# AUTOMATED POST-EARTHQUAKE PAVEMENT DAMAGE DETECTION USING DEEP LEARNING ON AERIAL DIGITAL IMAGES

# DETECCIÓN AUTOMATIZADA DE DAÑOS EN PAVIMENTOS POST-TERREMOTO MEDIANTE APRENDIZAJE PROFUNDO APLICADO A IMÁGENES AÉREAS

**Brigitte Puchoc[1*], Miguel Estrada[1,2], Italo Inocente[3], Fernando García[1]**

[1] *Centro Peruano Japonés de Investigaciones Sísmicas y Mitigación de Desastres, Universidad Nacional de Ingeniería, Lima, Perú*

[2] *Facultad de Ingeniería Civil, Universidad Nacional de Ingeniería, Lima, Perú*

[3] *Graduate School of Engineering, Chiba University, Chiba, Japan*

Recibido (Received): 28 / 02 / 2025 Publicado (Published): 30 / 12 / 2025

**ABSTRACT**

Perú is located in a highly seismic zone, making it vulnerable to infrastructure damage caused by earthquakes. For this reason, an early evaluation after a severe earthquake is important for mitigating impacts, particularly for decision-makers. Pavements often sustain significant damage among the most affected infrastructure, leading to crack formation. These cracks not only disrupt transportation networks but also pose safety hazards and hinder economic activities. Identifying pavement cracks is an important step in post-earthquake assessment; however, traditional inspection methods are typically slow, error-prone, labor-intensive, and often inaccessible in high-risk zones, limiting their effectiveness. This study applies deep learning techniques for automated pavement crack detection in post-earthquake scenarios using aerial images to address this issue. A DeepLabv3+ convolutional neural network was trained on 5600 labeled pavement crack images. The model achieved an Intersection over Union (IoU) of 65% on the validation set. It was subsequently applied to post-earthquake imagery from the 2007 Pisco earthquake in Peru. When evaluated against a manually segmented reference, the model yielded an IoU of 47.2% and an F1-score of 64.3%. These results indicate strong generalization capabilities despite the domain shift and resolution gap between training and testing data. The proposed method demonstrates the potential of deep learning models for rapid and scalable assessment of post-earthquake pavement damage, reducing reliance on manual inspections and supporting timely decision-making in disaster response contexts.

*Keywords: earthquake damage, pavement, deep learning, photogrammetry*

**RESUMEN**

Perú se encuentra en una zona altamente sísmica, lo que lo hace vulnerable a daños en la infraestructura provocados por terremotos, por ello una evaluación rápida después de un sismo de gran magnitud es importante para mitigar sus impactos, especialmente para los tomadores de decisiones. Entre las infraestructuras más afectadas, encontramos al pavimento, los cuales suelen sufrir daños significativos, generando la formación de grietas, estas no solo interrumpen la red de transporte, sino que también representan un riesgo para la seguridad y afectan las actividades económicas.
La detección de grietas en pavimentos es un paso clave en este proceso; sin embargo, los métodos de inspección tradicionales suelen ser lentos, propensos a errores, demandan un alto esfuerzo manual y, en muchas ocasiones, son inaccesibles en zonas de alto riesgo, lo que limita su efectividad. Para abordar este problema, en este estudio se aplican técnicas de aprendizaje profundo para la detección automatizada de grietas en pavimentos a partir de imágenes aéreas en escenarios post-sismo.
Se entrenó una red neuronal convolucional DeepLabV3+ con aproximadamente 5,600 imágenes etiquetadas de grietas en pavimentos. El modelo alcanzó un IoU de 65 % en el conjunto de validación, aplicándose a imágenes post-sismo del terremoto de Pisco de 2007. Al compararse con una segmentación manual de referencia, el modelo obtuvo un IoU de 47.2 % y un F1-score de 64.3 %. Estos resultados reflejan una buena capacidad de generalización pese a las diferencias entre los datos de entrenamiento y prueba. El método propuesto demuestra el potencial del aprendizaje profundo para realizar evaluaciones rápidas y escalables de daños en pavimentos post-terremoto, reduciendo la dependencia de inspecciones manuales y facilitando una toma de decisiones oportuna en contextos de emergencia.

