Enhancing Road Safety through Real-Time Unsupervised Anomaly Detection in Crash Barriers

Divya Varshney
Department of Civil Engineering
Indian Institute of Technology Roorkee
Roorkee, India
divya.ce@sric.iitr.ac.in

Avnish Panwar
MFSDAI
Indian Institute of Technology Roorkee
Roorkee, India
avnish_p@mfs.iitr.ac.in

Indrajit Ghosh
Department of Civil Engineering
Indian Institute of Technology Roorkee
Roorkee, India
indrajit.ghosh@ce.iitr.ac.in

***Abstract*— Real-time detection of damage in roadside crash barriers is vital for maintaining highway safety and enabling data-driven infrastructure management. Manual inspection methods are labor-intensive and impractical over long highway stretches. This paper presents a real-time unsupervised anomaly-detection framework based on deep convolutional autoencoders for identifying damaged crash barriers from roadway imagery. The proposed system integrates YOLOv10n for crash-barrier detection and the Segment Anything Model (SAM) for precise segmentation, isolating the barrier region before analysis. The autoencoder, trained exclusively on undamaged samples, classifies anomalies by comparing reconstruction errors against an adaptive threshold. A dataset comprising over 50,000 GoPro-captured images from Indian highways was used for training and validation. Experimental results demonstrate 97% accuracy and an AUC of 0.97, confirming the framework's ability to detect dents, cracks, and deformations under varying environmental conditions. The proposed approach eliminates the need for labeled data, offering a scalable, cost-effective solution for autonomous road-safety monitoring and predictive maintenance.**

Keywords*— *Unsupervised learning, Autoencoder, Anomaly detection, Road safety, YOLOv10n, Segment Anything Model (SAM)

I. INTRODUCTION

Autonomous vehicles and road surveillance systems depend on the accurate detection of road elements such as markings, signs, and crash barriers to ensure safe navigation and efficient monitoring. Effective identification of normal and damaged crash barriers that are broken, displaced, or deformed, is critical for maintaining road safety and enabling timely infrastructure maintenance. Crash barriers play a vital role in reducing the severity of accidents, preventing dangerous median crossovers, protecting vehicles from fixed objects, and ensuring smoother traffic flow while also safeguarding pedestrians and cyclists. India faces a significant road safety crisis, with 4,61,312 road accidents reported in 2022, resulting in 1,68,491 fatalities and 4,43,366 injuries [1]. Rollovers, often caused by speeding, distracted driving, sharp turns, and potholes, contribute heavily to these accidents. The government has set an ambitious target to reduce road accidents by 50% by 2024 through measures like improved infrastructure, stricter enforcement of traffic laws, and public awareness campaigns [2].

Integrating AI-driven technologies for anomaly detection in crash barriers offers a promising solution. Deep learning models, such as convolutional neural networks (CNNs), are trained on sensor data to detect patterns indicating damage, such as cracks, rust, and deformations, in real-time (7,8). This approach enables predictive maintenance by identifying potential issues before they escalate, ensuring continuous monitoring and enhancing road safety. The benefits include improved barrier reliability, cost efficiency through timely repairs, and data-driven decision-making for infrastructure management. By adopting these AI-powered solutions, transportation authorities can significantly enhance road safety, reduce fatalities, and ensure the long-term reliability of critical road assets.

II. LITERATURE REVIEW

A. Anomaly Detection

Anomaly detection techniques have been widely studied, particularly for applications in traffic monitoring, lane-changing behaviour, and road asset management. In [3], authors propose an efficient, unsupervised approach for detecting anomalous lane-changing events in driving behaviour. Their method combines a Recurrent-Convolutional Autoencoder to capture spatio-temporal features, a one-class Support Vector Machine (SVM) for anomaly identification, and t-SNE for visualization. The model, validated with data from 50 drivers and over 59,000 lane change events, successfully detects abnormal lane changes without labelled data, emphasizing the variability of abnormal behaviour. This personalized driving behaviour monitoring system enhances the ability to identify and analyze individual driving anomalies. Unsupervised anomaly detection is commonly applied to visual data, particularly in road asset monitoring. As discussed in the study by Rivera et al. [4], visual anomaly detection involves identifying unusual patterns in images. While supervised methods require labelled anomalies, unsupervised methods are more prevalent due to the rarity of anomalies. These methods are applied to various domains, such as defect detection, medical imaging, and traffic surveillance. The unsupervised approach often involves image-level and pixel-level anomaly detection, where pixel-level methods focus on detecting anomalies in specific regions of an image. However, these methods face challenges related to computational efficiency and the need for large, diverse datasets. Future research in this field focuses on improving these models' scalability and real-time application.

