

# DeepDamageNet: A two-step deep-learning model for multi-disaster building damage segmentation and classification using satellite imagery

Irene Alisjahbana    Jiawei Li    Ben Mullet    Yue Zhang  
Stanford University

{alisjahbana, jiaweili, bmullet, yzhang16}@stanford.edu

## Abstract

*Satellite imagery has played an increasingly important role in post-disaster building damage assessment. Unfortunately, current methods still rely on manual visual interpretation, which is often time-consuming and can cause very low accuracy. To address the limitations of manual interpretation, there has been a significant increase in efforts to automate the process. We present a solution that performs the two most important tasks in building damage assessment, segmentation and classification, through deep-learning models. We show our results submitted as part of the xView2 Challenge, a competition to design better models for identifying buildings and their damage level after exposure to multiple kinds of natural disasters. Our best model couples a building identification semantic segmentation convolutional neural network (CNN) to a building damage classification CNN, with a combined F1 score of 0.66, surpassing the xView2 challenge baseline F1 score of 0.28. We find that though our model was able to identify buildings with relatively high accuracy, building damage classification across various disaster types is a difficult task due to the visual similarity between different damage levels and different damage distribution between disaster types, highlighting the fact that it may be important to have a probabilistic prior estimate regarding disaster damage in order to obtain accurate predictions.*

## 1. Introduction

In the past two decades, the number of natural disasters that occur worldwide has increased dramatically, causing thousands of casualties and millions of dollars of loss to buildings every year. It is no surprise that rapid building damage assessment is considered to be one of the most important phases after major disasters, aiming to evaluate the state of the buildings after a given event. In addition to providing information to first responders for rescue operations and relief efforts, it also provides useful information for lo-

cal government and related organizations for the estimation of economic loss and damage distribution in the subsequent reconstruction and recovery phase [1]. Though significant efforts have been put into rapid post-disaster assessment, it remains a very challenging task.

Traditional methods of building damage assessments rely on ground-based surveys and field assessments to assess building damage, which involves dispatching experts to the field to manually inspect buildings individually. Though it is still perhaps the most reliable method of assessment, this method is often time-consuming to apply across large geographic regions, subjective and prone to inconsistencies in data quality [2].

Technological advancements in the remote sensing field have created new opportunities to address the limitations of field-based assessments and have been recognized to be useful for rapid post-disaster response [16, 5, 20, 6]. It enables the collection of data that is independent and rapidly available across large geographical regions. Unfortunately, most efforts for remote sensing-based building damage assessment today still rely on manual visual interpretation and volunteers, which can result in very low accuracy [13]. Though existing emergency mapping services such as the Copernicus Emergency Management Service (EMS) provide damage maps a few hours after a disaster, they are still typically carried out by human operators [15]. This involves manually detecting and segmenting the buildings before classifying them based on their damage level.

To address the limitations of manual interpretation, there has been a significant increase in efforts to automate the process. Research has shown that using machine learning and deep learning techniques, in particular convolution neural networks (CNN) can achieve high accuracy. In addition, the use of automated methods allows the independent and rapid assessment that can easily be applied across large geographical areas. Nonetheless, most of the work currently being done focuses on the two tasks (segmentation and classification) separately.

To that end, the xView2 Challenge aims to automate the process of building damage assessment from pre- and post-

disaster satellite imagery through the creation of accurate and efficient machine learning models. To support the advancement of the automation process, the xBD dataset [7], a new large-scale dataset containing pre- and post- disaster satellite imagery and ground truth labels from a variety of disasters and over 160,000 buildings was released.

Through this challenge, we present a framework to perform segmentation and classification of buildings from pairs of pre- and post- disaster images for the purposes of building damage assessment. In particular, we explored various methods and model architectures that would be generalizable across various types of natural disasters, such as earthquakes, tsunamis, floods, volcanic eruptions, wildfires, and hurricanes. We present our results and analyses on the advantages and limitations of using automated methods for the purposes of building damage assessment across different disasters.

## 2. Related Work

Given its obvious advantages, damage assessment via remote sensing has become a topic of much research in the last few decades. Verstappen [23] proposed one of the original frameworks for evaluating the use of remote sensing data in natural hazard risk reduction, and called for the use of remote sensing tools in combating both “instantaneous” disasters like cyclones, drought, floods, tsunamis, earthquakes, and volcanic hazards, and “creeping” disasters like desertification. Frameworks have also been proposed for the use of remote sensing data to predict high-risk areas, before disaster strikes [21].

In the pre-machine learning era, practitioners relied on classical segmentation and classification techniques for map production, *e.g.* a “region-growing and knowledge-based segmentation approach” that involves merging visually similar regions to build a disaster map [22]. It is only more recently that researchers have applied machine learning techniques to this problem, with increasing degrees of success.

