Cite this article
Kubo S, Zheng S and Chun P-J (2026)
Damage assessment of bridges using segmentation and shape analysis from aerial imagery.
Proceedings of the Institution of Civil Engineers – Forensic Engineering **179(2)**: 59–72,
https://doi.org/10.1680/jfoen.25.00025

Research Article
Paper 2500025
Received 16/05/2025; Accepted 09/10/2025

Forensic Engineering

Damage assessment of bridges using segmentation and shape analysis from aerial imagery

Shiori Kubo
Institute of Education, Research and Regional Cooperation for Crisis Management Shikoku, Kagawa University, Takamatsu, Japan
(Orcid:0000-0002-8416-3864) (corresponding author: kubo.shiori.a6@kagawa-u.ac.jp)

Shitao Zheng
Institute of Engineering Innovation, The University of Tokyo, Tokyo, Japan
(Orcid:0009-0008-1357-6080)

Pang-Jo Chun
Department of Civil Engineering, The University of Tokyo, Tokyo, Japan
(Orcid:0000-0002-9755-8435)

Rapid and quantitative assessment of bridge damage immediately after a disaster is critical for effective emergency response and early recovery. In this study, a method was developed to apply the YOLOv8 segmentation model to aerial imagery for detecting and classifying bridges based on the presence or absence of damage and to estimate bridge length and width from the detected regions to evaluate changes in shape dimensions. The results confirmed that for undamaged bridges, both bridge length and width could be estimated with high accuracy. In contrast, for damaged bridges, debris accumulation and inundation often resulted in only partial detection of the bridge, leading to a tendency for dimensional underestimation, particularly in the bridge length direction. Furthermore, it was confirmed that damage caused changes in the relative scaling relationship between bridge length and width, disrupting the original balance between these dimensions. These findings indicate that the observed dimensional shrinkage is not merely an estimation error but a quantitative characteristic of damage itself. This suggests that analysing changes in shape dimensions is a potentially effective method for detecting the presence of bridge damage.

Keywords: disaster/engineering bridges/UN SDG 9: Industry, innovation and infrastructure industry/UN SDG 11: Sustainable cities and communities

1. Introduction

Natural disasters such as earthquakes, tsunamis, and floods occur frequently around the world, causing damage to and functional disruptions of social infrastructure. In particular, damage to transportation infrastructure, including bridges and roads, significantly hinders rapid recovery and reconstruction efforts in disaster-affected areas. Therefore, promptly assessing the damage status of such infrastructure is critically important. Early assessment and evaluation of disaster damage enable prioritisation necessary for efficiently conducting emergency response and recovery activities, thereby contributing to the prevention of secondary disasters and the acceleration of recovery and reconstruction.

In recent years, many studies have been conducted to assess wide-area disaster-affected conditions, including bridges, roads, and entire urban areas, by applying deep learning techniques to satellite imagery, aerial photography, and unmanned aerial vehicle-captured images. Regarding damage detection for large-scale natural disasters, studies such as the detection of landslides using aerial imagery (Kubo *et al*., 2022) and the mapping of landslides and floods (Lang *et al*., 2024; Opara *et al*., 2024) have demonstrated the effectiveness of combining aerial imagery and deep learning techniques for wide-area disaster assessment. Many studies have also addressed building damage detection. For example, Miyamoto and Yamamoto (2020) detected buildings and identified the damage status using optical satellite imagery based on building information obtained from a geographic information system database. Furthermore, Bai *et al*. (2020) and Gholami *et al*. (2022) estimated building damage locations and classified damage levels based on pre- and post-disaster satellite imagery, while Alisjahbana *et al*. (2024) developed a damage detection model integrating building segmentation and damage classification. Regarding road infrastructure, Zhao *et al*. (2022) extracted and evaluated damaged road areas based on the tracking, learning, and detector framework using high-resolution post-disaster aerial imagery, and Ahmad *et al*. (2019) detected passable roads after floods based on satellite imagery and social media data. For bridges, Shukla *et al*. (2024) proposed a method for detecting changes in bridges using pre- and post-disaster satellite imagery, Liu *et al*. (2021) performed collapsed bridge detection by applying multiple regression models to change features derived from synthetic aperture radar (SAR) imagery, and Kopiika *et al*. (2025) performed damage detection and classification at the bridge component level by utilising SAR and high-resolution imagery.

These studies demonstrate that combining remote sensing technologies with deep learning methods is effective for the rapid assessment and evaluation of wide-area disaster effects, including damage to urban and transportation infrastructure. However, most previous studies have primarily focused on identifying the damage status of structures such as bridges, and methods for quantitatively evaluating damage conditions in relation to shape dimensions have not yet been sufficiently established. In this study, a segmentation approach is applied to aerial imagery to detect bridges according to their damage status, and for each detected bridge, the bridge length and width are calculated. The objective is to develop a new method for quantitatively assessing the disaster-affected conditions of bridges based on their shape dimensions.

2. Methods

2.1 Bridge segmentation model

In this study, a deep learning model from the YOLO (You Only Look Once) series, capable of addressing a wide range of image recognition tasks such as object detection and segmentation, is utilised to segment damaged and undamaged bridges from aerial imagery. For the segmentation, the segmentation model of YOLOv8 (Figure 1), released in January 2023 (Ultralytics, 2025), is adopted. YOLOv8 is designed based on the framework of the YOLO series, primarily YOLOv5, but incorporates enhancements such as improvements to the backbone architecture and loss functions, as well as the introduction of an anchor-free structure, achieving a balance between accuracy and inference speed. Furthermore, YOLOv8 maintains compatibility with previous versions of the YOLO series, allowing for easy switching between different versions, which is important from an implementation standpoint.

YOLOv8 has three key technical characteristics. First, by introducing an anchor-free structure, the output space is simplified, allowing the model to directly regress the centre coordinates and size of the target object. In contrast to conventional anchor-based structures, which predict the offsets from multiple predefined anchor boxes placed on feature maps, the anchor-free approach directly predicts the centre coordinates of objects and simultaneously estimates their size information, such as width and length, at those locations. For each grid cell, the following output vector is generated:

1. $z = (\hat{x}, \hat{y}, \hat{w}, \hat{h}, \hat{c}, m)$

where $(\hat{x}, \hat{y})$ represent the centre coordinates of the object, $(\hat{w}, \hat{h})$ denote its width and length, $\hat{c}$ is the class probability, and m is a coefficient vector for mask generation. This structure suppresses the generation of unnecessary candidate boxes, thereby improving

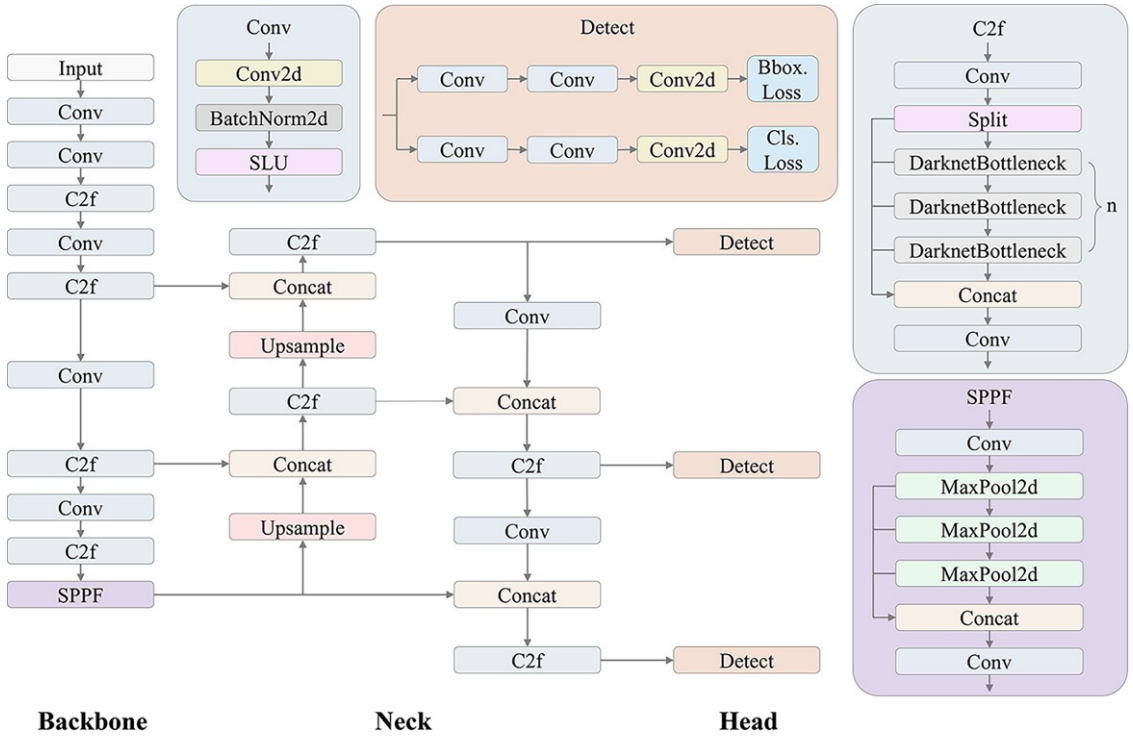

Figure 1. The network architecture of the YOLOv8 segmentation model is presented, including key components such as the backbone and mask prediction branch (created with reference to Wang *et al*. (2024))

computational efficiency and generalisation performance, and demonstrates high performance, particularly when applied to small datasets and real-world environments (Ge *et al*., 2021).

Second, in YOLOv8, the main component of the backbone has been changed from the C3 module to the C2f module (Figure 2). While C3 was based on repeated residual blocks following the design of ResNet, C2f adopts the design philosophy of DenseNet. It reduces the number of bottleneck layers by utilising skip connections and enhances feature reuse and gradient propagation efficiency through the splitting and merging of feature maps, resulting in a simpler and more efficient architecture (Sun *et al*., 2024; Wang *et al*., 2024).

