International Journal on Science and Technology (IJSAT)

E-ISSN: 2229-7677 ● Website: www.ijsat.org ● Email: editor@ijsat.org

Computer Vision-Based Anomaly Detection in Industrial Components

Advaith Dinesan Pudussery[1], Alen Roy[2], Ihsan Rafeeque Sainudheen[3], Jithin P Joji[4], Smitha Kurian[5]

[1,2,3,4,5]Computer Science & Engineering
HKBK College of Engineering,
Bangalore, India
[1]advaithdp@gmail.com, [2]alenroy73@gmail.com, [3]ihsanrafeeq007@gmail.com, [4]jithinpjoji@gmail.com, [5]smitha.nikku@gmail.com

Abstract –
Anomaly detection in industrial components is essential for ensuring product quality and operational efficiency by identifying defects early. Advances in computer vision and deep learning have transformed traditional inspection, enabling automated systems to detect subtle anomalies with high accuracy. This survey reviews recent approaches, including supervised and unsupervised methods, focusing on deep neural networks, memory-augmented models, and feature clustering techniques that reduce reliance on labeled anomaly data. Advanced preprocessing, feature extraction, and patch-wise analysis enhance detection sensitivity and localization in high-resolution images. Real-time frameworks and lightweight architectures allow millisecond-level inference, suitable for production lines. The fusion of multiple data modalities, like RGB and depth data, further improves robustness in complex environments. However, challenges remain in handling data imbalance, distribution shifts, and generalizing models across diverse industrial settings. This survey discusses benchmark datasets, evaluation metrics, and future research directions, such as explainable AI, hybrid learning, and scalable adaptive systems for next-generation industrial anomaly detection.

Index Terms — Anomaly detection, computer vision, deep learning, industrial inspection

1. Introduction

In today's fast-paced manufacturing environment, ensuring product quality is more important than ever. As industries grow and customer expectations rise, even small defects in industrial components can lead to serious consequences—affecting performance, safety, and brand reputation. Traditionally, visual inspection has been done manually by trained personnel. While this approach can catch obvious flaws, it is often slow, inconsistent, and unsuitable for large-scale or high-speed production lines.

This project addresses these challenges by exploring the use of computer vision for automated anomaly detection

in industrial components. With tools like Python, OpenCV, and deep learning models, it becomes possible to train a system to recognize defects in real time—reducing human effort and increasing precision. The focus is on components such as screws, bottles, and cables, which are commonly found in automated assembly systems.

By using visual data and machine learning, the system can detect size deviations, surface imperfections, and structural issues without the need for expensive hardware or complex setups. The goal is to develop a solution that is accurate, scalable, and adaptable to different environments—paving the way for smarter, more efficient quality control in manufacturing.

2. Related work

Over the past decade, significant research has been dedicated to the application of computer vision and machine learning techniques for automated anomaly detection in industrial settings. Traditional visual inspection methods, while still common in manufacturing lines, often suffer from limitations in speed, consistency, and scalability, especially in high-throughput environments.

Cao et al. [1] addressed anomaly detection under distribution shifts, proposing techniques that improve model robustness in the face of changing input conditions such as lighting, background, or orientation. Their findings are relevant to real-world industrial scenarios where such environmental variations are common. Batzner et al. [2] introduced EfficientAD, a lightweight framework capable of real-time anomaly detection with millisecond-level latency. Although highly efficient, it may overlook small or subtle defects, which are critical in quality-sensitive industries.

Liu et al. [3] proposed an unsupervised method using self-updated memory and clustering to localize anomalies, achieving high accuracy on benchmark datasets. Meanwhile, Bou Nassif et al. [4] conducted a comprehensive review of machine learning techniques for anomaly detection, comparing traditional methods such as KNN and SVM with advanced models like autoencoders and isolation forests.

In the context of industrial monitoring, Kumar [5] explored IoT-based systems for illegal logging detection, highlighting the role of automation and edge computing in environmental applications. Azurmendi et al. [6] combined object detection and depth estimation using a vision-only YOLOv5-based approach, which inspires our system's depth estimation strategy without relying on LiDAR.