XXX-X-XXXX-XXXX-X/XX/$XX.00 ©20XX IEEE

B. Pattern Recognition and Reconstruction Using Deep Learning

Pattern recognition, a critical aspect of machine learning, involves developing systems to categorize and recognize objects or events based on learned features. Key tasks within pattern recognition include clustering, classification, and dimensionality reduction. Bioinformatics uses these methods extensively for gene prediction, phenotype classification, and data analysis from high-throughput technologies [5]. This makes pattern recognition crucial in many bioinformatics applications, including predicting protein functions and understanding genetic interactions.

In the context of visual anomaly detection, both supervised and unsupervised methods play significant roles. Supervised methods require labelled data to train models for anomaly detection, while unsupervised methods are often employed due to the scarcity of labelled anomalous instances. These methods can be categorized into image-level detection, where an entire image is classified as anomalous, and pixel-level detection, which identifies anomalies in specific image regions [6], [7], [8]. Classical methods include density estimation, one-class classification, and image reconstruction, with pixel-level methods using advanced feature modelling for more precise anomaly detection.

Several deep learning techniques for traffic anomaly detection, emphasizing methods like Deep Semi-Supervised Anomaly Detection (Deep SAD) and Deviation Networks (DevNet), which are effective in detecting anomalies in various traffic datasets highlighted in [9]. Pre-processing steps, such as segmentation and localization, are critical for improving model performance in traffic-related scenes. Additionally, researchers demonstrate the use of transfer learning to enhance asset management in transportation, using pre-trained neural networks to achieve high accuracy despite a limited dataset [10]. Their approach achieved 97% accuracy for roadside barrier detection, illustrating the potential of transfer learning for improving the efficiency of asset data collection.

Building upon similar advancements in road-safety monitoring, a YOLOv8 + Segment Anything Model (SAM) pipeline for detecting and segmenting anomalous traffic signs under Indian conditions was developed in [11], achieving high precision through synthetic-data augmentation and autoencoder-based reconstruction. This work provides a methodological foundation for applying deep-learning frameworks to other road assets such as crash barriers.

A novel approach for anomaly detection in videos by combining future frame prediction and reconstruction methods is presented [12]. This approach, which involves sequentially predicting future frames and reconstructing images using two connected U-Net blocks, outperforms traditional methods that rely solely on one of these techniques. By combining both methods, their approach effectively detects anomalies by creating larger reconstruction errors for abnormal events while maintaining accurate predictions for normal events, making it suitable for real-time applications. Finally, in [13] a semi-supervised method for anomaly detection in High Performance Computing (HPC) systems has been introduced. Their approach uses autoencoders to detect anomalies based on the normal behaviour of supercomputer nodes. This method outperforms traditional semi-supervised techniques, offering a more efficient solution for anomaly detection in HPC environments.

III. Methodology

A. Data Collection

The data used to train an unsupervised model for anomaly detection was collected from mobile phones held by pillion riders on bikes and by installing advanced Go-Pro cameras in cars. Fig. 1 shows the samples for both regular and anomalous crash barrier data collected from the hilly highways of the lower Himalayas and the highways of the Gangetic plains. We gathered multiple videos from these different sites, carefully selecting diverse scenarios involving various guardrails to ensure a comprehensive dataset.

Fig. 1 Normal crash-barrier (Left), Anomalous crash-barrier (Right) from the Dataset

B. Data Pre-processing

We collected multiple videos and converted them into frames at a rate of 60 frames per second, resulting in approximately 54,000 images from a 30-minute video. To increase the size and variance of our dataset while preserving the shape and texture of the crash barriers, we applied various data augmentation techniques to these images. These techniques included flipping, applying Gaussian filters[14], and adjusting contrast and brightness. These augmentations helped our dataset encompass different daylight and seasonal scenarios, ensuring a more robust and comprehensive training set.

The video was converted into frames, from which we selected those containing crash barriers. To enhance dataset diversity, augmentations like flipping and Gaussian blur were applied. The YOLOv10n model processed these augmented frames, retaining only those with crash barrier detections. Annotations were done using Roboflow, and the annotated frames were used for segmentation tasks. Finally, the Segment Anything Model (SAM) was applied for precise image segmentation, dividing the images into relevant segments for further analysis [15].

