12870–12877, 2020.
- [20] H. Zhang, A. Sun, W. Jing, and J. T. Zhou, “Span-based localizing network for natural language video localization,” *arXiv preprint arXiv:2004.13931*, 2020.
- [21] J. Mun, M. Cho, and B. Han, “Local-global video-text interactions for temporal grounding,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10810–10819, 2020.
- [22] M. Jung, Y. Jang, S. Choi, J. Kim, J.-H. Kim, and B.-T. Zhang, “Background-aware moment detection for video moment retrieval,” in *Proceedings of the Winter Conference on Applications of Computer Vision*, pp. 8575–8585, February 2025.
- [23] W.-L. Chiang, Z. Li, Z. Lin, Y. Sheng, Z. Wu, H. Zhang, L. Zheng, S. Zhuang, Y. Zhuang, J. E. Gonzalez, *et al.*, “Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality,” See <https://vicuna.lmsys.org> (accessed 14 April 2023), 2023.
- [24] H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar, *et al.*, “Llama: Open and efficient foundation language models,” *arXiv preprint arXiv:2302.13971*, 2023.
- [25] M. Maaz, H. Rasheed, S. Khan, and F. S. 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.
- [26] K. Li, Y. He, Y. Wang, Y. Li, W. Wang, P. Luo, Y. Wang, L. Wang, and Y. Qiao, “Videochat: Chat-centric video understanding,” *arXiv preprint arXiv:2305.06355*, 2023.
- [27] E. Song, W. Chai, G. Wang, Y. Zhang, H. Zhou, F. Wu, H. Chi, X. Guo, T. Ye, Y. Zhang, *et al.*, “Moviechat: From dense token to sparse memory for long video understanding,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 18221–18232, 2024.
- [28] B. Huang, X. Wang, H. Chen, Z. Song, and W. Zhu, “Vtimellm: Empower llm to grasp video moments,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14271–14280, 2024.
- [29] Y. Guo, J. Liu, M. Li, X. Tang, X. Chen, and B. Zhao, “Vtg-llm: Integrating timestamp knowledge into video llms for enhanced video temporal grounding,” *arXiv preprint arXiv:2405.13382*, 2024.
- [30] S. Ren, L. Yao, S. Li, X. Sun, and L. Hou, “Timechat: A time-sensitive multimodal large language model for long video understanding,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14313–14323, 2024.
- [31] X. Zeng, K. Li, C. Wang, X. Li, T. Jiang, Z. Yan, S. Li, Y. Shi, Z. Yue, Y. Wang, *et al.*, “Timesuite: Improving mllms for long video understanding via grounded tuning,” *arXiv preprint arXiv:2410.19702*, 2024.
- [32] D. Guo, D. Yang, H. Zhang, J. Song, P. Wang, Q. Zhu, R. Xu, R. Zhang, S. Ma, X. Bi, *et al.*, “Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning,” *arXiv preprint arXiv:2501.12948*, 2025.
- [33] K. Feng, K. Gong, B. Li, Z. Guo, Y. Wang, T. Peng, J. Wu, X. Zhang, B. Wang, and X. Yue, “Video-r1: Reinforcing video reasoning in mllms,” *arXiv preprint arXiv:2503.21776*, 2025.
- [34] X. Li, Z. Yan, D. Meng, L. Dong, X. Zeng, Y. He, Y. Wang, Y. Qiao, Y. Wang, and L. Wang, “Videochat-r1: Enhancing spatio-temporal perception via reinforcement fine-tuning,” *arXiv preprint arXiv:2504.06958*, 2025.

