

# CoReS: Orchestrating the Dance of Reasoning and Segmentation

Xiaoyi Bao<sup>1,2,3,6</sup>, Siyang Sun<sup>3</sup>, Shuailei Ma<sup>4</sup>, Kecheng Zheng<sup>5</sup>, Yuxin Guo<sup>1,2,3</sup>,  
Guosheng Zhao<sup>1,2,6</sup>, Yun Zheng<sup>3</sup>, and Xingang Wang<sup>2,6\*</sup>

<sup>1</sup>School of Artificial Intelligence, University of Chinese Academy of Sciences

<sup>2</sup>Institute of Automation, Chinese Academy of Sciences

<sup>3</sup>Alibaba Group <sup>4</sup>Northeastern University <sup>5</sup>Ant Group

<sup>6</sup>Luoyang Institute for Robot and Intelligent Equipment

{baoxiaoyi2021,guoyuxin2021,zhaoguosheng2021,xingang.wang}@ia.ac.cn,  
{xiaomabufei,zkechengzk}@gmail.com, {siyang.ssy,zhengyun.zy}@alibaba-inc.com

<https://chain-of-reasoning-and-segmentation.github.io>

**Abstract.** The reasoning segmentation task, which demands a nuanced comprehension of intricate queries to accurately pinpoint object regions, is attracting increasing attention. However, Multi-modal Large Language Models (MLLM) often find it difficult to accurately localize the objects described in complex reasoning contexts. We believe that the act of reasoning segmentation should mirror the cognitive stages of human visual search, where each step is a progressive refinement of thought toward the final object. Thus we introduce the Chains of Reasoning and Segmenting (CoReS) and find this top-down visual hierarchy indeed enhances the visual search process. Specifically, we propose a dual-chain structure that generates multi-modal, chain-like outputs to aid the segmentation process. Furthermore, to steer the MLLM’s outputs into this intended hierarchy, we incorporate in-context inputs as guidance. Extensive experiments demonstrate the superior performance of our CoReS, which surpasses the state-of-the-art method by 6.5% on the ReasonSeg dataset.

**Keywords:** Reasoning Segmentation · Multi-Modal · Chain-of-Thought

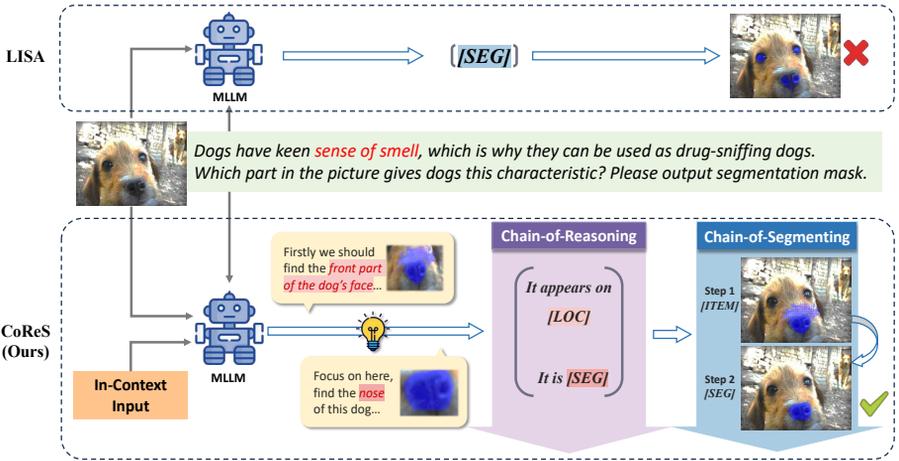
## 1 Introduction

Recently, multi-modal large language models (MLLM) have gained increasing recognition for their powerful capabilities in various tasks. By utilizing the knowledge repository contained in MLLM, humans can demand more complex multi-modal tasks than ever before. Among them, inference-based segmentation tasks can achieve more intelligent fine-grained understanding by combining traditional visual tasks with the reasoning process.

Currently, there are mainly two ways to handle this task: one is to equip MLLM with a segmentation decoder, and the other is to use LLM to output

---

\* Corresponding author.



**Fig. 1:** Comparison between our CoReS and LISA. **UP:** the process of LISA, **DOWN:** the diagram of CoReS. Given textual and visual inputs, LISA directly uses the `[SEG]` token output by MLLM to generate a mask. On the contrary, our CoReS involves breaking down the task of “finding the part that gives dogs keen sense of smell” into a logical chain such as “first find the front part of the dog’s face, then focus on this specific area, search for the nose of the dog.” It can be observed that LISA incorrectly segments the dog’s eyes, which are similarly round, dark, and important in sensory perception. In contrast, through in-context input and dual-chain structure, CoReS achieves the segmentation of the nose of the dog correctly.

the mask in text form directly. For example, LISA [19] equips SAM [15] decoder for LLaVA [26] and constructs an inference segment database for training. VistaLLM [37] directly uses LLM to generate text-formatted segmentation masks and designs an adaptive sampling algorithm to optimize the masks.

The existing MLLM can effectively segment objects at the object level but struggles to differentiate objects referred to in reasoning text accurately. As shown in Fig.1, when segmenting "part gives dogs sense of smell," the reasoning process directly searches for nose-like items that are round, dark, and important in sensory perception. Since the eyes of the dog have similar characteristics, LISA erroneously segments them as the object. The semantic similarity of such instances presents a substantial challenge to MLLM’s ability to accurately localize and segment complex reasoning tasks.

How can we complete complex reasoning segmentation tasks? From the way humans handle similar tasks, we gain inspiration. The brain realizes object localization and searching typically in a top-down, targeted manner, where each step is a progressive refinement of thought toward the final object. Specifically, it is guided by pre-existing knowledge including a broad understanding of the typical positions of objects in the scene. Then a detailed analysis and synthesis of the specific item is followed to gradually approach the precise segmentation of objects. For example, when asked to find “the wedding token usually exchanged

by the groom and the bride” in a picture, people draw upon their knowledge to presuppose that we should focus on the fingers where rings are usually worn. Therefore, people first identify the typical location of the hand, and then narrow the field of view to locate and segment the ring in the local area of the hand.

This top-down hierarchical structure can effectively help improve segmentation. Considering the disparities between modalities, how can we build such a visual logical hierarchy for multi-modal reasoning segmentation? Drawing inspiration from LLM’s chain of thought, we propose the Chain of Reasoning and Segmentation (CoReS), a multi-modal chain of thought for fine-grained tasks, to establish such a unified visual hierarchy. As is shown in Fig.1, it is a dual-chain structure aimed at decoupling the hierarchical thinking process. Specifically, the reasoning chain is reflected in the output of MLLM, injecting semantic information for different logical levels into different tokens in the chain. The segmentation chain utilizes the logic of the reasoning chain to iteratively optimize the segmentation results.

