

EIDSeg: A Pixel-Level Semantic Segmentation Dataset for Post-Earthquake Damage Assessment from Social Media Images

Huili Huang¹, Chengeng Liu², Danrong Zhang¹, Shail Patel³, Anastasiya Masalava³, Sagar Sadak³, Parisa Babolhavaej³, Weihong Low³, Max Mahdi Roozbahani⁴, J. David Frost²

¹School of Computational Science and Engineering, Georgia Institute of Technology, Atlanta, USA

²School of Civil and Environmental Engineering, Georgia Institute of Technology, Atlanta, USA

³School of Computer Science, Georgia Institute of Technology, Atlanta, USA

⁴School of Computing Instruction, Georgia Institute of Technology, Atlanta, USA

{hhaung413, cliu848, dzhang373, spatel1776, amasalava3, ssadak3, pbabolhavaej3, wlow7, mahdir,jf36}@gatech.edu

Abstract

Rapid post-earthquake damage assessment is crucial for rescue and resource planning. Still, existing remote sensing methods depend on costly aerial images, expert labeling, and produce only binary damage maps for early-stage evaluation. Although ground-level images from social networks provide a valuable source to fill this gap, a large pixel-level annotated dataset for this task is still unavailable. We introduce **EIDSeg**, the first large-scale semantic segmentation dataset specifically for post-earthquake social media imagery. The dataset comprises 3,266 images from nine major earthquakes (2008–2023), annotated across five classes of infrastructure damage: Undamaged Building, Damaged Building, Destroyed Building, Undamaged Road, and Damaged Road. We propose a practical three-phase cross-disciplinary annotation protocol with labeling guidelines that enables consistent segmentation by non-expert annotators, achieving over 70% inter-annotator agreement. We benchmark several state-of-the-art segmentation models, identifying Encoder-only Mask Transformer (EoMT) as the top-performing method with a Mean Intersection over Union (mIoU) of 80.8%. By unlocking social networks’ rich, ground-level perspective, our work paves the way for a faster, finer-grained damage assessment in the post-earthquake scenario.

Code and Datasets —

<https://github.com/HUILIHUANG413/EIDSeg>

Extended version — <https://arxiv.org/abs/2511.06456>

Introduction

Earthquakes cause an estimated \$100 billion in damage globally each year, posing a persistent challenge to the resilience of infrastructure and the response to disasters (GEM Foundation 2023). Rapid post-disaster analysis and response are the keys to rescue efforts, resource deployment, and reducing economic loss in the immediate aftermath of a major earthquake. However, damage assessment efforts often face a “cold start” problem: getting the analysis off the ground is delayed by challenges such as limited image availability, unpredictable weather conditions, and the need for labor intensive expert assessments, amongst others.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Conventional data collection is based on remote sensing imagery such as satellites and drones, combined with expert interpretation to delineate damaged areas and estimate severity (Gupta et al. 2019; Huang et al. 2025). Although widely adopted, these approaches face several well-documented limitations. First, acquiring high-resolution imagery is expensive and is often obstructed by weather conditions like clouds, fog, or rain, particularly in the critical early stages of a disaster (Gupta et al. 2019; Freddi et al. 2021; Munir et al. 2024). Second, damage analysis based on expert annotation is labor-intensive and requires substantial domain knowledge, making the process both time-consuming and resource-demanding. Recent studies have proposed automated tools for early disaster analysis, such as NASA’s Damage Proxy Maps (DPM) (NASA Jet Propulsion Laboratory (ARIA Team) 2020) and Microsoft’s Damage Model (Microsoft AI for Good Research Lab 2023), but they typically return only a binary label that indicates whether an area is damaged, without describing the category of infrastructure or severity level.

In contrast, mobile phones and social media platforms now provide an additional ground-level stream of photographs that are collected within minutes of an event. Social media platforms like X (formerly Twitter) expose street-scale damage invisible to aerial images. Despite this potential, social media imagery remains underutilized due to the lack of task-specific datasets with pixel-level annotations. While datasets such as EID (Huang et al. 2025), MEDIC (Alam et al. 2021a), and DAD (Nguyen et al. 2017a) offer multiclass labels for post-disaster damage, they are limited to image-level classification. Zhang et al. proposed the Damage Semantic Segmentation (DSS) dataset (Zhang et al. 2025), but it only contains 607 post-earthquake social media images without benchmark results. As a result, there are currently no established standards for semantic segmentation of ground-level post-earthquake imagery (Zhu, Liang, and Hauptmann 2020).

Developing a pixel-level segmentation benchmark for post-earthquake social media imagery presents two primary challenges. First, images vary widely in aspect ratio, resolution, and camera angle. Although annotators can typically assign a damage label to the most prominent building in the foreground, thoroughly segmenting all visible in-

frastructures, including distant or partially obstructed ones, requires significant manual effort. Second, neither existing standards in civil engineering for ground-level images nor semantic segmentation guidelines for aerial imagery are well suited to label infrastructure damage in ground-level social media images. Specifically, frameworks such as EMS-98 (Gru 1998) and HAZUS (Federal Emergency Management Agency (FEMA) 2020) evaluate structural damage based on construction type (e.g., masonry, reinforced concrete), which is often unrecognizable in user-generated content. Prior aerial image segmentation datasets such as EarthquakeNet(drone-based) (Jiang et al. 2024) and xBD(satellite-based) (Gupta et al. 2019) offer semantic segmentation annotations; however, their aerial perspectives do not apply directly to ground-level perspectives (Huang et al. 2025). These limitations underscore the need for an annotation protocol and dataset specifically designed for social media imagery, which is both scalable and accessible to non-expert annotators working under high-pressure early-stage post-earthquake scenarios.

