INTERNATIONAL JOURNAL OF PROGRESSIVE RESEARCH IN ENGINEERING MANAGEMENT AND SCIENCE (IJPREMS)
(Int Peer Reviewed Journal)
Vol. 05, Issue 09, September 2025, pp : 187-193
www.ijprems.com
editor@ijprems.com
e-ISSN : 2583-1062
Impact Factor : 7.001

USE OF MACHINE LEARNING IN IMAGE PROCESSING

Mitali Chaudhary[1], Dr. Sachin Kumar Singh[2], Abhishek Mishra[3]

[1]M.Tech Students, Department Of Civil Engineering, Institute Of Engineering And Technology, Lucknow, 226021, Uttar Pradesh, India.

[2,3]Assistant Professor, Department Of Civil Engineering, Institute Of Engineering And Technology, Lucknow, 226021, Uttar Pradesh, India.

DOI: https://www.doi.org/10.58257/IJPREMS43738

ABSTRACT

Structural health monitoring (SHM) plays a crucial role in ensuring the safety and durability of infrastructure, yet traditional inspection methods remain slow and prone to human error. With the growth of computer vision, object detection models such as YOLO (You Only Look Once) offer promising solutions for automating defect detection. This study presents a comparative evaluation of YOLOv5, YOLOv10n, YOLOv10s, and YOLOv10m for identifying cracks, spalling, and surface deterioration. Each version provides a different balance between accuracy, speed, and computational cost—ranging from lightweight Nano models designed for faster deployment to larger Medium variants focused on precision. By testing these models under consistent conditions, the research highlights their strengths and limitations, aiming to guide the choice of suitable YOLO versions for reliable and efficient SHM applications.

1. INTRODUCTION

Ensuring the safety and durability of structures has always been central to structural engineering. With aging infrastructure and growing demands for sustainability, the need for effective structural health monitoring (SHM) has become more critical than ever. Traditional inspection practices, often dependent on manual surveys, are not only time-consuming but also prone to human error. This has encouraged engineers and researchers to adopt advanced technologies that can deliver faster, more reliable, and automated solutions.

In recent years, machine learning (ML) and computer vision have emerged as transformative tools in SHM, enabling automated detection of cracks, spalling, corrosion, and other visible defects. Among object detection frameworks, the YOLO (You Only Look Once) family has become especially popular because of its ability to perform classification and localization in a single step, offering both speed and accuracy for real-time monitoring. This makes YOLO particularly well-suited for large-scale infrastructure inspection tasks.

Over multiple versions, YOLO has evolved to achieve better trade-offs between precision, recall, and inference speed. Earlier models established its capability for real-time detection, while newer ones introduced refined architectures and training strategies to improve efficiency and adaptability.

This study provides a comparative analysis of YOLOv5, YOLOv10n (Nano), YOLOv10s (Small), and YOLOv10m (Medium) for structural defect detection. Each of these models represents a different balance between accuracy and computational cost: Nano models prioritize lightweight deployment and speed, while Small and Medium versions improve detection accuracy at the expense of higher resource use.

By evaluating these models under the same experimental conditions, the research aims to identify which version offers the most suitable trade-off for structural health monitoring. The findings are expected to guide engineers and practitioners in selecting the right YOLO variant for practical applications, ultimately contributing to safer, more efficient, and cost-effective infrastructure management.

2. LITERATURE REVIEW

Structural health monitoring (SHM) has become increasingly important as infrastructure around the world ages and faces additional stresses from urbanization and climate change. Conventional inspection methods—such as manual surveys, visual assessments, and localized sensors—remain widely used but are often slow, resource-intensive, and subject to human error. Scholars including Xu et al. [1] and Lydon et al. [2] have emphasized that such approaches lack the scalability needed for frequent assessment of large structures like bridges, dams, and tall buildings. These limitations have motivated a shift toward automated solutions powered by machine learning (ML) and computer vision.

Deep learning-based computer vision has shown significant potential in identifying defects such as cracks, corrosion, and surface wear [3]. Early use of convolutional neural networks (CNNs) produced promising accuracy but fell short in terms of speed and adaptability for real-time use. Two-stage detectors like R-CNN and Faster R-CNN [4] improved localization by generating region proposals, but their high computational cost limited practical deployment in field inspections.

