

Article

A Structural–Visual Integrated Evaluation Framework for Seismic Damage Assessment of RC Double-Column Piers

Zhibing Yu ¹, Zixin Xu ², Chao Zhang ^{1,2,*}, Yongjun Zhou ¹, Xiaojun Hou ¹, Longqi Zhang ¹, Wang Liao ¹ and Haiou Li ²

¹ Sichuan Transportation Construction Group Co., Ltd., Chengdu 610041, China

² School of Civil Engineering, Chongqing Jiaotong University, Chongqing 400074, China

* Correspondence: zcbridge@163.com; Tel.: +86-18583760629

Abstract

The post-earthquake damage state of reinforced concrete (RC) double-column piers directly affects bridge traffic capacity and emergency response efficiency. To improve the interpretability of damage assessment, this study proposes a Structural–Visual Integrated Evaluation (SVIE) framework that combines structural response analysis with image-based damage evidence. Structural responses from quasi-static tests are used to define four baseline damage states: intact-to-slight, moderate, severe, and critical damage. An improved DeepLabv3+ model is then applied to 315 global-scene images for end-to-end semantic segmentation of background, concrete spalling, and reinforcement exposure. The extracted visual evidence is used to verify its consistency with the baseline structural states. On the test set, the model effectively identified concrete spalling regions, achieving an IoU, F1-score, Precision, and Recall of 78.86%, 88.18%, 89.93%, and 86.49%, respectively. For reinforcement exposure, although IoU and Recall were relatively low because of sample scarcity and small-target characteristics, Precision reached 70.40%, indicating that detected regions can provide supplementary evidence for severe local damage. The consistency analysis showed that the morphology of visual damage was generally compatible with the progression of structural damage states. The results provide a laboratory-based proof of concept for a mechanically grounded and visually interpretable framework for rapid post-earthquake assessment of RC double-column piers.

Keywords: reinforced concrete (RC) double-column piers; post-earthquake assessment; structural–visual integrated evaluation; semantic segmentation; visual evidence of damage

Academic Editor: Darius Bačinskas

Received: 26 May 2026

Revised: 16 July 2026

Accepted: 21 July 2026

Published: 4 August 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

As an important component of modern transportation networks, bridges serve as lifeline infrastructure for post-earthquake rescue, evacuation, emergency supply transport, and regional traffic recovery. Their post-earthquake functional state directly affects the efficiency of emergency response and the capacity for socioeconomic recovery [1,2]. As key load-resisting members that connect the bridge superstructure to the foundation, bridge piers transfer vertical loads and resist horizontal seismic actions. Their damage state, therefore, directly determines the post-earthquake traffic capacity of bridges and the priority of repair actions. Reinforced concrete (RC) double-column piers have been widely used in medium- and short-span girder bridges and urban viaducts

because of their well-defined structural configuration, favorable stability, and mature construction techniques [3]. Compared with single-column piers, RC double-column piers exhibit more pronounced spatially coupled seismic responses because of the interaction among the columns, cap beam, transverse tie beams, and joint regions [4]. Their damage may include failure in the plastic hinge regions of columns, local cracking in joint regions, concrete cover spalling, crushing of core concrete, and exposure of reinforcing bars. Therefore, rapid and objective assessment of the post-earthquake damage state of RC double-column piers is of great significance for maintaining bridge traffic capacity and ensuring the continuity of emergency rescue routes.

Conventional post-earthquake bridge assessment primarily relies on manual field inspection, experience-based judgment, and limited structural response data. In engineering practice, bridge assessment is commonly embedded within structured inspection and risk-classification procedures rather than being conducted as an isolated damage-recognition task. For example, the Italian Guidelines for existing bridges adopt a standardized multilevel and multi-hazard framework that links bridge census, inspection, risk classification, safety verification, and monitoring [5]. Comparable regulatory systems in the United States, the United Kingdom, and Japan also rely on standardized inspection, condition documentation, and subsequent detailed evaluation, although their organizational procedures and assessment emphases differ [6–8]. Within such frameworks, the proposed framework is intended to support rapid evidence extraction and preliminary prioritization rather than replace formal inspection ratings or safety verification. Manual visual inspection is intuitive and flexible; however, in large-scale earthquake disaster scenarios, it is inefficient, subjective, hazardous, and difficult to quantify rapidly. In particular, when the affected area is extensive, the number of bridges is large, or critical inspection regions are difficult to access, bridge-by-bridge and component-by-component manual inspection may delay post-earthquake traffic recovery and emergency response decisions [9–11]. The post-earthquake damage state of RC components is generally associated with deformation capacity, strength degradation, stiffness degradation, energy dissipation capacity, and cumulative damage [12]. Several studies have focused on failure-mode identification, performance-state classification, and post-earthquake residual capacity assessment of RC bridge piers. For example, Mangalathu and Jeon [13] developed a machine-learning-based method for identifying the failure modes of circular RC bridge piers using an experimental database and compared the classification performance of different models for flexural, flexure-shear, and shear failure modes. Ali et al. [14] further applied feature selection methods to identify key input parameters for failure-mode classification of RC bridge piers. Structural response data provide mechanically meaningful evidence for damage-state identification. However, in actual post-earthquake scenarios, complete sensor data, material properties, and detailed component information are often difficult to obtain promptly. Recent database-based studies on existing bridges have also shown that concrete and reinforcement properties may exhibit substantial variability, highlighting the importance of material characterization and knowledge uncertainty in structural safety assessment [15]. Therefore, assessments based solely on structural responses or design parameters may fail to reflect the spatial distribution and apparent severity of local pier damage.

Compared with structural response indicators, post-earthquake images can directly record visible damage phenomena, such as crack propagation, concrete spalling, reinforcement exposure, and local crushing. With the development of unmanned aerial vehicles, mobile devices, and high-resolution imaging technologies, the efficiency of post-earthquake bridge image acquisition has been substantially improved, providing an important source of information for rapid damage screening [16,17]. Ye et al. [18] proposed a post-earthquake bridge damage identification and condition assessment framework that

integrates satellite imagery, unmanned aerial vehicles, smartphones, and deep learning, demonstrating the potential of UAVs and intelligent image analysis for rapid post-earthquake inspection. Cheng et al. [19] developed an automated post-earthquake bridge safety assessment system based on STF-PointRend and EfficientNetB0, in which visual models were used for component recognition, damage-type detection, damage-level identification, and safety-state determination, thereby providing a feasible pathway for automated post-earthquake bridge assessment. Tazarv et al. [20] combined computer vision with seismic analysis for post-earthquake performance assessment of standard RC bridge piers and indicated that integrating visual inspection results with structural analysis results can accelerate post-earthquake bridge condition evaluation. Recent studies by Salamak and co-workers have further advanced the digitalization of bridge inspection through BrIM, augmented reality, mobile inspection, and digital twin technologies, providing integrated platforms that link field observations, structural information, and asset management decisions [21–23]. These developments also highlight the need to connect visually observed damage with mechanically interpretable structural states.

