

# The Potential of Copernicus Satellites for Disaster Response: Retrieving Building Damage from Sentinel-1 and Sentinel-2

Olivier Dietrich<sup>1</sup>, Merlin Alfredsson<sup>1</sup>, Emilia Arens<sup>2</sup>, Nando Metzger<sup>1</sup>,  
Torben Peters<sup>1</sup>, Linus Scheibenreif<sup>1</sup>, Jan Dirk Wegner<sup>2</sup>, Konrad Schindler<sup>1</sup>

<sup>1</sup> Photogrammetry & Remote Sensing Lab, ETH Zurich

<sup>2</sup> EcoVision Lab, Department of Mathematical Modeling and Machine Learning (DM3L), University of Zurich

**Keywords:** Deep Learning, Satellite Images, Building Damage Assessment, Disaster Response

## Abstract

Natural disasters demand rapid damage assessment to guide humanitarian response. Here, we investigate whether medium-resolution Earth observation images from the Copernicus program can support building damage assessment, complementing very-high resolution imagery with often limited availability. We introduce xBD-S12, a dataset of 10,315 pre- and post-disaster image pairs from both Sentinel-1 and Sentinel-2, spatially and temporally aligned with the established xBD benchmark. In a series of experiments, we demonstrate that building damage can be detected and mapped rather well in many disaster scenarios, despite the moderate 10 m ground sampling distance. We also find that, for damage mapping at that resolution, architectural sophistication does not seem to bring much advantage: more complex model architectures tend to struggle with generalization to unseen disasters, and geospatial foundation models bring little practical benefit. Our results suggest that Copernicus images are a viable data source for rapid, wide-area damage assessment and could play an important role alongside VHR imagery. We release the xBD-S12 dataset, code, and trained models to support further research at <https://github.com/prs-eth/xbd-s12>.

## 1. Introduction

Large-scale natural disasters like earthquakes, floods, and wildfires pose a persistent threat to large parts of humanity. In 2025 alone, events like the monsoon floods in Pakistan (United Nations, 2025b), wildfires in California (McKoy, 2025), and hurricanes in the Caribbean (United Nations, 2025a) have devastated entire communities in hours and have caused hundreds of casualties. Moreover, several types of disasters are induced by weather extremes and are intensifying as a consequence of climate change (Van Aalst, 2006; Banholzer et al., 2014; IPCC, 2023). Independent of their cause, all such events mandate an immediate humanitarian response. Yet, effective support requires information about the location and extent of damages. In recent years, satellite imagery has emerged as an important information source for disaster management. Beyond rapidly providing an overview of the affected region, very-high resolution (VHR) images with ground sampling distances  $\leq 2$  m have been shown to enable the identification of damaged assets, especially buildings, complementing sources on the ground (Kawasaki et al., 2013; Rolla et al., 2025).

Manually identifying damaged buildings in VHR imagery remains prohibitively time-consuming for fast disaster response. The machine learning revolution has naturally motivated efforts to automate this task (Sun et al., 2020; Braik and Koliou, 2024). Early approaches focused on single-hazard scenarios, e.g., tsunamis (Fujita et al., 2017), floods (Rudner et al., 2019), earthquakes (Xu et al., 2019), hurricanes (Cao and Choe, 2020), or wildfires (Galanis et al., 2021), successfully showing that deep neural networks can detect hazard-specific damage patterns in satellite imagery. Important progress came with the xView2 challenge and the accompanying xBD dataset (Gupta et al., 2019a,b), which enabled multi-hazard damage assessment for the first time. This benchmark catalyzed extensive research into automated damage assessment methodologies. The competition's winning solution employed a Siamese U-Net architecture with multiple encoder backbones, decomposing the



Figure 1. Copernicus satellites are surprisingly effective at mapping building damage, despite the moderate GSD of 10 m. Left: Sentinel-2 TCI product. Center: zero-shot prediction from Sentinel-1 and Sentinel-2. Right: reference map, adapted from California Dept. of Forestry and Fire Protection (2025).

problem into sequential building localization and damage classification stages (Durnov, 2019). This two-stage paradigm was subsequently adopted by several methods (Zhao and Zhang, 2020; Wu et al., 2021), while alternative approaches explored end-to-end joint optimization of both tasks through multi-task learning (Weber and Kané, 2020; Gupta and Shah, 2021; Hao et al., 2021; Zheng et al., 2021). More recently, the field has witnessed the integration of advanced architectures originally developed for other computer vision applications, including high-res networks (HRNet; Liu et al., 2022), transformers (Chen et al., 2022; Kaur et al., 2023), and deep state-space models (Chen et al., 2024).

Beyond architectural innovations, the research community has expanded the scope of available benchmarks to improve the diversity of imaging modalities and address operational constraints. Notable examples include the BRIGHT dataset (Chen et al., 2025), which combines pre-event VHR optical imagery with post-event VHR synthetic aperture radar (SAR) imagery, and the following DisasterM3, which combines xBD, BRIGHT, and ten new disaster events with textual descriptions for visual reasoning tasks (Wang et al., 2025). In parallel, the broader remote sensing community has embraced geospatial foundation models (GeoFMs). Large, pre-trained models like Clay (Clay Foundation, 2023), DOFA (Xiong et al., 2024), Prithvi (Szwarc-

cman et al., 2024), or AnySAT (Astruc et al., 2025) promise generalizable representations, learned from vast amounts of unlabeled satellite data. However, recent empirical studies suggest that they do not consistently attain the performance of task-specific, fully supervised baselines (e.g., Corley et al., 2024). Vision-language models (VLMs) are another emerging frontier that could change how we interact with Earth observation data. Early systems such as TEOChat (Irvin et al., 2024) demonstrate the potential of natural language interfaces to interpret satellite images, yet remote sensing VLMs remain in their infancy.