Third, the segmentation model of YOLOv8 adopts a prototype-based structure for mask prediction. This structure is based on the method introduced in YOLACT (Bolya *et al*., 2019) and enables a significant reduction in computational cost and inference time compared with conventional approaches that individually generate dedicated masks for each instance.

In YOLOv8, segmentation is performed using an instance segmentation approach, which partitions each object into pixel-level regions, and a multi-head structure is employed to output a corresponding mask for each detected object. The loss function adopts an integrated form that includes multiple components related to position, size, class, and mask and is defined as follows:

$$\mathcal{L}_{\text{total}} = \lambda_{\text{box}} \cdot \mathcal{L}_{\text{CIoU}} + \lambda_{\text{cls}} \cdot \mathcal{L}_{\text{cls}} + \lambda_{\text{mask}} \cdot \mathcal{L}_{\text{mask}} \quad (2)$$

where $\mathcal{L}_{\text{total}}$ is the total loss of the model, $\mathcal{L}_{\text{CIoU}}$ is the CIoU (Complete Intersection over Union) loss for bounding box regression, $\mathcal{L}_{\text{cls}}$ is the classification loss, and $\mathcal{L}_{\text{mask}}$ is the mask loss. By appropriately setting the weights λ_{box}, λ_{cls}, and λ_{mask} for each loss component, the overall training balance can be effectively adjusted (Zheng *et al*., 2019).

In this study, many bridges exhibit complex shapes and often have unclear boundaries with debris caused by earthquakes, tsunamis, and other disasters. Furthermore, the shape dimensions, such as bridge length and width, may vary depending on the damage status. To accurately capture such shape information and analyse its relationship with the damage status, it is necessary to perform pixel-level region extraction rather than relying solely on bounding boxes. Therefore, this study adopts the YOLOv8 segmentation model to enable the quantification of bridge attributes and to improve the accuracy of damage assessment.

2.2 Bridge shape calculation model

Based on the segmentation results of bridges obtained using YOLOv8, the shape dimensions of bridges are determined through three processes: (*i*) redefinition of segmentation areas, (*ii*) reclassification of segmentation area classes, and (*iii*) calculation of bridge shape dimensions.

(*i*) Redefinition of segmentation areas

The segmentation results may contain connection regions linking adjacent bridges or duplicate detections of the same bridge, and it is necessary to remove these artefacts. To address this, a binary mask of the same size as the input image is created, assigning a value of 1 to pixels corresponding to bridge regions and 0 to the background. A morphological operation is applied to the mask, specifically performing two rounds of opening operations to properly separate binarised regions and smooth their shapes. Subsequently, external contours are extracted from the morphologically processed mask, and spatially adjacent or partially overlapping regions are merged into single regions. This process enables the removal of noise and the elimination of duplicate detections from the segmentation results.

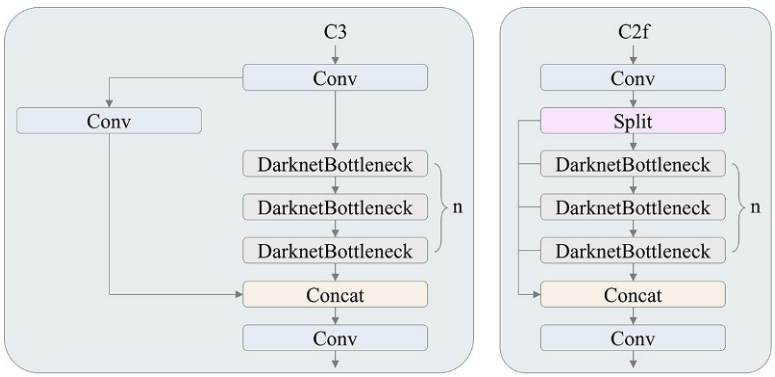

Figure 2. The structures of the C2f and C3 modules are illustrated (created with reference to Wang *et al*. (2024)). Compared with the C3 module used in YOLOv5, the C2f module adopted in YOLOv8 achieves greater simplicity and computational efficiency

(*ii*) Reclassification of segmentation area classes

Based on the damage status of pixels contained within each redefined segmentation area, a new class label is assigned. In the segmentation results, there are cases where, due to the influence of debris and other factors, only a portion of a bridge is classified as damaged while the rest is classified as undamaged. In such cases, inconsistent class assignments occur within a single bridge, which reduces the reliability of the classification. Therefore, in this study, undamaged bridges are assigned an ID of 0, and damaged bridges are assigned an ID of 1. The class label for each segmentation area is determined based on a thresholding process using the ratio r between the number of pixels with class ID 1, N_1, and the total number of pixels in the area, N_{total}, calculated as:

$$r = \frac{N_1}{N_{\text{total}}} \tag{3}$$

The class label C is assigned according to the following rule:

$$C = \begin{cases} 1 \text{ if } r \geq \tau \\ 0 \text{ otherwise} \end{cases} \tag{4}$$

where τ is a predefined threshold, set to 0.1 in this study. This threshold is set to a low value to prioritise detection sensitivity. For the purpose of rapid disaster assessment, this value was chosen to ensure that a bridge is classified as damaged even if only a small portion is identified as such, thereby minimising the risk of overlooking compromised structures. This process also serves to suppress local misclassifications, ensuring that a single class is consistently assigned to each redefined segmentation area.

(*iii*) Calculation of bridge shape dimensions

For each bridge region whose class and fine-grained shape have been redefined, shape dimensions are calculated. A minimum bounding rectangle is applied to the bridge region, and among the four edges of the resulting rectangle, the distance between the centre points of the two longer edges is defined as the bridge length h, while the distance between the centre points of the two shorter edges is defined as the bridge width w.

3. Bridge segmentation

3.1 Model configuration and training parameters

In this study, a pretrained YOLOv8x segmentation model from the YOLOv8 series is adopted for the purpose of bridge segmentation and damage status classification. The optimiser is Adam, with an initial learning rate of 0.002, a batch size of 16, and 2000 training epochs. All training images are resized to 600×600 pixels before being input to the model. Other parameters and data augmentation settings follow those of the original model. Under these settings, training was conducted, and the loss function reached its minimum at epoch 1968. The model weights at that point were adopted for validation.

3.2 Datasets

The data used for training and validation consist of bridge images collected from aerial imagery captured shortly after the 2024 Noto Peninsula Earthquake and the 2011 off the Pacific coast of Tohoku Earthquake (Figure 3). In this study, the classification of 'damaged' primarily refers to bridges that have suffered severe visual alterations, such as structural collapse, washout, or significant debris accumulation on the deck, as observed in the aftermath of the earthquakes and tsunamis. The training data include images containing both damaged and undamaged bridges from both earthquakes. The images from the 2024 Noto Peninsula Earthquake are 480×640 pixels, and those from the 2011 off the Pacific coast of Tohoku Earthquake are 480×600 pixels. Since the number of damaged bridges in the entire training dataset is limited, data augmentation was applied specifically to images containing damaged bridges from the 2011 off the Pacific coast of Tohoku Earthquake, which include a relatively higher number of damaged instances. For each original image, five augmentation operations were performed individually, including horizontal flipping, vertical flipping, and rotations of 90°, 180°, and 270°, resulting in five additional images. Including the original, each image yielded a total of six images. This procedure was applied to all such damaged-bridge images, thereby increasing the amount of training data by a factor of six. For validation, images captured in different regions during the 2011 off the Pacific coast of Tohoku Earthquake, which were not used in training, were employed. The validation images are 480×600 pixels, but all images, both for training and validation – are resized to 600×600 pixels before being input to the model. The composition of the dataset used in this study is summarised in Table 1.

3.3 Results of segmentation

This section presents the segmentation results for damaged and undamaged bridges. It is important to note that the visualisations in Figures 4–8 show the raw model output, which can include multiple, overlapping detections for a single object, as illustrated in Figure 4(c). In contrast, all subsequent quantitative analyses, including the accuracy evaluation in Section 3.4 and the shape estimation in Section 4, are based on the final, processed results after applying the reclassification method from Section 2.2. In these figures, red regions denote bridges classified as damaged, while blue regions denote those classified as undamaged.

Figure 4 shows the segmentation results for damaged bridges, while Figure 5 shows those for undamaged bridges. In all cases, regardless of bridge length, width, or damage status, the bridges were generally segmented accurately. Among Figures 4(a)–4(c), Figure 4(b) is particularly notable in that, despite the presence of many houses, vehicles, and debris accumulated on the bridge deck, the entire bridge was correctly segmented as damaged.

