10.48047/jocaaa.2024.33.08.517

Automated Industrial Defect Detection and Accurate Fractional-Grade Classification

M. Ramana Kumar[1*], S. Sreenath Kashyap[2]

[1]Associate Professor, Department of CSE, KPRIT College of Engineering, Ghatkesar, Telangana, 500088, India

[2]Professor & Head, Department of ECE, Kommuuri Pratap Reddy Institute of Technology, Ghatkesar, Telangana, 500088, India

[*]Corresponding Email: drramanakumar@kpritcoe.in

Abstract

In modern industrial manufacturing, ensuring the structural integrity of magnetic tiles is essential, as surface defects can severely compromise electromagnetic performance and operational reliability. Traditional manual visual inspection methods are inherently slow, labour-intensive, and prone to human error or subjectivity. To address these challenges, this paper presents an automated quality control framework for magnetic tile defect classification leveraging machine learning and computer vision architectures. The developed pipeline processes raw tile images by uniforming structural dimensions to a standard matrix, flattening the localized spatial arrays into high-dimensional numerical feature tensors. To evaluate the predictive efficiency of the system, two distinct classification paradigms, a parametric Support Vector Machine (SVM) utilizing a polynomial kernel and a non-parametric ensemble Random Forest Classifier (RFC) were trained and validated on partitioned imaging strata. A comprehensive evaluation framework was implemented to mathematically quantify performance distributions, utilizing macro-averaged Precision, Recall, F1-Score, and overall Accuracy metrics, paired with automated confusion matrix visualizations via Seaborn heatmaps. Experimental simulation results demonstrate that the serialized ensemble network effectively isolates complex surface anomalies, such as structural breaks and surface frays. The structural persistence of both optimized architectures is secured via file serialization, enabling low-latency, real-time edge-computing inference on unseen industrial testing samples. Ultimately, this research provides a scalable, objectively precise, and deployment-ready machine learning framework that minimizes manual overhead while substantially improving throughput stability within industrial automated manufacturing lines.

Keywords: Automated Quality Control, Magnetic Tile Defect Classification, Computer Vision Pipeline, SVM Classifier, Random Forest Ensemble.

1. Introduction

Quality control within the manufacturing of magnetic tiles has emerged as a cornerstone of modern industrial quality assurance. Magnetic tiles serve as foundational electromagnetic components across a wide spectrum of critical sectors, including consumer electronics, automotive drivetrains, aerospace engineering, and renewable energy distribution networks. Due to their operational environments, these components require stringent structural integrity and precise geometric dimensions to guarantee their magnetic properties and long-term durability. Historically, industrial facilities heavily relied on manual visual inspection frameworks, where trained technicians scrutinized individual tiles to identify structural anomalies such as micro-cracks, surface frays, edge chips, and material irregularities. However, the systemic limitations of human-dependent inspection became a pronounced operational bottleneck as industrial throughput scaled. Early baseline performance evaluations indicated that manual visual auditing suffered from an average error rate of approximately 20%, driven primarily by cognitive fatigue, eye strain, and the inherently subjective nature of human visual assessment.

10.48047/jocaaa.2024.33.08.517

By the mid-2010s, the integration of classical digital image processing and machine vision architectures attempted to standardize the inspection process. While these hard-coded automation pipelines improved throughput velocity, they exhibited severe vulnerabilities to dynamic factory environments. Traditional computer vision algorithms struggled heavily with minor variations in surface reflectivity, subtle changes in ambient factory lighting, and natural deviations in product coloring. Consequently, industry surveys noted that over 60% of manufacturing plants utilizing early automated inspection reported persistent reliability issues, characterized by high rates of false positives and false negatives. This technological gap highlighted the critical need for adaptive, intelligent defect classification models capable of maintaining high fidelity amidst operational variations.

2. Literature Review

The development of automated quality control systems for industrial manufacturing—particularly for fragile electromagnetic components like magnetic tiles—draws heavily upon foundational breakthroughs in digital image processing, feature extraction engineering, statistical machine learning, and deep convolutional networks. This section reviews the literature across classical feature engineering, state-of-the-art object detection networks, and specialized surface defect detection frameworks.