*Palabras Clave: peligro sísmico, pavimento, aprendizaje profundo, fotogrametría.*

[1] * Corresponding author:
E-mail: bpuchoce@uni.edu.pe

## 1. INTRODUCTION

Earthquakes are among the most destructive natural disasters, frequently causing extensive damage to critical infrastructure [1]. Although considerable research has focused on evaluating structural damage to buildings, damage to transportation infrastructure, especially road networks, has received comparatively less attention [2]. However, preserving road connectivity is essential in post-earthquake scenarios, as disrupted road networks hinder rescue, relief, and evacuation operations by limiting emergency vehicle mobility and increasing response times [3]. Additionally, compromised roads pose risks to public safety, intensify impacts on human lives and social stability, and impede long-term economic recovery and the restoration of other critical infrastructure systems [4]. Currently, road assessment methods rely primarily on visual inspections, which are costly, labor-intensive, and prone to unreliable results [5]. Furthermore, repairs to damaged road networks often take months or even years, resulting in prolonged disruptions during reconstruction [3], [6].

In recent years, remote sensing technologies have emerged as effective tools for assessing post-earthquake infrastructure damage [7]. Techniques such as aerial photogrammetry and unmanned aerial vehicles (UAVs) are increasingly employed to collect data rapidly over extensive areas [8]. These methods provide safer, faster, and more comprehensive evaluations than manual inspections [9]. However, interpreting these large datasets demands significant human effort and specialized expertise, which can cause delays and inconsistencies in damage assessment [10]. Recent studies have integrated artificial intelligence and machine learning techniques with remote sensing data to address these challenges. Convolutional neural networks (CNNs), in particular, have demonstrated considerable potential for automating the detection of infrastructure damage, significantly reducing both analysis time and required human resources [11]. Although several automated techniques exist for detecting general road damage, pavement failures induced by earthquakes, typically manifested as cracks, exhibit unique characteristics [12]. Existing methods often exhibit limitations identifying this damage type [13].

To overcome the limitations such as the dependency on manual interpretation and the inefficiencies associated with processing large volumes of geospatial data, this study introduces a dataset of flexible pavement distress in Peru, created from high-resolution aerial imagery. and implements a DeepLabv3+ architecture, selected for its high accuracy and fast inference [14], both critical for real-time applications. In addition, a tailored post-processing strategy is introduced to refine the delineation of pavement damage. To evaluate the effectiveness of the proposed approach in a real-world post-seismic scenario, the trained model is applied to four historical aerial imagery acquired after the 2007 Pisco earthquake. This methodology aims to improve the precision and robustness of automated road damage assessments, thereby supporting faster and more informed after the 2007 Pion-making in post-earthquake response efforts.

## 2. BACKGROUND

On August 15, 2007, an earthquake struck off the coast of Peru, causing severe damage throughout the region. Its epicenter was located offshore near Pisco, in the Ica region, registering a magnitude of 7.9 on the Richter scale. The event resulted in 519 fatalities, approximately 1,366 injuries, and the collapse of more than 58,000 buildings [15].

This research explores the potential of deep learning techniques as an alternative for automated post-earthquake pavement damage detection. The validation of the proposed approach is carried out using four historical aerial images captured over the district of San Clemente, an area of strategic importance as it encompasses the former Panamericana Sur highway, which connects Lima and Pisco [16], as illustrated in Fig. 1. This connection is particularly relevant given that Pisco was the region most severely affected by the earthquake. Moreover, San Clemente shares significant geographic and socio-economic ties with Pisco and plays a key role in the area's recovery efforts [16].

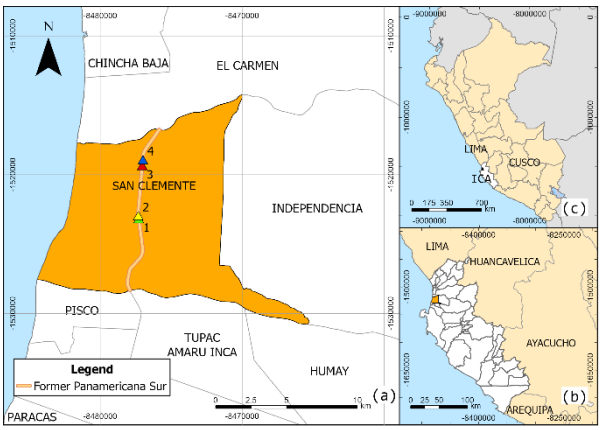


Fig. 1. (a) Study area: District of San Clemente, where the former Panamericana Sur passes through. The four colored markers indicate the locations of the historical aerial photographs used for analysis validation. (b) Location of the Ica Region within Peru, highlighting the province of Pisco and its connection to the capital city, Lima. (c) Reference map of Peru showing the position of Ica in the national context.