Thanks to the recent advancement of the object detection and segmentation algorithms in computer vision, many CNN models have been developed and applied for building footprint segmentation based on satellite data. Some of the most common models include U-Net [18], and the more recent feature pyramid network (FPN) [12]. U-Net has been successfully used in the satellite image building segmentation challenge SpaceNet 4. Chen et al. [4] used the FPN model for building segmentation using the Aerial Imagery for Roof Segmentation (AIRS) dataset. This dataset contains aerial imagery with 7.5cm resolution and contains over 220,000 buildings. The overall performance of the FPN model was evaluated through intersection over union (IoU) and F1-score metrics, in which they achieved a score of 0.882 and 0.937 respectively.

One limitation of using semantic segmentation for the

purposes of building damage assessment is the fact that semantic segmentation treats multiple objects of the same class as a single entity, whereas in reality, each building is required to be classified individually. As a result, instance segmentation has been increasingly used, as it identifies multiple objects of the same class as distinct individual objects [3, 10]. He et al. [8] first introduced Mask R-CNN, a framework for generating high-quality segmentation masks for each instance. Since then, Mask R-CNN has been used for a variety of different application as it is fast to implement and easy to generalize to different tasks including building footprint segmentation. One such application was for CrowdAI’s mapping challenge 2018 [14], aiming to segment buildings for humanitarian assistance. Using high resolution RGB satellite images, the challenge baseline achieved an IoU of 0.697.

In addition to building footprint segmentation, deep learning methods are widely used to classify the building damage level. Xu et al. [26] compared four architectures to perform binary classification (damage or not damaged) for a certain building cropped from a satellite imagery. Their results show that a twin-tower structure to fuse the pre-disaster and post-disaster images can improve the classification accuracy. In their experiments, the best model used two separate CNNs to extract features of the pre- and post-disaster images, before subtracting the extracted features along the channel dimension. They used the area under the RoC curve (AUC) as the primary metric of model, and achieved a score of 0.8302 AUC score on the Haiti earthquake dataset.

As seen above, most work focuses on either segmentation or classification. Gupta et al. [7] introduced a baseline model for the xView2 challenge that combines the segmentation and classification model. The segmentation model is based on an altered U-Net architecture, achieving an IoU of 0.66 and F1 score of 0.80. On the other hand, the classification model is based on a modified ResNet50. The weighted F1 is used as the primary metric for evaluating the overall dataset, obtaining a weighted F1 score of 0.2564.

To our knowledge, no previous work has been done to create a model that would be generalizable to all types of natural disasters. Most work have focused on a single disaster type. For example, Xu et al. [26] trained a model for earthquake disasters damage whereas Rudner et al. [19] for flooded buildings. Xu et al. [26] quantified how well the models will generalize to future disasters by training and testing models on different disaster events, and found that the AUC reduces for cross-region results. This suggests that cross-region prediction is difficult, even within the same disaster type (i.e earthquakes).

### 3. Dataset

#### 3.1. Dataset Description

The dataset used is the xBD dataset, released to supplement the xView2 Challenge [7]. A total of 2799 pairs of pre- and post-disaster RGB satellite images are provided in the training set, whereas 933 pairs of pre- and post-disaster satellite images are available in the test set. All imagery is obtained from DigitalGlobe satellite with a resolution of around 0.5m ground sample distance. The size of the images are  $1024 \times 1024$  pixels. An example of a pre- and post-disaster image can be seen in Figure 1.

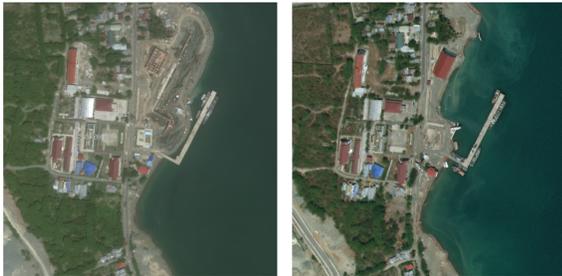


Figure 1: Pre-disaster (left) and its corresponding post-disaster satellite image

In addition to imagery, the xBD dataset provides ground truth data that includes the coordinates of each building polygons and the corresponding damage scale level for each building. The damage scale level that is used in the xBD dataset is a joint damage scale that ranges from no damage (0) to destroyed (3). The ground truth data is only available for the training set. A detailed description of the damage scale definition, image collection method, annotation process and quality control of the dataset is available in [7].

#### 3.2. Dataset Statistics

The training dataset consists of 10 different disasters, ranging from hurricanes, to earthquakes and is gathered from over 15 countries at various times of the year. It should be noted that the number of images available for each disaster event is highly unbalanced, with the Guatemala volcanic eruption having the least number of images.

