GLOBAL JOURNAL
OF COMPUTER SCIENCE AND TECHNOLOGY - D: NEURAL & AI

Volume 26 | 2026 | Article ID: GJCST-259447 | Language: English

A Deep Learning Framework for Industrial Pipeline Defect Detection using YOLOv8

Short Title: YOLOv8-OBB Pipeline Defect Detection

Online First

Manel Elleuchi [§] [*]

[*]Corresponding Author

Ahmed Fakhfakh

[§] Digital Research Centre of Sfax, Sfax, Tunisia [(OA)]

RECEIVED	ACCEPTED	ONLINE PUBLISHED	PEER REVIEW
2026-07-23	2026-08-01	2026-08-24	Double Blind

Abstract

Industrial pipeline inspection is a critical requirement in sectors such as water distribution, oil transportation, and manufacturing. Conventional inspection approaches remain largely manual, which leads to high operational costs, long inspection durations, and inconsistent defect detection due to human subjectivity. These limitations highlight the need for automated, accurate, and real-time inspection systems capable of reliable defect identification. In this work, we propose a robust real-time defect detection framework based on the YOLOv8 architecture enhanced with Oriented Bounding Boxes (YOLOv8-OBB). The proposed approach is specifically designed to handle elongated and arbitrarily oriented defects, such as cracks and leaks, which are common in industrial pipeline environments. The model is trained and evaluated on a custom dataset comprising 1,224 annotated images distributed across four defect categories: crack, dent, hole, and leak. Extensive experimental results demonstrate strong performance, achieving 83.92% precision, 86.06% recall, and 87.22% mAP@50, while maintaining real-time inference capability with processing times between 5 and 10 milliseconds per image. The results show the proposed system provides an efficient and scalable solution for intelligent industrial inspection and demonstrates strong potential for integration into real-world pipeline monitoring and maintenance platforms.

YOLOv8 | Industrial Pipeline Inspection | Deep Learning

AI USE STATEMENT
No generative AI was used for analysis or results.

FUNDING
No external funding was declared for this work.

CONFLICT OF INTEREST
The authors declare no conflict of interest.

DATA AVAILABILITY
Not applicable for this article.

ETHICS
No ethics committee approval was required for this article type.

CONSENT
Not applicable for this article.

TRIAL REG.
Not applicable.

Crossref DOI: 10.34257/GJCSTD259447

How to Cite: Elleuchi et al. (2026). A Deep Learning Framework for Industrial Pipeline Defect Detection using YOLOv8. Global Journal of Computer Science and Technology. Ahead of Print. DOI: 10.34257/GJCSTD259447

AR Experience Web Page DOI

Print ISSN 0975-4350

Online ISSN 0975-4172

Under the strict compliance and defined process of

METADATA CONTINUATION

AUTHOR CONTACT QR LEDGER

Manel Elleuchi[§*]

ARCHIVAL RECORD

GJCST · Vol 26 · 2026

Article ID GJCST-259447 · DOI 10.34257/GJCSTD259447

Print ISSN 0975-4350 · Online ISSN 0975-4172

Global Journal of Computer Science and Technology - Section D
GJCST

A Deep Learning Framework for Industrial Pipeline Defect Detection using YOLOv8

Manel Elleuchi [§*] and Ahmed Fakhfakh [§]

Affiliations
§ Digital Research Centre of Sfax, Sfax, Tunisia [(OA)]

Abstract

Industrial pipeline inspection is a critical requirement in sectors such as water distribution, oil transportation, and manufacturing. Conventional inspection approaches remain largely manual, which leads to high operational costs, long inspection durations, and inconsistent defect detection due to human subjectivity. These limitations highlight the need for automated, accurate, and real-time inspection systems capable of reliable defect identification. In this work, we propose a robust real-time defect detection framework based on the YOLOv8 architecture enhanced with Oriented Bounding Boxes (YOLOv8-OBB). The proposed approach is specifically designed to handle elongated and arbitrarily oriented defects, such as cracks and leaks, which are common in industrial pipeline environments. The model is trained and evaluated on a custom dataset comprising 1,224 annotated images distributed across four defect categories: crack, dent, hole, and leak. Extensive experimental results demonstrate strong performance, achieving 83.92% precision, 86.06% recall, and 87.22% mAP@50, while maintaining real-time inference capability with processing times between 5 and 10 milliseconds per image. The results show the proposed system provides an efficient and scalable solution for intelligent industrial inspection and demonstrates strong potential for integration into real-world pipeline monitoring and maintenance platforms.

