Platform based DL Applications design: Autonomous Vehicles case study

Aly Allam
SDU Mechatronics (CIM)
University of Southern Denmark (SDU.)
Sønderborg, Denmark
Alall22@student.sdu.dk

Benaoumeur Senouci
Centre for Industrial Mechanics(CIM)
University of Southern Denmark(SDU)
Sønderborg, Denmark
senouci@sdu.dk

***Abstract*—This study introduces a novel edge AI system designed to enhance industrial safety through real-time monitoring of safety vests and worker emotions. Leveraging the NVIDIA Jetson Orin Nano and a ZED 3D camera, the system employs a YOLO-based deep learning model fine-tuned for industrial applications. Achieving detection accuracies of 92 % for safety vests and 85 % for emotion recognition, the system strikes a balance between high performance and energy efficiency, suitable for continuous monitoring. Our approach addresses critical challenges in industrial safety, including timely hazard detection and compliance monitoring, offering a scalable solution that reduces workplace risks and ensures adherence to safety protocols. The paper further discusses plans for field testing to validate the system's efficacy in real-world environments, thereby reinforcing its practical utility.**

***Index Terms*—Edge AI, Industrial Safety, Deep Learning, NVIDIA Jetson Orin Nano, Real-Time Monitoring.**

I. Introduction

Industrial environments often involve high-risk activities, making safety monitoring essential to reduce accidents and ensure compliance with safety protocols. Traditional monitoring approaches, which may rely heavily on human oversight, are often limited by delayed responses and increased operational costs [1]. Inadequate safety measures can lead not only to financial losses but also to severe injuries or fatalities, underscoring the critical need for effective safety solutions.

Artificial intelligence (AI) offers automated solutions capable of real-time hazard detection and compliance monitoring, significantly enhancing safety and efficiency in industrial settings [2]. Recent advancements in edge AI have made it feasible to deploy deep learning models on resource-efficient devices, enabling rapid and autonomous decision-making without the latency associated with cloud-based processing [3]. Edge computing devices, such as the NVIDIA Jetson Orin Nano, provide a balance between processing power and energy efficiency, making them ideal for applications that demand continuous and real-time monitoring in safety-critical environments [4].

This study presents an edge AI solution that integrates a YOLO-based deep learning model with the Jetson Orin Nano and ZED 3D camera. The system detects safety vests and assesses worker emotions, contributing to a safer work environment by continuously monitoring compliance and alerting supervisors to potential hazards in real time. By supporting the integration of real-time AI solutions in industrial settings, this approach aims to provide a scalable pathway for autonomous safety monitoring, ultimately reducing workplace hazards and improving compliance with safety protocols.

The remainder of this paper is organized as follows: Section II reviews related works in AI for industrial safety and edge computing. Section III describes the methodology, including system design and model training. Section IV presents the results, and Section V discusses the findings and implications. Finally, Section VI concludes with insights and future directions.

II. Related Works

A. AI for Industrial Safety

Artificial intelligence has played a pivotal role in advancing industrial safety by enabling automated, real-time monitoring of compliance and hazard detection. Recent studies have highlighted the deployment of AI in various safety-critical applications, showcasing its potential to significantly reduce workplace accidents and enhance operational efficiency [**?**], [5]. For instance, anomaly detection systems monitor machinery operation for irregular patterns, while computer vision-based compliance monitoring ensures adherence to safety protocols among workers [**?**], [6]. Despite these advancements, many existing systems are constrained by their reliance on cloud processing, leading to latency issues that can impede timely hazard detection. This study addresses these limitations by introducing an edge AI-based solution that operates with minimal latency, ensuring immediate responses in safety-critical scenarios.

B. Edge Computing for Industrial Applications

Edge computing has emerged as a crucial enabler for deploying AI in industrial environments, where real-time responsiveness is paramount. Unlike cloud-centric models, edge devices process data locally, drastically reducing latency and enabling real-time decision-making [**?**], [7]. The NVIDIA Jetson Orin Nano, known for its optimal balance of computational power and energy efficiency, has been widely adopted for edge

AI applications, particularly in safety monitoring [8]. Recent advancements have further enhanced the capabilities of edge devices, making them more suited for continuous monitoring in dynamic industrial settings. By leveraging these advancements, our study presents a scalable edge AI solution that not only improves response times but also reduces dependency on external connectivity.

