

FMARS: Annotating Remote Sensing Images for Disaster Management using Foundation Models

*Original*

FMARS: Annotating Remote Sensing Images for Disaster Management using Foundation Models / Arnaudo, E., Lungo Vaschetti, J., Innocenti, L., Barco, L., Lisi, D., Fissore, V., Rossi, C.. - (2024), pp. 3920-3924. (IEEE International Symposium on Geoscience and Remote Sensing (IGARSS) 2024 Athens (GR) 7-12 July 2024) [10.1109/IGARSS53475.2024.10641130].

*Availability:*

This version is available at: 11583/2989784 since: 2024-06-21T16:00:36Z

*Publisher:*

IEEE

*Published*

DOI:10.1109/IGARSS53475.2024.10641130

*Terms of use:*

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

*Publisher copyright*

IEEE postprint/Author's Accepted Manuscript

©2024 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

# FMARS: ANNOTATING REMOTE SENSING IMAGES FOR DISASTER MANAGEMENT USING FOUNDATION MODELS

*E. Arnaudo*<sup>1,2</sup>, *J. L. Vaschetti*<sup>1,2</sup>, *L. Innocenti*<sup>1</sup>, *L. Barco*<sup>1,2</sup>, *D. Lisi*<sup>3</sup>, *V. Fissore*<sup>3</sup>, *C. Rossi*<sup>1</sup>

1. LINKS Foundation, *AI, Data & Space (ADS)*, Torino (TO), Italy

2. Politecnico di Torino, *Dipartimento di Automatica e Informatica (DAUIN)*, Torino (TO), Italy

3. Ithaca s.r.l., Torino (TO), Italy

## ABSTRACT


Very-High Resolution (VHR) remote sensing imagery is increasingly accessible, but often lacks annotations for effective machine learning applications. Recent foundation models like GroundingDINO [1] and Segment Anything (SAM) [2] provide opportunities to automatically generate annotations. This study introduces FMARS (Foundation Model Annotations in Remote Sensing), a methodology leveraging VHR imagery and foundation models for fast and robust annotation. We focus on disaster management and provide a large-scale dataset with labels obtained from pre-event imagery over 19 disaster events, derived from the Maxar Open Data initiative. We train segmentation models on the generated labels, using Unsupervised Domain Adaptation (UDA) techniques to increase transferability to real-world scenarios. Our results demonstrate the effectiveness of leveraging foundation models to automatically annotate remote sensing data at scale, enabling robust downstream models for critical applications. Code and dataset are available at <https://github.com/links-ads/igarss-fmars>.

**Index Terms**— Remote sensing, computer vision, machine learning, semantic segmentation.

## 1. INTRODUCTION

Remote Sensing (RS), and especially Very-High Resolution (VHR) images, represent a crucial resource for many real-world scenarios, including land use and land cover monitoring, urban planning, and disaster management. Recent advances in satellite technologies have allowed for increasingly accessible remote sensing data, also thanks to public and private programs such as the European Union’s Copernicus program [3] or the Maxar Open Data program [4], which helps to democratize the access to medium and high-resolution data for research purposes on a global scale. However, using such data with supervised machine learning might be challenging due to the limited availability of high-quality annotations.

This work was carried out in the context of the H2020 project SAFERS (GA n.869353), HEU project OVERWATCH (GA n.101082320) and Project NODES through the MUR—M4C2 1.5 of PNRR under Grant ECS00000036.



**Fig. 1:** Annotation workflow adopted for the three selected classes. Each class is treated separately, with its own prompt construction pipeline, while the segmentation masks are extracted from the same image embeddings, and merged together in a single output.

In parallel, the machine learning landscape is witnessing the emergence of foundation models, which are large, general-purpose models that can adapt to different downstream tasks with minimal to no effort. This also involves computer vision, with models such as CLIP [5], GroundingDINO [1], and Segment Anything [2], which are able to provide robust image classification, object detection and segmentation capabilities in a wide range of contexts without further finetuning. Nevertheless, despite their remarkable performance, these solutions have often been employed in the context of natural images, with only a few attempts at applying them extensively on RS data at scale [6].

