

Disaster assessment using computer vision and satellite imagery: Applications in water-related building damage detection

Danu Kim¹, Jeongkyung Won², Eunji Lee¹, Kyung Ryul Park¹, Jihee Kim¹
Sangyoon Park³, Hyunjoo Yang², Sungwon Park¹, Donghyun Ahn¹ and Meeyoung Cha^{4,1}

¹KAIST, ²Sogang University, ³The University of Hong Kong, ⁴Institute for Basic Science
{danu, mk35471, park.kr, jiheekim, psw0416, segaukwa}@kaist.ac.kr,
{jongrud, hyang}@sogang.ac.kr, sangyoon@hku.hk, mcha@ibs.re.kr

Abstract

The increasing frequency and severity of water-related disasters such as floods, tornadoes, hurricanes, and tsunamis in low- and middle-income countries exemplify the uneven effects of global climate change. The vulnerability of high-risk societies to natural disasters has continued to increase. To develop an effective and efficient adaptation strategy, local damage assessments must be timely, exhaustive, and accurate. We propose a novel deep-learning-based solution that uses pairs of pre- and post-disaster satellite images to identify water-related disaster-affected regions. The model extracts features of pre- and post-disaster images and uses the feature difference with them to predict damage in the pair. We demonstrate that the model can successfully identify local destruction using less granular and less complex ground-truth data than those used by previous segmentation models. When tested with various water-related disasters, our detection model reported an accuracy of 85.9% in spotting areas with damaged buildings. It also achieved a reliable performance of 97.5% in the case study on hurricane Iota. Our deep learning-based damage assessment model can help direct resources to areas most vulnerable to climate disasters, reducing their impacts while promoting adaptive capacities for climate-resilient development in the most vulnerable regions.

Introduction

The widespread impacts of human-induced climate change have been observed as the frequency and intensity of extreme events, including floods, tornadoes, hurricanes, and tsunamis, increase (IPCC 2022). Amid the growing climate risk, the global adaptive capacity to deal with disasters has not progressed accordingly, although the Sustainable Development Goals call for a worldwide response (Field et al. 2012). The lack of timely, comprehensive, and accurate data tracking damage at a fine-grained geographical level is one of the main reasons for such an adaptation deficit (Amundsen, Berglund, and Westskog 2010; Moser and Ekstrom 2010). For example, damage estimates are available only at the province level in the Emergency Events Database (EM-DAT), the largest international disaster database, making it difficult to pinpoint the worst-hit areas. Moreover, the conventional damage assessment using field surveys is resource-intensive

and time-consuming (Cao and Choe 2020), which hinders comprehensive regional coverage (Bakkensen, Shi, and Zurita 2018) and the rapid deployment of humanitarian assistance (Cao and Choe 2020). The field surveys may also suffer from cognitive biases such as reference dependence and recall errors (Guiteras, Jina, and Mobarak 2015).

Recent research in computer vision has combined high spatial resolution satellite images with machine learning to estimate disaster damage on a pixel level or an incident level (Ponnis et al. 2019; Weber and Kané 2020; Wu et al. 2021; Bai et al. 2020; Gupta and Shah 2021). These approaches use the XBD dataset created by (Gupta et al. 2019), the largest disaster damage dataset worldwide, providing pre- and post-disaster images with pixel-level damage labels. While the models developed using the dataset have made significant technical advances, they are fundamentally dependent on the existence of such fine-grained, complex damage labels in disaster-affected regions. Given that most high-quality ground-truth data comes from developed countries, a deep learning model that combines various forms of damage data and produces accurate local damage estimates would be helpful in actual disaster response efforts.

This paper presents a lightweight damage detection model based on deep learning and high spatial resolution satellite images. Our ground-truth data is less granular than the data employed by existing segmentation models. This feature is advantageous for developing countries lacking the statistical capacity and resources to produce quality local damage data. Evaluating our model on various water-related disasters from 2011 to 2019, our model achieved a performance of 90% in detecting disasters. Also, the case study on Providencia Island further demonstrates the generalizability of our model, as it successfully distinguished local destruction caused by Hurricane Iota in 2021.

We focused on water-related disasters such as floods, tornadoes, tsunamis, and hurricanes, among many others. We chose water-related disasters because their frequency and severity are increasing in the most vulnerable, low- and middle-income countries (Hallegatte et al. 2013; Rentschler, Salhab, and Jafino 2022; Edmonds et al. 2020). Despite the substantial losses caused by the disasters, a significant population is unable to leave disaster-prone regions (Tellman et al. 2021) for socioeconomic and political reasons (Lin, McDermott, and Michaels 2021; Raker 2020; Hunter 2005;

Henkel, Eunjee, and Magontier 2022). This necessitates an effective post-disaster response; therefore, local damage estimates can be especially helpful for prioritizing relief efforts and climate-resilient redevelopment.

The high-resolution satellite images have several advantages over other spatial data sets (i.e., Google Street views and aerial images) used in the related literature (Fujita et al. 2017; Vetrivel et al. 2018). First, satellite imagery guarantees extensive spatial and temporal coverage. Second, satellite images do not necessarily require resource-intensive damage labels for training, unlike other spatial input data used in the previous literature (Fujita et al. 2017; Vetrivel et al. 2018). They can combine various forms of damage labels (e.g., pixel-level labels, point coordinates of damaged properties, and district-level statistics) corresponding to their size.