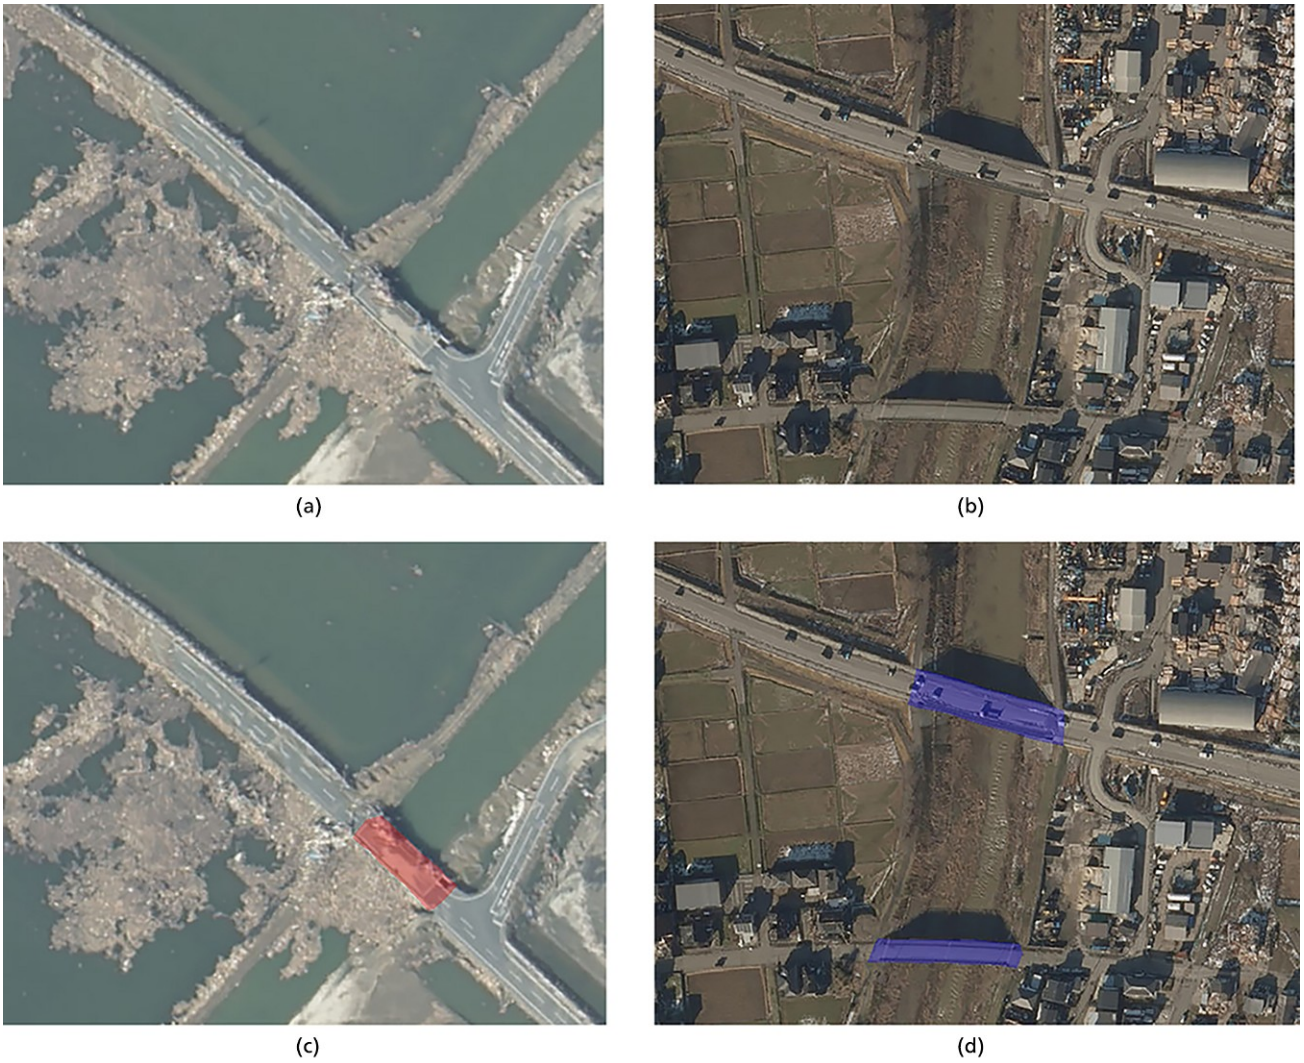

Figure 3. Examples of bridge data used for training and validation. The top row shows the original aerial images, and the bottom row presents the corresponding ground truth. In the annotations, red indicates damaged bridges and blue indicates undamaged bridges

Table 1. Summary of datasets used for training and validation

Purpose	Disaster name	Number of images	Undamaged bridge	Damaged bridge
Training	The 2024 Noto Peninsula Earthquake	2692	5437	34
	The 2011 off the Pacific Coast of Tohoku Earthquake	587*	272	468
Validation	The 2011 off the Pacific Coast of Tohoku Earthquake	67	34	41

The table lists the number of damaged and undamaged bridge images contained in aerial imagery captured during the 2024 Noto Peninsula Earthquake and the 2011 off the Pacific Coast of Tohoku Earthquake
*A total of 73 images of damaged bridges from the 2011 off the Pacific coast of Tohoku Earthquake were augmented sixfold. The asterisked values indicate the number of images after augmentation

Figure 4(c) illustrates a case where the model produced highly overlapping detections for a single bridge with conflicting classifications. The reclassification process described in Section 2.2 is applied to resolve such ambiguities. In this specific example, the union of all detected pixels constituted a total area of 16 844 pixels. Within this area, 16 773 pixels were classified as 'undamaged' and 16 799 pixels were classified as 'damaged', confirming a near-complete overlap. The damage ratio r is calculated using the 'damaged' pixel count, yielding approximately 0.9973. As this value significantly exceeds the threshold

