# DAMAGE EVALUATION FOLLOWING NATURAL DISASTERS USING DEEP LEARNING

Neha Gupta[1], Shikha Chadha[2], Rosey Chauhan[3], Pooja Singhal[4]

[1,4] ABES Engineering College, Ghaziabad, UP, India
[2,3]Sharda University, Greater noida

[1]123.neha.gupta@gmail.com, 0000-0003-4288-4521
[2]shikha1232@gmail.com, 0000-0002-5664-8522
[3]errose.it@gmail.com, 0000-0003-2557-1101
[4]pooja.singhal@abes.ac.in, 0000-0001-9768-8582

**Abstract**

Natural catastrophes including flooding, tornadoes, earthquakes, and wildfires have been occurring more frequently over the past few decades as a result of global warming and climate change. It is crucial that rescue workers are promptly notified of the location and extent of a building's destruction in order to maximise the effectiveness of their efforts. In this study, a potential deep learning-based method is put forth for identifying damaged buildings in high-resolution satellite photos. It solves the issue of limited training data common in many remote sensing applications by using generic data augmentation. It is suggested that a pretrained model be used in conjunction with transfer learning as a fine-tuning method for the relevant task. The trials with images of Port-au-Prince, Haiti showed that the suggested strategy works well with sparse training data. With enriched training data, the Convolutional Neural Network (CNN) model can detect damaged buildings with an accuracy of 83%, compared to only 53% with the original training data.

Keywords: Natural catastrophes, disaster response, high-resolution satellite photos, Convolutional Neural Network, data augmentation

## 1 Introduction

Since the First World War, when cameras were mounted on monoplanes, remote sensing for disaster management has advanced [1]. Natural catastrophes including flooding, tornadoes, earthquakes, and wildfires have been occurring more frequently over the past few decades as a result of global warming and climate change. The United States Geological Survey reports that between 2000 and 2012, there were 807 earthquakes in the United States that were greater than magnitude 5.0, and there were 23,608 earthquakes globally that resulted in an estimated 789,677 fatalities. It is, therefore, very important to give quick response to workers which provides accurate and timely information to help them to prevent themselves effectively. used an estimated $125 billion in damages in Texas [2]. Remote sensing data, due to its widespread coverage, cost-effectiveness, and regular updates, emerged as one of the most economical and precise data sources for assessing the approximately 250,000 residential and 30,000 commercial buildings affected by the Haiti earthquake [3].

Many studies have used satellite and aircraft imagery to assure damaged structures. A research project explored the use of both manned and un-manned multi-resolution satellite and air-borne photography. By a successful implementation method based on con-volutional neural networks, the research successfully attained an accuracy ranging from 89.8% to 94.4% in the identifying and classifying damage in building [4]. In another investigation that focused on the aftermath of the Bam Earth-quake in "Iran", "Janalipour" and "Mohammadzadeh" conducted a 4-stage process utilizing Quick-Bird satellite photos to pinpoint damaged structures. Initially it updated a vector-map with pre-event photos, geo-referencing a post-event image using the revised vector-map, and classifying and segmenting the post-event image based on pixels. After that, the segments were labeled, and various geometric propertieswere determined. Finally a decision-making system based on the ANFIS (adaptive network-based fuzzy inference system) was created, a hybrid learning technique combining fuzzy logic with neural-networks [5][6].

Larger areas of land can be gathered by satellites in a single path. The position of the satellite is a significant barrier for the collection of satellite data. If the satellite doesn't pass directly over the affected area, it will need to be directed there. Large acquisition angles from a satellite location may result in image distortion, poor georeferencing, and a user's inability to co-register the imagery with other significant data sets [7].

Response time, analysis depth, and mapping accuracy must all be compromised in a scenario with a short time period [8]. While high resolution data may be too large to analyse for a vast region, sub-meter imagery is frequently employed for assessment [9-10]. However, it is normal practise to immediately create a preliminary damage assessment map using the existing data, and subsequently the first estimates are improved using new information [10]. Social media data can help with situational awareness, rescue and relief efforts, and timely damage information [11]. The diverse character of the data, which exceeds the capabilities of the human analyst when dealing with social media data, is one of the biggest obstacles [11-12]. The first analysis of remote sensing data may be degraded due to the rapid requirement for damage information during a natural catastrophe, and more time will be required to undertake a more thorough analysis [13].

## 2 Literature Review

There are still two major obstacles to picture classification for damage assessment: the accessibility of data sources and the prompt collection of precise and useful data. This section reviews recent work that addresses these issues, from basic pixel-based solutions to cutting-edge deep learning techniques.