To fill this gap, we introduce **EIDSeg**, the first large-scale semantic segmentation dataset of post-earthquake social media images. We curate the dataset by filtering low-quality images from the EID dataset (Huang et al. 2025) and integrating high-quality samples from the DSS dataset (Zhang et al. 2025). To ensure our guideline is clear and accessible, six non-expert annotators with no civil engineering backgrounds labeled the dataset under the supervision of domain experts in post-earthquake engineering and infrastructure assessment. A three-phase quality control process was applied to enhance label consistency and reduce subjectivity.

The final EIDSeg dataset contains **3,266** images from nine major earthquakes (2008–2023), labeled with five pixel-level infrastructure damage classes: Undamaged Building (UD_Building), Damaged Building (D_Building), Destroyed Building (Debris), Undamaged Road (UD_Road), and Damaged Road (D_Road). To the best of our knowledge, EIDSeg is the largest publicly available semantic segmentation benchmark for post-earthquake social media imagery.

The inter-annotator analysis shows that the annotators achieved agreement over 70% in the final stage of the annotation, validating that the annotation guideline that we offer is clear and reproducible. We benchmark several state-of-the-art semantic segmentation models, including DeepLabV3+ (Chen et al. 2018), SegFormer (Xie et al. 2021), Mask2Former (Cheng et al. 2022), BEiT (Bao, Dong, and Wei 2022), OneFormer (Jain et al. 2023), and EoMT (Encoder-only Mask Transformer) (Kerssies et al. 2025). Of these, the EoMT achieves the best performance, with an mIoU of 80.8% and a pixel accuracy of 90.3%.

In summary, our main contributions are as follows.

- **Dataset:** We present **EIDSeg**, a benchmark of 3,266 post-earthquake social media images annotated with five-class pixel-level damage masks for infrastructure damage assessment.
- **Annotation guideline:** We design a three-phase cross-disciplinary annotation protocol that enables consistent labeling by non-expert annotators.

- **Benchmark:** We provide baseline performance using state-of-the-art semantic segmentation models to facilitate future comparisons.

Related Work

Earthquake Infrastructure Damage Assessment Scale

Civil engineering experts rely on drone images and field visits to post-earthquake sites to assess structural damage of infrastructures. A widely used standard, the European Macroseismic Scale (EMS-98) (Gru 1998), defines five damage levels for buildings. G1 (Negligible to slight), G2 (Moderate), G3 (Substantial to heavy), G4 (Very heavy), and G5 (Destruction). The manual includes illustrative sketches and example images to guide human assessors. However, EMS-98 is limited in scope; it focuses solely on reinforced concrete and masonry buildings, excluding other types of infrastructure, such as transportation networks. FEMA’s HAZUS (Federal Emergency Management Agency (FEMA) 2020) framework provides a more comprehensive classification system. It outlines 16 structural damage standards across 36 building categories, covering materials like wood, steel, and concrete. HAZUS also extends its damage evaluation to transportation and utility systems, using four severity levels: slight, moderate, extensive, and complete.

Despite their comprehensiveness, both EMS-98 and HAZUS present challenges when applied to social media image analysis. Many of the required attributes, such as type of construction, load-bearing wall failure, roof integrity, and the presence of hairline cracks, are not consistently visible in the images. Moreover, non-experts often struggle to identify specific infrastructure types based solely on partial or external views. This makes a direct application of detailed engineering taxonomies impractical.

Earthquake Infrastructure Damage Dataset

Table 1 shows an overview of the existing data for the assessment of post-earthquake damage. Most existing datasets for earthquake damage segmentation are either part of broader multi-disaster collections or limited to a single geographic event. For example, the widely used xBD dataset (Gupta et al. 2019) covers 19 disaster events, comprising 22,068 satellite images and 850,736 annotated building polygons. However, it includes only 192 pre-event and 192 post-event satellite images from the 2017 Mexico City earthquake, which is insufficient for large-scale earthquake-specific analysis. More recent segmentation datasets, such as TUE-CD (Liu et al. 2024) and TE23D (Ekkazan and Karşilgil 2025), focus specifically on the 2023 Türkiye earthquakes but are geographically constrained, limiting their generalizability to other seismic regions.

Several datasets focus on the drone imagery. For instance, EarthquakeNet (Jiang et al. 2024) provides high-resolution UAV-based images from Baoxing County with detailed annotations but includes only 69 images. Other datasets, such as AIDER (Kyrkou and Theocharides 2019) and AIDERv2 (Shianios, Kyrkou, and Kolios 2023), include earthquake-related classes such as *collapsed buildings*

Dataset	Data Source	Year	Resolution	Task	# Post-EQ Images	# Damaged Classes
xBD (Gupta et al. 2019)	Satellite	2019	1024×1024	Instance Seg./Cls.	192	4
TUE-CD (Liu et al. 2024)	Satellite	2024	256×256	Semantic Seg.	2338	2
TE23D (Ekkazan and Karshgil 2025)	Satellite	2025	256×256	Semantic Seg.	1,183	2
EarthquakeNet (Jiang et al. 2024)	UAV	2024	5616×3744	Semantic Seg.	69	8
MEDIC (Alam et al. 2021a)	Social Media	2021	Varies	Multi-Task Cls.	15,095	3
EID (Huang et al. 2025)	Social Media	2025	Varies	Cls.	13,513	4
DSS (Zhang et al. 2025)	Social Media	2025	Varies	Semantic Seg.	607	3
EIDSeg(Ours)	Social Media	2025	Varies	Semantic Seg.	3,266	5

Table 1: Overview of publicly available earthquake-damage datasets. “Seg. = Segmentation”, “Cls.” = classification, “Post-EQ” = post-earthquake.

to support the prediction of emergency response. However, these datasets lack damage severity labels, and the source events of the images are not specified, which limits their utility for systematic damage analysis.

