

Semantics from Space: Satellite-Guided Thermal Semantic Segmentation Annotation for Aerial Field Robots

Connor Lee, Saraswati Soedarmadji, Matthew Anderson, Anthony J. Clark, and Soon-Jo Chung

Abstract—We present a new method to automatically generate semantic segmentation annotations for thermal imagery captured from an aerial vehicle by utilizing satellite-derived data products alongside onboard global positioning and attitude estimates. This new capability overcomes the challenge of developing thermal semantic perception algorithms for field robots due to the lack of annotated thermal field datasets and the time and costs of manual annotation, enabling precise and rapid annotation of thermal data from field collection efforts at a massively-parallelizable scale. By incorporating a thermal-conditioned refinement step with visual foundation models, our approach can produce highly-precise semantic segmentation labels using low-resolution satellite land cover data for little-to-no cost. It achieves 98.5% of the performance from using costly high-resolution options and demonstrates between 70-160% improvement over popular zero-shot semantic segmentation methods based on large vision-language models currently used for generating annotations for RGB imagery. Code will be available at: <https://github.com/connorlee77/aerial-auto-segment>.

I. INTRODUCTION

Uninhabited Aerial Vehicles (UAVs) have been extensively used in field robotic applications, including precision agriculture [1], wildlife conservation [2], coastal mapping [3], and wildfire management [4]. To enable operations during nighttime and adverse weather conditions, UAVs can be equipped with long-wave thermal infrared cameras [5], [6] that provide dense scene perception in such settings. However, developing thermal scene perception for aerial field robotics requires ample data in order to train deep learning models for semantic segmentation [7]. This poses a challenge due to the scarcity of in-domain thermal data capturing typical aerial field robotic operational areas such as deserts [8], forests [4], and coastlines [9], [3].

Although several thermal semantic segmentation datasets of urban scenes have been curated for autonomous driving applications [10], [11], [12], few datasets exist that specifically target natural environments from an aerial viewpoint [13], [6]. To compensate for limited thermal data, existing works leverage large, annotated RGB datasets via domain adaptation techniques like image translation [10] and domain confusion [14], [15], as well as online learning [6] for thermal test-time adaptation. Despite reducing reliance on thermal training data, such methods still require annotated thermal data for comprehensive evaluation and robustness testing. While thermal datasets exist for field environments,

most lack annotations relevant for aerial semantic segmentation [4], [9], [16] besides [13]. As a result, collecting and annotating thermal datasets for semantic segmentation is still necessary to further improve thermal scene perception results via supervised training.

Capturing and annotating thermal field data presents unique challenges. Unlike in RGB, publicly-available thermal imagery is scarce due to the high costs and specialized nature of thermal sensors. Consequently, relevant thermal imagery cannot be scraped from the web and field roboticists must travel to various locations for data collection. This process incurs significant time and financial expenses, as it requires extensive travel and permits for flying and data capture. Moreover, annotating thermal imagery adds further costs and delays due to its distinct visual characteristics. This requires multiple rounds of attentive expert review and re-annotation [13], and adds more time to the curation process.

In this study, we propose a method to significantly reduce the time and cost of annotating aerial thermal field imagery for semantic segmentation. *Main contributions:* **1.** An algorithm that automatically generates high-quality segmentation labels for aerial thermal imagery using estimated camera pose and satellite-derived data. **2.** Experiments comparing segmentation labels generated from various satellite-derived data products, demonstrating competitive results with free options. **3.** Extensive ablation studies showcasing the robustness of our method to noisy camera pose estimation and temporal misalignments between thermal and satellite imagery. **4.** A demonstration for aerial field robot perception by training a semantic segmentation network solely on labels generated using our method, yielding promising results.

II. RELATED WORK

Semantic Segmentation: Semantic segmentation models perform per-pixel classification and are typically built upon convolutional neural networks and transformer architectures [18], [19], [20]. While conventional fully-supervised models achieve impressive results, they need large annotated training datasets for generalization. In applications like thermal semantic segmentation where labeled data is scarce, unsupervised domain adaptation (UDA) techniques are often employed. UDA methods like [10] synthesize labeled thermal training data from existing RGB datasets via image translation, while other works [14], [15] leverage RGB training for thermal inference by maximizing RGB-thermal domain confusion during training. However, UDA methods still face challenges: they require significant target domain data and thermal annotations for evaluation,

*This work was supported by the Office of Naval Research. The authors are with the California Institute of Technology. Email: {clee, ssoedarm, matta, ajc, sjchung}@caltech.edu. A. J. Clark is also with Pomona College.

