

LENS: Learning to Segment Anything with Unified Reinforced Reasoning

Lianghui Zhu^{1*}, Bin Ouyang^{1*†}, Yuxuan Zhang¹, Tianheng Cheng^{1‡}, Rui Hu¹, Haocheng Shen², Longjin Ran², Xiaoxin Chen², Li Yu¹, Wenyu Liu¹, Xinggang Wang^{1§}

¹School of EIC, Huazhong University of Science & Technology,
²vivo Mobile Communication Co., Ltd

Abstract

Text-prompted image segmentation enables fine-grained visual understanding and is critical for applications such as human-computer interaction and robotics. However, existing supervised fine-tuning methods typically ignore explicit chain-of-thought (CoT) reasoning at test time, which limits their ability to generalize to unseen prompts and domains. To address this issue, we introduce LENS, a scalable reinforcement-learning framework that jointly optimizes the reasoning process and segmentation in an end-to-end manner. We propose unified reinforcement-learning rewards that span sentence-, box-, and segment-level cues, encouraging the model to generate informative CoT rationales while refining mask quality. Using a publicly available 3-billion-parameter vision-language model, i.e., Qwen2.5-VL-3B-Instruct, LENS achieves an average cloU of 81.2% on the RefCOCO, RefCOCO+, and RefCOCOg benchmarks, outperforming the strong fine-tuned method, i.e., GLaMM, by up to 5.6%. These results demonstrate that RL-driven CoT reasoning significantly enhances text-prompted segmentation and offers a practical path toward more generalizable Segment Anything models (SAM).

Code & Models — <https://github.com/hustvl/LENS>

Project Page — <https://hustvl.github.io/LENS/>

Extended version — <https://arxiv.org/abs/2508.14153>

Introduction

Text-prompted segmentation takes a natural-language description and an image as input and returns a fine-grained segmentation mask. Unlike conventional semantic segmentation that relies on a fixed set of category labels, text-prompted segmentation must jointly interpret language and vision to localize arbitrary open-vocabulary objects. This requirement gives rise to three core challenges: (i) cross-modal localization of text-referenced objects, (ii) multi-step relational reasoning across modalities, and (iii) pixel-level

*These authors contributed equally.

†Work was done during Bin Ouyang’s internship at vivo Mobile Communication Co., Ltd.

‡Project leader

§Corresponding author (xgwang@hust.edu.cn)

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

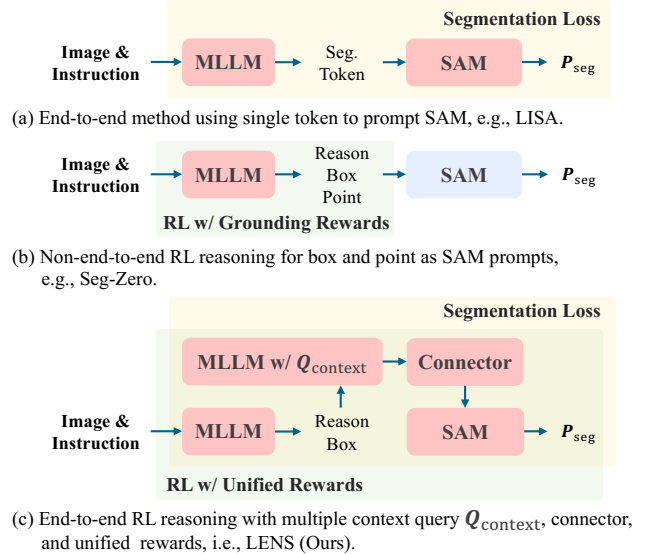


Figure 1: Framework Comparison between proposed LENS and other methods.

alignment between linguistic cues and image regions. Because of these capabilities, text-prompted segmentation is well-suited to real-world scenarios such as robotics, where an agent must understand its environment before acting.

Recent studies (Lai et al. 2024; Bai et al. 2024; Ren et al. 2024) incorporate multimodal large language models (MLLMs) to improve cross-modal localization. These methods drive the segmentation process with a single token “<seg>” and train the entire pipeline via supervised fine-tuning (SFT). However, this paradigm faces two major limitations: (i) it neglects the intermediate reasoning process that is essential for complex, reasoning-intensive tasks, and (ii) its heavy reliance on SFT often leads to overfitting and weak generalization. These issues motivate us to pursue a more robust and generalizable test-time reasoning framework.

Group Relative Policy Optimization (GRPO) (Shao et al. 2024) is a rule-based reinforcement-learning (RL) algorithm for post-training large language models. By optimizing policies with group-relative rewards, GRPO strengthens reasoning ability and empirically generalizes better than SFT. Mo-

tivated by GRPO, we present LENS, a unified, test-time reasoning framework for text-prompted segmentation. LENS adopts the GRPO strategy with the proposed unified rewards that simultaneously consider sentence-, box-, and segment-level cues. Furthermore, we introduce a context module to bridge the MLLM and the segmentation model, which extracts the reasoning and grounding information as a prior to guide the segmentation process.

Concurrent work SegZero (Liu et al. 2025) directly feeds the grounding output, i.e., the bounding box and points, of an MLLM into a frozen SAM, achieving strong reasoning segmentation ability. However, the final mask quality is fully determined by the MLLM’s grounding, and any errors in grounding will propagate downstream, leading to segmentation errors. In contrast, LENS offers an end-to-end solution that jointly optimizes language understanding and pixel-wise mask prediction (Fig. 1). We establish a tight coupling between the MLLM and SAM through the proposed context module and a pre-alignment stage. Moreover, the proposed unified reinforcement learning rewards, i.e., format reward, box-IoU reward, and segment-IoU reward, encourage the model to produce informative chain-of-thought rationales while refining mask quality. Extensive experiments show that LENS sets a new state-of-the-art on standard text-prompted segmentation benchmarks.