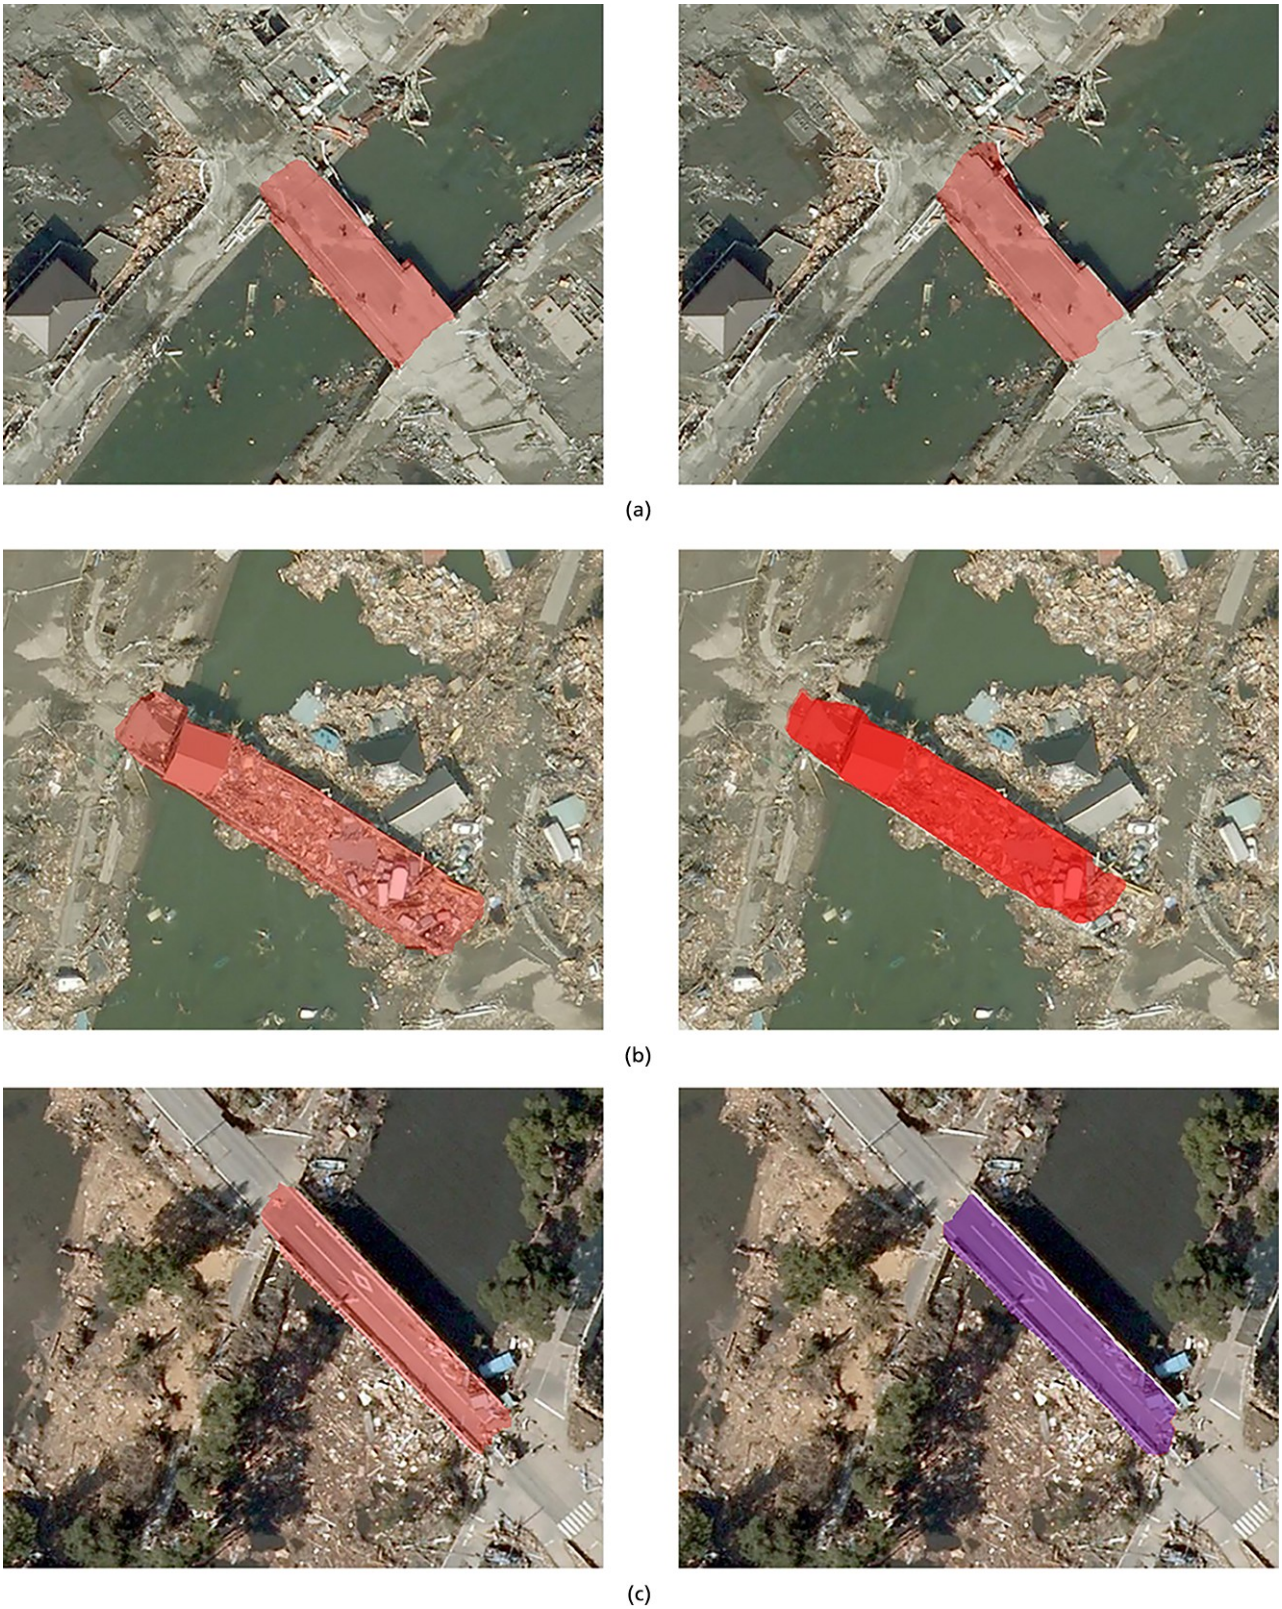

Figure 4. Segmentation results for damaged bridges. Even in the presence of debris or structural damage, damaged bridges were generally segmented appropriately

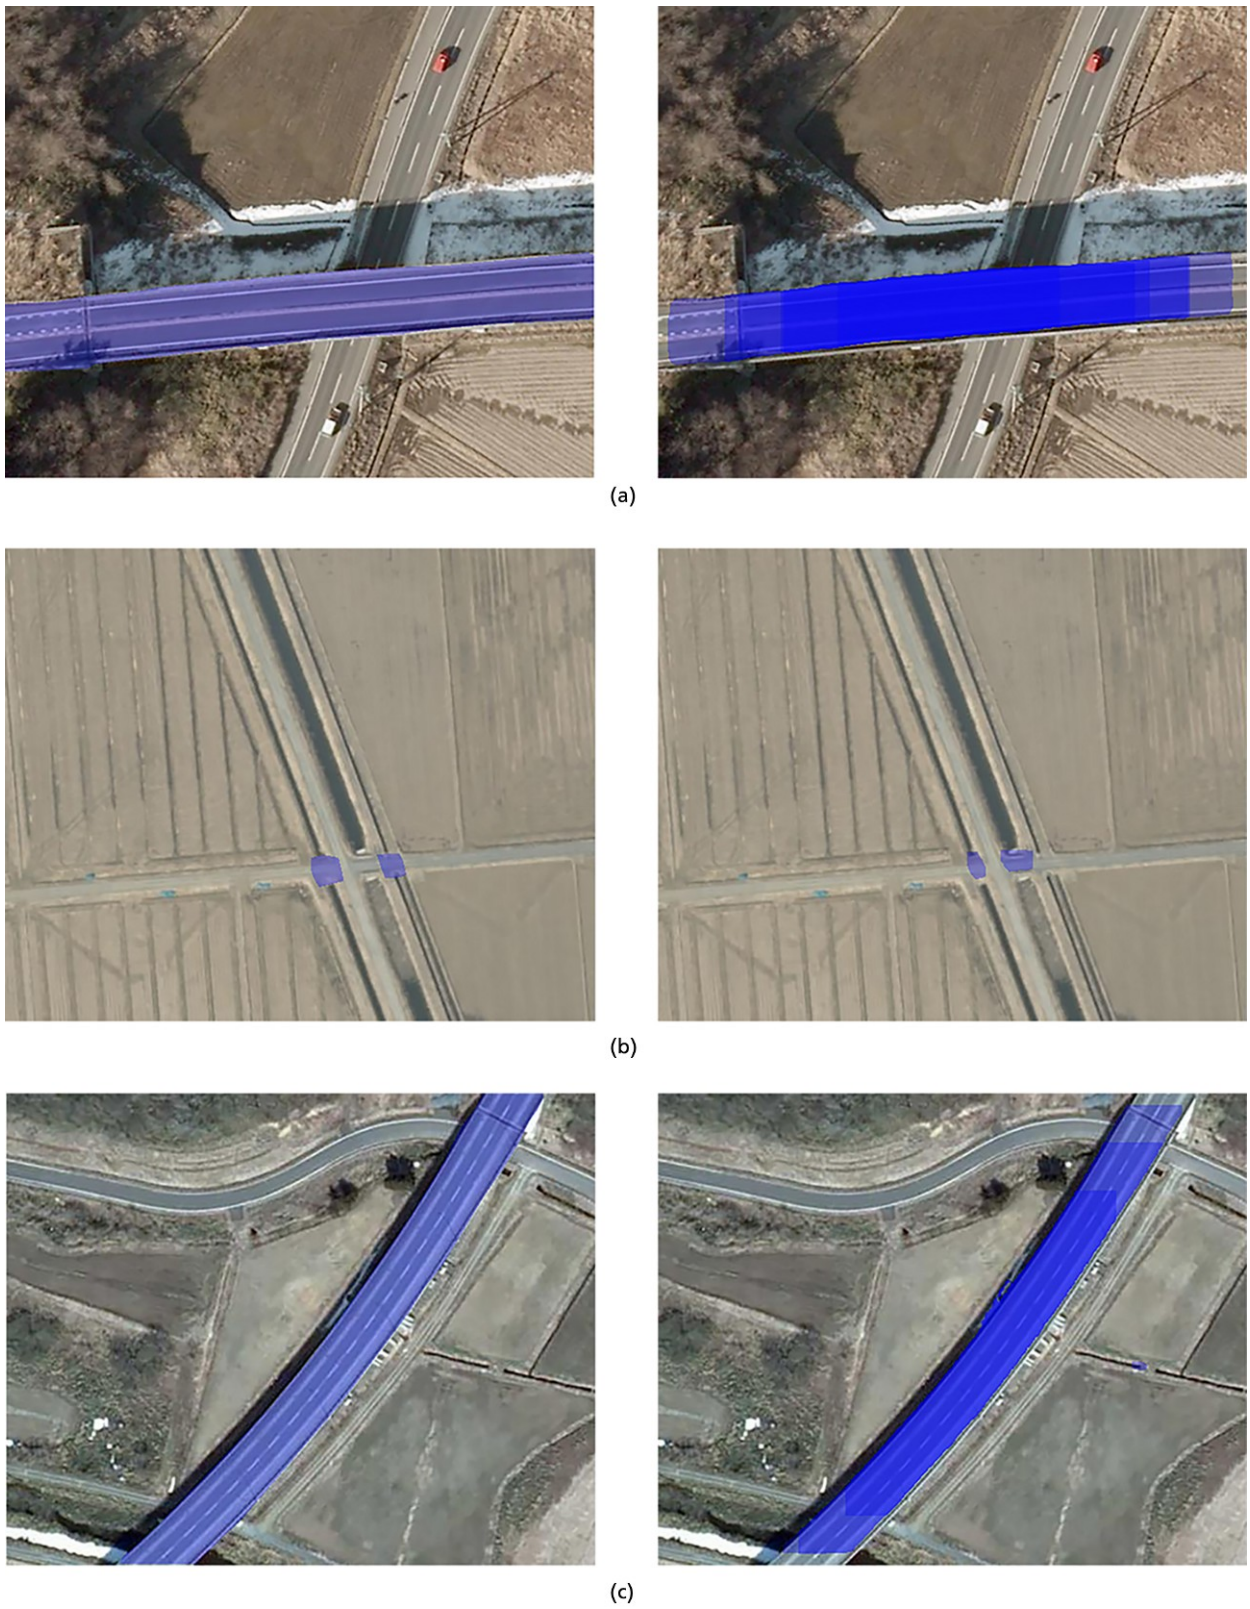

Figure 5. Segmentation results for undamaged bridges. Undamaged bridges were accurately detected regardless of their size

$\tau = 0.1$, the bridge region was correctly unified and assigned a final 'damaged' classification. In addition to the results for damaged bridges, Figure 5 presents examples of undamaged bridge segmentation. Figures 5(a) and 5(c) show long bridges, while Figure 5(b) shows a short bridge. In all cases, the entire bridge was appropriately segmented, indicating that the model demonstrates high segmentation performance regardless of bridge size.

Figures 6 and 7 show examples in which partial detection or missed detection occurred for damaged bridges. In particular, as shown in Figure 6, when a large amount of debris is present on the bridge deck or on the water surface near the bridge, or as shown in Figure 7, when part of the bridge has collapsed or the deck is obscured due to flooding, a decline in segmentation performance was observed. However, in cases where debris accumulation or flooding is limited, as in Figures 6(a), 6(b), 7(b), and 7(c), the damaged bridges can still be successfully detected. These observations suggest that the proposed method is effective for estimating the locations and number of damaged bridges.

Finally, Figure 8 presents examples of false detection for undamaged bridges. In Figures 6(a) and 6(b), not only are undamaged bridges mistakenly detected as damaged, but surrounding areas that are not part of a bridge are also misclassified as bridges. Similarly, in Figures 8(c), a road segment that is not actually a bridge is incorrectly identified as a bridge. These false detections may be attributed to the presence of bridges in the training data whose shapes and surrounding environments resemble those of ordinary roads.

Despite some cases of reduced accuracy and false detection, the results show that, for most bridges, either the entire structure or a substantial portion was appropriately detected regardless of damage status, suggesting the effectiveness of the proposed method.