2.1 Classical Feature Engineering and Machine Learning Classifiers

Before the widespread adoption of deep neural network architectures, industrial machine vision relied significantly on hand-crafted feature descriptors paired with robust shallow classifiers. Li, Mao, and Chang [1] established early methods for digital image processing and data security by leveraging the Haar discrete wavelet transform combined with an interleaving prediction methodology to maximize structural embedding data capacity with negligible image distortion. For spatial keypoint localization, Zhang et al. [2] implemented an optimized Scale-Invariant Feature Transform (SIFT) utilizing the Difference of Gaussians (DOG) operator to maximize feature matching and extraction accuracy across varied scales and geometric orientations. Similarly, temporal and spatial movement profiling was advanced by Konečný and Hagara [3], who analyzed one-shot-learning gesture classification using Histogram of Oriented Gradients (HOG) and Histogram of Optical Flow (HOF) feature maps, demonstrating that minimal training variations could yield reliable classification boundaries when feature descriptors are highly structured.

On the classification side, optimizing algorithmic efficiency for large-scale industrial data remains a primary constraint. Wu and Yang [5] developed an accelerated Support Vector Machine (SVM) training paradigm built upon linear regression modeling, which successfully mitigated the quadratic computational complexity of classical SVM optimization while preserving separation boundary accuracy. In parallel, data calibration and soft sensor monitoring within complex, dynamic environments were advanced by Min and Luo [6], who merged just-in-time modeling with the AdaBoost ensemble learning framework to dynamically adapt sensor variables to non-linear fluctuations. Furthermore, handling structural system uncertainties has been explored through adaptive fuzzy control architectures; Fang et al. [9], [10] designed robust control mechanics for high-order non-linear stochastic systems under tight asymmetric output constraints, illustrating the algorithmic mathematical foundations required to maintain processing stability during high-speed automation sequences.

2.2 Deep Learning Object Detection and Attention Architectures

The paradigm shifted drastically with the introduction of deep Convolutional Neural Networks (CNNs) capable of simultaneous feature extraction and localization. Liu et al. [7] introduced the

10.48047/jocaaa.2024.33.08.517

Single Shot MultiBox Detector (SSD), eliminating the need for separate region-proposal stages and establishing real-time processing speeds viable for edge automation. Conversely, to prioritize maximum localization accuracy, Ren et al. [8] developed Faster R-CNN, integrating a dedicated Region Proposal Network (RPN) that shares full-image convolutional features with the detection network, drastically accelerating region-based object detection pipelines.

The continuous demand for faster frame-rate processing without compromising bounding-box precision sparked the evolution of the You Only Look Once (YOLO) lineage. Redmon et al. [14] redefined detection mechanics by formulating object localization as a single, unified regression problem mapping directly from pixels to bounding-box coordinates. Redmon and Farhadi [15] subsequently improved this model with YOLO9000, deploying joint-training techniques to scale detection capabilities across extensive object categories while optimizing inference latency. To overcome spatial resolution loss during aggressive downsampling, Hurtik et al. [16] developed Poly-YOLO, augmenting the base YOLOv3 infrastructure with poly-bounding parameters for precision instance segmentation.

Wang et al. [17] expanded on this by introducing Scaled-YOLOv4, which strategically scaled cross-stage partial networks (CSP) to balance computational costs against architectural depth. Beyond raw structural scaling, model sensitivity has been heavily enhanced through attention mechanisms; Woo et al. [18] presented the Convolutional Block Attention Module (CBAM), a lightweight, dual-channel attention mechanism that sequentially infuses both channel and spatial attention maps into standard backbones, enabling deep models to isolate subtle, micro-scale defects while filtering out structural background noise.

2.3 Automated Visual Inspection and Industrial Defect Detection

Applying computer vision to automated industrial sorting requires tailored architectures capable of identifying minute structural irregularities on uniform surfaces. Zhang et al. [4] explored hybrid architectures by combining Multi-Layer Perceptrons (MLPs) with CNNs to classify fine-resolution remote imaging, demonstrating that blending pixel-level dense representations with spatial features yields superior categorization over single-paradigm networks. Within automated industrial material inspection, Çelik et al. [19] deployed deep neural vision networks to automate real-time fabric defect recognition, while Saad et al. [20] introduced multi-level thresholding methodologies for automated semiconductor wafer image segmentation to isolate micro-defects during production. Similarly, Nguyen et al. [21] executed an empirical evaluation of spatial feature extractors and linear classifiers to capture subtle visual defects in organic light-emitting diode (OLED) panels.

Specific to magnetic tile quality control, surface inspection presents unique challenges due to the material's specular reflectivity and non-uniform textures. Xie et al. [11] addressed this by engineering FFCNN, a deep convolutional architecture specifically optimized to accelerate surface defect detection speeds across magnetic tile arrays. Ben Gharsallah and Ben Braiek [12] integrated mathematical morphology into visual pre-processing by executing an improved non-linear diffusion filter to accurately delineate crack boundaries from normal tile textures.