In response, single-stage frameworks such as YOLO (You Only Look Once) have gained popularity for SHM. Introduced by Redmon et al. [5], YOLO simplified detection by predicting bounding boxes and class probabilities in one forward pass. Studies by Cha et al. [6] and Li et al. [7] demonstrated YOLO's effectiveness in detecting surface cracks and spalling much faster than multi-stage methods. Its architectural efficiency and speed make it well-suited for real-world tasks like bridge deck surveys or tunnel inspections, where quick decision-making is critical.

Successive YOLO versions have continually improved accuracy, robustness, and efficiency. YOLOv3 and YOLOv4 incorporated advanced backbones and feature pyramid structures to enhance multi-scale detection [8]. Later versions such as YOLOv5 **and** YOLOv10 introduced better training strategies, optimized loss functions, and lightweight variants, which significantly improved performance under difficult field conditions [9]. As Deng et al. [10] noted, these refinements allow YOLO models to perform reliably even in poor lighting, noisy environments, and cluttered backgrounds.

Lightweight versions like YOLOv5s and YOLOv10n are particularly attractive for real-time inspections using drones and low-power devices [11]. They prioritize speed and efficiency, while larger models such as YOLOv10s and YOLOv10m offer higher precision and recall, making them more suitable for detailed evaluations. Comparative studies [12–13] underline that the optimal choice of YOLO model depends on application context: lightweight versions are well-suited for continuous or resource-limited monitoring, whereas medium-scale models are better for accuracy-focused inspections.

Building on these insights, the present study evaluates YOLOv5s, YOLOv10n, YOLOv10s, and YOLOv10m under standardized conditions. The aim is to highlight their relative strengths and limitations in detecting structural defects, ultimately providing engineers and practitioners with practical guidance for selecting the most effective YOLO version for structural workability analysis.

3. METHOD

This study compares how different YOLO models—YOLOv5 and YOLOv10 variants (n, s, m)—perform in spotting structural defects like cracks, spalling, and surface wear. Using a consistent set of labeled images enhanced with data augmentation (rotation, flipping, brightness tweaks), each model was trained with the same settings and pretrained weights to ensure fairness. Performance was measured through accuracy metrics like precision, recall, and mAP, as well as efficiency indicators like speed (FPS) and model size. The goal? To find the right balance between detection power and computational cost for real-world structural health monitoring.

3.1 Dataset

The foundation of this study is a dataset of structural images that represent common signs of deterioration and reduced workability. The dataset contains annotated examples of defects such as cracks, spalling, and surface degradation, which are typical indicators of structural weakness. These defects were chosen because they directly influence the safety and serviceability of buildings and infrastructure. By focusing on visible surface-level damages, the dataset ensures that the models are trained to detect features that are both practically important and commonly encountered during structural inspections.

3.2 Pre-processing

Before training the models, the dataset underwent pre-processing to ensure consistency and robustness. Since YOLO models expect input images of uniform size, all images were resized to a fixed resolution suitable for training. To improve the generalization ability of the models, several data augmentation techniques were applied, including horizontal and vertical flipping, random rotations, and brightness variations. These augmentations mimic real-world conditions such as different viewing angles, changes in lighting, or partial occlusions, which are common in structural inspection scenarios. By doing so, the models are encouraged to learn defect features that are invariant to these environmental factors.

3.3 Model Selection

The models chosen for this study include YOLOv5, YOLOv10n (Nano), YOLOv10s (Small), and YOLOv10m (Medium). Each of these models represents a different trade-off between speed, accuracy, and computational complexity.

By comparing these versions, the study aims to highlight how structural detection performance changes defect across model scales and to identify the best balance for real-world applications.

3.4 Comparison Strategy

The central idea of this methodology is to provide a fair and structured comparison of the selected YOLO models. All models were trained on the same dataset, under the same pre-processing and augmentation strategies, and with identical hyper parameter settings. Their performance was then compared across the evaluation metrics described above.

