

# Indian Heritage Preservation using Crack detection Model

Ved Waje<sup>1</sup>, Maanav Valecha<sup>2</sup>, Nupur Pathare<sup>3</sup>

<sup>1</sup>Member, Vivekanand Education Society's Institute of Technology, Mumbai, India

**Abstract**—Preserving historical monuments is essential to protect cultural heritage, especially in a country such as India, which has a rich history of architectural landmarks. This research introduces an AI-based system for detecting cracks in heritage structures using deep learning and computer vision. The proposed system incorporates the Roboflow 3.0 Instance Segmentation Model, based on the COCO architecture, within a Streamlit web application. By applying advanced image processing and segmentation methods, the system accurately identifies structural cracks and marks them with bounding boxes and heatmaps. The application allows users to analyze multiple images at once, providing results in under one second per image. Users can adjust detection settings, select from different pretrained models, and download the processed images for further examination. A feedback mechanism helps improve the accuracy of the model over time, making the system adaptable to different types of structural damage. This approach improves the efficiency and reliability of heritage conservation by enabling early detection of cracks, supporting timely maintenance and restoration efforts.

**Index Terms**—Crack Detection, Heritage Preservation, Deep learning, COCO Architecture, Computer Vision.

## I. INTRODUCTION

The preservation of historical monuments is essential for maintaining cultural heritage and historical continuity. Over time, these structures deteriorate due to various factors, including environmental exposure, air pollution, natural disasters, and human activity [1]. Structural cracks are among the earliest indicators of damage, and their timely identification is crucial to prevent further degradation. Traditional inspection methods rely on manual surveys conducted by experts, which are often labor-intensive, time-consuming, and subject to human error. Furthermore, large-scale heritage sites require continuous monitoring, making manual assessments impractical for long-term conservation efforts [2], [3]. Recent advancements in artificial intelligence and computer vision have introduced data-driven approaches to structural damage detection,

significantly improving efficiency and accuracy. Deep learning-based models, particularly those utilizing instance segmentation techniques, provide a reliable method for identifying and classifying cracks in heritage structures [4]. Unlike conventional image processing techniques, which depend on predefined edge detection algorithms, deep learning models learn patterns of structural damage adaptively, allowing for more precise and extensive monitoring [5].

This research presents a web-based crack detection system that facilitates the analysis of structural damage in heritage sites using a pretrained instance segmentation model. The system uses the Roboflow 3.0 model, which is trained using the Common Objects in Context (COCO) architecture on the Darbhanga Fort dataset. The COCO architecture is widely used for object detection and instance segmentation tasks, making it well suited for identifying fine-grained structural details in historical monuments [6]. The Darbhanga Fort dataset consists of annotated images capturing cracks and damages in heritage structures, ensuring that the model is trained on real-world data. Users can upload images and receive real-time analysis with annotated outputs, including bounding boxes and heatmaps. Additionally, the system allows for the adjustment of detection parameters to accommodate specific conservation needs. A feedback mechanism is integrated to refine model performance over time, ensuring adaptability across different structural conditions.

By integrating artificial intelligence and deep learning techniques, this system provides an effective and efficient solution for structural damage assessment. The combination of the COCO architecture and the Darbhanga Fort dataset enhances the model's ability to detect cracks with high accuracy, making it a valuable tool for conservation professionals. The ability to identify structural weaknesses at an early stage supports heritage preservation efforts, allowing for timely

interventions and reducing the risk of further deterioration.

## II. RELATED WORKS

The application of advanced technologies in heritage conservation has increasingly focused on automating the detection and analysis of structural damage in historical monuments. Traditional methods, which rely heavily on manual inspections, are now being supplemented or replaced by automated systems that incorporate deep learning and computer vision techniques. One of the most effective approaches in this field is the use of deep learning models for object detection and segmentation to assess the condition of heritage sites. Recent studies have demonstrated the effectiveness of Convolutional Neural Networks (CNNs) for identifying structural damage in historical buildings [2], [7]. These models are trained on annotated datasets to detect cracks, erosion, and other forms of deterioration with high accuracy. Compared to manual inspections, deep learning-based techniques have shown significant improvements in both efficiency and precision [7], [8].