- [35] Z. Yan, X. Li, Y. He, Z. Yue, X. Zeng, Y. Wang, Y. Qiao, L. Wang, and Y. Wang, “Videochat-r1.5: Visual test-time scaling to reinforce multimodal reasoning by iterative perception,” *arXiv preprint arXiv:2509.21100*, 2025.
- [36] A. Y. Qwen, B. Yang, B. Zhang, B. Hui, B. Zheng, B. Yu, C. Li, D. Liu, F. Huang, H. Wei, *et al.*, “Qwen2.5 technical report,” *arXiv preprint*, 2024.
- [37] Z. Yue, K. Upasani, X. Yang, S. Ge, S. Nie, Y. Mao, Z. Liu, and D. Wang, “Dr. zero: Self-evolving search agents without training data,” *arXiv preprint arXiv:2601.07055*, 2026.
- [38] C. Huang, W. Yu, X. Wang, H. Zhang, Z. Li, R. Li, J. Huang, H. Mi, and D. Yu, “R-zero: Self-evolving reasoning llm from zero data,” *arXiv preprint arXiv:2508.05004*, 2025.
- [39] A. Zhao, Y. Wu, Y. Yue, T. Wu, Q. Xu, M. Lin, S. Wang, Q. Wu, Z. Zheng, and G. Huang, “Absolute zero: Reinforced self-play reasoning with zero data,” *arXiv preprint arXiv:2505.03335*, 2025.
- [40] Z. Li, H. Du, C. Huang, X. Wu, L. Yu, Y. He, J. Xie, X. Wu, Z. Liu, J. Zhang, *et al.*, “Mm-zero: Self-evolving multi-model vision language models from zero data,” *arXiv preprint arXiv:2603.09206*, 2026.
- [41] Y. He, C. Huang, Z. Li, J. Huang, and Y. Yang, “Visplay: Self-evolving vision-language models from images,” *arXiv preprint arXiv:2511.15661*, 2025.
- [42] O. Thawakar, S. Venkatraman, R. Thawkar, A. Shaker, H. Cholakkal, R. M. Anwer, S. Khan, and F. Khan, “Evolmm: Self-evolving large multimodal models with continuous rewards,” *arXiv preprint arXiv:2511.16672*, 2025.
- [43] S. Bai, K. Chen, X. Liu, J. Wang, W. Ge, S. Song, K. Dang, P. Wang, S. Wang, J. Tang, *et al.*, “Qwen2.5-vl technical report,” *arXiv preprint arXiv:2502.13923*, 2025.
- [44] M. Tschannen, A. Gritsenko, X. Wang, M. F. Naeem, I. Alabdulmohsin, N. Parthasarathy, T. Evans, L. Beyer, Y. Xia, B. Mustafa, *et al.*, “Siglip 2: Multilingual vision-language encoders with improved semantic understanding, localization, and dense features,” *arXiv preprint arXiv:2502.14786*, 2025.
- [45] S.-Y. Liu, X. Dong, X. Lu, S. Diao, P. Belcak, M. Liu, M.-H. Chen, H. Yin, Y.-C. F. Wang, K.-T. Cheng, *et al.*, “Gdpo: Group reward-decoupled normalization policy optimization for multi-reward rl optimization,” *arXiv preprint arXiv:2601.05242*, 2026.
- [46] Y. Wang, X. Meng, J. Liang, Y. Wang, Q. Liu, and D. Zhao, “Hawkeye: Training video-text llms for grounding text in videos,” *arXiv preprint arXiv:2403.10228*, 2024.
- [47] L. Qian, J. Li, Y. Wu, Y. Ye, H. Fei, T.-S. Chua, Y. Zhuang, and S. Tang, “Momentor: Advancing video large language model with fine-grained temporal reasoning,” *arXiv preprint arXiv:2402.11435*, 2024.
- [48] H. Wang, Z. Xu, Y. Cheng, S. Diao, Y. Zhou, Y. Cao, Q. Wang, W. Ge, and L. Huang, “Grounded-videollm: Sharpening fine-grained temporal grounding in video large language models,” *arXiv preprint arXiv:2410.03290*, 2024.
- [49] Y. Guo, J. Liu, M. Li, Q. Liu, X. Chen, and X. Tang, “Trace: Temporal grounding video llm via causal event modeling,” *arXiv preprint arXiv:2410.05643*, 2024.
- [50] S. Pramanick, E. Mavroudi, Y. Song, R. Chellappa, L. Torresani, and T. Afouras, “Enrich and detect: Video temporal grounding with multimodal llms,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 24297–24308, 2025.
- [51] X. Li, Y. Wang, J. Yu, X. Zeng, Y. Zhu, H. Huang, J. Gao, K. Li, Y. He, C. Wang, *et al.*, “Videochat-flash: Hierarchical compression for long-context video modeling,” *arXiv preprint arXiv:2501.00574*, 2024.
- [52] D.-A. Huang, S. Liao, S. Radhakrishnan, H. Yin, P. Molchanov, Z. Yu, and J. Kautz, “Lita: Language instructed temporal-localization assistant,” *arXiv preprint arXiv:2403.19046*, 2024.