For social media imagery, most studies have focused on image-level classification of damage severity (Nguyen et al. 2017a; Alam, Imran, and Ofli 2017; Nguyen et al. 2017b; Nia and Mori 2017; Alam et al. 2021a,b). Huang et al. developed the Earthquake Infrastructure Damage (EID) dataset, offering detailed labeling guidelines tailored for non-expert annotators and achieving strong performance in disaster analysis tasks (Huang et al. 2025). For segmentation, Zhang et al. introduced the DSS dataset, which provides pixel-level segmentation masks tailored for post-earthquake damage assessment (Zhang et al. 2025). However, DSS lacks a formal benchmark and contains only 607 annotated images, limiting its generalization potential.

Data Details

Data Collection and Cleaning

EIDSeg integrates images from both the EID dataset (Huang et al. 2025) and the DSS dataset (Zhang et al. 2025). It comprises images from nine major earthquakes that span 2008 to 2023, enhancing the generalizability of the dataset in diverse earthquake scenarios.

The EID dataset compiles social media images from multiple sources, applying strict quality control and a well-documented annotation protocol. It covers five major events: the 2015 Nepal, 2015 Illapel, 2016 Ecuador, 2017 Mexico, and 2017 Iran–Iraq earthquakes. Each image is categorized into one of four classes: Irrelevant or Non-informative (DS0), No Damage (DS1), Mild Damage (DS2), and Severe Damage (DS3). Compared to earlier image-level datasets such as MEDIC (Alam et al. 2021a), DAD (Nguyen et al. 2017a), and CrisisBenchmark (Alam et al. 2021b), EID provides finer class granularity and higher annotation quality. For EIDSeg, we include only DS1 to DS3 images, since DS0 images are mostly advertisements, maps, or logos that lack relevant infrastructure content.

The DSS dataset adds 607 images drawn from the 2008 Wenchuan, 2010 Haiti, 2015 Nepal, 2023 Morocco, and 2023 Türkiye earthquake, with additional examples retrieved through a Google search for “earthquake damage”.

DSS assigns three segmentation labels: Undamaged Structure, Damaged Structure, and Debris.

Labeling pixel-level masks for disaster damage is substantially more challenging than assigning image-level labels. Although annotators can often identify and categorize damage to prominent foreground structures (e.g., collapsed buildings or visible debris) as *severe*, labeling distant or partially occluded infrastructure, such as roads covered in dust or background buildings, requires finer visual detail and greater annotation effort. To ensure reliable pixel-level annotations, we implemented an automated quality control filter that developed in consultation with domain experts to retain only high-quality, information-rich images from the raw dataset. Specifically, we excluded images that failed the following quality criteria:

- **Sharpness.** We followed standard practice in blur detection by computing the variance of the Laplacian operator on grayscale-converted images using OpenCV. Images with a variance below 50 were flagged as blurry.
- **Resolution.** We enforced a minimum resolution of 512 × 256 pixels to eliminate low-resolution thumbnails and excessively cropped images. The distribution of image resolution are shown in Appendix C.

Images failing either criterion were moved to a separate directory for manual review and excluded from the primary annotation set. This lightweight filtering process ensured that the initial dataset used for labeling maintained a consistent level of visual clarity necessary for accurate and consistent pixel-level labeling. As a result, 6,646 images were left for annotation.

Damage Scale

Referring to previous civil engineering manuals and remote sensing studies (Gru 1998; Federal Emergency Management Agency (FEMA) 2020; Jiang et al. 2024; Zhang et al. 2025), the EIDSeg dataset focuses on two common infrastructure types frequently observed in social media imagery: **buildings** and **roads**.

Building structures refer to enclosed spaces with walls and a roof used for living, working, or storage, such as houses, offices, warehouses, dams, power plants, and similar stationary structures. In post-disaster scenarios, damage may leave these structures to exposed frameworks or partial

