

Generalizable Disaster Damage Assessment via Change Detection with Vision Foundation Model

Kyeongjin Ahn¹, Sungwon Han¹, Sungwon Park¹,
Jihee Kim^{2,1,3}, Sangyoon Park⁴, Meeyoung Cha^{5,1*}

¹School of Computing, Korea Advanced Institute of Science and Technology (KAIST), Daejeon, South Korea

²College of Business, Korea Advanced Institute of Science and Technology (KAIST), Daejeon, South Korea

³Graduate School of Data Science, Korea Advanced Institute of Science and Technology (KAIST), Daejeon, South Korea

⁴Division of Social Science, Hong Kong University of Science and Technology (HKUST), Kowloon, Hong Kong

⁵Data Science for Humanity, Max Planck Institute for Security and Privacy (MPI-SP), Bochum, Germany

Abstract

The increasing frequency and intensity of natural disasters call for rapid and accurate damage assessment. In response, disaster benchmark datasets from high-resolution satellite imagery have been constructed to develop methods for detecting damaged areas. However, these methods face significant challenges when applied to previously unseen regions due to the limited geographical and disaster-type diversity in the existing datasets. We introduce DAVI (Disaster Assessment with Vision foundation model), a novel approach that addresses domain disparities and detects structural damage at the building level without requiring ground-truth labels for target regions. DAVI combines task-specific knowledge from a model trained on source regions with task-agnostic knowledge from an image segmentation model to generate pseudo labels indicating potential damage in target regions. It then utilizes a two-stage refinement process, which operate at both pixel and image levels, to accurately identify changes in disaster-affected areas. Our evaluation, including a case study on the 2023 Türkiye earthquake, demonstrates that our model achieves exceptional performance across diverse terrains (e.g., North America, Asia, and the Middle East) and disaster types (e.g., wildfires, hurricanes, and tsunamis). This confirms its robustness in disaster assessment without dependence on ground-truth labels and highlights its practical applicability.

1 Introduction

Natural disasters such as wildfires, hurricanes, floods, and earthquakes are becoming increasingly prevalent and severe due to climate change (IPCC 2022). These catastrophic events significantly impact human lives and cause extensive economic losses, necessitating prompt responses in rescue operations and financial assistance. With limited resources available for life-saving relief and infrastructure recovery, assessing disaster damage at a fine-grained infrastructure-level becomes critical for effective resource distribution (Fiedrich, Gehbauer, and Rickers 2000; Field et al. 2012). Among the various efforts, multi-temporal remote sensing images have emerged as a valuable resource for disaster damage assessment, which can cover broad

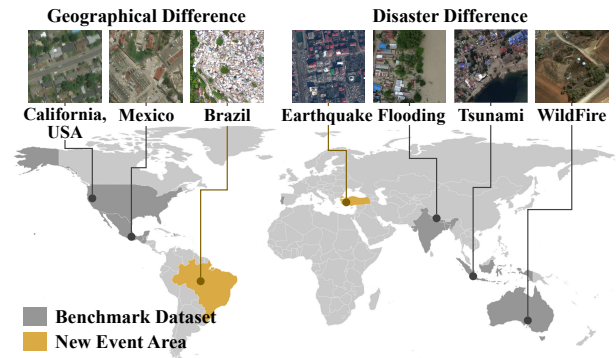


Figure 1: Collaborative efforts are underway to collect datasets on natural disasters, yet no single benchmark dataset covers the entire world. The yellow areas represent regions with recent disasters not included in the dataset.

geographic regions at a lower cost relative to traditional ground surveys (Weber and Kané 2020; Wu et al. 2021b; Gupta and Shah 2021).

Conventional methods rely on *change detection approaches* that compare visual differences in regional images taken before and after a disaster. Methods such as image differentiation (St-Charles, Bilodeau, and Bergevin 2014), image ratio (Skifstad and Jain 1989), principal component analysis (Kuncheva and Faithfull 2013), and change vector analysis (Bruzzone and Prieto 2000) are commonly applied to these problems. Recent works detect changes using deep features through thresholding techniques (Saha, Bovolo, and Bruzzone 2019), specialized convolution networks (Zhao et al. 2017), and contrastive learning (Wu et al. 2021a).

However, these techniques face significant challenges in real-world damage assessment due to substantial discrepancies between available training data (source data) and test data (target data). Figure 1 illustrates this issue, showing that the regions in the extensive disaster benchmark dataset, xBD (Gupta et al. 2019), span 45,000 square kilometers of annotated high-resolution satellite images of natural disasters, yet cover only a small part of the world. The limited coverage of these benchmark datasets makes it difficult to analyze features of previously unseen regions in terms of visual style. Additionally, the visual features of

*The corresponding author

damages vary by disaster type; for instance, wildfires can destroy entire buildings, hurricanes may damage only roofs, and floods may engulf structures without causing visible damage. Therefore, *domain adaptation* in images (i.e., shifts in visual styles) is required to handle diverse geographical landscapes, while *adjustment* in the damage types (i.e., shift in the decision boundary) must be supported.

We present DAVI (Disaster Assessment with VISION foundation model), designed to bridge domain disparities and detect structural damages at the building level in an unsupervised fashion. This is achieved through a test-time adaptation technique, which starts with a source model that is trained on independent source data and gradually updates the model to be used on target data without requiring target labels using pseudo labels (Liang, Hu, and Feng 2020; Wang et al. 2022; Niu et al. 2022; Chen et al. 2022). We minimize the entropy of pseudo labels to reduce uncertainty arising from domain shifts. Our approach also recognizes decision boundary shifts by employing an existing image segmentation model, SAM (Segment Anything Model) (Kirillov et al. 2023), and integrating the pseudo labels with general visual features. Image segmentation is especially useful for disaster assessment because it can generate semantic masks for various objects, including buildings, roads, vegetation, and water bodies. We demonstrate that this foundation model is also effective for investigating structural damage in disasters.

DAVI refines pseudo labels from the source model through a two-stage process. First, the pixel-level adjustment uses inconsistent change maps from temporally paired images and their augmented ones, then the image-level adjustment leverages coarse-grained damage presence information. These pseudo labels provide supervisory guidance to align model predictions in the absence of ground-truth labels. We evaluate the effectiveness of DAVI using real-world events from diverse regions and disaster types within the xBD benchmark dataset. The results show that DAVI successfully identifies structural damage in previously unseen areas, outperforming unsupervised change detection and domain adaptation baselines. Additionally, we validate its practical applicability with recent disaster samples not included in the benchmark dataset.

