Automated Welding Defect Detection Using YOLO: A Deep Learning Approach for Industrial Quality Assurance

G. Srujan Hrudhay *[1], CH. Bhavani Prasad[1], G. Vinay[1], and Mrs. M. Harika Chowdary[2]

[1]Student
[2]Assistant Professor
Department of Mechanical Engineering,
Godavari Institute of Engineering and Technology (A), Rajahmundry.

Abstract

Welding is a critical manufacturing process in which defects such as cracks, porosity, spatter, and incomplete fusion can significantly compromise structural integrity and product reliability. Conventional inspection methods are labor-intensive, subjective, and difficult to scale for high-throughput industrial environments. This paper presents an automated welding defect detection system based on a lightweight YOLO11s object detection model designed for real-time industrial quality assurance. The proposed system is trained and validated on a limited dataset of labeled weld images, reflecting realistic industrial data constraints. Experimental results demonstrate that the model achieves a mean Average Precision (mAP@0.5) of 46.23% with moderate recall, indicating its suitability as an automated screening and inspection assistance tool. The study highlights the feasibility of deploying lightweight deep learning models for visual inspection tasks under limited data availability and computational constraints.

Keywords: Welding defect detection, YOLO, deep learning, industrial inspection, computer vision, quality assurance

1 Introduction

Welding plays a vital role in industries such as automotive manufacturing, shipbuilding, aerospace, and infrastructure development. However, welding processes are inherently prone to defects arising from improper heat input, contamination, material inconsistencies, and operator variability. Common defects—including blowholes, cracks, incomplete

*srujanhrudhay@gmail.com

fusion, and excessive spatter—can severely reduce the mechanical strength and service life of welded joints.

Traditional welding inspection techniques, such as visual inspection, radiography, and ultrasonic testing, require skilled personnel and often involve high operational costs. Moreover, manual visual inspection is subjective and difficult to scale for continuous production lines. As a result, automated vision-based inspection systems have gained increasing attention in industrial quality assurance.

Recent advancements in deep learning, particularly convolutional neural networks (CNNs), have significantly improved object detection and classification performance. Among these, You Only Look Once (YOLO) models have emerged as a preferred choice for real-time detection due to their single-stage architecture and low inference latency. This work explores the application of a lightweight YOLO-based model for automated welding defect detection, focusing on practicality, computational efficiency, and deployment feasibility rather than large-scale benchmark optimization.

2 Related Work

Several studies have investigated automated welding defect detection using machine learning and deep learning techniques. Early approaches relied on handcrafted features combined with classical classifiers such as support vector machines (SVMs) and random forests. While effective under controlled conditions, these methods often lacked robustness to variations in lighting, surface texture, and defect morphology.

More recent works have employed deep CNN architectures for defect classification and segmentation. Two-stage detectors such as Faster R-CNN have demonstrated high detection accuracy but suffer from increased computational complexity and inference latency, limiting their suitability for real-time industrial deployment. Single-stage detectors, including SSD and YOLO variants, have shown improved speed–accuracy trade-offs, making them attractive for on-line inspection systems.

Although several studies report high accuracy using large annotated datasets, such datasets are rarely available in real industrial environments due to labeling cost, confidentiality, and process variability. Consequently, evaluating lightweight detection models under limited-data conditions remains an important and practical research direction.

3 System Architecture

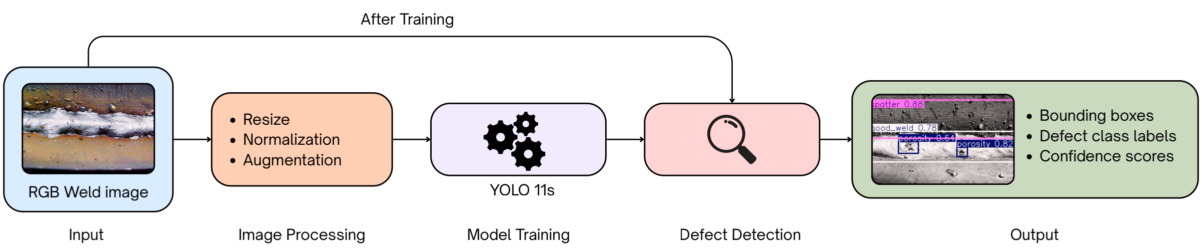

Figure 1: Block diagram of the proposed welding defect detection system.

The overall architecture of the proposed system consists of four main stages:

1. Input RGB weld image acquisition,
2. Image preprocessing,
3. YOLO-based defect detection,
4. Output visualization with detected defect classes and bounding boxes.

