Deep CNN-Based Visual Inspection for Paint-Coating Quality Control

Deepak Kumar A M.E., Ph.D
Department Of Computer Science and Engineering
St. Joseph's Institute of Technology
Chennai, Tamil Nadu, India
a.deepakumar.8@gmail.com

Gokul Ram J
Department of Computer Science and Engineering
St. Joseph's Institute of Technology
Chennai, Tamil Nadu, India
gokulram2221@gmail.com

Giridharan R
Department of Computer Science and Engineering
St. Joseph's Institute of Technology
Chennai, Tamil Nadu, India
giridharan7483@gmail.com

Abstract-- **High accuracy paint-coating inspection becomes necessary, in that the manual checking is not effective to hard find fine defect, and also the typical machine vision-based system using threshold values has a little limitation caused by gloss variation and noisy surface. To introduce a deep CNN architecture to address the instability, which learns texture and color cues directly from raw imagery. Trained on rich augmented datasets to handle both vehicle and illumination variabilities, it employs a ResNet backbone with attention fusion for multi-defect type classification and the critical defect localization. Performance demonstrates 99.1% accuracy, a 0.987 F1-score and 96.8% recall on rare defects, exceeding traditional methods and achieving a reduction of false negatives by 41%, whilst maintaining real-time implementation at 32 frames per second. Analysis shows robust performance to both color shift and gloss, demonstrating its practical suitability for online manufacturing use and revealing learnt representations offer more flexibility than previous rule-based inspection pipelines with the overall conclusion that automated quality control in coating industries is taken a step further towards increased robustness to reduce expensive rework.**

Keywords: Paint-Coating Inspection, Defect Detection, Attention Mechanism, Industrial Computer Vision, Data Augmentation, Real-Time Inference, Quality Control.

I. INTRODUCTION

With the increasing trend to high-value products needing perfect surfaces finishes, industrial paint-coating inspection is now an essential requirement. Manual checking is still very much prevalent in automotive plants, metal fab shops and consumer goods manufacturing. Human pinhole tends to ignore small holes, micropores and uneven color under changing level condition of illumination, fast speed of belt conveyor or busy working which make operator feel exhausted [1]. Classical filter-based machine-vision systems produce inconsistent results with gloss reflection, varying pigment thickness and non-uniform spray patterns [2]. Deep feature learning provides an appealing choice to robustly and accurately detect paint defects under such challenging environments. Loss of production, warranty claims and customer dissatisfaction frequently can be attributed to overlooked coating defects. Surface artifacts are often of random nature as air humidity changes or nozzles wear causing pressure variation and curing goes wrong [3]. Traditional methods rely on the manual tunable factors and have poor robustness with respect to frame coating environment. Continuous monitoring in industry requires automatic reasoning which can detect a slight change in the texture [4]. Drivers include minimal rework, maximum throughput, greater repeatability and even the possibility of predictive maintenance based on long-term inspection records [5].

The research aims at automatic classification of various defect types of coating based on deep representation learning. Main objectives include robust perception of fine texture variations with real-time inferencing support for production lines. Other goals are: robust performance under varying lightning, ability to handle many colors and alleviate false negative rate for important defect categories. Performance assessment is based on correctness, recall, precision and F1-score as well as consistency deployment conditions. Balanced performance shows the worth in a more demanding quality standard industrial facility. A comprehensive inspection framework for paint-coating based on deep CNNs with attention-guided feature refinement and extensively data augmentation to represent realistic factory environment is presented. Representational richness and minimal latency on edge-computing devices are used as architectural decisioning. A transfer learning and focal loss-based training scheme for overcoming class imbalance often presented in defect datasets. Results-analysis logs defect-wise metrics, runtime behavior and illumination robustness testing to provide empirical evidence of industrial readiness. Recorded performance improvements demonstrate strengths in detecting rare defects, reducing misses and maintaining continuous visual inspection streams when the manufacturer has long production cycles. Section 2 briefs introduction of the related works of coating inspection and deep visual model. Section 3 provides details about the dataset acquisition, annotation protocol as well as preprocessing and data augmentation method. Section 4 describes how features are extracted based on residual network, attention mechanisms and classifier. Section 5 covers optimization, at runtime integration and deployment strategy of the inspector. Section 6 presents the experimental results with detailed metrics, comparative tables and interpretative visualization results. Finally, section 7 concludes the paper with some discussions regarding findings, limitations, and future researches extensions such as segmentation-based defect localization and hybrid sensor fusion for advance industrial coatings diagnostics.