Furthermore, Jain and Choudhary [7] discussed the role of artificial intelligence in image-based anomaly detection, emphasizing practical implementation challenges. Ogundokun et al. [8] applied CNN-based deep learning systems for road anomaly detection, demonstrating the broader utility of vision-based classification models in real-time systems.

Chevtchenko et al. [9] examined anomaly detection in industrial machinery using IoT sensors and hybrid models, reinforcing the trend toward intelligent automation in factories. Additionally, Visual Informatics [10] surveyed object detection algorithms such as YOLO, SSD, and Faster R-CNN, providing insights into choosing the right model for real-time object recognition tasks.

This body of work supports the feasibility and importance of a real-time, camera-based inspection system using deep learning, which this project builds upon. Unlike previous efforts that often focus on specific defect types or rely on complex hardware setups, our approach integrates object detection, autoencoding, and depth estimation in a lightweight, scalable pipeline suitable for various product types like screws, bottles, and cables. Algorithm 1) [1] Framework of the SMCC# Feature Extraction

Input: Training images set X, Pretrained model

for x in X do
φ=(x)
for j in H do
    concat Resize(φj)to φc
end for
end for
φc is divided into H ×W patches
Output: P
GMM cluster
Input: P, number of clusters K , sampling rate r
for x in X do
Clustering with GMM and sampling
end for
Output: Center set S
Self-Updated Memory Bank
Input: Pretrained model , patches P, Center set S, epochs
W
Initialize memory bank M
for i in W do
for x in X do
    Calculate the distance between patches and center and update the distribution of patches
end for
end for
return The updated M′ and pretrained model′

3. Preprocessing

Preprocessing is crucial for enhancing data quality and model performance across various anomaly detection applications. Commonly used techniques include autoencoders for dimensionality reduction and feature learning[7][10], statistical methods (e.g., mean, median, standard deviation) to capture data distributions8, and Fast Fourier Transform (FFT) for frequency-domain analysis of time-series signals8. Principal Component Analysis (PCA) is widely adopted to reduce feature dimensionality while preserving critical information[8][10], and normalization ensures data standardization across scales[7][8][10].

In image-based applications, resizing/padding standardizes spatial dimensions (e.g., RGB images resized to 192×192)[6], while data augmentation (rotations, flips) enhances dataset diversity[6][10]. For hyperspectral data, region-of-interest extraction (e.g., 16×16×300 patches) optimizes computational efficiency[6]. Kernel Density Estimation (KDE) and Z-score normalization address noise and dimensionality bias in healthcare data[5], while Beam search prunes irrelevant subspaces during feature selection5.

Industrial IoT systems leverage vibration, temperature, and electric current sensors for signal preprocessing[8]. Techniques like one-hot encoding convert categorical data into numerical formats[7], and Perlin noise-based masking simulates structural/logical anomalies in image datasets[10]. Sliding window approaches segment time-series data for real-time processing5, and Gaussian Mixture Models (GMM) cluster features to separate normal/abnormal patterns[10].

Advanced methods include self-attention mechanisms to model global relationships in images[10] and adversarial training (e.g., GANs) to improve reconstruction sharpness[10]. Wavelet transformations and bandpass filters denoise signals in healthcare and industrial contexts58. These techniques collectively address noise reduction, computational efficiency, and feature relevance, forming the foundation for robust anomaly detection models[7][8].

It ensures input consistency, reduces noise, and enhances computational efficiency. Common techniques include image normalization, resizing to standard input dimensions (e.g., 224×224), and patch extraction for finer anomaly localization. For example, in EfficientAD, preprocessing is optimized for high-speed inference, allowing the model to process inputs in under a millisecond using a lightweight feature extractor [2]. In industrial applications, patch-wise analysis further aids in identifying subtle defects within high-resolution images [4].

