



DOI: 10.5281/zenodo.20609956

DEEP LEARNING-BASED CRACK DETECTION AND STRUCTURAL DEFECT MAPPING IN REINFORCED CONCRETE SYSTEMS

¹Dr. SriSudha Garugu, ²Dr. Tadi. Chandrasekhar, ³Vikas Gulati, ⁴Dr. Aasheesh Raizada, ⁵Anup Saxena, ⁶Dr. Vimal Bibhu, ⁷Ranjan Banerjee

¹Department of Computer Science & Engineering, ACE Engineering College, Anukhushapur Ghatkesar, JNTUH, Hyderabad, Email ID: srisudha.garugu@gmail.com, Orcid ID:0000-0002-6182-0155

²Assistant professor Department of AI & ML Aditya University, Surampalem, Andhra Pradesh
Email ID: ramyoga.2011@gmail.com Orcid ID: 0000-0002-1945-2977

³Associate Professor, Department of Mechanical Engineering, Specialization: Industrial Engineering, Lovely Professional University, Jalandhar-144022, India, Email ID: vikasgulati1979@gmail.com,
Orcid ID: 0000-0002-9917-7679

⁴Associate Professor, CDOE Mangalayatan University, Aligarh-202146, Uttar Pradesh, India,
Email ID: aasheesh.raizada@mangalayatan.edu.in

⁵Research Scholar, Department of Computer Science & Engineering, Specialization: Computer Science & Engineering, Noida International University, GB Nagar, Uttar Pradesh, India, Email ID: saxenaanup@gmail.com

⁶Professor Department of Computer Science & Engineering Specialization: Computer Science & Engineering, Noida International University, GB Nagar, Uttar Pradesh, India, Email ID: vimal.bibhu@niu.edu.in

⁷Assistant Professor, Department of Computer Science & Engineering, Brainware University, Barasat, Kolkata-700125, West Bengal, India, Email ID: rbkpcst@gmail.com, ORCID ID : 0009-0003-1950-7530

Abstract

The structural support of the built environment is made of reinforced concrete systems, but their overall performance is jeopardized by the progressive surface deterioration that is challenging to evaluate at any given time with the help of manual inspection. This paper will suggest a deep learning-based architecture of crack detection and mapping structural defects in reinforced concrete infrastructure with an emphasis on pixel-level semantic segmentation and spatial explanations. Multi-class annotation of high-resolution concrete images was done by a U-Net based convolutional neural network trained under class-imbalanced conditions. This method blends the processing of the data, optimization of weighted losses, quantitative results of Dice and Intersection-over-Union metrics, and mapping defect areas that can be interpreted as representations of the surface conditions. Experimental findings prove that the framework is effective to segment visually dominant defects like corrosion, spalling, and exposed reinforcement whereas fine cracks can be difficult to segment because of their sparse geometry. The systematic patterns of over- and under-estimation are seen in the analysis of defects areas, which demonstrates that predictions should be interpreted as relative values, but not as absolute ones. The research places semantic segmentation as a powerful decision support method. It aids architectural inspection, documentation of conditions, prioritization of maintenance, and development of trusted digital processes of reinforced concrete infrastructure system in the world.

Keywords: Reinforced concrete inspection, deep learning, semantic segmentation, structural defect mapping, crack detection.

1. Introduction

The construction of modern and historic built environments relies on reinforced concrete structures as the foundation of residential buildings, infrastructure, transport systems, and industrial

facilities. Their popularity is explained by the complementary mechanical action of concrete and steel reinforcement that can be used to achieve effective resistance to compressive and tensile forces. Although concrete systems are durable, there is

progressive deterioration due to exposure, ageing of materials, construction shortcomings, and continuous loading, which can affect the reinforced concrete systems. Superficial flaws like cracking, stain corrosion, spalling and bare reinforcement are observable signs of this corrosion and are the focus of structural condition evaluation.

Manual visual inspection of reinforced concrete buildings is mainly based on the traditional methods of visual inspection of the building by trained experts. This method is subjective, lacks uniformity among different inspectors, poses safety risks due to challenging access conditions, incurs high time and labor expenses, which, despite its widespread use, is a limitation of the method. Conventional inspection practices are becoming more and more challenged in scalability and effectiveness as the urban infrastructure networks grow and age at the same time. Such restrictions have prompted the accumulating enthusiasm of automated inspection methods that entail digital imaging with artificial intelligence to aid in structural evaluation and upkeep decision-making (Laxman et al., 2023).

The recent developments in deep learning have greatly increased the ability of the computer vision systems to detect concrete defects. Convolutional neural networks have been shown to be effective at learning the complex visual features directly on the image data and can be used to detect the cracks under different lighting, texture, and background conditions (Kolappan Geetha et al., 2023). Earlier researchers have effectively used deep learning to detect cracks, predict crack depth, and estimate the severity of the damage, which emphasizes the possibility of data-driven solutions to enhance the reliability and efficiency of the inspection process (Mir et al., 2022).

But most of the literature available is devoted to binary crack detection in which image locations are classified as either cracked or intact. Although useful in determining the existence of cracks, the methods give a poor reflection of the deterioration of reinforced concrete. As a matter of fact, cracks tend to be present with corrosion staining, spalling and reinforced opening, which are interconnected degradation processes in reinforced concrete structures. Some of them admit this complexity but still focus on defect presence instead of spatial extent, distribution, or interaction between different types of defects (Han et al., 2022; Xu et al., 2023).