The proposed model has several methodological contributions to environmental damage detection. First, the model effectively identifies local destruction by employing binary damage labels corresponding to satellite images' size. This approach reduces deep-learning models' reliance on fine-grained, complex ground-truth data, making the model more applicable to many developing countries that cannot afford such data. Second, to the best of our knowledge, our model is the first successful water-related disaster damage detection model. Third, our model is practical because it does not require extensive ground truth data specific to the damaged regions and its performance is robust to unseen regions.

Our algorithm can help policymakers by identifying the ideal location for humanitarian assistance deployment and minimizing the time lag between the onset of a disaster and assistance responses. The model provides the predicted labels to help determine some measure of centrality for the location where the resources should be concentrated. As this machine learning-based assessment can be implemented faster and cheaper than conventional on-site inspections, the development agencies would be able to deploy well-targeted humanitarian aids with less time and lower cost.

Related Work

Recent works have explored the potential of remote sensing data to assess societies' average exposure to disasters (Smith et al. 2019; Tellman et al. 2021). Addressing the limitations of conventional sources, another line of studies has combined remote sensing data with neural networks to measure disaster-incurred destruction directly. (Fujita et al. 2017; Amit and Aoki 2017; Duarte et al. 2018) applied deep-learning-based models to pre- and post-disaster imagery data to classify regions into damaged or undamaged areas. The proposed classification models were trained with a large number of damage labels specific to their input image data. With the methodological advances in machine learning techniques, semantic segmentation models built on detailed pixel-level ground-truth data also emerged (Potnis et al. 2019; Weber and Kané 2020; Wu et al. 2021; Bai et al. 2020; Gupta and Shah 2021).

Segmentation-based damage assessment using satellite imagery is an environmental application of change detection (Janalipour and Taleai 2017; Abbaszadeh Tehrani et al.

2022). Change detection is the process of identifying differences in the state of an object or phenomenon by observing it at different times (Singh 1989). Change detection models are either unsupervised or supervised. Unsupervised methods use clustering (Mehrotra et al. 2015; Celik 2009), thresholding (Khanbani, Mohammadzadeh, and Janalipour 2021; Ghanbari and Akbari 2015), or optimization (Kusetogullari, Yavariabdi, and Celik 2015) to find intrinsic differences between images without any prior information. On the other hand, the supervised methods are more common and can achieve higher accuracy than the unsupervised methods. However, they also require a large amount of pixel-level labels for training (Zhu 2017; Zou et al. 2022). Such types of labeling data are the most expensive labels in the computer vision field that are costly and time-consuming to gather.

The recent data revolution in the damage assessment field, the xBD dataset constructed by (Gupta et al. 2019), significantly extended the coverage of pixel-level damage labels worldwide. The objective of the xBD challenge is a particular semantic segmentation task: first to locate a building's footprint and then estimate the damage to each building. It contains pre- and post-disaster images along with pixel-level categorization for building damage. This dataset sparked the development of numerous segmentation-based damage detection models; for example, (Weber and Kané 2020; Wu et al. 2021) proposed U-Net-based models, and (Bai et al. 2020; Gupta and Shah 2021) used pyramid pooling modules. Both architectures are commonly used for image segmentation tasks. While the segmentation models crucially rely on the pixel-level ground-truth data as they exist in the original dataset, our grid-level prediction model did not require such granularity in the labeling data. Thus, we used a simplified version of the xBD dataset. We aimed to predict damage at rectangular-shaped grids of size $0.01km^2$, which can still incorporate the localized nature of disaster damages, but at much less computational cost and data constraint. This lightweight feature can be more helpful for rapid disaster responses in many countries.

One of the valuable global data sources in this aspect is provided by UNOSAT, the operational satellite applications program of UNITAR ¹. It provides building damage scales in 5-levels, ranging from *No-Damage* to *Destroyed*, along with their point locations. Despite their extensive coverage, especially in developing countries, the models developed for semantic segmentation tasks cannot utilize them as they do not provide pixel-level damage information. (Xu et al. 2019), the state-of-art model, utilized this data to build a binary classifier for detecting building damage at the grid level. We applied our disaster events to their model architecture and compared our model's performance relative to theirs.

Methods

We present a binary classification model that detects damage in a region from satellite image data. Our model identifies the rapid change mainly in building structures to distinguish the damage in target regions. Model input is a pair of images

¹<https://www.unitar.org/sustainable-development-goals/united-nations-satellite-centre-UNOSAT>

