

Automated Concrete Cube Failure Mode Classification Using Deep Learning and Engineering-Guided Decision Logic

Meenakshi Somnath Patil¹, Rajesh Ghongade², Hemant B. Dhonde³

¹Research Scholar, Bharati Vidyapeeth (Deemed to be University) College of Engineering, and Assistant Professor at Indira College of Engineering and Management, Pune, 412 115, India

²Professor, E&TC Department, Bharati Vidyapeeth (Deemed to be University) College of Engineering, Pune, 412 115, India

³Director, C-Probe Technologies Pvt. Ltd., Pune, 411016, India

Corresponding Author: Meenakshi Somnath Patil

Abstract: Compression strength of concrete on cubes is a traditional building quality control check. Nevertheless, the evaluation of cube failure mechanisms is still mainly based on visual inspection (manual), which is subjective and varies among different individuals with a high degree of misclassification, resulting in material waste, higher costs, and unnecessary environmental pollution. Although there has been study on crack detection for the surface of concrete, automatic categorization of engineering-valid cube failure modes is still scanty. This paper presents an automated fault mode classification framework which combines deep learning-based object detection with rule-based engineering decision logic. A hybrid YOLOv8 and convolutional neural network-based architecture is proposed to detect objects and cracks at multiple scales. The system provides a six-face assessment of each concrete cube and adopts morphological image processing to improve the rapid detection and estimation of crack boundaries along with noise reduction. Categorization of failure mode is done by decision rule on rule-based engineering based on the conventional method of testing. Additionally, a dataset including 100 concrete cube samples (and 600 images) was extracted and annotated by professionals. Severe class imbalance was dealt with by creating a balanced dataset via augmenting and synthesizing data. The experimental results show that the classification accuracy of all works are 94.3% on the original data and 96.8% on the balanced data, whilst cube detection has precision = 97.8%, recall = 96.5%, mAP@0.5 of 0.987. The system significantly shortens time of inspection by about 85% compared to manual assessment with high reliability. Results show that coupling deep learning with model-based decision logic yields robust scalable and objective failure mode classification paving the way for automated quality control in sustainable construction applications.

Keywords: Concrete cube failure mode classification, Deep learning-based inspection, YOLOv8 object detection, Multi-view image analysis, Engineering-guided decision rules, Construction quality control

1. Introduction

After water, concrete is the most consumed material on earth planet and still an essential element in the system of modern infrastructure. For this reason, its mechanical reliability and quality stability are a prerequisite for the safety and environmental sustainability of infrastructural applications [1–3]. Compressive strength testing of concrete cubical specimens is the most universally standardized and commonly applied test for quality control purposes [4]. In addition to the calculated value of compressive strength, a concrete cube failure mode provides vital qualitative information about how homogeneous the material is, whether curing was adequate, how well the sample was prepared or loading alignment during testing [5].

In routine operation, cube's failure modes are checked by visual inspection visual by laboratory staff. This is a subjective process that is heavily reliant on personal

experience and as such suffers from significant inter-observer variability and inconsistent definitions of acceptable versus non acceptable failures [6,7]. Typical acceptable and unacceptable cube failure modes during the uniaxial compression testing were illustrated in Fig. 1 and Fig. 2, respectively. Acceptable failures are normally in the form of uniform vertical cracking on the side faces with no damage or surface spalling occurring on loading surfaces, while unacceptable failure exhibits non-uniform cracking and premature damage appears on top and bottom faces.




Fig.1 Satisfactory Failures of Cube Inspection




Fig.2 Unsatisfactory Failure Mode of Cube Inspection

Manual assessment of failure mode could be misclassified and may in some cases reject a sound concrete or accept a defective one [8, 9]. Such inaccuracies inherently cause materials wastage, higher construction cost and worse environmental burdens, especially considering that the production of cement is responsible for about 7–8% of global anthropogenic CO₂ emissions [10,11].

In order to minimise the subjectivity involved in concrete evaluation, previous studies investigated non-destructive testing methods including ultrasonic pulse velocity, rebound hammer testing and vibration-based techniques [12,13]. Although these techniques increase objectivity compared to manual inspection, they are instrumental and operator-dependent methods, and thus too time-consuming for high-throughput routine cube testing [14]. Furthermore, this approach mainly focuses on predicting material properties rather than directly classifying failure modes.