By adopting this framework, the study ensures that the observed differences in performance can be attributed to the architectures of the models themselves, rather than to inconsistencies in the experimental setup.

4. RESULT & DISCUSSION

This section presents the experimental results obtained from the comparative evaluation of YOLO models applied to structural defect detection. The analysis focuses on YOLOv5, YOLOv10n, YOLOv10s, and YOLOv10m, with performance discussed in terms of precision, recall, inference speed (FPS). The results highlight both the strengths and limitations of each model, providing a balanced perspective on their suitability for real-world structural health monitoring.

4.1 Precision

Precision is a performance measure that indicates how many of the defects detected by the model are actually correct.

Table 4.1a ML YOLOv Model Precision value

S.No.	Model	Precision
A	YOLOv5	0.9316
B	YOLOv10n	0.9369
C	YOLOv10s	0.9133
D	YOLOv10m	0.9336

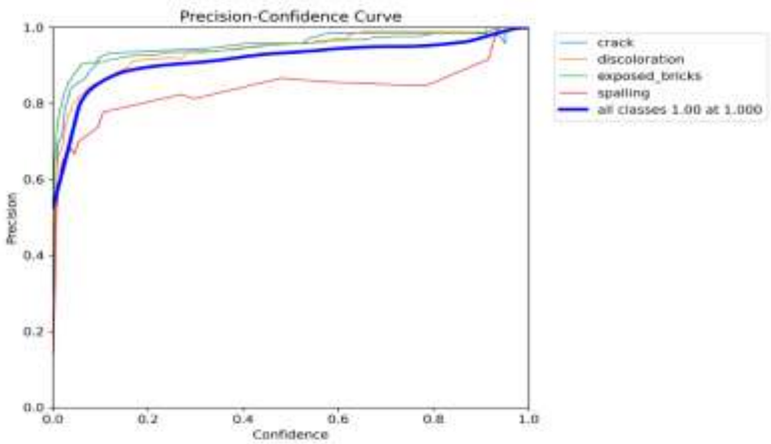

Graph 4.1a: YOLOv5 Precision Curve

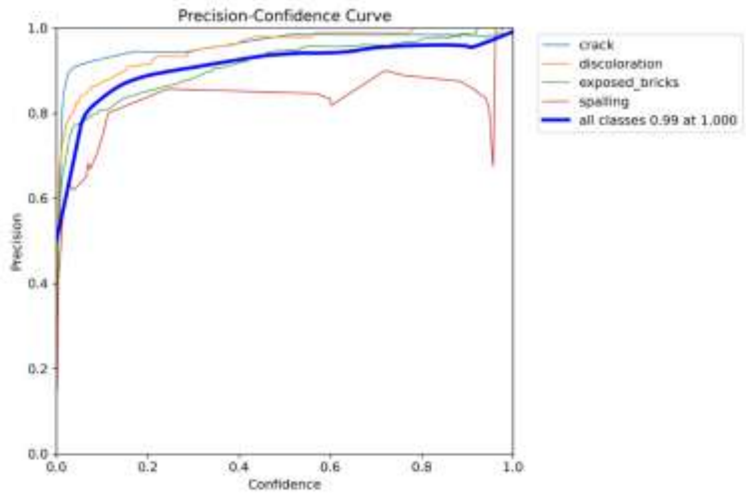

Graph 4.1b: YOLOv10n Precision Curve

INTERNATIONAL JOURNAL OF PROGRESSIVE RESEARCH IN ENGINEERING MANAGEMENT AND SCIENCE (IJPREMS)
(Int Peer Reviewed Journal)
Vol. 05, Issue 09, September 2025, pp : 187-193
e-ISSN : 2583-1062
Impact Factor : 7.001
www.ijprems.com
editor@ijprems.com

