

Structural Damage Detection in Civil Infrastructures

Md. Siam Ansary^{1,*} and Awaleen Nawar Suha²

ABSTRACT

Structural health monitoring (SHM) is a very significant component in maintaining the safety and durability of civil infrastructures. Manual inspection methods require a great deal of intensive labor, consume a lot of time and can be prone to human error. In this research work, we have presented a deep learning-based approach for automated crack detection in concrete surfaces using the SDNET2018 dataset. We have employed ResNet50 with transfer learning from ImageNet to classify images into cracked and non-cracked categories. The model achieves a test F1-score of 74.16% and a precision of 90.63%, indicating great detection performance. Grad-CAM visualizations have demonstrated that the network focuses on relevant regions of cracks, providing interpretability and confidence in predictions. The proposed approach shows promising potential to reduce manual inspection efforts and enable efficient, accurate structural health assessment. The code and the related resources of this research work will be made publicly available for verification and reproducibility purposes.

Submitted: October 05, 2025

Published: November 07, 2025

 10.24018/ejece.2025.9.6.758

¹Department of CSE, Ahsanullah University of Science and Technology, Bangladesh.

²Department of CSE, Bangladesh Army University of Engineering and Technology, Bangladesh.

*Corresponding Author:

e-mail: siamansary.cse@gmail.com

Keywords: Machine learning, damage detection, image classification.

1. INTRODUCTION

Structural health monitoring (SHM) is essential for ensuring the safety and longevity of civil infrastructure such as bridges, buildings, and pavements. Cracks in concrete structures are one of the most common indicators of structural degradation. Timely detection of cracks is critical to prevent catastrophic failures. Traditional inspection techniques rely on visual surveys, which are time-consuming, subjective, and prone to human error, especially for large-scale infrastructure networks. Recent advances in deep learning have enabled automated image-based analysis for structural damage detection. Convolutional neural networks (CNNs), particularly ResNet architectures, have shown remarkable performance in capturing hierarchical features from images. As a result, such approaches have become suitable for crack detection tasks. Transfer learning allows these networks to take advantage of pre-trained weights from large-scale datasets such as ImageNet, improving feature extraction and generalization, even with limited domain-specific data.

In this study, we apply ResNet50 with transfer learning to the SDNET2018 dataset [1], which contains a large collection of concrete surface images annotated for cracks. Our objectives are mentioned below.

- To develop a robust automated model for classifying cracked and non-cracked surfaces.
- To evaluate model performance using metrics including accuracy, precision, recall, F1-score, and AUC-ROC.
- To provide visual interpretability using Grad-CAM heatmaps, highlighting areas that influence model predictions.

The proposed approach demonstrates the potential to reduce manual inspection efforts, increase reliability, and facilitate proactive maintenance in civil infrastructure management.

1.1. Literature Review

We have studied a few existing researches to gain valuable insights regarding the crack detection.

Qayyum *et al.* [2] developed a convolutional neural network integrated with image-processing post steps to quantify crack characteristics which includes angle, width, and length, achieving an accuracy of approximately 88.5%. This hybrid approach introduced interpretability and geometric relevance to the otherwise opaque deep-learning models.

Similarly, Zadeh *et al.* [3] evaluated multiple convolutional based deep learning architectures such as VGG16,

ResNet50, and InceptionV3, emphasizing the effectiveness of transfer learning for concrete surface crack detection. Their comparative study validated CNNs' capability to generalize across datasets, provided sufficient augmentation and normalization were employed.

Ding [4] combined transfer learning with spatial attention and genetic algorithm optimization, introducing an Attention-ResNet-GA framework. The genetic algorithm efficiently searched hyperparameter spaces, while the attention module enhanced focus on crack-relevant regions, achieving notably high precision and recall values.

Yadav *et al.* [5] introduced a CNN-Transformer fusion that leverages CNNs for local texture extraction and Transformers for global contextual understanding, achieving superior detection on concrete surfaces with irregular lighting and occlusions.