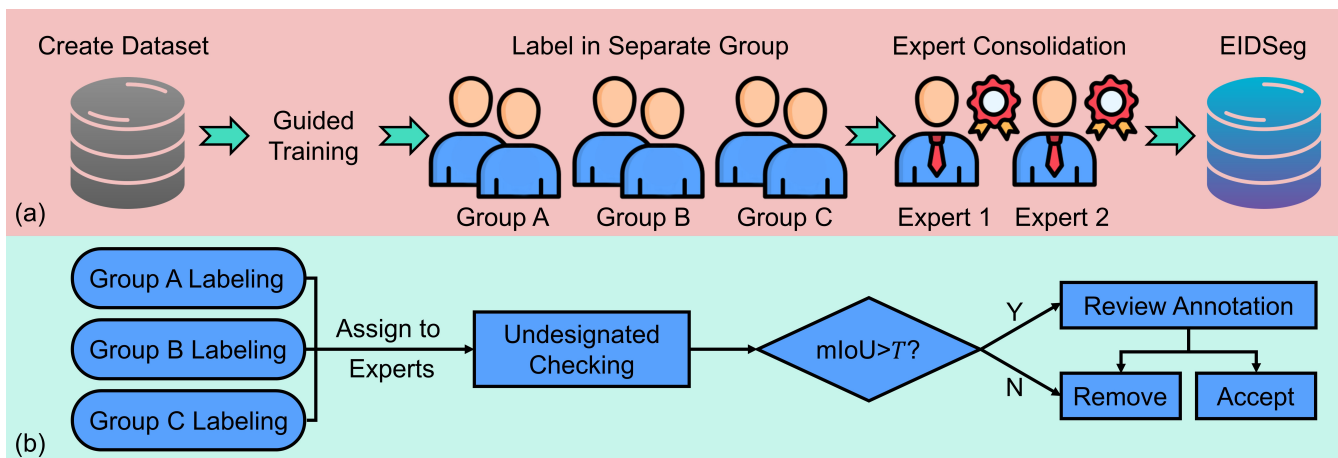


Figure 1: Annotation workflow of EIDSeg. (a): Dataset creation: images are collected and filtered from the EID and DSS datasets, then independently labeled by annotation groups (A–C). Experts evaluate label quality, and high-quality samples are curated into the final EIDSeg dataset. (b): Label evaluation: Annotations undergo experts review following an “Undesignated Checking” step, which address uncertain labels provided by annotators. If inter-annotator agreement (mIoU) exceeds a threshold T , the image proceeds to expert adjudication; otherwise, it is discarded to ensure dataset integrity.

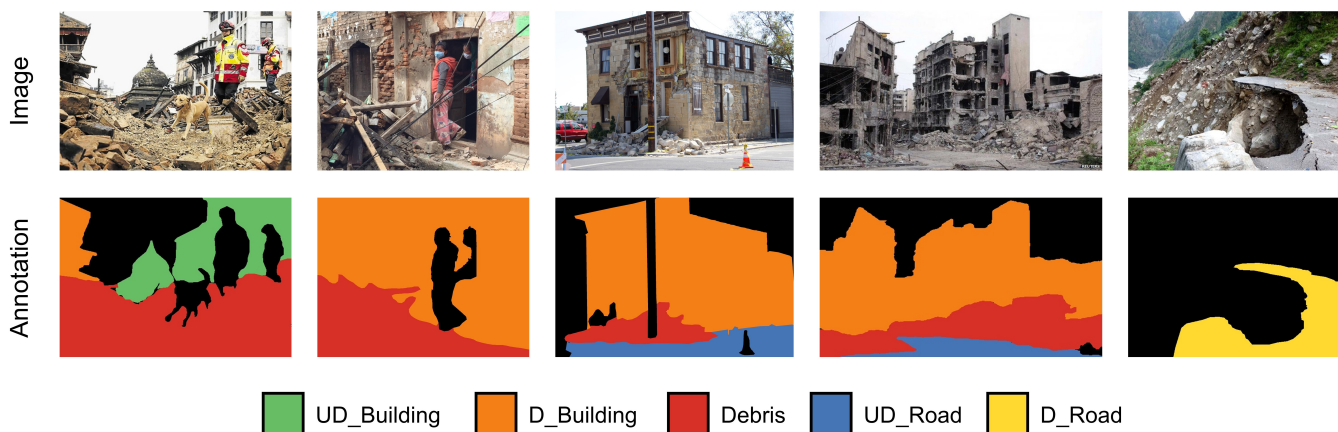


Figure 2: Damage Mask examples of EIDSeg.

ruins. These remnants are still classified as building structures since they remain identifiable.

Road structures are defined as any paths or routes designated for transportation. This includes bridges, highways, streets, sidewalks, and bike paths, with emphasis on linear pathways.

Each instance of the infrastructure is annotated into one of five damage-level classes: Undamaged Building, Damaged Building, Destroyed Building, Undamaged Road and Damaged Road. The detailed definitions of these classes are as follows:

- **Undamaged Building Structure (UD_Building):** Building structures that show no visible signs of disaster-related damage. Note that small cracks may not be visible in social media images due to resolution limits, and such cases are still considered undamaged.
- **Damaged Building Structure (D_Building):** Building

structures that exhibit partial to significant damage that requires repair or restoration. These structures remain identifiable and distinguishable from debris.

- **Destroyed Building Structure (Debris):** Structures that are no longer identifiable due to complete collapse. In cases like pancake collapses, where floors have compressed into layered debris, the entire area is labeled as debris.
- **Undamaged Road Structure (UD_Road):** Roads that appear to be free of visible cracks, potholes, or major surface disruptions. Minor imperfections, such as hairline cracks (e.g., less than 1/8 inch (Federal Emergency Management Agency (FEMA) 2020)), are acceptable if they do not compromise structural integrity or user safety.
- **Damaged Road Structure (D_Road):** Roads showing clear signs of damage, including large or extensive cracks, potholes, or surface deformations that affect func-

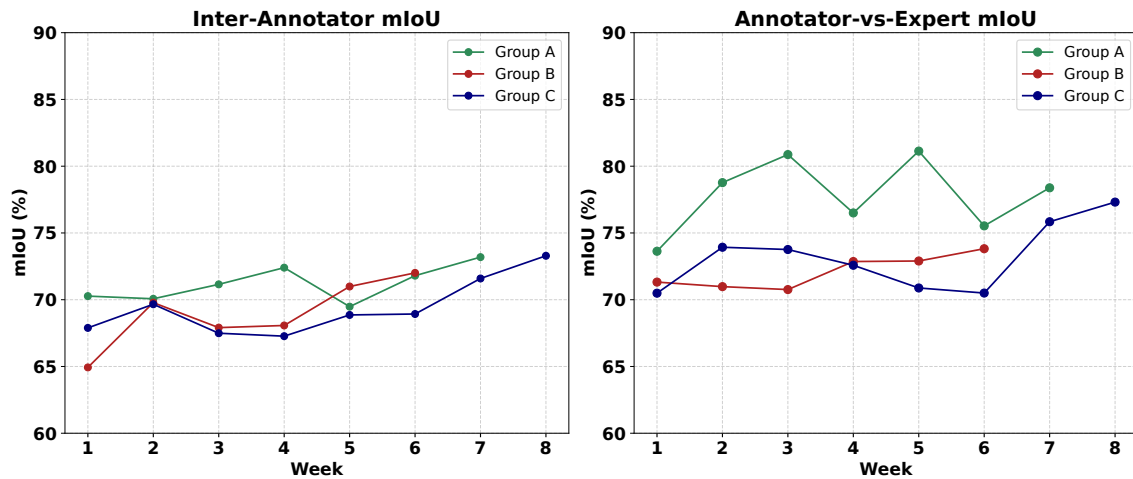


Figure 3: Weekly segmentation agreement across annotators. **Left:** Inter-annotator mIoU (%) among Groups A, B, and C. **Right:** Average annotator–expert mIoU (%) within each group over time. Note that some annotators completed the assigned tasks earlier, which explains the different end dates in the figure (Group A in Week 8, and Group B in Weeks 7–8).

tionality or safety and require repair.