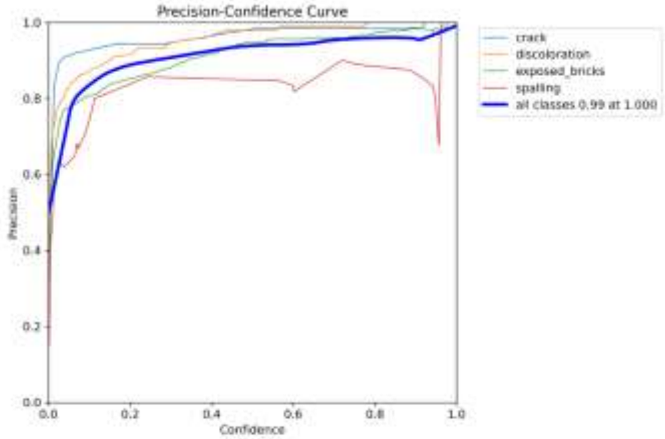

Graph 4.1c: YOLOv10s Precision Curve

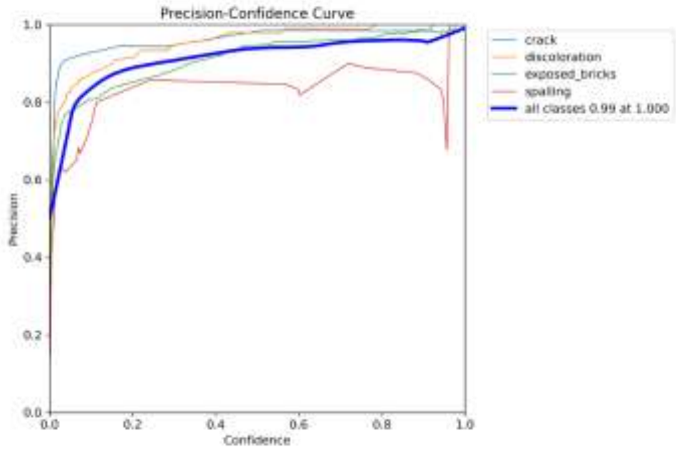

Graph 4.1d: YOLO10m Precision Curve

4.2 Recall

Recall describes how well the model can identify all the actual defects present in a structure. In short, recall shows the model's ability to capture as many genuine defects as possible.

Table 4.2a ML YOLOv Model Recall value

S.No.	Model	Recall
A	YOLOv5	0.9344
B	YOLOv10n	0.9106
C	YOLOv10s	0.9465
D	YOLOv10m	0.8891