Automated inspection of concrete surfaces has also been possible through the developments in computer vision techniques. Conventional image processing has been attempted for detection of surface cracks and damage patterns using techniques such as edge detection, thresholding and morphological operations [15,16]. Nonetheless they do not perform well in presence of changing illumination, surface texture variations and background noise that are widely occurred in laboratory and construction site conditions [17].

The robustness of image-based crack detection has been greatly improved with the recent advances in machine learning and deep learning. Convolutional neural networks (CNNs) have achieved the great success in learning and extracting hierarchical feature representations on crack index estimation, crack width measurement and surface damage localization [18–20]. Based on CNNs, object detection methods such as Faster R-CNN, RetinaNet, Mask R-CNN and YOLO (You Only Look Once) variants pushed forward automated inspection by allowing for multi-scale detection with accuracy and computational complexity improved aiding automatic

assignment of individual gross defects among multiple scales [21–23]. Of these, YOLO-based models are in particular well-suited for real-time implementation as they have a good trade-off between detection quality and speed [24].

Notwithstanding these developments, there is a clear gap in the literature. Most of the efforts are directed to crack detection or surface defect localization and usually applied not for engineering-valid of failure mode classification of cube shaped concrete specimens [25]. Crack presence is not the only criterion in deciding if a cube failure falls within criteria of standardized compressive testing. Engineering evaluation must consider distribution of cracks on various faces, condition of loading surfaces, consistency from set to set of the fracture patterns and correlation with standard failure characteristics.

Moreover, the poor asymmetry and lack of a balanced dataset of concrete cube analysis are another significant obstacle. The available datasets are usually small, suffer from serious class imbalance caused by the rarity of sufficient failure modes, and do not have standard multi-view annotations [26, 27]. Recent approaches have focused on the use of data augmentation and synthetic sample generation to address class imbalance above and improve generalization in DNN models [28,29]. Until now, however, such methods haven't been systematically used for classification of failure mode on concrete cubes.

To overcome these gaps, this study presents an automatic failure mode classification framework by combining deep learning-based visual perception with engineering-guided decision logic. The proposed system locally processes a six-face examination of each cube specimen, by an integrated framework realizes YOLOv8 for object and crack detection combined with CNN feature extraction and morphologic image processing to fine-tune the boundary.

Unlike full data-driven methods, the final failure mode classification is controlled by explicit rules based on cube testing practices standards allowing for interpretability and engineering adaptability. Incorporating deep learning models with domain-specific engineering expertise, the developed framework presents a solid and scalable objective automated approach to failure mode assessment of concrete cubes that aligns with INDUSTRY 4.0-driven quality control and sustainable construction methodologies.

2. Methodology

2.1 Overall System Architecture

The method proposed here is based on a structured pipeline in multiple stages for automatic classification of concrete cube failure mode that includes visual analysis in the form of deep learning and engineering-informed decision logic. An overall schematic of the system is shown in Fig. 3.




Fig.3 System architecture block diagram

The framework aims to simulate and automate the decision process which is conducted by domain experts: laboratory inspectors of significant experience, but without the subjectivity and the variability. All of the concrete cubes are inspected through a six-face orientation inspection procedure, guaranteeing that all surfaces are totally observed and responsible to interpret fracture patterns. The pipeline of the process is made up of a series of steps: image acquisition, pre-processing, object/crack detection, feature extraction and refining, failure mode classification and final decision output.

Modular design is adopted, and different components (e.g. the detection backbone or feature extraction module) can be replaced or will updated without changing the overall decision logic. This design results in increased scalability, reproducibility, and deployment potential to real-world settings to the construction quality control.

2.2 Image Acquisition and Preprocessing

Six high-definition images of each concrete cube sample are taken (from the four side surfaces, the upper surface, and the lower surface). This multi-view acquisition approach allows to assess the unity of fracture and the damage on loaded surfaces, which, in turn, are key parameters for non-acceptable failure modes.

All images are pre-processed before analysis to improve the visual quality and compatibility with the detection network. Preprocessing steps consist of image rescaling and normalization, noise reduction by applying Gaussian filter, and contrast enhancement to enhance crack visibility. These underlying image-processing steps minimize the effect of background artifacts, illumination non-uniformities and surface texture non-uniformities prevalent in laboratory-based and construction-site settings.

Data augmentation is further used during the model training to be robust and generalized. Augmentation phenomena comprise geometrical and photometrical variations as well as the introduction of a controlled amount of noise in order to represent invariant feature learning against varying imaging conditions.