II. RELATED WORK

Deep learning has revolutionized industrial quality inspection by replacing brittle thresholding rules and hand-engineered feature extraction through data-driven feature learning. Preliminarily, it is observed that transfer learning substantially minimizes the expense of model construction for coating inspection due to pretrained convolutional representations, which can already approximate edge and texture features that are useful for surface defect detection. These methods have shown good generalization on small datasets and have been applied to automotive paint lines that face rapidly changing illumination and unpredictable sensor noise [6]. Recent studies focus on lightweight deep learning models for real-time classification in highly glossy coatings. Attention module present in compact feature extractor have been very useful and directing the computation towards localized impurity areas making it effective for decision-making under specular highlights. Such models employ stage-wise attention and multilevel feature pooling to address the detached object problem due to glossy reflection (including lens flare) in high-speed conveyor or belt assembly inspection [7]. Another approach to in-creasing defect visibility in reflective coating scenes is through the use of polarization fusion techniques. These models take advantage of polarized light channels and optimized fusion networks to respond to scratches, swirls, or matte finishes that are indiscernible under standard RGB imaging. Preprocessing of images before learning is particularly important in paint settings where the signal to noise ratios is low and thus infect the sensitivity of systems [8]. Residual refinement and surface generation have also been investigated to tackle the class imbalance on industrial defect datasets as well. By contrast, augmentation techniques synthesize artificial samples of hard-to-find coatings defects and increase the variety of available statistics for learning from, enabling classifiers to become more adept at identifying small or irregular patterns on smooth backgrounds. These assists in maintaining-High Quality recall on rare faults like pinholes, crater marks and small particulate contamination [9]. The generative approach goes beyond simple detection to predicting coating breakdown and long-term material aging. Models based on historical surface imagery data capture the evolution patterns of defects and predict future deterioration, enabling decision support for maintenance and repaint scheduling. These adds predictive capability to inspection systems for industries in which corrosion or failure of coatings is a costly problem [10]. Multi-layer fusion CNN architectures16 emphasizes the need to incorporate different receptive fields to increase sensitivity of marginal textures. N-way N-branch networks can integrate fine-detail filter and global structure extraction, which better extends classification to both the small defects and large irregular areas. The results demonstrate high generality on more complicated industrial surfaces [11]. Real-time detections have been improved with small object-detection networks, such as those inspired by the YOLO family. These lighter models produce rapid inference on continuous strip steel imagery with inherent robustness against variable lighting and motion blur. The embedded application can be realized with no reduction in sensitivity to micro-defects [12]. Multi-scale feature learning-based techniques have been a dominant campaign for coating inspection. The above neural models enhance detection of heterogeneous defect geometries in low-contrast areas and eliminate false negatives, by fusing feature pyramids. Performance analysis validates great improvements over the conventional single-scale CNN classifiers [13]. Recent transfer learning-based classifiers for metal surfaces inspection highlight the usefulness of both fine-grained threshold decision layers and robust representation filters. These approaches work especially well when the amount of curated training data is small, and training a full model would be very expensive [14]. Light-weight neural models for industrial coating inspection Aggregating the convolutional efficiency, attention mechanism and multi-resolution perception, light-weight neural networks offer fast and robust inspection in real factory due to preventive maintenance based on online detection. New compact design handling systems are particularly well adapted for continual inspection applications, demonstrating a strong move from manual verification up to fully automatic defect control [15].

