

# Annotation Free Spacecraft Detection and Segmentation using Vision Language Models

Samet Hicsonmez, Jose Sosa, Dan Pineau, Inder Pal Singh, Arunkumar Rathinam,  
Abd El Rahman Shabayek, and Djamila Aouada

**Abstract**—Vision Language Models (VLMs) have demonstrated remarkable performance in open-world zero-shot visual recognition. However, their potential in space-related applications remains largely unexplored. In the space domain, accurate manual annotation is particularly challenging due to factors such as low visibility, illumination variations, and object blending with planetary backgrounds. Developing methods that can detect and segment spacecraft and orbital targets without requiring extensive manual labeling is therefore of critical importance. In this work, we propose an annotation-free detection and segmentation pipeline for space targets using VLMs. Our approach begins by automatically generating pseudo-labels for a small subset of unlabeled real data with a pre-trained VLM. These pseudo-labels are then leveraged in a teacher-student label distillation framework to train lightweight models. Despite the inherent noise in the pseudo-labels, the distillation process leads to substantial performance gains over direct zero-shot VLM inference. Experimental evaluations on the SPARK-2024, SPEED+, and TANGO datasets on segmentation tasks demonstrate consistent improvements in average precision (AP) by up to 10 points. Code and models are available at <https://github.com/gidddyupp/annotation-free-spacecraft-segmentation>.

## I. INTRODUCTION

Reliable detection and localization of spacecraft is a crucial step in achieving safe and autonomous operations in space [1] for missions such as in-orbit servicing and active debris removal. Current object detection and segmentation models [2], [3] for space targets rely on large amounts of annotated data to achieve high performance. However, in the space domain, collecting labeled data is highly challenging due to limited accessibility, high operational costs, and the difficulty of manual annotation under space conditions (see Figure 1). To address the scarcity of labeled data, synthetic datasets [4], [5] generated in simulation environments [6] are widely used as the primary training source. Such datasets are scalable and straightforward to generate and annotate. However, models trained exclusively on synthetic data often show a significant drop in performance when applied to real-world scenarios [7] due to the domain gap.

VLMs have shown unprecedented performance in open-world zero-shot visual recognition tasks such as classification [8], object detection [9], and semantic segmentation [10]–[13]. However, for particular domains, their zero-shot performance often falls short compared to models

The authors are with Interdisciplinary Centre for Security, Reliability and Trust (SnT), University of Luxembourg, Luxembourg. Email: {firstname.lastname}@uni.lu.  
The models and code are available at <https://gitlab.com/uniluxembourg/snt/cvi2/space/elite/vfm/afss>




Fig. 1. We present challenging examples taken from SPEED+ [4] (top) and TANGO [17] (bottom) where annotation is difficult due to the cluttered background and low visibility. Our method detects and segments the spacecraft object with high precision in all cases. Zoom in for details.

trained or fine-tuned with annotated data [14]. In the space domain, only a few studies have explored VLMs [15], [16] for spacecraft detection or segmentation. Nevertheless, these methods do not fully exploit the zero-shot capabilities of VLMs. Instead, they use them mainly as auxiliary tools for supervised training, providing limited analysis, and presenting sub-par zero-shot performances.

Beyond the space domain, several strategies have been proposed to improve model performance without requiring additional manual annotations. For instance, Test-Time Augmentation (TTA) [18]–[22] is a widely used method to improve model performance during inference by aggregating predictions coming from different views of the input image. For object detection, the predictions are usually aggregated using Non-Maximum Suppression (NMS) variants [23], [24] with a predefined overlap ratio. However, this approach eliminates many good candidates and considers only the most confident prediction. Recently, Weighted Box Fusion (WBF) [25] was proposed to adjust the final prediction based on the confidence of the overlapped detections. Another line of research explores Knowledge Distillation (KD) [26] techniques, which transfer the knowledge of large pre-trained teacher models into smaller and more efficient student models [27]–[33]. Usually, the distilled small student model achieves either on-par or better performance compared to the larger teacher model, thanks to the regularization effect of pseudo-labels and reduced overfitting of the small model size.

Building on these ideas, we propose an end-to-end annotation-free framework to improve the zero-shot performance of VLMs for spacecraft detection and segmentation. Our method begins with a small set of unlabeled real-world images, for which pseudo-labels are automatically generated using a pre-trained VLM. To enhance their reliability, we re-

fine these pseudo-labels through TTA and WBF. The refined pseudo-labels are then employed in a teacher–student distillation scheme, where compact student models are trained using the pseudo-labeled data. This process not only boosts performance over direct zero-shot inference, but also yields lightweight models that are efficient and practical for real-world deployment.

We conduct extensive experiments on three large-scale space datasets and validate that our approach significantly improves the initial zero-shot detection and segmentation performance. To the best of our knowledge, the proposed pipeline is the first to exploit VLMs for the challenging task of spacecraft detection and segmentation. This approach bypasses the need for extensive manual annotation, a major bottleneck in traditional computer vision methods for space applications. By leveraging VLMs, our pipeline can automatically create high-quality pseudo-labels from a small, unlabeled dataset, drastically reducing the time and cost associated with data preparation. The proposed method offers a scalable and adaptable solution for a wide range of computer vision tasks in the aerospace domain and beyond. By eliminating the reliance on manually labeled datasets, it enables rapid development and deployment of robust object detection and segmentation models for new and complex Space Situational Awareness (SSA) scenarios, such as tracking and monitoring space debris motion.