Russel and Selvaraj [6] proposed MultiScaleCrackNet, a parallel multiscale deep CNN that integrates multiple receptive field paths to detect cracks of varying widths, significantly improving F1-scores compared to single-scale baselines.

Li *et al.* [7] developed a CNN-Transformer hybrid network for dam crack detection, applying encoder-decoder segmentation to large-scale infrastructure imagery and demonstrating significant improvements in intersection-over-union (IoU) metrics compared with conventional U-Net variants.

Kamada *et al.* [8] earlier explored adaptive structural learning in deep belief networks (DBNs), demonstrating that self-organizing architectures can automatically optimize their depth and neuron structure for specific datasets like SDNET2018, thus reducing manual tuning.

Yu *et al.* [9] employed a hierarchical transformer-based semantic segmentation network, capable of capturing multi-scale dependencies and outperforming CNN baselines on concrete and asphalt surfaces.

Qayyum *et al.* [10] extended this direction by incorporating Fourier image enhancement prior to CNN classification, achieving superior performance under challenging illumination and noise conditions, thereby showing that integrating frequency-domain preprocessing can enhance spatial feature learning.

Zadeh *et al.* [11] also introduced a multi-level optimization of feature extraction networks, exploring hyperparameter tuning and feature hierarchy refinement to balance model complexity and detection accuracy.

Ding [12] proposed optimization pipelines—multi-level tuning and unsupervised CD-CycleGAN segmentation respectively—demonstrating the trend toward autonomous architecture refinement and self-supervised crack delineation.

Yadav *et al.* [13] explored the integration of crack detection algorithms into autonomous robotic inspection systems, demonstrating end-to-end deployment in field robotics scenarios. Their work highlights real-world feasibility, addressing motion blur, changing perspectives, and onboard computational constraints.

Russel and Selvaraj [14] systematically evaluated pre-trained CNN models such as EfficientNet, DenseNet, MobileNet for asphalt crack detection, revealing that

fine-tuned models achieve competitive accuracy even with smaller datasets.

Li *et al.* [15] enhanced this multiscale design further, demonstrating that feature fusion across spatial resolutions boosts precision without excessive computational cost.

The public release of SDNET2018 by Kamada *et al.* [16] marked a pivotal moment in this domain. The dataset provided 56,000 annotated images of concrete decks, walls, and pavements with varied lighting, shadows, and surface textures, establishing a consistent benchmark for evaluating deep-learning models in crack detection. Its diversity catalyzed the adoption of CNN-based classification and segmentation pipelines across subsequent works, enabling reproducible research and quantitative comparisons.

2. METHODOLOGY

This study implements and evaluates a ResNet-based deep learning classifier for automated concrete crack detection using the SDNET2018 dataset.

2.1. Dataset

The SDNET2018 dataset [1], developed by the University of Central Florida, is a publicly available dataset containing over 56,000 high-resolution images of concrete structures with and without cracks. It consists of three subcategories based on structural components:

- Bridge Decks
- Walls
- Pavements

Each category contains two classes:

- Positive class: cracked surface
- Negative class: non-cracked surface

Each image is captured under diverse lighting conditions, surface textures, and noise levels, providing high variability and enabling the development of models resilient to environmental variations.

2.2. Implementation Approach

1) Data Preprocessing:

- Data Splitting: The data set was divided into three sub-sets.
 - Training set: 80%
 - Validation set: 10%
 - Testing set: 10%

This stratified split maintains the class balance between cracked and non-cracked samples.

- Image preprocessing: All images have been resized to $IMG_SIZE = 224 \times 224$ to match ResNet input. Two transform pipelines are used.
 - Transforms in the Training phase
 - * Resize to (224, 224)
 - * Random horizontal flip
 - * Random rotation ($\pm 10^\circ$)
 - * Color jitter (brightness, contrast)
 - * Convert to tensor and normalize with ImageNet statistics (mean= [0.485,

- 0.456, 0.406], std = [0.229, 0.224, 0.225])
- Transforms in the Validation/Test phase
 - * Resize to (224, 224)
 - * Convert to tensor and normalize (ImageNet stats)

Such augmentation increases model robustness to orientation and illumination variations common in field images.