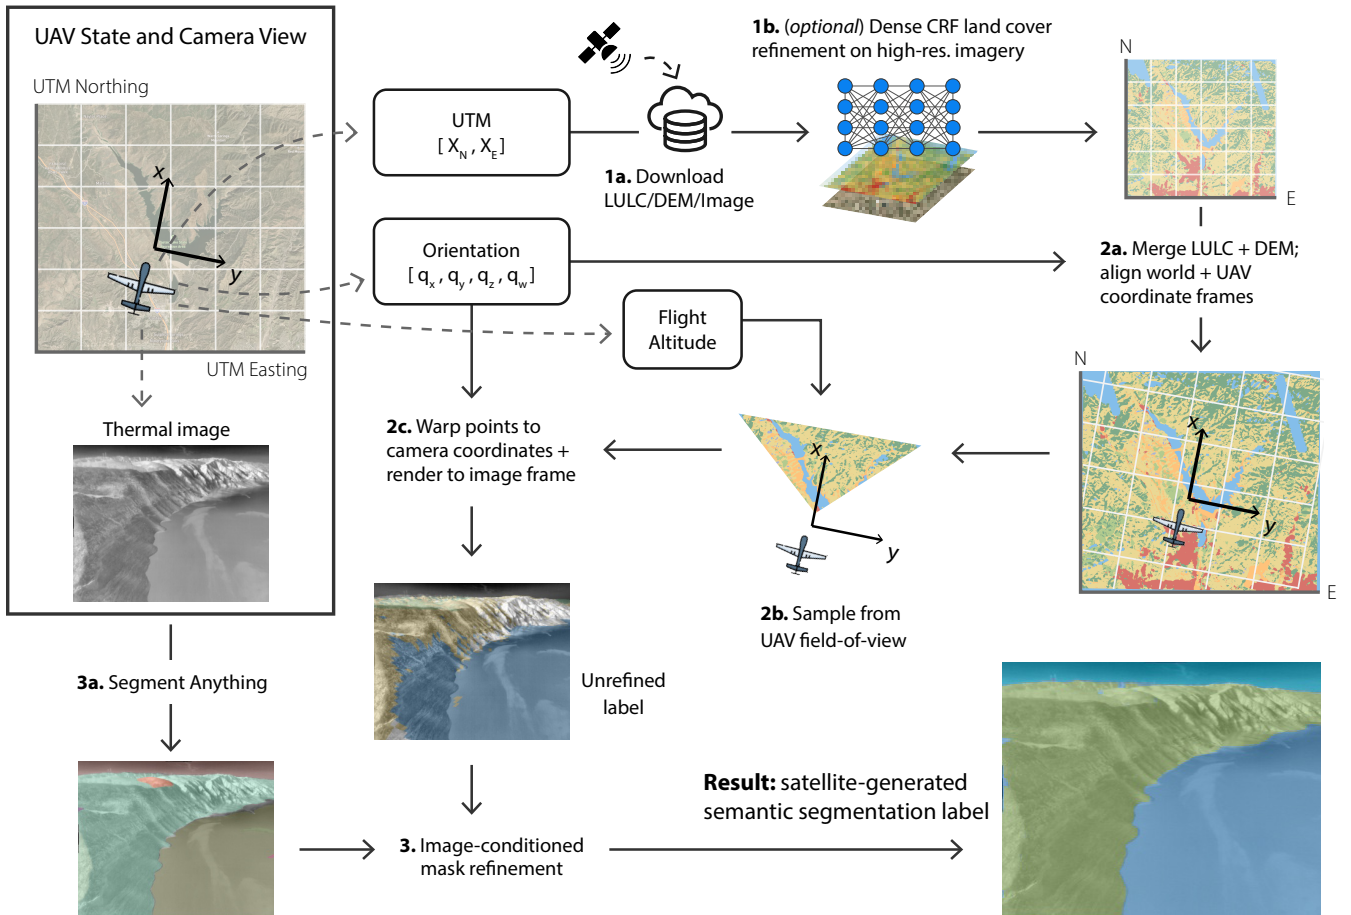


Fig. 1: Proposed pipeline for automatically generating semantic segmentation annotations from satellite-derived data. Coarse segmentation labels for thermal images are rendered from Land Use and Land Cover (LULC) datasets and Digital Elevation Maps (DEM). The labels are refined using Segment Anything [17] to capture fine details between segmentation instances.

and typically exhibit lower performance compared to fully-supervised methods [21].

Alternatively, recent large vision-language models like ODISE [22] and OV-Seg [23] can perform zero-shot semantic segmentation across the RGB spectrum by leveraging user-provided text prompts. Similarly, the Segment Anything Model [17] (SAM) can provide precise segmentations for any object but lacks semantic information. In general, the zero-shot semantic segmentation methods perform worse on thermal imagery compared to RGB [13]. Despite this, [13] finds that SAM can perform well in a semantic segmentation task if its segmentation outputs are assigned ground truth class labels. We leverage this finding in our approach.

Automatic Semantic Segmentation Annotation: Most works using automatic semantic segmentation labeling can be found in self-training and self-supervised learning literature. However, many focus on specialized applications with niche classes [6], [24] and are not relevant for general scene segmentation. For generalized semantic segmentation tasks, [25] self-trains their model using noisy labels predicted by their network for intra-RGB domain adaptation. In contrast, [26] adopts an incremental training approach and utilizes humans to select good network outputs as annotations and

manually correct bad ones before retraining. Other works manually annotate a subset of frames in video data, before propagating them to remaining frames using optical flow [27] or learned generative models [28].

As discussed, visual foundation models can also be used for annotation efforts. Zero-shot semantic segmentation models [22], [23], [29] are being used to provide labels, but do not transfer directly to non-RGB domains. [30] generates object detection labels for thermal imagery by using SAM on aligned RGB images and does not work in low-light settings.

In contrast to other works, [31] uses 3D information to generate semantic segmentation labels for construction sites and is most similar to our work. They register Building Information Models with point clouds from photogrammetry and render the labeled 3D points to an image frame. Unlike other works, methods like this operate independently of an image and can work for any imaging modality.

III. PRELIMINARIES

In this section, we briefly go over the different satellite-derived data products we use in our approach (Sec. IV).

Land Use and Land Cover Datasets: Publicly-available Land Use and Land Cover (LULC) datasets like Dynamic

World [32] and Impact Observatory [33] derive from satellite rasters obtained through the Sentinel-2 program. These datasets have a low spatial resolution of 10 m/pixel but have global coverage, and are updated using semantic segmentation networks that use multiple data bands for landcover classification. While daily coverage is possible, availability depends on factors like cloud coverage.

In contrast, high-resolution LULC like the Chesapeake Bay Program [34] and OpenEarthMap [35] offer sub-meter resolution but are limited in geographical and temporal coverage. While segmentation models can be trained on these datasets with high-resolution imagery, they may not generalize to different geographical areas.

High-Resolution Raster Imagery: These include imagery from aerial vehicles and satellites. Aerial imagery providers include the National Agricultural Imagery Program (NAIP) [36] while satellite imagery comes from providers like Planet, Maxar, and Airbus. Image resolutions range from 0.3 m/pixel to 3 m/pixel. Imagery can be available daily at a premium cost while free alternatives are captured triannually.

Lidar-Derived 3D Data Products: Digital surface (DSM) and digital elevation models (DEM) are raster data whose values denote the height at the corresponding geographic location. DSMs consider features above the ground like foliage and rocky terrain while DEMs report bare earth elevation. In this work, we use DEMs and DSMs with 1 m/pixel to 3 m/pixel resolution from the 3D Elevation Program (3DEP) from the United States Geological Survey [37].

IV. APPROACH

We present a 3 step method to automatically generate semantic segmentation annotations for thermal images captured from an aerial vehicle using satellite-derived data (Fig. 1).