In this work, we aim to bridge this gap between the growing availability of VHR remote sensing images and the potential of Vision Foundation Models (VFM) as robust annotators by designing an automated pipeline that combines open data sources with foundation models to generate either instance or semantic segmentation labels, starting from robust box prompts. Using this pipeline, named FMARS (i.e., Foundation Model Annotations in Remote Sensing), we automate the construction of a dataset designed for semantic labelling in damage assessment and disaster scenarios. The FMARS dataset includes 19 crisis events covered by Maxar Open Data imagery [4], and comprises more than 25M annotations over a surface of over 125.000  $km^2$ . On this data, we provide

instance-wise annotations subdivided into three example categories: buildings, roads, and high vegetation. To validate the effectiveness of the annotation approach, we train state-of-the-art models on the generated labels, employing Unsupervised Domain Adaptation (UDA) techniques [7, 8] for improved stability. Our results demonstrate the effectiveness of leveraging VFMs to automatically annotate remote sensing data at scale, enabling the development of smaller specific models for downstream applications.

## 2. RELATED WORK

Despite the recent advances in the computer vision field and the large data availability, remote sensing datasets remain limited in scope and scale compared to their natural images counterparts [9], as shown in Table 1. For instance, in downstream applications such as disaster management, the xBD dataset [10] focuses solely on building damage assessment, limiting its reuse in other contexts. On the other hand, general-purpose datasets such as DOTA [11] may not be easy to adapt to particular downstream tasks, due to the limitations of the available annotations. Considering models, Vision Foundation Models (VFM) have been successfully applied in several contexts, especially considering natural images [12, 2, 1], and downstream tasks in medical imagery [13, 14]. Previous works have already assessed the applicability of VFMs to remote sensing images, including SAM [15], applied to several semantic and instance segmentation tasks. Other attempts assessed the feasibility of using foundation models such as GroundingDINO [1] for annotation purposes [16], or the combination of SAM with text prompts, encoded via CLIP [5], to automatically generate segmentation masks for specific outputs [17]. However, to the best of our knowledge, despite the large availability of remote sensing imagery, only a few attempts have been made to provide automated annotations at scale. SAMRS is the prime example [6], providing an extended set of annotations over well-known datasets such as DOTA [11] and DIOR [18].

## 3. MATERIALS AND METHODS

### 3.1. Foundation Models

We adopt a combination of two large vision models for the annotation process, namely Segment Anything (SAM) [2], in its resource efficient variant [19], and GroundingDINO [1]. At its core, SAM and its derivatives are standard transformer-based segmentation networks that have been trained using *promptable segmentation*. In contrast with other segmentation objectives, this task receives two inputs: an image and a prompt. While the former is processed using a large and robust image encoder, the latter is embedded into the decoder using a prompt encoder, and exploited as query by a lightweight mask decoder that produces segmentation masks.

To resolve ambiguities, SAM can predict multiple outputs with its associated confidence for the same inputs. The prompts can be extremely flexible, ranging from sparse inputs such as a single point, a bounding box, or text, to dense arrays such as a binary mask. While points and boxes are encoded as simple positional embeddings, text is processed using off-the-shelf models such as CLIP [5], and masks are combined with the encoded image using a series of convolutions and element-wise sums. Following previous works [14], we adopt box prompts in our mask generation process, given its robustness and flexibility. This also combines naturally with the inputs at our disposal, comprising open data sources for buildings and roads (see Section 3.2), and box object detections derived from GroundingDINO. This model introduces cross-modal fusion between a text prompt and an image to provide open-set object detection capabilities, using BERT as text processor [20] and a Swin Transformer [21] as image encoder. While outputs may be approximate compared to human annotations, GroundingDINO provides a huge flexibility to generate bounding boxes for potentially any known object, given a text prompt. This allows us to obtain first estimates for objects not having a ground truth, such as vegetation, and thus exploit GroundingDINO as a prompt generator for the subsequent SAM masking phase [17].