To sum up, our contributions are as follows:

- We propose LENS, an end-to-end, test-time reasoning framework that jointly optimizes multimodal language understanding and pixel-level segmentation for text-prompted tasks.
- We introduce a context module, i.e., multiple context queries and a connector, to bridge the MLLM and the segmentation model. Through the pre-alignment stage, the context module can transform the chain-of-thought reasoning trace and grounding box into spatial priors that guide mask generation.
- Built upon GRPO, we propose a unified rewards scheme that simultaneously supervises sentence-level reasoning, object localization, and pixel-wise mask quality within a single reinforcement-learning objective.
- LENS achieves 81.2% average cIoU on RefCOCO-series benchmarks, 58.0% cIoU on ReasonSeg-Test, and 78.3% cIoU on GS-Eval, establishing new state-of-the-art performance for text-prompted segmentation.

Related Work

Text-prompted Segmentation Text-prompted segmentation, i.e., referring segmentation, aims to segment objects described by natural language expressions. This task requires both visual understanding and language comprehension capabilities. Early work includes ReferIt (Kazemzadeh et al. 2014) dataset and corresponding models that localize objects based on natural language descriptions.

Recent advances have focused on multimodal fusion architectures. LAVT (Yang et al. 2022) proposed a lightweight vision-transformer approach for referring segmentation. RefTR (Li and Sigal 2021) employed transformer-based cross-modal fusion. CRIS (Wang et al. 2022) introduced

contrastive learning for improved text-image alignment. More recent work includes X-Decoder (Zou et al. 2023a) and SEEM (Zou et al. 2023b), which unify various segmentation tasks including text-prompted segmentation.

The integration of large language models has opened new possibilities. LISA (Lai et al. 2024) combines SAM with large language models for reasoning segmentation. PerSAM (Zhang et al. 2023) enables few-shot personalization of SAM. However, these approaches primarily rely on supervised fine-tuning, which may limit their reasoning capabilities and generalization to unseen scenarios.

Reinforcement Learning in Large Language Models

Reinforcement Learning from Human Feedback (RLHF) has become a cornerstone technique for aligning large language models with human preferences. PPO (Proximal Policy Optimization) (Schulman et al. 2017a) serves as the foundational algorithm, enabling stable policy updates while preventing catastrophic forgetting.

Recent breakthroughs include InstructGPT (Ouyang et al. 2022) and ChatGPT, which demonstrated the effectiveness of RLHF in producing helpful and harmless responses. Constitutional AI (Bai et al. 2022) proposed training models to follow a set of principles. DPO (Direct Preference Optimization) (Rafailov et al. 2023) simplified the RLHF pipeline by directly optimizing on preference data.

The emergence of reasoning-focused RL approaches has shown remarkable success. DeepSeek-R1 (Guo et al. 2025) demonstrated that large-scale reinforcement learning can significantly improve reasoning capabilities, achieving competitive performance with leading models. Chain-of-thought reasoning with RL supervision has been explored in various contexts (Zelikman et al. 2022; Huang et al. 2022), showing that explicit reasoning processes can be learned and improved through reinforcement learning.

However, most existing RL approaches focus on text generation tasks. The application of RL to multimodal tasks, particularly vision-language understanding and segmentation, remains largely unexplored. Our work bridges this gap by introducing RL-based reasoning for text-prompted segmentation, enabling more robust and generalizable segmentation models.

LENS

In this section, we introduce the proposed LENS architecture and the proposed reinforcement learning training strategy. First, we introduce the proposed LENS architecture that enables end-to-end reasoning and segmentation. Then, we introduce the proposed pretraining alignment stage that establishes the foundational connection between the MLLM and SAM. Finally, we introduce the proposed reinforcement learning stage that jointly optimizes the model’s reasoning and segmentation capabilities.

LENS Architecture

The concurrent SegZero (Liu et al. 2025) triggers a frozen SAM with bounding boxes and points predicted by an MLLM, which achieves strong reasoning ability. However, its non-end-to-end nature hinders the joint alignment of

