

Locate Anything on Earth: Advancing Open-Vocabulary Object Detection for Remote Sensing Community

Jiancheng Pan^{1,2,*†}, Yanxing Liu^{3*}, Yuqian Fu^{4,5‡}, Muyuan Ma¹,
Jiahao Li¹, Danda Pani Paudel⁵, Luc Van Gool^{4,5}, Xiaomeng Huang^{1‡}

¹Tsinghua University

²Zhejiang University of Technology

³University of Chinese Academy of Sciences

⁴ETH Zürich

⁵INSAIT, Sofia University

jiancheng.pan.plus@gmail.com, liuyanxing21@mailsucas.ac.cn, yuqian.fu@insait.ai, hxm@tsinghua.edu.cn

Abstract

Object detection, particularly open-vocabulary object detection, plays a crucial role in Earth sciences, such as environmental monitoring, natural disaster assessment, and land-use planning. However, existing open-vocabulary detectors, primarily trained on natural-world images, struggle to generalize to remote sensing images due to a significant domain gap. Thus, this paper aims to advance the development of open-vocabulary object detection in remote sensing community. To achieve this, we first reformulate the task as **Locate Anything on Earth (LAE)** with the goal of detecting any novel concepts on Earth. We then developed the **LAE-Label Engine** which collects, auto-annotates, and unifies up to 10 remote sensing datasets creating the **LAE-1M** — the first large-scale remote sensing object detection dataset with broad category coverage. Using the LAE-1M, we further propose and train the novel **LAE-DINO Model**, the first open-vocabulary foundation object detector for the LAE task, featuring *Dynamic Vocabulary Construction (DVC)* and *Visual-Guided Text Prompt Learning (VisGT)* modules. DVC dynamically constructs vocabulary for each training batch, while VisGT maps visual features to semantic space, enhancing text features. We comprehensively conduct experiments on established remote sensing benchmark DIOR, DOTAv2.0, as well as our newly introduced 80-class LAE-80C benchmark. Results demonstrate the advantages of the LAE-1M dataset and the effectiveness of the LAE-DINO method.

Code — <https://github.com/jaychempan/LAE-DINO>

Introduction

As one of the most fundamental and important tasks in the field of computer vision, object detection (OD) and localization (Ren et al. 2015) has been extensively studied over the years, leading to the development of numerous detectors. In particular, open-vocabulary object detec-

tion (OVD) (Zareian et al. 2021) has been receiving increasing attention. OVD relaxes the limitation of close-set object categories in the traditional OD, allowing the detection of any novel concept during the testing time. Among various OVD methods, DINO (Zhang et al. 2023) based detectors, e.g., GroundingDINO (Liu et al. 2024), have recently shown promising performance on mainstream OVD benchmarks.

However, almost all of the state-of-the-art OVD methods are trained and tested on natural-world images. When applied to Earth science-related tasks, such as environmental monitoring, natural disaster assessment, land-use planning, these methods struggle to generalize due to the huge data domain gap. Unlike natural-world imagery, Earth science relies on remote sensing imagery, which exhibit much higher resolutions, distinct image styles (Ma, Pan, and Bai 2024; Pan et al. 2023b), and different semantic class concepts. This makes the direct transfer of current OVD models nontrivial. Therefore, in this paper, we are motivated to *advance open-vocabulary object detection for remote sensing community*.

To achieve this goal, we first reformulate the task of OVD for remote sensing filed as **Locate Anything on Earth (LAE)**¹. As illustrated in Figure 1, our aim is to enable LAE models could detect any novel concept on Earth. Our efforts are mainly made from two key aspects: first, a **LAE-Label Engine** is developed to construct the large-scale remote sensing training data; second, a novel **LAE-DINO Model** is proposed and trained to work as the first foundation models for the newly proposed LAE task.

More specifically, the LAE-Label engine is proposed to solve the lack of diverse object-level labeled data in the remote sensing community, which is essentially an indispensable part of training robust foundation models. To fully leverage the existing scattered remote sensing data which can be broadly grouped into labeled and unlabeled data, our LAE-Label engine proposes two distinct solutions. For labeled datasets, we focus on unifying them through image slicing, format alignment, and sampling, forming the fine-

*These authors contributed equally.

†Work is done during an internship at Tsinghua University.

‡Corresponding author.

¹Following (Zhang and Wang 2024), we use “localization” to describe detection tasks in the remote sensing domain.




Figure 1: Locate Anything on Earth (LAE) aims to detect any object on Earth and facilitate practical detection tasks, powered by LAE-Label Engine and LAE-DINO Model.

grained LAE-FOD dataset. For unlabeled datasets, we develop a semi-automated labeling pipeline using SAM (Kirillov et al. 2023), a large vision-language model, and rule-based filtering, resulting in the coarse-grained LAE-COD dataset. By combining LAE-FOD and LAE-COD, we ultimately construct the **LAE-1M** dataset with one million labeled objects across diverse categories. To our knowledge, LAE-1M is the first and largest remote sensing object detection dataset with broadest category coverage to date.

Technically, the LAE-DINO, a DINO-based OVD method, is proposed and trained on the LAE-1M dataset. The novel modules of LAE-DINO are designed to address two questions: 1) How to fit the OVD model in the training data that has around 1600 vocabularies? 2) How can the relationship between image and text be better utilized to achieve more effective vocabulary-conditioned object detection? As the answer of the first question, the Dynamic Vocabulary Construction (DVC) which dynamically selects the positive and negative vocabularies for each training batch is proposed. While the Visual-Guided Text Prompt Learning (VisGT) is presented to address the second issue. Based on the observation that different objects within a single image collectively define the scene, VisGT introduces the concept of “scene features” by averaging all object features. Through taking the scene features as a bridge, VisGT aligns visual features with text features, thereby enhancing the interaction between these two modalities. Extensive experiments are conducted on both open-set and close-set scenarios. Different models are compared taking different data as training data. Results reveal: 1) our proposed LAE-1M dataset significantly improves model performance, especially in open-set scenarios; and 2) our LAE-DINO model achieves state-of-the-art performance.

