

ZOOM-ZERO: REINFORCED COARSE-TO-FINE VIDEO UNDERSTANDING VIA TEMPORAL ZOOM-IN

000
001
002
003
004
005 **Anonymous authors**
006 Paper under double-blind review
007
008
009
010

ABSTRACT

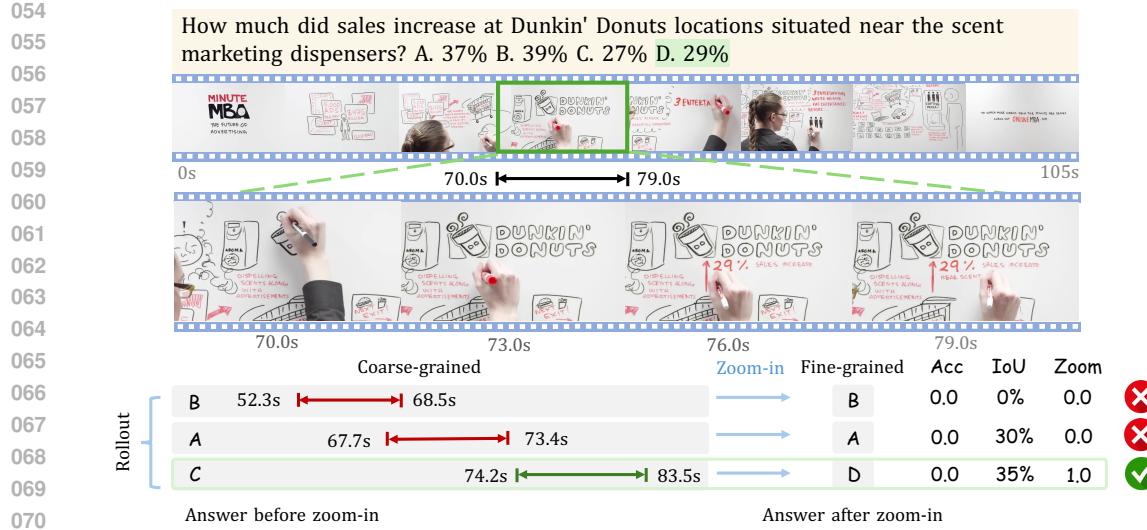
011 Grounded video question answering (GVQA) aims to localize relevant temporal
012 segments in videos and generate accurate answers to a given question; however,
013 large video-language models (LVLMs) exhibit limited temporal awareness. Al-
014 though existing approaches based on Group Relative Policy Optimization (GRPO)
015 attempt to improve temporal grounding, they still struggle to faithfully ground their
016 answers in the relevant video evidence, leading to temporal mislocalization and
017 hallucinations. In this work, we present **Zoom-Zero**, a coarse-to-fine framework
018 that first localizes query-relevant segments and then temporally zooms into the
019 most salient frames for finer-grained visual verification. Our method addresses
020 the limits of GRPO for the GVQA task with *two key innovations*: **(i)** a zoom-in
021 accuracy reward that validates the fidelity of temporal grounding prediction and
022 facilitates fine-grained visual verification on grounded frames; **(ii)** token-selective
023 credit assignment, which attributes rewards to the tokens responsible for temporal
024 localization or answer generation, mitigating GRPO’s issue in handling multi-
025 faceted reward signals. Our proposed method advances grounded video question
026 answering, improving temporal grounding by 5.2% on NExT-GQA and 4.6% on
027 ReXTime, while also enhancing average answer accuracy by 2.4%. Additionally,
028 the coarse-to-fine zoom-in during inference further benefits long-form video under-
029 standing by preserving critical visual details without compromising global context,
030 yielding an average improvement of 6.4% on long-video benchmarks. Our code
031 will be publicly available¹.
032

1 INTRODUCTION

033 Large video-language models (LVLMs) have achieved remarkable progress in video understanding (Li
034 et al., 2023b; 2024b; Cheng et al., 2024; Lin et al., 2024; Luo et al., 2023; Ataallah et al., 2024).
035 However, current LVLMs often struggle to remain faithfully grounded in key visual evidence, leading
036 to hallucinations when reasoning across video sequences. To evaluate this critical capability, video
037 temporal grounding (VTG) (Gao et al., 2017; Anne Hendricks et al., 2017; Lei et al., 2021) measures
038 how well models localize segments given an explicit event description, while the more comprehensive
039 task of grounded video question answering (GVQA) (Xiao et al., 2024) requires models to implicitly
040 infer the relevant moment from a general question for temporal localization and simultaneously
041 generating accurate answers.
042

043 The key challenge of GVQA lies in achieving precise temporal localization while maintaining general
044 video understanding capabilities. Reinforcement learning (RL) offers a promising solution for
045 sharpening specific capabilities while preserving generalization from a pretrained LVLM (Lai et al.,
046 2025). Recent efforts (Li et al., 2025b; Feng et al., 2025) have explored GRPO-based (Shao et al.,
047 2024) RL algorithm for video temporal grounding and reasoning. However, most approaches (Wang
048 et al., 2025b; Chen et al., 2025b) optimize with only format and Intersection over Union (IoU)
049 rewards, neglecting the quality of the generated answers. Although VideoChat-R1 (Li et al., 2025b)
050 incorporates an answer accuracy reward, these training objectives still cannot guarantee that localized
051 video segments actually contain the visual evidence required for correct reasoning. Moreover, limited
052 context budgets compel models to depend on coarse-grained representations, which overlook the
053

¹Please refer to the anonymous GitHub link for access to the code.



- 108 • We introduce a zoom-in accuracy reward that verifies localized segments contain the vi-
109 sual evidence required for correct reasoning in a finer-grained manner, enhancing both
110 localization precision and answer accuracy.
- 111 • We identify and address the limit of GRPO in handling multi-faceted reward signals by
112 selective token-level credit assignment, enabling effective learning from diverse reward
113 signals in GVQA.
- 114 • Our coarse-to-fine paradigm further enhances long-form video understanding by first
115 coarsely identifying key segments and then zooming into fine-grained details, preserving
116 global context while capturing critical information, resulting in an average 6.4% improve-
117 ment on long-video benchmarks.

119 2 RELATED WORK

120 **Large Video Language Models.** Multimodal large language models (MLLMs) (Zhu et al., 2024;
121 Liu et al., 2024b;a; Tong et al., 2024; Chen et al., 2023) have demonstrated remarkable progress
122 in vision-language tasks. Recent advancements have further extended their capabilities to video
123 understanding tasks (Li et al., 2023b; 2024b; Cheng et al., 2024; Lin et al., 2024; Luo et al., 2023;
124 Ataallah et al., 2024). Large Video Language Models (LVLMs) process videos by extracting and
125 encoding frames, and then rearranging them into final video representations. Some approaches (Li
126 et al., 2023b; 2024b; Cheng et al., 2024) leverage the Q-Former module from BLIP-2 (Li et al., 2023a)
127 to integrate visual and textual features, while others (Lin et al., 2024; Luo et al., 2023; Ataallah et al.,
128 2024) directly concatenate frame features. To address intensive video tokens for long videos, several
129 works train on sparsely sampled frames (Li et al., 2023b; Ataallah et al., 2024; Cheng et al., 2024;
130 Zhang et al., 2024b; Li et al., 2024a), while others try to handle long videos by token pooling (Maaz
131 et al., 2023; Li et al., 2024c; Song et al., 2024), token compression (Shen et al., 2025), memory
132 aggregation (He et al., 2024), or frame selection (Hu et al., 2025; Zhang et al., 2025; Wu et al., 2019;
133 Tang et al., 2025). Unlike frame-selection methods that search in the embedding space and select a
134 fixed set of frames, our approach tackles the long-video token challenge by explicitly enhancing the
135 model’s temporal grounding capability through reasoning over the user query.

136 **Grounded Video Question Answering.** Video Temporal Grounding (Gao et al., 2017; Anne Hen-
137 dricks et al., 2017; Lei et al., 2021) localizes relevant segments given an explicit event description.
138 The more advanced task of Grounded Video Question Answering (GVQA) (Xiao et al., 2024) re-
139 quires models to implicitly infer the relevant segment from a general question to perform localization
140 and question-answering jointly. Recent LVLM-based approaches reformulate grounding as text
141 generation (Nie et al., 2024; Ren et al., 2024; Huang et al., 2024a; Li et al., 2025b; Feng et al., 2025)
142 while other methods (Wang et al., 2025a; Huang et al., 2024b) expand vocabularies to learn temporal
143 embeddings for improved precision. Our approach leverages Qwen2.5VL (Bai et al., 2025) to predict
144 textual temporal spans and introduces a novel coarse-to-fine training paradigm: initially predicting
145 coarse timestamps for global localization, then dynamically zooming into identified segments for
146 high-resolution visual verification. In contrast with a concurrent work (Li et al., 2025c) that relies on
147 separate off-the-shelf VideoQA models to answer the query based on localized segments, our unified
148 framework seamlessly integrates temporal grounding with question-answering within a single model
149 for coherent video understanding.

150 **Reinforcement Learning for Grounded Video Question Answering.** Reinforcement learning (RL)
151 has emerged as a powerful paradigm for improving the reasoning ability of large language models.
152 Breakthroughs such as OpenAI-o1 (Jaech et al., 2024) and DeepSeek-R1 (Guo et al., 2025) have
153 demonstrated notable success in addressing complex problems. DeepSeek-R1 (Guo et al., 2025)
154 adopts group relative policy optimization (GRPO) to train LLMs to incentivize reasoning capability
155 at inference time. Recently, RL has been adapted to LVLMs with the goal of strengthening video
156 reasoning (Li et al., 2025b; Feng et al., 2025). Time-R1 (Wang et al., 2025b) and TVG-R1 (Chen
157 et al., 2025b) adopt a two-stage pipeline, beginning with supervised fine-tuning (SFT) as a cold
158 start, followed by GRPO-based RL training, while TimeZero (Wang et al., 2025b) demonstrates
159 that a purely GRPO approach can be more effective without an SFT stage. These methods leverage
160 only format and Intersection over Union (IoU) reward, whereas VideoChat-R1 (Li et al., 2025b)
161 further integrates answer accuracy into RL training. In this work, we enhance GRPO by decoupling
multi-faceted reward signals for selective token-level advantage estimation.

162 **3 PRELIMINARY**
 163

164 **GRPO.** Group Relative Policy Optimization (GRPO) (Shao et al., 2024) is a variant of Proximal
 165 Policy Optimization (PPO) (Schulman et al., 2017) for reinforcement learning. Unlike PPO, which
 166 relies on a critic model, GRPO directly compares groups of candidate responses. This design
 167 eliminates the dependency on a critic, thereby substantially reducing training costs. Given a question-
 168 answer pair (q, a) , policy $\pi_{\theta_{\text{old}}}$ generates G distinct candidate responses $o = o_1, \dots, o_G$ through
 169 policy sampling. Then, the verifiable reward(s) r_1, \dots, r_G is calculated for each response. GRPO
 170 normalizes the scores by computing their mean and standard deviation, and then evaluates the relative
 171 quality of the responses accordingly.