Keywords: *YOLOv8, Industrial Pipeline Inspection, Deep Learning*

*** Corresponding Author**
Manel Elleuchi

DOI
10.34257/GJCSTD259447

1. Introduction

Industrial pipeline inspection is a critical requirement in sectors such as water distribution, oil transport, and urban infrastructure. Structural defects in pipelines, including cracks, dents, holes, and leaks, can lead to severe environmental hazards, economic losses, and operational disruptions if left undetected. Consequently, regular and accurate defect assessment is essential to maintain infrastructure integrity and ensure operational safety.

Traditional inspection methods rely heavily on manual visual examination or remotely operated vehicles (ROVs) equipped with closed-circuit television (CCTV) cameras. However, manual inspection is labor-intensive, time-consuming, subjective, and prone to human error, particularly when dealing with long pipeline networks and complex defect patterns. Furthermore, hostile underground environments pose significant challenges to human inspectors.

To address these limitations, automated computer-vision-based inspection systems have attracted growing interest in recent years. In particular, deep learning models, such as convolutional neural networks (CNNs) and real-time object detectors, have demonstrated high potential for automated defect localization and classification. Despite these advances, detecting pipeline defects remains challenging due to complex background textures, non-uniform illumination, variable defect geometry, and low contrast on pipeline surfaces.

Conventional object detection frameworks typically employ axis-aligned bounding boxes (AABBs) to localize target instances. While effective for compact objects, axis-aligned boxes often include substantial background clutter when applied to elongated, rotated, or irregular pipeline defects such as longitudinal cracks and diagonal leaks. This background noise degrades bounding box precision and leads to inaccurate Intersection over Union (IoU) evaluations during model training.

To overcome these issues, oriented object detection methods utilize oriented bounding boxes (OBBs), which incorporate an additional orientation angle to align closely with defect boundaries. In this paper, we propose an industrial pipeline defect detection framework based on the state-of-the-art YOLOv8-OBB architecture. By leveraging oriented bounding box regression, the proposed approach achieves precise defect boundary delineation, improved feature representation, and robust detection performance under diverse industrial conditions.

The main contributions of this work are summarized as follows:

- We develop a YOLOv8n-OBB defect detection model tailored for industrial pipeline inspection, enabling accurate detection of oriented and irregular defects in real time.
- We curate and annotate a comprehensive dataset comprising 1,224 industrial pipeline images across four defect categories (cracks, dents, holes, and leaks) using OBB representations.

- We conduct extensive quantitative and qualitative evaluations, demonstrating the effectiveness, reliability, and real-time capability of the proposed approach.
- We evaluate dataset enrichment strategies to assess model generalization across diverse environmental and surface conditions.

2. Related Work

Object detection has experienced significant progress with the evolution of deep learning architectures. Single-stage detectors such as the YOLO (You Only Look Once) family [9, 10, 1, 13, 6] have gained widespread popularity in industrial applications due to their balance between detection accuracy and computational speed.

In the context of industrial inspection, object detection models have been extensively applied to surface defect recognition [7]. Similarly, Wang et al. [14] provide a comprehensive review of YOLO-based architectures for industrial defect detection.

Several studies have applied deep learning to sewer and pipeline inspection. Cheng et al. [2] proposed an automated defect detection system using deep learning, while Yin et al. [16] demonstrated strong classification performance on real-world sewer datasets.

However, standard object detectors rely on axis-aligned bounding boxes, which limits their ability to accurately detect rotated or elongated defects. To address this limitation, oriented object detection methods have been introduced. Xia et al. [15] and Zhou et al. [17] investigated oriented object detection methods, showing that rotated bounding boxes significantly improve localization accuracy. Additionally, RoI Transformer [3] further enhances detection of oriented objects in complex scenes.

More recently, transformer-based architectures such as Swin Transformer [8] have shown promising results in vision tasks, although they remain computationally heavier for real-time industrial deployment.

3. Proposed Methodology

3.1. Dataset Preparation

The dataset considered in this study contains 1,224 images collected from real industrial pipeline inspection environments. These images include various types of defects commonly encountered in pipeline systems, namely cracks, dents, holes, and leaks. The data were acquired under different environmental conditions, including variations in lighting, background complexity, and pipe materials, to ensure diversity and realism.

All images were manually annotated using the Roboflow platform, following the oriented bounding box (OBB) format. Unlike conventional axis-aligned annotations, OBB labeling allows each defect to be represented with its exact orientation, which is particularly important for elongated and rotated defects such as cracks and leaks.