- [53] B. Li, Y. Zhang, D. Guo, R. Zhang, F. Li, H. Zhang, K. Zhang, P. Zhang, Y. Li, Z. Liu, *et al.*, “Llava-onevision: Easy visual task transfer,” *arXiv preprint arXiv:2408.03326*, 2024.
- [54] H. Liu, C. Li, Y. Li, B. Li, Y. Zhang, S. Shen, and Y. J. Lee, “Llava-next: Improved reasoning, ocr, and world knowledge,” January 2024.
- [55] P. Zhang, X. Dong, Y. Zang, Y. Cao, R. Qian, L. Chen, Q. Guo, H. Duan, B. Wang, L. Ouyang, *et al.*, “Internlm-xcomposer-2.5: A versatile large vision language model supporting long-contextual input and output,” *arXiv preprint arXiv:2407.03320*, 2024.
- [56] B. Lin, B. Zhu, Y. Ye, M. Ning, P. Jin, and L. Yuan, “Video-llava: Learning united visual representation by alignment before projection,” *arXiv preprint arXiv:2311.10122*, 2023.
- [57] Y. Yao, T. Yu, A. Zhang, C. Wang, J. Cui, H. Zhu, T. Cai, H. Li, W. Zhao, Z. He, *et al.*, “Minicpm-v: A gpt-4v level mllm on your phone,” *arXiv preprint arXiv:2408.01800*, 2024.
- [58] M. Abdin, J. Aneja, H. Behl, S. Bubeck, R. Eldan, S. Gunasekar, M. Harrison, R. J. Hewett, M. Javaheripi, P. Kauffmann, *et al.*, “Phi-4 technical report,” *arXiv preprint arXiv:2412.08905*, 2024.
- [59] B. He, H. Li, Y. K. Jang, M. Jia, X. Cao, A. Shah, A. Shrivastava, and S.-N. Lim, “Mallm: Memory-augmented large multimodal model for long-term video understanding,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 13504–13514, 2024.
- [60] M. Cai, J. Yang, J. Gao, and Y. J. Lee, “Matryoshka multimodal models,” *arXiv preprint arXiv:2405.17430*, 2024.
- [61] R. Vedantam, C. Lawrence Zitnick, and D. Parikh, “Cider: Consensus-based image description evaluation,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4566–4575, 2015.
- [62] K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu, “Bleu: a method for automatic evaluation of machine translation,” in *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pp. 311–318, 2002.
- [63] C.-Y. Lin, “Rouge: A package for automatic evaluation of summaries,” in *Text summarization branches out*, pp. 74–81, 2004.
- [64] N. Reimers and I. Gurevych, “Sentence-bert: Sentence embeddings using siamese bert-networks,” in *Proceedings of the 2019 conference on empirical methods in natural language processing and the 9th international joint conference on natural language processing (EMNLP-IJCNLP)*, pp. 3982–3992, 2019.
- [65] S. Bai, Y. Cai, R. Chen, K. Chen, X. Chen, Z. Cheng, L. Deng, W. Ding, C. Gao, C. Ge, *et al.*, “Qwen3-vl technical report,” *arXiv preprint arXiv:2511.21631*, 2025.
- [66] W. Wang, Z. Gao, L. Gu, H. Pu, L. Cui, X. Wei, Z. Liu, L. Jing, S. Ye, J. Shao, *et al.*, “Internvl3.5: Advancing open-source multimodal models in versatility, reasoning, and efficiency,” *arXiv preprint arXiv:2508.18265*, 2025.
- [67] Y. Liu, S. Li, Y. Liu, Y. Wang, S. Ren, L. Li, S. Chen, X. Sun, and L. Hou, “Tempcompass: Do video llms really understand videos?,” *arXiv preprint arXiv:2403.00476*, 2024.
- [68] J. Skalse, N. Howe, D. Krasheninnikov, and D. Krueger, “Defining and characterizing reward gaming,” *Advances in Neural Information Processing Systems*, vol. 35, pp. 9460–9471, 2022.
- [69] J. Wei, X. Wang, D. Schuurmans, M. Bosma, F. Xia, E. Chi, Q. V. Le, D. Zhou, *et al.*, “Chain-of-thought prompting elicits reasoning in large language models,” *Advances in neural information processing systems*, vol. 35, pp. 24824–24837, 2022.