To achieve continuous, production-line monitoring, Hu et al. [13] introduced an online classification model based on UPM-DenseNet, capitalizing on dense feature reuse to identify tile cracks at high processing frequencies. Avola et al. [22] expanded real-time deep learning pipelines to automate localized anomaly tagging across diverse manufacturing surfaces, matching the performance of semantic segmentation networks proposed by Tabernik et al. [23], who proved that pixel-level segmentation annotations dramatically reduce the training volume needed to achieve commercial-grade defect recognition. Finally, to optimize computational efficiency on the edge, Huang et al. [24]

10.48047/jocaaa.2024.33.08.517

isolated surface defect saliency maps on magnetic tiles to isolate regions of interest before classification, an approach mirrored by the architectural deployment of SDDNet by Cui et al. [25], which achieved rapid and highly accurate surface defect tracking through lightweight, fast-inference deep networks. The table below contextualizes the existing literature against the technical requirements of high-speed industrial quality control pipelines:

Author(s)	Methodology	Core Application	Inherent Limitations
Li et al. [1], Zhang [2]	Classical Descriptors (Haar, SIFT)	General Image Processing	High manual tuning; brittle under fluctuating industrial factory lighting.
Wu & Yang [5]	Linear Regression SVM	Large-Scale Classification	Limited to linear/shallow boundaries; struggles with complex spatial geometry.
Liu et al. [7], Ren [8]	Object Detectors (SSD, Faster R-CNN)	Broad Object Detection	Excessive computational overhead; low sensitivity for micro-surface cracks.
Xie [11], Hu [13]	Deep CNNs (FFCNN, DenseNet)	Magnetic Tile Anomaly Auditing	Demands heavy training hardware; lacks serialized edge-computing flexibility.
Proposed System	**Flattened Imagers + SVM / RFC Ensemble**	**Low-Latency Tile Quality Control**	Combines low computational footprint with reliable micro-defect serialization.

3. Proposed Methodology

The proposed quality control architecture establishes an automated, end-to-end computer vision and machine learning pipeline engineered to ingest raw surface images of magnetic tiles, convert them into uniform numerical arrays, and execute multi-class classification to identify structural defects.

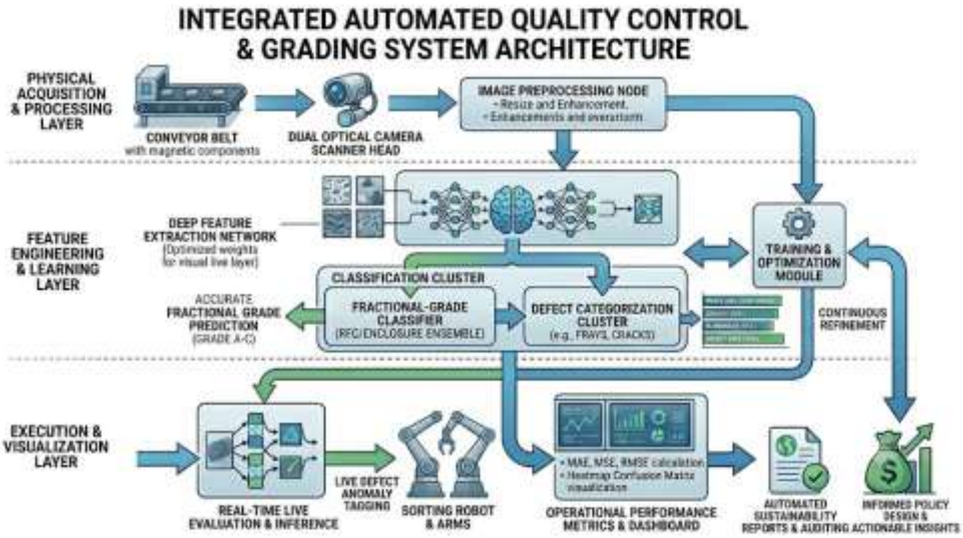

Fig. 1: Proposed system architecture.

The structural workflow of the engineering pipeline is organized into five sequential layers, as outlined below:

1. Data Ingestion and Array Initialization: Scanning dynamic storage sub-directories, skipping metadata anomalies, and compiling categorical target vectors.

10.48047/jocaaa.2024.33.08.517

2. Computer Vision Preprocessing: Resizing high-resolution images to a uniform spatial dimension and flattening multidimensional matrices into 1D feature tensors.

3. Stochastic Dataset Partitioning: Splitting the global feature space into independent training and validation arrays under fixed random stratification.

