OpenWorldSAM: Extending SAM2 for Universal Image Segmentation with Language Prompts

Shiting Xiao

Yale University ginny.xiao@yale.edu

Donghyun Lee

Yale University donghyun.lee@yale.edu

Rishabh Kabra

Google DeepMind rkabra@google.com

João CarreiraGoogle DeepMind

Google DeepMind joaoluis@google.com

Yuhang Li

Yale University yuhang.li@yale.edu

Priyadarshini Panda

Yale University priya.panda@yale.edu

Abstract

The ability to segment objects based on open-ended language prompts remains a critical challenge, requiring models to ground textual semantics into precise spatial masks while handling diverse and unseen categories. We present OpenWorldSAM, a framework that extends the prompt-driven Segment Anything Model v2 (SAM2) to open-vocabulary scenarios by integrating multi-modal embeddings extracted from a lightweight vision-language model (VLM). Our approach is guided by four key principles: i) Unified prompting: OpenWorldSAM supports a diverse range of prompts, including category-level and sentence-level language descriptions, providing a flexible interface for various segmentation tasks. ii) Efficiency: By freezing the pre-trained components of SAM2 and the VLM, we train only 4.5 million parameters on the COCO-stuff dataset, achieving remarkable resource efficiency. iii) *Instance Awareness*: We enhance the model's spatial understanding through positional tie-breaker embeddings and cross-attention layers, enabling effective segmentation of multiple instances. iv) Generalization: OpenWorldSAM exhibits strong zero-shot capabilities, generalizing well on unseen categories and an open vocabulary of concepts without additional training. Extensive experiments demonstrate that OpenWorldSAM achieves state-of-the-art performance in open-vocabulary semantic, instance, and panoptic segmentation across multiple benchmarks. Code is available at https://github.com/GinnyXiao/OpenWorldSAM.

1 Introduction

Image segmentation has long been constrained to closed-vocabulary settings, where models can only recognize objects from a predefined taxonomy [1, 2, 3, 4, 5, 6, 7, 8]. However, real-world applications, e.g., Embodied AI [9, 10], demand systems that can understand open-ended language descriptions (from single nouns like "pedestrian" to rich referring expressions such as "the man in a red shirt") and segment novel objects unseen during training. This open-vocabulary segmentation problem poses two core challenges: (1) Semantic grounding – mapping free-form text to visual entities, and (2) Instance awareness – distinguishing multiple objects that match the same description.

Detection-centric methods [11, 12] relied on two-stage pipelines, first detecting class-agnostic mask proposals then classifying them with vision-language models (VLMs), e.g., CLIP [13] and ALIGN [14]. While effective, such approaches struggle with complex queries and specialize exclusively in semantic segmentation, lacking versatility. Recent generalist models [15, 16] explore unified architectures that jointly handle vision and language, allowing a single model to perform detection, segmentation, and grounding tasks. These generalist models demonstrate impressive flexibility, but they typically involve resource-intensive pre-training. The emergence of promptable segmentation models like the Segment Anything Model (SAM) [17, 18] offered new possibilities – it introduced



Figure 1: Overview of the proposed framework. The green region highlights the SAM v2 baseline, supporting visual prompts (e.g., boxes, points) for interactive segmentation. Our OpenWorldSAM extension integrates open-vocabulary language understanding, enabling both category-level segmentation across semantic, instance, panoptic tasks and referring expression segmentation.

a paradigm shift by allowing users to segment arbitrary objects using simple visual prompts (*e.g.*, points, boxes). Trained on an extensive dataset, these models exhibit remarkable generalization and interactive capabilities. However, they inherently lack semantic understanding. Subsequent attempts to combine SAM with large language models (LLMs) [19, 20, 21] achieved language awareness but at prohibitive computational costs, imposing overwhelming overhead.

We posit that an ideal open-vocabulary segmenter should: (i) Natively support textual prompts without cascaded classification components, (ii) Preserve the knowledge of the vision foundation models like SAM without adding large overhead, and (iii) Segment multiple possible instances that could correspond to a single query. To this end, we propose OpenWorldSAM, an open-vocabulary extension to the SAM v2 (SAM2) architecture that satisfies these requirements. OpenWorldSAM injects language understanding while retaining SAM2's core strengths through a lightweight language adapter (\approx 4.5M trainable parameters), unifing category-level instance, semantic, and panoptic segmentation, and sentence-level referring expression segmentation (Figure 1).

Specifically, we feed the image and descriptive text input into a frozen multi-modal encoder and obtain fused semantic representations. These serve as prompts to SAM2' mask decoder that produces masks for any described object or region. We introduce a positional tie-breaker mechanism to resolve ambiguities when a text query could apply to multiple regions, allowing the model to perform multi-instance segmentation. Furthermore, our adapter employs a soft prompting technique that uses cross-attention between textual queries and image features, sharpening localization by allowing semantic contexts to focus toward relevant image areas. By combining these design innovations, OpenWorldSAM can accurately identify and segment arbitrary objects described by text, all while using only frozen pre-trained encoders and a tiny trainable adaptation module.

In summary, OpenWorldSAM represents a new paradigm of "segment anything in the open world". It inherits SAM's interactiveness while being guided by flexible language prompts. Our contributions include:

- 1. We introduce OpenWorldSAM, a unified interface that supports various open-vocabulary segmentation tasks. We propose an efficient language adapter with tie-breaker and cross-attention soft prompting, improving multi-object localization.
- 2. OpenWorldSAM achieves state-of-the-art zeroshot performance across six benchmarks (Figure 2), setting a new standard for open-vocabulary segmentation (*e.g.*, 60.4 mIoU on ADE20K [22]). OpenWorldSAM also acheives strong performance in referring expression segmentation (74.0 cIoU on RefCOCOg [23]) with substantially fewer resources compared to recent models.
- Our work demonstrates that lightweight architectural interventions can unlock zero-shot segmentation capabilities rivaling specialized models while preserving SAM2's efficiency and interactivity.

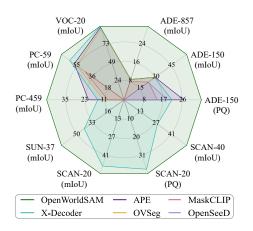
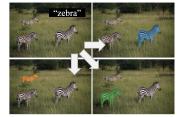


Figure 2: OpenWorldSAM achieves new state-of-the-art on six datasets with one suite of parameters.







(a) SAM's mask outputs with a point prompt

(b) OpenWorldSAM w/ tie-breaker

(c) wo/ tie-breake

Figure 3: (a) SAM takes a visual click and outputs 3 valid masks on the same person (the person, the backpack, and a backpack region) [17]. It will **not** output masks for the person standing next to her. (b) Tie-breakers shift the queries to distinct regions, enabling simultaneous segmentation of all three "zebra" instances. (c) Naïve approach [43]: A single language query for "zebra" causes SAM2 to segment only the most salient instance.

2 Related Work

Open-vocabulary segmentation. Recent advances in open-vocabulary segmentation have leveraged vision-language models (VLMs) [13, 14] to overcome the constraints of traditional closed-set segmentation models. Early approaches like LSeg [24], RegionCLIP [25] and OWL-ViT [26] established a baseline by introducing a contrastive learning framework to align image embeddings with CLIP-based text embeddings for zero-shot detection/segmentation. Subsequent methods [11, 27] scaled effectively by using weak supervision of large-scale images with captions (up to millions of regions) or text-only signals, enabling more flexible and broader semantic coverage. Two-stage approaches like MaskCLIP [28] and OVSeg [12] further refined this paradigm by generating mask proposals using MaskFormer [29] followed by CLIP-based classification, notably boosting accuracy through mask-adapted fine-tuning. Another line of works formulated this task as a visual grounding problem and established region-text fusion [30, 31, 32, 33]. More recently, unified architectures such as ODISE [34], X-Decoder [15], SEEM [35], OpenSeeD [36], HIPIE [37], Semantic-SAM [38] and APE [16] have integrated multiple segmentation tasks into a single framework, showing significant progress towards general-purpose models, but they typically required resource intensive pre-training.

Extending SAM for text-prompted segmentation. The Segment Anything Model (SAM) [17, 18] achieved a breakthrough in promptable segmentation by training on 1 billion masks, enabling it to generate high-quality masks for visual prompts. A flurry of recent works have explored infusing SAM with semantic or language understanding to move beyond its original prompt types. Grounded-SAM [39] is a pioneering effort that combines an open-vocabulary detector GroundingDINO [31] to generate bounding boxes from a text query, then feeds those boxes as prompts into SAM. Fast-SAM [40] matches CLIP embeddings with regions of interest. LLM-centric works [41, 20, 19, 21] attempt to map large LLMs or VLMs language embeddings into SAM or SAM-like decoder's prompt latent space to enable referring expression segmentation. Among these, LISA [19] pioneered "maskas-text-embedding" approach but was limited to single-object queries. LISA++ [42] introduced instance awareness through additional instruction-tuning data, though it requires LLMs to explicitly enumerate objects—a computationally expensive process. EVF-SAM [43] recently demonstrated a lightweight alternative, integrating SAM with a multi-modal BEiT-3 encoder [44] (673M parameters). While achieving state-of-the-art referring segmentation accuracy with minimal parameters, it remains constrained to single-object queries. Inspired by the success of EVF-SAM, we enhance SAM further into the domain of open-vocabulary segmentation, where the goal is to segment and label all objects ("things" and "stuff") in the scene with open-set categories.

