

BRIGHT: A globally distributed multimodal building damage assessment dataset with very-high-resolution for all-weather disaster response

Hongruixuan Chen^{1,2}, Jian Song^{1,2}, Olivier Dietrich³, Clifford Broni-Bediako², Weihao Xuan^{1,2}, Junjue Wang¹, Xinlei Shao¹, Yimin Wei¹, Junshi Xia², Cuiling Lan⁴, Konrad Schindler³, and Naoto Yokoya^{1,2}

¹Graduate School of Frontier Sciences, The University of Tokyo, Chiba, Japan

²RIKEN Center for Advanced Intelligence Project (AIP), RIKEN, Tokyo, Japan

³Department of Photogrammetry and Remote Sensing, ETH Zürich, Zürich, Switzerland

⁴Microsoft Research Asia, Beijing, China

Correspondence: Naoto Yokoya (yokoya@k.u-tokyo.ac.jp)

Abstract. Disaster events occur around the world and cause significant damage to human life and property. Earth observation (EO) data enables rapid and comprehensive building damage assessment, an essential capability crucial in the aftermath of a disaster to reduce human casualties and inform disaster relief efforts. Recent research focuses on developing artificial intelligence (AI) models to accurately map unseen disaster events, mostly using optical EO data. These solutions based on optical data are limited to clear skies and daylight hours, preventing a prompt response to disasters. Integrating multimodal EO data, particularly combining optical and synthetic aperture radar (SAR) imagery, makes it possible to provide all-weather, day-and-night disaster responses. Despite this potential, the lack of suitable benchmark datasets has constrained the development of robust multimodal AI models. In this paper, we present a Building damage assessment dataset using veRy-hIGH-resoluTion optical and SAR imagery (BRIGHT) to support AI-based all-weather disaster response. To the best of our knowledge, BRIGHT is the first open-access, globally distributed, event-diverse multimodal dataset specifically curated to support AI-based disaster response. It covers five types of natural disasters and two types of human-made disasters across 14 regions worldwide, focusing on developing countries where external assistance is most needed. The dataset’s optical and SAR images with spatial resolutions between 0.3 and 1 meters provide detailed representations of individual buildings, making it ideal for precise damage assessment. We train seven advanced AI models on BRIGHT to validate transferability and robustness. Beyond that, it also serves as a challenging benchmark for a variety of tasks in real-world disaster scenarios, including unsupervised domain adaptation, semi-supervised learning, unsupervised multimodal change detection, and unsupervised multimodal image matching. The experimental results serve as baselines to inspire future research and model development. The dataset (Chen et al., 2025), along with the code and pretrained models, is available at <https://github.com/ChenHongruixuan/BRIGHT> and will be updated as and when a new disaster data is available. BRIGHT also serves as the official dataset for the 2025 IEEE GRSS Data Fusion Contest Track II. We hope that this effort will promote the development of AI-driven methods in support of people in disaster-affected areas.

1 Introduction

A disaster is defined as a severe disruption in the functioning of a community or society due to the interaction between a hazard event and the conditions of exposure, vulnerability and capacity resulting in human, material, economic or environmental losses and impacts (Ge et al., 2020). According to the United Nations Office for Disaster Risk Reduction (UNDRR), between 1998 and 2017, natural disasters such as earthquakes, storms, and floods affected approximately 4.4 billion people and caused 1.3 million deaths. These disasters have also resulted in economic losses of 2,647 billion United States dollars (USD) in disaster-affected countries (UNDRR, 2018a). The threat of disasters is likely to increase due to global urbanization (Kreibich et al., 2022; Bastos Moroz and Thieken, 2024). Rapid and comprehensive damage assessment is crucial in the aftermath of a disaster to make informed and effective rescue decisions that minimize losses and impacts. Building damage assessment aims to provide information, including the area and amount of damage, the rate of collapsed buildings, and the type of damage each building has incurred. This information is critical in the early stages of a disaster, as the distribution of damaged buildings is closely related to life-saving efforts in an emergency response (Xie et al., 2016; Adriano et al., 2021). Conducting field surveys after a disaster can be difficult and dangerous, especially when transportation and communication systems are disrupted, making efficient on-site assessments challenging. Earth observation (EO) provides a safe and efficient way to obtain information on building damage in disaster areas due to its wide field of view and contactless operation.

The EO technologies commonly used for assessing building damage after disasters are optical and synthetic aperture radar (SAR). Optical imagery is a primary source for building damage assessment because of its intuitive and easy-to-interpret nature. For example, moderate-resolution optical data from the Landsat series and Sentinel-2 have been used to assess building damage (Yusuf et al., 2001; Fan et al., 2019; Sandhini Putri et al., 2022). Landsat and Sentinel-2 data are limited in spatial resolution and only provide broad approximations of affected areas, which lack precision for specific buildings, crucial for timely rescue. The new generation of very high-resolution (VHR) optical sensors, such as IKONOS and WorldView, provides EO data with spatial resolutions of a meter or less, enabling finer assessments at the level of individual buildings (Freire et al., 2014). These data have been used successfully in building damage assessment (Yamazaki and Matsuoka, 2007; Tong et al., 2012; Freire et al., 2014).

While accurate building damage maps can be obtained by visual interpretation of optical images by human experts, this process is time-consuming and labor-intensive for large-scale rapid assessments. In addition, it requires trained professionals. Therefore, recent studies have focused on developing automated methods for rapid building damage mapping (Tong et al., 2012; Xie et al., 2016; Gupta et al., 2019; Zheng et al., 2021). Among these, machine learning (ML) and deep learning (DL) techniques have significantly improved efficiency and accuracy in building damage assessment. Earlier work focused on a single disaster event with labels annotated for a specific disaster area to train a model. This model is then used to generate building damage maps for the same event (Xie et al., 2016; Xia et al., 2023). However, since training data were limited to a few building types, damage patterns, and background land cover distributions, the resulting models mostly lack generalizability and struggle to produce accurate building damage maps for new disaster events, which limits their practical use. Recent large-scale benchmark datasets, for example, the xBD dataset (Gupta et al., 2019) containing different types of disaster scenarios and



Figure 1. An example of the wildfire occurring in Maui, Hawaii, USA, August 2023. (a) Pre-event optical imagery (© Maxar). (b) Post-event optical image (© Maxar) with land-cover features obscured by wildfire smoke. (c) Post-event SAR imagery (© Capella Space) unaffected by smoke, showing the disaster area.

damages, have made it possible to adopt DL models to quickly and accurately map building damages after a newly occurred, previously unseen disaster (Zheng et al., 2021; Chen et al., 2022a; Shen et al., 2022; Kaur et al., 2023; Guo et al., 2024; Wang et al., 2024; Chen et al., 2024). For example, Zheng et al. (2021) trained DL models on the xBD dataset and applied them to map the damage to buildings in two unseen human-made disaster events. These studies have demonstrated the effectiveness of DL models for building damage mapping.

The optical EO technology uses a passive sensing technique, which requires solar illumination and cloud-free weather conditions. This severely limits the application of optical images in an emergency tool for all-weather disaster response (Adriano et al., 2021). In contrast, SAR sensors employ active illumination with longer microwaves and can acquire images in adverse weather conditions, offering great potential for all-weather disaster response. Most disaster events, especially wildfires, floods, and storms, are often accompanied by less-than-ideal imaging conditions. For example, Figure 1 shows EO imagery captured for a wildfire event that occurred in August 2023 in Hawaii, USA. The post-event optical image shown in Figure 1-(b) does not provide clear surface information due to the effects of the wildfire smoke. However, the SAR image illustrated in Figure 1-(c) is not affected by smoke and clearly shows the buildings damaged by the wildfire.

Due to the advantages of SAR imagery, various SAR-based methods have been proposed for building damage assessment. These methods utilize intensity (Matsuoka and Yamazaki, 2005, 2010; Matsuoka et al., 2010), coherence (Yonezawa and Takeuchi, 2001; Arciniegas et al., 2007; Watanabe et al., 2016; Liu and Yamazaki, 2017), and polarization features (Yamaguchi, 2012; Chen and Sato, 2013; Watanabe et al., 2016; Karimzadeh and Matsuoka, 2017) to assess building damage at a block unit level, depending on the acquisition mode. Several studies have attempted to extend the block-level approach and have explored new approaches at the building instance level using higher spatial resolution sensors such as COSMO-SkyMed and TerraSAR-X (Liu et al., 2013; Brett and Guida, 2013; Chini et al., 2015; Ge et al., 2019). DL-based methods have also been explored

with SAR data to assess building damage (Bai et al., 2018; Adriano et al., 2019; Bai et al., 2017; Li et al., 2023b). However, because of the lack of large-scale benchmark datasets, such as xBD in the optical domain, these methods have focused on local regions and single disaster events, and their ability to generalize to other disaster events remains largely unknown.

The inherent challenges of SAR data, such as oblique viewing angles, speckle noise, object occlusion, and geometric distortions, complicate the accurate mapping of building damage compared to optical imagery (Adriano et al., 2021; Xia et al., 2025). Furthermore, the limited availability of the VHR SAR data reduces its reliability as a source of pre-event data (Brunner et al., 2010; Adriano et al., 2021). Considering these practical limitations, the most effective strategy for rapid assessment of building damage in all weather could arguably be to combine pre-event optical images, which provide accurate localization and detailed building information in the visible spectrum, with post-event SAR images, which capture structural information as a cue for building damage (Adriano et al., 2019). Previous methods have attempted to align the two modalities with traditional statistical models (Stramondo et al., 2006; Chini et al., 2009; Brunner et al., 2010; Wang and Jin, 2012). These statistical models are sensor-specific and require dedicated modeling for each sensor. DL methods offer a promising solution by automatically learning a high-dimensional feature space that aligns the two modalities. However, to train a DL model, one must have access to a high-quality, large-scale dataset with comprehensive coverage of various disaster events and sufficient geographic diversity. This remains a significant challenge that needs to be addressed.

To support AI-based research aimed at all-weather building damage mapping, we present BRIGHT, the first open and globally distributed multimodal VHR dataset for building damage assessment. Advances in EO technology have enabled data providers like Capella¹ and Umbra² to offer VHR SAR imagery at a sub-meter level resolution per pixel. This allows for detailed building damage assessments at the individual building level, to guide targeted and effective rescue operations as required by emergency responders. Benefiting from the progress made in EO, BRIGHT incorporates both pre-event optical imagery and post-event SAR imagery with spatial resolutions ranging from 0.3 meters to 1 meter per pixel. The types of disaster events considered in BRIGHT are *earthquakes*, *storms* (e.g., *hurricane*, *cyclone*), *wildfires*, *floods*, and *volcanic eruptions*. These natural disasters accounted for 84% of the fatalities and 94% of the economic losses between 1998 and 2017 (UNDRR, 2018a). In addition to natural disasters, the BRIGHT dataset further considers disasters caused by human activity, such as *accidental explosions* and *armed conflicts*, which also pose significant threats to human life and infrastructure and can occur unexpectedly, requiring a rapid response (UNDRR, 2018b; Dietrich et al., 2025). The 14 disaster events cover 23 different regions distributed around the globe, with a focus on developing countries where external assistance is most urgently needed after a disaster. The labels are manually annotated with multi-level annotations that distinguish between damaged buildings and completely destroyed buildings.

1.1 Comparison with existing datasets

The comparison between BRIGHT and existing datasets for building damage assessment is summarized in Table 1. Most current building damage assessment datasets are limited in scale and scope due to the limited availability of disaster events

¹<https://www.capellaspace.com/>

²<https://umbra.space/>

Table 1. Comparison of BRIGHT with the existing building damage assessment datasets. The OA indicates whether the dataset is open access (OA) or not, and GSD is an acronym for ground sampling distance (GSD). Note that since some datasets integrate other datasets, we summarize only the largest one to avoid duplication here. For example, the BDD dataset (Adriano et al., 2021) includes the Tohoku-Earthquake-2011 dataset (Bai et al., 2018) and Palu-Tsunami-2018 dataset (Adriano et al., 2019).

Dataset	OA	Modality	GSD (m/pixel)	No. of events	Disaster type	No. of building	Granularity
ABCD (Fujita et al., 2017)	✓	Optical EO	0.4	1	Tsunami	N/A	Image-level
(Nguyen et al., 2017)	✓	Images on social media	N/A	4	3 natural disasters	N/A	Image-level
(Cheng et al., 2021)	✓	Optical EO	N/A	1	Hurricane	1,802	Image-level
(Xue et al., 2024)	✓	Street-view image	N/A	1	Hurricane	2,468	Image-level
FloodNet (Rahmemonfar et al., 2021)	✓	Optical EO	N/A	1	Flood	6,675	Pixel-level
RescueNet (Rahmemonfar et al., 2023)	✓	Optical EO	N/A	1	Hurricane	10,903	Pixel-level
Ida-BD (Kaur et al., 2023)	×	Optical EO	0.5	1	Hurricane	18,083	Pixel-level
CRASAR-U-DROIDS (Manzini et al., 2024)	✓	Optical EO	0.02-0.12	10	4 natural disasters 1 man-made disaster	21,716	Pixel-level
Noto-BDA-MV (Vescovo et al., 2025)	✓	Optical EO	N/A	1	Earthquake	140,208	Pixel-level
xBD (Gupta et al., 2019)	✓	Optical EO	<0.8	15	6 natural disasters	>700,000	Pixel-level
QQB (Sun et al., 2024b)	✓	Optical and SAR EO	<1	1	Earthquake	4,029	Pixel-level
BDD (Adriano et al., 2021)	×	Optical and SAR EO	1.2-3.3	9	3 natural disasters	123,453	Pixel-level
BRIGHT	✓	Optical and SAR EO	0.3-1	14	5 natural disasters 2 man-made disasters	384,596	Pixel-level

with corresponding open-source EO data and annotation efforts (Rahmemonfar et al., 2021; Gupta and Shah, 2021; Kaur et al., 2023). Because of the high cost and time required for pixel-level labeling, some of the existing datasets provide image-level
110 labeling, indicating only whether an image contains damaged buildings (Fujita et al., 2017; Nguyen et al., 2017; Cheng et al., 2021; Xue et al., 2024). Although these image-level labeling datasets have served the community well, they lack the spatial precision needed to guide specific rescue operations. The xBD dataset (Gupta et al., 2019) is currently the largest open data collection, covering six natural disasters in 15 regions with more than 700,000 building instances. However, the xBD includes only optical EO data. It does not support all-weather disaster response. Sun et al. (2024b) introduced a multimodal dataset, but
115 it is limited to a single disaster event and contains only about 4,000 building instances. The small size makes it challenging to train DL models and limits the transferability of the trained models.

The dataset most similar to BRIGHT is the BDD proposed by Adriano et al. (2021). The main differences between BDD and BRIGHT datasets are: (1) BRIGHT covers more disaster events and building instances, including both natural and human-made disasters. (2) BRIGHT has higher spatial resolution in both optical and SAR images. Whereas the highest resolution of
120 SAR images in BDD is 1.2 meters, BRIGHT provides finer detail with spatial resolutions ranging from 0.3 meters to 1 meter, enabling the detection of subtle structural damage in individual buildings. (3) Perhaps the most important difference is that whereas the re-distribution of BDD is restricted, BRIGHT is an open-source dataset publicly available to the global community. Apart from the datasets listed in Table 1, there are other datasets targeted at monitoring hazardous events related to disasters,

including landslides (Ghorbanzadeh et al., 2022; Meena et al., 2023), floods (Bonafilia et al., 2020; Zhang et al., 2023) and
125 wildfires (Artés et al., 2019; Huot et al., 2022; He et al., 2024), but are not related to building damage assessment.

1.2 Main contribution

The contributions of this paper are threefold:

1. We present BRIGHT, the first multimodal building damage dataset with sub-meter spatial resolution, which is publicly available to the community. BRIGHT employs a combination of pre-event optical imagery and post-event SAR imagery,
130 with various disaster events and rich geographic diversity, to support the study of AI-based multimodal building damage mapping, especially in developing countries.
2. We evaluate a suite of contemporary models on BRIGHT to establish robust baselines. Beyond supervised deep learning, BRIGHT can support a wide range of AI-based methods. It enables research in unsupervised domain adaptation (UDA), semi-supervised learning (SSL), unsupervised multimodal change detection (UMCD), and unsupervised multimodal
135 image matching (UMIM), among others. To demonstrate its utility, we benchmark a suite of representative models across several of these tasks. All experimental results, along with the source code and pretrained weights, are publicly released to provide strong baselines and accelerate future developments in disaster response in the community.
3. We provide an in-depth analysis that uncovers key challenges and mechanisms of the multimodal building damage assessment task. Through carefully designed experiments, we reveal the difficulties of cross-event generalization, in-
140 vestigate the role of pre-event optical data in aiding damage classification, and quantify the performance gaps between different post-event modalities. These findings offer valuable insights for the development of more robust and practical models for disaster response.

2 Dataset Description

2.1 Study areas and disaster events

145 We selected 14 disaster events across the globe for BRIGHT, as illustrated in Figure 2 and Table 2. Since both Capella Space and Umbra satellites were launched in 2020, we focused on study areas where disasters have occurred since then. The selected regions are primarily in developing countries, where public administration and disaster response capacities tend to be weaker compared to those in developed nations, making international assistance more critical. The dataset covers five major types of natural disasters: earthquakes, storms (including hurricanes and cyclones), wildfires, floods, and volcanic eruptions. Addi-
150 tionally, it includes human-made disasters, such as accidental explosions and armed conflicts. Detailed descriptions of the 14 disaster events are provided in Appendix A.




Figure 2. Geographic distribution of disaster events present in BRIGHT.

2.2 Construction of BRIGHT

Figure 3 shows the flowchart of developing BRIGHT. The optical EO data in the dataset are mainly from Maxar’s Open Data program³, while the SAR EO data are from Capella Space⁴ and Umbra⁵. Both Capella and Umbra data have two imaging modalities, *i.e.*, Spotlight and Stripmap, respectively. The Spotlight mode has a higher spatial resolution but less coverage. In the region of interest, we preferred Spotlight mode if suitable data was available in the data provider’s inventory. Otherwise, we chose Stripmap. The optical EO data consists of red, blue, and green bands, while the SAR EO data consists of amplitude data in the VV or HH bands. For optical EO data, the digital number was converted to reflectance and then standardized to an 8-bit data format. For SAR imagery, after the data had been terrain-corrected, we used the data provider’s recommended method for pre-processing the amplitude data to convert the data to an 8-bit data format. Although both optical and SAR images are geocoded, there are still pixel offsets between them. Therefore, multiple EO experts manually aligned the paired optical and SAR data and cross-checked their results to ensure the precise registration between the two modalities. Figure C1 in Appendix C shows the selected control points on three disaster scenes.

