

Traffic Sign Detection in Panoramic Imagery using SAM3 and YOLOv8 in Low-Data Scenarios

Khadija Ben Madane, Raphaël Couturier, Michel Salomon

Université Marie et Louis Pasteur

CNRS, Institut FEMTO-ST

F-90000 Belfort, France

{khadija.ben_madane, raphael.couturier, michel.salomon}@umlp.fr

Shanci Li, Rémy Marquis, Bertil Chapuis

INSIT, IICT

HEIG-VD

1400 Yverdon-les-Bains, Switzerland

{shanci.li, remy.marquis, bertil.chapuis}@heig-vd.ch

Abstract—Traffic sign detection and classification are essential tasks for intelligent transportation systems and road infrastructure monitoring. However, most existing approaches rely on fully supervised object detection models that require large annotated datasets and often struggle with complex scenarios such as panoramic imagery where image sizes are large.

In this work, we propose a hybrid pipeline that combines the Segment Anything Model (SAM3) for prompt-based object localization and YOLOv8m-cls for traffic sign classification. Unlike traditional methods, our approach does not require task-specific training for the detection stage, making it particularly suitable for low-data and real-world environments. The use of SAM enables flexible region extraction in high-resolution 360° images. To show the relevance of our proposal, we evaluate our method on a dataset of panoramic images in challenging conditions such as small and distant objects, scene complexity, and geometric distortions. Experimental results demonstrate strong performance, achieving a mAP@0.50 of 94.03%. Thus, this work highlights the potential of combining foundation models with specialized classifiers to develop scalable and practical solutions for real-world traffic sign analysis.

Index Terms—Traffic Sign Detection, Panoramic Imagery, Segment Anything Model, YOLOv8m-cls

I. INTRODUCTION

Traffic sign detection and recognition have become crucial tasks due to the rapid development of intelligent transportation systems and the increasing demand for road safety and infrastructure monitoring. Deep learning has improved object detection and image processing considerably, with Convolutional Neural Networks (CNNs) being widely used for tasks such as traffic sign recognition. Popular models include Faster R-CNN [1], single-shot detectors like SSD [2], YOLO, and more recent approaches like DETR [3], which achieve strong performance in terms of accuracy and efficiency. However, these methods usually require large annotated datasets and are designed for standard image sizes. As a result, they are less effective in challenging conditions such as high-resolution panoramic images.

This work was supported by the Interreg France–Switzerland program (2021–2027), co-financed by the European Regional Development Fund (ERDF) and the Swiss Confederation, under the project NearAI (2025–2028). Computations have been partially performed on the supercomputer facilities of the Mesocenter of Bourgogne Franche-Comté. This work was also granted access to the HPC resources of IDRIS under the allocations [AD010613582R3 and AD010617392] made by GENCI.

Traffic sign detection in panoramic images becomes significantly more challenging, as these images are characterized by very high resolution, geometric distortions, and the presence of small objects. Moreover, collecting and annotating large datasets in such environments is both costly and time-consuming, particularly for specific tasks such as French traffic sign recognition. This lack of annotated data represents a major obstacle for the deployment of supervised detection models in practical applications, especially in real-world scenarios where labeled data is limited or unavailable.

To effectively address the above challenges, we propose a hybrid approach that combines the Segment Anything Model (SAM3) for object localization using prompts and YOLOv8 for traffic sign classification. The key idea is to leverage the generalization capabilities of a foundation model to perform object detection without requiring task-specific training, thereby reducing the dependency on annotated datasets. The detected regions are then classified using a YOLOv8 model, a well-known computer vision model, enabling accurate recognition of traffic sign categories.

The main contributions of this work are as follows:

- Traffic sign detection using the well-known Segment Anything Model, a state-of-the-art model for many computer vision tasks, with tailored prompts.
- Demonstrating the ability of YOLOv8 model to classify objects detected by SAM3.
- An examination of several situations in which the proposed approach needs to be refined, even though its performance is already quite satisfactory.