Pixel-based methods rely on the surface features' spectral properties in multiple or hyperspectral pictures. The simplest unit of analysis for images is the pixel. A pixel-based technique primarily uses pixels without taking into account their geographical environment. In order to analyse a single pixel in a picture made up of potentially thousands or millions of pixels, statistical operators are used [14]. Unsupervised

classification and supervised classification are the two general categories into which pixel-based classification may be divided [15]. Unsupervised categorization divides an image's pixels into various classes according to their natural grouping. Without the aid of any training data or prior knowledge of the subject area, this procedure is carried out [16]. The Iterative Self-Organizing Data Analysis (ISODATA) technique and the k-means (and its variants) algorithm are two of the most popular algorithms [16-17]. The pixel spectral pattern 10 vector is to be classified into one of several classes using ISODATA. Similar algorithms like the k-means algorithm use additional heuristics to decide whether to split or combine groups. The method will separate a cluster if its variance is larger than a predetermined threshold. If not, the clusters will be merged [18-19]. A person must be in charge of supervised classification.

Object-based image categorization operates on homogeneous and spatially contiguous collections of pixels, or objects, rather than on individual pixels [20]. Pixel grouping is accomplished using an image segmentation procedure that, according to Haralick and Shapiro [21], may be divided into three main categories: spatial clustering, thresholding, and region growth. The ideal technique to assess the effectiveness of segmentation is for the human analyst to interpret the results (Estoque et al., 2015; Pal & Pal, 1993) because the accuracy of the image segmentation directly affects the image categorization [22]. A majority decision determines how an object is categorised based on the labels of its k closest neighbours. The procedure performs a Euclidean distance computation when the majority votes are tied and the object might belong to either of the two classes. When used with multimodal classes, KNN performs well. However, it makes categorization errors because it compares all the features equally [23]. Several studies have shown that object-based classifiers outperform pixel-based ones in terms of performance. In order to compare the MLC and KNN algorithms, Platt and Rapoza (Platt & Rapoza, 2008) used models that both included and excluded expert knowledge. They demonstrated that the best pixel-based strategy only managed an accuracy of 64%, while the object-based model employing the KNN algorithm plus expert knowledge produced the best accuracy results of 78%. [24] used QuickBird imagery to categorise urban land cover and compared MLC with KNN. On their original image, they achieved 90.4% accuracy with their KNN classifier and 63.3% accuracy with their pixel-based MLC classifier. They experimented with applying the same categorization techniques to a different image with various environmental factors. The pixel-based classifier in this instance had a significantly higher classification accuracy of 87.8%.

The Support Vector Machine (SVM) algorithm can only categorise data into two unique classes in its initial form, which makes it a two-class classification algorithm. It creates a high dimensional feature space from the input vectors. A hyperplane with qualities that guarantee the network's excellent generalisation ability is generated in this feature space [24]. It shows how the ideal separating hyperplane separates the data set 15 into discrete numbers of classes while minimising misclassification acquired during the training phase. [25] The use of a linear SVM assumes that the feature data can be separated linearly. However, in fact, data points from several classes frequently cross paths, rendering the fundamental linear decision limits inadequate. Consequently, kernel functions have been created [25].

With accuracy rates of 96.8% and 96.2%, respectively, the SVM algorithm performed somewhat better than the Bayesian classifier.

(ANN) is based on how the brain's neural networks are organised [25]. A typical neural network is made up of numerous straightforward, interconnected neurons, each of which generates a series of activations. Sensors that perceive the environment through weighted connections from previously engaged neurons trigger these neurons. The feedforward neural network was the first and most basic type of network.

In many applications involving object recognition, CNNs are quickly becoming a promising technology. Local receptive fields, shared weights, and spatial or temporal subsampling are three architectural concepts that CNNs combine. CNNs have several interconnected, multi-layered channels that are very capable of learning new features and classifiers. Additionally, they have the ability to simultaneously classify and change settings. Furthermore, this kind of ANN has the capacity to automatically embed both spectral and spatial data into the classification.

# 3 CNN'S INSIDE FOR DAMAGE DETECTION

In order to perform better than existing methods for damage identification using high-resolution satellite photos, CNN models that have been fine-tuned are adapted to computer vision datasets. Transfer learning, also known as the refining of models, is the application of learnt skills to a new situation. As a result, the goal of this work is to improve the ResNet152 model, one of most popular backbone framework in computer vision-related designs, so that it can more accurately distinguish between damaged and undamaged structures on satellite images. The ResNet34 and ResNet50 models, as well as additional ResNet models with different depths, were compared to this model. ResNet152 was selected because, when evaluated on the same sample set, it provided the best testing accuracy. Additionally, it is not by design that deeper ResNet models have lower training accuracy than their steeper counterpart. The four main steps of the suggested technique are depicted in Figure 1. After the landsat images has been pre-processed, the procedure begins by preparing the data by obtaining training data for the job of relevance (i.e., building footprints and manually labelling them as "Damaged" or "No Damage") (atmospheric correction and orthorectification).