The labels in BRIGHT consist of two components: building polygons and post-disaster building damage attributes. Expert annotators manually labeled the building polygons, then all labels underwent independent visual inspections of EO experts

³<https://www.maxar.com/open-data>

⁴<https://www.capellaspace.com/earth-observation/gallery>

⁵<https://umbra.space/open-data/>

Table 2. Summary of basic information of the BRIGHT dataset with disaster events listed in chronological order. GSI refers to the Geospatial Information Authority of Japan, and IGN refers to the Instituto Geográfico Nacional (National Geographic Institute) of Spain.

Disaster area	Type of disaster	Date	GSD (m/pixel)	Data provider / source	No. of tiles	No. of building
Beirut, Lebanon	Explosion (EP)	04 Aug. 2020	1	Maxar & Capella	133	25,496
Bata, Equatorial Guinea	Explosion (EP)	07 Mar. 2021	0.5	Maxar & Capella	107	8,893
Goma, DR Congo	Volcano eruption (VE)	22 May 2021	0.33	Maxar & Capella	123	18,741
Les Cayes, Haiti	Earthquake (EQ)	14 Aug. 2021	0.48	Maxar & Capella	73	18,918
La Palma, Spain	Volcano Eruption (VE)	19 Sept. 2021 - 13 Dec. 2021	0.3-0.35	IGN (Spain) & Capella	933	30,239
Boulder, USA	Wildfire (WF)	30 Dec. 2021 - 01 Jan., 2022	0.6	NAIP & Capella	77	8,365
Ukraine	Armed conflict (AC)	22 Mar. 2022 - 21 Sept. 2022	0.6	Google Earth & Capella	513	56,770
Turkey	Earthquake (EQ)	06 Feb. 2023	0.30-0.35	Maxar & Capella & Umbra	1,114	135,033
Kyaukpyu, Myanmar	Cyclone (CC)	14 May 2023	0.6	Google Earth & Capella	126	8,052
Maui, Hawaii, USA	Wildfire (WF)	08 Aug. 2023 - 09 Aug. 2023	0.6	NOAA & Capella	65	3,995
Morocco	Earthquake (EQ)	08 Sept. 2023	0.35-0.4	Maxar & Capella	567	6,269
Derna, Libya	Flood (FL)	10 Sept. 2023	0.35	Maxar & Capella	124	1,0979
Acapulco, Mexico	Hurricane (HC)	25 Oct. 2023	0.35-0.8	Google Earth & Capella	212	18,437
Noto, Japan	Earthquake (EQ)	01 Jan. 2024	0.5	GSI & Umbra	79	8,153
<i>Total</i>	-	-	0.3-1	-	4,246	384,596

Table 3. Definition of different categories in BRIGHT.

Category	Definition
Background (0)	All non-building pixels
Intact (1)	No visible signs of structural damage, water intrusion, shingle displacement, or burn marks.
Damaged (2)	Partial structural damage to the building, such as missing roof members, visible cracks, or partial wall/roof collapse. Buildings may be partially burned, surrounded by water or mud, or affected by nearby volcanic flows.
Destroyed (3)	Completely collapsed, burned, partially/completely covered by water/mud or no longer present




Figure 3. Overall flowchart of developing the BRIGHT dataset.

to ensure accuracy. Damage annotations were obtained from Copernicus Emergency Management Service⁶, the United Nations Satellite Centre (UNOSAT) Emergency Mapping Products⁷, and the Federal Emergency Management Agency (FEMA)⁸. These annotations were derived through visual interpretation of high-resolution optical imagery captured before and after the disasters by EO experts, supplemented by partial field visits. To harmonize these diverse annotations and ensure consistency across all 14 disaster events, we implemented a rigorous, multi-stage process. First, we established a single, standardized three-tier classification scheme, including Intact (with pixel value 1), Damaged (with pixel value 2), and Destroyed (with pixel value 3), with clear definitions provided in Table 3, drawing on the frameworks of FEMA’s Damage Assessment Operations Manual, EMS-98, the BDD dataset (Adriano et al., 2021), and the xBD dataset (Gupta et al., 2019). While the source agencies’ terminology can differ (e.g., “Severe Damage” vs. “Major Damage”), their underlying definitions for EO-based assessment are conceptually consistent. We leveraged this alignment for an initial rule-based mapping, where various intermediate damage tiers were conservatively aggregated into our single “Damaged” category. Second, our team of EO experts conducted a comprehensive manual verification and refinement of every annotation using multi-temporal VHR imagery on platforms like Google Earth Pro. This final stage served as the ultimate guarantor of consistency. We paid special attention to ambiguous source

⁶<https://emergency.copernicus.eu>

⁷<https://unosat.org/products>

⁸<https://www.fema.gov>




Figure 4. Thumbnails of local areas in 14 disaster events in the BRIGHT dataset. For visualization purposes, different events have different scales.

labels, such as “Possibly Damaged”. Adopting a conservative approach, these were re-classified as “Intact” if clear structural damage was not evident, thereby ensuring a high-confidence “Damaged” class. We also manually disaggregated all area-based annotations (i.e., where an entire block was assigned a single category). We re-processed these to assign a precise, building-wise damage label to each individual structure, ensuring instance-level consistency and granularity across the entire dataset. The damage annotations were provided as vector point files. The final building damage labels were generated by overlaying these points with the building polygons and assigning corresponding damage attributes. To prevent geographic misallocation due to possible coordinate offsets, the coordinate systems of the points and polygons were unified, with a visual inspection performed prior to the final allocation. Figure 4 presents thumbnails of selected local areas from the 14 disaster events.

2.3 Statistics of BRIGHT

The basic information about BRIGHT, including disaster events, EO data, the number of corresponding EO tiles, and the total number of building pixels, is summarized in Table 2. After cropping the EO data into 1024×1024 -pixel tiles, BRIGHT contains 4,246 multimodal image pairs.

The key statistics of BRIGHT are illustrated in Figure 5. Figure 5-(a) shows the pixel value distribution for optical and SAR images from one human-made disaster and three natural disasters. The varying geographical landscapes and land cover across different regions result in distinct means and standard deviations of pixel values. This highlights BRIGHT’s geographical diversity, which makes it a robust dataset for studying building damage assessment in diverse environments. To ensure that models trained on BRIGHT can accurately detect buildings and assess damage levels, it is crucial that the dataset includes a wide variety of building styles from different regions. Figure 5-(b) shows that BRIGHT covers buildings at multiple scales, exhibiting a “long-tail” distribution. This multi-scale representation challenges DL models to develop the ability to capture features at varying scales, enhancing robustness and accuracy.

Figure 5-(c) further illustrates the feature distribution of buildings in the optical and SAR images for the four events shown in Figure 5-(a), which demonstrates clear inter-event separability in both modalities. BRIGHT also faces a significant challenge of sample imbalance, as shown in Figure 5-(d). There is a notable imbalance between background pixels and foreground (building) pixels, with a ratio of approximately 7:1. The imbalance exists within the damage categories: about 6.5% of building pixels represent destroyed buildings, 10.7% correspond to damaged buildings, and 82.8% are intact buildings. This imbalance can complicate model training, necessitating careful strategies to develop robust DL models.

Moreover, since accurate registration ensures spatial consistency across modalities, the registration accuracy between optical and SAR EO data in BRIGHT was analyzed to provide a solid foundation for accurate building damage assessment. Due to the absence of real ground truth, a proxy method that leverages existing multimodal image descriptors was introduced to estimate registration errors. This approach is detailed in Appendix C. The Table 4 reports the mean registration errors, measured as the root mean square error (RMSE) of pixel displacements. It was obtained using three representative multimodal image registration methods: RIFT (Li et al., 2020), SRIF (Li et al., 2023a), and LNIFT (Li et al., 2022). The overall average RMSE is approximately 1.024 pixels, with a lower error of 1.006 pixels specifically within the building regions.




Figure 5. Statistics of the BRIGHT dataset. (a) Distributions of band values of samples from four study sites. (b) Distribution of building scales. (c) Feature distribution of buildings of four events under two imaging modalities. (d) Percentage of building and background pixels and percentage of different damage levels in building pixels.

Table 4. Proxy registration errors (in pixel) estimated using different multimodal image descriptors. The estimated registration errors for each event are reported in Table C1 in Appendix C.

Descriptor	RMSE (All)	RMSE (Building)
RIFT (Li et al., 2020)	1.125	1.138
LNIFT (Li et al., 2022)	0.857	0.825
SRIF (Li et al., 2023a)	1.090	1.054
Average	1.024	1.006

2.4 Dataset splitting strategy

To train DL models using BRIGHT and evaluate their generalizability, it is necessary to split the dataset into a training set, validation set, and test set. Gerard et al. (2024) suggested that dividing the dataset on an event-by-event basis, rather than randomly across the entire dataset, provides a more accurate reflection of a model’s generalizability. Therefore, for the 14 events listed in Table 2, we divide the corresponding data for each event into a ratio of 7:1:2 for training, validation, and test subsets, respectively. Then, the subsets obtained for each event are merged to create the final training, validation, and test sets. In the experiments, the baseline models are trained using the training set, and the optimal hyperparameters (e.g., learning rate) and checkpoints are selected based on performance on the validation set. The generalization capability of the baseline models is subsequently evaluated on the test set.

In addition to the above standard ML data splitting, we also introduce a cross-event transfer setup to better evaluate the ability of models to generalize across disaster events. This is a critical challenge in real-world applications where models are expected to handle unseen disaster types and locations. Two setups are established for cross-event transfer generalization:

- Zero-shot setup: This setting mimics a real-world scenario where a newly occurring disaster must be analyzed without any prior labeled data from the same event. We isolate one event as an unseen test set while using the remaining 13 events for training and validation. This setting evaluates the cross-event generalization ability of models, testing how well learned knowledge can be transferred from previous disasters to an entirely new disaster event. Due to the high variability of disaster types and geographies, this setup is inherently challenging, as models trained on past disasters may struggle to assess damage patterns in a previously unseen event accurately.
- One-shot setup: Recognizing the difficulty of the zero-shot setup, we introduce a one-shot setup. This setting simulates a realistic scenario where a single, representative sample from the new disaster can be quickly labeled to guide model adaptation. In this setting, a limited subset of labeled data (one pair for training and one pair for validation) from the target disaster event is incorporated into the training process. At the same time, the majority of the test set remains unseen. This setup evaluates the model’s ability to leverage a minimal amount of manually labeled data to improve disaster-specific adaptation.

It is worth noting that our cross-event transfer setup differs from classic few-shot learning tasks in the computer vision field (Amirreza Shaban and Boots, 2017; Wang et al., 2020). Our goal is not to recognize new classes, but to adapt the model’s knowledge of existing classes to a new domain, *i.e.*, an unseen disaster event.

3 Methodology

3.1 Problem statement

The objective of building damage assessment is to interpret EO data of areas affected by a disaster by generating a building damage map that reflects the extent of damage to buildings. To achieve this, two common approaches are typically employed.

One is to directly treat the building damage assessment task as a single semantic segmentation task (Adriano et al., 2021; Gupta and Shah, 2021). In this approach, the pre- and post-event images are taken as inputs of the model, and then the final damage map is directly predicted. This process can be formalized as $\mathbf{Y}^{dam} = \mathcal{M}^{seg}(\mathbf{X}^{T_1}, \mathbf{X}^{T_2})$, where $\mathbf{X}^{T_1}$ is the pre-event imagery, $\mathbf{X}^{T_2}$ is the post-event imagery, $\mathcal{M}^{seg}(\cdot)$ is a semantic segmentation model, $\mathbf{Y}^{dam}$ is the obtained damage map. In the context of this paper, $\mathbf{X}^{T_1}$ is VHR optical imagery and $\mathbf{X}^{T_2}$ is VHR SAR imagery.

The second adopts the task decoupling approach (Gupta et al., 2019; Zheng et al., 2021), which breaks down building damage assessment into two subtasks: the building localization task, *i.e.*, separating the building from the background, and the damage classification task, *i.e.*, focusing on the classification between different levels of damage. This approach can be formulated as $\mathbf{Y}^{loc} = \mathcal{M}^{loc}(\mathbf{X}^{T_1})$ and $\mathbf{Y}^{clf} = \mathcal{M}^{clf}(\mathbf{X}^{T_1}, \mathbf{X}^{T_2})$, where $\mathbf{Y}^{loc}$ is the building localization map, $\mathbf{Y}^{clf}$ is the damage classification map, $\mathcal{M}^{loc}(\cdot)$ and $\mathcal{M}^{clf}(\cdot)$ are models for building localization and damage classification, respectively. $\mathcal{M}^{loc}(\cdot)$ and $\mathcal{M}^{clf}(\cdot)$ can be two separate models (Gupta et al., 2019) or a unified multi-task learning model (Zheng et al., 2021; Chen et al., 2022a, 2024). The final building damage map is obtained by combining the two outputs using a simple mask operation: $\mathbf{Y}^{dam} = \mathbf{Y}^{loc} \odot \mathbf{Y}^{clf}$. Since this work aims not only to provide a large-scale multimodal dataset to support all-weather disaster response, but also to offer insights for designing appropriate methods in future research, both approaches are employed in the experiments to compare their results.

It is worth noting that in this work, we focus on the formulation of building damage assessment as a bi-temporal task, where both pre- and post-event images are used as inputs. This formulation aligns closely with generic change detection tasks, which aim to identify changes between two time points. Conceptually, building damage assessment can be viewed as a specialized “one-to-many” semantic change detection problem (Zheng et al., 2021, 2024; Lu et al., 2024), where the objective is not only to detect whether a change has occurred but also to categorize the type and severity of changes (damages) to buildings. Many existing methods are thus derived from or adapted versions of generic change detection frameworks (Chen et al., 2024; Zheng et al., 2024; Guo et al., 2024).

3.2 Benchmark suites

Several advanced deep network architectures from both the computer vision and EO communities are evaluated on BRIGHT. Since building damage assessment can be considered a specialized semantic segmentation task, we adopted two well-known segmentation networks from the computer vision field: UNet (Ronneberger et al., 2015) and DeepLabV3+ (Chen et al., 2018); and five state-of-the-art networks from the EO community: SiamAttnUNet (Adriano et al., 2021), SiamCRNN (Chen et al., 2020), ChangeOS (Zheng et al., 2021), DamageFormer (Chen et al., 2022a), and ChangeMamba (Chen et al., 2024). These seven networks encompass a broad range of representative DL architectures, including convolutional neural networks (CNNs), recurrent neural networks (RNNs), Transformers, and the more recent Mamba architecture. Among the seven networks, UNet, DeepLabV3+, and SiamAttnUNet adopt the first approach defined in Section 3.1, *i.e.*, directly treating building damage assessment as a single semantic segmentation task. In contrast, SiamCRNN, ChangeOS, DamageFormer, and ChangeMamba adopt the second approach by decoupling the task into building localization and damage classification tasks.

Beyond supervised DL models, BRIGHT also enables the evaluation of other learning strategies and methods commonly explored in the EO and computer vision communities:

- Unsupervised domain adaptation (UDA) methods for the zero-shot transfer setup, enabling models to transfer knowledge across disaster events with no labeled samples from the target event.
- 280 • Semi-supervised learning (SSL) approaches for the one-shot transfer setup, leveraging a small number of labeled samples and the remaining unlabeled samples from new disaster events to refine model adaptation.
- Unsupervised multimodal change detection (UMCD) methods, which exploit the modality-independent relationship in optical and SAR data to detect land-cover changes without requiring manual annotations.
- 285 • Unsupervised multimodal image matching (UMIM) methods, which aim to learn modality-independent features to enable automatic registration of multimodal data without relying on manual alignment.

3.3 Model training

To train the supervised models, we use a combination of cross-entropy loss and Lovasz softmax loss (Berman et al., 2018). Cross-entropy loss serves as the basic loss function for dense prediction tasks, while Lovasz softmax loss effectively addresses sample imbalance between non-building and building pixels and across different damage levels. For UNet, DeepLabV3+, 290 and SiamAttnUNet, which directly predict damage maps from the input multimodal image pairs, the training loss function is defined as:

$$\mathcal{L}_{coupled}^{bda} = \mathcal{L}_{ce}^{bda} + \mathcal{L}_{lov}^{bda}. \quad (1)$$

For SiamCRNN, ChangeOS, DamageFormer, and ChangeMamba, which decouple building damage assessment into building localization and damage classification subtasks, the training loss function is defined as:

$$295 \quad \mathcal{L}_{decoupled}^{bda} = \mathcal{L}_{ce}^{loc} + \mathcal{L}_{lov}^{loc} + \mathcal{L}_{ce}^{clf} + \mathcal{L}_{lov}^{clf}. \quad (2)$$

All the models are trained using the AdamW optimizer (Loshchilov and Hutter, 2017) with a learning rate of $1e^{-4}$ and a weight decay of $5e^{-3}$. The training process consists of 50,000 iterations, with a batch size of 16. To enhance sample diversity and improve model generalization, we apply several data augmentation techniques, including random flipping, random rotation (in 90-degree increments), and random cropping. For the zero-shot and one-shot setups, we also test UDA and SSL 300 techniques to better utilize unlabeled and limited target disaster data. The training settings for the UDA, SSL, UMCD, and UMIM approaches are provided in Appendix F, G, H, and I, respectively.

3.4 Accuracy assessment

We adopt overall accuracy (OA), F1 score (F1), and mean intersection over union (mIoU) to evaluate the performance of the models. These are commonly used metrics in building damage assessment (Zheng et al., 2021). Following the setup in previous

Table 5. Accuracy assessment for different DL models on the test set under the standard ML data split (set-level mIoU). The highest values are in **bold**, and the second-highest values are underlined.

Method	F_1^{loc} (%)	F_1^{clf} (%)	Final OA (%)	Final mIoU (%)	IoU per class (%)			
					Background	Intact	Damaged	Destroyed
UNet	87.97	72.24	95.47	64.94	96.19	71.27	39.13	53.17
DeepLabV3+	87.00	70.33	95.43	64.80	95.98	70.98	37.53	54.69
SiamAttnUNet	88.16	70.13	95.45	64.26	96.14	71.90	35.06	53.92
SiamCRNN	89.45	72.02	95.76	65.73	96.48	73.44	38.91	54.18
ChangeOS	89.60	71.88	95.84	65.98	96.54	73.85	38.99	54.53
DamageFormer	<u>90.29</u>	<u>72.51</u>	<u>96.13</u>	<u>67.09</u>	<u>96.87</u>	<u>75.04</u>	<u>39.86</u>	<u>56.59</u>
ChangeMamba	90.90	72.70	96.22	67.63	96.96	75.59	40.05	57.91

unimodal building damage assessment studies and the related xView2 Challenge (Gupta et al., 2019), the F1 score is used to assess the performance of the models in the building localization and damage classification subtasks. OA and mIoU are used to measure the overall quality of the building damage map, providing a comprehensive assessment of the models’ ability to localize buildings and classify damage levels accurately.