## 3. METHODOLOGY

This section presents the methodology used for the automated detection of post-earthquake pavement cracks using deep learning on aerial images. The workflow includes data acquisition, preprocessing, model training, and post-processing adaptations to optimize crack segmentation. Each stage is detailed in the following subsections, and the overall workflow is illustrated in Fig. 2.

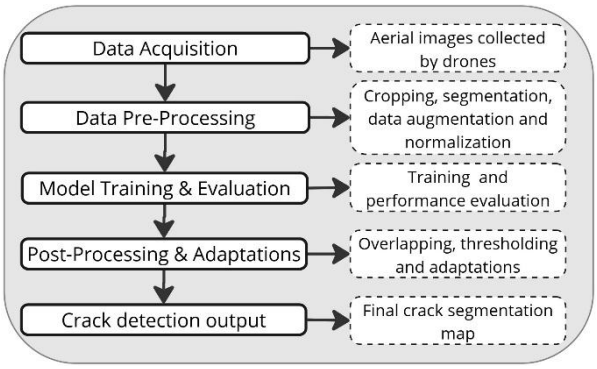


Fig. 2. Methodological pipeline for deep learning-based post-earthquake pavement crack detection.

### 3.1 Data acquisition

Acquiring post-earthquake pavement imagery poses significant challenges due to the rarity of such events and the historically limited availability of advanced technologies for systematic data collection. As an alternative, this study introduces a curated dataset of pavements exhibiting typical surface distress under non-disaster conditions. The dataset comprises four orthomosaics (O1, O2, O3, and O4) each corresponding to flexible pavement sections, with distinct ground sample distances (GSD) (see TABLE I), providing variability in image resolution. O1 corresponds to the San Martín de Porres district, specifically along Eduardo de Habich Avenue, while O2, O3, and O4 were captured in different sections of Maestro Peruano Avenue in the district of Comas. All images were acquired using a DJI Mavic 2 Pro drone and were processed and labeled by the authors as part of the research workflow developed in this article.

TABLE I
Orthomosaic data summary

| ID | GSD (cm/pixel) | Images | | |
|---|---|---|---|---|
| | | Obtained | Cracked | Non-Cracked |
| O1 | 0.52 | 4,039 | 636 | 204 |
| O2 | 1.39 | 1,536 | 73 | 23 |
| O3 | 1.40 | 930 | 47 | 15 |
| O4 | 1.41 | 1,501 | 64 | 23 |
| **TOTAL** | | **8,006** | **820** | **300** |

### 3.2 Data pre-processing

The pre-processing stage encompassed a series of transformations to structure the dataset for model training. The first step involved dividing high-resolution orthomosaics into 448×448 pixels patches, see Fig. 3, a standard image processing practice. This approach preserves crack details while maintaining a computationally efficient input size for deep learning models [17], [18]. Furthermore, this method simplifies manual annotation by reducing the scale of images requiring segmentation.

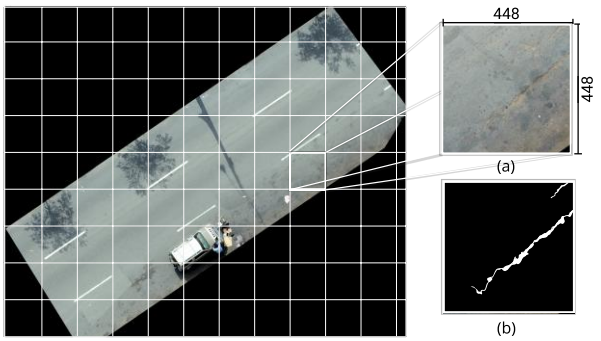


Fig. 3. The cropping process applied to the original orthomosaics, where each image was divided into 448×448 patches. Fig. 3(a) presents an example of an extracted image patch, while Fig. 3(b) shows its corresponding binary mask generated through manual crack segmentation.