C. Data Annotation and Segmentation

Although there was no requirement for data annotation for providing supervision to our anomaly detection system, the images that we presented to it needed to be devoid of the background which usually contains vegetation and road surface. That means first we had to pick out the region of interest in the images which could be fed to the anomaly

detection model. To pick regions of interest we used hand annotations of the bounding boxes on the crash barrier images. Utilizing the annotated data to Grounded SAM (Segment Anything Model). This model was essential to our methodology, enabling semantic segmentation for precise pixel-level classification.

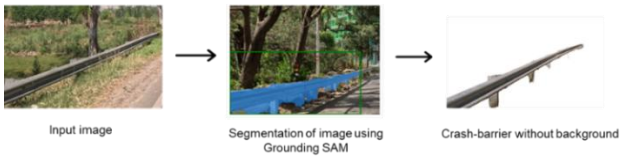

Fig. 2 Conversion of Input Image to Segmented Image

Grounding SAM is applied to the image of a crash-barrier shown in Fig. 2, and the segmentation-only Region of interest, i.e., crash-barrier, was taken for the further training process.

The entire process starts with the input video, which is first converted into multiple frames. These frames are then fed into an object detection model, specifically YOLOv10n [16], [17]. During the object detection phase, each frame is analyzed for the presence of crash barriers. If the crash barrier is detected (positive detection), the frame is saved; otherwise, it is discarded, as illustrated in Fig. 3.

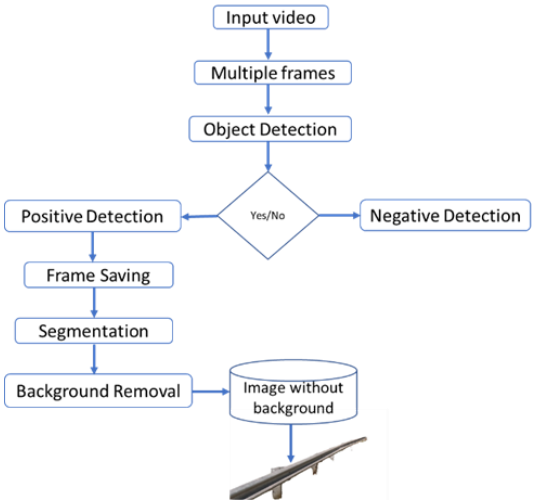

Fig. 3 Flow chart for Raw video to the Segmented Image

Once the relevant frames are saved, segmentation is performed using the Grounded Segment Anything Model (SAM). This segmentation step is crucial for background removal, ensuring that only the crash barrier regions are isolated. The final output is an image containing the crash barrier with the background removed, ready for further analysis and processing.

D. *Anomaly Detection*

This study employs a convolutional autoencoder, a neural network designed for unsupervised learning and image data compression. It consists of two components: an encoder that compresses input images into a compact latent-space representation by capturing spatial hierarchies and features [18] and a decoder that reconstructs the original images from this representation using deconvolutional layers (23) [16]. The model was trained to minimize reconstruction error, using Mean Squared Error (MSE) as the loss function, enabling efficient encoding and decoding for tasks like image denoising and dimensionality reduction.

The encoder processes images of size (3, 64, 64) through four convolutional layers, progressively reducing dimensions to (128, 4, 4), using kernel size 3, stride 2, and ReLU activation. The decoder reverses this process with transpose convolutional layers, reconstructing images back to (3, 64, 64). Sigmoid activation is applied in the final layer to ensure output values are within [0, 1]. The autoencoder effectively removed noise and irrelevant background, isolating regions of interest like crash barriers for further analysis.

The entire workflow, from processing raw images taken from a video to classifying them as anomalous or non-anomalous using an autoencoder, is illustrated in Fig. 4. Initially, raw images are extracted from the video. These images are then processed through a convolutional autoencoder, which attempts to reconstruct the original images. The reconstruction error is calculated by comparing the input images to their reconstructed counterparts[19]. If the reconstruction error exceeds a threshold of 0.01, the image is classified as an anomalous crash-barrier. Otherwise, it is classified as a normal, non-anomalous crash-barrier. This method leverages the autoencoder's ability to highlight discrepancies in the images, effectively identifying anomalies in the crash barrier.

