

# Guideline-Consistent Segmentation via Multi-Agent Refinement

Vanshika Vats, Ashwani Rathee, James Davis

University of California, Santa Cruz  
1156 High St, Santa Cruz, CA 95064  
{vvats, aratheel, davisje}@ucsc.edu

## Abstract

Semantic segmentation in real-world applications often requires not only accurate masks but also strict adherence to textual labeling guidelines. These guidelines are typically complex and long, and both human and automated labeling often fail to follow them faithfully. Traditional approaches depend on expensive task-specific retraining that must be repeated as the guidelines evolve. Although recent open-vocabulary segmentation methods excel with simple prompts, they often fail when confronted with sets of paragraph-length guidelines that specify intricate segmentation rules. To address this, we introduce a multi-agent, training-free framework that coordinates general-purpose vision-language models within an iterative Worker-Supervisor refinement architecture. The *Worker* performs the segmentation, the *Supervisor* critiques it against the retrieved guidelines, and a lightweight reinforcement learning stop policy decides when to terminate the loop, ensuring guideline-consistent masks while balancing resource use. Evaluated on the Waymo and Reason-Seg datasets, our method notably outperforms state-of-the-art baselines, demonstrating strong generalization and instruction adherence.

**Project page** — <https://guideline-seg.github.io/>

**Extended version** — <https://arxiv.org/abs/2509.04687>

## Introduction

Pixel-accurate semantic image segmentation is crucial for safety-critical domains such as autonomous driving, robotics, and medical imaging (Elhassan and Zhou 2024; Al-Tawil et al. 2024; Wang et al. 2022). Traditionally, segmentation models have relied heavily on supervised training on pre-defined classes that generalize poorly to novel categories, requiring costly re-annotation and retraining whenever new classes emerge (Ronneberger, Fischer, and Brox 2015; Chen et al. 2018). Recent advancements have introduced open-vocabulary semantic segmentation, leveraging large-scale Vision-Language Models (VLMs) to recognize arbitrary object categories via textual prompts (Liang et al. 2023; Ren et al. 2024).

Semantic segmentation is not merely about generating accurate masks, but producing the *correct* masks that align

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

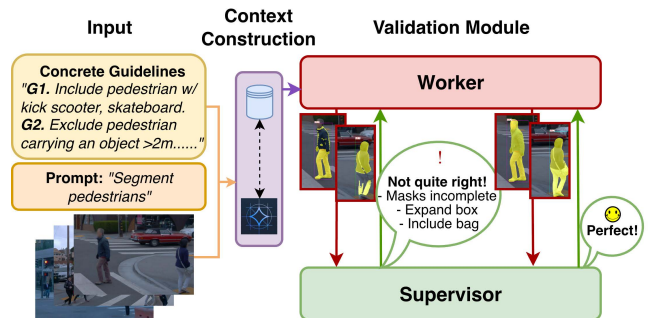


Figure 1: Method overview of our Worker-Supervisor iterative loop for guideline-consistent semantic segmentation.



Figure 2: Annotation inconsistencies in the Waymo segmentation dataset. Row (a) shows masks adhere to guideline G1 at timestamp  $t$ , while row (b), from the same scene at  $t + \Delta t$ , shows a clear violation of the same guideline.

with the intended labeling policy, *every time*. These decisions often rely on strict, text-defined rules. For instance, Waymo dataset (Sun et al. 2020) specifies that *pedestrian* class must ‘include skateboard riders, exclude mannequins’. Similarly, in land-cover segmentation (Demir et al. 2018), the class *range land* should cover ‘Any non-forest, non-farm, green land, grass’. When models ignore these intricate details, even pixel-accurate masks can be functionally wrong. Existing open-vocabulary systems excel on short prompts

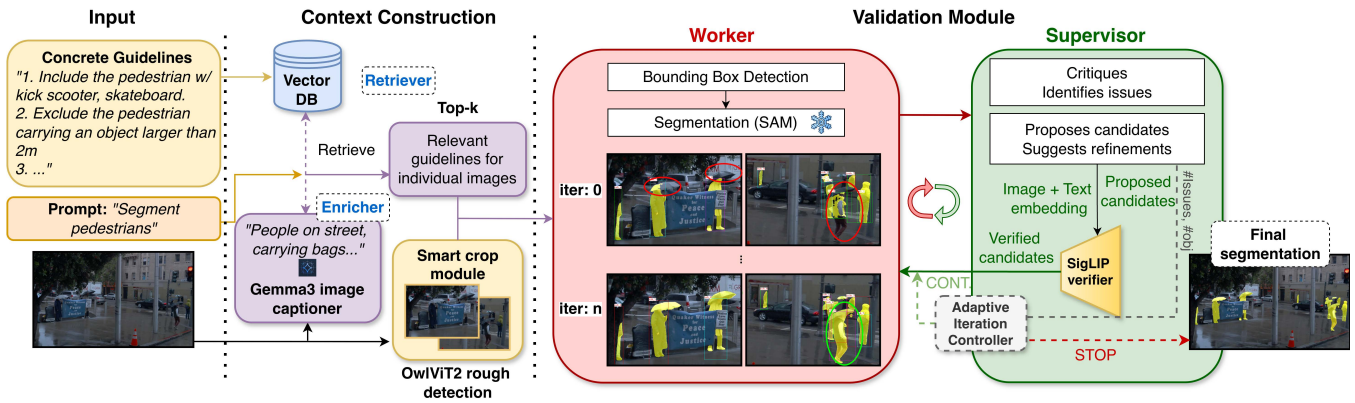


Figure 3: Overview of our pipeline. The input image and textual guidelines undergo context construction to extract scene-relevant rules. This focused context is processed through an iterative VLM loop where the *Worker* segments and the *Supervisor* critiques and suggests improvements. The number of iterations is controlled by an adaptive iteration controller. The final output is a segmentation that faithfully adheres to long and detailed guidelines.