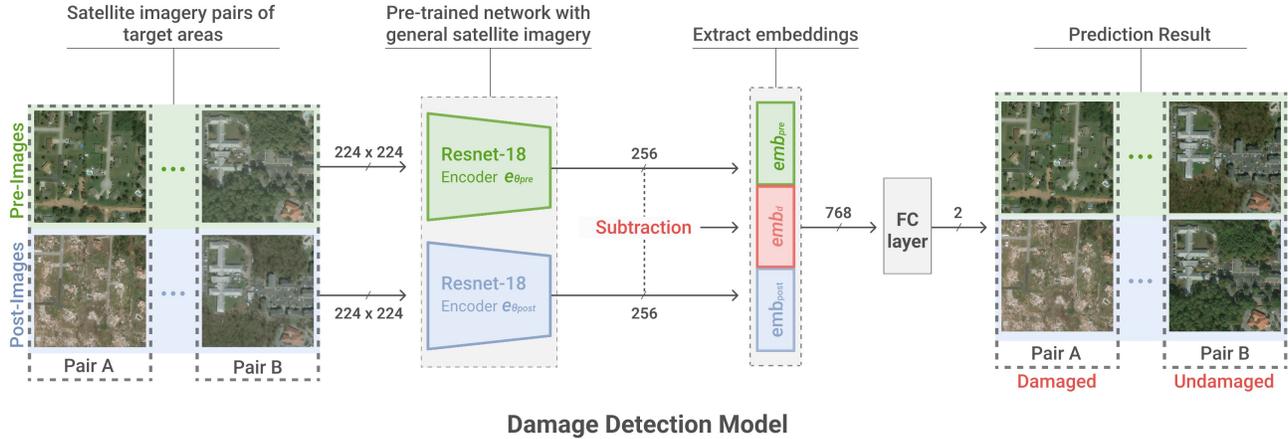


Figure 1: **The overall workflow of the proposed model.** Our damage detecting model takes pairs of pre- and post-disaster satellite images as inputs and classifies them into damaged pairs and undamaged pairs. We pre-trained the encoder with non-disaster satellite imagery before applying disaster satellite imagery. The pre-trained encoders extract embeddings from pre and post-disaster images of size 224×224 , positioning them in a 256-dimensional space. For the final classification, we used three embeddings: emb_{pre} , emb_{post} , emb_d , where emb_d is a subtraction of the first two. The fully connected layer takes them as input and predicts the binary label of each pair. The resulting binary label is assigned for each grid with a size of $0.093km^2$.

taken over the same geographic region before and after a disaster. Our model aims to determine whether a given region was substantially damaged due to a disaster. Due to the fixed temporal resolution of satellites, one great limit of the satellite-based approach is the difficulty of obtaining images right before and after a disaster. Such a time gap inevitably brings simple visual changes over the season and even general urban developments (e.g., the construction of buildings and roads). Our model needs to learn damage-specific features rather than those simple visual changes, to report disaster damage accurately.

We used transfer learning to train our model effectively, following the methods of (Jean et al. 2016; Xie et al. 2016). Transfer learning consists of two steps: pre-training and fine-tuning. Pre-training helps the model learn general low-level features of the image with more straightforward tasks on a large dataset. After the pre-training, the model is fine-tuned to fit the objective of the target task. We followed this step and first pre-trained the network with a simple classification task for one satellite image. Only satellite imagery of non-disaster situations was used during pre-training, which is much easier to gather. After the pre-training, we fine-tuned our model to detect the damage from a pair of satellite images. The model can learn to determine whether the given region is destroyed effectively, taking advantage of the learned general geographical features from the pre-training that are closely related to losses of properties and construction in the images.

Pre-training for General Feature Learning

Convolutional neural networks (CNN) for various tasks are often pre-trained with the ImageNet dataset (Deng et al. 2009). ImageNet-1000 is a large image classification dataset with over 1.2 million images with 1,000 classes. With such a large dataset, the model can learn mid-level visual features

such as edges and corners and be used as a generic feature extractor (Oquab et al. 2014). We also started from the ImageNet-1000 pre-trained ResNet-18 model. The ResNet-18 network is a convolutional neural network with 18 layers (He et al. 2016). We wanted our model to get familiar with the bird’s-eye viewpoint of satellite images before getting into the main task since it has never seen any satellite images. This approach is similar to that of (Xie et al. 2016), which used a chain of transfer learning to train a model for poverty mapping.

We used a set of satellite image \mathcal{I} from arbitrary regions and classified them into three classes according to their building density: more than 50%, under 50%, and 0%. We followed the method proposed by (Han et al. 2020a), considering the limited size of the labeled dataset. Human annotators labeled a total of 1,000 randomly chosen images such that every image in \mathcal{I} is labeled by five distinct annotators. For each image I_i^A , a label vector $y_i^A = [y_i^{high}, y_i^{low}, y_i^{zero}]$ ($0 \leq y_i^{high}, y_i^{low}, y_i^{zero} \leq 1$) is defined as the mean value of five one-hot building density response vectors from annotators. With this labeled pair set $L = \{(I_i^A, y_i^A)\}_{i=1}^{1000}$ and a larger set of unlabeled image set $U = \{I_j^U\}_{j=1}^M$ with $M = 145,921$, we pre-train ResNet-18 to conduct a three-class classification in a semi-supervised manner. Mean Teacher framework is utilized to penalize predictions of unlabeled data that are discrepancy between the student and teacher models (Tarvainen and Valpola 2017).

The pre-trained network is used as an encoder in our damage classification model, which extracts the embeddings that contain semantic information about the images as high-dimensional vectors.

Damage classification

After pre-training the model to understand the general features of satellite imagery with large non-disaster images, we fine-tuned the model to detect damage after a disaster. Figure 1 describes the overall workflow of our model. The disaster image pair set $D = \{(I_i^{pre}, I_i^{post}, \mathbf{y}_i)\}_{i=1}^N$ includes pairs with satellite images taken from the same pre-known damaged region before and after the disaster, and \mathbf{y}_i , the binary label which indicates damage to the region. The structure of our model resembles the Siamese network, which has two branches with identical structures and parameters (Chicco 2021). In our pseudo-siamese structure, however, the two identical branches do not share the same parameters, thereby giving the model greater flexibility (Zagoruyko and Komodakis 2015). In addition, damage assessment differs from conventional change detection tasks because the target of interest in pre- and post-disaster images differs. The post-disaster encoder should focus on distinct features related to disasters, such as building wreckages, whereas the pre-disaster encoder should focus on ordinary buildings. Our model could effectively learn such distinctive characteristics using a pseudo-siamese structure with two separate encoders.