A. Step 1: Generating 3D Semantic Maps from Satellite Data

We start by downloading relevant satellite data (LULC rasters, DEM or DSM, and high-resolution imagery) around the aerial vehicle’s global position and resample them to the highest resolution via bicubic interpolation. To simplify future calculations, we convert to UTM coordinates before merging the DEM and LULC rasters. Since current freely-available LULC data is low resolution (10 m), we optionally refine them by conditioning on high resolution imagery as described below. Alternatively, high-resolution LULC can also be created using a pretrained LULC segmentation network on high resolution imagery (see Sec. VI-A.2).

Land-Use-Land-Cover Refinement: We use dense conditional random fields [38] (CRF) to refine 10 m resolution LULC rasters with 1 m-3 m resolution aerial imagery (Fig. 2). To summarize, a dense CRF is defined by a Boltzmann distribution with energy function

$$E(\mathbf{X}|\mathbf{I}) = \sum_i \psi_u(x_i|I_i) + \sum_{i<j} \psi_p(x_i, x_j|I_i) \quad (1)$$

This function models the relationship between labels $\mathbf{x} \in \mathbf{X}$ and the conditioning image $I \in \mathbb{R}^{H \times W \times C}$. Here, ψ_u is a unary potential taken to be raw logits from a semantic

segmentation network and ψ_p is a pairwise potential that encourages label consistency among adjacent pixels with similar intensities.

In our method, we set ψ_u to be the logits from the model that generated our LULC labels. Like [39], we use a generalized ψ_p to condition on multi-band raster images:

$$\psi_p(\mathbf{f}_i, \mathbf{f}_j) = \mu \cdot \left[w^{(1)} \exp \left(-\frac{1}{2} \bar{\mathbf{p}}_{ij}^\top \Sigma_\alpha \bar{\mathbf{p}}_{ij} - \frac{1}{2} \bar{\mathbf{I}}_{ij}^\top \Sigma_\beta \bar{\mathbf{I}}_{ij} \right) + w^{(2)} \exp \left(-\frac{1}{2} \bar{\mathbf{p}}_{ij}^\top \Sigma_\gamma \bar{\mathbf{p}}_{ij} \right) \right] \quad (2)$$

Here, $\bar{\mathbf{p}}_{ij} = [\mathbf{p}_i - \mathbf{p}_j] \in \mathbb{R}^3$ is the difference between positions of pixel i and j , and $\bar{\mathbf{I}}_{ij} = [\mathbf{I}_i - \mathbf{I}_j] \in \mathbb{R}^C$ is the difference between image features at pixels i and j . We set μ as the standard Potts compatibility function [38].

We optimize the CRF by tuning weight parameters $w^{(1)}$ and $w^{(2)}$, and the Gaussian bandwidth parameters $\Sigma_\alpha = \theta_\alpha$, $\Sigma_\gamma = \theta_\gamma$, and $\Sigma_\beta = \text{diag}(\theta_\beta^{(1)}, \dots, \theta_\beta^{(C)})$. We minimize the boundary loss [40] and use this instead of cross entropy to account for imprecise labels at class boundaries due to low LULC resolution.

B. Step 2: LULC Projection to Aerial Camera Image Frame

To generate an LULC-derived semantic label for an image at time t , we start by transforming the world coordinates of each pixel $\mathbf{X}_w \in \mathbb{R}^3$ into the camera coordinate frame. This requires the position of the host vehicle $\mathbf{x}_t^{\text{uav}} \in \mathbb{R}^3$, taken from the onboard EKF-fused GPS position and barometric altitude, the orientation quaternion $\mathbf{q}_t \in \mathbb{H}$, taken from the EKF-fused IMU readings, and the offset between the aircraft and camera reference points. Using the calibrated camera intrinsic matrix \mathbf{K} , we can then project to image coordinates $\mathbf{x}_t^c \in \mathbb{Z}^2$. Formally, this is

$$\begin{bmatrix} \mathbf{x}_t^c \\ 1 \end{bmatrix} = \mathbf{K} \begin{bmatrix} \mathbf{R}(\mathbf{q}_t) & \mathbf{T}(\mathbf{x}_t^{\text{uav}}) \\ \mathbf{0}_{1 \times 3} & 1 \end{bmatrix} \begin{bmatrix} \mathbf{X}_w \\ 1 \end{bmatrix} \quad (3)$$

where \mathbf{R} is a rotation matrix and \mathbf{T} is a translation vector. We use OpenGL to render the projected LULC, using 3D coordinates as vertices, associated class labels as vertex colors, and depth-testing to avoid rendering occluded semantics.

To optimize memory and speed, we only consider 3D semantics within a specified distance in front of and on both sides of the camera. We also exponentially increase spacing between sampled vertices as distance from the camera increases, exploiting the compression of far-field points in the image frame. This enables us to use only 250×200 points when rendering within a $10 \text{ km} \times 8 \text{ km}$ bounding box.

C. Step 3: Rendered Label Refinement

Though the semantic segmentation labels have been rendered, they do not align well with the thermal images. This is primarily caused by poor spatial resolution and temporal misalignment, but could also stem from errors in LULC label generation and camera pose estimation. To improve alignment, we refine the labels by generating binary segmentation masks of the corresponding thermal image using the Segment




Fig. 2: Dense CRF refinement of Dynamic World land cover raster using NAIP and PlanetScope imagery of Castaic Lake, CA. Results via PlanetScope convey the actual scenery at time of thermal image capture due to its high revisit frequency but at a lower 3 m spatial resolution. NAIP refinement offers 1 m resolution but is susceptible to changes in the terrain (notably, water levels of lakes) due to its triennial capture cycle. Zoom in to see key differences (outlined in dashed boxes).

Algorithm 1 SAM-based Label Refinement

- 1: **Input:** Projected (unrefined) label mask $L \in \mathbb{N}^{H \times W}$,
- 2: Thermal image $I \in \mathbb{R}^{H \times W}$
- 3: **Output:** Refined semantic segmentation label M
- 4: **Initialize:** Segment Anything Model f_{sam}
- 5:
- 6: $\{M_{\text{sam}}^i\}_0^N \leftarrow f_{\text{sam}}(I)$ ▷ SAM produces binary masks
- 7: Initialize zero-array M of size $H \times W$
- 8: **for** $m_{\text{sam}} \in \{M_{\text{sam}}^i\}_0^N$ **do**
- 9: $x_{idx} \leftarrow [m_{\text{sam}} == 1]$ ▷ Get mask indices
- 10: $y_{\text{cls}} \leftarrow L[x_{idx}].\text{mode}()$ ▷ Find most freq. class in mask
- 11: $M[x_{idx}] = y_{\text{cls}}$
- 12: **end for**
- 13: **return** M

Anything Model. Then, we assign each mask a semantic class based on the most prevalent LULC class within it (Alg. 1).