3.4 Evaluation of segmentation results

The segmentation accuracy for bridges was evaluated based on accuracy, precision, recall, and F1-score, as defined in Equations 5–8. In this evaluation, the 'Damaged bridge' class (Class 1) was defined as the positive class, and the 'Undamaged bridge' class (Class 0) as the negative class. Based on these definitions, the terms TP, TN, FP, and FN in the equations denote true positive, true negative, false positive, and false negative, respectively. The values of these metrics for each class are shown in Table 2.

5. $$\text{Accuracy} = \frac{\text{TP} + \text{TN}}{\text{TP} + \text{TN} + \text{FP} + \text{FN}}$$

6. $$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}}$$

7. $$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}}$$

8. $$\text{F1} - \text{score} = \frac{2 \times (\text{Precision} \times \text{Recall})}{\text{Precision} + \text{Recall}}$$

For Class 0, which corresponds to undamaged bridges, the accuracy is 0.9748, indicating a low overall classification error. The high accuracy is a consequence of the pixel imbalance between the foreground representing bridges and the larger background area. This value is inflated by the vast number of true negative pixels from the correctly classified background. In contrast, precision and recall are not influenced by true negatives, thus providing a more focused assessment of performance on the foreground bridge classes. The recall is 0.5126 and the precision is 0.4706, meaning that more than half of the regions classified as undamaged bridges are actually undamaged. However, the relatively low precision suggests that some misclassifications are present. Since the recall value exceeds the precision value, it can be inferred that most undamaged bridges are successfully detected, although some classification errors may still be included. For Class 1, corresponding to damaged bridges, the recall is relatively low at 0.3558, indicating that some damaged bridges were not successfully detected. However, the precision is comparatively high at 0.5688, suggesting that most regions classified as damaged bridges are indeed damaged. This implies that false detections are relatively limited in the damaged class and that the detected information maintains a certain level of reliability. As discussed in the previous section, segmentation becomes difficult in areas where debris has accumulated on the bridge deck, flooding is present, or parts of the bridge have collapsed. Nevertheless, in portions where the bridge shape remains clearly visible, damaged bridges are detected appropriately. This indicates that the model has the potential to serve as an effective tool for prioritising the detection of affected areas during a disaster. Based on these results, the model demonstrates the capability to detect and classify both damaged and undamaged bridge regions with at least a moderate level of accuracy. In situations where it is necessary to assess the condition of a large number of bridges in a short time during a disaster, the proposed method may serve as an effective means of rapid and efficient visual support.

4. Estimation of bridge shape

Using the method described in Section 2.2, the bridge length and width were estimated for both damaged and undamaged bridges based on the ground truth and segmentation results from the validation data. The comparison results are shown in

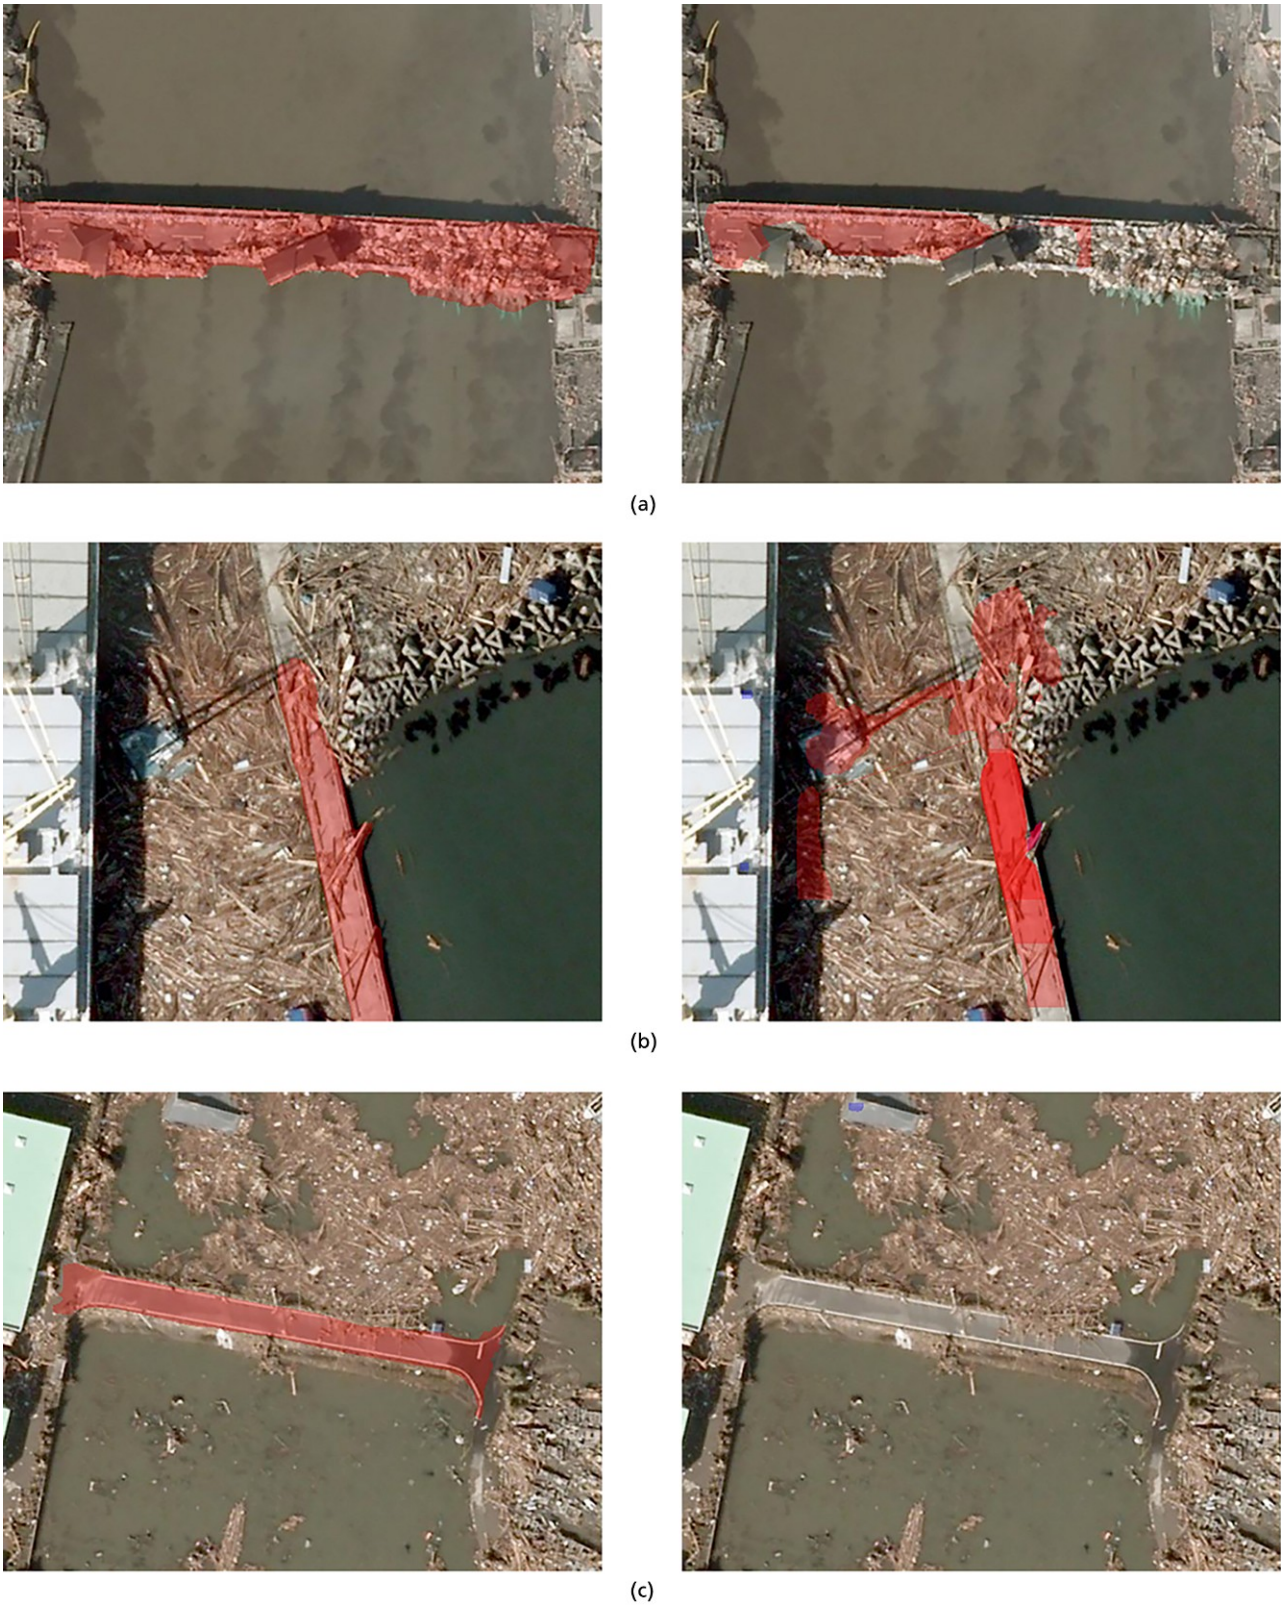

Figure 6. Segmentation results for damaged bridges were partially detected due to debris accumulation. In cases where large amounts of debris extend beyond the bridge deck or float on the surrounding water surface, parts of the bridge may not be detected

