Automated Flood Impact Detection from Remote Sensing Data using Transfer Learning: A New Zealand Case Study

Nethma Pathirana
Joint Centre for Disaster Research
Massey University
Wellington, New Zealand
ORCID: 0009-0007-1630-3144

Malintha Ranasinghe
Joint Centre for Disaster Research
Massey University
Wellington, New Zealand
ORCID: 0000-0002-1141-3367

Raj Prasanna
Joint Centre for Disaster Research
Massey University
Wellington, New Zealand
ORCID: 0000-0002-5608-6558

***Abstract*— Proper flood damage assessment is required for rapid emergency response and recovery planning. High-resolution drone imagery provides greater visual detail compared to satellite imagery and can support more accurate damage assessment after flood events. This work-in-progress study investigates the use of transfer learning with a pre-trained ResNet50 convolutional neural network to classify post-flood drone images. A manually annotated dataset of drone images representing damaged and non-damaged areas was used to fine-tune the pre-trained model and evaluate its performance using several metrics. The training process employed hyperparameter tuning to select the best model which achieved an accuracy of approximately 87%. These preliminary results demonstrate the potential of transfer learning for flood damage classification using a limited drone image dataset. Future work will focus on comparing the performance of additional deep learning models, incorporating satellite imagery and extending the approach from image-level classification to object-level damage detection.**

Keywords—Flood damage detection, Drone imagery, Transfer learning, ResNet50, Remote sensing, Disaster management

I. Introduction

Floods are one of the most high impact natural disasters. It causes heavy damages to houses, roads, bridges and buildings [1], [2]. New Zealand experiences frequent and often severe floods, driven by intense rainfall events, topography, and climate variability, with recent years showing an increase in both frequency and impacts. For instance, Cyclone Gabrielle has caused widespread damage to infrastructure, housing, and communities [3]. Earth Sciences New Zealand reveals that more than 750,000 New Zealanders live in locations exposed to flooding from one-in-100-year rainfall flooding events, highlighting the urgent need for efficient and scalable post-disaster damage assessment methods [4].

Damage assessment must be done directly after the flood for effective emergency response, recovery planning and resource allocation [5]. Traditional field assessments are usually time consuming and costly. Furthermore, it is dangerous to perform traditional field assessments in affected areas. However, in contrast, using aerial platforms like Unmanned Aerial Vehicles (UAVs) and satellites for remote sensing helps with faster and safer damage assessment [6], [7].

Although satellite imagery provides large-scale coverage, its spatial resolution and revisit frequency often limit its effectiveness in capturing fine-grained damage details [7]. Drones provide high resolution images that capture clear visual details that are helpful to identify damages from the flood. These include damaged buildings, houses, bridges, blocked roads and inundated areas [8].

While drone images provide detailed visual information for post-flood damage assessment, the manual interpretation of large volumes of high-resolution images remains a significant bottleneck [8]. Following major flood events, hundreds or thousands of drone images may be collected within a short time period, making manual inspection time-consuming, subjective and impractical for quick decision-making.

Moreover, analysis led by humans is difficult to scale during large-scale disasters where timely and consistent damage information is critical for emergency response and recovery planning. Therefore, there is a growing need for automated damage detection methods that can efficiently analyze drone imagery and provide rapid assessments of flood-induced impacts [7].

Several studies have explored automated post-event damage detection using remote sensing data. Early approaches relied on traditional image processing techniques or rule-based change detection methods that compare pre and post-disaster satellite imagery to identify damaged areas [9]. While these methods demonstrated the possibility of automated damage mapping, their performance often heavily depended on manual feature engineering. Furthermore, they were sensitive to variations in illumination, sensor characteristics and environmental conditions.

More recently, deep learning-based approaches have been widely used due to their ability to automatically extract complex spatial features from large-scale remote sensing imagery addressing challenges encountered in traditional methods [7], [10]. These studies primarily focus on satellite imagery and use large labeled datasets for training deep neural networks. However, the development of an automated post-event damage detection system using drone images is challenging. One major challenge is the the lack of large, publicly available labeled UAV image datasets for training deep learning models.

Transfer learning has emerged as a promising approach to address the challenge of limited labeled data in many fields including disaster damage assessment [11], [12], [13]. Instead of training deep neural networks from scratch, transfer learning in deep learning allows model which are previously trained using large datasets to be fine-tuned for specific applications with smaller datasets [13], [14].

While transfer learning has been explored by multiple studies for disaster damage detection using satellite imagery, comparatively fewer studies have explored the use of transfer learning for flood damage classification using high-resolution UAV images [7]. Furthermore, although the Land Information New Zealand Data Service provides access to high-resolution satellite and UAV imagery and demonstrates their value in

https://doi.org/10.54389/CYNC2230

disaster recovery [15], [16], there have been no notable attempts within the context of New Zealand to automate the extraction of flood damage information from these data sources.