(Lai et al. 2024; Sun et al. 2025) but crumble under long intricate rules. Worse, some public ‘ground-truth’ masks produced by costly human-annotators, themselves often violate their own manuals. As shown in Fig. 2, the same scene at different timestamps in the Waymo dataset exhibits inconsistent masks that ignore the guidelines. Domain-specific detectors and fine-tuned models might be adapted for guideline-aware segmentation, but only after costly retraining. Labeling rules can evolve, and relying on continuous retraining is impractical. These inconsistencies reveal the difficulty of the task of rule-adhering segmentation and undermine the reliability of conventional benchmarks.

In this study, we achieve guideline-consistent semantic segmentation via multi-agent, training-free framework. Leveraging VLMs (Gemini-2.5-flash-preview (Comanici et al. 2025)) as agents, our approach employs a hierarchical Worker-Supervisor system where each agent plays a distinct role, iteratively generating, critiquing, refining, and validating segmentation outputs (Fig. 1). Specifically, the *Worker* module first produces initial detections based on enriched image-specific contexts. These detections are then segmented by a frozen Segment Anything Model (SAM) (Kirillov et al. 2023; Ravi et al. 2024). Subsequently, the *Supervisor* module evaluates the Worker’s segmentations against the given guidelines, identifies errors, and proposes refinements. This loop continues until a reinforcement learning-based stop criterion terminates the process or a maximum iteration count is reached. Crucially, both the VLM and SAM remain frozen; all adaptation occurs through strategic prompting and context-aware control, making our method uniquely flexible, scalable, and capable of handling complex, paragraph-length instructions.

Our method outperforms existing state-of-the-art (SOTA) approaches, delivering a notable increase in guideline-consistent segmentation performance on the Waymo and ReasonSeg datasets. This highlights our framework’s superior ability to produce precise segmentations that closely adhere to intricate, lengthy guidelines without requiring model

retraining. Thus, the main contribution of our work is a novel framework for training-free guideline-consistent semantic segmentation, capable of interpreting and adhering to long and detailed labeling instructions.

## Related Works

**Vision-Language Models for Segmentation.** Vision-Language models (VLMs) align visual features with language (Radford et al. 2021; Li et al. 2022; Liu et al. 2023) and are increasingly being used in object detection (Liu et al. 2024b; Cheng et al. 2024) and segmentation (Liang et al. 2023; Xia et al. 2024). Recently, promptable segmentation has emerged as a flexible paradigm where an image and user prompts jointly determine the output mask. The prompts are typically either points/boxes (Kirillov et al. 2023) or text-conditioned (Ren et al. 2024). Unified VL segmenters (Lai et al. 2024; Zou et al. 2023; Sun et al. 2025) embed multi-modal prompts into a single transformer, enabling open-set and multi-prompt segmentation. Reasoning-based segmenters further combine language reasoning with segmentation, either through chain-of-thought style guidance or specialized `<seg>` tokens that steer mask prediction (Liu et al. 2025; Qian et al. 2025). However, these methods typically handle only single words or short sentences, often performing poorly on longer contexts. More recent models like Gemini-2.5 (Comanici et al. 2025), Qwen2.5-VL (Bai et al. 2025), and GPT-4V (OpenAI 2023) offer improved comprehension of longer textual inputs but still fall short of accurate detection and segmentation. Moreover, when operating in a single pass, they offer limited understanding of lengthy intricate rules, complex constraints, or structured verification guideline compliance. Using Gemini-2.5 as our baseline, we address these concerns by introducing an iterative-refinement and verification framework that yields guideline-consistent segmentation. This choice was driven by the need to effectively process longer contexts, a capability not yet reliably offered by current open-weight models.

**Agents and Self-Correction.** Recent research frames large models as smaller, specialized agents that collaborate to complete complex tasks (Wu et al. 2023; Shen et al. 2023). These agents iteratively refine their outputs based on mutual feedback without additional training (Madaan et al. 2023; Shinn et al. 2023). In vision, VLMs leverage self-correction for iterative re-grounding (Liao et al. 2025; Su et al. 2024; Sun, Xiao, and Lim 2021), training-free contrastive guidance (Wan et al. 2024), and visual prompting (Yang et al. 2023). Only a few agent-based approaches have addressed semantic segmentation (Sun et al. 2025), and existing ones neither produce accurate masks under long instruction nor include explicit self-critique mechanisms. To bridge these gaps, we propose a Worker-Supervisor agent framework that iteratively refines the segmentations, effectively addressing missed detections and pruning false positives under complex guidelines. Such agent-based workflows are well-suited to semantic segmentation tasks because they inherently require multiple steps: segmenting, verifying, refining, and adaptively determining when to stop, and thus greatly benefit from iterative planning and structured feedback under detailed, rule-rich instructions.

## Method

Fig. 3 summarizes our pipeline. Given the input image and a set of guidelines, we first build the context from the image description and retrieve relevant guidelines, followed by a loop of Worker  $\rightarrow$  Supervisor  $\rightarrow$  Worker ... till the Adaptive Iteration Controller (AiRC) decides when to stop or upon reaching a maximum number of predecided iterations.

### Inputs

We begin with an input image  $Img_i$  that must be segmented based on a prompt  $P$ , while strictly following a set of detailed guidelines. For instance, in the Waymo dataset (Sun et al. 2020), inclusion rules may state ‘*Include a small child or small items (e.g., umbrella, handbag, sign) with the pedestrian*’, and exclusion rules may require ‘*Exclude mannequins, statues, dummies, covered objects, billboards, posters, indoor pedestrians, and reflections*’. To enable structured interpretation, we use GPT-4o (Hurst et al. 2024) to convert them into a structured JSON format, assigning each a unique ID  $G = \{G_0, G_1, \dots, G_m\}$  and a brief description. This facilitates the context construction.

### Context Construction

Each image is different, and not all guidelines are relevant to every scene. For example, for an image without bicycles or cyclists, ‘ $G_3$ : *A person riding a bicycle is not labeled as a pedestrian, but labeled as a cyclist*’ becomes irrelevant. Including such unnecessary instructions can overwhelm the pipeline with irrelevant details. Thus, we construct image-specific contexts by retrieving only the most relevant guidelines.