To enhance the robustness and generalization ability of the proposed model, several data augmentation techniques [12] were applied during the training phase. These include horizontal and vertical flipping, random rotations, and scaling transformations. Such augmentations simulate real-world variations and help the model better handle unseen conditions.

The dataset was divided into three subsets: 70% for training, 20% for validation, and 10% for testing. This split ensures a balanced distribution of defect classes across all subsets, allowing reliable performance evaluation while avoiding data leakage.

3.2. Model Architecture

The proposed system is based on the YOLOv8n-OBB architecture [6], a recent state-of-the-art model developed by Ultralytics for real-time object detection. Compared with previous YOLO versions, YOLOv8 introduces several architectural refinements aimed at improving both detection accuracy and computational efficiency.

The architecture is organized into three main components: a backbone, a neck, and a detection head. The backbone, based on CSPDarknet, is responsible for extracting multi-scale feature representations from input images. The neck utilizes a Path Aggregation Network (PANet) to fuse feature maps at different resolutions, enabling better detection of objects at various scales. The detection head is decoupled, allowing separate optimization of classification and bounding box regression tasks, which improves overall performance.

The YOLOv8-OBB model was trained using a set of hyperparameters selected to ensure stable convergence and satisfactory generalization performance. Unlike conventional horizontal bounding boxes, OBB allows the model to predict not only the position and size of defects but also their orientation. This is particularly important for elongated and arbitrarily oriented defects such as cracks and leaks.

The use of OBB provides several advantages, including improved localization accuracy, reduced background noise in detection, and more precise Intersection over Union (IoU) computation. This makes the proposed approach well-suited for pipeline inspection scenarios where defects can appear under various orientations and perspectives.

3.3. Training Configuration

The YOLOv8-OBB model was trained under a carefully selected set of hyperparameters to ensure optimal convergence and generalization performance. The training process was conducted for 100 epochs, which was sufficient for full convergence without overfitting.

A batch size of 16 was used, optimized for a GPU environment with 8 GB of VRAM. All input images were resized to a resolution of 640×640 pixels, following the standard configuration for YOLOv8 models.

The model was trained using the Stochastic Gradient Descent (SGD) optimizer, which provides better generalization compared to adaptive optimizers in object detection tasks. The initial learning rate was set to 0.01, corresponding to the default YOLOv8 configuration.

To prevent overfitting and improve training efficiency, an early stopping strategy with a patience of 20 epochs was applied. In addition, data augmentation techniques such as Mosaic augmentation, flipping, and scaling were enabled during training to enhance model robustness.

Table 1 summarizes the training hyperparameters.

Table 1. Training Hyperparameters

Hyperparameter	Value	Justification
Epochs	100	Ensures full convergence
Batch size	16	Optimized for 8GB VRAM
Image size	640×640	Standard YOLOv8 input size
Optimizer	SGD	Better generalization
Learning rate	0.01	Default YOLOv8 value
Patience	20	Early stopping to avoid overfitting
Data augmentation	Enabled	Mosaic, flip, scale

Table 2. Overall Performance Metrics

Metric	Score (%)	Raw Value	Evaluation
Precision	83.92	0.83921	Excellent
Recall	86.06	0.86058	Excellent
mAP@50	87.22	0.87224	Remarkable
mAP@50–95	60.88	0.60882	Good
Inference Time	5–10 ms/image	–	Real-time

4. Experimental Results

4.1. Overall Performance

The proposed YOLOv8-OBB model was trained for 100 epochs using a custom dataset of 1,224 annotated images. The overall performance on the test set is summarized in Table 2.

The model achieves high precision and recall, with an mAP@50 of 87.22% and real-time inference capability, making it suitable for industrial deployment scenarios.

4.2. Training Dynamics

To better understand the learning behavior of the proposed model, we analyze the evolution of key training and validation metrics over the training process. This analysis provides insights into convergence stability, learning efficiency, and generalization capability.