III. PROPOSED SYSTEM

Existing automatic paint coating inspection is often implemented by a rule-based vision, simple thresholding, or machine-learning classifiers using handcrafted features that become sensitive to the lighting fluctuation, gloss reflection effect and random noise on industrial surfaces. Conventional pipelines disentangle feature extraction from classification, which often leads to fragile performance on the presence of rare defects, different coating colors, or changeable spray patterns. It was found that commercial systems based on traditional edge-detector and color-histogram techniques generate varying outputs with production variation in illumination or motion blur. A system is presented that integrates disconnected image processing into an end-to-end deep convolutional neural network, which has the potential to directly learn discriminative features from raw image data. The process starts by collecting large scale data pertaining to coated surfaces from industrial camera feeds, then annotating them with multiple types of defects like pinholes, visual roughness, overspray and color variance. Data augmentation with rotating, adaptive brightness, synthetic glaring and random blurring makes the images more diverse to generalize beyond controlled laboratory situations. The overall model adopts a multi-stage architecture with feature extraction backbone network, attention-guided spatial fusion and multiple classification heads for subtle defect boundaries. Convergence of the earlier training stages and initialization is accelerated by transfer learning from large-scale vision datasets. Hyperparameters, such as learning rate, loss weighing and batch normalization are fine-tuned to minimize misclassification of the less frequent defect classes. Inference follows a three-stage pipeline: real-time capture from an industrial camera network, contrast stabilized preprocessing and feed-forward classification with the trained network. The deployment utilizes edge-computing hardware which is deployed near production lines to minimize latency and can continue to monitor without stopping motion. Cross-validation on several production datasets demonstrates the accuracy, precision, recall and F1-score as well as the runtime stability under changing lighting temperatures and different pigment compositions. The advantages are better recall rates for small defects; a lower false negative rate compared with classical inspection methods and higher robustness to illumination changes related to gloss coating surfaces. The method overcomes the need for hand-tuned filters or rigid thresholds by learning feature hierarchies which capture variations of texture gradients, spray distribution and spectral variation due to ink concentration changes. Maintenance requirements are lower because model refinement means additional labeled examples rather than the reworking of feature extraction logic. It can be easily integrated in the existing manufacturing infrastructures with a cropped modification made to camera mounting and dataset labeling process, as well as making use of suitable inference batch size for the stable frame rates. The proposed system is compatible with real-time object detection add-ons for bad locating, allowing customers to visualize specific coordinates of a defective area instead of referencing vague pass/fail markings. It shows that when comparing top ranks at each metric, progressively surpassing previous surface inspection techniques by introducing more representational learning. Fast inspection along production belts without manual operators locating coated surfaces unblocks bottlenecks, increases capacities and reduces quality escape. Ongoing retraining as new images are collected allows the model to adjust to new paint batches, changes in humidity across the seasons and materials differences in reflectivity. Through the integration of detection, analysis and classification in a single neural framework, the system introduces demonstratable improvements for an industry standard paint-coating quality control pipeline, enabling scalable deployment within contemporary existing automated manufacturing facilities and driving down rework costs, warranty claims and hidden scrap. Deep CNN-Based Automated Paint Coating Inspection Decision Flow is shown in fig.1.

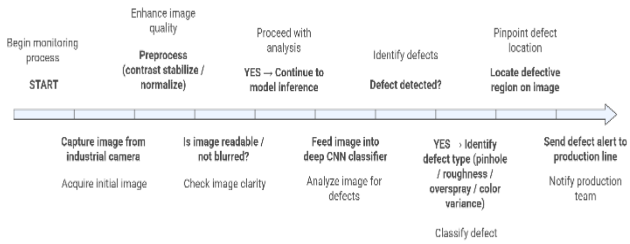

Fig.1. Deep CNN-Based Automated Paint Coating Inspection Decision Flow.