2.3 Dataset Description

The dataset used for this study consists of concrete cubes from laboratory and field testing sites. There are six images belonging to each specimen for all surfaces of a cube that forms this multi-view dataset appropriate for the comprehensive failure mode study. Table 1 Overview of the dataset composition (i.e., sample size, number of images and class distribution).

Table 1. Dataset Composition (Original Dataset)

Parameter	Details
Total number of concrete cube specimens	100
Images per cube	6 (six faces per cube)
Total images	600
Source – Construction sites	65 cubes (65%)
Source – Laboratory environments	35 cubes (35%)
Acceptable failure mode cubes	10 cubes (10%)
Non-acceptable failure mode cubes	90 cubes (90%)
Acceptable failure mode images	60 images
Non-acceptable failure mode images	540 images
Annotation type	Professional civil engineering annotation
Annotation standard	BS EN 12390-3
Failure mode labels	Acceptable / Non-acceptable

The dataset is a manifestation of testing phenomenon in the field, where unsatisfactory failure modes are orders of magnitude more common than satisfactory ones. Though this imbalance is desirably to increase the practical relevance, it is problematic for supervised learning and further requires certain balancing strategies during which system training.

2.4 Dataset Balancing and Augmentation

In order to minimize the negative influence of class imbalance, we augment the dataset following modern augmentation and synthetic data generation methods. For acceptable failure modes that are underrepresented, aggressive augmentation performs controlled geometric transformations, photometric variations and environmental simulation. These methods keep the cheap-to-implement structural attributes of admissible failures and improve sample diversity.

For unacceptable modes of failure, conservative enhancement techniques are utilized to prevent the loss of key cracking patterns. Moreover, synthetic image generation is applied as a complement to acceptable failure samples by procedurally generating realistic crack patterns in line with the standardised failure mode. The statistics of the generated balanced dataset are summarized in Table 2.

Table 2. Balanced Dataset Statistics (After Augmentation)

Parameter	Details
Total number of concrete cube specimens (balanced set)	200
Images per cube	6 (six faces per cube)
Total images	1200
Acceptable failure mode cubes	100 cubes (50%)
Non-acceptable failure mode cubes	100 cubes (50%)
Acceptable failure mode images	600 images
Non-acceptable failure mode images	600 images
Balancing technique	Data augmentation + synthetic sample generation
Augmentation applied to	Acceptable failure mode class
Augmentation operations	Rotation, flipping, scaling, brightness/contrast

	adjustment
Labeling consistency	Maintained from original dataset
Annotation standard	BS EN 12390-3

This balanced data facilitates better model stability, less classification bias and fair performance evaluation on both failure categories.

2.5 Object and Crack Detection Using YOLOv8

Object and crack detection is implemented with an optimized YOLOv8 architecture (as it has a high detection rate and is capable of real-time inference). The system architecture and detection process are shown in Fig. 4




Fig.4YOLOv8 architecture optimized for concrete cube detection

The object detection of YOLOv8 is anchor-free single-stage, and it predicts the object classes and bounding box coordinates at once. The network is composed of a convolutional backbone for hierarchical feature generation, a multi-scale feature fusion neck to process cracks of different scales, and a detection head for classification and localization.

This setting makes it possible to perform the simultaneous localization of the concrete cubes and their crack regions in a single forward pass for cost-effective inspection across multiple cube faces under deployment logistic constraints.

2.6 Morphological Refinement and Feature Extraction

By applying classical morphological image processing operations to raw detection results, the definition of crack boundaries is improved and spurious detections are suppressed. These involve erosion and dilation for noise removal, opening and closing operations to close gaps (which are commonly assumed to be the size or smaller than the smallest object in an image) and remove isolated objects, distance-based techniques for enhancing edge continuity.

Discriminative representations are then extracted from the fine detection results and then fed to the following classification. These features consist of geometric (crack length, width, area and orientation) spacial (arrangement of crack with respect to cubefaces), textural (roughness and variation in intensity) and morphological features (connectivity and structural integrity).

Combined feature calculation throughout all the six faces allows a global view of fracture characteristics instead of studying breaking sequence on an individual surface basis.

2.7 Failure Mode Classification Logic

Failure modes are classified according to explicit engineering criteria based on the standard stiffness test procedures. This rule-based methodology is to be adopted, so that a model is interpretable in terms of domain knowledge and makes up for the drawbacks of black-box classifiers.