The COCO (Common Objects in Context) architecture has been widely applied in instance segmentation tasks, particularly for heritage site preservation. Research has demonstrated that COCO-based models are highly effective in detecting and highlighting structural issues in monument images [6]. Studies have shown that these models can segment and classify different types of damage, providing conservationists with detailed visualizations for maintenance planning [9], [10]. The use of COCO-based instance segmentation techniques has also been explored in various applications, including real-time monitoring and structural assessment of cultural heritage sites [3].

In addition to deep learning-based segmentation models, researchers have explored the use of 3D point clouds and 2D visual data for structural damage assessment. For instance, an object detection approach utilizing a combination of 3D and 2D data has been proposed to enhance accuracy in detecting cracks and structural defects in heritage monuments [7]. This multi-modal approach allows for a more comprehensive understanding of monument degradation, particularly in complex architectural structures. Furthermore, transfer learning-based CNN models have been utilized for structural damage

detection, demonstrating high performance in crack classification tasks. Studies have implemented these models for wall crack detection, showing that transfer learning can significantly enhance detection accuracy even with limited training data [11]. Similarly, research has explored the use of preprocessing techniques and automated image analysis to improve detection reliability in heritage site preservation [12]. The integration of automated image processing tools into web-based applications has also emerged as a key advancement in heritage conservation. These applications pre-process images to enhance the quality and detail before analysis [4], [10]. The ability to automate the detection and highlighting of cracks, complete with confidence scores, provides a streamlined and efficient solution for conservationists. Research has shown that such web applications can significantly reduce the time required for damage assessment while improving accuracy [4], [13].

Additionally, machine learning-based classification and recognition systems have been applied to identify and categorize heritage sites and their structural conditions. Some studies have focused on transfer learning for cultural heritage classification, demonstrating that pre-trained models can effectively classify monuments based on architectural features [8]. Other research has examined IoT-based deep learning approaches for protecting historical buildings, highlighting the potential of combining sensor data with AI models to enable continuous monitoring of heritage sites [14]. Recent efforts have also focused on ringfort detection using aerial photography and machine learning, showcasing the versatility of AI-driven methodologies for heritage conservation beyond just crack detection [15]. Similarly, web-based applications for automated damage detection have been extensively reviewed, with findings emphasizing the importance of accessibility and real-time analysis for conservation professionals [13].

Overall, the integration of deep learning and computer vision technologies represents a major advancement in structural damage assessment for heritage preservation. By leveraging CNNs, instance segmentation models, 3D point clouds, and web-based AI applications, researchers have developed highly effective automated methods for detecting and analyzing damage in historical sites. These techniques not only improve accuracy and efficiency but also ensure that monuments are preserved more effectively for future generations [5].

### III. PROPOSED ARCHITECTURE

This research introduces a novel system for detecting and highlighting damages in historical monuments using advanced deep learning techniques. The proposed architecture leverages the Roboflow 3.0 Instance Segmentation Model, which is built on the widely recognized COCO (Common Objects in Context) standardized image annotation format [6]. The COCO architecture is widely adopted in the field of computer vision, particularly for tasks such as object detection and segmentation. It utilizes instance segmentation, which goes beyond simple object detection by not only identifying the presence of objects in an image but also delineating the precise

boundaries of each object. This is especially useful when dealing with multiple objects of the same class, allowing the model to distinguish between individual instances [6].

At the core of the COCO architecture is the Mask R-CNN model, an extension of the Faster R-CNN architecture, which integrates a branch for generating segmentation masks along with bounding boxes for object localization. The architecture begins with an image preprocessing stage, where adjustments such as contrast enhancement and sharpening are applied to highlight features critical for detecting cracks or structural damage. These adjustments help the model identify subtle cracks on heritage sites that might otherwise be difficult to detect.

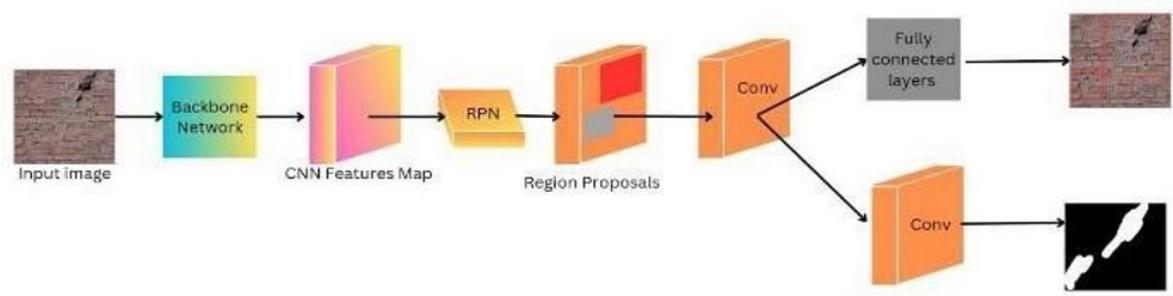


Fig. 1. Overview of the Mask R-CNN-based crack detection architecture.