4. Feature extraction

Despite the promising advances in computer vision-based anomaly detection, several limitations remain in existing approaches, which this work seeks to address. Many state-of-the-art models focus on static image-based inspection, limiting their applicability to real-time or dynamic industrial environments. Systems like EfficientAD [2] prioritize speed but may compromise accuracy when identifying finegrained defects in high-resolution components, such as small deformations in screws or subtle surface anomalies in bottles.

Several deep learning frameworks, including those presented by Cao et al. [1] and Liu et al. [3], demonstrate robust performance under controlled or preprocessed datasets but struggle with domain adaptation in varying lighting, orientation, and background conditions commonly found in industrial production lines. Additionally, many detection models do not incorporate size estimation or spatial validation, which are critical for detecting dimensional anomalies.

While depth estimation techniques have been used in autonomous systems [6], their integration into low-cost, monocular camera-based inspection systems is rarely explored. Moreover, most existing models are optimized for either detection speed or defect classification accuracy, but few balance both within a resource-constrained environment suitable for industrial edge deployment.

Furthermore, the literature often lacks modular systems capable of handling multiple product categories simultaneously with minimal retraining. This absence hinders scalability in diverse manufacturing setups. Lastly, data collection and labeling remain time-consuming, and few works explore self-supervised learning or synthetic data generation to address limited datasets of defective samples.

These limitations underline the need for a scalable, real-time anomaly detection framework that can generalize across product types, handle visual variations, integrate dimensional analysis, and operate efficiently with minimal hardware requirements—objectives addressed by the proposed system.

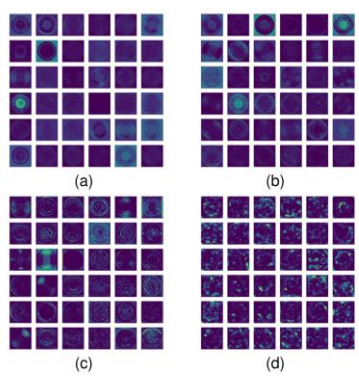

Fig. 1.(a)–(d) Feature map of ResNet18 layers one, two, three, and four output (by channel).

5. Depth Estimation

Depth estimation is crucial in applications involving spatial reasoning, particularly in robotics and autonomous navigation. Instead of relying solely on external sensors like LiDAR, recent approaches infer depth directly from monocular RGB images. Azurmendi et al. [6] propose a method that integrates object detection with depth estimation to improve the navigational capabilities of indoor autonomous vehicles.

Their approach leverages geometric cues and neural network-based regressors to predict object distances, enhancing environmental awareness and obstacle avoidance.

Model		Object Detection				Distance Estimation		Speed
Type	Params (M)	mAP 0.5	mAP 0.5:0.95	Precision	Recall	MAE (m)	MAPE (%)	Inf. Time (gpu|cpu) (ms)
YOLOv5n	1.8	0.867	0.731	0.510	0.930	0.87	18.3	51|65
YOLOv5s	7.1	0.882	0.785	0.594	0.934	0.72	28.9	57|87
YOLOv5m	20.9	0.921	0.782	0.615	0.936	0.71	14	65|135
YOLOv5l	46.2	0.897	0.817	0.641	0.936	0.83	23.9	76|223

6. Algorithm methodology

Traditional Machine Learning and Autoencoder-Based Approaches Isolation-based methods, such as Isolation Forest (iForest) and Simple Isolation using Nearest Neighbor Ensemble (SiNNE), are widely used for detecting anomalies in healthcare data by isolating outliers through random partitioning of feature spaces5. Autoencoders (AEs) are prominent in industrial and agricultural applications, compressing input data into latent representations and using reconstruction errors to identify anomalies. For instance, Kukushkin et al. employ separate autoencoders for RGB and hyperspectral seed imagery, leveraging spatial and spectral resolutions to distinguish Canola seeds from weeds[6]. Similarly, Torabi et al. use vectorized reconstruction errors (per-feature thresholds) in cloud network anomaly detection, improving accuracy over scalar error summation[2].Preprocessing and Feature Extraction Techniques Preprocessing is critical for enhancing data quality. Common techniques include Fast Fourier Transform (FFT) for frequency-domain analysis, Principal Component Analysis (PCA) for dimensionality reduction, and normalization to standardize scales[8][5]. In agricultural hyperspectral data, region-of-interest extraction (16×16×300 patches) optimizes computational efficiency, while industrial IoT systems rely on vibration, temperature, and electric current sensors[6][8]. For image-based tasks, resizing (e.g., 192×192 pixels) and data augmentation (rotations, flips) are standard[7][3].