V. LOW ALTITUDE AERIAL DATASET

We test our method using a thermal field robotics dataset, which includes off-nadir (20° - 45°) aerial views of rivers (Kentucky River, KY and Colorado River, CA), lakes (Castaic Lake, CA), and coastal (Duck, NC) areas across the United States [6], [13]. The dataset, captured from a multicopter, comprises 15 flight trajectories ranging from 40 m to 100 m in altitude and contains time-synchronized thermal imagery, GPS, and IMU measurements. Four trajectories are excluded from testing due to GPS data collection errors. While the dataset provides ground truth semantic segmentation annotations for 10 classes, we condense the classes into 6 categories in order to better conform with land cover classes. We end up with ground truth, 6-class semantic segmentation labels for 1304 sub-sampled images (CM-6) and further condense the classes again to create two additional class-sets, CM-5 (5 classes) and CM-3 (3 classes). A mapping of segmentation labels is shown in Fig. 3.

VI. RESULTS

A. Experimental Setup

1) *Raster Acquisition:* We acquired 10 m resolution Dynamic World LULC, 3D terrain data (3 m DEM, 1 m DEM,

TABLE I: Evaluation of dense CRF refinement of Dynamic World LULC on NAIP imagery with ground truth labels from Chesapeake Bay Program (see class mapping in Fig. 3).

CRF cond. source	Boundary Loss ↓	mIoU ↑	Dense CRF Parameters							
			w_1	w_2	θ_γ	θ_α	$\theta_\beta^{(0)}$	$\theta_\beta^{(1)}$	$\theta_\beta^{(2)}$	$\theta_\beta^{(3)}$
None	0.945	0.432	—	—	—	—	—	—	—	—
RGB [†]	0.914	0.441	1.00	1.00	200	195	7.00	7.00	7.00	—
RGB	0.777	0.452	47.2	2.63	33.3	149	1.14	1.14	1.14	—
RGB-NIR	0.749	0.453	47.4	0.14	61.5	194	128	0.22	125	2.71

[†]tuned by minimizing weighted cross entropy instead of boundary loss

2 m DSM) from USGS 3DEP, and high-resolution nadir imagery from NAIP (1 m) and Planet (3 m). Data was obtained via Microsoft Planetary Computer and Google Earth Engine.

2) *LULC from High-Resolution Imagery:* We used networks trained on Chesapeake Bay Program (CBP) and OpenEarthMap (OEM) datasets to produce two more high-resolution LULC sources alongside Dynamic World. For OEM-derived LULC, we used the pretrained U-Net model from [35]. To produce CBP-derived LULC, we fine-tuned a geospatial foundation model [41] on the CBP dataset, using the 7-class set from [34].

We trained for 1000 epochs with a batch size of 16, a $1e^{-3}$ learning rate, and RGB-NIR inputs of size 512×512 . To perform inference on large raster images, we use tiles with 50% overlap and applied flips for test-time augmentation.

3) *LULC Refinement with Dense CRFs:* We refined the 10 m Dynamic World LULC rasters on RGB-NIR imagery from NAIP and Planet (Fig. 2) using parameters from Tab. I. Parameters were found using Bayesian optimization with Optuna [42]. The search was done using NAIP as conditioning imagery and 1 m resolution labels from CBP as ground truth (see Fig. 3 for class mapping). For this use case, boundary loss was superior to standard cross-entropy loss (Tab. I).

4) *Rendered Label Refinement:* For SAM refinement of the projected LULC labels, we used the default ViT-H model. We prompted with 32×32 grid points and lowered the box non-maximum suppression threshold to 0.5.

5) *Thermal Image Preprocessing:* We rescaled raw 16-bit thermal pixel intensities to sit between the 2nd and 98th percentiles before applying a contrast-limited adaptive

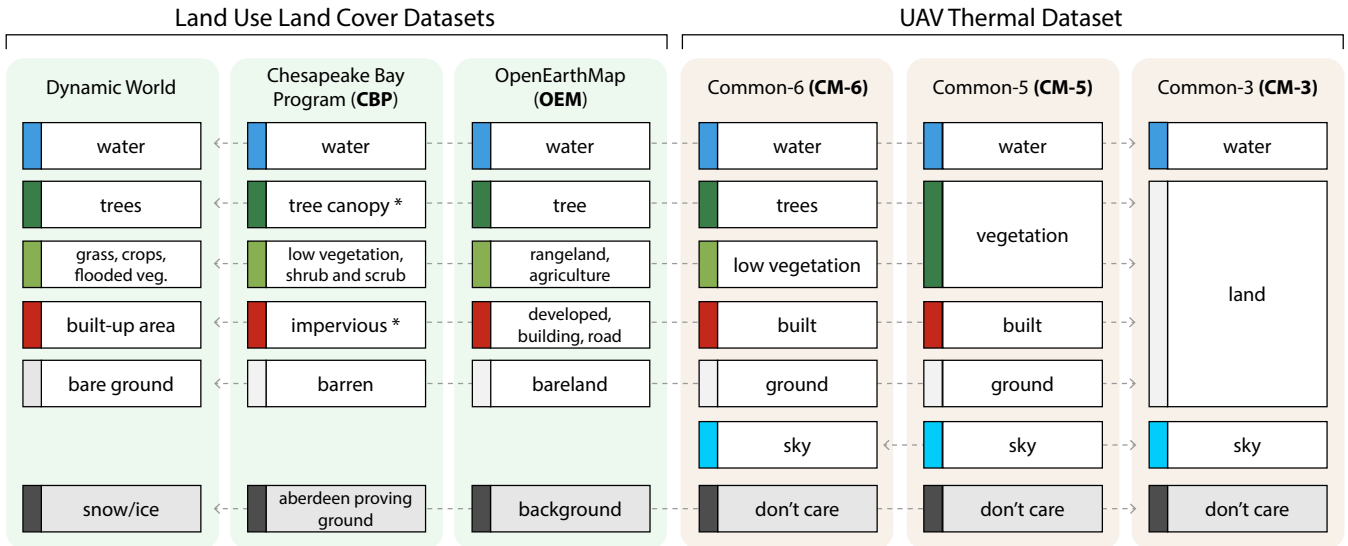


Fig. 3: Class mappings between LULC datasets and our ground truth evaluation set. The UAV thermal dataset is from [13].