Therefore, this study aims to bridge this gap by investigating the effectiveness of transfer learning for flood damage detection using drone images in New Zealand. We evaluate the use of transfer learning with a Convolutional Neural Network (ResNet50) [13] to classify (damaged vs non-damaged) post flood drone images captured following Cyclone Gabrielle which devastated the north island of New Zealand in 2023.

The main contributions of this work can be summarized as follows: (1) the creation and use of a manually labeled drone image dataset for flood damage classification; (2) the application of transfer learning using a ResNet50 model to create a fine-tuned damage classification model using a small dataset.

II. Methodology

This section briefly explains the dataset, preprocessing steps followed, model architecture and experimental setup used in this work in progress study.

The method consists of 3 main stages as illustrated in Figure 1: (1) Pre-processing the collected data (2) Fine-tunning a pre-trained ResNet50 model using transfer learning (3) Evaluating the fine-tuned model.

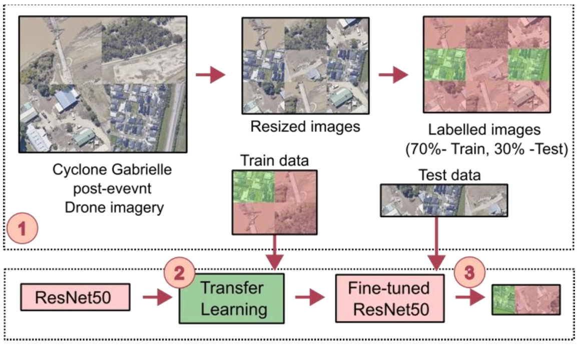

Figure 1: Overview of the proposed transfer learning-based methodology.

A. Dataset and Study Area

The dataset consists of drone images captured following Cyclone Gabrielle, which had significant impacts on the North Island of New Zealand in February 2023. Cyclone Gabrielle is considered one of the most destructive recent natural disasters in New Zealand, causing severe flooding, landslides and extensive damage to infrastructure including buildings, roads and bridges.

The most heavily affected regions included Hawke’s Bay, Tairāwhiti (Gisborne) and parts of the Tararua District. Major rivers such as the Tūtaekurī and Ngaruroro overflowed, causing widespread flooding in residential areas and agricultural land.

The drone images used in this study was obtained from publicly available datasets released by Toitū Te Whenua - Land Information New Zealand (LINZ) following Cyclone Gabrielle [16]. These datasets provide high-resolution (0.1m) aerial imagery that supports disaster response, recovery planning and environmental assessment.

From the available imagery, the initial study focused on 100 drone images to represent a balanced set of flood-affected and non-affected scenes. The image annotation process focused on images that clearly captured visible flood damage or undamaged areas across different urban and semi-urban locations within the affected regions as illustrated in Figure 2.

The final dataset contains 60 images representing damaged areas and 40 images representing undamaged areas. Images were manually labeled based on visible indicators of flood impact such as water inundation, debris and structural damage. The dataset was divided into 70% training data and 30% testing data, ensuring that both classes were represented in each split.

Figure 2: Damaged and undamaged drone images. The image on the left is manually labeled as an undamaged image and the image on the right is labeled as a damaged image based on their visual features.

B. Image Preprocessing

Before training the model, all images were resized to 224 × 224 pixels, which is the standard input size required by the ResNet50 architecture. The images were also normalized using the ImageNet mean and standard deviation values to ensure compatibility with the pre-trained model. These preprocessing steps help standardize the input data and improve model convergence during training.

C. Model Architecture

A ResNet50 convolutional neural network was used as the base model for this study. ResNet50 is a 50-layer deep residual network that introduces residual (skip) connections to address the vanishing gradient problem, enabling the training of deep neural networks [13].

The model used in this study was pre-trained on the ImageNet dataset, which is a large-scale image dataset containing millions of labeled images across different object categories [17]. Models trained on ImageNet learn rich feature representations that can be transferred to other computer vision tasks through transfer learning.

To adapt the model for the flood damage classification task, the final fully connected layer of ResNet50 was replaced with a two-node output layer corresponding to the Damaged and Non-Damaged classes as illustrated in Figure 3.

Since the dataset used in this study is relatively small, the convolutional layers of the network were frozen and only the final classification layer was trained. This approach helped to reduce the risk of overfitting while significantly decreasing training time.

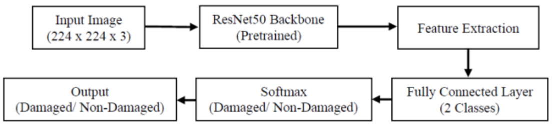

Figure 3: Architecture of the modified ResNet50 model used for flood damage classification.