Appendix

In this section, we provide details that are not included in the main manuscript due to the page limit.

A Details on EvoGround

A.1 Group reward-Decoupled Normalization Policy Optimization (GDPO)

We describe GDPO [45], an RL optimization algorithm derived from GRPO [32]. We first briefly review GRPO. Given a prompt p_i , the model samples G candidate responses $\{o_1, \dots, o_G\}$ from the behavior policy $\pi_{\theta_{\text{old}}}$. GRPO computes a group-normalized advantage $A_{\text{sum}}^{(i,j)}$ by first summing reward components into $R_{\text{sum}}^{(i,j)}$ and then normalizing:

$$A_{\text{sum}}^{(i,j)} = \frac{R_{\text{sum}}^{(i,j)} - \text{mean}\{R_{\text{sum}}^{(i,1)}, \dots, R_{\text{sum}}^{(i,G)}\}}{\text{std}\{R_{\text{sum}}^{(i,1)}, \dots, R_{\text{sum}}^{(i,G)}\}}, \quad R_{\text{sum}}^{(i,j)} = R_1^{(i,j)} + \dots + R_k^{(i,j)}. \quad (11)$$

GRPO then optimizes the policy to favor responses with higher relative advantage. However, such naively summing rewards overlooks differences in their difficulty and scale, often biasing optimization toward easier objectives while neglecting more challenging ones [68]. GDPO addresses this by reversing the order: it normalizes each reward component *individually before* aggregation, rather than after. Specifically, for each reward R_k , GDPO computes a per-reward normalized advantage $A_k^{(i,j)}$:

$$A_k^{(i,j)} = \frac{R_k^{(i,j)} - \text{mean}\{R_k^{(i,1)}, \dots, R_k^{(i,G)}\}}{\text{std}\{R_k^{(i,1)}, \dots, R_k^{(i,G)}\}}, \quad (12)$$

and then combines them with defined weights before a final normalization:

$$A_{\text{sum}}^{(i,j)} = w_1 A_1^{(i,j)} + w_2 A_2^{(i,j)} + \dots + w_k A_k^{(i,j)}, \quad (13)$$

$$\hat{A}_{\text{sum}}^{(i,j)} = \frac{A_{\text{sum}}^{(i,j)} - \text{mean}\{A_{\text{sum}}^{(i',j')} | i' \in \mathcal{B}, j' = 1, \dots, G\}}{\text{std}\{A_{\text{sum}}^{(i',j')} | i' \in \mathcal{B}, j' = 1, \dots, G\}}, \quad (14)$$

which applies batch-wise (\mathcal{B}) advantage normalization. This design enables prioritization of more critical objectives and is particularly suitable for multi-reward settings. We refer the reader to [32] and [45] for the full derivations of GRPO and GDPO, respectively.

A.2 Additional Implementation Details.

Hyperparameters. For both agents, we use the AdamW optimizer with a learning rate of 1e-6. We set the number of rollouts (G) to 2. We use a batch size of 1 with gradient accumulation of 4. We set the max token lengths to 512 for the proposer and to 200 for the solver. This considers their different role, as the proposer likely outputs longer answers than the solver since we require both natural language queries and corresponding timestamps. For the proposer, we set $\tau_{\text{length}} = 3$ and $\gamma = 30$.