3 Methodology

Motivation and key challenges. A fundamental limitation of SAM-like architectures is their inability to resolve multi-instance ambiguity from a single prompt. While visual prompts (e.g., points) may occasionally lack granularity specificity—for instance, a click on a backpack could imply segmentation of either the backpack or the entire person (Figure 3a)—they inherently localize to a single spatial region. Language prompts, however, introduce a distinct challenge: a text query like "zebra" may correspond to multiple spatially disjoint objects (Figure 3b), with no prior knowledge of instance counts. Prior attempts to add language capabilities either rely on segmentation-then-classification pipelines (losing end-to-end training) or require costly region-level text grounding

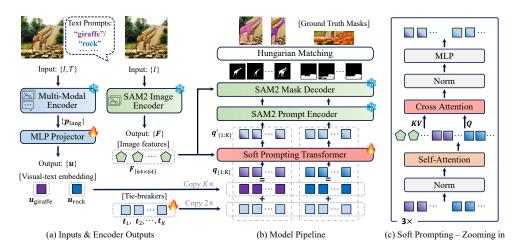


Figure 4: (a) Preliminaries on the inputs and outputs of the vision and multi-modal encoders. (b) OpenWorld-SAM pipeline. (c) Detailed soft prompting Transformer architecture.

during pre-training. Our key insight addresses this gap: SAM2's mask decoder can inherently segment multiple instances if equipped with *diverse positional guidance*, *i.e.*, learned cues that disentangle identical semantic queries into spatially distinct segmentation targets.

Architecture overview. Figure 4 depicts our framework which comprises: (i) a hierarchical SAM2 image encoder that extracts image features, (ii) a multi-modal vision-language encoder that jointly ingests the image and text prompt, (iii) a lightweight MLP projector, (iv) learnable positional tiebreakers for multi-instance queries, (v) a soft prompting Transformer block that aligns text-image features with SAM2's image features, and (vi) the SAM2 mask decoder producing final masks. Only a small language adapter with components (iii–v) is trained; all other backbones remain frozen.

Multi-modal encoder. We leverage BEiT-3 [44] to encode the input description into a semantic embedding. Given an image I and a text prompt T (e.g., a category name or a referring expression), we feed both modalities into BEiT-3's encoder to obtain joint visual—text embeddings. Concretely, tokens of T and patch embeddings of a downsampled I are concatenated and processed by BEiT-3, yielding a set of feature vectors $\{\mathbf{f}_{[\text{CLS}]}, \mathbf{f}_1, \dots\}$. We take the classification token $\mathbf{f}_{[\text{CLS}]}$ as a compact summary denoted as \mathbf{p}_{lang} of the prompt conditioned on the image content.

We adopt BEiT-3 because its early-fusion training on image-text pairs equips it with rich, bidirectional semantics—crucial for reasoning about unseen classes. Compared with CLIP-style contrastive image-text matching using only the features from the last encoder layers, BEiT-3 exposes finer cross-modal interactions. By embedding the text while it sees the image, the encoder already localizes the concept loosely (e.g., "giraffe" vs. "rock" in Figure 4) before any downstream segmentation, preventing the mask decoder from learning semantics from scratch.

Prompt projection. BEiT-3 emits 1,024-D tokens, whereas SAM's prompt channels are 256-D. A two-layer MLP acts as a projector that (i) preserves the coarse semantics of $\mathbf{p}_{lang} \in \mathbb{R}^{d_{1024}}$ and (ii) learns to highlight dimensions that are most useful for mask prediction: $\mathbf{u} = \mathrm{MLP}(\mathbf{p}_{lang}) \in \mathbb{R}^{d_{256}}$.

Positional tie-breaker and multi-instance queries generation. The projected visual-text embedding ${\bf u}$ captures what to segment but lacks awareness of how many instances exist and where they are. To enable multi-instance segmentation, we propose K learnable positional tie-breaker vectors $\{{\bf t}_1,\ldots,{\bf t}_K\}\subset \mathbb{R}^{d_{256}}$ that perturb ${\bf u}$ into K distinct queries:

$$\mathbf{q}_i = \mathbf{u} + \mathbf{t}_i, \quad i = 1, \dots, K. \tag{1}$$

These perturbations serve two purposes: 1) Positional disambiguation: Each t_i nudges the query towards different spatial regions (Figure 3b), mimicking how human annotators might click different points to segment each zebra. 2) Instance diversity: The tie-breakers are optimized during training to maximize coverage of distinct instances, preventing query collapse. Conceptually these queries play the role of the "object queries" in DETR [45]. Crucially, they impose segmentation distinction for the same language semantics, making positional tie-breaking a novel and key feature for OpenWorldSAM. In practice K = 20 covers >99% images in COCO [46]; for larger scenes K can be increased trivially.

Soft-prompting via cross-attention. The perturbed queries $\{\mathbf{q}_i\}$ interact with SAM2's image features through a 3-layer Transformer [47] in Figure 4c, which alternates self-attention (queries talk to each other, promoting diversity) and cross-attention (queries look at image features). Each language-aware query is refined on-the-fly by cross-attention with the frozen SAM2 features. SAM2's image encoder follows a hierarchical vision Transformer ("Hiera" [48, 49]) that outputs three features $\{\mathbf{F}_{256\times256},\mathbf{F}_{128\times128},\mathbf{F}_{64\times64}\}$ with $256^2,128^2$, and 64^2 spatial resolutions, respectively. We operate on the level-3 features with 64^2 resolution as they optimally balance precision for retaining boundary details and computational efficiency $(16\times$ cheaper than full-resolution attention). They are also used by SAM2 for mask decoding by default [18]. The soft prompting Transformer computes $\mathbf{q}_i' = \text{CrossAttn}(\mathbf{q}_i, \mathbf{F}_{64\times64}), \ i = 1, \ldots, K$, whose key/value inputs are the flattened level-3 features $\mathbf{F}_{64\times64} \in \mathbb{R}^{4096\times256}$. This step grounds the language-aware queries in SAM2's high-resolution visual features, resolving ambiguities (e.g., distinguishing adjacent zebras by stripe patterns).

Mask decoding and class assignment. The refined queries $\{\mathbf{q}_i'\}$ are input to SAM2's mask decoder alongside level-3 image features. We inject the queries as the prompt tokens in place of, e.g., point or box prompts in the original SAM2's prompt encoder to obtain prompt embeddings. The prompt embeddings are then passed to the mask decoder which outputs K masks and corresponding confidence scores. We assign each mask the original text prompt T as its class label, since the generation is fully conditioned on T and thus inherits the semantic identity.

Training. All heavy visual (Hiera) and vision-language (BEiT-3) encoders are kept frozen to preserve their pre-trained knowledge and avoid costly retraining. Only the MLP projector, tie-breakers, and the soft prompting Transformer are learnable. For each training sample and prompt, we match the K predicted masks to the ground-truth masks of class T via Hungarian matching [45], then apply a focal loss, encouraging precise segmentation of all instances described by the prompt. The tie-breakers $\mathbf{t}_i \in \mathbb{R}^{d_{256}}$ are implemented as learnable parameters randomly initialized from a normal distribution. During training, the Hungarian matching loss naturally encourages each \mathbf{t}_i to specialize in different spatial regions. Notably, this mechanism requires no explicit supervision about instance counts.

Inference. From the predicted K masks, we derive results for three segmentation tasks: semantic, instance, and panoptic. For semantic segmentation, we merge masks sharing the same class label, weighted by their confidence scores. For instance segmentation, we apply confidence-score filtering to remove masks below a certain threshold, followed by non-maximum suppression (NMS) to eliminate highly overlapping masks and retain distinct object instances. Similarly, for panoptic segmentation, we perform confidence-based filtering and NMS, ensuring each pixel is uniquely assigned to either a "thing" (instance) or "stuff" (semantic) label.

Optionally, we perform a two-stage inference. In this setup, masks obtained from the first inference stage are used as visual prompts fed back into SAM2's mask decoder, which refines mask contours. Qualitatively, two-stage inference improves the precision of mask boundaries for correct predictions (Appendix B). However, quantitative analysis (Table 1) reveals that the second inference stage provides minimal improvements in segmentation metrics, suggesting it mainly enhances visual quality rather than overall accuracy.

4 Experiments

Datasets and metrics. We train OpenWorldSAM on the COCO2017-Stuff [46] dataset with panoptic annotations, excluding the RefCOCOg UMD [23] validation set, following X-Decoder [15]. The training set contains 104k images. We evaluate the model in a zero-shot setting on eight segmentation tasks across five diverse datasets: ADE20K-150/857 [22], PASCAL VOC-20 [50], PASCAL Context-59/459 [51], ScanNet-20/40 [52], and SUN-RGBD-37 [53]. Evaluation metrics include panoptic quality (PQ), mean average precision (mAP), and mean intersection-over-union (mIoU), corresponding to panoptic, instance, and semantic segmentation tasks, respectively. For referring segmentation, we pre-train the model on COCO2017-stuff and finetune on RefCOCOg UMD training split. Following prior works, we report the cumulative intersection over the cumulative union (cIoU) metric on the RefCOCOg UMD validation split.

Implementation. We implement our model in PyTorch. We initialize the model with the public weights of SAM2-Hiera-Large and BEIT-3-Large. It is trained for 25 epochs on COCO-Stuff using the AdamW optimizer with a learning rate of 1e-4, batch size 8, on a single NVIDIA A100 GPU. Image resolution is set to 1024 for SAM2 and 224 for BEiT-3. Number of postional tie-breaks is set to 20 for COCO dataset. Our implementation details can be found in Appendix A.



Figure 5: Qualitative comparisons on ADE20K-857. In many cases, (e.g., (c) road, field), X-Decoder predicts semantically related but incorrect labels due to ambiguity in the category list. The final column shows X-Decoder predictions using oracle prompts, which reduces confusion. OpenWorldSAM, conditioned on the correct prompt, produces faithful masks and avoids semantic mismatches. Color maps for each model vary. Please refer to the predicted labels. Best viewed with zoom in. We use two-stage inference for the visualization.

4.1 Open-Vocabulary Segmentation Evaluation Protocols and Challenges

Ambiguity of open vocabulary evaluation. Most prior open-vocabulary segmentation methods—including X-Decoder [15], OVSeg [12], and MaskCLIP [28]—adopt a Global-Matching protocol: for each predicted mask, a model matches it against the entire dataset vocabulary using precomputed text embeddings and selects the best-aligned class. However, this strategy can be problematic when applied to datasets like ADE20K, which contain hundreds of fine-grained and overlapping labels. As observed in OVSeg [12], this leads to semantically reasonable predictions being marked incorrect under exact label matching: "The ground-truth category is 'building' while our model predicts 'skyscraper'." This ambiguity stems from the inherent subjectivity of language: synonymous or closely related concepts may be indistinguishable in a visual context, yet only one is accepted by the ground truth. We observe similar issues in our own qualitative analysis. As shown in Figure 5, X-Decoder predictions on ADE20K-857 often produce valid but non-canonical labels (e.g., 'road' instead of 'runway', or 'screen' instead of 'arcade machine'), resulting in unfair penalization.