The other weakness of most deep learning-based inspection research is that they focus on classification accuracy with little or no consideration of spatial interpretability. Architecturally and built-environmentally, automated inspection systems can be made useful not only at identifying defects, but also at displaying their spatial distributions on structural surfaces. The representations of the defects

at a pixel level including the location, extent, and co-occurrence of the defects are especially important in documentation, maintenance planning, and stakeholder communication. However, the research on semantic segmentation as a method of generating interpretable defect maps that can be used in architectural and infrastructure applications is comparatively small (Golding et al., 2022; Nguyen et al., 2023).

The recent surveys of research on crack detection using deep learning indicate the prevalence of strategies focused on crack detection, as well as a shortage of studies based on multi-defects segmentation (Golding et al., 2022; Nguyen et al., 2023). Although other studies have broadened the scope to encompass other types of defects, most studies are based on image-level classification or bounding-box localization, which fails to capture fine-grained spatial features of surface damage (Wang et al., 2022). There are new works on a higher representation like three-dimensional point clouds and multimodal sensing, but the high cost of these applications may require special devices and high-performance computers, which makes it difficult to apply them in everyday inspection settings (Chen et al., 2025).

Multi-defect mapping is also a difficult task because of the visual variability of the reinforced concrete surfaces. The variations in texture, weathering, staining, and the quality of construction add a lot of complexity, which influences the model generalization. Furthermore, some defects, specifically fine cracks, take a very small share of the image pixels, and thus causes a serious imbalance of classes in training the model. Unless it is taken into consideration, this imbalance may skew learning in favor of visually dominant defects, making it less sensitive to damage at an early stage (Hacıfendioglu et al., 2023).

The requirement of well-developed and easily understandable inspection tools is particularly acute in the framework of the architectural heritage and the aging building stock. The heterogeneous patterns of deterioration in older reinforced concrete structures are usually influenced by age-long exposure to environmental factors and history of maintenance. Non-invasive visual inspection systems with the ability to map various types of defects on the surface level provide a non-invasive method of recording the condition and help with conservation and maintenance decision-making (Hu et al., 2024). Simultaneously, these systems should be well positioned as a decision-support tool but not as an alternative to expert judgment.

To address them, the present study suggests a deep learning-based architecture of crack detection and structural defect mapping in reinforced concrete structures with a focus on pixel-level semantic

segmentation and visual representation. In contrast to the earlier crack-centered models, the framework concurrently subdivides several categories of defects that are typically related to the deterioration of reinforced concrete, such as cracks, corrosion, spalling, and uncovered reinforcement. This is a multi-class model that facilitates a more detailed description of surface condition and interaction of defects.

The study is deliberately narrowed down to the visual analysis of surface-levels based on RGB imagery. The suggested framework is not intended to identify internal or embedded damage, and such damage usually has to be evaluated using nondestructive methods like ultrasonic or tomographic (Alqurashi et al., 2025). Rather, it supplements the current inspection methods with surface defect maps that are consistent, repeatable and interpretable using easily accessible image data. The shortcomings of the dataset size, class imbalance and the natural challenge of fine-scale crack segmentation are recognized.

Irrespective of these shortcomings, the research is useful in the development of spatially interpretable defect mapping of reinforced concrete systems. The framework allows appropriating the algorithmic prediction to the architectural interpretation by converting the segmentation results into visual supermasters and quantitative defect indicators. These representations facilitate condition documentation, prioritization during maintenance and assist digital workflows in the built environment development. In a more general sense, the study is consistent with the interdisciplinary attempt to make artificial intelligence a visual mediation tool that complements the practice of assessing architecture and infrastructure (Amirkhani et al., 2024; Ali et al., 2022).

To this end, the proposed research aims at achieving three main goals, namely, creating a deep learning-based semantic segmentation system that can detect cracks on reinforced concrete surfaces at the pixel level; extending the same system to multi-class structural defect mapping, enabling the detection of cracks, corrosion, spalling, and exposed reinforcement simultaneously; and generating interpretable visual defect maps and quantitative measurements of defect areas to support architectural inspection and assessing the reinforced concrete systems (Alkannad et al., 2025).

2. Literature Review

The recent developments in the field of deep learning have had a great impact on automated inspection of reinforced concrete structures, especially in crack detection and surface damage analysis. Initial experiments with image-based deep learning showed that convolutional neural networks were

able to detect cracks on concrete bridge surfaces in different illumination and texture conditions, which was more effective than the traditional image processing methods (Wan et al., 2022). Such studies proved the viability of data-driven inspection methods, but most of them were done in binary crack detection and did not consider larger deterioration patterns found in reinforced concrete structures.

Later studies extended the use of deep learning to more complicated defects, such as detecting and predicting the defects in filling concrete. These research papers have emphasized the ability of the deep learning models to identify delicate visual features related to internal and surface-level anomalies, which means that they have the capacity to make improved damage assessments (Yang et al., 2025). However, these methods were typically limited to the particular defects and failed to broaden to detailed defect mapping on reinforced concrete systems.

