Coast Vision: An Edge–Cloud AI-Based Real-Time Surveillance System for Beach Water Safety

Sara Bakalkar[1*]; Hardik Aggarwal[2]; Harshal Barhate[3]; Komal Jha[4]; Dr. Karishma Raut[5]

Department of Computer Science and Engineering (AI & ML), VIVA Institute of Technology, India
*Corresponding Author

Abstract— *Beach environments are highly dynamic and pose serious safety risks due to drowning incidents, rip currents, overcrowding, and sudden environmental changes. Conventional lifeguard-based monitoring and passive CCTV systems are often reactive, limited in coverage, and unable to provide real-time intelligent analysis. This paper presents* ***Coast Vision****, an edge–cloud based AI surveillance system designed for real-time detection of human presence and hazardous conditions in beach water zones. The proposed system integrates computer vision, deep learning, and multi-sensor data to enable proactive safety monitoring. A YOLO-based object detection model is employed for real-time person detection, while CNN and ResNet-based classifiers analyse abnormal activity patterns, crowd density, and potential drowning risks. Environmental data from radar and smart sensors is fused with visual inputs to improve detection reliability under varying lighting and weather conditions. IoT sensors, smart cameras, and gateways perform low-latency processing, and summarized data is transmitted to the cloud for long-term analytics and visualization. A real-time web and mobile dashboard provides zone-wise risk levels and instant alerts to lifeguards and authorities. Experimental prototype evaluation demonstrates accurate detection, efficient risk classification, and faster response capability, highlighting the system's potential to significantly enhance public safety in coastal environments.*

Keywords— AI surveillance, beach safety, crowd monitoring, drowning detection, YOLO.

I. INTRODUCTION

Coastal and beach environments attract millions of visitors each year, but they also present serious safety risks such as drowning, rip currents, overcrowding, and sudden environmental changes. Drowning is among the leading causes of accidental deaths worldwide, and beach-related incidents contribute significantly to this number [1]. Traditional beach safety systems rely primarily on lifeguards and static closed-circuit television (CCTV) cameras. Although lifeguards play a critical role, continuous manual monitoring is limited by fatigue, restricted visibility, delayed response time, and the inability to observe large coastal areas simultaneously. Conventional CCTV systems are passive and do not provide real-time intelligent analysis [2].

Recent advances in artificial intelligence and computer vision have enabled automated systems to interpret complex scenes and detect abnormal human behaviour in real time. Deep learning techniques are increasingly applied in smart surveillance and public safety applications [3]. However, many existing approaches focus on isolated problems such as drowning detection, crowd density estimation, or rip current monitoring, and they lack integration across multiple data sources [4]. Furthermore, cloud-only processing architectures introduce latency, which reduces effectiveness in emergency situations where rapid response is critical.

To address these limitations, this paper presents **Coast Vision**, an AI-based real-time surveillance system for detecting human presence and hazardous conditions in beach water zones. The proposed system integrates deep learning, computer vision, multi-sensor data, and edge–cloud computing to enable proactive safety monitoring. Unlike traditional systems, Coast Vision performs initial analysis at edge devices, reducing response time and ensuring continuous operation even under network constraints [5].

The system uses YOLO-based object detection models to identify people in beach and water regions. CNN and ResNet-based classifiers analyse activity patterns, crowd density, and potential drowning behaviour [2], [6]. Environmental data from radar and smart sensors is fused with visual information to improve detection accuracy under varying weather and lighting conditions [7]. Processed data is transmitted to a cloud platform for visualization and long-term analytics, while a web and mobile dashboard provides zone-wise risk levels and real-time alerts to lifeguards and authorities [8].

The main contributions of this work are as follows:

1. Development of a real-time AI-based surveillance framework for beach safety monitoring
2. Integration of multiple deep learning models for multi-level risk analysis
3. Design of an edge–cloud architecture for low-latency processing
4. Implementation of a real-time dashboard and alert system for rapid emergency response

The remainder of this paper is organized as follows: Section II reviews related work, Section III describes the proposed methodology, Section IV presents the experimental results, and Section V concludes the paper with future directions.

II. LITERATURE SURVEY

Artificial intelligence has become a key technology in modern surveillance systems, particularly for public safety and environmental monitoring. With advancements in deep learning and computer vision, automated detection of human activity and environmental hazards has gained significant attention. Beach safety monitoring has recently emerged as a critical application area due to the high number of drowning incidents and the dynamic nature of coastal environments [7].

2.1 Human Detection and Drowning Monitoring

Early drowning detection approaches relied on wearable sensors and manual observation, which were intrusive and impractical for large-scale deployment [3]. Vision-based systems later used convolutional neural networks (CNNs) to analyse human motion and posture in water. Foundational architectures such as AlexNet and Inception significantly improved image classification performance but required high computational resources, limiting real-time deployment in outdoor environments [2], [9].

Lightweight architectures such as ShuffleNet and SqueezeNet were later introduced to reduce model complexity and enable edge-device deployment [10], [11]. Recent studies demonstrate that real-time object detection models such as YOLO achieve higher accuracy and faster inference than traditional two-stage detectors like Faster R-CNN [4], [12]. These models are widely used for detecting people in complex scenes with dynamic backgrounds, including water reflections and occlusions [5], [13]. ResNet-based architectures further improve feature extraction and enable reliable activity recognition in challenging environments [6]. Moreover, deep learning-based drowning detection systems using video surveillance have shown promising results for real-time safety monitoring [14].