II. RELATED WORKS

Object detection, and in particular traffic sign detection, is a major topic in computer vision. It is generally divided into two main categories: traditional methods and deep learning-based approaches. Traditional methods use handcrafted features, based on the color and shape of traffic signs. These methods work by looking for specific colors and shapes. They use basic tools to find edges, match patterns, or Hough transforms to extract geometric features [4], [5]. They are effective in simple and controlled environments, but they still have important limitations with complex backgrounds, changes in lighting,

and occlusions. In addition, their lack of adaptability makes them less suitable for modern traffic sign detection systems.

With the rapid development of deep learning, new approaches have been introduced and can generally be grouped into three main categories: two-stage detectors, one-stage detectors, and Transformer-based detectors.

A. Two-stage detectors

Two-stage detectors usually follow a simple idea. First, they try to find possible object regions in the image. Then, these regions are analyzed to decide what object they contain and to adjust their position. R-CNN [1] was one of the first models to use this idea and showed better results than earlier methods. Later, Fast R-CNN [6] worked on the whole image only once and then focused on regions of interest which made the process faster. After that, the Region Proposal Network (RPN) was introduced in Faster R-CNN [7] to generate these regions more efficiently and made the system further faster. Mask R-CNN [8] is a more recent model that predicts the shape of the objects. That means it can detect and segment objects. Even if these methods give very good results, they are usually slower and need more computation than one-stage detectors.

B. One-stage detectors

One-stage detectors perform object localization and classification in a single forward pass, enabling efficient and real-time detection. They can be categorized into anchor-based and anchor-free approaches. Early anchor-based approaches, such as YOLO [9]–[12] and SSD [13], formulate detection as a regression problem by directly predicting bounding boxes and class probabilities. Extensions such as DSSD [14] improve SSD by enhancing feature representations for small object detection, while RetinaNet [15] introduces focal loss to address class imbalance.

Anchor-free methods aim to avoid the complexity associated with predefined anchor boxes. Keypoint-based methods, such as CornerNet [16], detect objects by predicting paired keypoints (e.g., top-left and bottom-right corners), while CenterNet [17] improves this approach by representing objects through their center points. Other anchor-free methods, such as FCOS [18], perform dense pixel-wise predictions inspired by segmentation, enabling simpler and more flexible detection pipelines. Additionally, methods such as FSAF [19] further enhance feature selection in anchor-free detection.

More recent developments include anchor-free variants of YOLO, such as YOLOv8 [20] and later versions (e.g., YOLOv11 and YOLOv26 [21]), which improve efficiency and simplify the detection process. Furthermore, Transformer-based approaches have been introduced within the one-stage detection paradigm. DETR (Detection Transformer) [3] formulates object detection as a direct set prediction problem and eliminates the need for hand-crafted components such as anchor boxes and Non-Maximum Suppression (NMS). Variants such as RT-DETR [22], TS-DETR [23], and TSD-DETR [24] have been proposed to improve efficiency and performance, particularly for real-time and small object detection. However,

these methods generally require large annotated datasets and high computational resources.

C. Open-vocabulary detection

Unlike traditional supervised detectors, these methods leverage large-scale pretraining and vision-language models to generalize beyond predefined categories. Models such as OWL-ViT and GroundingDINO perform object detection using textual prompts, enabling the localization of objects not seen during training. Similarly, approaches such as LLMDet and APE further improve open-vocabulary detection by integrating multimodal representations and Large Language Models. In parallel, foundation models such as the Segment Anything Model (SAM3) [25] introduce prompt-based segmentation, allowing object localization with minimal supervision through points, bounding boxes, or masks. These methods provide flexible and scalable solutions for object localization in scenarios where labeled data is limited.

III. PROPOSED METHOD

In this work, we address the challenge of traffic sign detection in panoramic images under limited annotation conditions. While classification datasets are available, annotated data for object detection is scarce, making it difficult to train detection models. To overcome this limitation, we propose a hybrid approach that does not rely on large-scale detection training data, combining a prompt-based segmentation model for object localization with a lightweight YOLOv8-based classifier.