Subsequently, manual segmentation was performed to generate labeled data, providing ground truth masks for supervised learning. Based on the image distribution summarized in TABLE I, a total of 1,120 cropped images, with their corresponding binary masks, were obtained for initial training. To improve model generalization, data augmentation techniques were applied, including rotation, brightness adjustments, and contrast modifications, increasing the dataset size by a factor of five. As a result, a total of 5,600 training samples were generated. Finally, image normalization was conducted to standardize pixel intensity distributions, ensuring consistency across the dataset before training.

It is important to note that, although only 300 fully non-cracked images were selected, a considerable number of non-cracked pixels are present within the 820 cracked images, as surface cracking typically affects only a localized portion of each crop. As a result, the model was effectively exposed to a wide distribution of non-cracked pixels under varying lighting conditions, GSD values, and pavement appearances. The inclusion of the 300 entirely non-cracked images was therefore not intended to introduce a new class representation, but rather to reinforce the variability of non-cracked areas and improve class balance at the image level. This strategy contributed to a more robust training

process while maintaining computational efficiency and minimizing redundancy.

### 3.3 Model training and adaptations

Model training was conducted using the DeepLabV3+ architecture, a deep learning architecture widely employed for semantic segmentation tasks due to its accuracy and inference speed [19]. The dataset was divided into 80% for training, 10% for validation, and 10% for testing to ensure robust model evaluation. The model was trained for a maximum of 200 epochs, with early stopping implemented to prevent overfitting. Training was performed with a batch size of 8, using the Adam optimizer with a learning rate of 0.001. The training was performed on a high-performance workstation equipped with an NVIDIA RTX A4000 GPU (16 GB VRAM) and 128 GB of RAM, ensuring efficient processing of the deep learning model. The Intersection over Union (IoU) metric was employed to assess segmentation performance. It is a standard evaluation metric for semantic segmentation tasks, defined as the ratio between the area of overlap and the area of union between the predicted segmentation and the ground truth. Mathematically, it is expressed as equation (1)

$$IoU = \frac{Prediction \cap Ground\ Truth}{Prediction \cup Ground\ Truth} \quad (1)$$

This metric ranges from 0 to 1, where a value closer to 1 indicates higher segmentation accuracy. It provides a robust measure of how well the predicted mask aligns with the actual object boundaries in the image.

As a result of this training, Fig. 4 presents the IoU curves for both the training and validation datasets. The model exhibits a steady increase in performance during the initial epochs, showing a rapid improvement in IoU. Around epoch 20, the training IoU stabilized above 0.78, while the validation IoU plateaued around 0.65.

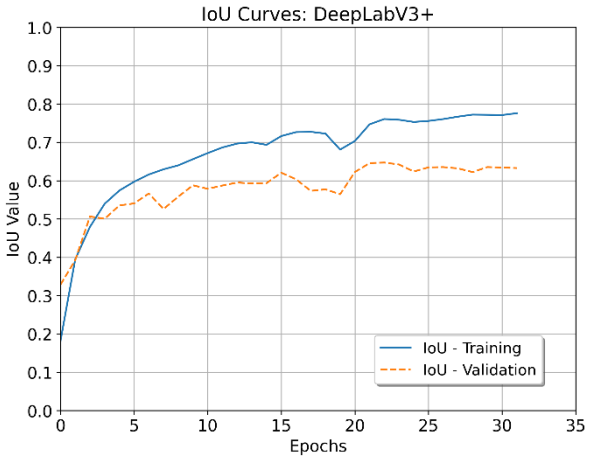


Fig. 4. Training and validation IoU curves for the DeepLabV3+ model.

This behavior indicates that the model effectively learns to crack patterns from the training set, though a small performance gap between training and validation suggests slight overfitting. However, the validation performance remains stable, implying that generalization is maintained within an acceptable range.