We encode all textual guidelines into a FAISS (Johnson, Douze, and Jégou 2019) vector database (DB), enabling fast similarity-based retrieval. Storing guidelines in a vector DB enables quick filtering of the most relevant rules for a given

image. This follows a context construction through an *Enricher*, which describes the scene and enriches the retrieval query, and a *Retriever*, which searches for the most relevant guidelines from the vector DB. All context modules are lightweight and add minimal computational overhead.

**Enricher.** To retrieve the most relevant guidelines for a given image, we first generate a concise yet informative textual description of the scene. Gemma3-4B (Gemma et al. 2025), a lightweight multimodal model, is used to caption the image based on the prompt. We then form a text-rich query  $Q = \{P, \langle caption \rangle, H \times W\}$ , where  $H \times W$  represents the image resolution. Including resolution helps account for scale-dependent rules like, ‘Ignore distant pedestrians’, ‘Exclude oversized objects’, enriching the semantic context. This query is encoded into a vector ( $Q_v$ ) using the SentenceTransformer (Reimers and Gurevych 2020) (all-MiniLM-L6-v2) to semantically fetch relevant guidelines from the FAISS DB.

**Retriever.** The Retriever module uses  $Q_v$  to perform a cosine similarity search over the FAISS index, retrieving the top- $k$  most relevant guidelines  $G_k \subset G$ . This selective retrieval avoids overloading the downstream Validation module with unnecessary rules, ensuring the context remains concise, focused, and aligned with the image content.

**Smart Crop.** Modern VLMs often struggle with small or distant objects due to limited attention resolution. To mitigate this, we employ a smart cropping strategy that divides the image into two focused regions based on rough object locations. We first downscale the image  $Img$  to  $0.8\times$  and pass it through OWLv2 (Minderer, Gritsenko, and Houlsby 2023) to obtain rough, noisy bounding boxes ( $B$ ) for the target class. This gives us focused regions of interest.

To avoid splitting objects and maintain balance, the module analyzes object distribution and spatial layout. It chooses a vertical split that keeps object counts roughly equal on both sides and maximizes the gap between clusters, ensuring clean separation. The resulting crops  $x_\ell$  and  $x_r$ , defined by  $\max_x(B_{\text{left}})$  and  $\min_x(B_{\text{right}})$ , preserve semantic coherence and are passed to the Validation module for processing.

### Validation Module

The Validation module handles segmentation through an iterative loop composed of two primary components: the **Worker**, which performs object detection and segmentation, and the **Supervisor**, which evaluates the Worker’s output against the retrieved guidelines and proposes corrections. An **Adaptive Iteration Controller (AiRC)** governs how many refinement rounds are performed. AiRC monitors factors such as object count and unresolved issues to decide whether to continue or stop the loop, up to a maximum iteration limit.

**Worker.** The Worker module is dedicated to executing the segmentation task and comprises two stages: an initial detection step and a refinement step across  $n$  iterations, incorporating Supervisor feedback.

At initialization, the Worker identifies bounding boxes for the target class based on the user prompt  $P$ , returning coord-

**G1:** "Include small items or child (< 2 m) with the pedestrian (umbrella, handbag, sign, etc.)."  
**G2:** "Exclude mannequins, statues, dummies, covered objects, billboards, posters, indoor pedestrians, and reflections."

**G3:** "Label a person riding a bicycle as a cyclist, not a pedestrian."  
**G4:** "Exclude large objects (> 2 m) or pushed items (bike, shopping cart, etc.) from the pedestrian class." ...



Figure 4: Qualitative comparison of our method with other state-of-the-art. Our method successfully complies with guidelines G1 [rows-(a),(c)], G2 [rows-(b),(d)], G3 and G4 [row-(e)]. More examples in supplementary material.

ordinates in the form  $[y_{min}, x_{min}, y_{max}, x_{max}]$ , class labels, and a unique ID  $Subjects : \{sub_0, sub_1, \dots, sub_m\}$  to each subject. These are passed to a frozen SAM model to produce segmentation masks. Note that this SAM model is not finetuned or retrained on any particular dataset to maintain its generalizability; we rely on carefully crafted prompt inputs (boxes, points) to guide it in producing accurate masks.

Since general-purpose VLMs are not optimized for coordinate-precise localization, the initial boxes are often coarse or misaligned, potentially leading to incorrect or incomplete masks. This inaccuracy propagates to the SAM segmentation outputs. Moreover, SAM, while powerful, is task-agnostic and may overlook rule-specific inclusion/exclusion criteria. To correct such issues, the Worker triggers the Supervisor, which audits the output and provides corrective feedback for iterative refinement.

**Supervisor.** The Supervisor is responsible for evaluating and critiquing the outputs generated by the Worker. Following the principle that a group of smaller, specialized agents often performs better than a single agent overloaded with responsibilities (Hong et al. 2023; Wu et al. 2024), we divide the Supervisor into two distinct components:

*Agent 1-Supervisor\_eval:* This agent has access to the image, the Worker’s output, and the relevant guidelines  $G_k$ . It identifies three types of issues: (i) missing objects that should have been segmented, (ii) false positives that violate exclusion rules, and (iii) refinement opportunities for imperfect masks. Each issue is recorded in a structured JSON, specifying object type, brief justification, and precise refinement suggestions such as ‘*expand box to the right to include hand*’ or ‘*shrink box to exclude the background*’. These are passed to Agent 2 for candidate box generation.

*Agent 2-Supervisor\_boxgen:* This agent receives

the missing\_objects, false\_positives, and corresponding critiques from Agent 1, and generates bounding boxes for each proposed candidate. However, since these items are often small and lack rich visual or textual context, the proposed boxes tend to be imprecise, undersized, or occasionally misplaced. We, thus, employ a SigLIP-based verifier.