## II. RELATED WORK

This section presents the relevant works on Vision Language Models and their applications in the space domain, followed by TTA methods on improving inference robustness and accuracy. Finally, recent developments in knowledge distillation to achieve an efficient student model are presented.

### A. Vision Language Models

Recent VLMs [8], [34] show strong zero-shot capabilities in image-level tasks. Open-vocabulary detection and segmentation models [9], [10], [12], [13], [35], [36] go further and extend this capability to dense pixel-level tasks.

OpenSeed [12] aligns textual embeddings from pre-trained language models with dense pixel-level features from a mask proposal backbone, enabling accurate segmentation of unseen categories. The method introduces a joint training strategy with contrastive objectives that bridge vision and language modalities, achieving competitive performance across open-vocabulary semantic and panoptic segmentation benchmarks. SEEM [13] extends this line of work by proposing a unified model capable of handling multiple segmentation tasks, i.e. semantic, instance, panoptic, and interactive, through flexible conditioning on prompts such as text, points, or boxes. This promptable design demonstrates strong generalization and adaptability, making SEEM a versatile framework for open-world scenarios.

Grounded-SAM [10], on the other hand, integrates the grounding capabilities of GroundingDINO [9], i.e. the ability to connect and align textual descriptions with specific objects or regions within an image, with the high-quality mask

generation of SAM [11], enabling open-vocabulary detection and segmentation directly from natural language queries. By leveraging the complementary strengths of grounding and segmentation, Grounded-SAM achieves robust zero-shot generalization across diverse object categories.

Despite the success of VLMs in natural images, their performance on space applications has not been fully explored. Recently, [15] experimented with different text prompts to detect spacecraft bounding boxes in a zero-shot manner on the SPEED+ [4] dataset using GroundingDINO [9]. However, their work is limited to a single dataset and task, and fails to exploit the full potential of VLMs (c.f. see Table III).

In [16], authors propose to use SAM [11] to automatically segment spacecraft. They train an object detection model within a supervised framework. Then, they use the predicted boxes to prompt the SAM. This approach assumes the availability of the bounding boxes to first train the object detection method. However, our method does not require any labeled data and performs detection and segmentation simultaneously.

### B. Test-time Augmentation (TTA)

TTA is a widely used technique in various visual detection tasks to improve robustness and accuracy at inference time [18]–[21], [25], [37]–[39]. The core idea is to apply a set of transformations such as horizontal or vertical flipping, scaling, and color perturbations to the input image, run the model on each augmented version, and then aggregate the predictions. While simple ensembling across augmentations can yield consistent gains, for object detection, it introduces the challenge of merging multiple overlapping detections.

Traditional strategies such as Non-Maximum Suppression (NMS) [23], [24] often discard potentially useful predictions when bounding boxes overlap, leading to suboptimal results. To address this, Weighted Boxes Fusion (WBF) [40] has emerged as a more effective merging strategy. Instead of suppressing overlapping boxes, WBF computes a weighted average of their coordinates, where the weights are determined by the associated confidence scores. This produces more accurate bounding boxes and preserves valuable information from multiple augmented predictions. Recent studies have shown that combining TTA with WBF provides consistent improvements in detection and segmentation benchmarks, striking a balance between robustness and precision.

### C. Label Distillation

Knowledge distillation [26] is a well-established approach for model compression, in which a high-capacity teacher network transfers its learned representations to a smaller and more efficient student model. The pseudo labels coming from the teacher model could be treated as ground truth by directly using them as hard labels, i.e. one-hot encoded. However, this approach overlooks the possible low-confidence detections. One approach to overcome this issue is using soft labels [41]–[44], i.e. incorporating the prediction confidences with specialized loss functions. Generalized




Fig. 2. The processing pipeline of our method. Given unlabeled training images, we first automatically annotate them using a pretrained VLM with a fixed prompt of “spacecraft”. In the next stage, we refine these pseudo-annotated images leveraging test-time augmentations. The refined labels are distilled to a shallow Student Model which is employed for the inference.

Focal Loss [45] extends focal loss by jointly modeling classification confidence and localization quality, while Varifocal Loss [46] builds on this idea by weighting high-quality samples more heavily, enabling detectors to better exploit predictions aligned with their IoU scores.

An alternative approach to reduce the pseudo label noise is applying confidence thresholding [27], [47], [48] to the detections. In this setting, only the high-quality detections of the teacher are used to train the student model. Due to its simplicity, we prefer to use this approach in our trainings.

### III. METHOD

We propose a pipeline for spacecraft detection and segmentation that does not assume the availability of manually annotated data. Our approach exploits the zero-shot capabilities of pre-trained VLMs to generate pseudo-labels, which are subsequently refined through an iterative teacher–student distillation framework. Then, the student model is deployed for inference. Figure 2 illustrates the overall workflow, which comprises four stages:

- 1) **Pseudo-labeling** stage utilizes pre-trained VLMs to generate necessary pseudo-labels.
- 2) **Label refinement** stage adopts TTA together with WBF and confidence-based thresholding to remove noisy or inaccurate predictions.
- 3) **Label distillation** stage implements an iterative distillation process involving a teacher and student model. The teacher produces pseudo-labels to train a simpler student model.
- 4) **Inference** stage employs solely the student model that has been trained with refined pseudo-labels.