The model then employs deep convolutional layers to extract hierarchical features at multiple levels, capturing fine details such as edges, textures, and structural patterns essential for detecting cracks. A Region Proposal Network (RPN) generates candidate regions in the image where objects might be present, focusing specifically on areas where cracks or other structural damages are likely to occur. Once these regions are identified, RoI Align is used to precisely extract features, ensuring accurate segmentation of the detected cracks. The Mask R-CNN then generates a segmentation mask for each region, outlining the exact shape and boundaries of the crack (Figure 1). Finally, the system assigns a confidence score to each detection, indicating the model's certainty in the accuracy of its predictions. This approach, when applied to the detection of cracks on Indian heritage sites, proves invaluable for automating the monitoring and preservation process, enabling early identification of damage while minimizing the potential for human error during manual inspections. The system is specifically designed to streamline the process of identifying and categorizing various types of cracks in images of historical monuments through an accessible and user-friendly web-based platform. By facilitating efficient and accurate detection of

structural damages, this system plays a crucial role in preserving cultural heritage and provides valuable insights for ongoing research in the field of conservation [9]. Additionally, the system's ability to adapt to different types of cracks and its integration of state-of-the-art AI technology make it a forward-looking tool that can evolve with future advancements in machine learning and image processing. This approach not only enhances the effectiveness of preservation efforts but also establishes a new standard in the application of AI for cultural heritage protection [3][11].

#### A. System Architecture

As shown in Figure 2, the architecture of the proposed system is divided into three primary modules: Input, Processing, and Output.

1) *Input Module:* The input module provides a web-based platform developed using Streamlit, designed to offer a userfriendly interface for users to upload images through a drag-and-drop feature or file selection component. The platform includes detailed instructions on acceptable file formats and recommended image quality. In addition, example images are provided to guide users in selecting suitable input images for analysis. These measures

ensure that the uploaded images adhere to the quality standards necessary for accurate processing.

2) *Processing Module:* The processing module consists of three key stages: preprocessing, damage detection, and postprocessing.

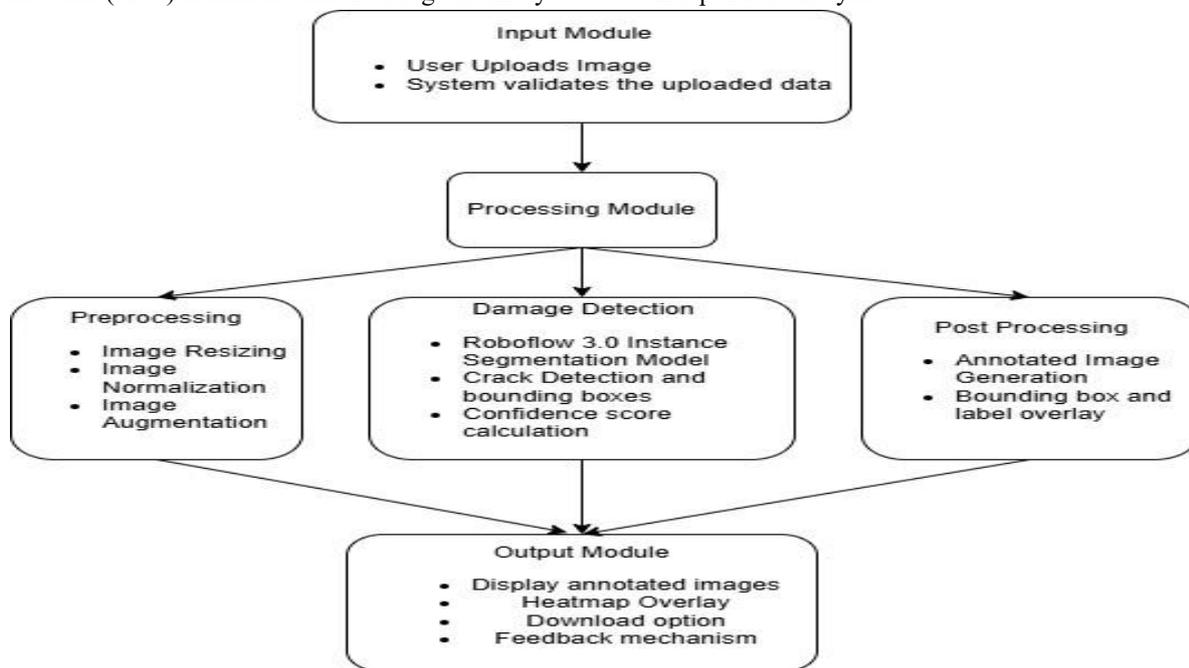
*Pre-processing:* Input images are first converted to grayscale to reduce computational complexity while preserving critical structural details. They are then resized to  $512 \times 512$  pixels to match the input dimensions expected by the Roboflow 3.0 model. Pixel values are normalized to a range of [0,1] to ensure consistency. To improve model generalization and detection accuracy, data augmentation techniques are applied, including rotation ( $\pm 15^\circ$ ), horizontal and vertical flipping, zooming (up to 20%), brightness modulation ( $\pm 10\%$ ), and Gaussian noise injection equal to 0.01.