172
$$A_{i,t} = \frac{r_i - \text{mean}(\{r_i\}_{i=1}^G)}{\text{std}(\{r_i\}_{i=1}^G)}, \quad (1)$$

 173
 174

175 where $A_{i,t}$ denotes the relative quality of the t -th token in i -th response. GRPO promotes higher-
 176 scoring answers within each group while regularizing the policy π_{θ} against the reference parameters
 177 π_{ref} via a KL-divergence penalty $D_{\text{KL}}(\cdot | \cdot)$, leading to the final objective:

178
$$\max_{\pi_{\theta}} \mathbb{E}_{(q, a), \{o_i\}_{i=1}^G \sim \pi_{\theta_{\text{old}}}(\cdot | q)} \left[\frac{1}{G} \sum_{i=1}^G \frac{1}{|o_i|} \sum_{t=1}^{|o_i|} \left(\frac{\pi_{\theta}(o_{i,t} | q, o_{i,<t})}{\pi_{\theta_{\text{old}}}(o_{i,t} | q, o_{i,<t})} \cdot A_{i,t} - \beta D_{\text{KL}}(\pi_{\theta} \| \pi_{\text{ref}}) \right) \right], \quad (2)$$

 179
 180
 181

182 where β is a regularization coefficient, preventing excessive deviation from the reference policy
 183 during optimization.

184 **Dynamic Spatiotemporal Resolution.** Qwen2.5-VL (Bai et al., 2025) dynamically adjusts tokens
 185 to determine the number of tokens per frame under a fixed video context budget. Specifically, the
 186 video context size is denoted as L_v , the maximum tokens per frame as V_{max} , the minimum tokens
 187 per frame as V_{min} , the video duration as F seconds, and the sampling rate as s frames per second.
 188 Based on these, the per-frame token resolution V_{res} is defined as follows:

189
$$N = \min \left(F * s, \frac{L_v}{V_{\text{min}}} \right), \quad V_{\text{res}} = \max \left(V_{\text{min}}, \min \left(\frac{L_v}{N}, V_{\text{max}} \right) \right) \quad (3)$$

 190
 191
 192

193 **4 ZOOM-ZERO**
 194

195 We propose a coarse-to-fine framework for grounded video question answering: a coarse-grained pass
 196 predicts query-conditioned intervals, followed by a fine-grained zoom-in that takes as input only the
 197 localized segments at higher per-frame token resolution (Section 4.1). Beyond standard format, IoU,
 198 and answer-accuracy rewards, we introduce a zoom-in accuracy reward to verify evidence within the
 199 localized span (Section 4.2). To overcome GRPO’s limit in uniform credit assignment, we develop
 200 token-selective credit assignment for finer-grained advantage estimation tailored to multi-faceted
 201 rewards in the GVQA task (Section 4.3).

202 **4.1 COARSE-TO-FINE VIDEO UNDERSTANDING VIA TEMPORAL ZOOM-IN**
 203

204 While dynamic token allocation offers flexibility, a fundamental trade-off remains: capturing long-
 205 range temporal context versus preserving fine-grained visual detail. Spatial or temporal downsampling
 206 inevitably discards critical information. This problem is exacerbated in longer videos, where preserving
 207 more frames often comes at the expense of per-frame spatial granularity. A coarse-to-fine strategy
 208 provides a principled remedy: a coarse pass preserves temporal context, followed by a fine-grained
 209 stage that processes evidence-bearing segments through temporal zoom-in².

210 More specifically, we leverage the model’s temporal grounding capability to perform a fine-grained
 211 zoom-in on relevant segments and recover the details of the video. From the coarse view of the video,
 212

213 ²We term it temporal zoom-in, not spatiotemporal, to avoid confusion, since no spatial regions are predicted
 214 from the model. However, the spatio-temporal grid size per token is increased, since salient frames are sampled
 215 at higher temporal resolution (if the full video was sparsely sampled in the coarse pass), and spatial resolution is
 dynamically increased.

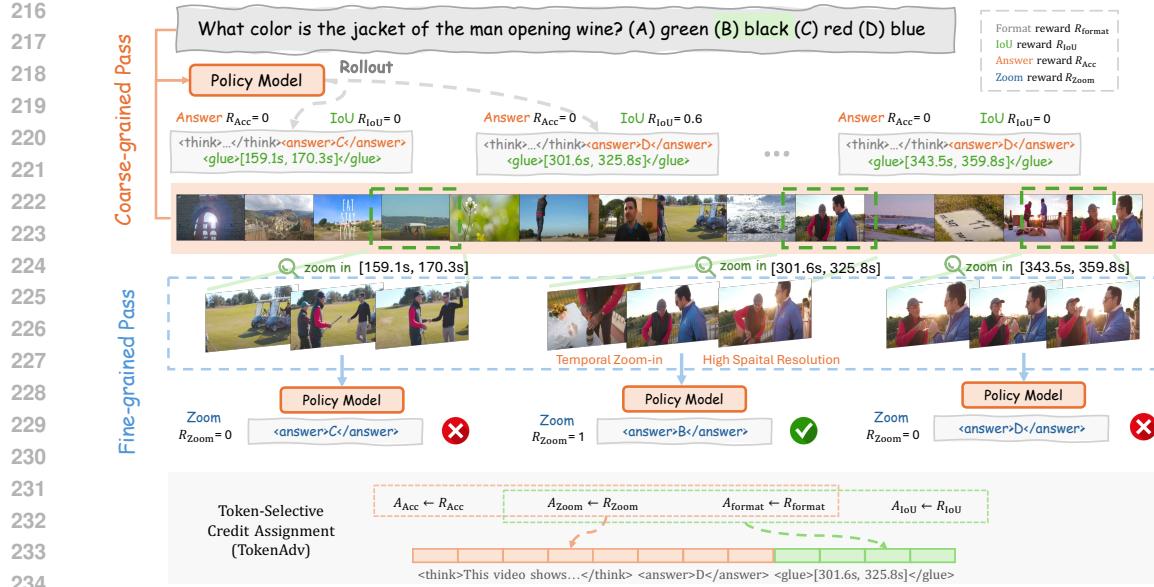


Figure 2: We present **Zoom-Zero**, a coarse-to-fine training pipeline that first rolls out samples to localize relevant segments with preliminary answers, followed by a fine-grained pass by zooming into spotlight segments and dynamically allocating high-resolution video tokens. The zoom reward enforces fine-grained visual verification of the predicted temporal span. In this example, only a faithful span prediction with the correct final answer yields the highest reward. Then we propose token-selective credit assignment (TokenAdv) for a finer-grained advantage estimation.

we obtain grounded start–end pairs $(s_1, e_1), (s_2, e_2), \dots, (s_n, e_n)$. We crop the video accordingly, yielding $N' < N$ frames. Under a fixed visual token budget L_v , the per-frame video tokens increases from $V_{\text{res}} = \frac{L_v}{N}$ to $V'_{\text{res}} = \frac{L_v}{N'} > V_{\text{res}}$, enabling more fine-grained visual verification of the selected segments. This coarse-to-fine temporal zoom-in preserves global context while concentrating high-resolution capacity on the frames that matter most.

Crucially, this paradigm hinges on accurate, query-conditioned temporal grounding. To this end, we leverage GRPO-based reinforcement learning with carefully designed rewards that jointly improve temporal grounding and question answering, as detailed in the following section.

4.2 REWARDS DESIGN

In this section, we first review the basic rewards used in GVQA, i.e., format, temporal grounding, and answer accuracy, and then introduce our proposed zoom-in reward for fine-grained visual verification.

Format Reward. To guide the model toward producing responses in the desired format, we require the output to follow the instructions below:

Format Prompt and Template

Answer the question: [QUESTION] according to the content of the video. Select the answer from: [OPTIONS]. Output key information relevant to the question and options, marking precise timestamps or time ranges in seconds within `<time> </time>` tags, and present them in an interleaved analysis format. Enclose the full analysis in `<think> </think>` tags. Then, provide your answer within the `<answer> </answer>` tags, output the corresponding letter of the option. At the same time, in the `<glue> </glue>` tags, include only the precise video segments (in seconds) that strongly support your answer, in the format of $[(s_1, e_1), (s_2, e_2), \dots]$. For example: `<answer>A</answer><glue>[(20.3, 30.8)]</glue>`.

We then apply regular expression matching to verify whether the model output conforms to this format. R_{format} is assigned as 1 if the format fully matches the template.

270 **Answer Accuracy Reward.** We define reward R_{Acc} to evaluate the correctness of the policy model’s
 271 answer in coarse understanding by taking as input the whole video.

272 **Temporal Grounding Reward.** For temporal grounding, the model is required to predict a timestamp
 273 interval that specifies the video segment relevant to the given textual query. To evaluate this prediction,
 274 we adopt the Intersection over Union (IoU) between the model-predicted interval (from `<glue>`
 275 `</glue>`) and the ground-truth interval as the reward function. $R_{\text{IoU}} = \frac{|\mathcal{I}_{\text{pred}} \cap \mathcal{I}_{\text{gt}}|}{|\mathcal{I}_{\text{pred}} \cup \mathcal{I}_{\text{gt}}|}$, where $\mathcal{I}_{\text{pred}}$
 276 and \mathcal{I}_{gt} are the predicted time intervals and the ground truth intervals.

277 **Zoom Accuracy Reward.** Based on the temporal grounding prediction (from `<glue>` `</glue>`)
 278 in the coarse pass, we can obtain a set of salient frames that enables fine-grained visual verification.
 279 In the finer-grained pass, the model takes as input the question and the zoomed-in frames from
 280 the coarse response to produce the final answer. The reward R_{Zoom} is assigned a value of 1 if the
 281 model produces an accurate final answer. This reward provides two key benefits: (1) enabling visual
 282 verification to ensure the predicted timestamp is accurately grounded in the relevant frames, and (2)
 283 facilitating a coarse-to-fine visual zoom-in to capture details within key frames.

284
 285 **4.3 TOKEN-SELECTIVE CREDIT ASSIGNMENT FOR ADVANTAGE ESTIMATION**

286 Since our approach involves multiple rewards, i.e., R_{format} , R_{Acc} , R_{Zoom} , and R_{IoU} , the key question
 287 becomes how to leverage them for the policy updates. Standard GRPO handles multi-faceted rewards
 288 by naively summing them into a single scalar, thereby collapsing the contributions of individual
 289 reward signals. The advantage is then estimated only from this aggregated value (Equation 1), which
 290 cannot be decoupled for gradient updates. As a result, the model receives no explicit guidance on
 291 which aspect of its behavior each reward reflects, making it difficult to attribute feedback to specific
 292 abilities. In addition, the same advantage is assigned uniformly across all tokens in a response, which
 293 hides the contribution of each token from its corresponding rewards. Appendix C provides a simple
 294 example illustrating this limitation.