Since MLLM cannot actively extract outputs conforming to this top-down hierarchy without guidance, we propose adding extra in-context inputs for MLLM. This involves providing randomly sampled textual examples that indicate the desired chain-like rules of output. Although unrelated to the user’s query input, these text-formed question-answer pairs contain the anticipated chain-like output rules. MLLM reads and transfers these rules to the output of multi-modal tasks, implicitly guiding the generation of the reasoning chain.

Our contributions can be summarized as follows:

- We propose CoReS, a multi-modal chain of thought. It provides a more accurate visual search for multi-modal fine-grained tasks through a top-down chain-like visual hierarchy.
- To form a chain-like multi-modal process, we propose a dual-chain structure, integrating both modalities in the same visual hierarchy. Moreover, to lead MLLM to generate such a hierarchy actively, we provide guidance through in-context input. With these randomly sampled text-based question-answer pairs, MLLM learns the chain-like rule and transfers it to multi-modal tasks.
- Extensive experiments demonstrate the effectiveness of our proposed method, which achieves state-of-the-art accuracy on ReasonSeg benchmarks.

## 2 Related Work

### 2.1 Interactive Segmentation

Traditional segmentation tasks use single images as input, requiring the segmentation of objects belonging to a set of predefined categories [2, 4, 6, 10–12, 20, 28, 31, 34, 39, 40, 42, 43, 50, 53, 56–58, 60]. These tasks are straightforward, with simple images, and distinct objects. Existing methods have already achieved impressive results in this area [5, 14, 23, 48]. To meet the demands of more complex task requirements, the field of segmentation research is gradually evolving from purely visual to multi-modal approaches.

One extension is the open-vocabulary segmentation task, which involves segmenting objects specified by class names that do not belong to a fixed set of categories [49]. Such tasks often employ visual language models like CLIP as a bridge between modalities [9, 24]. To facilitate more convenient interaction, the referring segmentation task [13, 33] enables interaction with human language rather than the semantic category list, aiming to leverage explicit text description to segment the target object.

To expand the diversity of interaction methods, existing works [16, 61, 62] design a variety of prompting methods including points, boxes, scribbles, noise masks, and so on. Kirillov et al [16] introduced SAM, trained with billions of high-quality masks, supporting bounding boxes and points as prompts while demonstrating exceptional segmentation quality. X-Decoder [61] and SEEM [62] further support various human interaction methods. Recently, Lai et al proposes LISA [19], simply integrated the MLLM with SAM [16] to tackle the reasoning segmentation task and enhance existing visual segments with self-reasoning abilities. Although LISA inherits the excellent abilities of large language models in text reasoning, its performance in segmentation, especially for hard-to-perceive part objects, is not satisfactory.

## 2.2 Multi-modal Large Language Model

The multi-modal large language models integrate the LLM [3, 7, 44] and vision encoder [38, 41] to transfer the reasoning ability and huge world knowledge for the vision tasks. Flamingo [1] proposes a cross-attention structure to integrate visual information into the nlp contexts, enabling visual in-context learning. Several works (such as BLIP-2 [22], mPLUG-OWL [52] Otter [21], LLaVA [26], MiniGPT-4 [59], FROMAGe [17] and so on) leverage the visual encoder and projection module (such as linear projection, Qformer, Casual Qformer and so on) to translate the visual features and directly feed them into the LLM along with the NLP token embeddings.

Recently, the use of MLLMs is no longer limited to understanding tasks, and there has been an attempt to utilize them for fine-grained visual tasks, such as grounding and segmentation. Kosmos-2 [35] constructs large-scale data of grounded image-text pairs, attempting to infuse grounding capabilities into LLMs. DetGPT [36] bridges the fixed MLLM and open-vocabulary detector. GPT4RoI [54] introduces RoI visual features as input and trains the model on region-text pairs. LISA [19] integrates SAM [15] as a segmentation decoder into MLLMs and proposes the reasoning segmentation task. VistaLLM [37] proposes a unified framework for MLLMs with single and multiple visual scene inputs and introduces an adaptive sampling algorithm to refine the NLP format mask of the MLLM’s output. In this paper, we design a novel chain-of-thought manner to enable MLLM to realize fine-grained reasoning tasks.

## 2.3 Chain of Thoughts

CoT, as a convenient and efficient method for enhancing complex reasoning in large language models (LLMs) [46], has significantly improved their performance in generating rationales and inferring accurate answers in numerous domains, including commonsense and arithmetic. Existing CoT prompting for LLMs is primarily used in inference and can be categorized into two major paradigms: Zero-shot-CoT and Few-shot-CoT. Zero-shot-CoT [18] directly leverages a single prompt like “Let’s think step by step” to generate reasoning chains. Few-shot-CoT [46] uses reasoning demonstrations one by one as the prompt information.

For MLLMs, due to limitations in model size and performance, directly employing these two kinds of CoT in inference does not lead to an effective enhancement. Therefore, multi-modal CoT [55] typically acquires this ability by fine-tuning the model on the constructed multi-modal CoT dataset. However, leveraging vision information effectively and fusing visual features with text representation in a multi-modal chain of thought (CoT) poses a significant challenge. Prior work [29] attempts to use image captions and incorporate them after text input, but this approach results in substantial information loss of images. To avoid the complexity of constructing extra multi-modal CoT dataset, KAM-CoT [32] proposes a two-stage training process with knowledge graphs grounding to generate effective rationales and answers for reducing the computational cost and substantial hardware resources. Utilizing the concept of a thought chain,  $V^*$  [47] combines LLM and MLLM to construct a visual search algorithm that surpasses the performance of GPT-4V. As far as we know, the multi-modal application of CoT has been confined to visual understanding tasks like visual question answering and has not extended to dense prediction tasks.