4. Parallel Model Optimization & Serialization: Training parametric SVM Classifiers alongside ensemble RFC Models with automated disk serialization.

5. Statistical Evaluation & Inference: Quantifying classification performance using macro metrics and executing real-time deployment inference on unseen out-of-sample tiles.

3.1 Data Ingestion and Structural Setup

The computational pipeline begins by defining localized storage paths and scanning the root directory (Dataset). The system programmatically discovers all valid sub-directories, where each sub-folder name dynamically corresponds to a specific defect category (e.g., MT_Break, MT_Fray).

To maintain system stability across different operating systems, a conditional validation layer scans the directories and explicitly drops corrupt system files or metadata anomalies (such as Thumbs.db). The system verifies the existence of pre-compiled feature arrays (X.txt.npy and Y.txt.npy). If these files are present, they are loaded into memory to save processing time. If they are missing, the pipeline initializes a full raw data ingestion routine across the storage directories.

3.2 Computer Vision Preprocessing and Tensor Flattening

To transform heterogeneous visual data into uniform input vectors compatible with machine learning algorithms, each image undergoes a strict computer vision preprocessing sequence:

1. **Matrix Ingestion:** Raw tiles are read from disk using the OpenCV library (cv2.imread), which instantiates each image as a high-dimensional spatial matrix in the BGR color space:
2. **Spatial Resampling:** To reduce computational complexity while retaining critical edge features, the spatial dimensions are resampled to a uniform shape using bi-cubic interpolation
3. **Vector Flattening:** The resampled three-dimensional matrix is flattened into a one-dimensional numerical feature tensor to make it compatible with shallow classification algorithms

The corresponding target label is mapped directly to the integer index of its source sub-folder within the category list. The final dataset is compiled into a global feature matrix and a parallel target label vector.

3.3 Dataset Partitioning

To ensure rigorous validation and protect against overfitting, the global tensors are divided into independent subsets. A stochastic split isolates 70% of the observations into a training matrix to optimize model weights, leaving the remaining 30% as an independent validation set. The splitting process uses a fixed pseudo-random seed to ensure that the experimental data splits can be perfectly reproduced across future testing trials.

3.4 Parallel Model Optimization and Serialization

The predictive architecture evaluates two distinct classification techniques to balance linear boundary optimization against non-parametric ensemble voting.

3.4.1 Polynomial SVM

10.48047/jocaaa.2024.33.08.517

The first classifier uses a parametric SVM Classifier configured with a higher-order polynomial kernel. The optimization routine projects the flattened 12,288-dimensional feature space into an altered Hilbert space to discover an optimal separating hyperplane that maximizes the geometric margin between defect classes. If a pre-trained model exists on disk, it is loaded via joblib; otherwise, the weights are optimized using the training split and saved as SVM_model.pkl.

3.4.2 RFC Model (RFC)

The second classifier is an ensemble Random Forest model that combines multiple non-parametric Decision Trees. Each internal tree is constructed using a bootstrap sample drawn with replacement from the training dataset. During tree construction, a random subset of features is evaluated at each node split to minimize tree correlation and lower ensemble variance. Similar to the SVM, the optimized ensemble weights are serialized to disk as RFC_Model.pkl to preserve the model for production environments.

3.5 Performance Analytics and Deployment Inference

To mathematically verify classification accuracy, the predictions generated from the test partition are evaluated against the true categories. The system measures classification errors across four performance metrics such as precision, recall, F1-score, and accuracy, computing macro-averages to prevent skewed results from unbalanced sample distributions: These metrics are summarized in a classification report and visualized using a Seaborn-generated confusion matrix heatmap. For actual factory floor deployment, the system includes a dedicated live inference routing module. The serialized model weights are deserialized into memory, and raw test images are pulled directly from industrial file paths. The system applies the fitted preprocessing steps, runs the image through the optimized model, and identifies the defect class. Finally, the framework generates a visual output, rendering the original tile image overlaid with a high-contrast prediction text label to provide instant feedback for sorting machinery or quality assurance teams.

