

Show or Tell? A Benchmark To Evaluate Visual and Textual Prompts in Semantic Segmentation

Gabriele Rosi^{1,2}, Fabio Cermelli²

¹ Politecnico di Torino, ² Focoos AI

¹ name.surname@polito.it, ² name.surname@focoos.ai

Abstract

Prompt engineering has shown remarkable success with large language models, yet its systematic exploration in computer vision remains limited. In semantic segmentation, both textual and visual prompts offer distinct advantages: textual prompts through open-vocabulary methods allow segmentation of arbitrary categories, while visual reference prompts provide intuitive reference examples. However, existing benchmarks evaluate these modalities in isolation, without direct comparison under identical conditions. We present Show or Tell (SoT), a novel benchmark specifically designed to evaluate both visual and textual prompts for semantic segmentation across 14 datasets spanning 7 diverse domains (common scenes, urban, food, waste, parts, tools, and land-cover). We evaluate 5 open-vocabulary methods and 4 visual reference prompt approaches, adapting the latter to handle multi-class segmentation through a confidence-based mask merging strategy. Our extensive experiments reveal that open-vocabulary methods excel with common concepts easily described by text but struggle with complex domains like tools while visual reference prompt methods achieve good average results but exhibit high variability depending on the input prompt. Through comprehensive quantitative and qualitative analysis, we identify the strengths and weaknesses of both prompting modalities, providing valuable insights to guide future research in vision foundation models for segmentation tasks. Code is available at <https://github.com/FocoosAI/ShowOrTell>.

1. Introduction

A long-standing goal of artificial intelligence is to create models that can generalize and adapt to multiple tasks without requiring complex and costly fine-tuning or retraining [6, 15, 24, 29, 42, 54, 67]. Prompt engineering has recently revolutionized large language models (LLMs) showing that optimizing input instructions can substantially alter model behavior and improve performance [9]. However, systematic investigations into prompt engineering for com-

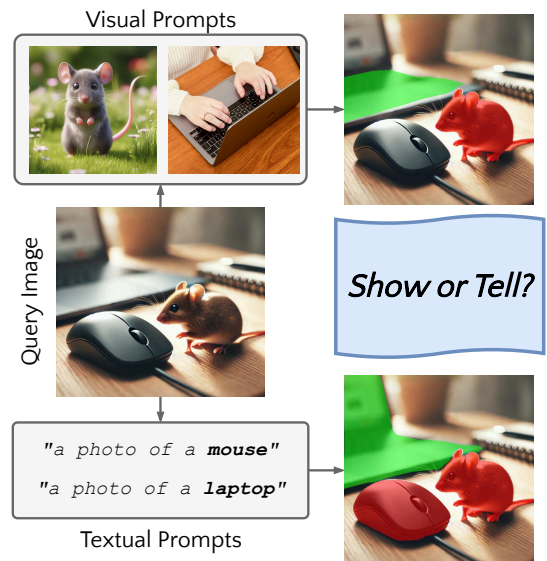


Figure 1. Our Show or Tell (SoT) benchmark evaluates the effectiveness of textual and visual prompts in semantic segmentation. Here the query image poses a challenge for textual prompts since “mouse” can refer to both an animal and a computer accessory, resulting in an inaccurate segmentation. Similarly, visual prompts also face difficulties in accurately identifying the “laptop” due to the limited information provided by a top-down view.

puter vision tasks remain in their early stages [23, 59].

Recent work in computer vision has begun exploring how insights from prompt engineering can enhance semantic segmentation. By integrating carefully designed prompts into segmentation pipelines, researchers are uncovering new ways to guide pixel-level predictions and improve model flexibility. Pioneering research in this direction are represented by open-vocabulary semantic segmentation [1, 8, 16, 30, 38, 44] which has initiated the exploration of textual prompting for vision tasks. These approaches leverage Vision Foundation Models (VFMs), particularly CLIP [21, 45], to extract and align textual embeddings from class names with visual features, thereby enabling the segmentation of arbitrary categories without prior training on specific classes. Despite their efficacy, these tex-

tual approaches encounter substantial limitations. Indeed, difficult concepts, such as the various types of cracks on a wall or a specific bird species, are hard to describe with text alone, as textual descriptions may fail to capture the unique visual cues needed for accurate recognition. Moreover, generic textual prompts frequently fail to distinguish between semantically distinct objects: as shown in Fig. 1, the prompt “a photo of a mouse” results in the segmentation of both the animal and the computer peripheral, highlighting the need for more precise contextual information.

To address these limitations, visual prompts have emerged as a compelling alternative due to their intuitive nature and straightforward implementation [23, 47, 51]. Human perception typically processes visual information before linguistic categorization, making visual prompting a natural approach to object identification. The use of masks, bounding boxes, or scribbles to indicate objects of interest provides an effective mean for identifying concepts that resist precise verbal description. Visual reference prompt methods [23, 32, 52, 68, 72] have yielded significant advancements in image segmentation. Recent approaches [32, 68] typically employ a two-stage pipeline leveraging VFMs: initially, a model such as DINO [5, 41] identifies pixels exhibiting similarity to the reference class, followed by prompting SAM [23] to generate the segmentation masks. However, relying on visual prompts poses challenges when objects vary in appearance across instances. As depicted in Fig. 1 choosing the right visual prompts could be challenging since they might contain limited visual cues, thus reducing the segmentation performance.

Based on these observations, each modality, textual and visual, offers distinct advantages and limitations. Comparing them provides two important benefits: it reveals scenarios where one type of prompt performs better than the other, and it helps determine which modality works best for semantic segmentation tasks. Notably, existing benchmarks [3, 75] typically evaluate visual and textual prompts separately, without directly comparing them under the same conditions. Additionally, previous studies have tested visual reference prompt methods only for single-class segmentation [32, 52, 68], failing to assess their performance in complex real-world scenarios.

To bridge this gap, we introduce a novel benchmark, *Show or Tell (SoT)*, specifically designed to evaluate both visual and textual prompts within the context of semantic segmentation. To properly assess the generalization ability, SoT spans 14 datasets across 7 domains (common, urban, food, waste, parts, tools, and land-cover). This benchmark enables a direct, head-to-head evaluation of different prompting methods in semantic segmentation, assessing their adaptability by only changing the prompts. To the best of our knowledge, this is the first effort toward an evaluation of multiple prompting techniques in semantic segmentation.

Our evaluation encompasses five open-vocabulary methods (MaskCLIP [10], TCL [31], CLIP-DiNoiser [63], NA-CLIP [19], and ProxyCLIP [25]) and four visual reference prompt approaches (SINE [33], PerSAM [72], Matcher [32], and GFSAM [68]), all assessed under identical experimental conditions. To facilitate a fair comparison of visual reference prompt methods, which were originally designed for binary segmentation tasks, we implement a novel adaptation for multi-class segmentation scenarios. This adaptation involves generating individual binary masks for each class and subsequently merging them according to the model’s confidence scores, enabling these methods to effectively segment scenes containing multiple objects.

Extensive experimental results highlight that open-vocabulary methods excel in domains where the classes represent common concepts that can be easily described by text, while visual reference prompt methods obtain good results on average but their results can significantly differ depending on the input prompt.

