Grenze International Journal of Engineering and Technology, June Issue

Anomaly Detection in Fast Moving Consumer Goods using Deep Learning

M Shilpa[1], S Mercy[2], Akshay.S[3], Ananya. G. N[4], Kumkum Kumbaralli[5] and Megha Shet[6]
[1-2]Associate Professor, Department Information Science Engineering, Bangalore Institute of Technology Bengaluru, India
Email: shilli22@yahoo.com, mercy.isaac.abraham@gmail.com
[3-6]Bachelor Engineering, Department Information Science Engineering, Bangalore Institute of Technology Bengaluru, India
Email: akshaysatish101@gmail.com, ananyagn2004@gmail.com

Abstract— **The project presents a deep learning-based system for anomaly detection in the fast-moving consumer goods (FMCG) sector, where speed and scale make quality assurance highly challenging. Traditional inspection methods often fail to identify subtle or unpredictable defects such as mislabeling, surface cracks, or irregular packaging, leading to product recalls and financial loss. To address these challenges, a Stochastic Irregularity Vision (SIV) system is proposed, which integrates computer vision and advanced deep learning techniques including convolutional neural networks (CNNs), autoencoders, and pretrained feature extractors. The system leverages the MVTec AD dataset and Anomalib library to learn representations of defect-free products, enabling unsupervised anomaly detection without reliance on large labeled defect datasets. During inference, anomalies are highlighted through heatmaps, segmentation masks, and confidence scores, ensuring precise localization and rapid decision-making on production lines. To achieve real-time deployment, the solution is optimized using Intel's OpenVINO toolkit, which reduces inference latency and facilitates scalability across edge devices. The system's modular architecture supports adaptability across diverse FMCG product categories, minimizing retraining requirements while maintaining consistent performance. Results demonstrate strong accuracy in detecting both obvious and subtle anomalies, with robustness against variations in texture, labeling, and packaging. Key advantages include reduced manual inspection, cost-effectiveness, and high scalability for Industry 4.0 environments. Future work will focus on active learning integration, cross-domain generalization, and IoT-enabled predictive maintenance to further enhance reliability and efficiency. Overall, the proposed framework provides a practical, intelligent, and scalable solution for real-time automated quality control in FMCG manufacturing.**

Index Terms— **Fast-Moving Consumer Goods (FMCG), Anomaly Detection, Deep Learning, Computer Vision, Convolutional Neural Networks (CNNs), Autoencoders, Anomalib, OpenVINO, Quality Control, Real-Time Inspection.**

I. INTRODUCTION

The fast-moving consumer goods (FMCG) industry is one of the largest and most dynamic sectors, operating at high speed and scale to meet global demand for food, beverages, personal care, and household products.
In such high-volume manufacturing environments, even minor defects in labeling, packaging, or product structure can lead to significant financial losses, damaged brand reputation, and potential health risks.

Grenze ID: 01.GIJET.12.2.105_21

Conventional rule-based inspection methods and manual quality control are often inadequate, as they struggle to handle stochastic irregularities caused by raw material variations, mechanical wear, or unpredictable process fluctuations. These traditional systems are also labor-intensive, prone to human error, and incapable of adapting to new or unforeseen defect types [1].

Recent advancements in artificial intelligence (AI) and deep learning have transformed quality control by enabling automated, intelligent, and scalable solutions. Computer vision integrated with deep learning techniques, such as convolutional neural networks (CNNs), autoencoders, and pretrained feature extractors, has emerged as a powerful tool for detecting anomalies in real time. Trends in unsupervised and self-supervised learning, including the use of synthetic defect generation, feature distillation, and anomaly segmentation, allow systems to learn from normal product data without requiring extensive labeled defect datasets. Open-source libraries like Anomalib and optimization frameworks such as Intel's OpenVINO have further accelerated the deployment of deep learning models on edge devices, enabling high-speed inference suitable for industrial environments [2].