Full prompts. In Figure 6, we provide a designed prompt for the proposer and solver, respectively. The proposer is encouraged to generate query–moment samples under valid conditions, and the solver aims to temporally ground the video from the given query sentence. Here, we set N to 4 for the proposer. If the proposer generates valid query–moment pairs less than N , the remaining entries are padded with zero moments to penalize insufficient sampling and encourage the generation of multiple pairs. For instance, if the proposer generates only one valid pair (*e.g.*, covering the entire video), the format reward is reduced to 0.25 since the remaining three pairs are assigned 0.

A.3 Details on the Reward Designs

SigLIP-2. SigLIP-2 [44] is a widely used vision-language encoder that extends the success of the original SigLIP with enhanced pre-training objectives. We use the 0.4B-scale model, which is lightweight and incurs negligible computations, for the consistency reward during training.

Intra-Consistency. In Equation 5, the exponential formulation is designed to map the variance of frame-level similarities to a sharply discriminative score. Since similarity values lie in $[-1, 1]$, the

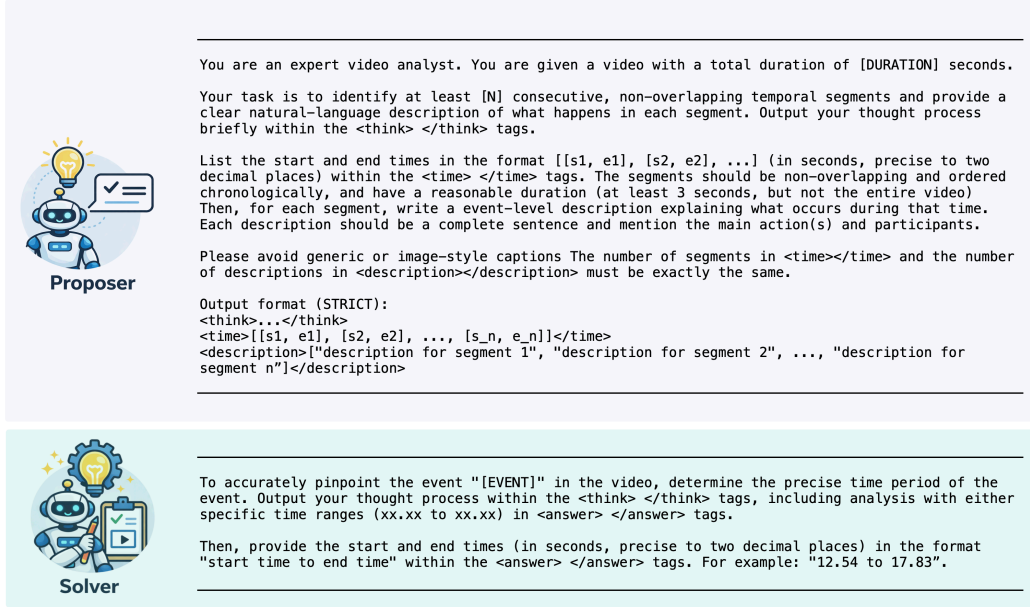


Figure 6: **Prompt designs of EvoGround.** We show the prompts used for the proposer (top) and solver (bottom). The proposer is instructed to generate consecutive, non-overlapping query–moment pairs from a raw video, while the solver is instructed to localize a given query within the video.

resulting variance σ_n is bounded within a limited range. In the ideal case, where all frames within a moment are perfectly aligned with the query, the variance becomes zero, yielding $\exp(0) = 1$. In contrast, highly inconsistent moments exhibit large fluctuations in similarity (*i.e.*, alternating between 1 and -1), resulting in $\sigma_n \approx 1$. Such cases are mapped to a near-zero value via the exponential function, $\exp(-1 \times 30) \approx 9.3 \times 10^{-14}$. This design ensures that well-aligned moments are preserved with high scores, while even moderate inconsistencies are strongly penalized, providing a clear separation between coherent and incoherent moment predictions. Additionally, we determine $|m_n|$ in Equation 5 as:

$$|m_n| = \text{clip}((e_n - s_n) \cdot \rho, T_{\min}, T_{\max}), \quad (15)$$

where $\rho = 1$, $T_{\min} = 8$, and $T_{\max} = 32$. This ensures that longer moments are represented with more sampled frames.