*Damage Detection:* The pretrained Roboflow 3.0 instance segmentation model processes the preprocessed images using deep convolutional layers to extract relevant features. A region proposal network (RPN) identifies candidate regions likely to

contain structural defects, which are then refined through instance segmentation. Each detected damage region is outlined using bounding boxes and segmentation masks, with confidence scores indicating detection reliability.

*Post-processing:* The system generates annotated images by overlaying bounding boxes, segmentation masks, and confidence scores on the original input images. A distinct color scheme is applied to differentiate damage regions, ensuring clear visualization and facilitating further analysis.

3) *Output Module:* The output module is responsible for presenting the results to the user in a structured and interpretable manner. The annotated images highlighting the detected damages are displayed within the web-based interface. Users can interact with the images through zoom and pan functionalities, allowing for a detailed examination of the detected structural defects. Additionally, an option is provided to download the annotated images, enabling users to retain records or conduct further independent analysis.



4) Fig. 2. General Framework Architecture

### B. Formulas

1) *Precision:* Precision is a critical evaluation metric in crack detection systems. It quantifies the system’s ability to correctly identify true cracks out of all the cracks it has labeled. Mathematically, it is the ratio of True Positives (correctly identified cracks) to the sum of True Positives and False Positives (incorrectly labeled non-cracks as cracks).

In the context of Indian heritage preservation, high precision is essential to prevent misdiagnosis of cracks. Monuments are intricate, and labeling non-damaged features (such as shadows, dirt, or patterns on surfaces) as cracks could mislead conservation efforts. This misdiagnosis not only wastes resources but might also lead to unnecessary repairs that could harm the authenticity of the monument.

$$\text{Precision} = \frac{\text{True Positives}(TP)}{\text{True Positives}(TP) + \text{False Positives}(FP)} \quad (1)$$

2) *Recall*: Recall, also known as sensitivity, measures the system's ability to identify all actual cracks present in an image. It is defined as the ratio of True Positives (correctly identified cracks) to the sum of True Positives and False Negatives (missed cracks). For Indian heritage sites, where structural integrity is paramount, recall plays a significant role in ensuring no damage goes unnoticed. Missing a crack, especially in a structurally critical area, could result in catastrophic failure over time. A low recall score might indicate that the system is too conservative in detecting cracks, which could be detrimental to long-term preservation efforts.

$$\text{Recall} = \frac{\text{True Positives}(TP)}{\text{True Positives}(TP) + \text{False Negatives}(FN)} \quad (2)$$

3) *F1 Score*: The F1 Score is a harmonic mean of precision and recall, providing a balanced measure of the system's performance. It is particularly useful in scenarios where there is an uneven distribution between actual cracks and non-crack features. By combining precision and recall, the F1 Score reflects the trade-off between avoiding false positives and ensuring comprehensive detection. For heritage conservation, an F1 Score ensures the system achieves both accuracy in detecting actual cracks and thoroughness in identifying all existing cracks. A high F1 Score indicates a well-rounded crack detection

$$F1 = \frac{2 \times \text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \quad (3)$$