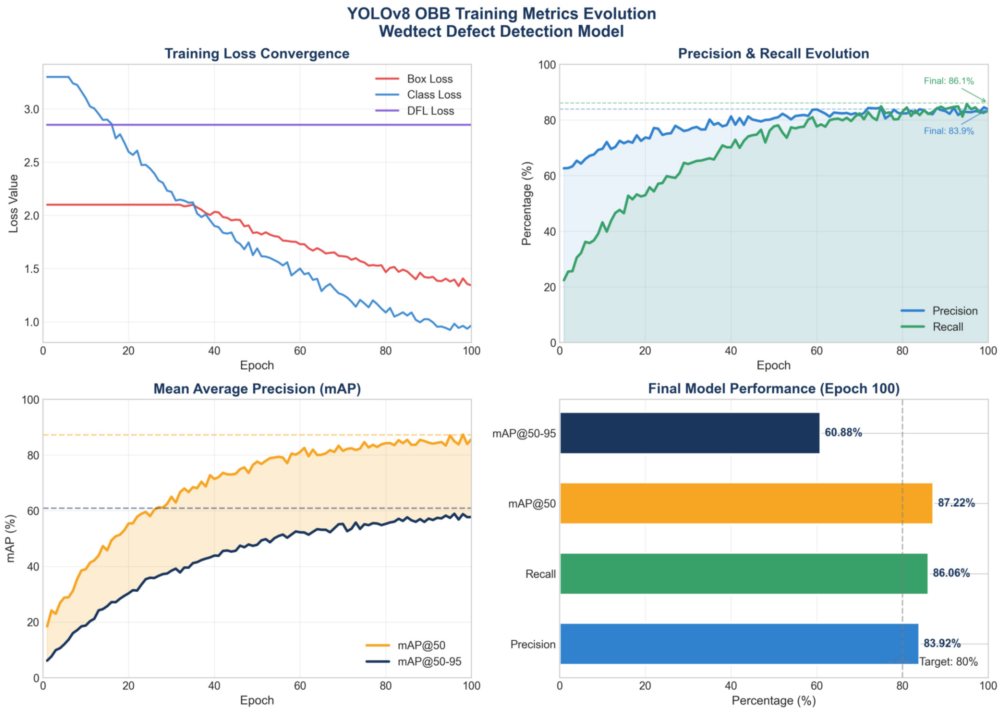

Figure 1. Training evolution over 100 epochs

Fig. 1 illustrates stable convergence, progressive loss decrease, and consistent improvement of precision and recall. Metrics stabilize after approximately 80 epochs, indicating convergence without overfitting.

4.3. Comprehensive Performance Evaluation

To provide a comprehensive evaluation of the proposed model, we analyze its performance from multiple perspectives, including class-wise discrimination ability, precision-recall trade-offs, and overall classification balance. These complementary evaluations help better understand the robustness of the YOLOv8-OBB framework in industrial pipeline inspection scenarios.

Figure 2. Normalized confusion matrix

The normalized confusion matrix presented in Fig. 2 highlights the strong classification capability of the proposed model. A clear diagonal dominance can be observed, indicating correct classification of most defect categories.

4.4. Qualitative and Visual Analysis

In addition to quantitative metrics, qualitative analysis is performed to evaluate the visual consistency and detection reliability of the proposed model.

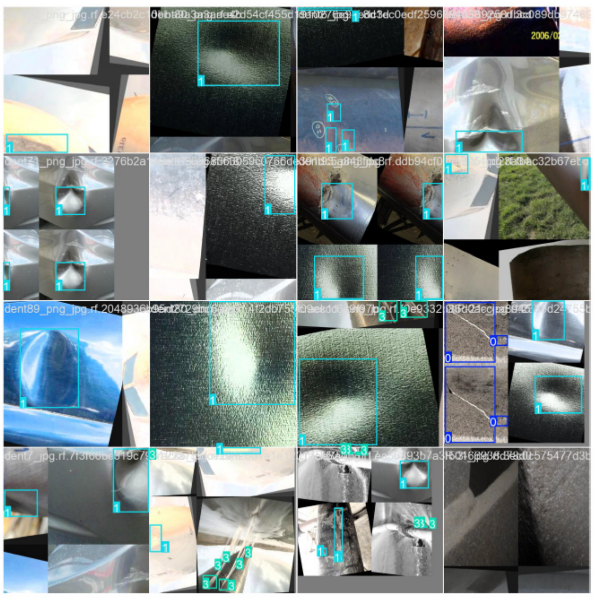

Figure 3. Training batch with OBB annotations

Fig. 3 shows an example of a training batch with oriented bounding box annotations. The annotations accurately capture defect orientation and shape.