The following subsections will describe in more detail each of the steps in our proposed pipeline.

#### A. Pseudo-label generation

The first stage of our pipeline employs a pre-trained VLMs to automatically annotate a given set of images. Rather than relying on manual labeling, the VLMs are prompted to produce pseudo-labels that capture instance segmentation masks and bounding box predictions.

Formally, for a given image  $I$  and a textual prompt  $T$ , the VLM  $\mathcal{V}$  produces 2D instance segmentation and bounding box predictions  $P_S, P_B = \mathcal{V}(I, T)$  respectively. Since the space datasets contain a single instance in images, we choose only the top detection and/or segmentation predictions of the models. This step generates a single pseudo-box and mask annotation for each training image.

To be concise, the subsequent sub-sections focus on the particulars of bounding box predictions. However, this approach is equally applicable to segmentation masks.

#### B. Label refinement

Given the domain gap between the data used to pre-train VLMs and the nature of space data, their zero-shot predictions might be inaccurate. To address this limitation, we introduce a refinement stage within our pipeline, where the initially generated pseudo-labels are improved to mitigate noise and enhance label quality. In particular, we follow standard refinement strategies, such as using TTA with WBF and confidence-based thresholding.

1) *Refinement using TTA*: We use test time augmentation strategy to enhance the quality of pseudo-labels by first compiling a set of augmentations applied to input images, followed by pooling the predictions using WBF.

Let  $\mathcal{A} = \{a_1, a_2, \dots, a_K\}$  represent the formally defined augmentation set. For an input image  $I$ ,  $K$  augmented versions are generated, each determined by an augmentation in  $\mathcal{A}$ :

$$I_k = a_k(I), \quad k = 1, \dots, K. \quad (1)$$

Then, each augmented image is passed through  $\mathcal{V}(\cdot)$ , which predicts a set of bounding boxes:

$$\mathcal{P}_{B_k} = \mathcal{V}(I_k, T) = \{(b_i^{(k)}, s_i^{(k)}, \ell_i^{(k)})\}_{i=1}^{N_k}, \quad (2)$$

where  $b_i^{(k)} \in [0, 1]^4$  are the normalized box coordinates,  $s_i^{(k)} \in [0, 1]$  is the confidence score, and  $\ell_i^{(k)}$  is the predicted class. Each box is then mapped back to the original image coordinates using the inverse transform  $a_k^{-1}$ :

$$\tilde{b}_i^{(k)} = a_k^{-1}(b_i^{(k)}). \quad (3)$$

Let  $\mathcal{B} = \{\tilde{b}_i^{(k)}, s_i^{(k)}, \ell_i^{(k)}\}$  be the union of all predictions from all the augmentations. WBF groups boxes with Intersection-over-Union (IoU) greater than a threshold  $\tau$ .

For each cluster  $C = \{(b_j, s_j)\}_{j=1}^M$ , the fused box  $b^*$  is computed as a weighted average:

$$b^* = \frac{\sum_{j=1}^M s_j b_j}{\sum_{j=1}^M s_j}, \quad s^* = \frac{1}{M} \sum_{j=1}^M s_j \quad (4)$$

where  $s^*$  is the fused confidence score. The assigned class label is determined by majority voting or weighted voting based on scores.

The final set of predictions after TTA and WBF is defined:

$$\mathcal{P}_B^* = \{(b_m^*, s_m^*, \ell_m^*)\}_{m=1}^{M^*}. \quad (5)$$

2) *Refinement using confidence based thresholding*: Upon implementing WBF, we acquire a set of fused predictions with associated confidence scores. However, when aggregated predictions show low IoU, WBF often leads to a significant reduction in confidence. To further enhance the quality of the predictions, we applied a confidence-based thresholding. This filtering procedure removes noisy, low-confidence predictions while retaining only those with high reliability.

$$\mathcal{P}_{B\theta}^* = \left\{ (b_m^*, s_m^*, \ell_m^*) \mid s_m^* \geq \theta \right\}, \quad (6)$$

where  $\theta \in [0, 1]$  is the confidence threshold.

### C. Label distillation

In this stage, the refined pseudo-labels are considered as outputs from a teacher model, and we train a shallow student model using iterative knowledge distillation. This phase offers three key advantages: i) label-distilled student models often outperform the teacher models, ii) a significant reduction in model size is achieved, and iii) the more compact model can operate in real-time and facilitate autonomous operations.

Given the confidence-filtered TTA pseudo box predictions  $\mathcal{P}_{B\theta}^*$  generated by the teacher model (VLM), the predictions  $\mathcal{Q}$  from the student model  $S_\phi$  for a specific image  $I$  are defined as:

$$\mathcal{Q} = S_\phi(I) = \{(b_j, p_j)\}_{j=1}^{N_S}, \quad (7)$$

where  $b_j$  and  $p_j$  are the predicted box coordinates and associated class labels, respectively.