While VHR imagery provides valuable detail for damage assessment and is often made public in the case of a natural disaster, its accessibility might be limited, which can be an issue especially during larger disasters (Ainscoe et al., 2025). Here, we investigate whether free, globally accessible satellite imagery can complement VHR data to perform building damage assessment at scale. Specifically, we explore the potential of the Copernicus missions, namely the multispectral imagery from Sentinel-2 and synthetic aperture radar (SAR) data from Sentinel-1, which both offer global coverage with sub-weekly revisits. To evaluate them for building damage assessment, we extend the established xBD dataset by pairing each VHR image with temporally aligned Sentinel-1/2 observations, enabling supervised learning of damage patterns at moderate resolution. Our contributions are: (1) **xBD-S12**, a novel dataset that spatiotemporally aligns each xBD image pair with corresponding Sentinel-2 and Sentinel-1 acquisitions. As a byproduct, we provide corrected geolocations for the original xBD dataset. (2) Deep learning models for damage assessment based on Sentinel imagery, accompanied by comprehensive per-disaster performance metrics to characterize model behavior across different event types. (3) Ablation studies to examine the utility of GeoFMs for the task, demonstrating that they offer limited benefit. (4) Independent, qualitative validation of the estimated damage maps for a disaster outside of xBD, by comparing to assessments from other sources. We release xBD-S12 as well as code and trained model weights, at <https://github.com/prs-eth/xbd-s12>.

## 2. Material and Methods

### 2.1 Data

We start by describing the construction of the xBD-S12 dataset. The data is summarized in Fig. 2, the geographical distribution is depicted in Fig. 3. and example patches are shown in Fig. 4.

**2.1.1 Original xBD dataset.** The basis of our work is the xBD dataset (Gupta et al., 2019a,b), originally developed for the xView2 challenge. It consists of 11,034 pre-/post-desaster VHR image pairs (each 1024×1024 pixels at ≈50 cm/pixel) and over 425,000 building polygons. The image pairs are derived from 66 tiles released publicly by MAXAR (either through their Open Data Program or specifically for the challenge) and span 19 disaster events across seven countries. Events are grouped into five categories: earthquakes, floods, storms, volcanic activity, and wildfires. From the original tiles, the authors selected the relevant patches and contracted a commercial service to manually annotate building footprints and corresponding damage levels. Each building footprint is assigned one of four discrete damage classes: no damage, minor damage, major damage, or destroyed. Owing to its scale and high-quality annotations, the xBD dataset has become a well-established benchmark for post-disaster building damage assessment and remains widely used in the field.




Figure 2. xBD-S12 comprises 10,315 image pairs across 16 disaster events. Colours denote the event-based split.

The original dataset is divided into four subsets: *train*, *tier3*, *test*, and *holdout*. The train, test, and holdout sets represent 80/10/10% random splits from the same ten disaster events, while the tier3 set contains data from nine additional disasters. During the challenge, participants trained their models on the train and tier3 sets, and used the test and holdout sets for evaluation. Due to the random split, the models see examples from every individual disaster location during training, with the associated danger of learning location- or event-specific patterns that do not generalize. To circumvent this, Hafner et al. (2025) proposed a different, event-based split that ensures geographical separation between the training and test sets, while maintaining at least one instance of each disaster type in the test set.

We systematically train and test our models with both splits to measure expected performance under two plausible scenarios: (1) *known distribution*, where a model is trained on a small labeled subregion of the specific event and then used to scale to the full region of interest, and (2) *unknown distribution*, where the model is applied to an unseen disaster for which no labels exist. Henceforth, we refer to the original split as *xView2*, using the combined train and tier3 sets for training and the test set for evaluation. The alternative split of Hafner et al. (2025) is termed as *event-based*.

Finally, the images in the xBD dataset suffer from georeferencing errors, likely due to an erroneous transformation. The misalignment does not matter for the original VHR task, which operates in image coordinates. It does, however, matter when aligning with Sentinel images for xBD-S12. We correct the coregistration, using the building footprints as ground control to fit an affine transformation. For details, see Appendix A<sup>1</sup>.

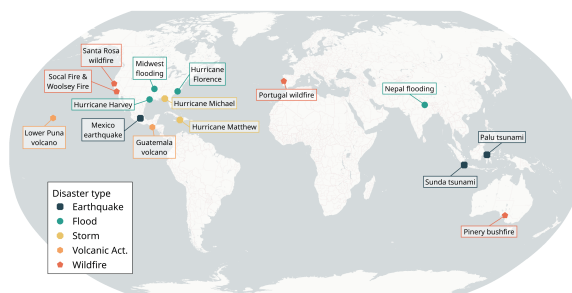


Figure 3. Geographical distribution of xBD-S12, grouped by disaster type as introduced by Hafner et al. (2025).

**2.1.2 Sentinel images.** Next, we determine the spatial extent of the coregistered xBD image patches, identify the corresponding Sentinel-1 and Sentinel-2 tiles, and download them, applying a consistent patch-level logic across all image pairs. Sentinel-2 tiles are selected based on cloud coverage and on

<sup>1</sup> Appendices available at <https://arxiv.org/abs/2511.05461>

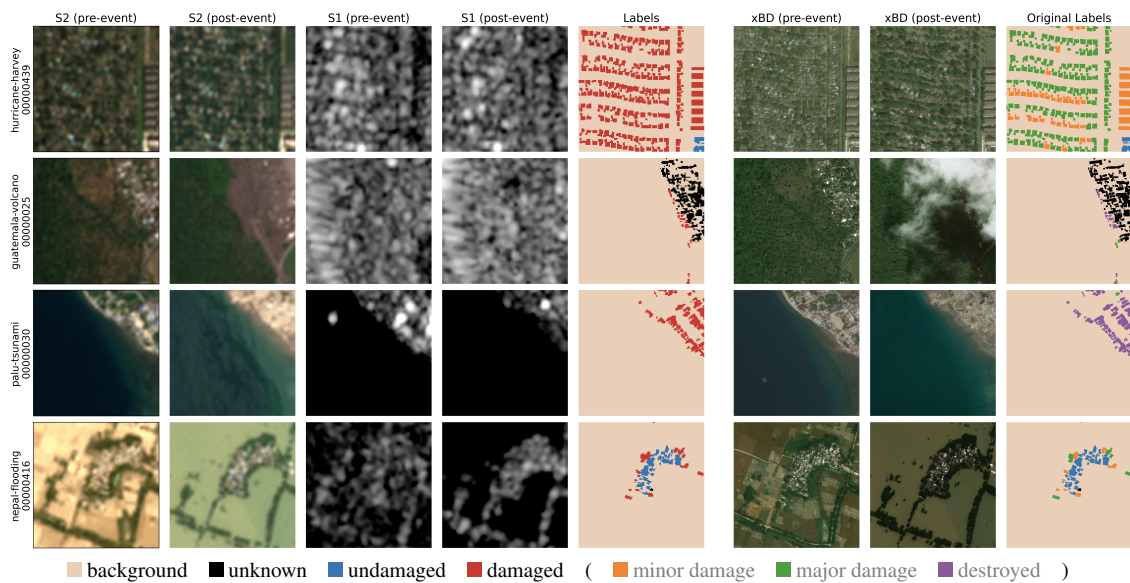


Figure 4. Example patches from xBD-S12. For visualisation purposes, we display the True Color Image product for Sentinel-2 and the VV-polarised (log-)amplitude for Sentinel-1. All tiles are  $128 \times 128$  px ( $\approx 4$  m GSD). On the right, VHR images ( $1024 \times 1024$  px,  $\approx 0.5$  m GSD) and labels from the original high-resolution xBD dataset are shown for reference. See Appendix E for more examples.