Disaster damage in uncharted territories exhibits immense diversity in visual and inherent characteristics. To effectively guide metrics until on-the-ground assessments, a detection method capable of handling such varied content is essential. DAVI, built on a vision foundation, represents a significant step forward in advancing this capability.

2 Related Work

2.1 Transfer Learning in Change Detection

Transfer learning methods focus on applying a pre-trained model from the source data to the target data. Given the high cost of acquiring ground-truth labels, especially in damage assessment, studies have utilized unsupervised methods. For example, several approaches use pre-trained weights on pretext tasks like segmentation or classification to detect changes in temporally paired images (Hou, Wang, and Liu 2017; de Jong and Bosman 2019; Zhang and Shi 2020;

Wu et al. 2021a; Leenstra et al. 2021). However, transfer learning algorithms often assume that the training and test sets are independent and identically distributed (*i.i.d.*), which degrades performance when domain discrepancies exist (Patricia and Caputo 2014). Domain adaptation techniques, an advanced version of transfer learning, partially address these discrepancies but assume that if pre-training and fine-tuning tasks share the same class label, the class’s semantic characteristics remain consistent, even if the visual styles differ (Wang et al. 2020; Liang, Hu, and Feng 2020; Wang et al. 2022). For example, an apple illustration and a real-life photo have different visual styles but share semantic traits that make them recognizable as apples. In our problem setting, however, the semantic differences caused by different types of disasters make domain adaptation methods difficult to apply, requiring a decision boundary shift.

2.2 Vision Foundation Model

Deep neural networks and self-supervised learning has significantly improved vision foundation models (VFMs) on large-scale vision data (Lüddecke and Ecker 2022; Wang et al. 2023b; Kirillov et al. 2023; Wang et al. 2023a; Zou et al. 2024; Wang et al. 2024). Examples of VFMs for image segmentation include, but are not limited to, CLIPSeg (Lüddecke and Ecker 2022), SegGPT (Wang et al. 2023b), SAM (Kirillov et al. 2023), and SEEM (Zou et al. 2024). These models use visual inputs and prompts — such as text, boxes, points, or masks — to specify target image segmentation and show remarkable adaptability to other domains, such as remote sensing (Osco et al. 2023; Chen et al. 2024), tracking (Yang et al. 2023a; Rajič et al. 2023), robotics (Yang et al. 2023b), and medical (Lei et al. 2023; Gong et al. 2023; Shaharabany et al. 2023; Wu et al. 2023; Ma et al. 2024), underscoring its effective capability to recognize various visual content. We employ SAM in our change detection method as an assistant due to its simplicity and ease of use. SAM uses an image encoder, a prompt encoder, and a mask decoder. It combines image and prompt embeddings through each encoder to guide object localization, and uses the mask decoder to generate masks from these embeddings for self-annotation.

2.3 Change Detection by Pseudo Labeling

When ground-truth labels are not available, one common strategy to provide supervisory signals is to introduce pseudo labels for training (Lee et al. 2013; Zou et al. 2020; Hu et al. 2023; Kim et al. 2022). This pseudo labeling has been used in change detection problems as well to achieve higher accuracy. Examples include Pix4Cap (Liu et al. 2023), which uses pseudo labels extracted from a pre-trained model to improve the model’s ability to change captioning, and PPL (Li et al. 2024), which generates pseudo labels through uncertainty estimation and progressive refinement on hyperspectral images. These change detection methods using pseudo labels rely on self-updating through their own predictions, making implementation easy but accumulating errors over time. In this work, we utilize pseudo-labeling for learning together with an external foundation model and refinement process to reduce error propagation.

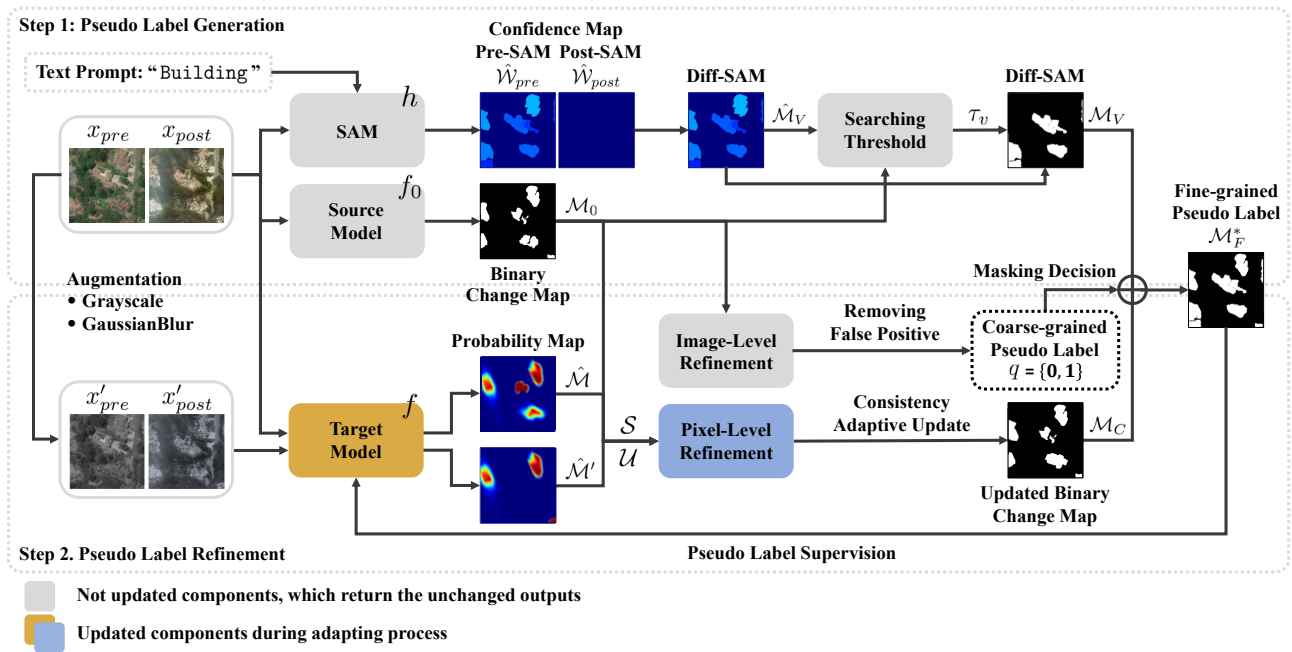


Figure 2: Model architecture. This figure illustrates how pre- and post-disaster images and their augmented versions are processed. The system has two phases: (Step 1) pseudo label generation, which leverages the source model and image segmentation foundation model (SAM) with a text prompt for generalizability across various disaster scenarios, and (Step 2) pseudo label refinement, which involves pixel- and image-level refinement to reduce noise in the fine-grained pseudo label. The final output is a fine-grained pseudo label indicating potential damage.