- **Undesignated:** Structures that are visually identifiable as roads or buildings but lack sufficient visual clarity to determine their damage levels with confidence. This uncertainty may result from image blur, occlusion, or limited visibility (e.g., only a portion of the structure is captured in the image).

Note that the EIDSeg damage scale defines only two categories for roads (*undamaged* and *damaged*) to align with what is visually structurally significant in post-disaster imagery. As shown in the rightmost column of Figure 2, roads generally do not undergo complete structural collapse as buildings do, which often results in clearly visible rubble or debris. Instead, road damage often appears as partial deterioration (Jiang et al. 2024). Additionally, the presence of a “debris-like” object on a roadway, such as collapsed structures or landslide material, does not necessarily imply that the road itself is damaged. In many cases, such rubbles originate from surrounding terrain (e.g., hillsides or mountains), rather than from the road infrastructure itself, and are therefore not included in our structural damage categories.

In addition to the structural damage classes, we introduce an *Undesignated* class to reduce annotation ambiguity and ease the labeling burden in challenging cases. To maintain consistency and quality, all *Undesignated* instances are reviewed by domain experts during each annotation round, allowing for a careful and conservative judgment. This class is excluded from the final EIDSeg dataset statistics. For all other annotation procedures, we follow the VOC2011 segmentation guidelines (Everingham et al. 2010).

Annotation Process

Figure 1 illustrates the overall pipeline for the data annotation process. We generated the EIDSeg mask through a three-phase annotation protocol designed with expert quality control. Six non-expert annotators with no civil engineering

backgrounds carried out the labeling under the supervision of two domain experts with expertise spanning both civil engineering and computer science.

Phase 1: Guided training. For the first three weeks, all six annotators underwent training using a 300-image practice set that had been thoroughly labeled by the experts. Each annotator worked independently and their results were evaluated weekly. We computed the mIoU between the masks of each annotator and the expert reference masks to assess progress and identify potential misunderstandings of the labeling guideline. This process allowed us to monitor progress and uncover any misinterpretations of the labeling guidelines. Based on these evaluations, the experts also refined the guidelines to improve clarity. This iterative feedback loop helped minimize misunderstandings and ensured that annotators could accurately identify infrastructure types and assess the severity of the damage before moving onto the full data set.

Phase 2: Production labeling with embedded gold images. Following the guided training phase, the six annotators were divided into three independent pairs (Groups A, B, and C) for full-scale labeling. Each week constituted one annotation round. In each round, the two annotators in a pair first labeled their assigned images independently. The experts then conducted review sessions with each group to address disagreements and clarify ambiguous cases. Annotators revised any identified errors, retained confident labels, and resubmitted the batch for final expert adjudication. The examples of annotator misalignment observed during Phase 2 are shown in Appendix A.

To ensure annotation quality, the experts prepared a curated set of 100 “gold standard” images. For each annotation round, they randomly selected 20–25 gold images and embedded them in the labeling batches assigned to each group. Two agreement metrics were tracked: (i) annotator-vs-expert mIoU (on gold standard images), and (ii) inter-annotator

mIoU (on jointly labeled images). Annotators with expert agreement below 0.60 were flagged for retraining. Images with inter-annotator agreement less than 0.60 were excluded from the final dataset to avoid the subjectivity issue (Huang et al. 2025). The 0.60 threshold was selected based on the highest annotator-expert agreement observed during Phase 1 training. This ensured that only consistent annotations were carried forward to Phase 3.

Phase 3: Expert consolidation. In the final phase, all images containing masks labeled as *Undesignated* underwent the Undesignated Checking process, during which experts manually reviewed the uncertain labels and removed images with excessive ambiguity. For example, if a building labeled as *Undesignated* remained indeterminate even after expert review, and the proportion of *Undesignated* pixels exceeded 2% of the image, the image was excluded from the dataset. Following this, experts conducted a comprehensive review of all remaining annotations, applying minor boundary refinements to improve precision and ensure consistency.

This multistage process ensures that each image in EID-Seg has been reviewed at least twice, first by trained annotators and then by expert reviewers, resulting in a high-quality, pixel-level segmentation benchmark with strong reliability.

Targeted Dataset Statistics

Figure 2 presents several representative examples from the EIDSeg dataset, along with their corresponding color-coded annotation masks. Figure 3 illustrates the average annotator-vs-expert mIoU and inter-annotator mIoU for each group over time during Phase 2. The agreement remained consistently high throughout the annotation process and exhibited an upward trend, indicating improved annotator consistency and better alignment with the expert-labeled ground truth over time. By the final round, all groups achieved agreement above 70%, a promising outcome for disaster assessment tasks. For comparison, the image-level EID dataset reached 86.6% agreement in its final round (Huang et al. 2025), while EIDSeg demands more complex pixel-level labeling across finer-grained classes.