4 Results and Analysis

4.1 Evaluation on standard machine learning data split

Table 5 shows the results for each model on the test set. We observe that ChangeMamba achieves the best overall performance, with an OA of 96.22%, a mIoU of 67.63%, and the highest F_1^{loc} and F_1^{clf} scores of 90.90% and 72.70%, respectively. DamageFormer also performs well, following ChangeMamba, with a mIoU of 67.09% and an OA of 96.13%. Both models demonstrate a strong capability in the building localization and damage classification tasks. The accuracy of ChangeMamba and DamageFormer underscores the importance of leveraging advanced DL architectures to improve performance in complex tasks such as building damage assessment. For models that use a direct prediction approach (UNet, DeepLabV3+, SiamAttnUNet), UNet achieves the best results, with a mIoU of 64.94% and an OA of 95.47%. However, its performance still lags behind the decoupled models, which emphasizes the advantage of task decoupling.

To ensure that the evaluation is not dominated by a few events with a large number of images, e.g., Turkey-EQ-2023, Table 6 presents the event-level mIoU for each model. ChangeMamba and DamageFormer achieve the highest average mIoU, with scores of 51.39% and 52.26%, respectively. DamageFormer performs very well on events such as Beirut-EP-2020, Marshall-WF-2021, and Derna-FL-2023. This shows its robustness across different types of disasters. Although performance varies across events, earthquake-related events such as Les Cayes-EQ-2021, Morocco-EQ-2023, and Noto-EQ-2024 present a greater challenge to all models, with a relatively low average mIoU. This highlights the need for further research to improve the

Table 6. The mIoU on different events for different DL models (event-level mIoU). The highest values are in **bold**, and the second-highest values are underlined.

Events	UNet	DeepLabV3+	SiamAttnUNet	SiamCRNN	ChangeOS	DamageFormer	ChangeMamba
Beirut-EP-2020	52.23	48.53	51.85	52.04	<u>52.57</u>	60.55	49.38
Bata-EP-2021	39.24	40.26	38.69	41.44	39.80	40.37	<u>40.47</u>
Goma-VE-2021	61.52	59.26	62.95	62.50	61.84	61.55	<u>62.89</u>
Les Cayes-EQ-2021	40.39	38.52	40.84	40.99	41.24	<u>41.75</u>	42.36
La Palma-VE-2021	62.72	64.31	65.56	64.24	<u>65.34</u>	64.27	65.13
Marshall-WF-2021	53.96	44.67	51.48	56.57	56.55	57.50	<u>57.23</u>
Ukraine-AC-2022	43.92	43.84	40.71	45.52	45.56	47.45	<u>47.43</u>
Turkey-EQ-2023	49.34	49.53	48.57	50.69	<u>52.34</u>	51.36	52.70
Kyaukpyu-CC-2023	38.11	38.56	40.49	<u>42.35</u>	42.03	48.42	41.05
Hawaii-WF-2023	53.02	57.81	57.92	55.34	58.40	60.82	<u>60.70</u>
Morocco-EQ-2023	42.62	42.05	42.58	43.17	43.05	<u>44.00</u>	44.06
Derna-FL-2023	59.47	56.08	66.82	59.48	57.73	62.24	<u>63.30</u>
Acapulco-HC-2023	49.21	46.66	42.15	49.98	51.08	<u>50.82</u>	47.01
Noto-EQ-2024	36.49	37.33	36.56	38.35	<u>42.84</u>	40.57	45.78
Average	48.73	47.67	49.08	50.19	50.74	52.26	<u>51.39</u>

robustness of earthquake damage assessment models, particularly where damage patterns are more complex and diverse. Figure 6 shows some building damage maps obtained by the seven models on the test set.

4.2 What have the models learned and what can they learn?

To better understand the models’ behavior beyond performance metrics, we explore the internal attention patterns of the trained ChangeOS using class activation maps (CAMs) (Selvaraju et al., 2017). Figure 7 presents the CAM responses of ChangeOS across six representative disaster events: volcano, explosion, flood, wildfire, earthquake, and hurricane. We observe that the attention distribution varies across disaster types. Taking Goma-VE-2021 as an example, for the “Destroyed” category, the encoder exhibits strong activations in nearly all the built-up regions in the optical images, accurately localizing individual buildings. This suggests that the model has effectively learned to extract detailed structural cues from pre-event optical imagery. In contrast, for the SAR images, the encoder shows intense activation over the lava-covered regions on the left. This indicates that the model has identified the lava-covered regions as a key signal for destruction, likely due to the significant backscatter changes caused by lava flow. In the “Damaged” category, the activations are more subtle. Attention is primarily focused near the boundary of the lava flow, where partial or ambiguous structural changes occur. In the decoder, the destroyed buildings in the lava-affected area are strongly activated, which aligns well with the reference labels. Conversely, for the “Damaged” class,




Figure 6. Damage maps predicted by different models on the test images of the 14 disaster events. The meaning of the color in reference maps and damage maps is consistent with Figure 4.




Figure 7. Visualization of feature response to “Damaged” and “Destroyed” categories in different layers of deep models over three event cases. (a) Goma-VE-2021. (b) Beirut-EP-2023. (c) Libya-FL-2023. (d) Hawii-WF-2023. (e) Turkey-EQ-2023. (f) Acapulco-HC-2023. In the visualization, closer to red indicates larger response values, and closer to dark blue the opposite. See Appendix B for implementation details.

only a few regions are activated. This suggests that inferring partial damage from SAR imagery in volcanic disaster scenarios remains a significant challenge, as subtle structural degradation is often not clearly reflected in SAR backscatter or texture.

Beyond understanding what the models have already learned, a more critical question is: What can they learn? More specifically, to what extent can optical–SAR modality improve the accuracy of building damage assessment across diverse disaster scenarios? To investigate this, we aggregate the IoUs of the seven models and calculate their average IoU across the seven major disaster types, as shown in Figure 8. The wildfire and volcano events exhibit the highest IoU scores for the “Destroyed” category, both exceeding 70%. This indicates that the model can effectively leverage SAR-based backscatter anomalies, such as lava flows or widespread debris fields, to detect fully destroyed structures. These results demonstrate a strong potential of optical–SAR fusion in such high-impact scenarios. For the “Damaged” category, the performance drops significantly. The




Figure 8. IoU distribution of deep models over seven disaster types. Each bar represents the average IoU of seven DL models for that specific category under each disaster type. The error bars indicate the standard deviation of IoU scores across the seven models.

average IoU for damaged buildings in wildfire events remains below 20%, while in volcano events, it falls to around 5%. This suggests that single-polarization SAR imagery lacks the fine-grained information needed to reliably distinguish partially damaged buildings, where structural integrity may still be partially preserved and backscatter signals remain ambiguous. In the earthquake events, both damaged and destroyed categories yield relatively low IoUs. This is likely due to the complex and heterogeneous patterns of structural collapse typical of seismic events, where damage is often subtle, partial, and highly variable. These conditions pose significant challenges for SAR-based assessment. Interestingly, the model achieves relatively high IoU scores in the flood and hurricane events for the “Damaged” category with approximately 50% and 60%, respectively. This indicates that SAR effectively captures contextual environmental changes, such as water inundation or terrain disruption, which indirectly aid in assessing building damage. In the case of the conflict event, the model’s performance on the “Destroyed” class is surprisingly low. This might be attributed to the limited number of destroyed samples in the dataset for this category, which leads to insufficient learning and poor generalization.

These quantitative limitations are vividly illustrated by the typical failure cases shown in Figure 9. In the Bata-Explosion-2021 event, models misclassify severely destroyed buildings as intact, reflecting the difficulty of interpreting heterogeneous debris patterns. Similarly, in the Noto-Earthquake-2024 event, large-scale collapses are largely missed, highlighting the challenge of diverse and subtle seismic damage. These examples visually confirm that the significant heterogeneity in damage patterns makes it challenging for models to learn a consistent and generalizable representation of damage.

In summary, these findings confirm both the promise and limitations of optical-SAR modality for all-weather, global-scale disaster response. Although this combination performs well in events characterized by large-scale surface disruption (e.g., wildfires, volcanoes), it struggles with subtle or localized damage patterns. Incorporating richer data sources, such as fully




Figure 9. Typical failure cases of different models on Bata-Explosion-2021 and Noto-Earthquake-2024 in BRIGHT.

Table 7. Performance comparison of UNet and DeepLabV3+ using only post-event SAR input and pre-event optical plus post-event SAR inputs for damage classification task. Here, accurate building masks are provided as the post-processing step to all models to isolate the effect of building localization task on the damage classification task.

Method	Modality	F_1^{clf} (%)	Final mIoU (%)	IoU per class (%)			
				Background	Intact	Damaged	Destroyed
UNet	Post-event SAR	68.71	69.84	100.0	88.19	35.83	55.35
	Pre-event optical + post-event SAR	73.59	72.41	100.0	89.38	44.83	55.42
DeepLabV3+	Post-event SAR	72.12	72.19	100.0	89.59	39.63	59.54
	Pre-event optical + post-event SAR	73.90	73.93	100.0	90.32	40.45	64.94

polarimetric SAR and LiDAR data, can further enhance the accuracy and reliability of future all-weather building damage assessments.

4.3 The role of optical pre-event data in multimodal building damage assessment

370 In the last section, CAM visualizations revealed that DL models also exhibit responses to disaster-specific patterns in pre-event optical imagery. This observation suggests that optical data may play a more complex role in multimodal building damage mapping than simply supporting building localization. In other words, in a multimodal bi-temporal setup, does pre-event

Table 8. Performance comparison of different post-event modalities on a subset of BRIGHT. Results are reported for UNet, DeepLabV3+, and DamageFormer on five disaster events where high-quality post-event optical imagery is available: Bata-Explosion-2020, Beirut-Explosion-2021, Hawaii-Wildfire-2023, Libya-Flood-2023, and Noto-Earthquake-2024.

Method	Post-event modality	F_1^{loc} (%)	F_1^{clf} (%)	Final mIoU (%)	IoU per class (%)			
					Background	Intact	Damaged	Destroyed
UNet	SAR	85.05	71.43	62.94	94.39	65.60	42.34	49.43
	Optical	86.46	75.64	65.96	94.56	68.62	45.27	55.36
	Optical+SAR	86.70	74.46	66.29	94.72	69.16	41.73	59.57
DeepLabV3+	SAR	83.55	67.52	60.57	93.86	65.62	35.64	47.16
	Optical	85.79	74.39	64.87	94.33	67.90	44.31	52.94
	Optical+SAR	85.90	74.87	65.84	94.48	69.60	44.68	54.60
DamageFormer	SAR	88.41	73.43	65.56	95.30	70.62	41.31	55.00
	Optical	88.32	78.04	69.76	95.37	72.72	47.26	63.68
	Optical+SAR	88.86	79.27	70.79	95.56	73.64	48.44	65.51

optical imagery act solely as a localization aid, or does it provide additional semantic cues that networks can exploit for more accurate damage classification?

375 To explore this, we conducted controlled experiments using UNet and DeepLabV3+. Both networks were trained under two configurations: (i) using post-event SAR imagery only, and (ii) using multimodal pre- and post-event inputs (optical-SAR). To isolate the contribution of pre-event optical data beyond building localization, we provided perfect building masks for postprocessing in both settings. This design ensures that any observed differences in performance are attributable to the additional information from pre-event optical imagery, rather than differences in network architecture or localization accuracy.

380 The results, summarized in Table 8, show that incorporating pre-event optical imagery leads to notable improvements in distinguishing building damage levels. For UNet, the IoU for the “Damaged” class increased from 35.83% (SAR only) to 44.83% (Optical-SAR), and for the “Destroyed” class from 55.35% to 55.42%. DeepLabV3+ exhibited significant gains also, with IoU improvements from 39.63% to 40.45% for “Damaged” category, and from 59.54% to 64.94% for “Destroyed” category. These results suggest that pre-event optical imagery contributes beyond mere building localization, enriching the feature space for
385 more effective semantic comparison for different building damage levels across modalities.

4.4 Impact of post-event modality on building damage assessment performance

Although the primary design of BRIGHT is to facilitate all-weather disaster response through the use of pre-event optical and post-event SAR imagery, it is also important to understand how these modalities compare when high-quality post-event optical imagery is available. To this end, we conducted supplementary experiments on a subset of events, including Bata-Explosion-
390 2020, Beirut-Explosion-2021, Hawaii-Wildfire-2023, Libya-Flood-2023, and Noto-Earthquake-2024, for which pre-processed post-event optical data were accessible. We evaluated three experimental setups: (i) optical-only (pre-event optical + post-event

Table 9. Further contributions to mIoU from post-processing algorithms. ChangeMamba (Chen et al., 2024) is used here as the baseline. Details on these algorithms are provided in Appendix E.

Method	mIoU (set)	mIoU (event)
Baseline	67.63	51.39
Test-time augmentation	68.50	51.95
Object-based majority voting	67.22	52.08
Ensembling multiple models	68.45	52.14
All	68.86	52.31

optical), (ii) SAR-only (pre-event optical + post-event SAR, i.e., the standard BRIGHT setting), and (iii) optical+SAR fusion (pre-event optical + post-event optical + post-event SAR).

Table 8 presents the experimental results. As expected, when ideal post-event optical imagery is available, the optical-only setup achieves higher performance than the SAR-only setup. For example, with DamageFormer, the optical-only configuration reaches a final mIoU of 69.76%, compared to 65.56% for SAR-only. Importantly, the performance gap between optical and SAR is not substantial, demonstrating that SAR alone provides a strong alternative in the absence of usable optical imagery. Moreover, the fusion of optical and SAR consistently yields the best results across all tested models. For instance, DamageFormer’s mIoU further increases to 70.79% with Optical+SAR fusion, indicating that SAR contributes complementary information that strengthens performance even under optimal optical conditions.

These findings underscore two important insights. First, multimodal fusion is beneficial even when high-quality optical data are available, as SAR provides unique structural information that enriches the optical signal. Second, the performance of the SAR-only approach, being reasonably close to the optical-only results, highlights the practical value of SAR in real-world disaster scenarios where post-event optical imagery is often unavailable. BRIGHT is therefore designed to advance the development of models for these realistic, often non-ideal, but operationally critical all-weather disaster response settings.

4.5 Effect of post-processing method

Post-processing techniques help refine raw predictions from DL models, to reduce noise, improve consistency, and ensure spatial coherence in damage maps (Zheng et al., 2021). Here, we explore the impact of post-processing algorithms. Table 9 presents the effect of three post-processing techniques applied to ChangeMamba. The post-processing methods evaluated include test-time augmentation, object-based majority voting, and model ensembling. The details of the methods are provided in Appendix E. As shown in Table 9, the test-time augmentation improves the mIoU by 0.87% at the set level and 0.56% at the event level, demonstrating its effectiveness in enhancing model robustness across diverse disaster scenarios. Object-based majority voting, which aggregates predictions at the building-object level to enforce spatial consistency, slightly reduces set-level mIoU (-0.41%) but improves event-level mIoU (+0.69%). Ensembling multiple models leads to a 0.82% increase in mIoU at the set level and a 0.75% increase at the event level, reinforcing its effectiveness in improving model performance across

Table 10. The mIoU on different events for different DL models in zero-shot and one-shot cross-event transfer setups. The highest values are in **bold**, and the second-highest values are underlined.

Events	UNet	DeepLabV3+	SiamAttnUNet	SiamCRNN	ChangeOS	DamageFormer	ChangeMamba
	zero- / one-shot	zero- / one-shot	zero- / one-shot	zero/one-shot	zero- / one-shot	zero- / one-shot	zero- / one-shot
Beirut-EP-2020	31.62 / 34.89	34.72 / 33.58	35.33 / 29.17	<u>39.43</u> / 40.94	38.48 / 38.21	38.71 / 39.93	40.30 / <u>40.92</u>
Bata-EP-2021	33.48 / 37.48	33.55 / 30.23	33.78 / 35.24	37.46 / <u>40.06</u>	38.00 / 38.18	<u>38.91</u> / 39.42	39.17 / 40.31
Goma-VE-2021	55.97 / 57.14	55.01 / 58.13	57.52 / 55.75	<u>58.91</u> / 59.54	58.66 / 57.85	55.86 / <u>59.89</u>	60.39 / 61.91
Les Cayes-EQ-2021	26.21 / 38.72	35.50 / 37.90	37.22 / 38.76	39.45 / 40.28	39.36 / 40.26	<u>40.41</u> / <u>41.07</u>	41.35 / 42.06
La Palma-VE-2021	32.74 / 35.29	30.86 / <u>36.25</u>	<u>34.00</u> / 34.99	32.90 / 33.32	29.96 / 32.60	32.28 / 37.62	34.18 / 34.93
Marshall-WF-2021	34.97 / 37.84	35.05 / 37.12	35.23 / 39.54	39.82 / <u>46.83</u>	36.31 / 43.72	40.93 / 51.15	<u>40.61</u> / 45.96
Ukraine-AC-2022	32.36 / 37.50	32.23 / 35.98	33.84 / 35.76	34.77 / <u>38.14</u>	36.30 / 39.28	33.84 / 36.51	<u>35.78</u> / 36.99
Turkey-EQ-2023	37.82 / 35.83	37.28 / <u>40.49</u>	36.09 / 37.49	38.32 / 40.11	<u>38.45</u> / 39.51	37.20 / 40.27	41.28 / 42.80
Kyaukpyu-CC-2023	31.46 / 31.71	34.01 / 34.75	33.77 / 35.93	<u>34.19</u> / 37.71	36.66 / <u>36.25</u>	32.99 / 33.54	33.85 / 36.08
Hawaii-WF-2023	31.73 / 39.80	31.32 / 45.44	<u>31.79</u> / 45.12	30.94 / 42.65	31.15 / 41.96	31.60 / <u>50.64</u>	36.14 / 50.89
Morocco-EQ-2023	38.24 / 39.81	40.48 / 41.23	40.78 / 40.69	40.89 / 42.89	41.40 / 41.89	<u>41.89</u> / <u>43.42</u>	42.86 / 43.51
Derna-FL-2023	33.25 / 38.44	35.07 / 40.02	33.76 / 42.15	35.88 / <u>44.77</u>	35.80 / 43.85	<u>36.27</u> / 44.69	37.02 / 45.47
Acapulco-HC-2023	26.99 / 38.76	27.78 / 33.24	30.93 / <u>38.94</u>	26.70 / 34.53	27.06 / 34.50	27.48 / 32.87	<u>28.45</u> / 38.96
Noto-EQ-2024	31.02 / 33.83	35.57 / 39.17	35.87 / 36.22	36.04 / 38.52	<u>39.14</u> / 39.15	36.37 / <u>43.93</u>	39.90 / 44.33
Average	34.13 / 38.36	35.60 / 38.82	36.42 / 38.98	37.55 / 41.45	<u>37.62</u> / 40.52	37.48 / <u>42.49</u>	39.38 / 43.23

different disaster events. Applying all post-processing techniques together yielded the highest performance improvement, with a 1.23% increase in set-level mIoU and a 0.92% increase in event-level mIoU. These results confirm that combining different post-processing methods can significantly enhance the reliability of AI-based damage assessments, ensuring better generalization across disaster types and locations. In summary, post-processing techniques are crucial in improving the accuracy of building damage maps. Future work can further explore adaptive post-processing strategies tailored to specific disaster types to enhance prediction reliability in multimodal EO data contexts.