Variants of Inter-Consistency. We consider two variants of inter-consistency: a hard ranking-based formulation and an exponential formulation:

$$\text{Inter}_n^{\text{hard}} = \frac{1}{N} \sum_{j=1}^N \mathbb{I}(\mu_n \geq \mu_{n,j}), \quad (16)$$

which measures the fraction of candidates that the n -th sample outperforms, capturing its relative ranking within the group. We also design the function with an exponential formulation:

$$\text{Inter}_n^{\text{exp}} = \begin{cases} 1 - \exp\left(-\gamma_{\text{inter}} \cdot \frac{1}{N-1} \sum_{j \neq n} \max(0, \mu_n - \mu_{n,j})\right), & \text{if } N > 1, \\ 0, & \text{otherwise.} \end{cases} \quad (17)$$

which provides a smooth and more sensitive measure of relative superiority, where larger margins between μ_n and other candidates lead to exponentially higher scores. We set γ_{inter} to 30. In Table 8, we compare the performance of applying these inter-consistency variants and confirm that computing the inter-consistency using a softmax demonstrates the best performance. This is likely due to its smooth and more discriminative treatment of relative differences among candidates compared to other variants.

Table 8: **Ablation on the variants of inter-consistency.**

Intra	TVGBench		
	R1@0.3	R1@0.5	R1@0.7
hard	41.7	27.3	15.0
exp	40.8	26.9	14.3
softmax	42.1	28.5	15.3

Table 9: **Results on Video-MME and TempCompass.** * denotes our reproduced results.

Method	Video-MME	TempCompass
	w/o subs	multi-choice QA
Qwen2.5-VL-7B* [43]	61.4	71.3
Time-R1* [14]	61.2	71.2
EvoGround (Ours)	62.3	71.5




Figure 7: **Query length distribution across iterations.** We visualize the query length distributions across different reward configurations and iterations. As previously discussed in Section 5, the feedback reward increases the length of queries compared to others.

A.4 Impact of Training Data and Comparison with Time-R1

A potential concern is that the performance of EvoGround may be influenced by the underlying video corpus used for training. We acknowledge that the choice of data can affect model performance. To control for this factor, we note that both EvoGround and Time-R1 [14] are trained on the same underlying video corpus (*i.e.*, TimeRFT). Therefore, any performance differences cannot be attributed solely to dataset selection.

Despite sharing the same video corpus, EvoGround differs from Time-R1 in several aspects. First, EvoGround does not rely on any human-annotated data, whereas Time-R1 is trained with labeled supervision. Second, EvoGround does not require a cold-start fine-tuning stage or the construction of curated chain-of-thought (CoT) [69] examples, both of which are involved in Time-R1 to guide training. Time-R1 further requires dynamic hard sampling during its training that mines hard samples (*i.e.*, low-IoU samples) on a curated dataset for each training epoch. In contrast, EvoGround learns directly from raw videos through a self-evolving proposer–solver loop driven by RL.

Beyond, EvoGround is more data-efficient than existing models, achieving strong performance with only a small amount of raw video data. As mentioned in Section 4.2, existing models typically rely on multi-stage training pipelines with curated datasets tailored for each stage. For example, VTimeLLM [28] adopts a three-stage framework (feature alignment, boundary perception, and instruction tuning), while Grounded-VideoLLM [48] employs a similar multi-stage pipeline that involves large-scale caption supervision (*i.e.*, 1.28M video-caption pairs) in early stages.

Taken together, these differences suggest that the improvement of EvoGround stems from the proposed self-evolution framework rather than reliance on curated video quality or dataset-specific advantages. While the underlying corpus still plays a role, our results demonstrate that strong temporal grounding performance can be achieved without annotated data or additional manual curation.