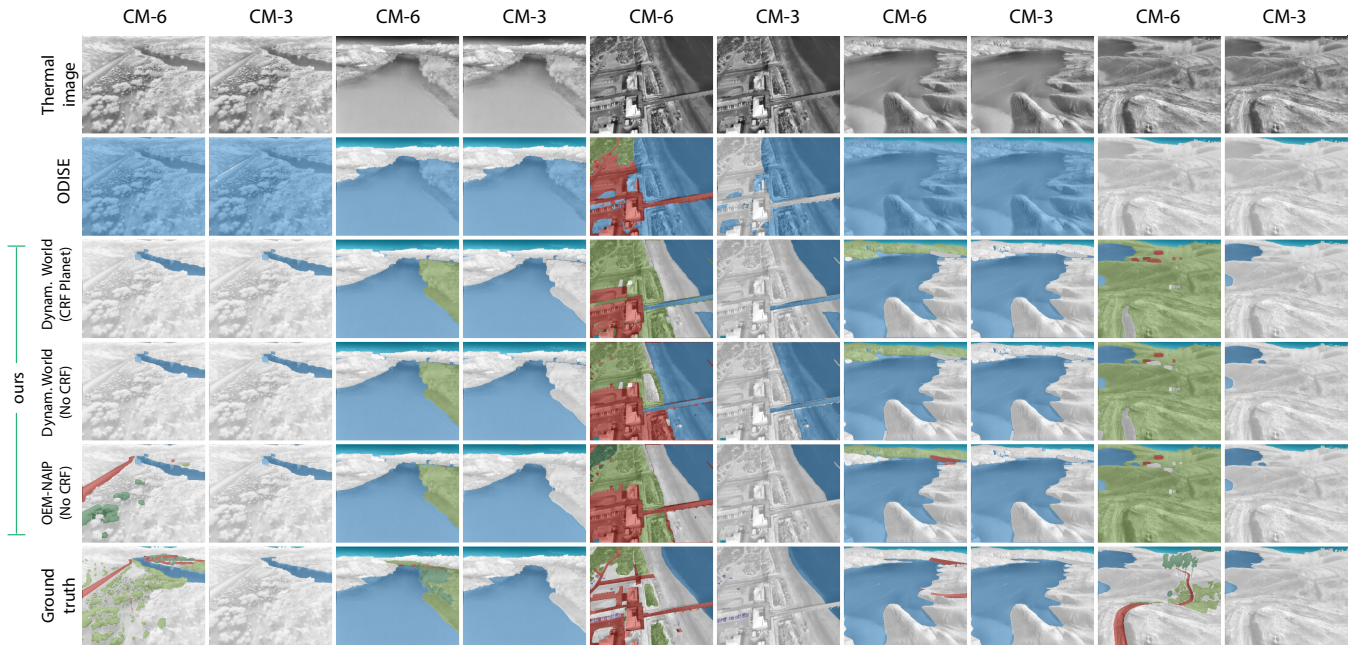


Fig. 4: Generated segmentations from the baseline (ODISE [22]), our methods, and the ground truth (GT) using class mappings and colors from Fig. 3. Mismatches between CM-6 labels and GT can occur depending on the LULC source used but are resolved with CM-3. Segmentations for classes containing small, sparse, and thin instances (CM-6), e.g. *low vegetation* and *built*, are hard to render due to low LULC resolution and low thermal contrast during label refinement.

histogram equalization with a 0.02 clip limit, following [6]. This was done for both visualization and algorithm input.

B. Satellite-based Semantic Segmentation Label Generation

We compare our LULC-generated semantic segmentation labels to manually-annotated, ground truth labels. Due to class differences between LULC data and ground truth, we evaluate on three ground truth-derived class sets of increasing generality (CM-6, CM-5, CM-3). We report the overall dataset mIoU and the trajectory-averaged mIoU in Tab. II.

Overall, our method delivers thermal semantic segmentation labels consistent with ground truth (Fig. 4). Notably, our best variants greatly outperform the zero-shot semantic segmentation models, ODISE [22], and OV-Seg [23], which were prompted with a list of classes present in the dataset. We note that ODISE and OV-Seg are occasionally effective on thermal images, but lack consistency.

Among our methods without LULC refinement, semantic segmentation label generation using Dynamic World and

TABLE II: LULC-generated semantic segmentation label assessment (mIoU) when compared to ground truth annotations, with comparisons against zero-shot visual foundation model baselines.

Method / LULC source	Dense CRF refinement src.	3D source	Dataset mIoU \uparrow			Trajectory avg. mIoU \uparrow		
			CM-6	CM-5	CM-3	CM-6	CM-5	CM-3
ODISE [22]	—	—	0.299	0.262	0.330	0.264	0.304	0.413
OV-Seg [23]	—	—	0.201	0.240	0.385	0.183	0.233	0.390
Chesapeake Bay (NAIP)	—	DEM (3m)	0.453	0.481	0.857	0.417	0.478	0.848
Chesapeake Bay (Planet)	—	DEM (3m)	0.236	0.305	0.657	0.201	0.251	0.555
Open Earth Map (NAIP)	—	DEM (3m)	0.549	0.562	0.868	0.440	0.528	0.864
Open Earth Map (Planet)	—	DEM (3m)	0.502	0.509	0.825	0.360	0.428	0.816
Dynamic World	—	DEM (3m)	0.577	0.572	0.876	0.450	0.518	0.860
Dynamic World	NAIP	DEM (3m)	0.556	0.535	0.868	0.441	0.504	0.865
Dynamic World	Planet	DEM (3m)	0.573	0.557	0.887	0.455	0.510	0.870

DEM (3 m) as a 3D source generally outperforms other variants using CBP- and OEM-derived LULC sources. LULC created from the OEM network on NAIP data provides improvements (0.005 - 0.01 mIoU) in trajectory-averaged mIoU over Dynamic World for the CM-5 and CM-3 class sets. Despite marginal gains, this is likely due to the higher resolution (1 m) of OEM/NAIP-derived LULC, which enables segmentation renderings of small or thin classes that are present in CM-5, such as roads (see Fig. 4). This is not possible with Dynamic World due to its lower 10 m resolution.