295 To overcome these limitations, we propose TokenAdv, a token-selective credit assignment for fine-
 296 grained token-level advantage estimation. Instead of summing up all rewards into one value for
 297 advantage estimation, we decouple advantage calculation separately for each reward type (Equation 4).
 298 In our case, since the outputs for answering and temporal grounding are explicitly formatted with
 299 task-specific tokens, it is feasible to distinguish the contribution of corresponding tokens to each
 300 aspect. Specifically, the token-level advantage is computed by averaging the relevant task-specific
 301 advantages for each token (Equation 5). This design allows the model to attribute feedback to specific
 302 rewards, improving its ability to learn from diverse, multi-faceted signals.

$$A_i^k = \frac{r_i^k - \text{mean}(\{r_i^k\}_{i=1}^G)}{\text{std}(\{r_i^k\}_{i=1}^G)}, \quad r^k \in \{R_{\text{format}}, R_{\text{Acc}}, R_{\text{Zoom}}, R_{\text{IoU}}\} \quad (4)$$

$$A_{i,t} = \begin{cases} \text{mean}(A_i^{\text{format}}, A_i^{\text{Zoom}}, A_i^{\text{IoU}}) & \text{if } o_{i,t} \in \text{<glue>} \dots \text{</glue>} \\ \text{mean}(A_i^{\text{format}}, A_i^{\text{Zoom}}, A_i^{\text{Acc}}) & \text{else} \end{cases} \quad (5)$$

311 By selectively assigning credits to task-specific tokens, we guide policy gradient updates toward the
 312 most influential parts of the output for each capability. This targeted credit assignment allows the
 313 model to effectively leverage diverse reward signals, leading to improved temporal grounding and
 314 question-answering performance, as shown in Figure 4.

315
 316 **5 EXPERIMENTS**

317
 318 **5.1 EXPERIMENTAL SETUPS**

319 **Benchmarks and Evaluation Metrics.** We evaluate on three GVQA benchmarks: NExT-GQA (Xiao
 320 et al., 2024), ReXTime (Chen et al., 2024a), and CG-Bench (Chen et al., 2025a). Temporal grounding
 321 is measured by mean Intersection-over-Union (mIoU), R@0.3 (IoU > 0.3), and R@0.5 (IoU >
 322 0.5); video understanding by multiple-choice question (MCQ) accuracy; **Acc@GQA** measures the
 323 percentages of questions that are correctly answered and visually grounded, i.e., IoP ≥ 0.5 , where

324
 325 Table 1: Grounded video question answering results on NExT-GQA (Xiao et al., 2024) and ReX-
 326 Time (Chen et al., 2024a). All models are of comparable scale (7B or 8B).

327 328 329 330 331 332 333 334 335 336 337 338 339 340 341 342 343 344 345 346 347 348 349 350 351 352 353 354 355 356 357 358 359 360 361 362 363 364 365 366 367 368 369 370 371 372 373 374 375 376 377	327 328 329 330 331 332 333 334 335 336 337 338 339 340 341 342 343 344 345 346 347 348 349 350 351 352 353 354 355 356 357 358 359 360 361 362 363 364 365 366 367 368 369 370 371 372 373 374 375 376 377 Models	327 328 329 330 331 332 333 334 335 336 337 338 339 340 341 342 343 344 345 346 347 348 349 350 351 352 353 354 355 356 357 358 359 360 361 362 363 364 365 366 367 368 369 370 371 372 373 374 375 376 377 NExT-GQA				327 328 329 330 331 332 333 334 335 336 337 338 339 340 341 342 343 344 345 346 347 348 349 350 351 352 353 354 355 356 357 358 359 360 361 362 363 364 365 366 367 368 369 370 371 372 373 374 375 376 377 ReXTime			
		327 328 329 330 331 332 333 334 335 336 337 338 339 340 341 342 343 344 345 346 347 348 349 350 351 352 353 354 355 356 357 358 359 360 361 362 363 364 365 366 367 368 369 370 371 372 373 374 375 376 377 Acc@GQA	327 328 329 330 331 332 333 334 335 336 337 338 339 340 341 342 343 344 345 346 347 348 349 350 351 352 353 354 355 356 357 358 359 360 361 362 363 364 365 366 367 368 369 370 371 372 373 374 375 376 377 mIoU	327 328 329 330 331 332 333 334 335 336 337 338 339 340 341 342 343 344 345 346 347 348 349 350 351 352 353 354 355 356 357 358 359 360 361 362 363 364 365 366 367 368 369 370 371 372 373 374 375 376 377 R@0.3	327 328 329 330 331 332 333 334 335 336 337 338 339 340 341 342 343 344 345 346 347 348 349 350 351 352 353 354 355 356 357 358 359 360 361 362 363 364 365 366 367 368 369 370 371 372 373 374 375 376 377 R@0.5	327 328 329 330 331 332 333 334 335 336 337 338 339 340 341 342 343 344 345 346 347 348 349 350 351 352 353 354 355 356 357 358 359 360 361 362 363 364 365 366 367 368 369 370 371 372 373 374 375 376 377 Acc	327 328 329 330 331 332 333 334 335 336 337 338 339 340 341 342 343 344 345 346 347 348 349 350 351 352 353 354 355 356 357 358 359 360 361 362 363 364 365 366 367 368 369 370 371 372 373 374 375 376 377 mIoU	327 328 329 330 331 332 333 334 335 336 337 338 339 340 341 342 343 344 345 346 347 348 349 350 351 352 353 354 355 356 357 358 359 360 361 362 363 364 365 366 367 368 369 370 371 372 373 374 375 376 377 R@0.3	327 328 329 330 331 332 333 334 335 336 337 338 339 340 341 342 343 344 345 346 347 348 349 350 351 352 353 354 355 356 357 358 359 360 361 362 363 364 365 366 367 368 369 370 371 372 373 374 375 376 377 R@0.5
SFT-based									
Qwen2.5-VL (Bai et al., 2025)	18.9	20.2	31.6	18.1	51.1	27.4	36.1	24.8	
RL-based									
TimeChat (Ren et al., 2024)	7.6	20.6	34.1	17.9	40.0	11.6	14.4	7.6	
VTimeLLM (Huang et al., 2024a)	17.4	24.4	36.1	20.1	36.1	20.1	28.8	17.4	
Grounded-VideoLLM (Wang et al., 2025a)	26.7	21.1	-	18.0	-	-	-	-	
CG-Bench (Chen et al., 2025a)									
Zoom-Zero (Ours)	29.0	37.6	55.6	33.8	62.0	43.2	56.5	44.1	

IoP is the intersection over prediction span. CG-Bench (Chen et al., 2025a), a long-form GQA benchmark, additionally introduces two metrics rec. @ IoU and acc. @ IoU: rec. @ IoU averages recall over IoU thresholds $\{0.1, 0.2, 0.3, 0.4, 0.5\}$ to estimate the probability of correctly retrieving clue intervals; acc. @ IoU counts a response as correct only if the predicted answer is accurate and its IoU exceeds the threshold, and is averaged over the same thresholds per the original protocol. We also assess general video understanding on four long-video benchmarks, VideoMME (Fu et al., 2025), MLVU (Zhou et al., 2025), LVBench (Wang et al., 2024) and CG-Bench (Chen et al., 2025a), and report MCQ accuracy.

Baselines. We compare our approach with several strong baselines, including SFT-based LVLMs with grounding capability such as VTimeLLM (Huang et al., 2024a), TimeChat (Ren et al., 2024) and Grounded-VideoLLM (Wang et al., 2025a), RL-based methods VideoChat-TPO (Li et al., 2025a), TVG-R1 (Chen et al., 2025b) VideoChat-R1 (Li et al., 2025b) as well as general-purpose LVLMs such as LLaVA-OneVision (Li et al., 2024a), Qwen2.5-VL (Bai et al., 2025) and InternVL2.5 (Chen et al., 2024b). All open-sourced models are of comparable scale (7B or 8B). All SFT-based models are evaluated in a zero-shot setting on NExT-GQA (Xiao et al., 2024). RL methods, VideoChat-R1 (Li et al., 2025b), and our model are trained on the NExT-GQA val split. For ReXTime (Chen et al., 2024a) (Table 1, right) and CG-Bench (Chen et al., 2025a) (Table 2, rightmost), all models are evaluated strictly in the zero-shot setting, ensuring a valid and fair comparison across methods.

Training Details. We adopt Qwen2.5-VL-7B (Bai et al., 2025) as the base model. The maximum number of video tokens is set to 8192, with videos sampled at 1 fps during training. The minimum video frame resolution is $16 \times 28 \times 28$ pixels and the maximum is $768 \times 28 \times 28$, allowing the number of tokens per frame to be adaptively adjusted under the video context budget. The maximum response length is capped at 512 tokens. The statistics of training data are shown in Appendix A and Table 7. All experiments are performed on NVIDIA A100 GPUs (80GB), with a global batch size of 64. Further implementation details are provided in the Appendix B.

5.2 MAIN RESULTS

Grounded Video Question Answering. We evaluate grounded video question answering on NExT-GQA (Xiao et al., 2024) and ReXTime (Chen et al., 2024a), reporting both answer accuracy and temporal grounding quality as shown in Table 1. Our model achieves state-of-the-art performance across all metrics on both benchmarks, surpassing strong RL-based baselines such as VideoChat-R1 (Li et al., 2025b). Notably, on NExT-GQA (Xiao et al., 2024), we improve mIoU by 5.2%, R@0.3 by 5.4%, and R@0.5 by 6.1% over the runner-up. On ReXTime (Chen et al., 2024a), our model also consistently yields an average improvement of 4.6% across all metrics.

We also report GQA performance on the challenging CG-Bench (Chen et al., 2025a) in Table 2, which contains very long videos where the answer-supporting clue typically occupies $\leq 1\%$ of the total duration. Beyond IoU, we evaluate how well the predicted segment covers the ground-truth

378
 379 Table 2: Performance on long video understanding (VideoMME (Fu et al., 2025), MLVU (Zhou et al.,
 380 2025), LVBench (Wang et al., 2024)) and long GVQA (CG-Bench (Chen et al., 2025a)) tasks. All
 381 open-sourced models are of comparable scale (7B or 8B).

382 Models	383 VideoMME (w/o & w/ sub.)		384 MLVU	385 LVBench	386 CG-Bench		
	387 Overall	388 Long	389 M-Avg	390 Avg	391 mIoU	392 rec. @ IoU	393 acc. @ IoU
384 Duration	385 1010s	386 2386s	387 651s	388 4101s	389 1624s		390
<i>Proprietary LVLMs</i>							
Gemini 1.5 Pro (Google, 2024)	75.0 / 81.3	67.4 / 77.4	-	33.1	3.85	5.61	2.64
GPT-4o (OpenAI, 2024)	65.3 / 77.2	65.3 / 72.1	64.6	30.8	5.73	8.12	4.33
<i>Open-Source LVLMs</i>							
LLaVA-OneVision (Li et al., 2024a)	58.2 / 61.5	- / -	64.7	-	1.56	1.19	1.72
LongVA (Zhang et al., 2024a)	52.6 / -	46.2 / -	56.3	-	2.91	3.15	1.32
InternVL2.5 (Chen et al., 2024b)	64.2 / 66.9	- / -	68.9	38.4	-	-	-
Qwen2.5-VL (Bai et al., 2025)	65.2 / 70.7	51.1 / 62.0	70.2	45.3	2.48	3.15	1.36
TVG-R1 (Chen et al., 2025b)	64.3 / 69.1	52.7 / 62.4	69.7	42.3	2.43	3.62	1.29
VideoChat-R1 (Li et al., 2025b)	64.3 / 69.1	53.4 / 62.3	69.5	43.7	5.91	8.38	2.56
Zoom-Zero (Ours)	66.0 / 71.2	54.8 / 64.2	70.8	45.7	6.68	9.30	3.62

394 clue span using Intersection-over-Ground truth (IoG; see Appendix D.1). As shown in Table 3, our
 395 model achieves a 7.7% gain of mIoG over the runner-up, validating that the zoom-in accuracy reward
 396 R_{Zoom} encourages predictions that not only localize the relevant temporal segments but also better
 397 cover the most salient frames containing key visual cues.