In summary, the contributions of the paper are three-fold:

- We introduce a novel benchmark, SoT, comprising 14 datasets in 7 domains and systematically comparing 4 visual and 5 textual prompting methods;
- We adapt visual reference prompt methods to operate in multi-class setting;
- Through an extensive quantitative and qualitative analysis, we identify the strengths and weaknesses of both prompting modalities, providing valuable insights for future research.

2. Related works

Prompt engineering for semantic segmentation. Prompt engineering, originating in natural language processing, has expanded to computer vision applications. Initially focused on text-to-image generation with models like DALL-E [46] and Stable Diffusion [48], this approach evolved with vision-language models such as CLIP [45], which demonstrated how textual prompts could shape visual representations for zero-shot recognition. In semantic segmentation, prompting techniques are used to guide foundation models including CLIP [45], DINO [5, 41], and SAM [23]. Most approaches rely exclusively on either text prompts (Open-Vocabulary Semantic Segmentation [18, 25, 38, 63]) or visual prompts (Visual Reference Prompt Segmentation [32, 52, 68]), with few exceptions [34, 39] combining both modalities. While previous benchmarks [3, 75] typically evaluate the approaches separately, our work directly compares them to assess their respective strengths and weaknesses in real-world scenarios.

Open-vocabulary segmentation. Open-vocabulary semantic segmentation extends traditional zero-shot semantic segmentation [4, 17, 43, 64] by enabling models to iden-

tify arbitrary object categories without class-specific training, leveraging the semantic understanding capabilities of vision-language foundation models [5, 23, 41, 45] to generalize beyond predefined class sets. The extension of CLIP’s aligned image-language representations [21, 45] to segmentation presents challenges, as CLIP’s architecture is not inherently designed for dense vision-language features. Initial approaches relied on fine-tuning CLIP to semantic segmentation by exploiting a large labeled set containing pixel-wise annotations from a large set of classes [1, 8, 16, 30, 38, 44]. However, these approaches require a large amount of annotated data and they alter the open-vocabulary capabilities of the original CLIP model, biasing its knowledge on the training dataset. For this reason, recent works have focused on training-free approaches [19, 22, 25, 50, 57, 63, 74]. Recent innovations include MaskCLIP [74], which revealed that value embeddings offer better localization than token embeddings, and approaches that refine CLIP’s attention mechanisms [19, 28, 55]. Other methods [25, 57, 63] combine CLIP’s semantic understanding with spatial consistency from models like DINO and SAM.

Visual Reference Segmentation. Visual Reference Segmentation employs annotated reference images to guide the segmentation of semantically similar regions in target images. Originating as Few-Shot Segmentation (FSS) [49], early approaches [7, 26, 60, 69] concentrated on training neural networks to extract prototypes from reference images and compute similarity for target image segmentation. Due to the limitations of prototype-based methods, subsequent research proposed extracting correlation maps to represent the relationship between reference and target images [36, 56] or utilizing attention maps to guide target image segmentation [20, 70]. The advent of vision foundation models (VFMs) has transformed the field, directing research toward the use of large-scale pre-trained models for target image segmentation. Several approaches [33, 61, 62, 65] have developed models with cross-task generalization capabilities. Painter [61] introduced an in-context learning framework wherein vision tasks are defined through exemplars. SINE [33] presented an encoder-decoder architecture that handles multiple tasks via in-context examples. However, the training of such models requires substantial computational resources and extensive datasets. Consequently, research has increasingly focused on employing existing VFMs such as DINO [5, 41] and SAM [23, 47] in a training-free manner, as these models offer superior generalization capabilities due to their comprehensive pretraining. PerSAM [72] adapts SAM for personalized segmentation with minimal parameter modifications. VRP-SAM [52] incorporates SAM with an external feature-matching encoder without fine-tuning. Matcher [32] and GFSAM [68] implement two-stage pipelines that employ DINOv2 [41] to compute cross-image similarities for extracting prompts for SAM.

| Name | Prompt type | Visual backbone | Trained |
|---------------------------------------------|-------------|--------------------------|---------|
| SINE [33] <small>[NeurIPS’24]</small> | | DINOv2 [41] | ✓ |
| PerSAM [72] <small>[ICLR’24]</small> | Visual | SAM [23] | ✗ |
| Matcher [32] <small>[ICLR’24]</small> | | DINOv2 [41] + SAM [23] | ✗ |
| GFSAM [68] <small>[NeurIPS’24]</small> | | DINOv2 [41] + SAM [23] | ✗ |
| TCL [8] <small>[CVPR’23]</small> | | CLIP [45] | ✓ |
| MaskCLIP [10] <small>[ECCV’22]</small> | Textual | CLIP [45] | ✗ |
| CLIP-DINoiser [63] <small>[ECCV’24]</small> | | CLIP [45] | ✗ |
| NACLIP [19] <small>[WACV’25]</small> | | CLIP [45] | ✗ |
| ProxyCLIP [25] <small>[ECCV’24]</small> | | CLIP [45] + DINO [5, 41] | ✗ |

Table 1. List of models analyzed in our benchmark. For each model we report the type of prompt employed, the visual backbone(s) used and if the model undergoes a training process.

3. Show or Tell Benchmark

To facilitate a rigorous comparison between visual and textual prompts for semantic segmentation, we introduce *Show or Tell (SoT)*. This novel benchmark evaluates 5 open-vocabulary and 4 visual reference prompt methods across 14 distinct datasets spanning 7 domains. In the following sections, we present an overview of the methods included in the benchmark (Sec. 3.1) and describe our approach for adapting visual reference prompts to generate predictions for multi-class semantic segmentation (Sec. 3.2). Subsequently, we describe the composition of the benchmark, including the datasets and their challenges (Sec. 3.3).

3.1. Models

A comprehensive evaluation of visual and textual prompts requires the selection of an appropriate subset of models for each prompt category. To accomplish this objective, we first conduct an analysis of the operational mechanisms of models employing these prompt types, followed by a justification of our selection criteria. A summary of the benchmarked methods is presented in Tab. 1.

Textual prompting based. Text serves as a powerful prompt, enabling the description of visual concepts using natural language. In recent years, this approach has been widely adopted by open-vocabulary segmentation (OVSS) models. Specifically, in the context of OVSS, the goal is to assign a semantic label (provided as a free-form text description) to each pixel (or region) in an image, without relying on a predefined set of labels. OVSS methods operate on an image and a vocabulary of textual class descriptions. These approaches extract dense visual features from the image using a visual encoder while simultaneously generating embeddings for each class in the vocabulary via a text encoder. The core challenge lies in aligning these visual and textual representations to produce accurate semantic segmentation maps for all prompted classes.

For our benchmark, we evaluate five representative methods: MaskCLIP [74], TCL [8], CLIP-DINoiser [63], NACLIP [19], and ProxyCLIP [25]. We deliberately em-

phasize training-free approaches that leverage VFMs like DINO [5, 41] and CLIP [45] rather than models specifically trained for segmentation tasks. This selection strategy serves dual purposes: it prevents potential bias from dataset contamination that might favor certain domains, and it better aligns with our cross-domain evaluation objectives. Training-free methods typically demonstrate stronger generalization capabilities due to their diverse pretraining on large-scale datasets [21, 45], allowing us to assess their intrinsic ability to transfer knowledge to novel visual concepts without task-specific optimization.