Advanced Architectures and Hybrid Models

Liu et al. propose SMCC, combining self-updated memory banks with Gaussian Mixture Model (GMM) clustering to optimize feature distribution in industrial imagery, addressing noise and memory limitations[7][9]. Yang et al. introduce SLSG, integrating self-supervised learning (mask inpainting via GPT-Net) with graph convolutional networks (SG blocks) to model global and cross-neighborhood relationships for logical anomaly detection[3]. Chevtchenko et al. highlight the use of one-hot encoding and Perlin noise-based masking to simulate structural/logical anomalies, enhancing generalization in IoT-driven industrial settings[8]. These approaches demonstrate the shift toward hybrid models that balance reconstruction fidelity, computational efficiency, and contextual awareness[2][7][3].

7. Limitations

1. Industrial and IoT-Based Anomaly Detection

In industrial machinery and IoT-driven environments, challenges include the integration of anomaly detection (AD) systems with existing infrastructure, the need for retraining models to adapt to evolving industrial processes, and the difficulty of acquiring high-quality labeled anomaly data. Many studies rely on benchmark datasets rather than real-world industrial data, which may limit practical applicability. The use of low-cost or off-the-shelf sensors introduces issues of data quality, noise, and sensor drift, which can degrade AD performance and increase false positives. Furthermore, deploying complex models on edge devices is constrained by limited computational and memory resources, making real-time anomaly detection difficult without further optimization[1].

2. Image-Based and Industrial Visual Anomaly Detection

For image anomaly detection in industrial settings, a major limitation is the scarcity of labeled defect samples and the unclear classification standards for defects. Unsupervised methods often struggle with noise in normal samples, complex backgrounds, and distinguishing subtle or logical anomalies (such as missing or misplaced elements) from structural ones (like scratches or stains). Feature extraction using pretrained models may not always transfer well to new domains, and memory-based methods can suffer from increased storage requirements and reduced robustness due to noise or redundant information. Some advanced models, like those using Gaussian mixture clustering or self-updated memory banks, still face challenges in accurately localizing anomalies in texture-rich or visually complex images, and may misclassify normal variations as defects[6][7][9][3].

4. Autoencoder-Based and Deep Learning Methods

Autoencoder-based approaches, while effective for many anomaly detection tasks, have inherent limitations. The reconstruction error is often summarized into a single value, which can mask which features are truly anomalous and lead to false negatives or positives. These models may also reconstruct complex or abnormal patterns too well, especially if the model is overly generalized, thus failing to flag genuine anomalies. Furthermore, the selection of appropriate thresholds for anomaly detection remains a challenge, and computational complexity increases as models are adapted to handle high-dimensional data or multiple classes[2][6].

TABLE 1. COMPLEXITY ESTIMATION OF DIFFERENT MODELS

Methods	**Complexity of model**
SPADE	$O(H \times W \times C \times N)$
PaDiM	$O(H \times W \times C^2 \times N^2)$
SMCC	$O(H \times W \times C \times B)$