In terms of image-processing methods, early object detection and instance segmentation models represented by Faster R-CNN, Mask R-CNN, and the YOLO series enabled automatic localization and category identification of damaged regions [24–27]. Semantic segmentation models, such as U-Net and DeepLabv3+, can further generate pixel-level damage masks, thereby providing a basis for analyzing damage morphology, spatial distribution, and geometric characteristics [28,29]. Hoskere et al. [30] developed MaDnet, a multitask semantic segmentation framework that enables the simultaneous identification of material categories, including concrete, steel, and asphalt, as well as multiple types of structural damage, such as cracks, spalling, corrosion, and exposed reinforcement. Their study demonstrated that deep learning models can identify both material attributes and damage categories. Kang et al. [31] proposed a pixel-level segmentation and quantification method for concrete cracks under complex backgrounds, demonstrating that the integration of deep learning and image processing can effectively extract crack geometric features. Zhao et al. [32] further proposed a refined crack assessment method for tunnel lining cracks, enabling the calculation of geometric parameters such as crack width and length. Wang et al. [33] developed an automatic detection method based on machine vision and deep learning for concrete spalling and reinforcement exposure in RC structures. These studies indicate that pixel-level recognition results can provide a database for advancing structural damage assessment from determining whether damage exists and where it is located to understanding how it develops and how severe it is.

Although significant progress has been made in computer-vision-based structural damage recognition, several challenges remain when these methods are directly applied to post-earthquake performance assessment of RC double-column piers. First, the joint regions of RC double-column piers have complex configurations and dense reinforcement layouts. In addition, post-earthquake field images are often affected by uneven illumination and background interference, making it difficult to ensure the accurate automated extraction of fine-grained damage features. Second, existing vision-based damage recognition studies have mainly focused on improving model accuracy and extracting damage geometry, whereas the mechanical relationship between visual damage features and structural performance states has not been sufficiently interpreted [34]. Furthermore, image-based damage recognition can characterize apparent damage but is not equivalent to changes in structural bearing capacity, deformation capacity, or post-earthquake trafficability. Conversely, assessments based solely on structural response indicators may overlook the spatial concentration and severity of local visible damage.

To address these limitations, this study proposes a Structural–Visual Integrated Evaluation (SVIE) framework for post-earthquake damage assessment of RC double-column

piers. The proposed framework integrates structural response analysis, image-based damage recognition, and consistency verification into a unified evaluation process. End-to-end semantic segmentation is performed on global scene images to extract pixel-level segmentation masks of concrete spalling and reinforcement exposure, which are then introduced into the structural performance interpretation process as visual damage evidence. On this basis, the correspondence between baseline structural damage states and visual damage evidence is analyzed using quasi-static test results. The novelty of this study lies in transforming segmentation-derived damage morphology from a standalone visual recognition result into supplementary evidence within a mechanically interpretable structural–visual consistency framework. The proposed framework provides a technical pathway for rapid post-earthquake assessment of RC double-column piers with both structural-mechanics-based evidence and visual interpretability, and it may also serve as a reference for intelligent post-earthquake bridge inspection and engineering decision-making.

2. Structural–Visual Integrated Evaluation Framework

2.1. Framework Architecture

To enable interpretable post-earthquake damage assessment of RC double-column piers, a Structural–Visual Integrated Evaluation (SVIE) framework is developed in this study, as shown in Figure 1. The SVIE framework consists of four components: the input module, structural stream, visual stream, and decision and output module. These components constitute an assessment workflow of structural identification, visual verification, and result interpretation.




Figure 1. The proposed Structural–Visual Integrated Evaluation (SVIE) framework.

- Input module.** This module receives both structural response data and visual image data. The structural information may include load, displacement, and strain measurements at critical regions, together with available bridge records, geometric characteristics, reinforcement details, and material properties. These data are used to interpret hysteretic characteristics, deformation capacity, and mechanical states of the components. The visual images are obtained from post-earthquake field scenes or experimental damage records and are used to capture the spatial distribution of apparent pier damage. Together, the structural and image data provide the basis for the subsequent dual-stream analysis.

- **Structural stream.** This module serves as the primary analytical pathway in the SVIE framework and is used to determine the baseline damage state of RC double-column piers. Based on force–displacement responses, strain development, and deformation characteristics, key structural stages can be identified, including cracking, yielding, peak resistance, strength degradation, and final failure. The structural performance level is then classified accordingly. The baseline damage state generated by the structural stream has clear mechanical significance and serves as the primary reference for subsequent verification using visual evidence.
- **Visual stream.** This module serves as an auxiliary analytical pathway and is used to extract visual evidence related to damage morphology from images. An end-to-end semantic segmentation strategy is adopted, in which global scene images are directly input into the improved DeepLabV3+ network to generate pixel-level segmentation masks for the background, concrete spalling regions, and reinforcement exposure regions. In the segmentation results shown in Figure 1, yellow denotes concrete spalling, whereas red denotes reinforcement exposure. Based on the segmentation results, surface deterioration evidence associated with concrete spalling and reinforcement-exposure evidence associated with advanced local damage can be obtained. These visual indicators are then used to support the interpretation, confirmation, and consistency verification of the structural damage state [35,36].
- **Decision and output.** This module compares the baseline structural state with the visual damage evidence. Consistent evidence supports acceptance of the baseline state, whereas transition-sensitive or conflicting evidence triggers a damage evolution or inconsistency flag. When both streams indicate advanced deterioration, the severe damage state is confirmed. The structural stream remains the primary basis for classification, while the visual stream provides supplementary verification and warning information.

2.2. Evidence-Based Decision Logic

The SVIE framework adopts an evidence-based decision logic in which structural responses are used as the baseline, while visual evidence is used as supplementary information. The structural stream determines the baseline damage state based on force–displacement responses, strain development, deformation capacity, and experimental observations. The visual stream extracts visual evidence, including concrete spalling and reinforcement exposure, from the end-to-end semantic segmentation results of global scene images. The decision layer then evaluates the correspondence between the structural state and the visual damage morphology and provides the corresponding interpretation.

Structural–visual consistency is evaluated by comparing the baseline structural state with the expected visual damage morphology defined for the four damage levels. The two streams are considered consistent when the observed visual evidence falls within the expected damage characteristics of the corresponding structural state, and no feature associated with a higher damage level is detected. If the visual evidence is either more severe or less severe than expected, the case is classified as inconsistent and is processed through Path B. According to the relationship between the baseline structural state and visual evidence, the decision logic is divided into three paths, as shown in Figure 1 and summarized in Table 1.

Table 1. Decision rules of the SVIE framework.

Path	Decision Interpretation	Decision Output
A	Structural and visual evidence are compatible.	Baseline damage state accepted
B	Transition-sensitive or inconsistent evidence is identified.	Damage evolution or inconsistency flag
C	Both streams indicate severe damage.	Severe damage stage confirmed

Path A represents consistent damage interpretation. When the observed visual evidence matches the expected damage morphology of the baseline structural state, the output is “Baseline damage state accepted”.

Path B addresses damage-transition and inconsistency cases. When the visual evidence indicates more advanced damage than the baseline structural state, a damage evolution flag is issued without directly upgrading the structural classification. Conversely, when the structural response indicates more severe deterioration than is visible in the images, the structural state is retained, and an inconsistency flag is issued. In both cases, further inspection or monitoring is recommended.

Path C represents severe-damage confirmation. When pronounced structural deterioration is accompanied by reinforcement exposure, severe concrete spalling, core crushing, or highly concentrated local failure, the output is “Severe damage stage confirmed”.

Visual evidence does not independently determine or override the structural damage state. Instead, it is used to verify the baseline assessment, identify potential damage transitions, and indicate cases that require further inspection.