### 3.2. Data Sources

Considering fine-grained segmentation in disaster management contexts, VHR imagery becomes necessary since lower-resolution satellite sources such as Copernicus Sentinel-2 do not provide enough image content to characterize objects of interest, such as buildings, or roads. To this date, the largest source of disaster-related VHR imagery is represented by the Maxar Open Data Program [4]. This initiative provides pre- and post-event RGB images from more than 100 major crisis events since 2017 worldwide, with a total surface coverage of more than  $2.6M km^2$ . We select a subset of resources containing RGB imagery, obtaining 19 events, spanning from 2022 to 2023, as displayed in Table 2, and summing up to an area of  $127,134 km^2$ . Inspired by current state-of-the-art disaster management datasets [10], we focus our dataset construction process on infrastructures, namely buildings and roads, which are often the focus in post-event damage assessment, and high vegetation, which usually occludes the underlying surface. Among open resources providing infrastructure information, we select the Microsoft’s Building Footprints and Road Detection datasets <sup>1</sup>, which contain building footprints polygons and road graphs on a global scale generated by applying deep learning models on VHR satellite imagery, respectively. For buildings, we do not directly adopt them as ground truth labels, but rather we exploit them as trustworthy yet approximate prompts for the SAM model. Lacking a reliable source to derive high vegetation prompt from, for such

<sup>1</sup><https://github.com/microsoft/GlobalMLBuildingFootprints>

Dataset	# Images	Image size	Resolution (cm)	Bands	# Instances	# Categories	Area ( $km^2$ )
Vaihingen	33	$2,500 \times 2,500$	9	IRRG	None	6	1.33
Potsdam	38	$6,000 \times 6,000$	5	RGBIR	None	6	11.08
iSAID	2,806	$4000 \times 4000$	$\geq 50$	RGB	655,451	15	11,224
xBD	9,168	$1024 \times 1024$	$\geq 50$	RGB	$> 700,000$	4	45,000
SAMRS	105,090	Mixed	$\geq 50$	RGB	$> 1.6M$	Mixed	Unknown
FMARS	6,896	$17,408 \times 17,408$	$\geq 30$	RGB	$> 25M$	3	$> 125,000$

**Table 1:** Brief comparison between FMARS and similar VHR datasets available in literature.

Event name	Year	Area ( $km^2$ )	Event name	Year	Area ( $km^2$ )
Cyclone Mocha	2023	3,446.4	Morocco earthquake	2023	49,901.9
Italy (Emilia) flooding	2023	1,519.1	Canada (NWT) wildfires	2023	468.6
Gambia flooding	2022	391.2	Sudan flooding	2022	249.3
Hurricane Fiona	2022	1,341.8	Afghanistan earthquake	2022	4,180.6
Hurricane Ian	2022	30,743.2	Cyclone Emnati	2022	8,506.0
Hurricane Idalia	2023	12,156.4	Kentucky flooding	2022	1,641.6
India floods	2023	496.3	Pakistan flooding	2022	7,528.7
Indonesia earthquake	2022	1,011.3	Georgia landslide	2023	157.4
Turkey earthquake	2023	2,745.7	South Africa flooding	2022	559.7
Kalehe flooding	2022	89.9			

**Table 2:** List of events included in the FMARS dataset, including its year and total surface coverage derived from VHR imagery.

category we adopt GroundingDINO as our bounding box generator.