Though the distribution of images based on each disaster is highly imbalanced, the number of buildings that are present in each image also varies significantly between each disaster (Figure 2). For some disasters, there are images where a single image contains over 1000 buildings. On the other hand, there are disasters with images that contain less than 10 buildings (Figure 3).

Furthermore, the damage level for the buildings in the training dataset is also imbalanced. Non-damaged buildings takes up over 70% of the training data. However, even

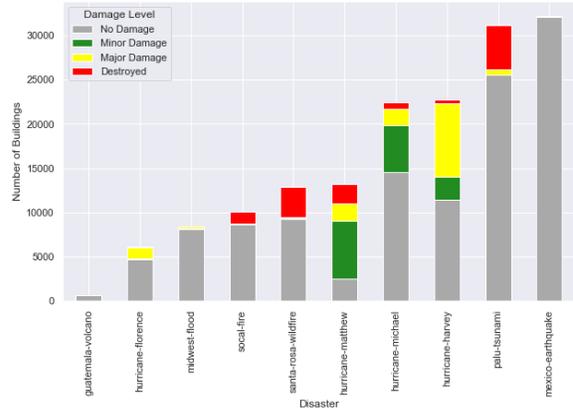


Figure 2: Distribution of building damage level in training set



Figure 3: Differences in building densities

though there are mostly non-damage buildings in the complete training data, this distribution varies significantly between each disaster as seen in Figure 3. For example, there are more damaged buildings compared to non-damaged buildings for Hurricane Matthew.

### 4. Methods

#### 4.1. Pre-Processing

Before splitting the xBD dataset into train/val set, we manually looked through the dataset for features that might be problematic for our model. For example, there were some images that are covered by clouds, causing occlusion to some buildings. However, the statistical analysis showed that less than 1% of the images were covered by clouds. As a result, we decided to leave those images in our dataset.

We then split the data into training and validation with 80% to 20% ratio. Since the hold-out test set used by the challenge is drawn from the same disaster type distribution as the data released, our validation set was created by preserving this distribution. For each disaster, we took 20% of the images for the validation set, and 80% for the training

Disaster	Total	Train. Set	Val. Set
Guatemala Volcano	18	15	3
Hurricane Florence	319	256	63
Hurricane Harvey	319	256	63
Hurricane Matthew	238	191	47
Hurricane Michael	343	191	68
Mexico Earthquake	121	97	24
Midwest Flood	279	224	55
Palu Tsunami	113	91	22
Santa Rosa Wildfire	226	181	45
Socal Fire	823	659	164
Total Images	2799	2245	554

Table 1: Number of images in the training and validation sets based on disaster

set. The number of images in the training and validation set for each disaster can be seen in Table 1.

We recognize that even with the split generated above, the building damage class imbalance persists. A majority of the buildings in our split dataset will still be non-damaged buildings. In our current approach, we did not do any further distribution of data in order to reduce this imbalance.

## 4.2. Model Architecture

We experimented with two techniques for performing the objective task. First, we built and trained an end-to-end model, which was designed to simultaneously perform both segmentation of building footprints and classification of damage extent. Second, we implemented a series of two-step models, which treated building segmentation and damage classification as two distinct tasks, with separate models optimized for each of these tasks (Figure 4).

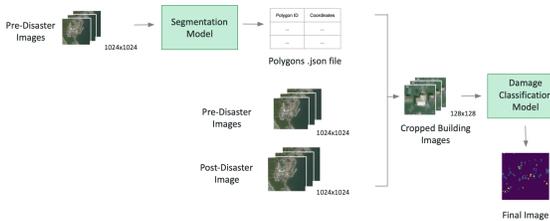


Figure 4: Architecture of two-step model. The pre-disaster image is fed into a segmentation model to obtain building polygon coordinates. This result in addition to pre- and post- disaster cropped building image is then fed into a damage classification model to obtain the final image.

### 4.2.1 End-to-End Model

For the end-to-end model, the inputs were paired pre-disaster and post-disaster images both of size  $1024 \times 1024 \times 3$  and the output shows the pixel-level classification results of size  $1024 \times 1024$ . The classification values ranged within 0 to 4, where 0 represents no-building and 1-4 represented different building damage levels from no-damage to destroyed, as discussed above. The model architecture chosen here was U-Net with a ResNet-34 encoder pre-trained on ImageNet dataset. To incorporate both pre-disaster image and post-disaster image, two architectures were proposed (Figure 5).