3 Methodology

3.1 Problem Formulation

We present an efficient method for disaster assessment, DAVI, which compares pre- and post-disaster images to detect structural damage (e.g., building, infrastructure level). DAVI runs in a test-time adaptation setting, using change detection models trained on unveiled disaster-related data (i.e., source data), adapted to target data. Below is the problem formulation of our setting with an overview of DAVI.

Let us define the source data D_s as unveiled data from the source domain and the target data D_t as veiled data from the target domain. The source data D_s contain temporally paired images and their corresponding ground truth labels, indicated as $\{\bar{x}_{pre}, \bar{x}_{post}, y\} \in D_s$, where $\bar{x}_{pre} \in \mathbb{R}^{H \times W \times 3}$ represents a pre-disaster image, $\bar{x}_{post} \in \mathbb{R}^{H \times W \times 3}$ represents a post-disaster image and $y \in \{0, 1\}^{H \times W}$ represents a binary change map, respectively. The target data D_t consist of temporally paired images, indicated as $\{x_{pre}, x_{post}\} \in D_t$ without access to ground-truth labels. Here, we ensure that there is no overlap between the source and target data. The model trained with the source data (i.e., source model) is denoted as f_0 . The target model, denoted as f , starts with the same weights as the source model f_0 , and is then updated to adapt to the target data during training. Our goal is to develop a domain-shifting strategy that relies solely on the source model and unlabeled target data, ultimately achieving high-accuracy predictions for changed areas within the target data using the target model f .

Figure 2 illustrates the model architecture. We design pseudo labels $\mathcal{M}_F^* \in \{0, 1\}^{H \times W}$ that accommodate domain shifts across various types of disaster occurring globally. These pseudo labels substitute for unavailable ground-truth labels in the target data. The intuition behind using pseudo labels is that they provide supervisory guidance, helping the model detect changes in the target data. The generation phase (Step 1) leverages both the source model f_0 with task-specific knowledge based on the source domain and SAM h with task-agnostic knowledge from extensive image corpus. This is followed by a refinement process (Step 2) that improves the accuracy of the pseudo labels by removing noisy parts within them using consistency-based information and coarse-grained information. Ultimately, in addition to supervised learning on the generated pseudo labels, we minimize the entropy of the predictions to learn domain-specific features for the target data. Our objective of training the target model f can be formulated as:

$$H(\hat{\mathcal{M}}) = - \sum_i \sum_j p^{ij} \log(p^{ij}) \quad (1)$$

where $H(\hat{\mathcal{M}})$ refers to Shannon entropy loss function (Shannon 1948) frequently employed in one of the test-time adaptation approaches, which tries to connect errors arising from domain discrepancies and domain shifts (Wang et al. 2020). Here, $\hat{\mathcal{M}}$ indicates the change maps computed by the target model f , and p^{ij} denotes the probability value within $\hat{\mathcal{M}}$, where (i, j) represents the coordinates of

individual pixels. Through this objective function, the target model f eliminates generalization errors by minimizing the entropy of its predictions on the target data.

$$\mathcal{L} = CE(\mathcal{M}_F^*, \hat{\mathcal{M}}) + \lambda H(\hat{\mathcal{M}}) \quad (2)$$

where \mathcal{M}_F^* is generated fine-grained pseudo labels and λ is an adjusting hyperparameter, which is set to 0.1 as default.

3.2 Step 1: Pseudo Label Generation

We generate the fine-grained pseudo labels \mathcal{M}_F that represent pixel-level binary change maps. This information combines two types of pseudo labels: \mathcal{M}_0 from the source model f_0 and \mathcal{M}_V from SAM, h . This enhances model generalizability across various disaster scenarios by utilizing complementary knowledge. Relying solely on pseudo labels from the source model can lead to inaccurate signals due to mismatches between source and target distributions, ultimately hindering the model’s applicability. If the model overfits the source data, it may exhibit undue uncertainty or certainty when applied to the target data. This affects the interpretation of the model’s predictions, potentially requiring a recalibration of confidence thresholds or adjustments to output probabilities. Trained on billion-scale vision data, SAM can analyze visual features across various domains. By integrating this with the source model’s knowledge, the combined approach adeptly manages decision boundary shifts, enhancing adaptation in diverse scenarios.

Pseudo label generation with task-specific knowledge.

Given the paired images x_{pre} and x_{post} , the source model f_0 predicts the probability p_0^{ij} of the change for each pixel in the change maps $\hat{\mathcal{M}}_0$. (i.e., $\hat{\mathcal{M}}_0^{ij} = p_0^{ij}$)

$$\hat{\mathcal{M}}_0^{ij} = f_0(x_{pre}^{ij}, x_{post}^{ij}) \quad (3)$$

We convert the probability maps $\hat{\mathcal{M}}_0$ into the binary change maps \mathcal{M}_0 as follows:

$$\mathcal{M}_0^{ij} = \begin{cases} 1 & \text{if } p_0^{ij} \geq 0.5 \\ 0 & \text{otherwise} \end{cases}, \quad m_0^{ij} \in \mathcal{M}_0 \quad (4)$$

Pseudo label generation with task-agnostic knowledge.