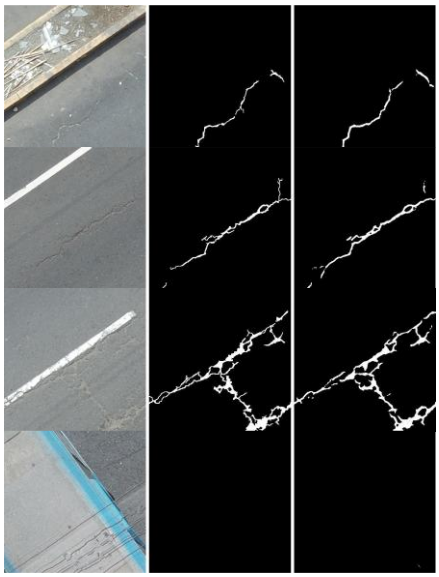

(a) (b) (c)

Fig. 5. Inference results of DeepLabV3+ trained model on test set images.

The inference analysis on the test set confirms the effectiveness of the DeepLabV3+ model in segmenting pavement cracks from aerial images. Fig. 5 illustrates the results, where Fig. 5(a) shows the original pavement images, Fig. 5(b) corresponds to the manually segmented ground truth, and Fig. 5(c) presents the crack masks predicted by the model.

These results demonstrate that the model accurately detects most cracks, maintaining continuity and characteristics. However, some limitations are evident, including minor segmentation

discontinuities and false positives in visually complex pavement regions, particularly near road markings and structural joints. Despite these challenges, the predicted masks exhibit a high consistency with ground truth annotations, indicating the model's strong generalization ability for pavement imagery under normal conditions.

### 3.4 Post Processing & Adaptations

The trained model was applied in the aerial images obtained by the earthquake damage from 2007 Pisco Earthquake on the road former "*Panamericana Sur*" highway, representing data not included in the training dataset. These images, like presented in Fig. 6, were taken on August 17, 2007, at 2:00 PM approximately using an aircraft and has a resolution of 2870×2896 pixels with an approximate GSD of 4 cm/pixel.

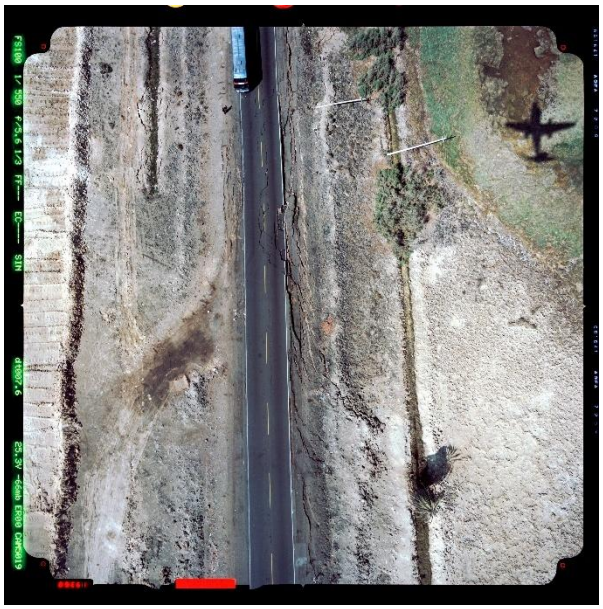

Fig. 6. Image from former Panamericana Sur, Post-earthquake

The trained model was applied to the aerial images using a sliding window approach with 448×448 patches and a 50% overlap, given that the aerial imagery had significantly larger dimensions than the training images. Additionally, a probability threshold of 0.5 was applied to generate a binary crack mask capable of distinguishing between cracked and non-cracked pavement areas.

However, this approach exhibited limitations when applied to post-earthquake images, requiring further adaptations to optimize segmentation performance. This was primarily due to the higher capture altitude, which resulted in patches covering both sides of the road, as shown in Fig. 5, with pavement occupying only one-third of each patch. At the same time, the remaining area consisted mainly of soil and debris. Since the model was trained on images where the background primarily consisted of pavement, as the original patches had a limited field of view, as shown in Fig. 3(a), this discrepancy negatively affected segmentation performance. To address this issue, cropping the pavement region, as shown in Fig. 7 (a).