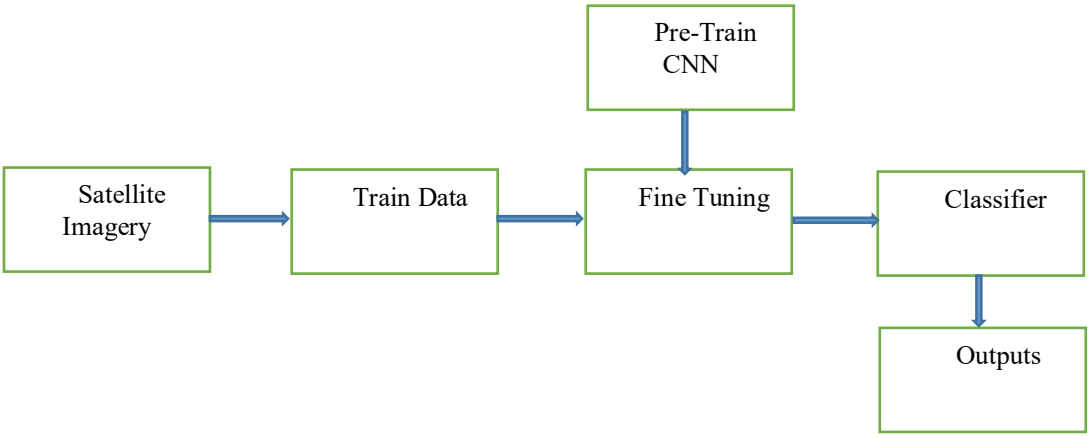


Fig.1. Methodology

Then, training, validation, and testing data are sorted into their respective categories. The next step is model fine-tuning, which adjusts the classification layers of the current architecture and trains the neural network's filter weights. Finally, metrics will be used to judge the output accuracy of the improved classifier.

# 4 IMPLEMENTATION

The effectiveness of the suggested strategy is illustrated in this section. The Jupyter Notebook environment has been used to implement the experiments. After the magnitude 7.0 Haitian earthquake of January 12, 2010, Maxar Technologies provided WorldView-2 (WV2) satellite imagery of Port-au-Prince, Haiti. It is freely accessible through the Maxar open data initiative and has a spatial resolution of 1.84 metres (Satellite Imaging Corporation, n.d.). The red, green, and blue bandwidths of four photos that have been plaid together to form the data each include three bands. The World Bank, the European Commission (EC), the Operational Satellite Applications Programme (UNOSAT), and the United Nations Institute for Training and Research (UNITAR) Joint Research Centre (JRC) collaborated to produce the Disaster Response Needs Assessment and Recovery Framework and ground truth data was gathered from these products.

Based on the ground truth information, building footprints were manually categorised as "Damage" or "No Damage" after being recovered from the WV2 satellite picture. They were divided into 93 testing samples, 322 training samples, and 46 validation samples at random. Figure 2 displays the locations of these sample datasets for each category. These locations were picked because 1) comprehensive building damage information is available to the public and 2) the buildings there are typical of the Port-au-Prince commune in terms of size, height, and structure.

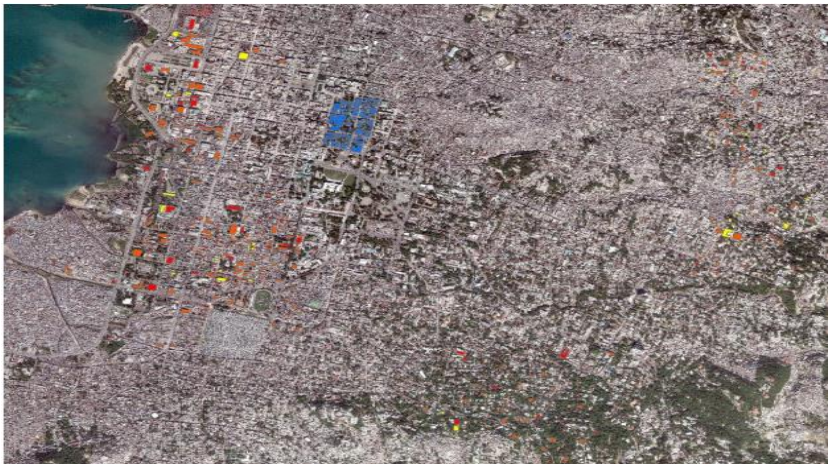

Fig.2. (Orange - training, yellow - validation, red - testing A, blue - testing B) Data locations and classifications.

The collected 322 training examples were used to first perfect the ResNet152 model. With the enhanced training samples (training+) produced by the rotation transformation, a second experiment was carried out. The zero-padded training images that were taken from the WV2 satellite imagery were used in both tests. Each of these 322 buildings was rotated 72 degrees constantly until 360 degrees, as illustrated in Figure 3, as part of the data augmentation (DA) process, which increased the training data 41 by a factor of five.