4) *Intersection Over Union(IoU)*: Intersection over union is a spatial metric used to evaluate the accuracy of bounding boxes drawn around detected cracks. It measures the overlap between the predicted bounding box and the ground truth (actual crack location), relative to their union. In crack detection for Indian heritage sites, IoU is crucial for validating the exactness of the detection. A high IoU ensures that the detected bounding box closely matches the actual crack, reducing ambiguity in assessments.

$$\text{IoU} = \frac{\text{Area of Overlap}}{\text{Area of Union}} \quad (4)$$

5) *Heatmap Intensity*: Heatmaps are generated to visually represent the severity and density of cracks. The intensity of the heat map at any point is based on a Gaussian distribution centered on the coordinates of the detected crack. This approach is beneficial for Indian heritage preservation, as it allows conservationists to quickly assess areas of high damage density. By visualizing the heatmap, conservation teams can prioritize intervention in critical zones, ensuring effective resource allocation.

$$H(x, y) = \exp\left(-\frac{(x - x_0)^2 + (y - y_0)^2}{2\sigma^2}\right) \quad (5)$$

where  $(x_0, y_0)$  represents the crack center,  $\sigma$  determines the spread, and  $H(x, y)$  represents the intensity at any given point. Larger  $\sigma$  values result in broader heatmaps, which are useful for highlighting larger affected areas.

6) *Loss Function in the model (L)*: The crack detection model minimizes an overall loss function comprising classification, bounding box regression, and mask prediction losses. This combination ensures the system identifies cracks accurately, localizes them precisely, and generates high-quality visualizations.

$$L = L_{cls} + L_{bbox} + L_{mask} \quad (6)$$

where:

- *Classification Loss ( $L_{cls}$ )*: Ensures the model correctly classifies areas as cracks or non-cracks.
- *Bounding Box Loss ( $L_{bbox}$ )*: Refines the location of the bounding boxes around cracks.
- *Mask Loss ( $L_{mask}$ )*: Enhances the segmentation quality of crack regions.

These loss components collectively contribute to improving detection accuracy, which is essential for preserving the intricate details of Indian heritage structures.

### C. Implementation of Proposed System

To begin, users select a pre-trained crack detection model from the dropdown menu, choosing the one that best suits their requirements (Figure 3).

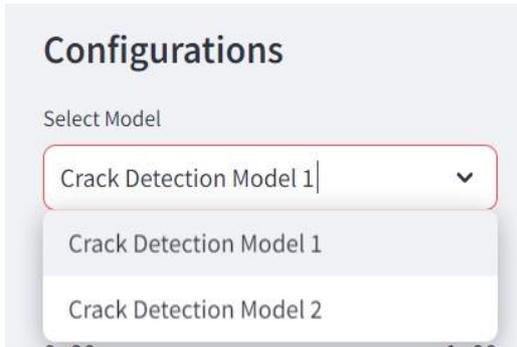


Fig. 3. Select model dropdown.

Next, they adjust the detection sensitivity using the threshold slider, which ranges from 0.00 to 1.00, allowing them to balance sensitivity and precision in crack detection (Figure 4).

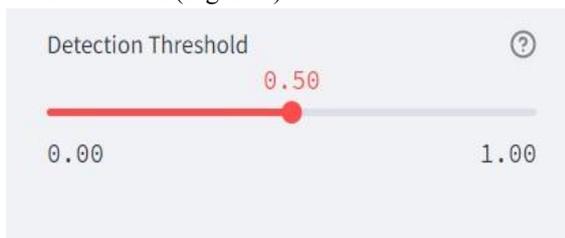


Fig. 4. Adjust detection threshold.

The model information panel provides details about the selected model, including its architecture, training dataset, accuracy, and checkpoint details (Figure 5).

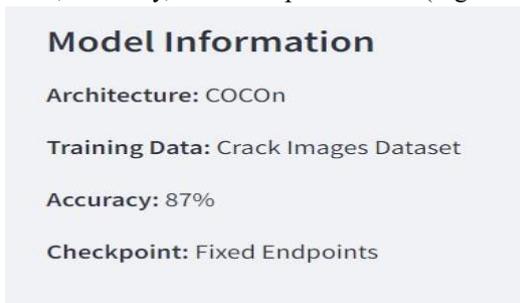


Fig. 5. Model Information.

Users then upload images through the browse and upload images feature, where pre-processing is applied, including grayscale conversion, resizing to 512x512 pixels, normalization, and contrast enhancement (Figure 6).