In the distillation process, the student model undergoes training with the standard supervised loss:

$$\mathcal{L}_{\text{distill}} = \alpha \mathcal{L}_{\text{cls}}(p, \hat{y}) + \beta \mathcal{L}_{\text{reg}}(b, \hat{b}) \quad (8)$$

where  $\hat{b}, \hat{y}$  are the box and class predictions of the teacher VLM, respectively.  $p, b$  are the student's predicted class probability and box, and finally,  $\alpha, \beta$  are the weighting factors between the classification loss  $\mathcal{L}_{\text{cls}}$  and the regression loss  $\mathcal{L}_{\text{reg}}$ .

Upon training the student model, its predictions are employed to assign new labels to the training images, acting as pseudo-labels. These are used to retrain the student model from the ground up. This process, referred to as iterative knowledge distillation, serves to improve pseudo-label accuracy and enhance the student model's performance. Our approach employs a single iteration of this distillation process.

### D. Inference

We utilized the distilled student model for inference. Owing to its lightweight design, the student facilitates real-time operation with minimal computational and memory requirements, ideal for in-orbit platform deployment. At the same time, it preserves much of the accuracy of the teacher model, offering a favorable trade-off between efficiency and performance.

## IV. EXPERIMENTS

In this section, we introduce the curated datasets employed in the experiments and summarize the upper bound performances (Oracle) achieved through fully-supervised training. Subsequently, we present the zero-shot baseline performance of VLMs in spacecraft detection and segmentation, followed by the impact of the label refinement and label distillation steps on enhancing this baseline.

### A. Datasets

In our experiments, we use widely adopted SPARK-2024 [5], SPEED+ [4] and TANGO [17] datasets.

**SPARK-2024** dataset [5] was part of the Spacecraft Pose Estimation Challenge at the AI4SPACE workshop held during CVPR'24. Stream-2 of the dataset is dedicated to spacecraft trajectory estimation and comprises sequences of both simulated images and real images acquired in the SnT Zero-G lab [6]. The training set includes 100 groups of trajectories, each containing 300 labeled synthetic images with ground-truth pose and bounding box annotations, but missing segmentation annotations. The test set consists of four real trajectory sequences comprising 2100 images in total. In our experiments, we used only the test sequences and created a training split with 500 randomly selected images and used the remaining images (1600) for evaluation.

**SPEED+** dataset [4] was part of the SPEED + challenge hosted by the European Space Agency (ESA) and provides both synthetic images and real images with high-fidelity images of the Tango spacecraft model captured under two distinct illumination conditions (Lightbox and Sunlamp). The Lightbox domain contains 6,740 images of the spacecraft model illuminated by diffuse albedo lightboxes, which simulate Earth's orbital lighting conditions. In contrast, the Sunlamp domain consists of 2,791 images of the same model illuminated by a metal-halide arc lamp, which approximates direct sunlight. Both splits contain box and segmentation annotations. Normally, these subsets are only used for testing, i.e., there is no official training split. However, since our aim is to distill from large VLMs, we create custom training and test datasets on these two datasets using the test splits. We follow the same procedure as SPARK-2024 and create training splits with 500 randomly selected images and use the remaining images (6200 and 2200) for evaluation.

**TANGO** dataset [17] is introduced to support monocular vision-based spacecraft pose estimation. The dataset includes laboratory images of the Tango spacecraft model generated under diverse lighting conditions, backgrounds, and camera poses, with accurate bounding box, segmentation, and 6-DoF pose annotations. The dataset has been widely used to validate detection, segmentation, and pose estimation algorithms in spaceborne vision applications. The Tango dataset comprises 30,002 training images and 3,002 test images. We used the test split as is and randomly selected 500 images from the training set.

TABLE I

ZERO-SHOT DETECTION AND SEGMENTATION PERFORMANCES OF VLMs ON THE SPARK-2024, SPEED+ AND TANGO TEST SETS.

	VFM	Detection			Segmentation		
		$AP$	$AP_{50}$	$AP_{75}$	$AP$	$AP_{50}$	$AP_{75}$
Spark	SEEM	17.7	49.8	2.6	N/A	N/A	N/A
	OpenSEED	8.5	28.0	0.8	N/A	N/A	N/A
	GSAM-2	53.3	97.9	61.7	N/A	N/A	N/A
SLamp	SEEM	25.3	45.4	23.7	28.0	43.6	31.1
	OpenSEED	33.1	63.3	32.1	40.2	62.8	45.3
	GSAM-2	73.9	97.9	92.1	72.5	94.7	85.0
LBox	SEEM	11.4	19.6	11.3	37.1	52.6	41.8
	OpenSEED	35.0	59.3	36.9	37.1	54.8	41.6
	GSAM-2	65.8	91.4	74.0	66.2	83.9	73.5
Tango	SEEM	25.7	60.5	7.7	41.4	69.7	41.5
	OpenSEED	23.8	87.7	0.4	62.2	91.6	77.9
	GSAM-2	58.2	91.8	75.7	58.2	85.9	69.0

### B. Metrics

In line with previous related work [49] on object detection and instance segmentation, we use Average Precision (AP) for evaluation. AP is defined as the area under the precision–recall curve, averaged over IoU thresholds from 0.50 to 0.95 with a step size of 0.05. In addition,  $AP_{50}$  and  $AP_{75}$  are commonly reported for reference.  $AP_{50}$  uses a loose matching criterion ( $\text{IoU} \geq 0.50$ ), while  $AP_{75}$  applies a stricter one ( $\text{IoU} \geq 0.75$ ).  $AP_{50}$  mainly reflects detection capability, while  $AP_{75}$  emphasizes localization accuracy.