3. Data Source and Image Acquisition

The data used in this study were obtained from laboratory quasi-static cyclic tests on seven 1/5.6-scale RC double-column pier specimens. The specimens were designed based on a typical highway double-column bridge pier, with a model-column diameter of 250 mm. The experimental program considered variations in axial compression ratio, shear-span ratio, and volumetric transverse-reinforcement ratio to produce different deformation capacities and damage-development characteristics. All specimens were subjected to constant axial loading and reversed cyclic lateral loading. The data types and their roles in the SVIE framework are summarized in Table 2.

Table 2. Summary of data used for SVIE framework validation.

Data Type	Source	Information Included	Use in This Study
Structural response	Quasi-static cyclic test	Load, displacement response	Determination of baseline structural damage state.
Strain/deformation response	Test measurements	Reinforcement strain, concrete strain, and deformation development.	Verification of key mechanical states.
Damage observations	Test records	Cracking, concrete spalling, local crushing, reinforcement exposure.	Physical interpretation of structural stages.
Global-scene images	Test image acquisition	Full-view damage images with complex backgrounds.	Input for end-to-end semantic segmentation.
Pixel-level labels	Manual annotation	Background, concrete spalling, reinforcement exposure.	Training and validation of the segmentation model.
Segmentation masks	DeepLabv3+ output	Pixel-level visual damage regions.	Extraction of visual evidence for consistency analysis.

During the tests, load, pier-top displacement, strain at critical regions, and full-process damage images were collected. The structural response data were mainly used to identify key states of the specimens during loading, including concrete cracking, longitudinal bar yielding, concrete spalling, strength degradation, and final failure. These data provided the mechanical basis for classifying the baseline structural damage state. Detailed specimen configurations and experimental response analyses are beyond the scope of the present study and will be reported separately. The visual data were obtained from global scene images collected during the loading process. Unlike locally cropped images, global scene images include not only the bridge pier components but also complex background information, such as the loading device, ground background, marking lines,

shadows, and adjacent components. Therefore, they are closer to the image forms that may be acquired during actual post-earthquake inspection. In this study, such images were used as the input to the visual stream to examine the capability of the end-to-end semantic segmentation method to extract apparent damage evidence under complex background conditions.

For image annotation, two types of visible damage closely related to the post-earthquake performance state of RC double-column piers were considered: concrete spalling and reinforcement exposure. The pixel-level annotation categories included the background, concrete spalling regions, and reinforcement exposure regions. The complete dataset contained 315 global-scene images, including 252 for training, 31 for validation, and 32 for testing. The original images were resized while preserving their aspect ratios and were subsequently padded to 512×512 pixels before being input into the network.

A pronounced pixel-level class imbalance was observed in the dataset. Background pixels accounted for approximately 98.75% of all pixels, whereas concrete spalling and reinforcement exposure together accounted for only approximately 1.25%. Among the two damage classes, reinforcement exposure was the sparsest because it generally occupied only small and fragmented regions in the global-scene images. This imbalance was addressed during training through class weighting, a combined segmentation loss, and auxiliary supervision.

4. Baseline Damage-State Definition

In the SVIE framework, structural responses are used to provide the mechanical baseline for damage-state identification. The baseline structural states are mainly defined according to the force–displacement responses, strain development at critical regions, and observed damage phenomena obtained during the tests. To improve the objectivity and reproducibility of the classification, the four damage states were further associated with measurable response indicators. The reference variables included the displacement ductility ratio $\mu_\Delta = \Delta/\Delta_y$ and the strength retention ratio $F/F_{\max}$, where Δ_y is the yield displacement, and $F_{\max}$ is the peak lateral force.

The damage states of RC double-column piers are categorized into four levels. DS1 corresponds to the pre-yield stage, with $\mu_\Delta < 1.0$, no pronounced strength degradation, and no obvious damage beyond minor cracking or slight surface deterioration. DS2 begins at reinforcement yielding and covers moderate nonlinear deformation without evident post-peak strength loss. DS3 is associated with advanced nonlinear deformation, substantial stiffness deterioration, expanded concrete spalling, and a lateral resistance that remains above 85% of $F_{\max}$. DS4 is reached when the post-peak resistance decreases to 85% of $F_{\max}$, or when severe material damage, such as reinforcement exposure, bar buckling, or core concrete crushing, is observed. These quantitative indicators are used together with the corresponding physical damage phenomena rather than as isolated criteria.

The four levels are termed intact-to-slight damage, moderate damage, severe damage, and critical damage, respectively. The intact and slightly damaged states are combined into DS1 because the primary load-resisting and service functions are generally maintained. Critical damage denotes a near-failure level, but it does not necessarily indicate global structural collapse. In the SVIE framework, these four levels serve only as baseline structural states for evaluating whether the visual evidence extracted from concrete spalling and reinforcement exposure is compatible with the corresponding mechanical stage. Their structural basis, visual characteristics, and role in the framework are summarized in Table 3.

Table 3. Baseline damage states used in the SVIE framework.

Baseline Damage State	Structural Basis	Typical Visual Evidence	Role in SVIE
DS1: Intact-to-slight damage	Elastic or near-elastic response; no obvious strength degradation.	No obvious damage, hair-line cracks, or slight local surface deterioration.	Baseline state for minor-damage confirmation.
DS2: Moderate damage	Crack development, yielding tendency, and moderate deformation demand.	Local cover deterioration or limited concrete spalling.	State for moderate-damage interpretation.
DS3: Severe damage	Significant nonlinear deformation, local crushing, or a tendency toward strength degradation.	Expanded concrete spalling, main cracks, or concentrated damage zones.	State for identifying damage-transition evidence.
DS4: Critical damage	Obvious strength degradation, large deformation demand, or severe strain response.	Reinforcement exposure, bar buckling, or core concrete crushing.	State for severe-damage confirmation.

5. Segmentation Results and Consistency Analysis

5.1. End-to-End Segmentation Model

In the SVIE framework, the core task of the visual stream is to extract pixel-level visual evidence from global scene images to characterize apparent damage morphology. An end-to-end semantic segmentation strategy is adopted to avoid dependence on an additional localization module. This strategy also preserves the spatial relationships among damaged regions, bridge pier components, loading devices, and complex backgrounds, making it more consistent with the practical requirements of post-earthquake field image analysis.

As shown in Figure 2, the DeepLabv3+ semantic segmentation model is used as the basic network architecture, and appropriate improvements are introduced in terms of the encoder, feature enhancement, and training supervision. First, MobileNetv3-Large is adopted as the feature extraction backbone to reduce computational complexity and improve the potential for deployment in engineering inspection scenarios. The encoder extracts low-, mid-, and high-level features from the input image, which provide boundary details, local structural information, and high-level semantic context, respectively. Subsequently, the Atrous Spatial Pyramid Pooling (ASPP) module extracts contextual features under different receptive fields through multiscale atrous convolutions, thereby enhancing the ability of the model to represent damaged regions at different scales. The decoder further fuses multilevel features and restores the spatial resolution through upsampling. Finally, three segmentation masks corresponding to the input image are generated, including the background, concrete spalling, and reinforcement exposure classes.