We summarize the main contributions as follows,

- We advocate the Locate Anything on Earth (LAE) task for remote sensing and pave the way for LAE by contributing the LAE-1M data with one-million instances.
- We propose a novel LAE-DINO detector for LAE, with dynamic vocabulary construction (DVC) and Visual-Guided Text Prompt Learning (VisGT) as novel modules.

- Extensive experimental results on several different testing benchmarks demonstrate the advantages of the LAE-1M dataset and the effectiveness of the LAE-DINO.

Related Work

Generic Object Detection for Remote Sensing. Object detection (OD) is one of the classical vision tasks in computer vision, which is to obtain the locations of regions of interest from a given image. Object detection methods can be mainly divided into single-stage and two-stage object detection. Single-stage methods (e.g. YOLO Family (Redmon and Farhadi 2017)) perform classification and regression directly on a predefined mass of anchor boxes. The two-stage methods (e.g. Faster R-CNN (Ren et al. 2015)) fine-tune bounding boxes based on the single-stage, which are usually more accurate than single-stage methods but are slower. Some representative works, such as DINO-based detector (Zhang et al. 2023) based on Transformer, explore the trade-off between performance and computational cost for more robust detectors. In the remote sensing community, extensive research efforts (Tao et al. 2023; Cong et al. 2022; Reed et al. 2023; Bastani et al. 2023; Sun et al. 2022a; Guo et al. 2024) are concentrated on the extraction of fundamental imagery knowledge from large volumes of unlabeled data, utilizing advanced self-supervised or unsupervised methodologies. Some methods (e.g. CALNet (He et al. 2023)) work on involving visible (RGB) and infrared (IR) images to enhance detection performance. While these methods are broadly applicable, they exhibit limited effectiveness in enhancing detection capabilities.

Object Detection from Few-Shot Learning to Open-Vocabulary Learning. Few-Shot Object Detection (FSOD) (Chen et al. 2018) aims to detect unseen objects using only a few labeled examples. FSOD approaches are divided into fine-tuning-based (Chen et al. 2018), which transfers knowledge (Hospedales et al. 2021) from base to novel classes, and meta-learning-based, which uses “learning to learn” to generalize across novel classes. CD-FSOD (Fu et al. 2024)) explores cross-domain FSOD (e.g., from natural to remote sensing images), yet it relies on visual

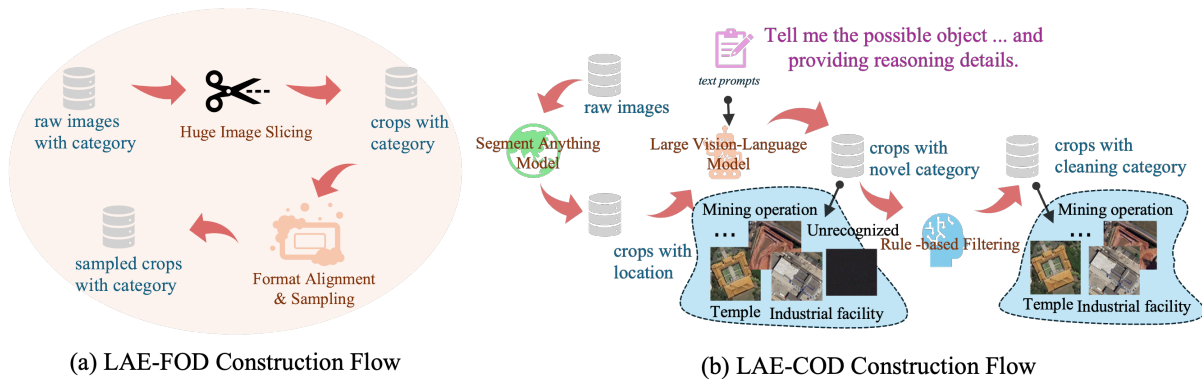


Figure 2: The pipeline of our LAE-Label Engine.

images, offering limited support for new vocabulary.

Therefore, Open-Vocabulary Object Detection (OVD) adopts a more practice-oriented learning paradigm (Zareian et al. 2021) compared with FSOD, aiming to construct an open visual-semantic space to enhance out-of-category identification and localisation. OVR-CNN (Zareian et al. 2021) first proposed to acquire knowledge from natural language vocabularies by pre-training the backbone with image-caption data. After that, RegionCLIP (Zhong et al. 2022) and GLIP (Li et al. 2022) became unified with the image-text matching task, expanding the visual-semantic space with more powerful flooding capabilities. While the previous work mainly improves zero-shot recognition with the help of vision-language pre-training, Grounding-DINO (Liu et al. 2024) obtained a more robust grounding capability by introducing a stronger detector structure and fine-grained multimodal feature fusion. CasDet (Li et al. 2024b) combines semi-supervised learning and OVD to augment aerial detection. Due to insufficient domain annotation data, these works are weaker in open-set detection, although some show promising results in closed-set detection.

Locate Anything on Earth Task

Task: Locate Anything on Earth. Locate Anything on Earth (LAE) draws inspiration from the Open-Vocabulary Object Detection (OVD) task but is specifically tailored for the remote sensing field. Given remote-sensing imagery as input, LAE aims to achieve robust object recognition and localization based on provided text prompts.

LAE maintains a base training dataset \mathcal{D}_{base} and any potential testing dataset \mathcal{D}_{test} . Formally, the base dataset is represented as $\mathcal{D}_{base} = \{I, \{(b, y)_r\}\}$, where I denotes a remote sensing image, and each image comprises r objects with corresponding localization annotations b and category annotations y . Specifically, I is defined as $I \in \mathbb{R}^{H \times W \times C}$, b as $b \in \mathbb{R}^4$, and y as an element of \mathcal{V}_{base} , where \mathcal{V}_{base} is the set of vocabularies present in \mathcal{D}_{base} . A large \mathcal{V}_{base} is generally preferable for training foundational LAE models effectively. Moreover, we define \mathcal{V}_{Ω} as the entire language vocabulary and \mathcal{V}_{test} as the testing vocabulary within \mathcal{D}_{test} . Consistent with the fundamental settings of OVD (Zareian

et al. 2021), no constraints are imposed on \mathcal{V}_{test} , indicating that it can be any subset of \mathcal{V}_{Ω} .