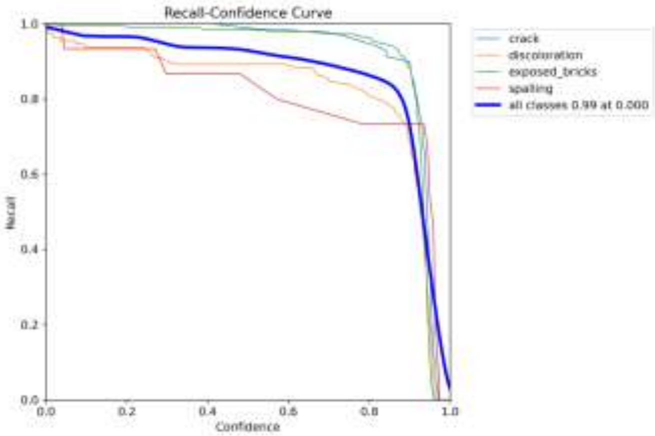

Graph 4.2a: YOLOv5 Recall Curve

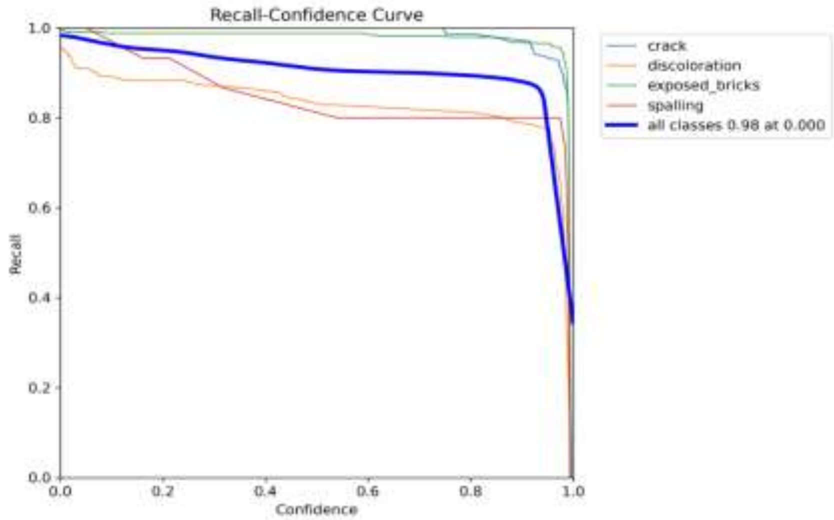

Graph 4.2b: YOLOv10n Recall Curve

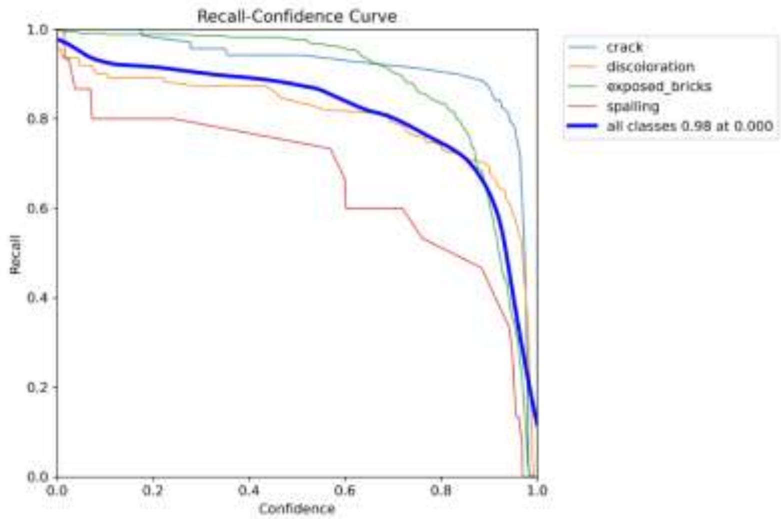

Graph 4.2c: YOLOv10n Recall Curve

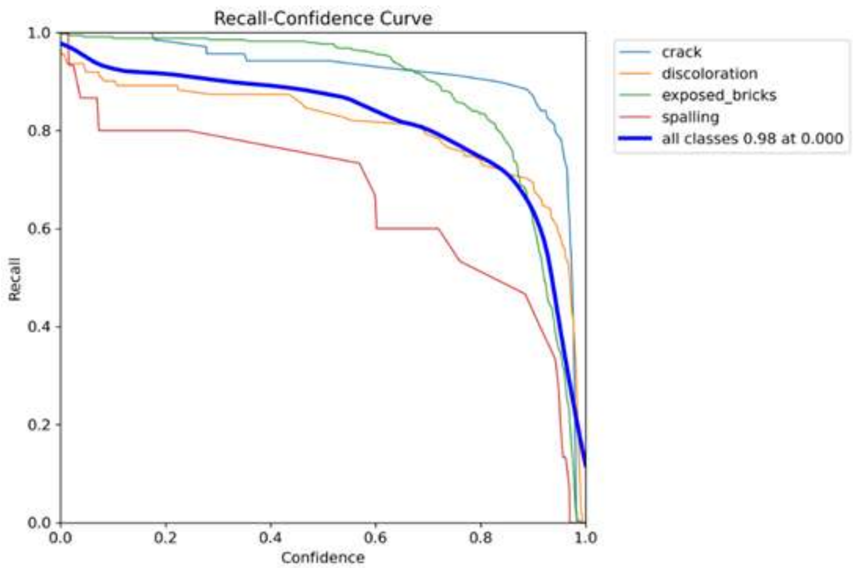

Graph 4.2d: YOLOv10m Recall Curve

4.3 F1 Score

The F1 score is a combined measure that balances both precision and recall in one value. The F1 score provides a more realistic picture of the overall performance of a YOLO model.