The decision logic is simply stated as:

- If cracks or damage are observed at the top and bottom, this failure mode is considered unacceptable because it represents incorrect load processing or preparation of the test specimen.
- If the upward and downward sides are intact and all \(\times\) four lateral ones show a unified liner crack pattern giving an even observation, failure mode is rated as pass.
- All other conditions are not acceptable.
- By integrating engineering into the classification procedure, the proposed system is able to close the loop between a (model based) action and an interpretable decision.

2.8 Detection Fusion and System Implementation

For the combination of detection and morphological processing outputs, an Intersection over Union (IoU)-based fusion approach is used. IoU measures the spatial overlap between bounding boxes, and thus permits robust matching of detections from different processing paths.

The overall procedure that involves detection and refinement, fusion and classification stages can be represented by Fig. 5.




Fig. 5: System implementation details

This fusion scheme strengthens the robustness to localization error and existence of consistent crack representation before classification.

2.9 Training Configuration and Reproducibility

The model is optimized hyper-parameters through empirical tuning to ensure convergence and generalization. The full training configuration, learning rate, batch size and optimizer selection for each model is described in Table 3.

Table 3. Training Configuration and Hyperparameters

Parameter	Value / Description
Deep learning framework	YOLOv8
Backbone architecture	CSPDarknet
Input image size	640 × 640
Batch size	16
Number of epochs	300
Learning rate	0.001
Optimizer	Adam

Loss function	CloU loss + classification loss
IoU threshold	0.5
Confidence threshold	0.25
Data split ratio	70% training, 20% validation, 10% testing
Hardware accelerator	GPU
Training mode	Supervised learning
Augmentation applied	Yes (training set only)
Early stopping	Not applied

All the experiments are performed with fixed random seed and same data partitions for reproducibility. The modular structure of the framework means that future investigators can reproduce or expand upon this methodology without having to redefine the core decision logic.

3. Results and Discussion

3.1 Object Detection Performance

3.1.1 Concrete Cube Detection Accuracy

The YOLOv8-based detection framework was tested under the standard object detection metrics that include precision, recall, F1-score and mean average precision (mAP). The experiment results prove that the detection module exhibits stable performance on both of original and balanced datasets, indicating that it is robust to detect the concrete cube specimens with different surface texture and under different illuminations.

The high precision shows that the system has little probability of creating false detect, and high recall also indicates that cube specimens can be detected without missing. The achieved mAP@0.5 demonstrate also a good spatial localization accuracy since they are close to unity. This level of performance is especially critical for further failure mode classification, because if cubes are localized incorrectly, errors may get propagated until the final stages of crack analysis and decision-making.

It is shown in Table 4 that performance as compared with manual inspection, traditional computer vision methods, and base deep learning models. Our framework significantly outperforms all of the reference approaches, confirming that subjective manual examination is limited and the robustness to complex fracture patterns of classical image processing methods is not enough.

Table 4. Performance Comparison of Different Inspection Methods

Method	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	Processing Time per Cube
Manual visual inspection	68.0	65.2	66.5	65.8	~5–7 min
Classical image processing	75.4	73.1	74.6	73.8	~2–3 min
CNN-based crack detection	88.6	87.2	86.9	87.0	~45–60
YOLOv8-based detection (without rules)	91.2	90.4	89.7	90.0	~18 s

Proposed framework (original dataset)	94.3	94.8	93.9	94.3	~12 s
Proposed framework (balanced dataset)	96.8	97.1	96.5	96.8	~12 s

3.1.2 Training Convergence and Stability Analysis

The training and validation performance of the YOLOv8 detection model on the entire period of training is shown in Fig. 6.




Fig. 6 Training and validation metrics over 300 epochs

As shown in Fig. 6, the curves of training and validation loss keep decreasing monotonically, which confirms that discriminative visual features are well learned. The similarity between the training and validation curves indicates that the model does not overfit to the training data. Value of performance metrics including precision, recall, and mAP converges in the first third or so of training epochs which means fast convergence and good tuning of learning parameters.

Lack of a significant sharp decline or divergence between training and validation curves shows that the chosen hyperparameters (along with optimization methodology) are stable. This stability is particularly important in practical deployments, as model retraining should generate predictable and reproducible performance results.