### Crack Detection

This app allows you to upload images and detect cracks using a pre-trained model.

Upload one or more images to begin.

Choose images...

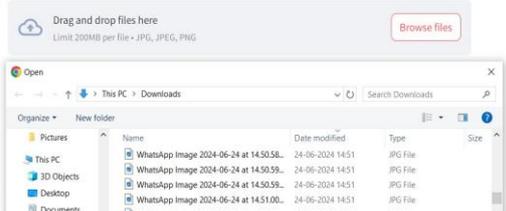


Fig. 6. Browse and upload images.

Once preprocessing is complete, the crack detection algorithm processes the images, generating bounding boxes



Processed image: crack1.PNG

Download Processed Image

Fig. 7. Crack detection with bounding boxes.

help improve the detection algorithms and user interface (Figure 10,11).

that highlight detected cracks along with confidence scores

(Figure 7).

A heatmap overlay is also generated, visually representing the concentration and confidence levels of detected cracks within the image. The processed images, including detected cracks and heatmaps, can be downloaded for further analysis or documentation (Figure 8).



Heatmap Overlay

Download Heatmap Image

Fig. 8. Heatmap Overlay.

The inference result panel provides detailed insights into detected cracks, including coordinates, dimensions, and confidence levels, enabling users to make informed decisions (Figure 9).

### Inference Result

Time: 0.10105412000029901 seconds  
 Image Width: 768  
 Image Height: 768

**Predictions**

**Detection 1**  
 Class: crack  
 Confidence: 0.57  
 Coordinates (x, y, w, h): 369.0, 385.0, 628.0, 762.0  
 Detection ID: e0878885-b537-493f-b4a1-bc701b736c09

Fig. 9. Inference result.

Fig. 10. Feedback and detection statistics

timestamp	image_name	model_id	rating	comments
2024-08-16 17:58:43	crack2.jpg	crack-detection-x9bz6/1	4	Detections made by this system is accurate. I am satisfied by the results

1 entry

Show 10 per page

Fig. 11. Feedback collection

### Detection Statistics

Number of Cracks Detected: 1  
 Average Confidence Score: 0.57

Processing completed for all uploaded files.

### Provide Feedback

Feedback for crack2.jpg

Rate the accuracy of detections for crack2.jpg



Additional comments for crack2.jpg

Detections made by this system is accurate. I am satisfied by the results

Submit Feedback for crack2.jpg

## IV. RESULT

Finally, users can view detection statistics and submit feedback, which is systematically stored in a CSV file to be used using the Darbhanga Fort dataset comprising 3,102 im-

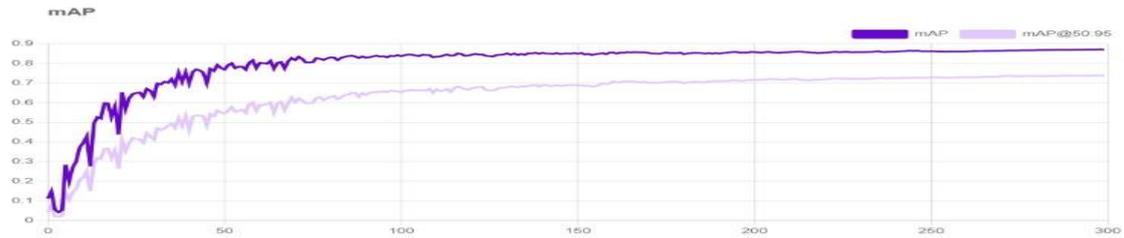


Fig. 12. Mean Average Precision graph.