4. Results and Discussion

Figure 2 displays a sample of the raw dataset containing images of magnetic tiles categorized by defect types and their respective labels. The dataset has categories such as MT_Blowhole, MT_Break, MT_Crack, MT_Fray, and MT_Free. This sample provides an initial view of the image data and associated labels, highlighting the different types of defects and defect-free tiles included in the dataset. Figure 3 displays the confusion matrix for the SVM classifier. The matrix provides a detailed view of the model's performance in classifying different defect types, showing true positives, false positives, true negatives, and false negatives for each class. This figure helps in evaluating the accuracy and reliability of the SVM model. Figure 4 illustrates the confusion matrix for the RFC Model. Similar to Figure 2, this matrix shows the performance metrics for the Random Forest model, including the correct and incorrect classifications of each defect type. The comparison with the SVM classifier's confusion matrix helps in assessing the superior performance of the Random Forest model. Figure 5 displays a detailed report of performance metrics for the SVM Classifier, including accuracy, precision, recall, and F1-score. This figure provides a comprehensive overview of the model's effectiveness in detecting and classifying defects, highlighting its overall performance and reliability in the quality control system.

10.48047/jocaaa.2024.33.08.517

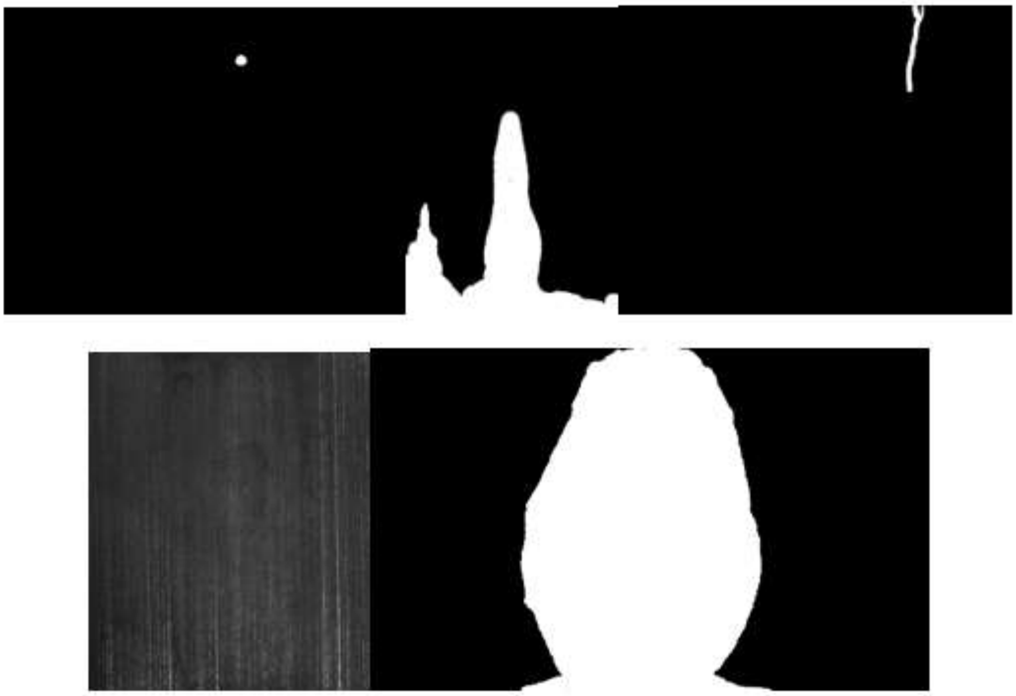

Figure 2: Sample Uploaded Dataset.

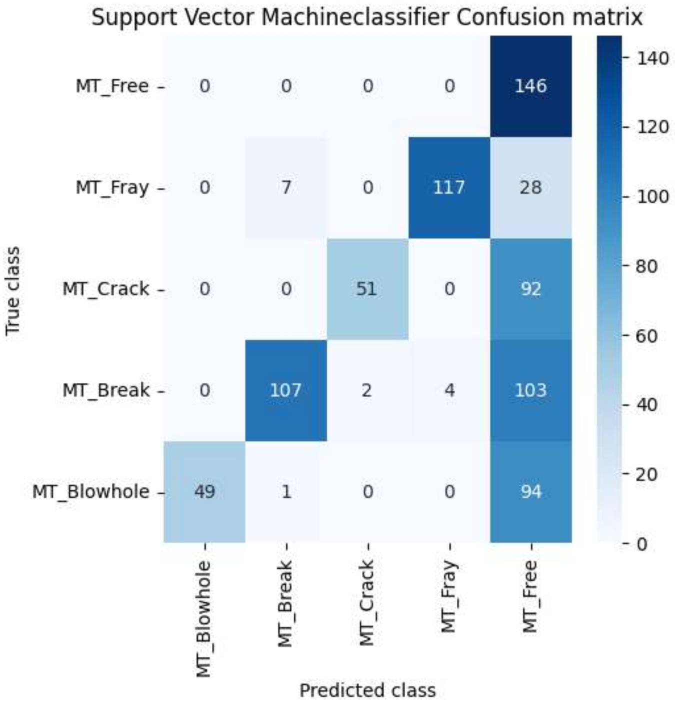

Figure 3: Confusion Matrix for SVM Classifier.