3.2 Crack Detection and Failure Mode Classification Performance

3.2.1 Multi-Class Crack Detection Results

The crack detection device in the proposed setup was tested under a variety of surface conditions (cracked or smooth surfaces). The model exhibits a good degree of discrimination, with high precision and recall for all the considered categories. This affirms the performance of the approach in detecting crack patterns, including for prominent and stippled cracks that are important for accurate failure mode determination.

Reliability of detecting smooth surfaces which can be false cracks on sound loading faces is especially critical; such detection could result in an incorrect categorization of otherwise acceptable failure modes. The class-averaged balanced performance of the model suggests that it is not biased towards crack-dominant surfaces, a common issue in crack-centric detection systems.

In general, these findings illustrate that the fell detection system is a robust visual platform which may be used to enable higher level engineering interpretation rather than being a stand-alone crack detection tool.

3.2.2 Confusion Matrix Analysis

The confusion matrix for failure mode classification is presented in **Fig. 7**.




Fig. 7 Confusion matrix - crack cube detection

The correct classification of acceptable and non-acceptable failure modes is high based on confusion matrix. True positives and true negatives are much greater than false classifications, suggesting the high reliability of decisions. The system is characterised by an outstanding accuracy with respect to detecting non-acceptable failure modes, which are of primary concern for the structural safety.

Small false-positive ratio implies few good concrete samples are misclassified as bad ones, which minimizes material rejections and the corresponding cost. On the other hand, the low false negative rate confirms that the system can efficiently recognize defective specimens and reduce risk of structurally compromised concrete in construction workflows.

These findings justify the synergy of deep learning-based detection and explicit engineering decision rules.

3.3 Precision-Recall Trade-off Analysis

The precision-recall relationship for cube detection is illustrated in **Fig. 8**.




Fig.8 Precision-Recall Curve for Cube Detection

This means that the detection model can uniformly exchange sensitivity and specificity for each recall values. The area under the precision–recall curve is close to unity, indicating the stability of this approach for detection over a range of confidence thresholds.

This directly indicates that the operational thresholds are set to their most optimal point in terms of false guild detections and missed organisms. This level of stability is needed for integration into automatic inspection systems, for instance due to a divergent threshold performance volatile quality control results may be obtained.

3.4 Comparative Performance Analysis

A full comparative analysis has been carried out to compare the proposed framework with other inspection techniques. It can be seen from Table 4 that the performance results of the proposed method are better than those of other baselines in terms of all evaluation measures including accuracy, precision, recall and time.

The baselines of the framework outperform previous approach through balanced dataset training and refined decision engine. The processing time per pipeline is significantly faster, on the order of 12 secs/cube, speeding up manual inspection workflows that often take minutes for each specimen.

This efficiency increase allows for feasible large-scale ring deployment without interfering with ongoing lab or site testing.

3.5 Statistical Validation and Reliability Assessment

Statistical validation of the developed method was achieved through repeated experimental runs to evaluate its ruggedness and robustness. The corresponding confidence intervals and measures of variability are also presented in Table 5.

Table 5. Statistical Validation and Reliability Analysis

Metric	Original Dataset	Balanced Dataset
Mean classification accuracy (%)	94.3	96.8
Standard deviation (%)	1.12	0.84
Precision (%)	94.8	97.1
Recall (%)	93.9	96.5
F1-score (%)	94.3	96.8
Confidence interval (95%)	± 2.1	± 1.6
Number of experimental runs	5	5

The small confidence intervals and standard deviation values as shown in Table 5 indicate that system performance is stable and repeatable. This consistency is essential in quality control tasks where automated decisions should be reliable across lots, test environments and shifts.

The numerical results validate that the superior performance is not fortuitous but indicative of a consistent edge of our approach.

3.6 Discussion

The results demonstrated also that the proposed method overcomes some of the common drawbacks associated with conventional compression tests of concrete cubes. A deep learning based visual analysis with engineering driven decision rules allows the system to be both accurate and interpretable (which does not necessarily

goes together in an automated inspection).

In contrast to the studies available in the literature (in which only detection of a surface crack is considered, with or without taking into account through-thickness cracks) this approach deals with failure modes in an engineering oriented way by testing on close packet of cube all faces and it explicitly tests deformability of loading surfaces. This multi-view framework provides a robust classification of acceptable and unacceptable failures that is congruent to standard practice.

It is critical to balance the two during model learning and enhancements in order to be robust in classification, especially for those acceptable failure modes that are underrepresented. In addition, the less computation time and convergence stability illustrate that the system is suitable for both laboratory and construction site.