**SigLIP Verifier.** To verify candidates from Agent 2, we crop image regions with a small buffer to retain context (e.g., a tightly cropped bag alone may be unclear, but showing it being held helps in recognition). These image crops are passed through a SigLIP (Zhai et al. 2023) image encoder, while the associated labels are encoded using a text encoder. SigLIP outputs a logit for each image-text pair, which we convert to a probability using a sigmoid function. If this probability exceeds a threshold, the candidate is considered a valid suggestion by the Supervisor; otherwise, it is discarded.

The verified candidates and refinement instructions are returned to the Worker for final adjustments. For missing objects, the Worker accepts verified bounding boxes and uses SAM to generate the corresponding masks. False positives are removed by prompting SAM with the identified box and a negative point label to erase the mask. For refinements, the Worker modifies existing bounding boxes based on the Supervisor’s suggestions before re-generating the masks.

### Adaptive Iteration Controller

Each iteration in our loop requires three API calls, so the overall cost scales with the number of iterations. A fixed count could be used, but it would not adapt well across varying scenes, i.e., too many iterations may waste resources, while too few might cut off meaningful improvements. To address this, we implement a reinforcement learning (RL) policy that dynamically determines when to stop the loop. This adaptive iteration controller (AiRC) tracks scene density and unresolved guideline violations, and the policy selects CONTINUE or STOP accordingly. We frame the iteration-control problem as a finite-horizon Markov Decision Process (MDP) with the following elements.

**Issue count.** We quantify Supervisor’s feedback as issue count  $I$ , which shapes the reward function. Each genuine guideline error (missed or false object) counts 1 point, driving the learning signal. Refinements usually suggest low-impact mask tweaks and can be repetitive and noisy, and thus count as 0.1 points. Thus, the issue count  $I_t$  at the end of each iteration becomes

$$I_t = I_{miss} + I_{false} + 0.1I_{ref} \quad (1)$$

**State space.** The state representation captures both the scene complexity and the presence of unresolved guideline-violating issues  $I$ . Each  $Img$  crop is mapped to one of six abstract states, defined as  $s = 2d + v$  where  $d \in \{0, 1, 2\}$  is a coarse scene-complexity density bucket (few/medium/crowd) computed from the initial object count in the crop, and  $v \in \{0, 1\}$  flags whether residual guideline violations remain after the current pass (clean=0, dirty=1).

**Actions.** There exist two possible actions at each step: either to stop the iterative process or proceed to the next step, up to a predefined MAX\_ITERS.  $a \in \{\text{STOP}, \text{CONTINUE}\}$ .

**Immediate reward.** The immediate reward reflects the change in issue count  $I_t$  across iterations. At iteration  $t$ :

$$r(s, a, s') = \begin{cases} \underbrace{(I_t - I_{t+1})}_{\text{issues fixed}} - \underbrace{c}_{\text{step cost}} + \underbrace{b[I_{t+1} = 0]}_{\text{early-resolve bonus}}, & a_t = \text{CONTINUE}, \\ 0, & a_t = \text{STOP}, I_t = 0, \\ - \underbrace{p}_{\text{early-stop penalty}}, & a_t = \text{STOP}, I_t > 0. \end{cases} \quad (2)$$

The reward  $r$  is defined by the drop in issue count between consecutive steps, with a small step cost  $c$  per iteration. This encourages the AiRC to continue only when further refinement is worthwhile. An early-stop penalty  $p$  is applied if STOP is chosen with unresolved issues, while a clean-scene bonus  $b$  is given if all issues are resolved before reaching MAX\_ITERS. Action values are updated with standard tabular Q-learning (Watkins and Dayan 1992). To put simply, AiRC chooses CONTINUE if

$$Q(s_t, \text{CONTINUE}) > Q(s_t, \text{STOP}) \quad (3)$$

The module selects CONTINUE whenever the Q-value (expected discounted return) for CONTINUE in the current state exceeds the Q-value for STOP; otherwise, it stops. (More in supplementary)

We use RL instead of a fixed heuristic because the value of an additional refinement emerges only after that pass, creating a delayed, sequential reward structure that RL’s temporal-difference updates naturally handle. RL learns exactly when the expected error reduction justifies the per-iteration cost, adapting to each crop’s complexity without manual thresholds. In contrast, static rules like a fixed pass count or issue-count cutoff either leave errors in dense scenes or waste effort on simple ones. A compact six-state Q-table delivers this adaptability with minimal overhead.

## Experiments

### Datasets

We evaluate our method on two datasets, the Waymo Perception dataset (Sun et al. 2020) and ReasonSeg (Lai et al. 2024). For Waymo, we focus on the *Pedestrian* class, which poses a greater challenge due to its subtle object boundaries and frequent segmentation errors. In contrast, larger objects like *Car* can inflate performance metrics and shadow fine-grained issues. Additionally, as noted in Fig. 2, the Waymo ground truth often deviates from its own labeling guidelines. Since our method explicitly aims to enforce guideline consistency, relying on such noisy annotations could lead to misleading results. To address this, we manually curate a set of 101 samples that adhere strictly to the defined guidelines. Creating this smaller, high-quality set was a necessary and labor-intensive step to enable a fair and precise evaluation

Method Metric	LISA-7B (Lai et al. 2024)			LISA-13B (Lai et al. 2024)			GroundedSAM (Ren et al. 2024)			READ (Qian et al. 2025)			Gemini-2.5 (Comanici et al. 2025)			SegZero (Liu et al. 2025)			Ours
	S	C	F	S	C	F	S	C	F	S	C	F	S	C	F	S	C	F	F
gIOU	43.42	36.47	23.41	66.74	40.96	20.16	75.00	20.84	20.43	50.90	47.79	43.94	67.84	71.39	69.02	73.49	73.72	71.96	<b>80.57</b>
cIOU	42.73	34.53	18.32	73.21	34.04	14.21	79.78	20.91	19.36	53.73	54.16	45.55	72.01	76.38	74.24	75.98	76.91	73.72	<b>86.70</b>
mPr	85.11	40.27	24.38	80.77	46.62	22.41	88.12	33.16	29.35	75.63	67.57	51.74	79.05	84.18	80.91	89.80	88.38	86.34	<b>91.06</b>
mRec	45.55	74.86	79.79	77.52	74.91	62.41	81.51	32.52	28.11	60.82	55.55	69.75	79.75	80.78	79.89	78.46	77.91	78.45	<b>84.78</b>
mDice	57.14	36.47	33.17	77.92	51.60	28.30	83.90	27.61	25.63	63.86	58.19	59.41	77.75	80.76	78.75	83.75	82.15	80.88	<b>87.20</b>

Table 1: Segmentation results on the Waymo *guideline-consistent* dataset. As the input guideline length increases from single-word (S) to condensed (C) to full-length (F), the performance of other SOTA methods declines, with the exception of Gemini-2.5. Our method consistently outperforms all methods, even with full-length guidelines.