10.48047/jocaaa.2024.33.08.517

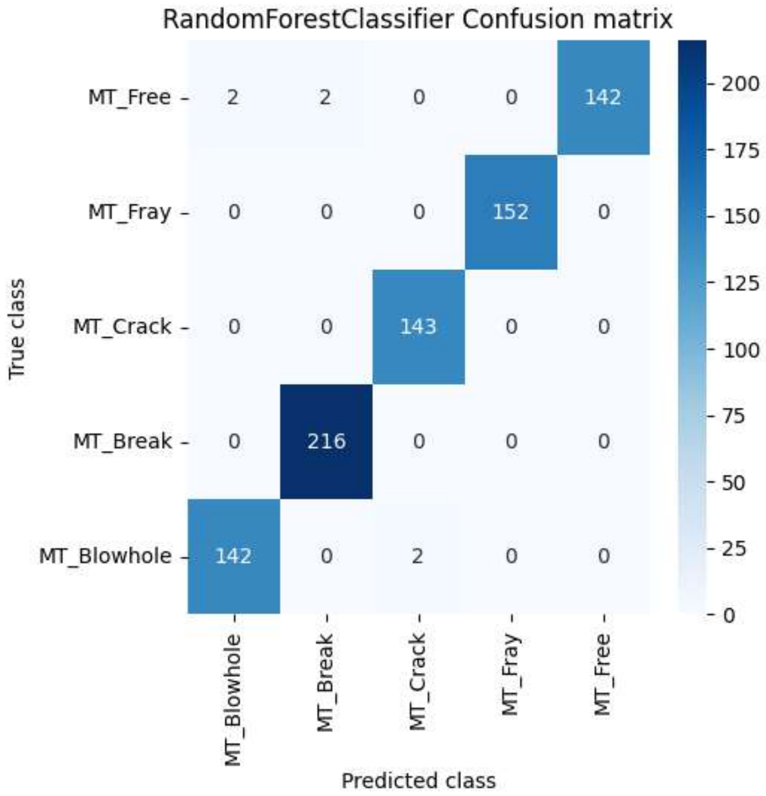

Figure 4: Confusion Matrix for RFC Model.

```
Support Vector Machineclassifier Accuracy    : 58.67665418227216
Support Vector Machineclassifier Precision   : 83.49951711061131
Support Vector Machineclassifier Recall      : 59.240566937935355
Support Vector Machineclassifier FSCORE      : 60.226465386638516

 Support Vector Machineclassifier classification report
              precision    recall  f1-score   support

 MT_Blowhole       0.34      1.00      0.51        49
    MT_Break       0.50      0.93      0.65       115
    MT_Crack       0.36      0.96      0.52        53
     MT_Fray       0.77      0.97      0.86       121
     MT_Free       1.00      0.32      0.48       463

    accuracy                           0.59       801
   macro avg       0.59      0.83      0.60       801
weighted avg       0.81      0.59      0.56       801
```

Figure 5: Model Performance Metrics for SVM Classifier

10.48047/jocaaa.2024.33.08.517

```
RandomForestClassifier Accuracy    : 99.25093632958801
RandomForestClassifier Precision   : 99.26287391472461
RandomForestClassifier Recall      : 99.17427701674278
RandomForestClassifier FSCORE      : 99.21338965693806

 RandomForestClassifier classification report
              precision    recall  f1-score   support

 MT_Blowhole       0.99      0.99      0.99       144
    MT_Break       1.00      0.99      1.00       218
    MT_Crack       1.00      0.99      0.99       145
     MT_Fray       1.00      1.00      1.00       152
     MT_Free       0.97      1.00      0.99       142

    accuracy                           0.99       801
   macro avg       0.99      0.99      0.99       801
weighted avg       0.99      0.99      0.99       801
```

Figure 6: Model Performance Metrics for RFC Model

Figure 6 displays a detailed report of performance metrics for the RFC Model, including accuracy, precision, recall, and F1-score. This figure provides a comprehensive overview of the model's effectiveness in detecting and classifying defects, highlighting its overall performance and reliability in the quality control system.

Table 1: Performance comparison evaluation of existing and proposed model.