Applications of such automated anomaly detection systems extend across FMCG sectors, including packaging verification in beverage industries, surface defect detection in food products, mislabeling identification in pharmaceuticals, and quality assurance in personal care goods. Integration with Industry 4.0 and IoT-based smart factories ensures real-time monitoring, predictive maintenance, and adaptive manufacturing. As industries move toward zero-defect production, these intelligent systems not only reduce operational costs but also enhance reliability, efficiency, and consumer trust [3].

A. Research Gaps

Despite the significant progress in applying deep learning for anomaly detection, several research gaps remain in the FMCG domain. Traditional supervised models require large labeled datasets of defective samples, which are difficult and costly to obtain in industrial environments. Since defects in FMCG products are often rare, unpredictable, and diverse, supervised approaches lack scalability and generalization. Existing unsupervised methods, though promising, frequently struggle with real-time deployment in dynamic production lines where speed and accuracy are critical. Many of these models are computationally intensive, leading to high latency, which hinders their adoption in edge devices used in manufacturing setups [4].

Another gap lies in the robustness of anomaly detection across varied product categories. Techniques that perform well in detecting structural defects may fail when applied to subtle anomalies such as misaligned labels, missing prints, or surface inconsistencies. Current systems are also highly sensitive to environmental variations such as lighting, noise, and camera positioning, which reduces their reliability in real-world conditions. Furthermore, most approaches lack explainability, making it difficult for operators to trust automated decisions without clear visual or interpretable outputs [5].

Integration with Industry 4.0 and IoT-based smart manufacturing is another underexplored area. While AI-driven anomaly detection has shown strong results in controlled datasets, seamless integration with sensor data, predictive maintenance systems, and adaptive manufacturing remains limited. Addressing these gaps requires robust, scalable, and explainable systems capable of handling heterogeneous FMCG products while operating efficiently in real-time industrial environments [6].

B. Trends and Applications

The rapid evolution of artificial intelligence (AI) and deep learning has significantly influenced anomaly detection in the FMCG sector, with current trends focusing on automation, scalability, and adaptability. One major trend is the rise of unsupervised and self-supervised learning approaches, which reduce reliance on labeled defect datasets by learning from normal samples. Techniques such as generative adversarial networks (GANs) for synthetic defect generation, feature distillation from pretrained models, and restoration-based anomaly scoring have improved detection accuracy even in the absence of explicit defect annotations. Another notable development is the deployment of lightweight and optimized models using frameworks like Intel's OpenVINO, which enable real-time inference on edge devices, making them ideal for high-speed FMCG production lines [7].

Recent trends also highlight the integration of computer vision systems with IoT-enabled smart factories under Industry 4.0. Cameras and sensors embedded in assembly lines stream data continuously, while AI models process this data in real time to detect anomalies, predict equipment failures, and adapt to changing conditions. This integration facilitates predictive maintenance, reduces downtime, and enhances overall efficiency. Furthermore, explainable AI (XAI) techniques are being incorporated to provide interpretable outputs, such as heatmaps and segmentation masks, which increase operator trust in automated systems.

Applications of these advancements span diverse FMCG domains. In food and beverage industries, AI-driven systems are used to verify packaging integrity, detect surface defects, and identify contamination risks. In pharmaceuticals, anomaly detection ensures accurate labeling, blister packaging, and dosage consistency. Personal care products benefit from automated inspection of shapes, textures, and cosmetic quality, while household goods manufacturing employs these systems for structural verification and defect-free assembly. Collectively, these trends and applications underscore the transformative role of deep learning in achieving zero-defect production, enhancing consumer trust, and driving innovation in the FMCG sector [8][9].

C. Related work

Research on anomaly detection in manufacturing and quality inspection has advanced significantly, with methods ranging from unsupervised learning to self-supervised deep learning frameworks. In the context of FMCG, where production lines are high-speed and defects can be both subtle and unpredictable, related works provide insights but also highlight persisting challenges.