4.6 Evaluation on cross-event transfer setup

4.6.1 Baseline methods

Cross-event transfer, especially under zero-shot settings, poses a significant challenge for building damage assessment. As shown in Table 10, the average mIoU of all baseline models in the zero-shot setting is below 40%, a noticeable drop compared to their performance under the fully supervised standard ML data split in Table 6, where the models typically achieve 48% to 52% mIoU. This performance gap underscores the difficulty of generalizing to unseen disaster events without access to any target domain supervision due to substantial domain shifts in imaging conditions, damage patterns, urban morphology, and sensor response. Despite this, all models exhibit clear performance gains in the one-shot setting, where a small number of

Table 11. Results of unsupervised domain adaptation methods adopted for zero-shot cross-event transfer setup. DeepLabV3+ (Chen et al., 2018) is used as the baseline here. Detailed mIoU on each event is listed in Table F1.

Method	Avg. mIoU
Source-Only	35.60
AdaptSeg (Tsai et al., 2018)	36.05
AdvEnt (Vu et al., 2019)	35.43
CLAN (Luo et al., 2019)	<u>36.01</u>
PyCDA (Lian et al., 2019)	33.19
FDA (Yang and Soatto, 2020)	36.29

Table 12. Results of semi-supervised learning methods adopted for one-shot cross-event transfer setup. DeepLabV3+ (Chen et al., 2018) is used as the baseline here. Detailed mIoU on each event is listed in Table G1.

Method	Avg. mIoU
Baseline	38.82
MT (Tarvainen and Valpola, 2017)	40.00
CCT (Ouali et al., 2020)	<u>39.49</u>
GCT (Ke et al., 2020)	39.34
CPS (Chen et al., 2021)	37.06

430 labeled samples from the target event are available. This suggests that even minimal supervision can significantly aid adaptation to new disaster contexts.

Among the evaluated models, ChangeMamba consistently achieves the highest overall performance, with an average mIoU of 39.38% in zero-shot and 43.23% in one-shot settings, followed by DamageFormer. This highlights the strength of recent Transformer- and Mamba-based architectures in transferring learned knowledge under complex multimodal and disaster
435 scenarios. Models using a decoupled architecture generally outperform direct segmentation models, which confirms that separating building localization and damage classification improves generalization in cross-event transfer tasks. In contrast, the Acapulco-HC-2023 and Ukraine-AC-2022 events show weaker results across all models, reflecting the difficulty in transferring to domains with limited or inconsistent destruction patterns.

4.6.2 Unsupervised domain adaptation and semi-supervised learning methods

440 In the context of cross-event transfer, the UDA and SSL methods naturally emerge as promising strategies to bridge the domain gap between the source and the target disaster events. In the zero-shot setting, UDA methods aim to improve model generalization by aligning the source and target domains without requiring any target labels. In the one-shot setting, SSL methods leverage a small number of labeled samples with abundant unlabeled data from the target event, making them especially appealing for real-world disaster scenarios where rapid and comprehensive annotation is infeasible. It remains an open

question whether these UDA and SSL methods, originally developed for natural image domains, can effectively handle the challenges of multimodal EO data in complex disaster scenarios. We evaluated several representative UDA and SSL methods using DeepLabV3+ as the baseline model to examine their capabilities in complex disaster scenarios. The results for UDA and SSL are reported in Tables 11 and 12, respectively.

Table 11 presents the performance of five classical UDA methods. Compared to the source-only baseline (35.60% mIoU), most methods achieve modest improvements, with FDA (+0.69%) and AdaptSeg (+0.45%) showing the most consistent gains. CLAN also performs slightly better (+0.41%), suggesting that category-level alignment contributes positively even under a large domain shift. In contrast, AdvEnt shows negligible change, and PyCDA significantly underperforms, dropping 2.41% mIoU below the baseline. This performance degradation indicates that approaches relying on pseudo-label refinement or curriculum learning may struggle in multimodal imagery under disaster scenarios, where spatial layout and damage semantics vary drastically across events. Overall, while UDA methods show some promise, their improvements are relatively minor and not robust across all events, underscoring the difficulty of domain alignment in EO-based damage mapping.

In the one-shot setting, we evaluated several popular SSL methods, as shown in Table 12. With only minimal supervision in the target event, all methods except CPS improved over the one-shot baseline (38.82%). The best-performing method is Mean Teacher (MT), which yields a gain of +1.18%, followed by CCT (+0.67%) and GCT (+0.52%). These results show that simple consistency-based teacher-student frameworks are particularly effective in leveraging unlabeled data under limited supervision, likely due to their robustness to noisy or class-imbalanced targets.

4.6.3 Why is cross-event transfer challenging?

To better understand why cross-event generalization is difficult, we explore two fundamental factors rooted in the nature of real-world disaster response.

1. **Inconsistent damage signatures across events.** Figure 10 presents violin plots of SAR backscatter values for intact, damaged, and destroyed buildings across multiple event pairs. These plots reveal two key observations. First, even within the same disaster category (the first row), the pixel intensity distributions for damaged and destroyed buildings differ significantly between events. This indicates that SAR-based damage signatures are inconsistent across locations, possibly due to differences in urban layout, building materials, or sensor incidence angles. Secondly, the distribution shift becomes more pronounced across different types of disasters (the second row). For example, the signatures of destroyed buildings in wildfires notably differ from those in volcano eruptions, which also differ from floods or hurricanes. These variations reflect fundamental differences in damage mechanisms: buildings burned in wildfires, submerged in floods, or collapsed in earthquakes leave very different patterns in SAR backscatter. Such distributional discrepancies make it extremely difficult for models to generalize from one event to another. A model trained on one disaster might learn class boundaries (e.g., between “Damaged” and “Destroyed”) that do not transfer well to another disaster, especially when the visual and physical properties of damage are fundamentally different.




Figure 10. The comparison of pixel distribution of different categories in SAR images for different events. The first row compares different events with the same disaster type. (a) Explosion. (b) Volcano eruption. (c) Wildfire. (d) Cyclone/Hurricane (e) Earthquake. The second row compares different disaster types. (f) Explosion vs Conflict. (g) Wildfire vs Volcano Eruption. (h) Cyclone vs Flood. (i) Explosion vs Wildfire. (j) Conflict vs Earthquake.

2. **Lack of target domain supervision for model selection.** Another critical challenge in cross-event setup is the absence of labeled target domain samples for model selection. In typical domain adaptation benchmarks in computer vision, a validation set of target domain is available to tune hyperparameters or select the best checkpoint. However, no such validation data can be assumed in disaster response scenarios. Figure 11 illustrates the resulting issue: the gap between the actual best mIoU on the test event and the mIoU of the checkpoint selected using a source domain validation set for several representative models under both zero-shot and one-shot settings. In the zero-shot setting, this performance gap is substantial across most models. This indicates that relying on source domain validation leads to suboptimal model selection due to poor reflection of the target distribution. The one-shot setting helps reduce this gap by enabling limited target-aware selection. However, the issue is not entirely resolved. Even with a few labeled samples, model instability and domain shift still make selection challenging.

Together, these two findings highlight that cross-event transfer is hindered not only by semantic and statistical shifts across disaster types but also by operational constraints that prevent ideal training and tuning. Future studies should consider both aspects, *i.e.*, developing models that are robust to distributional variance and designing selection strategies that do not depend on target supervision, such as self-validation, early stopping heuristics, or domain-agnostic performance proxies (Yang et al., 2023).




Figure 11. Comparison of models’ best performance (mIoU) on test events versus the best checkpoints selected on validation sets under the cross-event transfer setting. Each point represents the performance on a single test event under cross-event transfer. The farther a point lies from the diagonal line, the larger the gap between the model’s selected performance and its true upper bound.

4.7 Evaluation of unsupervised multimodal change detection methods

Unsupervised multimodal change detection (UMCD) plays a crucial role in post-disaster assessment by enabling rapid analysis of affected areas without requiring labeled data or prior model training (Lv et al., 2022). Current studies are mainly conducted on toy datasets with limited geographic diversity and scene complexity. These datasets are for general land cover changes. Whether these methods work in real-world disaster occurrence/building damage scenarios is still unknown. BRIGHT offers a new opportunity to evaluate UMCD methods to provide insights into their robustness and scalability in disaster response scenarios. Six representative UMCD methods (Sun et al., 2021; Chen et al., 2022b, 2023; Han et al., 2024; Sun et al., 2024a; Liu et al., 2025) are evaluated in this work. The experimental setup is described in Appendix H.

Table 13 presents the performance of UMCD methods on BRIGHT and UMCD benchmark datasets. To provide a baseline for reference, we also include a random guessing result, representing the performance floor under the UMCD setup. The UMCD benchmark datasets⁹ are detailed in Figure H1 in Appendix H. Although these methods achieved considerable performance on existing UMCD benchmark datasets, their performance suffers noticeable declines on BRIGHT. For instance, while they achieved F1 scores between 70-85% on existing benchmarks, their F1 scores dropped to 20% on BRIGHT. This dramatic

⁹These benchmark datasets are open access at <http://www-labs.iro.umontreal.ca/~mignotte/>

Table 13. Results of representative unsupervised multimodal change detection methods. KC is the acronym of kappa coefficient. The highest values are highlighted in **bold**, and the second-highest results are highlighted in underline. The accuracies on the UMCD benchmark dataset are the accuracies on the four datasets presented in Figure H1, obtained from their literature. Details of methods and benchmark datasets are presented in Appendix H. The random guessing baseline is included to indicate the performance floor under the UMCD setup. The “-” symbol indicates that the corresponding method did not report results on that dataset in their original publications.

Method	UMCD benchmark datasets			BRIGHT			
	OA	F1	KC	OA	F1	IoU	KC
<i>Random guessing</i>	50.0	8.4 / 6.0 / 11.0 / 11.4	0.0	50.00	7.83	4.08	0.00
IRG-McS (Sun et al., 2021)	98.3 / - / 97.1 / 97.2	80.4 / - / 75.4 / 73.7	79.4 / - / 73.9 / 75.1	90.03	12.65	6.75	7.65
SR-GCAE (Chen et al., 2022b)	98.6 / 98.5 / - / -	82.9 / 77.6 / - / -	82.1 / 76.9 / - / -	77.83	14.35	7.73	5.64
FD-MCD (Chen et al., 2023)	98.2 / 97.8 / - / 96.7	81.4 / 72.2 / - / 73.2	82.3 / 71.1 / - / 71.4	80.96	15.84	8.60	<u>7.94</u>
AOSG (Han et al., 2024)	- / - / - / 96.4	- / - / - / 77.7	- / - / - / 75.9	77.93	10.75	5.68	3.98
AGSCC (Sun et al., 2024a)	98.3 / - / 95.9 / 97.6	78.2 / - / 68.0 / 77.9	77.3 / - / 65.8 / 76.6	<u>88.49</u>	<u>14.82</u>	<u>8.00</u>	9.54
AEKAN (Liu et al., 2025)	98.7 / - / - / 98.3	83.8 / - / - / 84.7	83.1 / - / - / 83.9	81.60	13.09	7.00	3.56

505 performance gap underscores the limitations of current UMCD research and highlights the challenges posed by real disaster scenarios. We identify three primary reasons for this decline:

1. **Limitations in traditional UMCD datasets.** The existing UMCD datasets consist of only a handful of image pairs, often depicting simple land cover changes, such as urban expansion and deforestation, with relatively low-resolution imagery and limited geographic diversity, as shown in Figure H1 in Appendix H. These datasets fail to capture the complexity and variability found in real-world disaster scenarios. In contrast, BRIGHT provides thousands of VHR multimodal image pairs across different types of disasters, significantly increasing the diversity of test cases and making it a more challenging benchmark for UMCD model evaluation.
2. **Interference from non-disaster changes.** Unlike previous UMCD benchmarks, where general land cover changes are the primary detection target, changes such as vegetation growth, water body shifts, and urban development may introduce noise and interfere with damage detection in real-world disaster scenarios. While prior UMCD studies treat such changes as valid detection targets, BRIGHT requires methods to differentiate true building structural damage from irrelevant land cover variations, posing a unique challenge to current approaches.
3. **Problematic evaluation protocol.** Current UMCD research follows a problematic evaluation protocol where models are trained, hyperparameters are tuned, and then validated on the same dataset. This approach can actually lead to overfitting, artificially making the model look very accurate but not generalizable to real-world scenarios. In a real-world disaster scenario, there are no labeled samples available for hyperparameter tuning, which makes the current practice unrealistic. BRIGHT exposes this limitation, as models now need to be trained or tuned before being tested directly on new, unseen data, requiring truly generalizable and adaptive learning strategies.

Table 14. Registration performance of different UMIM methods on the scene of Noto-Earthquake-2024. Note that the offsets here are in meters and are calculated based on the geo-coordinates of control points manually selected by EO experts. N/A means that the method is unable to complete the registration task.

Method	Before (meters)	After (meters)
FLSS (Ye et al., 2017)	16.97	10.01
HOPC (Ye et al., 2019)		9.37
LNIFT (Li et al., 2022)		N/A
SRIF (Li et al., 2023a)		N/A

The results have highlighted the challenges of applying existing UMCD models to real-world disaster scenarios. BRIGHT reveals significant limitations in current methodologies and presents new opportunities for future UMCD research.

4.8 Evaluation of unsupervised multimodal image matching methods

Precise image alignment is a critical prerequisite for any multimodal EO application. BRIGHT offers a unique opportunity to evaluate the performance of existing UMIM algorithms under realistic, large-scale disaster conditions. Due to the lack of real pixel-level ground truth correspondences, we adopt a proxy evaluation strategy using manually selected control points (as shown in Figure C1 in Appendix C) as references. These points were selected by EO experts to represent identifiable and stable features across modalities. While this does not constitute absolute ground truth, using such human-verified correspondences provides a valuable reference. This allows us to assess how closely automated methods approximate human matching ability under multimodal and disaster conditions.

Table 14 presents the quantitative results of four UMIM methods on the Noto-Earthquake-2024 scene. These methods fall into two main categories: feature-based methods (LNIFT (Li et al., 2022) and SRIF (Li et al., 2023a)), which rely on sparse keypoint detection and matching, and area-based methods (FLSS (Ye et al., 2017) and HOPC (Ye et al., 2019)), which operate on local regions to estimate correspondences. A detailed description of each method is provided in Appendix I. The experimental results reveal that traditional feature-based methods fail to achieve successful automatic registration. We attribute this limitation to the large spatial extent of the scenes and the drastic cross-modal differences, which make direct keypoint matching highly error-prone and unreliable.

In contrast, area-based methods, which rely on the similarity in a local region rather than point correspondence at the global level, can partially mitigate registration errors and achieve relatively more stable performance. However, these methods still face a critical limitation: they typically perform local matching within a fixed search window. This strategy is inherently inadequate for handling large displacements, especially in SAR imagery over mountainous or uneven terrain, where terrain-induced distortions can cause substantial pixel shifts that far exceed the local search range. Thus, automatic methods may serve as a useful auxiliary tool for preliminary matching or candidate generation, but human expertise remains indispensable for ensuring precise and reliable alignment in operational scenarios. We hope that BRIGHT and its challenging real-world

scenarios will inspire the development of new UMIM methods that are more robust to large-scale scene variations, sensitive to terrain-induced distortions, and ultimately capable of reducing human efforts in future operational EO-based disaster response workflows.

5 Discussion

5.1 Limitation of BRIGHT

We begin this subsection by acknowledging that the composition of the BRIGHT dataset is fundamentally shaped by practical constraints in data availability. While BRIGHT represents a significant step forward in assembling a large-scale, multimodal, and globally distributed dataset for disaster response, it is important to recognize several inherent limitations. These limitations arise not only from the scarcity of open-access VHR SAR imagery, especially over disaster-affected regions, but also from the challenges of manual annotation and the uneven distribution of events. To provide a clearer picture for potential users, we summarize these constraints in four aspects below.

1. **Registration error.** BRIGHT dataset consists of optical and SAR images covering the same locations. SAR images, in particular, can be distorted and stretched in certain areas. Despite thorough preprocessing, including manual alignment and cross-checking by multiple experts in EO data processing, minor alignment errors may persist, as Table 4 suggests.
2. **Label quality.** The building polygons in BRIGHT were manually annotated by expert annotators. Although manual labeling generally ensures high accuracy, minor errors in polygon boundaries are inevitable due to the complexity of building shapes and the variability in image resolution. These inaccuracies may slightly affect the performance of the models trained on BRIGHT. Furthermore, experts assessed the extent of damage to buildings through visual interpretation of optical EO data. This process is susceptible to occasional misjudgments, contributing to label noise.
3. **Sample and regional imbalance.** Although BRIGHT is rich in geographically diverse data, it has the problem of regional imbalance in the number of labels. Some of its events have more tiles and building numbers, thus, are more dominant in training and evaluation, e.g., Turkey-Earthquake-2023 (1,114 tiles) v.s. Hawaii-Wildfire-2023 (65 tiles) in Table 2. To address this, we used an additional event-level evaluation method. However, the effect of events, which account for a large percentage of the sample, on the model during the training phase is still not negligible. This may affect the generalizability of the trained model in real-world disaster scenarios. In addition, all the disaster events in BRIGHT are located near the equator or in the northern hemisphere, with no events from the southern hemisphere. This spatial bias could potentially limit the applicability of trained models to regions with different building styles, land cover patterns, or SAR imaging characteristics prevalent in the southern part of the globe.
4. **Modality and temporal scope.** The dataset’s scope is defined by two key characteristics of the available data. First, it exclusively utilizes single-polarization SAR imagery. The current version lacks the more informative multi-polarization or dense time-series SAR data, which, if available, could enable more nuanced damage characterization and long-term

recovery monitoring, respectively. Second, the dataset’s temporal coverage is concentrated on events from 2020 onwards.