Oracle-Prompts evaluation. To address this, we introduce an alternative evaluation strategy: **Oracle Prompts**—during evaluation, we explicitly provide the ground-truth class names as prompts. This mimics the intended use case of **prompt-based models** like SAM, which are inherently interactive and conditioned on user input. Under this protocol, the model does not have to resolve linguistic ambiguity across the full label space; it segments *what the user asks for.* We report results under both settings: Table 1 shows baseline performance using the global matching protocol, consistent with prior works. Table 2 revisits X-Decoder under the oracle-prompt protocol for a more equitable comparison to OpenWorldSAM, which by design is evaluated under oracle prompts. We believe this approach provides a more fair assessment of SAM-style models in open-vocabulary segmentation.

4.2 Open-Vocabulary Segmentation Performance Analysis

Zero-shot open-vocabulary transfer. OpenWorldSAM generalizes out-of-the-box to a broad set of segmentation tasks without any weight adaptation. As shown in Table 1, it achieves state-of-the-art performance across almost all datasets and evaluation metrics. Its performance consistently surpasses strong baselines such as X-Decoder and APE, despite using only 4.5M trainable parameters. On ADE20K-857, OpenWorldSAM achieves 33.1% mIoU, outperforming the previous best (X-Decoder) by +23.9 absolute points ($9.2 \rightarrow 33.1$). On PASCAL Context-459, it achieves 47.5% mIoU,

Table 1: Zero-shot performance of open-vocabulary segmentation models across multiple benchmarks. For COCO, different methods use different supervisions of mask (m), class label (cls) and caption (cap). "ITP" indicates whether model uses image-text pairs/referring data. "DET" indicates extra detection data (*e.g.*, Objects365, LVIS, OpenImages, *etc.*) "*" denotes the model has the capability for the task but does not have number reported. "-" means the model does not have the ability for the specific task. Purple color means a fully supervised approach, and tan a semi-supervised learning approach. Two-stage inference means we refine mask contours by re-prompting SAM using the raw mask predictions. Bold entries indicate the best performance.

Model	Train Params			(p/s)	ITP	DET	_	ADE-1		ADE-857	VOC-20	PC-59	PC-459	SUN-37	SCA		SCAN-40
		m	cls	cap			PQ	mAP	mIoU	mIoU	mIoU	mIoU	mIoU	mIoU	mIoU	PQ	mIoU
MSeg (B) [54]	70 (M)	1	~	~	×	×	33.7	32.6	19.1	*	73.4	43.4	*	29.6	33.4	*	*
GroupViT (S) [27]	44 (M)	×	X	×	~	×	-	-	*	*	52.3	22.4	*	*	*	-	*
LSeg+ (B) [24]	112 (M)	~	~	~	×	×	-	-	18.0	3.8	*	46.5	7.8	*	*	-	*
ZegFormer (B) [55]	60 (M)	~	~	~	~	×	-	-	*	8.1	80.7	*	*	*	*	-	-
OpenSeg (B) [11]	86 (M)	~	X	~	~	×	-	-	26.4	8.1	70.2	44.8	11.5	*	*	-	*
OVSeg (B) [12]	0.6 (M)	V	~	×	×	×	-	-	29.6	9.0	94.5	55.7	12.4	*	*	-	*
MaskCLIP (L) [28]	428 (M)	~	~	×	×	×	15.1	6.0	23.7	8.2	*	45.9	10.0	*	*	*	*
OpenSeeD (L) [36]	39 (M)	~	X	~	~	~	19.7	15.0	23.4	*	*	*	*	*	*	*	*
X-Decoder-Seg ⁺ (B) [15]	28 (M)	~	~	×	×	×	16.9	9.5	23.8	4.6	97.8	64.7	12.1	32.2	35.1	33.8	18.5
X-Decoder (L) [15]	38 (M)	V	~	~	~	×	21.8	13.1	29.6	9.2	97.7	64.0	16.1	43.0	49.5	39.5	29.7
APE-B (L) [16]	42 (M)	~	~	~	~	~	26.4	23.5	29.0	9.2	95.8	58.3	21.0	*	*	*	*
ESC-Net [56]	451 (M)	•	•	×	×	×	-	-	41.8	18.1	98.3	65.6	27.0	*	*	-	*
OpenWorldSAM	4.5 (M)	~	~	×	×	×	35.2	16.9	60.4	33.1	98.0	73.7	47.5	67.7	65.0	41.9	55.6
+ two-stage inference	4.5 (M)	~	~	×	×	x	36.3	15.6	58.0	32.6	97.6	72.6	45.8	68.2	64.8	39.9	54.1

Table 2: Oracle-Prompts evaluation of open-vocabulary segmentation models. We report the state-of-the-art (SOTA) model X-Decoder [15]'s performance under both evaluation protocols. Other methods are omitted either because: 1) they are not SOTA, or 2) they do not support oracle-prompts evaluation.

Model	Evaluation Protocol	ADE-150 mIoU	ADE-857 mIoU	VOC-20 mIoU	PC-59 mIoU	PC-459 mIoU	SUN-37 mIoU	SCAN-40 mIoU
X-Decoder (L) [15]	Global-Matching (default)	29.6	9.2	97.7	64.0	16.1	43.0	29.7
X-Decoder (L)	Oracle-Prompts	51.5	29.1	98.1	75.5	42.3	67.1	49.1
OpenWorldSAM	Oracle-Prompts (default)	60.4	33.1	98.0	73.7	47.5	67.7	55.6

improving over APE's 21.8% by +25.7 points, and on ScanNet-40, it reaches 55.6% mIoU, a +25.9 point improvement over X-Decoder's 29.7%. On AP score we under-perform APE, which included extra detection datasets, *e.g.*, Objects 365 [57], in their training recipe for better localization.

We attribute our strong performance to the model's prompt-conditioned decoding mechanism, which directly leverages language input to guide mask prediction. This is particularly advantageous when the target concept is known at query time. In contrast, global retrieval-based models such as X-Decoder must resolve ambiguity across the entire vocabulary space, which introduces classification error. While one might argue that differing evaluation protocols confound the comparison, it's important to note that both families of models require the same semantic input—the only difference lies in **when** and **how** that input is used.

Oracle-Prompts evaluation. As SAM-style models are designed for interactive segmentation, oracle prompts closely reflect practical use cases—such as human-in-the-loop annotation, robotic object search, or dynamic UI feedback. To fairly compare with the state-of-the-art generalist model X-Decoder [15], we also evaluate it under oracle prompts: we restrict its vocabulary to the ground-truth classes for each image. As shown in Table 2, OpenWorldSAM continues to outperform even under these controlled conditions. Notably, on large-vocabulary datasets such as ADE20K-857 and PASCAL Context-459, OpenWorldSAM achieves 33.1% and 47.5% mIoU, surpassing X-Decoder by +4.0 and +5.2 points, respectively. This highlights our model's superior language grounding ability in long-tailed, fine-grained category distributions. On smaller datasets like PASCAL Context-59 and PASCAL VOC-20, where most categories overlap with COCO, X-Decoder slightly outperforms our model (75.5% vs. 73.7% mIoU and 98.1% vs. 98.0%), suggesting it benefits more from class memorization in such settings. Moreover, Figure 5 illustrates that global matching often fails despite producing correct masks. Conditioning on oracle prompts significantly reduces this ambiguity, highlighting the robustness of our evaluation protocol and the effectiveness of prompt-based segmentation.

Qualitative Results. Figure 5 presents example outputs of OpenWorldSAM on challenging scenes, with comparisons to X-Decoder under both evaluation protocols. In one example 5(a), an image from ADE20K-857 containing a game room scene is segmented by our model using prompts for



Figure 6: Qualitative results on RefCOCOg. OpenWorldSAM is capable of understanding spatial relationship, colors, actions, and shapes, *etc*.

Table 3: Referring segmentation performance (cIoU) comparison on RefCOCOg benchmark validation set between our proposed OpenWorldSAM and prior SOTA methods. We abbreviate the datasets: C (COCO), RC (RefCOCO/+), RCg (RefCOCOg), PL (PACO-LVIS), O365 (Objects365), V (Video segmentation datasets), OID (OpenImages Detection), VG (Visual Genome), ADE (ADE20K), PP (PASCAL-Part), PC (PASCAL-VOC). We compare model trainable parameters, model capabilities (OV seg (open-vocabulary segmentation) and Inter Seg (interactive segmentation)), and training data required. "*" denotes an estimate of the trainable parameters, since these models use LoRA [58] with rank-8/16 adapters for finetuning.