### 3.3. Annotation Workflow

We aim to generate segmentation labels for three classes: buildings, roads, and high vegetation. While disaster risk management mainly focuses on damage assessment by comparing pre- and post-event images, we first concentrate our efforts on delineating infrastructures on pre-event acquisitions only. In fact, damage assessment frameworks typically delineate relevant entities in pre-event images, using the post-event image to determine the sustained damage [10, 22]. We argue that the first phase (i.e., identifying the exposed elements before the event) is crucial for any subsequent analysis, while the damage assessment could be carried out considering the output of the first phase and the post-event image using an ad-hoc model. Considering buildings, we generate box prompts by simply extracting axis-aligned bounding boxes (AABB) from each footprint polygon. On the other hand, road graphs represent a challenge for prompt-based segmentation because their sparse lattice does not allow for fine contour generation. In this case, the point-based prompts did not yield satisfactory results, therefore we opted to simply rasterize the available vector lines with a predefined buffer radius of 5m. For vegetation, we derive boxes using GroundingDINO with simple text queries like *green trees* or *bushes*, observing better performance on trees with the latter, likely due to the aerial viewpoint occluding the tree trunk, which is uncommon in natural images. In order to ensure a certain degree of confidence for the generated outputs, we use a minimum box threshold of 0.12 and a text threshold of 0.3. We further filter out noisy outputs by applying non-maxima suppression (NMS) at 0.5, removing boxes with aspect ratio

lower than 1:2, and maximum area over  $7000 m^2$ . Similar to buildings, we then use the generated boxes as prompts for SAM to extract segmentation labels. Last, we store the resulting delineation and its class as a single vector polygon to allow for both instance or semantic segmentation tasks.

### 3.4. Experiments


Method	Background		Roads		High Veg.		Buildings		mAcc.	mIoU
	Acc.	IoU	Acc.	IoU	Acc.	IoU	Acc.	IoU		
SegFormer (base)	72.91	61.41	0.11	0.10	7.60	1.33	0.00	0.00	20.15	15.71
MIC	44.79	42.47	55.94	29.89	64.45	10.56	82.47	21.33	61.91	26.06
DAFormer	53.06	50.14	55.44	31.79	64.61	16.80	79.91	17.29	<b>63.26</b>	<b>33.07</b>

**Table 3:** Performance comparison of DAFormer and MIC on the FMARS dataset across different classes.

Method	Background		Roads		High Veg.		Buildings		mAcc.	mIoU
	Acc.	IoU	Acc.	IoU	Acc.	IoU	Acc.	IoU		
FMARS labels	71.34	41.16	68.72	47.03	69.37	58.54	59.47	54.14	67.23	50.22
SegFormer (base)	97.40	27.90	0.06	0.06	8.24	7.68	0.00	0.00	26.44	8.91
MIC	76.59	36.21	44.84	40.15	51.78	48.52	63.54	56.41	59.19	45.32
DAFormer	70.56	38.02	65.97	54.77	56.57	52.64	69.10	60.20	<b>65.55</b>	<b>51.41</b>

**Table 4:** Performance comparison of DAFormer and MIC on a manually labelled Maxar partition across different classes.

Using the dataset annotated with our FMARS approach, we can train standard semantic segmentation models to evaluate the knowledge transfer ability to smaller and more deployable models. In our experiments, we adopt state-of-the-art solutions based on Segformer [23]. To counteract the inherent inaccuracies in fully automated labeling and the consequently lower recall for categories such as high vegetation, we apply UDA techniques for improved stability during training. Specifically, we adopt DAFormer ([7] and Masked Image Consistency (MIC) [8], both based on self-training in a *teacher-student* framework paradigm. We select a separate full-size image from each event as our test set based on the average information content, for a total of 19 images, and we conduct a full training using pretrained ImageNet weights for the backbone components. To address the high precision and lower recall of the generated labels, as well as missing categories, we train the models in an open set context, ignoring the background class [24]. As simple baseline, we apply a confidence threshold to the Softmax outputs, empirically evaluating the optimal cutoff threshold  $\tau = 0.9$  for both



**Fig. 2:** Qualitative results obtained over two example areas, namely USA (top) and Gambia (bottom). from left to right: RGB image, DAFormer, MIC, and FMARS ground truth. Best viewed zoomed in.

models. We conduct every experiment using a tile size of  $512 \times 512$  and random sampling, weighted by the entropy (i.e., information content) of the available label, for 30,000 iterations using AdamW as optimizer. For the UDA components, we maintain their original configuration, except for the removal of the ImageNet feature distance.