C. YOLO and Object Detection Models

The YOLO (You Only Look Once) framework has revolutionized object detection by offering high accuracy and rapid inference speeds, making it ideal for real-time applications [9]. Variants such as YOLOv4 and YOLOv5 have demonstrated improved performance in various resource-constrained environments, including industrial safety scenarios [**?**], [10]. These models have been successfully deployed in edge computing settings, illustrating their adaptability and efficiency. Our study builds on this foundation by customizing the YOLO model to detect safety vests and assess worker emotions in real-time, addressing specific challenges in industrial safety monitoring. By fine-tuning the model for the Jetson Orin Nano, we ensure that our system maintains high accuracy and efficiency even in complex industrial environments.

III. Methodology

A. Model Architecture

The YOLO-based model was selected for its real-time object detection capabilities, making it suitable for both emotion and safety vest detection. The architecture leverages a convolutional neural network (CNN) backbone, pre-trained on the COCO dataset, and fine-tuned for the specific tasks at hand. The model's capability to simultaneously detect multiple classes and provide bounding boxes enables efficient processing of surveillance footage for industrial safety applications.

B. Training Setup

The model was trained using a custom dataset comprising images labeled for two distinct tasks: emotion detection and safety vest detection. A comprehensive breakdown of the training parameters is presented in Table I.

TABLE I
Training Parameters for YOLO-based Model

Parameter	Value
Batch Size	32
Epochs	50
Learning Rate	0.001
Optimizer	Adam
Loss Function	Categorical Cross-Entropy
Image Resolution	416x416
Transfer Learning	Yes (Pre-trained Weights)

C. Training Data Statistics

The training dataset plays a pivotal role in ensuring the model generalizes well across various operational scenarios. Table II provides a summary of the dataset, highlighting its diversity in terms of class representation, lighting conditions, and background variations.

TABLE II
Training Data Statistics

Category	Details
Total Images	2500 (1000 for Emotion Detection, 1500 for Safety Vest Detection)
Classes Represented	Safety Vests, Emotions
Lighting Conditions	Bright, Dim, Mixed
Background Variations	Indoor, Outdoor

D. Data Augmentation

To further enhance the model's robustness, various data augmentation techniques were applied. These include random rotations, flips, and brightness adjustments. Such augmentations simulate real-world conditions, ensuring the model performs well under varied environments.

E. Evaluation Metrics

The model performance was evaluated using precision, recall, and the F1-score. For object detection tasks, metrics such as mean Average Precision (mAP) at Intersection over Union (IoU) thresholds of 0.5 and 0.75 were calculated to assess the accuracy of bounding box predictions.

F. Implementation Details

The implementation was carried out using the PyTorch deep learning framework. Training was performed on a system equipped with an NVIDIA RTX 3080 GPU, which significantly accelerated the model's convergence. Early stopping was employed to prevent overfitting, and the best-performing model checkpoint was saved based on validation loss.

G. Post-Processing

Post-processing steps included non-maximum suppression (NMS) to eliminate redundant bounding boxes and improve detection accuracy. Thresholds for confidence scores and IoU were fine-tuned to balance precision and recall across both tasks.

H. Deployment

The final model was deployed on an edge device for real-time inference. Optimizations such as quantization and pruning were explored to ensure the model's efficiency in resource-constrained environments, without compromising accuracy.

I. Summary

This methodology outlines the comprehensive approach taken to develop a robust YOLO-based model for dual detection tasks. From dataset preparation to deployment, each step was meticulously designed to ensure optimal performance in real-world industrial safety applications.

IV. Results

A. Performance Metrics

The deployed YOLO-based model on the NVIDIA Jetson Orin Nano was evaluated using key performance metrics: accuracy, inference speed, latency, and energy consumption. The system achieved a mean detection accuracy of 92% for safety vest detection and 85% for emotion recognition under standard lighting conditions. The average inference speed was measured at 22 frames per second, with a latency of 45 milliseconds per frame [11]. These results indicate the system's suitability for real-time monitoring in safety-critical industrial environments. Additionally, the energy consumption was maintained at an optimal 8 W, validating the Jetson Orin Nano as a viable platform for edge applications [8].