temporal proximity to both the VHR images and the dates of the disasters. We rely on Google’s CloudScore+ (Pasquarella et al., 2023) for cloud coverage, which has been precomputed for the entire Sentinel-2 archive and is accessible through Google Earth Engine (GEE, Gorelick et al., 2017). We make sure that post-event imagery clearly captures the disaster effects (e.g., flooding is only visible on a few days after the event), whereas we allow for a broader date range in the pre-event imagery and optimize for minimal cloud cover. To facilitate access to True Color Image (TCI) products, we download the original .SAFE Sentinel-2 Level-2A files using  $\Phi$ -down (Del Prete, 2025). The level-2A product consists of 12 multispectral bands with native ground sampling distances (GSDs) ranging from 10 m to 60 m.

Once Sentinel-2 images are selected, we identify the closest Sentinel-1 acquisitions, again with respect to both the VHR images and the disaster dates. We also require the pre- and post-event Sentinel-1 images to share the same orbit direction, since ascending and descending orbits lead to very different viewing geometries of the same area. We do not prioritize one orbit direction over the other, resulting in a mix of both across the dataset. For one disaster (Hurricane Matthew), pre- and post-disaster images from the same orbit direction are not available, so we allow images with different directions. We use the Ground Range Detected (GRD) product available on GEE, which provides log-amplitudes in the VV and VH polarizations, resampled to 10 m. GEE automatically performs standard preprocessing for each tile, namely thermal noise removal, radiometric calibration, and terrain correction. We do not perform any additional preprocessing. For more details on dataset creation, see Appendix B.

**2.1.3 Preprocessing.** We discard three disasters, namely the three tornado events, as they occurred before the launch of the Sentinel missions. For the remaining 16 disasters, we extract Sentinel-1 and Sentinel-2 patches that match the spatial extent of the xBD patches, resulting in a total of 10,315 image pairs.

**Spatial resampling.** The xBD dataset features VHR images with slightly varying GSDs. We find that, translated to Sentinel’s 10 m resolution, this would result in patch sizes vary-

ing between 48 and 52 pixels. To obtain consistent input dimensions across our dataset, we resample all Sentinel patches to a fixed size of  $128 \times 128$  pixels with Lanczos interpolation, corresponding to an effective GSD of  $\approx 4$  m. This resampling serves three purposes: (1) it standardizes the patch dimensions in the xBD-S12 dataset, (2) it provides a more suitable input size for modern deep learning architectures, and (3) there is evidence that building detection in Sentinel data benefits from moderate upsampling (Sirko et al., 2023).

**Label simplification.** The original fine-grained classification task is to some degree ill-posed, because the delimitation of the four damage classes is somewhat ambiguous and depends on the disaster type (Gupta et al., 2019b; Hafner et al., 2025). What is more, at 10 m GSD individual buildings occupy only few pixels, so that the distinction becomes impractical. We therefore merge all damage classes (minor, major, destroyed) into a single *damaged* class, resulting in three final classes: *background*, *intact*, and *damaged*. This simplification aligns better with what can realistically be expected from damage mapping with Sentinel data, where the aim is to identify affected buildings rather than to quantify the exact degree of damage.

**Invalid mask.** We create an *invalid* mask by keeping track of unclassified pixels in the original xBD dataset. These pixels occur either when the VHR image does not cover the full patch or when the post-event image is cloudy. These pixels are excluded from training and evaluation.

## 2.2 Models

**Problem formulation.** We formulate building damage assessment as a multiclass semantic segmentation problem. Our input consists of pre- and post-disaster images from Sentinel-1 and Sentinel-2:  $\mathbf{X}_{S1,pre}, \mathbf{X}_{S1,post} \in \mathbb{R}^{2 \times 128 \times 128}$  and  $\mathbf{X}_{S2,pre}, \mathbf{X}_{S2,post} \in \mathbb{R}^{12 \times 128 \times 128}$ . The task is to predict a damage map  $\hat{\mathbf{Y}} \in \mathbb{R}^{128 \times 128}$  where each pixel belongs to one of the three classes: *background*, *intact*, or *damaged*.

**2.2.1 Architecture.** Our baseline is a standard U-Net with ResNet34 backbone, pretrained on ImageNet, followed by a

lightweight segmentation head that predicts per-pixel class logits. Through ablation studies, we found that an encoder depth of three layers is sufficient and that removing all batch normalization layers from the encoder improves performance.

**Early vs. late fusion.** We test two fusion strategies. In the *early fusion*, all pre- and post-event bands are concatenated channel-wise and fed directly to a single U-Net. In the *late fusion*, we use a Siamese design to separately encode the pre- and post-event images through a shared U-Net (without the segmentation head). The resulting feature maps are then concatenated before the final segmentation head.

**Two-step vs. joint optimization.** We test two variants: either a single model that directly predicts the three classes (*joint*); or two separate models (*2-step*), one that localizes buildings via binary segmentation into building vs. background, and a second one that performs binary damage classification. In the latter case, the final damage map is obtained by masking the damage predictions to the estimated building pixels.

**Ensembling.** For robustness, we perform model ensembling by averaging the logits of three networks trained with different seeds before converting them to class probabilities with the final activation function.

**2.2.2 Comparison methods.** We compare our straightforward segmentation network against multiple dedicated architectures developed for damage mapping in VHR imagery:

- *Strong baseline* (Hafner et al., 2025): A simplified version of the xView2 winning solution, framed as multiple independent binary semantic segmentations.
- *DisasterAdaptiveNet* (Hafner et al., 2025): An extension of the strong baseline that adds a FiLM (Perez et al., 2018) module to modulate the predictions according to the disaster type.
- *ChangeOS* (Zheng et al., 2021): A deep object-based semantic change detection framework using a partial Siamese encoder with task-aware encoders and decoders.
- *ChangeMamba* (Chen et al., 2024): Change detection based on the Mamba (Gu and Dao, 2024) deep state space architecture, designed to enable the global receptive field of Transformers with a lower memory footprint.