As illustrated in Fig. 1, the proposed method follows a three-stage architecture. In the first stage, the input panoramic image is processed by SAM3 using a set of predefined textual prompts, namely “road sign” and “traffic light”, to guide the detection process. Each prompt is applied independently, and SAM3 generates a set of candidate regions for each prompt in the form of bounding boxes, associated confidence scores, and segmentation masks. In particular, using only the prompt “road sign” does not allow SAM3 to reliably detect traffic lights. Therefore, we added the prompt “traffic light” to improve detection coverage. Based on these prompts, SAM3 produces candidate bounding boxes with confidence scores, and their segmentation masks. To reduce noise, only detections with a confidence score above a predefined threshold are retained. The resulting bounding boxes are then used to extract the corresponding regions from the original image. Finally, the extracted regions are passed to a YOLOv8 classification model, which assigns each region to a traffic sign class.

The considered Segment Anything Model (SAM3) has shown strong performance in recent studies, outperforming several methods such as OWLv2, GroundingDINO, and previous versions like SAM and SAM2 [25]. It works well on both images and videos, which is useful in real-world applications. Another advantage of SAM3 is that it does not require task-specific training. It can detect and segment objects using simple prompts, which makes it suitable for situations where labeled data is limited, as in our case. We use SAM3 in a zero-shot setting, without modifying its architecture, and




Fig. 1. Proposed method for traffic sign detection and classification in panoramic images.




Fig. 2. Architecture of the prompt-based detection model (SAM3).

rely on textual prompts to detect traffic-related objects. As shown in Fig. 2, SAM3 is based on a transformer architecture using DETR, which stands for DEtection TRansformer, as decoder in output. It includes an image encoder to extract visual features and a prompt encoder to process text inputs.

These features are combined and passed to a decoder that predicts segmentation masks, bounding boxes, and confidence scores. Since our work focuses on panoramic images, we do not use image exemplars in SAM3 due to the large number of classes (more than 280). Instead, we rely on a classification

model. We use YOLOv8m-cls provided by Ultralytics. It is based on the standard YOLOv8 backbone with a classification head for predicting the class, which allows us to recognize traffic sign classes under different conditions, as each class is represented by multiple images captured from different angles, under varying lighting conditions, and in both clear and challenging situations such as rain or blur. This is particularly important, as traffic signs can appear under different conditions from one image to another.

IV. EXPERIMENTAL SETUP

A. Datasets

This study focuses on the detection and classification of French road signs. To achieve this objective, we rely on a dataset specifically tailored to the French traffic sign system. The primary dataset used in this work is France road signs classification dataset, publicly available on Hugging Face [26]. This dataset contains 67567 annotated instances of road signs extracted from street-level imagery provided by Panoramax. The data covers a wide range of real-world conditions, including different weather situations and daytime scenarios. It is composed of 258 classes, all labeled according to the official French traffic sign catalogue. However, despite its richness, certain classes in this dataset are underrepresented or missing. To address this limitation and improve class balance, we augmented our dataset with additional images sourced from the French Traffic Sign Classification Dataset (FTSC) [27]. This complementary dataset also contains French traffic signs captured under diverse environmental conditions and annotated following the same labeling standard. By combining these two datasets, we obtain a more comprehensive and balanced dataset, enhancing the model’s ability to generalize across different classes and real-world scenarios. The final dataset therefore includes a diverse set of French road signs with variations in lighting, weather, scale, and perspective, making it suitable for robust detection and classification tasks.

After merging the two datasets, data augmentation was applied to reduce class imbalance. Classes with fewer samples were augmented until reaching a target of 500 images per class, for a total of 140,755 images across 280 classes. This procedure generated synthetic variations of existing samples using only moderate transformations in order to preserve the semantic content of the traffic signs. The applied transformations are summarized in Table I. They include: random small-angle rotations in the range of $\pm 20^\circ$, mild perspective transformations, photometric adjustments such as brightness, contrast, color, and sharpness modifications, light Gaussian blur and salt-and-pepper noise. These transformations were designed to simulate realistic variations observed in road scenes, which helps the model generalize better to real-world driving conditions.

B. Test dataset

To evaluate the proposed model, a test dataset consisting of 420 manually annotated panoramic images from France was used. These images are 360 degree panoramas with a high

resolution of 5760×2880 pixels, providing detailed visual information suitable for both detection and classification tasks.

The dataset includes 116 traffic sign classes. This number is lower than in the training set due to the rarity of certain traffic sign categories, which are difficult to capture in real-world environments. Nevertheless, these rare classes are included in the training dataset to ensure better model generalization.