Visual prompting based. Visual prompts represent concepts through explicit visual cues such as masks, bounding boxes, or scribbles, offering an alternative to textual descriptions. These prompts are predominantly employed in few-shot segmentation (FSS) [32, 49, 52], where the objective is to generate a segmentation mask for a query image that identifies the same class annotated in a small support set of images. While FSS methods demonstrate adaptability to novel classes, they typically operate in a single-class segmentation paradigm. Although some approaches have been extended to generalized and incremental few-shot segmentation [7, 64], these methods are trained on specific datasets (e.g., PASCAL VOC [14], ADE20K [73]) and lack the generalization capabilities required for our evaluation.

For our benchmark, we specifically selected visual reference prompt methods that leverage vision foundation models (VFMs) such as DINO [5, 41] and SAM [23], prioritizing their inherent generalizability across diverse visual domains. The benchmark includes SINE [33], PerSAM [72], Matcher [32], and GFSAM [68]. Among these, SINE [33] employs DINO [41] for visual prompting, and fine-tunes a transformer decoder to produce mask at different granularity able to perform in-context segmentation. Despite their effectiveness, these methods are designed to segment only one class at a time, necessitating adaptation for multi-class semantic segmentation in our benchmark.

3.2. Adapting Visual Reference Prompt Methods

Visual reference prompt methods are designed to generate a single mask when provided with an annotated support set. These methods accept a support set containing images with masked annotations of a target class and subsequently produce a binary segmentation map identifying the corresponding class in a query image. For comprehensive evaluation in SoT, however, these methods require adaptation to accommodate multi-class scenarios. We implement this adaptation through a two-stage process: first, computing individual masks for each class, and second, integrating these masks based on the model’s confidence for each prediction.

Producing multiple masks. To adapt visual reference prompt methods for multi-class segmentation, we first create a semantic support set \mathcal{S}^{sem} by randomly sampling k

| Domain | Dataset | Classes | Train images | Val. images |
|-------------------|----------------------|---------|--------------|-------------|
| Common | ADE20K [73] | 150 | 20,000 | 2,000 |
| | PASCAL VOC 2012 [14] | 21 | 1,464 | 1,449 |
| Urban | Cityscapes [12] | 19 | 2,975 | 500 |
| | UAVID [35] | 7 | 200 | 70 |
| Waste | Trash [40] | 12 | 832 | 92 |
| | ZeroWaste [2] | 4 | 3,002 | 572 |
| Food | Pizza [66] | 5 | 437 | 122 |
| | UECFood [13] | 102 | 9,000 | 1,000 |
| Tools | Toolkits [37] | 8 | 48 | 6 |
| | PIDray [71] | 12 | 29,454 | 3,733 |
| Parts | House-Parts [53] | 22 | 700 | 201 |
| | MHPv1 [27] | 17 | 4,000 | 980 |
| Land-Cover | LoveDA-Rural [58] | 6 | 1,366 | 992 |
| | LoveDA-Urban [58] | 6 | 1,156 | 667 |

Table 2. Our benchmark SoT is composed of 14 datasets divided into 7 different domains. For each dataset, we report the number of classes and the number of train and validation images.

images for each of the n classes in the dataset. This results in a total of $n \times k$ support images. When processing a query image q , we run the model n separate times, once for each class, to generate a binary mask for each class. While this approach is straightforward to implement, it requires n separate forward passes, making it computationally expensive. We analyze this efficiency issue in detail in Sec. 4.4.

Merging predictions. After processing a query image, visual reference prompt methods produce n independent binary masks, one for each class. To create a unified semantic segmentation map, we must assign a confidence score to each mask. We standardize this process across methods: for training-free approaches (such as PerSAM [72], Matcher [32], and GFSAM [68]), we use the mean confidence score from the visual backbone (DINOv2 [41] or SAM [23]) within each predicted mask. For trained methods like SINE [33], we use the output probability of the decoder. Once these confidence scores are assigned, we apply an argmax operation across all class masks to produce the final semantic segmentation prediction.

3.3. Datasets

Our benchmark, SoT, aims to rigorously evaluate how textual and visual prompting methods perform across diverse real-world scenarios. To achieve this, we carefully selected 14 datasets spanning 7 distinct domains. This selection was motivated by the need to assess these methods beyond controlled environments, following recent benchmarks that test segmentation approaches “in the wild” [3, 75]. We deliberately included both common scenes and specialized applications to ensure comprehensive coverage of visual contexts that practitioners might encounter. In the following, we describe each domain and its constituent datasets, emphasizing their unique visual characteristics that might differently affect the performance of textual versus visual prompting



Figure 2. Sample images drawn from the datasets that compose our SoT benchmark.

approaches. Table 2 provides a summary of all datasets included in our benchmark, while Figure 2 reports some example images from each dataset.

Common scenes represent foundational benchmarks in segmentation research. ADE20K [73] encompasses 150 classes categorized as stuff (e.g. *sky, grass, road*) or things (e.g. *car, person, chair*), providing a comprehensive framework. In contrast, PASCAL VOC 2012 [14] adopts an object-focused approach with 20 thing classes distributed across varied indoor and outdoor environments.

Urban scenes are critical for autonomous vehicles. Cityscapes [12] provides high-resolution street-level imagery with 19 classes with dynamic objects and varying weather conditions. UAVid [35] captures aerial perspectives with 7 categories, introducing complexities due to scale variations, occlusions, and perspective distortions.

Waste domain datasets capture litter in diverse environments. Trash [40] and ZeroWaste [2] contain 12 and 4 waste classes respectively. The variability in shape makes certain classes particularly challenging: for instance, *Styrofoam pieces* can resemble plastic and appear in unpredictable forms, complicating both textual and visual prompting.

Food domain presents challenges due to preparation and presentation variability. Pizza [66] includes categories for pizza base and 4 dressings, while UECFood [13] comprises 102 dishes. Classes like *tomato* may appear whole or crushed, while dishes with similar appearance may be difficult to identify (e.g., *miso soup vs chinese soup*).

Tools domain is important for security and robotics applications. Toolkit [37] contains 8 common tools on a work-

bench with many occlusions, while PIDray [71] presents 12 classes in X-ray scans. Objects like *scissors* or *lighters* are particularly challenging to identify in X-ray imagery due to material composition variations and occlusions.

Parts segmentation demands understanding of object structures. House-Parts [53] includes 22 architectural elements where features like *wooden doors* may lack visual distinction. MHPv1 [27] contains 17 human body parts, introducing complexity in distinguishing spatially related elements such as *right leg* versus *left shoe*.

Land-Cover applications are increasingly important for satellite imagery analysis. LoveDA [58] presents a dataset with 6 classes distinguishing rural and urban areas. This domain is challenging as classes like *buildings* are not usually seen from above, while *agricultural* terrain varies significantly based on crop types and seasonal changes.

4. Results

4.1. Implementation details

For visual reference prompting methods, the semantic support set S^{sem} is computed once at the beginning of the evaluation and remains unchanged. The images that constitute S^{sem} are drawn from the training split of each dataset, while the evaluation is conducted on the validation set when available, or on the test set otherwise.

We adapted visual reference prompting methods by reimplementing them within our codebase, while evaluating open-vocabulary methods using MMSegmentation [11]. Due to the inherent variability in visual prompt selection, we conducted each experiment by randomly sampling five different support sets. We report both the mean and standard deviation of the results. For OVSS methods, we follow their protocol and use a vocabulary containing the dataset classes and, when required, also the *background* class. For our evaluation metric, we used the standard mean Intersection over Union (mIoU), with the background class excluded from computation when present.