Conversely, LULC generated from Planet imagery provides poor results due to domain differences between OEM/CBP training images and Planet. When used for dense CRF refinement of 10 m Dynamic World rasters, however, Planet imagery uniquely provides ~ 0.01 boost for both mIoU metrics on the most general CM-3 class set. This behavior is absent when refining on NAIP imagery due to terrain changes between thermal and NAIP acquisition dates.

Furthermore, we note that our method can handle temporal mismatches between satellite and thermal data even as environments naturally evolve. For example, coastal tide patterns and varying lake levels (Fig. 5, Castaic Lake) may shift class boundaries within short time periods. Due to SAM’s ability to capture entire class instances, our rendered label refinement step (Sec. IV-C) is notably able to overcome such changes as long as most of the true class is still rendered.

Due to its accessibility and competitive performance, we advocate using Dynamic World LULC for satellite-based semantic segmentation label generation efforts, with potential refinement via temporally-relevant, high resolution imagery. However, this will inevitably change with advancements in LULC creation and as sub-meter data products with high temporal coverage become more freely-accessible.




Fig. 5: Rendered label refinement process with SAM [17].

TABLE III: Ablation studies

(a) 3D source ablation

Method	3D source	Traj. avg. mIoU		
		CM-6	CM-5	CM-3
Dynamic World + SAM	DEM (3m)	0.450	0.518	0.860
	DEM (1m)	0.441	0.507	0.842
	DSM (2m)	0.439	0.504	0.848

(b) Label refinement ablation

Method	Projected label refine method	Traj. avg. mIoU		
		CM-6	CM-5	CM-3
Dynamic World + DEM (3m)	SAM	0.450	0.518	0.860
	SLIC	0.392	0.452	0.711
	Felzenszwalb	0.369	0.426	0.677

C. Ablation Study

In these ablations, we use Dynamic World as our semantic source. Unless otherwise specified, we use 3 m DEMs to add 3D context and do not use any CRF refinement.

1) *3D Data Source*: First, we compare LULC-based semantic segmentation label generation with 3 m DEMs against two additional 3D data sources: 2 m DSMs and 1 m DEMs. Due to limited coverage, we lack DSMs and 1 m DEMs over the thermal data capture areas of Colorado River and Duck, respectively, and resort to 3 m DEMs in those areas. Our results show that 3 m DEMs provide consistently higher trajectory-averaged mIoU across all three class sets, despite the other two sources supposedly providing more accurate and precise 3D terrain data (Tab. IIIa). Reasons for this may include temporal differences or spatial misalignment during orthorectification. Nonetheless, all 3D sources generally perform well and any one of these 3D products can be used for our method when the other two are unavailable.




Fig. 6: Effect of LULC spatial resolution on semantic segmentation label generation.

2) *Raster Spatial Resolution*: To assess the impact of LULC spatial resolution on label generation, we generate labels from Dynamic World LULC rasters resampled to 10 m (native), 5 m, and 1 m resolution. We use nearest neighbor interpolation on the LULC directly, and CRF refinement on NAIP and Planet rasters (resampled to 10 m, 5 m, and 1 m resolutions via bicubic interpolation).

Our results (Fig. 6) suggest that LULC spatial resolution matters more for more specific class sets (CM-6/CM-5), and becomes less critical as class sets generalize (CM-3). Moreover, we find greater benefits from CRFs when conditioning on higher-resolution imagery, especially when dealing with the larger and more specific class sets (CM-6/CM-5). This is

likely due to smoothing over small or thin class instances that comprise of a few pixels when refining at lower resolutions.

3) *Segment Anything vs. Classical Methods for Projected LULC Refinement*: We compare our choice of SAM for projected LULC label refinement against SLIC [43] and Felzenszab [44] superpixels. We use implementations from `scikit-image` [45], setting SLIC’s number of segments to 100 and compactness to 10, and Felzenszab’s scale parameter to $1e^4$. We select these parameters to maximize segmentation area while remaining within class instances.

Overall, SAM consistently outperforms the other methods, with mIoU margins increasing from 0.06 (Tab. IIIb). This is because SAM can produce semantically distinct masks in the thermal domain, albeit less reliably than in RGB. This allows minor imperfections to be ignored through majority vote (Alg. 1). In contrast, classical methods produce fragmented, semantic-agnostic masks which offer little benefit.




Fig. 7: Effect of global pose estimate precision on semantic segmentation label generation with SAM refinement.

4) *State Estimate Precision*: To quantify the effect of global pose estimation precision on our label generation process, we systematically perturb these measurements by sampling from a normal distribution with increasing variance. Our analysis reveals that, with 95% confidence, label generation remains robust for global positioning and altitude estimates within roughly 4 m and for orientations within roughly 3.5° (Fig. 7). These findings are consistent across class sets. During development, both synchronizing the timing of image capture to the IMU data was shown to be critical, as was the SAM refinement stage for compensation for attitude estimate errors (see Kentucky River in Fig. 5).

D. Application: Semantic Segmentation Model Training

To demonstrate our method for field robot perception, we trained an EfficientViT-B0 semantic segmentation network [46] using the aerial thermal dataset and general train/val/test split from [13]. Three sets of labels (CM-6, CM-5, and CM-3) were generated for training and validation using our method, with ground truth labels converted accordingly for testing and baseline training. All networks were trained following the thermal training procedure from [13].

Our semantic segmentation results (Tab. IVa) closely match the mIoU of the generated annotations (Tab. II). Networks trained with CM-3 classes resulted in 0.889 mIoU during testing, compared to 0.962 mIoU for those trained with ground truth labels. Networks trained on CM-5 and CM-6 show larger gaps (Tab. IVa) but still show the benefit of our method. We find this is largely due to difficulties in

accurately rendering land-based classes, specifically *low vegetation* and *built* (Tab. IVb). These classes contain small and thin entities like sparse shrubs or roads, and are not always precisely shown in LULC data. Also, they can be missed during rendered label refinement (Sec. IV-C) due to blurred and low-contrast appearance in thermal imagery. Despite this, our method can effectively train semantic segmentation models, particularly with the CM-3 class set, and support field robotic applications like nighttime river navigation [6].

TABLE IV: Test results (mIoU) of semantic segmentation networks trained on LULC-generated labels and networks trained on manually-annotated ground truth.