Method	Foundation Model	Train Params	w/ SAM?	OV Seg?	Inter Seg?	Training Data	cIoU
X-Decoder (L) [15]	CLIP-B [13] (63M)	38 (M)	×	~	×	C, RCg, Cap4M	64.6
SEEM (L) [35]	CLIP-B [13] (63M)	39 (M)	×	~	~	C, RC, RCg, PL	65.6
PolyFormer (L) [59]	BERT-B [60] (104M)	450 (M)	×	×	×	RC, RCg	71.2
UNINEXT (H) [32]	BERT-B [60] (104M)	673 (M)	×	~	×	C, RC, O365, V	74.4
APE-B (L) [16]	CLIP-L [13] (123M)	42 (M)	×	~	×	C, PC, O365, OID, VG, RC, RCg	63.5
PixelLM [61]	LLaMA2 [62] (13B)	29 (M)*	×	~	×	C, RC, ADE, PL, MUSE	69.3
LISA [19]	Vicuna [63] (7B)	32 (M)*	~	×	×	C, RC, ADE, PL, PP	66.4
GLaMM [20]	Vicuna [63] (7B)	$40 (M)^*$	✓	~	×	RC, GranD	74.2
u-LLaVA [21]	Vicuna [63] (7B)	$44 (M)^*$	~	×	×	C, RC, ADE, PL, PC	71.6
u-LLaVA [21]	Vicuna [63] (7B)	7 (B)	~	×	×	C, RC, ADE, PL, PC	74.8
Sa2VA [64]	InternVL2 [65] (1B)	$22 (M)^*$	~	×	~	RC, RCg, V, GranD	72.3
Sa2VA [64]	InternVL2 [65] (4B)	35 (M)*	~	×	~	RC, RCg, V, GranD	74.1
EVF-SAM [43]	BEIT-3-L [44] (673M)	674 (M)	~	×	×	RC	76.8
OpenWorldSAM	BEIT-3-L [44] (673M)	4.5 (M)	~	~	•	C, RCg	74.0

various objects ("ceiling, light, seat, person, arcade machine"). OpenWorldSAM accurately masks each object and stuff region, whereas X-Decoder misclassifies the "arcade machine" due to confusion between similar semantic objects under Global-Matching, and produces fragmented masks for the person and seat under Oracle-Prompts. Similarly in example 5(b), X-Decoder misclassifies the "wall" and proposes object masks for prompts that did not exist in the ground truth (e.g., "window glass") under Global-Matching, and failed to segment "plant" under Oracle-Prompts. This showcases our model's clear understanding of category semantics (thanks to the VLM prompt) combined with precise mask delineation (thanks to SAM2's capability). More qualitative results in Appendix C.

4.3 Referring Expression Segmentation Performance Analysis

Performance. As shown in Table 3 and Figure 6, OpenWorldSAM achieves strong performance on the RefCOCOg validation set, obtaining a cIoU of 74.0%, significantly outperforming earlier generalist models like SEEM and X-Decoder (≈65%), and competitive with specialized models such as GLaMM (71.2%) and UNINEXT (74.4%). Notably, OpenWorldSAM reaches this accuracy using just BEiT-3 encoder with 673M parameters and an additional 4.5M trainable parameters, substantially fewer than recent large-scale models like LISA, GLaMM, and u-LLaVA, which rely on much larger vision-language foundations (7B+ parameters) and multiple additional datasets. While EVF-SAM achieves higher cIoU (76.8%), this advantage stems from training on twice the referring data (full RefCOCO series vs. our RefCOCOg subset). Crucially, OpenWorldSAM inherits SAM's interactive features, offering unique flexibility in multiple segmentation tasks, which distinguishes it from higher-scoring yet narrower models.

4.4 Ablation Studies

We systematically validate OpenWorldSAM's design through zero-shot transfer on ADE20K-150/857 benchmark and fintuning on RefCOCOg benchmark.

Table 4: Ablation on the VLM choice, *e.g.*, CLIP [13] model from OpenAI. '✔' denotes modality used during training, and 'X' unused. Only the adapter modules are trainable, and the VLMs are kept frozen. Late fusion means we concatenate text/image features from the last layers of CLIP's text/image encoder. Early fusion means BEiT-3 processes both modalities in all 24 Transformer layers.

Encoder	Params	Text	Image	Modality Fusion	PQ	ADE-1 AP	50 mIoU		RefCOCOg cIoU
CLIP-Large	123 (M)	~	×	_	13.5	2.9	25.7	12.8	25.2
CLIP-Large	428 (M)	~	~	Late (Last-layer Concat)	14.0	3.6	26.5	14.0	25.3
BEiT-3-Large	370 (M)	~	×	-	13.6	3.1	26.3	13.3	26.1
BEiT-3-Large	673 (M)	~	~	Early (All-layer Attention)	35.2	16.9	60.4	33.1	74.0

Table 5: Ablation on trainable and inference modules. For training, '✔' denotes trainable, and 'X' denotes frozen. For inference, '✔' denotes activate, and 'X' denotes non-activate.

Evn		Tra	in Modules		Train Params	Infer	ence Mo	lules	A	ADE-1	50	ADE-857
Exp	Tie-breaker	BEiT-3	Cross-Attn	MLP Projector	Hain Faranis	Tie-breaker	BEiT-3	Cross-Attn	PQ	AP	mIoU	mIoU
E1	×	×	×	v	1.2 (M)	×	~	×	0.4	1.0	1.2	0.2
E2	~	×	×	~	1.3 (M)	~	×	×	-	9.5	-	-
E3	~	×	×	~	1.3 (M)	·	~	×	35.1	17.1	56.8	32.2
E4	~	~	×	~	674.0 (M)	·	~	×	13.6	3.5	24.4	10.6
E5	~	×	~	~	4.5 (M)	~	~	/	35.2	16.9	60.4	33.1
E6	~	~	~	✓	677.2 (M)	~	~	~	15.9	3.8	23.6	11.2

Multi-modal encoder analysis. In Table 4, we compares performances using different VLM encoders and fusion methods (early fusion vs. late fusion). BEiT-3's early cross-modal fusion (joint text-image processing across all layers) outperforms CLIP's late fusion (last-layer concatenation) by +33.9 mIoU, +21.2 PQ, and +13.3 AP on ADE-150, demonstrating that deep semantic integration is critical for aligning language concepts with visual regions, echoing findings by EVF-SAM [43].

Visual Context Matters. Table 4 demonstrates that removing visual inputs to BEiT-3 (text-only) causes catastrophic performance collapse (-34.4 mIoU on ADE-150). This confirms that SAM's segmentation backbone cannot ground textual semantics without explicit visual-textual co-encoding.

Optimal Training Strategy. In Table 5, we varied the trainable modules in OpenWorldSAM (thus varying total new parameters from 1.2M to 770M). We found in E5 that freezing BEiT-3 and training only the language adapter module (tie-breaker + cross-attention, 4.5M parameters) yields optimal performance (60.4 mIoU ADE-150). Notably, comparing E6 vs E5 and E4 vs E3, we found fine-tuning the entire BEiT-3 encoder (673M parameters) significantly degrades accuracy (mIoU drops from 60.4 to 23.6), likely due to underfitting on sparse category label prompts compared to its original web-scale pretraining.

Positional tie-breaker vs. none. Comparing E3 vs E1 in Table 5, positional tie-breaker boosts AP from 1.0% to 17.1%. As shown in Figure 3, without the tie-breaker, the model usually collapses on one instance of the class (especially if the one was particularly salient among others). This confirms the necessity of this component for reliable instance segmentation.

Cross-Attention layer removal. As shown in Table 5, E5 vs E3, removing the cross-attention layers expectedly led to inferior performance (-1.5 mIoU on ADE-150 and -0.9 mIoU on ADE-857). This indicates that cross-attention helps align prompts to the intended visual regions.

5 Conclusion

OpenWorldSAM bridges the gap between promptable segmentation and open-vocabulary understanding by unifying SAM's segmentation prowess with vision-language models' semantic grounding. This approach generalizes across tasks (semantic/instance/panoptic) and prompts (nouns/sentences), offering practitioners a unified tool for real-world scenarios where novel objects and ambiguous queries are the norm. Three innovations drive this success: (1) Positional tie-breakers enable multi-instance segmentation from single-text queries, resolving a critical limitation of SAM-like architectures. (2) Cross-modal soft prompting dynamically aligns language semantics with SAM's visual space, ensuring precise localization without costly LLMs. (3) Frozen foundation synergy leverages pretrained knowledge from SAM and BEiT-3, proving that dense prediction tasks benefit as much as

classification from parameter-efficient adaptation. Beyond technical contributions, OpenWorldSAM advances a paradigm for extending segmentation foundations: instead of training monolithic models, strategic adaptation of frozen components achieves open-world readiness at minimal cost.

Acknowledgement. This work was supported in part by CoCoSys, a JUMP2.0 center sponsored by DARPA and SRC, the National Science Foundation (CAREER Award, Grant #2312366, Grant #2318152), the DARPA Young Faculty Award and the DoE MMICC center SEA-CROGS (Award #DE-SC0023198).

References

- [1] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015.
- [2] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention–MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18*, pages 234–241. Springer, 2015.
- [3] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):834–848, 2017.
- [4] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 801–818, 2018.
- [5] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017.
- [6] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. *Advances in neural information processing systems*, 34:12077–12090, 2021.
- [7] Bowen Cheng, Ishan Misra, Alexander G Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1290–1299, 2022.
- [8] Feng Li, Hao Zhang, Huaizhe Xu, Shilong Liu, Lei Zhang, Lionel M Ni, and Heung-Yeung Shum. Mask dino: Towards a unified transformer-based framework for object detection and segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3041–3050, 2023.
- [9] Lei Fan, Jianxiong Zhou, Xiaoying Xing, and Ying Wu. Active open-vocabulary recognition: Let intelligent moving mitigate clip limitations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16394–16403, 2024.
- [10] Sebastián Barbas Laina, Simon Boche, Sotiris Papatheodorou, Simon Schaefer, Jaehyung Jung, and Stefan Leutenegger. Findanything: Open-vocabulary and object-centric mapping for robot exploration in any environment. *arXiv preprint arXiv:2504.08603*, 2025.
- [11] Golnaz Ghiasi, Xiuye Gu, Yin Cui, and Tsung-Yi Lin. Scaling open-vocabulary image segmentation with image-level labels. In *European conference on computer vision*, pages 540–557. Springer, 2022.
- [12] Feng Liang, Bichen Wu, Xiaoliang Dai, Kunpeng Li, Yinan Zhao, Hang Zhang, Peizhao Zhang, Peter Vajda, and Diana Marculescu. Open-vocabulary semantic segmentation with mask-adapted clip. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7061–7070, 2023.