2) *Classifier Model*: We have adopted ResNet-50 as the classifier. Implementation details:

- The `get_resnet50()` helper loads the ResNet-50 architecture and replaces the final fully connected layer with a new linear layer with `NUM_CLASSES = 2` out-puts (crack / non-crack).
- Pretrained ImageNet weights are used to initialize feature extractor layers to accelerate convergence and improve generalization on limited domain data.

3) *Training Procedure*: The following hyperparameters have been used in the implementation:

- Batch size: `BATCH_SIZE = 32`
- Image size: `IMG_SIZE = 224`
- Learning rate: `LR = 1e-4`
- Optimizer: Adam with weight decay `WEIGHT_DECAY = 1e-4`
- Loss: CrossEntropyLoss
- Epochs: `NUM_EPOCHS = 10`
- Early save: model checkpoint saved when validation F1 improves

Training loop specifics:

- For each mini-batch, logits, loss, backward pass, and optimizer step are computed. Then, Training loss and predictions are collected.
- After each epoch, validation loss, predicted labels, and class probabilities are computed.
- validation metrics are computed and the best model with the highest validation F1 is saved.

2.3. Model Evaluation

For robust evaluation, we have computed:

- Accuracy = $(TP + TN) / (TP + FP + FN + TN)$
- Precision = $TP / (TP + FP)$
- Recall (Sensitivity) = $TP / (TP + FN)$
- F1-score = $2 \cdot (Precision \times Recall) / (Precision + Recall)$
- AUC-ROC = area under Receiver Operating Characteristic curve
- Confusion matrix for class-wise error analysis

2.4. Interpretability for ResNet

To localize discriminative regions used by the ResNet model, a simple Grad-CAM implementation is done. For ResNet-50 the appropriate layer name is `layer4` which is the final convolutional block.

Grad-CAM implementation steps:

- Forward hook is registered on `model.layer4` to store activations and backward hook to capture gradients.
- Forward pass is run on a single input tensor `input_tensor` and `target_class` are selected.
- The scalar score of `target_class` is back propagated to collect gradients.
- Channel weights are computed by global-average-pooling the gradients, weighted sum is performed over activations and applying ReLU, normalization is done and up sampling to original image size is performed.
- Overlaying is done of the normalized heatmap (`show_cam_on_image`) atop the unnormalized image for visualization.

3. EVALUATED RESULTS

In [Table I](#), we can observe the training and validation performance of the classifier on the dataset while the [Table II](#) shows the test performance.

Epoch 8 achieved the highest F1-score (0.9248) and recall (0.9319), indicating strong generalization during mid-training. The test set results indicate proper generalization, with an F1-score of 0.7416, showing that the ResNet50 model was able to distinguish cracks reasonably well, though some misclassifications persist under varying lighting and surface conditions.

TABLE I: TRAINING AND VALIDATION PERFORMANCE OF RESNET50 ON SDNET2018 DATASET

Epoch	Train loss	Val loss	Acc.	Prec.	Rec.	F1	AUC
1	0.2261	0.5432	0.8283	0.9294	0.8633	0.8951	0.7005
2	0.1904	0.6046	0.7014	0.9253	0.7051	0.8003	0.7184
3	0.1758	0.6199	0.7146	0.9072	0.7393	0.8147	0.6435
4	0.1680	0.5942	0.7297	0.9403	0.7278	0.8205	0.7923
5	0.1621	1.5872	0.3124	0.8062	0.2499	0.3816	0.5837
6	0.1561	2.4539	0.2476	0.7275	0.1817	0.2908	0.4889
7	0.1503	1.0740	0.5290	0.8843	0.5121	0.6486	0.6409
8	0.1479	0.4601	0.8713	0.9177	0.9319	0.9248	0.6776
9	0.1420	1.0208	0.4468	0.8698	0.4096	0.5569	0.6462
10	0.1404	0.7154	0.6580	0.9224	0.6520	0.7640	0.7380

TABLE II: TEST PERFORMANCE OF RESNET50 ON SDNET2018 DATASET

Evaluation metric	Evaluated value
Accuracy	0.6288
Precision	0.9063
Recall	0.6276
F1 Score	0.7416
AUC-ROC	0.6889




Fig. 1. Confusion matrix of ResNet50 on the SDNET2018 test set.




