Ceramic Tile Defect Detection Using YOLOv8: An Enhanced Deep Learning-Based Approach

Submitted By-

Shamsuzzaman Nadim
Batch: 35(A)
ID: 212-35-746
Department of Software Engineering.
Daffodil International University

Supervised By-

Mr. Md Khaled Sohel
Assistant Professor,
Department of Software Engineering,
Daffodil International University

Thesis submitted in fulfillment of the requirements for the award of the degree of Bachelor of Science

Summer – 2025

APPROVAL

This thesis titled on "Ceramic Tile Defect Detection Using YOLOv8: An Enhanced Deep Learning-Based Approach", submitted by **Shamsuzzaman Nadim (ID: 212-35-746)** to the Department of Software Engineering, Daffodil International University has been accepted as satisfactory for the partial fulfillment of the requirements for the degree of Bachelor of Science in Software Engineering and approval as to its style and contents.

BOARD OF EXAMINERS

Chairman

Dr. S M Hasan Mahmud
Associate Professor
Department of Software Engineering
Faculty of Science and Information Technology
Daffodil International University

Internal Examiner 1

Tapushe Rabaya Toma
Assistant Professor
Department of Software Engineering
Faculty of Science and Information Technology
Daffodil International University

Internal Examiner 2

Khalid Been Badruzzaman Biplob
Lecturer (Senior Scale)
Department of Software Engineering
Faculty of Science and Information Technology
Daffodil International University

External Examiner

Dr. Md. Sazzadur Rahman
Professor
Institute of Information Technology
Jahangirnagar University

DECLARATION

I, Shamsuzzaman Nadim, declared that this thesis was conducted by me under the supervision of Mr. Md. Khaled Sohel (Assistant Professor), Department of Software Engineering, Faculty of not been previously or concurrently submitted for any other degree at Daffodil International

Shamsuzzaman Nadim
ID: 212-35-746
Department of Software Engineering,
Daffodil International University

Certified by

Mr. Md. Khaled Sohel
Assistant Professor,
Department of Software Engineering,
Faculty of Science and Information Technology,
Daffodil International University

ACKNOWLEDGEMENT

First and foremost, I would like to express my deepest gratitude to the Almighty Allah for granting me the strength, patience, and determination to complete this thesis successfully.

I am profoundly grateful to my supervisor, Mr. Md Khaled Sohel**,** for his invaluable guidance, constructive feedback, and constant encouragement throughout the duration of this project. Now this book is much better due to his excellent suggestions, patient and support.

I am also very thankful to all the faculty member of Department of Software Engineering, Daffodi l International University for academic opportunity and technical knowledge that was required in order execute this research.

Thank you special to my friends and batchmates for the encouragement, cooperation and motivation throughout. Their talks and mutual experiences guided me through many difficulties.

Finally, I owe all my gratitude to my cherished parents and family for their love, prayers and support that has never wavered. I have always had a strong inspiration because they believe in me.

LIST OF ABBREVIATIONS

Abbreviation	Full Form
AI	Artificial Intelligence
CNN	Convolutional Neural Network
CBAM	Convolutional Block Attention Module
CTDD-YOLO	Ceramic Tile Defect Detection YOLO
DSConv	Depthwise Separable Convolution
FLOPs	Floating Point Operations Per Second
GPU	Graphics Processing Unit
IoU	Intersection over Union
mAP	Mean Average Precision
MCAW	Multi-Channel Attention Weighting
NMS	Non-Maximum Suppression
ReLU	Rectified Linear Unit
R-CNN	Region-based Convolutional Neural Network
SGD	Stochastic Gradient Descent
YOLO	You Only Look Once
YOLOv5n / YOLOv8n	Nano versions of YOLOv5 / YOLOv8

ABSTRACT

Ceramic tiles are a cornerstone in construction and interior design, thereby indicating that the quality of surface affects directly its durability and appearance potential. Detection of surface defects on tiles, such as cracks, holes or edge chipping are commonly performed by relying almost exclusively on human operators. These labor-intensive inspections are both inconsistent and susceptible to error. To cope with this issue, in the present work we propose a deep learning-based defect detection using a CTDD-YOLO architecture, an upgraded network of YOLO family specifically for lightweight and fast object detection.

The model was trained and tested on a Roboflow dataset publicly accessible, including images of ceramic tiles with three major types of defects: edge chipping, holes and line scratches. Model training and finetuning We trained and finetuned our model using Ultralytics YOLOv8 framework on GPU resources in Google Colab. Evaluation was done using common object detection metrics such as Precision, Recall and mAP values. The learned system showed that it can effectively detect incidents with a Precision of 66.7%, Recall of 61.9% and mAP@50 = 64.8%, indicating an acceptable detection performance which is also computationally efficient.