- [13] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PmLR, 2021.
- [14] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International conference on machine learning*, pages 4904–4916. PMLR, 2021.
- [15] Xueyan Zou, Zi-Yi Dou, Jianwei Yang, Zhe Gan, Linjie Li, Chunyuan Li, Xiyang Dai, Harkirat Behl, Jianfeng Wang, Lu Yuan, et al. Generalized decoding for pixel, image, and language. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 15116–15127, 2023.
- [16] Yunhang Shen, Chaoyou Fu, Peixian Chen, Mengdan Zhang, Ke Li, Xing Sun, Yunsheng Wu, Shaohui Lin, and Rongrong Ji. Aligning and prompting everything all at once for universal visual perception. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13193–13203, 2024.
- [17] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In Proceedings of the IEEE/CVF international conference on computer vision, pages 4015–4026, 2023.
- [18] Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, et al. Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*, 2024.
- [19] Xin Lai, Zhuotao Tian, Yukang Chen, Yanwei Li, Yuhui Yuan, Shu Liu, and Jiaya Jia. Lisa: Reasoning segmentation via large language model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9579–9589, 2024.
- [20] Hanoona Rasheed, Muhammad Maaz, Sahal Shaji, Abdelrahman Shaker, Salman Khan, Hisham Cholakkal, Rao M Anwer, Eric Xing, Ming-Hsuan Yang, and Fahad S Khan. Glamm: Pixel grounding large multimodal model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13009–13018, 2024.
- [21] Jinjin Xu, Liwu Xu, Yuzhe Yang, Xiang Li, Fanyi Wang, Yanchun Xie, Yi-Jie Huang, and Yaqian Li. u-llava: Unifying multi-modal tasks via large language model. In *ECAI 2024*, pages 618–625. IOS Press, 2024.
- [22] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 633–641, 2017.
- [23] Varun K Nagaraja, Vlad I Morariu, and Larry S Davis. Modeling context between objects for referring expression understanding. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part IV 14*, pages 792–807. Springer, 2016.
- [24] Boyi Li, Kilian Q Weinberger, Serge Belongie, Vladlen Koltun, and René Ranftl. Language-driven semantic segmentation. *arXiv preprint arXiv:2201.03546*, 2022.
- [25] Yiwu Zhong, Jianwei Yang, Pengchuan Zhang, Chunyuan Li, Noel Codella, Liunian Harold Li, Luowei Zhou, Xiyang Dai, Lu Yuan, Yin Li, et al. Regionclip: Region-based language-image pretraining. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16793–16803, 2022.
- [26] Matthias Minderer, Alexey Gritsenko, Austin Stone, Maxim Neumann, Dirk Weissenborn, Alexey Dosovitskiy, Aravindh Mahendran, Anurag Arnab, Mostafa Dehghani, Zhuoran Shen, et al. Simple open-vocabulary object detection. In *European conference on computer vision*, pages 728–755. Springer, 2022.

- [27] Jiarui Xu, Shalini De Mello, Sifei Liu, Wonmin Byeon, Thomas Breuel, Jan Kautz, and Xiaolong Wang. Groupvit: Semantic segmentation emerges from text supervision. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 18134–18144, 2022.
- [28] Xiaoyi Dong, Jianmin Bao, Yinglin Zheng, Ting Zhang, Dongdong Chen, Hao Yang, Ming Zeng, Weiming Zhang, Lu Yuan, Dong Chen, et al. Maskclip: Masked self-distillation advances contrastive language-image pretraining. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10995–11005, 2023.
- [29] Bowen Cheng, Alex Schwing, and Alexander Kirillov. Per-pixel classification is not all you need for semantic segmentation. *Advances in neural information processing systems*, 34:17864–17875, 2021.
- [30] Liunian Harold Li, Pengchuan Zhang, Haotian Zhang, Jianwei Yang, Chunyuan Li, Yiwu Zhong, Lijuan Wang, Lu Yuan, Lei Zhang, Jenq-Neng Hwang, et al. Grounded language-image pre-training. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10965–10975, 2022.
- [31] Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Qing Jiang, Chunyuan Li, Jianwei Yang, Hang Su, et al. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. In *European Conference on Computer Vision*, pages 38–55. Springer, 2024.
- [32] Bin Yan, Yi Jiang, Jiannan Wu, Dong Wang, Ping Luo, Zehuan Yuan, and Huchuan Lu. Universal instance perception as object discovery and retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15325–15336, 2023.
- [33] Lewei Yao, Jianhua Han, Youpeng Wen, Xiaodan Liang, Dan Xu, Wei Zhang, Zhenguo Li, Chunjing Xu, and Hang Xu. Detclip: Dictionary-enriched visual-concept paralleled pre-training for open-world detection. Advances in Neural Information Processing Systems, 35:9125–9138, 2022.
- [34] Jiarui Xu, Sifei Liu, Arash Vahdat, Wonmin Byeon, Xiaolong Wang, and Shalini De Mello. Open-vocabulary panoptic segmentation with text-to-image diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2955–2966, 2023.
- [35] Xueyan Zou, Jianwei Yang, Hao Zhang, Feng Li, Linjie Li, Jianfeng Wang, Lijuan Wang, Jianfeng Gao, and Yong Jae Lee. Segment everything everywhere all at once. *Advances in neural information processing systems*, 36:19769–19782, 2023.
- [36] Hao Zhang, Feng Li, Xueyan Zou, Shilong Liu, Chunyuan Li, Jianwei Yang, and Lei Zhang. A simple framework for open-vocabulary segmentation and detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1020–1031, 2023.
- [37] Xudong Wang, Shufan Li, Konstantinos Kallidromitis, Yusuke Kato, Kazuki Kozuka, and Trevor Darrell. Hierarchical open-vocabulary universal image segmentation. *Advances in Neural Information Processing Systems*, 36:21429–21453, 2023.
- [38] Feng Li, Hao Zhang, Peize Sun, Xueyan Zou, Shilong Liu, Jianwei Yang, Chunyuan Li, Lei Zhang, and Jianfeng Gao. Semantic-sam: Segment and recognize anything at any granularity. arXiv preprint arXiv:2307.04767, 2023.
- [39] Tianhe Ren, Shilong Liu, Ailing Zeng, Jing Lin, Kunchang Li, He Cao, Jiayu Chen, Xinyu Huang, Yukang Chen, Feng Yan, et al. Grounded sam: Assembling open-world models for diverse visual tasks. arXiv preprint arXiv:2401.14159, 2024.
- [40] Xu Zhao, Wenchao Ding, Yongqi An, Yinglong Du, Tao Yu, Min Li, Ming Tang, and Jinqiao Wang. Fast segment anything. *arXiv preprint arXiv:2306.12156*, 2023.
- [41] Yonglin Li, Jing Zhang, Xiao Teng, Long Lan, and Xinwang Liu. Refsam: Efficiently adapting segmenting anything model for referring video object segmentation. *arXiv* preprint *arXiv*:2307.00997, 2023.

- [42] Senqiao Yang, Tianyuan Qu, Xin Lai, Zhuotao Tian, Bohao Peng, Shu Liu, and Jiaya Jia. An improved baseline for reasoning segmentation with large language model. *CoRR*, 2023.
- [43] Yuxuan Zhang, Tianheng Cheng, Lianghui Zhu, Rui Hu, Lei Liu, Heng Liu, Longjin Ran, Xiaoxin Chen, Wenyu Liu, and Xinggang Wang. Evf-sam: Early vision-language fusion for text-prompted segment anything model. *arXiv* preprint arXiv:2406.20076, 2024.
- [44] Wenhui Wang, Hangbo Bao, Li Dong, Johan Bjorck, Zhiliang Peng, Qiang Liu, Kriti Aggarwal, Owais Khan Mohammed, Saksham Singhal, Subhojit Som, et al. Image as a foreign language: Beit pretraining for vision and vision-language tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19175–19186, 2023.
- [45] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European conference on computer vision*, pages 213–229. Springer, 2020.
- [46] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer vision–ECCV 2014: 13th European conference, zurich, Switzerland, September 6-12, 2014, proceedings, part v 13*, pages 740–755. Springer, 2014.
- [47] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [48] Chaitanya Ryali, Yuan-Ting Hu, Daniel Bolya, Chen Wei, Haoqi Fan, Po-Yao Huang, Vaibhav Aggarwal, Arkabandhu Chowdhury, Omid Poursaeed, Judy Hoffman, et al. Hiera: A hierarchical vision transformer without the bells-and-whistles. In *International conference on machine learning*, pages 29441–29454. PMLR, 2023.
- [49] Daniel Bolya, Chaitanya Ryali, Judy Hoffman, and Christoph Feichtenhofer. Window attention is bugged: how not to interpolate position embeddings. *arXiv preprint arXiv:2311.05613*, 2023.
- [50] Mark Everingham and John Winn. The pascal visual object classes challenge 2012 (voc2012) development kit. *Pattern Analysis*, *Statistical Modelling and Computational Learning*, *Tech. Rep*, 8(5):2–5, 2011.
- [51] Roozbeh Mottaghi, Xianjie Chen, Xiaobai Liu, Nam-Gyu Cho, Seong-Whan Lee, Sanja Fidler, Raquel Urtasun, and Alan Yuille. The role of context for object detection and semantic segmentation in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 891–898, 2014.
- [52] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5828–5839, 2017.
- [53] Shuran Song, Samuel P Lichtenberg, and Jianxiong Xiao. Sun rgb-d: A rgb-d scene understanding benchmark suite. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 567–576, 2015.
- [54] John Lambert, Zhuang Liu, Ozan Sener, James Hays, and Vladlen Koltun. Mseg: A composite dataset for multi-domain semantic segmentation. In *Proceedings of the IEEE/CVF conference* on computer vision and pattern recognition, pages 2879–2888, 2020.
- [55] Jian Ding, Nan Xue, Gui-Song Xia, and Dengxin Dai. Decoupling zero-shot semantic segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11583–11592, 2022.
- [56] Minhyeok Lee, Suhwan Cho, Jungho Lee, Sunghun Yang, Heeseung Choi, Ig-Jae Kim, and Sangyoun Lee. Effective sam combination for open-vocabulary semantic segmentation. In Proceedings of the Computer Vision and Pattern Recognition Conference, pages 26081–26090, 2025.