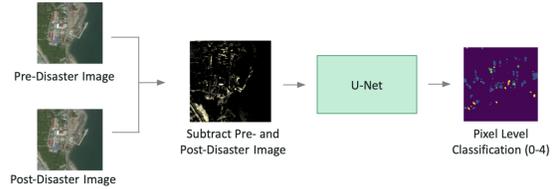


Figure 5: Architecture of end-to-end model. The pre- and disaster image is subtracted before feeding into a U-Net to perform simultaneous segmentation and classification.

The intuitive way to make use of the pre- and post-disaster images is to compare the difference between them. Thus, we subtracted the pre-disaster image pixel values from the post-disaster image to obtain a new image of size  $1024 \times 1024 \times 3$ . This was then treated as a new input to be fed directly to the ResNet without making any further changes.

### 4.2.2 Two-Step Model: Building Segmentation

In the two-step model, the first step is to segment the building footprint and extract the polygon coordinates of each building. We tried two different model architectures for this task, semantic segmentation and instance segmentation, which are commonly used for building footprint segmentation.

**Instance Segmentation:** The objective of the instance segmentation is to label each foreground pixel with a corresponding object and instance - a combination of semantic segmentation and object detection. For this, we used Mask R-CNN [8]. For ease of implementation, we followed the implementation from the Facebook AI research platform Detectron2 [25]. The backbone model of the Mask R-CNN is a ResNet-50 [9] with Feature Pyramid Network (FPN) [12]. A faster R-CNN [17] model is combined with the backbone model to perform the semantic segmentation and object detection.

**Semantic Segmentation:** A ResNet-50 Feature Pyramid Network (the same backbone model as the Mask R-CNN) was used to conduct the semantic segmentation of the building footprint. The model input is the pre-disaster image and the output is the binary mask of each image with 0 (background) and 1 (building). After obtaining the semantic mask of each image, we outlined every single building to extract their polygon coordinates using Imantics [11], a Python package.

### 4.2.3 Two step model: Building Classification

After segmentation maps were generated using the techniques outlined in the previous section, building-by-building damage classification was performed. The model takes as input a close cropping of an individual building in both pre- and post-disaster images and predicts the damage level score (1 for no damage, 4 for destroyed). The basic structure of the damage classification model includes two CNN feature extractors which are used to pre-process the pre- and post-disaster building images. Then the features maps can be used alone or combined with other features to conduct the prediction.

#### Pre & Post Concatenation:

The damage classification model concatenating pre- and post- feature maps can be seen in Figure 6a. Similar to what was done by Xu et al. [26], we implemented a twin-tower model architecture which concatenates the output features of a ResNet-50 model (pre-trained on ImageNet). This concatenated feature map is then flattened and fed to several fully connected layers, which outputs a damage level classification score ranging from 1 (No damage) to 4 (Destroyed).

#### Pre & Post Concatenation with Disaster Prediction:

As mentioned before, different disasters in different regions will be covered in this single prediction model. We found that the building damage characteristics of different disasters are significantly different, which motivates us to incorporate the disaster-type as an extra features in the Pre & Post Concatenation model in Figure 6b.

The xview2 dataset provides a unique disaster-type label for each training image, which enables us to train a separate disaster type prediction models. The disaster-type will then be represented as an one-hot vector and concatenated with the pre- and post- features extracted from CNN and the combined features are used for the final prediction.

**Pre & Post Concatenation with SSIM:** As an extension, we also experimented by adding a third branch with additional features before concatenating and feeding to the fully connected layers. The structural similarity index (SSIM) was added as the additional feature, a metric that measures the similarity between two images [24].

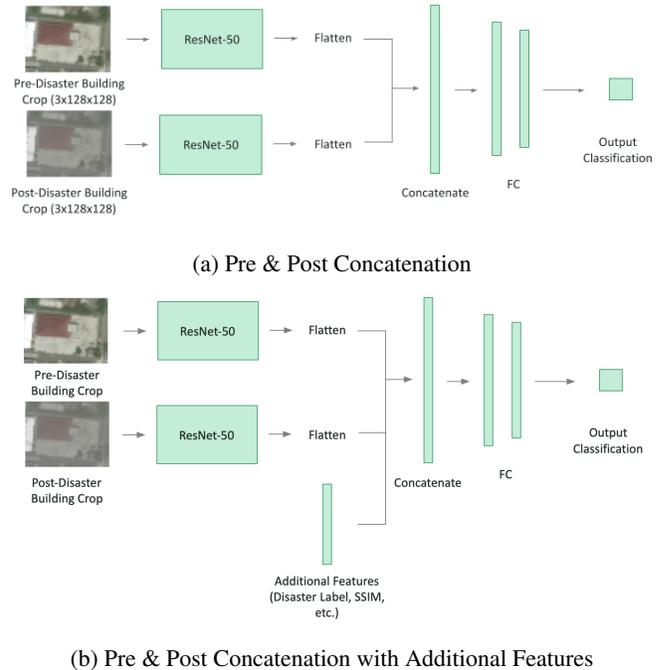


Figure 6: Architecture of classification model. a) A twin-tower architecture concatenates the output features of a ResNet-50 model from pre- and post- disaster cropped building images. b) An additional third branch can be added to include additional features.