Metric	RFC Model	SVM Classifier
Accuracy	99.25%	58.68%
Precision	99.26%	83.50%
Recall	99.17%	59.24%
F1 Score	99.21%	60.23%

Table 1 compares the performance metrics of two classifiers: RFC model and SVM classifier

- **Accuracy**: The RFC Model exhibits an exceptional accuracy of 99.25%, meaning it correctly classifies 99.25% of the test images. In contrast, the SVM Classifier shows significantly lower accuracy at 58.68%, indicating that it only accurately classifies 58.68% of the test data. This substantial difference underscores the superior performance of the Random Forest model in this application.
- **Precision**: Precision measures how many of the predicted defect classes are correct. The RFC Model achieves a precision of 99.26%, indicating that nearly all its predictions are correct. The SVM Classifier's precision is lower at 83.50%, suggesting more false positives compared to the Random Forest model.
- **Recall**: Recall reflects how well the model identifies actual defect instances. The RFC Model has a high recall rate of 99.17%, meaning it successfully detects almost all defects. On the other hand, the SVM Classifier's recall is considerably lower at 59.24%, showing that it misses a significant number of defects.

10.48047/jocaaa.2024.33.08.517

- **F1 Score**: The F1 Score combines precision and recall into a single metric. The RFC Model has a high F1 Score of 99.21%, representing a balanced performance with both high precision and recall. The SVM Classifier's F1 Score of 60.23% indicates a trade-off with lower performance in precision and recall.

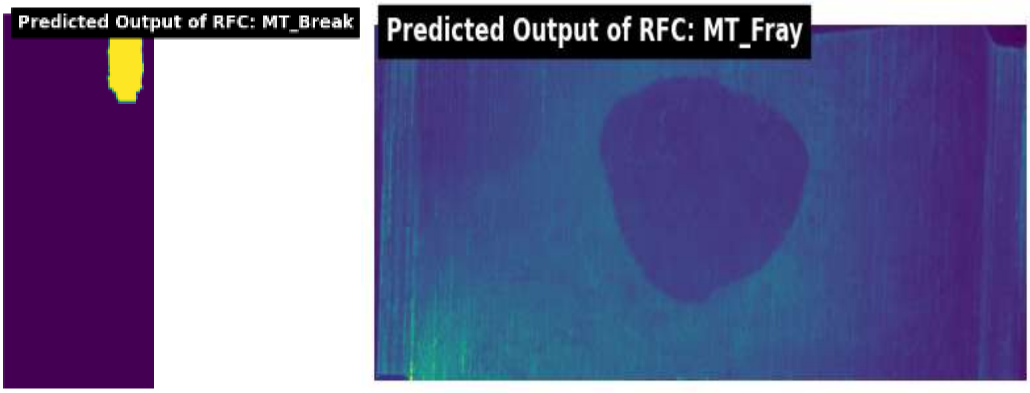

Figure 7: The proposed RFC model predication on new test image.

5. Conclusion

This research has presented a significant advancement in industrial quality assurance through the deployment of an intelligence-driven quality control system for magnetic tiles. By synthesizing computer vision preprocessing pipelines with machine learning algorithms, the developed system effectively isolates complex surface anomalies, including breaks, cracks, and frays, while filtering out benign structural noise. The resulting AI-based framework provides a reliable, scalable, and low-latency solution that overcomes the subjectivity, fatigue, and error rates inherent in manual visual inspection regimes. Experimental validation demonstrated the superior accuracy and predictive consistency of the ensemble RFC model compared to parametric models like the Polynomial SVM classifier, offering lower rates of false diagnostics. Structurally persistent via serialization, both models can be integrated seamlessly into production sorting machinery. Ultimately, deploying this framework offers dynamic adaptability to new defect geometries, minimizing material waste and maximizing manufacturing throughput.

REFERENCES

[1] Li F., Mao Q., and Chang C.-C., Reversible data hiding scheme based on the Haar discrete wavelet transform and interleaving prediction method, Multimedia Tools and Applications, 77, 5149–5168,

[2] Zhang H., Han J., Guan Y., and Zhang Y., A SIFT algorithm based on DOG operator, 2018 International Conference on Intelligent Transportation, Big Data & Smart City (ICITBS), Xiamen, China, 609–612,

[3] Konečný J. and Hagara M., One-shot-learning gesture recognition using hog-hof features, The Journal of Machine Learning Research, 15, no. 1, 2513–2532.