Although cropping improved segmentation, the results remained inconsistent, as shown in Fig. 7(b), which illustrates the inference using a 50% overlap and a 0.5 threshold on the cropped image. This inconsistency occurred because, compared to the test images, the post-earthquake image exhibited fewer false positives in mask probabilities, meaning that even regions with a low predicted probability of being a crack were actual cracks; this allowed for a lower probability threshold of 0.025, which, while maintaining the 50% overlap, improved crack visibility, as seen in Fig. 7(c). However, even with this adjustment, segmentation results were not fully optimized. The overlap was increased to 70% to improve consistency, leading to a significant enhancement in segmentation performance, as illustrated in Fig. 7 (d).

Despite these improvements, segmentation still faced challenges due to ground sample distance (GSD) discrepancies. The post-earthquake image had a GSD of approximately 4 cm/pixel, whereas the highest GSD in the training set was 1.4 cm/pixel, leading to a loss of detail in crack patterns. To compensate for this difference, bicubic rescaling technique was applied to artificially increase pixel density while maintaining spatial consistency. A factor of 2.5 would have resulted in an artificial GSD of approximately 1.5 cm/pixel, making it more comparable to the training data. However, after evaluating segmentation performance across different rescaled images, a scaling factor 1.5 was selected as it provided the most accurate and reliable results, as shown in Fig. 7(e). The segmentation results of the rescaled image, using a 0.025 threshold and 70% overlap, are presented in Fig. 7(f).

Although the inference process initially employed a 70% overlap, segmentation accuracy decreased near the road edges due to border patches lacking overlapping predictions, limiting the model's ability to infer crack patterns accurately. A 300-pixel reflection padding was applied to mitigate this, artificially extending the pavement beyond its original boundaries, as seen in Fig. 7(g). This allowed edge pixels to interact with a more representative context, improving segmentation along the road edges, as illustrated in Fig. 7(h). Finally, the padded borders were cropped after segmentation, restoring the image to its original dimensions and producing the final crack segmentation mask, as shown in Fig. 7(i).

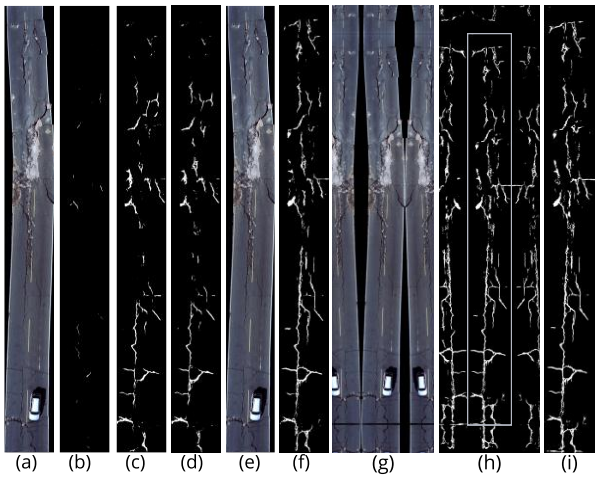

Fig. 7. Post-Processing Techniques Applied to Crack Segmentation in Post-Earthquake Images

## 4. RESULTS AND DISCUSSION

The DeepLabV3+ trained model was applied on four aerial imageries of earthquake-damaged pavement to assess its performance under post-disaster conditions. Fig. 8 presents the segmentation results, where cracks detected by the model are highlighted in green.


Fig. 8. Post-Earthquake Pavement Crack Detection Results.

The model effectively identifies major crack patterns along the damaged roadway, particularly in areas where pavement displacement is evident. A quantitative pixel-wise evaluation against a manually segmented reference resulted in an Intersection over Union (IoU) of 47.2%, a precision of 49.0%, and a recall of 93.1%, resulting in an F1-score of 64.3%. These results confirm that the model successfully detects most crack regions, but tends to over-segment them by generating masks that are visibly thicker than the actual cracks, see Fig. 9 This behavior increases the number of false positives and reduces spatial accuracy, but ensures a high sensitivity level, an advantageous trait in rapid post-disaster screening scenarios.

The model effectively identifies major crack patterns along the damaged roadway, particularly in areas where pavement displacement is evident. However, some limitations were observed, including the incomplete detection of finer cracks and the misclassification of terrain irregularities as pavement damage. These errors arise from the high GSD of the post-earthquake imagery and variations in lighting conditions, both of which impact the segmentation accuracy. Another limitation of the proposed approach is that it is not fully automated, as it requires manual cropping of road areas from the full aerial images prior to inference.