402
 403 Table 3: Temporal grounding coverage ratio. For ReXTIME (Chen et al., 2024a), results (IoU and
 404 accuracy) are only obtainable via server submission without access to ground-truth spans; therefore,
 405 to report mIoG (mean Intersection-over-Ground truth), we use the validation set for comparison.

406 Models	407 NExT-GQA			408 ReXTIME val			409 CG-Bench		
	410 mIoU	411 mIoG	412 mIoP	413 mIoU	414 mIoG	415 mIoP	416 mIoU	417 mIoG	418 mIoP
Qwen2.5-VL (Bai et al., 2025)	20.2	56.8	29.5	31.6	54.3	43.2	2.48	10.35	4.16
VideoChat-R1 (Li et al., 2025b)	32.4	93.5	39.1	43.5	64.3	52.8	5.91	18.44	7.34
Zoom-Zero (Ours)	37.6	94.7	43.2	44.7	67.6	53.5	6.68	26.15	8.25

420 **Long Video Understanding.** We compare against general-purpose LVLMs with temporal grounding
 421 capability and RL-based models explicitly optimized for grounding (TVG-R1 (Chen et al., 2025b),
 422 VideoChat-R1 (Li et al., 2025b)). While RL approaches that prioritize grounding can trade off
 423 general GQA accuracy, our method proposes token-selective credit assignment that decouples reward
 424 signals from answer accuracy and temporal grounding, assigning credit to the appropriate tokens.
 425 This mitigates the accuracy–grounding trade-off and yields stronger temporal localization without
 426 degrading question answering, as shown in Table 2.

427 **Qualitative Results.** We provide qualitative results in Figures 5 to 8 in the Appendix to demonstrate
 428 the model’s performance on the GVQA task. For example, as shown in Figure 5, the model can
 429 localize each event mentioned in the question and arrange them in the correct chronological order.

430 Table 4: Long video understanding via temporal zoom-in evaluated with MCQ accuracy.

431 Models	432 VideoMME (long w/ sub.)	433 MLVU	434 LVBench	435 CG-Bench
Qwen2.5-VL (Bai et al., 2025)	62.0	70.2	45.3	29.2
Zoom-Zero (Ours)	64.2	70.8	45.7	36.1
Zoom-Zero + Coarse-to-fine (Ours)	66.2	71.4	46.3	39.0
Zoom-Zero + Divide-and-conquer (Ours)	68.7	73.4	48.1	42.2

432 5.3 LONG VIDEO UNDERSTANDING VIA TEMPORAL ZOOM-IN
433434 The above experiments demonstrate our model’s ability to answer questions while faithfully localizing
435 relevant video segments. Although our primary goal is to enhance GVQA, we further present two
436 strategies that further benefit long-video understanding through temporal zoom-in.437 **Coarse-to-Fine.** In long-video scenarios, we first let the model trade spatial resolution for broad
438 temporal coverage to obtain a global overview. Once it has a coarse understanding and localizes the
439 query-relevant interval, we enable a fine-grained pass at higher spatial resolution for frames of interest
440 as mentioned in Section 4.1. As Table 4 (Coarse-to-fine) shows, it consistently improves performance
441 by providing targeted visual verification on a small set of salient frames with higher spatial resolution,
442 thus enhancing fine-grained understanding. We provide spatial and temporal resolution before and
443 after zoom-in in Table 13 and qualitative results in Figures 9 to 11 in the Appendix.444 **Divide-and-Conquer.** Another strategy is to partition a long video into non-overlapping windows
445 and perform a temporal search over them. For each window, the model predicts a query-relevant
446 temporal span and an answer. We then aggregate frames from spans with high-confidence answers
447 and apply a fine-grained zoom-in. Answer confidence is computed as the probability of the predicted
448 answer token, where $c = p_{\pi_\theta}(t)$ for token t strictly between `<answer>` and `</answer>`. We
449 select the top spans based on answer confidence and aggregate those frames as input to obtain the
450 final answer. As shown in Table 4 (Divide-and-conquer), it yields an average +6.4% improvement
451 over the baseline Qwen2.5-VL (Bai et al., 2025). Please refer to Appendix D.3 for ablation studies
452 on the window size and the number of aggregated predicted temporal spans.453
454 Table 5: Ablation studies.
455
456 457 458 459 460 461 462 463 464 465 466 467 468 469 470 471 472 473 474 475 476 477 478 479 480 481 482 483 484 485 486 487 488 489 490 491 492 493 494 495 496 497 498 499 500 501 502 503 504 505 506 507 508 509 510 511 512 513 514 515 516 517 518 519 520 521 522 523 524 525 526 527 528 529 530 531 532 533 534 535 536 537 538 539 540 541 542 543 544 545 546 547 548 549 550 551 552 553 554 555 556 557 558 559 560 561 562 563 564 565 566 567 568 569 570 571 572 573 574 575 576 577 578 579 580 581 582 583 584 585 586 587 588 589 590 591 592 593 594 595 596 597 598 599 600 601 602 603 604 605 606 607 608 609 610 611 612 613 614 615 616 617 618 619 620 621 622 623 624 625 626 627 628 629 630 631 632 633 634 635 636 637 638 639 640 641 642 643 644 645 646 647 648 649 650 651 652 653 654 655 656 657 658 659 660 661 662 663 664 665 666 667 668 669 670 671 672 673 674 675 676 677 678 679 680 681 682 683 684 685 686 687 688 689 690 691 692 693 694 695 696 697 698 699 700 701 702 703 704 705 706 707 708 709 710 711 712 713 714 715 716 717 718 719 720 721 722 723 724 725 726 727 728 729 730 731 732 733 734 735 736 737 738 739 740 741 742 743 744 745 746 747 748 749 750 751 752 753 754 755 756 757 758 759 759 760 761 762 763 764 765 766 767 768 769 770 771 772 773 774 775 776 777 778 779 779 780 781 782 783 784 785 786 787 788 789 789 790 791 792 793 794 795 796 797 798 799 800 801 802 803 804 805 806 807 808 809 809 810 811 812 813 814 815 816 817 818 819 819 820 821 822 823 824 825 826 827 828 829 829 830 831 832 833 834 835 836 837 838 839 839 840 841 842 843 844 845 846 847 848 849 850 851 852 853 854 855 856 857 858 859 859 860 861 862 863 864 865 866 867 868 869 869 870 871 872 873 874 875 876 877 878 879 879 880 881 882 883 884 885 886 887 888 889 889 890 891 892 893 894 895 896 897 898 899 900 901 902 903 904 905 906 907 908 909 909 910 911 912 913 914 915 916 917 918 919 919 920 921 922 923 924 925 926 927 928 929 929 930 931 932 933 934 935 936 937 938 939 939 940 941 942 943 944 945 946 947 948 949 950 951 952 953 954 955 956 957 958 959 959 960 961 962 963 964 965 966 967 968 969 969 970 971 972 973 974 975 976 977 978 979 979 980 981 982 983 984 985 986 987 988 989 989 990 991 992 993 994 995 996 997 998 999 1000 1001 1002 1003 1004 1005 1006 1007 1008 1009 1009 1010 1011 1012 1013 1014 1015 1016 1017 1018 1019 1019 1020 1021 1022 1023 1024 1025 1026 1027 1028 1029 1029 1030 1031 1032 1033 1034 1035 1036 1037 1038 1039 1039 1040 1041 1042 1043 1044 1045 1046 1047 1048 1049 1049 1050 1051 1052 1053 1054 1055 1056 1057 1058 1059 1059 1060 1061 1062 1063 1064 1065 1066 1067 1068 1069 1069 1070 1071 1072 1073 1074 1075 1076 1077 1078 1079 1079 1080 1081 1082 1083 1084 1085 1086 1087 1088 1089 1089 1090 1091 1092 1093 1094 1095 1096 1097 1098 1099 1099 1100 1101 1102 1103 1104 1105 1106 1107 1108 1109 1109 1110 1111 1112 1113 1114 1115 1116 1117 1118 1119 1119 1120 1121 1122 1123 1124 1125 1126 1127 1128 1129 1129 1130 1131 1132 1133 1134 1135 1136 1137 1138 1139 1139 1140 1141 1142 1143 1144 1145 1146 1147 1148 1149 1149 1150 1151 1152 1153 1154 1155 1156 1157 1158 1159 1159 1160 1161 1162 1163 1164 1165 1166 1167 1168 1169 1169 1170 1171 1172 1173 1174 1175 1176 1177 1178 1179 1179 1180 1181 1182 1183 1184 1185 1186 1187 1188 1189 1189 1190 1191 1192 1193 1194 1195 1196 1197 1198 1199 1199 1200 1201 1202 1203 1204 1205 1206 1207 1208 1209 1209 1210 1211 1212 1213 1214 1215 1216 1217 1218 1219 1219 1220 1221 1222 1223 1224 1225 1226 1227 1228 1229 1229 1230 1231 1232 1233 1234 1235 1236 1237 1238 1239 1239 1240 1241 1242 1243 1244 1245 1246 1247 1248 1249 1249 1250 1251 1252 1253 1254 1255 1256 1257 1258 1259 1259 1260 1261 1262 1263 1264 1265 1266 1267 1268 1269 1269 1270 1271 1272 1273 1274 1275 1276 1277 1278 1279 1279 1280 1281 1282 1283 1284 1285 1286 1287 1288 1289 1289 1290 1291 1292 1293 1294 1295 1296 1297 1298 1299 1299 1300 1301 1302 1303 1304 1305 1306 1307 1308 1309 1309 1310 1311 1312 1313 1314 1315 1316 1317 1318 1319 1319 1320 1321 1322 1323 1324 1325 1326 1327 1328 1329 1329 1330 1331 1332 1333 1334 1335 1336 1337 1338 1339 1339 1340 1341 1342 1343 1344 1345 1346 1347 1348 1349 1349 1350 1351 1352 1353 1354 1355 1356 1357 1358 1359 1359 1360 1361 1362 1363 1364 1365 1366 1367 1368 1369 1369 1370 1371 1372 1373 1374 1375 1376 1377 1378 1379 1379 1380 1381 1382 1383 1384 1385 1386 1387 1388 1389 1389 1390 1391 1392 1393 1394 1395 1396 1397 1398 1399 1399 1400 1401 1402 1403 1404 1405 1406 1407 1408 1409 1409 1410 1411 1412 1413 1414 1415 1416 1417 1418 1419 1419 1420 1421 1422 1423 1424 1425 1426 1427 1428 1429 1429 1430 1431 1432 1433 1434 1435 1436 1437 1438 1439 1440 1441 1442 1443 1444 1445 1446 1447 1448 1449 1450 1451 1452 1453 1454 1455 1456 1457 1458 1459 1460 1461 1462 1463 1464 1465 1466 1467 1468 1469 1470 1471 1472 1473 1474 1475 1476 1477 1478 1479 1480 1481 1482 1483 1484 1485 1486 1487 1488 1489 1489 1490 1491 1492 1493 1494 1495 1496 1497 1498 1499 1499 1500 1501 1502 1503 1504 1505 1506 1507 1508 1509 1509 1510 1511 1512 1513 1514 1515 1516 1517 1518 1519 1519 1520 1521 1522 1523 1524 1525 1526 1527 1528 1529 1529 1530 1531 1532 1533 1534 1535 1536 1537 1538 1539 1539 1540 1541 1542 1543 1544 1545 1546 1547 1548 1549 1549 1550 1551 1552 1553 1554 1555 1556 1557 1558 1559 1559 1560 1561 1562 1563 1564 1565 1566 1567 1568 1569 1569 1570 1571 1572 1573 1574 1575 1576 1577 1578 1579 1579 1580 1581 1582 1583 1584 1585 1586 1587 1588 1589 1589 1590 1591 1592 1593 1594 1595 1596 1597 1598 1599 1599 1600 1601 1602 1603 1604 1605 1606 1607 1608 1609 1609 1610 1611 1612 1613 1614 1615 1616 1617 1618 1619 1619 1620 1621 1622 1623 1624 1625 1626 1627 1628 1629 1629 1630 1631 1632 1633 1634 1635 1636 1637 1638 1639 1639 1640 1641 1642 1643 1644 1645 1646 1647 1648 1649 1649 1650 1651 1652 1653 1654 1655 1656 1657 1658 1659 1659 1660 1661 1662 1663 1664 1665 1666 1667 1668 1669 1669 1670 1671 1672 1673 1674 1675 1676 1677 1678 1679 1679 1680 1681 1682 1683 1684 1685 1686 1687 1688 1689 1689 1690 1691 1692 1693 1694 1695 1696 1697 1698 1699 1699 1700 1701 1702 1703 1704 1705 1706 1707 1708 1709 1709 1710 1711 1712 1713 1714 1715 1716 1717 1718 1719 1719 1720 1721 1722 1723 1724 1725 1726 1727 1728 1729 1729 1730 1731 1732 1733 1734 1735 1736 1737 1738 1739 1739 1740 1741 1742 1743 1744 1745 1746 1747 1748 1749 1749 1750 1751 1752 1753 1754 1755 1756 1757 1758 1759 1759 1760 1761 1762 1763 1764 1765 1766 1767 1768 1769 1769 1770 1771 1772 1773 1774 1775 1776 1777 1778 1779 1779 1780 1781 1782 1783 1784 1785 1786 1787 1788 1789 1789 1790 1791 1792 1793 1794 1795 1796 1797 1798 1799 1799 1800 1801 1802 1803 1804 1805 1806 1807 1808 1809 1809 1810 1811 1812 1813 1814 1815 1816 1817 1818 1819 1819 1820 1821 1822 1823 1824 1825 1826 1827 1828 1829 1829 1830 1831 1832 1833 1834 1835 1836 1837 1838 1839 1839 1840 1841 1842 1843 1844 1845 1846 1847 1848 1849 1849 1850 1851 1852 1853 1854 1855 1856 1857 1858 1859 1859 1860 1861 1862 1863 1864 1865 1866 1867 1868 1869 1869 1870 1871 1872 1873 1874 1875 1876 1877 1878 1879 1879 1880 1881 1882 1883 1884 1885 1886 1887 1888 1889 1889 1890 1891 1892 1893 1894 1895 1896 1897 1898 1899 1899 1900 1901 1902 1903 1904 1905 1906 1907 1908 1909 1909 1910 1911 1912 1913 1914 1915 1916 1917 1918 1919 1919 1920 1921 1922 1923 1924 1925 1926 1927 1928 1929 1929 1930 1931 1932 1933 1934 1935 1936 1937 1938 1939 1939 1940 1941 1942 1943 1944 1945 1946 1947 1948 1949 1949 1950 1951 1952 1953 1954 1955 1956 1957 1958 1959 1959 1960 1961 1962 1963 1964 1965 1966 1967 1968 1969 1969 1970 1971 1972 1973 1974 1975 1976 1977 1978 1979 1979 1980 1981 1982 1983 1984 1985 1986 1987 1988 1989 1989 1990 1991 1992 1993 1994 1995 1996 1997 1998 1999 1999 2000 2001 2002 2003 2004 2005 2006 2007 2008 2009 2009 2010 2011 2012 2013 2014 2015 2016 2017 2018 2019 2019 2020 2021 2022 2023 2024 2025 2026 2027 2028 2029 2029 2030 2031 2032 2033 2034 2035 2036 2037 2038 2039 2039 2040 2041 2042 2043 2044 2045 2046 2047 2048 2049 2049 2050 2051 2052 2053 2054 2055 2056 2057 2058 2059 2059 2060 2061 2062 2063 2064 2065 2066 2067 2068 2069 2069 2070 2071 2072 2073 2074 2075 2076 2077 2078 2079 2079 2080 2081 2082 2083 2084 2085 2086 2087 2088 2089 2089 2090 2091 2092 2093 2094 2095 2096 2097 2098 2099 2099 2100 2101 2102 2103 2104 2105 2106 2107 2108 2109 2109 2110 2111 2112 2113 2114 2115 2116 2117 2118 2119 2119 2120 2121 2122 2123 2124 2125 2126 2127 2128 2129 2129 2130 2131 2132 2133 2134 2135 2136 2137 2138 2139 2139 2140 2141 2142 2143 2144 2145 2146 2147 2148 2149 2149 2150 2151 2152 2153 2154 2155 2156 2157 2158 2159 2159 2160 2161 2162 2163 2164 2165 2166 2167 2168 2169 2169 2170 2171 2172 2173 2174 2175 2176 2177 2178 2179 2179 2180 2181 2182 2183 2184 2185 2186 2187<br