B Further Analysis

Results on VideoQA benchmarks. We evaluate EvoGround on representative video question answering (QA) benchmarks, including Video-MME [5] and TempCompass [67] in Table 9. We




Figure 8: **Per-sample IoU improvements under different thresholds δ .** The dashed line represents no change, and the upper-left regions represent improved cases.

Table 10: **Results on Charades-STA, ActivityNet, and TVGBench.**

Method	Charades-STA				ActivityNet-Captions				TVGBench			
	R1@0.3	R1@0.5	R1@0.7	mIoU	R1@0.3	R1@0.5	R1@0.7	mIoU	R1@0.3	R1@0.5	R1@0.7	mIoU
EvoGround (2.5K)	77.2	60.5	35.5	53.1	62.9	43.6	25.0	43.9	42.1	28.5	15.3	29.4
EvoGround (7.5K)	81.7	67.7	45.6	60.3	64.8	44.3	25.8	44.7	41.7	28.6	15.1	29.5
Δ	+4.5	+7.2	+10.1	+7.2	+1.9	+0.7	+0.8	+0.8	-0.4	+0.1	-0.2	+0.1

confirm that EvoGround does not disrupt the backbone (*i.e.*, Qwen2.5-VL-7B) and demonstrates a slight improvement.

Query length distributions across iterations and different reward configurations. In Figure 7, we visualize the query length distribution for each reward configuration across iterations. As previously mentioned, the feedback reward enriches query descriptiveness, showing the highest average query length.

Per-sample comparison across different δ . In Figure 8, we compare per-sample improvements across different δ on TVGBench. $\delta = 0.5$ results in a denser distribution in the high-IoU region (*i.e.*, upper-left), compared to $\delta = 0.3$, along with higher average improvements (+1.29% vs. +0.21%).

Evaluation on TemporalBench. As shown in Figure 9, to evaluate on TemporalBench [17], we first generate query–moment pairs from given videos and concatenate them into a single paragraph. We then compare it with the ground-truth sentence and confirm that EvoGround effectively captures consecutive video events.

Visualization of generated data. We also visualize generated data from the proposer at the third iteration in Figure 10. Each query is well-aligned with the corresponding frames within the moment and densely covers throughout the video.

Results of EvoGround with scaling up video data. Extending the results in Table 4, we report the performance of EvoGround (7.5K) on additional benchmarks including Charades-STA in Table 10. While we observe some improvements across benchmarks and metrics, the gains are modest, particularly on out-of-distribution datasets such as TVGBench and ActivityNet-Captions. This is expected, as the additional training videos are sourced from Charades-STA, which primarily consists of short videos. Consequently, incorporating such data helps the model better handle similar domains, but is less effective for broader generalization.



Ground-truth: Women dressed in blue uniforms and hats walk in two lines while holding hands. They are smiling and waving at the crowd of people who are cheering them on by clapping their hands on the side. A group of men wearing heavy dark jackets and hat walk through while shaking hands and smiling at the crowd.

EvoGround: Students in robes walk through a grand hall, greeted by applause from others. The group of students continues walking, with some students in uniform hats and robes. The students pass by a group of adults in robes, who are clapping and smiling. The students continue walking, with the crowd of students and adults still applauding and smiling.

Figure 9: Captioning results on TemporalBench.

gA5sihQiKrE_00:00:02:700_00:00:40:600.mp4 (37.9 s)



[0.0s, 3.9s] The video begins with a collage of cooking scenes, introducing the topic of selecting the right cookware for an induction cooktop.



[3.9s, 9.9s] A black induction cooktop appears with various cookware items, including pots and pans, being arranged on it, demonstrating the compatibility of different cookware.



[9.9s, 19.8s] A close-up of a black frying pan with a handle is shown, followed by a shot of a stainless steel pan being turned, highlighting the features of the cookware.



[19.8s, 37.9s] The video concludes with a demonstration of the pans and pots being placed on the induction cooktop, showcasing the process of using the cookware effectively.

Figure 10: Generated query–moment pairs from EvoGround.