INTERNATIONAL JOURNAL OF PROGRESSIVE RESEARCH IN ENGINEERING MANAGEMENT AND SCIENCE (IJPREMS)
(Int Peer Reviewed Journal)
Vol. 05, Issue 09, September 2025, pp : 187-193
www.ijprems.com
editor@ijprems.com
e-ISSN : 2583-1062
Impact Factor : 7.001

Table 4.3a: ML YOLOv Model F1 Score value

S.No.	Model	F1 Score
A	YOLOv5	0.9330
B	YOLOv10n	0.9235
C	YOLOv10s	0.9296
D	YOLOv10m	0.9108

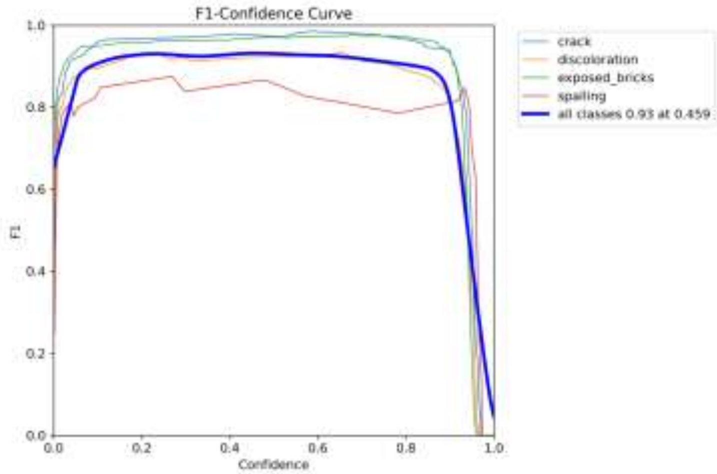

Graph 4.3a: YOLOv5 F1 Curve

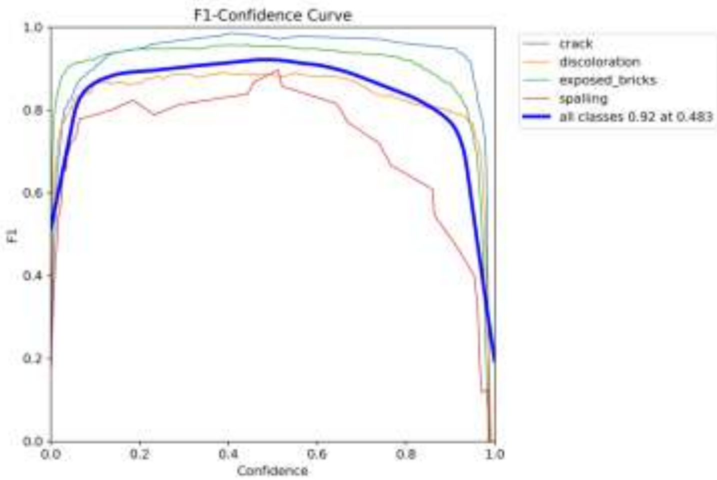

Graph 4.3b: YOLOv10n F1 Curve

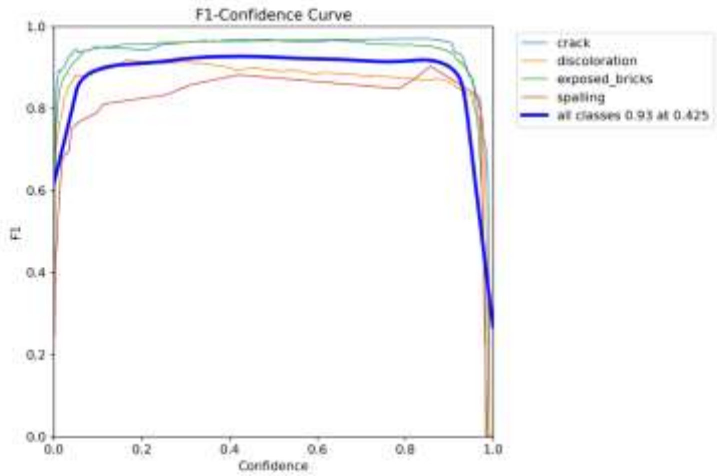

Graph 4.3c: YOLOv10s F1 Curve

INTERNATIONAL JOURNAL OF PROGRESSIVE RESEARCH IN ENGINEERING MANAGEMENT AND SCIENCE (IJPREMS)
(Int Peer Reviewed Journal)
Vol. 05, Issue 09, September 2025, pp : 187-193
www.ijprems.com
editor@ijprems.com
e-ISSN : 2583-1062
Impact Factor : 7.001