In conclusion, the results verify that combining domain expertise with deep learning architectures provides a feasible and scalable path to achieve objective, consistent, and sustainable assessment of concrete quality.

4. Conclusions

This paper introduced an automated solution to classify the failure modes of concrete cubes, which combines deep learning based visual analysis with engineering-guided decision logic. The proposed method aims to overcome a fundamental shortcoming in existing concrete quality evaluation methodologies, for which failure mode interpretation continues to be primarily based on human judgment, among other factors associated with standardized compressive testing.

Integrating YOLOv8-based object and crack detection with a six-face inspection protocol and explicit engineering guidelines, the system progresses beyond generic crack detection to engineering-valid failure mode recognition. Experimental results showed that the system has high classification accuracy, precision and recall for identifying acceptable and non-acceptable failures so that it can consistently discriminate whether conductive particles are developed on misaligned substrates under realistic test conditions. The results also indicated that the proposed method is greatly superior to hand inspection, classical image processing methods and baseline DNN models as well as saving a lot of time for visual inspection.

One of major advantage of the proposed model is its interpretability. Instead of straight black-box predictions, the ultimate decision is guided by rules resulting from accepted compressive testing practices. Such integration of domain expertise guarantees that fully autonomous decisions still conform to engineering principles and build trust encouraging the employment in laboratory as well as field applications.

The inclusion of multi-view six-face analysis was found to be crucial for accurate failure characterization. Machining all faces of the cube allowed a standard interpretation of fracture patterns, and to make sure damage on loading sides (a telltale sign of not-acceptable failure) was not missed. Moreover, the employment of morphological operation refinement and IoU-based detection fusion achieved stiffness crack locations that are consistent and reliable for the subsequent classification.

It is concluded that the proposed model is an objective, consistent, and expeditious method for determination of failure mode in concrete cube. The method has direct applications in automated quality control procedures and the improvement of robustness and durability of constructions.

5. Future Research Directions

Although the proposed model shows promising results, there are some issues that can be addressed in future to improve its performance and generalization capability.

First, two-dimensional image analysis underlies the current implementation. Generalizing the framework to include 3D surface reconstruction or depth-sensitive imaging, would also allow for volumetric characterization of fractures and more complete examination of complicated failure modes that might be obscured in 2D.

Second, while the rule-based decision logic guarantees interpretability, hybrid strategies which mix engineering rules and adaptive learning should be considered in future works. Such methods, for example, may enable the system to compensate decision thresholds according to different standards or material properties or regional test strategies and still present transparency.

Third, the dataset employed in this study is of standard cube shapes and controlled imaging configuration. Larger variations in cube sizes, surface norms and the environment used for training can counteract overfitting and enhance robustness when deployed on real-world systems.

Fourth, the method of lightweight model optimization and edge computing could be explored, for making it deployable on mobile devices and embedded systems. It would enable on-site inspection using handheld cameras or smart testing devices not dependent on high-level computing infrastructure.

Last but not least, there is potential for integration of the presented theory with digital construction environments. Integrating automated failure mode identification with internet of things (IoT) enabled testing devices, centralized data sets and secure record keeping would allow for end-to-end quality monitoring, traceability, and data-informed decision-making throughout the construction process.

6. References

1. Scrivener, K.L., John, V.M., & Gartner, E.M. (2018). Eco-efficient cements: Potential economically viable solutions for a low-CO₂ cement-based materials industry. *Cement and Concrete Research*, 114, 2–26.
2. Mehta, P.K., & Monteiro, P.J.M. (2019). *Concrete: Microstructure, Properties, and Materials* (4th ed.). McGraw-Hill Education.
3. Neville, A.M. (2020). *Properties of Concrete* (6th ed.). Pearson Education.
4. Sadouskaya, A., Leanovich, S., Budrevich, N.A., & Sadovskaya, E.A. (2021). Compressive strength of steel fibre reinforced concrete obtained by testing cubes and cylinders. *MATEC Web of Conferences*, 350, 00008.
5. Halstead, P.E. (1969). The significance of concrete cube tests. *Magazine of Concrete Research*, 21(69), 187–194.
6. Mukhopadhyay, T., Naskar, S., Dey, S., & Chakrabarti, A. (2019). Condition assessment and strengthening of aged structures. *Practice Periodical on Structural Design and Construction*, 24(3), 04019013.
7. Kim, J., Cho, Y., Lee, J., & Kim, Y.H. (2022). Defect detection and characterization in concrete based on FEM and ultrasonic techniques. *Materials*, 15(22), 8160.
8. Shahrokhinasab, E., Hosseinzadeh, N., Monirabbasi, A., & Torkaman, S. (2020). Performance evaluation of image-based crack detection systems in concrete structures. *Journal of Infrastructure Systems*, 26(4), 04020039.