Overall, LAE necessitates that models learn from \mathcal{D}_{base} and subsequently identify the correct object localizations b and categories y for images in \mathcal{D}_{test} based on the provided text prompt \mathcal{T} .

Engine: LAE-Label Engine. As widely recognized, one of the essential requirements for training foundational models is the availability of large amounts of training data. Thus, naturally, this paper also aims to construct a dataset that could support the training of foundational LAE models. However, in the remote sensing community, the existing datasets show such limitations: 1) the human-labeled datasets are small-scale and have different sizes and data format; 2) the large-scale image-text pairs which could be easily obtained from the Internet lacks well annotations.

To tackle these two limitations, we propose the LAE-Label data engine which makes use of both the well-labeled data and the massive unlabeled data. More specifically, as shown in Figure 2(a), for those well-labeled datasets, we first slice the huge image of different datasets and then unify the format as the same. This part results in our fine-grained LAE-FOD dataset; For the unlabeled data, as in Figure 2(b), we build a comprehensive semi-automated data construction flow based on SAM (Kirillov et al. 2023) and Large Vision-Language Model (LVLM).

We begin by extracting the location information of Regions of Interest from remote sensing seed datasets using SAM. The detailed information of the seed datasets is listed in Table 1. Next, we obtain the categories of the zoomed-in ROI areas by taking advantage of the LVLM, i.e., InternVL (Chen et al. 2024) which has a powerful zero-shot recognition as learned on huge amounts of data with the text prompt as demonstrated in Figure 2(b). Finally, we filter out invalid and irrelevant categories using a rule-based method. In this way, our coarse-grained LAE-COD dataset is constructed, offering a rich vocabulary for open-vocabulary pre-training.

LAE-DINO Open-Vocabulary Detector

Overview. Due to the huge success of DINO (Zhang et al. 2023), the recent DINO-based detector e.g., GroundingDINO

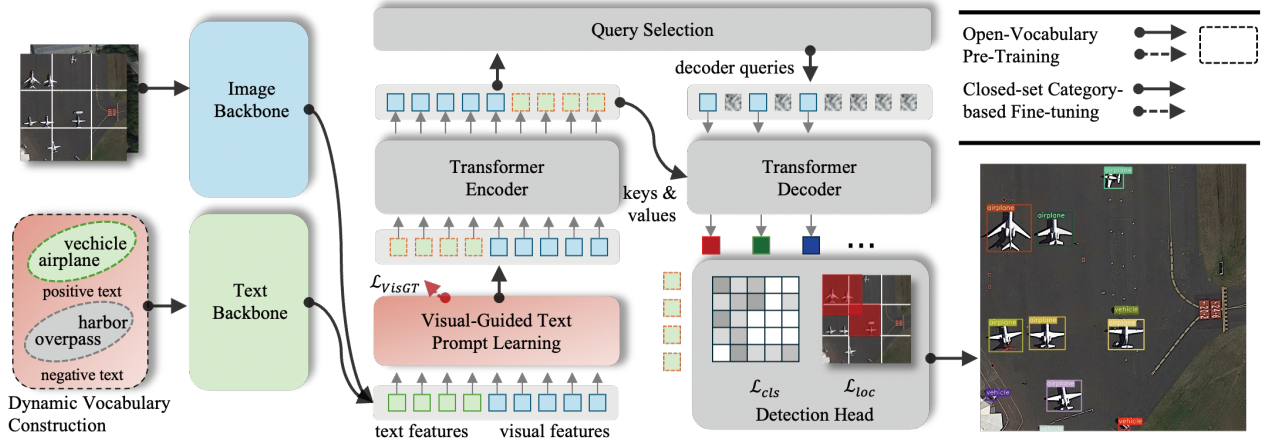


Figure 3: The pipeline for LAE-DINO.

(Liu et al. 2024) and VideoGrounding-DINO (Wasim et al. 2024), show promising detection performance on open-vocabulary detection scenarios. Thus, in this paper, we also build our method upon the DINO and form our novel **LAE-DINO** detector. As illustrated in Figure 3, except for the data engine part, our LAE-DINO mainly contains the Dynamic Vocabulary Construction (DVC), the Image Backbone E_{img} , the Text Backbone E_{text} , the Visual-Guided Text Prompt Learning (VisGT), the Transformer Encoder E_{T_E} , the Query Selection M_{qs} , the Transformer Decoder E_{T_D} , and the Detection Head M_{det} . Note that the E_{img} , E_{text} , E_{T_E} , M_{qs} , E_{T_D} , and M_{det} are basic and common modules in DINO-based detectors, thus we keep them same with the former GroundingDINO. While the DVC and the VisGT are newly proposed in this paper. Typically, the DVC is proposed to tackle the large vocabulary set issue posed by our constructed training data, and the VisGT is a novel method that uses the visual information to further guide and transform the text features.

In the following paragraphs, we will first introduce the basic pipeline of DINO-based Detector and then present our two novel modules.

DINO-based Detectors. Though developed in different directions and with different new modules, the DINO-based detectors basically share the same core pipeline: Given the training dataset \mathcal{D}_{base} , the first thing is to construct the vocabulary set \mathcal{V}_{base} by simply merging all the existing vocabularies. The vocabulary set includes positive vocabularies for categories in the images and negative words for those not seen during training.

For each batch training iteration, as indicated in Figure 3, the image backbone E_{img} and the text backbone E_{text} are used to extract the visual features $F_I \in \mathbb{R}^{n_I \times d}$, the text features $F_T \in \mathbb{R}^{n_T \times d}$ from the input image I and vocabulary set \mathcal{V}_{base} , respectively. The n_I and n_T mean the number of image and text tokens, while the d denotes the dimension of features. Usually, the Swin-Transformer (Liu et al. 2021) is used as the E_{img} and the BERT (Devlin et al. 2018) is used as the E_{text} . In addition, since the \mathcal{V}_{base} contains both

the positive and negative vocabularies, we further denote the text features generated from n_{T_p} positive vocabularies as $F_T^P = [\tilde{F}_{T_1}, \tilde{F}_{T_2}, \dots, \tilde{F}_{T_{n_{T_p}}}] \in \mathbb{R}^{n_{T_p} \times d}$.