### C. Oracle

To set the upper bound (Oracle) performance, we train two lightweight object detection models and one semantic segmentation model using ground-truth labels and evaluate them on the test sets. In our experiments, we adopt Efficient-Det [50] for object detection and YOLOv11 [51] for object detection and segmentation, which have around 7M parameters and run in real time, i.e. both have FPS more than 60. Both models are trained following the official implementations and the default training parameters for 300 epochs. During training and inference, the images are resized to  $640 \times 640$  pixels for both models.

Table III presents the results of the Oracle experiment (see the first row of each dataset). Since the SPARK-2024 dataset has no segmentation labels, we train only the object detection models. Overall, both detectors achieve strong results across all datasets by achieving more than 95% points on the challenging  $AP_{75}$  metric, which is crucial for performing pose estimation or instance segmentation on cropped objects. YOLOv11 achieves the best overall performance on all datasets except for the SPEED+ Lightbox subset.

### D. Analysis of VLMs on zero-shot detection

We started our experiments by first evaluating the performance of VLMs on three datasets. To this end, we selected three recent open-source VLMs [10], [12], [13] which achieve state-of-the-art results on open vocabulary detection and segmentation tasks on different benchmarks, and evaluated their zero-shot object detection performance

TABLE II

THE EFFECT OF UTILIZING TTA AND WBF WITH GROUNDED SAM-2.

Dataset	Detection			Segmentation		
	$AP$	$AP_{50}$	$AP_{75}$	$AP$	$AP_{50}$	$AP_{75}$
Spark	54.4	98.0	64.0	N/A	N/A	N/A
SLamp	74.6	98.9	93.7	73.2	95.8	86.3
LBox	66.9	91.7	75.6	66.8	85.0	73.8
Tango	58.5	92.0	76.3	58.7	86.4	69.6

on the curated test sets. Note that no fine-tuning is performed on these selected VLMs using any form of labeled data.

Table I presents the object detection and segmentation performances on the SPARK-2024, SPEED+, and TANGO datasets. We used *spacecraft* as the textual prompt  $T$  for each VLM and gathered only the top prediction for each image. Prediction confidences are used directly to calculate the metrics for detection and segmentation.

Except for the TANGO dataset, the SEEM and OpenSEED models achieve significantly low performance, especially on the detection task. OpenSEED achieves the best zero-shot segmentation result in the TANGO dataset, but its detection performance falls behind, especially on the  $AP_{75}$  metric. On the other hand, the GroundedSAM-2 [10] model shows promising results, e.g., 90%+ on  $AP_{50}$  detection task and 85%+ on  $AP_{50}$  segmentation task on each dataset. Hence, we only used the GroundedSAM-2 model as the baseline VLM in the next stages of our method.

### E. Effect of TTA and WBF

In order to improve the zero-shot detection performance of GroundedSAM-2, we experimented with different TTA settings. In natural images, the most common setting is to use the horizontal flip. Since space images have no notion of up and down, we added vertical flip to the augmentation list. Considering the nature of space images, we further introduced pixel manipulation augmentations such as brightness, saturation, and color to improve the overall contrast of the images.

We ran extensive experiments on various combinations of augmentations on all datasets. In each case, the object detection predictions are merged using WBF with a threshold  $\tau$  of 0.55. The best results are achieved using only the vertical flip as an augmentation. Increasing the number of augmentations significantly decreased the detection and segmentation performance, most likely due to the possible low-confidence detections.

Table II presents the results when TTA and WBF are employed for zero-shot detection. In each dataset, we observed more than 1.5 point improvement in the challenging  $AP_{75}$  metric for the detection task. Similarly, an improvement of up to 1.0 points is observed in the segmentation  $AP_{50}$  metric.

### F. Distillation

In order to train the shallow detectors, we first extract the pseudo-labels (i.e. bounding boxes and segmentation masks) for the training sets following the same pipeline with TTA and WBF. Then, we use these pseudo-labels

TABLE III

ORACLE AND DISTILLATION RESULTS ON THE SPARK, SPEED+ AND TANGO TEST SETS. CF: CONFIDENCE FILTERING, TR: TRAIN RELABELING.