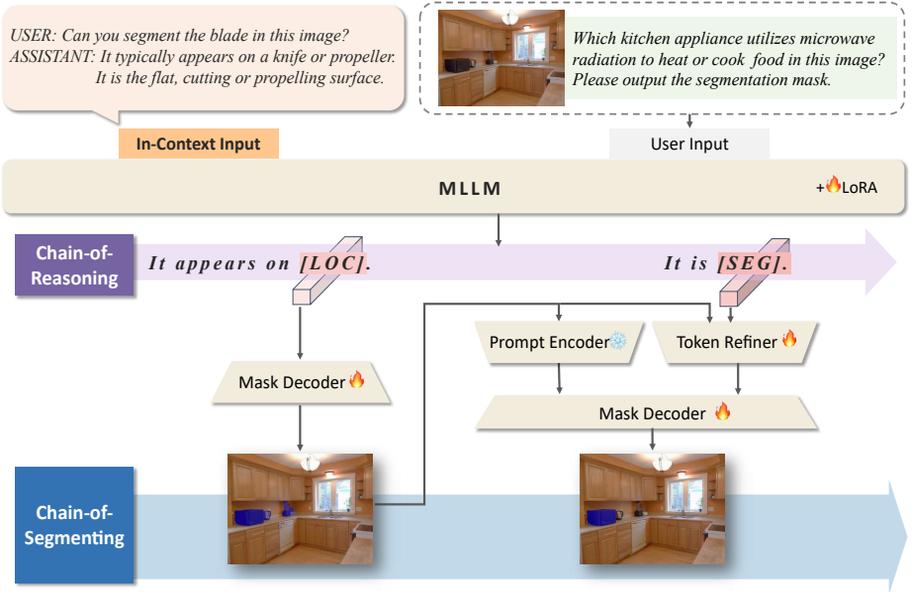
## 3 Method

In line with the process of human visual search, a top-down visual hierarchy can aid the MLLM in progressively pinpointing objects referred to in reasoning texts. Inspired by this concept, we introduce CoReS, which applies the idea of the chain of thought from natural language processing to the execution of fine-grained tasks in a multi-modal context. As illustrated in Fig.2, the core of CoReS lies in its dual-modal, dual-chain structure, which consists of a chain-of-reasoning and a chain-of-segmenting. In addition to this, the extra in-context guidance plays a critical role in the formation of a chain-like hierarchy. Detailed explanations of these two components are provided in Sec.3.1 and Sec.3.2.

### 3.1 Dual-Modal Dual-Chain

**Chain-of-Reasoning:** To achieve consistency between the modalities in the multi-modal chain of thought, the output of MLLM needs to conform to the top-down semantic logic required for fine-grained visual tasks.

To this end, we use the kind of response templates like “It appears on [LOC]. It is [SEG].” By constraining the sentence structure of the MLLM



**Fig. 2:** Overall architecture of CoReS. The input of MLLM consists of the user input in gray and the extra in-context input in orange, which consists of question-answer examples unrelated to the user query. MLLM generates output at the logical level of chain-of-reasoning, where the token embeddings of [LOC] and [SEG] serve as prompt inputs for different positions of the segmentation chain, guiding the chain to generate segmentation results progressively. We exclude the visual feature input to the mask decoder extracted by the extra vision backbone for conciseness here.

outputs, implicit constraints are imposed on the logical chain output  $H$  by MLLM. Given the image  $Q_{img}$  and query input  $Q_{text}$ , this template of output compels MLLM to incorporate different information at different positions of the token, thus forming a semantic-level chain of reasoning. Specifically, MLLM injects the information about the scene or item that the object most likely to appear on into the [LOC] token while the information about the object itself is injected into the [SEG] token of the following sentence.

We fine-tune MLLM using LoRA and employ cross-entropy loss for the supervision of the chain of reasoning. The specific formula of  $\mathcal{L}_{CoR}$  is as follows.

$$\mathcal{L}_{CoR} = \mathcal{L}_{CE}(\mathbf{p}(H|Q_{img}, Q_{text}), \mathbf{t}). \quad (1)$$

$\mathbf{p}(H|Q_{img}, Q_{text})$  and  $\mathbf{t}$  refer to output logits of MLLM and the token embedding of our chain-like template, respectively.  $\mathcal{L}_{CE}$  is the cross-entropy loss.

**Chain-of-Segmenting:** Upon obtaining text-form outputs that align with the visual hierarchy, the next task is to apply them as guidance for the visual modal. Thus we propose the Chain-of-Segmenting, iteratively generating segmentation

results in a unified hierarchical manner. Obtaining mask annotations for intermediate logical layers is impractical due to the logic subjectivity, so MLLM with world knowledge is used for a unified unsupervised representation. The scene-level information injection of C1 is analyzed in the appendix.

During the hierarchical generation of segmentation, the segmentation result from the previous level  $m^{t-1}$  serves as a cue. Since the [LOC] and  $m^{t-1}$  in the first logical layer focus on common locations and do not entirely represent inclusion relationships, we choose its logical form to serve as a soft mask prompt, processed by the prompt encoder  $\theta$  of the SAM framework.

$$\mathcal{M}^t = \theta(m^{t-1}), t \geq 1. \quad (2)$$

$\mathcal{M}^t$  is the output of prompt encoder from the  $t$ -th logical level.

The implementation of text guidance follows the embedding-based approach in LISA. From the last layer feature output by MLLM, the embedding ( $\mathbf{h}^0$ ,  $\mathbf{h}^1 \in \mathbb{R}^{1 \times 256}$ ) are extracted at positions of [LOC] and [SEG]. They serve as a prompt in the text modality to guide the corresponding segmentation process.

To further enhance the inter-modal connections in a top-down hierarchy, we propose to refine the token embedding of the reasoning chain with the help of the visual output  $m^{t-1}$  from the previous logical level. Specifically, we use masked-average-pooling (MAP) to create a prototype. This prototype corresponds to the visual features activated at the current stage. We then use this prototype to calibrate text tokens at the next logical level. This process allows for iterative optimization of tokens generated in a single forward pass of the reasoning chain, aligned with the hierarchy.

$$\hat{\mathbf{h}}^t = \mathcal{R}(\beta(\mathbf{h}^t), F_v(Q_{img}), m^{t-1}), t \geq 1. \quad (3)$$

$\hat{\mathbf{h}}^t$  is the refined token embedding from the  $t$ -th logical level while  $\mathbf{h}^t$  is the unrefined one in the chain-of-reasoning.  $\beta$  is a multilayer perceptron (MLP) to project output features. We employ the SAM image encoder  $F_v$  as the visual backbone.  $\mathcal{R}$  refers to the token refiner whose working function can be expressed by the following formula.

$$\mathcal{R}(h, i, m) = h + CA(h, MAP(i, m)). \quad (4)$$

$CA$  and  $MAP$  refer to cross-attention and masked average pooling, respectively.

Subsequently, the extracted features of the query image  $Q_{img}$  are passed to the mask generator  $\gamma$  as input. For brevity, the aforementioned process is omitted in Fig.2. In summary, the generation process of the segmentation chain can be expressed as follows.

$$m^t = \gamma(F_v(Q_{img}), \mathcal{M}^t, \hat{\mathbf{h}}^t). \quad (5)$$

where  $m^t$  is the  $t$ -th level output in the segmenting chain. In the dual-level setting of the segmentation chain,  $t = 0/1$ .