After that, the Transformer encoder E_{T_E} which takes both the image features F_I and the text features F_T are applied to fuse the multi-modal features. Then, the query selection M_{qs} is used to initialize the region queries which consists of the learnable content queries and dynamic positional queries. Finally, the Transformer decoder E_{T_D} and the detection head M_{det} output the both the location and category predictions $\{(\hat{b}, \hat{y})_r\}$ for modality alignment.

Upon the predictions $\{(\hat{b}, \hat{y})_r\}$ and the ground truth $\{(b, y)_r\}$, two classical losses are calculated. One is the standard Cross Entropy (CE) loss \mathcal{L}_{cls} (Li et al. 2022; Liu et al. 2024) for evaluating the classification results between \hat{y} and y , another is the Generalized Intersection over Union (GIoU) loss \mathcal{L}_{loc} (Rezatofighi et al. 2019) for evaluating the locations. The detailed calculation method for \mathcal{L}_{cls} and \mathcal{L}_{loc} are as follows,

$$\mathcal{L}_{cls} = \sum_{i=1}^r \mathcal{L}_{CE}(\hat{y}_i, y_i), \quad (1)$$

$$\mathcal{L}_{loc} = \lambda_{L_1} \sum_{i=1}^r \mathcal{L}_{L_1}(\hat{b}_i, b_i) + \lambda_{GIoU} \sum_{i=1}^r \mathcal{L}_{GIoU}(\hat{b}_i, b_i). \quad (2)$$

Dynamic Vocabulary Construction. Current OVD detectors (Li et al. 2022; Liu et al. 2024) like Grounding DINO rely on fixed-token-length text encoders (e.g., BERT (Devlin et al. 2018) or CLIP (Radford et al. 2021)) that concatenate all categories into a “extremely long text”, which is non-trivial for datasets with numerous categories. For example, while BERT allows a maximum of 256 tokens, our dataset includes about 1600 categories, exceeding this limit. This motivated us to develop the dynamic vocabulary construction (DVC), which reduces the number of input categories.

To tackle with such a “extremely long text”, APE (Shen et al. 2024) tries to blend the individual concepts of vocabularies as independent text prompts but discarding the corre-

lation among vocabularies. Our DVC sets a dynamic vocabulary length $N_{\mathcal{D}\mathcal{V}}$, for each training iteration, several positive and negative vocabularies will be selected to form the $N_{\mathcal{D}\mathcal{V}}$ vocabulary set. Concretely, if the number of base vocabulary \mathcal{V}_{base} is larger than $N_{\mathcal{D}\mathcal{V}}$, i.e., $\|\mathcal{V}_{base}\| > N_{\mathcal{D}\mathcal{V}}$. In each training batch, DVC ensures the input category length is fixed at $N_{\mathcal{D}\mathcal{V}}$. All current batch categories are considered positive categories (for example N_{pos}) are included in the input. The remaining $(N_{\mathcal{D}\mathcal{V}} - N_{pos})$ part is filled by randomly sampling negative categories from the rest of the whole category set \mathcal{V}_{base} . DVC can effectively reduce the number of iterations for text encoder inference.

Visual-Guided Text Prompt Learning. OVD models primarily rely on the relationship between image and text to achieve the open-vocabulary learning. Current DINO-based detectors, including our LAE-DINO, utilize this relations through textual prompt learning.

However, a picture paints a thousand word which means that sparse and limited categories are hard to fully represent an image. Also inspired by MQ-Det (Xu et al. 2024), incorporating visual prompts from additional supported images with text prompts, we propose the VisGT module which aims at leveraging the visual information to further improve the semantic representation. Notably our VisGT does not utilise visual prompts like MQ-Det, but rather visual-guided text prompts to compensate for the lack of single text prompt.

Specifically, as in Figure 3, VisGT is not an object-level alignment but an image-level alignment that represents the overall objects of the scene, preserving the knowledge of vocabulary to fine-grained detection across different categories.

The detailed architecture of our VisGT is shown in Figure 4. First of all, we propose the ‘‘scene features’’ by fusing different text features. The observation behind this is that the different object categories together could convert some useful scene information. For example, *airplane* and *vehicle* are two typical concepts that are strongly related to the *airport* scene. Thus, given the textual features F_T with its positive textual features F_{T_p} , we define the scene feature s as,

$$s = \frac{1}{n_{T_p}} \sum_{i=1}^{T_p} L_i \tilde{F}_{T_i}, \quad (3)$$

where L_i is the token length of the i -th category, which corresponds to the T_i -th token.

By combining different instance-agnostic positive text features F_T^P , our scene features s could be regarded as some special feature that contains both the instance-level and category-relative features. This scene feature s works as the ground truth when we try to map the visual information also into the semantic space.

As for the mapping of visual feature to semantic feature \hat{s}^l , we introduce the Multi-scale Deformable Self-Attention (MDSA) (Zhu et al. 2021) as a tool as follows,

$$\hat{s}^l = \begin{cases} \sum_{i=1}^{n_I} \frac{F_{I_{j,:}}}{n_I}, & (l = 1), \\ FFN^l(MDSA^l(\hat{s}^{l-1})), & (l > 1), \end{cases} \quad (4)$$




Figure 4: VisGT maps visual features into semantic space. The scene features are instance-level and category-relative features from different textual features in an image, which represents the scenographic information from the image. For example, *airplane* and *vehicle* belong to the *airport*.

where $FFN^l(\cdot)$ is the l -th layer of the Feed-Forward Network, and $MDSA^l(\cdot)$ means the l -th of the MDSA module.

We denote the transformed visual features as \hat{S}_{v2t} where the ‘‘v2t’’ shows that our expectation of transferring the feature from visual space to textural space.