Method	Detection						Segmentation			
	Efficient-Det			YOLOv11			YOLOv11			
	$AP$	$AP_{50}$	$AP_{75}$	$AP$	$AP_{50}$	$AP_{75}$	$AP$	$AP_{50}$	$AP_{75}$	
Spark	Oracle	98.7	99.0	99.0	99.2	99.5	99.5	N/A	N/A	N/A
	TTA+WBF	65.0	98.9	78.9	68.9	99.5	89.4	N/A	N/A	N/A
	TTA+WBF+CF	66.8	98.9	80.4	70.8	99.5	97.4	N/A	N/A	N/A
	TTA+WBF+CF+TR	68.7	98.9	80.5	76.4	99.5	98.2	N/A	N/A	N/A
Sunlamp	Oracle	81.6	99.0	97.8	85.0	99.5	98.8	81.5	94.3	94.1
	TTA+WBF	75.4	99.0	94.9	74.2	95.4	92.0	73.0	94.4	91.0
	TTA+WBF+CF	75.0	99.0	94.7	74.6	95.3	91.9	73.0	94.8	91.6
	TTA+WBF+CF+TR	74.9	99.0	94.7	75.0	95.3	92.6	74.0	94.3	91.0
Lightbox	Oracle	84.7	99.0	97.6	80.3	99.4	93.4	82.3	99.5	95.7
	TTA+WBF	66.3	94.9	76.0	69.1	96.4	79.3	61.6	91.1	72.6
	TTA+WBF+CF	66.4	97.0	76.8	69.6	96.5	80.1	62.5	92.1	72.6
	TTA+WBF+CF+TR	68.2	97.2	79.1	71.0	96.8	81.6	63.5	93.6	73.7
Tango	Oracle	87.4	94.8	93.0	92.1	99.4	97.7	79.1	99.1	96.2
	TTA+WBF	66.0	92.2	87.0	85.0	93.5	93.1	67.0	93.4	84.1
	TTA+WBF+CF	66.2	92.6	89.0	86.4	94.9	93.7	69.7	94.5	87.1
	TTA+WBF+CF+TR	70.6	92.1	90.4	86.5	95.6	93.8	70.0	94.7	86.9

as ground-truth annotations and train the Efficient-Det and YOLOv11 models for both object detection and semantic segmentation tasks. We use the same hyper-parameters as *Oracle* training, and similarly train the models for 300 epochs. Table III summarizes the experimental results using TTA and WBF for each dataset (see the second rows as TTA+WBF). On the SPARK-2024 dataset, the YOLOv11 model significantly improves detection performance; notably, the metrics  $AP$  and  $AP_{75}$  are increased by more than 14 and 25, respectively. Similarly, TANGO and Lightbox detection performances have improved considerably on every metric. Regarding segmentation results, except for the Lightbox dataset, baseline zero-shot performance is greatly improved; e.g., on the TANGO dataset  $AP$  and  $AP_{75}$  metrics are increased by almost 10 points.

Although the TTA improves the overall quality of the predictions, we observe that some low-confidence detections remain. Hence, before training the shallow models, we apply confidence filtering. We chose 0.5 for the SPARK-2024 and TANGO datasets and 0.6 for the SPEED+ dataset as confidence thresholds to remove low-quality detections. These values are selected on the basis of the mean confidences of the predictions for each dataset. After this step, the number of training images dropped to 395, 451, 417, and 471 for SPARK-2024, Sunlamp, Lightbox, and TANGO, respectively. We retrain detection and segmentation models in these confidence-filtered training sets. The results are presented in Table III in the third rows (marked as TTA+WBF+CF) of each dataset. We observe a consistent improvement on all datasets for both tasks. Notably, the detection performance of the SPARK-2024 dataset is increased by 8 points on the challenging  $AP_{75}$  metric. Similarly, TANGO segmentation

metrics are improved up to 3 points.

Here we take a step forward and employ the iterative knowledge distillation, i.e. using the current student model as the next teacher. We relabel the training sets using the trained student model and train the student model from scratch. Table III presents the results of this iterative distillation experiment in the last rows (denoted TTA + WBF + CF + TR). In the SPARK-2024 dataset, the detection  $AP$  is improved by more than 5 points. On other datasets, we observe up to 3 point improvements for different metrics both for detection and segmentation tasks. However, in rare cases, there happens a small drop in performance, e.g., segmentation results of the Sunlamp dataset, most likely due to overfitting.

Compared to the initial VLM predictions, our proposed pipeline significantly improves (e.g., more than 10 points on TANGO segmentation  $AP$ ) the detection and segmentation results on three challenging datasets, even with very shallow student models.

### G. Visual Results

In Figure 3, we present visual results containing the input image, ground-truth labels, zero-shot VLM predictions, and our student model predictions for each dataset. Our method improves both detection and segmentation performance even when the pseudo-labels contain significant noise, as in the cases of the TANGO dataset. These results show the importance of label distillation in improving performance.

## V. CONCLUSION

In this work, we present an end-to-end spacecraft detection and segmentation pipeline that does not rely on any form




Fig. 3. Visual comparison of zero-shot VLM and our student model predictions. Our method greatly improves VLM predictions.

of manually labeled data. Our method exploits the zero-shot capabilities of recent VLMs. It first extracts object bounding boxes and instance masks using a VLM, then utilizes these pseudo-labels in an iterative teacher-student distillation framework. Extensive experiments on three recent large-scale space datasets show that our method significantly improves initial zero-shot VLM results and reduces the gap

to fully supervised training.