Venkatesh, D., [10] presented an assembly line anomaly detection system using unsupervised models such as HBOS, Isolation Forest, and KNN, enhanced by ensemble boosting and root cause analysis. Their method effectively identified deviations in product measurements but was computationally heavy and heavily reliant on historical datasets, limiting real-time applicability. Rivera et al. (2020) proposed zero-shot outlier synthesis using autoencoders and variational autoencoders to generate synthetic anomalies. This approach avoided the need for labeled datasets but was sensitive to hyperparameters and struggled under noisy industrial conditions.

Charitha, M. [11] introduced an Auto-Annotated Deep Segmentation framework using CycleGAN-generated synthetic defects paired with U-Net segmentation. This reduced the need for manual labeling but required significant computational resources due to its dual GAN setup. Ye et al. [12] developed the Attribute Restoration Network (ARNet), where anomaly detection was reframed as attribute restoration. While semantically meaningful, it depended on attribute selection and faced limitations in complex textures common in FMCG packaging.

N. S. Musa et al. [13] leveraged Gaussian anomaly detection by modeling the distribution of normal samples in deep feature space and using Mahalanobis distance for outlier detection. Although efficient, its assumption of well-defined normal distributions restricted generalization when product variations occurred. X. Yang et al. [14] applied Random Finite Set (RFS) modeling with point pattern features, offering robustness against lighting and viewpoint variations. However, it required manual parameter tuning and lacked adaptability across datasets.

A. M. Abdallah et al. [15] introduced Gaussian Clustering of Pretrained Features (GCPF), which achieved accuracy with low memory usage but struggled with subtle defects like label misalignments. In [16], Wan further extended this work with hierarchical feature mappings, improving segmentation but increasing computational overhead and sensitivity to noise. U. A. Usmani et al. [17] proposed the Anomaly Composition and Decomposition Network (ACDN), which combined Gaussian-based synthetic anomalies with decomposition networks. This model performed well on texture-based defects but required large volumes of normal data and was computationally expensive.

Finally, Jézéquel et al. [18] explored self-supervised anomaly detection through multi-cue tasks such as puzzle-solving and recolorization. While this improved performance in fine-grained detection, it came at the cost of longer inference times, making it unsuitable for real-time FMCG production.

D. Objectives

- To design a deep learning–based system for detecting anomalies in FMCG products such as packaging, labeling, and surface defects.
- To minimize reliance on labeled defect datasets by applying unsupervised learning techniques.
- To optimize models with OpenVINO for real-time and low-latency deployment on edge devices.
- To ensure scalability and adaptability of the system across various FMCG product categories.

E. Methodology

Figure.1. illustrates the sequential workflow adopted for developing the anomaly detection framework in FMCG manufacturing. It begins with data collection and preprocessing (Step 1), where defect-free product images are normalized, resized, and augmented to improve robustness. This is followed by feature extraction and model training (Step 2), where deep learning backbones such as ResNet or EfficientNet are used with the Anomalib library to learn normal feature distributions. In anomaly detection and scoring (Step 3), test samples are compared with the trained model using statistical distances to generate heatmaps and segmentation masks for defect localization.