4.2. Quantitative results

Table 3 reveals that visual prompting methods generally outperform textual approaches across our benchmark, with GFSAM [68] achieving the highest average performance (38.7 mIoU on average with five prompts). The performance gap between modalities dramatically varies across domains, with visual methods showing particular strength in specialized contexts while textual approaches remain competitive in common scenes. Increasing visual prompts from one to five consistently improves performance for all methods, with GFSAM [68] showing the most substantial average gain (+8.7 mIoU). The high standard deviation observed in some domains (± 14.1 for GFSAM [68] on Pizza) confirms our hypothesis about prompt selection sen-

| MODEL | Prompt | Common Scene | | Urban | | Waste | | Food | |
|--------------------------------------------------------|-----------------|--------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|
| | | PASCAL VOC | ADE20K | Cityscapes | UAVid | Trash | ZeroWaste | Pizza | UECFood |
| Open-vocabulary methods (textual prompt) | | | | | | | | | |
| MaskCLIP [10] | <i>Textual</i> | 38.8 | 9.8 | 12.6 | 23.5 | 4.3 | 6.6 | 25.5 | 4.5 |
| TCL [8] | <i>Textual</i> | 51.2 | 14.9 | 23.1 | 11.2 | 19.0 | 9.6 | 26.1 | 10.3 |
| CLIP-DINoiser [63] | <i>Textual</i> | 62.1 | 20.0 | 31.7 | 25.4 | 13.8 | 14.5 | 37.6 | 13.9 |
| NACLIP [19] | <i>Textual</i> | 52.1 | 17.3 | 31.4 | 25.8 | 18.5 | 10.4 | 43.5 | 16.6 |
| ProxyCLIP (DINOv2) [25] | <i>Textual</i> | 58.6 | 21.6 | 35.2 | 29.3 | 18.7 | 14.3 | 49.2 | <u>22.7</u> |
| ProxyCLIP (DINO) [25] | <i>Textual</i> | <u>60.6</u> | <u>22.6</u> | 40.1 | 30.2 | 21.1 | <u>15.1</u> | <u>52.5</u> | 22.6 |
| Visual reference prompt methods (visual prompt) | | | | | | | | | |
| SINE [33] | <i>1 prompt</i> | 31.2±2.6 | 15.9±0.7 | 42.7±2.2 | <u>31.0±6.1</u> | 9.4±1.5 | 8.3±0.8 | 15.2±1.9 | 1.3±0.2 |
| | <i>5 prompt</i> | 37.2±1.7 | 18.3±0.2 | <u>44.1±1.0</u> | 36.9±1.6 | 14.4±2.4 | 8.1±1.5 | 19.5±2.6 | 2.7±0.4 |
| PerSAM [72] | <i>1 prompt</i> | 10.1±1.2 | 2.7±0.2 | 15.1±0.4 | 6.6±1.8 | 3.6±1.0 | 3.4±1.1 | 16.8±1.1 | 1.9±0.2 |
| | <i>5 prompt</i> | - | - | - | - | - | - | - | - |
| Matcher [32] | <i>1 prompt</i> | 35.3±1.9 | - [†] | 30.6±3.7 | 14.0±2.0 | 32.6±3.7 | 10.5±3.0 | 36.0±9.8 | - [†] |
| | <i>5 prompt</i> | 42.9±1.6 | - [†] | 37.7±1.8 | 20.5±1.1 | <u>46.6±3.2</u> | 13.1±1.6 | 40.2±2.8 | - [†] |
| GFSAM [68] | <i>1 prompt</i> | 34.0±3.6 | 17.4±1.0 | 36.5±3.3 | 21.7±3.8 | 38.6±4.4 | 14.2±3.4 | 51.6±14.1 | 17.9±1.4 |
| | <i>5 prompt</i> | 45.0±1.4 | 23.4±0.6 | 44.6±0.9 | 29.7±1.6 | 52.5±5.6 | 16.8±3.7 | 62.2±2.7 | 26.8±1.1 |

| MODEL | Prompt | Land-Cover | | Tools | | Parts | | AVG |
|--------------------------------------------------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|------|
| | | LoveDA-Rural | LoveDA-Urban | Toolkits | PIDray | House-Parts | MHPv1 | |
| Open-vocabulary methods (textual prompt) | | | | | | | | |
| MaskCLIP [10] | <i>Textual</i> | 14.0 | 26.1 | 2.0 | 2.0 | 4.6 | 12.4 | 13.3 |
| TCL [8] | <i>Textual</i> | 12.3 | 14.9 | 6.5 | 4.8 | 7.8 | 11.5 | 15.9 |
| CLIP-DINoiser [63] | <i>Textual</i> | 25.3 | 35.3 | 6.0 | 3.5 | 7.3 | 17.3 | 22.4 |
| NACLIP [19] | <i>Textual</i> | 24.2 | 31.1 | 12.9 | 6.4 | 7.4 | 20.9 | 22.8 |
| ProxyCLIP (DINOv2) [25] | <i>Textual</i> | 31.4 | <u>40.3</u> | 3.4 | 7.4 | 3.4 | 21.5 | 25.5 |
| ProxyCLIP (DINO) [25] | <i>Textual</i> | <u>32.7</u> | 35.4 | 6.4 | 7.5 | 3.6 | <u>28.1</u> | 27.1 |
| Visual reference prompt methods (visual prompt) | | | | | | | | |
| SINE [33] | <i>1 prompt</i> | 14.5±5.0 | 26.7±4.7 | 43.7±8.6 | 1.8±0.3 | 12.3±1.8 | 8.3±1.1 | 18.7 |
| | <i>5 prompt</i> | 20.3±3.9 | 31.9±4.2 | 64.0±9.2 | 3.2±0.3 | 14.7±1.7 | 10.0±0.5 | 23.2 |
| PerSAM [72] | <i>1 prompt</i> | 8.1±2.1 | 8.7±3.0 | 21.7±7.5 | 2.3±0.4 | 4.0±1.6 | 8.3±0.8 | 8.1 |
| | <i>5 prompt</i> | - | - | - | - | - | - | - |
| Matcher [32] | <i>1 prompt</i> | 18.3±3.1 | 22.7±6.2 | 85.6±9.7 | 6.4±1.8 | 16.0±2.4 | 16.3±2.0 | 27.2 |
| | <i>5 prompt</i> | 25.4±2.6 | 30.5±4.8 | 88.9±2.9 | 10.3±1.1 | 25.0±5.0 | 19.3±2.2 | 33.5 |
| GFSAM [68] | <i>1 prompt</i> | 30.5±8.2 | 36.4±6.0 | 71.6±3.7 | <u>14.0±4.5</u> | 18.8±3.2 | 21.4±1.8 | 30.3 |
| | <i>5 prompt</i> | 35.6±1.4 | 43.4±3.1 | <u>86.9±1.4</u> | 21.0±2.6 | <u>24.7±2.6</u> | 28.6±1.8 | 38.7 |

Table 3. Benchmark results for the selected model. We present the mIoU for each model, and for visual reference prompt methods, we also include the standard deviation across different seeds. Note that for PerSAM [72], only the configuration with one prompt is reported due to its limitation in handling multiple prompts. [†] denotes experiments that required excessive time to yield results.