Figure 4 displays the pixel distribution in EIDSeg. In particular, the dataset contains a higher proportion of pixels labeled damaged compared to undamaged. This trend reflects the underlying class distribution of the EID image corpus, which contains more disaster-affected scenes (52.1%) than undamaged ones (26.86%) (Huang et al. 2025). This suggests that social media users are more likely to capture and share images of visible damage than undisturbed scenes after an earthquake.

Experiment

Experimental Setup

Dataset Preparation: We follow the data splitting strategy adopted in the xBD paper (Gupta et al. 2019) and divide the EIDSeg dataset into 80% for training, 10% for validation and 10% for testing.

Models and Training: We evaluated eight state-of-the-art semantic segmentation models: DeepLabV3+(ResNet-101) (Chen et al. 2018), SegFormer (B5) (Xie et al. 2021),

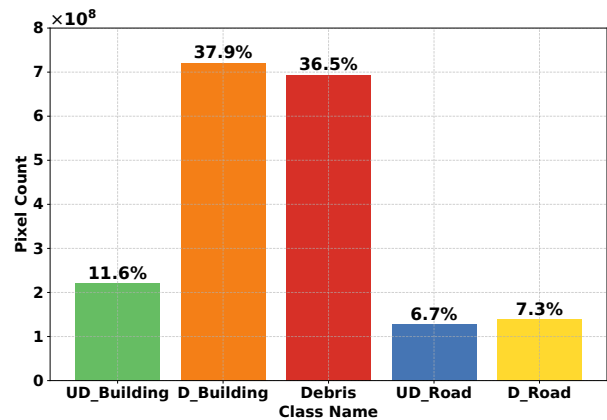


Figure 4: Pixel distribution across segmentation classes in the EIDSeg dataset. Each bar shows the total pixel count per class, along with the percentage it contributes to the full dataset. Class colors are chosen to reflect semantic segmentation label: *Undamaged Building*, *Damaged Building*, *Debris*; and *Undamaged Road*, *Damaged Road*.

Mask2Former (Swin-Base and Swin-L) (Cheng et al. 2022), BEiT (ViT-B and ViT-L) (Bao, Dong, and Wei 2022), OneFormer (Swin-L) (Jain et al. 2023), and EoMT (ViT-L) (Kerssies et al. 2025). All models are initialized with Cityscapes-pre-trained weights except for BEiT, which uses the ADE20K-pre-trained checkpoint due to the absence of a Cityscapes version.

To accommodate the architectural constraints of each model, we adopt model-specific input resolutions: 640×640 for BEiT; 1024×1024 for EoMT; and 512×512 for other models. During training, images are augmented with random horizontal flips, color jittering, and random cropping. All models are trained using the AdamW optimizer with a learning rate of 1×10^{-5} , $\beta_1 = 0.9$, $\beta_2 = 0.999$, and a weight decay of 0.01. Pixel-wise cross-entropy loss is used for supervision. The learning rate is scheduled using the ReduceLROnPlateau strategy with patience of 10 epochs and a reduction factor of 0.5. Training runs for up to 100 epochs, with early stopping triggered after 30 epochs of no improvement on the validation set. All experiments are run on a single NVIDIA A100-80GB GPU and a batch size of 4. The same training protocol is used for all models to ensure a fair and consistent comparison.

Evaluation: We use standard evaluation metrics that include mIoU, frequency-weighted intersection over union (FWIoU), pixel accuracy (PA), floating point operations (FLOPs), and the number of model parameters. FLOPs are computed using the `fvcore` library¹, and reported in terms of GFLOPs (i.e., FLOPs × 10⁹).

Benchmark Analysis

As discussed in Related Work, the DSS dataset is the only post-earthquake social media dataset for semantic segmentation. The original DSS dataset reported only one baseline:

¹<https://github.com/facebookresearch/fvcore>

Model	Backbone	Pre-train	Input	mIoU (%)	FWIoU (%)	PA (%)	FLOPs (G)	Params (M)
DeepLabV3 (Chen et al. 2018)	ResNet-101	Cityscapes	512 ²	67.1	68.2	80.6	79.29	58.75
SegFormer (Xie et al. 2021)	MiT-B5	Cityscapes	512 ²	74.4	75.2	86.9	110.16	84.60
Mask2Former (Cheng et al. 2022)	Swin-S	Cityscapes	512 ²	76.1	77.1	87.5	93.21	68.72
Mask2Former	Swin-L	Cityscapes	512 ²	77.4	78.4	88.7	250.54	215.45
BEiT (Bao, Dong, and Wei 2022)	ViT-B	ADE20K	640 ²	77.7	78.1	88.8	890.06	163.26
BEiT	ViT-L	ADE20K	640 ²	78.7	79.6	89.5	1823.53	441.09
OneFormer (Jain et al. 2023)	Swin-L	Cityscapes	512 ²	79.2	78.3	88.2	1042.14	218.77
EoMT (Kerssies et al. 2025)	ViT-L	Cityscapes	1024 ²	80.8	80.9	90.3	1341.85	319.02

Table 2: Semantic Segmentation Benchmark of *EIDSeg*. Best scores are **bold**.

Model	UD_Building	D_Building	Debris	UD_Road	D_Road	mIoU (%)
DeepLabV3+	43.5	65.4	77.3	75.7	73.7	67.1
SegFormer	54.9	73.5	82.3	79.9	81.3	74.4
Mask2Former-S	58.9	76.3	83.4	79.3	82.8	76.1
Mask2Former-L	63.5	76.9	84.3	79.6	82.7	77.4
BEiT-B	63.8	76.6	83.7	80.0	84.3	77.7
BEiT-L	66.4	77.9	85.1	82.9	81.3	78.7
OneFormer	61.0	75.7	83.9	84.1	91.5	79.2
EoMT	70.1	80.0	84.6	82.0	87.3	80.8

Table 3: Class-wise IoU and mIoU (%) for each model on *EIDSeg*. Best scores are **bold**.