Parameters in the projection layer and mask decoder are set trainable, while the SAM image encoder and prompt encoder are kept frozen to keep the generalization ability. For the segmentation chain, we only provide supervision for

*A chat between a curious human and an artificial intelligence assistant. The assistant gives helpful, detailed, and polite answers to the human's questions. To realize segmentation, the assistant guesses all possible object it typically appears on and its distinct appearance.*

*USER: What is the dog paw in this image? Please output segmentation mask.*

*ASSISTANT: Sure, it appears on the end of each leg. It is tough and cushioned with pads and claws.*

*USER:  What part of a car is essential for motion and supports the vehicle's weight in this image? Please output the segmentation mask.*



**Fig. 3:** Examples of MLLM input. The brown section represents the system prompt commonly used for MLLM. The green section is the in-context input proposed in this paper, and the black section is the user input for questioning the image.

the final mask generated at the last level of the chain. We choose dice loss and cross-entropy loss as the loss function, which can be represented as:

$$\mathcal{L}_{CoS} = \lambda_d \mathcal{L}_{DICE}(m^T, M_{gt}) + \lambda_c \mathcal{L}_{CE}(m^T, M_{gt}). \quad (6)$$

where  $m^T$  and  $M_{gt}$  are the predicted final mask and the ground truth, respectively. In the dual-level setting of the segmentation chain,  $T = 1$ . The analysis for the case when  $T > 1$  can be found in Sec.4.4. The weights  $\lambda_d$  and  $\lambda_c$  are set to 0.5 and 2.0 empirically.

As the following equation shows, the final loss is a weighted sum of the textual loss from the chain-of-reasoning  $\mathcal{L}_{CoR}$  and the visual loss  $\mathcal{L}_{CoS}$  from the chain-of-segmenting.

$$\mathcal{L}_{total} = \lambda_R \mathcal{L}_{CoR} + \lambda_S \mathcal{L}_{CoS}. \quad (7)$$

where the weights  $\lambda_R$  and  $\lambda_S$  are set to 0.5 empirically. The impact of these weights can be found in the Appendix.

### 3.2 In-Context Guidance

The described hierarchical dual-chain structure progressively identifies the objects referred to in the reasoning text. However, if only sentence templates are used as the supervisory signal, the MLLM cannot actively uncover and output hierarchical logical relationships.

We thus propose a training paradigm involving in-context guidance to provide logical cues to the output of MLLM. To present the required top-down output logic rules to the MLLM, we introduce some contextually provided examples as guidance, before the initial user input of image-query pairs, as depicted in Fig.3.

Specifically, we pre-construct a context library in a pure text format. This process is to induce ChatGPT to generate more question-answer pairs with several manually written examples. For the question part of the examples, we utilize the same questioning format. For the sake of conciseness, CoReS does not

include corresponding images and directly poses questions based on the name of the segmented object. For the answer format, no special tokens are used in the examples; instead, normal replies are given based on the logical aspects of where the questioning object usually appears and its common features.

During each forward pass in CoReS, we randomly draw examples from the context library to serve as in-context input. The reason we opt not to generate query-specific textual prompts lies in the high computational and parameter demands. The two-stage method requires forward passes through two different MLLMs, which is quite inefficient.

Examples extracted in this manner are unrelated to the current query image but implicitly embed the logic rules for output. As a rule provider, this in-context input serves as a prompt for MLLM to inject specific information into specific token positions. For instance, the token at the [LOC] position should be injected with location-related information such as “a knife or propeller”, while the token at the [SEG] position should be injected with information regarding “the flat, cutting or propelling surface” of the object. In this way, MLLM extracts the top-down rules from pure textual context and transfers it to the output for new instances referred to by multi-modal queries.

Unlike in-context learning traditionally used solely to enhance the inference performance, our in-context question-answer pairs also play a role in the training process. Through this contextual prompting, CoReS teaches the model what semantic logic to follow in generating answers during the tuning phase.

## 4 Experiments

### 4.1 Implementation Details

**For training datasets**, we follow LISA to utilize a combination of datasets for semantic, referring, and reasoning segmentation while abandoning the use of the visual-question-answering dataset. Related reasons and details can be found in the appendix. **For the test dataset**, we conduct experiments on the ReasonSeg, which comprises image-question pairs with reasoning difficulty and the ground truth segmentation masks. **For evaluation metrics**, we use graph-averaged Intersection-over-Union (gIoU), which is the most commonly used metric in referring segmentation. The cumulative Intersection-over-Union (cIoU) is used as an auxiliary indicator, despite its bias to large-area targets of images. We conduct experiments in the Appendix to clarify this bias.

The multi-modal LLM, unless specifically indicated otherwise, refers to the LLaVA-7B-v0. Moreover, we use SAM-ViT-H as the image backbone. Other specific implementation details are presented in the Appendix.

### 4.2 Comparison with State-of-the-Arts

We compare the performance of CoReS with other methods on the ReasonSeg dataset, as shown in Tab.1. Our approach improves by approximately 35%

**Table 1:** Performance comparison between CoReS and other methods on ReasonSeg. “ft” denotes using the reasoning segmentation dataset to fine-tune the model.

Method	Text Model	ReasonSeg val overall	ReasonSeg test		
			short	long	overall
OVSeg [24]	CLIP ViT-L	28.5	18.0	28.7	26.1
GRES [25]	BERT	22.4	17.6	22.6	21.3
X-Decoder [61]	UniCL	22.6	20.4	22.2	21.7
SEEM [62]	UniCL	25.5	20.1	25.6	24.3
LISA [19]	LLaVA-7B	44.4	37.6	36.6	36.8
CoReS		54.8	41.0	50.9	48.7
LISA(ft)		52.9	40.6	49.4	47.3
CoReS(ft)		<b>59.4</b>	<b>44.2</b>	<b>55.0</b>	<b>52.4</b>
LISA(ft)	LLaVA-13B	56.2	44.3	54.0	51.7
CoReS(ft)		<b>61.8</b>	<b>49.7</b>	<b>58.3</b>	<b>55.9</b>
LISA(ft)	LLaVA-v1.5-13B	65.0	55.4	63.2	61.3
CoReS(ft)		<b>68.1</b>	<b>57.9</b>	<b>66.4</b>	<b>65.5</b>

compared to the multi-modal grounding models that do not use LLM. Furthermore, we also compare the performance of LISA, a baseline method that similarly incorporates world knowledge by MLLM. This comparison with CoReS is made using MLLM of different scales. Without fine-tuning the ReasonSeg dataset, our method outperforms LISA by 10.4%, and this margin remains 6.5% after fine-tuning. It is worth noting that the un-tuned version of CoReS outperforms the fine-tuned version of LISA by 1.9%. This demonstrates that the performance improvement brought by CoReS does not rely on provided training data but fundamentally leverages the potential of the MLLM more efficiently. When using LLaVA-13B and LLaVA-v1.5-13B as MLLMs, CoReS also outperforms LISA by a considerable margin. Achieving the same level of improvement becomes more challenging on a higher-performing base than LISA, thus the slight decrease in the magnitude of improvement is reasonable. Overall, these significant performance improvements prove the efficacy of CoReS.