- [57] Shuai Shao, Zeming Li, Tianyuan Zhang, Chao Peng, Gang Yu, Xiangyu Zhang, Jing Li, and Jian Sun. Objects365: A large-scale, high-quality dataset for object detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 8430–8439, 2019.
- [58] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3, 2022.
- [59] Jiang Liu, Hui Ding, Zhaowei Cai, Yuting Zhang, Ravi Kumar Satzoda, Vijay Mahadevan, and R Manmatha. Polyformer: Referring image segmentation as sequential polygon generation. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pages 18653–18663, 2023.
- [60] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, pages 4171–4186, 2019.
- [61] Zhongwei Ren, Zhicheng Huang, Yunchao Wei, Yao Zhao, Dongmei Fu, Jiashi Feng, and Xiaojie Jin. Pixellm: Pixel reasoning with large multimodal model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 26374–26383, 2024.
- [62] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. arXiv preprint arXiv:2307.09288, 2023.
- [63] Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. Judging Ilm-as-a-judge with mt-bench and chatbot arena. Advances in Neural Information Processing Systems, 36:46595–46623, 2023.
- [64] Haobo Yuan, Xiangtai Li, Tao Zhang, Zilong Huang, Shilin Xu, Shunping Ji, Yunhai Tong, Lu Qi, Jiashi Feng, and Ming-Hsuan Yang. Sa2va: Marrying sam2 with llava for dense grounded understanding of images and videos. *arXiv preprint arXiv:2501.04001*, 2025.
- [65] Zhe Chen, Weiyun Wang, Hao Tian, Shenglong Ye, Zhangwei Gao, Erfei Cui, Wenwen Tong, Kongzhi Hu, Jiapeng Luo, Zheng Ma, et al. How far are we to gpt-4v? closing the gap to commercial multimodal models with open-source suites. *Science China Information Sciences*, 67(12):220101, 2024.
- [66] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2. https://github.com/facebookresearch/detectron2, 2019.
- [67] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Scharwächter, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset. In *CVPR Workshop on the Future of Datasets in Vision*, volume 2, page 1, 2015.
- [68] Fisher Yu, Wenqi Xian, Yingying Chen, Fangchen Liu, Mike Liao, Vashisht Madhavan, Trevor Darrell, et al. Bdd100k: A diverse driving video database with scalable annotation tooling. *arXiv preprint arXiv:1805.04687*, 2(5):6, 2018.
- [69] Petra Bevandić, Marin Oršić, Ivan Grubišić, Josip Šarić, and Siniša Šegvić. Multi-domain semantic segmentation with overlapping labels. In *Proceedings of the IEEE/CVF winter* conference on applications of computer vision, pages 2615–2624, 2022.
- [70] Hangbo Bao, Wenhui Wang, Li Dong, Qiang Liu, Owais Khan Mohammed, Kriti Aggarwal, Subhojit Som, Songhao Piao, and Furu Wei. Vlmo: Unified vision-language pre-training with mixture-of-modality-experts. *Advances in Neural Information Processing Systems*, 35:32897–32912, 2022.
- [71] Peng Wang, An Yang, Rui Men, Junyang Lin, Shuai Bai, Zhikang Li, Jianxin Ma, Chang Zhou, Jingren Zhou, and Hongxia Yang. Ofa: Unifying architectures, tasks, and modalities through a simple sequence-to-sequence learning framework. In *International conference on machine learning*, pages 23318–23340. PMLR, 2022.

- [72] Bin Xiao, Haiping Wu, Weijian Xu, Xiyang Dai, Houdong Hu, Yumao Lu, Michael Zeng, Ce Liu, and Lu Yuan. Florence-2: Advancing a unified representation for a variety of vision tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4818–4829, 2024.
- [73] Quan Sun, Yuxin Fang, Ledell Wu, Xinlong Wang, and Yue Cao. Eva-clip: Improved training techniques for clip at scale. *arXiv preprint arXiv:2303.15389*, 2023.
- [74] Michael Tschannen, Alexey Gritsenko, Xiao Wang, Muhammad Ferjad Naeem, Ibrahim Alabdulmohsin, Nikhil Parthasarathy, Talfan Evans, Lucas Beyer, Ye Xia, Basil Mustafa, et al. Siglip 2: Multilingual vision-language encoders with improved semantic understanding, localization, and dense features. *arXiv preprint arXiv:2502.14786*, 2025.
- [75] Jiahui Yu, Zirui Wang, Vijay Vasudevan, Legg Yeung, Mojtaba Seyedhosseini, and Yonghui Wu. Coca: Contrastive captioners are image-text foundation models. arXiv preprint arXiv:2205.01917, 2022.
- [76] Xi Chen, Xiao Wang, Soravit Changpinyo, AJ Piergiovanni, Piotr Padlewski, Daniel Salz, Sebastian Goodman, Adam Grycner, Basil Mustafa, Lucas Beyer, et al. Pali: A jointly-scaled multilingual language-image model. *arXiv preprint arXiv:2209.06794*, 2022.
- [77] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, pages 19730–19742. PMLR, 2023.
- [78] Zhiliang Peng, Wenhui Wang, Li Dong, Yaru Hao, Shaohan Huang, Shuming Ma, and Furu Wei. Kosmos-2: Grounding multimodal large language models to the world. *arXiv preprint arXiv:2306.14824*, 2023.
- [79] Shilong Zhang, Xinjiang Wang, Jiaqi Wang, Jiangmiao Pang, and Kai Chen. What are expected queries in end-to-end object detection? *arXiv* preprint arXiv:2206.01232, 2022.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: In our abstract and introduction, we talk about our contribution on proposing a novel framework for open-vocabulary segmentation. We provide extensive experiments on comprehensive datasets to support this claim. We also conduct in-depth ablation studies that verifies the effectiveness of our model design.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We discuss the limitations in Appendix D, which is about the model generalization quality to outdoor scenes and self-driving scenes.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: Our work does not include theoretical assumptions and proofs.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We provide the detailed methodology and experimental setup in Sections 3 and 4. Moreover, we provide all source codes to reproduce the results, including training scripts (detailed configurations included) and evaluation scripts (model checkpoints included). We will open source the code on GitHub after acceptance.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We provide the source codes as supplementary material. We provide instructions that contain the exact command and environment needed to run to reproduce the results. We provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: Justification: Training and test details can be found in Sections 3 and 4, and Appendix A.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: We did not include the error bars as we fix the random seed for every experiment, reducing the impact from data loader and other parameters initialization.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We provide the details of computer resources in Section 4. All experiments can be run on a single A100 GPU. We also provide analysis on trainable parameters in Section 4.

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

Answer: [Yes]

Justification: Our experiments conform to the NeurIPS Code of Ethics.

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: There is no social impact of this work.

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The paper poses no such risks.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: Our model and its code development are based on baseline works which are credited in the paper. Our datasets are the standard benchmarks that are widely used in academia.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: The paper does not release new assets.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: The core method development in this research does not involve LLMs as any important, original, or non-standard components.

A Experimental Settings

A.1 Pre-training

We implement our model in PyTorch, building on the Detectron2 [66] framework. We initialize the base models with the public weights of SAM2-Hiera-Large¹ and BEIT-3-Large². The model is pre-trained for 25 epochs on the COCO-2017 training split (104K images) [46], excluding the RefCOCOg-UMD validation subset [23]. We use the *panoptic* annotations, which provide pixel-accurate masks and category labels for all 132 *thing* and *stuff* classes. Training is conducted on a single NVIDIA A100 (80 GB) GPU with a batch size of 8. Optimization employs AdamW (learning rate 1e-4). A step decay scheduler drops the learning rate by a factor of 0.1 at 89% and 96% of the total iterations. Compared with recent generalist models, our recipe is markedly more data-efficient (see Table 6).

Table 6: A detailed list of training data for generalist models and OpenWorldSAM. O365: Objects365. OID: OpenImages Detection. VG: Visual Genome. INB: ImageNetBoxes. RefC: RefCOCO/+/g.

Method	Train Data (grouped by annota	tion types)		Image Consumption
	Instance-level	Image-level	Batch Size	#Epoch \times #Image / Batch Size \times #Iter
X-Decoder [15]	COCO, RefC	Cap4M	32, 1024	200M (50 Ep × 4M Img)
OpenSeeD [36]	COCO, O365	_	32, 64	$48M (30 \text{ Ep} \times 1.8 \text{M Img})$
APE (B) [16]	COCO, LVIS, O365, OID, VG, RefC	-	16	17.28M (16 Bs × 1.08M Iter)
OpenWorldSAM	COCO	_	8	2.50M (25 Ep × 0.104M Img)

A.2 Zero-Shot Evaluation

We evaluate semantic, instance, and panoptic segmentation in a *zero-shot* setting. For instance segmentation and panoptic segmentation, we apply confidence-score filtering to remove masks with scores below 0.7, followed by non-maximum suppression (NMS) with IoU threshold 0.5 to remove duplicate detections and retain distinct object instances. The confidence scores, originally termed "estimated IoU scores" in SAM [17, 18], are direct outputs from SAM2's mask decoder. These scores were optimized during SAM2's pre-training to select high-quality (*i.e.*, confident) mask outputs.

 Table 7: Open-Vocabulary Segmentation Benchmark Statistics.

Evaluation Dataset	Coope type	A	Annotation	s	# Images	# Classes
Evaluation Dataset	Scene type	Semantic	Instance	Panoptic	# Illiages	# Classes
ADE-150	common	~	~	~	2000	150
ADE-847	common	✓	×	×	2000	847
Pascal Voc	common	✓	×	×	1449	20
Pascal Context-59	common	✓	×	×	5105	59
Pascal Context-459	common	✓	×	×	5105	459
SUN RGB-D	in-door	✓	×	×	5050	37
ScanNet-20	in-door	✓	×	✓	5436	20
ScanNet-40	in-door	~	×	×	5436	40

The open-vocabulary benchmark comprises 5 datasets covering 8 different segmentation tasks; statistics are summarized in Table 7. We show a comprehensive evaluation protocol for open-vocabulary segmentation in various vocabulary sizes and image domains.

A.3 Finetuning

For referring-expression segmentation we fine-tune the pre-trained checkpoint for 10 epochs. Because images from RefCOCOg were seen during pre-training (with category labels substituted for referring expressions ground truth), we adopt a conservative learning rate of 1e-5. We use a batch size of 8 during training.