The pre-trained ResNet-18 in the previous step was fine-tuned as image encoders $e_{\theta_{pre}}$ and $e_{\theta_{post}}$ for pre-event and post-event image sets, respectively, to learn specialized features separately. Each encoder reduces the dimensions of given satellite imagery, transforming it into high-dimensional embeddings representing the area’s visual features. Then, we generate an embedding difference vector $emb_{di} \in R^d$ and the final concatenated vector v_i for the i -th pair as follows:

$$emb_{di} = e_{\theta_{pre}}(I_i^{pre}) - e_{\theta_{post}}(I_i^{post}) \quad (1)$$

$$v_i = \text{Concat}(e_{\theta_{pre}}(I_i^{pre}), emb_{di}, e_{\theta_{post}}(I_i^{post})). \quad (2)$$

We used the difference of embedding in the final concatenated vector to consider the change caused by the disaster in imagery. By feeding the embedding difference to the classifier, the model can learn the relationship between the features more effectively.

Finally, a damage classifier is trained to minimize the loss \mathcal{L} defined between \mathbf{y}_i and predicted value as follows:

$$\mathcal{L} = \frac{1}{|D|} \sum_{I_i^{pre}, I_i^{post}, \mathbf{y}_i \in D} H(\mathbf{y}_i, \hat{\mathbf{y}}_i), \quad (3)$$

where H is a binary cross entropy loss function. The damage classifier derives the predicted value $\hat{\mathbf{y}}_i = W \cdot \mathbf{v}_i$, where $W \in R^{d \times \dim(\mathbf{y}_i)}$ is a trainable weight matrix of the classifier.

Data

Satellite imagery Dataset

We use the zoom level coordinate system to define satellite images’ size, resolutions, and alignment. The purpose of using the system is to maintain consistency with other studies that combine satellite images and machine learning techniques (Han et al. 2020b; Jean et al. 2016). The system is a tile-based coordinate system that divides the entire world into non-overlapping square-shaped images. At z of 0, the entire world map is fitted to a single image tile and an increase

of the zoom level by one results in half-sized image tiles. Thus, at z of 1, the world map is divided into 2×2 image tiles and hence has four times the resolution compared to $z=0$. A higher zoom level divides the world map into more tiles, and each tile will cover a smaller geospatial area at a higher resolution. The zoom levels for input images of machine learning models are selected considering the task and data availability. In the case of damage detection models, the target area is small and requires detailed information. Thus, the existing damage detection models use relatively high spatial resolution images (z of 16 to 19) compared to the other models, such as poverty mapping or land classification models. This paper also employs satellite images at a zoom level of 17 (tile size of $0.093km^2$, $1.193m/\text{pixel}$), allowing the model to consider the wreckage of buildings and roads. The RGB spectral bands are present in all of the images utilized in this research. For model pre-training, images of arbitrary regions were employed. We collected 146,921 satellite images spanning 2017 and 2018 from the ArcGIS World Imagery Wayback resource² for pre-training.

Disaster Dataset

The model was trained with the xBD dataset (Gupta et al. 2019), which is the largest building damage assessment dataset. The dataset includes pre- and post-images of various natural disasters, along with building annotations and damage scale labels. Damage labels span five types; four are related to damage scales of buildings (i.e., no-damage, minor-damage, major-damage, and destroyed), and one is non-buildings. The dataset covers 22 natural disaster events of seven different categories. In this study, we targeted water-related disasters, including hurricanes, tornadoes, tsunamis, and floods. By focusing on them, our model can better learn the characteristics of water-related disasters, which are quite different from other disasters, like geological events. We chose water-related disasters, as their occurrences and impacts have been observed to grow among the most vulnerable countries to climate change.

The satellite imagery in the xBD collection has a zoom level of 16, each image tile covering $0.373km^2$. We cropped each image tile into four half-sized image tiles. The resulting image tiles cover $0.093km^2$ each, which fits the zoom level of our interest, $z = 17$.

Since our target is classification at an image level, the xBD dataset with labels at the pixel level cannot be directly applied to our model. Reducing the complexity of the data also brings positive effects such as minimizing noises in the original data, including the mismatch of building boundaries between the pre- and post-disaster images and the uneven distribution of damage class labels. We used a simple method to aggregate the information from each building polygon to derive the binary disaster label for each image. The image is classified as damaged if the maximum damage level of buildings in the image is greater than or equal to major-damaged.

After applying this rule, we acquired 12,241 damaged pairs and 25,109 undamaged pairs of satellite images. Among the total of 37,350 images, 11,018 did not include any build-

²<https://livingatlas.arcgis.com/wayback/>

ings. The number of images, including major-damage, minor-damage, and destroyed, is 10,209, 10,046, and 6,413, respectively. The number of images with maximum damage levels of minor-damage, major-damage, and destroyed was 3,923, 5,828, and 6,413, respectively.

Results

Comparison with Baseline models

The xView2 challenge³ held in 2019 used the xBD dataset as a benchmark, and many models were proposed for damage assessment. However, our model cannot be compared to these models because the tasks and evaluation metrics are different. The xView2 challenge is defined at the building level, whereas our model detects damage at the grid level.

Before the xBD dataset was created, created a binary classification model which uses satellite imagery of pre- and post-disaster to detect the damage level. They built their own dataset covering only three disasters to train the model. Their model comes in four different versions: Channel Concatenate (CC), Post-disaster Only (PO), Twin-tower Concatenate (TTC), and Twin-tower Subtract (TTS) (Xu et al. 2019). To the best of our knowledge, this work is the only research that has tried to detect the damage at a grid level. Therefore, we compared their models as a baseline.

The pre- and post-images are concatenated in the CC model before being fed to the AlexNet (Krizhevsky, Sutskever, and Hinton 2012). Only post-disaster satellite images were used as input in the PO model. TTC and TTS models input both images to AlexNet’s first convolution layer to extract the activation map at the lower level. TTC employs the concatenated activation map of two images as an input for the remaining convolution layers, while TTS uses the subtraction of the two activation maps. We implemented the model and trained it with our simplified xBD dataset.