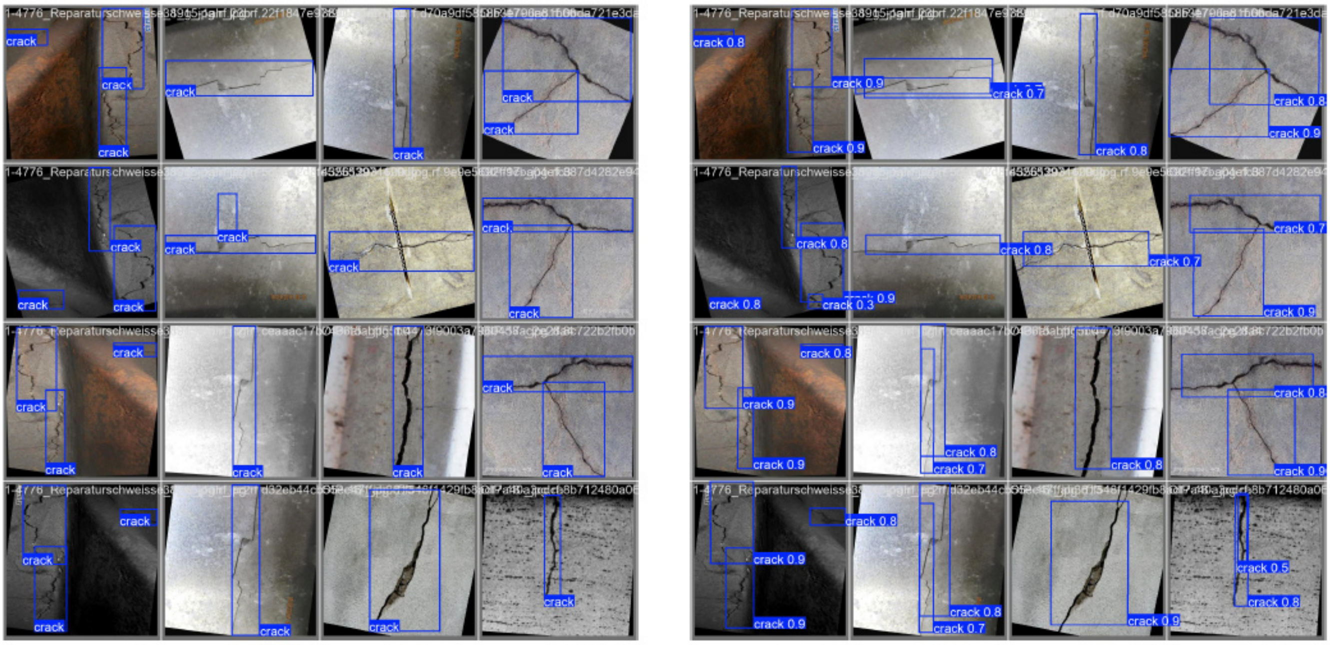

Figure 4. Ground truth vs predictions

Fig. 4 compares ground truth annotations with model predictions on the validation set. The results show strong alignment between predicted and actual defect locations.

Figure 5. Example prediction on test image

Fig. 5 presents predictions on unseen test data, demonstrating strong generalization ability. Overall, the qualitative results confirm that the model is capable of accurately localizing and classifying defects under different conditions, demonstrating strong generalization ability.

4.5. Dataset Enrichment and Impact Analysis

To further enhance model robustness, we investigate the impact of dataset enrichment through additional scraped images.

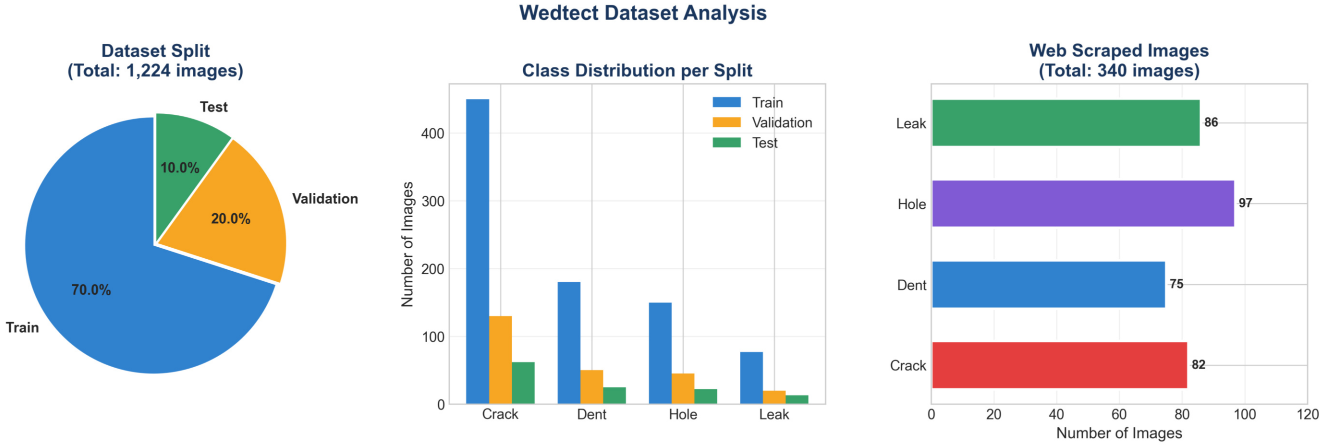

Figure 6. Dataset distribution (original vs scraped)

Fig. 6 shows the distribution of original and scraped datasets, introducing higher variability in environmental conditions and defect appearances.