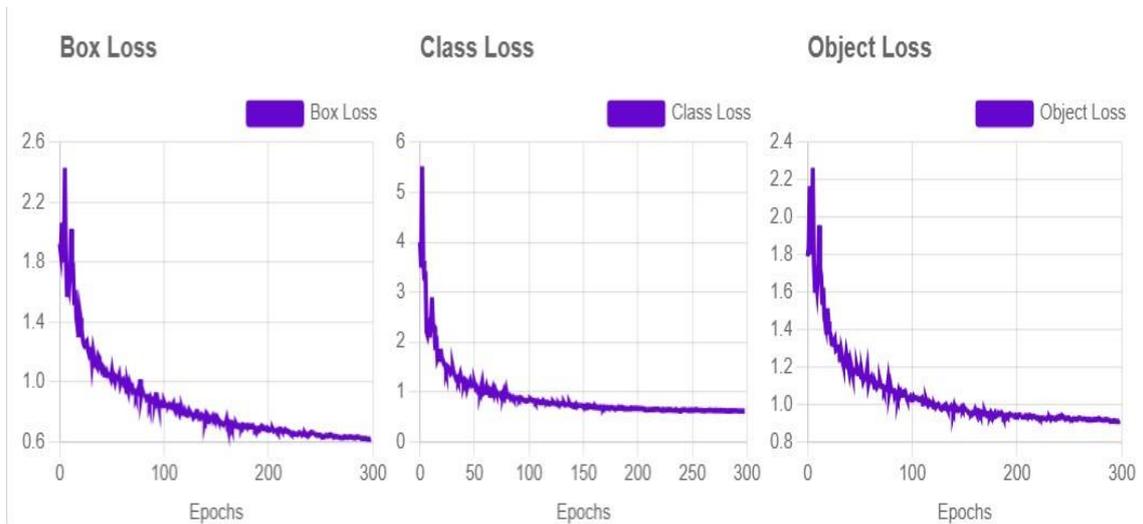


Fig. 13. Box loss, Class loss and Object loss.

ages, divided into 70% training (2,172 images), 20% validation (617 images), and 10% testing (313 images). The dataset included diverse crack patterns, ensuring relevant real-world applicability.

#### A. Training Performance Analysis

Figures 12 and 13 illustrate the model's training progress over 300 epochs, highlighting the convergence of Box Loss, Class Loss, and Object Loss (Figure 13), as well as the progression of Mean Average Precision (mAP) (Figure 12). The loss metrics exhibit a sharp decline within the initial 50 epochs, indicating rapid feature learning, followed by stabilization around epoch 100. By epoch 300, the final loss values remain consistently low, with Box Loss, Class Loss, and Object Loss converging to approximately 0.6, 0.8, and 0.8, respectively. This trend suggests effective feature extraction while minimizing overfitting.

The mAP progression further validates the model's performance, with Figure 12 demonstrating both mAP@0.5 and mAP@50:95 trends. The mAP@0.5 metric rapidly increases to approximately 85% within the first 100 epochs and stabilizes at 87.8% by the final epoch, indicating high detection accuracy. Additionally, the mAP@50:95 metric follows a steady upward trajectory, reaching approximately 79%, underscoring the model's robustness in detecting cracks across varying intersection over union (IoU) thresholds. These results collectively affirm the model's strong generalization capability and effectiveness in real-world crack detection scenarios.

#### B. Quantitative Performance Metrics

Table 1 summarizes the model's evaluation metrics on the test set (313 images).

TABLE I  
EVALUATION METRICS ON TEST SET

Metric	Value (%)
mAP@0.5	87.8
mAP@50:95	79.2
Precision	87.0
Recall	87.6
F1-Score	87.3

The model achieves a Mean Average Precision (mAP@0.5) of 87.8%, with a high F1-score (87.3%), ensuring a balance between Precision (87.0%) and Recall (87.6%).

The model achieves an average inference time of 0.1 seconds per image in a GPU-powered environment, making it highly suitable for real-time structural monitoring. This efficiency ensures quick detection and analysis, which is essential for timely maintenance and safety assessments.

While the model effectively detects cracks, it can be further enhanced to identify other structural issues such as erosion, broken components, and surface degradation. However, challenges remain in handling complex backgrounds and extreme lighting conditions, where minor misdetections may occur. To improve performance, future work will focus on adaptive data augmentation techniques to enhance the model's ability to generalize across different environments. Additionally, finetuning with larger and more diverse real-world datasets will help improve detection accuracy and reliability for various structural defects.

## V. CONCLUSION

This research presents a deep learning-based system for detecting structural damage in heritage sites using the COCO architecture and the Roboflow 3.0 instance segmentation model. With an accuracy of 87.8%, the system is deployed as a Streamlit web application, allowing users to upload images and receive detailed analysis with annotated outputs and heatmap overlays. Traditional crack detection methods, such as edge detection and thresholding, often struggle with complex textures and lighting changes. While CNN-based classifiers and Faster R-CNN models improve accuracy, they do not provide precise instance segmentation. In contrast, this study uses Mask R-CNN, which not only detects cracks but also outlines their exact shape and size for better damage assessment. The system can also be extended to identify other structural issues, such as erosion and surface degradation. By improving detection accuracy and offering a user-friendly web interface, this system builds on existing methods for structural health monitoring and helps conservation professionals detect damage early, ensuring better preservation of historical monuments.