## 5. Experiments

### 5.1. Evaluation Metrics

The evaluation metrics used to reflect the overall performance of our models are the mean Intersection-Over-Union (mIoU) and weighted F1-score. mIoU is specifically used for our segmentation model, as it is commonly used in literature to evaluate segmentation models. A weighted F1 score was also used for both the segmentation and classification models. This metric balances precision and recall in a harmonic mean and is well-suited for imbalanced datasets such as the xBD dataset [7].

For the purposes of the xView2 challenge, the final evaluation metric is a combined F1 score: 30% segmentation F1 score + 70% classification F1 score. This metric is calculated over the hold-out test dataset.

### 5.2. Results

#### 5.2.1 End-to-End Model

Examples of results obtained with our end-to-end model can be seen in Figure 7. It can be seen that our model does relatively well for the segmentation task. Our best model obtained a 0.795 F1 score on the test dataset for the seg-

mentation task, similar to what was obtained in the xView2 baseline model. Our error analysis showed that we missed some small buildings and those buildings partially hidden in the forest/water.

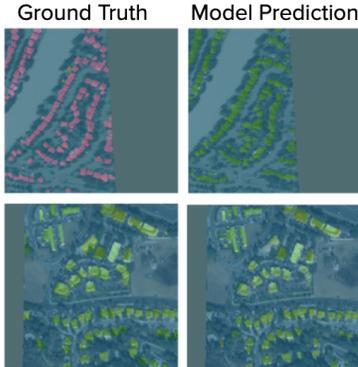


Figure 7: Ground truth data (left column) and model predictions (right column) for end-to-end model. Our model accurately predicts minor damage buildings (green) in some images while completely missing destroyed buildings (red) in others.

However, the end-to-end model was not able to accurately classify the damage level of each building and performed significantly worse than the baseline model. We realized that an end-to-end model is unable to accurately capture all the complexities of building damage levels, and therefore we proceeded to focus on the two-step model.

### 5.2.2 Building Segmentation

As described in Section 4.2.2, we experimented with two different segmentation architectures.

**Instance segmentation:** We trained a mask R-CNN model with two P100 GPUs using a batch size of 4 and the initial learning rate of 0.01. The original high-resolution image with dimensions  $1024 \times 1024 \times 3$  was directly used as the training images. The number of regions of interest (ROIs) per image was set to be 1000. COCO instance segmentation challenge pre-trained weight was used as the initial weight. We trained the model for 28,000 iterations.

**Semantic Segmentation:** We trained a ResNet-50 FPN model with two P4 GPUs. The batch size was 16 and the initial learning rate was 0.02. Again, the original high-resolution image with dimensions  $1024 \times 1024 \times 3$  was directly used as the training images. ImageNet pre-trained weight was used as the initial weight. This time, we trained the model for 50,000 iterations.

The prediction result of the instance segmentation and semantic segmentation model is shown in Figure 8a and

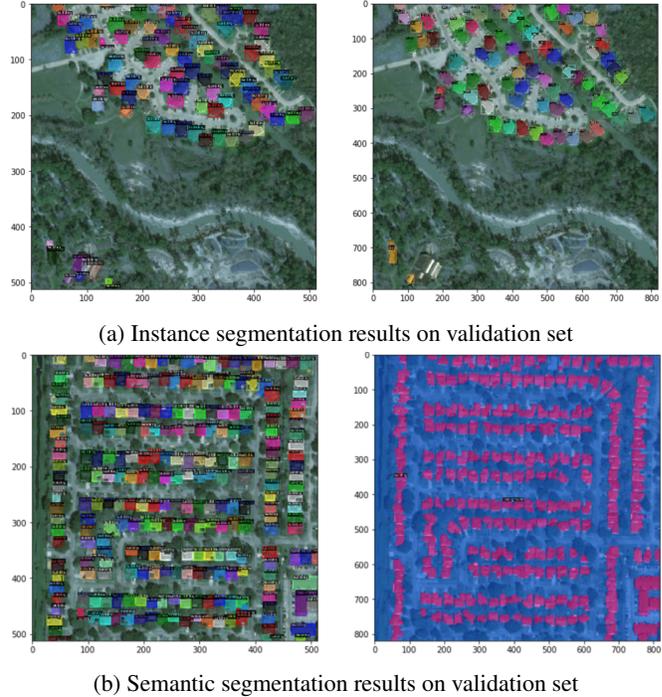


Figure 8: Building segmentation results on validation set. Left column is the ground truth data. The different colors represent each labeled individual building. a) right is the instance segmentation prediction. The different colors represent each individually detected instance of buildings. b) right is the semantic segmentation prediction. The prediction is a binary mask, with background (blue) and buildings (red) is obtained.