Suppose that we have already learned good \hat{S}_{v2t} , to facilitate the enhancement of visual and textural features, we combine the original text features F_T together with the \hat{S}_{v2t} as the input of the Transformer encoder E_{T_E} as,

$$E_{T_E}([F_T + \hat{S}_{v2t}, F_I]), \quad (5)$$

Constraint Loss of VisGT. To supervise the learning of \hat{S}_{v2t} , we propose to use the contrastive loss (Hadsell et al. 2006) as the constraint (Pan et al. 2023a) between the predicted scene features \hat{s}^l and predefined scene features s . Formally, given a batch data with n images, we have the VisGT constraint loss as below,

$$\mathcal{L}_{VisGT} = p(s = \hat{s}_i) = \frac{\exp(\phi_{i,i}/\tau)}{\sum_{j=1}^n \exp(\phi_{i,j}/\tau)}, \quad (6)$$

where τ is the temperature parameter and $\phi_{i,j} = \hat{s}_i^{lT} s_j$ denotes the similarity matrix.

with the \mathcal{L}_{VisGT} and the classical classification loss \mathcal{L}_{cls} , localization loss \mathcal{L}_{loc} , our final loss function is as,

$$\mathcal{L} = \mathcal{L}_{cls} + \alpha \mathcal{L}_{loc} + \beta \mathcal{L}_{VisGT}, \quad (7)$$

where α and β are the weight factors.

Experiments

Experimental Setup

LAE-1M Dataset. We constructed a large-scale remote sensing object detection dataset by using our LAE-Label Engine pipeline as in Figure 2. As a brief recall, our dataset contains the fine-grained LAE-FOD and the coarse-grained LAE-COD. The final constructed LAE-1M dataset covered **one million** instances.

Table 1 summarizes the sub-datasets used for building the LAE-1M dataset. Specifically, for most of the datasets, a 0.4 random sampling rate is adopted if the number of instance

of same class as larger than 100. Xview is the only exception, for which we sample 0.2 to eliminate the duplicate instances. The purpose of sampling instances from different classes across all datasets is to maximize the learning of each class’s features while preserving the original dataset’s data distribution.

Evaluation Benchmarks. To evaluate the validity of our LAE-1M dataset and LAE-DINO model, DIOR (Li et al. 2020) and DOTAv2.0 (Xia et al. 2018) which are commonly used in the remote sensing community are used as benchmarks as in MTP(Wang et al. 2024a). Note that the results of DOTAv2.0 are all based on horizontal detection boxes for building a foundational location detector. In addition, to better validate the open-set detectors, we constructed LAE-80C containing 80 classes as a new remote sensing OVD benchmark. More details are included in Appendix. Based on the above three benchmarks, both the open-set and closed-set detection capabilities are evaluated. Specifically, we introduce the HRRSD (Zhang et al. 2019) dataset with a total of thirteen classes, which contains ten base classes appearing in LAE-1M dataset and three novel classes that do not, to perform the few-shot detection experiments. The mAP , AP_{50} , and AP_{75} are used as the evaluation metrics.

Implementation Details. We conducted all pre-training experiments on four A100 GPUs. To avoid memory overflow caused by having too many objects in a single image during batch training, we split image annotations with over 200 objects into smaller groups, ensuring the number of instances remains unchanged. Additionally, the alignment heads’ categories are set to 1600 for open-vocabulary pre-training. During training, key parameters are carefully set: the length of the dynamic vocabulary number in DVC module i.e., N_{DV} is set to 60, the number of layers l for MDSA and FFN is set to 7, and the hyper-parameters α and β of the loss function are set to 1 and 10, respectively. The open-vocabulary pre-training of LAE-DINO lasts approximately 180K steps, spanning about 48 GPU hours with a batch size of 2 per GPU. More details are provided in *More Implementation Details* section in Appendix.

	Datasets	Instances
LAE-COD	AID (Xia et al. 2017)	34,214
	NWPU-RESISC45 (Hichri 2021)	28,906
	SLM (Yuan et al. 2022)	106
	EMS (From Google Earth)	39,013
LAE-FOD	DOTA (Xia et al. 2018)	188,282
	DIOR (Li et al. 2020)	192,472
	FAIR1M (Sun et al. 2022b)	1.02 M
	NWPU VHR-10 (Cheng et al. 2014)	3,651
	RSOD (Long et al. 2017)	6,950
	Xview (Lam et al. 2018)	~ 1 M
	HRSC2016 (Liu et al. 2017)	2,976
Condensing-Tower (Zhang and Deng 2019)	2,382	

Table 1: LAE-1M dataset contains abundance categories composed of coarse-grained LAE-COD and fine-grained LAE-FOD.

Detection Results

Open-Set Detection. We compare the open-set detection results with two effective OVD methods, GLIP (Li et al. 2022) and GroundingDINO (Liu et al. 2024), trained on natural and remote sensing scenes datasets as shown in Table 2. Note that due to the natural difference of tasks, those CLIP-based (Wang et al. 2024b) and grounding (Mall et al. 2024; Kuckreja et al. 2024; Li et al. 2024a) methods are not considered as competitors. To train on LAE-1M dataset, we similarly introduce DVC on the GLIP and GroundingDINO. First of all, the detection results find that the OVD method of pre-training on natural scene dataset hardly works on remote sensing open-set detection, indicating a substantial gap between remote sensing and natural scene. Secondly, GroundingDINO has a more powerful open-set detection capability compared to GLIP from DIOR and LAE-80C. Clearly, our LAE-DINO has a better open-set detection compared with GroundingDINO, with an increase of 1.9%, 0.8%, and 2.5% on DIOR, DOTAv2.0, and LAE-80C benchmarks, respectively. These detection results show that our LAE-DINO has a more robust open-set detection in the remote sensing field.