(a) Segmentation results after training on CM-6 (least inclusive), CM-5, and CM-3 (most inclusive) class sets.

Annotation Method	Class set		
	CM-6	CM-5	CM-3
LULC-generated	0.542	0.547	0.889
Ground truth	0.819	0.836	0.962

(b) Per-class IoU for networks trained using the CM-6 class set.

Annot. Method	water	trees	low veg.	built	ground	sky
Generated	0.880	0.529	0.165	0.289	0.521	0.868
Ground truth	0.963	0.787	0.702	0.653	0.854	0.955

E. Computational Costs and Pricing

Our method annotates a single image in 3 s, 2.86 s of which is due to SAM. Annotations are *free* when using only Dynamic World LULC but cost $\sim \$10$ USD/km² with CRF refinement due to the price of realtime, high-resolution satellite imagery. With our method, annotating 2 000 images takes 1.6 hours on a single workstation, in contrast with the usual 2-4 week timeframe and \$3 000 to \$8 000 USD outsourcing cost¹. We note that CRF refinement can be cost-effective for large data volumes in a concentrated area due to its one-time cost, but 98.5% of its performance (CM-6, CM-3) is achievable with free 10 m resolution LULC (Sec. VI-B).

VII. CONCLUSION

We presented a novel method for automatically generating high-quality semantic segmentation annotations for classes often encountered by aerial robots in field settings. Our approach leverages satellite data products and employs refinement steps to achieve fine precision at class boundaries even with low-resolution satellite data, achieving 98.5% of the performance of costly high-resolution options. We demonstrated the robustness of our method to global positioning and attitude estimation errors, indicating that it can provide good segmentations even with inexpensive sensors and slight data desynchronization, and identified limitations due to small and thin class instances. Lastly, we demonstrated its application to field robot perception by successfully training a semantic segmentation network solely with generated labels. This method enables rapid training of thermal perception stacks using incremental learning as new field data is collected.

¹\$1.50 to \$4.00 per image for 1-10 semantic segmentation classes, based on pricing from Scale AI at the time of writing: <https://scale.com/pricing>