A number of review articles summarised the ever-increasing amount of literature on deep learning-based crack detection and found that there are still significant methodological shortcomings. Systematic reviews highlighted the fact that encoder-decoder networks and fully convolutional networks were better at localization error, but the majority of studies were limited to the detection of a single defect, usually cracks (Hamishebahar et al., 2022). Further, performance measurement was often restricted to measures of accuracy, and little attention was paid to spatial interpretability or usefulness of the measure to structural evaluation.

Attempts to include quantitative damage assessment were made in research that concentrated on automated detection and measurement of the damage in concrete pavements. These systems integrated segmentation results with damage quantification values, including defect area and severity indices, which showed the advancement in the direction of infrastructure management applications (Garita-Duran et al., 2025). Regardless of these developments, these frameworks tended to be asset-focused and failed to capture the overall complexity of reinforced concrete deterioration where there are many interacting defect types.

The ability of deep learning to describe crack morphology was further investigated in material-oriented studies. The results of crack texture feature identification in fiber-reinforced concrete demonstrated that deep learning models were able to discriminate the material-specific crack patterns, which highlights the significance of the domain-sensitive training data (Zhou et al., 2022). On the same note, multispectral imagery, combined with deep learning methods, helped to better identify cracks and their severity, especially fine cracks that are hard to image with standard RGB imagery (Fan,

2025). Although efficient, such approaches have provided new sensing needs that can demarcate scalability and mass usage.

Another important development in automated inspection was the combination of deep learning and unmanned aerial vehicles. UAV platform-based real-time inspection systems proved to be effective in detecting cracks on large concrete structures, which underscores the viability of the operations in infrastructure monitoring (Kim et al., 2024). These systems were however mainly focused on crack presence and lacked detailed multi-defect spatial mapping and full representation of surface conditions.

In a broader inspection point of view, systematic analysis of reinforced concrete bridge defects has highlighted the interdependence of cracking, corrosion, spalling and exposure to reinforcement. Through such reviews, it was noted that structural deterioration was to be viewed as a whole and not as individual defects, and it showed that there was a lack of connection between the practice of inspection and most AI-based research (Abdelkader et al., 2023). This finding was further supported by surveys of crack detection methods based on image processing and machine learning, which indicated that most of the methods focused on the detection accuracy, rather than the spatial representation and interpretability (Kirthiga & Elavenil, 2024). Deep learning models using the vision were very good at controlling datasets, but the results were hard to translate into practical condition measurements that could be applied to real-life scenarios (Nabizadeh and Parghi, 2023).

Recent methodological developments tried to combine detection and assessment by including crack identification and damage assessment measures. Segmentation-based deep learning models with damage assessment showed better structural understanding but were still mostly concentrated on cracks and not various defect types (Guo et al., 2024). Comparative analysis of convolutional neural network models proved that encoder decoder models are always better at pixel level crack segmentation and especially in high-resolution images (Seol et al., 2025).

The transfer learning methods enhanced generalization in the detection of concrete bridge defects, which points to the advantages of using pre-trained models on different datasets (Zoubir et al., 2022). However, these studies often used classification or coarse localization methods, and thus their ability to reflect defect extent and interaction. In the same way, the damage classification of reinforced concrete bridges using deep learning had high classification accuracy but had poor spatial resolution of defect distribution (Abubakr et al., 2024).

In general, the literature shows that there has been significant advancement in automated crack detection but there remains a gap in multi-defect, pixel-level mapping structures that can generate interpretable depictions of reinforced concrete deterioration. This gap is what directly drives the current study in terms of interest in integrated structural defect mapping through semantic segmentation.

3. Methodology

3.1 Research Design

The proposed research design is a quantitative, experimental research design, based on the computer vision and architectural visual analysis, with the aim of creating and proving a deep learning-based framework of crack detection and multi-class structural defect mapping in reinforced concrete structures. The study lies between artificial intelligence, architectural imaging, and infrastructure inspection, and focuses on the pixel-level spatial representation as opposed to binary defect classification.

The general methodological approach takes a pipeline approach that starts with the acquisition and preparation of high-resolution reinforced concrete surface imagery and then uses supervised deep learning-based semantic segmentation to detect various defect classes. The results of the segmentation are then assessed quantitatively with standard measures of performance to determine the accuracy of segmentation and class behavior. Lastly, the findings are converted to spatial defect maps which can be used to interpret the distribution, extent and co-occurrence of defects to facilitate structural assessment and architectural inspection processes.

In contrast to the previous literature, which is mostly aimed at detecting the presence of cracks, the study is specifically aimed at generating interpretable surface-level defect maps that visually and quantitatively map various deterioration phenomena, including cracks, corrosion, spalling, and uncovered reinforcement, in reinforced concrete systems. The design thus focuses on the spatial precision, visual readability, and mapping in accordance with the architectural inspection and built-environment management processes.

The controlled conditions of the research are also experimental in nature, which facilitates reproducibility and systematic comparison of model behavior when preprocessing, training, and validation are done under uniform conditions, which also allows controlled evaluation of model behavior.