**2.2.3 Foundation models.** We also test two recent geospatial *foundation* models: Prithvi-EO-2.0 with 300M parameters (Szwarcman et al., 2024), as available in Terratorch (Gomes et al., 2025), and DOFA (Xiong et al., 2024) with ViT-Base backbone, as implemented in TorchGeo (Stewart et al., 2024). In both cases, we keep the backbone frozen and only finetune a UperNet (Xiao et al., 2018) to decode the features into segmentation masks. We experiment with training either on the full dataset or on a small subset to simulate a low-data regime.

## 2.3 Training and evaluation

**2.3.1 Building footprint buffering.** We apply a morphological dilation to obtain a 3-pixel buffer around the ground truth building footprints and add the buffer pixels to the *invalid* mask during training. This strategy (1) mitigates potential misregistration between MAXAR and Sentinel imagery, and (2) encourages bolder predictions by not penalizing outputs extending slightly beyond ground truth footprints. Given Sentinel’s medium resolution, precise boundary delineation is challenging;

we therefore prioritize recall over boundary alignment and accepting slightly inflated predictions. During inference, we evaluate both with and without the buffer. A performance drop when removing the buffer from the evaluation indicates that the model indeed inflated the footprints, in practice an acceptable price to pay for more complete damage detection.

**2.3.2 Performance metrics.** Following the xView2 Challenge, we separately evaluate the results w.r.t. building localization and damage detection, in both cases using the F1 score (the harmonic mean between precision and recall). We report two scores: the localization score  $F1_{loc}$  and the damage score  $F1_{dmg}$ . Unless stated otherwise,  $F1_{loc}$  is evaluated without a buffer around the building footprints. We additionally show  $F1_{loc, B=X}$ , evaluated with an X-pixel buffer around the footprints. The damage score  $F1_{dmg}$  is the harmonic mean between the F1 scores of classes *intact* and *damaged*. Following the xView2 challenge, it is computed only over pixels within the ground truth building footprints, so  $F1_{dmg}$  is not impacted by the buffer. Again following xView2, the localization and damage scores are linearly combined into an overall “competition score”:  $F1_{comp} = 0.3 \cdot F1_{loc} + 0.7 \cdot F1_{dmg}$ .

**2.3.3 Training details.** Sentinel-1 and Sentinel-2 data are normalized channel-wise using percentile-based min-max scaling:  $x_{norm} = \text{clip}((x - p_1) / (p_{99} - p_1), 0, 1)$ , where  $p_1$  and  $p_{99}$  are computed from the training set. During training, we apply geometric augmentations (random horizontal/vertical flips and 90-degree rotations) jointly to imagery and labels.

To address class imbalance (background pixels vastly outnumber building pixels, and intact buildings outnumber damaged ones), we employ biased random sampling, where the probability of drawing a sample is inversely proportional to the frequency of its pixel classes in the training set.

All models are trained for 40 epochs using the AdamW optimizer (Loshchilov and Hutter, 2017) with learning rate  $5 \times 10^{-4}$ , weight decay  $1 \times 10^{-4}$ , and cosine annealing schedule with 3 epochs of linear warm-up. Unlike prior works that employ specialized loss functions, we found the standard cross-entropy loss to perform best, likely due to our simplified label set. As the final model to be evaluated on the test set, we retain the checkpoint with highest  $F1_{comp}$  on the validation set (respectively higher  $F1_{loc}$  for localization-only experiments).

## 3. Results

Table 1 presents the detailed results of our approach against several SOTA models. We report metrics both with and without a 3-pixel buffer around the building footprints. From these metrics, we make several key observations:

- ChangeMamba achieves the best performance on the original xView2 split, but exhibits poor generalization to unseen disasters in the event-based split. This suggests that the model memorizes event- or location-specific characteristics from the training data, which then hamper generalization to unseen events and locations at inference.
- In contrast, our simpler two-step approach works best for the event-based split. We hypothesize that, with its lower capacity, it operates more like a conventional change detection model and captures fewer intricate, event-specific patterns; thus generalizing better to unseen events. The result is in line with the general trend that architectural sophistication becomes less important as the GSD increases.

Model	<i>xView2 split</i>			<i>event-based split</i>		
	F1 <sub>comp</sub> (B=3)	F1 <sub>loc</sub> (B=3)	F1 <sub>dmg</sub>	F1 <sub>comp</sub> (B=3)	F1 <sub>loc</sub> (B=3)	F1 <sub>dmg</sub>
U-Net (2-step, early fusion)	0.760 (0.848)	0.597 (0.891)	0.830	0.690 (0.756)	<b>0.502</b> (0.722)	0.771
U-Net (2-step, late fusion)	<b>0.761</b> (0.849)	0.597 (0.891)	0.831	<b>0.710</b> (0.775)	<b>0.502</b> (0.722)	<b>0.799</b>
U-Net (joint, early fusion)	<b>0.764</b> (0.843)	<b>0.611</b> (0.873)	0.830	0.687 (0.738)	0.491 (0.662)	0.771
U-Net (joint, late fusion)	0.760 (0.845)	<b>0.606</b> (0.888)	0.826	<b>0.692</b> (0.743)	0.455 (0.622)	<b>0.794</b>
StrongBaseline	0.760 (0.844)	0.595 (0.876)	<b>0.831</b>	<b>0.690</b> (0.740)	0.444 (0.609)	<b>0.796</b>
DisasterAdaptiveNet	0.734 (0.815)	0.599 (0.869)	0.792	0.636 (0.655)	0.358 (0.419)	0.756
ChangeOS	0.718 (0.801)	0.488 (0.767)	0.816	0.589 (0.644)	0.432 (0.615)	0.656
ChangeMamba	<b>0.800</b> (0.831)	<b>0.651</b> (0.755)	<b>0.863</b>	0.655 (0.667)	0.362 (0.402)	0.780

Table 1. F1 scores for different approaches on both the original xView2 split and the event-based split proposed by Hafner et al. (2025). Values are computed on an ensemble of three runs where logits were averaged before thresholding. Values in parentheses are computed excluding a 3-pixel buffer around buildings during evaluation. Since F1<sub>dmg</sub> is computed only on the building pixels in the groundtruth, it is not affected by the buffer. We highlight the top three results for each metric as **first**, **second**, and **third**.