C. System configuration and parameters

The system configuration used for the experiments is summarized in Table II. For the SAM3 stage, the prompts “road sign” and “traffic light” were used. Images are processed by SAM3 at their original resolution and are internally resized to a default resolution of approximately 1008×1008 during processing. The resulting crops are directly passed to YOLOv8m-cls, which automatically resizes them to 224×224 . For classification, we fine-tuned YOLOv8m-cls model from Ultralytics, with an input resolution of 224×224 , a batch size of 64, and 300 training epochs. Training was performed using 8 workers on a single GPU. The minimum classification confidence threshold was fixed at 0.60 to filter out low-confidence predictions. The YOLOv11m-cls model was also trained and tested using the same parameters and environment as YOLOv8m-cls, to ensure a fair comparison between the two classification models within the proposed pipeline.

V. RESULTS AND DISCUSSION

The proposed model is evaluated on the test dataset to assess its performance using Precision, Recall, F1-score, and mean Average Precision at an IoU threshold of 0.5 (mAP@50). Precision and Recall measure the accuracy and completeness of the predictions, while the F1-score provides a balance between them. The mAP@50 metric evaluates detection performance based on the overlap between predicted and ground-truth bounding boxes.

Table III shows the quantitative results of the proposed pipeline. The SAM3 model, which is used only for traffic sign detection, achieves a very high Recall (99.90%), meaning that almost all traffic signs are detected. However, it also produces many false positives (FP = 225), which reduces the Precision (81.12%). The AP@0.50 value of 90.60% reflects its detection capability. Since both detection and classification are required, we combine SAM3 with classification models YOLOv8m-cls and YOLOv11m-cls. The combination of SAM3 with YOLOv8m-cls gives the best results, with a Precision of 92.45%, Recall of 96.07%, and an F1-score of 94.22%. The number of false positives is greatly reduced (from 225 to 76), showing that the classifier effectively removes incorrect detections. The Recall slightly decreases (from 99.90% to 96.07%), because some correct detections are rejected when the predicted class is incorrect or when the confidence is below the threshold (0.6). The mAP@0.50 value of 94.03% confirms the good performance of the full pipeline. In comparison, the SAM3 + YOLOv11m-cls model performs worse, with lower Precision (83.69%), Recall (89.05%), F1-score (86.29%), and mAP@0.50 (87.38%). This indicates that YOLOv11m-cls is

TABLE I
SUMMARY OF DATA AUGMENTATION TECHNIQUES APPLIED TO UNDERREPRESENTED CLASSES

Category	Transformation	Parameters / Range	Probability
Geometric	Rotation	Angle $\in [-20^\circ, +20^\circ]$	0.9
	Perspective transform	Max shift = 6% of width/height	0.4
Photometric	Brightness	Factor $\in [0.75, 1.25]$	0.9
	Contrast	Factor $\in [0.75, 1.30]$	0.9
	Color	Factor $\in [0.75, 1.25]$	0.6
	Sharpness	Factor $\in [0.5, 1.6]$	0.3
Image Degradation	Gaussian blur	Radius $\in [0.2, 1.0]$	0.25
	Salt & pepper noise	Noise ratio $\in [0.002, 0.01]$	0.25




Fig. 3. Examples of detection and classification results in some challenging panoramic scenes.

TABLE II
EXPERIMENTAL ENVIRONMENT CONFIGURATION

Component	Specification
CPU	Intel Xeon Gold 6252 (2.10 GHz), 96 cores
GPU	NVIDIA A100-PCIE
VRAM	40 GB
RAM	503 GB
Operating System	Rocky Linux 8.5
Python Version	3.9
PyTorch Version	2.4.1
CUDA Version	CUDA 12.4

less suitable for this task or dataset. The two pipeline configurations have similar inference times. SAM3 + YOLOv8m-cls takes 0.29 seconds per image, and SAM3 + YOLOv11m-cls takes 0.30 seconds. while SAM3 alone takes about 0.28 seconds per image. This shows that adding the classification stage only slightly increases the computation time. Overall, the

proposed method provides a good balance between detection and classification performance and computational efficiency, especially in low-data scenarios. While SAM3 alone is effective for detection, adding a classification stage is essential for traffic sign classification and recognition.