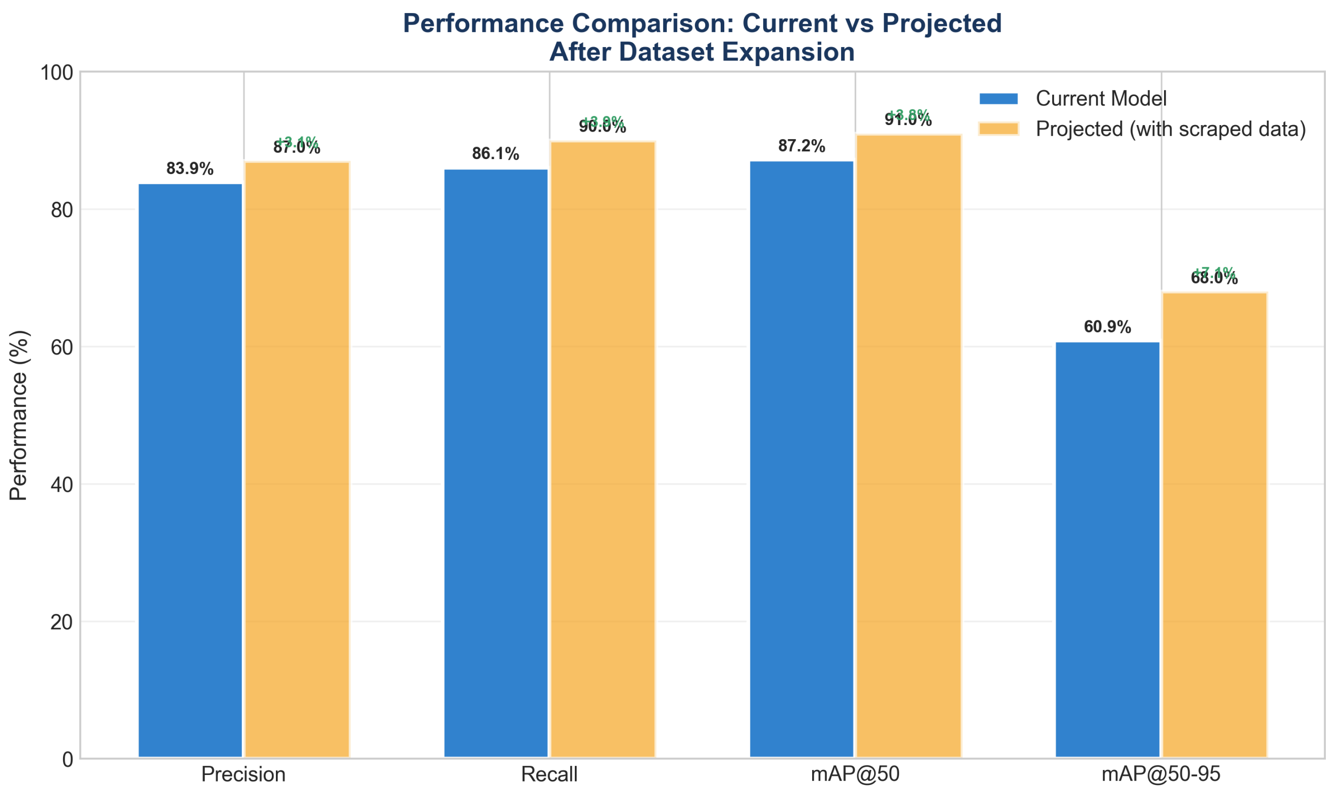

Figure 7. Performance comparison: original vs augmented dataset

Fig. 7 illustrates the impact of dataset enrichment. The augmented dataset improves generalization, robustness, and detection stability. The results clearly indicate that dataset diversity plays a critical role in improving generalization and reducing sensitivity to environmental variations.

4.6. Confidence and Training Stability Analysis

To further assess model reliability, we analyze prediction confidence distribution along with training stability behavior.