D. *Training and Evaluation Setup*

The model was trained using the cross-entropy loss function and the Adam optimizer. Key hyperparameters, including learning rate, batch size, and number of epochs, were selected through hyperparameter tuning to optimize classification performance.

The experiments were conducted on the Google Colab platform with an NVIDIA Tesla T4 GPU (16 GB VRAM) and approximately 12 GB of system memory.

The performance of the trained model was evaluated using several standard classification metrics including accuracy, precision, recall and F1-score, along with confusion matrix analysis [18]. These metrics provide a comprehensive evaluation of the model's ability to correctly classify damaged and non-damaged images.

III. Results

The proposed model demonstrates effective performance in classifying damaged and non-damage images. Following hyperparameter tuning, the optimal model configuration selected based on its superior performance on the test set was used for final evaluation..

On the test set of 30 drone images, the fine-tuned model achieved an accuracy of 0.87, precision of 0.87, recall of 0.87, and an F1-score of 0.86, indicating a balanced performance across all evaluation metrics. The precision and recall values are well balanced at 0.87, indicating that the model is equally effective in classifying flood affected images while minimizing false positives. The resulting F1-score of 0.86 indicates stable and reliable performance across both metrics, with no significant bias toward over- or under-prediction.

In addition, the model maintains high computational efficiency, with an average training time of approximately 3 minutes and a mean inference time of around 20 milliseconds per image, highlighting its suitability for near real-time applications.

Further, the confusion matrix for the fine-tuned ResNet50 model is illustrated in Figure 4. The confusion matrix indicates that 17 damaged images were correctly classified, while 1 damaged image was misclassified. Similarly, 9 non-damaged images were correctly classified and 3 non-damaged images were incorrectly classified.

Overall, 26 out of the 30 test images were correctly classified, demonstrating the ability of the model to effectively distinguish between damaged and non-damaged areas.

The results indicate that the model performs well for both classes, although the classification performance is slightly better for the damaged class, which may be due to the stronger visual indicators present in flooded scenes.

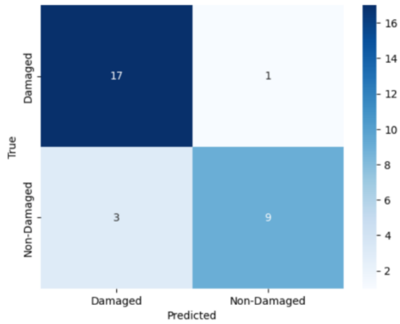

Figure 4: Confusion matrix of the fine-tuned ResNet50 model for flood damage classification

IV. Discussion

The results demonstrate that transfer learning using the ResNet50 architecture can effectively be applied to classify flood-related drone imagery. Hyperparameter tuning process revealed that the fine-tuned model trained for 20 epochs achieved better and most stable performance. This improvement is likely due to an optimal balance between sufficient learning and generalization, allowing the model to capture relevant features without overfitting, thereby producing more reliable results on the test set. However, the model showed slightly stronger performance in detecting damaged images, which may be due to the presence of clear visual indicators such as standing water, debris and structural damage within the affected areas.

Despite these promising findings, several limitations should be acknowledged. The dataset used in this study is relatively small and geographically limited, consisting of only 100 manually annotated drone images from the Cyclone Gabrielle event. Therefore, the results should be interpreted as preliminary findings rather than a fully generalized solution for flood damage detection.

Furthermore, the current approach focuses only on binary image-level classification, where each image is categorized as either damaged or non-damaged. This approach does not provide detailed information about the location or severity of damage within the image, which may be important for practical disaster response applications.

In addition to the overall accuracy, the confusion matrix analysis provides further insight into the model performance. The model correctly classified most of both damaged and non-damaged images, with only a small number of misclassifications. This indicates that the model is capable of effectively distinguishing between the two classes. The results also show that the model performs slightly better in identifying damaged areas, likely due to the presence of clear visual features such as water and debris. Overall, these findings demonstrate the effectiveness of the proposed approach for flood damage classification using a limited dataset.

V. Future Work

This study represents an initial step toward automated flood impact detection using drone imagery. Several extensions are planned to further expand and improve the current work.

Firstly, future experiments will compare the performance of other deep learning architectures with the ResNet50 model used in this study. In particular, models such as U-Net and YOLOv8 will be evaluated. This comparison will help to identify which deep learning architecture is most suitable for flood damage detection using aerial imagery.

Secondly, the study will be extended to incorporate satellite imagery from Sentinel-1 and Sentinel-2 satellites. The results obtained from drone imagery and satellite imagery will be compared in order to better understand the strengths and limitations of each data source. Such a comparison can help determine which type of remote sensing imagery is more effective for accurate flood impact detection.