- Joint optimization is beneficial when training and testing in the same distribution (xView2 split), but the less elegant 2-step approach is more robust when mapping unseen disasters (event-based split).
- The 3-pixel buffer greatly increases the localization score of all models, indicating that they detect buildings correctly and inflate their outline, as expected. In other words, training with the buffer has the desired effect of increasing per-building recall at the cost of per-pixel precision, so as to avoid missing possible damages.
- The FiLM layer in DisasterAdaptiveNet provides no improvement over the baseline (StrongBaseline). It appears that conditioning on the disaster type either benefits only VHR image analysis or, more likely, that it improves the discrimination of detailed damage classes like *minor* vs. *major* damage, which we merged and whose definitions are disaster-dependent.
- ChangeOS markedly underperforms compared to the other tested models. The reason is likely its object-based approach, which becomes less effective at the 10 m resolution of Sentinel, where buildings and other man-made urban structures no longer form well-delineated, separable “objects”.

For the remainder of this section, we continue with our best-performing model, i.e., the 2-step strategy with the Siamese (late fusion) architecture.

### 3.1 Performance per disaster event

Table 2 shows the metrics per event, for both test splits. There are substantial performance variations, highlighting both the capabilities and the limitations of medium-resolution satellite imagery for damage assessment.

In the in-domain xView2 split, we achieve relatively high scores on all disasters. The highest scores are reached for wildfires, floods, and tsunamis, which can be explained by the nature of the associated damage, where large contiguous areas are affected in a way that significantly alters their spectral signatures (burn scars, respectively water). A notable exception is the *Mexico earthquake*, for which we cannot detect any damage. This disaster exposes a fundamental limitation of the lower spatial resolution: Damages are sparse and only affect individual, small buildings dispersed across a densely built-up area, rendering them invisible at the Sentinels’ 10 m GSD (also for human interpreters). See Appendix D for more details on this disaster.

As expected, performance is in general lower for the event-based split, since the training did not include data from the

vicinity of the test regions. Only for the *Guatemala volcano*, performance increases, which may at first seem surprising. It turns out that the discrepancy stems from the particular choice of test data in the xView2 split. The test set for this small event consists of only five patches, all of which happen to depict challenging regions, with small buildings at the boundaries of lava streams and significant distribution shifts compared to the other patches from the same disaster.

	F1 <sub>comp</sub> (B=3)	F1 <sub>loc</sub> (B=3)	F1 <sub>dmg</sub>	
<i>xView2 split</i>	Guatemala volcano	0.611 (0.671)	0.484 (0.684)	0.665
	Hurricane Florence	0.785 (0.852)	0.485 (0.708)	0.913
	Hurricane Harvey	0.749 (0.839)	0.616 (0.917)	0.806
	Hurricane Matthew	0.557 (0.628)	0.423 (0.659)	0.615
	Hurricane Michael	0.596 (0.683)	0.566 (0.853)	0.609
	Mexico earthquake	0.194 (0.287)	0.648 (0.958)	0.000
	Midwest flooding	0.612 (0.692)	0.530 (0.795)	0.647
	Palu tsunami	0.777 (0.863)	0.643 (0.929)	0.834
	Santa Rosa wildfire	0.842 (0.942)	0.575 (0.909)	0.956
	Socal fire	0.763 (0.843)	0.548 (0.815)	0.855
<i>event-based</i>	Guatemala volcano	0.721 (0.809)	0.467 (0.760)	0.830
	Hurricane Matthew	0.403 (0.474)	0.444 (0.683)	0.385
	Nepal flooding	0.616 (0.655)	0.441 (0.572)	0.690
	Santa Rosa wildfire	0.747 (0.828)	0.547 (0.817)	0.833
	Sunda tsunami	0.336 (0.396)	0.536 (0.739)	0.249

Table 2. Per-disaster performance breakdown for both test splits. Metrics are computed using the same methodology as Table 1.

### 3.2 Input data for building localization

When using the 2-step approach, the pre-disaster image alone would, in principle, be sufficient for the building localization task. Nevertheless, feeding also post-event images—which must be available for damage mapping—should improve performance, if only by providing redundancy for undamaged (or lightly damaged) buildings. Table 3 shows F1<sub>loc</sub> scores for both the validation and test subsets. They confirm that access to both pre- and post-event views consistently improves performance. The most significant gain occurs for the test set of the event-based split. We hypothesize that, by seeing views of the same buildings captured on different dates, the model learns some degree of invariance to atmospheric and lighting conditions, which is beneficial when applied to unseen geographic regions.

### 3.3 Effect of buffer size

Masking a buffer around building footprints during training significantly improves mapping performance. To determine the right size for the buffer, we train models with buffers ranging

Input	<i>xView2 split</i>		<i>event-based split</i>	
	Valid	Test	Valid	Test
Pre	0.833	0.868	0.869	0.566
Pre+Post	0.865	0.891	0.891	0.722

Table 3.  $F1_{loc}(B=3)$  for different input configurations. Multi-temporal inputs improve generalization to unseen locations, respectively events.

from 0 to 5 pixels and compute the corresponding localization scores (Fig. 5). As expected, larger buffers yield increasingly diffuse predictions, trading pixel-level accuracy for higher building-level recall. Depending on the data split, localization scores peak around 2 to 3 pixels. On average, the 3-pixel buffer enjoys a slight edge, hence we adopt it as our default setting. I.e., we accept inflated building outlines to minimize the risk of missing damaged buildings.

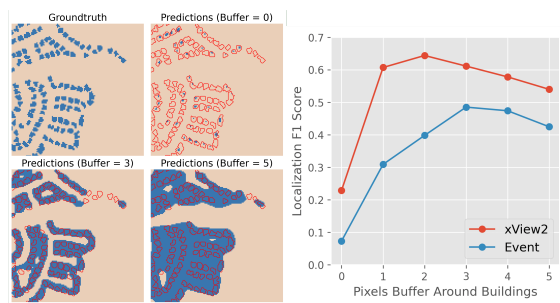


Figure 5. Effect of buffer size on building localization. Left: Predicted footprints with different training buffer sizes. Right:  $F1_{loc}$  score as a function of buffer size.

### 3.4 Geospatial foundation models

We additionally evaluate whether pretrained GeoFMs can improve damage assessment. We test two recent models, Prithvi-EO-2.0-300M (Szwarcman et al., 2024) and DOFA-Base (Xiong et al., 2024). We treat the GeoFM as a frozen feature extractor and finetune only the UperNet decoder (see Appendix F). Table 4 summarizes results for four different scenarios.