Closed-Set Detection. To prove the benefits of OVD, we perform fine-tuning experiments in remote sensing scenes, comparing some generic detectors (GD) on DIOR and DOTAv2.0 datasets as shown in Table 3. Most previous GDs are fine-tuned to object detection datasets after pre-training on remote-sensing images using a self-supervised approach. We directly cite the original paper results due to the lack of open source for these generic detectors. For the OVD methods, we also provide results of fine-tuning experiments based on pre-training on natural scene datasets. Comparing the GD and OVD methods, it shows that the OVD method, which introduces textual prompts, is significantly higher than the GD method with a raise of about 6% at least in the AP_{50} on DIOR’s closed-set detection. The results of LAE-DINO fine-tuned on DIOR demonstrates that an outstanding performance on the DOTAv2.0, with a mAP of 57.9, an increase of 2.8% compared with GroundingDINO.

Table 4 shows that the closed-set detection results on DIOR test set with the fine-tuning data randomly sampled at different scales DIOR-*full*, DIOR- $\frac{1}{2}$, and DIOR- $\frac{1}{4}$ from the DIOR train set. We find that with just half of the DIOR train set, the AP_{50} could reach 89.1. This detection results shows that only a small amount of data is needed to fine-tune after open-vocabulary pre-training, which can achieve satisfactory results in real-world detection tasks.

Ablation Studies

VisGT Analysis. We perform ablation experiments on DIOR test set to explore the specific role of VisGT as shown in Table 5. *LAE-1M Pre-Training* is the open-set detection results, and *DIOR Fine-Tuning* is closed-set detection results that are directly fine-tuned on DIOR training dataset.

From the *LAE-1M Pre-Training* experiment, the group with VisGT achieved a 1.9% increase in AP_{50} for DIOR’s open-set detection. This result indicates that our VisGT enhances the understanding of complex remote sensing scenes by incorporating visual-guided text prompts. We also found

Method	Pre-Training Data	DIOR	DOTAv2.0	LAE-80C
		AP_{50}	mAP	mAP
GLIP (Li et al. 2022)	O365,GoldG,CC3M,SBU	1.1	0.2	0.1
GLIP with <i>DVC</i> (Li et al. 2022)	LAE-1M	82.8	43.0	16.5
GroundingDINO (Li et al. 2022)	O365,GoldG,Cap4M	0.3	0.3	0.1
GroundingDINO with <i>DVC</i> (Li et al. 2022)	LAE-1M	83.6	46.0	17.7
LAE-DINO (Ours)	LAE-1M	85.5	46.8	20.2

Table 2: The open-set detection results on DIOR, DOTAv2.0 and LAE-80C benchmarks. All models in the table are based on Swin-T and BERT backbones. O365, GoldG, CC3M, SBU and Cap4M are natural scene datasets.

Method	Backbone	Pre-Training Data	Fine-Tuning	
			DIOR(AP_{50})	DOTAv2.0(mAP)
<i>Generic Object Detection</i>				
GASSL (Ayush et al. 2021)	ResNet-50	-	67.40	-
CACO (Mall et al. 2023)	ResNet-50	Sentinel-2	66.91	-
TOV (Tao et al. 2023)	ResNet-50	TOV-NI,TOV-R	70.16	-
Scale-MAE (Reed et al. 2023)	ViT-L	FMoW	73.81	-
SatLas (Bastani et al. 2023)	Swin-B	SatlasPretrain	74.10	-
RingMo (Sun et al. 2022a)	Swin-B	RingMoPretrain	75.90	-
SkySense (Guo et al. 2024)	Swin-H	multi-modal RSI	78.73	-
MTP (Wang et al. 2024a)	Swin-H	MillionAID	81.10	-
<i>Open-Vocabulary Object Detection</i>				
GLIP-FT (Li et al. 2022)	Swin-T	O365,GoldG,CC3M,SBU	87.8	50.6
GroundingDINO-FT (Liu et al. 2024)	Swin-T	O365,GoldG,Cap4M	90.4	54.0
GroundingDINO-FT (Liu et al. 2024)	Swin-T	LAE-1M	91.1	55.1
LAE-DINO-FT (Ours)	Swin-T	O365,GoldG,Cap4M	92.0	55.5
LAE-DINO-FT (Ours)	Swin-T	LAE-1M	92.2	57.9

Table 3: The closed-set detection results on on DIOR and DOTAv2.0 test set. The results of DOTAv2.0 are all based on horizontal detection boxes. GeoImageNet, Sentinel-2, TOV-NI,TOV-R, FMoW, SatlasPretrain, MillionAID, RingMoPretrain and multi-modal RSI are remote sensing datasets.

Method	Fine-Tuning Data	DIOR AP_{50}
LAE-DINO	DIOR- <i>full</i>	92.2
LAE-DINO	DIOR- $\frac{1}{2}$	89.1
LAE-DINO	DIOR- $\frac{1}{4}$	85.6

Table 4: The closed-set detection results on the DIOR test set with the fine-tuning data randomly sampled at different scales from the DIOR train set.

a further improvement after *DIOR fine-tuning*, with an increase to 92.2 at AP_{50} , and further support for VisGT.

LAE-1M Analysis. To explore how LAE-COD and LAE-FOD of LAE-1M work, we set up two sets of comparison experiments on our LAE-DINO as shown in Appendix. We find that the detection of base classes in LAE-FOD can be improved by adding additional LAE-COD for pre-training, where the mAP of DOTAv2.0 test set can be improved by 2.3%. This also implies the feasibility of our LAE-Label to help interpret common categories of remote sensing imagery. As for the annotation quality and novel class detection of the LAE-Label engine, its survey report is in Appendix.

VisGT Reanalysis & Visualisation

We set different weights to \mathcal{L}_{VisGT} to observe its impact on detection performance. The reanalysis of VisGT and the visualisation of detection results are in Appendix.

Method	Pre-Training Data	DIOR AP_{50}
<i>LAE-1M Pre-Training</i>		
<i>PT-baseline</i>	LAE-1M	83.6
+ VisGT	LAE-1M	85.5
<i>DIOR Fine-Tuning</i>		
<i>FT-baseline</i>	-	89.9
+ VisGT	-	92.0
+ VisGT	LAE-1M	92.2

Table 5: The ablation results on the DIOR test set. PT-baseline denotes the Pre-Training baseline, and FT-baseline denotes the Fine-Tuning baseline.