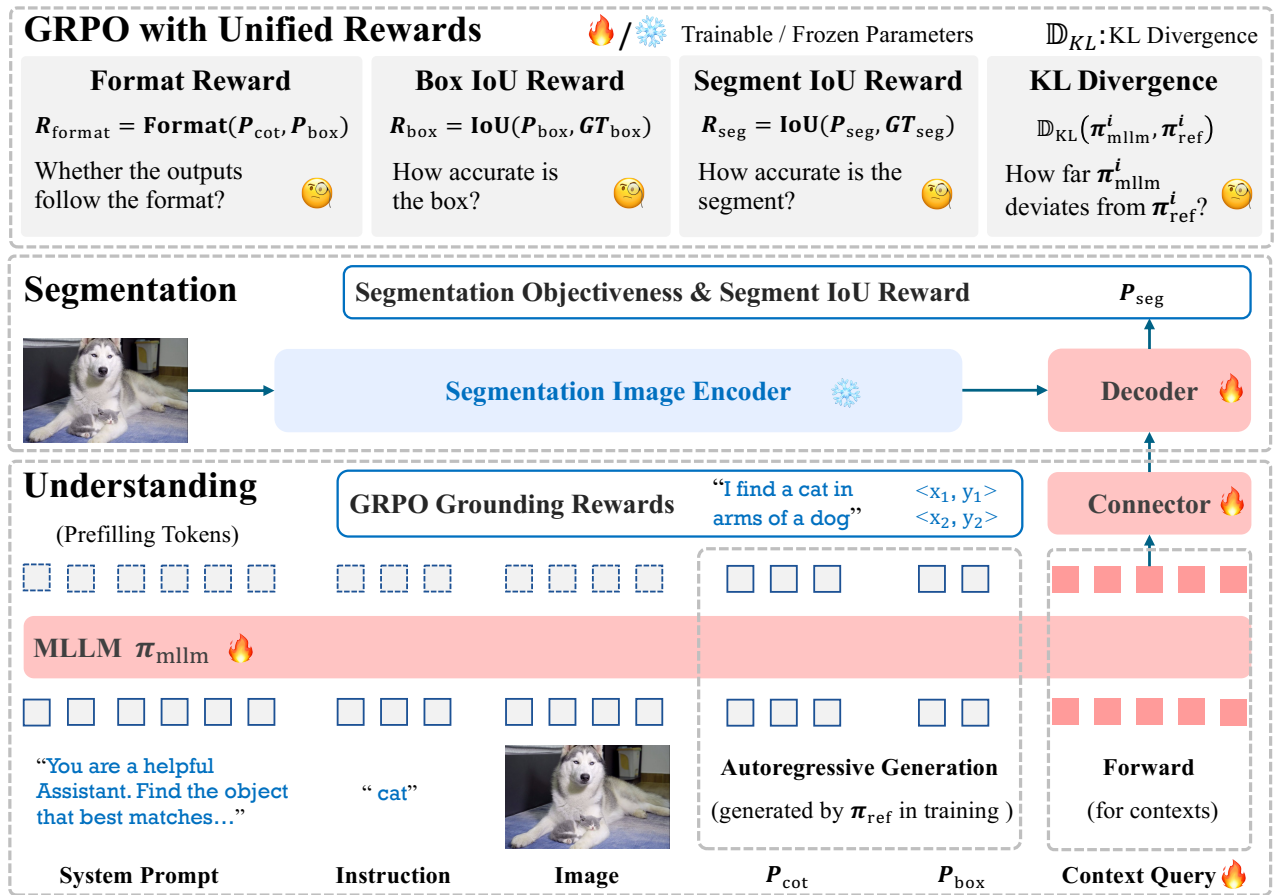


Figure 2: An Overview of LENS framework. In the pretraining alignment stage, we only train the context query and connector with the segmentation objectiveness. In the reinforcement learning stage, we train all the parts except for the segmentation image encoder with the multi-grained objectiveness, i.e., the unified GRPO rewards and segmentation loss.

understanding and segmentation representations and ultimately yields sub-optimal segmentation masks. To overcome this limitation, we introduce LENS, an end-to-end reinforcement-learning framework that couples an MLLM with SAM through a context module. Thereby, the model can jointly optimize the reasoning and segmentation capabilities under the multi-grained unified rewards and pixel-level entropy loss.

Reasoning Model. We use MLLMs, i.e., Qwen2.5-VL (Bai et al. 2025), as the reasoning model π_{mllm} . Previous MLLMs (Liu et al. 2023; Bai et al. 2023) show promising performance in object-centric localization, but still lack pixel-level perception ability. Some work (Lai et al. 2024; Bai et al. 2024) incorporates pre-trained segmentation models into MLLMs to compensate for the lack of pixel-level perception ability. However, these methods simply use a segmentation token in short response, i.e., ‘It is $\langle \text{seg} \rangle$ ’, to guide the segmentation, which neglects the reasoning process, e.g., Chain-of-Thought (CoT) (Wei et al. 2022). To address this, we propose to utilize the reasoning process to guide the segmentation, which brings more robust and effective segmentation ability. We input the system prompt p_{sys} ,

textual instruction T and image I to the reasoning model, and obtain the CoT P_{CoT} and box prediction P_{box} .

$$P_{\text{CoT}}, P_{\text{box}} = \pi_{\text{mllm}}.\text{generate}(p_{\text{sys}}, T, I). \quad (1)$$

Context Module. Concurrent method (Liu et al. 2025) only rely detokenized bounding box and points to prompt the SAM, which leads to suboptimal performance and eliminates the end-to-end optimization. To address this, we propose a context module to bridge the gap between understanding and segmentation during the reinforcement learning process. The context module contains the context query and the connector. After the MLLM generates the CoT and box prediction, we use the context query to extract the context information from the MLLM. Then, we use the connector to project the context information as the prompt.

We randomly initialize the context query $Q \in \mathbb{R}^{M \times C}$, where M is the number of context queries and C is the hidden dimension of the MLLM. Notably, the context query is appended to the end of the input, and generated sequence, gathering the information through a forward pass.

$$Q' = \pi_{\text{mllm}}.\text{forward}(p_{\text{sys}}, T, I, P_{\text{CoT}}, P_{\text{box}}, Q), \quad (2)$$

$$Q_{\text{seg}} = \pi_{\text{connector}}(Q') \quad (3)$$

Next, we feed the context query embeddings $Q' \in \mathbb{R}^{M \times C}$ output by the MLLM to the connector, which is a shallow transformer to project the context query embeddings to the prompt space of the SAM, i.e., Q_{seg} .

Segmentation Model. SAM (Kirillov et al. 2023) is the most popular segmentation model (Cheng et al. 2025), which accepts prompts including point and box and produces precise segmentation results. Some methods (Zhang et al. 2024b; Lai et al. 2024) incorporate MLLMs, which enable SAM to segment objects with text prompts. However, these methods output the single segmentation token ‘< seg >’ directly, which neglects the reasoning process and token capacity. To address this, we propose to use multiple queries to extract reasoning context and ensure enough token capacity. Moreover, we incorporate pixel-level entropy loss to reinforcement learning training, which provides more fine-grained optimization guidance. We feed the projected segmentation prompt Q_{seg} to the SAM model π_{sam} for segmentation P_{seg} .

$$P_{\text{seg}} = \pi_{\text{sam}}(Q_{\text{seg}}, I). \quad (4)$$

Pretraining Alignment Stage

In this stage, we establish the foundational connection between the MLLM and SAM to enable end-to-end segmentation. Since understanding and segmentation models have different pretrained representations, we introduce the context module to bridge this gap while preserving the original capabilities of both models.