Qualitative results also prove the effectiveness of CoReS in Fig.4, Fig.5, and the appendix. As shown in Fig.4, when facing complex reasoning problems, LISA exhibits errors in grounding the exact instance referred by the reasoning query, while CoReS achieves the correct answers. Fig.5 provides a more detailed demonstration of the effects of the multi-modal chain of thought. Whether it is complex reasoning, such as “moth evading predators,” or segmentation difficulties caused by underwater reflections, CoReS gives the right mask. It is the chain-like hierarchy that results in correct and fine-grained segmentation, which cannot be achieved by LISA.

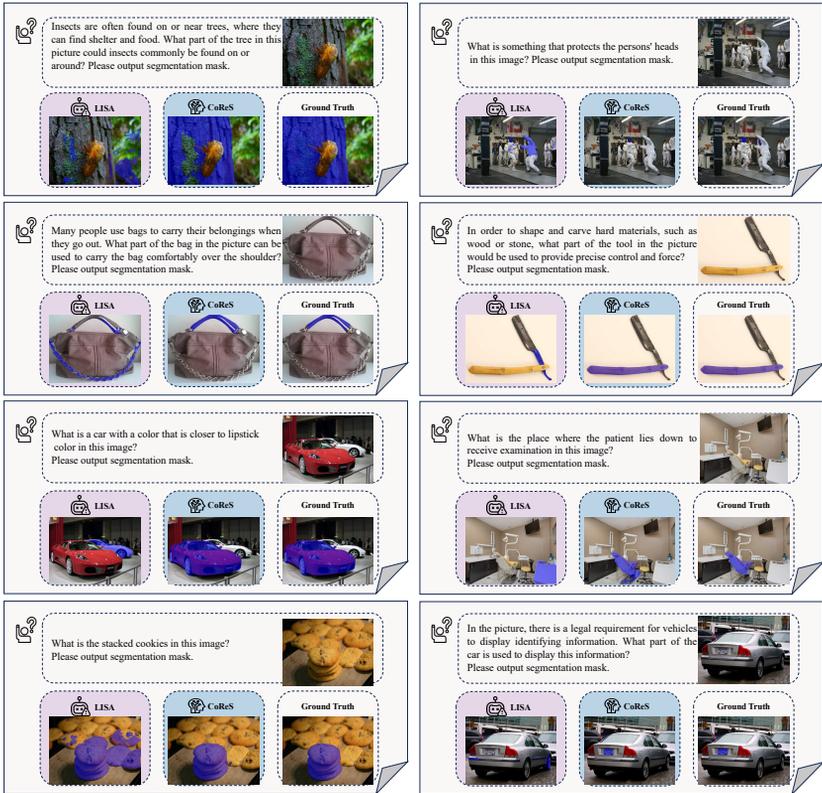
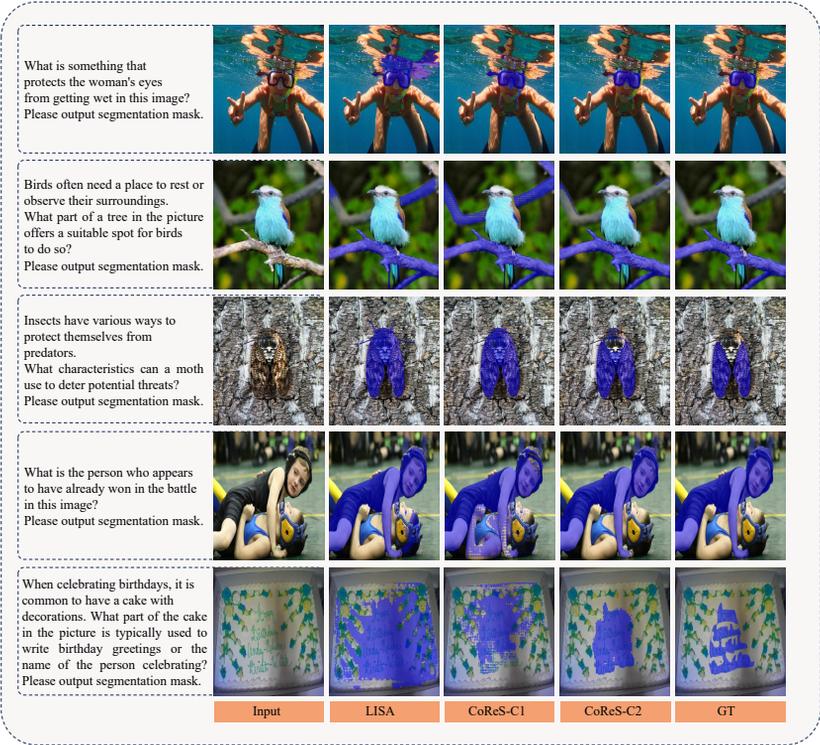


Fig. 4: Visual comparison of CoReS and LISA.

### 4.3 Comparison on other benchmarks

We also evaluate the performance of CoReS on the referring segmentation task, as shown in Tab.2. As a degraded form of reasoning segmentation, the results of the referring segmentation are compared on refCOCO, refCOCO+, and refCOCog, and its performance is measured using cIoU. For general models, we compare with X-Decoder and SEEM. For specialized models, we compare CoReS with methods like PolyFormer. Without using their REC pretraining and mass data from tasks like object detection and caption for targeted training, CoReS still achieves competitive results, as shown in Tab.2. However, these methods could not handle the ReasonSeg tasks that CoReS is designed for.