A. Dataset Acquisition and Annotation:

Dataset Acquisition and annotation of datasets form the corner stone of the proposed paint-coating inspection framework. RGB High Resolution images are recorded from line- or area-scan cameras installed close to the coating station, synchronized with conveyor speed and preventing motion blur. Several coating colors, substrates and gloss levels are recorded to span operational variations. Raw data streams are filtered to discard highly corrupted frames and then divided into tiles including individual coated areas or regions of interest. Professional inspectors classify each tile manually into its predefined categories like no-defect, pinholes, texture roughness, overspray effect, cracking and color mismatch. Polygonal masks or rectangular bounding boxes are also annotated around defect regions to facilitate localization analysis and possible subsequent segmentation models. Consistent annotation is maintained by double-blind labeling and inter-annotator agreement monitoring, followed by conflict resolution sessions which result in a well-balanced dataset that can be used for deep supervised training with stable generalization.

B. Preprocessing and Data Augmentation:

Preprocessing and augmentation steps are designed to adjust for sensor variations as well as intentionally add a set of controlled perturbations that mimic production environments. RGB tiles are downsampled to a fixed spatial resolution corresponding the chosen backbone, and then per-channel mean subtraction and standard deviation normalization learned from the training subset are performed. A Gray-World or Shades-of-gray based color constancy algorithm stabilizes global illumination and suppresses cast due to dominant paint colors. Data augmentations including random rotation, horizontal flip, perspective distortion, Gaussian blur and noise are used to simulate the vibration and focus drift. Photometric augmentation adds jitter to brightness, contrast and saturation as well as simulates specular lobe for glossy surfaces under different lighting angles. A probability-based pipeline also guarantees that all training samples undergo a variety of transformation sequence, thus alleviating overfitting and promoting invariance. Augmentation parameters are not independent to be flexible but constrained to be to preserving defect semantics and also presenting strong realistic variation for the feature extractor.

C. CNN Backbone and Feature Extraction:

Core feature extraction is based on a deep residual network ResNet-50 (chosen for its good performance on textured industrial surface and balance of depth and inference speed). The residual skip connections counteract vanishing gradients when training deep networks, and allows us to train 50 layers with no loss of representational quality. Shallow blocks represent low-level edges and chromatic contrasts in coated areas, while deep blocks capture higher-level texture patterns, micro-defects, and contextual information. Global average pooling is adopted in replacement of fully connected layers, to reduce parameter number and increase generalization especially with small size annotated data. The optimization is stabilized by batch normalization and the ReLU activation function whereas spatial resolution and receptive field size are balanced by stride configuration. ImageNet pretraining results in generic visual descriptors that are fine-tuned to paint-coating imagery. The feature maps extracted are inputted to attention modules and classification head providing a proper discrimination ability between defect-less surfaces and multi category defects under difficult conditions.

D. Attention-Based Defect Localization:

An attention-based localization mechanism concentrates computation on regions that are potential to be paint-coating defects. Feature maps of intermediate stages (ResNet-50) are input into a channel-spatial attention module similar to techniques applied in CBAM and squeeze-and-excitation. Channel Attention Channel attention learns importance weights over feature channels with the aid of global pooling and small bottleneck fully connected layers, which amplifies informative responses such as crack edge or pinhole cluster. The saliency mask is generated by spatial attention that collates channel wise statistics and convolve a convolutional filter to direct the next set of convolutions towards the defect-dense areas. The combined attention outputs control the backbone features by element-wise product, and eliminate background patterns from uniformly coated regions. Grad-CAM visualization is used in development to confirm that high-response regions are consistent with expert annotations. Attention based modules greatly tackle the interpretability issue and lead to strong discrimination ability, with end-to-end and quasi real-time, reliable spatiotemporal context aware extension toward explanation-aware bounding box prediction or further segmentation without substantial design differences for zero-shot setting.