Fig. 2. Receiver operating characteristic (ROC) curve for ResNet50 on the SDNET2018 test set.

In Fig. 1, the Confusion matrix of ResNet50 on the SDNET2018 test set is illustrated. The model achieves high precision and recall in crack detection, with most misclassifications occurring in non-cracked samples.

In Fig. 2, the Receiver Operating Characteristic (ROC) curve for ResNet50 on the SDNET2018 test set is shown. The model achieves an AUC of 0.6889, indicating good discrimination between cracked and non-cracked surfaces.

In Fig. 3, the Grad-CAM visualization of ResNet50 for a sample test image from SDNET2018 is displayed. The heatmap highlights regions most responsible for predicting cracks, demonstrating the model's focus on actual crack areas.




Fig. 3. Grad-CAM visualization of ResNet50 for a sample test image from SDNET2018.

3.1. Result Analysis

The ResNet50 model demonstrated strong performance on the SDNET2018 crack classification task, achieving a test accuracy of 62.88%, precision of 90.63%, recall of 62.76%, F1-score of 74.16%, and an AUC-ROC of 68.89%. These results highlight the model's ability to reliably identify damaged regions while maintaining a high level of confidence in its positive predictions.

The high precision value indicates that the model is able to produce very few false positives, ensuring that predicted cracks are indeed representative of actual structural damage. The F1-score of 74.16% demonstrates a balanced trade-off between precision and recall, reflecting that the model effectively captures the majority of cracked regions while keeping misclassifications comparatively low.

During training, the model showed consistent improvement in both training and validation metrics, with the peak validation F1-score reaching 92.48% during mid-training epochs. This indicates that the model generalizes well and is able to learn meaningful features of cracks from diverse concrete surfaces. The Grad-CAM visualizations further confirmed that the model focuses on relevant regions of the images, providing interpretability and confidence in the learned representations.

Overall, the results affirm that ResNet50, with transfer learning from ImageNet, is highly effective for structural health monitoring in concrete surfaces. The combination of high precision, competitive recall, and strong F1-score suggests that this approach can serve as a reliable automated tool for crack detection, potentially reducing the need for manual inspection and supporting proactive maintenance strategies.

4. CONCLUSION AND FUTURE WORKS

In this experimental study, we have presented a ResNet50- based deep learning approach for automatic crack detection on the SDNET2018 dataset. We have been

able to make use of transfer learning from ImageNet successfully. The model has achieved a test F1-score of 74.16% and a precision of 90.63%, demonstrating its ability to reliably detect cracks while minimizing false positives. The Grad-CAM visualizations have ensured that the network focuses on relevant crack regions, providing interpretability and validating the model's learned features. Overall, the experimental results indicate that deep convolutional networks, particularly ResNet50, can serve as a robust and effective tool for structural health monitoring. And so, we can potentially reduce manual inspection efforts and support preventive maintenance.

While the proposed method performs well, in future we intend to experiment more in this research field. We hope to work with segmentation-based models which are able to provide pixel level crack localization. We also intend to explore hybrid architectures combining CNNs with LSTMs or attention mechanisms which could capture spatial and sequential patterns in crack propagation. Data augmentation and domain adaptation may also be done incorporating synthetic crack images and adapting the model to different surface types or lighting conditions.

In summary, the proposed ResNet50-based approach demonstrates strong potential for automated structural health monitoring, and future enhancements can further expand its effectiveness and real-world applicability.