Table 1: Comparison of precision, recall, F1 score and accuracy with baselines.

	Prec	Rec	F1	Acc
CC	<u>0.762</u>	0.641	0.696	<u>0.817</u>
PO	0.716	0.621	0.663	0.794
TTC	0.741	<u>0.669</u>	<u>0.703</u>	0.815
TTS	0.759	0.648	0.697	0.816
Ours	0.796	0.767	0.781	0.859

Table 1 shows the baselines’ worst and best performance compared to our model. The experiments were conducted using five different train and test splits. The bold and underlined text indicates the best and second-best performance, respectively. Both twin-tower baseline models, TTC and TTS, showed better performance than the single-tower ones, CC and PO. This result indicates that the model can capture damage better when pre- and post-event images are embedded separately. TTC and TTS showed similar performances, but

³<https://xview2.org/>

TTC detected the damage slightly better since concatenation tends to lose less information in training than subtraction. Compared to all of the four baseline models, the result from our model showed higher performance in all four evaluation metrics. In particular, the large gap in recall showed our model’s ability to produce fewer false negatives on the damaged region. The proposed model utilizes embeddings of both images with a full encoder, which better preserves the information of the original image. We also took advantage of the benefits of TTC and TTS by putting together embeddings of pre- and post-images and subtracting them.

Ablation Study

To check the role of each component, we conducted an ablation study where we removed each component from the model and evaluated the model performance. In this manner, we can check which components contribute the most and which may be removed.

- **w/o emb_d .** To understand the impact of the embedding difference vector, emb_d , on model performance, we removed the emb_d and used only the two embedding vectors.
- **w/o sep .** The encoder for pre and post-images is trained separately in the full model and does not share the parameters. We used the identical network for both images to test whether separating the two networks helps the model performance.
- **w/o emb_d, sep .** We removed emb_d and shared one network for pre and post-images. This model can measure how the two components affect each other.
- **w/o pre-training.** We trained the model from the random initialized state to understand how the model gets general information from the building density classification task.

Table 2 compares the performance of each model. The experiments were conducted using five different train and test splits. The full model outperformed the others, demonstrating the value of each component. Interestingly, without-pretraining model had the best overall performance among the ablation models, followed by without emb_d model. While the without emb_d model shows relatively high performance, the without emb_d, sep showed the poorest performance, demonstrating the need for independent encoders.

This contradicts the findings of (Weber and Kané 2020), which claimed that the model performs better when one network is shared for pre and post-disaster images.

Table 2: Precision, recall, F1 score, and accuracy of the ablation models.

Ablation Model	Prec	Rec	F1	Acc
w/o emb_d	0.788	<u>0.765</u>	0.776	0.855
w/o sep	<u>0.797</u>	0.752	0.774	0.856
w/o emb_d, sep	0.784	0.761	0.771	0.852
w/o pre-training	0.807	0.755	<u>0.780</u>	0.860
Full model	0.796	0.767	0.781	<u>0.859</u>