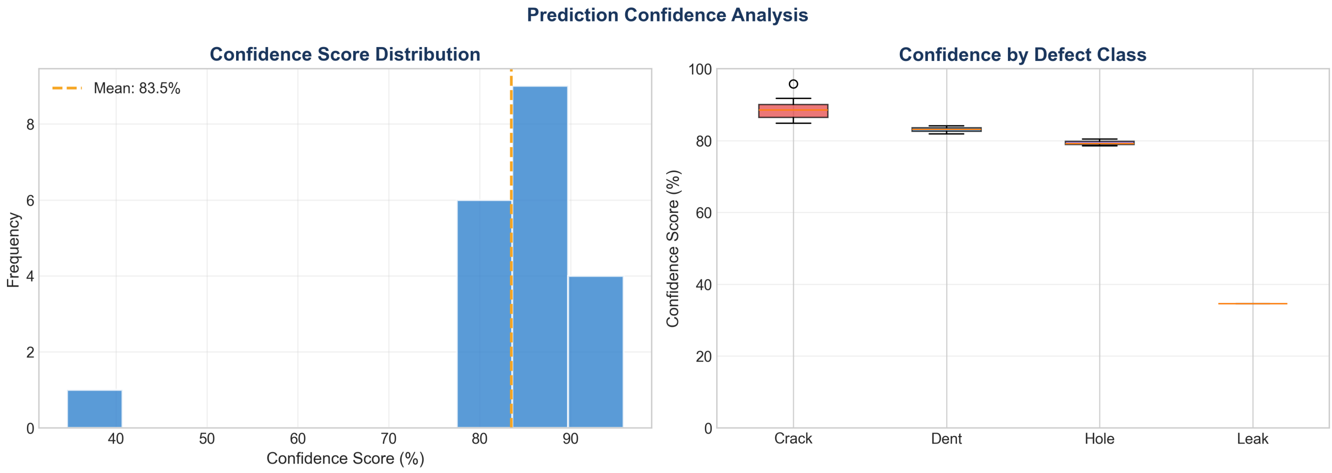

Figure 8. Confidence score distribution

Fig. 8 shows the confidence score distribution with an average of approximately 83.5%, indicating reliable predictions.

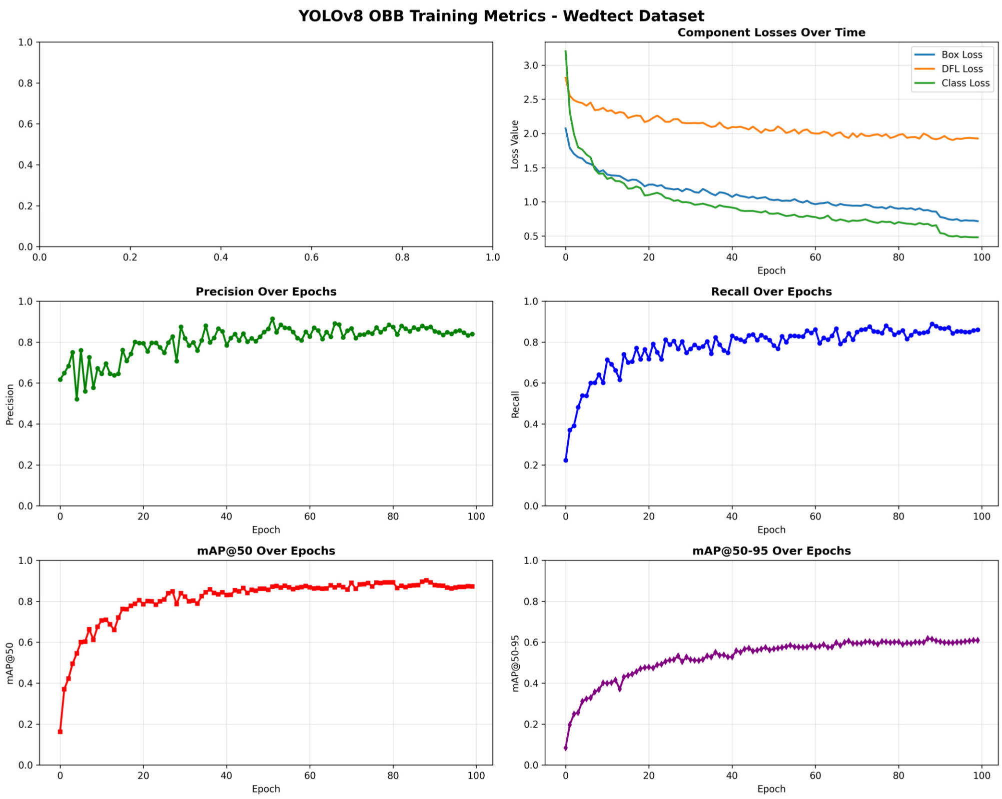

Figure 9. Detailed training and performance metrics

Fig. 9 illustrates stable convergence of loss functions and consistent improvement of performance metrics during training. The analysis confirms stable training behavior and high prediction reliability, with consistent confidence levels across defect categories.