Given the automated pipeline and the low reliability of the obtained labels for performance measurement, we validate results against the FMARS labels, as well as a small sample of 45 manually labelled tiles, derived from crops of each image in the test set. Tables 3 and 4 present the numerical results in terms of accuracy and Intersection over Union (IoU), evaluated against the FMARS test set and manual labels, respectively. The scores highlight the challenge of the problem at hand, where the baseline solution, without any precautions, collapses during training. Both UDA solutions appear very effective, with DAFormer even surpassing the original FMARS labels on the manually produced ground truth. This is evident in the qualitative results shown in Fig. 2 that demonstrate a high fidelity between the model predictions and the ground truth, with DAFormer exhibiting more robustness to the challenging *high vegetation* class. These findings highlight the effectiveness of FMARS as an automated technique to leverage foundation models for annotation tasks in remote sensing. With the necessary precautionary measures such as UDA techniques, the proposed pipeline allows the generation of accurate labels on a large scale, and even the knowledge transfer to smaller yet accurate downstream segmentation models in the absence of manually labeled datasets.

#### 4. CONCLUSIONS

In this work, we propose FMARS, a pipeline for automated large-scale annotation of VHR imagery leveraging VFM. FMARS exploits the increasing availability of open-source images and the flexibility of promptable large models to automatically generate high-quality instance and semantic segmentation labels. As an example, we focus on the critical domain of disaster management and construct the FMARS dataset, providing over 25 million annotations across 19 disaster events around the globe, derived from pre-event imagery from the Maxar Open Data initiative. As representative downstream application, we train state-of-the-art segmentation frameworks on the generated labels. Due to the potential inaccuracies in the automated labels, we employ UDA techniques to increase the robustness of the learned features to real-world scenarios. While the proposed approach demonstrates promising results, it has certain limitations. First, the dataset currently focuses only on pre-event images. Second, we limited the taxonomy to three classes in these tests. Finally, achieving high recall requires significant computational effort, especially in downstream training. Future work may address these limitations by developing a more robust automated annotation pipeline, or by exploring zero-shot learning approaches for open-set or panoptic segmentation tasks, with a potentially boundless taxonomy.