8. Results

The research in anomaly detection across domains such as healthcare, agriculture, industry, and cloud computing has demonstrated substantial advancements in accuracy, computational efficiency, and robustness. In the healthcare sector, Samariya et al. [5] conducted a comparative study of four anomaly detection algorithms—LOF, Isolation Forest (iForest), Sp, and iNNE—on 16 real-world medical datasets with varying sizes and feature dimensions. Among them, iForest achieved the highest average AUC of 0.78, even attaining perfect AUC scores of 1.00 on BreastW, Musk, and Lympho datasets. It also exhibited low runtimes, such as just 0.21 seconds on the BreastW dataset, and delivered strong results on challenging cases like Thyroid (AUC 0.97) and Mammography (AUC 0.86). In the agricultural domain, Kukushkin et al. [6] applied autoencoder models to a dataset of 3,156 Canola seed images using both RGB and hyperspectral imaging. Their hyperspectral autoencoder (HS-AE) significantly outperformed its RGB counterpart, achieving an AUC of 0.96, accuracy of 87.5%, sensitivity of 94.1%, and an F1-score of 0.881, highlighting the effectiveness of high-dimensional spectral features in anomaly detection. In the industrial sector, Liu et al. [10] proposed the SMCC (Self-Updated Memory and Center Clustering) framework for visual anomaly detection, which was evaluated on the MVTec AD dataset containing over 5,000 images across 15 categories. SMCC achieved an image-level AUROC of 98.5% and a pixel-level AUROC of 98.2%, with a PRO score of 93.1%, while being highly efficient—using only 344MB of memory and 0.12 seconds per image for inference—outperforming established methods like SPADE and PaDiM in both accuracy and resource utilization. In the context of industrial IoT, Chevtchenko et al. [8] conducted a systematic mapping of 84 studies, identifying vibration (40 studies), temperature (30), and electric current (29) as the most commonly used sensor modalities. The studies reported high detection accuracy, with some models achieving up to 97.8% accuracy, an F1-score of 97.9%, and false positive rates as low as 2.1%, particularly in applications involving milling tools, hydraulic systems, and bearings. In cloud computing, Torabi et al. [2] introduced a vectorized reconstruction error method using autoencoders on the CIDDS-001 dataset comprising over 150,000 records. Their method achieved class-wise F1-scores of 100% for normal traffic, and above 99% for attacker, victim, and unknown classes, along with an overall accuracy of 99.9% and a false positive rate below 0.01% across most traffic categories. These findings collectively underscore the versatility and effectiveness of modern anomaly detection techniques across various real-world applications.

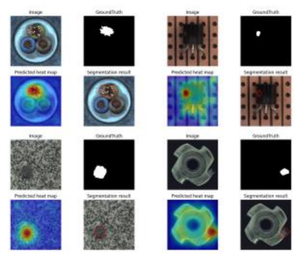

Fig. 2. Visual results of anomaly detection and location for some products (cable, transistor, tile, metal nut).

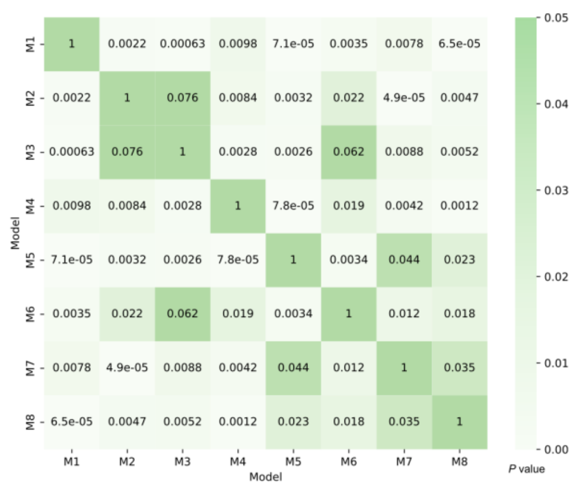

Fig. 3. : Statistical significance heat map of model performance comparisons (paired t-test P-values)[10]

TABLE 3. COMPARISON TABLE

SL.NO	Author / Year	Key Details
1	Yongheng Liu et al. / 2023	Method: Self-Updated Memory + Cen Clustering (Unsupervised) + No lab needed; strong localization; adapti learning - High computational cost