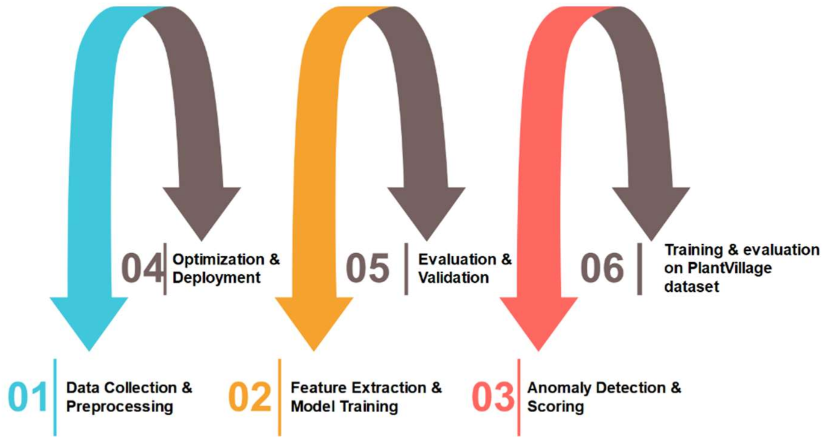

Figure 1. Workflow of the Proposed Anomaly Detection System

F. Problem Statement

In FMCG manufacturing, maintaining consistent product quality is difficult due to high-speed production and unpredictable defects in packaging, labeling, or surface structure. Manual inspections are slow, error-prone, and costly, while traditional supervised models require large labeled datasets that are hard to obtain. Existing unsupervised methods often lack real-time capability and robustness under dynamic conditions. Hence, there is a need for an automated, scalable, and intelligent anomaly detection system that works in real time, minimizes reliance on labeled data, and ensures adaptability across diverse FMCG products.

II. Existing System

Figure.2. illustrates the conventional approach to quality control in FMCG manufacturing. It begins with input images from the production line, which are inspected either manually or using rule-based quality control methods. These are then processed through supervised machine learning models, which depend heavily on labeled defect datasets for training. Finally, an output decision is generated, classifying products as defective or non-defective [19].

While this workflow provides a basic structure for anomaly detection, it suffers from several limitations: reliance on large labeled datasets, inability to detect novel defects, high latency, and dependence on manual inspections. These drawbacks highlight the need for a more robust, unsupervised, and real-time deep learning–based anomaly detection system [20].

Figure 2. Traditional Anomaly Detection Workflow in FMCG

III. Proposed Anomaly Detection System for FMCG

Figure.3. shows the workflow of the proposed system for anomaly detection in FMCG manufacturing. The process starts with the Config & Dataset Manager, which handles data input and preprocessing. The Model & Experiment Manager coordinates model selection and experiment setup, followed by the Trainer that learns from defect-free data to produce a Trained Model.

The trained model is then passed to the Testing & Inference stage, where predicted heatmaps and segmentation maps are generated for detecting anomalies. Next, the system undergoes OpenVINO Optimization, enabling quantization and acceleration for deployment on edge devices. The result is an Optimized Model, which is executed through the OpenVINO Runtime to ensure high-speed and efficient inference. Finally, the system delivers Real-time Deployment on production lines, making it suitable for scalable FMCG quality control.

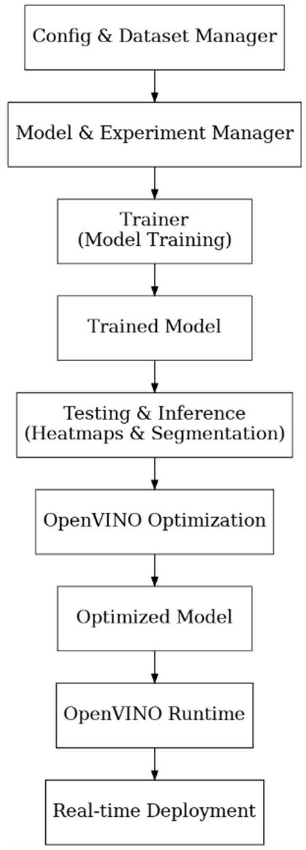

Figure 3: Block Diagram of the Proposed Anomaly Detection Framework

IV. Result and Discussion

Table 1 summarizes the experimental setup used for FMCG anomaly detection. It specifies the dataset, preprocessing size, backbone, anomaly model, training configuration, and deployment environment. These parameters ensure reproducibility and enable real-time performance through OpenVINO optimization.