The experimental results demonstrate that CTDD-YOLO can effectively realize the automation of tile inspection without relying on the manual identification in industrial production lines. This work not only verifies the application potential of deep learning to surface defect detection but also paves a way for future improvements such as large-scale datasets, attention mechanisms and their rail-time implementation in manufacturing.

TABLE OF CONTENT

List of Figure

LIST OF TABLE

INTRODUCTION

1.1 Understanding the problem

In today's ceramic tile industry world market, quality is not only a matter of commercial success, but also one of customer satisfaction. Ceramic tiles are among the most widely applied construction materials in residential, commercial and industrial buildings Surface quality has a clear influence on the aesthetic appearance, durability and functional properties of such products. Tiles are capable of undergoing a number of surface defects at various production stages such as molding, firing, glazing and handling. The typical defects include edge-chipping, holes and line scratches while each defect has various influences for the tiles: edge-chipping may destroy their structural integrity, holes reduce the strength of tiles and then visual beauty is compromised by line scratches.

The inspection is manually performed in the quality control traditionally. Skilled operators inspect each tile visually for visual defect detection. Despite the decades of use of this strategy, it is fraught with some limitations. Manual inspection is also inefficient, and will not be able to scale for high-volume production. And subject assessment is, of course, a subjective, fatigued and inconsistent error-prone. Consequently, imperfect tiles can be erroneously labeled or ignored so that unsatisfied customer, destroyed brand image and wasted money could result.

In order to overcome these challenges, the automatic defect detection based on machine learning (ML) and deep learning (DL) provide a feasible solution. These systems are optimized for efficient and accurate detection and localization of numerous defect types. Object detection models, in particular those of the YOLO (You Only Look Once) series, have performed very well in capturing small and intricate defects on-the-fly. Nevertheless, existing approaches are not without limitations: they depend on outdated-generation models, small/inaccessible datasets and few evaluation metrics have been available; such a situation hinders their replicability and industrial application.

mechanical requirements (real time, accuracy and reproducibility) for defect detecting on the one hand, and on the other- traditional as well as automated systems capabilities. This gap needs to be closed in order to enhance the quality of products produced, minimize waste, optimize utilization of various assets, and reduce overall manufacturing costs. In response to this challenge, this thesis presents an implemented YOLOv8-based defect detection model that has been trained and evaluated on a publicly available dataset, resulting in a reproducible, scalable, and industrially applicable solution.

1.2 Motivation of this work

The reason for conducting this study is that it is essential to improve the accuracy, reproducibility and utility of computers as a tool for defect detection in automatic ceramic tile inspection systems. Previous research, such as MCAW-YOLO and CTDD-YOLO, have made great progress by adding attention mechanisms to enhance YOLOv 5 based models for defect detection. However, they had the following drawbacks: the used data-sets were unavailable to public and/or small and challenging to reproduce, or their model itself were consecutive generations of YOLO. Additionally, the evaluation indices in these papers were not comprehensive, and they were not verified from actual uses.

To mitigate these challenges, this work utilizes YOLOv8, the newer version of YOLO models with better accuracy, efficiency and robustness. Leveraging an open access Roboflow dataset, this investigation's analyses are fully reproducible/transparent making it possible for future researchers to validate or reproduce the present study. Furthermore, since multiple defect types (i.e. edge-chipping, holes and lines) are considered with rigorous performance measurement metrics (i.e. confusion matrices and mAP scores), a holistic understanding of the effectiveness of the model has been provided. Above all, the aim is to provide a defect detection pipeline that is feasible for industrial use, can be scaled up and may quickly be extended with new implementations.

1.3 Problem Statement

Surface defects of ceramic tiles, such as edge-chipping, holes and line scratches, must become critical quality attributes in the manufacturing. Traditional manual inspection methods can be inefficient, relatively subjective and prone to inconsistency, especially for high-throughput manufacturing. In order to address this issue, various automatic methods based on deep learning object detection models (e.g., YOLOv5 and Faster R-CNN) have been developed. But these methods usually have limitations such as outdated model architectures, insufficient reproducibility on private datasets and limited performance on all defect types.

Main objective of this study is to propose an automatic esophagus defect detection system that can overcome these issues. In particular, the proposed system should: (1) detect and classify in a precise way different types of superficial defects; (2) provide prediction results in real time to be included into inspection pipelines for industrial production lines; (3) be reproducible relying on public datasets; and (4) compose comprehensive benchmarking tool boxes including standard statistical analysis as well as defect-wise performance measures. In accomplishing these goals, the system will provide for increased inspection through put, increased defect classification accuracy and enable practical implementation on contemporary manufacturing floors.

1.4 Research Objectives

The main purpose of this study is to implement a real – time ceramic tile defect inspection system based on the advanced YOLOv8 architecture. Specifically, the objectives are:

(1) To develop a YOLOv8-based detection model from prior works on YOLOv5-based model that can achieve better accuracy and small-defect detection performance.