Fig. 4 Workflow of Extracted Crash-barrier to the Reconstructed crash-barrier

IV. Results and Discussion

This section evaluates the performance of the autoencoder-based approach in detecting anomalies. The accuracy is measured by comparing the autoencoder's reconstruction error against a predefined threshold, allowing us to distinguish between anomalous and non-anomalous data. By analyzing the results, we can determine the effectiveness of the autoencoder in identifying deviations from normal patterns and its potential for real-world anomaly detection applications.

A. *Evaluation Matrices*

We utilized the evaluation metrics presented in the Table to assess the model's performance. Below is a brief explanation of each metric:

TABLE I PERFORMANCE METRICS FOR EVALUATION

Metric	Significance	Equation
Accuracy	It tells about the overall performance of the model.	$\frac{TP+TN}{TP+TN+FP+FN}$
Precision	It tells about the accuracy of the positive predictions made by the model	$\frac{TP}{TP+TN}$

Recall	Ability of a classification model to correctly identify all positive instances from the total number of actual positive instances	$\frac{TP}{TP+FN}$
F-1 Score	It provides a balance between precision and recall, making it useful for evaluating the overall performance of a model.	$2\left(\frac{Precision \times Recall}{Precision + Recall}\right)$
mAP (Mean Average Precision)	Comprehensive metric that evaluates the accuracy of object detection models by considering both precision and recall across multiple object classes	$\sum \frac{AP_i}{N}$

Fig. 5 shows the training performance of the YOLOv10n model for crash-barrier detection. As the number of training epochs increases, both mAP50 and mAP50–95 progressively improve, approaching values close to 0.9, indicating strong detection accuracy. Simultaneously, the loss metrics decrease, reflecting reduced prediction error and improved model convergence. These trends confirm that the model effectively learns discriminative features over time, leading to enhanced object detection performance.

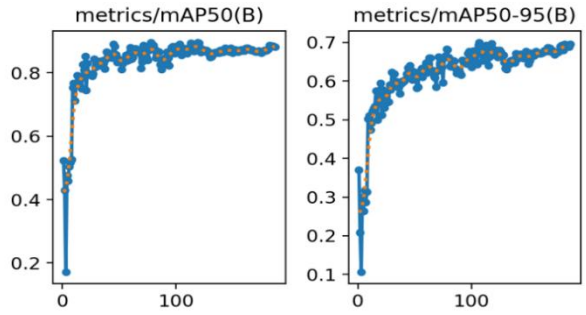

Fig. 5 Training performance of YOLOv10n showing mAP50 and mAP50–95 convergence with increasing epochs

These metrics in Fig. 6 indicate that the convolutional autoencoder is highly effective at identifying both normal and anomalous images with a high degree of precision and recall, resulting in an overall accuracy of 97%.

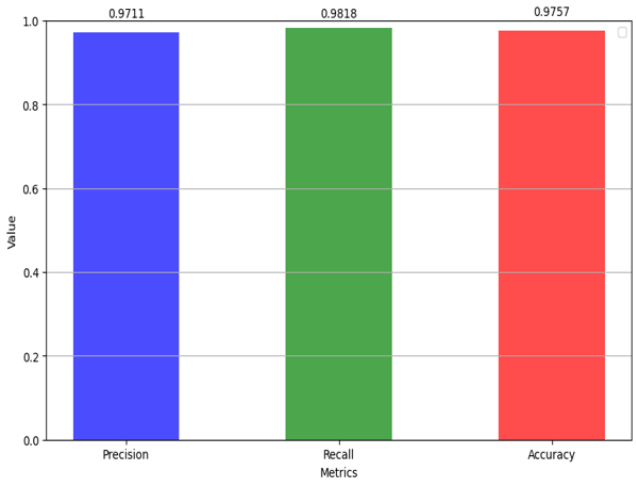

Fig. 6 Metrics of Convolutional Autoencoder

B. *Reconstruction Error Distribution and Reconstruction Results*

The crash barrier was reconstructed using a convolutional autoencoder to assess potential anomalies in its structure. The autoencoder learns to compress and reconstruct normal barrier images, allowing deviations in reconstruction errors to indicate anomalies. During evaluation, a threshold of 0.01 was set based on the reconstruction error, meaning any deviation exceeding this value is considered an anomaly.