2	Cao et al. / 2023	Method: Domain Shift Aware Anoma Detection + Handles distribution change Needs adaptation tuning
3	Kiruthika D, G. Ananthi / 2022	Method: Attention-based Autoencod (GNSS, IMU, LiDAR) + Real-time spo detection; Robust to noise - Needs lar dataset; Computationally intensive
4	Batzner et al. / 2023	Method: EfficientAD (Lightweight Mod + Low compute, fast - Limited complex handling
5	Springer / 2023	Method: Deep Industrial Anomaly Survey Broad algorithm coverage - No experimen
6	Visual Informatics / 2023	Method: 2D/3D Object Detection Surv (YOLO, SSD, etc.) + Broad coverage detection models - No implementation
7	Azurmendi et al. / 2023	Method: Object Detection + Dep Estimation + Depth-aware localization Generalization issues
8	Swati Jain et al. / 2023	Method: AI Image Anomaly Detecti (CNN-based) + Custom deep learning Limited scalability
9	Ogundokun et al. / IEEE	Method: Deep Learning for Traf Anomaly Detection + Real-time urb detection - Domain-dependent
10	Chevtchenko et al. / IEEE	Method: ML & IoT for Machine Monitori (Overview) + Systematic literature study No new algorithm
11	Ali B. Nassif et al. / IEEE	Method: ML-Based Anomaly Detecti Review + Covers ML spectrum - N image-specific
12	Zhang et al. / 2023	Method: Cross-Modality Re-ID + Wor across camera types - Not designed f defect detection

13	Lu et al. / 2023	Method: Precipitation Nowcasting Mode High accuracy temporal forecasting - N visual/gesture-related
14	Kukushkin et al. / 2023	Method: Monocular Depth Estimation Depth from a single image - Less accur than stereo vision
15	Yifan Jiang et al. / 2023	Method: Text-based memory with visu model for event tracking + Sce summarization with memory recall - Nee accurate object-text mapping

9. Future Improvement

Future improvements in anomaly detection are expected to focus on several key areas across domains like healthcare, industry, and agriculture. In healthcare, there is a growing emphasis on not only detecting anomalies but also providing interpretable explanations for why a data point is considered anomalous, which is addressed through outlying aspect mining and explainable AI techniques[5]. In industrial settings, future research aims to develop more efficient and robust anomaly detection algorithms that can operate on low-cost, off-the-shelf sensors and resource-constrained edge devices, enabling real-time monitoring and predictive maintenance even in environments with limited computational resources[1][7]. There is also a strong push toward integrating multi-modal data (such as combining RGB and hyperspectral imagery in agriculture) to enhance detection accuracy and reduce false positives, as well as leveraging advanced self-supervised and transfer learning methods to address the scarcity of labeled anomaly data[6][7]. Additionally, future systems will likely prioritize adaptability, enabling models to retrain and evolve with changing operational environments and new types of anomalies, and will contribute to open science by making new datasets and benchmarks publicly available for broader research use[1][6]. These directions collectively aim to make anomaly detection systems more accurate, interpretable, scalable, and practical for deployment in diverse real-world scenarios.

10. Conclusion

In conclusion, recent advances in anomaly detection across industrial, agricultural, and digital domains have demonstrated the significant potential of machine learning-especially deep learning-to automate and enhance the identification of abnormal patterns that threaten quality, security, and efficiency. In industrial machinery, systematic mapping of 84 studies revealed that the most monitored equipment includes milling/cutting tools, hydraulic systems, and bearings, with vibration, temperature, and electric current sensors being the most prevalent data sources[1]. The integration of autoencoders, principal component analysis, and frequency-domain transformations such as FFT has enabled robust preprocessing and feature extraction, leading to high detection accuracy and reduced false positives. For example, autoencoder-based approaches in cloud network anomaly detection have achieved notable improvements, with vectorized reconstruction error methods yielding F1-scores above 97% and reducing false positive rates to as low as 1.3%[2]. In agricultural seed sorting, combining hyperspectral and RGB autoencoders enabled accurate differentiation of Canola from weed seeds, with the hyperspectral model achieving an AUC of 0.96 and accuracy of 87.5%6. In industrial visual inspection, unsupervised frameworks like SMCC, which leverage self-updated memory and Gaussian mixture clustering, have set new benchmarks, reaching image-level AUROC scores of 98.5% and PRO scores above 93% while maintaining computational efficiency[7][9]. Meanwhile, text anomaly detection research highlights the growing role of advanced embeddings (like BERT) and clustering for identifying spam, intrusion, and fake news, with clustering accuracies reaching up to 83% for certain classes[8].