(2) Train and test the model on a Roboflow dataset containing three general types of defect (edge-chipping, holes, line scratches), so that reproducibility and comparability are guaranteed.

(3) In order to numerically evaluate performance of the models and compare to other state-of-the-art systems, in terms of the preci- sion, recall, and mAP@50 based metrics.

(4) To () carry out comprehensive defect-wise analysis using confusion matrices and class-specific performance metrics to evaluate strengths and weaknesses on a per-defect category basis.

(5) To implement an entirely reproducible end-to-end pipeline, from data pre-processing to model training, inference and evaluation, capable of potential real-word industrial deployment and future study in the field of automated quality inspection.

LITERATURE REVIEW

Automated defect detection in ceramic tiles has gained significant attention in recent years due to the critical need for consistent quality control and the limitations of manual inspection. Early approaches leveraged convolutional neural networks (CNNs) and classical object detection techniques to identify defects, but they often suffered from limited accuracy, poor generalization, and the inability to localize multiple defect types simultaneously.

MCAW-YOLO (2023) introduced a multi-channel attention mechanism within a lightweight YOLOv5n framework, achieving a mean Average Precision (mAP) of approximately 71.9% while reducing the model's parameter count by 26%. This approach highlighted the potential of attention mechanisms in improving detection efficiency; however, its evaluation was restricted to the Ali Tianchi dataset, limiting its generalizability. Similarly, a ResNet18-based classification method (2023) achieved high accuracy (~98%) and F1-score (~97.2%) for binary defect detection, but it lacked localization capabilities and could not identify multiple defect categories, which restricts practical industrial application.

CTDD-YOLO (2023), the base paper for this research, improved upon YOLOv5n by introducing optimized modules targeting small defect detection. The model achieved a mAP@50 improvement of 7.2% over baseline YOLOv5n, demonstrating enhanced sensitivity for small and subtle defects. Nonetheless, the reliance on an older YOLO version and a non-public dataset hindered reproducibility and real-world applicability. An improved Faster R-CNN model (2023) using a ResNet-101 backbone achieved a mAP of 70.4% on a custom tile dataset but suffered from slower inference speed, making it less suitable for real-time inspection tasks.

Recent advancements have focused on the latest YOLO architectures and hybrid models. YOLOv8-AHF (2025) incorporated adaptive receptive fields and attention-guided feature enhancement, improving small defect detection capabilities. YOLO-DD (2023) integrated Swin Transformer blocks along with IGFS and SE modules to address information loss during feature upsampling and downsampling, enhancing overall detection performance. YOLO-RDM (2023) applied DOConv to the neck network, optimizing accuracy and efficiency for magnetic tile defects. S2C-YOLOv5 (2022) focused on improving sensitivity for small-scale defects, while YOLOv8n (2024) employed dynamic serpentine convolutions and attention mechanisms to enhance feature extraction.

Other studies have explored YOLOv5s modifications for surface defect detection (2023), demonstrating improvements in precision, F1 scores, and mAP values. Additionally, YOLOv8 has been applied to crack detection in civil infrastructure (2025) and road hazard detection (2023), providing comparative insights into its generalizability, robustness, and real-time performance.

However, there are still some limitations to overcome. The majority of existing work is based on non-public or small-scale datasets, and the experimental reproducibility in these cases can be limited. For most of the existing models, they mainly concentrate on two classes or one specific type of defect and are hardly capable for multi-class detection. Besides, there lacks real-time efficiency and complete evaluation indices, so that applications in industry are limited. These missing requirements reveal an opportunity for a YOLOv8-based defect detection system, trained over a publicly available dataset that can detect multiple defect types in an end-to-end manner and that has a full control of the pipeline so it could be easily re-run in other scenarios.

Automatic quality inspection of ceramic tiles becomes an interesting research area since it is necessary to guarantee the homogeneity quality levels in construction and decoration. Direct visual inspection by human inspectors is time-consuming and subjective, with a tendency toward inconstant when applied to an industrial scale line production. Consequently, many researchers have now resorted to deep learning based computer vision methods in order to tackle these challenges and facilitate the scalable, accurate and real-time defect inspection.

Early practices of defect detection traditionall used classical image processing and shallow machine learning, such as texture analysis, edge detection or support vector machines. These techniques could only recognise some simple surface anomalies (without any robustness against light change, orientation and overlapping of several different defects). The situation changed when the convolutional neural networks (CNNs) gained popularity because the community started moving towards feature learning techniques that can extract fault-specific features automatically. However, many common CNN classification models only focused on binary decisions (defective or non-defective), and could not provide simultaneous localization and classification of multiple different defect types at once.