Compared to LISA, CoReS does not exhibit a significant improvement on refCOCO and refCOCO+, whose average query lengths are 3.6 and 3.5, respectively. However, on refCOCog with an average query length of 8.4, our method outperforms LISA by about 2%. We attribute this to the fact that CoReS’s multi-modal chain of thought primarily enhances the performance under complex and challenging queries. Therefore, the improvement is less significant for simpler referential objects in the former two datasets, but more pronounced



**Fig. 5:** Qualitative interpretation of the advantages of the multi-modal chain-of-thought over LISA. From left to right are the input image, LISA result, CoReS first logic layer segmentation result, CoReS final result, and ground truth mask.

**Table 2:** Performance comparison between CoReS and other methods on referring segmentation datasets. The LISA and CoReS here refer to their LLaVA-7B-v0 version. We use cIoU as the metric here. ‘S’ and ‘G’ refer to general/specialized types of models. ‘RS’ refers to the ability to handle reasoning segmentation tasks.

Method	T	RS	refCOCO			refCOCO+			refCOCOg	
			val	testA	testB	val	testA	testB	val(U)	testU
MCN [30]	S		62.4	64.2	59.7	50.6	55.0	44.7	49.2	49.4
VLT [8]	S		67.5	70.5	65.2	56.3	61.0	50.1	55.0	57.7
CRIS [45]	S		70.5	73.2	66.1	62.3	68.1	53.7	59.9	60.4
LAVT [51]	S		72.7	75.8	68.8	62.1	68.4	55.1	61.2	62.1
ReLA [25]	S		73.8	76.5	70.2	66.0	71.0	57.7	65.0	66.0
PolyFormer [27]	S		74.8	76.6	71.1	<b>67.6</b>	<b>72.9</b>	<b>59.3</b>	67.8	69.1
X-Decoder [61]	G		-	-	-	-	-	-	64.6	-
SEEM [62]	G		-	-	-	-	-	-	65.7	-
LISA [19]	G	✓	74.9	<b>79.1</b>	72.3	65.1	70.8	58.1	67.9	70.6
CoReS	G	✓	<b>76.0</b>	78.6	<b>72.5</b>	65.1	70.0	58.6	<b>69.0</b>	<b>70.7</b>

for datasets involving more complex object references such as refCOCOg. This

**Table 3:** Ablation studies on the key components of CoReS. “InC” and “CoR” refer to the in-context input and the chain-of-reasoning. “CoS” and “CoS-R” denote the chain-of-segmenting without and with the token refinement, respectively.

InC	Components			ReasonSeg	
	CoR	CoS	CoS-R	gIoU	cIoU
				52.9	54.0
✓				53.2	55.4
	✓			53.9	55.6
		✓		54.0	54.1
✓	✓			55.3	56.9
✓		✓		55.4	58.0
	✓	✓		56.9	59.8
	✓		✓	57.5	60.4
✓	✓	✓		58.4	59.3
✓	✓		✓	<b>59.4</b>	<b>62.1</b>

demonstrates that the proposed multi-modal chain of thought can indeed reduce the difficulty of finding the referred instance by complex reasoning query, proving the effectiveness of CoReS. We include performance results on other benchmarks in the appendix.

#### 4.4 Ablation Studies

**Different design choices of CoReS.** To demonstrate the effectiveness of our proposed in-context input and the dual-chain-of-thought structure, we conduct ablation experiments on the model design details mentioned for the CoReS in Tab.3 and different design alternatives in the Appendix.

In the case of using chain-of-segmenting without chain-of-reasoning, we use the same [SEG] embedding as the prompt input for each level, resulting in only a 1.1% improvement in gIoU. This demonstrates that the improvement brought by our method is not due to an increase in computational power. When using chain-of-reasoning without the segmentation chain, to avoid information loss, we take the average of the two token embeddings as the input for the mask decoder, resulting in a 1.0% boost, demonstrating that the decomposition of thinking levels can indeed reduce the difficulty and improve performance.

The CoS guided by CoR leads to a 4.6% performance increase, proving the effective interaction and guidance between the two modalities through chain thinking. By adding in-context input to the dual-chain structure, CoReS can continue to gain a 1.9% improvement. This demonstrates that the strategy of in-context training can provide prompts and guidance for the dual-chain structure. Consistent with Sec.3, this proves that our practices are reasonable and reliable.

**Different numbers of the in-context input.** We conduct experiments with different numbers of the in-context input to further explore their effect. As shown in Tab.4, adding one in-context example can lead to a 1.9% improvement.

However, as the number of examples continues to increase, performance improvement is slightly decreased. This may be due to the limited quality of the context library, which may not be conducive to the performance. In general, regardless of the number of in-context examples added, compared to not using in-context input, it always leads to improved results. This proves the effectiveness of the logical guidance of in-context inputs.

**Different logical levels.** We also look into the effect of different logical levels of the dual chain of thought. In Sec.3 and Tab.1-Tab.4, CoReS selects a logical depth of 2, resulting in a 5.8% improvement in gIoU, where both the in-context input and the dual-modal chain adhere to a two-level semantic structure from [LOC] to [SEG].

**Table 4:** Ablation studies of different numbers of the in-context input.

num	ReasonSeg	
	gIoU	cIoU
0	57.5	60.4
1	<b>59.4</b>	<b>62.1</b>
2	59.4	61.9
4	59.3	62.0

**Table 5:** Ablation studies of different logical levels. [S], [P], and [L] refer to the [SEG], [POS], and [LOS] token.

token	ReasonSeg	
	gIoU	cIoU
[S]	53.2	55.4
[L] + [S]	59.4	<b>62.1</b>
[L] + [P] + [S]	<b>59.7</b>	60.8

Similarly, we test the performance of CoReS with a three-level semantic hierarchy. For the corresponding levels, we heuristically select “commonly occurring categories”, “specific locations in the image” and “intrinsic features”. The template of the chain-of-reasoning, the mask proposal process, and the in-context inputs are also adjusted accordingly to the three levels. The performance does not increase much. We deem the reason to be that questions in ReasonSeg are not so particularly complex that require three levels of thinking and grounding. Furthermore, the specific content of each level of the three-level logic is heuristically, and its scientificity and reasonableness need to be considered. Besides, the longer logic chains itself is associated with increased backpropagation difficulty and slower convergence speed. However, overall, multi-level logic does bring a significant improvement in performance compared to single-plane logic, which is all consistent with the theoretical analysis in Sec.3.

## 5 Conclusion

We propose CoReS, a dual-modal chain-of-thought framework for fine-grained reasoning tasks. Adopting a top-down logic for visual search, we propose a dual-chain hierarchical structure that assists the MLLM in accurately localizing objects referred to in reasoning texts. Additionally, the integration of in-context inputs enables the MLLM to achieve rule transfer, guiding the generation of chain-like outputs in multi-modal tasks. Experimental results confirm that our approach yields a substantial improvement over state-of-the-art methods in reasoning segmentation tasks. We hope that CoReS will inspire future research on a broader range of complex multi-modal tasks.