In the in-domain *xView2* split, Prithvi achieves a slightly higher overall score than our 2-step model (but still significantly lower than ChangeMamba, Table 1), with DOFA not too far behind. However, that performance plummets as one shifts to the event-based split. It appears that, when faced with a generic feature extractor not tuned to damage mapping, the decoder has a stronger tendency to overfit to the specific conditions of the training set, thus generalizing poorly to unseen events (respectively, locations).

To probe the value of GeoFMs in the low-data regime, we also train only on a single event, then test on another of the same type. The reasoning is that with so little training data, pretrained features should be particularly helpful. Somewhat unexpectedly, for wildfires, this is not the case. All methods achieve high scores. In other words, burnt areas are well detectable at Sentinel resolution, but even a small amount of data is sufficient to learn that detection as well as a GeoFM. Overall, the U-Net maintains the best overall performance, driven by superior building localization. For hurricanes, the two GeoFMs do achieve notable improvements over the U-Net baseline. Yet, absolute mapping performance remains very low for all models, indicating that more data is needed to train a meaningful build-

ing detector or damage classifier—even if the feature extractor has been pretrained on massive data.

Overall, we conclude that pretrained GeoFMs are suitable for detecting broad environmental changes like burnt or flooded areas, but struggle with more fine-grained tasks like building localization. Consequently, existing GeoFMs bring only small, if any practical benefits for damage mapping. We point out that finetuning also the encoder (feature extractor) of a GeoFM might improve its performance. However, this would incur substantially higher computational cost than training the lightweight U-Net model from scratch.

		$F1_{comp}(B=3)$	$F1_{loc}(B=3)$	$F1_{dmg}$
<i>xView2 split</i>	U-Net	0.761 (0.849)	<b>0.597</b> (0.891)	0.831
	Prithvi	<b>0.764</b> (0.840)	0.571 (0.821)	<b>0.847</b>
	DOFA	0.723 (0.797)	0.557 (0.805)	0.794
<i>event-based split</i>	U-Net	<b>0.710</b> (0.775)	<b>0.502</b> (0.722)	<b>0.799</b>
	Prithvi	0.634 (0.695)	0.437 (0.641)	0.718
	DOFA	0.533 (0.594)	0.436 (0.642)	0.574
<i>Woolsey fire → Santa Rosa fire</i>	U-Net	<b>0.746</b> (0.784)	<b>0.462</b> (0.589)	0.868
	Prithvi	0.725 (0.745)	0.312 (0.378)	<b>0.902</b>
	DOFA	0.698 (0.760)	0.405 (0.615)	0.823
<i>hurricane Michael → hurricane Matthew</i>	U-Net	0.143 (0.178)	<b>0.349</b> (0.464)	0.055
	Prithvi	<b>0.271</b> (0.316)	0.309 (0.460)	<b>0.255</b>
	DOFA	0.255 (0.293)	0.300 (0.427)	0.236

Table 4. Performance of GeoFMs across different training scenarios. Values in parentheses show scores when excluding a 3-pixel buffer around buildings, also during evaluation.

### 3.5 Inference on unseen disasters

To test generalization to real-world examples outside of *xBD-S12*, we apply our trained U-Net model to the 2025 Palisades wildfire, a large recent event that occurred in January 2025 in the western suburbs of Los Angeles. We acquire Sentinel imagery following the same protocol as for *xBD-S12*, with post-event acquisitions on February 1 for Sentinel-2 and on February 7 for Sentinel-1.

Fig. 6 compares our damage predictions with the official damage assessment from the California Dept. of Forestry and Fire Protection (2025). Although we cannot carry out a quantitative assessment, as the raw data is not publicly available, we observe an excellent correlation between our damage map and the dashboard visualizations of the official assessment.

## 4. Conclusion

We have shown that freely available medium-resolution satellite imagery from the Copernicus missions can support rapid damage assessment for disaster response. To that end, we have constructed *xBD-S12*, a dataset of paired pre- and post-event images from Sentinel-1 and Sentinel-2, with damage annotations derived from the earlier VHR *xBD* dataset.

We find that, in many scenarios, building damage is detectable despite the 10 m GSD. For disasters that induce large-scale, spectrally distinct changes, such as wildfires or floods, performance is consistently high. Events with more localized damage, such as hurricanes, are more challenging. For small, scattered damages like those from weak earthquakes, Copernicus data are not suitable due to insufficient resolution, at least under our current methodological framework.

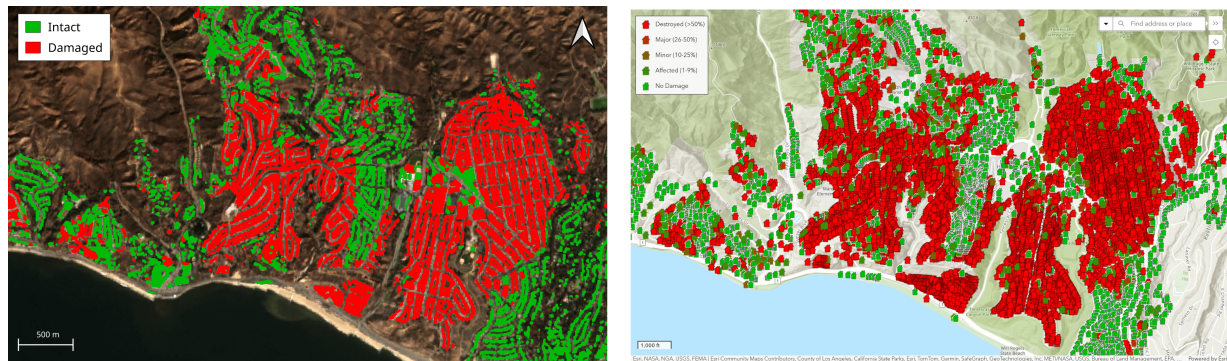


Figure 6. Left: output of our models when trained on the entire dataset. Right: reference map adapted from the damage assessment conducted by the [California Dept. of Forestry and Fire Protection \(2025\)](#). Base maps: Sentinel-2 TCI / ESRI World Terrain Base.