SegFormer-B5 trained for three-class segmentation, which achieved an mIoU of 72% (Zhang et al. 2025). In contrast, Table 2 presents a comprehensive evaluation of eight modern architectures on the proposed *EIDSeg* benchmark, which adopts a more fine-grained five-class labeling. Under these conditions, SegFormer-B5 attains an mIoU of **74.4%**, yielding a **2.4%** improvement over the DSS result while addressing a more detailed five-class labeling that includes two additional categories. This gain highlights the advantages of our expanded image corpus and refined annotation strategy, which collectively improve class balance and reduce labeling noise.

EoMT achieves the highest overall performance, with a FWIoU of **80.9%** and pixel accuracy of **90.3%**. Its superior results can be attributed to training in larger input resolutions (1024 × 1024) that preserve spatial detail, as well as recent architectural advances in transformer design. Other transformer-based models, including Mask2Former, BEiT, and OneFormer, exhibit strong performance, demonstrating that *EIDSeg* effectively supports generalization across diverse model families.

Table 3 presents the class-wise IoU for each model. Among all classes, *Undamaged Building* and *Damaged Building* consistently exhibit lower IoU scores. This can be attributed to two main factors. First, during the annotation process for *EIDSeg*, the annotators followed a conservative guideline: if a building appeared structurally intact but exhibited visible cracks not attributable to wear and tear, it was labeled as *Damaged Building* due to lack of information about its interior condition. Although this labeling strategy supports cautious decision-making in post-disaster analysis, it may lead to confusion for models when encountering visually similar undamaged structures. Second, since the dataset

is collected from social media, buildings in the background often appear blurry or lack sufficient detail, making it difficult to assess structural damage. This reflects an inherent trade-off in the use of social media imagery. Although it enables access to timely data from disaster-stricken areas where professionals may not be present, it also limits control over image quality and resolution. Visual examples of misclassified images and the leave-one-event-out results are presented in Appendix B and Appendix D, respectively.

Conclusion

In this work, we introduce *EIDSeg*, the largest pixel-level semantic segmentation dataset for post-earthquake damage analysis based on social media images. The dataset comprises 3,266 images collected from nine major earthquakes that occurred between 2008 and 2023, covering a diverse range of real-world disaster scenarios. Annotations span five semantic classes: Undamaged Building, Damaged Building, Destroyed Building, Undamaged Road, and Damaged Road, allowing a fine-grained evaluation of damage to built infrastructure. Through a carefully designed annotation guideline and expert-led quality control, *EIDSeg* provides a reliable foundation for training vision models capable of automating post-earthquake damage assessment using ground-view imagery.

Acknowledgements

The research described herein was supported in part by the School of Computing Instruction and the Elizabeth and Bill Higginbotham Professorship at Georgia Tech. This support is gratefully acknowledged. Special thanks are due to Mohammad Taher for helping with the annotation. Support for the undergraduate students participating in the project was

provided by the US National Science Foundation through the Geotechnical Engineering Program under Grant No. CMMI-1826118. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of NSF.