## REFERENCES

- [1] M. Zaccarini, A. Iannucci, M. Orlandi, M. Vandini, and S. Zambruno, "A multi-disciplinary approach to the preservation of

- cultural heritage: A case study on the piazzetta degli ariani, ravenna,” in *2013 Digital Heritage International Congress (DigitalHeritage)*, Marseille, France, 2013, pp. 337–340.
- [2] J. Smith and A. Johnson, “Deep learning approaches for damage detection in structural engineering,” *Transactions on Neural Networks and Learning Systems*, vol. 30, no. 5, pp. 1234–1245, May 2019.
- [3] L. Wang, Y. Zhang, and Q. Liu, “Instance segmentation for historical conservation: Methods and applications,” *Journal of Computational Vision and Imaging Systems*, vol. 22, no. 2, pp. 234–245, Feb. 2021.
- [4] A. Belhi, H. Gasmi, A. K. Al-Ali, A. Bouras, and S. Fougou, “Deep learning and cultural heritage: The ceproqha project case study,” *Transactions on Neural Networks and Learning Systems*, vol. 33, no. 8, pp. 4560–4572, Aug. 2022.
- [5] J. Paul, S. Ghose, K. Aggarwal, N. Nethaji, S. Pal, and A. D. Purkayastha, “Machine learning advances aiding recognition and classification of indian monuments and landmarks,” in *IEEE 8th Uttar Pradesh Section International Conference on Electrical, Electronics and Computer Engineering (UPCON)*, Dehradun, India, 2021, pp. 1–8.
- [6] A. Patel, R. Kumar, and V. Singh, “Coco-based instance segmentation techniques for heritage site preservation,” *Transactions on Image Processing*, vol. 28, no. 11, pp. 4567–4580, Nov. 2019.
- [7] R. Pathak, A. Saini, A. Wadhwa, H. Sharma, and D. Sangwan, “An object detection approach for detecting damages in heritage sites using 3-d point clouds and 2-d visual data,” in *Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2022, pp. 789–800.
- [8] U. Kulkarni, M. S. Meena, S. V. Gurlahosur, and U. Mudengudi, “Classification of cultural heritage sites using transfer learning,” in *2021 International Conference*, vol. 9, 2021, pp. 45678–45688.
- [9] P. Y. Pawar and B. S. Ainapure, “Image dataset of pune city historical places for degradation detection, classification, and restoration,” *Journal of Image and Vision Computing*, vol. 38, no. 5, pp. 123–134, May 2023.
- [10] Anshul, A. B. Ashish, A. N. Yeole, G. P. M. S., and A. P. B. P., “Cnnpowered monument detection webapp for preserving india’s cultural heritage,” *Transactions on Image Processing*, vol. 29, no. 9, pp. 4561–4572, Sep. 2021.
- [11] S. B. Ali, R. Wate, S. Kujur, A. Singh, and S. Kumar, “Wall crack detection using transfer learning-based cnn models,” in *IEEE 17th India Council International Conference (INDICON)*, New Delhi, India, 2020, pp. 1–7.
- [12] K. Patel, N. Sharma, and R. Gupta, “Preprocessing and automated image analysis techniques for heritage site preservation,” *Transactions on Computer Vision and Pattern Recognition*, vol. 32, no. 7, pp. 3456–3468, Jul. 2021.
- [13] M. Brown and C. Davis, “Web-based applications for automated damage detection: A review,” *Transactions on Automation Science and Engineering*, vol. 17, no. 4, pp. 1012–1023, Oct. 2020.
- [14] M. Casillo, A. Lorusso, F. Colace, F. Marongiu, B. B. Gupta, and D. Santaniello, “A deep learning approach to protecting cultural heritage buildings through iot-based systems,” *Transactions on Industrial Informatics*, vol. 16, no. 7, pp. 3124–3134, Jul. 2020.
- [15] K. Phelan and D. Riordan, “Detection of ringforts from aerial photography using machine learning,” in *31st Irish Signals and Systems Conference (ISSC)*, Letterkenny, Ireland, 2020, pp. 1–6.