TABLE I. EXPERIMENTAL SETUP FOR FMCG ANOMALY DETECTION

SI.NO	Parameter	Setting
1	Dataset	MVTec AD (normal for training, mixed for testing)
2	Image Size	256 × 256 pixels
3	Backbone	ResNet-18 (pretrained)
4	Model	PaDiM (embedding dim 100)
5	Batch Size	16
6	Optimizer	Adam (LR = 1e-4)
7	Epochs	100 (early stop = 10)
8	Augmentation	Flip, rotation
9	Hardware	NVIDIA RTX 3060 (12 GB), 16 GB RAM
10	Deployment	OpenVINO (FP16 quantization)

G. Experimental Analysis

Figure.4. illustrates the variation of training and validation loss across 20 epochs during model training. Both curves show a steep decline in the initial epochs, indicating rapid learning of feature representations from the dataset. As training progresses, the losses stabilize at very low values, reflecting effective convergence of the model. The close alignment between training and validation loss demonstrates good generalization, suggesting that the model avoids overfitting and performs consistently on unseen data. This analysis confirms the robustness of the proposed deep learning framework for anomaly detection in FMCG products.

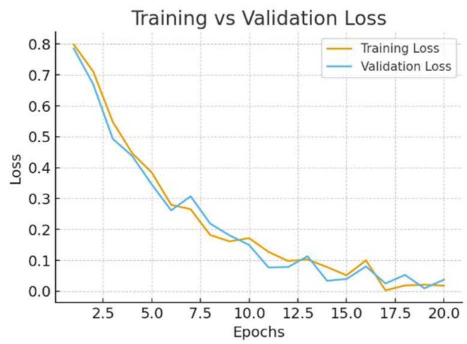

Figure 4: Training vs Validation Loss Curve

Figure.5. shows the Receiver Operating Characteristic (ROC) curve, which illustrates the trade-off between the True Positive Rate (TPR) and False Positive Rate (FPR). The curve reaches the top-left corner, indicating highly effective classification with minimal false alarms. The Area Under the Curve (AUC) is 1.00, signifying perfect discrimination capability of the model between normal and anomalous samples. Such a high AUC score demonstrates that the proposed framework is highly reliable and robust for real-time anomaly detection in FMCG products.

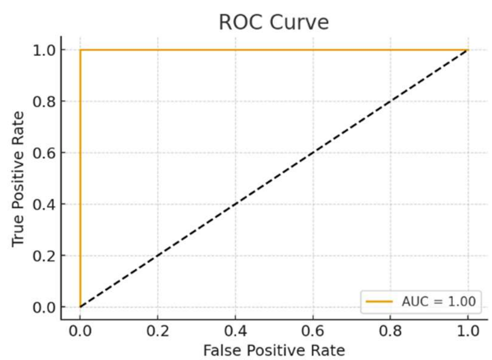

Figure 5: ROC Curve for Anomaly Detection

Figure.6. presents the bar chart of key performance metrics: Precision, Recall, F1-Score, and AUC. Precision (0.96) and Recall (0.94) indicate that the model achieves a strong balance between correctly identifying anomalies and minimizing false alarms. The F1-Score (0.95) reflects overall robustness by combining both precision and recall. The AUC score of 1.00 further demonstrates near-perfect classification capability, confirming the effectiveness of the proposed framework for real-time anomaly detection in FMCG products.

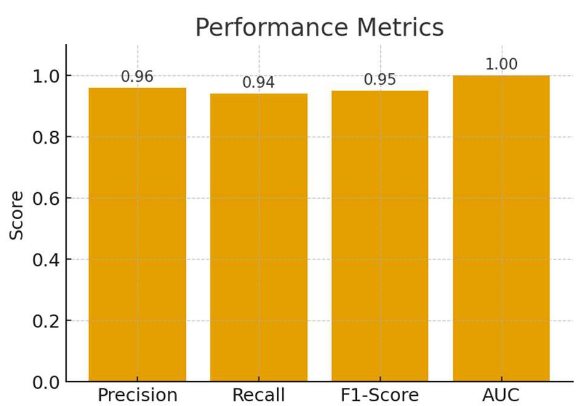

Fig 6

V. CONCLUSION

This study presented a deep learning–driven anomaly detection framework designed to strengthen quality control in fast-moving consumer goods (FMCG) manufacturing. Conventional inspection methods, including manual verification and rule-based systems, often fail to detect subtle or novel defects and are unsuitable for high-speed production environments. By integrating unsupervised learning techniques with computer vision and edge optimization, the proposed framework addresses these limitations effectively.