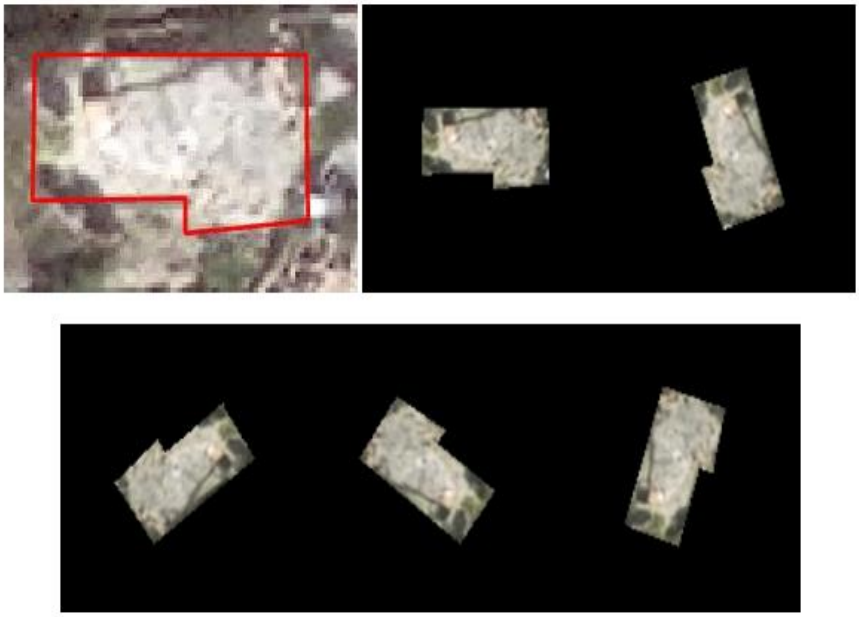

Fig.3. The top-left image is the WV2 satellite's original image, and the boundary was used to extract the image chip. The extracted and padded image may be seen in the upper middle image. The further photos display the DA outcomes through a series of 72-degree rotations.

The quantity of labelled sample data used for each stage of the process is summarised in Table 1. The ResNet152 classifier's weights are fitted using training data, which is also used to learn how the process works. The ResNet152 classifier's parameters can be adjusted with the help of validation data, which can also be used to check whether the training data is adequate, for instance. The performance of the completely trained classifier, the final model, is next evaluated using the testing data.

Table 1 Summary of sample for fine tuning

| Sample Types | Total | Damage | No Damage |
|---|---|---|---|
| Training (Original) | 322 | 182 | 140 |
| Training+ (after DA) | 1610 | 910 | 700 |
| Validation | 46 | 26 | 20 |
| Testing | 93 | 53 | 40 |

Preventing a model from being either overfitted or underfitted is a crucial aspect of machine learning. The aim of training is to achieve 0% validation loss. However, overfitting is suggested if the validation loss is more than the training loss, and underfitting is indicated when the reverse is true. Robust fitting is achieved if the validation loss and the training loss are equal (Brownlee, 2019b). The use of an

approach known as "early stopping" was made. The validation loss stops improving for five consecutive epochs at 43 using the early stopping technique, at which point the model is finished being trained. The model was trained for a total of 20 epochs, as shown in Table 2, but due to early stopping, it ceased improving validation loss on the 15th epoch. The model was trained for 6 epochs only while using the original training data in testing. Therefore, the 20-epoch DA comparison is made to the six founding training data result in order for the two studies to have a consistent framework of early ending. The outcomes of the ResNet152 model epochs comparing the initial training data to the training+ data are displayed in Figure 4.

Table 2 Training return for the DA model demonstrating training and validation loss

| Epoch | train_loss | valid_loss | Accuracy | time |
|---|---|---|---|---|
| 1 | 1.042376 | 0.903653 | 0.515528 | 0:36 |
| 2 | 0.936403 | 0.731321 | 0.720497 | 0:23 |
| 3 | 0.874362 | 0.575533 | 0.770186 | 0:21 |
| 4 | 0.783001 | 0.508756 | 0.776398 | 0:21 |
| 5 | 0.767506 | 0.537253 | 0.782609 | 0:21 |
| 6 | 0.751107 | 0.613045 | 0.782609 | 0:21 |
| 7 | 0.729267 | 0.525315 | 0.751553 | 0:21 |
| 8 | 0.694654 | 0.46184 | 0.807453 | 0:21 |
| 9 | 0.660684 | 0.433513 | 0.807453 | 0:21 |
| 10 | 0.653519 | 0.441762 | 0.801242 | 0:21 |
| 11 | 0.605963 | 0.488908 | 0.813665 | 0:21 |
| 12 | 0.583887 | 0.500495 | 0.782609 | 0:21 |
| 13 | 0.562885 | 0.400379 | 0.819876 | 0:21 |
| 14 | 0.534634 | 0.310225 | 0.869565 | 0:21 |
| 15 | 0.523058 | 0.288835 | 0.850932 | 0:21 |
| 16 | 0.521369 | 0.366022 | 0.826087 | 0:21 |
| 17 | 0.498879 | 0.51325 | 0.807453 | 0:21 |
| 18 | 0.474929 | 0.376632 | 0.832298 | 0:21 |
| 19 | 0.440805 | 0.320855 | 0.869565 | 0:21 |
| 20 | 0.406367 | 0.316084 | 0.863354 | 0:21 |