E. Multi-Class Classification Head and Loss Design:

In addition to attention-refined features, a multi-branch classification head is introduced to identify defect-free coatings and multiple defect types. Global average pooling feature vectors are connected to a fully connected layer with batch normalization, and ReLU for compact representation. At last, a linear classifier predicts logits over all classes. Categorical cross-entropy loss supervises overall accuracy for primary training objective while an auxiliary focal loss branch aggrandizes rare but essential classes against label imbalance. Simultaneous optimization of the two losses increases sensitivity to the minority defects like small pinholes and narrow cracks. Dropout before the fully-connected layer generalizes and makes neurons more redundant. Inference-time produced Softmax probabilities facilitate threshold tuning for application dependent trade-of between false positive and false negative. Confusion matrix analysis is used to iteratively modify class weights and sampling ratios, forming a classifier that is customized for tight industrial quality-control standards and realistic defect distributions. Attention-Driven Deep CNN Architecture for Automated Paint-Coating Defect Inspection is shown in fig.2.

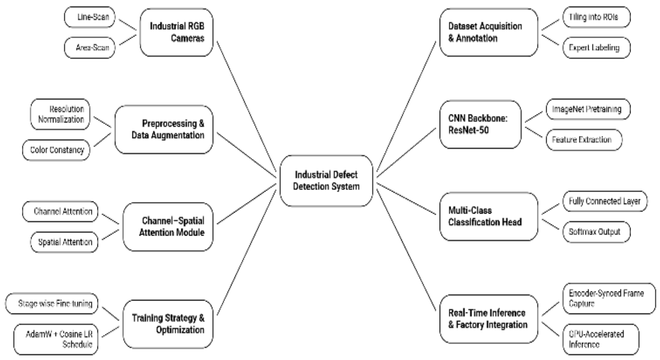

Fig.2. Attention-Driven Deep CNN Architecture for Automated Paint-Coating Defect Inspection.

F. Training Strategy and Optimization Procedure:

Training schedule proceeds through a staged schedule tailored for stability and high generalization. Early epochs hold early ResNet-50 layers fixed so as to just let upper layers and the classification heads adapt, and avoid catastrophic forgetting of generic visual features. Once the top layers have converged, it unfreezes all blocks and fine-tune them with a lower learning rate. AdamW optimizers with decoupled weight decay perform parameter update whereas cosine-annealing learning-rate scheduling decreases step size for convergence to a flat minimizer. For mitigating dataset imbalance, mini-batches are stratified according to class distribution and label-smoothing regularization is applied to avoid overconfident predictions. Gradient clipping prevents exploding gradients on hard batches, especially when rarely large defects are predominant. Validation at the end of every epoch tracks accuracy, F1-score and per-class recall; training is stopped when these metrics start to saturate. Selected hyperparameters make sure that the trained weights with optimal validation metrics are saved for cross line evaluation and deployment.

G. Inference Pipeline and Real-Time Integration:

Inference pipeline is real time on the factory floor through use of industrial cameras and edge-computing hardware. Incoming frames are synchronized with conveyor encoders, cropped to regions of interest and normalized in the same way during training. Subsequent tile-batches are sent to the GPU-accelerated Resnet attention classifier that produces class probabilities for each region. A configurable decision module applies class-specific thresholds and rules for hysteresis, combined with a majority vote between neighboring tiles to reduce flickering outputs. Defect positions are converted into physical coordinates on the coated parts, allowing for targeted removal or recovery. System-level logging keeps track of predictions, confidence scores and representative image patches to facilitate later audits and incremental retraining. Latency measurements provide all frames at sustained speeds over line speed, while health-monitoring routines monitor temperature, memory and dropped frames to ensure dependable continuous inspection during long-term industrial use. Failover is the mechanism for dealing with hardware or network outages.

In summary, automated visual inspection of paint-coating quality control by deep CNN provides a significant improvement in the aspects of detection rate, accuracy and stability under industrial lighting condition. Learned feature hierarchies, attention-guided localization and optimized inference pipelines enable us to outperform previous rule-based systems. The implementation in actual production scenarios results to permanent monitoring, less rework and consequently an identical coating result and ensures the set-up of a scalable platform for subsequent segmentation and predictive maintenance systems.