486
 487 **Table 6: Trade-off between inference speed and accuracy gain.** *Time* denotes the average inference
 488 time per video.

489 490 491 492 493 494 495 496 497 498 499 500 501 502 503 504 505 506 507 508 509 510 511 512 513 514 515 516 517 518 519 520 521 522 523 524 525 526 527 528 529 530 531 532 533 534 535 536 537 538 539	Models	MLVU		LVBench		VideoMME long (w/ sub.)		
		Acc	Time	Acc	Time	Acc	Time	
Duration		651s		4101s		2386s		
Qwen2.5-VL (Bai et al., 2025)		70.2	18.5s	45.3	39.7s	62.0	25.6s	
Zoom-Zero (Ours)		70.8	18.7s	45.7	40.6s	64.2	25.8s	
Zoom-Zero + Coarse-to-fine (Ours)		71.4	31.1s	46.3	55.5s	66.2	35.2s	
Zoom-Zero + Divide-and-conquer (Ours)		73.4	33.5s	48.1	110.5s	68.7	59.7s	

500 **Inference speed analysis.** We provide a clearer breakdown of the effectiveness–latency trade-off in
 501 the table below, and report three inference scenarios in Table 6

502 **(i) One-stage inference:** Zoom-Zero (second row) uses the same one-stage inference pipeline as
 503 the baseline, resulting in nearly identical inference time (it might vary a little due to the number of
 504 generated tokens). Trained with our proposed method, this setting yields an average improvement
 505 of +1.0 over the baseline without introducing additional latency. Please kindly note that the main
 506 experimental results as shown in Table 1 and Table 2 only have one-stage inference. **(ii) Two-
 507 stage inference (Coarse-to-fine):** The coarse-to-fine variant adds a fine-grained pass on grounded
 508 frames. This introduces a moderate increase in computation, approximately 1.4 \times inference time,
 509 while delivering a higher average absolute improvement of +2.1 over the baseline. **(iii) Two-stage
 510 inference (Divide-and-conquer):** The divide-and-conquer scheme is an optional test-time scaling
 511 strategy designed to further push performance. While it increases inference time to around 2.3 \times , it
 512 also achieves the largest gain, improving the baseline by +4.3 on average.

513 6 CONCLUSION

514 We introduce Zoom-Zero, a coarse-to-fine framework for grounded video question answering that
 515 first localizes query-relevant segments, then zooms into salient frames to capture fine-grained details.
 516 Our approach enhances GRPO for GVQA with two key contributions: (i) a zoom-in accuracy reward
 517 for evidence-faithful temporal grounding and fine-grained visual verification, and (ii) token-selective
 518 credit assignment for advantage estimation, assigning credit to the tokens responsible for localization
 519 or answer generation, respectively, addressing GRPO’s limits under multi-faceted reward signals.
 520 Our method improves both temporal grounding and answer accuracy, raising temporal grounding by
 521 5.2% on NExT-GQA and 4.6% on ReXTIME. Its coarse-to-fine paradigm boosts long-form video
 522 understanding by an average of 6.4%, preserving critical detail without sacrificing global context.

523 ETHICS STATEMENT

524 Our work builds on large video-language models (LVLMs) and reinforcement learning for grounded
 525 video question answering. We do not collect or annotate any human subject data; all experiments
 526 use publicly available datasets under research licenses. We adhere to the terms of use specified by
 527 the original dataset creators and provide appropriate citations. Our approach does not introduce
 528 additional risks of data misuse or privacy leakage.

529 530 REPRODUCIBILITY STATEMENT

531 We make every effort to ensure reproducibility of our results. Full implementation details are
 532 provided in Appendix B. All datasets used in our experiments are publicly accessible and described
 533 in Appendix A. We provide the evaluation protocols and metrics in Section 5.1, and present ablation
 534 studies to analyze the effect of key components in Section 5.4 and Appendix D.

540 REFERENCES
541

542 Lisa Anne Hendricks, Oliver Wang, Eli Shechtman, Josef Sivic, Trevor Darrell, and Bryan Russell.
543 Localizing moments in video with natural language. In *CVPR*, 2017.

544 Kirolos Ataallah, Xiaoqian Shen, Eslam Abdelrahman, Essam Sleiman, Deyao Zhu, Jian Ding, and
545 Mohamed Elhoseiny. Minigpt4-video: Advancing multimodal llms for video understanding with
546 interleaved visual-textual tokens. *arXiv preprint arXiv:2404.03413*, 2024.

547

548 Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang,
549 Shijie Wang, Jun Tang, et al. Qwen2.5-vl technical report. *arXiv preprint arXiv:2502.13923*, 2025.

550

551 Guo Chen, Yicheng Liu, Yifei Huang, Yuping He, Baoqi Pei, Jilan Xu, Yali Wang, Tong Lu, and Limin
552 Wang. Cg-bench: Clue-grounded question answering benchmark for long video understanding. In
553 *ICLR*, 2025a.

554

555 Jr-Jen Chen, Yu-Chien Liao, Hsi-Che Lin, Yu-Chu Yu, Yen-Chun Chen, and Frank Wang. Rextime:
A benchmark suite for reasoning-across-time in videos. *NeurIPS*, 2024a.

556

557 Jun Chen, Deyao Zhu, Xiaoqian Shen, Xiang Li, Zechun Liu, Pengchuan Zhang, Raghuraman
558 Krishnamoorthi, Vikas Chandra, Yunyang Xiong, and Mohamed Elhoseiny. Minigpt-v2: large
559 language model as a unified interface for vision-language multi-task learning. *arXiv preprint
arXiv:2310.09478*, 2023.

560

561 Ruizhe Chen, Zhiting Fan, Tianze Luo, Heqing Zou, Zhaopeng Feng, Guiyang Xie, Hansheng Zhang,
562 Zhuochen Wang, Zuozhu Liu, and Huaijian Zhang. Datasets and recipes for video temporal
563 grounding via reinforcement learning. *arXiv preprint arXiv:2507.18100*, 2025b.

564

565 Zhe Chen, Weiyun Wang, Yue Cao, Yangzhou Liu, Zhangwei Gao, Erfei Cui, Jinguo Zhu, Shenglong
566 Ye, Hao Tian, Zhaoyang Liu, et al. Expanding performance boundaries of open-source multimodal
567 models with model, data, and test-time scaling. *arXiv preprint arXiv:2412.05271*, 2024b.