The objective of this study is to address low-data scenarios, where detection annotations are limited or unavailable. The proposed method is not designed to compete with fully supervised detectors. Instead, it shows that combining a prompt-based model (SAM3) with a classification model (YOLOv8-cl) can give good results on panoramic images without training a detector.

Figure 3 presents qualitative examples of the proposed method in challenging panoramic scenes. The results show that the method is able to detect and classify small and distant traffic signs, as illustrated in Fig. 3(d). The classifier (YOLOv8m-cl) also provides correct predictions for several small detected regions, which confirms its robustness under

TABLE III
COMPARISON OF SAM3 COMBINED WITH YOLOV8M-CLS AND YOLOV11M-CLS

Method pipeline		Outcomes confusion matrix				Evaluation metrics			Time (s/img)
Detection	Classification	TP	FP	FN	Precision	Recall	F1-score	Score	
SAM3		967	225	1	0.8112	0.9990	0.8954	AP@0.50 (0.9060)	0.2840
SAM3	YOLOv8m-cls	930	76	38	0.9245	0.9607	0.9422	mAP@0.50 (0.9403)	0.2974
SAM3	YOLOv11m-cls	862	168	106	0.8369	0.8905	0.8629	mAP@0.50 (0.8738)	0.3000

difficult conditions. However, some limitations can also be observed. In certain cases, two nearby traffic signs with similar shapes may be grouped and detected as a single object, especially when they are small and far from the camera, as shown in Fig. 3(a). In contrast, when the signs are clearer and more distinguishable, the method is able to detect them separately, as illustrated in Fig. 3(b). In addition, shape diversity appears to help the detection stage: when adjacent traffic signs have different shapes, they are more likely to be detected as distinct objects, even when they are relatively small, as shown in Fig. 3(c). Furthermore, the method is prone to false positives in some situations. As shown in Fig. 3(f), the model correctly detects a real traffic sign; however, this sign does not belong to any of the classes included in the training dataset. As a result, the classifier assigns it to the most similar known class. Similarly, in Fig. 3(e), a non-traffic object (commercial sign or logo) is incorrectly detected and classified as a traffic sign due to visual similarity. These observations highlight the limitations of the model in handling unseen classes and visually similar non-traffic objects.

VI. CONCLUSION

This paper introduces a hybrid approach for French traffic sign detection and classification in panoramic images by combining SAM3 with a YOLOv8 based classification model. This method does not require training for detection, which makes it suitable when annotated data is limited. The results show that the model performs well across different evaluation metrics, even in challenging conditions like small traffic signs. The results indicate that combining a prompt-based segmentation model with a classification model works well in this situation. However, there are still some limitations especially when dealing with very small or partially occluded traffic signs. Future work will focus on improving detection in these cases, for example by fine-tuning SAM3. Overall, the approach gives reliable results for traffic sign recognition without requiring large annotated datasets, which makes it suitable for real-world applications.