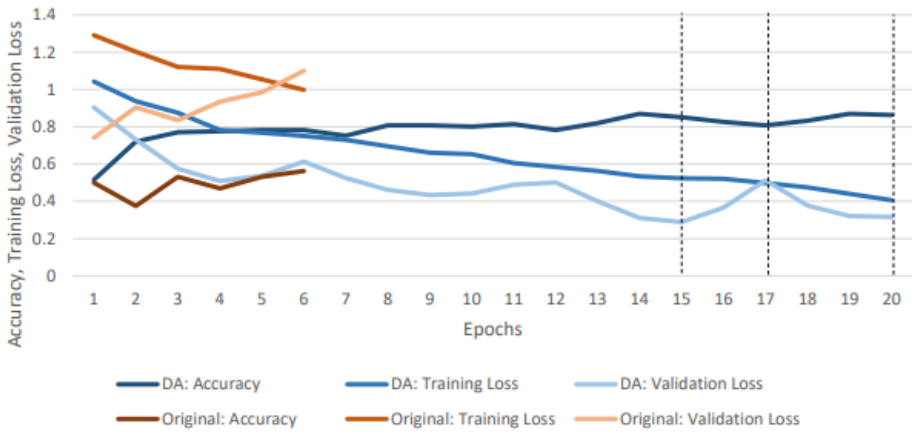


Fig. 4. DA vs. Original Training Epoch Results Comparison

According to the graph, the validation loss for the DA model is lowest at the 15th epoch, which is the desired outcome; nevertheless, the gap between validation loss and training loss at that time is one of the largest in the model. Validation loss has maximized at this time even though it is still larger than training loss by.01 at the 17th epoch. Due to early halting, the model terminates at the 20th epoch. It is the epoch with the second-lowest difference and a region with the third-lowest validation loss. There is a difference of.09 between validation and training at this time.

The ideal model is one that uses the initial random result of 20 epochs with the training+ data. The 20-epoch model can be trained in a much less amount of time. 20 epochs were trained in 7.27 minutes on training+ photos as opposed to 34 and 50 epochs in 12.48 and 18.37 minutes, respectively. Second, even with the modest validation loss taken into account, the training and validation losses were the closest at 20 iterations. Thirdly, the early halting feature will improve generalisation for yet-to-occur situations by preventing overfitting as well as overfitting (Brownlee, 2019a). The 34 and 50-epoch approaches only slightly outperform the 20, 34, and 50-epoch model in terms of accuracy, but at the expense of a large amount of time.

# 4 RESULTS

Figure 5 compares the outcomes of the two studies as well as the accuracy of training, validation, and testing. The training, validation, and testing accuracy are 92%, 85%, and 83%, respectively, when DA is applied to the training images. The training, validation, and testing accuracies using the initial training data are 50%, 43%, and 53%, respectively. The DA accuracies are, on average, 38% more accurate than the model using the original training samples.

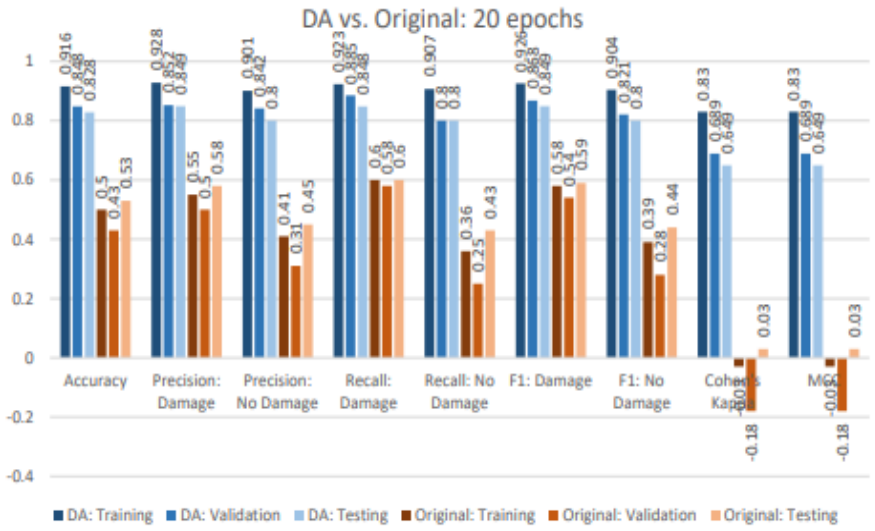


Fig.5. Analyzing the statistics from the original model and the DA.