sitivity. Interestingly, the relative performance of prompting modalities dramatically shifts across domains. In common scenes, textual prompts excel on PASCAL VOC [14] (CLIP-DINoiser: 62.1 mIoU vs. GFSAM: 45.0 mIoU), yet this advantage reverses on ADE20K [73] where visual prompting slightly prevails. This pattern contrasts sharply in specialized domains like waste recognition, where GFSAM (52.5 mIoU) outperforms ProxyCLIP (21.1 mIoU) by an impressive 31.4 points on the Trash dataset, and tools, where Matcher achieves an exceptional 88.9 mIoU on Toolkits versus NACLIP’s 12.9 mIoU. The urban do-

main falls between these extremes, with visual methods maintaining a moderate but consistent advantage on both Cityscapes [12] (GFSAM: 44.6 mIoU vs. ProxyCLIP: 40.1 mIoU) and UAVid [35] (SINE: 36.9 mIoU vs. ProxyCLIP: 30.2 mIoU). Perhaps, the most revealing is the parts domain, where the performance gap varies dramatically between datasets: visual methods vastly outperform textual approaches on House-Parts [53] (Matcher: 25.0 mIoU vs. TCL: 7.8 mIoU), yet the difference nearly vanishes on MHPv1 [27] (GFSAM: 28.6 mIoU vs. ProxyCLIP: 28.1 mIoU). This suggests textual descriptions adequately cap-

ture human anatomy but struggle with architectural elements that benefit from visual reference. Similarly, on PIDray’s [71] X-ray scans, GFSAM (21.0 mIoU) triples the performance of textual methods (7.5 mIoU), demonstrating textual prompts’ inability to capture the unique visual characteristics of tools in security screening contexts.

Our results further validate our strategic emphasis on training-free methods, as trained approaches like SINE [33] and TCL [8] exhibit limited cross-domain adaptability. Despite SINE’s [33] strong performance on Cityscapes [12] (44.1 mIoU), it performs poorly on specialized domains like waste (14.4 mIoU and 8.1 mIoU) and food (19.5 mIoU and 2.7 mIoU). This stark contrast highlights the superior adaptability of training-free methods across diverse visual domains, particularly when leveraging visual prompts that can capture domain-specific visual characteristics that textual descriptions often fail to convey.

4.3. Qualitative results

In Figure 3, we compare the segmentation prediction obtained on our benchmark via textual prompts and a single visual prompts using ProxyCLIP [25] and GFSAM [68].

Segmentation masks generated using visual prompts demonstrate superior boundary refinement and shape definition compared to those produced with textual prompts. This phenomenon is particularly pronounced in the Trash [40] and Toolkits [37] datasets, where visual prompts effectively capture intricate object details, such as the cardboard (*yellow class, Trash*) and the Allen key (*red class, Toolkits*), with significantly higher precision than their textual counterparts. In more challenging datasets such as PIDray [71] and MHPv1 [27], both prompting approaches demonstrate limitations in generating comprehensive masks for the specified classes. Although visual prompting successfully segments the sprayer in the X-ray scans of the PIDray [71] dataset (*blue class*), it exhibits reduced efficacy in delineating the various body parts in the MHPv1 [27] dataset. Conversely, in datasets such as ADE20K [73], Pizza [66], UECFOOD [13], Cityscapes [12], and PASCAL VOC [14], both methods produce accurate results. However, in PASCAL VOC [14], the limitations of a single prompt become evident: despite accurate segmentation of target classes, the final prediction contains extraneous classes not present in the image. A significant qualitative observation emerges from the UAVid [35] dataset. In this instance, textual prompts generate an approximate segmentation of the entire scene, whereas visual prompting demonstrates reduced efficacy in capturing numerous buildings and distant vegetation. This observation reveals a fundamental limitation of visual prompting: its diminished capacity to effectively segment small objects in complex, cluttered environments. Textual prompts, despite their inherent imprecision, provide more comprehensive segmentation in such scenarios.

| | ADE20K | | Cityscapes | | PIDray | |
|--------------------|----------|----------|------------|--------|----------|---------|
| | #Forward | Time | #Forward | Time | #Forward | Time |
| CLIP-DiNoiser [63] | 1 | ~ 0.08s | 1 | ~ 0.1s | 1 | ~ 0.07s |
| NACLIP [19] | 1 | ~ 0.05s | 1 | ~ 0.2s | 1 | ~ 0.05s |
| ProxyCLIP [25] | 1 | ~ 0.6s | 1 | ~ 0.5s | 1 | ~ 0.2s |
| PerSAM [72] | 150 | ~ 320s | 19 | ~ 40s | 12 | ~ 25s |
| Matcher [32] | 150 | ~ 1,800s | 19 | ~ 240s | 12 | ~ 150s |
| GFSAM [68] | 150 | ~ 220s | 19 | ~ 30s | 12 | ~ 19s |

Table 4. Inference time and the number of forward required by visual reference (1 prompt) and open-vocabulary methods to produce a segmentation mask for a **single image**. Inference time is measured on single NVIDIA L4.

4.4. Computational analysis

Table 4 presents the number of forward passes required by each method to generate a segmentation mask for all classes in a given image, along with the corresponding inference time (in seconds) per image. Open-vocabulary methods (upper part of the table) only require a single forward pass and produce segmentation masks in less than a second. CLIP-DiNoiser [63] and NACLIP [19] demonstrate particularly efficient performance, with inference times below 0.2 seconds per image. Conversely, ProxyCLIP [25], which achieves superior performance on average, exhibits longer inference times ranging from 0.2 to 0.6 seconds. This increased computational demand can be attributed to its integration of DINO [5, 41].

As detailed in Section 3.2, visual reference methods (bottom part of the table) requires n forward passes, where n represents the number of classes. Consequently, inference times for these methods are substantially larger than textual prompting and range from 8 to over 1800 seconds per image. Matcher [32] exhibits the highest computational demand among visual reference methods, with inference times exceeding 30 minutes on ADE20K [73], primarily due to its dependence on SAM’s computationally intensive automatic mask generator [23]. In contrast, GFSAM [68] demonstrates the most efficient performance, with an inference time of 220 seconds on ADE20K [73].

5. Conclusions

We presented Show or Tell (SoT), a comprehensive benchmark evaluating both visual and textual prompts for semantic segmentation across 14 datasets spanning 7 diverse domains. Our experiments with 5 open-vocabulary and 4 visual reference prompt methods revealed distinct strengths and limitations of each prompting modality.