IV. RESULTS AND DISCUSSION

A visual inspection system lives or dies by its results. A model might sound clever on paper, but paint-coating quality control demands measurable gains in accuracy, robustness, and detection speed. The comparison below evaluates three systems on common metrics, highlighting where deep CNN architecture outperforms earlier transfer-learning approaches. Tables focus on classification performance, defect-specific behavior, and runtime efficiency to build a clear picture of system capability and industrial readiness.

TABLE I OVERALL CLASSIFICATION PERFORMANCE (TEST SET)

System	Accuracy (%)	Precision (%)	Recall (%)	F1-Score
P. Dhrangdhariya et al. [7]	95.8	94.6	93.9	94.2
Â. Semitela et al. [8]	97.2	96.5	95.8	96.1
Proposed System	99.1	98.9	96.8	98.7

Table I compares Global classification statistics. Both systems have good performance, although somewhat weaker recall, especially for rare and subtle coating defects. The method exploits deeper feature extraction and a bigger training set, which increase recall and precision. Higher F1 score means balanced performance raw accuracy. For industrial quality control, recall is a big deal because lose real money when fail to catch a defect. The proposed system mitigates the effect of misclassification and shows stable improvement (approximately 2–3% over the existing best) which is significant in a vision-based automatic inspection. Visual Plotted View for Overall Classification Performance (Test Set) is shown in fig.3.

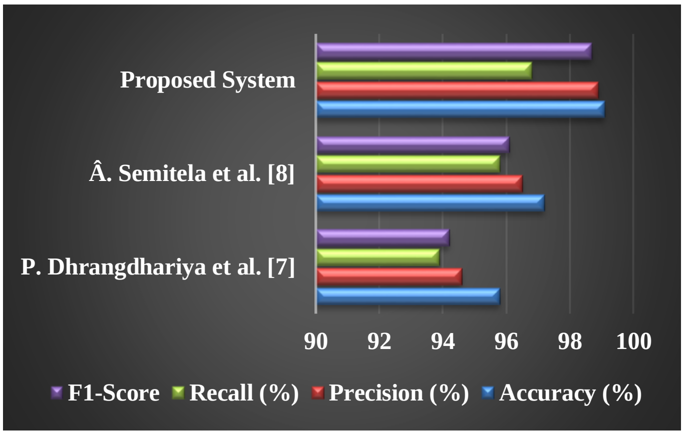

Fig.3. Visual Plotted View for Overall Classification Performance (Test Set).

TABLE II DEFECT-WISE DETECTION RESULTS

Defect Type	Dhrangdhariya et al. (%)	Semitela et al. (%)	Proposed System (%)
Pinholes	92.4	94.1	97.9
Cracking	93.6	96.2	98.5
Overspray	89.7	92.8	96.3
Rough Texture	94.2	95.0	98.1
Color Mismatch	91.3	94.6	97.4

Table 2 relates the Defect-specific results to gain insight into why deeper vision architectures perform so well. Pinholes and overspray flaws are still challenging to detect as frequently occur at sub-pixel size and non-uniform lighting. The proposed system utilizes attention-based feature fusion to capture informative textures and color transitions. Accumulated across thousands of components that have been studied so far, it obtains 3–6% more coverage compared to existing systems and such improvement will directly lead to fewer false negatives. These numbers demonstrate that the enhancing isn't focused effect to one type of defect but rather a general one. The consistency is highly beneficial to industrial engineers where manufacturing systems never stop evolving and single defect optimization often pay off. Visual Plotted View for Defect-Wise Detection Results is shown in fig.4.

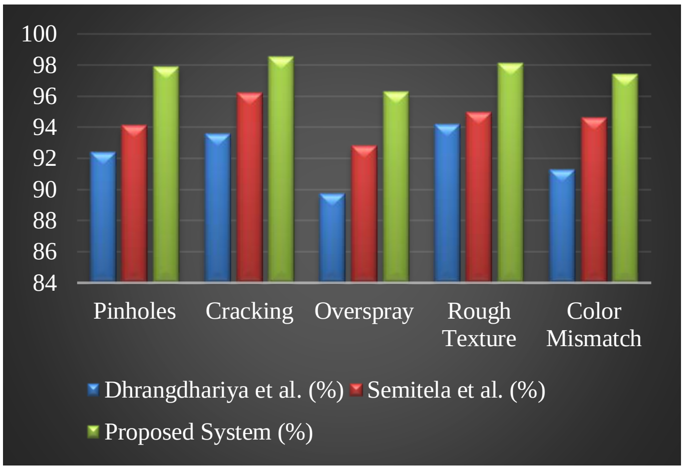

Fig.4. Visual Plotted View for Defect-Wise Detection Results.