¹https://github.com/facebookresearch/sam2

²https://github.com/microsoft/unilm/tree/master/beit3

B Qualitative Comparison on Two-Stage-Inference

During inference, we perform an optional two-stage inference. First, the model predicts *multi-instance* masks. These masks are then fed back as visual prompts, and SAM2's mask decoder is run a second time to refine the contours. Figure 7 illustrates the visual improvement. However, quantitative gains are marginal across segmentation metrics (see Sec. 4.2 of the main paper), suggesting it mainly enhances visual quality rather than overall accuracy. The reasons are twofold: (1) Two-stage inference only refines mask contours; IoU-style metrics saturate once coarse localization is accurate, so small contour tweaks seldom raise mIoU/PQ/AP; (2) Errors will be amplified on hard examples. On incorrectly localized masks from stage 1, refinement anchored to incorrect regions can further degrade metrics. Given that the two-stage inference serves as an optional, low-cost post-processing step, users can conveniently enable or disable it based on their preference.

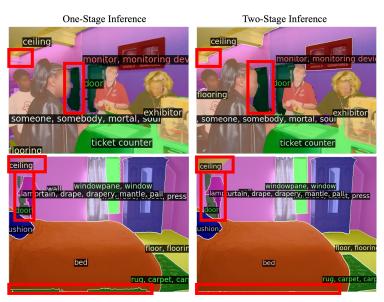


Figure 7: Qualitative results comparisons on using two-stage inference refinement on ADE20K-857.

C Additional Zero-Shot Qualitative Results

Figure 8 showcases multiple challenging indoor scenes drawn from ADE20K-150/857 [22], and PASCAL Context-459 [51]. In each sub-panel, we compare example outputs of OpenWorldSAM with comparisons to X-Decoder under both global-matching and oracle-prompts evaluation protocols.

Panel (a) (ADE20K-150) top row depicts a cluttered bedroom. OpenWorldSAM cleanly delineates thin structures such as the "closet" edge and the narrow "lamp stem", and assigns a single coherent mask to the "cushion". X-Decoder fragments the closet and mis-classifies the cushion as a generic "pillow" under global matching. Under oracle-prompts, X-decoder fails to predict "cushion". Similarly, the bottom row depicts an airport conveyor belt. X-Decoder mis-classifies the "bulletin board" as the "crt screen", the "box" as the "trade name" under global matching, and still mis-classifies the "box" under oracle-prompts.

Panel (b) (ADE20K-857) top row shows a dining area. Under the global-matching protocol, X-Decoder hallucinates "rug"/"rocking chair" labels and fragments the "sofa bed" pixels. The bottom row shows a cluttered living room where X-Decoder outputs fragmented low-quality masks and false predictions under both evaluation protocols. In comparison, our model preserves category fidelity—introducing no extra labels—and produces noticeably cleaner chair boundaries, illustrating the synergy between BEiT-3 language grounding and SAM2's high-resolution masks.

Panel (c) (PASCAL Context-459) top row shows that X-Decoder fails to predict the "cloth" object. The bottom row is an indoor scene crowded with small objects ("cd", "speaker", "chair"). OpenWorld-SAM retrieves almost every queried category (except for "cd") and suppresses false positives such as

"calendar" and "ladder" that appear in X-Decoder's output, demonstrating stronger open-vocabulary grounding and sharper instance separation.

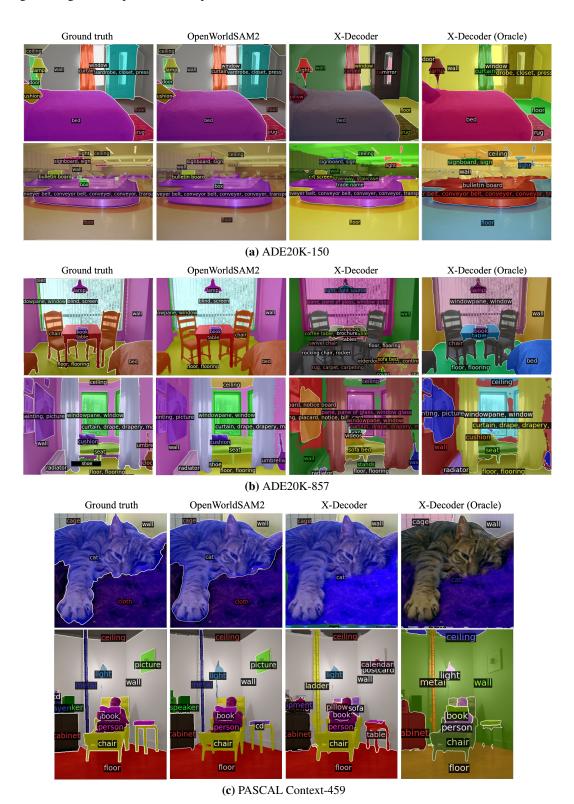


Figure 8: Qualitative comparisons between X-Decoder [15] and OpenWorldSAM on ADE20K-150, ADE20K-857, and PASCAL Context-459.

D Limitations - Outdoor Generalization

Despite strong results on indoor and everyday photographs, OpenWorldSAM under-performs on driving datasets such as Cityscapes [67] and BDD10K [68] (Table 8). Fine-tuning on Cityscapes narrows the gap, yet performance still trails methods explicitly exposed to multi-domain data. Understanding the source of this shortfall is essential for future extensions.

Observed failure modes. Figure 9 shows high IoU for broad *stuff* regions (e.g. *road*, *sky*), but a sharp drop for small or elongated *thing* instances. Correspondingly, AP remains low for *motorcycle*, *person*, *bicycle*, etc.

Table 8: Outdoor performance. Open-vocabulary models are evaluated zero-shot on Cityscapes and BDD10K; the last row is fine-tuned on Cityscapes.

Model	Evaluation	Ci	tyscape	es	BDD10K		
Model	Evaluation	mIoU	AP	PQ	mIoU	PQ	
X-Decoder (L) [15]	zero-shot	52.0	24.9	38.1	47.2	17.8	
OpenWorldSAM	zero-shot	39.4	10.1	26.4	31.3	15.6	
OpenWorldSAM	Finetune on Cityscapes	57.4	12.0	36.1	38.0	17.4	

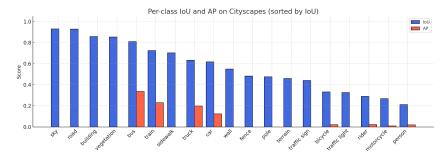


Figure 9: Per-class IoU And AP on Cityscapes (sorted by IoU). Performance collapses on thin or distant *thing* classes (e.g. *person*, *traffic light*).

Hypotheses.

- (i) *Domain shift*. COCO images are mostly handheld and indoor, whereas Cityscapes/BDD10K contain forward-looking dash-cam frames with motion blur, glare and night scenes. X-Decoder was co-trained on web-scale image-text pairs that include many outdoor photos, so its visual encoder has wider coverage. Large-scale multi-domain training is known to mitigate domain shift [69].
- (ii) Resolution bottleneck. Cityscapes frames are 2048×1024. Rescaling to 1024×1024 (SAM default) reduces poles and traffic lights to nearly one pixel at the feature stride of 16×. X-Decoder keeps an FPN branch at 8×, preserving thin structures.

Take-away. COCO-only pre-training for OpenWorldSAM leaves a blind spot for urban driving imagery—particularly for distant, thin or cluttered objects under challenging lighting. Bridging the gap likely requires (i) explicit exposure to outdoor domains and (ii) higher-resolution feature branches. We leave large-scale outdoor pre-training and depth-aware augmentation for future work.

E Model Structure Details

Table 9 summarizes the architectural differences between OpenWorldSAM and competing models, detailing each method's visual backbone, segmentation head, text encoder, and training-image resolution.

Table 9: Architectural choices for recent open-vocabulary models: visual backbone, base segmentation model, text/multi-modal encoder, and training image size.

Method	Visual Backbone	Base Model	Text Encoder	Image S	Size
				Short	Long
MSeg (B) [54]	HRNet-W48 (65 M)	HRNet-Seg	_	1024	1024
GroupViT (S) [27]	ViT-S/16 (22 M)	GroupViT	Transformer	224	224
LSeg+(B)[24]	CLIP ViT-B/16 (86 M)	DenseCLIP	CLIP	512	512
ZegFormer (B) [55]	CLIP ViT-B/16 (86 M)	ZegFormer	CLIP	640	640
OpenSeg (B) [11]	ResNet-101 (45 M)	OpenSeg	CLIP/ALIGN	640	640
OVSeg (B) [12]	CLIP ViT-B/16 (86 M)	MaskFormer	CLIP	640	640
MaskCLIP (L) [28]	CLIP ViT-L/14 (307 M)	MaskCLIP	CLIP	1024	1024
X-Decoder [15]	DaViT-L (196 M)	X-Decoder	CLIP	1024	1024
OpenSeeD [36]	Swin-L (197 M)	MaskDINO	UniCL	1024	1024
SEEM [35]	DaViT-L (196 M)	X-Decoder	CLIP	800	1333
APE (B) [16]	ViT-L (307 M)	DETA	CLIP	1024	1024
PolyFormer (L) [59]	Swin-L (197 M)	PolyFormer	BERT	1024	1024
UNINEXT (H) [32]	ViT-H (632 M)	DINO	BERT	$320 \sim 800$	1333
PixelLM [61]	CLIP ViT-L/14 (307 M)	PixelLM	LLaMA2-13B	448	448
LISA [19]	SAM ViT-H (636 M)	SAM	Vicuna-7B	1024	1024
GLaMM [20]	SAM ViT-H (636 M)	SAM	Vicuna-7B	1024	1024
u-LLaVA [21]	SAM ViT-H (636 M)	SAM	Vicuna-7B	1024	1024
EVF-SAM [43]	SAM ViT-H (636 M)	SAM	BEiT-3	1024	1024
EVF-SAM2 [43]	SAM2 Hiera-L (224 M)	SAM2	BEiT-3	1024	1024
OpenWorldSAM	SAM2 Hiera-L (224 M)	SAM2	BEiT-3	1024	1024

E.1 Possible Text Encoder Alternatives

We argue that the key ingredients for open-vocabulary segmentation are backbone-agnostic: any strong interactive segmenter can supply high-resolution mask decoding, while any pretrained vision-language encoder can provide semantics. What is missing is a lightweight adaptor that (i) aligns the two embedding spaces, (ii) scales to multiple object instances from a single text query, and (iii) preserves the efficiency that makes interactive segmentation attractive in the first place.