To improve the recognition capability of the model for fine-grained damaged regions, a convolutional block attention module (CBAM) is introduced during feature extraction. This module enhances the response of key damaged regions from both the channel and spatial dimensions, enabling the model to focus more effectively on damaged areas, such as concrete spalling and reinforcement exposure, rather than irrelevant textures in complex backgrounds. In addition, given that concrete spalling boundaries are usually irregular and that regions of reinforcement exposure occupy a very small proportion of global images, an edge supervision branch and an auxiliary classification branch are introduced during training. The edge supervision branch is used to enhance the ability of the model to delineate damage contours, whereas the auxiliary classification branch provides additional supervision signals for intermediate features to improve training stability. It should be noted that these two branches are mainly used for feature constraint and gradient

guidance during training, whereas the main segmentation output is used as the visual evidence extraction result during inference.




Figure 2. Architecture of the improved DeepLabv3+ model for end-to-end visual evidence extraction.

During training, the images were resized while preserving their aspect ratio and padded to an input size of 512×512 . Data augmentation methods, including flipping, brightness, and contrast adjustment, were adopted to improve the adaptability of the model to complex backgrounds and illumination variations [37]. The model was trained using the AdamW optimizer with a batch size of 2. To alleviate the pronounced class imbalance, class weights of 0.3, 1.0, and 15.0 were assigned to the background, concrete spalling, and reinforcement exposure classes, respectively. The main segmentation loss was formulated as a weighted combination of cross-entropy, Dice, and Focal losses, with weights of 0.4, 0.4, and 0.2, respectively. The edge supervision and auxiliary classification branches were used only during training, whereas only the main segmentation output was retained during inference. The overall loss can be expressed as follows:

$$L_{\text{main}} = 0.4L_{\text{CE}} + 0.4L_{\text{Dice}} + 0.2L_{\text{Focal}} \quad (1)$$

$$L_{\text{total}} = L_{\text{main}} + \lambda_{\text{edge}}L_{\text{edge}} + \lambda_{\text{aux}}L_{\text{aux}} \quad (2)$$

where L_{main} denotes the main segmentation loss, L_{edge} denotes the edge supervision loss, and L_{aux} denotes the auxiliary classification loss. The parameters λ_{edge} and λ_{aux} are the weights of the corresponding supervision branches. The main segmentation loss consists of cross-entropy loss, Dice loss, and Focal loss, with weights of 0.4, 0.4, and 0.2, respectively. The edge supervision and auxiliary classification branches are used for training the Full model.

5.2. Segmentation Performance

To verify the capability of the improved DeepLabv3+ model to extract concrete spalling and reinforcement exposure regions from global-scene images, the segmentation performance was evaluated through test-set quantitative analysis, ablation experiments,

comparisons with representative models, computational-efficiency evaluation, and training-convergence analysis. The evaluation metrics included IoU, F1-score, Precision, and Recall [38,39]. Among them, IoU was used to measure the overlap between the predicted regions and the ground-truth annotations. F1-score was used to comprehensively reflect the balance between Precision and Recall, whereas Precision and Recall were used to evaluate the reliability of positive predictions and the degree of missed detection, respectively.

Table 4 presents the segmentation performance of the Full model on the test set. The model achieved the best performance for the background class, with an IoU of 99.82% and a Dice score of 99.91%, indicating that non-damaged regions in global scene images could be effectively distinguished. For the concrete spalling class, the model achieved an IoU of 78.86%, an F1-score of 88.18%, a Precision of 89.93%, and a Recall of 86.49%, indicating that regions of concrete surface deterioration could be extracted relatively stably. In contrast, the IoU and Recall of the reinforcement exposure class were 24.71% and 27.58%, respectively, which were considerably lower than those of the concrete spalling class. However, its Precision reached 70.40%, suggesting that the model tended to make conservative predictions for reinforcement exposure. In other words, once a reinforcement exposure region was detected, the prediction was generally reliable.

Table 4. Test-set performance of the Full segmentation model.

Class	IoU	F1-Score	Precision	Recall
Background	99.82%	99.91%	99.88%	99.94%
Spalling	78.86%	88.18%	89.93%	86.49%
Rebar	24.71%	39.63%	70.40%	27.58%
Mean	67.79%	75.90%	86.74%	71.33%

It should be noted that the lower performance for the reinforcement exposure class was mainly related to sample scarcity and fine-grained visual characteristics. In global scene images, reinforcement exposure typically accounts for only a very small proportion of pixels and is often mixed with concrete spalling, shadows, fragmented boundaries, and background textures, making it difficult for the model to completely recall all small reinforcement regions. Therefore, in this study, the reinforcement exposure results were mainly considered as auxiliary evidence of severe local damage. The relatively high Precision indicates that detected reinforcement-exposure regions can provide reliable supplementary evidence for severe local damage. However, because the low Recall may result in missed damaged regions, the absence of a detected reinforcement region should not be interpreted as evidence that severe damage is absent.

To further analyze the contribution of different improvement modules, four ablation variants were designed, including the baseline, +CBAM, +CBAM+Edge, and Full model. Table 5 compares the performance of these variants on the validation set. When CBAM was introduced alone, the Mean IoU decreased slightly from 66.62% to 66.23%, while the Spalling IoU and Rebar IoU decreased from 75.04% and 25.06% to 74.09% and 24.86%, respectively. These results indicate that CBAM alone did not improve the segmentation performance on the present dataset. This may be because feature recalibration by itself is insufficient to provide stable guidance for highly sparse and fragmented damage regions without additional boundary or semantic constraints. After edge supervision was further introduced together with CBAM, the Mean IoU, Spalling IoU, and Rebar IoU increased to 67.01%, 74.38%, and 26.92%, respectively. Compared with the CBAM-only model, these values increased by 0.78, 0.29, and 2.06 percentage points, respectively. In particular, the Rebar IoU was 1.86 percentage points higher than that of the baseline. These results indicate that edge supervision provides useful boundary information for small and irregularly shaped damage regions and improves the recognition of fragmented reinforcement-

exposure areas. After the auxiliary classification branch was further incorporated, the Full model achieved the highest Mean IoU, Spalling IoU, and Rebar IoU, reaching 68.16%, 75.42%, and 29.32%, respectively. Compared with the baseline, the Full model improved Mean IoU by 1.54 percentage points, Spalling IoU by 0.38 percentage points, and Rebar IoU by 4.26 percentage points.

The largest observed gain occurred for reinforcement exposure, indicating that the combined modules mainly benefited the recognition of the sparsest and most fragmented damage class. Nevertheless, the overall Mean IoU improvement remained moderate and should not be interpreted as evidence of broad model superiority.

Table 5. Ablation study of the segmentation model on the validation set.

Model Variant	CBAM	Edge	Aux cls	Mean IoU	Spalling IoU	Rebar IoU
Baseline	×	×	×	66.62%	75.04%	25.06%
+CBAM	√	×	×	66.23%	74.09%	24.86%
+CBAM+Edge	√	√	×	67.01%	74.38%	26.92%
Full	√	√	√	68.16%	75.42%	29.32%

Note: √ indicates that the corresponding module was included, whereas × indicates that it was not included.

It is worth noting that the ablation results do not indicate a simple linear superposition of the effects of the individual modules. CBAM alone did not improve the segmentation performance relative to the baseline. However, when combined with edge supervision and the auxiliary classification branch, the Full model achieved the best overall performance. This suggests that the effectiveness of the attention mechanism depends on complementary boundaries and semantic constraints. Specifically, CBAM recalibrates the responses of damage-related features, edge supervision strengthens the representation of irregular damage boundaries, and the auxiliary classification branch provides additional semantic constraints on intermediate features. Their combined use improves the recognition of fine-grained damage regions under complex background conditions.