Next, our method segments objects of interest (e.g., buildings, roads, vegetation, and water bodies) based on the vision foundation model like SAM to adjust the source model’s pseudo label. It segments specific objects using a text prompt and assigns a confidence score to each object:

$$\hat{\mathcal{W}} = h(x; \text{prompt}), \quad w^{ij} \in \hat{\mathcal{W}} \text{ and } w^{ij} \in [0, 1]. \quad (5)$$

To obtain semantic changes between the paired images, we subtract two segmented maps with confidence scores — specifically, Pre-SAM and Post-SAM. When using the text prompt “Building”, SAM has high confidence for identifying building areas in pre-disaster images, while it would have decreased confidence for damaged building areas in post-disaster images because it lacks the essential features for recognizing buildings. Therefore, we leverage the confidence difference between pre-disaster and post-disaster images as an indicator of building destruction, utilizing SAM h with the text prompt “Building”:

$$\hat{\mathcal{M}}_V^{ij} = \max(0, w_{pre}^{ij} - w_{post}^{ij}), \quad d_v^{ij} \in \hat{\mathcal{M}}_V \quad (6)$$

where $\hat{\mathcal{M}}_V$ represents the confidence difference maps from SAM, referred to as Diff-SAM. We clip the confidence difference maps to be greater than zero to focus on the changes in destruction rather than construction. Then, we transform the confidence difference maps $\hat{\mathcal{M}}_V$ into the binary change maps \mathcal{M}_V as follows:

$$\mathcal{M}_V^{ij} = \begin{cases} 1 & \text{if } d_v^{ij} \geq \tau_v \\ 0 & \text{otherwise} \end{cases} \quad (7)$$

where τ_v represents a threshold determined through optimal threshold searching. For simplicity, we also hereafter refer to \mathcal{M}_V as Diff-SAM. Finally, we combine two binary change maps, using a max function to generate the fine-grained pseudo labels $\hat{\mathcal{M}}_F$:

$$\hat{\mathcal{M}}_F^{ij} = \max_{ij}(\mathcal{M}_0^{ij}, \mathcal{M}_V^{ij}) \quad (8)$$

This combination enables the target model to focus on damaged patterns that the source model could not capture, thereby enhancing recall for affected areas.

Searching Optimal Threshold. The threshold τ_v sets the confidence level at which segmented buildings are identified as damaged. The appearance of damaged buildings can differ depending on factors such as regional characteristics and the type of disaster, meaning the optimal threshold can also vary accordingly. We determine the optimal threshold by exploring various candidate values that maximize the F1-score between \mathcal{M}_0 and \mathcal{M}_V on the target data. This approach is based on the intuition that finding the precise threshold requires disaster-specific knowledge, which can be aided by the source model, even if it pertains to a different domain.

3.3 Step 2: Pseudo Label Refinement

We refine the fine-grained pseudo labels through in two processes: **Consistency Adaptive Update** and **Coarse-grained Pseudo Labels** to minimize noise in these labels. Despite utilizing knowledge from the source model and SAM, the noise remains in the fine-grained pseudo labels. This issue can be attributed to the distortion of knowledge adapted from the source model to the target distribution due to the domain discrepancy, or to SAM’s insufficient capability to adequately consider features associated with changes in the target data. We first focus on pixel-level adjustments, targeting areas identified as changed within the fine-grained pseudo labels, guided by confidence-based consistency from the data augmentation strategy. We then engage in image-level adjustments to further enhance noise reduction effectively. This approach leverages on broadly unchanged information between the paired images, from the coarse-grained pseudo labels. These labels are used to mask the generated fine-grained pseudo labels, rendering them as all zeros.

Consistency Adaptive Update. Unsupervised or semi-supervised learning can lead to overconfident and poorly calibrated predictions (Park et al. 2021; Gao et al. 2020). Consequently, pseudo-labeling-based on confidence scores can be suboptimal. To address this, we propose using inconsistencies arising from data augmentation variations to

assess uncertainty for pixel-level refinement. These inconsistencies highlight the model's uncertainty in its decisions. Provided with the paired images x_{pre} and x_{post} , along with their augmented counterparts x'_{pre} and x'_{post} , we adaptively update the predicted binary change maps \mathcal{M}_0 from the source model f_0 using the confidence metrics obtained from these original and augmented paired images. Here, the target model f predicts the probability p^{ij} of the change for each pixel in the change maps $\hat{\mathcal{M}}$ as same as the source model f_0 .

$$\hat{\mathcal{M}}^{ij} = f(x_{pre}^{ij}, x_{post}^{ij}) \quad (9)$$

Similarly, we obtain probability maps from another augmentation, denoted as $\hat{\mathcal{M}}' = f(x'_{pre}, x'_{post})$. Using these maps, we then compute mean $\hat{\mathcal{U}}$ and standard deviation \mathcal{S} from $\hat{\mathcal{M}}$ and $\hat{\mathcal{M}}'$ to quantify confidence based on pixel-wise variability. As in Eq. (4), the averaged probability maps $\hat{\mathcal{U}}$ are converted into the binary change maps \mathcal{U} .

Our refinement operates under a strategy using $\mu^{ij} \in \mathcal{U}$ and $\sigma^{ij} \in \mathcal{S}$. If variation σ^{ij} is marginal (i.e., it has low uncertainty), we use the value m_0^{ij} in the binary change map \mathcal{M}_0 from the source model f_0 (see Eq. (4)) to its corresponding value μ^{ij} in the binary change map \mathcal{U} from the target model f . Conversely, if variation σ^{ij} is high (i.e., it has high uncertainty), we use the original value m_0^{ij} of the binary change map \mathcal{M}_0 . In each iteration, the updated binary change maps \mathcal{M}_C are generated from \mathcal{M}_0 in Eq. (8), guided by the evolving \mathcal{U} . This process ensures they progressively reflect the target knowledge as follows:

$$\mathcal{M}_C^{ij} = \begin{cases} \mu^{ij} & \text{if } \sigma^{ij} < \tau_r \text{ and } \mu^{ij} = 1 \\ m_0^{ij} & \text{otherwise} \end{cases} \quad (10)$$

where τ_r is the threshold for pixel update decisions, set to 0.001. Using these updated maps, we can derive the refined fine-grained pseudo labels.

$$\mathcal{M}_F^{ij} = \max_{ij}(\mathcal{M}_C^{ij}, \mathcal{M}_V^{ij}) \quad (11)$$