When analysing the damage and no damage accuracies from the conventional training data and training+ data, there is also an aggregate of a 33% and a 46% increased accuracy result, respectively. The recall evaluation for damage and no damage both demonstrate improvement for the DA result, rising by an average of 29% and 49%, respectively. The test for classifying damaged buildings is more accurate than the test for classifying non-damaged structures when the F1 score, which is a measure of test accuracy, is taken into account for all tests. The F1 score also reveals that the two models' average difference ranges from 31% for damage to 47% for no damage. The model with DA has, on average,.78 more inter-relatability than the original model, according to Cohen's Kappa, a measurement of the inter-relatability here between prediction classes and regression coefficients classes. Last but not least, MCC values can range from -1 to 1, with -1 denoting the greatest possible discrepancy between prediction and classification algorithm, 0 denoting randomness in the model, and 1 denoting the model's flawless prediction. The model with DA performs significantly better than the model with original samples, as shown by the MCC score average of.72 for the model with DA and -.06 for the model with original training data.

The first testing data set, designated as "Testing A," is produced at random. The other is the brand-new set, known as "Testing B," which is used to evaluate the model's capacity for generalisation. According to Figure 6, Testing B consists of 189 samples drawn from seven city blocks. These samples are neither dispersed across a wide area nor sparsely sampled, in contrast to the 93 samples gathered for Testing A. Furthermore, the Testing B examples have much smaller building footprints. Table 3 demonstrates that, on average, the samples gathered for Testing 49 B are much smaller than the samples gathered for training in Testing A.

Fig.6. A test of sample B. Ones with damage are shown in red, whereas unharmed buildings are shown in blue.

Testing B's accuracy was assessed to be 63%, while Testing A's accuracy was 83%. Precision, or the percentage of accurate damage or no damage forecasts, came out to be 50% and 80%, respectively. In comparison to Testing A's 85%, the damage precision is only 50%. But interestingly, the classification of no harm for Testing A and Testing B was discovered to be the same, 80%.

According to the recall metrics for Testing A and Testing B, 85% and 76%, respectively, of actual positives (also known as ground truth positives or "Damage"), or correct identifications, are found. Comparing Testing A and Testing B, the recall was calculated to be 80% and 55%, respectively, for the no damage categorization. Damage from both testing areas had a better recall than no damage. For the classification of damaged buildings in Testing A and Testing B, the F1 score, which is a gauge of the model's robustness and precision, came out to be 85% and 60%, respectively.

Table 3 Building footprint sizes are compared

| Area (sqm) | Training | Testing A | Testing B |
|---|---|---|---|
| Average | 291.7 | 360.1 | 197.0 |
| Median | 197.0 | 200.0 | 150.3 |
| Min | 30.1 | 18.9 | 33.1 |
| Max | 2058.0 | 1912.2 | 1092.7 |

Additionally, it displays an F1 score of 80% and 65% for no damage. These measurements demonstrate that the trained ResNet152 model is not robust enough to generalise on new examples. Finally, testing A performs much better than testing b according to Cohen's Kappa and the Matthews Correlation Coefficient. The degree to which the forecast and the ground truth labels agree is shown by Cohen's Kappa. Cohen's Kappa for Testing A is.65, but Kappa for Testing B is.28. The MCC is a measurement of the correlation between the predicted binary classification and the

ground truth. Testing B scored.30, whereas Testing A obtained a rating of.65. The association between the actual truth and predicted classifiers is more erratic according to Testing B's MCC value of.30 than it is according to Testing A's MCC value.