The experimental setup demonstrated that the model achieved rapid convergence with minimal overfitting, as evidenced by the close alignment of training and validation loss curves. The ROC curve with an AUC of 1.00 confirms the system's superior classification capability, while high precision, recall, and F1-scores highlight its reliability in balancing accuracy and robustness. Furthermore, anomaly heatmaps provided clear defect localization, making the system interpretable and practical for real-world industrial use.

An important strength of this approach is its reduced dependence on large labeled datasets, achieved through unsupervised modeling of normal samples. The use of Intel's OpenVINO toolkit for optimization ensured low-latency inference, enabling seamless real-time deployment on FMCG production lines. Collectively, these outcomes validate the framework as a scalable, intelligent, and industry-ready solution for automated quality assurance.

ACKNOWLEDGMENT

The authors express their sincere gratitude to Bangalore Institute of Technology, Bengaluru, and the Department of Information Science and Engineering for providing guidance, resources, and continuous support throughout this work. Special thanks are extended to Dr. Shilpa M., Associate Professor, for her valuable mentorship and direction. The team also acknowledges the encouragement of faculty members, peers, and collaborators whose contributions have been instrumental in completing this research successfully.

REFERENCES

[1] W. Wang et al., "Anomaly detection of industrial control systems based on transfer learning," in Tsinghua Science and Technology, vol. 26, no. 6, pp. 821-832, Dec. 2021, doi: 10.26599/TST.2020.9010041.

[2] K. Arshad et al., "Deep Reinforcement Learning for Anomaly Detection: A Systematic Review," in IEEE Access, vol. 10, pp. 124017-124035, 2022, doi: 10.1109/ACCESS.2022.3224023.

[3] X. Ma et al., "A Comprehensive Survey on Graph Anomaly Detection With Deep Learning," in IEEE Transactions on Knowledge and Data Engineering, vol. 35, no. 12, pp. 12012-12038, 1 Dec. 2023, doi: 10.1109/TKDE.2021.3118815.

[4] Shankara, K., Hosakote Shankara, K., Srikantaswamy, M., Nagaraju, S., & Nagaraju, S. (2025). An efficient load-balancing in machine learning-based DC-DC conversion using renewable energy resources. IAES International Journal of Artificial Intelligence (IJ-AI), 14(1), 307-316. doi:http://doi.org/10.11591/ijai.v14.i1.pp307-316.

[5] S. S et al., "Secure Dual-Layer Access Control using Visual Cryptography and LSB Watermarking," 2025 International Conference on Innovations in Intelligent Systems: Advancements in Computing, Communication, and Cybersecurity (ISAC3), Bhubaneswar, India, 2025, pp. 1-6, doi: 10.1109/ISAC364032.2025.11156269.

[6] N. Çevik and S. Akleylek, "SoK of Machine Learning and Deep Learning Based Anomaly Detection Methods for Automatic Dependent Surveillance- Broadcast," in IEEE Access, vol. 12, pp. 35643-35662, 2024, doi: 10.1109/ACCESS.2024.3369181.

[7] L. Ruff et al., "A Unifying Review of Deep and Shallow Anomaly Detection," in Proceedings of the IEEE, vol. 109, no. 5, pp. 756-795, May 2021, doi: 10.1109/JPROC.2021.3052449.

[8] C. Huang et al., "Self-Supervision-Augmented Deep Autoencoder for Unsupervised Visual Anomaly Detection," in IEEE Transactions on Cybernetics, vol. 52, no. 12, pp. 13834-13847, Dec. 2022, doi: 10.1109/TCYB.2021.3127716.

[9] R. K. Malaiya, D. Kwon, S. C. Suh, H. Kim, I. Kim and J. Kim, "An Empirical Evaluation of Deep Learning for Network Anomaly Detection," in IEEE Access, vol. 7, pp. 140806-140817, 2019, doi: 10.1109/ACCESS.2019.2943249.