Conclusion

In this paper, we introduced the Locate Anything on Earth (LAE) task, focusing on achieving open-vocabulary object detection for remote sensing. To advance the development of LAE, we concentrated on two key areas: 1) **Data:** We developed the LAE-Label Engine, a semi-automated labeling pipeline that collects and annotates data from up to 10 datasets. Using the LAE-Label Engine, we constructed LAE-1M, the first large-scale remote sensing object detection dataset. 2) **Model:** We presented LAE-DINO, a foundational open-vocabulary object detector for the LAE task, validated for its robust and generalizable detection capabilities. We believe our work will greatly advance Earth science applications by defining a clear task, providing large-scale training data, and offering a foundation model.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (42125503, 42430602).

References

- Ayush, K.; Uz Kent, B.; Meng, C.; Tanmay, K.; Burke, M.; Lobell, D.; and Ermon, S. 2021. Geography-aware self-supervised learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*.
- Bastani, F.; Wolters, P.; Gupta, R.; Ferdinando, J.; and Kembhavi, A. 2023. Satlaspretrain: A large-scale dataset for remote sensing image understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*.
- Chen, H.; Wang, Y.; Wang, G.; and Qiao, Y. 2018. Lstd: A low-shot transfer detector for object detection. In *Proceedings of the AAAI conference on artificial intelligence*.
- Chen, Z.; Wu, J.; Wang, W.; Su, W.; Chen, G.; Xing, S.; Zhong, M.; Zhang, Q.; Zhu, X.; Lu, L.; et al. 2024. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Cheng, G.; Han, J.; Zhou, P.; and Guo, L. 2014. Multi-class geospatial object detection and geographic image classification based on collection of part detectors. *ISPRS Journal of Photogrammetry and Remote Sensing*.
- Cong, Y.; Khanna, S.; Meng, C.; Liu, P.; Rozi, E.; He, Y.; Burke, M.; Lobell, D.; and Ermon, S. 2022. Satmae: Pre-training transformers for temporal and multi-spectral satellite imagery. *Advances in Neural Information Processing Systems*.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Fu, Y.; Wang, Y.; Pan, Y.; Huai, L.; Qiu, X.; Shangguan, Z.; Liu, T.; Kong, L.; Fu, Y.; Van Gool, L.; et al. 2024. Cross-Domain Few-Shot Object Detection via Enhanced Open-Set Object Detector. *arXiv preprint arXiv:2402.03094*.
- Guo, X.; Lao, J.; Dang, B.; Zhang, Y.; Yu, L.; Ru, L.; Zhong, L.; Huang, Z.; Wu, K.; Hu, D.; et al. 2024. Skysense: A multi-modal remote sensing foundation model towards universal interpretation for earth observation imagery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Hadsell, R.; Chopra, S.; LeCun, Y.; and LeCun, Y. 2006. Dimensionality reduction by learning an invariant mapping. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*.
- He, X.; Tang, C.; Zou, X.; and Zhang, W. 2023. Multi-spectral Object Detection via Cross-Modal Conflict-Aware Learning. In *Proceedings of the 31st ACM International Conference on Multimedia*.
- Hichri, H. 2021. NWPU-RESISC45 Dataset with 12 classes.
- Hospedales, T.; Antoniou, A.; Micaelli, P.; and Storkey, A. 2021. Meta-learning in neural networks: A survey. *IEEE transactions on pattern analysis and machine intelligence*.
- Kirillov, A.; Mintun, E.; Ravi, N.; Mao, H.; Rolland, C.; Gustafson, L.; Xiao, T.; Whitehead, S.; Berg, A. C.; Lo, W.-Y.; et al. 2023. Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*.
- Kuckreja, K.; Danish, M. S.; Naseer, M.; Das, A.; Khan, S.; and Khan, F. S. 2024. Geochat: Grounded large vision-language model for remote sensing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Lam, D.; Kuzma, R.; McGee, K.; Dooley, S.; Laielli, M.; Klaric, M.; Bulatov, Y.; and McCord, B. 2018. xView: Objects in Context in Overhead Imagery. *arXiv:1802.07856*.
- Li, K.; Wan, G.; Cheng, G.; Meng, L.; and Han, J. 2020. Object detection in optical remote sensing images: A survey and a new benchmark. *ISPRS journal of photogrammetry and remote sensing*.
- Li, K.; Wang, D.; Xu, H.; Zhong, H.; and Wang, C. 2024a. Language-guided progressive attention for visual grounding in remote sensing images. *IEEE Transactions on Geoscience and Remote Sensing*.
- Li, L. H.; Zhang, P.; Zhang, H.; Yang, J.; Li, C.; Zhong, Y.; Wang, L.; Yuan, L.; Zhang, L.; Hwang, J.-N.; et al. 2022. Grounded language-image pre-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Li, Y.; Guo, W.; Yang, X.; Liao, N.; He, D.; Zhou, J.; and Yu, W. 2024b. Toward Open Vocabulary Aerial Object Detection with CLIP-Activated Student-Teacher Learning. *arXiv:2311.11646*.
- Liu, S.; Zeng, Z.; Ren, T.; Li, F.; Zhang, H.; Yang, J.; Li, C.; Yang, J.; Su, H.; Zhu, J.; and Zhang, L. 2024. Grounding DINO: Marrying DINO with Grounded Pre-Training for Open-Set Object Detection.
- Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*.
- Liu, Z.; Yuan, L.; Weng, L.; and Yang, Y. 2017. A high resolution optical satellite image dataset for ship recognition and some new baselines. In *International conference on pattern recognition applications and methods*.
- Long, Y.; Gong, Y.; Xiao, Z.; and Liu, Q. 2017. Accurate object localization in remote sensing images based on convolutional neural networks. *IEEE Transactions on Geoscience and Remote Sensing*.
- Ma, Q.; Pan, J.; and Bai, C. 2024. Direction-Oriented Visual-Semantic Embedding Model for Remote Sensing Image-Text Retrieval. *IEEE Transactions on Geoscience and Remote Sensing*.
- Mall, U.; Hariharan, B.; Bala, K.; and Bala, K. 2023. Change-aware sampling and contrastive learning for satellite images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Mall, U.; Phoo, C. P.; Liu, M. K.; Vondrick, C.; Hariharan, B.; and Bala, K. 2024. Remote Sensing Vision-Language