To examine the robustness of the observed improvement, a paired image-level bootstrap analysis with 10,000 resamples was conducted on the 31-image validation set. Based on the unrounded confusion-matrix results, the Full model improved Mean IoU by 1.55 percentage points and Rebar IoU by 4.29 percentage points relative to the baseline. However, the corresponding 95% confidence intervals were [−1.44, 5.70] and [−4.30, 15.21], respectively, and neither difference was statistically significant at the 0.05 level. Therefore, the observed gains should be interpreted as modest empirical improvements on the present validation set rather than conclusive evidence of statistically robust superiority. The wide confidence interval for Rebar IoU reflects the limited and highly imbalanced reinforcement-exposure samples. The relatively high Precision. The analysis quantifies sampling uncertainty on the fixed validation set but not variability across independent training runs.

To further evaluate the performance of the proposed model, comparative experiments were conducted using U-Net, PSPNet, and the baseline DeepLabv3+ model. All models were trained and evaluated using the same dataset split, input resolution, and evaluation procedure. The comparison results are summarized in Table 6.

Table 6. Comparison with representative semantic segmentation models on the validation set.


Model	Mean IoU	F1-Score	Spalling IoU	Rebar IoU
U-Net	67.71%	76.25%	76.75%	26.58%
PSPNet	60.26%	69.53%	61.94%	19.25%
Baseline	66.62%	75.23%	75.04%	25.06%
Full	68.16%	77.07%	75.42%	29.32%

Note: Bold values indicate the best performance for each metric among the compared models.

The Full model achieved the highest Mean IoU and F1-score, reaching 68.16% and 77.07%, respectively. Although its overall improvement over U-Net was moderate, the Full model achieved the highest Rebar IoU of 29.32%, exceeding U-Net, PSPNet, and the baseline by 2.74, 10.07, and 4.26 percentage points, respectively. This result indicates that the proposed modules primarily benefited the recognition of sparse and fragmented reinforcement-exposure regions, which represented the most challenging damage class in the present dataset. U-Net achieved a slightly higher Spalling IoU than the Full model, with values of 76.75% and 75.42%, respectively. Nevertheless, the Full model maintained comparable performance for concrete spalling while achieving the highest overall metrics and a more balanced recognition capability across the two damage classes.

The computational efficiency of the proposed model was also evaluated using an NVIDIA GeForce RTX 4060 Laptop GPU with an input resolution of 512×512 pixels. After 20 warm-up runs, the average pure inference time over 32 timed runs was 14.05 ± 1.14 ms per image, corresponding to 71.2 FPS. When image preprocessing was included, the average processing time was 21.50 ms per image, corresponding to 46.5 FPS. Image-file reading was excluded from the timing procedure. These results indicate that the visual stream can support rapid image-based screening without introducing a substantial computational delay.

Furthermore, Figure 3 shows the variation curves of validation loss, Mean IoU, Spalling IoU, and Rebar IoU for different model variants during training. The validation loss of each model decreased rapidly at the early training stage and gradually stabilized in the later stage, indicating that the models could quickly learn the main damage-related features from global scene images. Both Mean IoU and Spalling IoU showed a trend of rapid improvement followed by gradual stabilization. Among the variants, the Full model generally achieved better performance in the later training stage, suggesting that the simultaneous introduction of CBAM, edge supervision, and the auxiliary classification branch was beneficial for improving feature representation and overall segmentation performance.






Figure 3. Training convergence and validation performance curves of different model variants: (a) validation loss; (b) Mean IoU; (c) Spalling IoU; (d) Rebar IoU.

Compared with the concrete spalling class, Rebar IoU showed more obvious fluctuations and remained relatively low. This was mainly because reinforcement exposure regions occupied a small pixel proportion in global images and were often mixed with concrete spalling, shadows, fragmented boundaries, and background textures. As a result, a small pixel-level prediction error could lead to a considerable change in IoU. Nevertheless, Rebar IoU still showed an overall increasing trend, and the Full model generally maintained a higher level in the later training stage, indicating that the complete model had a better capability for recognizing fine-grained reinforcement exposure features.

5.3. Structural–Visual Consistency

The segmentation results were further used for structural–visual consistency analysis. Figure 4 presents representative segmentation results under different baseline structural states and their corresponding SVIE interpretations. For DS1: intact-to-slight damage, the structural response was generally in the elastic or near-elastic stage, and no obvious concrete spalling or reinforcement exposure regions were observed in the segmentation results. This indicates that the visual evidence was consistent with the slight damage state. For DS2: moderate damage, local concrete surface deterioration began to appear in the images, and the segmentation results captured local damaged regions near the plastic hinge region. This suggests that the visual evidence could support the structural judgment of a moderate damage state. For DS3: severe damage, the concrete spalling regions further expanded and became concentrated in local critical regions, indicating that the apparent damage had developed from localized deterioration to more extensive surface failure. In the SVIE framework, such results can be regarded as visual indications of damage-stage transition or accelerated damage development. For DS4: critical damage, obvious reinforcement exposure, or severe local spalling regions appeared in the segmentation results, indicating that cover failure and serious damage had occurred. These features can be used as auxiliary confirmation information for the severe damage state.






Baseline structure state	DS1 Intact-to-slight damage		DS2 Moderate damage		DS3 Severe damage		DS4 Critical damage	
	• Near-elastic response • No obvious strength degradation		• Crack development • Yielding tendency		• Nonlinear deformation • Local crushing tendency		• Strength degradation • Severe strain response	
Representative global-scene image and segmentation result	Global-scene image	Segmentation result	Global-scene image	Segmentation result	Global-scene image	Segmentation result	Global-scene image	Segmentation result
								
SVIE interpretation	 State confirmed		 Moderate damage supported		 Damage transition flag		 Severe damage confirmed	

Figure 4. Structural–visual consistency analysis under different baseline damage states.

As shown in Figure 4, as the baseline structural state developed from DS1 to DS4, the visual damage morphology reflected by the segmentation results showed a progressive trend from no obvious damage or slight local damage to local spalling, extensive surface deterioration, reinforcement exposure, and severe local failure. This trend was generally consistent with the damage evolution characteristics of RC double-column piers under cyclic loading, indicating that end-to-end semantic segmentation results can provide intuitive visual support for structural damage-state interpretation.

It should be emphasized that visual evidence in the SVIE framework is mainly used for verification, interpretation, and supplementation, rather than as a direct replacement for structural assessment. When the structural state is consistent with the visual damage morphology, the baseline damage state can be accepted. When the visual results show obvious expansion of the damaged region or high concentration of local damage, a damage evolution flag can be output. When reinforcement exposure or severe crushing is observed together with corresponding structural response deterioration, the severe damage state can be confirmed. In this way, semantic segmentation results are transformed into visual evidence for structural interpretation, thereby improving the intuitiveness and interpretability of post-earthquake assessment results for RC double-column piers.

6. Discussion

6.1. Validation and Practical Use

The proposed framework was validated using quasi-static cyclic tests on seven RC double-column pier specimens and a dataset of 315 global-scene images, including 252 training images, 31 validation images, and 32 test images. The specimens exhibited progressive damage ranging from minor cracking and local surface deterioration to concrete spalling, reinforcement exposure, and severe local failure. Because both the structural responses and images were obtained under controlled laboratory conditions, the present study should be regarded as a proof of concept rather than direct validation under real post-earthquake inspection conditions.