580 This is a direct consequence of its reliance on modern commercial VHR SAR providers (Capella Space and Umbra), whose open-data initiatives largely commenced around that time.

Overall, despite these limitations, it is the first time such an open multimodal VHR dataset has been constructed for multimodal EO research with a large-scale and diverse disaster context.

5.2 Significance of BRIGHT

585 Delays in both EO data acquisition and damage interpretation workflows often hinder timely disaster response (Ye et al., 2024). Traditional expert-driven building damage mapping is time-consuming and not scalable. While ML and DL methods offer automated alternatives, their effectiveness remains limited by the scope and quality of available training data. Most existing open-source datasets are optical, restricting the models’ operational applicability in adverse weather and low-light conditions. As the first globally distributed multimodal dataset, BRIGHT encompasses pre-event optical images and post-event
590 SAR images. This unique combination overcomes the limitations of optical EO data by enabling models trained on BRIGHT to monitor disaster-stricken areas regardless of weather conditions or daylight. Compared to existing building damage datasets, BRIGHT offers several distinct characteristics: multimodal data, VHR imagery with sub-meter spatial resolution, coverage of five types of natural disasters and two human-made disasters, rich geographic diversity, and open access to the community. Due to these features, BRIGHT is anticipated to serve as a benchmark for many future studies and practical disaster relief
595 applications.

Beyond building damage assessment, BRIGHT can also support several research directions within the EO and vision community. In this work, we have applied BRIGHT to evaluate the performance of several UDA and SSL methods for cross-event transfer. We also demonstrated its applicability to UMCD and UMIM, showcasing its versatility as a benchmark for multiple EO challenges under real-world constraints. Furthermore, the dataset provides a strong foundation for broader multimodal EO
600 research. Its high-quality annotations and geographic diversity directly apply to tasks, such as building footprint extraction, land cover mapping, height estimation, and EO-based visual question answering (VQA). Researchers can also repurpose or extend BRIGHT to create task-specific benchmarks, enabling flexible experimentation across tasks. BRIGHT is also well-positioned to support the development of EO-based foundation models, large-scale pre-trained models designed to generalize across sensors, tasks, and regions (Wang et al., 2023; Hong et al., 2024). Its rich combination of modalities, spatial detail, and contextual di-
605 versity provides the data diversity required to build such general-purpose foundation models. This contribution is significant as the field moves toward creating versatile, scalable AI models that can be applied across different types of EO data and disaster scenarios (Li et al., 2024).

Looking ahead, BRIGHT can be further enhanced by incorporating additional modalities. For example, the inclusion of fully polarimetric SAR data would enable more nuanced damage classification than current single polarization SAR data by
610 characterizing the different scattering properties of building materials and debris. Meanwhile, LiDAR data would offer precise 3D information to directly quantify structural collapse and enable a truly terrain-aware analysis. Future versions may also

help fill current geographic gaps, including southern hemisphere regions, to ensure more globally representative coverage. Ultimately, we envision that BRIGHT, true to its name, will bring even a glimmer of brightness to people in disaster-stricken areas by enabling more prompt and effective disaster response and relief.

615 6 Data and Code Availability

The BRIGHT dataset is available at <https://doi.org/10.5281/zenodo.14619797> (Chen et al., 2025). The code for training and testing benchmark methods (including code related to IEEE GRSS DFC 2025) is accessible at <https://github.com/ChenHongruixuan/BRIGHT>. Models' checkpoints can be downloaded at <https://doi.org/10.5281/zenodo.15349461>.

7 Conclusions

620 In this paper, we introduced BRIGHT, the first globally distributed multimodal dataset with open access to the community, covering 14 natural and human-made disaster events. BRIGHT includes pre-event optical and post-event SAR images with sub-meter spatial resolution. Beyond introducing the dataset, we conducted a comprehensive series of experiments to validate its utility. We benchmarked several state-of-the-art supervised learning models under a standard machine learning data split. Moreover, we extended the evaluation to a cross-event transfer setting, simulating real-world scenarios where no or
625 limited target annotations are available. Furthermore, we assessed the performance of unsupervised domain adaptation, semi-supervised learning methods, unsupervised multimodal change detection, and image matching techniques. The findings serve as performance baselines and provide valuable insights for future research in DL model design for real-world disaster response. BRIGHT is an ongoing project, and we remain committed to continuously enhancing its diversity and quality by incorporating new disaster events and refining the existing data. Our objective is to improve BRIGHT's utility for practical disaster response
630 applications at all levels (regional, national, and international) and research in the community.

Appendix A: Details of disaster events

A1 Explosion in Beirut, 2020

On August 4, 2020, a massive explosion occurred at the Port of Beirut in Lebanon, caused by the improper storage of 2,750 tons of ammonium nitrate. The explosion caused widespread damage within a radius of several kilometers, significantly impacting
635 the port and surrounding neighborhoods, including areas such as Gemmayzeh, Mar Mikhael, and Achrafieh. It resulted in 218 deaths and more than 7,000 injuries, and left approximately 300,000 people homeless (Fakih and Majzoub, 2021). Economic losses were estimated to be 15 billion USD (Valsamos et al., 2021). The disaster compounded Lebanon's ongoing economic challenges and contributed to political instability and social unrest.

A2 Explosion in Bata, 2021

640 On March 7, 2021, a series of four explosions occurred at the Cuartel Militar de Nkoantoma in Bata, Equatorial Guinea, caused by improperly stored explosives. The blasts led to at least 107 deaths and over 615 injuries, and widespread destruction throughout the city (OCHA, 2021a). A total of 243 structures were destroyed or severely damaged, displacing many residents. Around 150 families sought refuge in temporary shelters, while others stayed with relatives. Local hospitals treated more than 500 injured individuals, and the economic impact was severe, underscoring the dangers associated with the improper storage
645 of hazardous materials.

A3 Volcano Eruption in DR Congo and Rwanda, 2021

On May 22, 2021, Mount Nyiragongo in the Democratic Republic of the Congo erupted, causing widespread devastation. The eruption resulted in 32 deaths and the destruction of 1,000 homes. The displacement of thousands as lava flows threatened the city of Goma (IFRC, 2021). Nearly 400,000 people were evacuated due to the risk of further volcanic activity, including
650 potential magma flow beneath Goma and nearby Lake Kivu. Despite continued seismic activity, life in Goma largely returned to normal by August 2021. However, plans to relocate parts of the city remain under consideration due to the ongoing threat from the volcano.

A4 Earthquake in Haiti, 2021

On August 14, 2021, a magnitude 7.2 earthquake struck Haiti's Tiburon Peninsula, primarily affecting the Nippes, Sud, and
655 Grand'Anse departments. The disaster caused over 2,200 deaths and more than 12,200 injuries, and left thousands homeless. The economic losses were significant, estimated at over USD 1.5 billion. Approximately 137,500 buildings, including homes, schools, and hospitals, were damaged or destroyed (OCHA, 2021b). As the deadliest natural disaster of 2021, the earthquake exacerbated Haiti's existing challenges, including widespread poverty and political instability.

A5 Volcano Eruption in La Palma, 2021

660 On September 19, 2021, the Cumbre Vieja volcano on La Palma, part of Spain's Canary Islands, erupted following several days of seismic activity. The eruption primarily impacted the island's western side, covering over 1,000 hectares with lava and destroying more than 3,000 buildings, including the towns of Todoque and La Laguna. The lava flow, measuring about 3.5 kilometers wide and 6.2 kilometers long, reached the sea, cutting off the coastal highway and forming a new peninsula with extensive lava tubes. Although the timely evacuation of around 8,000 people prevented major casualties, one person died
665 from inhaling toxic gases (Troll et al., 2024). It caused significant damage to arable land and affected livelihoods, displacing thousands of residents. Economic losses exceeded 800 million Euros.

A6 Wildfire in Colorado, 2021

The Marshall Fire, which started on December 30, 2021, in Boulder County, Colorado, became the most destructive wildfire in the State's history in terms of destroyed buildings. Fueled by dry grass from an unusually warm and dry season, and winds up to 185 km/h, the fire killed two people and injured eight. It destroyed 1,084 structures, including homes, a hotel, and a shopping center, causing over 2 billion USD in damage (Forrister et al., 2024). More than 37,500 residents were evacuated, and significant damage to public drinking water systems occurred (Forrister et al., 2024).

A7 Armed Conflict in Ukraine, 2022

In February 2022, Russian forces launched a full-scale invasion of Ukraine, resulting in widespread destruction and displacement. By November 2024, total damages to Ukraine's infrastructure had reached \$170 billion. More than 236,000 residential buildings were damaged or destroyed, including 209,000 private houses, 27,000 apartment buildings, and 600 dormitories. Over 4,000 educational institutions and 1,554 medical facilities were also affected, with extensive damage to transport, energy, and telecommunications infrastructure. The conflict has inflicted more than 40,000 civilian casualties, displaced four million people internally, and forced 6.8 million to flee. As of late 2024, approximately 14.6 million Ukrainians require humanitarian assistance (Andrienko et al., 2025).

A8 Earthquake in Turkey, 2023

On February 6, 2023, a magnitude 7.8 earthquake struck southeastern Turkey near Gaziantep, followed by a magnitude 7.7 aftershock. The disaster, the most powerful earthquake in Turkey since 1939, caused widespread destruction across approximately 350,000 km^2 , affecting 14 million people and displacing 1.5 million. The death toll reached 53,537, with 107,213 injuries (STL, 2024), making it one of the deadliest earthquakes in modern history. Economic losses were estimated at 148.8 billion USD, with over 518,000 houses and 345,000 apartments destroyed (Government of Türkiye, 2023). The earthquake caused severe damage to infrastructure, agriculture, and essential services, further worsening the region's economic challenges. International aid was mobilized to support the affected populations.

A9 Cyclone in Myanmar, 2023

In May 2023, Cyclone Mocha, a Category 5 hurricane, struck Myanmar, causing widespread devastation in the country's coastal regions, particularly in Rakhine State. According to official estimates, at least 148 people were killed and 132 injured, although other sources suggest higher figures (GFDRR and World Bank, 2023). The cyclone affected around 1.2 million people in Rakhine alone, with over 200,000 buildings reportedly damaged or destroyed, making it one of the most destructive cyclones in the region in the past 15 years. The direct economic damage was estimated at 2.24 billion USD in damages, equivalent to 3.4% of Myanmar's GDP (GFDRR and World Bank, 2023).

A10 Wildfire in Hawaii, 2023

In August 2023, a series of wildfires broke out on the island of Maui, Hawaii, causing widespread destruction and significant impacts on the local population and environment. The fires, fueled by dry conditions and strong winds, primarily affected the town of Lahaina, where at least 102 people were killed and two remain missing (Hedayati et al., 2024). Over 2,200 buildings
700 were destroyed, including many historic landmarks, resulting in estimated damages of 5.5 billion USD (Jones et al., 2024; NCEI, 2025). The fires prompted evacuations and led to the displacement of thousands of residents, with significant economic losses in the tourism and agriculture sectors.

A11 Earthquake in Morocco, 2023

On September 8, 2023, a 6.9 magnitude earthquake struck Morocco’s Al Haouz Province near Marrakesh, causing widespread
705 devastation. The earthquake killed nearly 3,000 people, injured more than 5,500, and displaced more than half a million (OCHA, 2023). The earthquake damaged or destroyed nearly 60,000 houses, with the heaviest losses reported in rural communities of the Atlas Mountains (International Federation of Red Cross and Red Crescent Societies, 2024). Overall damage is estimated at about USD 7 billion (National Geophysical Data Center / World Data Service, 2025), while direct economic losses of roughly USD 30 million amount to approximately 0.24% of Morocco’s GDP. Cultural heritage in Marrakesh was also
710 hard-hit, sections of the UNESCO-listed Medina and several historic mosques sustained severe damage.

A12 Flood and Storm in Libya, 2023

In September 2023, Storm Daniel brought catastrophic flooding to northeastern Libya, particularly in Derna, when two dams collapsed and released an estimated 30 million cubic meters of water. At least 5,923 people were killed, though local officials warned fatalities could reach 18,000–20,000 (USAID Bureau for Humanitarian Assistance, 2024). Post–event analyses revealed
715 that approximately 10% of Derna’s housing stock was destroyed and a further 18.5% sustained damage. Across the wider coastal belt—from Benghazi, through Jabal Al Akhdar and Al Marj, to Derna—an estimated 18,838 dwellings were damaged or obliterated (Normand and Heggy, 2024). The disaster, considered the second deadliest dam failure in history, destroyed the city’s infrastructure; four bridges collapsed, and entire neighborhoods were washed out to sea. Long-standing neglect of dam maintenance, compounded by Libya’s political turmoil, contributed significantly to the scale of the tragedy.

720 A13 Hurricane in Mexico, 2023

In September 2023, Hurricane Norma, a Category 4 hurricane, struck the western coast of Mexico, severely affecting Sinaloa and Baja California Sur. This was followed by Hurricane Otis in October, which made landfall near Acapulco as a Category 5 hurricane. Otis was the strongest Pacific hurricane to hit Mexico, causing at least 52 deaths and leaving 32 missing. The storm caused unprecedented destruction, with more than 51,864 homes destroyed and damages estimated at 12-16 billion USD,
725 surpassing Hurricane Wilma as the costliest Mexican hurricane (Reinhart and Reinhart, 2024).

A14 Earthquake in Noto, 2024

On January 1, 2024, a magnitude 7.5 earthquake struck the Noto Peninsula in Ishikawa Prefecture, Japan, reaching a maximum JMA seismic intensity of Shindo 7. The earthquake caused widespread destruction, particularly in the towns of Suzu, Wajima, Noto, and Anamizu, and triggered a 7.45-meter tsunami along the coast of the Sea of Japan. The disaster resulted in 572 deaths and over 1,300 injuries, with 193,529 structures damaged across nine prefectures (Fire and Disaster Management Agency, 2025). It was the deadliest earthquake in Japan since the 2011 Tōhoku disaster. As of February 20 2024, 12,929 people remained in 521 evacuation centers (Japanese Red Cross Society, 2024). The event prompted Japan’s first major tsunami warning since 2011. Total damage is estimated at 16.9 trillion Japanese Yen, reflecting substantial economic losses with significant impacts on infrastructure and communities across the affected regions (Kaneko, 2024).

Appendix B: Feature visualization

For the feature visualisation of buildings in different events in Figure 5-(c), we specifically employ DINOv2 (Oquab et al., 2024) to extract the high-dimensional features of buildings in different events. The corresponding reference maps are used to exclude background features. The features are then visualized using the t-SNE algorithm (van der Maaten and Hinton, 2008).

To visualize the spatial focus of models during building damage mapping, we apply Grad-CAM technique¹⁰, a gradient-based technique for generating class-specific activation maps (Selvaraju et al., 2017). We perform Grad-CAM analysis on the ChangeOS model. Specifically, we compute the activation responses for the “Damaged” and “Destroyed” categories from the following components: the third or fourth ResNet block of the two branches of encoder, capturing mid- to high-level modality-specific features; the fusion module of the damage classification branch in the decoder, reflecting early-stage fusion and decision-making. These activations help reveal which spatial regions the model relies on when assessing different damage levels and how it utilizes information from optical and SAR modalities differently, depending on the disaster context.

Appendix C: Manual registration and estimating registration errors

We performed the manual registration process using QGIS, with the “Georeferencer” plugin to align SAR images to the optical imagery as the reference. The transformation type was set to “Thin Plate Spline”, and “Lanczos resampling (6×6 kernels)” was applied to achieve high-quality interpolation. The manually selected control points by EO experts on some disaster scenes are shown in Figure C1.

Despite multiple rounds of meticulous registration and cross-validation by several EO experts, registration errors between optical and SAR images cannot be completely eliminated. Since we cannot directly obtain ground truth (i.e., actual ground control points), we propose a proxy method for registration error estimation:

¹⁰<https://github.com/jacobgil/pytorch-grad-cam>




Figure C1. Overview of manually selected control points used for registration on some events. The number in parentheses of each event is the number of control points selected on that view of the image. The labeled yellow boxes highlight specific zoom-in regions that are shown in below. Data source: SAR images of Bata-Explosion-2020 and La Palma-Volcano-2020 are sourced from Capella Space Open Data Gallery while the SAR image of Noto-Earthquake-2024 is from Umbra Open Data Program.

1. Feature point selection: Select many feature points (e.g., 3000 points) from both the optical and SAR tiles using a keypoint detection algorithm.
2. Feature extraction: Use validated and well-performing multimodal registration descriptors (Li et al., 2020, 2022, 2023a) to extract modality-independent features for the selected points.

Table C1. The estimated proxy registration error (in pixel) for each event.

Event	RMSE (All)	RMSE (Building)
Beirut-EP-2020	1.365	1.365
Bata-EP-2021	1.226	1.160
Goma-VE-2021	0.966	0.975
Les Cayes-EQ-2021	1.392	1.377
La Palma-VE-2021	0.976	0.944
Marshall-WF-2021	1.263	1.266
Ukraine-AC-2022	0.899	0.882
Turkey-EQ-2023	0.937	0.901
Kyaukpyu-CC-2023	0.970	0.998
Hawaii-WF-2023	1.124	1.124
Morocco-EQ-2023	1.088	1.050
Derna-FL-2023	1.267	1.255
Acapulco-HC-2023	1.166	1.148
Noto-EQ-2024	1.072	1.043

3. Feature matching: Perform feature matching between the extracted feature points across the optical and SAR images.
4. Registration error computation: After completing the matching process, compute the pixel distances between the registered point pairs. These distances serve as a proxy for registration error.
5. Outlier removal: To improve the reliability of the estimation, we exclude matched points with a root mean square error (RMSE) greater than a certain threshold (e.g., 20 pixels). This threshold is set because after multiple rounds of expert corrections and cross-checking, large registration errors have already been eliminated. Points exceeding this threshold can be considered outliers.
- By following this method, we can estimate the registration errors between optical and SAR images in the dataset, providing a reliable foundation for further damage assessment analysis.

Appendix D: Details of benchmark deep learning models

Here, we present some implementation details of our baseline model.

- For UNet, the encoder has five convolutional blocks, each consisting of two 3×3 convolutional layers and two batch normalization layers. The number of channels in the five blocks are 64, 128, 256, 512, and 1024, respectively. The decoder contains four convolutional blocks, and the number of channels is set accordingly to the number of channels of the skip-connected features extracted by the encoder.