Figure 8b respectively. To conduct a consistent comparison of the two models, we computed the metric scores only based on their predicted masks. The predicted mask of Instance segmentation are a integration of all the individual masks of each building. We compared the mIoU score of our models on our validation set in Table 2. In addition, we compared our results to the xView2 challenge baseline. Though it seems that both models performed relatively well through visual inspection (Figure 8a and Figure 8b), the results show that our semantic segmentation model performs significantly better than the xView2 challenge baseline as well as our instance segmentation model.

The inferior performance of the instance segmentation model was surprising. We realized that one of the major issue of using Mask R-CNN for instance segmentation is its inability to detect all the buildings, particularly where there are more than 1000 buildings in a single image. We believe that using a larger and more complex model can improve the performance of the instance segmentation model. However, given the difficulties of optimizing the instance segmentation model as well as the fact that our semantic

Method	mIoU
xView2 Baseline	0.78
Instance Segmentation	0.70
Semantic Segmentation	0.85

Table 2: Building segmentation results on the validation set. mIoU is calculated based on the predicted binary mask of building footprint

segmentation model was able to achieve relatively accurate results, we chose the semantic segmentation model for our final model.

### 5.2.3 Building Classification

As discussed in Section 4.2.3, we implemented several different architectures for our damage classification models.

**Pre & Post Concatenation:** Inspired by the work done by Xu et al. [26], who studied a task highly analogous to the one we perform here, we implemented a twin-tower model which concatenates the output features of a ResNet-50 model from pre- and post-disaster building cropped images.

**Pre & Post Concatenation with Disaster Labels:** Acting off the knowledge that expected damage level distribution is different depending on the disaster, we added an additional feature as a third branch before concatenating and feeding into the fully connected layers. The additional feature is a label denoting the disaster to which a building belonged (e.g. “1” for Hurricane Matthew). Note that these labels are not available in the xView2 test set. To appropriately classify test set images with this model, we built an additional CNN model to identify image disaster name based on pre- and post-disaster imagery. This disaster name classification task was accomplished with this CNN to a high level of accuracy, achieving 95% accuracy on the validation set.

**Pre & Post Concatenation with SSIM:** We experimented with adding hand-engineered features to our models to boost performance. Here we used SSIM, a metric that is commonly used in image processing to measure the similarity between two images. Similar to the model with disaster label, this additional feature is added as a third branch of our architecture before concatenation.

All models were trained using cross-entropy loss for damage classification score. The models were trained using one NVIDIA Tesla T4 GPU with a batch size of 500 and initial learning rate of  $8 \times 10^{-4}$  and were trained until model convergence. These models were first tested for appropriate implementation by over-training on a smaller subset of the data.

Method	Train Acc.	Val.Acc.
Pre & Post Concat.	0.77	0.80
Pre & Post Concat. + Disaster	0.91	0.86
Pre & Post Concat. + SSIM	0.83	0.82

Table 3: Damage classification task results

The results of our experiments are summarized in Table 3. Counter-intuitively, Pre & Post Concatenation model scored higher on the validation set than on the training set. We attribute this to the model simply memorizing the damage distribution of the dataset, rather than learning anything intrinsic about the input image, and then benefiting from a validation set that had a slightly more favorable distribution of building damage scores. As anticipated, the model with the additional feature of disaster labels show a higher accuracy, which proves that providing the model with information about which disaster a given building came from dramatically improves performance. The model with SSIM as additional feature also performs better than the basic Pre & Post concatenation model, though still lower than the model with the disaster label. This shows that though knowing the disaster improves the performance, adding additional hand-engineered features can also help with performance.

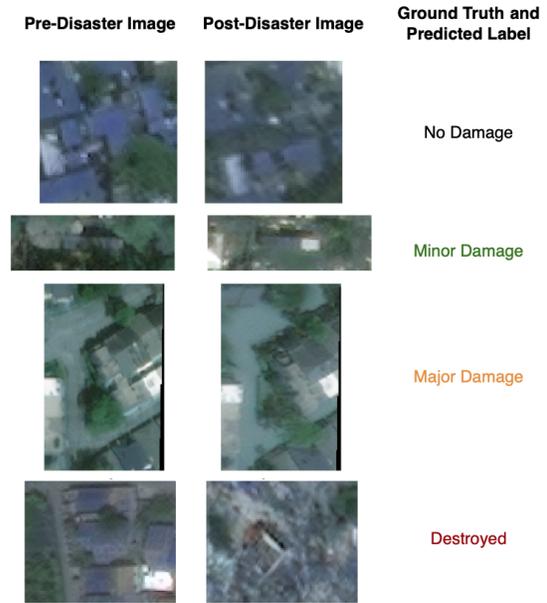


Figure 9: Validation results of correct damage classification predictions

When visualizing our damage classification predictions, we noticed that our model performs well if both the pre- and post-disaster images are identical in terms of color and