Method	gIOU	cIOU
SEEM (Zou et al. 2023)	21.2	25.5
GroundedSAM (Ren et al. 2024)	26.0	14.5
LISA-7B (Lai et al. 2024)	46.0	44.4
FaST (Sun et al. 2025)	48.7	47.6
LISA-7B(ft) (Lai et al. 2024)	52.9	54.0
Gemini-2.5 (Comanici et al. 2025)	55.5	44.9
LISA-13B-LLava-1.5(ft) (Lai et al. 2024)	57.7	60.3
READ (Qian et al. 2025)	59.8	<b>67.6</b>
SegZero (Liu et al. 2025)	62.6	62.0
Ours	<b>68.1</b>	66.4

Table 2: Segmentation results on the ReasonSeg *val* dataset. Our method outperforms the other SOTA gIOU by a good margin.

of our core contribution. The data was randomly selected from 55 diverse scenes, spanning various conditions such as sunny, rainy, and nighttime settings, ensuring broad coverage and variability in the evaluation (See supplementary).

To further evaluate the generalizability of our method, we also test it on the ReasonSeg validation set, which contains 200 challenging in-the-wild images annotated with pixel-level masks and textual rationales that span from concise phrases to more complex reasoning prompts. While the dataset does not include full-length guidelines, its intricate queries simulate guideline-driven segmentation scenarios and support the study of explainable segmentation.

## Implementation

Given an image  $I$ , prompt  $P$ , and guidelines  $G$ , we embed  $G$  and index with FAISS (Johnson, Douze, and Jégou 2019). A lightweight captioner (Gemma3-4B) (Gemma et al. 2025) forms the query  $Q = \{P, \langle \text{caption} \rangle, H \times W\}$  to retrieve the top- $k=8$  rules. We resize  $I$  to  $0.8 \times$  and obtain coarse boxes with OWLv2 to seed Smart-Crop.

We employ Gemini-2.5-flash-preview-05-20 API for both the Worker and Supervisor VLMs, using it as our baseline. The Worker runs at temperature  $T = 0.5$ , as recommended for object detection flexibility.

Supervisor\_eval critiques with  $T = 0.3$  for more deterministic reasoning, using dynamic thinking mode for improved contextual judgment. Supervisor\_boxgen has  $T = 0.5$  for bounding box suggestions. Masks are produced by SAM2.1\_hiera\_large. Supervisor proposals are validated with SigLIP; crops are accepted when sigmoid is  $\geq 0.5$ .

To account for VLM non-determinism, we allow MIN\_ITERs=2 and MAX\_ITERs=4, with early stopping controlled by AiRC. The stop policy is a  $3 \times 2$  tabular Q-learner ( $\alpha = 0.3, \gamma = 0.9$ ) with fixed  $\epsilon$ -greedy exploration ( $\epsilon = 0.02$ ). The reward uses step cost  $c = (-)0.02$ , early-stop penalty  $p = (-)2.0$ , and clean-scene bonus  $b = +1.0$ . The Q-table is persisted across runs. The inference is performed on a single RTX3080 GPU, but has the flexibility to be split across  $3 \times$  GPUs for efficiency. To validate the robustness of our framework against VLM non-determinism, we perform three independent runs with different random seeds and report means and standard deviation.

## Evaluation Metrics

We use global intersection over union (gIOU) and complete IoU (cIOU) as primary metrics (Kazemzadeh et al. 2014; Mao et al. 2016; Lai et al. 2024). gIOU is the per-image IoU averaged over the dataset; cIOU is the ratio of the cumulative intersection to the cumulative union across all images. Because cIOU disproportionately weighs large objects more, gIOU offers a more balanced assessment. We also report mean Precision (mPr), mean Recall (mRec), and mean Dice/F1 (mDice), computed per image and averaged, to quantify false positives, false negatives, and their trade-off.

## Results and Discussion

Figure 4 and Tables 1–2 summarize our main results. In Table 1, we evaluate the Waymo *guideline-consistent* set under three text settings: (i) the full Waymo guidelines, (ii) a condensed set of short, crisp phrases of guidelines, and (iii) a single-word main class name ‘Pedestrian’. Since phrase-conditioned SOTA methods like LISA and GroundedSAM are limited to handling short phrases and sentences, we expect performance to drop as the referring text length grows, a trend confirmed in Table 1. An exception is seen in Gemini-2.5 having enhanced long-context

Dataset	W	C	S	gIOU	cIOU	mPr	mRec
Waymo	✓	-	-	69.02	74.24	80.91	79.89
	✓	✓	-	73.87	78.40	88.36	80.81
	✓	✓	✓	<b>80.57</b>	<b>86.70</b>	<b>91.06</b>	<b>84.78</b>
ReasonSeg	✓	-	-	55.56	44.98	70.66	61.78
	✓	-	✓	<b>68.12</b>	<b>66.35</b>	<b>77.82</b>	<b>76.06</b>

Table 3: Ablation table showing the benefit of iterative critiquing. W-Worker, C- Context construction, S-Supervisor.