#### ACKNOWLEDGMENT

This work is supported by the Luxembourg National Research Fund (FNR) under the project reference C21/IS/15965298/ELITE. The experiments were performed on the Luxembourg National Supercomputer

## REFERENCES

- [1] H. A. Dung, B. Chen, and T.-J. Chin, "A spacecraft dataset for detection, segmentation and parts recognition," in *CVPR*, 2021.
- [2] G. Murray, T. Bourlai, and M. Spolaor, "Mask r-cnn: Detection performance on speed spacecraft with image degradation," in *IEEE International Conference on Big Data*, 2021.
- [3] H. Vighh, S. Loughran, Y. Rachlin, R. Allen, and J. Ruprecht, "Training deep learning spacecraft component detection algorithms using synthetic image data," in *IEEE Aerospace Conference*, 2023.
- [4] T. H. Park, M. Märtens, G. Lecuyer, D. Izzo, and S. D'Amico, "Speed+: Next-generation dataset for spacecraft pose estimation across domain gap," in *IEEE Aerospace Conference*, 2022.
- [5] "Spark 2024: Datasets for spacecraft semantic segmentation and spacecraft trajectory estimation." [Online]. Available: <https://doi.org/10.5281/zenodo.10908214>
- [6] M. Olivares-Mendez, M. R. Makhdoomi, B. C. Yalçın, Z. Bokal, V. Muralidharan, M. O. Del Castillo, V. Gaudilliere, L. Pauly, O. Borgue, M. Alandihallaj, *et al.*, "Zero-g lab: A multi-purpose facility for emulating space operations," *Journal of Space Safety Engineering*, vol. 10, no. 4, pp. 509–521, 2023.
- [7] S. Hashimoto, D. Hirano, H. Higuchi, K. Fujimoto, M. Ito, S. Iizuka, Y. Sugimoto, R. Proffitt, Y. Kobayashi, and N. Ishihama, "Domain adaptation for 6-dof pose estimation: Deep learning with uncertainty," in *IEEE Aerospace Conference*, 2022.
- [8] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, *et al.*, "Learning transferable visual models from natural language supervision," in *ICML*, 2021.
- [9] S. Liu, Z. Zeng, T. Ren, F. Li, H. Zhang, J. Yang, Q. Jiang, C. Li, J. Yang, H. Su, *et al.*, "Grounding dino: Marrying dino with grounded pre-training for open-set object detection," in *ECCV*, 2024.
- [10] T. Ren, S. Liu, A. Zeng, J. Lin, K. Li, H. Cao, J. Chen, X. Huang, Y. Chen, F. Yan, Z. Zeng, H. Zhang, F. Li, J. Yang, H. Li, Q. Jiang, and L. Zhang, "Grounded sam: Assembling open-world models for diverse visual tasks," 2024.
- [11] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo, *et al.*, "Segment anything," in *ICCV*, 2023.
- [12] H. Zhang, F. Li, X. Zou, S. Liu, C. Li, J. Yang, and L. Zhang, "A simple framework for open-vocabulary segmentation and detection," in *ICCV*, 2023.
- [13] X. Zou, J. Yang, H. Zhang, F. Li, L. Li, J. Wang, L. Wang, J. Gao, and Y. J. Lee, "Segment everything everywhere all at once," *NeurIPS*, 2023.
- [14] Y. Feng, Y. Liu, S. Yang, W. Cai, J. Zhang, Q. Zhan, Z. Huang, H. Yan, Q. Wan, C. Liu, *et al.*, "Vision-language model for object detection and segmentation: A review and evaluation," *arXiv preprint arXiv:2504.09480*, 2025.
- [15] M. Ulmer, L. Klüpfel, M. Durner, and R. Triebel, "How important are data augmentations to close the domain gap for object detection in orbit?" in *IEEE Aerospace Conference*, 2025.
- [16] Z. Zhang and Z. Dang, "Automatic segmentation annotation of space target using segment anything model and object detection prompts," *IEEE Transactions on Aerospace and Electronic Systems*, 2024.
- [17] M. Bechini, M. Lavagna, and P. Lunghi, "Dataset generation and validation for spacecraft pose estimation via monocular images processing," *Acta Astronautica*, vol. 204, pp. 358–369, 2023.
- [18] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," *NeurIPS*, 2012.
- [19] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *CVPR*, 2016.
- [20] M. Kimura, "Understanding test-time augmentation," in *NeurIPS*, 2021.
- [21] T.-A. Vu, S. Sarkar, Z. Zhang, B.-S. Hua, and S.-K. Yeung, "Test-time augmentation for 3d point cloud classification and segmentation," in *International Conference on 3D Vision*, 2024, pp. 1543–1553.
- [22] N. Samet, G. Puy, and R. Marlet, "Losc: Lidar open-voc segmentation consolidator," *arXiv preprint arXiv:2507.07605*, 2025.
- [23] A. Neubeck and L. Van Gool, "Efficient non-maximum suppression," in *ICPR*, 2006.
- [24] N. Bodla, B. Singh, R. Chellappa, and L. S. Davis, "Soft-nms – improving object detection with one line of code," in *ICCV*, 2017.