The structural stream is intended to provide mechanically interpretable baseline states rather than to reproduce complete earthquake loading histories. In practice, it may be constructed from monitoring records, displacement measurements, rapid structural analysis, digital-twin estimates, or other available response indicators. The reliability of the structural stream also depends on the available level of knowledge regarding bridge geometry, reinforcement details, and the mechanical properties of concrete and reinforcing steel. For existing bridges, these properties may be obtained from design and

maintenance records, material testing, nondestructive evaluation, or representative database information. When reliable values are unavailable, conservative parameter ranges or uncertainty-aware structural analyses should be used because material variability may affect the estimated stiffness, strength, yielding point, and corresponding damage-state boundaries [12,15]. In the present laboratory study, the specimens were characterized using measured material properties; therefore, the influence of incomplete material knowledge was not explicitly investigated. The framework is intended for preliminary post-earthquake screening after image acquisition and before detailed inspection or repair decisions. Within existing regulatory procedures, SVIE could be applied after rapid or special inspection data are collected and before detailed safety verification. Its outputs would not replace official condition ratings or engineering calculations but could help standardize visual evidence, identify inconsistencies between apparent damage and structural information, and prioritize bridges or components for detailed inspection, monitoring, load restriction, or analytical assessment. Potential users include bridge owners, maintenance agencies, emergency inspection teams, and structural engineers. Under the reported hardware configuration, the visual stream achieved a preprocessing–inference speed of 46.5 FPS, indicating that model computation is unlikely to be the main bottleneck in rapid image-based screening. However, the overall assessment time will also depend on image acquisition, data transmission, image-quality control, and the availability of structural information.

When structural information is unavailable, the visual stream may still support rapid screening and prioritization. However, such results should be regarded as visual-only screening because images alone cannot fully characterize hidden deterioration, stiffness loss, or residual load-carrying capacity. The framework has not yet been implemented in real post-earthquake bridge inspection, and further field validation is required before routine deployment. Although the current study focuses on RC double-column piers, the general dual-stream architecture may be extended to other bridge components. Such applications would require recalibration of the structural indicators, visual damage categories, training data, and decision rules for the target component.

6.2. Value of Structural–Visual Integration

The segmentation model is used primarily to extract visible damage evidence rather than to establish a new state-of-the-art architecture. Existing approaches can be broadly divided into three categories. The first focuses on visual-only detection or segmentation of cracks, concrete spalling, and reinforcement exposure [30–33,35,36]. The second combines visual observations with seismic analysis or broader digital inspection platforms, but the two information sources are often interpreted in parallel rather than through an explicit decision logic [18–23]. The proposed SVIE framework differs from these approaches by using mechanically defined structural states as the baseline and incorporating segmentation-derived damage morphology through explicit consistency-based decision paths. Visual evidence is therefore used for baseline-state verification, transition or inconsistency warning, and severe-damage confirmation rather than for direct damage-state assignment.

Visual-only assessment can identify the type, location, and distribution of visible damage but cannot independently determine stiffness degradation, deformation demand, strength loss, or residual mechanical condition. Structural-only assessment provides a mechanically meaningful baseline state but does not directly describe the spatial concentration and morphology of local surface damage. Their integration therefore adds decision-oriented information, including state confirmation, identification of transition-sensitive or inconsistent cases, and severe-damage confirmation.

Visible damage and structural deterioration do not necessarily evolve synchronously. Internal yielding or stiffness degradation may occur before pronounced surface damage becomes visible, whereas localized spalling may develop while the global response remains relatively stable. Such discrepancies are therefore not treated as framework failure but are processed through the inconsistency pathway defined in Section 2.2, where further inspection or monitoring is recommended.

6.3. Segmentation Limitations

Reinforcement exposure segmentation was less accurate than concrete spalling segmentation, particularly in terms of IoU and Recall. This limitation is mainly attributed to the extremely small pixel proportion, fragmented geometry, weak contrast, and overlap with spalling, shadows, debris, and background textures. As a result, narrow or partially occluded reinforcement regions may be missed.

From a practical perspective, missed reinforcement regions may lead to underestimation of severe local damage if the segmentation output is interpreted independently. This limitation is particularly important because reinforcement exposure is commonly associated with cover loss and advanced material deterioration. To reduce this risk, reinforcement exposure is not used as a standalone criterion for damage-state assignments in the SVIE framework. Severe damage is confirmed through the combined interpretation of structural deterioration, severe concrete spalling, reinforcement exposure, and the consistency between the two evidence streams. When structural deterioration is evident, but reinforcement exposure is not detected, the baseline structural state is retained, and the case is treated as potentially inconsistent rather than downgraded.

In practical applications, low-confidence predictions, severe spalling without detected reinforcement, or images affected by occlusion and insufficient resolution should trigger additional image acquisition, close-range inspection, or manual verification. Future improvement may be achieved through targeted sample enrichment, damage-focused augmentation, patch-based or multiscale segmentation, class-balanced losses, boundary-aware supervision, and multi-view or close-range image acquisition.

6.4. Field Applicability

Although the global-scene images used in this study contained shadows, loading devices, marking lines, and adjacent components, they were still collected under controlled experimental conditions and do not fully represent real post-earthquake scenes. The framework has not yet been evaluated using images collected after actual seismic events or publicly available field benchmark datasets; therefore, its cross-domain generalization capability remains to be demonstrated.

In field applications, illumination variation, occlusion, debris, oblique viewing angles, camera distance, motion blur, and surface contamination may reduce segmentation performance, particularly for small and fragmented damage regions. Future datasets should therefore include broader environmental and viewing conditions, while stronger augmentation, image-quality screening, multi-view acquisition, and domain adaptation may improve robustness. In practice, low-quality or highly occluded images should trigger additional image acquisition or manual inspection rather than being used directly for final damage-state interpretation.

7. Conclusions

To meet the need for rapid, intuitive, and interpretable post-earthquake damage assessment of RC double-column piers, this study proposed a Structural–Visual Integrated Evaluation (SVIE) framework. The framework integrates structural response analysis, semantic segmentation of global scene images, and structural–visual consistency

verification into a unified workflow. In this framework, the structural stream is used to determine the baseline damage state, the visual stream is used to extract apparent damage evidence, and the decision layer is used to evaluate the correspondence between the two streams. The main conclusions are summarized as follows:

- (1) The proposed SVIE framework integrates structural response data and visual image information to form an assessment logic of structural identification, visual verification, and result interpretation. In this framework, structural responses serve as the primary basis for damage-state identification, whereas visual segmentation results are used as apparent damage evidence to interpret, verify, and supplement the baseline structural assessment. This strategy avoids the limitation of directly assigning damage levels based solely on image recognition results and improves the interpretability of post-earthquake assessment outcomes.
- (2) The improved DeepLabv3+ model achieved reliable segmentation of concrete spalling in global-scene images, whereas reinforcement exposure remained more challenging because of its small pixel proportion, fragmented geometry, and background interference. The relatively high Precision for reinforcement exposure indicates that detected regions can still provide supplementary evidence for severe local damage, although missed detections remain an important limitation.
- (3) The ablation and comparative experiments showed that the proposed improvements produced a modest empirical improvement in overall segmentation performance, with the largest observed gain occurring for reinforcement exposure, while maintaining comparable performance for concrete spalling. However, the paired bootstrap analysis indicated that the improvements in Mean IoU and Rebar IoU were not statistically significant on the present 31-image validation set. Therefore, these gains should be interpreted as preliminary empirical evidence rather than conclusive proof of general model superiority. The Full model also retained sufficient computational efficiency for rapid image-based screening.
- (4) The consistency analysis showed that, as the baseline structural state developed from DS1 to DS4, the visual segmentation results exhibited a progressive trend from no obvious damage or slight local damage to local spalling, extensive surface deterioration, reinforcement exposure, and severe local failure. This trend was generally consistent with the damage evolution characteristics of RC double-column piers under cyclic loading, indicating that end-to-end semantic segmentation results can provide intuitive visual support for structural damage-state interpretation.