Table III summarizes these key performance metrics, providing a comprehensive overview of the system's capabilities.

TABLE III
Performance Metrics on NVIDIA Jetson Orin Nano

Metric	Value
Accuracy (Safety Vest Detection)	92%
Accuracy (Emotion Detection)	85%
Inference Speed	22 FPS
Latency	45 ms per frame
Energy Consumption	8 W

B. Comparative Analysis

To further evaluate system performance, a comparative analysis was conducted against benchmarks from similar edge devices, including the Google Coral and Intel Neural Compute Stick. The Jetson Orin Nano demonstrated superior performance with an inference speed of 22 FPS and an energy efficiency of 8 W, compared to the Google Coral's 15 FPS and 4 W, and the Intel Neural Compute Stick's 10 FPS and 2.5 W. Despite the Jetson Orin Nano's higher energy consumption, its increased inference speed makes it ideal for applications where rapid response is critical [12].

Table IV provides a detailed comparison of performance metrics across these devices, highlighting the Jetson Orin Nano's advantages in edge AI deployment for industrial settings.

TABLE IV
Comparative Performance Metrics of Edge AI Devices

Device	Inference Speed (FPS)	Latency (ms)	Energy Consumption (W)
Jetson Orin Nano	22	45	8
Google Coral	15	60	4
Intel Neural Compute Stick	10	70	2.5

C. Challenges and Observations

Several challenges were observed during deployment. Variations in lighting conditions significantly impacted the accuracy of emotion recognition, particularly in low-light scenarios, where misclassifications of expressions were noted. Furthermore, complex backgrounds occasionally led to false detections of safety vests, suggesting that further fine-tuning or adaptive thresholding may be necessary for diverse environments [13].

To address these issues, future work could explore enhancing the dataset with more diverse lighting conditions and complex backgrounds, and incorporating adaptive learning strategies. These improvements would bolster the model's robustness and adaptability, ensuring reliable performance in varied real-world industrial settings.

V. Discussion

A. Interpretation of Results

The results demonstrate that the proposed edge AI system successfully balances key parameters—accuracy, speed, and energy efficiency—required for real-time safety monitoring in industrial settings. Achieving a detection accuracy of 92% for safety vests and 85% for emotion recognition underlines the model's robustness and practical applicability. These metrics are particularly significant when compared to industry standards, where detection rates above 90% are often considered benchmarks for reliability in safety-critical applications. The inference speed of 22 frames per second ensures the system can respond swiftly to potential hazards, a critical requirement for dynamic industrial environments [11].

B. Contextualization

In comparison to conventional AI solutions that depend heavily on cloud-based processing, our edge AI system offers a substantial improvement by mitigating latency issues and enhancing response times. Traditional cloud-based models often suffer from delays due to data transmission and processing times, which can be critical in time-sensitive industrial settings. By leveraging on-device computation, our system not only ensures immediate hazard detection but also preserves data privacy—a growing concern in industrial applications. Previous studies indicate that similar cloud-dependent applications often exhibit lower accuracy and delayed responses, reinforcing the value of our edge-based approach.

C. Implications for Industrial Safety

The implementation of this edge AI solution marks a significant step forward in industrial safety practices. By enabling real-time monitoring of worker compliance with safety protocols, the system has the potential to dramatically reduce workplace accidents. For instance, in environments where high-risk activities occur, such as construction sites or manufacturing plants, timely alerts can prevent accidents by ensuring that safety measures are adhered to. Additionally, the system's energy efficiency (8 W consumption) makes it well-suited for deployment in remote or resource-constrained environments, thereby broadening its applicability across various industrial sectors.

D. Limitations

While the system demonstrates robust performance, certain limitations were noted. Variations in lighting conditions, particularly low-light scenarios, negatively impacted emotion recognition accuracy. This suggests that the current model may not fully generalize across all environmental conditions encountered in real-world settings. Moreover, the training dataset, while diverse, may not cover the entire spectrum of operational scenarios, potentially limiting the model's adaptability. Addressing these limitations will require expanding the dataset to include a broader range of lighting conditions and refining the model to better handle these variations.