568

569 Zesen Cheng, Sicong Leng, Hang Zhang, Yifei Xin, Xin Li, Guanzheng Chen, Yongxin Zhu, Wenqi
570 Zhang, Ziyang Luo, Deli Zhao, et al. Videollama 2: Advancing spatial-temporal modeling and
571 audio understanding in video-llms. *arXiv preprint arXiv:2406.07476*, 2024.

572

573 Jang Hyun Cho, Andrea Madotto, Effrosyni Mavroudi, Triantafyllos Afouras, Tushar Nagarajan,
574 Muhammad Maaz, Yale Song, Tengyu Ma, Shuming Hu, Suyog Jain, et al. Perceptionlm: Open-
575 access data and models for detailed visual understanding. *arXiv preprint arXiv:2504.13180*,
2025.

576

577 Kaituo Feng, Kaixiong Gong, Bohao Li, Zonghao Guo, Yibing Wang, Tianshuo Peng, Junfei Wu,
578 Xiaoying Zhang, Benyou Wang, and Xiangyu Yue. Video-r1: Reinforcing video reasoning in
579 mllms. *arXiv preprint arXiv:2503.21776*, 2025.

580

581 Chaoyou Fu, Yuhan Dai, Yondong Luo, Lei Li, Shuhuai Ren, Renrui Zhang, Zihan Wang, Chenyu
582 Zhou, Yunhang Shen, and et al. Zhang, Mengdan. Video-mme: The first-ever comprehensive
583 evaluation benchmark of multi-modal llms in video analysis. In *CVPR*, 2025.

584

585 Jiyang Gao, Chen Sun, Zhenheng Yang, and Ram Nevatia. Tall: Temporal activity localization via
586 language query. In *CVPR*, 2017.

587

588 Gemini Team Google. Gemini 1.5: Unlocking multimodal understanding across millions of tokens
589 of context. *arXiv preprint arXiv:2403.05530*, 2024.

590

591 Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu,
592 Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms
593 via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.

594

595 Bo He, Hengduo Li, Young Kyun Jang, Menglin Jia, Xuefei Cao, Ashish Shah, Abhinav Shrivastava,
596 and Ser-Nam Lim. Ma-lmm: Memory-augmented large multimodal model for long-term video
597 understanding. In *CVPR*, 2024.

594 Kai Hu, Feng Gao, Xiaohan Nie, Peng Zhou, Son Tran, Tal Neiman, Lingyun Wang, Mubarak Shah,
 595 Raffay Hamid, Bing Yin, et al. M-llm based video frame selection for efficient video understanding.
 596 In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 13702–13712,
 597 2025.

598 Bin Huang, Xin Wang, Hong Chen, Zihan Song, and Wenwu Zhu. Vtimellm: Empower llm to grasp
 599 video moments. In *CVPR*, 2024a.

601 De-An Huang, Shijia Liao, Subhashree Radhakrishnan, Hongxu Yin, Pavlo Molchanov, Zhiding Yu,
 602 and Jan Kautz. Lita: Language instructed temporal-localization assistant. In *ECCV*, 2024b.

603 Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec
 604 Helyar, Aleksander Madry, Alex Beutel, Alex Carney, et al. Openai o1 system card. *arXiv preprint*
 605 *arXiv:2412.16720*, 2024.

606 Ranjay Krishna, Kenji Hata, Frederic Ren, Li Fei-Fei, and Juan Carlos Niebles. Dense-captioning
 607 events in videos. In *ICCV*, 2017.

608 Song Lai, Haohan Zhao, Rong Feng, Changyi Ma, Wenzhuo Liu, Hongbo Zhao, Xi Lin, Dong Yi,
 609 Min Xie, Qingfu Zhang, Hongbin Liu, Gaofeng Meng, and Fei Zhu. Reinforcement fine-tuning
 610 naturally mitigates forgetting in continual post-training. *arXiv preprint arXiv:2507.05386*, 2025.

611 Jie Lei, Tamara L Berg, and Mohit Bansal. Detecting moments and highlights in videos via natural
 612 language queries. *NeurIPS*, 2021.

613 Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Yanwei
 614 Li, Ziwei Liu, and Chunyuan Li. Llava-onevision: Easy visual task transfer. *arXiv preprint*
 615 *arXiv:2408.03326*, 2024a.

616 Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image
 617 pre-training with frozen image encoders and large language models. In *International conference*
 618 *on machine learning*, pp. 19730–19742. PMLR, 2023a.

619 KunChang Li, Yinan He, Yi Wang, Yizhuo Li, Wenhui Wang, Ping Luo, Yali Wang, Limin Wang, and
 620 Yu Qiao. Videochat: Chat-centric video understanding. *arXiv preprint arXiv:2305.06355*, 2023b.

621 Kunchang Li, Yali Wang, Yinan He, Yizhuo Li, Yi Wang, Yi Liu, Zun Wang, Jilan Xu, Guo Chen,
 622 Ping Luo, et al. Mvbench: A comprehensive multi-modal video understanding benchmark. In
 623 *CVPR*, 2024b.

624 Rui Li, Xiaohan Wang, Yuhui Zhang, Zeyu Wang, and Serena Yeung-Levy. Temporal preference
 625 optimization for long-form video understanding. *arXiv preprint arXiv:2501.13919*, 2025a.

626 Xinhao Li, Ziang Yan, Desen Meng, Lu Dong, Xiangyu Zeng, Yinan He, Yali Wang, Yu Qiao,
 627 Yi Wang, and Limin Wang. Videochat-r1: Enhancing spatio-temporal perception via reinforcement
 628 fine-tuning. *arXiv preprint arXiv:2504.06958*, 2025b.

629 Yanwei Li, Chengyao Wang, and Jiaya Jia. Llama-vid: An image is worth 2 tokens in large language
 630 models. In *ECCV*, 2024c.

631 Zeqian Li, Shangzhe Di, Zhonghua Zhai, Weilin Huang, Yanfeng Wang, and Weidi Xie. Universal
 632 video temporal grounding with generative multi-modal large language models. *arXiv preprint*
 633 *arXiv:2506.18883*, 2025c.

634 Bin Lin, Bin Zhu, Yang Ye, Munan Ning, Peng Jin, and Li Yuan. Video-llava: Learning united visual
 635 representation by alignment before projection. In *EMNLP*, 2024.

636 Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee.
 637 Llava-next: Improved reasoning, ocr, and world knowledge, January 2024a. URL <https://llava-vl.github.io/blog/2024-01-30-llava-next/>.

638 Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *NeurIPS*,
 639 2024b.

648 Ruipu Luo, Ziwing Zhao, Min Yang, Junwei Dong, Minghui Qiu, Pengcheng Lu, Tao Wang, and
 649 Zhongyu Wei. Valley: Video assistant with large language model enhanced ability. *arXiv preprint*
 650 *arXiv:2306.07207*, 2023.

651 Muhammad Maaz, Hanoona Rasheed, Salman Khan, and Fahad Shahbaz Khan. Video-chatgpt:
 652 Towards detailed video understanding via large vision and language models. *arXiv preprint*
 653 *arXiv:2306.05424*, 2023.

654 Ming Nie, Dan Ding, Chunwei Wang, Yuanfan Guo, Jianhua Han, Hang Xu, and Li Zhang. Slowfocus:
 655 Enhancing fine-grained temporal understanding in video LLM. In *NeurIPS*, 2024.

656 OpenAI. Gpt-4o system card, 2024. URL <https://openai.com/index/hello-gpt-4o/>.

657 Shuhuai Ren, Linli Yao, Shicheng Li, Xu Sun, and Lu Hou. Timechat: A time-sensitive multimodal
 658 large language model for long video understanding. In *CVPR*, 2024.

659 John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy
 660 optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

661 Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang,
 662 Mingchuan Zhang, YK Li, Y Wu, et al. Deepseekmath: Pushing the limits of mathematical
 663 reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.

664 Xiaoqian Shen, Yunyang Xiong, Changsheng Zhao, Lemeng Wu, Jun Chen, Chenchen Zhu, Zechun
 665 Liu, Fanyi Xiao, Balakrishnan Varadarajan, Florian Bordes, et al. Longvu: Spatiotemporal adaptive
 666 compression for long video-language understanding. In *ICML*, 2025.

667 Enxin Song, Wenhao Chai, Guanhong Wang, Yucheng Zhang, Haoyang Zhou, Feiyang Wu, Xun
 668 Guo, Tian Ye, Yan Lu, Jenq-Neng Hwang, et al. Moviechat: From dense token to sparse memory
 669 for long video understanding. In *CVPR*, 2024.

670 Xi Tang, Jihao Qiu, Lingxi Xie, Yunjie Tian, Jianbin Jiao, and Qixiang Ye. Adaptive keyframe
 671 sampling for long video understanding. In *Proceedings of the Computer Vision and Pattern
 672 Recognition Conference*, pp. 29118–29128, 2025.

673 Shengbang Tong, Ellis Brown, Penghao Wu, Sanghyun Woo, Manoj Middepogu, Sai Charitha
 674 Akula, Jihan Yang, Shusheng Yang, Adithya Iyer, Xichen Pan, et al. Cambrian-1: A fully open,
 675 vision-centric exploration of multimodal llms. In *NeurIPS*, 2024.

676 Haibo Wang, Zhiyang Xu, Yu Cheng, Shizhe Diao, Yufan Zhou, Yixin Cao, Qifan Wang, Weifeng
 677 Ge, and Lifu Huang. Grounded-videollm: Sharpening fine-grained temporal grounding in video
 678 large language models. In *EMNLP*, 2025a.

679 Weihan Wang, Zehai He, Wenyi Hong, Yean Cheng, Xiaohan Zhang, Ji Qi, Xiaotao Gu, Shiyu Huang,
 680 Bin Xu, Yuxiao Dong, et al. Lvbench: An extreme long video understanding benchmark. *arXiv*
 681 *preprint arXiv:2406.08035*, 2024.

682 Ye Wang, Ziheng Wang, Boshen Xu, Yang Du, Kejun Lin, Zihan Xiao, Zihao Yue, Jianzhong Ju,
 683 Liang Zhang, Dingyi Yang, et al. Time-r1: Post-training large vision language model for temporal
 684 video grounding. *arXiv preprint arXiv:2503.13377*, 2025b.

685 Zuxuan Wu, Caiming Xiong, Chih-Yao Ma, Richard Socher, and Larry S Davis. Adaframe: Adaptive
 686 frame selection for fast video recognition. In *Proceedings of the IEEE/CVF Conference on
 687 Computer Vision and Pattern Recognition*, pp. 1278–1287, 2019.

688 Junbin Xiao, Angela Yao, Yicong Li, and Tat-Seng Chua. Can i trust your answer? visually grounded
 689 video question answering. In *CVPR*, 2024.

690 Peiyuan Zhang, Kaichen Zhang, Bo Li, Guangtao Zeng, Jingkang Yang, Yuanhan Zhang, Ziyue
 691 Wang, Haoran Tan, Chunyuan Li, and Ziwei Liu. Long context transfer from language to vision.
 692 *arXiv preprint arXiv:2406.16852*, 2024a.