Coarse-grained Pseudo Labels. We employ the coarse-grained pseudo labels generated from the binary change maps \mathcal{M}_0 of the source model f_0 for the image-level refinement. These labels provide image-level information indicating the absence (0) or presence (1) of changes in the paired images. We use this information to guide the model more effectively by masking the generated fine-grained pseudo labels \mathcal{M}_F in instances where the binary change maps \mathcal{M}_0 show no changes. This refinement process helps suppress false positives (FP) in the fine-grained pseudo labels rather than relying on the unmasked fine-grained pseudo labels. We can formulate the coarse-grained pseudo labels q as follows:

$$q = \begin{cases} 1 & \text{if } \sum_{i,j} m_0^{ij} > 0 \\ 0 & \text{otherwise} \end{cases} \quad (12)$$

Using these labels, we mask the fine-grained pseudo labels $\hat{\mathcal{M}}_F$.

$$\mathcal{M}_F^* = q \cdot \mathcal{M}_F, \quad m_f^{ij} \in \mathcal{M}_F^* \quad (13)$$

where \mathcal{M}_F^* is the final fine-grained pseudo labels after the image-level refinement. Using these fine-grained pseudo

labels, we proceed to train the target model f as follows:

$$CE(\mathcal{M}_F^*, \hat{\mathcal{M}}) = - \sum_i \sum_j m_f^{ij} \log(p^{ij}) \quad (14)$$

4 Experiments

4.1 Experimental Setup

Datasets. Our evaluation utilizes the well-known change detection benchmark dataset, xBD. This dataset is widely used for change detection problems and building damage assessment, encompassing different regions affected by 11 various disasters, including wildfires, hurricanes, floods, and earthquakes. For training the source model, we specifically selected data from wildfires in Woolsey, California, USA, and floods in Monsoon, Nepal. Our analysis targets wildfires in Santa Rosa, California, USA, hurricanes in Texas, USA, and tsunamis in Sulawesi, Indonesia. To ensure the model's broader applicability for disaster assessment across diverse terrains, we selected source and target data that differ in both geographical locations and disaster types.

Ground-truth Labels. Due to the deviation in disaster level information across different types of disasters, we convert the four-category disaster level information (0: *no damage*, 1: *minor damage*, 2: *major damage*, 3: *destroyed*) into the binary information (0: *no damage*, 1: *damage*) at each pixel within ground-truth labels.

Baselines. Given that ground-truth labels in target regions are not accessible, we employ four traditional unsupervised change detection baselines — CVA (Bruzzone and Prieto 2000), IRMAD (Nielsen 2007), PCAKmeans (Celik 2009), and ISFA (Wu, Du, and Zhang 2013) — as well as two deep learning-based unsupervised change detection baselines, DCVA (Saha, Bovolo, and Bruzzone 2019) and SCCN (Zhao et al. 2017). We also adopt three relevant domain adaptation baselines — TENT (Wang et al. 2020), SHOT (Liang, Hu, and Feng 2020), and CoTTA (Wang et al. 2022) — and fine-tune them on the same backbone, following the same pre-training process.

Implementation Details. Our experiments utilize SNUNet-CD (Fang et al. 2021) as a backbone for both the source model f_0 and the target model f , which predict the change maps $\hat{\mathcal{M}}_0$ and $\hat{\mathcal{M}}$, respectively. Additionally, BIT-CD (Chen, Qi, and Shi 2021) is employed as another backbone for both models. For all experiments, the batch size is 8. On SNUNet-CD, the source and target model learning rates are $1e-3$ and $1e-5$, respectively, while on BIT-CD, they are $1e-2$ and $1e-4$, both with a weight decay of 0.01. A StepLR scheduler (step size 8, gamma 0.5) is used. All models are trained for 50 epochs using the AdamW optimizer. To regulate the degree of adaptation across experiments, λ is set to 0.1 in Eq. (2).

Performance Evaluation. Our proposed model is compared with various baselines, including traditional and deep learning based unsupervised methods, as well as domain adaptation techniques, using the following common assessment metrics: macro averages of precision, recall, F1-score, and accuracy.

Methods	Backbones	Wildfires (California, USA)				Hurricanes (Texas, USA)				Tsunamis (Sulawesi, Indonesia)			
		Precision	Recall	F1-score	Accuracy	Precision	Recall	F1-score	Accuracy	Precision	Recall	F1-score	Accuracy
Oracle	SNUNet-CD	0.8476	0.7733	0.8087	0.9939	0.8075	0.6190	0.7008	0.9767	0.7402	0.5025	0.5986	0.9878
	BIT-CD	0.8015	0.8272	0.8142	0.9938	0.8162	0.7075	0.7580	0.9801	0.7524	0.6436	0.6937	0.9897
Source	SNUNet-CD	0.8110	0.5119	0.6276	0.9899	0.0466	0.9338	0.0888	0.1542	0.4546	0.0272	0.0513	0.9818
Source†	SNUNet-CD	0.6566	0.8109	0.7256	0.9899	0.6053	0.0414	0.0775	0.9565	0.5200	0.0191	0.0369	0.9819
	BIT-CD	0.8227	0.5063	0.6268	0.9900	0.7317	0.0168	0.0328	0.9563	0.7046	0.0834	0.1491	0.9828
CVA	-	0.0343	0.4583	0.0638	0.7774	0.0527	0.2544	0.0873	0.7653	0.0304	0.3450	0.0559	0.7888
IRMAD	-	0.0290	0.4885	0.0547	0.7090	0.0436	0.2688	0.0750	0.6948	0.0407	0.4293	0.0743	0.7908
PCAKmeans	-	0.0322	0.5501	0.0609	0.7191	0.0306	0.2408	0.0543	0.6299	0.0197	0.3207	0.0371	0.6986
ISFA	-	0.0336	0.4603	0.0626	0.7563	0.0471	0.2163	0.0773	0.7569	0.0384	0.4194	0.0703	0.7855
DCVA	CNN	0.0320	0.5552	0.0605	0.7028	0.0430	0.2643	0.0740	0.6970	0.0023	0.8394	0.0045	0.6994
SCCN	CNN	0.0383	0.4821	0.0709	0.7772	0.0500	0.2317	0.0822	0.7551	0.0030	0.7444	0.0061	0.7853
TENT	SNUNet-CD	0.8810	0.1864	0.3077	0.9861	0.5994	0.0053	0.0104	0.9559	0.4449	0.0180	0.0346	0.9818
SHOT	SNUNet-CD	0.6824	0.8025	0.7376	0.9906	0.6057	0.0413	0.0774	0.9565	0.5218	0.0191	0.0369	0.9819
CoTTA	SNUNet-CD	0.4432	0.2772	0.3209	0.9897	0.3568	0.1006	0.1135	0.9562	0.2232	0.0355	0.0488	0.9817
DAVI	SNUNet-CD	0.7206	0.7767	0.7476	0.9913	0.5211	0.3552	0.4225	0.9571	0.5640	0.3389	0.4234	0.9833
	BIT-CD	0.7075	0.6715	0.6890	0.9900	0.5968	0.5925	0.5947	0.9643	0.3701	0.4187	0.3929	0.9766

† This is the fine-tuned source model that was trained using binary change maps from the pre-trained source model as fine-grained pseudo labels.