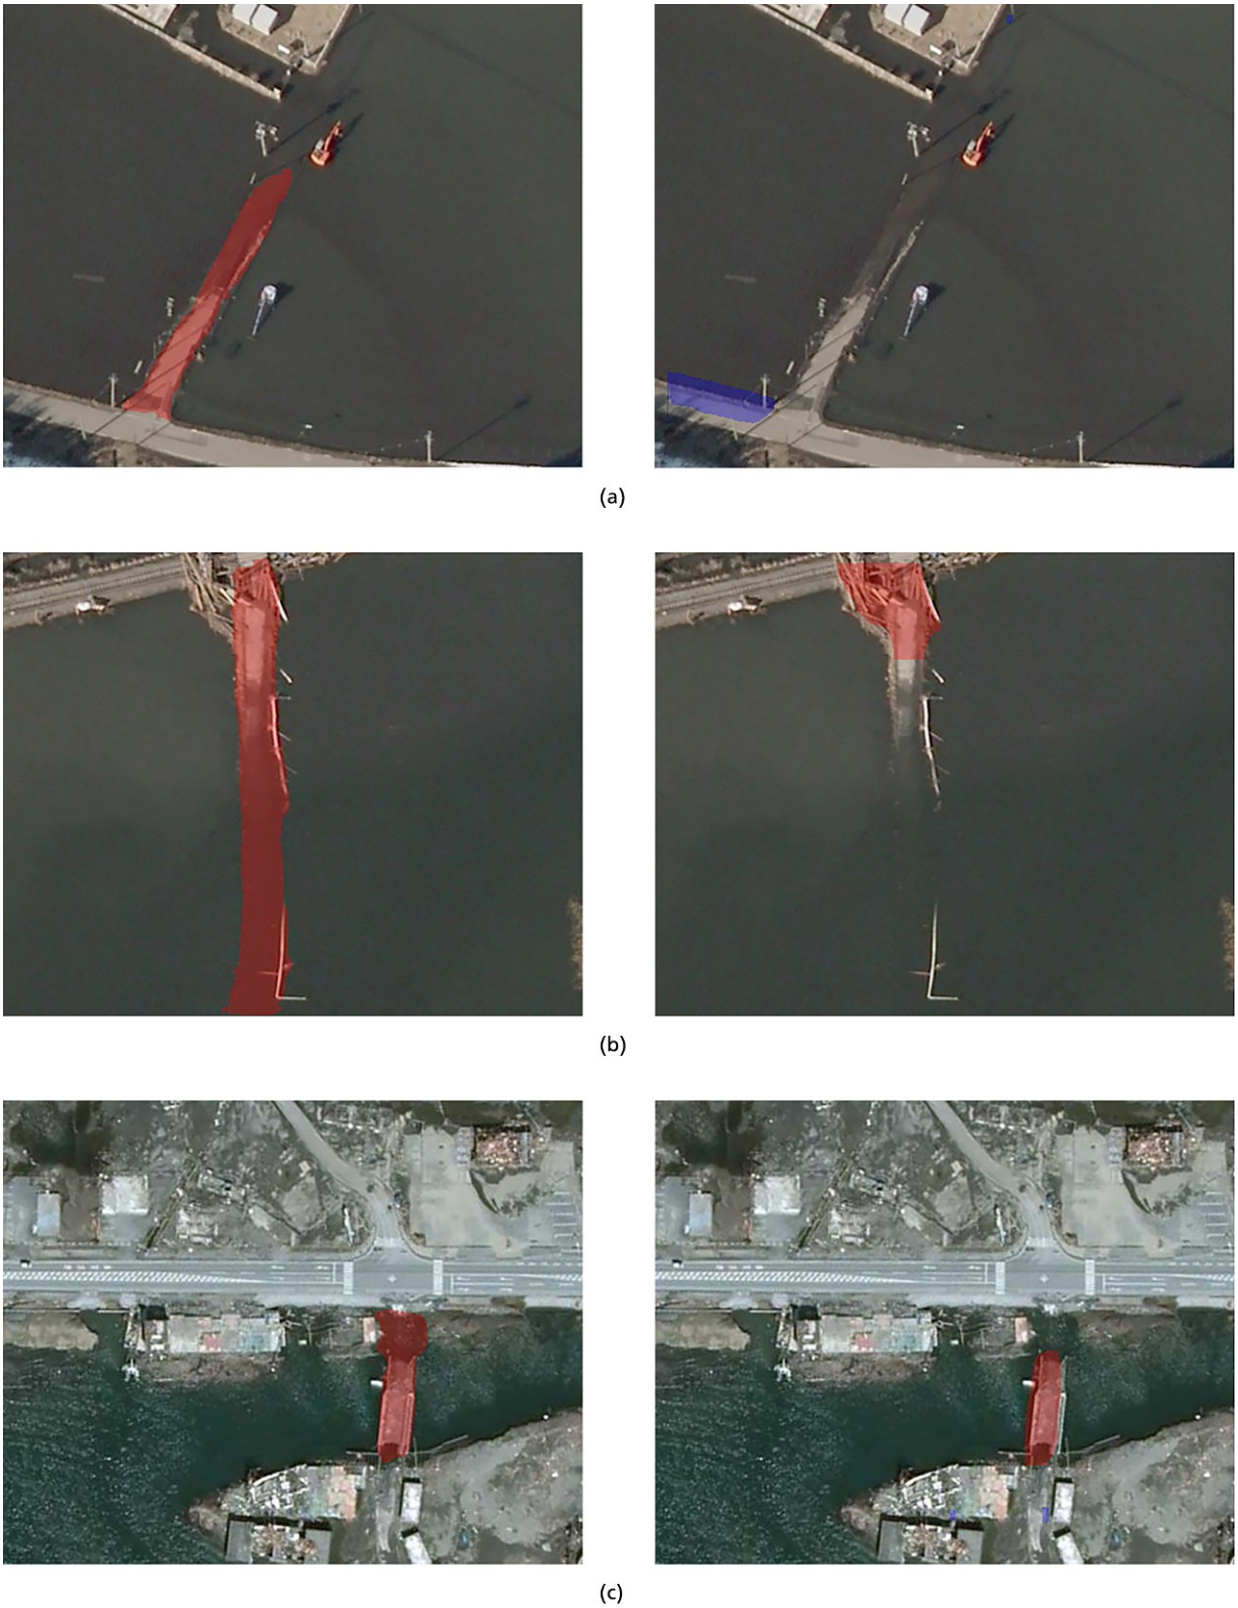

Figure 7. Segmentation results for damaged bridges with detection failures caused by flooding or structural collapse. When parts of a bridge become unobservable due to inundation or the loss of structural components, those regions may not be detected