Method		Segmentation F1 Score	Classification F1 Score	Overall F1 Score
xView2 Baseline		0.80482	0.06091	0.28408
End-to-end Model		0.79584	/	0.23875
Two-Step Model	Semantic Segmentation + Pre & Post Concatenation	0.84330	0.51	0.61
	<b>Semantic Segmentation + Pre &amp; Post Concatenation + Disaster</b>	<b>0.84330</b>	<b>0.5873</b>	<b>0.6641</b>
	Semantic Segmentation + Pre & Post Concatenation + SSIM	0.84330	0.54679	0.63574

Table 4: Summary of results submitted to the xView2 challenge Leaderboard



Figure 10: Validation results of incorrect damage classification predictions

nadir angle. This can be seen in Figure 9. On the other hand, the model has issues dealing with very small buildings, and when the pairs of images were taken in very different times, causing differences in color, shadows, and nadir angles (Figure 10).

### 5.3. Test Results

For final analysis, we performed inference on the xView2 challenge test dataset and submitted the results to the xView2 leaderboard for scoring. These test dataset results were generated for the End-to-End Model and the two step models. For all two-step models, we used the output of the semantic segmentation model, which performed better in our experiments. Because the test set images do not include information about source disaster, we used the separate CNN model discussed earlier to classify test images by disaster name. These names were then included in the model which used disaster name as a feature. Results are summarized in Table 4.

The End-to-End model performed only slightly worse than the xView2 baseline on the segmentation part, a somewhat surprising result given the model’s significantly less complex design. The classification f1 score of the End-to-End model is not available, which we attribute to a submission error which could not be resolved by the time of

writing. However, due to its relatively poor performance in training, we expect that the End-to-End model would not perform well on the test set. Both two-step models performed better in localization and classification compared to baseline. Our best model, a two-step model with Pre & Post Concatenation and disaster label, scored significantly better than baseline (overall F1 score of 0.66 versus 0.28 for baseline).

## 6. Conclusion

In this paper, we presented a framework for building damage assessment after natural disasters. We experimented with several different architectures that focused on building segmentation and classification, in particular an end-to-end model that simultaneously performs those two tasks and a two-step model. Our approach presents one of the first attempts at creating a model that is able to generalize to all types of natural disasters.

We submitted our results as part of the xView2 Challenge, a competition to design better models for identifying buildings and their damage level after exposure to multiple kinds of natural disasters. Our best model couples a building identification semantic segmentation convolutional neural network (CNN) to a building damage classification neural network, with a combined F1 score of 0.66, surpassing the xView2 challenge baseline F1 score of 0.28.

Our results point to several important conclusions. We found that two-step models (segmentation and classification) that are trained individually works better than an end-to-end model that performs the two task simultaneously.

Our best performing model included an additional feature representing the disaster name. This result indicates two possibilities: (1) that similar levels of damage are visually different in different disasters, or (2) given the difficulty of the damage classification tasks, a best performing model will rely heavily on a learned prior probability distribution for damage, which will differ by disaster.

Though this is a huge step towards the automation of building damage assessments for various types of natural disasters, it is clear that there is still a lot of work to be done to increase the accuracy in the building damage classification. Though our results suggest that training a model for each individual disaster might boost performance in the

xView2 challenge, we realize that this will not be applicable for new disasters and real-world applications. Instead, our results suggest developing methods that incorporate disaster severity, region vulnerability, and other factors in the form of a prior estimate of damage.

## 7. Acknowledgments

We would like to acknowledge the Stanford CS325B (Data for Sustainable Development) teaching team for their continuous feedback and ideas. In particular, we would like to thank Professor Marshall Burke, Professor David Lobell and Stefano Ermon, as well as Robin Cheong, Matthew Tan and Burak Uz kent.