During pretraining alignment, we freeze the weights of both the MLLM (π_{mllm}) and SAM (π_{sam}), training only the lightweight context module, i.e., the connector and the context queries. This approach ensures that the rich pretrained knowledge in both models is preserved while enabling them to work together effectively.

The training objective focuses on learning segmentation-aware representations through the context query and task connector. We train the model on segmentation datasets where the context query learns to extract relevant visual-semantic information from the MLLM’s representations, and the task connector learns to transform this information into effective prompts for SAM.

We use segmentation objectives in this stage:

$$\mathcal{L}_{\text{align}} = \mathcal{L}_{\text{seg}}(P_{\text{seg}}, M_{\text{gt}}), \quad (5)$$

where M_{gt} represents the ground truth segmentation mask and \mathcal{L}_{seg} is the segmentation loss (e.g., dice loss, focal loss).

Reinforcement Learning Stage

While the pretraining alignment stage enables basic multimodal understanding and segmentation, the quality of reasoning process significantly impacts segmentation performance. Existing methods (Liu et al. 2025; Shen et al. 2025) typically optimize only the understanding part, overlooking the crucial segmentation decoder.

To address this limitation, we propose unified rewards ranging from sentence-, box-, and segment-level to jointly optimize the reasoning quality and segmentation accuracy.

In this stage, we unfreeze the MLLM and segmentation decoder parameters while keeping segmentation image encoder frozen, allowing the model to learn enhanced reasoning strategies. The improved Chain-of-Thought (CoT) reasoning P_{CoT} serves as richer contextual information to strengthen the context query representations, ultimately leading to better segmentation results.

Group Relative Policy Optimization (GRPO). The GRPO (Shao et al. 2024) introduces a new rule-based reinforcement learning algorithm for post-training optimization of large language models. Traditional reinforcement learning algorithms, i.e., PPO (Schulman et al. 2017b), use separate critic models to evaluate the quality of the policy model, leading to high memory and computational costs. To alleviate this problem, GRPO eliminates the separate critic model and only uses group relative rewards to guide the optimization of the policy model. GRPO employs reward function $R(\cdot)$ to evaluate the outputs and Kullback-Leibler (KL) divergence penalty \mathbb{D}_{KL} to regularize the policy update. Due to the page limitation, we leave more details in the appendix.

GRPO with Unified Rewards. Reward function is the key to the success of reinforcement learning, as it guides the model to learn the desired behavior. Text-driven segmentation tasks encompass both understanding and segmentation, which is more complex than standard MLLM tasks.

Unlike standard VLM-R1 methods (Shen et al. 2025; Yu et al. 2025) that rely on purely understanding reward signal, i.e., format reward and box IoU reward, we introduce a comprehensive, multi-faceted reward system specifically designed for both understanding and segmentation. Our reward function encompasses three complementary components:

- **Format Reward (R_{format}):** Ensures the MLLM output adheres to the expected structure and format consistency. The format template requires the MLLM to output the reasoning process in ‘<thinking>’ ... ‘</thinking>’ tag pair and localization results in ‘<answer>’ ... ‘</answer>’ tag pair. If the format is correct, the reward is 1, otherwise, the reward is 0.
- **Box IoU Reward (R_{box}):** Measures localization accuracy through IoU between predicted and ground truth bounding boxes. The range is $[0, 1]$.
- **Segment IoU Reward (R_{seg}):** Evaluates overall segmentation quality using mask-level IoU. We introduce the segment IoU reward to evaluate the segmentation quality, whose range is $[0, 1]$.

Additionally, we use a KL Divergence (\mathbb{D}_{KL}) regularization term to prevent the context query representations from deviating significantly from the pretrained representations.

The unified reward function is formulated as:

$$R_{\text{unified}} = \lambda_1 R_{\text{format}} + \lambda_2 R_{\text{box}} + \lambda_3 R_{\text{seg}}, \quad (6)$$

where $\{\lambda_i\}_{i=1}^3$ are balancing hyperparameters.

Training Objective. To achieve comprehensive optimization, we combine the unified-reward-based GRPO objective with supervised segmentation loss:

$$\mathcal{J}_{\text{LENS}}(\theta) = \mathcal{J}(\theta; R_{\text{unified}}, \mathbb{D}_{\text{KL}}) + \alpha \mathcal{L}_{\text{seg}}(P_{\text{seg}}, M_{\text{gt}}) \quad (7)$$