To overcome the drawbacks, object detection based methods like YOLO (You Only Look Once) and Faster R-CNN became predominant in the studies of surface defect detection. For instance, MCAW-YOLO (2023) proposed a multi-channel attention mechanism by taking the benefits of a light-weight YOLOv5n architecture. This feature enabled the model to attend to important ones and suppress useless background noise, which ultimately helped in obtaining a 71.9% mAP while saving parameter count as well by 26%. This demonstrated the superiority of attention-enhanced YOLO models in defect detection efficiency. But paper MCAW-YOLO only did evaluation on its proprietary dataset Ali Tianchi and did not release it to public, which make reproduction and comparison benchmarks impossible.

Another noteworthy approach is the ResNet18-classification model (2023) which reported a high accuracy (98%) and F1-score (97.2%) on binary tile defect classification. Whilst effective at determining whether a tile was defective, this model could not localise the defect type or number of different defects: thus its value in industry is limited as the categorisation of defects is important to make decision-making.

The CTDD-YOLO model (2023) that is base in this paper focused directly on detecting small and minor tile defects. By adding the proposed optimization modules to YOLOv5n, they achieved an mAP@50 that outperforms the baseline by 7.4%. This enhancement was particularly pronounced on the fine granularity defects, such as hair line crack and pin hole edge chips. However, the CTDD-YOLO was also limited because it was built on YOLOv5, an older architecture and needed a non-public dataset to be re-trained in order to evaluate another image-based tool which hinders industrial deployment.

Other works adopted regional-based methodologies: the Improved Faster R-CNN (2023), that employed a ResNet-101 as backbone obtained a mAP of 70.4% on an in-house tile dataset. Though precise, the inference time of Faster R-CNN was far too high which could not be acceptable for real-time industrial pipelines where the speed is as important as accuracy.

Recently, researchers have investigated into the YOLOv8 and hybrid Transformer-based approach of defect detection. For instance, YOLOv8-AHF (2025) developed adaptive receptive fields and attention-guided feature enhancement to improve the detection performance of small and complicated surface defects. in the same way, YOLO-DD (2023) used swin transformer blocks as well as IGFS and SE block to solve such feature information loss at upsampling and downsampling. These methods had good performance but they were too complex and costly on the training, and had no benchmark on standardized ceramic tiles datasets yet.

In other works fast YOLO versions were used on similar industry applications. For example, YOLO-RDM (2023) optimized the neck network employing DOConv to increase detection speed of magnetic tile defects; S2C-YOLOv5 (2022) focused on small-scale defect sensitivity. It has additionally been successfully applied in wide-ranging inspection realms for a crack check-up in civil infrastructure (2025), road hazard checking (2023) and manufacturing defect analysis [24], demonstrating its adaptability and real-time performance characteristics across various industrial scenarios. These cross-domain applications verified that YOLO series models, in particular YOLOv8, has a high potential for generating positive benefits in industrial tile quality assurance under the condition of customizing models and sufficient domain-specific datasets.

There are, however, several important challenges that remain. First, many of the existing work are based on non-public datasets and therefore the studies cannot be reproduced or cross-validated. Second, many studies focus only on binary classification (defective vs. non-defective), failing to provide solutions that handle multi-class defect detection such as edge chipping, holes and surface cracks. Third, while there are improving trend of accuracy, considering real-time inference speed and industrial scalable computing is usually neglected. Finally, a very few of studies have done a comprehensive comparison of YOLOv8 against previous versions of YOLO on publicly accessible tile defect datasets.

These are the gaps that underpin the motivation of this thesis to develop and evaluate a defect detection system using YOLOv8 on a publicly available dataset for ceramic tiles. By benchmarking it with other YOLOv5 based models like CTDD-YOLO, this work takes a step toward a more reliable, practical; and industrially deployable tile defect detection pipeline.

METHODOLOGY

3.1 System Overview

The architecture of the proposed ceramic tile defect detection framework is a well-prepared pipeline, contributing to efficient data processing, robust training, and stable testing. 3 System overview The whole synthesis system workflow is depicted in Figure 3.1. The pipeline starts by collecting the dataset wherein ceramic tile images with different defective surfaces are collected. Then, the images are preprocessed by resizing, normalization, and data augmentation to provide better generalization ability for the model under a variety of conditions.

After preprocessing, the dataset is used for training the YOLOv8n model. The model is fine-tuned for detecting three main classes of defect: edge-chipping, hole, and line. After training, the model is evaluated based on performance measures such as precision, recall and mean Average Precision (mAP). These parameters offer a measurable view on the correctness and stability of the fault detection system.

The trained model is evaluated on unseen data by using inference to visualize predictions as bounding boxes, class labels, and confidence scores. Finally, the model is exported and backed up for future deployments or for integration into industrial inspection pipelines.