- [25] D. Shanmugam, D. Blalock, G. Balakrishnan, and J. Guttag, "Better aggregation in test-time augmentation," in *ICCV*, 2021.
- [26] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," *arXiv preprint arXiv:1503.02531*, 2015.
- [27] K. Sohn, Z. Zhang, C.-L. Li, H. Zhang, C.-Y. Lee, and T. Pfister, "A simple semi-supervised learning framework for object detection," *arXiv preprint arXiv:2005.04757*, 2020.
- [28] M. Xu, Z. Zhang, H. Hu, J. Wang, L. Wang, F. Wei, X. Bai, and Z. Liu, "End-to-end semi-supervised object detection with soft teacher," in *ICCV*, 2021.
- [29] H. Touvron, M. Cord, M. Douze, F. Massa, A. Sablayrolles, and H. Jégou, "Training data-efficient image transformers & distillation through attention," in *ICML*, 2021.
- [30] Q. Xie, M.-T. Luong, E. Hovy, and Q. V. Le, "Self-training with noisy student improves imagenet classification," in *CVPR*, 2020.
- [31] V. Sanh, L. Debut, J. Chaumond, and T. Wolf, "Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter," *arXiv preprint arXiv:1910.01108*, 2019.
- [32] X. Jiao, Y. Yin, L. Shang, X. Jiang, X. Chen, L. Li, F. Wang, and Q. Liu, "Tinybert: Distilling bert for natural language understanding," in *EMNLP*, 2020.
- [33] K. Wu, H. Peng, Z. Zhou, B. Xiao, M. Liu, L. Yuan, H. Xuan, M. Valenzuela, X. S. Chen, X. Wang, H. Chao, and H. Hu, "Tinyclip: Clip distillation via affinity mimicking and weight inheritance," in *ICCV*, 2023.
- [34] C. Jia, Y. Yang, Y. Xia, Y.-T. Chen, Z. Parekh, H. Pham, Q. Le, Y.-H. Sung, Z. Li, and T. Duerig, "Scaling up visual and vision-language representation learning with noisy text supervision," in *ICML*, 2021.
- [35] T. Ren, Q. Jiang, S. Liu, Z. Zeng, W. Liu, H. Gao, H. Huang, Z. Ma, X. Jiang, Y. Chen, Y. Xiong, H. Zhang, F. Li, P. Tang, K. Yu, and L. Zhang, "Grounding dino 1.5: Advance the "edge" of open-set object detection," 2024.
- [36] F. Li, H. Zhang, P. Sun, X. Zou, S. Liu, J. Yang, C. Li, L. Zhang, and J. Gao, "Semantic-sam: Segment and recognize anything at any granularity," *ECCV*, 2024.
- [37] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.
- [38] I. Kim, Y. Kim, and S. Kim, "Learning loss for test-time augmentation," in *NeurIPS*, 2020.
- [39] N. Moshkov, B. Mathe, A. Kertesz-Farkas, R. Hollandi, and P. Horvath, "Test-time augmentation for deep learning-based cell segmentation on microscopy images," *Scientific reports*, vol. 10, no. 1, 2020.
- [40] R. Solovyev, W. Wang, and T. Gabruseva, "Weighted boxes fusion: Ensembling boxes from different object detection models," *Image and Vision Computing*, vol. 107, p. 104117, 2021.
- [41] Y. Li, J. Yang, Y. Song, L. Cao, J. Luo, and L.-J. Li, "Learning from noisy labels with distillation," in *ICCV*, 2017.
- [42] L. Yuan, F. E. Tay, G. Li, T. Wang, and J. Feng, "Revisiting knowledge distillation via label smoothing regularization," in *CVPR*, 2020.
- [43] B. Zi, S. Zhao, X. Ma, and Y.-G. Jiang, "Revisiting adversarial robustness distillation: Robust soft labels make student better," in *ICCV*, 2021.
- [44] Z. Yang, A. Zeng, Z. Li, T. Zhang, C. Yuan, and Y. Li, "From knowledge distillation to self-knowledge distillation: A unified approach with normalized loss and customized soft labels," in *ICCV*, 2023.
- [45] X. Li, W. Wang, L. Wu, S. Chen, X. Hu, J. Li, J. Tang, and J. Yang, "Generalized focal loss: Learning qualified and distributed bounding boxes for dense object detection," in *NeurIPS*, 2020.
- [46] H. Zhang, Y. Wang, F. Dayoub, and N. Sunderhauf, "Varifocalnet: An iou-aware dense object detector," in *CVPR*, 2021.
- [47] D.-H. Lee *et al.*, "Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks," in *Workshop on challenges in representation learning, ICML*, 2013.
- [48] Y.-C. Liu, C.-Y. Ma, Z. He, C.-W. Kuo, K. Chen, P. Zhang, B. Wu, Z. Kira, and P. Vajda, "Unbiased teacher for semi-supervised object detection," in *ICLR*, 2021.
- [49] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft coco: Common objects in context," in *ECCV*, 2014.
- [50] M. Tan, R. Pang, and Q. V. Le, "Efficientdet: Scalable and efficient object detection," in *CVPR*, 2020.
- [51] R. Khanam and M. Hussain, "Yolov11: An overview of the key architectural enhancements," *arXiv preprint arXiv:2410.17725*, 2024.