693 Shaojie Zhang, Jiahui Yang, Jianqin Yin, Zhenbo Luo, and Jian Luan. Q-frame: Query-aware frame
 694 selection and multi-resolution adaptation for video-llms. *arXiv preprint arXiv:2506.22139*, 2025.

702 Yuanhan Zhang, Bo Li, haotian Liu, Yong jae Lee, Liangke Gui, Di Fu, Jiashi Feng, Ziwei Liu, and
703 Chunyuan Li. Llava-next: A strong zero-shot video understanding model, April 2024b. URL
704 <https://llava-vl.github.io/blog/2024-04-30-llava-next-video/>.
705

706 Junjie Zhou, Yan Shu, Bo Zhao, Boya Wu, Shitao Xiao, Xi Yang, Yongping Xiong, Bo Zhang,
707 Tiejun Huang, and Zheng Liu. Mlvu: A comprehensive benchmark for multi-task long video
708 understanding. In *CVPR*, 2025.

709 Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. Minigpt-4: Enhancing
710 vision-language understanding with advanced large language models. In *ICLR*, 2024.

711

712

713

714

715

716

717

718

719

720

721

722

723

724

725

726

727

728

729

730

731

732

733

734

735

736

737

738

739

740

741

742

743

744

745

746

747

748

749

750

751

752

753

754

755

APPENDIX

A TRAINING DATA

Table 7 summarizes the statistics of the training datasets. For the QVHighlights (Lei et al., 2021) training split, which contains 9,996 examples, we only keep videos longer than 120 seconds. For the PLM-Video (Cho et al., 2025) multiple-choice split, we perform a quality check to remove examples that cannot be correctly answered using the full video, but can be correctly answered when restricted to the cropped segment defined by the clue duration. This ensures that the clue duration indeed provides sufficient information for identifying the correct video segment.

After Stage I training, we employ the model trained from Stage I for offline data filtering. Specifically, we generate $n = 8$ responses per example and discard those without meaningful reward signals. For question-answering ability, we exclude examples for which all generated responses answer the question correctly, as they lack discriminative signals. For temporal grounding ability, we filter out examples with low response variance. In particular, we retain only examples where the responses yield a sufficiently strong relative reward signal, quantified by the difference between the maximum IoU and the mean IoU across responses:

$$\delta = \max_{1 \leq i \leq n} \text{IoU}_i - \frac{1}{n} \sum_{i=1}^n \text{IoU}_i \quad (6)$$

We filter out examples with $\delta < 0.1$.

Table 7: **Statistics of training data.** NExT-GQA and ActivityNet in seconds stage are sampled from the first stage by filtering reward variation based on the first-stage model.

	Dataset	#Queries	Video Len.	Moment Len.
Stage I	NExT-GQA (Xiao et al., 2024)	3,358	43.9s	8.5s
	ActivityNet (Krishna et al., 2017)	4,727	177.3s	48.35
	QVHighlights (Lei et al., 2021)	7,218	150.0s	34.1s
Stage II	ActivityNet (Krishna et al., 2017)	1,395	220.87	88.4s
	NExT-GQA (Xiao et al., 2024)	1,004	50.2s	7.1s
	PLM-Video (Cho et al., 2025)	5,333	808.6s	26.1s

B IMPLEMENTATION DETAILS

We adopt Qwen2.5-VL-7B (Bai et al., 2025) as the base model. The maximum number of video tokens is set to 8,192, with videos sampled at 1 fps during training. The minimum video frame resolution is $16 \times 28 \times 28$ pixels, and the maximum is $768 \times 28 \times 28$, allowing the number of tokens per frame to be adaptively adjusted under the video context budget. The maximum response length is capped at 512 tokens. Due to computational resource limitations, we conduct RL training in two stages. In the first stage, we train on 20K short-video GQA examples from NExT-GQA (Xiao et al., 2024), ActivityNet (Krishna et al., 2017), and QVHighlights (Lei et al., 2021). In the second stage, we train on the *yt1b_mcqa* split from PLM-Video (Cho et al., 2025), combined with the short-video data sampled from the first stage, for a total of 7K examples. The statistics of training data are shown in Appendix A and Table 7. All experiments are performed on NVIDIA A100 GPUs (80GB), with a global batch size of 64.

During inference, we evaluate all models at 1 FPS with a context size of 8,192 on the short-video benchmarks NExT-GQA (Xiao et al., 2024) and ReXTime (Chen et al., 2024a). For long-video benchmarks: CG-Bench (Chen et al., 2025a), VideoMME (Fu et al., 2025), MLVU (Zhou et al., 2025), and LVBench (Wang et al., 2024). We also uniformly sampled a maximum of 256 frames and set the context size to 16,384.

810

811 Table 8: A simple example to demonstrate the GRPO’s uniform credit assignment problem.

Response	R_{IoU}	R_{Acc}	A_{IoU}	A_{Acc}	R_{Sum}	A_{Sum}
1	0.0	1.0	-1.40	+0.82	1.0	+0.06
2	0.5	0.0	+0.44	-1.22	0.5	-1.54
3	0.4	1.0	+0.07	+0.82	1.4	+1.34
4	0.8	0.0	+1.55	-1.22	0.8	-0.58
5	0.2	1.0	-0.66	+0.82	1.2	+0.70

812

820 Table 9: Ablation on the number of generated responses G per prompt during GRPO training.

G	NExT-GQA				ReXTime			
	Acc	mIoU	R@0.3	R@0.5	Acc	mIoU	R@0.3	R@0.5
2	69.6	33.7	50.3	27.5	58.5	36.9	49.0	37.1
4	69.8	35.2	52.6	29.6	58.8	40.1	53.2	40.6
8	70.7	37.6	55.6	33.8	62.0	43.2	56.5	44.1

821

822 C LIMITATION OF GRPO IN UNIFORM CREDIT ASSIGNMENT

831 In Table 8, we present a simple example with two rewards, R_{IoU} and R_{Acc} , across five responses to
 832 illustrate the limitations of GRPO arising from naïve reward summation and uniform credit assignment.
 833 For example, the first response attains the lowest temporal grounding reward, $R_{\text{IoU}}^{(1)} = 0$, yet its
 834 overall advantage under standard GRPO is positive, $A_{\text{Sum}}^{(1)} = +0.06$. In contrast, response 4 achieves
 835 much better temporal grounding, $R_{\text{IoU}}^{(4)} = 0.8$, but receives a lower advantage, $A_{\text{Sum}}^{(4)} = -0.58$.
 836 Due to uniform credit assignment, all tokens in response 1 are reinforced by the positive advantage,
 837 while all tokens in response 4 are penalized. This hides the contribution of tokens that support more
 838 accurate temporal grounding.

839 In contrast, computing separate advantages for each reward, A_{IoU} and A_{Acc} , provides a clearer view
 840 of each task’s contribution. By selectively assigning these decoupled advantages to the corresponding
 841 tokens, our approach TokenAdv, updates the policy to increase the probability of tokens that positively
 842 impact their respective tasks.

843 D EXPERIMENTS

844 D.1 TEMPORAL GROUNDING COVERAGE

845 In addition to IoU, we evaluate how well the predicted segment covers the ground-truth clue span using
 846 Intersection-over-Ground truth (IoG), defined as $\text{IoG} = \frac{|\mathcal{I}_{\text{pred}} \cap \mathcal{I}_{\text{gt}}|}{|\mathcal{I}_{\text{gt}}|}$, where $\mathcal{I}_{\text{pred}}$ is the predicted
 847 temporal span and \mathcal{I}_{gt} is the ground-truth clue span. We report mean IoG (mIoG) as the average
 848 IoG across instances. IoG directly measures coverage of the ground truth and thus verifies whether
 849 temporal grounding captures the key frames relevant to the query, particularly informative for the
 850 finer-grained zoom-in.

851 We present results in Table 3. For ReXTime (Chen et al., 2024a), only IoU and accuracy are available
 852 via server-side evaluation, and the ground-truth clue spans are not released; consequently, we compute
 853 and report mIoG on the validation set for comparison. Our model improves mIoG by 1.2% on NExT-
 854 GQA (Xiao et al., 2024) and by 3.3% on ReXTime (Chen et al., 2024a). For very long videos such
 855 as CG-Bench (Chen et al., 2025a), mIoU can be less informative because the larger denominator
 856 depresses scores. Considering both mIoU and mIoG shows that our model not only localizes the
 857 relevant moments but also achieves strong coverage of key frames.

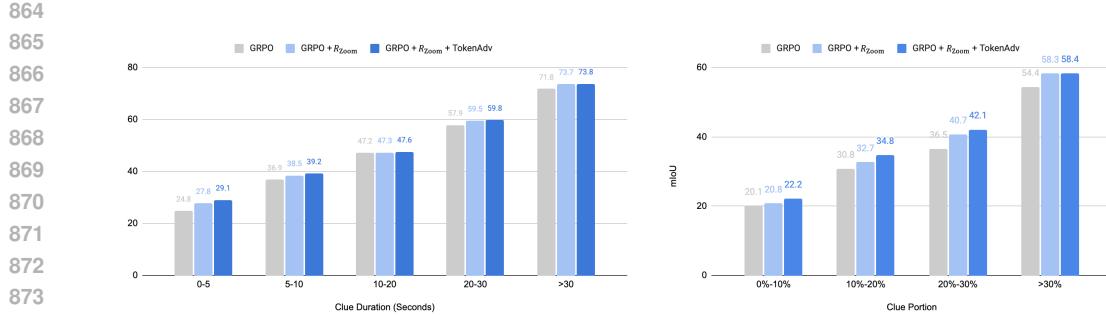


Figure 3: **Temporal grounding robustness analysis on NExT-GQA.** Left: mIoU results across different ground-truth clue durations. Right: mIoU results across different clue proportions (ground-truth clue duration relative to the total video duration).

D.2 THE NUMBER OF GENERATED RESPONSES

We investigate the impact of the number of generated responses G per prompt during GRPO training, as this hyperparameter directly influences the diversity and quality of the policy optimization process. As presented in Table 9, increasing G from 2 to 8 consistently improves performance across both datasets and all evaluation metrics. Based on these results, we adopt $G = 8$ for all main experiments, as it provides the optimal balance between computational efficiency and performance gains.

D.3 DIVIDE-AND-CONQUER

We study the impact of window size (Table 10) and the number of predicted temporal spans aggregated in the divide-and-conquer strategy. Because this approach requires scanning every sliding window during the coarse-grained pass, it introduces an average $\times 2.3$ increase in inference cost. Nevertheless, it improves performance across all long-video benchmarks by an average of +6.4%, demonstrating that our temporal zoom-in with higher spatial resolution provides substantial benefits for long video understanding. (Table 11) shows the impact of number of aggregated temporal spans with top answer confidence.

Table 10: Window size ablation.

Window Size	VideoMME (long w sub.)	MLVU	LVBench
128	67.4	72.1	47.6
256	68.7	73.4	48.1
384	68.4	72.4	48.5

Table 11: Number of aggregated temporal spans with top answer confidence.