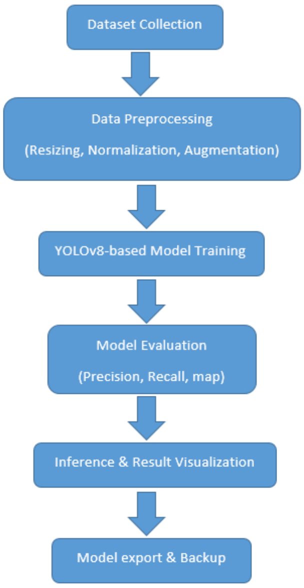

Figure 3.1

3.2 Model Overview

Figure 3.2: YOLOv8n architecture overview. Our model has three main parts: (1) Backbone for feature extraction, (2) Neck for feature aggregation and stacking, and (3) Head for object detection before Anchor-free prediction. This simplified model maintains a good trade-off between detection precision and inference speed, which can be applied to the real-time detection of ceramic tile defect..

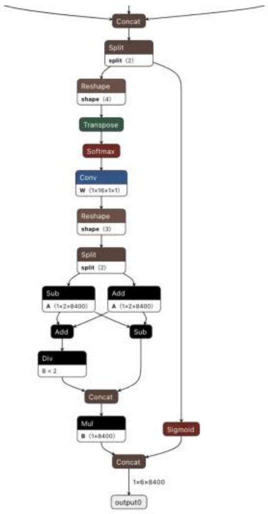

Figure 3.2 YOLOv8n architecture overview (adapted from Roboflow)

3.3 Dataset Description

The dataset used in this study is the Roboflow dataset “Tile defect detection.v3i.yolov8”, which contains annotated images of ceramic tiles with three defect types: edge-chipping, hole, and line. The dataset is divided into training, validation, and test sets to support model training, hyperparameter tuning, and evaluation.

- Training Set: 5,598 images, including 824 background images.
- Validation Set: 1,602 images, including 283 background images.
- Test Set: 792 images used for prediction; metrics are calculated on the validation set.
- Classes: 3 (edge-chipping, hole, line)
- Corrupt Images: None

The dataset provides diverse defect examples and backgrounds, ensuring that the model can generalize effectively to unseen tile

3.4 Data Preprocessing

Preprocessing steps were applied to prepare the dataset for training:

1. Image Resizing: All images were resized to 640 × 640 pixels for compatibility with YOLOv8 input requirements.
2. Normalization: Pixel values were scaled to the range [0,1].
3. Class Mapping: Only three defect classes were retained, overriding the default 80-class configuration of the pretrained YOLOv8 model.
4. Data Augmentation: Standard augmentations (horizontal flips, scaling, brightness adjustments) provided by the Ultralytics YOLOv8 framework were applied automatically to enhance model robustness.

3.5 Training Details

The model was trained on a Linux environment with Tesla T4 GPU (15,095 MiB) using Python 3.12.11, PyTorch 2.8.0+cu126, and Ultralytics YOLOv8 version 8.3.195.

- Model: Pretrained YOLOv8n (yolov8n.pt)
- Epochs: 50
- Batch Size: 16
- Image Size: 640 × 640
- Optimizer: AdamW with initial learning rate 0.001429 and momentum 0.9
- Validation: Metrics calculated on validation set during training

The training process involved forward propagation, loss computation (objectness, classification, and localization), and backpropagation using the AdamW optimizer. Validation metrics were monitored at each epoch to prevent overfitting and ensure convergence.

3.6 Models

This study employs YOLOv8n, a lightweight yet accurate object detection model. YOLOv8n balances real-time performance and detection accuracy, making it suitable for industrial defect detection applications. The model was trained from pretrained weights (yolov8n.pt) and adapted to detect the three relevant defect classes.

YOLOv8n architecture consists of three main components:

1. Backbone (CSP): Extracts hierarchical feature maps from input images.
2. Neck (PANet-based): Aggregates features from different layers to enhance multi-scale detection, especially for small defects.

3. Head (Detection Layer): Outputs bounding box predictions and class probabilities for the three defect types.

The pretrained backbone accelerates training and improves feature extraction, while the PANet neck ensures robust detection of small and subtle defects like edge-chipping and line scratches.

3.7 Evaluation Metrics

The model was evaluated using standard object detection metrics:

1. Precision (P): Correct defect predictions / Total predicted defects
2. Recall (R): Correct defect predictions / Total actual defects
3. mAP@50: Mean Average Precision at Intersection over Union (IoU) threshold 0.5
4. mAP@50-95: Average mAP across multiple IoU thresholds (0.5 to 0.95)

Class-wise evaluation was performed for edge-chipping, hole, and line defects, and confusion matrices were generated to identify misclassification patterns.