## 5. REFERENCES

- [1] Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Chunyuan Li, Jianwei Yang, Hang Su, Jun Zhu, et al., “Grounding dino: Marrying dino with grounded pre-training for open-set object detection,” *arXiv preprint arXiv:2303.05499*, 2023.
- [2] 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,” *arXiv preprint arXiv:2304.02643*, 2023.
- [3] European Union, “Copernicus programme,” 2014.
- [4] Maxar, “Maxar open data program,” 2017.
- [5] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al., “Learning transferable visual models from natural language supervision,” in *International conference on machine learning*. PMLR, 2021, pp. 8748–8763.
- [6] Di Wang, Jing Zhang, Bo Du, Minqiang Xu, Lin Liu, Dacheng Tao, and Liangpei Zhang, “Samrs: Scaling-up remote sensing segmentation dataset with segment anything model,” in *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2023.
- [7] Lukas Hoyer, Dengxin Dai, and Luc Van Gool, “Daformer: Improving network architectures and training strategies for domain-adaptive semantic segmentation,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 9924–9935.
- [8] Lukas Hoyer, Dengxin Dai, Haoran Wang, and Luc Van Gool, “Mic: Masked image consistency for context-enhanced domain adaptation,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2023, pp. 11721–11732.
- [9] Pedram Ghamisi and Naoto Yokoya, “Img2dsm: Height simulation from single imagery using conditional generative adversarial net,” *IEEE Geoscience and Remote Sensing Letters*, vol. 15, no. 5, pp. 794–798, 2018.
- [10] Ritwik Gupta, Bryce Goodman, Nirav Patel, Ricky Hosfelt, Sandra Sajeev, Eric Heim, Jigar Doshi, Keane Lucas, Howie Choset, and Matthew Gaston, “Creating xbd: A dataset for assessing building damage from satellite imagery,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, 2019, pp. 10–17.
- [11] Gui-Song Xia, Xiang Bai, Jian Ding, Zhen Zhu, Serge Belongie, Jiebo Luo, Mihai Datcu, Marcello Pelillo, and Liangpei Zhang, “Dota: A large-scale dataset for object detection in aerial images,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 3974–3983.
- [12] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin, “Emerging properties in self-supervised vision transformers,” in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 9650–9660.
- [13] Can Cui, Ruining Deng, Quan Liu, Tianyuan Yao, Shunxing Bao, Lucas W Remedios, Yucheng Tang, and Yuankai Huo, “All-in-sam: from weak annotation to pixel-wise nuclei segmentation with prompt-based finetuning,” *arXiv preprint arXiv:2307.00290*, 2023.
- [14] Maciej A Mazurowski, Haoyu Dong, Hanxue Gu, Jichen Yang, Nicholas Konz, and Yixin Zhang, “Segment anything model for medical image analysis: an experimental study,” *Medical Image Analysis*, vol. 89, pp. 102918, 2023.
- [15] Simiao Ren, Francesco Luzi, Saad Lahrichi, Kaleb Kassaw, Leslie M Collins, Kyle Bradbury, and Jordan M Malof, “Segment anything, from space?,” in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2024, pp. 8355–8365.
- [16] Jieliu Zhang, Zhongliang Zhou, Gengchen Mai, Lan Mu, Mengxuan Hu, and Sheng Li, “Text2seg: Remote sensing image semantic segmentation via text-guided visual foundation models,” *arXiv preprint arXiv:2304.10597*, 2023.
- [17] Lucas Prado Osco, Qiusheng Wu, Eduardo Lopes de Lemos, Wesley Nunes Gonçalves, Ana Paula Marques Ramos, Jonathan Li, and José Marcato Junior, “The segment anything model (sam) for remote sensing applications: From zero to one shot,” *International Journal of Applied Earth Observation and Geoinformation*, vol. 124, pp. 103540, 2023.
- [18] Yang Zhan, Zhitong Xiong, and Yuan Yuan, “Rsvg: Exploring data and models for visual grounding on remote sensing data,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 61, pp. 1–13, 2023.
- [19] Yunyang Xiong, Bala Varadarajan, Lemeng Wu, Xiaoyu Xiang, Fanyi Xiao, Chenchen Zhu, Xiaoliang Dai, Dilin Wang, Fei Sun, Forrest Iandola, Raghuraman Krishnamoorthi, and Vikas Chandra, “Efficientsam: Leveraged masked image pretraining for efficient segment anything,” 2023.
- [20] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova, “Bert: Pre-training of deep bidirectional transformers for language understanding,” 2019.
- [21] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo, “Swin transformer: Hierarchical vision transformer using shifted windows,” 2021.
- [22] Yu Shen, Sijie Zhu, Taojiannan Yang, Chen Chen, Delu Pan, Jianyu Chen, Liang Xiao, and Qian Du, “Bdanet: Multiscale convolutional neural network with cross-directional attention for building damage assessment from satellite images,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 60, pp. 1–14, 2021.
- [23] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo, “Segformer: Simple and efficient design for semantic segmentation with transformers,” *Advances in Neural Information Processing Systems*, vol. 34, pp. 12077–12090, 2021.
- [24] Hugo Oliveira, Caio Silva, Gabriel LS Machado, Keiller Nogueira, and Jefersson A Dos Santos, “Fully convolutional open set segmentation,” *Machine Learning*, pp. 1–52, 2023.