Method	RefCOCO			RefCOCO+			RefCOCOg		AVG
	val	testA	testB	val	testA	testB	val	test	
<i>without active chain-of-thought (CoT) reasoning</i>									
LAVT (Yang et al. 2022)(CVPR 22)	72.7	75.8	68.8	62.1	68.4	55.1	61.2	62.1	65.8
ReLA (Liu, Ding, and Jiang 2023)(CVPR 23)	73.8	76.5	70.2	66.0	71.0	57.7	65.0	66.0	68.3
X-Decoder (Zou et al. 2023a)(CVPR 23)	-	-	-	-	-	-	64.6	-	-
SEEM (Zou et al. 2023b)(NeurIPS 23)	-	-	-	-	-	-	65.7	-	-
LISA (Lai et al. 2024) (CVPR 24)	74.1	76.5	71.1	62.4	67.4	56.5	66.4	68.5	67.9
PixelLM (Ren et al. 2024)(CVPR 24)	76.9	78.5	74.4	69.2	72.1	64.5	70.7	72.4	72.3
PerceptionGPT-7B (Pi et al. 2024)(CVPR 24)	75.1	78.6	71.7	68.5	73.9	61.3	70.3	71.7	71.4
OMG-LLaVA (Zhang et al. 2024a)(NeurIPS 24)	78.0	80.3	74.1	69.1	73.1	63.0	72.9	72.9	72.9
VISA (Yan et al. 2024)(ECCV 24)	72.4	75.5	68.1	59.8	64.8	53.1	65.5	66.4	65.7
GLaMM (Rasheed et al. 2024)(CVPR 24)	79.5	<u>83.2</u>	76.9	72.6	<u>78.7</u>	64.6	74.2	74.9	75.6
SAM3-Agent-Gemini2.5-Pro (OpenReview 2025)	74.9	77.8	69.9	66.9	71.1	62.4	73.3	73.6	74.2
<i>with active chain-of-thought (CoT) reasoning</i>									
Seg-Zero-3B (Liu et al. 2025) (arXiv 25)	-	79.3	-	-	73.7	-	71.5	-	-
Seg-Zero-7B (Liu et al. 2025) (arXiv 25)	-	80.3	-	-	76.2	-	72.6	-	-
LENS-2B (Ours)	<u>80.7</u>	82.7	<u>77.1</u>	<u>73.8</u>	78.0	<u>67.3</u>	<u>75.7</u>	<u>76.4</u>	<u>76.5</u>
LENS-3B (Ours)	84.2	85.3	81.0	79.4	82.8	74.3	81.2	81.0	81.2

Table 1: Comparison on the Referring Expression Segmentation (RES) task using cIoU as the evaluation metric. Our method achieves state-of-the-art performance across all RefCOCO, RefCOCO+, and RefCOCOg benchmarks. **Bold** indicates the best (SoTA) performance, and underlined denotes the second-best. The “active chain-of-thought (CoT) reasoning” means using explicit chain-of-thought at test time.

where $\mathcal{J}(\theta; R_{\text{unified}}, \mathbb{D}_{\text{KL}})$ represents the standard GRPO objective using our unified rewards and KL divergence regularization, α controls the balance between reinforcement learning and supervised learning signals, and $\mathcal{L}_{\text{seg}}(P_{\text{seg}}, M_{\text{gt}})$ is the segmentation loss.

This joint optimization enables the model to benefit from both reward-driven reasoning improvements and direct segmentation supervision, resulting in enhanced performance across reasoning and segmentation capabilities.

Experiment

Implementation Details We employ Qwen-VL series models (Bai et al. 2025; Wang et al. 2024) as the reasoning model and SAM2-Large (Ravi et al. 2024) as the segmentation model. Specifically, LENS-2B utilizes Qwen2-VL and LENS-3B uses Qwen2.5-VL. Training is conducted on 16 NVIDIA L40S GPUs, with our pipeline built upon the DeepSpeed engine (Rasley et al. 2020). Detailed experimental settings are summarized in the Appendix.

Benchmarks. We evaluate the proposed framework on multiple benchmarks: Referring Expression Segmentation (RES) benchmarks (RefCOCO series) (Nagaraja, Morariu, and Davis 2016; Yu et al. 2016), GroundingSuite benchmarks (Hu et al. 2025), and Reasoning Segmentation benchmarks (Lai et al. 2024). The quantitative comparisons are presented in Table 1 and Table 2, respectively.

Evaluation Metrics. Following previous works on referring segmentation (Kazemzadeh et al. 2014), we adopt two commonly used evaluation metrics: generalized IoU (gIoU) and cumulative IoU (cIoU). Specifically, gIoU is computed

as the average of per-image Intersection-over-Union (IoU) scores, while cIoU is defined as the IoU of the cumulative predicted and ground-truth masks across the entire dataset.

Main Results

Note that PSALM (Zhang et al. 2024c), HyperSeg (Wei et al. 2024a), and InstructSeg (Wei et al. 2024b) utilize pre-trained Mask2Former (Cheng et al. 2022) models, whose training data, i.e., COCO 2017 (Lin et al. 2014), has overlap with the RefCOCO benchmarks. Therefore, we compare with them in Table 2, evaluating their performance on the Reasoning Segmentation and GroundingSuite-Eval benchmarks.

Referring Expression Segmentation. As shown in Table 1, the proposed LENS achieves state-of-the-art performance across all splits of the RefCOCO, RefCOCO+, and RefCOCOg benchmarks. Compared to prior models such as GLaMM (Rasheed et al. 2024), our method delivers significant improvements, especially on RefCOCO+ and RefCOCOg, which are known to require a more precise understanding of compositional and spatial language. Notably, LENS-3B surpasses the best previous method by 6.1% on RefCOCOg-test, highlighting the model’s strength in handling complex referring expressions.

ReasonSeg and GroundingSuite-Eval Benchmarks. As shown in Table 2, we evaluate the proposed LENS on the ReasonSeg and GroundingSuite-Eval benchmarks. GroundingSuite (GS) is designed for challenging scenarios, and none of the methods in Table 2 are trained on GS-Train. On the zero-shot GS-Eval benchmark, LENS achieves a 9.4% improvement over the second-best method, demonstrating strong out-of-domain (OOD) generalization.