Fig. 7 illustrates that the reconstruction error for the convolutional autoencoder is approximately 0.01 to 0.09. The lower reconstruction error in the convolutional autoencoder suggests that it is effective in accurately reconstructing the images, making it a more suitable choice for tasks requiring high-quality image reconstruction and anomaly detection. This also shows the better feature extraction capacity of convolution based encoder as compared to a linear layer based autoencoder.

Fig. 7 Reconstruction error results

C. *Confusion Matrix and ROC(Receiver Operating Characteristic) Curve*

a) Confusion Matrix: The performance of the autoencoder was also assessed using a confusion matrix, which provides a detailed breakdown of prediction accuracy across different categories. For this evaluation, a total of 532 good images and 548 anomalous images were used, Fig. 8 shows the confusion matrix.

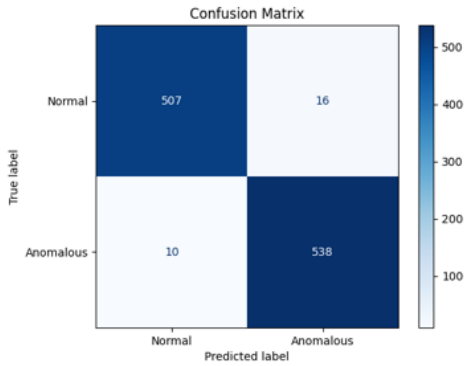

Fig. 8 Confusion matrix of Convolutional Autoencoder

b) ROC Curve: The performance of the autoencoder was further evaluated using the ROC (Receiver Operating Characteristic) curve, a graphical representation of the true

positive rate against the false positive rate. The area under the ROC curve (AUC) is a key metric indicating the overall ability of the model to distinguish between anomalous and non-anomalous data.

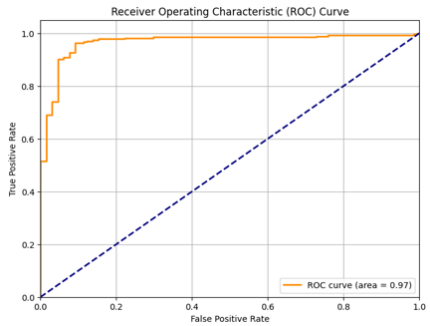

Fig. 9 ROC curve of Convolutional Autoencoder

For the convolutional autoencoder, the AUC (Area Under Curve) is 0.97 shown in Fig. 9, demonstrating an excellent capability to identify anomalies with a high degree of accuracy correctly.

V. Conclusions

This study introduced a convolutional autoencoder-based approach for real-time anomaly detection in crash barriers. By leveraging unsupervised deep learning, the system effectively reconstructs images and identifies deviations, achieving a high accuracy of 97%. The model's ability to capture spatial features enables precise anomaly detection, outperforming traditional methods. With its strong performance in reconstruction error analysis, confusion matrix evaluation, and ROC curve assessment, the convolutional autoencoder proves to be a reliable solution for automated road infrastructure monitoring. This approach has significant implications for improving road safety by enabling early detection of barrier damages, facilitating timely maintenance, and enhancing the effectiveness of autonomous road surveillance systems.

References

[1] "Press Release: Press Information Bureau." Accessed: Jul. 30, 2024. [Online]. Available: https://pib.gov.in/PressReleaseIframePage.aspx?PRID=1973295

[2] Y. Zou and A. P. Tarko, "An insight into the performance of road barriers − redistribution of barrier-relevant crashes," Accid Anal Prev, vol. 96, pp. 152–161, Nov. 2016, doi: 10.1016/J.AAP.2016.07.022.

[3] P. Fan, J. Guo, Y. Wang, and J. S. Wijnands, "A hybrid deep learning approach for driver anomalous lane changing identification," Accid Anal Prev, vol. 171, p. 106661, Jun. 2022, doi: 10.1016/J.AAP.2022.106661.

[4] M. Riveiro, M. Lebram, and M. Elmer, "Anomaly detection for road traffic: A visual analytics framework," IEEE Transactions on Intelligent Transportation Systems, vol. 18, no. 8, pp. 2260–2270, Aug. 2017, doi: 10.1109/TITS.2017.2675710.

[5] D. De Ridder, J. De Ridder, and M. J. T. Reinders, "Pattern recognition in bioinformatics," Brief Bioinform, vol. 14, no. 5, pp. 633–647, Sep. 2013, doi: 10.1093/BIB/BBT020.