Overall, the SVIE framework provides a preliminary technical pathway for integrating mechanically interpretable structural information with image-based damage evidence during rapid post-earthquake screening of RC double-column piers. Its practical value lies in supporting baseline-state verification, identifying inconsistent or transition-sensitive cases, and prioritizing components for further inspection or structural evaluation. However, the present validation remains limited to seven scaled specimens and laboratory-acquired images, and the framework has not yet been verified under real post-earthquake bridge-inspection conditions. In addition, the relatively low Recall for reinforcement exposure may lead to missed severe local damage when visual information is used alone. Future research should therefore focus on real post-earthquake image collection, cross-domain validation, reinforcement-exposure sample enrichment, uncertainty-aware integration of material and structural information, and the incorporation of UAV, monitoring, and digital-twin data. These developments are necessary before routine engineering deployment.

Author Contributions: Conceptualization, Z.Y. and C.Z.; methodology, Z.X.; software, Z.X.; validation, Y.Z., X.H. and L.Z.; data curation, W.L. and H.L.; writing—original draft preparation, Z.X.; writing—review and editing, Z.Y. and C.Z.; visualization, Y.Z. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the National Natural Science Foundation of China, grant number 52378482; the Sichuan Province “Challenge-Based Approach” Project for Science and Technology Empowering Disaster Prevention, Mitigation, and Relief, project number Sichuan Transportation Construction (2025)-JZ-16; and the Chongqing Talent Plan Project, grant number cstc2022ycjh-bgzxm0133.

Data Availability Statement: Data will be made available on request. The source code is not publicly available at this stage because it forms part of an ongoing project and is subject to project-related restrictions.

Conflicts of Interest: Authors Zhibing Yu, Chao Zhang, Yongjun Zhou, Xiaojun Hou, Longqi Zhang, Wang Liao, and Haiou Li were employed by Sichuan Transportation Construction Group Co., Ltd. The remaining author, Zixin Xu, declares that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

References

1. Liu, L.; Yang, D.Y.; Frangopol, D.M. Network-level risk-based framework for optimal bridge adaptation management considering scour and climate change. *J. Infrastruct. Syst.* **2020**, *26*, 04019037.
2. Zhou, Y.; Banerjee, S.; Shinozuka, M. Socio-economic effect of seismic retrofit of bridges for highway transportation networks: A pilot study. *Struct. Infrastruct. Eng.* **2010**, *6*, 145–157. <https://doi.org/10.1080/15732470802663862>.
3. Shi, Y.; Li, J.; Qin, H.G.; Wang, D.S.; Wang, J.W.; Sun, Z.G. Review on seismic performance of bridge double-column bents. *China J. Highway Transp.* **2021**, *34*, 134–154. <https://doi.org/10.19721/J.CNKI.1001-7372.2021.02.004>.
4. Kawashima, K.; Takahashi, Y.; Ge, H.; Wu, Z.; Zhang, J. Reconnaissance report on damage of bridges in 2008 Wenchuan, China, earthquake. *J. Earthq. Eng.* **2009**, *13*, 965–996. <https://doi.org/10.1080/13632460902859169>.
5. Salvatore, W.; Uva, G.; Venanzi, I. Structural and Seismic Risk Classification of Bridges According to the Italian Guidelines: A Critical Perspective from the Jointed Experience of Fabre and ANAS. *Int. J. Bridge Eng. Manag. Res.* **2026**, *3*, 21426006–1:18. <https://doi.org/10.70465/ber.v3i1.72>.
6. Federal Highway Administration. *National Bridge Inspection Standards*, 23 CFR Part 650, Subpart C; U.S. Department of Transportation: Washington, DC, USA, 2022. Available online: <https://www.fhwa.dot.gov/bridge/nbis2022/qanda.cfm> (accessed on 20 July 2026).
7. National Highways. *CS 450: Inspection of Highway Structures*; Design Manual for Roads and Bridges; National Highways: Birmingham, UK, 2021. Available online: <https://www.standardsforhighways.co.uk/search/c5c2c3e5-f7f3-4c94-8254-184e41ccd1a0> (accessed on 20 July 2026).
8. Ministry of Land, Infrastructure, Transport and Tourism. *Periodic Inspection Guidelines for Road Bridges*; MLIT: Tokyo, Japan, 2024. Available online: https://www.mlit.go.jp/road/road_e/s3_maintenance.html (accessed on 20 July 2026).
9. Park, Y.J.; Ang, A.H.S. Mechanistic seismic damage model for reinforced concrete. *J. Struct. Eng.* **1985**, *111*, 722–739. [https://doi.org/10.1061/\(ASCE\)0733-9445\(1985\)111:4\(722\)](https://doi.org/10.1061/(ASCE)0733-9445(1985)111:4(722)).
10. Cha, Y.J.; Ali, R.; Lewis, J.; Büyüköztürk, O. Deep learning-based structural health monitoring. *Autom. Constr.* **2024**, *161*, 105328. <https://doi.org/10.1016/j.autcon.2024.105328>.
11. Zhang, L.; Shen, J.; Zhu, B. A review of the research and application of deep learning-based computer vision in structural damage detection. *Earthq. Eng. Eng. Vib.* **2022**, *21*, 1–21. <https://doi.org/10.1007/s11803-022-2074-7>.
12. Dong, H.; Feng, W.; Zheng, Z.; Han, Q.; Du, X. Experimental investigation on failure mechanism and hysteresis behavior of RC single- and double-column piers. *J. Bridge Eng.* **2024**, *29*, 04024094. <https://doi.org/10.1061/JBENF2.BEENG-6912>.
13. Mangalathu, S.; Jeon, J.S. Machine learning-based failure mode recognition of circular reinforced concrete bridge columns: Comparative study. *J. Struct. Eng.* **2019**, *145*, 04019104. [https://doi.org/10.1061/\(ASCE\)ST.1943-541X.0002402](https://doi.org/10.1061/(ASCE)ST.1943-541X.0002402).