Importantly, we do not regard moderate-resolution imagery as a replacement for VHR products, but rather as a complementary data source in the disaster response workflow. Especially in the absence of VHR imagery, damage assessment with Copernicus data can go a surprisingly long way. Its main value lies in rapid, coarse coverage, whereas a higher resolution is better suited for detailed damage inventories.

We have followed the task definition of xBD, but we note that, in operational settings, building footprints are often available from national inventories or global datasets such as Overture Maps ([Overture Maps Foundation](#)). Where such data exists, one would typically prefer it and employ satellite imagery only for damage classification. Especially at 10 m GSD, building localization remains the weakest component of the pipeline.

Our approach deliberately prioritizes simplicity and accessibility, using only standard GRD products from Sentinel-1 and bi-temporal acquisitions. Future work could explore richer temporal information from longer time series (e.g., [Dietrich et al., 2025](#)), or incorporate InSAR products such as coherence maps or damage proxy maps, which have shown great promise for detecting structural changes (e.g., [Ainscoe et al., 2025](#); [Scher and Van Den Hoek, 2025](#)). Although geospatial foundation models were of limited use in our experiments, we emphasize that they continue to improve, and future versions may well become effective tools for damage assessment.

## References

Ainscoe, E. A., Swaminathan, R., Way, L., Modugno, S., Chin, S. T., Panta, N., Crevoisier, T., Yun, S.-H., 2025. Earthquake damage mapped more comprehensively and accurately by radar satellites than optical imagery. *Nat. Comm. Earth Env.*, 6(1), 631.

Astruc, G., Gonthier, N., Mallet, C., Landrieu, L., 2025. AnySat: One earth observation model for many resolutions, scales, and modalities. *Comput. Vis. Pat. Rec.*

Banholzer, S., Kossin, J., Donner, S., 2014. The impact of climate change on natural disasters. Z. Zommers, A. Singh (eds), *Reducing disaster: Early warning systems for climate change*, Springer, 21–49.

Braik, A. M., Koliou, M., 2024. Automated building damage assessment and large-scale mapping by integrating satellite imagery, GIS, and deep learning. *Computer-Aided Civil Infrastruct. Eng.*, 39(15), 2389–2404.

California Dept. of Forestry and Fire Protection, 2025. Palisades fire damage maps. <https://recovery.lacounty.gov/palisades-fire/>. (accessed 2025-11-04).

Cao, Q. D., Choe, Y., 2020. Building damage annotation on post-hurricane satellite imagery based on convolutional neural networks. *Natural Hazards*, 103(3), 3357–3376.

Chen, H., Nemni, E., Vallecorsa, S., Li, X., Wu, C., Bromley, L., 2022. Dual-tasks siamese transformer framework for building damage assessment. *Int. Geosci. Remote Sens. Symp.*

Chen, H., Song, J., Dietrich, O., Broni-Bediako, C., Xuan, W., Wang, J., Shao, X., Wei, Y., Xia, J., Lan, C. et al., 2025. BRIGHT: A globally distributed multimodal building damage assessment dataset with very-high-resolution for all-weather disaster response. *Earth Syst. Sci. Data*.

Chen, H., Song, J., Han, C., Xia, J., Yokoya, N., 2024. ChangeMamba: Remote sensing change detection with spatio-temporal state space model. *IEEE T. Geosci. Remote Sens.*, 62, 1–20.

Clay Foundation, 2023. The Clay foundation model – an open source AI model and interface for Earth. GitHub repository, <https://github.com/Clay-foundation/model>. (accessed 2025-11-04).

Corley, I., Robinson, C., Ortiz, A., 2024. A change detection reality check. *preprint arXiv:2402.06994*.

Del Prete, R., 2025. phidown: A python tool for streamlined data downloading from CDSE. Zenodo, <https://doi.org/10.5281/zenodo.15332053>. (accessed 2025-11-04).

Dietrich, O., Peters, T., Sainte Fare Garnot, V., Sticher, V., Ton-That Whelan, T., Schindler, K., Wegner, J. D., 2025. An open-source tool for mapping war destruction at scale in Ukraine using Sentinel-1 time series. *Nat. Comm. Earth Env.*, 6(1).

Durnov, V., 2019. xView2: 1st place solution. GitHub repository, [https://github.com/vdurnov/xview2\\_1st\\_place\\_solution](https://github.com/vdurnov/xview2_1st_place_solution). (accessed 2025-11-04).

Fujita, A., Sakurada, K., Imaizumi, T., Ito, R., Hikosaka, S., Nakamura, R., 2017. Damage detection from aerial images via convolutional neural networks. *IAPR Conf. Machine Vision Appl.*

Galanis, M., Rao, K., Yao, X., Tsai, Y.-L., Ventura, J., Fricker, G. A., 2021. DamageMap: A post-wildfire damaged buildings classifier. *Int. J. Disaster Risk Red.*, 65, 102540.

Gomes, C., Blumenstiel, B., de Sousa Almeida, J. L., de Oliveira, P. H., Fraccaro, P., Escofet, F. M., Szwarcman, D., Simumba, N., Kienzler, R., Zadrozny, B., 2025. TerraTorch: The Geospatial Foundation Models Toolkit. *preprint arXiv:2503.20563*.

Gorelick, N., Hancher, M., Dixon, M., Ilyushchenko, S., Thau, D., Moore, R., 2017. Google Earth Engine: Planetary-scale geospatial analysis for everyone. *Remote Sens. Env.*, 202, 18–27.