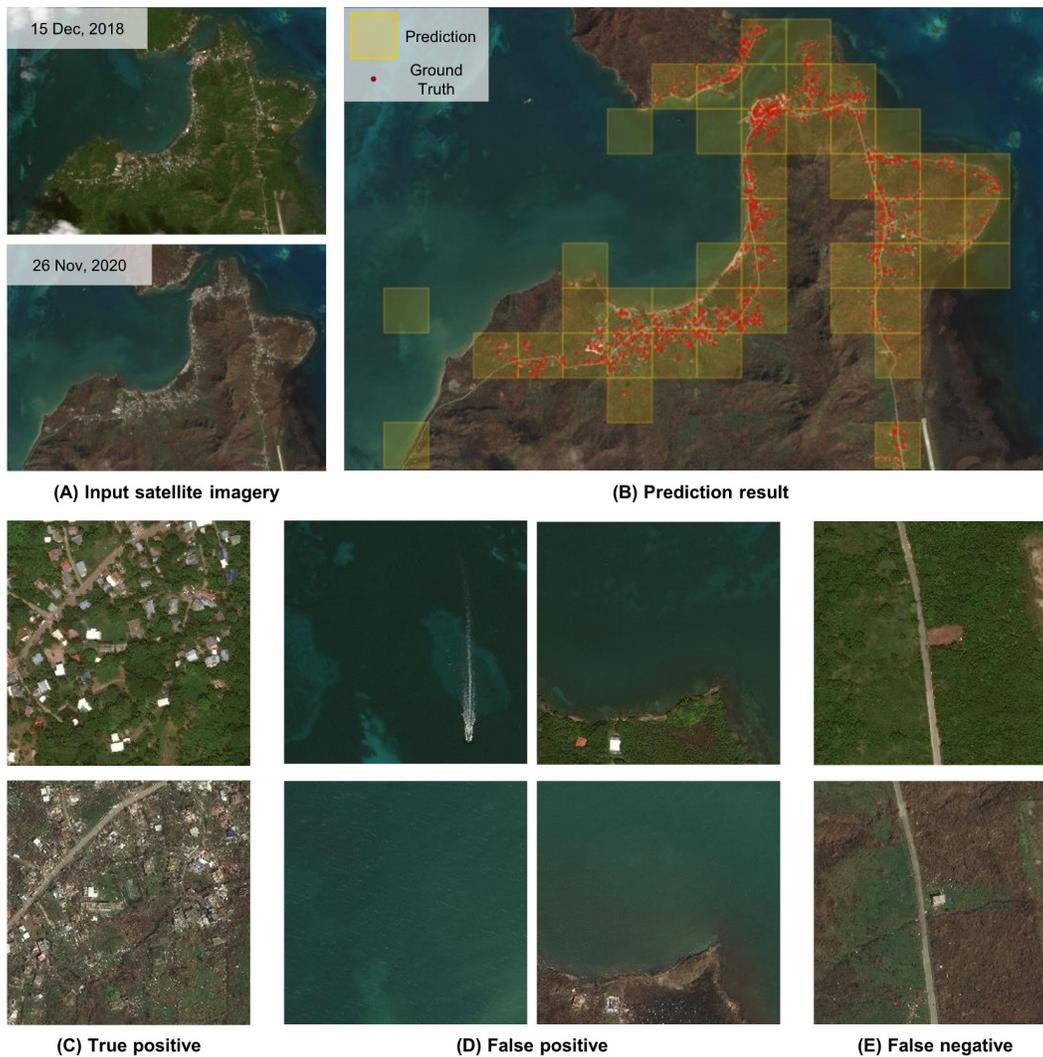


Figure 2: **The prediction result of our model on Providencia Island, Colombia.** (A) shows the input satellite imagery captured on December 15, 2018, and November 26, 2020. (B) shows the result of our model. The red dot is the point of damaged buildings in the UNOSAT ground truth data, and the yellow box indicates the grid where our model reports damage. The zoomed grids of true positive, false positive, and false negative cases are shown in (C), (D), and (E), respectively.

The cause is likely due to the difference in task objective — classification vs. object identification. The network sees the doubled training data when two images are shared, which can help with detailed object detection learning. Our model, on the other hand, is a classification model that considers the entire context rather than just a few data points. Separately focusing on the before and after contexts can help the model perform better. Also, the *without sep* model has the second best precision with the poorest recall among all ablation models. When the pre- and post-embeddings are extracted using the same encoder, the model is more likely to report observations as undamaged more frequently, resulting in higher precision with poor recall.

Case Study on Hurricane Iota

Having seen the potential for using computer vision techniques to assess disaster damage, we now introduce one case study of Hurricane Iota, which hit Colombia’s Providencia Island in November 2020. The xBD dataset does not include Hurricane Iota. Therefore, it is valuable to test our detection model for the northern part of the island to demonstrate the applicability of the model. According to (IFRC 2021), a substantial proportion of the island’s infrastructure, an estimated 98%, was destroyed, and 95% of its population was affected. We compared our model’s prediction result to the ground truth label generated by (UNITAR 2020).

Figure 2 shows the input images and the prediction results. We utilized Maxar SecureWatch to access pre- and post-disaster satellite imagery of the target area. The pre-

disaster image was taken on 15 December 2018, and the post-disaster image was taken on 26 November 2020. The disaster occurred in November 2020, yet we could only find a cloud-free pre-disaster image of the region from December 2018, as shown in Figure 2(A). Both images contain RGB spectral bands, and we processed the images to match the specifications for the model training dataset described in Section 4.1. We cropped the images into tiles that do not overlap, each covering 0.093km^2 (zoom level of 17). The model successfully detected damage in the disaster-affected area as shown in Figure 2(B), which shows ground truth labels (red dots) and damage prediction (shaded tiles). Even though the two images have different color compositions, the overlap between ground truth labels and predictions indicates that the model successfully distinguished the characteristics of damaged buildings from other changes (with the 2020 image missing green space). Our model had a 97.5% accuracy and an F1 score of 0.851, with 43 true positives, six false positives, and nine false negatives out of 595 grids.

An example pre- and post-image pair of a damaged region is shown in Figure 2(C). We also find that some regions were falsely incorrectly labeled as damaged when they were not affected or vice versa (i.e., a false positive) and that some damage was overlooked (i.e., a false negative). The left image pair in Figure 2(D) shows a false positive case due to a modeling error. The majority of the training data comes from inland images. Hence, a passing ship captured in the pre-disaster image served as a noise source. Filtering out the sea using shoreline data could be one way to avoid faulty detection caused by moving objects. The right image pair in Figure 2(D) shows another false positive case due to erroneous human labels. The two buildings in the image appear to be damaged, but the ground truth labels were missing in the region due to a lack of human resources. Our model correctly identified damage in the area, demonstrating its ability to recognize damaged structures.

The image pair in Figure 2(E) shows a false negative case due to limited temporal resolution. The only satellite imagery available in the area prior to the disaster was from December 15, 2018, nearly two years before the disaster. During the period the two satellite images were taken, a new building was being built that was destroyed by the hurricane. Given that there is no visible building structure in the pre-disaster image, the model failed to detect damage in the area, resulting in a false negative case. This limitation in temporal resolution, however, is likely to be resolved thanks to the increasing availability of small satellites.

Conclusion

This study demonstrated how computer vision techniques could be used to develop a data-driven strategy for disaster response. Our lightweight model could successfully identify the damaged areas from water-related disasters with only pre- and post-satellite images and simple damage labels. Since our model has little data constraint regarding both the input and the label data, it can help disaster responses in many settings where the previous deep-learning-based detection models were not applicable. For example, the development agency

could prioritize their resources based on the sum of our binary labels within its administrative units, and also locate precisely where to deploy the tents and shelters within the units. This timely and accurate damage estimates provided by our model could help promote the adaptive capacities of many vulnerable countries.