Captured weld images are preprocessed to ensure consistent input resolution and normalization before being passed to the detection network. The trained model outputs bounding boxes, class labels, and confidence scores for identified welding defects.

The single-stage detection framework enables efficient inference while maintaining acceptable localization performance for industrial inspection tasks.

4 Dataset and Experimental Setup

4.1 Dataset Description

The dataset used in this study consists of 299 labeled RGB weld images collected from industrial and experimental welding scenarios. Each image was manually annotated with bounding boxes corresponding to visible welding defects.

The dataset was divided as follows:

- Training set: 241 images (80.6%)
- Validation set: 58 images (19.4%)
- Testing: Unseen images used for qualitative evaluation

The YOLO11s model was trained using an 80:20 train–validation split, while unseen images were used to assess post-training generalization. Although limited in size, this dataset reflects realistic industrial conditions where collecting and annotating large volumes of defect data is often impractical.

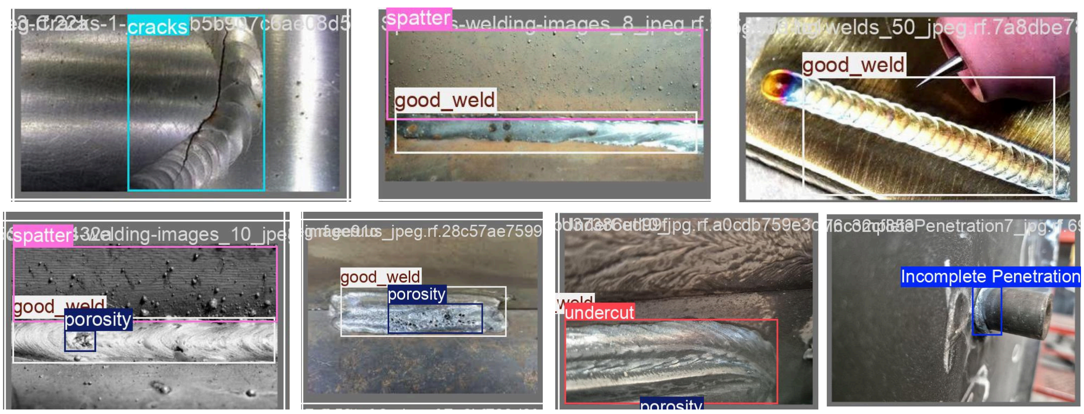

Figure 2: Sample weld images with annotated defect labels.

4.2 Preprocessing and Training Configuration

All images were resized to match the input resolution required by the YOLO architecture. Standard data augmentation techniques, including random scaling and horizontal flipping, were applied during training to improve robustness. The model was trained for 50 epochs using stochastic gradient descent with default YOLO hyperparameters, and validation performance was monitored to ensure stable convergence and prevent severe overfitting.

5 Detection Model and Design Considerations

The proposed system employs YOLO11s, a lightweight variant of the YOLO architecture optimized for fast inference and reduced computational overhead. YOLO performs object detection as a single regression problem, allowing simultaneous localization and classification of defects.

The selection of YOLO11s was guided by practical deployment considerations. Lightweight architectures reduce the risk of overfitting when training data is limited and are better suited for real-time industrial environments with constrained computational resources. Evaluating such models under small-data conditions also provides insights applicable to other industrial inspection tasks, including surface defect detection and manufacturing quality monitoring.

6 Results and Discussion

6.1 Overall Detection Performance

Table 1: Overall performance metrics of the proposed YOLO11s-based system.

Metric	Value
Precision	42.98%
Recall	54.68%
mAP@0.5	46.23%
mAP@0.5–0.95	20.96%

The achieved mAP@0.5 of 46.23% demonstrates that the model can localize and identify welding defects with reasonable accuracy despite the limited dataset size. The lower mAP@0.5–0.95 reflects sensitivity to stricter localization thresholds, which is expected in fine-grained defect detection tasks where defect boundaries are visually ambiguous.

6.2 Training and Validation Behavior

The training curves show stable convergence without abrupt divergence between training and validation losses. While validation losses remain higher than training losses, this behavior is consistent with the limited dataset size and class imbalance, indicating controlled generalization rather than overfitting.