3.2 Data Collection Methods

3.2.1 Dataset Source

The High-Resolution Concrete Defect Segmentation (HRCDS) dataset, which is a publicly-available

benchmark dataset, was used to extract data to this study specifically designed to segment defects on surfaces of reinforced concrete at the pixel-level (Guo & Bao, 2025). The data set consists of RGB images of concrete elements at high resolution with a variety of surface defects that may occur in structural and architectural practice. Each of the images is provided with a manually annotated ground-truth mask, in which pixels are marked depending on the type of defect. There are grayscale and color-coded segmentation masks in the dataset, so the dataset needs thorough preprocessing to make sure that the classes are labeled consistently across the samples.

3.2.2 Image Characteristics

The images of the HRCDS represent actual reinforced concrete surfaces in the real world with different light conditions, surface texture, and the degree of damage. The common resolutions of images are more than 700x 1000 pixels, and fine-scale defects like hairline cracks and corrosion staining can be represented visually. To make computations and consistency possible, both images and matching masks were resized to 256x256 pixels when training and evaluating the models.

3.2.3 Defect Classes

In this study, the semantic classes of reinforced concrete elements surface conditions were taken into account and amounted to five semantic classes. These were background, which signified intact concrete surfaces with no visible defects, crack, which signified linear or branching cracks that signified stress or material breakdown, exposed reinforcement, which signified visible steel bars that were caused by the loss of concrete cover, corrosion, which signified rust stains and surface discoloration associated with the oxidation of steel and spalling, which signified localized break-off of concrete exposing underlying layers of material.

A combination of these classes is a way to interpret the condition of the systems based on reinforced concrete, which is progressive deterioration, and can be viewed on a system-level, as opposed to the local identification of defects.

3.3 Population and Sampling Strategy

3.3.1 Population Definition

The sample of the study includes surface images of reinforced concrete elements, which have visible deterioration, which is typical of buildings, infrastructure, and civil assets in the actual working conditions. Even though the dataset itself does not reflect a single geographic location, it represents a wide spectrum of concrete surface conditions that can be experienced in practice, which makes it possible to generalize it at the level of surface-inspection.

3.3.2 Sampling Approach

To divide the dataset into the training, validation, and test sets, a random stratified sampling approach was used. Out of total 1200 confirmed image-mask pairs, 70 percent (867 images) of the data were assigned to the training set, 15 percent (153 images) to the validation set, and the rest 15 percent (180 images) to the test set with even distribution of the defect classes across all subsets.

A fixed seed was used to randomize in order to achieve reproducibility. All the splits were mutually exclusive, so that no image was represented in more than one subset. This sampling approach enables objective assessment of the generalization of the models and sufficient representation of all the defect classes in subsets.

Because of the high class imbalance of the real concrete inspection data (especially the low frequency of cracks), sampling was done on the image level, not pixel balancing, and kept realistic distributions of defects.

3.4 Data Preprocessing and Preparation

Before training the models, all the images and masks were standardized through preprocessing. First, images and masks were downsampled to 256x256 pixels, with bilinear interpolation applied to RGB pictures and nearest neighbour interpolation applied to masks in order to keep discrete class values. Second, RGB images were brought to floating-point tensors in [0, 1] range to guarantee the numerical stability of the training process. Lastly, simple geometric data augmentation methods, such as horizontal and vertical flipping, were used to augment the training data with a 0.5 probability, and no data augmentation was used on the validation or test data.

Special focus was on subtle mask forms. Classes IDs represented as grayscale masks and RGB masks represented as color palettes were automatically transformed into single representations as class-IDs, using a pre-defined color-to-class mapping. Those pixels whose color was uncertain or unrecognizable were assigned to the background category conservatively, which does not add false labels of defects.

3.5 Data Analysis Techniques

3.5.1 Model Architecture

The U-Net convolutional neural network was chosen as the main part of the segmentation model because it has proven to be effective in pixel-level semantic segmentation problems especially in situations where little training data is available and the spatial detail is fine. The architecture put in place consists of encoder path which is followed by successive convolutional blocks and max-pooling layers to

extract hierarchical features, a bottleneck layer to extract global contextual information and a decoder path that uses transposed convolutions and skip connections to merge high level semantic features with low level spatial information in order to localize precisely.

It was set up with around 483,000 trainable parameters, which makes the model computationally efficient and can be trained on a CPU-based training environment, although it has enough representational capacity.

3.5.2 Training Strategy

The training of the model was done with a supervised learning method using a weighted cross-entropy loss to overcome the imbalance of the classes of defects. The Adam optimizer was used with a learning rate of 1×10^{-3} to optimize. The batch size of two was used to fit the CPU memory constraints and a total of five epochs was used to train the model. In order to deal with extreme class imbalance, especially under-representation of cracks, pixel-wise frequencies of classes were determined using the training data, and inverse-frequency weights on classes were added to the loss. This approach is the one that promotes the model to focus on the occurrence of the rare but structurally important defects without the artificial manipulation of the data distributions.

The model checkpoint with the largest validation mean Intersection-over-Union (mIoU) without background was kept to be evaluated finally.

3.5.3 Evaluation Metrics

The quantitative method of segmentation performance was based on the Dice coefficient to assess the overlap between the predicted and ground-truth regions and the Intersection over Union (IoU) to determine the spatial agreement. Besides that, the mean Dice and mean IoU values were also calculated without the background class to summarize the defect-specific performance. All the measures were estimated on the held-out test set to provide an unbiased assessment of the overall generalization power of the model.