Experimental results demonstrated that our model outperformed existing baselines in detecting water-related damage, achieving high accuracy of 91.4%. Our model successfully identified the damaged areas even with sparse damage labels. The ablation study confirmed that the unique embedding method applied on pre- and post-disaster images via two separate encoders was critical to the performance. Moreover, pre-training our model with non-disaster satellite imagery before learning disaster-specific features was a critical factor in our model's success in damage detection, which is far more challenging. Lastly, evaluations of our model in a case study demonstrated its robustness. Specifically, the case study result suggested that our model is applicable to real-world responses with high accuracy.

Future work may improve our work in several aspects. First, integrating other socioeconomic indicators into a model can help find socioeconomically more vulnerable regions. For example, if two areas with the same magnitude of destruction have starkly different socioeconomic vulnerabilities, the model should indicate more severe damage for the more vulnerable area. Another extension of our approach can be made to identify damage in man-made disasters. Since they typically target more densely populated civilian buildings relative to natural disasters, future work could focus on producing damage estimates at a finer scale. Lastly, future innovations could work on harmonizing the increasingly diverse satellite sources and producing robust damage estimates across the different sources, as in real-world responses, it is most often infeasible to choose a particular type of satellite imagery for the damage assessment.

References

- Abbaszadeh Tehrani, N.; Mohd Shafri, H. Z.; Salehi, S.; Chanussot, J.; and Janalipour, M. 2022. Remotely-Sensed Ecosystem Health Assessment (RSEHA) model for assessing the changes of ecosystem health of Lake Urmia Basin. *International Journal of Image and Data Fusion*, 13(2): 180–205.
- Amit, S. N. K. B.; and Aoki, Y. 2017. Disaster detection from aerial imagery with convolutional neural network. In *2017 international electronics symposium on knowledge creation and intelligent computing (IES-KCIC)*, 239–245. IEEE.
- Amundsen, H.; Berglund, F.; and Westskog, H. 2010. Overcoming barriers to climate change adaptation—a question of multilevel governance? *Environment and Planning C: Government and Policy*, 28: 276–289.
- Bai, Y.; Hu, J.; Su, J.; Liu, X.; Liu, H.; He, X.; Meng, S.; Mas, E.; and Koshimura, S. 2020. Pyramid pooling module-based semi-siamese network: A benchmark model for assessing building damage from xBD satellite imagery datasets. *Remote Sensing*, 12(24): 4055.
- Bakkensen, L. A.; Shi, X.; and Zurita, B. D. 2018. The impact of disaster data on estimating damage determinants and

- climate costs. *Economics of Disasters and Climate Change*, 2(1): 49–71.
- Cao, Q. D.; and Choe, Y. 2020. Building damage annotation on post-hurricane satellite imagery based on convolutional neural networks. *Natural Hazards*, 103: 3357—3376.
- Celik, T. 2009. Unsupervised change detection in satellite images using principal component analysis and k -means clustering. *IEEE geoscience and remote sensing letters*, 6(4): 772–776.
- Chicco, D. 2021. *Siamese Neural Networks: An Overview*, 73–94. New York, NY: Springer US.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255. Ieee.
- Duarte, D.; Nex, F.; Kerle, N.; and Vosselman, G. 2018. Satellite Image Classification of Building Damages Using Airborne and Satellite Image Samples in a Deep Learning Approach. *ISPRS Annals of Photogrammetry, Remote Sensing & Spatial Information Sciences*, 4(2).
- Edmonds, D. A.; Caldwell, R. L.; Brondizio, E. S.; and Siani, S. M. 2020. Coastal flooding will disproportionately impact people on river deltas. *Nature communications*, 11(1): 1–8.
- Field, C. B.; Barros, V.; Stocker, T. F.; and Dahe, Q. 2012. *Managing the risks of extreme events and disasters to advance climate change adaptation: special report of the intergovernmental panel on climate change*. Cambridge University Press.
- Fujita, A.; Sakurada, K.; Imaizumi, T.; Ito, R.; Hikosaka, S.; and Nakamura, R. 2017. Damage detection from aerial images via convolutional neural networks. In *2017 Fifteenth IAPR international conference on machine vision applications (MVA)*, 5–8. IEEE.
- Ghanbari, M.; and Akbari, V. 2015. Generalized minimum-error thresholding for unsupervised change detection from multilook polarimetric SAR data. In *2015 IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, 1853–1856. IEEE.
- Guiteras, R.; Jina, A.; and Mobarak, A. M. 2015. Satellites, Self-Reports, and Submersion: Exposure to Floods in Bangladesh. *American Economic Review*, 105(5): 232–36.
- Gupta, R.; Goodman, B.; Patel, N.; Hosfelt, R.; Sajeev, S.; Heim, E.; Doshi, J.; Lucas, K.; Choset, H.; and Gaston, M. 2019. Creating xBD: A dataset for assessing building damage from satellite imagery. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, 10–17.
- Gupta, R.; and Shah, M. 2021. Rescuenet: Joint building segmentation and damage assessment from satellite imagery. In *2020 25th International Conference on Pattern Recognition (ICPR)*, 4405–4411. IEEE.
- Hallegatte, S.; Green, C.; Nicholls, R. J.; and Corfee-Morlot, J. 2013. Future flood losses in major coastal cities. *Nature Climate Change*, 3: 802–806.
- Han, S.; Ahn, D.; Cha, H.; Yang, J.; Park, S.; and Cha, M. 2020a. Lightweight and robust representation of economic scales from satellite imagery. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 428–436.
- Han, S.; Ahn, D.; Park, S.; Yang, J.; Lee, S.; Kim, J.; Yang, H.; Park, S.; and Cha, M. 2020b. Learning to score economic development from satellite imagery. In *Proceedings of the CM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD)*, 2970–2979.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 770–778.
- Henkel, M.; Eunjee, K.; and Magontier, P. 2022. The Unintended Consequences of Post-Disaster Policies for Spatial Sorting.
- Hunter, L. M. 2005. Migration and environmental hazards. *Population and environment*, 26(4): 273–302.
- IFRC. 2021. *DREF Final Report Colombia - Hurricane Iota (MDR0017)*. International Federation of Red Cross and Red Crescent Societies.
- IPCC. 2022. *Climate Change 2022: Impacts, Adaptation, and Vulnerability*. Cambridge University Press, In Press.
- Janalipour, M.; and Taleai, M. 2017. Building change detection after earthquake using multi-criteria decision analysis based on extracted information from high spatial resolution satellite images. *International journal of remote sensing*, 38(1): 82–99.
- Jean, N.; Burke, M.; Xie, M.; Davis, W. M.; Lobell, D. B.; and Ermon, S. 2016. Combining satellite imagery and machine learning to predict poverty. *Science*, 353(6301): 790–794.
- Khanbani, S.; Mohammadzadeh, A.; and Janalipour, M. 2021. A novel unsupervised change detection method from remotely sensed imagery based on an improved thresholding algorithm. *Applied Geomatics*, 13(1): 89–105.
- Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25.
- Kusetogullari, H.; Yavariabdi, A.; and Celik, T. 2015. Unsupervised change detection in multitemporal multispectral satellite images using parallel particle swarm optimization. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 8(5): 2151–2164.
- Lin, Y.; McDermott, T. K.; and Michaels, G. 2021. Cities and the sea level.
- Mehrotra, A.; Singh, K. K.; Nigam, M. J.; and Pal, K. 2015. Detection of tsunami-induced changes using generalized improved fuzzy radial basis function neural network. *Natural Hazards*, 77(1): 367–381.
- Moser, S. C.; and Ekstrom, J. A. 2010. A framework to diagnose barriers to climate change adaptation. *PNAS*, 107(51): 22026–22031.
- Oquab, M.; Bottou, L.; Laptev, I.; and Sivic, J. 2014. Learning and transferring mid-level image representations using convolutional neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1717–1724.