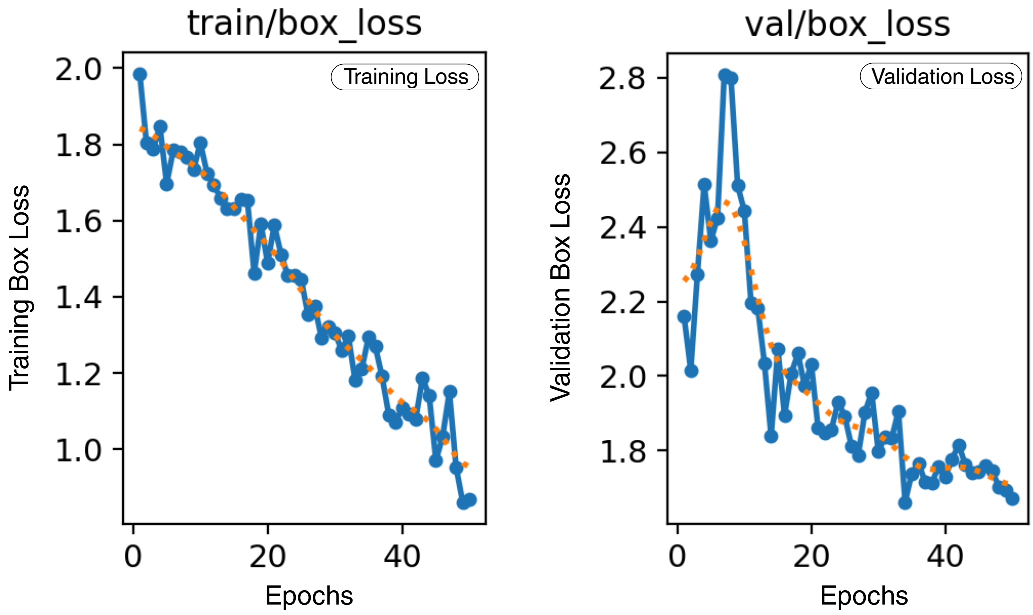

Figure 3: Training and validation box loss curves across epochs.

6.3 Per-Class Performance Analysis

Table 2: Per-class evaluation metrics for welding defect detection.

Class	GT	Pred	Precision	Recall	F1-score	AP@0.5
All	58	96	0.42	0.551	0.457	0.206
Incomplete Penetration	1	1	0.509	1.000	0.995	0.199
Cracks	7	8	0.217	0.375	0.283	0.156
Good Weld	46	59	0.692	0.661	0.684	0.408
Incomplete Fusion	1	1	0.464	1.000	0.497	0.151
Porosity	8	9	0.233	0.222	0.184	0.100
Spatter	9	14	0.353	0.351	0.271	0.163
Undercut	4	4	0.473	0.250	0.282	0.266

Performance variations across defect classes are primarily influenced by class imbalance and limited ground truth samples. The Good Weld category achieves the highest AP@0.5, indicating effective discrimination between defect-free and defective welds. Lower performance for visually subtle defects such as porosity and cracks highlights the inherent difficulty of fine-grained welding defect detection.

6.4 Qualitative Evaluation

Qualitative results demonstrate that the model can meaningfully localize defect regions and classify multiple defect types under varying surface conditions. These results support the quantitative findings and indicate suitability for inspection assistance applications.

Figure 4: Qualitative detection results on unseen weld images.

6.5 Discussion

The results indicate that lightweight object detection models such as YOLO11s can provide meaningful inspection performance even under limited-data conditions. While the achieved metrics do not match large-scale benchmark studies, they are realistic and representative of industrial environments. The moderate recall suggests suitability for automated screening and operator assistance rather than fully autonomous inspection.

7 Conclusion

This paper presented an automated welding defect detection system based on a YOLO11s deep learning framework. The approach achieves real-time performance with stable accuracy under constrained data conditions, demonstrating practical applicability in industrial inspection scenarios.

8 Reproducibility

To support reproducibility and further research, the implementation, configuration files, and trained model weights are made publicly available at :
`https://github.com/SrujanHrudhay/automatic-weld-defect-detection-using-yolo11s`

References

[1] Glenn Jocher, Ayush Chaurasia, and Jing Qiu. Ultralytics yolov8. *GitHub Repository*, 2023. https://github.com/ultralytics/ultralytics.

[2] NVIDIA Corporation. Cuda toolkit documentation, 2023. https://developer.nvidia.com/cuda-toolkit.

[3] Adam Paszke et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.

[4] Python Software Foundation. Python programming language, 2023. https://www.python.org.

[5] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 779–788, 2016.

[6] Roboflow Inc. Roboflow: Dataset management and annotation platform, 2023. https://roboflow.com.