3.5.4 Structural Defect Mapping Analysis

In addition to the conventional segmentation measures, the analysis presents a defect-area mapping analysis that can be used to interpret the structure. The percentage of pixels of each defect type in the ground truth and predicted mask of each test image was computed. These values were summed up to obtain mean and median defect-area percentages in each of the classes.

Such an analysis offers the comparison of actual and predicted spatial extents of defects and can help to understand systematic over- or under-estimation patterns and justify the idea of structural defect mapping as opposed to detection..

3.6 Ethical Considerations

The study will be based only on publicly available, non-identifiable image data of concrete surfaces. There were no human subjects, personal information, or sensitive infrastructure data. In this regard, the research cannot raise any issues of privacy, consent, or human-subject ethics. In terms of methodological ethics, the transparent reporting of the model limitations and variations of the performances by the classes were carefully considered. The research did not exaggerate the predictive accuracy or purported practical readiness and a clear line was drawn between experimental validation in controlled settings and possible real-world implementation settings.

The use of all datasets is in accordance with the licensing terms of the original dataset and proper attribution is done. The study is aimed as a decision-support and visualization tool rather than structural assessment by professionals.

4. Experimental Results

4.1 Overview of Experimental Evaluation

This part provides the experimental findings of the implementation of the proposed deep learning architecture on the reinforced concrete dataset on HRCDS. The analysis is designed based on three complementary dimensions, namely, quantitative segmentation performance, qualitative spatial defect mapping, and defect-area distribution analysis, which allows assessing both predictive performance and spatial interpretability in a comprehensive manner. The results are discussed in the context of visual interpretability, spatial consistency, and mapping behavior instead of focusing on the accuracy at the benchmark level, which reflects the goal of the study to facilitate the workflow of the architectural inspection and reinforced concrete condition assessment.

4.2 Quantitative Segmentation Performance

4.2.1 Per-Class Dice and IoU Performance

Table 1 shows the per-class Dice coefficient and Intersection over Union (IoU) values calculated on the test set held out on each defect category without including the background class in the mean aggregation.

Table 1. Per-class segmentation performance on the test set

Defect Class	Dice	IoU
--------------	------	-----

Crack	0.086	0.070
Exposed reinforcement	0.360	0.272
Corrosion	0.388	0.357
Spalling	0.292	0.246
Mean (no background)	0.281	0.236

The findings show a definite variation in performance based on the classes. Corrosion has the highest overlap scores, then exposed reinforcement and spalling, and finally, crack detection has significantly lower accuracy because of the thinness, discontinuity, and low contrast of crack patterns.

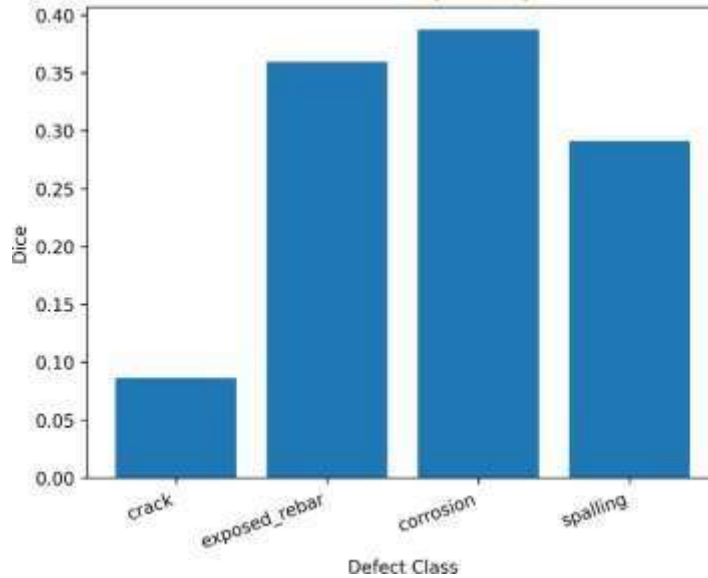


Figure 1. Per-class Dice coefficients on the test set.

Figure 1 demonstrates the significant difference in the performance of segmentation of defects of different types, where corrosion and exposed reinforcement scored significantly higher in Dice compared to cracks. The low Dice score of cracks is illustrated by the fact that it is harder to segment thin, discontinuous features correctly, and visually dominant defects have more consistent spatial overlap.




Figure 2. Per-class Intersection over Union (IoU) scores on the test set.

Figure 2 indicates that corrosion has the highest level of IoU followed by exposed reinforcement and spalling, so that visually dominant defects will be reliable in terms of spatial overlap. Conversely, the significantly

smaller IoU of cracks shows the difficulty of the correct segmentation of the thin and discontinuous crack patterns in the high-resolution concrete images.

4.2.2 Learning Behavior and Model Convergence

The learning property of the segmentation model is described as a continuous decrease in training loss with moderate changes in validation loss, which implies that it has stable optimization in the case of class-imbalanced and limited training cycles.




Figure 3. Training and validation loss curves across epochs.

Figure 3 shows that the optimization of the stable model has a decreasing training loss with the number of epochs. The average change in the validation loss is due to the effect of the class imbalance and the number of training iterations whereas the general trend reflects successful learning without serious overfitting. The loss of training has a steady decreasing pattern, which indicates the optimization of network

parameters. Conversely, the validation loss has a medium variance, and it can be explained by the high degree of class imbalance, small batch sizes, and few training epochs. In order to further evaluate the generalization behavior, the development of validation mean IoU (without background) in epochs was evaluated, which showed gradual improvement and then a minor performance saturation.