5. Discussion

The experimental results demonstrate that the proposed YOLOv8-OBB model achieves strong and consistent performance across all evaluation metrics. In particular, the integration of oriented bounding boxes significantly improves localization accuracy for elongated and arbitrarily oriented defects such as cracks and leaks, compared to standard axis-aligned detectors.

Quantitatively, the model achieves a precision of 83.92% and a recall of 86.06%, confirming its reliability in complex industrial environments. The confusion matrix further supports these findings, showing strong classification performance with minimal inter-class confusion. Slight performance degradation is observed for the dent class, which can be attributed to its high variability and weak or ambiguous visual structure.

The precision-recall curves indicate a stable trade-off between sensitivity and specificity across different confidence thresholds. Similarly, the confidence score distribution confirms consistent prediction reliability across all defect categories.

From a learning perspective, the training process exhibits stable convergence without signs of overfitting. This stability is reinforced by the use of data augmentation and the introduction of oriented bounding boxes, which improve the model's ability to capture spatial characteristics of irregular and rotated defects.

Despite these promising results, challenges remain in cases involving low-contrast surfaces and highly textured backgrounds, where defect boundaries become less distinguishable. Future work may explore attention mechanisms and transformer-based architectures to further enhance feature representation and robustness in such difficult scenarios.

Overall, the proposed system provides an efficient, accurate, and real-time solution for industrial pipeline inspection, demonstrating strong potential for real-world deployment.

6. Conclusion

This work presented a YOLOv8-OBB-based framework for industrial pipeline defect detection. By integrating oriented bounding boxes into the detection pipeline, the proposed approach significantly improves the localization of elongated and rotated defects while maintaining real-time inference capability.

Extensive experiments conducted on a custom dataset demonstrate strong performance in terms of precision, recall, and mAP across multiple defect categories. The results also confirm stable training behavior and good generalization ability under diverse environmental conditions.

Furthermore, dataset enrichment through additional web-scraped images contributed to improving model robustness and detection consistency. Although the proposed framework achieves encouraging results, there is still room for improvement.

Future work will focus on expanding the dataset, exploring domain adaptation techniques, and deploying the model on embedded edge devices to enable real-time industrial monitoring in operational environments.

REFERENCES

[1] Bochkovskiy, A., Wang, C.Y., Liao, H.Y.M.: Yolov4: Optimal speed and accuracy of object detection. arXiv preprint arXiv:2004.10934 (2020)

[2] Cheng, J., et al.: Automatic sewer defect detection based on deep learning. Automation in Construction 123, 103533 (2021)

[3] Ding, J., et al.: Learning roi transformer for oriented object detection in aerial images. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 2849–2858 (2019)

[4] Girshick, R.: Fast r-cnn. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV). pp. 1440–1448 (2015)

[5] He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask r-cnn. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV). pp. 2961–2969 (2017)

[6] Jocher, G., et al.: Yolov8 by ultralytics. GitHub Repository (2023)

[7] Liu, W., et al.: Surface defect detection using yolo-based deep learning models: A review. Sensors 23(5), 2345 (2023)

[8] Liu, Z., et al.: Swin transformer: Hierarchical vision transformer using shifted windows. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 10012–10022 (2021)

[9] Redmon, J., Divvala, S., Girshick, R., Farhadi, A.: You only look once: Unified, real-time object detection (2016)

[10] Redmon, J., Farhadi, A.: Yolov3: An incremental improvement. arXiv preprint arXiv:1804.02767 (2018)

[11] Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection. In: Advances in Neural Information Processing Systems (NeurIPS). pp. 91–99 (2015)

[12] Shorten, C., Khoshgoftaar, T.: A survey on image data augmentation. Journal of Big Data (2019)

[13] Wang, C.Y., Bochkovskiy, A., Liao, H.Y.M.: Yolov7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors. arXiv preprint arXiv:2207.02696 (2022)

[14] Wang, T., et al.: A sewer pipeline defect detection method based on improved yolov5. Processes 11(8), 2508 (2023)

[15] Xia, G.S., et al.: Oriented r-cnn for object detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2021)

[16] Yin, X., et al.: Deep learning for sewer defect detection. Automation in Construction 115, 103146 (2020)

[17] Zhou, X., et al.: Oriented object detection in aerial images with box boundary-aware vectors. Pattern Recognition Letters 146, 175–181 (2021)