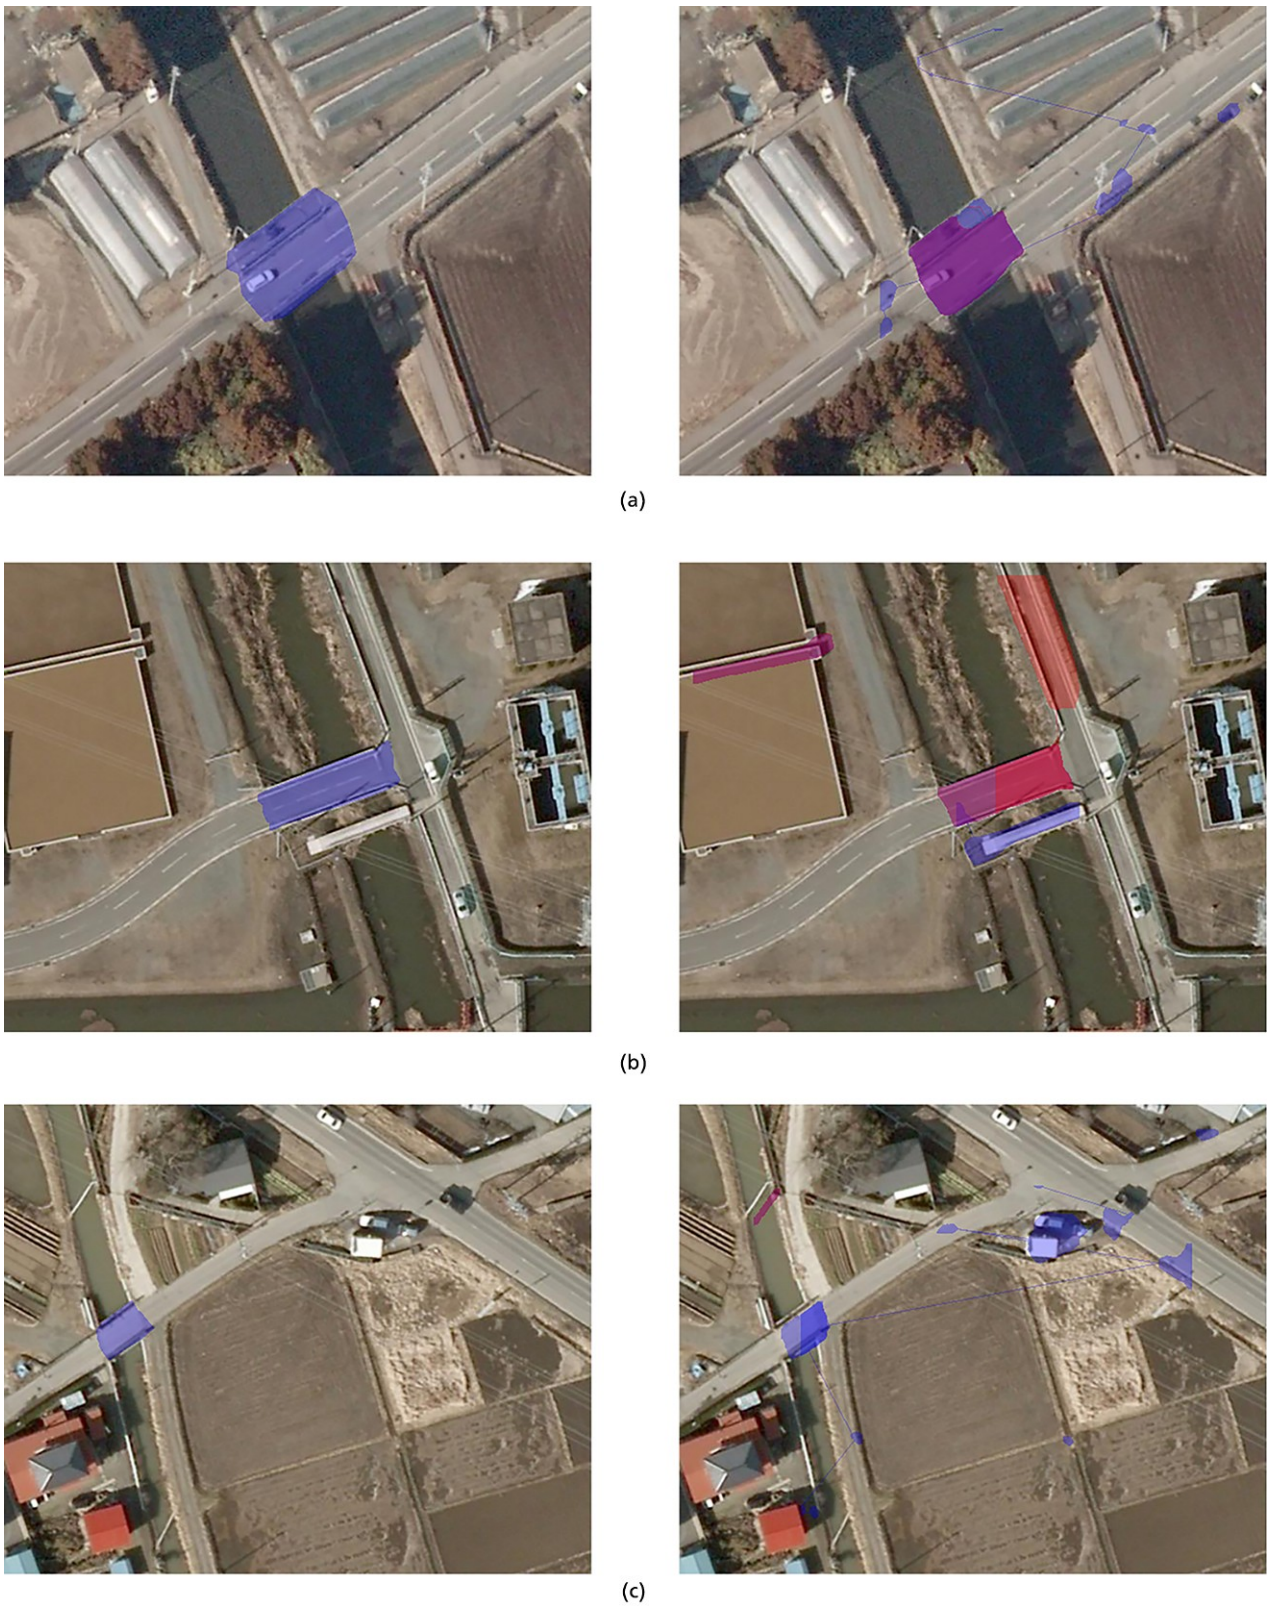

Figure 8. Examples of false detections involving undamaged bridges. These include cases where undamaged bridges were misclassified as damaged, or road segments were incorrectly detected as undamaged bridges

Table 2. Detection and classification accuracy for each class

Class	Accuracy	Precision	Recall	F1-score
Class 0 (Undamaged bridge)	0.9748	0.4706	0.5126	0.4907
Class 1 (Damaged bridge)	0.9596	0.5688	0.3558	0.4378

The YOLOv8 segmentation model was evaluated for undamaged bridges (Class 0) and damaged bridges (Class 1). For each class, accuracy, precision, recall, and F1-score were calculated to quantitatively assess the model's performance

Figures 9(a) and 9(b), and the visualised relationship between bridge length and width for damaged and undamaged bridges is presented in Figures 9(c) and 9(d), respectively.

In the width estimation results for undamaged bridges shown in Figure 9(a), there is little difference between the ground truth and the estimated values, indicating that the bridge widths were estimated with high accuracy in most cases. As described in the previous chapter, since the bridges were correctly detected in their entirety, the estimated values closely match the ground truth, resulting in low estimation error.

In contrast, the width estimation results for damaged bridges (Figure 9(b)) show slightly greater variation compared with those for undamaged bridges. This can be attributed to the partial detection of bridge regions due to factors such as debris accumulation and inundation, leading to an underestimation of the bridge width. However, in many cases, the detection range in the width direction remains largely preserved, and thus the impact on

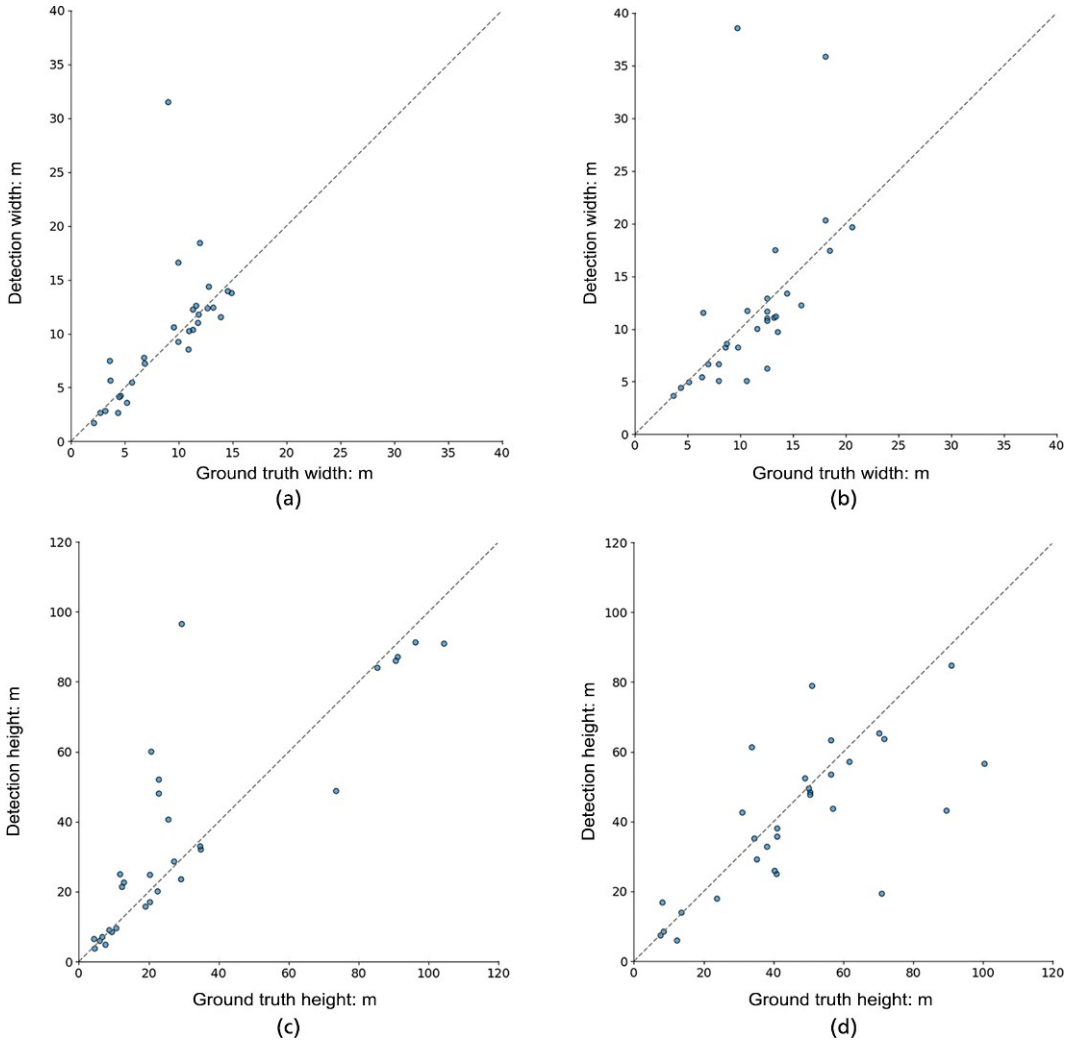

Figure 9. Scatter plots showing the relationship between ground truth and estimated values of bridge length and width for damaged and undamaged bridges. For undamaged bridges, the estimated dimensions closely match the ground truth, indicating high estimation accuracy. In contrast, for damaged bridges, the estimated bridge length tends to be underestimated, suggesting that reduced detection coverage due to damage affects the accuracy of dimensional estimation: (a) width of undamaged bridges; (b) width of damaged bridges; (c) length of undamaged bridges; (d) length of damaged bridges

dimensional estimation is limited. As a result, the width of damaged bridges is still estimated with high accuracy in many instances.

In the bridge length estimation for undamaged bridges (Figure 9(c)), there is little difference between the ground truth and the estimated values, and the estimation error is small for most bridges. However, for a subset of bridges with actual lengths around 20 m, the estimated lengths are occasionally overestimated to be 40–60 m. This overestimation is considered to be due to the similarity in shape between the bridges and adjacent roads, which caused portions of the road to be incorrectly detected as part of the bridge.

In the bridge length estimation for damaged bridges (Figure 9(d)), as noted in the previous chapter, there are several cases in which parts of a bridge are covered by debris, submerged, or washed away due to the damage. In such cases, those regions are not detected as part of the bridge, resulting in an underestimation of the actual bridge length. Figure 9(d) clearly shows that this contributes significantly to the underestimation of bridge length.

To evaluate the reproducibility of dimensional relationships, an analysis was conducted focusing on the correlation between bridge length and width. As shown in Figure 10, for undamaged bridges, the distribution trends of the ground truth and the detection results were consistent, and no significant differences were observed in the slope and intercept of the regression lines. This indicates not only that the entire bridges were accurately detected but also that the inherent structural relationship between bridge length and width was successfully reproduced in the estimation results.