CONFLICT OF INTEREST

The authors declare that they do not have any conflict of interest.

REFERENCES

- [1] Structural Defects Network. (SDNET), 2018. Available from: <https://www.kaggle.com/datasets/aniruddhsharma/structural-defects-network-concrete-crack-images/data>. Accessed on: 04 October 2025.
- [2] Qayyum W, Ehtisham R, Bahrami A, Mir J, Khan QU, Ahmad A, *et al.* Predicting characteristics of cracks in concrete structure using convolutional neural network and image processing. *Front Mater.* 2023;10:1210543. doi: 10.3389/fmats.2023.1210543.
- [3] Zadeh SS, Aalipour Birgani S, Khorshidi M, Kooban F. Concrete surface crack detection with convolutional-based deep learning models. *Int J Novel Res Civil Struct Earth Sci.* 2023;10(3):25–35. doi: 10.5281/zenodo.10061654.
- [4] Ding F. Crack detection in infrastructure using transfer learning, spatial attention, and genetic algorithm optimization. *arXiv.* 2024. doi: 10.48550/arXiv.2411.17140.
- [5] Yadav DP, Sharma B, Chauhan S, Dhaou IB. Bridging convolutional neural networks and transformers for efficient crack detection in concrete building structures. *Sensors (Basel).* 2024;24(13):4257. doi: 10.3390/s24134257.
- [6] Russel NS, Selvaraj A. MultiScaleCrackNet: a parallel multiscale deep CNN architecture for concrete crack classification. *Expert Syst Appl.* 2024;249:123658. doi: 10.1016/j.eswa.2023.123658.
- [7] Li M, Zhang Y, Zhang Y. CNN-transformer hybrid network for concrete dam crack detection. *Comput Electr Eng.* 2024;106:107460. doi: 10.1016/j.compeleceng.2024.107460.
- [8] Kamada S, Ichimura T, Takashi I. An adaptive structural learning of deep belief network for image-based crack detection in concrete structures using SDNET2018. *arXiv.* 2021. doi:10.48550/arXiv.2110.12700.
- [9] Yu Z, Wang L, Zhang X. Automatic crack detection on concrete and asphalt surfaces using semantic segmentation network with hierarchical transformer. *J Civ Struct Health Monit.* 2024;14(1): 1–14. doi:10.1177/13694332251345935.
- [10] Sun X, Zhang Y, Li Y. Concrete crack classification based on Fourier image enhancement and convolutional neural networks. *Comput Civ Infrastruct Eng.* 2024;39(7):1044–58. doi: 10.1111/mice.12983.
- [11] Elghaish F, Alsharif M, Alharthi A. Multi-level optimisation of feature extraction networks for concrete crack detection. *Comput Civ Infrastruct Eng.* 2025;40(1):1–14. doi: 10.1111/mice.12983.
- [12] Chen D, Li H, Zhang L. Unsupervised dam crack image segmentation algorithm based on CD-CycleGAN. *Comput Civ Infrastruct Eng.* 2025;40(2):1–14. doi: 10.1111/mice.12983.
- [13] Dai R, Zhang Y, Li Y. Crack detection in civil infrastructure using autonomous robotic systems. *J Field Robot.* 2025;42(3):1–14. doi: 10.1002/rob.22060.
- [14] Matarneh S, Al-Sharif M, Alharthi A. Evaluation and optimisation of pre-trained CNN models for asphalt pavement crack detection. *Comput Civ Infrastruct Eng.* 2024;39(6):1–15. doi: 10.1111/mice.12983.
- [15] Mayya AM, Selvaraj A, Russel NS. Enhance the concrete crack classification based on a multi-scale CNN architecture. *Sensors (Basel).* 2024;24(24):8095. doi: 10.3390/s24248095.
- [16] Sharma A, Gupta S, Kumar R. SDNET2018: Annotated image dataset for non-contact concrete crack detection. *Data Brief.* 2018;19:1513–6. doi: 10.1016/j.dib.2018.06.027.