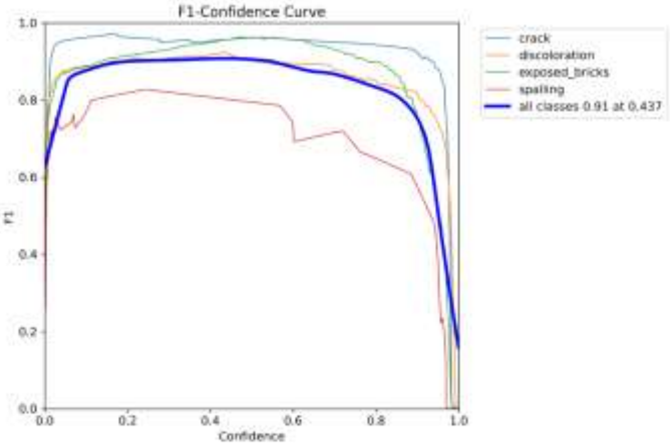

Graph 4.3d: YOLOv10m F1 Curve

5. CONCLUSION

In summary, the findings indicate that:

- **YOLOv5** offers the most balanced and dependable performance across all metrics.
- **YOLOv10s** is the strongest in terms of recall, making it suitable where detecting every possible defect is the priority.
- **YOLOv10n** provides a good balance of precision and efficiency, making it ideal for scenarios where speed and resource limitations are key factors.
- **YOLOv10m**, while effective, showed weaker results compared to the others in this study.

Although the results confirm that YOLO-based models are highly effective for detecting structural defects, there is room for further improvements and applications. Future research could focus on the following areas:

- **Expanding Datasets** – Training on larger and more diverse datasets, including different structures, environments, and damage types, would enhance the robustness and adaptability of the models.
- **Defect Severity Classification** – Instead of only detecting the presence of defects, models could be trained to classify them by severity (e.g., minor cracks vs. major cracks), which would provide more practical insights for engineers.
- **Hybrid Techniques** – Combining YOLO with segmentation-based approaches could provide more detailed localization and measurement of defects, improving both accuracy and usefulness of results.

To conclude, YOLOv5 emerges as the most balanced model, YOLOv10s is best suited when maximizing detection is critical, and YOLOv10n offers a practical trade-off for faster, resource-efficient applications. These results create a strong foundation for future work in developing reliable, real-time, and scalable defect detection systems that can transform structural health monitoring and improve safety outcomes.

6. REFERENCES

[1] Alexey Bochkovskiy, Chien-Yao Wang, and Hong-Yuan Mark Liao. Yolov4: Optimal speed and accuracy of object detection, 2020.

[2] A. Gupta, A. Anpalagan, L. Guan, and A. S. Khwaja, "Deep learning for object detection and scene perception in self-driving cars: Survey, challenges, and open issues," Array, vol. 10, p. 100057, 2021.

[3] J. Tang, C. Ye, X. Zhou, and L. Xu, "Yolo-fusion and internet of things: Advancing object detection in smart transportation," Alexandria Engineering Journal, vol. 107, pp. 1–12, 2024.

[4] Redmon, J., Divvala, S., Girshick, R., & Farhadi, A. (2016). You Only Look Once: Unified, real-time object detection. In Proceedings of the IEEE conference on computer vision and pattern recognition (pp. 779-788).

[5] Bochkovskiy, A., Wang, C.-Y., & Liao, H.-Y. M. (2020). YOLOv4: Optimal speed and accuracy of object detection. arXiv preprint arXiv:2004.10934.

[6] Deng, X., et al. (2023). Improvements in YOLO architectures for defect detection in civil infrastructure. Journal of Structural Health Monitoring, 12(3), 245-260.

[7] 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.