Figure 4. Validation mean IoU (excluding background) across training epochs.

Figure 4 indicates that the validation mIoU at first varies irregularly because of the imbalance of classes and insufficient training data, but it steadily increases, which reaches its maximum at the fourth epoch. This trend shows that there is good learning of multi-class defect representations with slight performance saturation thereafter.

The validation mIoU reaches its highest point at the fourth epoch (≈ 0.236) and then it decreases slightly, indicating the slight overfitting. Based on this, the epoch four model checkpoint was chosen to final test-set test.

4.3 Qualitative Visual Results and Spatial Defect Mapping

4.3.1 Visual Comparison of Ground Truth and Predicted Maps

The qualitative analysis of the predicted results of segmentation techniques shows that the framework presented yields consistent and decipherable maps of spatial defects. The results of visual overlaid of predicted masks of original images have a high correlation with ground-truth corrosion, exposed

reinforcement, and spalling annotations, whereas the corpus predictions of cracks are more fragmented and diffuse.

Even though the cracks are not outlined with a high boundary precision, the predicted crack areas are always used to indicate areas of surface discontinuity. These qualitative observations are used to complement the quantitative results and emphasize on the significance of spatial interpretation in assessing the outputs of segmentation.

4.4 Defect-Area Mapping and Distribution Analysis

4.4.1 Mean Defect Area Comparison

One of the key contributions of the study is the investigation of the distribution of defect-areas based on segmentation results. Table 2 gives a summary of the mean percentages of defect areas of each defect class between ground-truth annotations and model predictions.

Table 2. Mean defect-area percentages (test set)

Defect Class	GT Mean (%)	Pred Mean (%)
Crack	0.24	2.85
Exposed reinforcement	17.82	14.35
Corrosion	1.98	6.23
Spalling	1.38	2.04

The differences between the estimated extent of defects of the classes are compared by comparing the areas of defects found by ground-truth annotations and model predictions.




Figure 5. Comparison of mean defect-area percentages between ground truth and predicted results on the test set.

Figure 5 demonstrates an overall bias in the annotated and predicted defects extents, whereby crack and corrosion regions are overestimated, and

exposed reinforcement is underestimated slightly. These patterns suggest a steady spatial mapping behaviour and not a random error, which validates

the application of predicted defect areas in relative condition evaluation of reinforced concrete systems.

4.4.2 Observed Mapping Patterns and Biases

A number of uniform mapping patterns can be observed as a result of the defect-area analysis. The weaknesses of segmentation networks are heavily over-forecasting the areas of crack compared to the ground truth, which is an indication of the bias towards predicting larger areas around slender linear structures. There is slight underestimation at the exposed reinforcement areas, which points to the conservative boundary delineation. Corrosion and spalling are strongly over-predicted because they are sensitive to variations on the surface and texture. These behaviors are systematic mapping behaviors and not random error and thus important in the interpretation of architectural and infrastructure analysis.

4.5 Interpretation of Structural Defect Mapping Results

The defect-area analysis shows that the framework proposed can be used to examine surface conditions relative and not with metrological accuracy. The created spatial maps disclose areas of the high concentration of deterioration, the presence of cracks and the corrosion staining, and the gradual shift of cracking to spalling and reinforcement exposure.

These trends represent familiar deterioration procedures in reinforced concrete systems and condition documentation of maintenance, maintenance priorities and enhancement of visual inspection.

5. Discussion

This study has shown that semantic segmentation based on deep learning can be successfully applied to detect cracks and map structural defects of multiple classes in a reinforced concrete structure, especially where the goal is not to detect cracks and structural defects in binary but to map them spatially. The performance tendencies, visual representations of the qualitative results, and the defect-region analysis offer the understanding of the strengths and limitations of the use of pixel-based deep learning models to evaluate the surface conditions in the built environment.

Quantitative analysis demonstrated the evident performance differences based on the classes. Corrosion and exposed reinforcement were visually salient defects that scored higher on Dice and IoU scores than cracks, which suggests that spatially continuous and high-contrast defects are more dependable to segment. This acts in line with the previous evidence that demonstrates lower accuracy of fine cracks because they are sparsely

geometric and low-density pixels on high-resolution imagery (Laxman et al., 2023). The crack IoU was therefore relatively low in this study, and it cannot be attributed to model inadequacy, but to the fact that this is a difficult task at the fine scale of crack segmentation.

Qualitative analysis also revealed that the suggested framework is able to produce spatially consistent and interpretable defect maps with only moderate overlap-based metrics. Graphical overlaid images revealed that crack forecasts tend to coincide with real crack paths but are spatially enlarged leading to overestimation of the area. Other studies on semantic segmentation have also made the same observations, with the networks in question considering continuity more important than preciseness in linear features (Golding et al., 2022). The architectural inspection viewpoint would find such behavior beneficial, since it draws attention to areas of possible degradation instead of single pixels, and thus assists in conservative condition determination.