Method	ReasonSeg-Val		ReasonSeg-Test		GroundingSuite-Eval	
	gIoU	cIoU	gIoU	cIoU	gIoU	cIoU
PSALM (Zhang et al. 2024c)(ECCV 24)	-	-	-	-	39.2	34.4
HyperSeg-3B [‡] (Wei et al. 2024a)(CVPR 25)	59.2	56.7	-	-	58.2	62.8
InstructSeg-3B [‡] (Wei et al. 2024b)(ICCV 25)	61.9	65.2	-	-	55.7	57.0
LISA-7B(ft) (Lai et al. 2024) (CVPR 24)	52.9	54.0	55.6	<u>56.9</u>	60.9	68.6
Seg-Zero-3B (Liu et al. 2025) (arXiv 25)	58.2	53.1	56.1	48.6	67.3	<u>68.9</u>
SAM3-Agent-Qwen2.5-VL-7B (OpenReview 2025)	65.4	50.5	62.6	56.2	65.7	58.3
LENS-3B (Ours)	<u>62.1</u>	<u>64.9</u>	<u>57.2</u>	58.0	<u>67.0</u>	78.3

Table 2: Performance comparison on the ReasonSeg and GroundingSuite-Eval benchmarks, evaluated using gIoU and cIoU. **Bold** and underlined entries indicate the best and second-best results, respectively. [‡] marks models trained on both train and test splits of ReasonSeg. All models are not trained on GroundingSuite. LENS shows superior reasoning and generalization ability.




Figure 3: Qualitative results on the Referring Expression Segmentation task. The proposed LENS can accurately segment partially obscured objects. Benefit from the proposed unified framework, even if there is an error in the context box, the segmentation module can correct it based on the rich context in the multiple queries.

Query Num.	1	16	32	64	128
cIoU (%)	85.9	87.4	87.8	87.9	87.2

Table 3: Ablation study on the number of queries used in context module, evaluated cIoU metric on RefCOCO testA in pretraining alignment stage. Performance improves with more queries up to 64, after which it slightly drops, indicating a saturation point.

In terms of segmentation quality, LENS achieves competitive results in gIoU and exhibits a clear lead in cIoU, which captures holistic segmentation performance across objects and scenes. For example, on GS-Eval, LENS attains a cIoU of 78.3%, surpassing Seg-Zero-3B by 9.4%.

Ablation Studies

Number of Queries. We examine the impact of varying the number of queries in the reasoning module. As shown in Table 3, performance improves steadily from 1 to 64 queries, suggesting that a larger query set enables the model to capture more contextual reasoning cues and deliver to the seg-

Connector	MLP	ViT
cIoU (%)	72.8	87.8

Table 4: Ablation study on the connector module between MLLM and segmentation model, evaluated cIoU metric on RefCOCO testA in pretraining alignment stage. Using a ViT connector significantly outperforms a simple MLP, demonstrating the benefit of attention-based context fusion.

mentation task. However, using 128 queries leads to a slight drop, likely due to increased redundancy and computational overhead. This indicates that 64 queries offer a good trade-off between accuracy and efficiency.

Connector Design. We further investigate the effectiveness of our connector. The connector is implemented as a 4-layer Qwen-2.5 Transformer with a hidden dimension of 2048, matching the MLLM output. As shown in Table 4, replacing a simple MLP with this Transformer significantly improves performance. Overall, these studies underscore the importance of architectural choices in enabling effective reasoning and accurate segmentation within our framework.

Reasoning segmentation for LENS

System Prompt:

You are a helpful assistant. Find the object that best matches "{object to segment}" and provide its bounding box.

Please:

1. Analyze all objects in the image carefully
2. Compare candidates against the target description
3. Select the most closely matching object
4. Provide precise bounding box coordinates

Format your response as:

<think>

[Your step-by-step analysis and reasoning]

</think>

<answer>

```json

```
[
 {"bbox_2d": [x1,y1,x2,y2], "label":
question}
]
```

...

</answer>

### User Instruction:

**the car that may need repair**

Assistant:

<think>

The task is to identify **the car that may need repair** among the vehicles shown in the image. The key indicators for this include visible damage or wear, open hoods, or any other signs of mechanical issues.

Looking at the image:

- The black car on the left has its hood open, which could indicate it needs repair or maintenance.

- The blue car on the right appears to be in good condition with no visible damage or open hoods.

Based on these observations, the black car on the left is the one that may need repair.

</think>

<answer>

```json

```
[
  {"bbox_2d": [0, 129, 365, 347], "label": "the car that may
need repair"}
]
```

...

</answer><context query>



Input image



Output segmentation

Figure 4: Visualization of reasoning segmentation. MLLM first generates a CoT reasoning process and a probable box as priors. Then the context queries extract messages from priors, and prompt segmentation module for an accurate mask.

| Context Module | Align. Stage | RL Stage | | | cIoU |
|----------------|--------------|------------------|---------------------|------------------|-------------|
| | | R_{box} | R_{format} | R_{seg} | |
| × | × | × | × | × | 79.3 |
| ✓ | × | ✓ | × | × | – |
| ✓ | ✓ | × | × | × | 80.9 |
| ✓ | × | ✓ | ✓ | × | 72.3 |
| ✓ | ✓ | ✓ | ✓ | × | 81.9 |
| ✓ | ✓ | ✓ | ✓ | ✓ | 82.7 |

Table 5: Ablation study of components on RefCOCO testA.

Framework Analysis. To investigate the contribution of each component in our framework, we conduct an ablation study on the RefCOCO testA split, as shown in Table 5. We start with a Lisa-like baseline (Lai et al. 2024) model without the context module, alignment stage, or reinforcement learning (RL), which achieves a cIoU of 79.3%. Directly adding the context module and applying the box reward makes the training unstable. After enabling format reward, the model captures basic localization cues, but the lack of explicit alignment between textual and visual features leads to sub-optimal performance. To address this issue, we enable the alignment stage, where the context module is pre-trained to align MLLM and SAM. This significantly improves the segmentation accuracy, boosting cIoU to 81.9%, and demonstrates the critical importance of semantic alignment in the proposed unified framework. Finally, introducing the segmentation-level reward (R_{seg}) further refines the

mask quality, pushing the performance to 82.7% cIoU.