Open-vocabulary methods excelled in common scenes and urban environments where concepts are readily described textually. ProxyCLIP [25] variants achieved the highest scores on PASCAL VOC [14], ADE20K [73], and Cityscapes [12]. However, these methods performed poorly

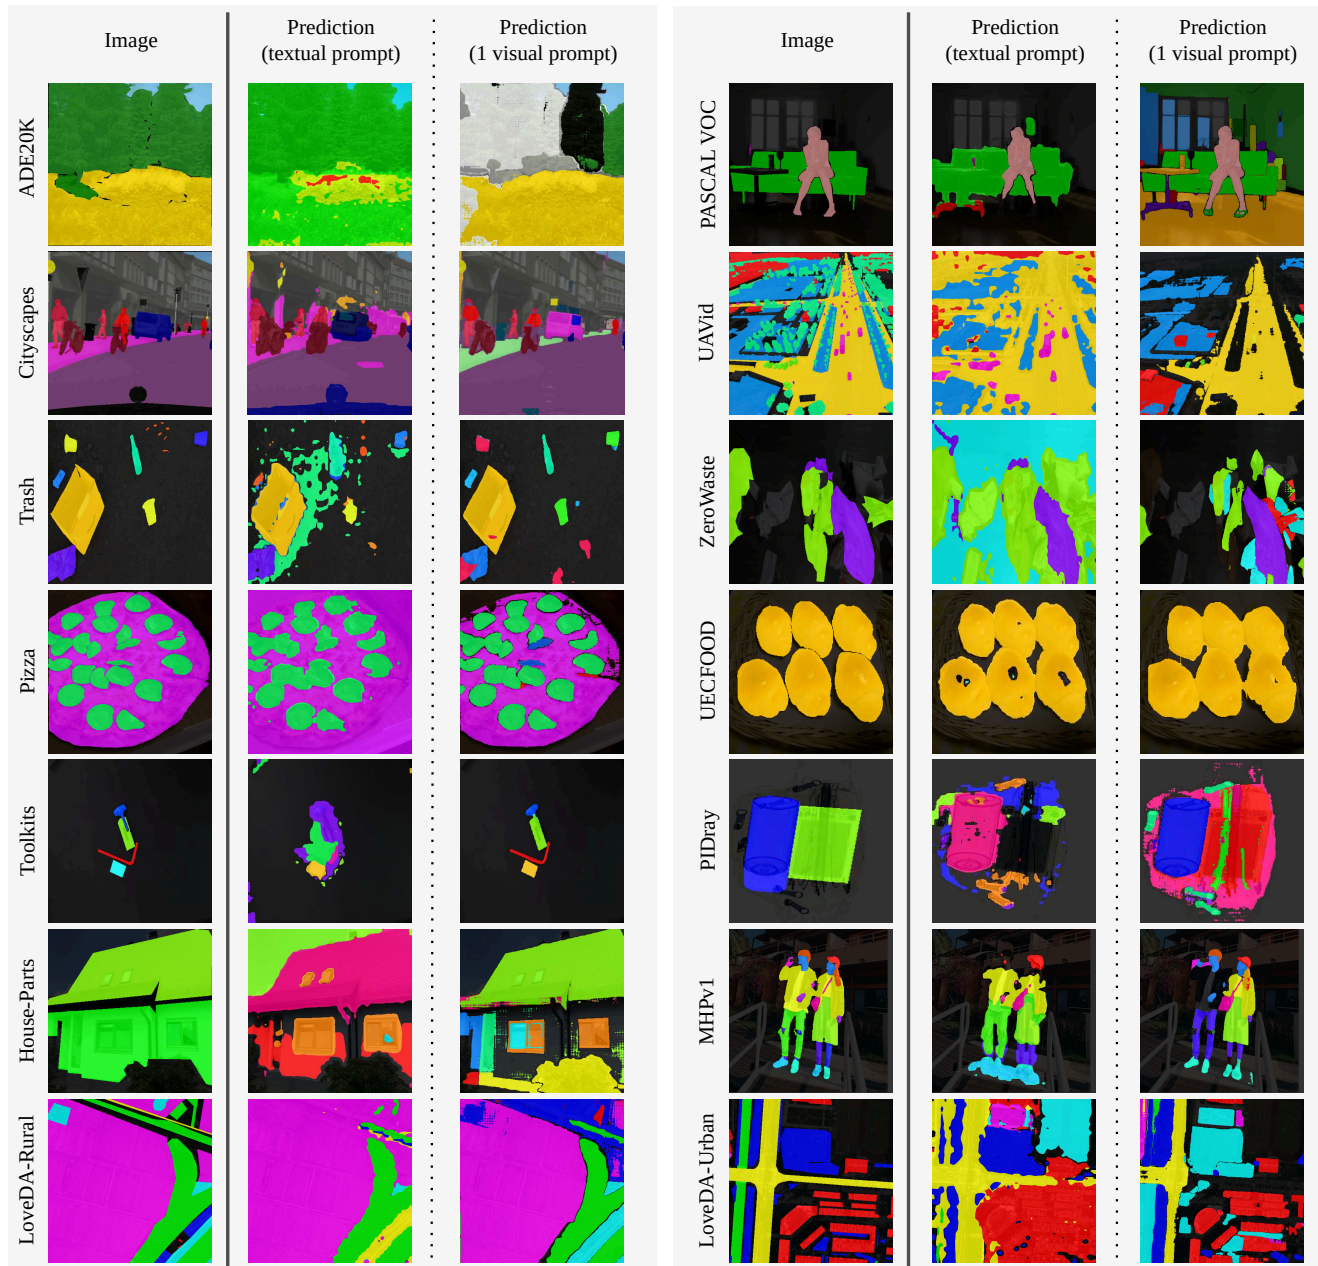


Figure 3. Qualitative results of ProxyCLIP [25] (*textual prompt*) and GFSAM [68] (*visual prompt*) across all the dataset in our benchmark.

in specialized domains like tools and parts, where textual descriptions inadequately capture complex visual characteristics. Visual reference prompt methods demonstrated more consistent performance across domains, with GFSAM [68] showing strong results when provided with multiple visual prompts. These methods performed very well in specialized categories like tools (88.9% on Toolkits) where visual examples offer more effective guidance than text descriptions. Their performance, however, varied considerably depending on the quality and representativeness of prompts, revealing a significant sensitivity to prompt selection.

Open challenges Our benchmark identifies several key challenges for future research.

- **Computational efficiency:** Visual reference prompt methods process binary class-by-class, resulting in slower inference compared to open-vocabulary approaches.
- **Prompt selection sensitivity:** Visual in-context learning methods show high sensitivity to prompt selection, undermining their reliability in practical applications.
- **Domain specialization:** Open-vocabulary methods struggle with specialized domains due to poor text-to-visual feature alignment.

Acknowledgements This publication is part of the project PNRR-NGEU which has received funding from the MUR – DM 117/2023. We acknowledge IS CRA for awarding this project access to the LEONARDO supercomputer, owned by the EuroHPC Joint Undertaking, hosted by CINECA (Italy).