- For DeepLabV3+, we applied ResNet-50 initialized with ImageNet pretraining weights (He et al., 2016) as the encoder. We then modified the input channel of the encoder so that the network can predict building damage maps directly from the optical and SAR stacked images.
- For SiamAttnUNet, we followed the settings in its original literature (Adriano et al., 2021).
- For SiamCRNN, we applied pseudo-siamese ResNet-18 as encoder and four ConvLSTM layers (Shi et al., 2015) with 3×3 convolutional kernel size and a hidden dimension of 128, as the decoder.
- For ChangeOS, we used the ResNet-18 initialized with ImageNet pretraining weights as the encoder. We modified the encoder to a pseudo-siamese structure to extract features from pre-event and post-event images with different modalities.
- For DamageFormer, we used pseudo-siamese Swin-Transformer-Tiny (Liu et al., 2021) as the encoder instead of pure-siamese MixFormer (Xie et al., 2021) as in the original literature (Chen et al., 2022a).
- For the ChangeMamba family of models (Chen et al., 2024), we chose one of them, MambaBDA-Tiny, for our experiments. We modified the encoder to a pseudo-siamese structure to extract features from pre-event and post-event images with different modalities.

Appendix E: Details of post-processing methods

To further refine the raw predictions from DL models and enhance the quality of building damage maps, we apply three post-processing techniques: test-time augmentation (TTA), object-based majority voting, and model ensembling in Section 4.5. This section provides an overview of these methods and details the specific implementation used in our study.

- Test-time augmentation is a widely used technique to improve model robustness by applying transformations to the input images at inference time. Instead of making a single prediction per image, multiple augmented versions of the same input are passed through the network, and the resulting predictions are aggregated. This reduces model sensitivity to spatial variations and increases prediction stability. In our study, we apply rotation (90° , 180° , 270°) and horizontal/vertical flipping to the input images during inference. The network produces a set of predictions for each augmented version, which are then aggregated at the logit level by summing the outputs before applying the softmax function.
- Object-based majority voting aims to enforce spatial consistency by considering entire building instances rather than making independent pixel-wise predictions. Instead of classifying each pixel separately, the final label is determined based on a majority vote across the entire building object, leading to more coherent and reliable results. We follow the setup from (Zheng et al., 2021), where each building is treated as an independent object, and a weighted majority voting scheme is applied. The weighting is determined by the inverse of class proportions in the training set, meaning that underrepresented damage categories (e.g., destroyed buildings) have higher voting weights compared to dominant classes (e.g., intact buildings).

Table F1. The mIoU on different events for different UDA methods in Table 11.

Events	AdaptSeg	AdvEnt	CLAN	PyCDA	FDA
Beirut-EP-2020	35.55	33.07	37.25	32.15	35.89
Bata-EP-2021	37.54	35.20	36.98	30.01	36.31
Goma-VE-2021	52.49	51.54	56.91	49.67	52.46
Les Cayes-EQ-2021	37.11	36.86	39.47	33.40	33.90
La Palma-VE-2021	30.54	30.04	30.10	32.76	30.50
Marshall-WF-2021	37.20	35.01	34.35	30.14	33.93
Ukraine-AC-2022	32.36	32.65	32.21	32.92	35.03
Turkey-EQ-2023	33.69	34.54	36.68	34.58	36.82
Kyaukpyu-CC-2023	34.62	29.22	30.83	31.31	36.71
Hawaii-WF-2023	36.52	42.17	33.36	34.30	36.19
Morocco-EQ-2023	37.93	37.42	39.04	30.55	38.76
Derna-FL-2023	36.69	35.31	35.24	37.31	38.53
Acapulco-HC-2023	27.05	27.32	27.73	30.11	31.68
Noto-EQ-2024	35.49	35.65	33.99	25.44	31.39
Average	36.05	35.43	36.01	33.19	36.29

- 805 – Model ensembling is a technique that combines predictions from multiple models to reduce uncertainty and improve generalization. By leveraging diverse models with different architectures, ensembling helps smooth out individual model biases and enhances the overall robustness of predictions. We perform ensembling by combining the results of three top-performing models: ChangeOS (Zheng et al., 2021), DamageFormer (Chen et al., 2022a), and ChangeMamba (Chen et al., 2024). Each model independently generates a damage proxy map, and the final prediction is obtained by averaging their logits before applying the softmax function.

Appendix F: Details of unsupervised domain adaptation methods

810 To evaluate the effectiveness of UDA techniques in the cross-event zero-shot setting, we selected several representative methods from the computer vision literature originally designed for semantic segmentation tasks. The following methods are implemented:

- AdaptSeg (Tsai et al., 2018): an adversarial training framework that aligns output space distributions between source and target domains.
- 815 – AdvEnt (Vu et al., 2019): a variant of adversarial adaptation focusing on entropy minimization and structured output alignment.
- CLAN (Luo et al., 2019): a category-level alignment method that selectively aligns class-specific features.

- PyCDA (Lian et al., 2019): a curriculum domain adaptation framework that iteratively refines pseudo-labels across domains.
- 820 – FDA (Yang and Soatto, 2020): a simple style-transfer-based approach that transfers source images into the target domain appearance via frequency mixing.

All methods are implemented using DeepLabV3+ as the backbone architecture. This choice is motivated by the fact that DeepLab-based networks are commonly used in the official implementations of these UDA methods, and many codebases are tightly coupled with DeepLab’s structure. Therefore, adopting DeepLabV3+ as the shared backbone ensures compatibility with
825 existing implementations and avoids the need for extensive code modification. It also provides a fair and consistent basis for comparison across all methods.

In our experimental setup, each test disaster event is treated as the target domain, and only its unlabeled optical-SAR image pairs are used for adaptation. Importantly, no labeled target domain data is used for training or model selection, in line with our zero-shot assumption. This setup is fundamentally different from conventional UDA protocols in the computer vision
830 community, where a target domain validation set is typically available to select the best model checkpoint, and the target test set is held out purely for evaluation. In contrast, our setting is designed to reflect real-world disaster response scenarios, where it is realistic to obtain unlabeled imagery, but no ground-truth annotations are available at adaptation time. Therefore, model selection must rely solely on source domain feedback, which introduces additional challenges and more closely aligns with the operational constraints of emergency response.

835 The hyperparameters for each method closely follow the original papers to preserve method fidelity. Below are additional implementation details specific to certain methods:

- FDA: For each training iteration, we generate frequency-transferred source images using FDA and mix them 1:1 with the original source images. The combined dataset is used to train the model from scratch.
- PyCDA: Training proceeds in two stages. First, the model is trained on the source domain for 30,000 iterations. In
840 the second stage, the adaptation module is applied to the target domain using pseudo-labeling and uncertainty-aware refinement as proposed in the original work.

For all methods, the same data preprocessing pipeline and training schedule (e.g., learning rate, batch size) are used unless otherwise specified in the original method, like those supervised models. The final evaluation is conducted on the target domain, and performance is reported in terms of mIoU.

845 **Appendix G: Details of semi-supervised learning methods**

In the one-shot cross-event transfer setting, we evaluate several representative SSL methods for semantic segmentation. These methods were originally developed for natural image domains and are adapted here to the multimodal EO-based building damage assessment scenario.

Table G1. The mIoU on different events for different semi-supervised learning methods in Table 12.

Events	MT	CCT	GCT	CPS
Beirut-EP-2020	32.91	35.68	35.48	35.46
Bata-EP-2021	37.94	36.61	38.25	36.96
Goma-VE-2021	56.46	57.68	53.52	47.49
Les Cayes-EQ-2021	40.20	39.57	38.83	39.41
La Palma-VE-2021	33.52	33.03	36.90	33.29
Marshall-WF-2021	39.73	40.12	38.74	32.94
Ukraine-AC-2022	39.11	36.34	36.44	31.34
Turkey-EQ-2023	38.96	39.72	36.50	37.15
Kyaukpyu-CC-2023	36.72	35.84	33.56	30.22
Hawaii-WF-2023	54.10	48.93	49.99	47.64
Morocco-EQ-2023	37.18	40.54	41.10	37.54
Derna-FL-2023	37.55	41.07	40.45	40.28
Acapulco-HC-2023	32.30	30.16	33.36	31.51
Noto-EQ-2024	43.40	37.61	37.74	37.65
Average	40.00	39.49	39.34	37.06

- Mean Teacher (MT) (Tarvainen and Valpola, 2017) maintains an exponential moving average (EMA) of model weights to generate stable predictions for unlabeled data, enforcing consistency between the student and teacher outputs.
- CCT (Ouali et al., 2020) applies perturbations to the input data and enforces consistency between different views of the same image using multiple decoder branches.
- GCT (Ke et al., 2020) further enhances this idea by introducing guidance from a pre-trained model to regularize learning.
- CPS (Chen et al., 2021) trains two separate networks that generate pseudo-labels for each other, encouraging cross-supervised learning on unlabeled data.

All methods also use DeepLabV3+ as the backbone network. This decision is consistent with their original implementations and allows for seamless integration with official codebases without the need for architectural modifications. It also ensures a fair and consistent comparison across methods.

In the setup, 13 labeled disaster events serve as the training and validation sets, while a single labeled sample from the test (target) event is provided to simulate a realistic one-shot adaptation scenario. In addition, the remaining unlabeled samples from the test event are made available and used to facilitate SSL.

Each model is trained in two stages. In the first stage, the model is trained for 10,000 iterations using only standard supervised losses, including cross-entropy and Lovász-Softmax on the labeled data from the source domain and the single labeled sample from the target domain. This stage serves to initialize the model with a stable representation before semi-supervised training




Figure H1. Illustration of benchmark datasets commonly used in unsupervised multimodal change detection. The order from left to right is the corresponding order in the Table 13. In reference maps, white color indicates changed pixels, and black color indicates unchanged pixels. Shuguang, River, and Italy datasets are sourced from (Zhang et al., 2016), while the Gloucester dataset is from (Prendes et al., 2015).

865 begins. In the second stage, each SSL method applies its respective semi-supervised objective using the unlabeled data from the target domain. These include consistency regularization (e.g., Mean Teacher), cross-view perturbation training (e.g., CCT and GCT), or dual-network pseudo supervision (e.g., CPS). Hyperparameters for each method follow the configurations suggested in their original publications.

Final performance is evaluated on the target disaster event using mIoU as the primary evaluation metric.

870 Appendix H: Unsupervised multimodal change detection: methods and evaluation protocol

H1 Evaluated UMCD methods

We further evaluate several recent UMCD methods on BRIGHT. These methods are designed to detect binary changes from paired multimodal imagery without relying on manual labels.

- *Random guessing* assumes each pixel has an equal probability (50%) of being classified as changed or unchanged. This baseline provides a reference performance floor for evaluating the relative effectiveness of other UMCD methods.

875

- IRG-McS (Sun et al., 2021) constructs superpixel-based graphs for each modality and iteratively matches their structural relationships to detect changes across heterogeneous image pairs, followed by Markov co-segmentation for final change map generation.
- 880 – SR-GCAE (Chen et al., 2022b) learns robust graph-based representations of local and nonlocal structural relationships in multimodal images using a graph convolutional autoencoder, enabling effective change detection without supervision.
- FD-MCD (Chen et al., 2023) proposes a Fourier domain framework that analyzes both local and nonlocal structural relationships using graph spectral convolution and adaptive fusion, enabling robust unsupervised multimodal change detection.
- 885 – AOSG (Han et al., 2024) constructs an adaptively optimized structured graph to capture patch-level structural features in multimodal images, iteratively refining change intensity measures by fusing self-change and cross-domain structural differences for accurate unsupervised change detection.
- AGSCC (Sun et al., 2024a) translates multimodal images via a structure cycle-consistent image regression framework that enforces similarity in structural graphs across domains, using adaptive graphs and multiple regularization terms to robustly detect changes without supervision.
- 890 – AEKAN (Liu et al., 2025) utilizes a superpixel-based Siamese AutoEncoder built on Kolmogorov–Arnold Networks (KAN) to extract latent commonality features between multimodal images, using reconstruction and hierarchical consistency losses to detect changes in an unsupervised manner.

H2 A more practical evaluation protocol on BRIGHT

The training and evaluation protocols commonly adopted in the UMCD literature are often limited to individual image pairs
 895 (shown in Figure H1), using the same scene for representation learning, hyperparameter tuning, and evaluation. While this is understandable due to the lack of large-scale public datasets in this field, such a setup fails to reflect real-world use cases and often leads to overfitting and overestimated performance. To address this, we introduce a standardized and practical evaluation protocol using the BRIGHT dataset. Specifically, we use the validation set from the standard ML split in Section 2.4 as the training set for UMCD methods, including any hyperparameter or threshold tuning. These models are then evaluated on the
 900 combined training and test sets, which are strictly held out during adaptation. This avoids overlap between tuning and evaluation scenes and better simulates the real-world deployment setting, where a model is expected to generalize to new, unseen data.

For evaluation, it is important to note that most of these methods were originally developed for land cover change detection, where changes correspond to transitions between semantic categories such as vegetation, water, or built-up areas. Applying them directly to disaster damage detection poses certain challenges. In particular, “Damaged” buildings do not always exhibit
 905 strong spatial or spectral signals, making them hard to distinguish in an unsupervised setting. However, “Destroyed” buildings often result in a complete change of land cover appearance (e.g., collapsed structures, debris, or ground exposure), which aligns better with the assumptions of existing UMCD methods. Therefore, in our evaluation, we treat buildings labeled as “Destroyed”

in BRIGHT as the positive (changed) class, and all other regions (including intact, damaged, and background) as the unchanged class. This ensures fair adaptation to disaster scenarios while respecting the original design of these methods.

910 Appendix I: Details of unsupervised multimodal image matching methods

In our extended evaluation of UMIM methods, we categorized the selected algorithms into two groups based on their matching strategies: feature-based and area-based methods.

Feature-based methods, such as LNIFT (Li et al., 2022) and SRIF (Li et al., 2023a), aim to detect sparse keypoints independently in both optical and SAR images and compute modality-invariant feature descriptors. The matching is then performed
915 globally by comparing descriptor similarity between the two modalities. These methods are implemented using publicly available code¹¹.

In contrast, area-based methods, including FLSS (Ye et al., 2017) and HOPC (Ye et al., 2019), begin by identifying keypoints only in the optical image. These points are then matched by sliding a template window across a local search region in the SAR image, relying on local patch similarity. In our experiments, we used a template size of 120×120 pixels, and a search window
920 of 200×200 pixels for both methods. This category is also implemented using open-source code¹².

We performed registration on the entire original EO image pair from the Noto-Earthquake-2024 event. Because of the large spatial size of the input images, images were downsampled to half their original resolution to improve computational efficiency and maintain feasibility for all methods.

For quantitative evaluation, we adopted a control-point-based proxy metric. A set of manually selected control points, validated by EO experts shown in Figure C1, is used to represent stable, cross-modal, and clearly identifiable features (e.g., building corners, road intersections). The average spatial offset is computed using their coordinates in a common projected coordinate system (e.g., UTM), both before and after registration. Before registration, the offset can be calculated as:

$$E_{\text{before}} = \frac{1}{N} \sum_{i=1}^N \sqrt{(x_i^{\text{opt}} - x_i^{\text{sar}})^2 + (y_i^{\text{opt}} - y_i^{\text{sar}})^2}. \quad (I1)$$

where $(x_i^{\text{opt}}, y_i^{\text{opt}})$ and $(x_i^{\text{sar}}, y_i^{\text{sar}})$ are the projected coordinates of the i -th control point in the optical and SAR images, respectively.
930

After registration, SAR points are transformed using the estimated mapping $\mathcal{T}$, and the post-registration error is computed as:

$$E_{\text{after}} = \frac{1}{N} \sum_{i=1}^N \sqrt{(x_i^{\text{opt}} - \hat{x}_i^{\text{sar}})^2 + (y_i^{\text{opt}} - \hat{y}_i^{\text{sar}})^2}, \quad (I2)$$

where $(\hat{x}_i^{\text{sar}}, \hat{y}_i^{\text{sar}}) = \mathcal{T}(x_i^{\text{sar}}, y_i^{\text{sar}})$ are transformed SAR coordinates.

This approach provides an alternative approximation of alignment accuracy in the absence of ground truth. Although this
935 metric is not a substitute for full correspondence maps, it offers a practical and interpretable measure of registration perfor-

¹¹<https://github.com/LJY-RS/SRIF>

¹²<https://github.com/yeyuanxin110/CFOG>

mance, and enables us to assess how well unsupervised methods can approach human-level matching under complex, real-world disaster conditions.

Author contributions. **HC**: conceptualization (lead), data curation (lead), funding acquisition (equal), methodology (lead), project administration (lead), investigation (lead), software (lead), visualization (lead), writing – original draft preparation (lead), writing – review and editing (equal). **JS**: conceptualization (support), data curation (equal), funding acquisition (support), methodology (support), investigation (support), software (support), writing – original draft preparation (equal), visualization (support). **OD**: conceptualization (support), data curation (equal), software (support), visualization (equal), writing – review and editing (support). **CBB**: conceptualization (support), data curation (support), writing – review and editing (lead). **WX**: data curation (support), project administration (support), writing – review and editing (equal). **JW**: data curation (support), writing – review and editing (support). **XS**: funding acquisition (support), writing – review and editing (support). **YW**: writing – review and editing (support). **JX**: conceptualization (support), data curation (support), supervision (support), writing – review and editing (support). **CL**¹³: conceptualization (support), supervision (support), writing – linguistic refinement (equal). **KS**: conceptualization (equal), resources (support), supervision (equal), writing – review and editing (equal). **NY**: conceptualization (equal), funding acquisition (lead), resources (lead), supervision (lead), writing – review and editing (equal).

Competing interests. The contact author has declared that none of the authors has any competing interests.

Acknowledgements. This work was supported in part by the JSPS, KAKENHI under Grant Number 24KJ0652 and 22H03609, the Council for Science, Technology and Innovation (CSTI), the Cross-ministerial Strategic Innovation Promotion Program (SIP), Development of a Resilient Smart Network System against Natural Disasters (Funding agency: NIED), JST, FOREST under Grant Number JPMJFR206S, Microsoft Research Asia, Next Generation AI Research Center of The University of Tokyo, and Young Researchers Exchange Programme between Japan and Switzerland under the Japanese-Swiss Science and Technology Programme.

The authors would also like to give special thanks to Sarah Preston of Capella Space, Capella Space’s Open Data Gallery, Maxar Open Data Program, and Umbra’s Open Data Program for providing the valuable data.

¹³Cuiling Lan did not participate in any activities related to the acquisition, processing, utilization, or distribution of the datasets. Her access to the data was solely limited to the information presented in the current paper, similar to that of the readers.