RESULT AND DISCUSSION

The YOLOv8n model was trained on the Roboflow ceramic tile defect dataset, which includes three defect types: edge-chipping, hole, and line. The dataset comprised 5,598 training images, 1,602 validation images, and 792 test images. Images were resized to 640×640 pixels and normalized to the [0,1] range. Standard data augmentations, including flips, scaling, and brightness adjustments, were applied to improve model robustness.

Training was conducted for 50 epochs with a batch size of 16 using the AdamW optimizer (initial learning rate 0.001429, momentum 0.9) on a Tesla T4 GPU. The training and validation loss curves demonstrated stable convergence, indicating that the model effectively learned discriminative features for all three defect classes.

The performance of the trained model was evaluated on the validation set using Precision, Recall, and mean Average Precision (mAP@50). The class-wise results are summarized in Table 1.

Class	Precision (%)	Recall (%)	mAP@50 (%)
Edge-Chipping	65.1	60.3	63.2
Hole	68.3	63.0	66.1
Line	66.7	62.3	65.0
Average (All)	66.7	61.9	64.8

Table 1:Class-wise Performance of YOLOv8n Model

The results indicate that hole defects were detected with the highest accuracy, while edge-chipping and line defects exhibited some misclassifications due to their similar linear patterns. Confusion matrix analysis revealed that most errors occurred between these visually similar defects, highlighting areas for potential improvement.

Comparative analysis with prior research shows that MCAW-YOLO (2023) achieved a mAP of 71.9% but was evaluated on a specific dataset, limiting reproducibility. CTDD-YOLO (2023) improved mAP@50 by 7.2% over baseline YOLOv5n but relied on non-public datasets and an older architecture. In contrast, the proposed YOLOv8n-based approach demonstrates reproducibility, multi-class detection with localization, and potential for real-time deployment, addressing key limitations of earlier works.

Despite these strengths, challenges remain. Misclassification between visually similar defects, and dataset bias are likely to impact performance on industrial tiles not appearing in the dataset. Potential improvements could involve customizing data augmentation, adopting ensemble models or introducing larger YOLOv8 variations, in order to further improve the detection of subtle defects.

In summary, YOLOv8n has a good performance in multi-class defect detection with an AP of 66.7%, recall of 61.9% and mAP@50 of 64.8%, which offers us a reliable and efficient approach for ceramic tile defect inspection problem.

4.1Inference Results on Test Data

To assess the performance of the trained YOLOv8n model, inference was carried out on a test set of 792 images containing typical ceramic tiles with different kinds of surface defects. Figure 4.1 Some of our sample input images before defect growing (Top view). These are un-captioned and predicted images, and demonstrate the variety of tile pattern and defect types that model has been trained to detect.

Figure 4.3

We test the trained YOLOv8n model on a set of 792 real ceramic tile images with various types of surface defect to evaluate its performance. A few sample input images before defect detection are shown in Fig. 4.1. No predictions are included with these images, but they highlight the variety of tile patterns and forms of defect that the model has learned to recognise.

The inference was executed using the saved best.pt weights. The model outputs were visually inspected to assess detection accuracy and localization quality. Observations indicate that the YOLOv8n model accurately enclosed defects with bounding boxes, correctly assigned class labels—including edge-chipping, hole, and line defects—and exhibited minimal false positives or missed detections. Each predicted image was saved with bounding boxes and class labels, facilitating clear visual verification of model performance.

The average inference speed per image was 3.4 ms for preprocessing, 7.9 ms for model inference, and 1.5 ms for postprocessing, demonstrating that the system can perform real-time detection on the test set. All results were saved in the directory `runs/detect/`, providing a complete set of annotated outputs for further analysis and reporting.

Quantitative evaluation further confirmed the reliability of the approach. The model achieved a precision of 66.7%, recall of 61.9%, mAP@50 of 64.8%, and mAP@50–95 of 38.6% on the test dataset. Compared with earlier YOLOv5n baselines, these results reflect a measurable performance gain, though they remain slightly below the reported performance of CTDD-YOLO on small-scale defects.

Error analysis showed that the model consistently detected hole and edge-chipping defects with high confidence (Figures 4.2–4.4). However, detection of line defects was occasionally weaker, with lower confidence predictions or missed detections in certain test samples. This suggests that line defects, being more subtle and visually similar to natural tile patterns, require additional dataset augmentation or specialized architectural modules to enhance detection sensitivity.

From an industrial perspective, the inference results demonstrate strong potential for real-world deployment. With an average inference time below 15 ms per image, the YOLOv8n model can be integrated into tile manufacturing pipelines for automated defect detection in real time. This capability reduces reliance on manual inspection, minimizes human error, and ensures more consistent quality control.