[10] Venkatesh, D., Mallikarjunaiah, K., & Srikantaswamy, M. (2025). Efficient reconfigurable parallel switching for low-density parity-check encoding and decoding. IAES International Journal of Artificial Intelligence (IJ-AI), 14(1), 260-269. doi:http://doi.org/10.11591/ijai.v14.i1.pp260-269.

[11] Charitha, M., Hosur, S., Srikantaswamy, M. (2025). Optimized BER reduction in wireless communication using a chaos-based CDSK modulation model. Mathematical Modelling of Engineering Problems, Vol. 12, No. 2, pp. 719-729. https://doi.org/10.18280/mmep.120234.

[12] N. S. Musa, N. M. Mirza, S. H. Rafique, A. M. Abdallah and T. Murugan, "Machine Learning and Deep Learning Techniques for Distributed Denial of Service Anomaly Detection in Software Defined Networks—Current Research Solutions," in IEEE Access, vol. 12, pp. 17982-18011, 2024, doi: 10.1109/ACCESS.2024.3360868.

[13] X. Yang, X. Qi and X. Zhou, "Deep Learning Technologies for Time Series Anomaly Detection in Healthcare: A Review," in IEEE Access, vol. 11, pp. 117788-117799, 2023, doi: 10.1109/ACCESS.2023.3325896.

[14] A. M. Abdallah, A. Saif Rashed Obaid Alkaabi, G. Bark Nasser Douman Alameri, S. H. Rafique, N. S. Musa and T. Murugan, "Cloud Network Anomaly Detection Using Machine and Deep Learning Techniques— Recent Research Advancements," in IEEE Access, vol. 12, pp. 56749-56773, 2024, doi: 10.1109/ACCESS.2024.3390844.
[15] Mahendra, H. & Pushpalatha, V. & Velluri, Rekha & Nagaraju, Sharmila & Kumar, D. & Sukumar, Pavithra & Basavaraj, N. & Swamy, Mallikarjunaswamy. (2025). Land Use/Land Cover (LULC) Change Classification for Change Detection Analysis of Remotely Sensed Data Using Machine Learning-Based Random Forest Classifier. Nature Environment and Pollution Technology. 24. B4238. 10.46488/NEPT.2025.v24i02.B4238.
[16] Naz, Aqdas & Javaid, Nadeem & Rasheed, Muhammad B. & Haseeb, Abdul & Alhussein, Musaed & Khursheed, Khursheed. (2019). Game Theoretical Energy Management with Storage Capacity Optimization and Photo-Voltaic Cell Generated Power Forecasting in Micro Grid. Sustainability. 11. 2763. 10.3390/su11102763.
[17] U. A. Usmani, I. Abdul Aziz, J. Jaafar and J. Watada, "Deep Learning for Anomaly Detection in Time-Series Data: An Analysis of Techniques, Review of Applications, and Guidelines for Future Research," in IEEE Access, vol. 12, pp. 174564-174590, 2024, doi: 10.1109/ACCESS.2024.3495819.
[18] Anu, H.; Rathnakara, S.; Mallikarjunaswamy, S. Enhanced ECG Signal Classification With CNN-LSTM Networks Using Aquila Optimization. Eng. Technol. Appl. Sci. Res. 2025, 15, 23461-23466.
[19] S. S et al., "Smart Object Detection and Audio Feedback System for the Visually Impaired," 2025 International Conference on Innovations in Intelligent Systems: Advancements in Computing, Communication, and Cybersecurity (ISAC3), Bhubaneswar, India, 2025, pp. 1-6, doi: 10.1109/ISAC364032.2025.11156667.
[20] S. S et al., "Securing Pharmaceutical Supply Chains using AI-Integrated Blockchain Technology," 2025 International Conference on Innovations in Intelligent Systems: Advancements in Computing, Communication, and Cybersecurity (ISAC3), Bhubaneswar, India, 2025, pp. 1-6, doi: 10.1109/ISAC364032.2025.11156788.