Table 1: Performance of DAVI, other unsupervised change detection, and domain adaptation baselines, tested across different geographic regions (North America and Asia) and disaster types (wildfires, hurricanes, and tsunamis). Top results are bolded.

4.2 Comparison with Baselines

Table 1 compares DAVI with other relevant baselines. Traditional methods (CVA, IRMAD, PCAKmeans, and ISFA) and deep learning-based methods (DCVA and SCCN) perform poorly overall compared to the proposed method due to their lack of disaster-specific knowledge. Although these methods are applied directly to the target data without relying on a source model trained on change detection tasks, and thus are not impacted by domain differences, their effectiveness is significantly hindered by the absence of specific knowledge. Secondly, the baselines (Source and Source†) that directly utilize the source model benefit from the disaster-related knowledge in the source data, similar to our method, which aids in analyzing the target data. However, their performance declines as domain discrepancies between the source and target increases. For example, while the performance decrease for wildfires is minimal, the performance for other disasters is significantly affected. Thirdly, only a few results from domain adaptation methods (TENT, SHOT, and CoTTA) surpass other baselines, despite these approaches containing mechanisms to address domain discrepancies. Specifically, as the visual characteristics of a disaster decline further from the source data, performance tends to decline, indicating that adapting visual style shifts based on the source model solely is not enough.

The proposed method, addressing both visual style shifts and decision boundary shifts, consistently achieves superior results across all scenarios. This is particularly evident for disaster types not present in the source data, such as hurricanes and tsunamis, where image segmentation demonstrates highly benefits due to their distinct semantic characteristics. We also confirm the robustness of our method against variations in data quantity and imbalance.

4.3 Component Analysis

Ablation Study. We validate the effectiveness of each component of the proposed method through an ablation study

Component	Wildfires (California, USA)			
	Precision	Recall	F1-score	Accuracy
①: Source Only	0.6824	0.8025	0.7376	0.9906
②: ① w. coarse-grained	0.7964	0.6903	0.7396	0.9920
③: Diff-SAM Only	0.7177	0.3486	0.4693	0.9870
④: ③ w. coarse-grained	0.7504	0.3365	0.4647	0.9872
⑤: Source + Diff-SAM	0.8068	0.6860	0.7415	0.9921
⑥: ⑤ w. coarse-grained	0.7260	0.7679	0.7463	0.9914
⑦: ⑤ w. Refinement	0.7082	0.7819	0.7432	0.9911
⑧: ⑥ w. Refinement (DAVI)	0.7206	0.7767	0.7476	0.9913

Table 2: Ablation study result on key components.

in Table 2. The table shows the degree to which the critical components of the model. We make several observations on the results. When the domain difference between the source and target is minimal, using the source model’s outputs (based on Eq. (4)) is more effective than relying solely on SAM’s image segmentation (based on Eq. (7)); this comparison is evident in lines denoted in ① and ③. Nevertheless, line ⑤ presents that integrating image segmentation with the source model’s outputs (based on Eq. (8)) improves performance. This finding suggests that utilizing image segmentation knowledge in unexpected regions could significantly enhance disaster assessment. Furthermore, the results demonstrate that using coarse-grained pseudo labels mostly leads to better outcomes in terms of F1-scores, as evident in lines (②, ⑥, and ⑧) compared to cases without them (①, ⑤, and ⑦). We can confirm that masking effectively reduces false positives (FP) in the fine-grained pseudo labels.

Hyperparameter Analysis. The optimal threshold, τ_v in Eq. (7) for generating the binary change maps from the confidence difference maps produced by SAM, is reported in Figure 4. Our results confirm that optimal performance is achieved by selecting the threshold based on F1-score criteria, comparing binary change maps from the source model and SAM across various thresholds.




Figure 3: Step-by-step visualization of DAVI. From left to right, images represent pre- and post-disasters, corresponding confidence maps, confidence difference maps discretized by the optimal threshold, binary change maps from the source model, fine-grained pseudo labels without refinement, our pseudo labels, and ground-truth labels. Structures accurately identified through image segmentation appear with a cyan box, while those recognized through refinement appear with a yellow box.




Figure 4: Hyperparameter analysis of τ_v . Solid line shows F1-score of DAVI, while dashed line shows F1-score between the source model's and SAM's binary change maps across thresholds.

4.4 Qualitative Analysis

In addition to the quantitative results, we also visually inspect the impact of each component within DAVI for the case of wildfires. Figure 3 illustrates examples where our method benefits from the source model and SAM, refining the pseudo labels based on the target model outputs. The cyan boxes highlight areas where SAM has positively influenced detection by identifying regions missed by the source model. The yellow boxes indicate areas improved through the refinement process, either by correcting recognized buildings (b) or identifying new structures (c).

4.5 Real-World Case Study

We present representative cases and their evaluation results using the 2023 Türkiye earthquake dataset, freely provided by Maxar, to demonstrate our method's applicability to recent real-world disasters, as shown in Figure 5. The dataset includes 1,648 pairs of pre- and post-disaster images.

Despite differences in geography and disaster types between the source data and the Türkiye dataset, DAVI successfully identifies disaster-affected areas with greater accuracy than other baselines, while minimizing false identifications of undamaged areas. For quantitative evaluation, we conducted human annotations on a randomly selected subset




Figure 5: Evaluation on the 2023 Türkiye earthquake. Columns show pre- and post-disaster images, predictions from baselines and our method. The top row shows damaged cases, while the bottom row shows undamaged cases.

of images (10% of the total pairs). Our approach achieved an F1-score of 0.68, significantly outperforming the baselines (0.36 for Source† and 0.42 for SHOT). These results underscore the model's effectiveness in rapid assessment for real-world disaster scenarios where obtaining ground-truth labels for target regions is not feasible. Application to the Türkiye scenario highlights the potential of our method to provide timely and accurate disaster response insights.