## References

- [1] L. Barrington, S. Ghosh, M. Greene, S. Har-Noy, J. Berger, S. Gill, A. Y.-M. Lin, and C. Huyck. Crowdsourcing earthquake damage assessment using remote sensing imagery. *Annals of Geophysics*, 54(6), 2012.
- [2] D. Brown, S. Platt, J. Bevington, K. Saito, B. Adams, T. Chenvidyakarn, R. Spence, R. Chuenpagdee, and A. Khan. Monitoring and evaluating post-disaster recovery using high-resolution satellite imagery—towards standardised indicators for post-disaster recovery. *Martin Centre: Cambridge, UK*, 2015.
- [3] L.-C. Chen, A. Hermans, G. Papandreou, F. Schroff, P. Wang, and H. Adam. Masklab: Instance segmentation by refining object detection with semantic and direction features. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4013–4022, 2018.
- [4] Q. Chen, L. Wang, Y. Wu, G. Wu, Z. Guo, and S. L. Waslander. Aerial imagery for roof segmentation: A large-scale dataset towards automatic mapping of buildings. *arXiv preprint arXiv:1807.09532*, 2018.
- [5] L. Dong and J. Shan. A comprehensive review of earthquake-induced building damage detection with remote sensing techniques. *ISPRS Journal of Photogrammetry and Remote Sensing*, 84:85–99, 2013.
- [6] S. Ghaffarian, N. Kerle, and T. Filatova. Remote sensing-based proxies for urban disaster risk management and resilience: A review. *Remote sensing*, 10(11):1760, 2018.
- [7] R. Gupta, B. Goodman, N. Patel, R. Hosfelt, S. Sajeev, E. Heim, J. Doshi, K. Lucas, H. Choset, and M. Gaston. Creating xbd: A dataset for assessing building damage from satellite imagery. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 10–17, 2019.
- [8] K. He, G. Gkioxari, P. Dollár, and R. Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017.
- [9] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [10] V. Iglovikov, S. S. Seferbekov, A. Buslaev, and A. Shvets. Ternaunetv2: Fully convolutional network for instance segmentation. In *CVPR Workshops*, pages 233–237, 2018.
- [11] Imantics. Image semantics documentation. <https://imantics.readthedocs.io/en/latest/#>, 2019.
- [12] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2117–2125, 2017.
- [13] S. Loos, K. Barns, G. Bhattacharjee, R. Soden, B. Herfort, M. Eckle, C. Giovando, B. Girardot, K. Saito, G. Deierlein, A. Kiremidjian, J. W. Baker, and D. Lallemand. The Development and Uses of Crowdsourced Building Damage Information based on Remote-Sensing. Technical Report 1, 2018.
- [14] S. P. Mohanty. Crowdai mapping challenge 2018 : Baseline with mask rcnn. <https://github.com/crowdai/crowdai-mapping-challenge-mask-rcnn>, 2018.
- [15] F. Nex, D. Duarte, A. Steenbeek, and N. Kerle. Towards real-time building damage mapping with low-cost UAV solutions. *Remote Sensing*, 11(3):1–14, 2019.
- [16] E. M. Rathje and B. J. Adams. The role of remote sensing in earthquake science and engineering: Opportunities and challenges. *Earthquake Spectra*, 24(2):471–492, 2008.
- [17] S. Ren, K. He, R. Girshick, and J. Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015.
- [18] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [19] T. G. Rudner, M. Rußwurm, J. Fil, R. Pelich, B. Bischke, V. Kopačková, and P. Biliński. Multi3net: segmenting flooded buildings via fusion of multiresolution, multisensor, and multitemporal satellite imagery. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 702–709, 2019.
- [20] H. Taubenböck, C. Geiß, M. Wieland, M. Pittore, K. Saito, E. So, and M. Eineder. The capabilities of earth observation to contribute along the risk cycle. In *Earthquake Hazard, Risk and Disasters*, pages 25–53. Elsevier, 2014.
- [21] D. M. Tralli, R. G. Blom, V. Zlotnicki, A. Donnellan, and D. L. Evans. Satellite remote sensing of earthquake, volcano, flood, landslide and coastal inundation hazards. *ISPRS Journal of Photogrammetry and Remote Sensing*, 59(4):185–198, 2005.
- [22] C. Van der Sande, S. De Jong, and A. De Roo. A segmentation and classification approach of ikonos-2 imagery for land cover mapping to assist flood risk and flood damage assessment. *International Journal of applied earth observation and geoinformation*, 4(3):217–229, 2003.
- [23] H. T. Verstappen. Aerospace technology and natural disaster reduction. *Advances in Space Research*, 15(11):3–15, 1995.
- [24] Z. Wang, A. C. Bovik, H. R. Sheikh, E. P. Simoncelli, et al. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.

- [25] Y. Wu, A. Kirillov, F. Massa, W.-Y. Lo, and R. Girshick. Detectron2. <https://github.com/facebookresearch/detectron2>, 2019.
- [26] J. Z. Xu, W. Lu, Z. Li, P. Khaitan, and V. Zaytseva. Building damage detection in satellite imagery using convolutional neural networks. *arXiv preprint arXiv:1910.06444*, 2019.