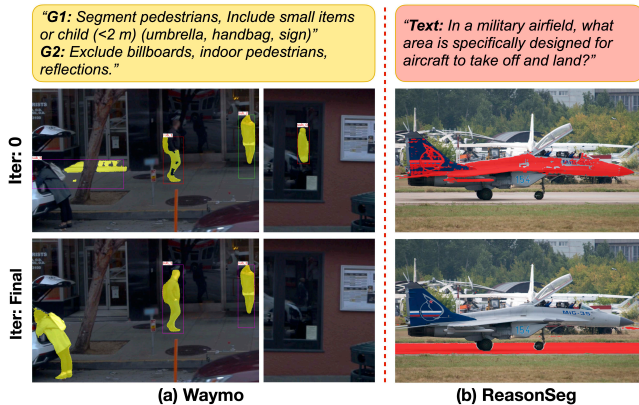


Figure 5: Examples showing how our iterative loop produces guideline-adhering segmentations over just relying on the VLM initial outputs.

understanding and achieves better results when given the full set of guidelines. Adding our image-specific guideline retrieval with iterative refinement yields +57.16 gIoU over the strongest phrase-conditioned full-guideline baseline, +11.55 gIoU over the baseline Gemini-2.5, and +8.61 gIoU ( $80.57 \pm 0.6$ ) over SOTA. Results are stable over three runs: high recall ( $84.78 \pm 0.9$ ) and precision ( $91.06 \pm 0.5$ ) show that our method recovers missed items (e.g., umbrella, backpack) and removes false positives (e.g., cyclist, reflections) that competing methods overlook.

On the ReasonSeg dataset (Table 2), we outperform the strongest baseline (SegZero) by +5.5 gIoU. This demonstrates generalization to in-the-wild scenes with rationale-style instructions.

Our results follow directly from the design. First, the context construction via image-specific guideline retrieval prevents instruction overload, so the VLM reasons over only the rules relevant to the scene, explaining robustness on long guidelines. Smart crop improves initial detections of small or distant instances. Second, the Worker-Supervisor decomposition allows iterative correction of missed items and false positives. Third, AiRC adaptively stops refinement to balance accuracy and resource usage, avoiding premature stops or unnecessary iterations. We prioritize segmentation accuracy over latency, so the loop architecture may add VLM calls, but the controller bounds iterations to keep cost reasonable (2.6 iters on average,  $\approx \$0.0088$  per sample with Gemini-2.5-flash; See supplementary). Fi-

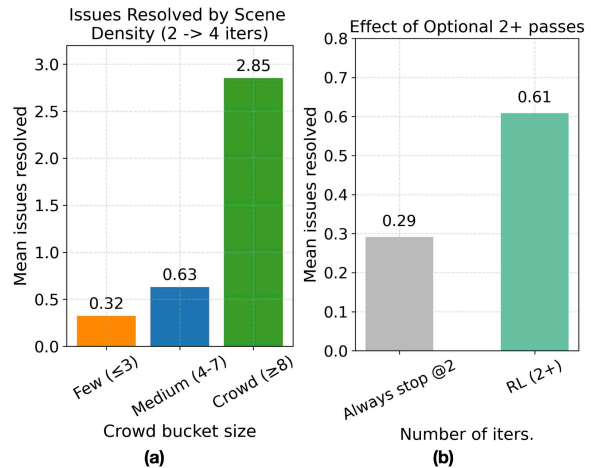


Figure 6: Effect of using RL controller measured in mean issues resolved over the dataset by (a) scene density and (b) letting AiRC dynamically make the stop decision and not hard-stopping at 2 iterations.

nally, keeping VLM and SAM frozen avoids dataset-specific overfitting, consistent with the gains we observe on ReasonSeg. Together, these choices address the weaknesses of open-vocabulary methods under complex policies, resulting in substantial accuracy improvements and stable performance even on long prompt lengths.

## Ablation

We ablate the iterative Worker-Supervisor loop by comparing pure Worker-only VLM detection against our context-enhanced and Supervisor-refined approaches (Table 3). Providing targeted context boosts performance, with further improvement from Supervisor feedback. On Waymo, increases in IoU and recall indicate effective recovery of missed objects while the precision gains reflect pruning of false positives (Fig. 5-a). For ReasonSeg, context construction is unnecessary due to shorter prompts; however, Supervisor refinement still notably improves segmentation (Fig. 5-b). These results confirm that iterative refinement substantially enhances guideline adherence compared to one-shot VLM inference.

We also evaluate the advantage of using an adaptive iteration controller instead of a fixed stop heuristic. Fig.6-a shows the gains when we do not hard-stop at `MIN_ITER=2`. The largest gains on crowded scenes indicate that extra effort is directed where violations are most likely. Fig.6-b shows that, on average, our AiRC module resolves 0.61 violations per crop versus 0.29 for a hard stop at two passes, a 110% improvement at the cost of running extra passes on 48% of crops. Both plots effectively show that dynamic stopping allocates compute to difficult scenes and increases overall quality at modest additional cost.

## Limitations and Future Work

Our approach depends on clear, comprehensive guidelines; ambiguity or conflicts in the instructions can degrade its sta-

bility. Also, our current system’s dependency on a proprietary VLM API constrains portability and on-premises deployment. This was a necessary trade-off, as the Supervisor’s critiquing role demands a level of long-context reasoning and instruction adherence that open-weight models do not yet reliably offer. We also prioritize accuracy over latency, so multiple API calls limit real-time use. In future work, we will extend the framework to instance segmentation alongside semantic masks, enabling explicit association of missed/false objects with the correct subject instances. We also plan to reduce the dependency on proprietary VLM APIs by evaluating open weight, locally deployable models such as LLaVA-NeXT (Liu et al. 2024a) and Qwen2.5-VL (Bai et al. 2025) as they become increasingly capable.