References

- [1] Luca Barsellotti, Roberto Amoroso, Lorenzo Baraldi, and Rita Cucchiara. Enhancing open-vocabulary semantic segmentation with prototype retrieval. In *International Conference on Image Analysis and Processing*, 2023. 1, 3
- [2] Dina Bashkirova, Mohamed Abdelfattah, Ziliang Zhu, James Akl, Fadi Alladkani, Ping Hu, Vitaly Ablavsky, Berk Calli, Sarah Adel Bargal, and Kate Saenko. Zerowaste dataset: Towards deformable object segmentation in cluttered scenes. In *CVPR*, 2022. 4, 5
- [3] Benedikt Blumenstiel, Johannes Jakubik, Hilde Kühne, and Michael Vössing. What a mess: Multi-domain evaluation of zero-shot semantic segmentation. *Advances in Neural Information Processing Systems*, 36:73299–73311, 2023. 2, 4
- [4] Maxime Bucher, Tuan-Hung Vu, Matthieu Cord, and Patrick Pérez. Zero-shot semantic segmentation. *NeurIPS*, 32, 2019. 2
- [5] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *ICCV*, 2021. 2, 3, 4, 7
- [6] Fabio Cermelli, Massimiliano Mancini, Samuel Rota Bulo, Elisa Ricci, and Barbara Caputo. Modeling the background for incremental learning in semantic segmentation. In *CVPR*, 2020. 1
- [7] Fabio Cermelli, Massimiliano Mancini, Yongqin Xian, Zeynep Akata, and Barbara Caputo. Prototype-based incremental few-shot semantic segmentation. *arXiv preprint arXiv:2012.01415*, 2020. 3, 4
- [8] Junbum Cha, Jonghwan Mun, and Byungseok Roh. Learning to generate text-grounded mask for open-world semantic segmentation from only image-text pairs. In *CVPR*, 2023. 1, 3, 6, 7
- [9] Banghao Chen, Zhaofeng Zhang, Nicolas Langrené, and Shengxin Zhu. Unleashing the potential of prompt engineering in large language models: a comprehensive review. *arXiv preprint arXiv:2310.14735*, 2023. 1
- [10] Jun Chen, Deyao Zhu, Guocheng Qian, Bernard Ghanem, Zhicheng Yan, Chenchen Zhu, Fanyi Xiao, Sean Chang Culatana, and Mohamed Elhoseiny. Exploring open-vocabulary semantic segmentation from clip vision encoder distillation only. In *ICCV*, 2023. 2, 3, 6
- [11] MMSegmentation Contributors. MMSegmentation: Openmmlab semantic segmentation toolbox and benchmark. <https://github.com/open-mmlab/mms Segmentation>, 2020. 5
- [12] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *CVPR*, 2016. 4, 5, 6, 7
- [13] Takumi Ege, Wataru Shimoda, and Keiji Yanai. A new large-scale food image segmentation dataset and its application to food calorie estimation based on grains of rice. In *Proceedings of the 5th international workshop on multimedia assisted dietary management*, 2019. 4, 5, 7
- [14] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88:303–338, 2010. 4, 5, 6, 7
- [15] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*, 2017. 1
- [16] Golnaz Ghiasi, Xiuye Gu, Yin Cui, and Tsung-Yi Lin. Scaling open-vocabulary image segmentation with image-level labels. In *ECCV*, 2022. 1, 3
- [17] Zhaxuan Gu, Siyuan Zhou, Li Niu, Zihan Zhao, and Liqing Zhang. Context-aware feature generation for zero-shot semantic segmentation. In *Proceedings of the 28th ACM International Conference on Multimedia*, 2020. 2
- [18] Jie Guo, Qimeng Wang, Yan Gao, Xiaolong Jiang, Shaohui Lin, and Baochang Zhang. Mvp-seg: Multi-view prompt learning for open-vocabulary semantic segmentation. In *Chinese Conference on Pattern Recognition and Computer Vision (PRCV)*, 2023. 2
- [19] Sina Hajimiri, Ismail Ben Ayed, and Jose Dolz. Pay attention to your neighbours: Training-free open-vocabulary semantic segmentation. *arXiv preprint arXiv:2404.08181*, 2024. 2, 3, 6, 7
- [20] Sunghwan Hong, Seokju Cho, Jisu Nam, Stephen Lin, and Seungryong Kim. Cost aggregation with 4d convolutional swin transformer for few-shot segmentation. In *ECCV*. Springer, 2022. 3
- [21] Gabriel Ilharco, Mitchell Wortsman, Ross Wightman, Cade Gordon, Nicholas Carlini, Rohan Taori, Achal Dave, Vaishaal Shankar, Hongseok Namkoong, John Miller, Hananeh Hajishirzi, Ali Farhadi, and Ludwig Schmidt. Openclip. https://github.com/mlfoundations/open_clip, 2021. 1, 3, 4
- [22] Laurynas Karazija, Iro Laina, Andrea Vedaldi, and Christian Rupprecht. Diffusion models for zero-shot open-vocabulary segmentation. *arXiv e-prints*, pages arXiv–2306, 2023. 3
- [23] 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 *ICCV*, 2023. 1, 2, 3, 4, 7
- [24] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017. 1
- [25] Mengcheng Lan, Chaofeng Chen, Yiping Ke, Xinjiang Wang, Litong Feng, and Wayne Zhang. Proxyclip: Proxy attention improves clip for open-vocabulary segmentation. In *ECCV*, 2024. 2, 3, 6, 7, 8