9. Dorafshan, S., Maguire, M., & Hoffer, N. (2018). Challenges in automated crack detection of concrete structures. *Computer-Aided Civil and Infrastructure Engineering*, 33(8), 642–660.
10. Andrew, R.M. (2019). Global CO₂ emissions from cement production. *Earth System Science Data*, 11, 1675–1710.
11. Miller, S.A., Horvath, A., & Monteiro, P.J.M. (2021). Readily implementable techniques can cut annual CO₂ emissions from the global cement industry by at least 20%. *Environmental Science & Technology*, 52(9), 5560–5569.
12. Malhotra, V.M., & Carino, N.J. (2019). *Handbook on Nondestructive Testing of Concrete* (2nd ed.). CRC Press.
13. Lorenzi, A., Tisbieriek, F., & Dal Bó, R. (2020). Ultrasonic testing for damage assessment of concrete structures. *Construction and Building Materials*, 247, 118517.
14. Taffese, W.Z., & Sistonen, E. (2017). Machine learning for durability and service-life assessment of reinforced concrete structures. *Automation in Construction*, 77, 1–14.
15. Abdel-Qader, I., Abudayyeh, O., & Kelly, M.E. (2020). Analysis of edge-detection techniques for crack identification in concrete. *Journal of Computing in Civil Engineering*, 34(2), 04019045.
16. Savio, A.A.D., Torres, A.L., Salas, D., Olivera, M.A.V., & Ibarra, G.T.U. (2023). Detection and evaluation of construction cracks through image analysis using computer vision. *Applied Sciences*, 13(17), 9662.
17. Cha, Y.J., Choi, W., & Büyüköztürk, O. (2017). Deep learning-based crack damage detection using convolutional neural networks. *Computer-Aided Civil and Infrastructure Engineering*, 32(5), 361–378.
18. Xu, Y., Bao, Y., Chen, J., Zuo, W., & Li, H. (2019). Surface crack detection using deep learning-based fully convolutional networks. *Computer-Aided Civil and Infrastructure Engineering*, 34(7), 567–582.
19. Li, S., Zhao, X., Zhou, G., & Du, K. (2021). Deep learning-based crack detection for concrete structures: A review. *Automation in Construction*, 125, 103606.
20. Zhang, L., Yang, F., Zhang, Y.D., & Zhu, Y.J. (2020). Road crack detection using deep convolutional neural networks. *IEEE Transactions on Intelligent Transportation Systems*, 21(4), 1525–1535.
21. Zou, Q., Cao, Y., Li, Q., Mao, Q., & Wang, S. (2021). CrackTree: Automatic crack detection from pavement images. *Neurocomputing*, 408, 161–171.
22. Lin, T.Y., Dollár, P., Girshick, R., He, K., Hariharan, B., & Belongie, S. (2020). Feature pyramid networks for object detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(5), 2733–2747.
23. Bochkovskiy, A., Wang, C.Y., & Liao, H.Y.M. (2020). YOLOv4: Optimal speed and accuracy of object detection. *IEEE Conference on Computer Vision and Pattern Recognition Workshops*.
24. Redmon, J., & Farhadi, A. (2018). YOLOv3: An incremental improvement. *arXiv preprint* → (used only conceptually; implementation basis)
25. Downey, A., D'Alessandro, A., Ubertini, F., & Laflamme, S. (2018). Automated crack detection in conductive smart-concrete structures. *Measurement Science and Technology*, 29(3), 035107.
26. Deng, L., Wang, J., Xu, S., & Zhang, Y. (2022). Dataset challenges and solutions for concrete crack detection using deep learning. *Automation in Construction*, 136, 104158.
27. Shorten, C., & Khoshgoftaar, T.M. (2019). A survey on image data augmentation for deep learning. *Journal of Big Data*, 6, 60.
28. Frid-Adar, M., et al. (2018). GAN-based synthetic medical image augmentation. *Neurocomputing*, 321, 321–331.

29. BS EN 12390-3:2019. Testing hardened concrete – Compressive strength of test specimens. British Standards Institution.