REFERENCES

- [1] Ren, S., He, K., Girshick, R., Sun, J., Faster, R.-C.N.N.: Towards Real-Time object detection with region proposal networks. *IEEE Trans. Pattern Anal. Mach. Intell.* 39(6), 1137–1149 (2017).
- [2] Liu, W., et al.: SSD: Single shot multibox detector. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) *Computer Vision – ECCV 2016*. ECCV 2016. Lecture Notes in Computer Science, vol. 9905. Springer, Cham (2016).
- [3] Zhu, X., Su, W., Lu, L., et al., 2020. Deformable detr: Deformable transformers for end-to-end object detection. *arxiv preprint arxiv:2010.04159*.
- [4] Takaki, M.; Fujiyoshi, H. Traffic sign recognition using SIFT features. *IEEJ Trans. Electron. Inf. Syst.* 2009, 129, 824–831.
- [5] Ruta, A.; Li, Y.; Liu, X. Real-time traffic sign recognition from video by class-specific discriminative features. *Pattern Recognit.* 2010, 43, 416–430.
- [6] Girshick, R. Fast r-cnn. *arXiv* 2015, arXiv:1504.08083.
- [7] Ren, S.; He, K.; Girshick, R.; Sun, J. Faster R-CNN: Towards real-time object detection with region proposal networks. *IEEE Trans. Pattern Anal. Mach. Intell.* 2016, 39, 1137–1149.
- [8] He, K.; Gkioxari, G.; Dollár, P.; Girshick, R. Mask r-cnn. In *Proceedings of the IEEE International Conference on Computer Vision, Venice, Italy, 22–29 October 2017*; pp. 2961–2969.
- [9] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 779–788, 2016.
- [10] Joseph Redmon and Ali Farhadi. YOLO9000: better, faster, stronger. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7263–7271, 2017.
- [11] Joseph Redmon and Ali Farhadi. YOLOv3: An incremental improvement. *arXiv preprint arXiv:1804.02767*, 2018.
- [12] Alexey Bochkovskiy, Chien-Yao Wang, and Hong-Yuan Mark Liao. YOLOv4: Optimal speed and accuracy of object detection. *arXiv preprint arXiv:2004.10934*, 2020.
- [13] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg, “SSD: Single Shot MultiBox Detector,” in *Proc. European Conf. Computer Vision (ECCV)*, 2016, pp. 21–37.
- [14] C.-Y. Fu, W. Liu, A. Ranga, A. Tyagi, and A. C. Berg, “DSSD: Deconvolutional Single Shot Detector,” *arXiv:1701.06659*, 2017.
- [15] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, “Focal Loss for Dense Object Detection,” in *Proc. IEEE Int. Conf. Computer Vision (ICCV)*, 2017, pp. 2980–2988.
- [16] H. Law and J. Deng, “CornerNet: Detecting objects as paired keypoints,” in *Proc. European Conf. Computer Vision (ECCV)*, 2018, pp. 765–781.
- [17] K. W. Duan, et al., “CenterNet: Keypoint triplets for object detection,” in *Proc. IEEE Int. Conf. Computer Vision (ICCV)*, 2019, pp. 6568–6577.
- [18] Z. Tian, C. Shen, H. Chen, and T. He, “FCOS: Fully Convolutional One-Stage Object Detection,” in *Proc. IEEE Int. Conf. Computer Vision (ICCV)*, 2019, pp. 9626–9635.
- [19] C. Zhu, Y. He, and M. Savvides, “Feature Selective Anchor-Free Module for Single-Shot Object Detection,” *arXiv:1903.00621*, 2019.
- [20] Jocher Glenn, YOLOv8, 2023, <https://github.com/ultralytics/tree/main>
- [21] R. Sapkota, R. H. Cheppally, A. Sharda, and M. Karkee, “YOLOv26: Key Architectural Enhancements and Performance Benchmarking for Real-Time Object Detection,” *arXiv preprint arXiv:2509.25164*, 2025.
- [22] Yian Zhao, Wenyu Lv, Shangliang Xu, Jinman Wei, Guanzhong Wang, Qingqing Dang, Yi Liu, Jie Chen, DETRs beat YOLOs on real-time object detection, 2024, *arXiv:2304.08069*.
- [23] Y. Niu et al., “TS-DETR: A Small Object Detection Model in Autonomous Driving Systems,” 2024 IEEE International Conference on Systems, Man, and Cybernetics (SMC), Kuching, Malaysia, 2024, pp. 1558–1563, doi: 10.1109/SMC54092.2024.10831658.
- [24] L. Zhang et al., “TSD-DETR: A lightweight real-time detection transformer of traffic sign detection for long-range perception of autonomous driving,” *Engineering Applications of Artificial Intelligence*, vol. 139, p. 109536, Jan. 2025.
- [25] N. Carion, L. Gustafson, Y.-T. Hu, S. Debnath, R. Hu, D. Suris, “SAM 3: Segment Anything with Concepts,” *arXiv preprint arXiv:2511.16719*, 2025.
- [26] Panoramax, “Classified French Road Signs Dataset,” Hugging Face. [Online]. Available: https://huggingface.co/datasets/Panoramax/classified_fr_road_signs.
- [27] A. Caunes, “French Traffic Sign Classification (FTSC),” GitHub repository. <https://github.com/andrewcaunes/FTSC>.