REFERENCES

- [1] A. Pretto, *et al.*, “Building an aerial-ground robotics system for precision farming: an adaptable solution,” *IEEE Robot. Automat. Mag.*, vol. 28, no. 3, pp. 29–49, 2020.
- [2] E. Bondi *et al.*, “Birdsai: A dataset for detection and tracking in aerial thermal infrared videos,” in *Proc. IEEE Winter Conf. Appl. Comput. Vis.*, 2020, pp. 1747–1756.
- [3] K. L. Brodie, B. L. Bruder, R. K. Slocum, and N. J. Spore, “Simultaneous mapping of coastal topography and bathymetry from a lightweight multicamera uas,” *IEEE Trans. Geosci. Remote Sensing*, vol. 57, no. 9, pp. 6844–6864, 2019.
- [4] A. Jong *et al.*, “WIT-UAS: A wildland-fire infrared thermal dataset to detect crew assets from aerial views,” in *Proc. IEEE/RSJ Int. Conf. Intell. Robots and Syst.*, 2023, pp. 11 464–11 471.
- [5] J. Delaune, R. Hewitt, L. Lytle, C. Sorice, R. Thakker, and L. Matthies, “Thermal-inertial odometry for autonomous flight throughout the night,” in *Proc. IEEE/RSJ Int. Conf. Intell. Robots and Syst.*, 2019, pp. 1122–1128.
- [6] C. Lee, J. G. Frennert, L. Gan, M. Anderson, and S.-J. Chung, “Online self-supervised thermal water segmentation for aerial vehicles,” in *Proc. IEEE/RSJ Int. Conf. Intell. Robots and Syst.*, 2023, pp. 7734–7741.
- [7] S. Minaee, Y. Boykov, F. Porikli, A. Plaza, N. Kehtarnavaz, and D. Terzopoulos, “Image segmentation using deep learning: A survey,” *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 44, no. 7, pp. 3523–3542, 2021.
- [8] G. Loianno, *et al.*, “Localization, grasping, and transportation of magnetic objects by a team of mavs in challenging desert-like environments,” *IEEE Robotics and Automation Letters*, vol. 3, no. 3, pp. 1576–1583, 2018.
- [9] S. Nirgudkar, M. DeFilippo, M. Sacarny, M. Benjamin, and P. Robinette, “Massmind: Massachusetts maritime infrared dataset,” *Int. J. Robot. Res.*, vol. 42, no. 1-2, pp. 21–32, 2023.
- [10] C. Li, W. Xia, Y. Yan, B. Luo, and J. Tang, “Segmenting objects in day and night: Edge-conditioned cnn for thermal image semantic segmentation,” *IEEE Transactions on Neural Networks and Learning Systems*, vol. 32, no. 7, pp. 3069–3082, 2020.
- [11] Q. Ha, K. Watanabe, T. Karasawa, Y. Ushiku, and T. Harada, “Mfnet: Towards real-time semantic segmentation for autonomous vehicles with multi-spectral scenes,” in *Proc. IEEE/RSJ Int. Conf. Intell. Robots and Syst.*, 2017, pp. 5108–5115.
- [12] J. Vertens, J. Zürn, and W. Burgard, “Heatnet: Bridging the day-night domain gap in semantic segmentation with thermal images,” in *Proc. IEEE/RSJ Int. Conf. Intell. Robots and Syst.*, 2020, pp. 8461–8468.
- [13] C. Lee, *et al.*, “CART: Caltech aerial RGB-thermal dataset in the wild,” *arXiv preprint arXiv:2403.08997*, 2024.
- [14] L. Gan, C. Lee, and S.-J. Chung, “Unsupervised RGB-to-thermal domain adaptation via multi-domain attention network,” in *Proc. IEEE Int. Conf. Robot. and Automation*, 2023, pp. 6014–6020.
- [15] Y.-H. Kim, U. Shin, J. Park, and I. S. Kweon, “Ms-uda: Multi-spectral unsupervised domain adaptation for thermal image semantic segmentation,” *IEEE Robotics and Automation Letters*, vol. 6, no. 4, pp. 6497–6504, 2021.
- [16] S. S. Shivakumar, N. Rodrigues, A. Zhou, I. D. Miller, V. Kumar, and C. J. Taylor, “Pst900: RGB-thermal calibration, dataset and segmentation network,” in *Proc. IEEE Int. Conf. Robot. and Automation*, 2020, pp. 9441–9447.
- [17] A. Kirillov *et al.*, “Segment anything,” *arXiv:2304.02643*, 2023.
- [18] L.-C. Chen, G. Papandreou, F. Schroff, and H. Adam, “Rethinking atrous convolution for semantic image segmentation,” *arXiv preprint arXiv:1706.05587*, 2017.
- [19] K. He, G. Gkioxari, P. Dollár, and R. Girshick, “Mask R-CNN,” in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 2961–2969.
- [20] Z. Liu *et al.*, “Swin transformer: Hierarchical vision transformer using shifted windows,” in *Proc. IEEE Int. Conf. Comput. Vis.*, 2021, pp. 10 012–10 022.
- [21] A. Mehra, B. Kailkhura, P.-Y. Chen, and J. Hamm, “Understanding the limits of unsupervised domain adaptation via data poisoning,” *Proc. Advances Neural Inform. Process. Syst. Conf.*, vol. 34, pp. 17 347–17 359, 2021.
- [22] J. Xu, S. Liu, A. Vahdat, W. Byeon, X. Wang, and S. De Mello, “Open-vocabulary panoptic segmentation with text-to-image diffusion models,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2023, pp. 2955–2966.
- [23] F. Liang *et al.*, “Open-vocabulary semantic segmentation with mask-adapted clip,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2023, pp. 7061–7070.
- [24] S. Daftry, Y. Agrawal, and L. Matthies, “Online self-supervised long-range scene segmentation for MAVs,” in *Proc. IEEE/RSJ Int. Conf. Intell. Robots and Syst.*, 2018, pp. 5194–5199.
- [25] N. Araslanov and S. Roth, “Self-supervised augmentation consistency for adapting semantic segmentation,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2021, pp. 15 384–15 394.
- [26] B. Yu, R. Tibbetts, T. Barna, A. Morales, I. Rekleitis, and M. J. Islam, “Weakly supervised caveline detection for auv navigation inside underwater caves,” in *Proc. IEEE/RSJ Int. Conf. Intell. Robots and Syst.*, 2023, pp. 9933–9940.
- [27] A. Marcu, V. Licaret, D. Costea, and M. Leordeanu, “Semantics through time: Semi-supervised segmentation of aerial videos with iterative label propagation,” in *Proc. Asian Conf. Comput. Vis.*, 2020.
- [28] A. Berg, J. Johnander, F. Durand de Gevigney, J. Ahlberg, and M. Felberg, “Semi-automatic annotation of objects in visual-thermal video,” in *Proc. IEEE Int. Conf. Comput. Vis. Workshop*, Oct 2019, pp. 2242–2251.
- [29] Scale AI, “Introducing scale’s automotive foundation model,” <https://scale.com/blog/afm1>, 2023.
- [30] J. E. Gallagher, A. Gogia, and E. J. Oughton, “A multispectral automated transfer technique (matt) for machine-driven image labeling utilizing the segment anything model (sam),” *arXiv preprint arXiv:2402.11413*, 2024.
- [31] A. Braun and A. Borrmann, “Combining inverse photogrammetry and bim for automated labeling of construction site images for machine learning,” *Automation in Construction*, vol. 106, p. 102879, 2019.
- [32] C. F. Brown *et al.*, “Dynamic world, near real-time global 10 m land use land cover mapping,” *Scientific Data*, vol. 9, no. 1, p. 251, 2022.
- [33] K. Karra, C. Kontgis, Z. Statman-Weil, J. C. Mazzariello, M. Mathis, and S. P. Brumby, “Global land use / land cover with sentinel 2 and deep learning,” in *Proc. IEEE Int. Geosci. and Remote Sensing Symp.*, 2021, pp. 4704–4707.
- [34] C. Robinson *et al.*, “Large scale high-resolution land cover mapping with multi-resolution data,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2019, pp. 12 726–12 735.
- [35] J. Xia, N. Yokoya, B. Adriano, and C. Broni-Bediako, “Openearthmap: A benchmark dataset for global high-resolution land cover mapping,” in *Proc. IEEE Winter Conf. Appl. Comput. Vis.*, 2023, pp. 6254–6264.
- [36] U.S. Department of Agriculture Farm Service Agency, Aerial Photography Field Office, “National Agricultural Imagery Program,” <https://earthexplorer.usgs.gov>, 2012-2023.
- [37] U.S. Geological Survey, “3D Elevation Program,” <https://usgs.gov/3d-elevation-program>.
- [38] P. Krähnenbühl and V. Koltun, “Efficient inference in fully connected crfs with gaussian edge potentials,” *Proc. Advances Neural Inform. Process. Syst. Conf.*, vol. 24, 2011.
- [39] K. Kamnitsas *et al.*, “Efficient multi-scale 3d cnn with fully connected crf for accurate brain lesion segmentation,” *Medical Image Analysis*, vol. 36, pp. 61–78, 2017.
- [40] A. Bokhovkin and E. Burnaev, “Boundary loss for remote sensing imagery semantic segmentation,” in *Proc. Int. Symp. on Neural Networks*. Springer, 2019, pp. 388–401.
- [41] M. Mendieta, B. Han, X. Shi, Y. Zhu, and C. Chen, “Towards geospatial foundation models via continual pretraining,” in *Proc. IEEE Int. Conf. Comput. Vis.*, 2023, pp. 16 806–16 816.
- [42] T. Akiba, S. Sano, T. Yanase, T. Ohta, and M. Koyama, “Optuna: A next-generation hyperparameter optimization framework,” in *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2019.
- [43] R. Achanta, A. Shaji, K. Smith, A. Lucchi, P. Fua, and S. Süsstrunk, “Slic superpixels compared to state-of-the-art superpixel methods,” *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 34, no. 11, pp. 2274–2282, 2012.
- [44] P. F. Felzenszwalb and D. P. Huttenlocher, “Efficient graph-based image segmentation,” *Int. J. Comput. Vis.*, vol. 59, pp. 167–181, 2004.
- [45] S. van der Walt *et al.*, “scikit-image: image processing in python,” *PeerJ*, vol. 2, p. e453, jun 2014.
- [46] H. Cai, J. Li, M. Hu, C. Gan, and S. Han, “Efficientvit: Lightweight multi-scale attention for high-resolution dense prediction,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 17 302–17 313.