- Foundation Models without Annotations via Ground Remote Alignment. In *The Twelfth International Conference on Learning Representations*.
- Pan, J.; Ma, Q.; Bai, C.; and Bai, C. 2023a. A Prior Instruction Representation Framework for Remote Sensing Image-text Retrieval. In *Proceedings of the 31st ACM International Conference on Multimedia*.
- Pan, J.; Ma, Q.; Bai, C.; and Bai, C. 2023b. Reducing Semantic Confusion: Scene-aware Aggregation Network for Remote Sensing Cross-modal Retrieval. In *Proceedings of the 2023 ACM International Conference on Multimedia Retrieval*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*.
- Redmon, J.; and Farhadi, A. 2017. YOLO9000: better, faster, stronger. In *Proceedings of the IEEE conference on computer vision and pattern recognition*.
- Reed, C. J.; Gupta, R.; Li, S.; Brockman, S.; Funk, C.; Clipp, B.; Keutzer, K.; Candido, S.; Uyttendaele, M.; and Darrell, T. 2023. Scale-mae: A scale-aware masked autoencoder for multiscale geospatial representation learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*.
- Ren, S.; He, K.; Girshick, R.; and Sun, J. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*.
- Rezatofghi, H.; Tsoi, N.; Gwak, J.; Sadeghian, A.; Reid, I.; and Savarese, S. 2019. Generalized intersection over union: A metric and a loss for bounding box regression. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*.
- Shen, Y.; Fu, C.; Chen, P.; Zhang, M.; Li, K.; Sun, X.; Wu, Y.; Lin, S.; and Ji, R. 2024. Aligning and prompting everything all at once for universal visual perception. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Sun, X.; Wang, P.; Lu, W.; Zhu, Z.; Lu, X.; He, Q.; Li, J.; Rong, X.; Yang, Z.; Chang, H.; et al. 2022a. RingMo: A remote sensing foundation model with masked image modeling. *IEEE Transactions on Geoscience and Remote Sensing*.
- Sun, X.; Wang, P.; Yan, Z.; Xu, F.; Wang, R.; Diao, W.; Chen, J.; Li, J.; Feng, Y.; Xu, T.; et al. 2022b. FAIR1M: A benchmark dataset for fine-grained object recognition in high-resolution remote sensing imagery. *ISPRS Journal of Photogrammetry and Remote Sensing*.
- Tao, C.; Qi, J.; Zhang, G.; Zhu, Q.; Lu, W.; and Li, H. 2023. TOV: The original vision model for optical remote sensing image understanding via self-supervised learning. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*.
- Wang, D.; Zhang, J.; Xu, M.; Liu, L.; Wang, D.; Gao, E.; Han, C.; Guo, H.; Du, B.; Tao, D.; et al. 2024a. MTP: Advancing remote sensing foundation model via multi-task pretraining. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*.
- Wang, Z.; Prabha, R.; Huang, T.; Wu, J.; and Rajagopal, R. 2024b. Skyscript: A large and semantically diverse vision-language dataset for remote sensing. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Wasim, S. T.; Naseer, M.; Khan, S.; Yang, M.-H.; and Khan, F. S. 2024. VideoGrounding-DINO: Towards Open-Vocabulary Spatio-Temporal Video Grounding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Xia, G.-S.; Bai, X.; Ding, J.; Zhu, Z.; Belongie, S.; Luo, J.; Dacu, M.; Pelillo, M.; and Zhang, L. 2018. DOTA: A Large-Scale Dataset for Object Detection in Aerial Images. In *The IEEE Conference on Computer Vision and Pattern Recognition*.
- Xia, G.-S.; Hu, J.; Hu, F.; Shi, B.; Bai, X.; Zhong, Y.; Zhang, L.; and Lu, X. 2017. AID: A benchmark data set for performance evaluation of aerial scene classification. *IEEE Transactions on Geoscience and Remote Sensing*.
- Xu, Y.; Zhang, M.; Fu, C.; Chen, P.; Yang, X.; Li, K.; and Xu, C. 2024. Multi-modal queried object detection in the wild. *Advances in Neural Information Processing Systems*.
- Yuan, Z.; Zhang, W.; Li, C.; Pan, Z.; Mao, Y.; Chen, J.; Li, S.; Wang, H.; and Sun, X. 2022. Learning to Evaluate Performance of Multimodal Semantic Localization. *IEEE Transactions on Geoscience and Remote Sensing*.
- Zareian, A.; Rosa, K. D.; Hu, D. H.; and Chang, S.-F. 2021. Open-vocabulary object detection using captions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Zhang, C.; and Wang, S. 2024. Good at captioning, bad at counting: Benchmarking GPT-4V on Earth observation data. arXiv:2401.17600.
- Zhang, H.; and Deng, Q. 2019. Deep learning based fossil-fuel power plant monitoring in high resolution remote sensing images: A comparative study. *Remote Sensing*.
- Zhang, H.; Li, F.; Liu, S.; Zhang, L.; Su, H.; Zhu, J.; Ni, L.; and Shum, H.-Y. 2023. DINO: DETR with Improved De-Noising Anchor Boxes for End-to-End Object Detection. In *The Eleventh International Conference on Learning Representations*.
- Zhang, Y.; Yuan, Y.; Feng, Y.; and Lu, X. 2019. Hierarchical and Robust Convolutional Neural Network for Very High-Resolution Remote Sensing Object Detection. *IEEE Transactions on Geoscience and Remote Sensing*.
- Zhong, Y.; Yang, J.; Zhang, P.; Li, C.; Codella, N.; Li, L. H.; Zhou, L.; Dai, X.; Yuan, L.; Li, Y.; et al. 2022. Regionclip: Region-based language-image pretraining. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*.
- Zhu, X.; Su, W.; Lu, L.; Li, B.; Wang, X.; and Dai, J. 2021. Deformable DETR: Deformable Transformers for End-to-End Object Detection. In *International Conference on Learning Representations*.