- [26] Gen Li, Varun Jampani, Laura Sevilla-Lara, Deqing Sun, Jonghyun Kim, and Joongkyu Kim. Adaptive prototype learning and allocation for few-shot segmentation. In *CVPR*, 2021. 3
- [27] Jianshu Li, Jian Zhao, Yunchao Wei, Congyan Lang, Yidong Li, Terence Sim, Shuicheng Yan, and Jiashi Feng. Multiple-human parsing in the wild. *arXiv preprint arXiv:1705.07206*, 2017. 4, 5, 6, 7
- [28] Yi Li, Hualiang Wang, Yiqun Duan, and Xiaomeng Li. Clip surgery for better explainability with enhancement in open-vocabulary tasks. *arXiv e-prints*, pages arXiv-2304, 2023. 3
- [29] Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2935–2947, 2017. 1
- [30] Feng Liang, Bichen Wu, Xiaoliang Dai, Kunpeng Li, Yanan Zhao, Hang Zhang, Peizhao Zhang, Peter Vajda, and Diana Marculescu. Open-vocabulary semantic segmentation with mask-adapted clip. In *CVPR*, 2023. 1, 3
- [31] Yuanwei Liu, Nian Liu, Qinglong Cao, Xiwen Yao, Junwei Han, and Ling Shao. Learning non-target knowledge for few-shot semantic segmentation. In *CVPR*, 2022. 2
- [32] Yang Liu, Muzhi Zhu, Hengtao Li, Hao Chen, Xinlong Wang, and Chunhua Shen. Matcher: Segment anything with one shot using all-purpose feature matching. *arXiv preprint arXiv:2305.13310*, 2023. 2, 3, 4, 6, 7
- [33] Yang Liu, Chenchen Jing, Hengtao Li, Muzhi Zhu, Hao Chen, Xinlong Wang, and Chunhua Shen. A simple image segmentation framework via in-context examples. In *NeurIPS*, 2024. 2, 3, 4, 6, 7
- [34] Timo Lüddecke and Alexander Ecker. Image segmentation using text and image prompts. In *CVPR*, 2022. 2
- [35] Ye Lyu, George Vosselman, Gui-Song Xia, Alper Yilmaz, and Michael Ying Yang. Uavid: A semantic segmentation dataset for uav imagery. *ISPRS journal of photogrammetry and remote sensing*, 165:108–119, 2020. 4, 5, 6, 7
- [36] Juhong Min, Dahyun Kang, and Minsu Cho. Hypercorrelation squeeze for few-shot segmentation. In *ICCV*, 2021. 3
- [37] mst. mask dataset. <https://universe.roboflow.com/mst/mask-2ihnt>, 2022. visited on 2025-03-06. 4, 5, 7
- [38] Jishnu Mukhoti, Tsung-Yu Lin, Omid Poursaeed, Rui Wang, Ashish Shah, Philip HS Torr, and Ser-Nam Lim. Open vocabulary semantic segmentation with patch aligned contrastive learning. In *CVPR*, 2023. 1, 2, 3
- [39] Savinay Nagendra, Kashif Rashid, Chaopeng Shen, and Daniel Kifer. Samic: Segment anything with in-context spatial prompt engineering. *arXiv preprint arXiv:2412.11998*, 2024. 2
- [40] Sara Najafi. Trash (v2). <https://universe.roboflow.com/sara-najafi/trash-segmentation2/dataset/2>, 2022. 4, 5, 7
- [41] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. 2, 3, 4, 7
- [42] German I Parisi, Ronald Kemker, Jose L Part, Christopher Kanan, and Stefan Wermter. Continual lifelong learning with neural networks: A review. *Neural networks*, 113:54–71, 2019. 1
- [43] Giuseppe Pastore, Fabio Cermelli, Yongqin Xian, Massimiliano Mancini, Zeynep Akata, and Barbara Caputo. A closer look at self-training for zero-label semantic segmentation. In *CVPR*, 2021. 2
- [44] Jie Qin, Jie Wu, Pengxiang Yan, Ming Li, Ren Yuxi, Xuefeng Xiao, Yitong Wang, Rui Wang, Shilei Wen, Xin Pan, et al. FreeSeg: Unified, universal and open-vocabulary image segmentation. In *CVPR*, 2023. 1, 3
- [45] 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 *ICML*, 2021. 1, 2, 3, 4
- [46] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *ICML*, 2021. 2
- [47] 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. 2, 3
- [48] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. 2
- [49] Amirreza Shaban, Shray Bansal, Zhen Liu, Irfan Essa, and Byron Boots. One-shot learning for semantic segmentation. *arXiv preprint arXiv:1709.03410*, 2017. 3, 4
- [50] Gyungin Shin, Weidi Xie, and Samuel Albanie. Reco: Retrieve and co-segment for zero-shot transfer. *NeurIPS*, 35: 33754–33767, 2022. 3
- [51] Aleksandar Shtedritski, Christian Rupprecht, and Andrea Vedaldi. What does clip know about a red circle? visual prompt engineering for vlms. In *ICCV*, 2023. 2
- [52] Yanpeng Sun, Jiahui Chen, Shan Zhang, Xinyu Zhang, Qiang Chen, Gang Zhang, Errui Ding, Jingdong Wang, and Zechao Li. Vrp-sam: Sam with visual reference prompt. In *CVPR*, 2024. 2, 3, 4
- [53] TestCoco. abc dataset. <https://universe.roboflow.com/testcoco/abc-fqun0>, 2022. visited on 2025-03-06. 4, 5, 6
- [54] Sebastian Thrun and Lorien Pratt. *Learning to Learn*. Springer, 1998. 1
- [55] Feng Wang, Jieru Mei, and Alan Yuille. Sclip: Rethinking self-attention for dense vision-language inference. In *ECCV*, 2024. 3
- [56] Haochen Wang, Xudong Zhang, Yutao Hu, Yandan Yang, Xianbin Cao, and Xiantong Zhen. Few-shot semantic segmentation with democratic attention networks. In *ECCV*, 2020. 3
- [57] Haoxiang Wang, Pavan Kumar Anasosalu Vasu, Fartash Faghri, Raviteja Vemulapalli, Mehrdad Farajtabar, Sachin Mehta, Mohammad Rastegari, Oncel Tuzel, and Hadi Pouransari. Sam-clip: Merging vision foundation models

- towards semantic and spatial understanding. In *CVPR*, 2024. 3
- [58] Junjue Wang, Zhuo Zheng, Ailong Ma, Xiaoyan Lu, and Yanfei Zhong. Loveda: A remote sensing land-cover dataset for domain adaptive semantic segmentation. In *NeurIPS*, 2021. 4, 5
- [59] Jiaqi Wang, Zhengliang Liu, Lin Zhao, Zihao Wu, Chong Ma, Sigang Yu, Haixing Dai, Qiushi Yang, Yiheng Liu, Songyao Zhang, et al. Review of large vision models and visual prompt engineering. *Meta-Radiology*, 1(3):100047, 2023. 1
- [60] Kaixin Wang, Jun Hao Liew, Yingtian Zou, Daquan Zhou, and Jiashi Feng. Panet: Few-shot image semantic segmentation with prototype alignment. In *ICCV*, 2019. 3
- [61] Xinlong Wang, Wen Wang, Yue Cao, Chunhua Shen, and Tiejun Huang. Images speak in images: A generalist painter for in-context visual learning. In *CVPR*, 2023. 3
- [62] Xinlong Wang, Xiaosong Zhang, Yue Cao, Wen Wang, Chunhua Shen, and Tiejun Huang. Seggpt: Towards segmenting everything in context. In *ICCV*, 2023. 3
- [63] Monika Wysoczańska, Oriane Siméoni, Michaël Ramamonjisoa, Andrei Bursuc, Tomasz Trzciniński, and Patrick Pérez. Clip-dinoiser: Teaching clip a few dino tricks for open-vocabulary semantic segmentation. In *ECCV*, 2024. 2, 3, 6, 7
- [64] Yongqin Xian, Subhabrata Choudhury, Yang He, Bernt Schiele, and Zeynep Akata. Semantic projection network for zero-and few-label semantic segmentation. In *CVPR*, 2019. 2, 4
- [65] Aoran Xiao, Weihao Xuan, Heli Qi, Yun Xing, Ruijie Ren, Xiaoqin Zhang, Ling Shao, and Shijian Lu. Cat-sam: Conditional tuning for few-shot adaptation of segment anything model. In *ECCV*, 2024. 3
- [66] y. piiz dataset. <https://universe.roboflow.com/y-rgb4q/piiz>, 2023. visited on 2025-03-06. 4, 5, 7
- [67] Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *ICML*, 2017. 1
- [68] Anqi Zhang, Guangyu Gao, Jianbo Jiao, Chi Liu, and Yunchao Wei. Bridge the points: Graph-based few-shot segment anything semantically. *NeurIPS*, 37:33232–33261, 2024. 2, 3, 4, 5, 6, 7, 8
- [69] Chi Zhang, Guosheng Lin, Fayao Liu, Jiushuang Guo, Qingyao Wu, and Rui Yao. Pyramid graph networks with connection attentions for region-based one-shot semantic segmentation. In *ICCV*, 2019. 3
- [70] Gengwei Zhang, Guoliang Kang, Yi Yang, and Yunchao Wei. Few-shot segmentation via cycle-consistent transformer. *NeurIPS*, 34:21984–21996, 2021. 3
- [71] Libo Zhang, Lutao Jiang, Ruyi Ji, and Heng Fan. Pidray: A large-scale x-ray benchmark for real-world prohibited item detection. *International Journal of Computer Vision*, 131(12):3170–3192, 2023. 4, 5, 7
- [72] Renrui Zhang, Zhengkai Jiang, Ziyu Guo, Shilin Yan, Junting Pan, Xianzheng Ma, Hao Dong, Peng Gao, and Hongsheng Li. Personalize segment anything model with one shot. *arXiv preprint arXiv:2305.03048*, 2023. 2, 3, 4, 6, 7
- [73] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In *CVPR*, 2017. 4, 5, 6, 7
- [74] Chong Zhou, Chen Change Loy, and Bo Dai. Extract free dense labels from clip. In *ECCV*, 2022. 3
- [75] 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 *CVPR*, 2023. 2, 4