[6] Z. Zheng et al., "A Fused Method of Machine Learning and Dynamic Time Warping for Road Anomalies Detection," IEEE Transactions on Intelligent Transportation Systems, vol. 23, no. 2, pp. 827–839, Feb. 2022, doi: 10.1109/TITS.2020.3016288.

[7] Y. Bao, Z. Tang, H. Li, and Y. Zhang, "Computer vision and deep learning–based data anomaly detection method for structural health monitoring," Struct Health Monit, vol. 18, no. 2, pp. 401–421, Mar. 2019, doi: 10.1177/1475921718757405.

[8] K. Gopalakrishnan, H. Gholami, A. Vidyadharan, A. Choudhary, and A. Agrawal, "CRACK DAMAGE DETECTION IN UNMANNED AERIAL VEHICLE IMAGES OF CIVIL INFRASTRUCTURE USING PRE-TRAINED DEEP LEARNING MODEL," International Journal for Traffic and Transport Engineering, vol. 8, no. 1, pp. 1–14, 2018, doi: 10.7708/ijtte.2018.8(1).01.

[9] A. Alloqmani, Y. B. Abushark, A. I. Khan, and F. Alsolami, "Deep Learning based Anomaly Detection in Images: Insights, Challenges and Recommendations," IJACSA) International Journal of Advanced Computer Science and Applications, vol. 12, no. 4, p. 2021, Accessed: Jul. 31, 2024. [Online]. Available: www.ijacsa.thesai.org

[10] C. Ryan, F. Murphy, and M. Mullins, "End-to-End Autonomous Driving Risk Analysis: A Behavioural Anomaly Detection Approach," IEEE Transactions on Intelligent Transportation Systems, vol. 22, no. 3, pp. 1650–1662, Mar. 2021, doi: 10.1109/TITS.2020.2975043.

[11] D. Varshney, T. Pal and I. Ghosh, "Detection and Segmentation of Anomalous Traffic Signs Using Yolov8 and Segment Anything Model (SAM) for Indian Specific Conditions," *2024 IEEE International Conference on Electronics, Computing and Communication Technologies (CONECCT)*, Bangalore, India, 2024, pp. 1-6, doi: 10.1109/CONECCT62155.2024.10677306

[12] Y. Tang et al., "Integrating prediction and reconstruction for anomaly detection," PaReL, vol. 129, pp. 123–130, Jan. 2020, doi: 10.1016/J.PATREC.2019.11.024.

[13] A. Borghesi, A. Bartolini, M. Lombardi, M. Milano, and L. Benini, "A semisupervised autoencoder-based approach for anomaly detection in high performance computing systems," Eng Appl Artif Intell, vol. 85, pp. 634–644, Oct. 2019, doi: 10.1016/J.ENGAPPAI.2019.07.008.

[14] M. Kristan, A. Leonardis, and D. Skočaj, "Multivariate online kernel density estimation with Gaussian kernels," Pattern Recognit, vol. 44, no. 10–11, pp. 2630–2642, Oct. 2011, doi: 10.1016/J.PATCOG.2011.03.019.

[15] A. Kirillov et al., "Segment Anything," Proceedings of the IEEE International Conference on Computer Vision, pp. 3992–4003, Apr. 2023, doi: 10.1109/ICCV51070.2023.00371.

[16] J. K. Chow, Z. Su, J. Wu, P. S. Tan, X. Mao, and Y. H. Wang, "Anomaly detection of defects on concrete structures with the convolutional autoencoder," Advanced Engineering Informatics, vol. 45, p. 101105, Aug. 2020, doi: 10.1016/J.AEI.2020.101105.

[17] R. K. Megalingam, K. Thanigundala, S. R. Musani, H. Nidamanuru, and L. Gadde, "Indian traffic sign detection and recognition using deep learning," International Journal of Transportation Science and Technology, vol. 12, no. 3, pp. 683–699, Sep. 2023, doi: 10.1016/J.IJTST.2022.06.002.

[18] A. B. Amjoud and M. Amrouch, "Object Detection Using Deep Learning, CNNs and Vision Transformers: A Review," IEEE Access, vol. 11, pp. 35479–35516, 2023, doi: 10.1109/ACCESS.2023.3266093.

[19] C. Zhou and R. C. Paffenroth, "Anomaly detection with robust deep autoencoders," Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, vol. Part F129685, pp. 665–674, Aug. 2017, doi: 10.1145/3097983.3098052/SUPPL_FILE/ZHOU_ANOMALY_DETECTION.MP4.