Visualizations Fig. 3 shows examples from the RefCOCOg benchmark. LENS accurately segments target objects based on referring expressions, demonstrating strong performance in scenes with partially obscure objects and multiple similar instances. Notably, the proposed unified framework and end-to-end optimization make LENS robust against potential errors in the box prior. Fig. 4 presents results on the reasoning segmentation benchmark. LENS exhibits robust test-time reasoning capabilities, effectively handling tasks that require spatial understanding, multi-step inference, and comparative reasoning. These visualizations highlight the LENS’s ability to integrate high-level textual reasoning with fine-grained visual perception.

Conclusion

In this paper, we introduce LENS, a unified test-time reasoning framework for text-prompted segmentation. By coupling a multimodal LLM with a segmentation model through a context module and a pretraining alignment stage, LENS treats language understanding and mask prediction as equally critical components. The extracted chain-of-thought rationales serve as spatial priors that directly guide segmentation. In addition, we design a unified reinforcement-learning (RL) objective with unified rewards that simultaneously supervise sentence-level reasoning, object localization, and pixel-accurate masking. We believe LENS offers fresh insights into the seamless integration of RL and visual segmentation and will spur further research toward more general, robust, and intelligent vision–language systems.

Acknowledgements

This work was partially supported by the National Natural Science Foundation of China (No. 62276108).

References

- Bai, J.; Bai, S.; Yang, S.; Wang, S.; Tan, S.; Wang, P.; Lin, J.; Zhou, C.; and Zhou, J. 2023. Qwen-VL: A Versatile Vision-Language Model for Understanding, Localization, Text Reading, and Beyond. *arXiv:2308.12966*.
- Bai, S.; Chen, K.; Liu, X.; Wang, J.; Ge, W.; Song, S.; Dang, K.; Wang, P.; Wang, S.; Tang, J.; et al. 2025. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*.
- Bai, Y.; Kadavath, S.; Kundu, S.; Askell, A.; Kernion, J.; Jones, A.; Chen, A.; Goldie, A.; Mirhoseini, A.; McKinnon, C.; et al. 2022. Constitutional ai: Harmlessness from ai feedback. *arXiv preprint arXiv:2212.08073*.
- Bai, Z.; He, T.; Mei, H.; Wang, P.; Gao, Z.; Chen, J.; Zhang, Z.; and Shou, M. Z. 2024. One token to seg them all: Language instructed reasoning segmentation in videos. *Advances in Neural Information Processing Systems*, 37: 6833–6859.
- Cheng, B.; Misra, I.; Schwing, A. G.; Kirillov, A.; and Girshick, R. 2022. Masked-attention mask transformer for universal image segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 1290–1299.
- Cheng, T.; Wang, X.; Liao, J.; and Liu, W. 2025. Cross-layer attentive feature upsampling for low-latency semantic segmentation. *Machine Vision and Applications*, 36(1): 18.
- Guo, D.; Yang, D.; Zhang, H.; Song, J.; Zhang, R.; Xu, R.; Zhu, Q.; Ma, S.; Wang, P.; Bi, X.; et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Hu, R.; Zhu, L.; Zhang, Y.; Cheng, T.; Liu, L.; Liu, H.; Ran, L.; Chen, X.; Liu, W.; and Wang, X. 2025. GroundingSuite: Measuring Complex Multi-Granular Pixel Grounding. *arXiv preprint arXiv:2503.10596*.
- Huang, J.; Gu, S. S.; Hou, L.; Wu, Y.; Wang, X.; Yu, H.; and Han, J. 2022. Large language models can self-improve. *arXiv preprint arXiv:2210.11610*.
- Kazemzadeh, S.; Ordonez, V.; Matten, M.; and Berg, T. 2014. Referitgame: Referring to objects in photographs of natural scenes. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, 787–798.
- Kirillov, A.; Mintun, E.; Ravi, N.; Mao, H.; Rolland, C.; Gustafson, L.; Xiao, T.; Whitehead, S.; Berg, A. C.; Lo, W.-Y.; et al. 2023. Segment anything. In *Proceedings of the IEEE/CVF international conference on computer vision*, 4015–4026.
- Lai, X.; Tian, Z.; Chen, Y.; Li, Y.; Yuan, Y.; Liu, S.; and Jia, J. 2024. Lisa: Reasoning segmentation via large language model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9579–9589.
- Li, M.; and Sigal, L. 2021. Referring transformer: A one-step approach to multi-task visual grounding. *Advances in neural information processing systems*, 34: 19652–19664.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *European conference on computer vision*, 740–755. Springer.
- Liu, C.; Ding, H.; and Jiang, X. 2023. Gres: Generalized referring expression segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 23592–23601.
- Liu, H.; Li, C.; Wu, Q.; and Lee, Y. J. 2023. Visual instruction tuning. *Advances in neural information processing systems*, 36: 34892–34916.
- Liu, Y.; Peng, B.; Zhong, Z.; Yue, Z.; Lu, F.; Yu, B.; and Jia, J. 2025. Seg-zero: Reasoning-chain guided segmentation via cognitive reinforcement. *arXiv preprint arXiv:2503.06520*.
- Nagaraja, V. K.; Morariu, V. I.; and Davis, L. S. 2016. Modeling context between objects for referring expression understanding. In *European Conference on Computer Vision*, 792–807. Springer.
- Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35: 27730–27744.
- Pi, R.; Yao, L.; Gao, J.; Zhang, J.; and Zhang, T. 2024. Perceptiongpt: Effectively fusing visual perception into llm. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 27124–27133.
- Rafailov, R.; Sharma, A.; Mitchell, E.; Manning, C. D.; Ermon, S.; and Finn, C. 2023. Direct preference optimization: Your language model is secretly a reward model. *Advances in neural information processing systems*, 36: 53728–53741.
- Rasheed, H.; Maaz, M.; Shaji, S.; Shaker, A.; Khan, S.; Cholakkal, H.; Anwer, R. M.; Xing, E.; Yang, M.-H.; and Khan, F. S. 2024. Glamm: Pixel grounding large multi-modal model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13009–13018.
- Rasley, J.; Rajbhandari, S.; Ruwase, O.; and He, Y. 2020. Deepspeed: System optimizations enable training deep learning models with over 100 billion parameters. In *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*, 3505–3506.
- Ravi, N.; Gabeur, V.; Hu, Y.-T.; Hu, R.; Ryali, C.; Ma, T.; Khedr, H.; Rädle, R.; Rolland, C.; Gustafson, L.; Mintun, E.; Pan, J.; Alwala, K. V.; Carion, N.; Wu, C.-Y.; Girshick, R.; Dollár, P.; and Feichtenhofer, C. 2024. SAM 2: Segment Anything in Images and Videos. *arXiv:2408.00714*.
- Ren, Z.; Huang, Z.; Wei, Y.; Zhao, Y.; Fu, D.; Feng, J.; and Jin, X. 2024. Pixellm: Pixel reasoning with large multi-modal model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 26374–26383.
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017a. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.

- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017b. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Shao, Z.; Wang, P.; Zhu, Q.; Xu, R.; Song, J.; Bi, X.; Zhang, H.; Zhang, M.; Li, Y.; Wu, Y.; et al. 2024. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*.
- Shen, H.; Liu, P.; Li, J.; Fang, C.; Ma, Y.; Liao, J.; Shen, Q.; Zhang, Z.; Zhao, K.; Zhang, Q.; et al. 2025. Vlm-r1: A stable and generalizable r1-style large vision-language model. *arXiv preprint arXiv:2504.07615*.
- Wang, P.; Bai, S.; Tan, S.; Wang, S.; Fan, Z.; Bai, J.; Chen, K.; Liu, X.; Wang, J.; Ge, W.; et al. 2024. Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*.
- Wang, Z.; Lu, Y.; Li, Q.; Tao, X.; Guo, Y.; Gong, M.; and Liu, T. 2022. Cris: Clip-driven referring image segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11686–11695.
- Wei, C.; Zhong, Y.; Tan, H.; Liu, Y.; Zhao, Z.; Hu, J.; and Yang, Y. 2024a. Hyperseg: Towards universal visual segmentation with large language model. *arXiv preprint arXiv:2411.17606*.
- Wei, C.; Zhong, Y.; Tan, H.; Zeng, Y.; Liu, Y.; Zhao, Z.; and Yang, Y. 2024b. Instructseg: Unifying instructed visual segmentation with multi-modal large language models. *arXiv preprint arXiv:2412.14006*.
- Wei, J.; Wang, X.; Schuurmans, D.; Bosma, M.; Xia, F.; Chi, E.; Le, Q. V.; Zhou, D.; et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35: 24824–24837.
- Yan, C.; Wang, H.; Yan, S.; Jiang, X.; Hu, Y.; Kang, G.; Xie, W.; and Gavves, E. 2024. Visa: Reasoning video object segmentation via large language models. In *European Conference on Computer Vision*, 98–115. Springer.
- Yang, Z.; Wang, J.; Tang, Y.; Chen, K.; Zhao, H.; and Torr, P. H. 2022. Lavt: Language-aware vision transformer for referring image segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 18155–18165.
- Yu, E.; Lin, K.; Zhao, L.; Yin, J.; Wei, Y.; Peng, Y.; Wei, H.; Sun, J.; Han, C.; Ge, Z.; et al. 2025. Perception-r1: Pioneering perception policy with reinforcement learning. *arXiv preprint arXiv:2504.07954*.
- Yu, L.; Poirson, P.; Yang, S.; Berg, A. C.; and Berg, T. L. 2016. Modeling context in referring expressions. In *European conference on computer vision*, 69–85. Springer.
- Zelikman, E.; Wu, Y.; Mu, J.; and Goodman, N. 2022. Star: Bootstrapping reasoning with reasoning. *Advances in Neural Information Processing Systems*, 35: 15476–15488.
- Zhang, R.; Jiang, Z.; Guo, Z.; Yan, S.; Pan, J.; Ma, X.; Dong, H.; Gao, P.; and Li, H. 2023. Personalize segment anything model with one shot. *arXiv preprint arXiv:2305.03048*.
- Zhang, T.; Li, X.; Fei, H.; Yuan, H.; Wu, S.; Ji, S.; Loy, C. C.; and Yan, S. 2024a. Omg-llava: Bridging image-level, object-level, pixel-level reasoning and understanding. *Advances in neural information processing systems*, 37: 71737–71767.
- Zhang, Y.; Cheng, T.; Zhu, L.; Hu, R.; Liu, L.; Liu, H.; Ran, L.; Chen, X.; Liu, W.; and Wang, X. 2024b. Evf-sam: Early vision-language fusion for text-prompted segment anything model. *arXiv preprint arXiv:2406.20076*.
- Zhang, Z.; Ma, Y.; Zhang, E.; and Bai, X. 2024c. Psalm: Pixelwise segmentation with large multi-modal model. In *European Conference on Computer Vision*, 74–91. Springer.
- Zou, X.; Dou, Z.-Y.; Yang, J.; Gan, Z.; Li, L.; Li, C.; Dai, X.; Behl, H.; Wang, J.; Yuan, L.; et al. 2023a. Generalized decoding for pixel, image, and language. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 15116–15127.
- Zou, X.; Yang, J.; Zhang, H.; Li, F.; Li, L.; Wang, J.; Wang, L.; Gao, J.; and Lee, Y. J. 2023b. Segment everything everywhere all at once. *Advances in neural information processing systems*, 36: 19769–19782.