References

- Adriano, B., Xia, J., Baier, G., Yokoya, N., and Koshimura, S.: Multi-Source Data Fusion Based on Ensemble Learning for Rapid Building
960 Damage Mapping during the 2018 Sulawesi Earthquake and Tsunami in Palu, Indonesia, *Remote Sensing*, 11, 2019.
- Adriano, B., Yokoya, N., Xia, J., Miura, H., Liu, W., Matsuoka, M., and Koshimura, S.: Learning from multimodal and multitemporal earth
observation data for building damage mapping, *ISPRS J. Photogramm.*, 175, 132–143, 2021.
- Amirreza Shaban, Shray Bansal, Z. L. I. E. and Boots, B.: One-Shot Learning for Semantic Segmentation, in: *Proceedings of the British
Machine Vision Conference (BMVC)*, pp. 167.1–167.13, 2017.
- 965 Andrienko, D., Goriunov, D., Grudova, V., Markuts, J., Marshalok, T., Neyter, R., Piddubnyi, I., Studennikova, I., and Topolskov, D.: Report
on Damages to Infrastructure from the Destruction Caused by Russia’s Military Aggression Against Ukraine as of November 2024, Tech.
rep., Kyiv School of Economics (KSE) Institute, 2025.
- Arciniegas, G. A., Bijker, W., Kerle, N., and Tolpekin, V. A.: Coherence- and Amplitude-Based Analysis of Seismogenic Damage in Bam,
Iran, Using ENVISAT ASAR Data, *IEEE T. Geosci. Remote*, 45, 1571–1581, 2007.
- 970 Artés, T., Oom, D., de Rigo, D., Durrant, T. H., Maianti, P., Libertà, G., and San-Miguel-Ayanz, J.: A global wildfire dataset for the analysis
of fire regimes and fire behaviour, *Scientific Data*, 6, 296, 2019.
- Bai, Y., Adriano, B., Mas, E., and Koshimura, S.: Machine Learning Based Building Damage Mapping from the ALOS-2/PALSAR-2 SAR
Imagery: Case Study of 2016 Kumamoto Earthquake, *Journal of Disaster Research*, 12, 646–655, 2017.
- Bai, Y., Gao, C., Singh, S., Koch, M., Adriano, B., Mas, E., and Koshimura, S.: A Framework of Rapid Regional Tsunami Damage Recog-
975 nition From Post-event TerraSAR-X Imagery Using Deep Neural Networks, *IEEE Geosci. Remote S.*, 15, 43–47, 2018.
- Bastos Moroz, C. and Thieken, A. H.: Urban growth and spatial segregation increase disaster risk: lessons learned from the 2023 disaster on
the North Coast of São Paulo, Brazil, *Natural Hazards and Earth System Sciences*, 24, 3299–3314, 2024.
- Berman, M., Triki, A. R., and Blaschko, M. B.: The Lovász-Softmax Loss: A Tractable Surrogate for the Optimization of the Intersection-
Over-Union Measure in Neural Networks, in: *Proc. CVPR IEEE*, 2018.
- 980 Bonafilia, D., Tellman, B., Anderson, T., and Issenberg, E.: Sen1Floods11: A Georeferenced Dataset to Train and Test Deep Learning Flood
Algorithms for Sentinel-1, in: *IEEE Comput. Soc. Conf.*, 2020.
- Brett, P. T. B. and Guida, R.: Earthquake Damage Detection in Urban Areas Using Curvilinear Features, *IEEE T. Geosci. Remote*, 51,
4877–4884, 2013.
- Brunner, D., Lemoine, G., and Bruzzone, L.: Earthquake damage assessment of buildings using VHR optical and SAR imagery, *IEEE T.*
985 *Geosci. Remote*, 48, 2403–2420, 2010.
- Chen, H., Wu, C., Du, B., Zhang, L., and Wang, L.: Change Detection in Multisource VHR Images via Deep Siamese Convolutional
Multiple-Layers Recurrent Neural Network, *IEEE T. Geosci. Remote*, 58, 2848–2864, 2020.
- Chen, H., Nemni, E., Vallecorsa, S., Li, X., Wu, C., and Bromley, L.: Dual-Tasks Siamese Transformer Framework for Building Damage
Assessment, in: *International Geoscience and Remote Sensing Symposium (IGARSS)*, pp. 1600–1603, 2022a.
- 990 Chen, H., Yokoya, N., Wu, C., and Du, B.: Unsupervised Multimodal Change Detection Based on Structural Relationship Graph Represen-
tation Learning, *IEEE T. Geosci. Remote*, pp. 1–18, 2022b.
- Chen, H., Yokoya, N., and Chini, M.: Fourier domain structural relationship analysis for unsupervised multimodal change detection, *ISPRS
J. Photogramm.*, 198, 99–114, 2023.

- Chen, H., Song, J., Han, C., Xia, J., and Yokoya, N.: ChangeMamba: Remote Sensing Change Detection With Spatiotemporal State Space Model, *IEEE T. Geosci. Remote*, 62, 1–20, 2024.
- Chen, H., Song, J., Dietrich, O., Broni-Bediako, C., Xuan, W., Wang, J., Shao, X., Yimin, W., Xia, J., Lan, C., Schindler, K., and Yokoya, N.: BRIGHT: A globally distributed multimodal building damage assessment dataset with very-high-resolution for all-weather disaster response, <https://doi.org/10.5281/zenodo.15335889>, 2025.
- Chen, L.-C., Zhu, Y., Papandreou, G., Schroff, F., and Adam, H.: Encoder-Decoder with Atrous Separable Convolution for Semantic Image Segmentation, in: *Computer Vision – ECCV 2018*, pp. 833–851, 2018.
- Chen, S.-W. and Sato, M.: Tsunami Damage Investigation of Built-Up Areas Using Multitemporal Spaceborne Full Polarimetric SAR Images, *IEEE T. Geosci. Remote*, 51, 1985–1997, 2013.
- Chen, X., Yuan, Y., Zeng, G., and Wang, J.: Semi-Supervised Semantic Segmentation with Cross Pseudo Supervision, in: *Proc. CVPR IEEE*, pp. 2613–2622, 2021.
- Cheng, C.-S., Behzadan, A. H., and Noshadravan, A.: Deep learning for post-hurricane aerial damage assessment of buildings, *Computer-Aided Civil and Infrastructure Engineering*, 36, 695–710, 2021.
- Chini, M., Pierdicca, N., and Emery, W. J.: Exploiting SAR and VHR Optical Images to Quantify Damage Caused by the 2003 Bam Earthquake, *IEEE T. Geosci. Remote*, 47, 145–152, 2009.
- Chini, M., Anniballe, R., Bignami, C., Pierdicca, N., Mori, S., and Stramondo, S.: Identification of building double-bounces feature in very high resolution SAR data for earthquake damage mapping, in: *Int. Geosci. Remote Se.*, pp. 2723–2726, 2015.
- Dietrich, O., Peters, T., Garnot, V. S. F., Sticher, V., Whelan, T. T.-T., Schindler, K., and Wegner, J. D.: An Open-Source Tool for Mapping War Destruction at Scale in Ukraine Using Sentinel-1 Time Series, *Communications Earth & Environment*, 6, 215, 2025.
- Fakih, L. and Majzoub, A.: "They Killed Us from the Inside": An Investigation Into the August 4 Beirut Blast, Human Rights Watch, 2021.
- Fan, X., Nie, G., Deng, Y., An, J., Zhou, J., Xia, C., and Pang, X.: Estimating earthquake-damage areas using Landsat-8 OLI surface reflectance data, *International Journal of Disaster Risk Reduction*, 33, 275–283, 2019.
- Fire and Disaster Management Agency: Damage Caused by the 2024 Noto Peninsula Earthquake and Response Status of Fire Departments, etc. (117th Report), Situation Report 117, Fire and Disaster Management Agency, Tokyo, Japan, 2025.
- Forrister, A., Kuligowski, E. D., Sun, Y., Yan, X., Lovreglio, R., Cova, T. J., and Zhao, X.: Analyzing Risk Perception, Evacuation Decision and Delay Time: A Case Study of the 2021 Marshall Fire in Colorado, *Travel Behaviour and Society*, 35, 100 729, 2024.
- Freire, S., Santos, T., Navarro, A., Soares, F., Silva, J., Afonso, N., Fonseca, A., and Tenedório, J.: Introducing mapping standards in the quality assessment of buildings extracted from very high resolution satellite imagery, *ISPRS J. Photogramm.*, 90, 1–9, 2014.
- Fujita, A., Sakurada, K., Imaizumi, T., Ito, R., Hikosaka, S., and Nakamura, R.: Damage detection from aerial images via convolutional neural networks, in: *2017 Fifteenth IAPR International Conference on Machine Vision Applications (MVA)*, pp. 5–8, 2017.
- Ge, P., Gokon, H., Meguro, K., and Koshimura, S.: Study on the Intensity and Coherence Information of High-Resolution ALOS-2 SAR Images for Rapid Massive Landslide Mapping at a Pixel Level, *Remote Sensing*, 11, 2019.
- Ge, P., Gokon, H., and Meguro, K.: A review on synthetic aperture radar-based building damage assessment in disasters, *Remote Sens. Environ.*, 240, 111 693, 2020.
- Gerard, S., Borne-Pons, P., and Sullivan, J.: A simple, strong baseline for building damage detection on the xBD dataset, *arXiv preprint arXiv:2401.17271*, 2024.
- GFDRR and World Bank: Extremely Severe Cyclonic Storm Mocha, May 2023, Myanmar: Global Rapid Post-Disaster Damage Estimation (GRADE) Report, Tech. rep., Global Facility for Disaster Reduction and Recovery (GFDRR) and World Bank, 2023.

- Ghorbanzadeh, O., Xu, Y., Ghamisi, P., Kopp, M., and Kreil, D.: Landslide4Sense: Reference Benchmark Data and Deep Learning Models for Landslide Detection, *IEEE T. Geosci. Remote*, 60, 1–17, 2022.
- Government of Türkiye: Türkiye Earthquakes Recovery and Reconstruction Assessment, Tech. rep., Government of Türkiye, <https://reliefweb.int/report/turkiye/turkiye-earthquakes-recovery-and-reconstruction-assessment>, 2023.
- Guo, H., Su, X., Wu, C., Du, B., and Zhang, L.: SAAN: Similarity-Aware Attention Flow Network for Change Detection With VHR Remote Sensing Images, *IEEE T. Image Process.*, 33, 2599–2613, 2024.
- Gupta, R. and Shah, M.: RescueNet: Joint Building Segmentation and Damage Assessment from Satellite Imagery, in: 2020 25th International Conference on Pattern Recognition (ICPR), pp. 4405–4411, 2021.
- Gupta, R., Goodman, B., Patel, N., Hosfelt, R., Sajeev, S., Heim, E., Doshi, J., Lucas, K., Choset, H., and Gaston, M.: Creating xBD: A Dataset for Assessing Building Damage from Satellite Imagery, in: *IEEE Comput. Soc. Conf.*, 2019.
- Han, T., Tang, Y., Zou, B., and Feng, H.: Unsupervised multimodal change detection based on adaptive optimization of structured graph, *Int. J. Appl. Earth Obs.*, 126, 103 630, 2024.
- He, K., Zhang, X., Ren, S., and Sun, J.: Deep residual learning for image recognition, *Proc. CVPR IEEE*, pp. 770–778, 2016.
- He, K., Shen, X., and Anagnostou, E. N.: A global forest burn severity dataset from Landsat imagery (2003–2016), *Earth Syst. Sci. Data*, 16, 3061–3081, 2024.
- Hedayati, F., Monroy, X., Sluder, E., Fallahian, H., and Shabanian, M.: The 2023 Lahaina Conflagration, <https://ibhs.org/wp-content/uploads/FINAL-Lahaina-Conflagration.pdf>, 2024.
- Hong, D., Zhang, B., Li, X., Li, Y., Li, C., Yao, J., Yokoya, N., Li, H., Ghamisi, P., Jia, X., Plaza, A., Gamba, P., Benediktsson, J. A., and Chanussot, J.: SpectralGPT: Spectral Remote Sensing Foundation Model, *IEEE T. Pattern Anal.*, 46, 5227–5244, 2024.
- Huot, F., Hu, R. L., Goyal, N., Sankar, T., Ihme, M., and Chen, Y.-F.: Next Day Wildfire Spread: A Machine Learning Dataset to Predict Wildfire Spreading From Remote-Sensing Data, *IEEE T. Geosci. Remote*, 60, 1–13, 2022.
- IFRC: Emergency Appeal: Democratic Republic of Congo and Rwanda, Report, International Federation of Red Cross and Red Crescent Societies (IFRC), <https://reliefweb.int/report/democratic-republic-congo/emergency-appeal-democratic-republic-congo-and-rwanda-mount>, 2021.
- International Federation of Red Cross and Red Crescent Societies: Morocco Earthquake 2023: Operation Update #4 (12 months) (MDRMA010), Situation Report MDRMA010OU5, International Federation of Red Cross and Red Crescent Societies, Geneva, <https://reliefweb.int/report/morocco/morocco-earthquake-2023-operation-update-4-12-months-mdrma010>, published 16 November 2024; accessed 18 April 2025, 2024.
- Japanese Red Cross Society: Operation Update No. 30: 2024 Noto Peninsula Earthquake—The Japanese Red Cross Society’s Response, <https://www.jrc.or.jp/english/relief/2024NotoPeninsulaEarthquake.html>, web page, 2024.
- Jones, M. W., Kelley, D. I., Burton, C. A., Di Giuseppe, F., Barbosa, M. L. F., Brambleby, E., Hartley, A. J., Lombardi, A., Mataveli, G., McNorton, J. R., Spuler, F. R., Wessel, J. B., Abatzoglou, J. T., Anderson, L. O., Andela, N., Archibald, S., Armenteras, D., Burke, E., Carmenta, R., Chuvieco, E., Clarke, H., Doerr, S. H., Fernandes, P. M., Giglio, L., Hamilton, D. S., Hantson, S., Harris, S., Jain, P., Kolden, C. A., Kurvits, T., Lampe, S., Meier, S., New, S., Parrington, M., Perron, M. M. G., Qu, Y., Ribeiro, N. S., Saharjo, B. H., San-Miguel-Ayanz, J., Shuman, J. K., Tanpipat, V., van der Werf, G. R., Veraverbeke, S., and Xanthopoulos, G.: State of Wildfires 2023–2024, *Earth Syst. Sci. Data*, 16, 3601–3685, 2024.

- Kaneko, Y.: Prioritizing Recovery: Expected Roles of Public Support during the Noto Peninsula Earthquake, Research Report of the Research Center for Urban Safety and Security, Kobe University, 28, 31–39, http://www.rcuss.kobe-u.ac.jp/publication/Year2024/pdfEach28/28_05.pdf, article in English; RCUSS Research Report No. 28, 2024.
- Karimzadeh, S. and Mastuoka, M.: Building Damage Assessment Using Multisensor Dual-Polarized Synthetic Aperture Radar Data for the 2016 M 6.2 Amatrice Earthquake, Italy, *Remote Sensing*, 9, 2017.
- Kaur, N., Lee, C.-C., Mostafavi, A., and Mahdavi-Amiri, A.: Large-scale building damage assessment using a novel hierarchical transformer architecture on satellite images, *Computer-Aided Civil and Infrastructure Engineering*, 38, 2072–2091, 2023.
- Ke, Z., Qiu, D., Li, K., Yan, Q., and Lau, R. W. H.: Guided Collaborative Training for Pixel-Wise Semi-Supervised Learning, in: *European Conference Computer Vision*, p. 429–445, 2020.
- Kreibich, H., Loon, A. F. V., Schröter, K., Ward, P. J., Mazzoleni, M., Sairam, N., Abeshu, G. W., Agafonova, S., AghaKouchak, A., Aksoy, H., Alvarez-Garretón, C., Aznar, B., Balkhi, L., Barendrecht, M. H., Biancamaria, S., Bos-Burgering, L., Bradley, C., Budiyo, Y., Buytaert, W., Capewell, L., Carlson, H., Cavus, Y., Couasnon, A., Coxon, G., Daliakopoulos, I., de Ruiter, M. C., Delus, C., Erfurt, M., Esposito, G., François, D., Frappart, F., Freer, J., Frolova, N., Gain, A. K., Grillakis, M., Grima, J. O., Guzmán, D. A., Huning, L. S., Ionita, M., Kharlamov, M., Khoi, D. N., Kieboom, N., Kireeva, M., Koutroulis, A., Lavado-Casimiro, W., Li, H.-Y., Llasat, M. C., Macdonald, D., Mård, J., Mathew-Richards, H., McKenzie, A., Mejia, A., Mendiola, E. M., Mens, M., Mobini, S., Mohor, G. S., Nagavciuc, V., Ngo-Duc, T., Huynh, T. T. N., Nhi, P. T. T., Petrucci, O., Nguyen, H. Q., Quintana-Seguí, P., Razavi, S., Ridolfi, E., Riegel, J., Sadik, M. S., Savelli, E., Sazonov, A., Sharma, S., Sörensen, J., Souza, F. A. A., Stahl, K., Steinhausen, M., Stoelzle, M., Szalińska, W., Tang, Q., Tian, F., Tokarczyk, T., Tovar, C., Tran, T. V. T., Huijgevoort, M. H. J. V., van Vliet, M. T. H., Vorogushyn, S., Wagener, T., Wang, Y., Wendt, D. E., Wickham, E., Yang, L., Zambrano-Bigiarini, M., Blöschl, G., and Baldassarre, G. D.: The challenge of unprecedented floods and droughts in risk management, *Nature*, 608, 80–86, 2022.
- Li, C., Hong, D., Zhang, B., Liao, T., Yokoya, N., Ghamisi, P., Chen, M., Wang, L., Benediktsson, J. A., and Chanussot, J.: Interpretable foundation models as decryptors peering into the Earth system, *The Innovation*, 5, 100682, 2024.
- Li, J., Hu, Q., and Ai, M.: RIFT: Multi-Modal Image Matching Based on Radiation-Variation Insensitive Feature Transform, *IEEE T. Image Process.*, 29, 3296–3310, 2020.
- Li, J., Xu, W., Shi, P., Zhang, Y., and Hu, Q.: LNIFT: Locally Normalized Image for Rotation Invariant Multimodal Feature Matching, *IEEE T. Geosci. Remote*, 60, 1–14, 2022.
- Li, J., Hu, Q., and Zhang, Y.: Multimodal image matching: A scale-invariant algorithm and an open dataset, *ISPRS J. Photogramm.*, 204, 77–88, 2023a.
- Li, T., Wang, C., Zhang, H., Wu, F., and Zheng, X.: DDFormer: A Dual-Domain Transformer for Building Damage Detection Using High-Resolution SAR Imagery, *IEEE Geosci. Remote S.*, 20, 1–5, 2023b.
- Lian, Q., Duan, L., Lv, F., and Gong, B.: Constructing Self-Motivated Pyramid Curriculum for Cross-Domain Semantic Segmentation: A Non-Adversarial Approach, in: *IEEE I. Conc. Comp. Vis.*, pp. 6757–6766, 2019.
- Liu, T., Xu, J., Lei, T., Wang, Y., Du, X., Zhang, W., Lv, Z., and Gong, M.: AEKAN: Exploring Superpixel-Based AutoEncoder Kolmogorov-Arnold Network for Unsupervised Multimodal Change Detection, *IEEE T. Geosci. Remote*, 63, 1–14, 2025.
- Liu, W. and Yamazaki, F.: Extraction of Collapsed Buildings in the 2016 Kumamoto Earthquake Using Multi-Temporal PALSAR-2 Data, *Journal of Disaster Research*, 12, 241–250, <https://doi.org/10.20965/jdr.2017.p0241>, 2017.
- Liu, W., Yamazaki, F., Gokon, H., and ichi Koshimura, S.: Extraction of Tsunami-Flooded Areas and Damaged Buildings in the 2011 Tohoku-Oki Earthquake from TerraSAR-X Intensity Images, *Earthquake Spectra*, 29, 183–200, 2013.

- Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., and Guo, B.: Swin Transformer: Hierarchical Vision Transformer using Shifted Windows, in: IEEE I. Conc. Comp. Vis., pp. 9992–10 002, 2021.
- Loshchilov, I. and Hutter, F.: Decoupled weight decay regularization, arXiv preprint arXiv:1711.05101, 2017.
- Lu, W., Wei, L., and Nguyen, M.: Bitemporal Attention Transformer for Building Change Detection and Building Damage Assessment, 1110 IEEE J-STARS., 17, 4917–4935, 2024.
- Luo, Y., Zheng, L., Guan, T., Yu, J., and Yang, Y.: Taking a Closer Look at Domain Shift: Category-Level Adversaries for Semantics Consistent Domain Adaptation, in: Proc. CVPR IEEE, pp. 2502–2511, 2019.
- Lv, Z., Huang, H., Li, X., Zhao, M., Benediktsson, J. A., Sun, W., and Falco, N.: Land Cover Change Detection With Heterogeneous Remote Sensing Images: Review, Progress, and Perspective, P. IEEE, 110, 1976–1991, 2022.
- 1115 Manzini, T., Perali, P., Karnik, R., and Murphy, R.: Crasar-u-droids: A large scale benchmark dataset for building alignment and damage assessment in georectified suas imagery, arXiv preprint arXiv:2407.17673, 2024.
- Matsuoka, M. and Yamazaki, F.: Building Damage Mapping of the 2003 Bam, Iran, Earthquake Using Envisat/ASAR Intensity Imagery, Earthquake Spectra, 21, 285–294, 2005.
- Matsuoka, M. and Yamazaki, F.: Comparative analysis for detecting areas with building damage from several destructive earthquakes using 1120 satellite synthetic aperture radar images, Journal of Applied Remote Sensing, 4, 041 867, 2010.
- Matsuoka, M., Koshimura, S., and Nojima, N.: Estimation of building damage ratio due to earthquakes and tsunamis using satellite SAR imagery, in: Int. Geosci. Remote Se., pp. 3347–3349, 2010.
- Meena, S. R., Nava, L., Bhuyan, K., Puliero, S., Soares, L. P., Dias, H. C., Floris, M., and Catani, F.: HR-GLDD: a globally distributed dataset using generalized deep learning (DL) for rapid landslide mapping on high-resolution (HR) satellite imagery, Earth Syst. Sci. Data, 1125 15, 3283–3298, 2023.
- National Geophysical Data Center / World Data Service: NCEI/WDS Global Significant Earthquake Database, <https://doi.org/10.7289/V5TD9V7K>, accessed: 2025-04-18, 2025.
- NCEI: U.S. Billion-Dollar Weather and Climate Disasters, Tech. rep., NOAA National Centers for Environmental Information (NCEI), <https://www.ncei.noaa.gov/access/billions/>, 2025.
- 1130 Nguyen, D. T., Ofli, F., Imran, M., and Mitra, P.: Damage Assessment from Social Media Imagery Data During Disasters, in: Proceedings of the 2017 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining 2017, p. 569–576, 2017.
- Normand, J. C. and Heggy, E.: Assessing flash flood erosion following storm Daniel in Libya, Nature Communications, 15, 6493, 2024.
- OCHA: Equatorial Guinea: Bata Explosions - Flash Update No. 3 (As of 11 March 2021), Report, United Nations Office for the Coordination of Humanitarian Affairs (OCHA), [https://reliefweb.int/report/equatorial-guinea/](https://reliefweb.int/report/equatorial-guinea/equatorial-guinea-bata-explosions-flash-update-no3-11-march-2021) 1135 [equatorial-guinea-bata-explosions-flash-update-no3-11-march-2021](https://reliefweb.int/report/equatorial-guinea/equatorial-guinea-bata-explosions-flash-update-no3-11-march-2021), 2021a.
- OCHA: Haiti: Earthquake Situation Report No. 4, Report, United Nations Office for the Coordination of Humanitarian Affairs (OCHA), <https://reliefweb.int/report/haiti/haiti-earthquake-situation-report-no-4-7-september-2021>, 2021b.
- OCHA: Morocco Earthquake: Situation Report #5, Tech. rep., United Nations Office for the Coordination of Humanitarian Affairs (OCHA), <https://reliefweb.int/report/morocco/morocco-earthquake-situation-report-5-november-14-2023>, 2023.
- 1140 Oquab, M., Darcet, T., Moutakanni, T., Vo, H. V., Szafraniec, M., Khalidov, V., Fernandez, P., HAZIZA, D., Massa, F., El-Nouby, A., Assran, M., Ballas, N., Galuba, W., Howes, R., Huang, P.-Y., Li, S.-W., Misra, I., Rabbat, M., Sharma, V., Synnaeve, G., Xu, H., Jegou, H., Mairal, J., Labatut, P., Joulin, A., and Bojanowski, P.: DINOv2: Learning Robust Visual Features without Supervision, Transactions on Machine Learning Research, 2024.

- Ouali, Y., Hudelot, C., and Tami, M.: Semi-Supervised Semantic Segmentation With Cross-Consistency Training, in: Proc. CVPR IEEE, pp. 12 671–12 681, 2020.
- Prendes, J., Chabert, M., Pascal, F., Giros, A., and Tourneret, J.-Y.: Change detection for optical and radar images using a Bayesian nonparametric model coupled with a Markov random field, in: 2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 1513–1517, 2015.
- Rahnemounfar, M., Chowdhury, T., Sarkar, A., Varshney, D., Yari, M., and Murphy, R. R.: FloodNet: A High Resolution Aerial Imagery Dataset for Post Flood Scene Understanding, IEEE Access, 9, 89 644–89 654, 2021.
- Rahnemounfar, M., Chowdhury, T., and Murphy, R.: RescueNet: A High Resolution UAV Semantic Segmentation Dataset for Natural Disaster Damage Assessment, Scientific Data, 10, 913, <https://doi.org/10.1038/s41597-023-02799-4>, 2023.
- Reinhart, B. J. and Reinhart, A.: Hurricane Otis (EP182023): Tropical Cyclone Report, Tropical Cyclone Report EP182023, National Hurricane Center, National Oceanic and Atmospheric Administration, Miami, FL, https://www.nhc.noaa.gov/data/tcr/EP182023_Otis.pdf, issued 7 March 2024; PDF archived and retrieved the same day, 2024.
- Ronneberger, O., Fischer, P., and Brox, T.: U-Net: Convolutional Networks for Biomedical Image Segmentation, in: Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015, pp. 234–241, 2015.
- Sandhini Putri, A. F., Widyatmanti, W., and Umarhadi, D. A.: Sentinel-1 and Sentinel-2 data fusion to distinguish building damage level of the 2018 Lombok Earthquake, Remote Sensing Applications: Society and Environment, 26, 100 724, 2022.
- Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., and Batra, D.: Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization, in: 2017 IEEE International Conference on Computer Vision (ICCV), pp. 618–626, 2017.
- Shen, Y., Zhu, S., Yang, T., Chen, C., Pan, D., Chen, J., Xiao, L., and Du, Q.: BDANet: Multiscale Convolutional Neural Network With Cross-Directional Attention for Building Damage Assessment From Satellite Images, IEEE T. Geosci. Remote, 60, 1–14, 2022.
- Shi, X., Chen, Z., Wang, H., Yeung, D.-Y., Wong, W.-k., and Woo, W.-c.: Convolutional LSTM Network: a machine learning approach for precipitation nowcasting, in: Adv. Neur. In., vol. 1, p. 802–810, 2015.
- STL: Turkey Earthquake: Emergency Situation Report (05.02.2024), Tech. rep., Support to Life (STL), <https://reliefweb.int/report/turkiye/turkey-earthquake-emergency-situation-report-05022024>, 2024.
- Stramondo, S., Bignami, C., Chini, M., Pierdicca, N., and Tertulliani, A.: Satellite radar and optical remote sensing for earthquake damage detection: results from different case studies, Int. J. Remote Sens., 27, 4433–4447, 2006.
- Sun, Y., Lei, L., Guan, D., and Kuang, G.: Iterative Robust Graph for Unsupervised Change Detection of Heterogeneous Remote Sensing Images, IEEE Trans. Image Process., 30, 6277–6291, 2021.
- Sun, Y., Lei, L., Guan, D., Wu, J., Kuang, G., and Liu, L.: Image Regression With Structure Cycle Consistency for Heterogeneous Change Detection, IEEE T. Neur. Net. Lear., 35, 1613–1627, 2024a.
- Sun, Y., Wang, Y., and Eineder, M.: QuickQuakeBuildings: Post-Earthquake SAR-Optical Dataset for Quick Damaged-Building Detection, IEEE Geosci. Remote S., 21, 1–5, 2024b.
- Tarvainen, A. and Valpola, H.: Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results, in: Adv. Neur. In., edited by Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R., vol. 30, 2017.
- Tong, X., Hong, Z., Liu, S., Zhang, X., Xie, H., Li, Z., Yang, S., Wang, W., and Bao, F.: Building-damage detection using pre- and post-seismic high-resolution satellite stereo imagery: A case study of the May 2008 Wenchuan earthquake, ISPRS J. Photogramm., 68, 13–27, 2012.

- Troll, V. R., Aulinas, M., Carracedo, J. C., Geiger, H., Perez-Torrado, F. J., Soler, V., Deegan, F. M., Bloszies, C., Weis, F., Albert, H., Gisbert, G., Day, J. M. D., Rodríguez-Gonzalez, A., Gazel, E., and Dayton, K.: The 2021 La Palma eruption: social dilemmas resulting from life close to an active volcano, *Geology Today*, 40, 96–111, 2024.
- 1185 Tsai, Y.-H., Hung, W.-C., Schuster, S., Sohn, K., Yang, M.-H., and Chandraker, M.: Learning to Adapt Structured Output Space for Semantic Segmentation, in: 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 7472–7481, 2018.
- UNDRR: Economic Losses, Poverty & Disasters: 1998-2017, Report, United Nations Office for Disaster Risk Reduction (UNDRR), <https://www.undrr.org/publication/economic-losses-poverty-disasters-1998-2017>, 2018a.
- UNDRR: Words into Action guideline: Man-made/technological hazards, Report, United Nations Office for Disaster Risk Reduction (UN-
1190 DRR), <https://www.undrr.org/publication/words-action-guideline-man-made/technological-hazards>, 2018b.
- USAID Bureau for Humanitarian Assistance: Libya Assistance Overview, Fact sheet, U.S. Agency for International Development, Washington, DC, <https://reliefweb.int/report/libya/libya-assistance-overview-april-2024>, pDF fact sheet (posted 17 April 2024), 2024.
- Valsamos, G., Larcher, M., and Casadei, F.: Beirut explosion 2020: A case study for a large-scale urban blast simulation, *Safety Science*, 137, 105 190, 2021.
- 1195 van der Maaten, L. and Hinton, G.: Visualizing Data using t-SNE, *J. Mach. Learn. Res.*, 9, 2579–2605, 2008.
- Vescovo, R., Adriano, B., Wiguna, S., Ho, C. Y., Morales, J., Dong, X., Ishii, S., Wako, K., Ezaki, Y., Mizutani, A., Mas, E., Tanaka, S., and Koshimura, S.: The 2024 Noto Peninsula earthquake building damage dataset: Multi-source visual assessment, *Earth System Science Data Discussions*, 2025, 1–24, <https://doi.org/10.5194/essd-2024-363>, 2025.
- Vu, T.-H., Jain, H., Bucher, M., Cord, M., and Pérez, P.: ADVENT: Adversarial Entropy Minimization for Domain Adaptation in Semantic
1200 Segmentation, in: Proc. CVPR IEEE, pp. 2512–2521, 2019.
- Wang, D., Zhang, Q., Xu, Y., Zhang, J., Du, B., Tao, D., and Zhang, L.: Advancing Plain Vision Transformer Toward Remote Sensing Foundation Model, *IEEE T. Geosci. Remote*, 61, 1–15, 2023.
- Wang, J., Guo, H., Su, X., Zheng, L., and Yuan, Q.: PCDASNet: Position-Constrained Differential Attention Siamese Network for Building Damage Assessment, *IEEE T. Geosci. Remote*, 62, 1–18, 2024.
- 1205 Wang, T.-L. and Jin, Y.-Q.: Postearthquake Building Damage Assessment Using Multi-Mutual Information From Pre-Event Optical Image and Postevent SAR Image, *IEEE Geosci. Remote S.*, 9, 452–456, 2012.
- Wang, Y., Yao, Q., Kwok, J. T., and Ni, L. M.: Generalizing from a Few Examples: A Survey on Few-shot Learning, *ACM Comput. Surv.*, 53, 2020.
- Watanabe, M., Thapa, R. B., Ohsumi, T., Fujiwara, H., Yonezawa, C., Tomii, N., and Suzuki, S.: Detection of damaged urban areas using
1210 interferometric SAR coherence change with PALSAR-2, *Earth, Planets and Space*, 68, 131, 2016.
- Xia, H., Wu, J., Yao, J., Zhu, H., Gong, A., Yang, J., Hu, L., and Mo, F.: A Deep Learning Application for Building Damage Assessment Using Ultra-High-Resolution Remote Sensing Imagery in Turkey Earthquake, *Int. J. Disast. Risk Sc.*, 14, 947–962, 2023.
- Xia, J., Chen, H., Broni-Bediako, C., Wei, Y., Song, J., and Yokoya, N.: OpenEarthMap-SAR: A Benchmark Synthetic Aperture Radar Dataset for Global High-Resolution Land Cover Mapping, *arXiv preprint arXiv:2501.10891*, 2025.
- 1215 Xie, E., Wang, W., Yu, Z., Anandkumar, A., Alvarez, J. M., and Luo, P.: SegFormer: Simple and Efficient Design for Semantic Segmentation with Transformers, in: *Adv. Neur. In.*, vol. 34, pp. 12 077–12 090, 2021.
- Xie, S., Duan, J., Liu, S., Dai, Q., Liu, W., Ma, Y., Guo, R., and Ma, C.: Crowdsourcing Rapid Assessment of Collapsed Buildings Early after the Earthquake Based on Aerial Remote Sensing Image: A Case Study of Yushu Earthquake, *Remote Sensing*, 8, 2016.

- Xue, Z., Zhang, X., Prevatt, D. O., Bridge, J., Xu, S., and Zhao, X.: Post-hurricane building damage assessment using street-view imagery and structured data: A multi-modal deep learning approach, arXiv preprint arXiv:2404.07399, 2024.
- 1220 Yamaguchi, Y.: Disaster Monitoring by Fully Polarimetric SAR Data Acquired With ALOS-PALSAR, P. IEEE, 100, 2851–2860, 2012.
- Yamazaki, F. and Matsuoka, M.: Remote Sensing Technologies in Post-Disaster Damage Assessment, Journal of Earthquake and Tsunami, 01, 193–210, 2007.
- Yang, J., Qian, H., Xu, Y., Wang, K., and Xie, L.: Can We Evaluate Domain Adaptation Models Without Target-Domain Labels?, arXiv preprint arXiv:2305.18712, 2023.
- 1225 Yang, Y. and Soatto, S.: FDA: Fourier Domain Adaptation for Semantic Segmentation, in: Proc. CVPR IEEE, pp. 4084–4094, 2020.
- Ye, S., Zhu, Z., and Suh, J. W.: Leveraging past information and machine learning to accelerate land disturbance monitoring, Remote Sens. Environ., 305, 114 071, 2024.
- Ye, Y., Shen, L., Hao, M., Wang, J., and Xu, Z.: Robust Optical-to-SAR Image Matching Based on Shape Properties, IEEE Geosci. Remote Sens., 14, 564–568, 2017.
- 1230 Ye, Y., Bruzzone, L., Shan, J., Bovolo, F., and Zhu, Q.: Fast and Robust Matching for Multimodal Remote Sensing Image Registration, IEEE T. Geosci. Remote, 57, 9059–9070, 2019.
- Yonezawa, C. and Takeuchi, S.: Decorrelation of SAR data by urban damages caused by the 1995 Hyogoken-nanbu earthquake, Int. J. Remote Sens., 22, 1585–1600, 2001.
- 1235 Yusuf, Y., Matsuoka, M., and Yamazaki, F.: Damage assessment after 2001 Gujarat earthquake using Landsat-7 satellite images, Journal of the Indian Society of Remote Sensing, 29, 17–22, 2001.
- Zhang, J., Liu, K., and Wang, M.: Flood detection using Gravity Recovery and Climate Experiment (GRACE) terrestrial water storage and extreme precipitation data, Earth Syst. Sci. Data, 15, 521–540, 2023.
- Zhang, P., Gong, M., Su, L., Liu, J., and Li, Z.: Change detection based on deep feature representation and mapping transformation for multi-spatial-resolution remote sensing images, ISPRS J. Photogramm., 116, 24–41, 2016.
- 1240 Zheng, Z., Zhong, Y., Wang, J., Ma, A., and Zhang, L.: Building damage assessment for rapid disaster response with a deep object-based semantic change detection framework: From natural disasters to man-made disasters, Remote Sens. Environ., 265, 112 636, 2021.
- Zheng, Z., Zhong, Y., Zhao, J., Ma, A., and Zhang, L.: Unifying remote sensing change detection via deep probabilistic change models: From principles, models to applications, ISPRS J. Photogramm., 215, 239–255, 2024.