5 Conclusion

This paper introduces a test-time adaptation strategy that leverages a vision foundation model to detect fine-grained structural disaster damages. Experimental results on a real-world disaster benchmark and a newly collected disaster dataset demonstrate significant improvements in damage detection compared to other baselines including unsupervised and domain adaptation methods, particularly to unseen regions. To contribute to social good, we plan to make our model and implementation codes available to the research community and NGOs for humanitarian purposes. Future work will involve testing the approach in recently occurring disasters and enhancing the method by adaptively determining the extent to which the source model should be utilized, based on the differences between the source and target data.

Acknowledgments

This research was partly supported by the Institute for Basic Science (IBS-R029-C2). Sangyoon thanks the HKUST Institute of Emerging Market Studies with support from EY (IEMS24HS02).

References

- Bruzzone, L.; and Prieto, D. F. 2000. Automatic analysis of the difference image for unsupervised change detection. *IEEE Transactions on Geoscience and Remote Sensing*, 38(3): 1171–1182.
- Celik, T. 2009. Unsupervised change detection in satellite images using principal component analysis and k -means clustering. *IEEE Geoscience and Remote Sensing Letters*, 6(4): 772–776.
- Chen, H.; Qi, Z.; and Shi, Z. 2021. Remote sensing image change detection with transformers. *IEEE Transactions on Geoscience and Remote Sensing*, 60: 1–14.
- Chen, K.; Liu, C.; Chen, H.; Zhang, H.; Li, W.; Zou, Z.; and Shi, Z. 2024. RSPrompter: Learning to prompt for remote sensing instance segmentation based on visual foundation model. *IEEE Transactions on Geoscience and Remote Sensing*.
- Chen, W.; Lin, L.; Yang, S.; Xie, D.; Pu, S.; and Zhuang, Y. 2022. Self-supervised noisy label learning for source-free unsupervised domain adaptation. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 10185–10192. IEEE.
- de Jong, K. L.; and Bosman, A. S. 2019. Unsupervised change detection in satellite images using convolutional neural networks. In *2019 International Joint Conference on Neural Networks*, 1–8. IEEE.
- Fang, S.; Li, K.; Shao, J.; and Li, Z. 2021. SNUNet-CD: A densely connected Siamese network for change detection of VHR images. *IEEE Geoscience and Remote Sensing Letters*, 19: 1–5.
- Fiedrich, F.; Gehbauer, F.; and Rickers, U. 2000. Optimized resource allocation for emergency response after earthquake disasters. *Safety science*, 35(1-3): 41–57.
- Field, C. B.; Barros, V.; Stocker, T. F.; and Dahe, Q. 2012. *Managing the risks of extreme events and disasters to advance climate change adaptation: special report of the intergovernmental panel on climate change*. Cambridge University Press.
- Gao, M.; Zhang, Z.; Yu, G.; Arik, S. Ö.; Davis, L. S.; and Pfister, T. 2020. Consistency-based semi-supervised active learning: Towards minimizing labeling cost. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part X 16*, 510–526. Springer.
- Gong, S.; Zhong, Y.; Ma, W.; Li, J.; Wang, Z.; Zhang, J.; Heng, P.-A.; and Dou, Q. 2023. 3dsam-adapter: Holistic adaptation of sam from 2d to 3d for promptable medical image segmentation. arXiv:2306.13465.
- Gupta, R.; Goodman, B.; Patel, N.; Hosfelt, R.; Sajeev, S.; Heim, E.; Doshi, J.; Lucas, K.; Choset, H.; and Gaston, M. 2019. 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*, 10–17.
- Gupta, R.; and Shah, M. 2021. Rescuenet: Joint building segmentation and damage assessment from satellite imagery. In *2020 25th International Conference on Pattern Recognition*, 4405–4411. IEEE.
- Hou, B.; Wang, Y.; and Liu, Q. 2017. Change detection based on deep features and low rank. *IEEE Geoscience and Remote Sensing Letters*, 14(12): 2418–2422.
- Hu, M.; Wu, C.; Du, B.; and Zhang, L. 2023. Binary change guided hyperspectral multiclass change detection. *IEEE Transactions on Image Processing*, 32: 791–806.
- IPCC. 2022. *Climate Change 2022: Impacts, Adaptation, and Vulnerability*. Cambridge University Press.
- Kim, D.; Won, J.; Lee, E.; Park, K. R.; Kim, J.; Park, S.; Yang, H.; and Cha, M. 2022. Disaster assessment using computer vision and satellite imagery: Applications in detecting water-related building damages. *Frontiers in Environmental Science*, 10: 969758.
- 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*, 4015–4026.
- Kuncheva, L. I.; and Faithfull, W. J. 2013. PCA feature extraction for change detection in multidimensional unlabeled data. *IEEE Transactions on Neural Networks and Learning Systems*, 25(1): 69–80.
- Lee, D.-H.; et al. 2013. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. *ICML Workshop on Challenges in Representation Learning*, 3(2): 896.
- Leenstra, M.; Marcos, D.; Bovolo, F.; and Tuia, D. 2021. Self-supervised pre-training enhances change detection in Sentinel-2 imagery. In *Pattern Recognition. ICPR International Workshops and Challenges: Virtual Event, January 10–15, 2021, Proceedings, Part VII*, 578–590. Springer.
- Lei, W.; Wei, X.; Zhang, X.; Li, K.; and Zhang, S. 2023. Medlsam: Localize and segment anything model for 3d medical images. arXiv:2306.14752.
- Li, Q.; Mu, T.; Tuniyazi, A.; Yang, Q.; and Dai, H. 2024. Progressive pseudo-label framework for unsupervised hyperspectral change detection. *International Journal of Applied Earth Observation and Geoinformation*, 127: 103663.
- Liang, J.; Hu, D.; and Feng, J. 2020. Do we really need to access the source data? source hypothesis transfer for unsupervised domain adaptation. In *International Conference on Machine Learning*, 6028–6039. PMLR.
- Liu, C.; Chen, K.; Qi, Z.; Zhang, H.; Zou, Z.; and Shi, Z. 2023. Pixel-level change detection pseudo-label learning for remote sensing change captioning. arXiv:2312.15311.
- Lüddecke, T.; and Ecker, A. 2022. Image segmentation using text and image prompts. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7086–7096.