## References

1. Alayrac, J.B., Donahue, J., Luc, P., Miech, A., Barr, I., Hasson, Y., Lenc, K., Mensch, A., Millican, K., Reynolds, M., et al.: Flamingo: a visual language model for few-shot learning. *NeurIPS* (2022) [4](#)
2. Badrinarayanan, V., Kendall, A., Cipolla, R.: Segnet: A deep convolutional encoder-decoder architecture for image segmentation. *TPAMI* (2017) [3](#)
3. Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al.: Language models are few-shot learners. *Advances in neural information processing systems* **33**, 1877–1901 (2020) [4](#)
4. Chen, L., Papandreou, G., Kokkinos, I., Murphy, K., Yuille, A.L.: Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *TPAMI* (2018) [3](#)
5. Cheng, B., Collins, M.D., Zhu, Y., Liu, T., Huang, T.S., Adam, H., Chen, L.C.: Panoptic-deeplab: A simple, strong, and fast baseline for bottom-up panoptic segmentation. In: *CVPR* (2020) [3](#)
6. Cheng, B., Schwing, A., Kirillov, A.: Per-pixel classification is not all you need for semantic segmentation. *NeurIPS* (2021) [3](#)
7. Chowdhery, A., Narang, S., Devlin, J., Bosma, M., Mishra, G., Roberts, A., Barham, P., Chung, H.W., Sutton, C., Gehrmann, S., et al.: Palm: Scaling language modeling with pathways. *Journal of Machine Learning Research* **24**(240), 1–113 (2023) [4](#)
8. Ding, H., Liu, C., Wang, S., Jiang, X.: Vlt: Vision-language transformer and query generation for referring segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2022) [12](#)
9. Ding, Z., Wang, J., Tu, Z.: Open-vocabulary panoptic segmentation with maskclip. *arXiv preprint arXiv:2208.08984* (2022) [4](#)
10. Fu, J., Liu, J., Tian, H., Li, Y., Bao, Y., Fang, Z., Lu, H.: Dual attention network for scene segmentation. In: *CVPR* (2019) [3](#)
11. Guo, Y., Ma, S., Su, H., Wang, Z., Zhao, Y., Zou, W., Sun, S., Zheng, Y.: Dual mean-teacher: An unbiased semi-supervised framework for audio-visual source localization. *Advances in Neural Information Processing Systems* **36** (2024) [3](#)
12. Huang, Z., Wang, X., Huang, L., Huang, C., Wei, Y., Liu, W.: Ccnet: Criss-cross attention for semantic segmentation. In: *ICCV* (2019) [3](#)
13. Kazemzadeh, S., Ordonez, V., Matten, M., Berg, T.: Referitgame: Referring to objects in photographs of natural scenes. In: *EMNLP* (2014) [4](#)
14. Kirillov, A., He, K., Girshick, R., Rother, C., Dollár, P.: Panoptic segmentation. In: *CVPR* (2019) [3](#)
15. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., et al.: Segment anything. *arXiv preprint arXiv:2304.02643* (2023) [2](#), [4](#)
16. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., et al.: Segment anything. *arXiv:2304.02643* (2023) [4](#)
17. Koh, J.Y., Salakhutdinov, R., Fried, D.: Grounding language models to images for multimodal generation. *arXiv preprint arXiv:2301.13823* **1**(2) (2023) [4](#)
18. Kojima, T., Gu, S.S., Reid, M., Matsuo, Y., Iwasawa, Y.: Large language models are zero-shot reasoners. *CoRR* **abs/2205.11916** (2022). <https://doi.org/10.48550/arXiv.2205.11916>, <https://doi.org/10.48550/arXiv.2205.11916> [5](#)

19. Lai, X., Tian, Z., Chen, Y., Li, Y., Yuan, Y., Liu, S., Jia, J.: Lisa: Reasoning segmentation via large language model. arXiv preprint arXiv:2308.00692 (2023) [2](#), [4](#), [10](#), [12](#)
20. Lai, X., Tian, Z., Jiang, L., Liu, S., Zhao, H., Wang, L., Jia, J.: Semi-supervised semantic segmentation with directional context-aware consistency. In: CVPR (2021) [3](#)
21. Li, B., Zhang, Y., Chen, L., Wang, J., Yang, J., Liu, Z.: Otter: A multi-modal model with in-context instruction tuning. arXiv:2305.03726 (2023) [4](#)
22. Li, J., Li, D., Savarese, S., Hoi, S.: Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. arXiv:2301.12597 (2023) [4](#)
23. Li, Y., Zhao, H., Qi, X., Wang, L., Li, Z., Sun, J., Jia, J.: Fully convolutional networks for panoptic segmentation. In: CVPR (2021) [3](#)
24. Liang, F., Wu, B., Dai, X., Li, K., Zhao, Y., Zhang, H., Zhang, P., Vajda, P., Marculescu, D.: Open-vocabulary semantic segmentation with mask-adapted clip. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7061–7070 (2023) [4](#), [10](#)
25. Liu, C., Ding, H., Jiang, X.: Gres: Generalized referring expression segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 23592–23601 (2023) [10](#), [12](#)
26. Liu, H., Li, C., Wu, Q., Lee, Y.J.: Visual instruction tuning. arXiv:2304.08485 (2023) [2](#), [4](#)
27. Liu, J., Ding, H., Cai, Z., Zhang, Y., Satzoda, R.K., Mahadevan, V., Manmatha, R.: Polyformer: Referring image segmentation as sequential polygon generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18653–18663 (2023) [12](#)
28. Liu, W., Rabinovich, A., Berg, A.C.: Parsenet: Looking wider to see better. arXiv (2015) [3](#)
29. Lu, P., Mishra, S., Xia, T., Qiu, L., Chang, K., Zhu, S., Tafjord, O., Clark, P., Kalyan, A.: Learn to explain: Multimodal reasoning via thought chains for science question answering. CoRR **abs/2209.09513** (2022). <https://doi.org/10.48550/arXiv.2209.09513>, <https://doi.org/10.48550/arXiv.2209.09513> [5](#)
30. Luo, G., Zhou, Y., Sun, X., Cao, L., Wu, C., Deng, C., Ji, R.: Multi-task collaborative network for joint referring expression comprehension and segmentation. In: Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition. pp. 10034–10043 (2020) [12](#)
31. Ma, S., Wang, Y., Wang, S., Wei, Y.: Fgahoi: Fine-grained anchors for human-object interaction detection. IEEE Transactions on Pattern Analysis and Machine Intelligence pp. 1–16 (2023). <https://doi.org/10.1109/TPAMI.2023.3331738> [3](#)
32. Mondal, D., Modi, S., Panda, S., Singh, R., Rao, G.S.: Kam-cot: Knowledge augmented multimodal chain-of-thoughts reasoning. arXiv preprint arXiv:2401.12863 (2024) [5](#)
33. Nagaraja, V.K., Morariu, V.I., Davis, L.S.: Modeling context between objects for referring expression understanding. In: ECCV (2016) [4](#)
34. Noh, H., Hong, S., Han, B.: Learning deconvolution network for semantic segmentation. In: ICCV (2015) [3](#)
35. Peng, Z., Wang, W., Dong, L., Hao, Y., Huang, S., Ma, S., Wei, F.: Kosmos-2: Grounding multimodal large language models to the world. arXiv:2306.14824 (2023) [4](#)