Finally, the current approach focuses on image-level binary classification. Future work will explore more detailed object-level damage assessment, where individual damaged structures such as buildings, roads and bridges can be detected and analyzed separately rather than assigning a single damage label to the entire image. This would significantly improve the practical usefulness of the system for disaster response, damage assessment and recovery planning.

VI. Conclusion

This work-in-progress study demonstrates the potential of using transfer learning for flood damage classification using drone imagery. In this study, drone images captured after Cyclone Gabrielle in New Zealand were used to develop an automated damage classification model. A pre-trained ResNet50 model was fine-tuned using a manually annotated image dataset consisting of 100 drone images, which were labeled as either Damaged or Non-Damaged.

The experimental results showed that the fine-tuned model achieved an accuracy of approximately 87% on the test dataset, indicating that transfer learning can be effectively applied to disaster damage assessment tasks using relatively small datasets. These results demonstrate the potential of deep learning-based approaches to support rapid post-disaster damage assessment using aerial imagery.

However, the dataset size and geographic scope used in this study are limited. Therefore, the findings presented in this work should be considered as preliminary results. Future work will include further model comparisons, integration of satellite imagery and the extension of the current approach toward object-level damage detection to provide more detailed information for disaster response and recovery planning.

References

[1] CRED, "2024 Disasters in Numbers," 2025.

[2] UNDRR, *The Human Cost of Disasters: An Overview of the Last 20 Years*. 2020.

[3] NASA, "Cyclone Gabrielle Lashes New Zealand." Accessed: Mar. 27, 2026. [Online]. Available: https://science.nasa.gov/earth/earth-observatory/cyclone-gabrielle-lashes-new-zealand-150972/

[4] ESNZ, "Nationwide study reveals escalating flood risk." Accessed: Mar. 27, 2026. [Online]. Available: https://www.earthsciences.nz/news/nationwide-study-reveals-escalating-flood-risk

[5] O. D. Aribisala, S.-G. Yum, M. Das Adhikari, and M.-S. Song, "Flood damage assessment: a review of microscale methodologies for residential buildings," *Sustainability*, vol. 14, no. 21, p. 13817, 2022.

[6] N. Kerle, F. Nex, M. Gerke, D. Duarte, and A. Vetrivel, "UAV-based structural damage mapping: A review," *ISPRS Int. J. Geoinf.*, vol. 9, no. 1, p. 14, 2019.

[7] X. X. Zhu *et al.*, "Deep learning in remote sensing: A comprehensive review and list of resources," *IEEE Geosci. Remote Sens. Mag.*, vol. 5, no. 4, pp. 8–36, 2017.

[8] D. Whitehurst, K. Joshi, K. Kochersberger, and J. Weeks, "Post-flood analysis for damage and restoration assessment using drone imagery," *Remote Sens. (Basel).*, vol. 14, no. 19, p. 4952, 2022.

[9] J. Sublime and E. Kalinicheva, "Automatic post-disaster damage mapping using deep-learning techniques for change detection: Case study of the Tohoku tsunami," *Remote Sens. (Basel).*, vol. 11, no. 9, p. 1123, 2019.

[10] R. Gupta *et al.*, "xbd: A dataset for assessing building damage from satellite imagery," *arXiv preprint arXiv:1911.09296*, 2019.

[11] S. Kaur *et al.*, "Transfer learning-based automatic hurricane damage detection using satellite images," *Electronics (Basel).*, vol. 11, no. 9, p. 1448, 2022.

[12] Q. Lin, T. Ci, L. Wang, S. K. Mondal, H. Yin, and Y. Wang, "Transfer learning for improving seismic building damage assessment," *Remote Sens. (Basel).*, vol. 14, no. 1, p. 201, 2022.

[13] 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*, 2016, pp. 770–778.

[14] J. Jiang, Y. Shu, J. Wang, and M. Long, "Transferability in deep learning: A survey," *arXiv preprint arXiv:2201.05867*, 2022.

[15] Land Information New Zealand (LINZ), "Hawke's Bay 0.1m Cyclone Gabrielle Aerial Photos (2023)," 2023, Accessed: Mar. 27, 2026. [Online]. Available: https://data.linz.govt.nz/layer/112726-hawkes-bay-01m-cyclone-gabrielle-aerial-photos-2023/

[16] Land Information New Zealand (LINZ), "High resolution imagery of flood-hit areas supports cyclone recovery." Accessed: Mar. 27, 2026. [Online]. Available: https://www.linz.govt.nz/news/2023-03/high-resolution-imagery-flood-hit-areas-supports-cyclone-recovery

[17] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *2009 IEEE conference on computer vision and pattern recognition*, Ieee, 2009, pp. 248–255.

[18] Ž. Vujović, "Classification model evaluation metrics," *International Journal of Advanced Computer Science and Applications*, vol. 12, no. 6, pp. 599–606, 2021.