References

1998. European Macroseismic Scale 1998 (EMS-98). Technical report, Centre Européen de Géodynamique et de Séismologie, Luxembourg, Luxembourg. Provides standardized macroseismic intensity guidelines for Europe :contentReference[oaicite:1]index=1.
- Alam, F.; Alam, T.; Hasan, M. A.; Hasnat, A.; Imran, M.; and Ofli, F. 2021a. MEDIC: A Multi-Task Learning Dataset for Disaster Image Classification. *arXiv preprint arXiv:2108.12828*.
- Alam, F.; Imran, M.; and Ofli, F. 2017. Image4Act: Online Social Media Image Processing for Disaster Response. *2017 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, 601–604.
- Alam, F.; Sajjad, H.; Imran, M.; and Ofli, F. 2021b. Crisis-Bench: Benchmarking Crisis-related Social Media Datasets for Humanitarian Information Processing. In *Proceedings of the 15th International AAAI Conference on Web and Social Media (ICWSM)*.
- Bao, H.; Dong, L.; and Wei, F. 2022. BEiT: BERT Pre-Training of Image Transformers. In *International Conference on Learning Representations (ICLR)*.
- Chen, L.-C.; Zhu, Y.; Papandreou, G.; Schroff, F.; and Adam, H. 2018. Encoder-Decoder with Atrous Separable Convolution for Semantic Image Segmentation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 801–818.
- Cheng, B.; Misra, I.; Schwing, A. G.; Kirillov, A.; and Girdhar, R. 2022. Masked-Attention Mask Transformer for Universal Image Segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 1290–1299.
- Ekkazan, C.; and Karşlıgil, M. E. 2025. TE23D: A Dataset for Earthquake Damage Assessment and Evaluation. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 18: 3852–3863.
- Everingham, M.; Gool, L. V.; Williams, C. K. I.; Winn, J.; and Zisserman, A. 2010. The PASCAL Visual Object Classes (VOC) Challenge. *International Journal of Computer Vision*, 88(2): 303–338.
- Federal Emergency Management Agency (FEMA). 2020. *HAZUS Earthquake Model Technical Manual*. FEMA, U.S. Department of Homeland Security, Washington, DC. HAZUS 4.2 SP3; provides detailed methodology for earthquake loss estimation.
- Freddi, F.; Galasso, C.; Cremen, G.; Dall’Asta, A.; DiSarno, L.; Giaralis, A.; Gutiérrez-Urzúa, F.; Málaga-Chuquitaype, C.; Mitoulis, S.; Petrone, C.; Sextos, A.; Sousa, L.; Tarbali, K.; Tubaldi, E.; Wardman, J.; and Woo, G. 2021. Innovations in earthquake risk reduction for resilience: Recent advances and challenges. *International Journal of Disaster Risk Reduction*, 60: 102267. Open access under CC BY-NC-ND.
- GEM Foundation. 2023. Assessing Seismic Hazard and Risk Globally for an Earthquake Resilient World. Available at: <https://www.globalquakemodel.org/news/global-earthquake-risk-report-2023>.
- Gupta, R.; Henry, C. S.; Kuo, A. V.; Chandrasekhar, V.; Madinaveitia, A. V.; Adams, J. H.; Hinton, R.; Zink, M.; Burke, D.; Pham, J.; Sadowski, R. S.; and Rapoport, B. L. 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 (CVPRW)*, 10–17.
- Huang, H.; Zhang, D.; Masalava, A.; Roozbahani, M. M.; Roy, N.; and Frost, J. D. 2025. Enhancing the fidelity of social media image data sets in earthquake damage assessment. *Earthquake Spectra*.
- Jain, S.; Girdhar, R.; Misra, I.; Teterwak, P.; Kirillov, A.; Chen, X.; and Gkioxari, G. 2023. OneFormer: One Transformer to Rule Universal Image Segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 16408–16418.
- Jiang, S.; Bian, Y.; Wang, Y.; Li, X.; Liu, Z.; Ren, Y.; and Zhao, Y. 2024. EarthquakeNet: A High-Resolution UAV-Based Dataset for Earthquake Damage Assessment. In *Proceedings of the 2024 IEEE International Conference on Image Processing (ICIP)*, 55–61.
- Kerssies, T.; Cavagnero, N.; Hermans, A.; Norouzi, N.; Averta, G.; Leibe, B.; Dubbelman, G.; and de Geus, D. 2025. Your ViT is Secretly an Image Segmentation Model. *arXiv preprint arXiv:2503.19108*. Accepted to CVPR 2025.
- Kyrkou, C.; and Theocharides, T. 2019. Deep-Learning-Based Aerial Image Classification for Emergency Response Applications Using Unmanned Aerial Vehicles. *arXiv preprint arXiv:1906.08716*.
- Liu, Y.; Zhang, K.; Guan, C.; Zhang, S.; Li, H.; Wan, W.; and Sun, J. 2024. Building Change Detection in Earthquake: A Multiscale Interaction Network With Offset Calibration and a Dataset. *IEEE Transactions on Geoscience and Remote Sensing*, 62: 1–17.
- Microsoft AI for Good Research Lab. 2023. Turkey Building Damage Assessment: AI-Based Satellite Imagery Analysis after the February 6, 2023 Earthquake. Technical report, Microsoft Research, Redmond, WA. Uses high-resolution satellite imagery and AI in partnership with AFAD to estimate building damage over four cities.
- Munir, A.; Siddiqui, A. J.; Anwar, S.; El-Maleh, A.; Khan, A. H.; and Rehman, A. 2024. Impact of Adverse Weather and Image Distortions on Vision-Based UAV Detection: A Performance Evaluation of Deep Learning Models. *Drones*, 8(11).
- NASA Jet Propulsion Laboratory (ARIA Team). 2020. *ARIA Damage Proxy Maps for Rapid Earthquake Damage Assessment*. NASA / JPL – Advanced Rapid Imaging and Analysis (ARIA), Pasadena, CA. Based on SAR-derived coherence change; quantitative validation in O’Donnell (2024).

Nguyen, D.; Ofli, F.; Imran, M.; and Mitra, P. 2017a. Damage assessment from social media imagery data during disasters. In Diesner, J.; Ferrari, E.; and Xu, G., eds., *ASONAM*, Proceedings of the 2017 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining, ASONAM 2017, 569–576. Association for Computing Machinery, Inc. 9th IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining, ASONAM 2017 ; Conference date: 31-07-2017 Through 03-08-2017.

Nguyen, T. D.; Alam, F.; Ofli, F.; and Imran, M. 2017b. Automatic Image Filtering on Social Networks Using Deep Learning and Perceptual Hashing During Crises. *ArXiv*, abs/1704.02602.

Nia, K. R.; and Mori, G. 2017. Building Damage Assessment Using Deep Learning and Ground-Level Image Data. *2017 14th Conference on Computer and Robot Vision (CRV)*, 95–102.

Shianios, D.; Kyrkou, C.; and Kolios, P. S. 2023. A Benchmark and Investigation of Deep-Learning-Based Techniques for Detecting Natural Disasters in Aerial Images. In *Computer Analysis of Images and Patterns*, volume 14185 of *Lecture Notes in Computer Science*, 321–332. Springer, Cham.

Xie, E.; Wang, W.; Yu, Z.; Anandkumar, A.; Alvarez, J. M.; and Luo, P. 2021. SegFormer: Simple and Efficient Design for Semantic Segmentation with Transformers. *arXiv preprint arXiv:2105.15203*.

Zhang, D.; Huang, H.; Smith, N. S.; Roy, N.; and Frost, J. D. 2025. From Pixels to Damage Severity: Estimating Earthquake Impacts Using Semantic Segmentation of Social Media Images. *arXiv preprint arXiv:2507.02781*. Preprint posted July 3, 2025; under review.

Zhu, X.; Liang, J.; and Hauptmann, A. 2020. MSNet: A Multilevel Instance Segmentation Network for Natural Disaster Damage Assessment in Aerial Videos. *arXiv preprint arXiv:2006.16479*. Also presented at WACV 2021.