# Aggregated Spans	VideoMME (long w sub.)	MLVU	LVBench
3	73.2	66.8	47.5
4	73.4	68.7	48.1
5	73.2	68.4	47.9

D.4 COARSE-TO-FINE VIDEO UNDERSTANDING ON GVQA

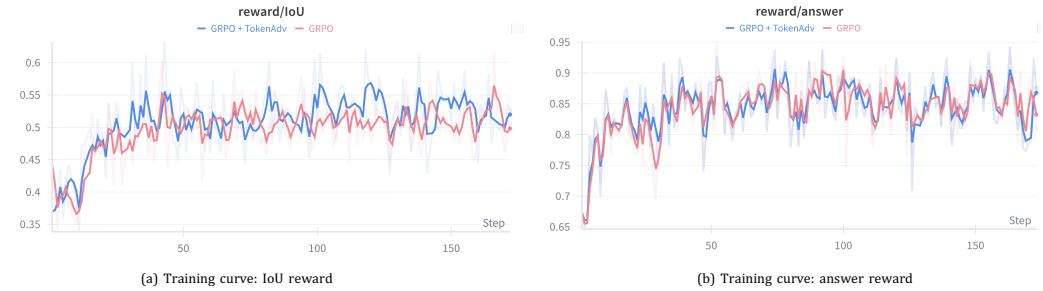
We further evaluate short-form GVQA answer accuracy on NExT-GQA and ReXTIME through temporal zoom-in, as reported in Table 12. Both benchmarks consist of short videos, where the model

918
919 Table 12: Grounded question answering (GQA) results on NExT-GQA (Xiao et al., 2024) and
920 ReXTime (Chen et al., 2024a) with temporal zoom-in.

Models	NExT-GQA				ReXTime			
	Acc	mIoU	R@0.3	R@0.5	Acc	mIoU	R@0.3	R@0.5
Zoom-Zero	70.7	37.6	55.6	33.8	62.0	43.2	56.5	44.1
Zoom-Zero + Coarse-to-fine	71.4	N/A	N/A	N/A	62.8	N/A	N/A	N/A

928 Table 13: Spatial and temporal resolution in coarse-to-fine video understanding.

Benchmark	VideoMME Long		MLVU		LVBench	
	Duration	2386s	651s	4101s	Coarse	Fine
Avg frames	256	136	253	86	256	154
Avg FPS	0.1	1.0	0.2	1.0	0.06	1.0
Avg tokens/frame	54	76	64	190	64	62



950 Figure 4: Training curve: IoU reward R_{IoU} and answer reward R_{Acc} comparison with baseline
951 GRPO and our improved GRPO with TokenAdv.

952
953
954 **Question:** Arrange the following events from the video in the correct chronological order: (1)The action of the
955 lures is shown underwater as several different fish go after the lures; (2)Several men show off the different lures
956 they are using for ice fishing; (3)The video ends with the closing credits and Graphics shown on the screen; (4)An
957 introduction comes onto the screen for a video about fishing lures.

958 **Options:** (A) 3->2->1->4 (B) 4->2->1->3 (C) 1->2->3->4 (D) 2->1->3->4



962 **Output:** <think>First, the video introduces the topic of fishing lures with an opening screen <time> (20.3, 49.7)</time>. Next, several men showcase the different lures they are using for ice fishing <time> (283.5, 304.2)</time>. Then, the action of the lures is demonstrated underwater as various fish chase after them <time> (593.1, 611.4)</time>. Finally, the video concludes with closing credits and graphics <time> (673.8, 680.2)</time>. </think>\n<answer>B</answer><glue>[(20.3, 49.7), (283.5, 304.2), (593.1, 611.4), (673.8, 680.2)]</glue>

967 Figure 5: A qualitative example for long video understanding.

968
969
970
971 can preserve most temporal context and details within the context budget. In this setting, the zoom-in paradigm improves performance by 0.7% on NExT-GQA and 0.8% on ReXTime.

972

Question: Are there any irregularities in this surveillance video? If there are, what sort are they?
Options: (A) Abuse (B) Arrest (C) Arson (D) Stealing



977

Output: <think>During the video, a person is seen entering a room with a bag <time>(30.8, 46.9) </time>. They then proceed to light a fire <time>(170.0, 230.0) </time> using a lighter, which spreads rapidly across the floor. This indicates arson activity. </think>\n<answer>C</answer><glue>[(170.0, 230.0)]</glue>

980

Figure 6: A qualitative example for long video understanding.

981

Question: Which of the following options correctly sequences the order in which the competition items appear in the video?

Options: A. 400m run, shot put, 110m hurdles. B. 100m run, shot put, long jump.
 C. 100m run, long jump, shot put. D. 400m run, 110m hurdles, shot put.



988

Output: <think><time>(78.00, 150.00)</time> The video starts with the 100m Decathlon event, where athletes compete in the 100m run. <time>(150.00, 220.00)</time> Next, the Long Jump Decathlon is shown, where athletes perform the long jump. <time>(380.00, 490.00)</time> Finally, the Shot Put Decathlon is displayed, where athletes throw the shot put. This sequence matches the order in the options provided. </think>\n<answer>C</answer><glue>[(78.00, 150.00), (150.00, 220.00), (380.00, 490.00)]</glue>



995

Figure 7: A qualitative example for long video understanding.

1001

Question: What color suit is the female anchor who also has a chest bracelet and gold bracelet?

Options: (A) Grey (B) Pink (C) Purple (D) Black



1006

Output: <think>The female anchor is wearing a pink suit with a chest bracelet and gold bracelet <time>(130.67, 140.07) </time>. This matches the description provided in the question. </think>\n<answer>C</answer><glue>[(130.67, 140.07)]</glue>



1012

Figure 8: A qualitative example for long video understanding.

1014

E LIMITATION AND FUTURE DIRECTION

1019

Our current approach performs only a single round of zoom-in during both training and inference. We did not explore iterative or recursive zooming due to computational constraints. However, multi-stage zooming could further refine temporal grounding by progressively narrowing the search space and focusing on increasingly fine-grained visual cues.

1024

Another limitation is that the zoom-in process is enforced rather than adaptive. Ideally, the model itself should decide whether, when, and how many times to zoom in, guided by the task objective. A goal-oriented, multi-step zooming policy could potentially yield more efficient and faithful grounding.

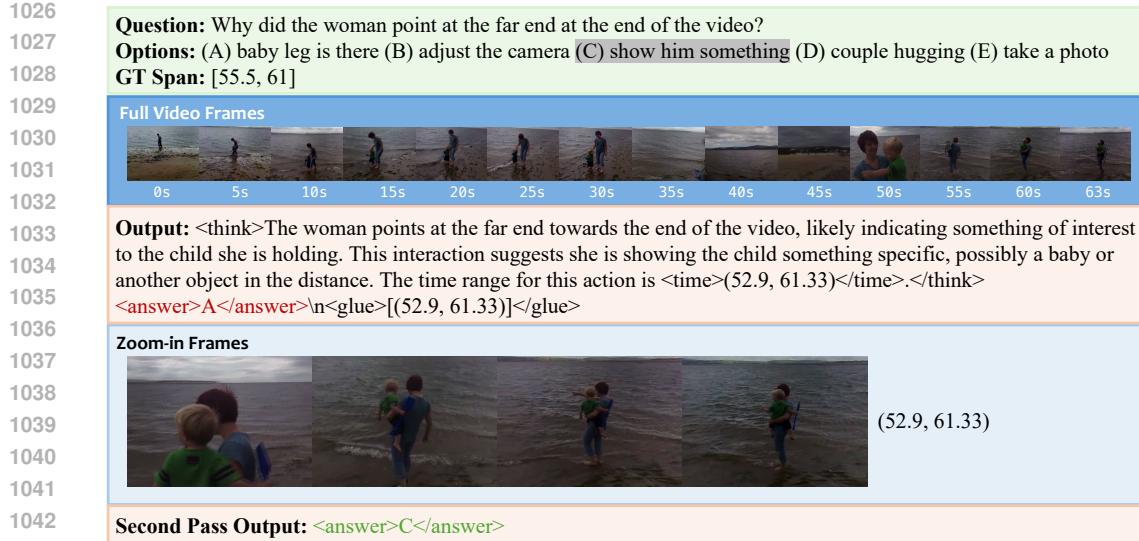


Figure 9: A qualitative example for grounded videoQA with temporal zoom-in.

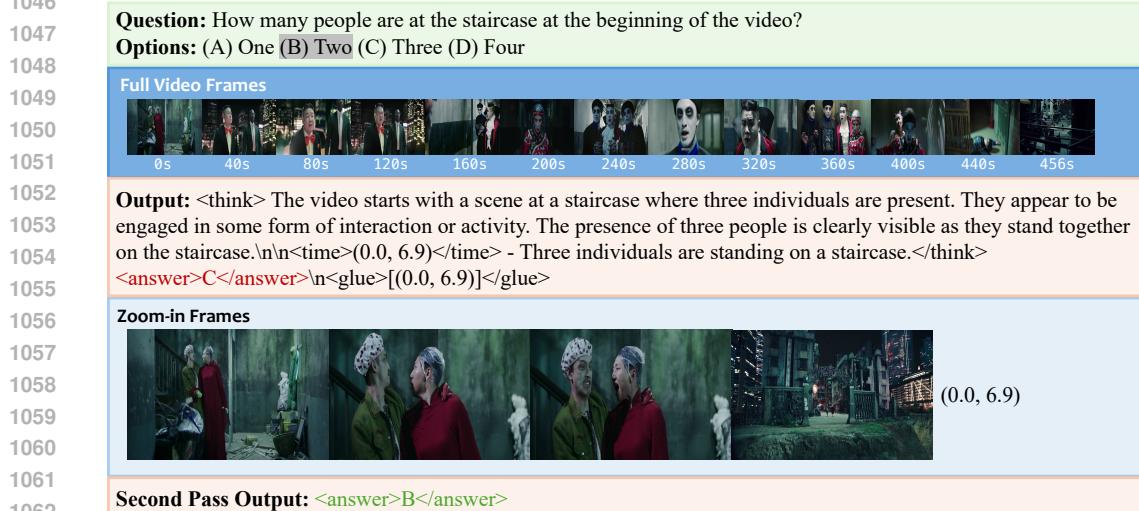


Figure 10: A qualitative example for long video understanding with coarse-to-fine zoom-in.

1063
1064
1065
1066 If a strong pretrained model with reasonable temporal grounding ability and exploration samples is
1067 sufficient, our framework could also be trained without explicit temporal interval annotations. Instead
1068 of relying on rule-based reward, i.e., R_{IoU} , the model could learn to verify whether key visual clues
1069 exist within its predicted temporal segments. This self-verification mechanism has the potential to
1070 mutually enhance answer accuracy and temporal grounding, especially in long-video scenarios where
1071 temporal annotations are often scarce to obtain.
1072
1073
1074
1075
1076
1077
1078
1079

1080
1081
1082
1083
1084
1085
1086
1087
1088
1089
1090
1091
1092
1093
1094
1095
1096
1097

1098 **Question:** What did I take out of the cupboard?

1099 **Options:** (A) cotton hand gloves I took out of the cupboard (B) wool sweater
1100 (C) plastic hangers (D) metal spoon

1101 **Full Video Frames**



1102 **Output:** <think> The video shows a person opening a cupboard and taking out items. The specific item taken out is
1103 not clearly visible, but it appears to be a piece of clothing based on the context of the action. The options provided
1104 suggest that the item could be either a wool sweater or cotton hand gloves. However, without a clear view of the item,
1105 it's difficult to determine which one it is</think>\n<answer>B</answer>\n<glue>[(224.0, 251.0)]</glue>

1106 **Zoom-in Frames**



1107 **Second Pass Output:** <answer>A</answer>

1108
1109
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1130
1131
1132
1133