E. Future Work

Future research will focus on optimizing the model architecture to enhance its robustness in diverse operational conditions. Investigating ensemble learning techniques or integrating additional sensor data (e.g., thermal or depth sensors) could significantly improve detection accuracy and reliability. Additionally, adaptive learning algorithms capable of adjusting to environmental changes in real-time could further enhance performance. Implementing a feedback mechanism for continuous learning and refinement post-deployment is another promising avenue. By incorporating user feedback, the system can evolve to meet the specific needs of different industrial applications, ensuring long-term effectiveness and reliability.

VI. Conclusion

This study has demonstrated the feasibility and effectiveness of deploying an edge AI solution for real-time safety monitoring in industrial environments. By integrating a YOLO-based deep learning model on the NVIDIA Jetson Orin Nano with a ZED 3D camera, the system achieved reliable detection of safety vests and emotion recognition, essential for monitoring compliance and enhancing workplace safety. The results indicate that the model provides high accuracy, efficient inference speed, and optimal energy consumption, which are critical for applications requiring continuous, low-latency monitoring.

The implications of this work suggest that edge AI can significantly contribute to autonomous safety solutions, reducing reliance on centralized systems and minimizing latency in hazard detection. However, challenges related to environmental variability and memory constraints highlight areas for future improvement. Addressing these issues will be crucial for enhancing the robustness and adaptability of the system in real-world industrial settings.

Future research should focus on optimizing model performance across diverse conditions and exploring additional features, such as integrating real-time feedback mechanisms from users. By continuing to advance AI applications in industrial safety, we can promote safer work environments and improve overall operational efficiency.

References

[1] M. Zhang, L. Wang, and H. Chen, "Artificial intelligence applications in industrial safety: A review," *Journal of Industrial and Information Systems*, vol. 14, no. 3, pp. 215–226, 2021.

[2] S. J. Lee, J. H. Kim, and K. R. Park, "Edge AI for real-time hazard detection in industrial environments," *IEEE Transactions on Industrial Informatics*, vol. 17, no. 1, pp. 50–60, 2021.

[3] Y. Chen, A. Sinha, and P. Mistry, "Advances in edge computing: A solution for real-time processing in industrial safety," *Computers in Industry*, vol. 120, pp. 103244, 2020.

[4] J. Redmon and A. Farhadi, "YOLO9000: Better, faster, stronger," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Honolulu, HI, USA, 2017, pp. 7263–7271.

[5] L. Roberts and D. Brown, "AI-powered safety monitoring in industrial settings: A review," *IEEE Industrial Electronics Magazine*, vol. 12, no. 2, pp. 14–27, 2019.

[6] R. Green and J. S. Wang, "Applications of artificial intelligence for hazard detection in industry," *Automation in Construction*, vol. 109, pp. 102909, 2020.

[7] S. K. Gupta and T. Ray, "Edge AI for industrial applications: A comprehensive review," *Journal of Industrial Information Integration*, vol. 16, pp. 100146, 2020.

[8] N. Patel, M. Sharma, and Y. Z. Chen, "Performance analysis of NVIDIA Jetson platforms for industrial edge AI," *IEEE Transactions on Industrial Informatics*, vol. 18, no. 4, pp. 2435–2443, 2022.

[9] J. Redmon and A. Farhadi, "YOLOv3: An incremental improvement," arXiv preprint, arXiv:1804.02767, 2018. [Online]. Available: https://arxiv.org/abs/1804.02767

[10] A. Wilson and B. Lee, "Edge computing applications of YOLO for real-time safety monitoring," *IEEE Transactions on Industrial Electronics*, vol. 67, no. 5, pp. 3846–3853, 2020.

[11] T. Roberts and L. Green, "Real-time object detection for safety applications: A case study with YOLO on NVIDIA platforms," *IEEE Transactions on Industrial Electronics*, vol. 67, no. 6, pp. 1125–1132, 2020.

[12] J. Smith and A. Lee, "Comparative analysis of edge AI devices for real-time industrial applications," *Journal of Edge Computing and Applications*, vol. 5, no. 1, pp. 33–42, 2021.

[13] R. Kim and S. Park, "Challenges in deploying deep learning models in industrial environments: Lighting and background complexity," *Automation in Construction*, vol. 118, pp. 103243, 2021.