14. Ali, N.M.; Farouk, A.I.B.; Haruna, S.I.; Alanazi, H.; Adamu, M.; Ibrahim, Y.E. Feature selection approach for failure mode detection of reinforced concrete bridge columns. *Case Stud. Constr. Mater.* **2022**, *17*, e01383. <https://doi.org/10.1016/j.cscm.2022.e01383>.
15. D'Amato, M.; Ranaldo, A.; Rosciano, M.; Zona, A.; Morici, M.; Gioiella, L.; Micozzi, F.; Poeta, A.; Quaglini, V.; Cattaneo, S.; et al. The Development and Statistical Analysis of a Material Strength Database of Existing Italian Prestressed Concrete Bridges. *Infrastructures* **2025**, *10*, 203. <https://doi.org/10.3390/infrastructures10080203>.
16. Lotfi Karkan, N.; Shakeri, E.; Sadeghi, N.; Banihashemi, S. Smart surveillance of structural health: A systematic review of deep learning-based visual inspection of concrete bridges using 2D images. *Infrastructures* **2025**, *10*, 338. <https://doi.org/10.3390/infrastructures10120338>.
17. Villarino, A.; Valenzuela, H.; Antón, N.; Domínguez, M.; Méndez Cubillos, X.C. UAV applications for monitoring and management of civil infrastructures. *Infrastructures* **2025**, *10*, 106. <https://doi.org/10.3390/infrastructures10050106>.
18. Ye, X.W.; Ma, S.Y.; Liu, Z.X.; Ding, Y.; Li, Z.X.; Jin, T. Post-earthquake damage recognition and condition assessment of bridges using UAV integrated with deep learning approach. *Struct. Control Health Monit.* **2022**, *29*, e3128. <https://doi.org/10.1002/stc.3128>.
19. Cheng, M.Y.; Sholeh, M.N.; Harsono, K. Automated vision-based post-earthquake safety assessment for bridges using STF-PointNet and EfficientNetB0. *Struct. Health Monit.* **2024**, *23*, 776–795. <https://doi.org/10.1177/14759217231168709>.
20. Tazarv, M.; Won, K.; Jang, Y.; Hart, K.; Greenway, E. Post-earthquake serviceability assessment of standard RC bridge columns using computer vision and seismic analyses. *Eng. Struct.* **2022**, *272*, 115002. <https://doi.org/10.1016/j.engstruct.2022.115002>.
21. Salamak, M.; Januszka, M. BrIM bridge inspections in the context of Industry 4.0 trends. In Proceedings of the 9th International Conference on Bridge Maintenance, Safety and Management (IABMAS 2018), Melbourne, Australia, 9–13 July 2018; pp. 2260–2267. <https://doi.org/10.1201/9781315189390-307>.
22. Jasiński, M.; Salamak, M.; Łaziński, P.; Kuras, P.; Owerko, T. The concept of creating digital twins of bridges using load tests. *Sensors* **2023**, *23*, 7349. <https://doi.org/10.3390/s23177349>.
23. Fawad, M.; Salamak, M.; Chen, Q.; Uscilowski, M.; Koris, K.; Jasinski, M.; Lazinski, P.; Piotrowski, D. Development of immersive bridge digital twin platform to facilitate bridge damage assessment and asset model updates. *Comput. Ind.* **2025**, *164*, 104189. <https://doi.org/10.1016/j.compind.2024.104189>.
24. He, K.; Gkioxari, G.; Dollár, P.; Girshick, R. Mask R-CNN. In Proceedings of the IEEE International Conference on Computer Vision, Venice, Italy, 22–29 October 2017; pp. 2961–2969. <https://doi.org/10.1109/ICCV.2017.322>.
25. Ren, S.; He, K.; Girshick, R.; Sun, J. Faster R-CNN: Towards real-time object detection with region proposal networks. In *Advances in Neural Information Processing Systems*; Neural Information Processing Systems Foundation, Inc.: San Diego, CA, USA, 2015; Volume 28.
26. Wang, C.Y.; Bochkovskiy, A.; Liao, H.Y.M. YOLOv7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Vancouver, BC, Canada, 18–22 June 2023; pp. 7464–7475. <https://doi.org/10.1109/CVPR52729.2023.00721>.
27. Cha, Y.J.; Choi, W.; Büyüköztürk, O. Deep learning-based crack damage detection using convolutional neural networks. *Comput.-Aided Civ. Infrastruct. Eng.* **2017**, *32*, 361–378. <https://doi.org/10.1111/mice.12263>.
28. Chen, L.C.; Zhu, Y.; Papandreou, G.; Schroff, F.; Adam, H. Encoder-decoder with atrous separable convolution for semantic image segmentation. In Proceedings of the European Conference on Computer Vision, Munich, Germany, 8–14 September 2018; pp. 801–818. https://doi.org/10.1007/978-3-030-01234-2_49.
29. Ronneberger, O.; Fischer, P.; Brox, T. U-Net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015*; Springer: Cham, Switzerland, 2015; pp. 234–241. https://doi.org/10.1007/978-3-319-24574-4_28.
30. Hoskere, V.; Narazaki, Y.; Hoang, T.A.; Spencer, B.F., Jr. MaDnet: Multi-task semantic segmentation of multiple types of structural materials and damage in images of civil infrastructure. *J. Civ. Struct. Health Monit.* **2020**, *10*, 757–773. <https://doi.org/10.1007/s13349-020-00409-0>.
31. Kang, D.; Benipal, S.S.; Gopal, D.L.; Cha, Y.J. Hybrid pixel-level concrete crack segmentation and quantification across complex backgrounds using deep learning. *Autom. Constr.* **2020**, *118*, 103291. <https://doi.org/10.1016/j.autcon.2020.103291>.
32. Zhao, S.; Zhang, D.; Xue, Y.; Zhou, M.; Huang, H. A deep learning-based approach for refined crack evaluation from shield tunnel lining images. *Autom. Constr.* **2021**, *132*, 103934. <https://doi.org/10.1016/j.autcon.2021.103934>.
33. Wang, S.; Wan, J.; Zhang, S.; Du, Y. Automatic detection method for concrete spalling and exposed steel bars in reinforced concrete structures based on machine vision. *Buildings* **2024**, *14*, 1580. <https://doi.org/10.3390/buildings14061580>.

34. Deng, H.R.; Li, H.J.; Xu, L.Q.; Khudhair, A.; Gao, Y. Real-time bridge disaster management: Enabling technology and application framework. *Autom. Constr.* **2025**, *174*, 106150. <https://doi.org/10.1016/j.autcon.2025.106150>.
35. Yasmin, T.; La, D.; La, K.; Nguyen, M.T.; La, H.M. Concrete spalling detection system based on semantic segmentation using deep architectures. *Comput. Struct.* **2024**, *300*, 107398. <https://doi.org/10.1016/j.compstruc.2024.107398>.
36. Santos, R.; Ribeiro, D.; Lopes, P.; Cabral, R.; Calçada, R. Detection of exposed steel rebars based on deep-learning techniques and unmanned aerial vehicles. *Autom. Constr.* **2022**, *139*, 104324. <https://doi.org/10.1016/j.autcon.2022.104324>.
37. Buslaev, A.; Iglovikov, V.I.; Khvedchenya, E.; Parinov, A.; Druzhinin, M.; Kalinin, A.A. Albumentations: Fast and flexible image augmentations. *Information* **2020**, *11*, 125. <https://doi.org/10.3390/info11020125>.
38. Andrushia, A.D.; Anand, N.; Lubloy, E.; G, P.A. Deep learning based thermal crack detection on structural concrete exposed to elevated temperature. *Adv. Struct. Eng.* **2021**, *24*, 1896–1909. <https://doi.org/10.1177/1369433220986637>.
39. Munawar, H.S.; Hammad, A.W.A.; Haddad, A.; Soares, C.A.P.; Waller, S.T. Image-based crack detection methods: A review. *Infrastructures* **2021**, *6*, 115. <https://doi.org/10.3390/infrastructures6080115>.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.