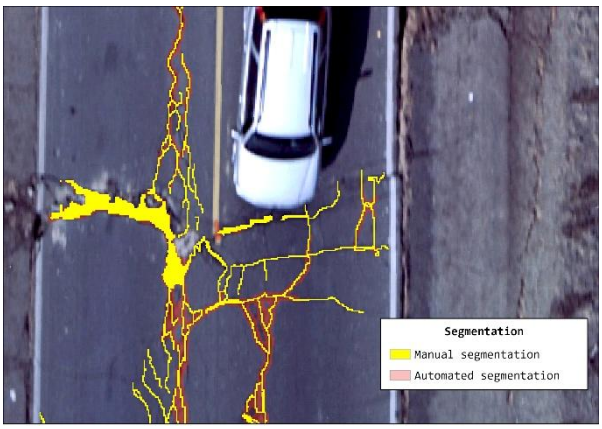

Fig. 9. Visual comparison between manual and automated pavement crack segmentation

## CONCLUSIONS

The proposed model, based on the DeepLabV3+ architecture, demonstrated the ability to identify major crack patterns despite the challenges posed by aerial imagery, such as high ground sampling distance (GSD), variable lighting conditions, and surface irregularities. To address these limitations, a tailored post-processing strategy was implemented, incorporating bicubic rescaling, adaptive thresholding, and a 70% overlap in inference tiles, which improved the reliability of the segmentation outputs.

A quantitative pixel-wise evaluation against manually segmented references yielded an Intersection over Union (IoU) of 47.2%, a precision of 49.0%, and a recall of 93.1%, resulting in an F1-score of 64.3%. These results confirm that the model successfully detects most crack regions, although it tends to over-segment, as reflected in the lower precision.

Overall, the findings highlight the potential of deep learning, when combined with appropriate pre- and post-processing techniques, to support rapid and

scalable assessments of pavement damage following seismic events. This automated approach reduces dependence on manual inspections and facilitates timely decision-making in emergency response.

For future research, additional architectures will be explored, and earthquake damage data from roads in different countries will be collected to assess the generalization and robustness of the trained model across diverse geographic regions and pavement conditions. Additionally, a study will be conducted on the influence of GSD on inference results, focusing on its impact on segmentation accuracy and model performance in varying post-earthquake environments.

## ACKNOWLEDGMENTS

The authors would like to express their gratitude to the Geomatic Laboratory of CISMID for providing support and access to the aerial imagery dataset used in this study.

## REFERENCES

[1] F. Kadri, B. Birregah, and E. Châtelet, "The Impact of Natural Disasters on Critical Infrastructures: A Domino Effect-based Study," *J. Homel. Secur. Emerg. Manag.*, vol. 0, no. 0, Jan. 2014, doi: 10.1515/jhsem-2012-0077.

[2] H. McCullough, "Earthquake damage to transportation systems," National Oceanic and Atmospheric Administration, USA, 1994.

[3] S. Karimzadeh, M. Ghasemi, M. Matsuoka, K. Yagi, and A. C. Zulfikar, "A Deep Learning Model for Road Damage Detection After an Earthquake Based on Synthetic Aperture Radar (SAR) and Field Datasets," *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.*, vol. 15, pp. 5753–5765, 2022, doi: 10.1109/JSTARS.2022.3189875.

[4] M. Sathurshan, A. Saja, J. Thamboo, M. Haraguchi, and S. Navaratnam, "Resilience of Critical Infrastructure Systems: A Systematic Literature Review of Measurement Frameworks," *Infrastructures*, vol. 7, no. 5, p. 67, May 2022, doi: 10.3390/infrastructures7050067.

[5] A. Ragnoli, M. R. De Blasiis, and A. Di Benedetto, "Pavement Distress Detection Methods: A Review," *Infrastructures*, vol. 3, no. 4, p. 58, Dec. 2018, doi: 10.3390/infrastructures3040058.