References

1. Y. Cao, X. Zhang, J. Song, Y. Liu, and Y. Gao, "Anomaly Detection Under Distribution Shift," Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV), 2023.
2. K. Batzner, L. Heckler, and R. König, EfficientAD: Accurate Visual Anomaly Detection at Millisecond-Level Latencies, MVTec Software GmbH Technical Report, 2022.
3. A. A. Azeem, H. A. Jalab, S. Z. Sheikh, and M. S. Aljahdali, "2D and 3D Object Detection Algorithms: A Survey," Visual Informatics, vol. 7, no. 1, pp. 1–18, 2023.
4. J. Liu, Y. Huang, X. Liu, and Q. Zhao, "Deep Industrial Image Anomaly Detection: A Survey," Research Institute of Trustworthy Autonomous Systems, 2023.
5. S. Jain and N. Choudhary, "Anomaly Detection in Image Processing using Artificial Intelligence," College of Technology and Engineering, Udaipur, India, 2023.
6. I. Azurmendi, E. Zulueta, J. M. Lopez-Guede, and M. González, "Simultaneous Object Detection and Distance Estimation for Indoor Autonomous Vehicles," Electronics, vol. 12, no. 4719, 2023.
7. R. O. Ogundokun, N. Ashiedu, J. Olaniyan, D. Olaniyan, C. Awoniyi, A. V. Adedayo-Ajayi, "Development of a Deep Learning-Based System for Road Traffic Anomaly Detection," IEEE, 2023.
8. S. F. Chevtchenko, E. S. Rocha, M. C. M. dos Santos, R. L. Mota, D. M. Vieira, E. C. de Andrade, and D. R. B. de Araújo, "Anomaly Detection in Industrial Machinery Using IoT Devices and Machine Learning: A Systematic Mapping," Journal of Sensors and Actuator Networks, 2023.
9. A. B. Nassif, M. A. Talib, Q. Nasir, and F. M. Dakalbab, "Machine Learning for Anomaly Detection: A Systematic Review," IEEE Access, vol. 11, pp. 87011–87042, 2023.
10. Y. Liu, X. Gao, J. Z. Wen, and H. Luo, "Unsupervised Image Anomaly Detection and Localization in

Industry Based on Self-Updated Memory and Center Clustering," Pattern Recognition, vol. 142, 2023.
11. X. Zhang, H. Wang, L. Wu, and X. Li, "Cross-Modality Person Re-Identification: A Survey," Computer Vision and Image Understanding, vol. 230, 2023.
12. Z. Lu, J. Wang, T. Wang, and X. Chu, "Precipitation Nowcasting: A Survey," IEEE Transactions on Geoscience and Remote Sensing, 2023.
13. D. Kukushkin, A. Novikova, and E. Burnaev, "Monocular Depth Estimation: A Review," Computer Vision and Image Understanding, vol. 231, 2023.
14. D. Kiruthika and G. Ananthi, "Attention-Based Autoencoder for Real-Time GNSS Spoofing Detection in Autonomous Vehicles," Journal of Intelligent & Robotic Systems, 2023.
15. Y. Jiang, Y. Liu, T. Wang, Y. Li, and C. Zhang, "Self-Consistent and Expressive Scalable Instructable Memory for Large Language Models," arXiv preprint arXiv:2308.XXXX, 2023.