## Conclusion

We introduce guideline-consistent semantic segmentation, emphasizing strict adherence to detailed textual guidelines crucial in real-world applications. Our novel, training-free framework employs a multi-agent Worker-Supervisor architecture enhanced by a lightweight reinforcement learning stop policy, iteratively refines segmentation outputs to ensure compliance with complex, paragraph-length instructions without retraining. Evaluations on the Waymo and ReasonSeg datasets demonstrate that our approach notably surpasses existing state-of-the-art methods, highlighting its robustness, generalization capability, and efficiency in handling intricate guideline sets.

## References

- Al-Tawil, B.; Hempel, T.; Abdelrahman, A.; and Al-Hamadi, A. 2024. A review of visual SLAM for robotics: Evolution, properties, and future applications. *Frontiers in Robotics and AI*, 11: 1347985.
- Bai, S.; Chen, K.; Liu, X.; Wang, J.; Ge, W.; Song, S.; Dang, K.; Wang, P.; Wang, S.; Tang, J.; Zhong, H.; Zhu, Y.; Yang, M.; Li, Z.; Wan, J.; Wang, P.; Ding, W.; Fu, Z.; Xu, Y.; Ye, J.; Zhang, X.; Xie, T.; Cheng, Z.; Zhang, H.; Yang, Z.; Xu, H.; and Lin, J. 2025. Qwen2.5-VL Technical Report. *arXiv preprint arXiv:2502.13923*.
- Chen, L.-C.; Zhu, Y.; Papandreou, G.; Schroff, F.; and Adam, H. 2018. Encoder-Decoder with Atrous Separable Convolution for Semantic Image Segmentation. In *European Conference on Computer Vision (ECCV)*, 833–851. Springer.
- Cheng, T.; Song, L.; Ge, Y.; Liu, W.; Wang, X.; and Shan, Y. 2024. YOLO-World: Real-Time Open-Vocabulary Object Detection. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*.
- Comanici, G.; Bieber, E.; Schaekermann, M.; Pasupat, I.; Sachdeva, N.; Dhillon, I.; Blistein, M.; Ram, O.; Zhang, D.; Rosen, E.; et al. 2025. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*.
- Demir, I.; Koperski, K.; Lindenbaum, D.; Pang, G.; Huang, J.; Basu, S.; Hughes, F.; Tuia, D.; and Raskar, R. 2018. DeepGlobe 2018: A Challenge to Parse the Earth through Satellite Images. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 172–17209.
- Elhassan, M. A. M.; and Zhou, C. 2024. Real-time Semantic Segmentation for Autonomous Driving: A Review of CNNs, Transformers, and Beyond. *Journal of King Saud University - Computer and Information Sciences*.
- Gemma, T.; Kamath, A.; Ferret, J.; Pathak, S.; Vieillard, N.; Merhej, R.; Perrin, S.; Matejovicova, T.; Ramé, A.; Rivière, M.; et al. 2025. Gemma 3 technical report. *arXiv preprint arXiv:2503.19786*.
- Hong, S.; Zhuge, M.; Chen, J.; Zheng, X.; Cheng, Y.; Wang, J.; Zhang, C.; Wang, Z.; Yau, S. K. S.; Lin, Z.; et al. 2023. MetaGPT: Meta programming for a multi-agent collaborative framework. In *The Twelfth International Conference on Learning Representations*.
- Hurst, A.; Lerer, A.; Goucher, A. P.; Perelman, A.; Ramesh, A.; Clark, A.; Ostrow, A.; Welihinda, A.; Hayes, A.; Radford, A.; et al. 2024. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*.
- Johnson, J.; Douze, M.; and Jégou, H. 2019. Billion-scale similarity search with GPUs. *IEEE Transactions on Big Data*, 7(3): 535–547.
- Kazemzadeh, S.; Ordonez, V.; Matten, M.; and Berg, T. 2014. ReferItGame: Referring to Objects in Photographs of Natural Scenes. In Moschitti, A.; Pang, B.; and Daelemans, W., eds., *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, 787–798. Association for Computational Linguistics.
- Kirillov, A.; Mintun, E.; Ravi, N.; Mao, H.; Rolland, C.; Gustafson, L.; Xiao, T.; Whitehead, S.; Berg, A. C.; Lo, W.-Y.; Dollár, P.; and Girshick, R. 2023. Segment Anything. *arXiv:2304.02643*.
- 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, J.; Li, D.; Xiong, C.; and Hoi, S. 2022. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International conference on machine learning*, 12888–12900. PMLR.
- Liang, F.; Wu, B.; Dai, X.; Li, K.; Zhao, Y.; Zhang, H.; Zhang, P.; Vajda, P.; and Marculescu, D. 2023. Open-vocabulary semantic segmentation with mask-adapted CLIP. In *Proceedings of the IEEE/CVF CVPR*, 7061–7070.
- Liao, Y.-H.; Mahmood, R.; Fidler, S.; and Acuna, D. 2025. Can Large Vision–Language Models Correct Semantic Grounding Errors By Themselves? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Liu, H.; Li, C.; Li, Y.; Li, B.; Zhang, Y.; Shen, S.; and Lee, Y. J. 2024a. LLaVA-NeXT: Improved reasoning, OCR, and world knowledge.
- Liu, H.; Li, C.; Wu, Q.; and Lee, Y. J. 2023. Visual instruction tuning. *NeurIPS*, 36: 34892–34916.