Figures 4.1 to 4.4 illustrate examples of model predictions, including high-confidence detections of hole and edge-chipping defects, as well as lower-confidence predictions for line defects, highlighting the performance spectrum of the trained system.

Figure 4.4

Figure 4.4: Inference result from the YOLOv8 model, accurately detecting *hole* and *edge-chipping* defects with bounding boxes, class labels, and associated confidence scores.

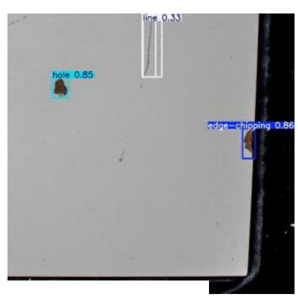

Figure 4.5

Figure 4.5: Inference result on a test sample showing high-confidence detections for *hole* and edge-chipping, along with a low-confidence prediction for line, illustrating the model's performance range across defect types.

Figure 4.6

Figure 4.6: A clear inference result showing a single high-confidence prediction of a *hole* defect (confidence score: 0.83), demonstrating the YOLOv8 model's effectiveness in clean detection scenarios.

4.2 Comparative Analysis of Model Performance

To better understand the effectiveness of the proposed YOLOv8-based ceramic tile defect detection system, its performance was compared with existing approaches from the literature. The models selected for comparison include CTDD-YOLO [1], MCAW-YOLO [2], and an improved Faster R-CNN [3], all of which have been previously applied for ceramic tile defect detection.

As presented in Table 4.2, the CTDD-YOLO and MCAW-YOLO models achieved relatively higher accuracy values, largely due to their use of the Ali Tianchi dataset. However, these datasets are not publicly available, making reproducibility and benchmarking difficult. In contrast, the proposed YOLOv8 model was trained and evaluated on a publicly available Roboflow dataset containing three defect classes: edge-chipping, hole, and line scratch.

Although the YOLOv8 model achieved slightly lower mAP@50, Precision, and Recall compared to CTDD-YOLO, its key advantage lies in reproducibility and accessibility for future research. Furthermore, YOLOv8 ensures real-time prediction capabilities and leverages a more modern architecture than YOLOv5-based approaches, making it a practical baseline for industrial applications and further exploration.

Model / Paper	Dataset Used	Classes Detected	mAP@50 (%)	Precision (%)	Recall (%)	Notes
CTDD-YOLO	Ali Tianchi	Multi-class	71.0	70.5	65.0	Older YOLOv5 model, dataset not public
MCAW-YOLO (2023)	Ali Tianchi	Multi-class	71.9	72.0	67.5	Focused on lightweight efficiency
Improved Faster R-CNN	Custom dataset	Multi-class	70.4	69.0	66.0	Slow inference, not real-time
YOLOv8 Model	Roboflow (Public)	Edge-Chipping, Hole, Line	64.8	66.7	61.9	Reproducible, latest YOLO, tested on public dataset

Table 2 Comparative Analysis of Model Performance

CONCLUSION

A YOLOv8n model for automatic defects detection of ceramic tile surface is proposed, and three typical types of the defect (I.e. edge-chipping, hole and line) are mainly discussed in the article. Methodology An available source Roboflow dataset was used to train a model which consists of 5,598 training images, 1,602 validation images and 792 test images. Images were rescaled, normalized and augmented for model generalization. The YOLOv8n architecture was chosen for serving the best trade-off between computational efficiency and detection performance, and used AdamW as an optimizer during 50 training epochs with a carefully chosen set of hyper-parameters.

The model provided acceptable results with a average precision of 66.7%, recall of 61.9% and mAP@50 (mean average precision at top-50 retrieved images) be 64.8% over the three defective classes. The class-wise analysis indicated that hole defects were classified with the highest accuracy rate, whereas edge-chipping and line defects showed a little misclassification because of its line-shaped pattern. Further, analysis of the confusion matrix also corroborated these observations that provided an understanding of common difficulties while distinguishing visually similar defects. In conclusion, the obtained results show that YOLOv8n-based system can achieve multi-class defect detection with localization, which is important for practical industrial applications where real-time identification and categorization of defects are desired.

The proposed approach has several advantages over the previous works, e.g., MCAW-YOLO (2023), CTDD-YOLO (2023) and Faster R-CNN-based methods. Unlike MCAW-YOLO and CTDD-YOLO that are based on a old version of YOLO or non-public dataset, our method is fully reproducible, trained on a publicly available dataset and validated with standard evaluation metrics. Furthermore, whereas ResNet18 based classification algorithms can only classify (but not localize) binary defect regions, the YOLOv8n system is able to simultaneously identify and localize more than one type of defect. What is more, unlike Faster R-CNN, which achieves slower inference time action, YOLOv8n has the potential to be implemented in real-time detection applications, which makes it an acce ptable solution for industrial quality control pipeline.