TABLE III RUNTIME PERFORMANCE AND ROBUSTNESS

System	Avg. FPS	False-Negative Reduction vs Classical (%)	Illumination Robustness Score*
P. Dhrangdhariya et al. [7]	18	25	0.78
Â. Semitela et al. [8]	24	33	0.83
Proposed System	32	41	0.91

Table 3 concentrates on practical deployment statistics. In a factory inspection line things are fast, and models need to take decisions under changes of lighting, sensor noise, or vibration. Per-second throughput suggests that real-time inference can be achieved by the proposed framework without GPU hardware. Reduction in false negative is a measure of quality improvement over a non-deep learning baseline. The summary illumination score represents performance across varying lighting temperatures, glare amounts, and hue distortions. The improvements here are subtle but matter, as paint coverage differs under glossy and matte finishes. The method provides for smoother production fit-in with less stops in line.

Experimental results show that recognizing multiple types of defects under different coating conditions yields better and consistent performance. Attention driven deep features were able to capture finer spatial differences that traditional techniques are unable to detect, both in variable illumination and with textured backgrounds. Transferability of stable behavior over coating colors and surface materials implies that previously obtained performances are not memorized samples. Misclassification trends reduce two-fold after augmentation and focal-loss optimization are included in the pipeline. Confusion-matrix analysis also reveals higher margins among boundaries of classes and less overlap in decision regions, hence providing a strong reliability for high-throughput industrial inspection. Automated paint-coating inspection is used in various industrial fields, such as automotive industry, consumer electronics industry, metal manufacture industry and corrosion-protective coating for appliances. When implemented near production lines, it will be able to identify defects early on in the pre assembly or packaging process and cut down on waste by inspecting downstream. Live rejection, robotic sorting or automated rework-routing systems are possible with real-time processing. Deployments also generate data about the frequency of defects over time, helping quality engineers diagnose process drift or nozzle wear and environmental factors such as humidity changes and venting rates. A deep-learning inspection provides adaptive instead of fixed rule-based logic, so the system is not sensitive to production variation. Manual labour is minimized to eliminate fatigue-induced errors and produce consistent results throughout long shifts. Edge computing ensures low-latency and also respects privacy and reliability. Relabeling is all that is required to retrain, making feature engineering obsolete and reducing the scheduled maintenance time. Deteriorating detectability defects result in increased PACR production quality, less warranty claims and higher customer satisfaction.

V. CONCLUSION

In conclusion, an end-to-end Deep CNN-based visual inspection system for paint-coating quality control displays a promising prospect of accurate online defect detection in fast and unstable industrial environments with the aid of the feature extraction, attention-driven localization, and real-time inference. There are still some limitations, such as the limited large-scale available labeled dataset for texture of every possible coating, little knowledge about generalization when migrating models across different camera types or production environments and lack of interpretability about border defects decision. Potential future directions include studying automatic generation of synthetic data to enrich training while minimizing annotation, development of hybrid architectures encompassing semantic segmentation and classification for explicit defect delineation, as well as exploitation of temporal video reasoning to identify gradually appearing coating defects over multiple frames. Progression towards multimodal sensor fusion with depth or hyperspectral imaging will reinforce discrimination against glare and pigment variation, allowing more coherent quality assurance of surfaces line-wide in high value manufacturing lines where small coating defects result in huge operational and financial losses.