The defect-area mapping analysis is one of the contributions of this work. The comparison of the areas of predicted and ground-truth defects was systematic with similar trends such as overestimation of cracks and conservative segmentation of the exposed reinforcement. These trends are consistent with the findings in damage quantification studies, in which segmentation models are likely to smooth or broaden uncertain boundaries when there is an imbalance in classes (Garita-Duran et al., 2025). Instead of making practical usefulness invalid, these trends underline the value of using model outputs as relative measures of surface condition but not as an exact measure, intensive to realistic application as an inspection procedure.

The current framework is an improvement of the current state of practice in terms of deep learning-based inspection studies, as it incorporates various defects of reinforced concrete into a single segmentation model. Although a significant number of the existing studies are based solely on crack detection, commonly assuming that cracks are distinct phenomena, the proposed method explicitly addresses the fact that cracking, corrosion, spalling, and exposure to reinforcement are all interconnected, and patterns of deterioration can be interpreted on a system-wide level. This comprehensive description is closer to the professional structural assessment practices (Abdelkader et al., 2023).

Regardless of these strengths, there is a number of limitations that should be identified. The research is also restricted to the surface RGB imagery and fails to record the internal or subsurface damages, which need non-destructive methods of sensing.

Moreover, the imbalance of classes (especially cracks) limits the accuracy of achievable segmentation even when weighted loss functions are used. Lastly, the model training was performed with limited computational resources and limited number of epochs, which can limit additional performance optimization.

The future research efforts should then focus on the expansion of datasets with more variation in surface conditions and defect severity, multi scale architecture to enhance sensitivity to fine cracks, and temporal data to enhance deterioration monitoring. The integration of defect mapping results into electronic inspection systems or building information modeling software is one of the promising prospects of the further development of data-driven management of reinforced concrete structures. In general, the offered framework can be considered as a decision-support and visualization tool that enhances expert judgment, which is consistent with developing interdisciplinary thinking of artificial intelligence in the built environment (Amirkhani et al., 2024).

6. Conclusion

This paper introduced a deep learning framework to crack detection and map structural defects to reinforced concrete systems, paying a lot of attention to pixel level semantic segmentation and spatial interpretability. The proposed solution, which was not based on binary crack detection, allowed detecting and mapping several surface flaws at the same time, such as cracks, corrosion, spalling, and exposed reinforcement, offering a more complete view of reinforced concrete degradation. It was experimentally shown that the framework is capable of producing coherent and

interpretable defect maps, and that the performance of the framework in terms of class-dependencies is strongly tied to the visual nature of particular defects. Although the fine-scale cracks were less accurately segmented by their geometry and contrast, visually salient defects like corrosion and bare reinforcement were better segmented, which validates the potential of deep learning to be used in conjunction with surface condition evaluation. The paper also emphasized the usefulness of defect-area mapping as an additional analysis of conventional overlap metrics. The comparison of the predicted defects and the actual defects made it possible to detect systematic patterns, such as overestimates in the crack and underestimates in the segmentation of the exposed reinforcement, which highlights the need to consider the output of models as a relative measure of the surface condition instead of an absolute one. Architecturally and in terms of infrastructure management, such spatial representations are useful in condition documentation, prioritization of maintenance and informed decision-making. On the findings, the best place of deep learning based inspection tools is as decision-support systems that supplement professional judgment, not substitute it, and should have visual interpretability, produce clear reports of limitations, and should be integrated with current inspection processes. Future studies ought to be improved on sensitivity to fine-scale cracks using multi-scale architectures and larger and more varied datasets, use of time analysis to monitor deterioration, and combine segmentation results with digital systems like building information modeling systems to improve data-driven management of reinforced concrete assets.

References

- [1]. Laxman, K. C., Tabassum, N., Ai, L., Cole, C., & Ziehl, P. (2023). Automated crack detection and crack depth prediction for reinforced concrete structures using deep learning. *Construction and Building Materials*, 370, 130709.
- [2]. Kolappan Geetha, G., Yang, H. J., & Sim, S. H. (2023). Fast detection of missing thin propagating cracks during deep-learning-based concrete crack/non-crack classification. *Sensors*, 23(3), 1419.
- [3]. Mir, B. A., Sasaki, T., Nakao, K., Nagae, K., Nakada, K., Mitani, M., ... & Jindai, M. (2022). Machine learning-based evaluation of the damage caused by cracks on concrete structures. *Precision Engineering*, 76, 314-327.
- [4]. Han, X., Zhao, Z., Chen, L., Hu, X., Tian, Y., Zhai, C., ... & Huang, X. (2022). Structural damage-causing concrete cracking detection based on a deep-learning method. *Construction and Building Materials*, 337, 127562.
- [5]. Xu, Y., Li, Y., Zheng, X., Zheng, X., & Zhang, Q. (2023). Computer-vision and machine-learning-based seismic damage assessment of reinforced concrete structures. *Buildings*, 13(5), 1258.
- [6]. Golding, V. P., Gharineiat, Z., Munawar, H. S., & Ullah, F. (2022). Crack detection in concrete structures using deep learning. *Sustainability*, 14(13), 8117.
- [7]. Nguyen, S. D., Tran, T. S., Tran, V. P., Lee, H. J., Piran, M. J., & Le, V. P. (2023). Deep learning-based crack detection: A survey. *International Journal of Pavement Research and Technology*, 16(4), 943-967.
- [8]. Wang, W., Su, C., & Fu, D. (2022, September). Automatic detection of defects in concrete structures based on deep learning. In *Structures* (Vol. 43, pp. 192-199). Elsevier.