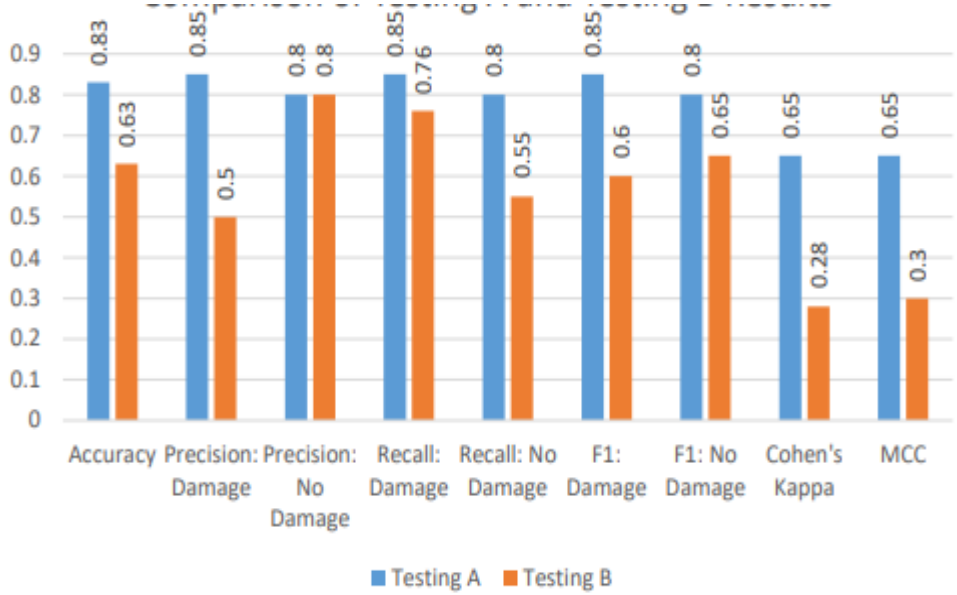


Fig.7. Comparison of Testing A and Testing B Results

There are a few reasons why the trained ResNet152 model's generalisation performance is relatively poor. The size of Testing B's building footprint in comparison to the footprints of the training buildings is one of the considerations. According to Table 5 above, the average building footprint area for Testing B is 197.0 sqm, whereas it is 291.7 sqm for training and 360.1 sqm for Testing A. The tiny area of the Testing B footprints could significantly affect the generalisation findings since the smaller textures provided would leave the model with less features to classify. The broken construction texture, which is particularly evident in the precision scores, is a contributing factor to the generally subpar generalisation performance. The no damage building footprints were taken from structures that had sustained no damage; hence, whether the buildings are from Testing A or Testing B, their rooftops are largely comparable throughout the area in the satellite image. Consequently, the accuracy rate for the no damage categorization is nearly the same, at 80%.

# 4 CONCLUSIONS AND FUTURE WORK

The training set can be artificially increased while keeping labels intact using the suggested generic data augmentation based on geometric transformation. The DA technique further improves CNN performance and prevents overfitting, as shown by the studies. Additionally, a transfer learning fine-tuning technique is suggested, which

modifies the high-level layers of the pre-trained ResNet152 model while leaving the low-level layers untouched. Results show that pre-trained architectures are extremely effective and produce viable classification (85% and 80% in Testing A precision and recall for classifying damaged and no-damaged buildings, respectively), even though they are not pre-trained on the application of this work (i.e., damaged or no-damaged buildings). The success of our strategy is further demonstrated by the fact that the fine-tuning process converged after 20 epochs of training rounds where the validation loss is near to the training loss.

Future research would concentrate on examining model generalisation strategies, labelling training data, and automating approaches for extracting building footprints. There is no question that a larger training dataset would enhance model performance even further. Manual extraction, however, is time-consuming, expensive, and difficult. Building detection ML models, such as Ren et alFast-RCNN, .'s have been created to recognise buildings from high resolution satellite photos. Automating their labelling continues to be a challenge. The task of model generalisation is challenging. Research should continue on methods to strengthen the model's resistance to data errors. For instance, further data augmentation methods must to be researched and used in the process of fine-tuning.

# REFERENCES


1. Cable, S. (2015). Aerial photography and the First World War. The National Archives. https://blog.nationalarchives.gov.uk/aerial-photography-first-world-war/
2. Kolbe, A. R., Hutson, R. A., Shannon, H., Trzcinski, E., Miles, B., Levitz, N., Puccio, M., James, L., Noel, J. R., & Muggah, R. (2010). Mortality, crime and access to basic needs before and after the Haiti earthquake: a random survey of Port-au-Prince households. Medicine, Conflict and Survival, 26(4), 281–297. https://doi.org/10.1080/13623699.2010.535279
3. N. Gupta, R. Chauhan, and S. Chadha, "Unmanned Aerial Vehicle (UAV) for parcel delivery," Int. J. Eng. Res. Technol., vol. 13, no. 10, pp. 2824–2830, 2020, doi: 10.37624/IJERT/13.10.2020.2824-2830.
4. Duarte, D., Nex, F., Kerle, N., & Vosselman, G. (2018). SATELLITE IMAGE CLASSIFICATION OF BUILDING DAMAGES USING AIRBORNE AND SATELLITE IMAGE SAMPLES IN A DEEP LEARNING APPROACH. ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences, 4(2), 89–96. https://doi.org/10.5194/isprs-annals-IV-2-89-2018
5. Janalipour, M., & Mohammadzadeh, A. (2016). Building Damage Detection Using Object-Based Image Analysis and ANFIS From High-Resolution Image (Case Study: BAM Earthquake, Iran). IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing, 9(5), 1937–1945. https://doi.org/10.1109/JSTARS.2015.2458582
6. Rathje, E. M., & Adams, B. J. (2008). The Role of Remote Sensing in Earthquake Science and Eengineering: Opportunities and Challenges. Earthquake Spectra, 24(2), 471–492. https://doi.org/10.1193/1.2923922
7. Voigt, S., Schneiderhan, T., Twele, A., Gähler, M., Stein, E., & Mehl, H. (2011). Rapid Damage Assessment and Situation Mapping: Learning from the 2010 Haiti Earthquake. Photogrammetric Engineering & Remote Sensing, 77(9), 923–931. https://doi.org/10.14358/PERS.77.9.923
8. A. Suri, R. V. S. Bhadauria, and L. K. Bansal, "Survey on Methods of Face Mask Detection System," 2022 International Mobile and Embedded Technology Conference (MECON), Mar. 2022, doi: 10.1109/mecon53876.2022.9751815.
9. 
10. S. Chadha, M. R. Chauhan, and M. N. Gupta, "Flood Prediction And Rainfall Analysis Using Light Gradient Boosted Machine," vol. 20, no. 6, pp. 1–6, 2022.