- Gu, A., Dao, T., 2024. Mamba: Linear-time sequence modeling with selective state spaces. *Conf. on Language Modeling*.
- Gupta, R., Goodman, B., Patel, N., Hosfelt, R., Sajeev, S., Heim, E., Doshi, J., Lucas, K., Choset, H., Gaston, M., 2019a. Creating xBD: A dataset for assessing building damage from satellite imagery. *Comput. Vis. Pat. Rec. Workshops*.
- Gupta, R., Hosfelt, R., Sajeev, S., Patel, N., Goodman, B., Doshi, J., Heim, E., Choset, H., Gaston, M., 2019b. xBD: A dataset for assessing building damage from satellite imagery. *preprint arXiv:1911.09296*.
- Gupta, R., Shah, M., 2021. RescueNet: Joint building segmentation and damage assessment from satellite imagery. *Int. Conf. Pat. Rec.*
- Hafner, S., Gerard, S., Sullivan, J., Ban, Y., 2025. DisasterAdaptiveNet: A robust network for multi-hazard building damage detection from very-high-resolution satellite imagery. *Int. J. Appl. Earth Obs. Geoinf.*, 143, 104756.
- Hao, H., Baireddy, S., Bartusiak, E. R., Konz, L., LaTourette, K., Gribbons, M., Chan, M., Delp, E. J., Comer, M. L., 2021. An attention-based system for damage assessment using satellite imagery. *Int. Geosci. Remote Sens. Symp.*
- IPCC, 2023. *Weather and Climate Extreme Events in a Changing Climate*. Cambridge University Press, 1513–1766.
- Irvin, J. A., Liu, E. R., Chen, J. C., Dormoy, I., Kim, J., Khanna, S., Zheng, Z., Ermon, S., 2024. TEOChat: A large vision-language assistant for temporal earth observation data. *preprint arXiv:2410.06234*.
- Kaur, N., Lee, C.-C., Mostafavi, A., Mahdavi-Amiri, A., 2023. Large-scale building damage assessment using a novel hierarchical transformer architecture on satellite images. *Computer-Aided Civil Infrastruct. Eng.*, 38(15), 2072–2091.
- Kawasaki, A., Berman, M. L., Guan, W., 2013. The growing role of web-based geospatial technology in disaster response and support. *Disasters*, 37(2), 201–221.
- Liu, C., Sepasgozar, S. M., Zhang, Q., Ge, L., 2022. A novel attention-based deep learning method for post-disaster building damage classification. *Expert Syst. Appl.*, 202, 117268.
- Loshchilov, I., Hutter, F., 2017. Decoupled weight decay regularization. *preprint arXiv:1711.05101*.
- McKoy, J., 2025. Death count for 2025 LA county wildfires likely higher than records show, BU research finds. Boston University, <https://www.bu.edu/articles/2025/death-count-california-wildfires-higher-than-recorded/>. (accessed 2025-11-04).
- Overture Maps Foundation, 2025. Open Map Data. <https://overturemaps.org>. (accessed 2025-11-04).
- Pasquarella, V. J., Brown, C. F., Czerwinski, W., Rucklidge, W. J., 2023. Comprehensive quality assessment of optical satellite imagery using weakly supervised video learning. *Comput. Vis. Pat. Rec.*
- Perez, E., Strub, F., De Vries, H., Dumoulin, V., Courville, A., 2018. FiLM: Visual reasoning with a general conditioning layer. *AAAI Conf. on Artif. Intell.*
- Rolla, J., Khuller, A., An, K., Emberson, R., Fielding, E., Schultz, L., Miner, K., 2025. Satellite-Aided Disaster Response. *AGU Advances*, 6(1).
- Rudner, T. G., Rußwurm, M., Fil, J., Pelich, R., Bischke, B., Kopačková, V., Biliński, P., 2019. Multi3Net: segmenting flooded buildings via fusion of multiresolution, multisensor, and multitemporal satellite imagery. *AAAI Conf. on Artif. Intell.*
- Scher, C., Van Den Hoek, J., 2025. Nationwide conflict damage mapping with interferometric synthetic aperture radar: A study of the 2022 Russia–Ukraine conflict. *Sci. Remote Sensing*, 11, 100217.
- Sirko, W., Brempong, E. A., Marcos, J. T., Annkah, A., Korme, A., Hassen, M. A., Sapkota, K., Shekel, T., Diack, A., Nevo, S. et al., 2023. High-resolution building and road detection from Sentinel-2. *preprint arXiv:2310.11622*.
- Stewart, A. J., Robinson, C., Corley, I. A., Ortiz, A., Lavista Ferres, J. M., Banerjee, A., 2024. TorchGeo: Deep Learning With Geospatial Data. *ACM Trans. Spatial Alg. & Systems*.
- Sun, W., Bocchini, P., Davison, B. D., 2020. Applications of artificial intelligence for disaster management. *Natural Hazards*, 103(3), 2631–2689.
- Szwarcman, D., Roy, S., Fraccaro, P., Gíslason, Þ. E., Blumentiel, B., Ghosal, R., de Oliveira, P. H., Almeida, J. L. d. S., Sedona, R., Kang, Y. et al., 2024. Prithvi-EO-2.0: A versatile multi-temporal foundation model for Earth observation applications. *preprint arXiv:2412.02732*.
- United Nations, 2025a. Jamaica: International support 'crucial' to hurricane recovery says Guterres. UN News, <https://news.un.org/en/story/2025/11/1166248>. (accessed 2025-11-04).
- United Nations, 2025b. 'The needs are huge': Pakistan reels from floods as millions left homeless. UN News, <https://news.un.org/en/story/2025/09/1165864>.
- Van Aalst, M. K., 2006. The impacts of climate change on the risk of natural disasters. *Disasters*, 30(1), 5–18.
- Wang, J., Xuan, W., Qi, H., Liu, Z., Liu, K., Wu, Y., Chen, H., Song, J., Xia, J., Zheng, Z. et al., 2025. DisasterM3: A Remote Sensing Vision-Language Dataset for Disaster Damage Assessment and Response. *preprint arXiv:2505.21089*.
- Weber, E., Kané, H., 2020. Building disaster damage assessment in satellite imagery with multi-temporal fusion. *preprint arXiv:2004.05525*.
- Wu, C., Zhang, F., Xia, J., Xu, Y., Li, G., Xie, J., Du, Z., Liu, R., 2021. Building damage detection using U-Net with attention mechanism from pre-and post-disaster remote sensing datasets. *Remote Sensing*, 13(5), 905.
- Xiao, T., Liu, Y., Zhou, B., Jiang, Y., Sun, J., 2018. Unified perceptual parsing for scene understanding. *Europ. Conf. Comput. Vis.*
- Xiong, Z., Wang, Y., Zhang, F., Stewart, A. J., Hanna, J., Borth, D., Papoutsis, I., Saux, B. L., Camps-Valls, G., Zhu, X. X., 2024. Neural plasticity-inspired multimodal foundation model for earth observation. *preprint arXiv:2403.15356*.
- Xu, J. Z., Lu, W., Li, Z., Khaitan, P., Zaytseva, V., 2019. Building damage detection in satellite imagery using convolutional neural networks. *preprint arXiv:1910.06444*.
- Zhao, F., Zhang, C., 2020. Building damage evaluation from satellite imagery using deep learning. *Int. Conf. Inform. Reuse & Integ. Data Sci.*
- Zheng, Z., Zhong, Y., Wang, J., Ma, A., Zhang, L., 2021. 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. Env.*, 265, 112636.