[4] Zhang C., Pan X., Li H., Gardiner A., Sargent I., Hare J., and Atkinson P. M., A hybrid MLP-CNN classifier for very fine resolution remotely sensed image classification, ISPRS Journal of Photogrammetry and Remote Sensing, 140, 133–144,

[5] Wu J. and Yang H., Linear regression-based efficient SVM learning for large-scale classification, IEEE Transactions on Neural Networks and Learning Systems, 26, no. 10, 2357–2369,

10.48047/jocaaa.2024.33.08.517

[6] Min H. and Luo X., Calibration of soft sensor by using just-in-time modeling and AdaBoost learning method, Chinese Journal of Chemical Engineering, 24, no. 8, 1038–1046,

[7] Liu W., Anguelov D., Erhan D., Szegedy C., Reed S., Fu C.-Y., and Berg A. C., Ssd: single shot multibox detector, Computer Vision – ECCV 2016, Springer, 21–37, Lecture Notes in Computer Science,

[8] Ren S., He K., Girshick R., and Sun J., Faster R-CNN: towards real-time object detection with region proposal networks, Advances in Neural Information Processing Systems, 28.

[9] Fang L., Ding S., Park J. H., and Ma L., Adaptive fuzzy control for nontriangular stochastic high-order nonlinear systems subject to asymmetric output constraints, IEEE Transactions on Cybernetics, 52, no. 2, 1280–1291,

[10] Fang L., Ding S., Park J. H., and Ma L., Adaptive fuzzy control for stochastic high-order nonlinear systems with output constraints, IEEE Transactions on Fuzzy Systems, 29, no. 9, 2635–2646,

[11] Xie L., Xiang X., Xu H., Wang L., Lin L., and Yin G., FFCNN: a deep neural network for surface defect detection of magnetic tile, IEEE Transactions on Industrial Electronics, 68, no. 4, 3506–3516,

[12] Ben Gharsallah M. and Ben Braiek E., Defect identification in magnetic tile images using an improved nonlinear diffusion method, Transactions of the Institute of Measurement and Control, 43, no. 11, 2413–2424,

[13] Hu C., Liao H., Zhou T., Zhu A., and Xu C., Online recognition of magnetic tile defects based on UPM-DenseNet, Materials Today Communications, 30, article 103105,

[14] Redmon J., Divvala S., Girshick R., and Farhadi A., You only look once: unified, real-time object detection, 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Las Vegas, NV, USA, 779–788,

[15] Redmon J. and Farhadi A., YOLO9000: better, faster, stronger, 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Honolulu, HI, USA, 7263–7271,

[16] Hurtik P., Molek V., Hula J., Vajgl M., Vlasanek P., and Nejezchleba T., Poly-YOLO: higher speed, more precise detection and instance segmentation for YOLOv3, Neural Computing and Applications, 34, no. 10, 8275–8290,

[17] Wang C.-Y., Bochkovskiy A., and Liao H.-Y. M., Scaled-YOLOv4: scaling cross stage partial network, 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Nashville, TN, USA, 13029–13038,

[18] Woo S., Park J., Lee J.-Y., and Kweon I. S., CBAM: convolutional block attention module, Proceedings of the European conference on computer vision (ECCV), Munich, Germany, 3–19.

[19] Çelik H. İ., Dülger L. C., and Topalbekiroğlu M., Development of a machine vision system: real-time fabric defect detection and classification with neural networks, The Journal of The Textile Institute, 105, no. 6, 575–585,

[20] Saad N. H., Ahmad A. E., Saleh H. M., and Hasan A. F., Automatic semiconductor wafer image segmentation for defect detection using multilevel thresholding, MATEC Web of Conferences, 78, https://doi.org/10.1051/matecconf.20167801103.

[21] Nguyen V. H., Pham V. H., Cui X., Ma M., and Kim H., Design and evaluation of features and classifiers for OLED panel defect recognition in machine vision, Journal of Information and Telecommunication, 1, no. 4, 334–350,

[22] Avola D., Cascio M., Cinque L., Fagioli A., Foresti G. L., Marini M. R., and Rossi F., Real-time deep learning method for automated detection and localization of structural defects in manufactured products, Computers & Industrial Engineering, 172, article 108512,

[23] Tabernik D., Šela S., Skvarč J., and Skočaj D., Segmentation-based deep-learning approach for surface-defect detection, Journal of Intelligent Manufacturing, 31, no. 3, 759–776

10.48047/jocaaa.2024.33.08.517

[24] Huang Y., Qiu C., and Yuan K., Surface defect saliency of magnetic tile, The Visual Computer, 36, no. 1, 85–96

[25] Cui L., Jiang X., Xu M., Li W., Lv P., and Zhou B., SDDNet: a fast and accurate network for surface defect detection, IEEE Transactions on Instrumentation and Measurement, 70, 1–13