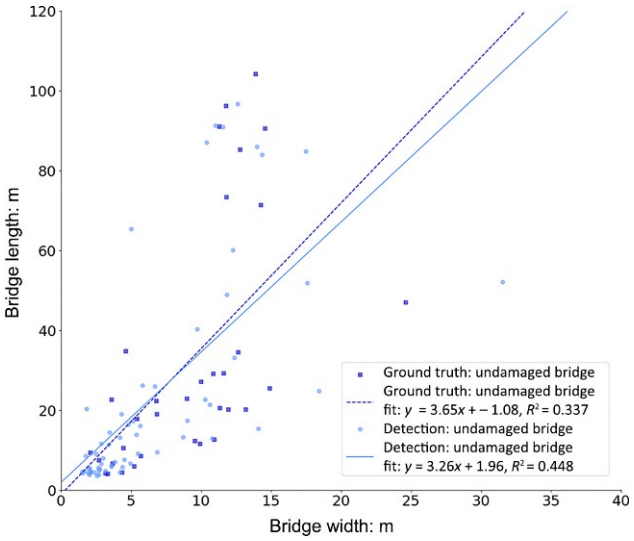

Figure 10. Scatter plot showing the relationship between bridge length and width for undamaged bridges, based on both ground truth and detection results. The two sets show a high degree of correlation, and the similarity in regression slope and intercept indicates that the structural dimensional relationship is accurately preserved

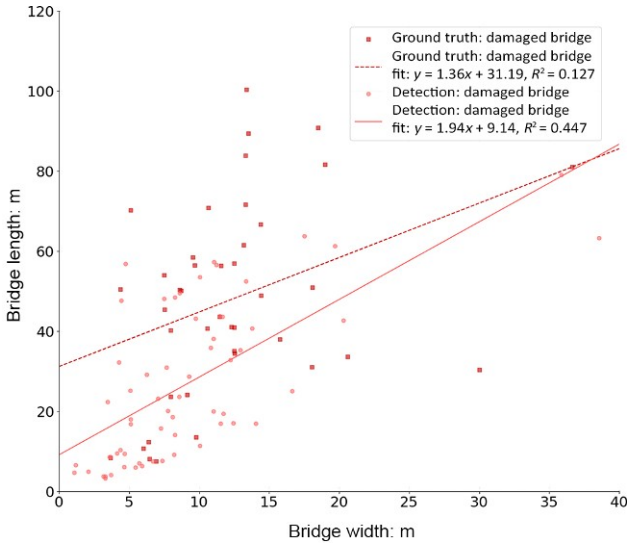

Figure 11. Scatter plot showing the relationship between bridge length and width for damaged bridges, based on both ground truth and detection results. The regression line for the detection results tends to have a steeper slope and smaller intercept than that of the ground truth, suggesting that partial detection due to damage alters the dimensional relationship. This implies that damage can affect the relative scale of bridge dimensions

In contrast, for damaged bridges, Figure 11 shows that the regression lines derived from the detection results tended to have larger slopes and smaller intercepts compared with those from the ground truth. This trend is considered to be due to partial detection of bridges caused by damage, especially leading to the underestimation of dimensions in the length direction. In other words, damage not only caused overall dimensional shrinkage but also altered the relative scaling between bridge length and width, indicating that the impact of damage was reflected in the interdimensional relationship.

These findings reveal that the shrinkage tendency caused by damage is not merely an estimation error but rather a dimensional characteristic specific to damaged bridges. The goal of this analysis was to evaluate the utility of such dimensional changes as a damage indicator, where the required accuracy is defined by the ability to capture this statistical trend rather than by absolute measurements. The capability to identify this change in the dimensional pattern demonstrates that the method meets the necessary accuracy to detect the presence of damage based on shape analysis.

5. Conclusion

In this study, a method was developed to detect bridges from aerial imagery using deep learning, classify them according to the presence or absence of damage, and estimate their shape dimensions. Validation results confirmed that for undamaged bridges, both bridge length and width could be estimated with high accuracy,

showing a high degree of consistency with the ground truth. In contrast, for damaged bridges, the detection of the entire bridge was often difficult due to factors such as debris accumulation, inundation, and bridge collapse, leading to a tendency for dimensions, particularly bridge length, to be underestimated. Furthermore, focusing on the relationship between bridge length and width, it was found that while the distribution trends between the ground truth and detection results were consistent for undamaged bridges, damaged bridges exhibited biases in the slope and intercept of the regression lines. These biases reflect the dimensional shrinkage tendency resulting from partial missed detections. This dimensional shrinkage tendency observed in damaged bridges is considered not merely as an estimation error but as a structural feature associated with damage, suggesting its potential utility as an auxiliary indicator for detecting the presence of damage.

In the future, we aim to develop models capable of accommodating a wider variety of disaster types and bridge structures, while also utilising satellite imagery with geospatial information to estimate the distribution of damaged bridges over large areas. In addition, we plan to establish advanced damage estimation methods by integrating damage information from the surrounding environments of the target bridges.

Funding

This work was supported by the Council for Science, Technology and Innovation (CSTI), Cross-Ministerial Strategic Innovation Promotion Program (SIP), and the 3rd period of SIP 'Smart Infrastructure Management System' (Grant No. JPJ012187, Funding agency: Public Works Research Institute).

REFERENCES

Ahmad K, Pogorelov K, Riegler M *et al.* (2019) Automatic detection of passable roads after floods in remote sensed and social media data. *Signal Processing: Image Communication* **74**: 110–118, 10.1016/j.image.2019.02.002.

Alisjahbana I, Li J, Strong B and Zhang Y (2024) A two-step deep-learning model for multi-disaster building damage segmentation and classification using satellite imagery. *arXiv Preprint arXiv:2405.04800*, 10.48550/arXiv.2405.04800.

Bai Y, Hu J, Su J *et al.* (2020) Pyramid pooling module-based semi-Siamese network: a benchmark model for assessing building damage from xBD satellite imagery datasets. *Remote Sensing* **12(24)**: 4055, 10.3390/rs12244055.

Bolya D, Zhou C, Xiao F and Lee YJ (2019) YOLACT: real-time instance segmentation. *arXiv Preprint arXiv:1904.02689*, 10.48550/arXiv.1904.02689.

Ge Z, Liu S, Wang F, Li Z and Sun J (2021) YOLOX: exceeding YOLO series in 2021. *arXiv Preprint arXiv:2107.08430*.

Gholami S, Robinson C, Ortiz A, Yang S, Margutti J, Birge C, Dodhia R and Lavista Ferres J (2022) On the deployment of post-disaster building damage assessment tools using satellite imagery: a deep learning approach. *In Proc. 2022 IEEE Intl. Conf. on Data Mining Workshops (ICDMW)*, Orlando, FL, USA, 1029–1036, 10.1109/ICDMW58026.2022.00134.

Kopiika N, Karavias A, Krassakis P *et al.* (2025) Rapid post-disaster infrastructure damage characterisation using remote sensing and deep learning technologies: a tiered approach. *Automation in Construction* **170**: 105955, 10.1016/j.autcon.2024.105955.

Kubo S, Yamane T and Chun PJ (2022) Study on accuracy improvement of slope failure region detection using mask R-CNN with augmentation method. *Sensors (Basel, Switzerland)* **22(17)**: 6412, 10.3390/s22176412.

Lang F, Zhu Y, Zhao J *et al.* (2024) Flood mapping of synthetic aperture radar (SAR) imagery based on semi-automatic thresholding and change detection. *Remote Sensing* **16(15)**: 2763, 10.3390/rs16152763.

Liu W, Maruyama Y and Yamazaki F (2021) Detection of collapsed bridges from multi-temporal SAR intensity images. *Remote Sensing* **13(17)**: 3508, 10.3390/rs13173508.

Miyamoto T and Yamamoto Y (2020) Using multimodal learning model for earthquake damage detection based on optical satellite imagery and structural attributes. *In IGARSS 2020 – 2020 IEEE International Geoscience and Remote Sensing Symposium*, 6623–6626. 10.1109/IGARSS39084.2020.9324464.

Opara JN, Moriwaki R and Chun PJ (2024) Delineating landslide and debris flow detection in Japan through aerial photography: a YOLO v8 approach to disaster management. *Intelligence, Informatics and Infrastructure* **5(1)**: 111–123, 10.11532/jsceiiai.5.1_111.

Shukla R, Pabbisetty SK, Jayanthi S and Janardhanan K (2024) Evaluating the damage of collapsed bridges using remote sensing technologies: case study: Baltimore's Francis Scott Key Bridge. *Eco Cities* **5(2)**: 1181, 10.59400/be.v2i2.1811.

Sun S, Yang J, Chen Z, Li J and Sun R (2024) Tibia-YOLO: an assisted detection system combined with industrial CT equipment for leg diseases in broilers. *Applied Sciences* **14(3)**: 1005, 10.3390/app14031005.

Ultralytics (2025) YOLOv8: Next-Generation Object Detection and Segmentation, See https://github.com/ultralytics/ultralytics (accessed 18/10/2025).

Wang Y, Zhang K, Wang L and Wu L (2024) An improved YOLOv8 algorithm for rail surface defect detection. *IEEE Access* **12**: 44984–44997, 10.1109/ACCESS.2024.3380009.

Zhao K, Liu J, Wang Q, Wu X and Tu J (2022) Road damage detection from post-disaster high-resolution remote sensing images based on TLD framework. *IEEE Access* **10**: 43552–43561, 10.1109/ACCESS.2022.3169031.

Zheng Z, Wang P, Liu W *et al.* (2019) Distance-IoU loss: faster and better learning for bounding box regression. *In Proceedings of the AAAI Conference on Artificial Intelligence* **34(7)**: 12993–13000.