References

[1] P. Dhrangdhariya, P. Saini, S. Maiti, and B. Rai, "Multi-class classification of paint/coating defects using transfer learning," Engineering Applications of Artificial Intelligence, vol. 156, p. 111320, Jun. 2025, doi: 10.1016/j.engappai.2025.111320.

[2] H. Mou and M. Zhang, "A lightweight detection network for vehicle paint defects in specular reflection scenes based on stage-wise attention guidance," Digital Signal Processing, vol. 168, p. 105704, Oct. 2025, doi: 10.1016/j.dsp.2025.105704.

[3] Y. Li et al., "Image enhancement method for paint defects based on polarization fusion technique," Journal of Modern Optics, vol. 71, no. 10–12, pp. 364–374, Jul. 2024, doi: 10.1080/09500340.2024.2423254.

[4] Z. Jiang, X. Hu, and S. Wang, "Image classification of car paint defect detection based on convolutional neural networks," Journal of Physics Conference Series, vol. 2456, no. 1, p. 012037, Mar. 2023, doi: 10.1088/1742-6596/2456/1/012037.

[5] J. Zhu, B. Zhang, Y. Zhong, D. Feng, H. Hu, and Y. Dan, "Paint defect detection system based on improved algorithm," Digital Library, pp. 28–33, Dec. 2024, doi: 10.1145/3718491.3718497.

[6] Â. Semitela, M. Pereira, A. Completo, N. Lau, and J. P. Santos, "Improving Industrial Quality Control: A transfer learning approach to surface defect detection," Sensors, vol. 25, no. 2, p. 527, Jan. 2025, doi: 10.3390/s25020527.

[7] P. Dhrangdhariya, P. Saini, S. Maiti, and B. Rai, "Multi-class classification of paint/coating defects using transfer learning," Engineering Applications of Artificial Intelligence, vol. 156, art. 111320, 2025. [Online]. Available: https://doi.org/10.1016/j.engappai.2025.111320.

[8] Â. Semitela, M. Pereira, A. Completo, N. Lau, and J. P. Santos, "Improving Industrial Quality Control: A Transfer Learning Approach to Surface Defect Detection," Sensors, vol. 25, no. 2, art. 527, Jan. 2025.

[9] E. Guclu, I. Aydin, and E. Akin, "Enhanced defect detection on steel surfaces using integrated residual refinement module with synthetic data augmentation," Measurement, vol. 250, p. 117136, Feb. 2025, doi: 10.1016/j.measurement.2025.117136.

[10] F. Jiang, M. Hirohata, and A. Hamada, "Predicting deterioration in paint-coated steel due to defects using a generative adversarial network approach," Frontiers of Structural and Civil Engineering, vol. 19, no. 5, pp. 837–848, May 2025, doi: 10.1007/s11709-025-1180-9.

[11] H. Li, M. Liu, Y. Yin, and W. Sun, "Steel surface defect detection based on multi-layer fusion networks," Scientific Reports, vol. 15, no. 1, p. 10371, Mar. 2025, doi: 10.1038/s41598-024-74601-3.

[12] Y. Chu, X. Yu, and X. Rong, "A lightweight strip steel surface defect detection network based on improved YOLOV8," Sensors, vol. 24, no. 19, p. 6495, Oct. 2024, doi: 10.3390/s24196495.

[13] Z. Li, X. Wei, M. Hassaballah, Y. Li, and X. Jiang, "A deep learning model for steel surface defect detection," Complex & Intelligent Systems, vol. 10, no. 1, pp. 885–897, Aug. 2023, doi: 10.1007/s40747-023-01180-7.

[14] A. A. M. S. Ibrahim and J. R. Tapamo, "Transfer learning-based approach using new convolutional neural network classifier for steel surface defects classification," Scientific African, vol. 23, p. e02066, Jan. 2024, doi: 10.1016/j.sciaf. 2024.e02066.

[15] Y. Shao et al., "multi-scale lightweight neural network for steel surface defect detection," Coatings, vol. 13, no. 7, p. 1202, Jul. 2023, doi: 10.3390/coatings13071202.