Our OpenWorldSAM is a general plug-in architecture that satisfies these desiderata while keeping all heavy backbones frozen. Although we instantiate the framework with SAM2 and BEiT-3 in this paper, neither component is required by design; alternative interactive decoders or vision-language encoders can be swapped in with only minor re-training of the adapter.

Table 10 surveys representative VLM encoders that could replace BEiT-3 in OpenWorldSAM with \leq 5 M adaptor parameters. All rows assume the heavy backbone is **frozen**; only the 256-D projector and tie-breakers are fine-tuned.

Adaptor fine-tuning recipe (all encoders). Freeze all VLM weights and SAM2 decoder; initialize a $d_{\rm in} \times 256$ MLP projector and K 256-D tie-breaker embeddings (default K=20, total ≈ 5 M params). For training, one could use unchanged Hungarian matching loss on COCO.

Takeaway. Early-fusion encoders (VLMo, OFA, Florence-2) require zero architectural change beyond projector resizing and are therefore the most promising immediate swaps. Dual-encoders (CLIP family) need a shallow cross-attention adaptor to overcome missing image context. Larger hybrids (BLIP-2, Kosmos-2, PaLI) open research directions (multi-query tie-breakers, OCR) at the cost of real-time guarantees.

F Additional Ablation Studies

We provide additional ablation studies on the number of tie-breaker tokens and the number of cross-attention layers.

Table 10: Candidate vision–language encoders. "TFM" stands for Transformer. "Pooled dim" is the size of the single semantic vector exposed to the adaptor; "GFLOPs/Img" computed at 224² resolution for the visual branch.

Family / Exemplars	Arch. type	Pooled dim	Params	Pros for OpenWorldSAM	Adaptor-specific tweaks	GFLOPs/Img
			Early-fi	Early-fusion Transformers (drop-in closest to BEiT-3)		
VLMo-B/L [70]	joint enc. TFM	768	230/341M	same interface as BEiT-3; smaller model; multi-lingual	closest to BEiT-3 \rightarrow just replace tokenizer + dimension in the projector; keep tie-breakers unchanged	18.6/25.4
OFA-B/L [71]	joint enc. TFM	768	184/312M	instruction-tuned; handy if we ever expose captioning	adjust tokenizer and change input dim in projector; reports slightly weaker alignment than BEiT-3	17.9/24.7
Florence-2-Base [72]	joint enc. TFM	1024	230M	SOTA zero-shot retrieval; 10-lang support	none beyond changing tokenizer and input dim in projector	26.3
			Dual-encod	Dual-encoder Contrastive (text vector not image-conditioned)		
CLIP-ViT-L/14 [13]	ViT+Text enc.	768	304M	unlimited vocabulary; tiny latency; many open checkpoints	semantic vector is not image-conditioned → our ablation saw weaker performance. Mitigation: add a 2-layer cross-attn adapter that re-injects image tokens before the projector; expect AP drop if no cross-attn	19.0
EVA-CLIP-E [73] SigLIP-2-S [74]	ViT-G/14 + Text ViT/Text	1024 512	610M 86M	stronger semantics than CLIP-L edge-friendly; multilingual	memory heavy; expect AP drop if no cross-attn expect AP drop if no cross-attn	37.2
		E	coder–Deco	Encoder-Decoder w/ Contrastive Head (pooled vector from decoder)	rr)	
CoCa-Base [75] PaLJ-3B [76]	ViT enc. + TFM dec. ViT-E enc. + T5 dec.	768 1024	365 M 3.0 B	better long-tail semantics 100-lang OCR; robust semantics	need to tap the unimodal decoder hidden state memory heavy; need to tap the unimodal decoder hidden state	23.7 56.4
			3	Query-former Hybrids (multiple vectors)		
BLIP-2-OPT-2.7B [77] Kosmos-2 [78]	ViT + Q-Former + LLM ViT enc. + LLM dec.	32×256 768	1.1B 1.6B	native multi-query optional box tokens for UX studies	pool/average queries or extend SAM prompt len. Requires a one-step decode per prompt (latency) and an additional MLP to strip location bias	31.5 34.8

F.1 Effect of varying tie-breaker tokens

We set the hyperparameter K=20, meaning for each prompt (e.g., a category name), our model can identify up to 20 distinct objects. For crowded scenes containing more than 20 objects per category, increasing K is straightforward and advisable. In practice, COCO images typically contain a moderate number of distinct categories and instances (the original COCO paper reports "on average, our dataset contains 3.5 categories and 7.7 instances per image." [46]). The chosen value should match or exceed the maximum expected number of objects per category. For reference, DETR [45] used 100 total queries, aligning roughly with the maximum number of objects per image. Our choice (K=20) results, on average, in approximately 70 queries per image (20 queries \times 3.5 categories), providing ample coverage for typical scenes.

Further, [79] observed that increasing queries initially improved Average Precision (AP), but then plateaued or even slightly declined when queries became excessive, indicating redundancy in higher query counts. However, recall does improve with more queries, since more detection slots increase the chance to find each object.

We conducted additional ablation experiments in Table 11 by varying K, pretrained on COCO and evaluated on ADE20K instance segmentation.

Table 11: Ablation on the number of tie-breakers K.

Metric	K = 10	K = 20	K = 30
Average Precision (AP)	14.2	16.9	16.5
Average Recall@100 (AR)	21.6	28.8	29.4

Observations. (1) Increasing K from 10 to 20 improves recall and AP; beyond 20 gains saturate, mirroring the behavior reported for DETR-style object queries; (2) Average Recall with max 100 detections per image (AR@100) improve when increasing K from $10 \rightarrow 20 \rightarrow 30$; (3) K=20 is optimal for balancing precision and recall in standard datasets.

F.2 Effect of varying number of cross-attention layers

In Table 12, we observe consistently higher accuracy with 3-layer cross-attention across datasets, confirming the importance of multi-layer cross-attention. However, a single-layer variant significantly narrows the gap with fewer parameters (2.4M vs. 4.5M), suggesting a practical compromise between parameter count and accuracy.

Table 12: Ablation on cross-attention depth across datasets. Metrics are PQ/AP/mIoU for ADE-150 and mIoU for the others.

Variant	Params	ADE-150 (PQ/AP/mIoU)	ADE-857 (mIoU)	PC-59	PC-459	VOC-20	SUN-37	SCAN-40
no cross-attn	1.7 (M)	35.1/17.1/56.8	32.2	70.4	44.2	97.3	63.6	53.8
1-layer cross-attn	2.4 (M)	35.1/16.8/59.0	32.8	72.6	46.3	97.5	66.4	54.0
3-layer cross-attn	4.5 (M)	35.2/16.9/60.4	33.1	73.7	47.5	98.0	67.7	55.6

G Inference Speed Analysis

We conducted detailed profiling to quantify the impact of adding the VLM and our adapter modules to SAM. In Table 13 and 14, we present inference timing breakdowns for processing a single 1024×1024 image on an NVIDIA A5000 GPU, averaged over five independent runs.

Summary (single prompt). SAM modules total time: 373.35 ms (81.0%), Non-SAM modules total time: 87.21 ms (18.9%), Non-SAM overhead: 87.21 ms.

Summary (six prompts). SAM modules total time: 539.72 ms (78.4%), Non-SAM modules total time: 148.65 ms (21.6%), Non-SAM overhead: 148.65 ms.

Table 13: Inference timing breakdown for a single text prompt (20 queries).

Module	Time (ms)	Percentage	Category
sam_backbone_feature_prep	329.83	71.6%	SAM
<pre>prompt_tokenization</pre>	0.43	0.1%	Non-SAM
beit3_forward	70.84	15.4%	Non-SAM
mlp_projection_layer	6.68	1.4%	Non-SAM
prepare_batched_tie_breaker_tokens	0.13	0.0%	Non-SAM
cross_attention	8.45	1.8%	Non-SAM
sam_prompt_encoder	0.11	0.0%	SAM
sam_mask_decoder	43.41	9.4%	SAM
postprocessing	0.68	0.1%	Non-SAM
TOTAL TIME	460.69	100.0%	_

Table 14: Inference timing breakdown for six text prompts (120 queries).

Module	Time (ms)	Percentage	Category
sam_backbone_feature_prep	334.42	48.6%	SAM
prompt_tokenization	1.02	0.1%	Non-SAM
beit3_forward	123.73	18.0%	Non-SAM
mlp_projection_layer	4.48	0.6%	Non-SAM
prepare_batched_tie_breaker_tokens	0.20	0.0%	Non-SAM
cross_attention_layers	18.17	2.6%	Non-SAM
sam_prompt_encoder	0.12	0.0%	SAM
sam_mask_decoder	205.18	29.8%	SAM
postprocessing	1.06	0.2%	Non-SAM
TOTAL TIME	688.50	100.0%	_

Takeaway. The profiling results show that adding the VLM and adapter modules results in only a moderate increase in inference time (approximately 19–22% overhead). Most computational cost remains within SAM's backbone and mask decoder.

Mask Decoder scaling. sam_mask_decoder cost grows almost linearly with $(K \times P)$.

- Going from $1 \rightarrow 20$ queries (same prompt) adds \sim 41 ms.
- Going from 1 prompt \rightarrow 6 prompts (120 queries) adds a further \sim 162 ms.

Note that one text prompt mimics user clicks 20 times on an image. If automatic mask generation is desired without user intervention, SAM's built-in auto-mask generator uses a dense 32×32 grid of point prompts, incurring significantly higher costs compared to our text-based prompting approach.

Overall overhead. Relative to one vanilla SAM2 call, our pipeline is approximately 39% slower for a single prompt $(332 \rightarrow 461 \text{ ms})$. However, it becomes approximately $3 \times$ more efficient when handling three or more prompts, as the backbone and VLM overhead are amortized. Thus, our enhancements introduce manageable overhead, maintaining practical usability in real-world applications.