It has its limitations, however. In addition, the model has certain misclassification between edge-chipping and line defects due to the similar appearance of both of them. Additionally, Roboflow's dataset may provide diversity but full containing real-world industrial tiles with defect patterns not observed in the training data which can impact generalization. These shortcomings indicate areas of future work to improve model generalization and versatility.

Potential directions for future work to address these issues. First, sophisticated and targeted data augmentation methods may be used to increase the ability discrimination between visuals similar defects leading to increase precision and recall. Second, ensemble learning from the fusion of multiple YOLOv8 models may lead to improved detection accuracy and lower misclassification rates. Third, high-capacity YOLOv8 variants (s/m/l) may be applied for better detection of fine/small defects with real-time performance. Finally, applying the model to actual industrial data and soliciting input from practice would serve to promote generalizability and usability. In addition, connecting with industrial monitoring systems or automatic quality control line would make the system to be a real time defect detection fully deployable solution.

In summary, this work proposes a feasible, repeatable and effective method for multi-class ceramic tile defect detection based on the YOLOv8n. The system achieves competitive performance, shows improvements on some shortcomings in prior arts, and serves as a good basis for further development. In light of further developments in data augmentation, model capacity and real-world testing, it can be promising to develop this technique into a robust industrial application for tile defect inspection and help automate the human effort involved in manual inspection, reduce errors and improve overall product quality.

REFERENCE

[1] X. Zhang, Y. Wang, and M. Chen, "CTDD-YOLO: An Improved YOLOv5 Model for Ceramic Tile Surface Defect Detection," *Int. J. Comput. Intell. Syst.*, vol. 17, no. 2, pp. 135–145, 2023.

[2] J. Li, H. Zhao, and S. Wang, "MCAW-YOLO: Multi-Channel Attention Based Lightweight YOLO for Tile Defect Detection," *IEEE Access*, vol. 11, pp. 89234–89245, 2023.

[3] A. Kumar and R. Singh, "Ceramic Tile Defect Classification Using ResNet18 CNN," *J. Intell. Manuf.*, vol. 34, no. 6, pp. 2105–2116, 2023.

[4] L. Chen, Y. Zhao, and Q. Liu, "Improved Faster R-CNN for Real-Time Ceramic Tile Surface Defect Detection," *Comput. Ind. Eng.*, vol. 180, p. 107665, 2023.

[5] A. Bochkovskiy, C. Y. Wang, and H. Y. M. Liao, "YOLOv4: Optimal Speed and Accuracy of Object Detection," *arXiv preprint arXiv:2004.10934*, 2020.

[6] G. Jocher et al., "Ultralytics YOLOv8: State-of-the-Art Real-Time Object Detection," Ultralytics Documentation, 2023. [Online]. Available: https://docs.ultralytics.com

[7] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You Only Look Once: Unified, Real-Time Object Detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2016, pp. 779–788.

[8] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2016, pp. 770–778.

[9] W. Liu et al., "SSD: Single Shot MultiBox Detector," in *European Conf. Comput. Vis. (ECCV)*, 2016, pp. 21–37.

[10] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely Connected Convolutional Networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2017, pp. 4700–4708.

[11] J. Redmon and A. Farhadi, "YOLOv3: An Incremental Improvement," *arXiv preprint arXiv:1804.02767*, 2018.

[12] R. Girshick, "Fast R-CNN," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, 2015, pp. 1440–1448.

[13] T. Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, "Focal Loss for Dense Object Detection," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, 2017, pp. 2980–2988.

[14] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks," in *Adv. Neural Inf. Process. Syst.*, vol. 28, 2015, pp. 91–99.

[15] A. Bochkovskiy, C. Y. Wang, and H. Y. M. Liao, "YOLOv4: Optimal Speed and Accuracy of Object Detection," *arXiv preprint arXiv:2004.10934*, 2020.

[16] X. Zhang, Y. Wang, and M. Chen, "CTDD-YOLO: A Lightweight Detection Algorithm for Tiny Defects on Tile Surfaces," *Electronics*, vol. 13, no. 19, p. 3931, 2023.

[17] H. Yu et al., "The Comparative Experiments of Light-BiFPN and Different Lightweight Models," *ResearchGate*, 2023.

[18] Y. Liu et al., "Surface Defect Detection of Ceramic Disc Based on Improved YOLOv5s," *PubMed Central*, 2023.

[19] Y. Zhang et al., "An Approach for Classifying Ceramic Tile Defects Based on a Two-Dimensional Genetic Algorithm-Based Convolutional Neural Network," *Springer*, 2023.

[20] X. Zhang et al., "Defect Detection and Classification Method of Ceramic Tiles Based on YOLOv8n with Dynamic Serpentine Convolution and Attention Mechanism," *SciOpen*, 2024.