11. N. Gupta and K. K. Rana, “Disaster Prediction And Post Disaster Management Using Machine Learning And Bluetooth,” vol. 18, no. 5, pp. 274–292, 2021.
12. Li, M., Zang, S., Zhang, B., Li, S., & Wu, C. (2014). A Review of Remote Sensing Image Classification Techniques: the Role of Spatio-contextual Information. European Journal of Remote Sensing, 47(1), 389–411. https://doi.org/10.5721/EuJRS20144723
13. Ouyang, Z. T., Zhang, M. Q., Xie, X., Shen, Q., Guo, H. Q., & Zhao, B. (2011). A comparison of pixel-based and object-oriented approaches to VHR imagery for mapping saltmarsh plants. Ecological Informatics, 6(2), 136–146. https://doi.org/10.1016/j.ecoinf.2011.01.002
14. M. Rahul, N. Tiwari, R. Shukla, D. Tyagi, and V. Yadav, “A New Hybrid Approach for Efficient Emotion Recognition using Deep Learning,” International Journal of Electrical & Electronics Research, vol. 10, no. 1, pp. 18–22, Mar. 2022, doi: 10.37391/ijeer.100103.
15. Qian, J., Zhou, Q., & Hou, Q. (2007). COMPARISON OF PIXEL-BASED AND OBJECTORIENTED CLASSIFICATION METHODS FOR EXTRACTING BUILT-UP AREAS IN ARIDZONE. ISPRS Workshop on Updating Geo-Spatial Databases with Imagery & The 5th ISPRS Workshop on DMGISs, 163–171.
16. https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.221.8137&rep=rep1&type=pdf
17. Haralick, R. M., & Shapiro, L. G. (1985). Image Segmentation Techniques. Computer Vision, Graphics, and Image Processing, 29(1), 100–132. https://www.sciencedirect.com/science/article/abs/pii/S0734189X85901537
18. Estoque, R. C., Murayama, Y., & Akiyama, C. M. (2015). Pixel-based and object-based 57 classifications using high- and medium-spatial-resolution imageries in the urban and suburban landscapes. Geocarto International, 30(10), 1113–1129. https://doi.org/10.1080/10106049.2015.1027291
19. P. K. Chaubey et al., “Sentiment Analysis of Image with Text Caption using Deep Learning Techniques,” Computational Intelligence and Neuroscience, vol. 2022, pp. 1–11, Jun. 2022, doi: 10.1155/2022/3612433.
20. Platt, R. V., & Rapoza, L. (2008). An Evaluation of an Object-Oriented Paradigm for Land Use/Land Cover Classification. The Professional Geographer, 60(1), 87–100. https://doi.org/10.1080/00330120701724152
21. Myint, S. W., Gober, P., Brazel, A., Grossman-Clarke, S., & Weng, Q. (2011). Per-pixel vs. objectbased classification of urban land cover extraction using high spatial resolution imagery. Remote Sensing of Environment, 115(5), 1145–1161. https://doi.org/10.1016/j.rse.2010.12.017
22. Cortes, C., & Vapnik, V. (1995). Support-vector networks. Machine Learning, 20(3), 273–297. https://doi.org/10.1007/BF00994018
23. Kavzoglu, T., & Colkesen, I. (2009). A kernel functions analysis for support vector machines for land cover classification. International Journal of Applied Earth Observation and Geoinformation, 11(5), 352–359. https://doi.org/10.1016/j.jag.2009.06.002
24. Maulik, U., & Chakraborty, D. (2017, March). Remote Sensing Image Classification: A survey of support-vector-machine-based advanced techniques. IEEE Geoscience and Remote Sensing Magazine, 5(1), 33–52. https://doi.org/10.1109/MGRS.2016.2641240
25. P. Aggarwal, P. Jain, J. Mehta, R. Garg, K. Makar, and P. Chaudhary, “Machine learning, data mining, and big data analytics for 5G-Enabled IoT,” in Springer eBooks, 2021, pp. 351–375. doi: 10.1007/978-3-030-67490-8_14.
26. P. Agarwal, N. Garg, and P. Singh, “Predicting Poverty Index using Deep Learning on Remote Sensing and Household Data,” International Journal of Recent Technology and Engineering, vol. 8, no. 3, pp. 164–168, Sep. 2019, doi: 10.35940/ijrte.c3918.098319.
27. S. Shaikh, D. Waghole, P. Kumbhar, V. Kotkar, and P. Awaghade, “Patient monitoring system using IoT,” International Journal of Recent Scientific Research, Dec. 2017, doi: 10.1109/bid.2017.8336594.