[6] D. Blackman, H. Nakanishi, and A. M. Benson, "Disaster resilience as a complex problem: Why linearity is not applicable for long-term recovery," *Technol. Forecast. Soc. Change*, vol. 121, pp. 89–98, Aug. 2017, doi: 10.1016/j.techfore.2016.09.018.

[7] F. Yamazaki and W. Lui, "Remote sensing technologies for post-earthquake damage assessment: a case study on the 2016 kumamoto earthquake," presented at the 6th Asia Conf. on Earthquake Engg., Asia, Sept. 2016.

[8] V. Croce and I. Zaragoza, "UAV-based 3D Photogrammetry for post-Earthquake Studies on Seismic damaged Cities – A Case Study: Castelluccio di Norcia," presented at the International Conference on Education and Information Systems, Technologies and Applications EISTA, Orlando, USA, July 2018.

[9] N. Kerle, F. Nex, M. Gerke, D. Duarte, and A. Vetrivel, "UAV-Based Structural Damage Mapping: A Review," *ISPRS Int. J. Geo-Inf.*, vol. 9, no. 1, p. 14, Dec. 2019, doi: 10.3390/ijgi9010014.

[10] V. Hoskere, Y. Narazaki, T. A. Hoang, and B. F. S. Jr, "Towards Automated Post-Earthquake Inspections with Deep Learning-based Condition-Aware Models," Sept. 24, 2018, *arXiv*: arXiv:1809.09195. doi: 10.48550/arXiv.1809.09195.

[11] N. S. Gulgec, M. Takáč, and S. N. Pakzad, "Structural Damage Detection Using Convolutional Neural Networks," in *Model Validation and Uncertainty Quantification, Volume 3*, R. Barthorpe, R. Platz, I. Lopez, B. Moaveni, and C. Papadimitriou, Eds., in Conference Proceedings of the Society for Experimental Mechanics Series. , Cham: Springer International Publishing, 2017, pp. 331–337. doi: 10.1007/978-3-319-54858-6_33.

[12] A. Banerji, P. Topdar, and A. Datta, "A Critical Review of Road Failures due to Major Earthquakes in Asian Continent," presented at the Sixth International Conference on Advances in Design and Construction of Structures, Bangalore, India, Oct. 2017.

[13] P. Dondi *et al.*, "Improving Post-Earthquake Crack Detection using Semi-Synthetic Generated Images," 2025, pp. 19–35. doi: 10.1007/978-3-031-91907-7_2.

[14] Y. J. Lee, H. G. Jung, and J. K. Suhr, "Semantic Segmentation Network Slimming and Edge Deployment for Real-Time Forest Fire or Flood Monitoring Systems Using Unmanned Aerial Vehicles," *Electronics*, vol. 12, no. 23, p. 4795, Nov. 2023, doi: 10.3390/electronics12234795.

[15] J. Alarcón, F. Taucer, and E. So, "THE 15 AUGUST 2007 PISCO, PERU, EARTHQUAKE - POST-EARTHQUAKE FIELD SURVEY," presented at the The 14 th World Conference on Earthquake Engineering, Beiking, China, Oct. 2008.

[16] J. O'Connor, L. Mesa, and M. Nykamp, "Damage to the Highway System from the Pisco, Peru Earthquake of August 15, 2007."

[17] S. Han, K. Phasouk, J. Zhu, and Y. Fong, "Optimizing deep learning-based segmentation of densely packed cells using cell surface markers," *BMC Med. Inform. Decis. Mak.*, vol. 24, no. 1, p. 124, May 2024, doi: 10.1186/s12911-024-02502-6.

[18] L. Hou, D. Samaras, T. M. Kurc, Y. Gao, J. E. Davis, and J. H. Saltz, "Patch-based Convolutional Neural Network for Whole Slide Tissue Image Classification," Mar. 09, 2016, *arXiv*: arXiv:1504.07947. doi: 10.48550/arXiv.1504.07947.

[19] Y. Chen, C. Zhou, Z. Yan, T. Huang, G. Wang, and J. Hu, "Lightweight Semantic Segmentation Network Based on DeepLabV3+," in *2022 International Conference on Artificial Intelligence and Computer Information Technology (AICIT)*, Yichang, China: IEEE, Sept. 2022, pp. 1–5. doi: 10.1109/AICIT55386.2022.9930215.