- Potnis, A. V.; Shinde, R. C.; Durbha, S. S.; and Kurte, K. R. 2019. Multi-class segmentation of urban floods from multi-spectral imagery using deep learning. In *IGARSS 2019-2019 IEEE International Geoscience and Remote Sensing Symposium*, 9741–9744. IEEE.
- Raker, E. J. 2020. Natural Hazards, Disasters, and Demographic Change: The Case of Severe Tornadoes in the United States, 1980–2010. *Demography*, 57(2): 653–674.
- Rentschler, J.; Salhab, M.; and Jafino, B. A. 2022. Flood exposure and poverty in 188 countries. *Nature Communications*, 13(1): 3527.
- Singh, A. 1989. Review article digital change detection techniques using remotely-sensed data. *International journal of remote sensing*, 10(6): 989–1003.
- Smith, A.; Bates, P. D.; Wing, O.; Sampson, C.; Quinn, N.; and Neal, J. 2019. New estimates of flood exposure in developing countries using high-resolution population data. *Nature communications*, 10(1): 1–7.
- Tarvainen, A.; and Valpola, H. 2017. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *Advances in neural information processing systems*, 30.
- Tellman, B.; Sullivan, J.; Kuhn, C.; Kettner, A.; Doyle, C.; Brakenridge, G.; Erickson, T.; and Slayback, D. 2021. Satellite imaging reveals increased proportion of population exposed to floods. *Nature*, 596(7870): 80–86.
- UNITAR. 2020. *Damage Assessment in Providencia Island, Providencia and Santa Catalina Department, Colombia as of 21 November 2020*. United Nations Institute for Training and Research.
- Vetrivel, A.; Gerke, M.; Kerle, N.; Nex, F.; and Vosselman, G. 2018. Disaster damage detection through synergistic use of deep learning and 3D point cloud features derived from very high resolution oblique aerial images, and multiple-kernel-learning. *ISPRS journal of photogrammetry and remote sensing*, 140: 45–59.
- Weber, E.; and Kané, H. 2020. Building disaster damage assessment in satellite imagery with multi-temporal fusion. *arXiv preprint arXiv:2004.05525*.
- Wu, C.; Zhang, F.; Xia, J.; Xu, Y.; Li, G.; Xie, J.; Du, Z.; and 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.
- Xie, M.; Jean, N.; Burke, M.; Lobell, D.; and Ermon, S. 2016. Transfer learning from deep features for remote sensing and poverty mapping. In *Thirtieth AAAI Conference on Artificial Intelligence*.
- Xu, J. Z.; Lu, W.; Li, Z.; Khaitan, P.; and Zaytseva, V. 2019. Building damage detection in satellite imagery using convolutional neural networks. *arXiv preprint arXiv:1910.06444*.
- Zagoruyko, S.; and Komodakis, N. 2015. Learning to compare image patches via convolutional neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4353–4361.
- Zhu, Z. 2017. Change detection using landsat time series: A review of frequencies, preprocessing, algorithms, and applications. *ISPRS Journal of Photogrammetry and Remote Sensing*, 130: 370–384.
- Zou, L.; Li, M.; Cao, S.; Yue, F.; Zhu, X.; Li, Y.; and Zhu, Z. 2022. Object-Oriented Unsupervised Change Detection Based on Neighborhood Correlation Images and k-Means Clustering for the Multispectral and High Spatial Resolution Images. *Canadian Journal of Remote Sensing*, 48(3): 441–451.