- Ma, J.; He, Y.; Li, F.; Han, L.; You, C.; and Wang, B. 2024. Segment anything in medical images. *Nature Communications*, 15(1): 654.
- Nielsen, A. A. 2007. The regularized iteratively reweighted MAD method for change detection in multi-and hyperspectral data. *IEEE Transactions on Image processing*, 16(2): 463–478.
- Niu, S.; Wu, J.; Zhang, Y.; Chen, Y.; Zheng, S.; Zhao, P.; and Tan, M. 2022. Efficient test-time model adaptation without forgetting. In *International Conference on Machine Learning*, 16888–16905. PMLR.
- Osco, L. P.; Wu, Q.; de Lemos, E. L.; Gonçalves, W. N.; Ramos, A. P. M.; Li, J.; and Junior, J. M. 2023. The segment anything model (sam) for remote sensing applications: From zero to one shot. *International Journal of Applied Earth Observation and Geoinformation*, 124: 103540.
- Park, S.; Han, S.; Kim, S.; Kim, D.; Park, S.; Hong, S.; and Cha, M. 2021. Improving unsupervised image clustering with robust learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12278–12287.
- Patricia, N.; and Caputo, B. 2014. Learning to learn, from transfer learning to domain adaptation: A unifying perspective. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1442–1449.
- Rajič, F.; Ke, L.; Tai, Y.-W.; Tang, C.-K.; Danelljan, M.; and Yu, F. 2023. Segment anything meets point tracking. arXiv:2307.01197.
- Saha, S.; Bovolo, F.; and Bruzzone, L. 2019. Unsupervised deep change vector analysis for multiple-change detection in VHR images. *IEEE Transactions on Geoscience and Remote Sensing*, 57(6): 3677–3693.
- Shaharabany, T.; Dahan, A.; Giryas, R.; and Wolf, L. 2023. Autosam: Adapting sam to medical images by overloading the prompt encoder. arXiv:2306.06370.
- Shannon, C. E. 1948. A mathematical theory of communication. *The Bell System Technical Journal*, 27(3): 379–423.
- Skifstad, K.; and Jain, R. 1989. Illumination independent change detection for real world image sequences. *Computer Vision, Graphics, and Image Processing*, 46(3): 387–399.
- St-Charles, P.-L.; Bilodeau, G.-A.; and Bergevin, R. 2014. SuBSENSE: A universal change detection method with local adaptive sensitivity. *IEEE Transactions on Image Processing*, 24(1): 359–373.
- Wang, D.; Shelhamer, E.; Liu, S.; Olshausen, B.; and Darrell, T. 2020. Tent: Fully Test-Time Adaptation by Entropy Minimization. In *International Conference on Learning Representations*.
- Wang, Q.; Fink, O.; Van Gool, L.; and Dai, D. 2022. Continual test-time domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7201–7211.
- Wang, W.; Chen, Z.; Chen, X.; Wu, J.; Zhu, X.; Zeng, G.; Luo, P.; Lu, T.; Zhou, J.; Qiao, Y.; et al. 2024. Visionllm: Large language model is also an open-ended decoder for vision-centric tasks. *Advances in Neural Information Processing Systems*, 36.
- Wang, X.; Wang, W.; Cao, Y.; Shen, C.; and Huang, T. 2023a. Images speak in images: A generalist painter for in-context visual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6830–6839.
- Wang, X.; Zhang, X.; Cao, Y.; Wang, W.; Shen, C.; and Huang, T. 2023b. Seggpt: Segmenting everything in context. arXiv:2304.03284.
- Weber, E.; and Kané, H. 2020. Building disaster damage assessment in satellite imagery with multi-temporal fusion. arXiv:2004.05525.
- Wu, C.; Chen, H.; Du, B.; and Zhang, L. 2021a. Unsupervised change detection in multitemporal VHR images based on deep kernel PCA convolutional mapping network. *IEEE Transactions on Cybernetics*, 52(11): 12084–12098.
- Wu, C.; Du, B.; and Zhang, L. 2013. Slow feature analysis for change detection in multispectral imagery. *IEEE Transactions on Geoscience and Remote Sensing*, 52(5): 2858–2874.
- Wu, C.; Zhang, F.; Xia, J.; Xu, Y.; Li, G.; Xie, J.; Du, Z.; and Liu, R. 2021b. Building damage detection using U-Net with attention mechanism from pre-and post-disaster remote sensing datasets. *Remote Sensing*, 13(5): 905.
- Wu, J.; Fu, R.; Fang, H.; Liu, Y.; Wang, Z.; Xu, Y.; Jin, Y.; and Arbel, T. 2023. Medical sam adapter: Adapting segment anything model for medical image segmentation. arXiv:2304.12620.
- Yang, J.; Gao, M.; Li, Z.; Gao, S.; Wang, F.; and Zheng, F. 2023a. Track anything: Segment anything meets videos. arXiv:2304.11968.
- Yang, J.; Tan, W.; Jin, C.; Liu, B.; Fu, J.; Song, R.; and Wang, L. 2023b. Pave the way to grasp anything: Transferring foundation models for universal pick-place robots. arXiv:2306.05716.
- Zhang, M.; and Shi, W. 2020. A feature difference convolutional neural network-based change detection method. *IEEE Transactions on Geoscience and Remote Sensing*, 58(10): 7232–7246.
- Zhao, W.; Wang, Z.; Gong, M.; and Liu, J. 2017. Discriminative feature learning for unsupervised change detection in heterogeneous images based on a coupled neural network. *IEEE Transactions on Geoscience and Remote Sensing*, 55(12): 7066–7080.
- Zou, X.; Yang, J.; Zhang, H.; Li, F.; Li, L.; Wang, J.; Wang, L.; Gao, J.; and Lee, Y. J. 2024. Segment everything everywhere all at once. *Advances in Neural Information Processing Systems*, 36.
- Zou, Y.; Zhang, Z.; Zhang, H.; Li, C.-L.; Bian, X.; Huang, J.-B.; and Pfister, T. 2020. Pseudoseg: Designing pseudo labels for semantic segmentation. arXiv:2010.09713.