- Liu, S.; Zeng, Z.; Ren, T.; Li, F.; Zhang, H.; Yang, J.; Jiang, Q.; Li, C.; Yang, J.; Su, H.; et al. 2024b. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. In *ECCV*, 38–55. Springer.
- 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*.
- Madaan, A.; Tandon, N.; Gupta, P.; Hallinan, S.; Gao, L.; Wiegrefe, S.; Alon, U.; Dziri, N.; Prabhume, S.; Yang, Y.; Gupta, S.; Majumder, B. P.; Hermann, K.; Welleck, S.; Yazdanbakhsh, A.; and Clark, P. 2023. SELF-REFINE: iterative refinement with self-feedback. In *Proceedings of the 37th International Conference on NeurIPS*. Red Hook, NY, USA: Curran Associates Inc.
- Mao, J.; Huang, J.; Toshev, A.; Camburu, O.; Yuille, A.; and Murphy, K. 2016. Generation and Comprehension of Unambiguous Object Descriptions. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 11–20.
- Medeiros, L. 2023. Language Segment-Anything. <https://github.com/luca-medeiros/lang-segment-anything>.
- Minderer, M.; Gritsenko, A.; and Hounsby, N. 2023. Scaling open-vocabulary object detection. *Advances in Neural Information Processing Systems*, 36: 72983–73007.
- OpenAI. 2023. GPT-4V(ision) System Card. <https://openai.com/index/gpt-4v-system-card/>. Accessed: 2025-11-15.
- Qian, R.; Yin, X.; Dou, D.; and Dou, D. 2025. Reasoning to attend: Try to understand how SEG token works. In *Proceedings of the CVPR Conference*, 24722–24731.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *ICML*, 8748–8763. PmLR.
- 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 preprint arXiv:2408.00714*.
- Reimers, N.; and Gurevych, I. 2020. Making Monolingual Sentence Embeddings Multilingual using Knowledge Distillation. In *Proceedings of EMNLP*, 849–862. Association for Computational Linguistics.
- Ren, T.; Liu, S.; Zeng, A.; Lin, J.; Li, K.; Cao, H.; Chen, J.; Huang, X.; Chen, Y.; Yan, F.; Zeng, Z.; Zhang, H.; Li, F.; Yang, J.; Li, H.; Jiang, Q.; and Zhang, L. 2024. Grounded SAM: Assembling Open-World Models for Diverse Visual Tasks. *arXiv:2401.14159*.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, volume 9351, 234–241. Springer.
- Shen, Y.; Song, K.; Tan, X.; Li, D.; Lu, W.; and Zhuang, Y. 2023. Hugginggpt: Solving ai tasks with chatgpt and its friends in hugging face. *Advances in Neural Information Processing Systems*, 36: 38154–38180.
- Shinn, N.; Cassano, F.; Gopinath, A.; Narasimhan, K.; and Yao, S. 2023. Reflexion: language agents with verbal reinforcement learning. In *Proceedings of the NeurIPS*. Red Hook, NY, USA: Curran Associates Inc.
- Su, W.; Miao, P.; Dou, H.; and Li, X. 2024. ScanFormer: Referring Expression Comprehension by Iteratively Scanning. In *2024 IEEE/CVF CVPR*, 13449–13458. Los Alamitos, CA, USA: IEEE Computer Society.
- Sun, G.; Jin, M.; Wang, Z.; Wang, C.-L.; Ma, S.; Wang, Q.; Geng, T.; Wu, Y. N.; Zhang, Y.; and Liu, D. 2025. Visual Agents as Fast and Slow Thinkers. In *The Thirteenth International Conference on Learning Representations*.
- Sun, M.; Xiao, J.; and Lim, E. G. 2021. Iterative Shrinking for Referring Expression Grounding Using Deep Reinforcement Learning. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 14055–14064.
- Sun, P.; Kretschmar, H.; Dotiwala, X.; Chouard, A.; Patnaik, V.; Tsui, P.; Guo, J.; Zhou, Y.; Chai, Y.; Caine, B.; Vasudevan, V.; Han, W.; Ngiam, J.; Zhao, H.; Timofeev, A.; Ettinger, S.; Krivokon, M.; Gao, A.; Joshi, A.; Zhang, Y.; Shlens, J.; Chen, Z.; and Anguelov, D. 2020. Scalability in Perception for Autonomous Driving: Waymo Open Dataset. In *Proceedings of the IEEE/CVF CVPR*.
- Wan, D.; Cho, J.; Stengel-Eskin, E.; and Bansal, M. 2024. Contrastive Region Guidance: Improving Grounding in Vision-Language Models Without Training. In *Computer Vision – ECCV 2024: 18th European Conference, Milan, Italy, September 29–October 4, 2024, Proceedings, Part LXXIX*, 198–215. Berlin, Heidelberg: Springer-Verlag. ISBN 978-3-031-72985-0.
- Wang, R.; Lei, T.; Cui, R.; Zhang, B.; Meng, H.; and Nandi, A. K. 2022. Medical image segmentation using deep learning: A survey. *IET image processing*, 16(5): 1243–1267.
- Watkins, C.; and Dayan, P. 1992. Technical Note: Q-Learning. *Machine Learning*, 8: 279–292.
- Wu, C.; Yin, S.; Qi, W.; Wang, X.; Tang, Z.; and Duan, N. 2023. Visual chatgpt: Talking, drawing and editing with visual foundation models. *arXiv preprint arXiv:2303.04671*.
- Wu, Q.; Bansal, G.; Zhang, J.; Wu, Y.; Li, B.; Zhu, E.; Jiang, L.; Zhang, X.; Zhang, S.; Liu, J.; et al. 2024. Autogen: Enabling next-gen LLM applications via multi-agent conversations. In *First Conference on Language Modeling*.
- Xia, Z.; Han, D.; Han, Y.; Pan, X.; Song, S.; and Huang, G. 2024. GSVA: Generalized segmentation via multimodal large language models. In *Proceedings of the IEEE/CVF CVPR*, 3858–3869.
- Yang, J.; Zhang, H.; Li, F.; Zou, X.; Li, C.; and Gao, J. 2023. Set-of-mark prompting unleashes extraordinary visual grounding in gpt-4v. *arXiv preprint arXiv:2310.11441*.
- Zhai, X.; Mustafa, B.; Kolesnikov, A.; and Beyer, L. 2023. Sigmoid loss for language image pre-training. In *Proceedings of the IEEE/CVF ICCV*, 11975–11986.
- Zou, X.; Yang, J.; Zhang, H.; Li, F.; Li, L.; Wang, J.; Wang, L.; Gao, J.; and Lee, Y. J. 2023. Segment everything everywhere all at once. In *Proceedings of NeurIPS*. Red Hook, NY, USA: Curran Associates Inc.