- [9]. Chen, S., Fan, G., Li, J., & Hao, H. (2025). Automatic complex concrete crack detection and quantification based on point clouds and deep learning. *Engineering Structures*, 327, 119635.
- [10]. Haciefendioğlu, K., Altunışık, A. C., & Abdioğlu, T. (2023). Deep learning-based automated detection of cracks in historical masonry structures. *Buildings*, 13(12), 3113.
- [11]. Hu, D., Yee, T., & Goff, D. (2024). Automated crack detection and mapping of bridge decks using deep learning and drones. *Journal of Civil Structural Health Monitoring*, 14(3), 729-743.
- [12]. Alqurashi, I., Alsulami, M., Alver, N., & Catbas, N. (2025). Ultrasonic tomography with deep learning for detecting embedded components and internal damage of concrete structures. *Developments in the Built Environment*, 100742.
- [13]. Amirkhani, D., Allili, M. S., Hebbache, L., Hammouche, N., & Lapointe, J. F. (2024). Visual concrete bridge defect classification and detection using deep learning: A systematic review. *IEEE Transactions on Intelligent Transportation Systems*, 25(9), 10483-10505.
- [14]. Ali, L., Alnajjar, F., Khan, W., Serhani, M. A., & Al Jassmi, H. (2022). Bibliometric analysis and review of deep learning-based crack detection literature published between 2010 and 2022. *Buildings*, 12(4), 432.
- [15]. Alkannad, A. A., Al Smadi, A., Yang, S., Al-Smadi, M. K., Al-Makhlafi, M., Feng, Z., & Yin, Z. (2025). Crackvision: Effective concrete crack detection with deep learning and transfer learning. *IEEE Access*.
- [16]. Wan, C., Xiong, X., Wen, B., Gao, S., Fang, D., Yang, C., & Xue, S. (2022). Crack detection for concrete bridges with imaged based deep learning. *Science progress*, 105(4), 00368504221128487.
- [17]. Yang, Y., Yin, Y., & Zhou, J. (2025). Enhanced detection and prediction of concrete filling defects with deep learning. *Construction and Building Materials*, 477, 141180.
- [18]. Hamishebahar, Y., Guan, H., So, S., & Jo, J. (2022). A comprehensive review of deep learning-based crack detection approaches. *Applied Sciences*, 12(3), 1374.
- [19]. Garita-Durán, H., Stöcker, J. P., & Kaliske, M. (2025). Deep learning-based system for automated damage detection and quantification in concrete pavement. *Results in Engineering*, 25, 104546.
- [20]. Zhou, S., Pan, Y., Huang, X., Yang, D., Ding, Y., & Duan, R. (2022). Crack texture feature identification of fiber reinforced concrete based on deep learning. *Materials*, 15(11), 3940.
- [21]. Fan, C. L. (2025). Concrete crack detection and severity assessment using deep learning and multispectral imagery analysis. *Measurement*, 247, 116825.
- [22]. Kim, B., Natarajan, Y., Preethaa, K. S., Song, S., An, J., & Mohan, S. (2024). Real-time assessment of surface cracks in concrete structures using integrated deep neural networks with autonomous unmanned aerial vehicle. *Engineering Applications of Artificial Intelligence*, 129, 107537.
- [23]. Abdelkader, E. M., Zayed, T., & Faris, N. (2023). Synthesized evaluation of reinforced concrete bridge defects, their non-destructive inspection and analysis methods: a systematic review and bibliometric analysis of the past three decades. *Buildings*, 13(3), 800.
- [24]. Kirthiga, R., & Elavenil, S. (2024). A survey on crack detection in concrete surface using image processing and machine learning. *Journal of Building Pathology and Rehabilitation*, 9(1), 15.
- [25]. Nabizadeh, E., & Parghi, A. (2023). Vision-based concrete crack detection using deep learning-based models. *Asian Journal of Civil Engineering*, 24(7), 2389-2403.
- [26]. Guo, F., Cui, Q., Zhang, H., Wang, Y., Zhang, H., Zhu, X., & Chen, J. (2024). A new deep learning-based approach for concrete crack identification and damage assessment. *Advances in Structural Engineering*, 27(13), 2303-2318.
- [27]. Seol, D., Kim, H., Park, S., & Lee, J. (2025). Comparative study of cnn-based deep learning models for concrete crack detection. *Journal of Building Engineering*, 113651.
- [28]. Zoubir, H., Rguig, M., El Aroussi, M., Chehri, A., Saadane, R., & Jeon, G. (2022). Concrete bridge defects identification and localization based on classification deep convolutional neural networks and transfer learning. *Remote Sensing*, 14(19), 4882.
- [29]. Abubakr, M., Rady, M., Badran, K., & Mahfouz, S. Y. (2024). Application of deep learning in damage classification of reinforced concrete bridges. *Ain Shams engineering journal*, 15(1), 102297.
- [30]. Guo, P., & Bao, Y. (2025). HRCDS: A benchmark dataset for high-resolution concrete damage
- [31]. segmentation (Version 1) [Data set]. Mendeley Data. <https://doi.org/10.17632/6x4dzzrs2h.1>