36. Pi, R., Gao, J., Diao, S., Pan, R., Dong, H., Zhang, J., Yao, L., Han, J., Xu, H., Kong, L., Zhang, T.: Detgpt: Detect what you need via reasoning. arXiv:2305.14167 (2023) [4](#)
37. Pramanick, S., Han, G., Hou, R., Nag, S., Lim, S.N., Ballas, N., Wang, Q., Chellappa, R., Almahairi, A.: Jack of all tasks, master of many: Designing general-purpose coarse-to-fine vision-language model. arXiv preprint arXiv:2312.12423 (2023) [2](#), [4](#)
38. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International conference on machine learning. pp. 8748–8763. PMLR (2021) [4](#)
39. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: MICCAI (2015) [3](#)
40. Shelhamer, E., Long, J., Darrell, T.: Fully convolutional networks for semantic segmentation. TPAMI (2017) [3](#)
41. Sun, Q., Fang, Y., Wu, L., Wang, X., Cao, Y.: Eva-clip: Improved training techniques for clip at scale. arXiv preprint arXiv:2303.15389 (2023) [4](#)
42. Tian, Z., Chen, P., Lai, X., Jiang, L., Liu, S., Zhao, H., Yu, B., Yang, M.C., Jia, J.: Adaptive perspective distillation for semantic segmentation. TPAMI (2022) [3](#)
43. Tian, Z., Cui, J., Jiang, L., Qi, X., Lai, X., Chen, Y., Liu, S., Jia, J.: Learning context-aware classifier for semantic segmentation. AAAI (2023) [3](#)
44. Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., et al.: Llama: Open and efficient foundation language models. arXiv preprint arXiv:2302.13971 (2023) [4](#)
45. Wang, Z., Lu, Y., Li, Q., Tao, X., Guo, Y., Gong, M., Liu, T.: Cris: Clip-driven referring image segmentation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 11686–11695 (2022) [12](#)
46. Wei, J., Wang, X., Schuurmans, D., Bosma, M., Chi, E.H., Le, Q., Zhou, D.: Chain of thought prompting elicits reasoning in large language models. CoRR [abs/2201.11903](#) (2022), <https://arxiv.org/abs/2201.11903> [5](#)
47. Wu, P., Xie, S.: V\*: Guided visual search as a core mechanism in multimodal llms. arXiv preprint arXiv:2312.14135 [17](#) (2023) [5](#)
48. Xiong, Y., Liao, R., Zhao, H., Hu, R., Bai, M., Yumer, E., Urtasun, R.: Upsnet: A unified panoptic segmentation network. In: CVPR (2019) [3](#)
49. Xu, J., Liu, S., Vahdat, A., Byeon, W., Wang, X., De Mello, S.: Open-vocabulary panoptic segmentation with text-to-image diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2955–2966 (2023) [4](#)
50. Yang, M., Yu, K., Zhang, C., Li, Z., Yang, K.: Densenet for semantic segmentation in street scenes. In: CVPR (2018) [3](#)
51. Yang, Z., Wang, J., Tang, Y., Chen, K., Zhao, H., Torr, P.H.: Lavt: Language-aware vision transformer for referring image segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18155–18165 (2022) [12](#)
52. Ye, Q., Xu, H., Xu, G., Ye, J., Yan, M., Zhou, Y., Wang, J., Hu, A., Shi, P., Shi, Y., et al.: mplug-owl: Modularization empowers large language models with multimodality. arXiv:2304.14178 (2023) [4](#)
53. Yu, F., Koltun, V.: Multi-scale context aggregation by dilated convolutions. In: ICLR (2016) [3](#)

54. Zhang, S., Sun, P., Chen, S., Xiao, M., Shao, W., Zhang, W., Chen, K., Luo, P.: Gpt4roi: Instruction tuning large language model on region-of-interest. arXiv:2307.03601 (2023) [4](#)
55. Zhang, Z., Zhang, A., Li, M., Zhao, H., Karypis, G., Smola, A.: Multimodal chain-of-thought reasoning in language models. CoRR **abs/2302.00923** (2023). <https://doi.org/10.48550/arXiv.2302.00923>, <https://doi.org/10.48550/arXiv.2302.00923> [5](#)
56. Zhao, H., Qi, X., Shen, X., Shi, J., Jia, J.: Icnnet for real-time semantic segmentation on high-resolution images. In: ECCV (2018) [3](#)
57. Zhao, H., Shi, J., Qi, X., Wang, X., Jia, J.: Pyramid scene parsing network. In: CVPR (2017) [3](#)
58. Zhao, H., Zhang, Y., Liu, S., Shi, J., Loy, C.C., Lin, D., Jia, J.: Psanet: Point-wise spatial attention network for scene parsing. In: ECCV (2018) [3](#)
59. Zhu, D., Chen, J., Shen, X., Li, X., Elhoseiny, M.: Minigpt-4: Enhancing vision-language understanding with advanced large language models. arXiv:2304.10592 (2023) [4](#)
60. Zhu, Z., Xu, M., Bai, S., Huang, T., Bai, X.: Asymmetric non-local neural networks for semantic segmentation. In: ICCV (2019) [3](#)
61. Zou, X., Dou, Z.Y., Yang, J., Gan, Z., Li, L., Li, C., Dai, X., Behl, H., Wang, J., Yuan, L., et al.: Generalized decoding for pixel, image, and language. In: CVPR (2023) [4](#), [10](#), [12](#)
62. Zou, X., Yang, J., Zhang, H., Li, F., Li, L., Gao, J., Lee, Y.J.: Segment everything everywhere all at once. arXiv:2304.06718 (2023) [4](#), [10](#), [12](#)