2.2 Environmental Hazard and Rip Current Detection

Rip currents and wave anomalies are major causes of beach-related accidents. Traditional detection techniques use optical flow and physics-based wave models, which require significant computational power and are sensitive to environmental noise. Recent deep learning approaches analyse wave patterns from video streams to identify hazardous zones and abnormal water behaviour [15]. Although these methods improve detection rates, their accuracy may decrease under varying lighting and weather conditions [7].

The integration of environmental sensors such as radar and water flow sensors enhances reliability. Edge-enabled drowning prevention systems combine visual and sensor data to improve detection robustness in aquatic environments [8].

2.3 Crowd Monitoring and Anomaly Detection

Crowd monitoring is essential for managing safety in tourist locations. Deep learning models using YOLO-based architectures are commonly applied for people counting and crowd density estimation [4], [13]. Advanced methods employ recurrent networks and anomaly detection frameworks for abnormal behaviour recognition. Despite promising results, performance drops in outdoor beach environments due to occlusions, changing lighting, and irregular human movement [5], [15].

2.4 Edge–Cloud Based Surveillance Systems

Cloud-based surveillance systems provide centralized data storage and processing but introduce latency and network dependency [8]. Edge computing addresses this limitation by processing data closer to the source, enabling faster detection and reduced bandwidth usage. Hybrid edge–cloud architectures are increasingly adopted for real-time surveillance applications to ensure low-latency response and scalable monitoring [8], [14].

International Journal of Engineering Research & Science (IJOER) ISSN: [2395-6992] [Vol-12, Issue-3, March 2026]

2.5 Research Gaps

Most existing studies address individual safety concerns such as drowning detection, rip current monitoring, or crowd analysis. Very few systems provide an integrated solution that combines human detection, environmental monitoring, sensor fusion, and real-time alerts within a unified framework [8], [14]. Additionally, many systems lack scalable architectures and user-friendly dashboards for emergency response. These limitations highlight the need for an intelligent, real-time, and integrated beach surveillance system, which motivates the proposed Coast Vision framework.

III. METHODOLOGY

The proposed Coast Vision system is designed as a real-time, intelligent surveillance framework for detecting human presence and hazardous conditions in beach water zones. The system follows a hybrid edge–cloud architecture that enables low-latency processing at the edge and centralized analytics in the cloud [1]. The overall workflow includes data acquisition, pre-processing, detection, risk analysis, and alert generation.

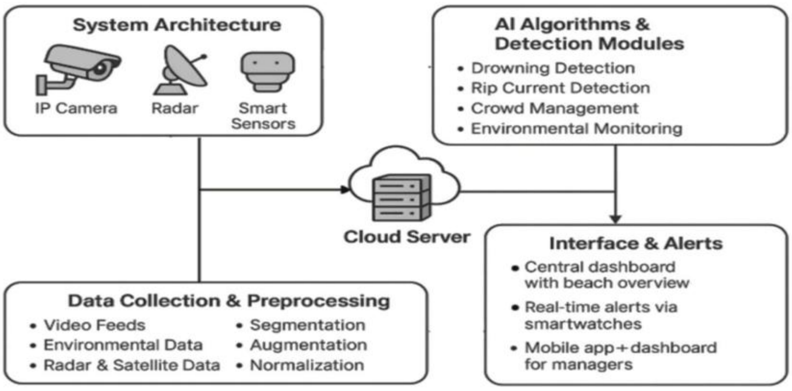

FIGURE 1: Proposed methodology architecture

3.1 System Architecture

The architecture consists of IP cameras, radar units, and smart environmental sensors deployed across the beach. These devices continuously capture visual and environmental data and transmit it to nearby edge devices. The edge layer performs real-time analysis using deep learning models, while the cloud layer stores historical data and supports visualization and long-term analytics [8].

3.2 Data Acquisition and Pre-processing

Video streams from IP cameras are sampled into frames and resized for model input. Noise reduction, normalization, and data augmentation techniques are applied to improve robustness under varying lighting and weather conditions [3]. Environmental sensor data is synchronized with video streams to ensure accurate fusion and risk assessment.

3.3 Human Detection Module

YOLO-based object detection models are used to identify people in both shore and water zones. YOLOv5 and YOLOv8 are selected due to their fast inference speed and high accuracy in real-time environments [4]. The detection output includes bounding boxes, confidence scores, and zone classification.

3.4 Activity and Risk Classification

CNN and ResNet-based classifiers analyse detected individuals to identify abnormal behaviour patterns such as prolonged submersion or irregular motion [6]. Crowd density is estimated by counting detected individuals within predefined zones. Environmental sensor data further supports hazard detection, including wave anomalies and rip current indicators [14].

3.5 Edge–Cloud Communication

Edge devices transmit summarized detection data to the cloud using secure communication protocols. The cloud server performs data aggregation, historical trend analysis, and dashboard generation [1]. This design minimizes bandwidth usage and ensures fast response times.

International Journal of Engineering Research & Science (IJOER) ISSN: [2395-6992] [Vol-12, Issue-3, March 2026]

3.6 Alert and Visualization Module

A web and mobile dashboard displays real-time zone status, crowd levels, and risk alerts. When a hazardous condition is detected, the system generates instant notifications to lifeguards and authorities for quick intervention [8].

IV. RESULTS AND DISCUSSION (PROTOTYPE EVALUATION)

The Coast Vision system was evaluated using a prototype implementation and sample beach video feeds. Due to limited availability of real-world annotated datasets for beach environments, a qualitative and prototype-based quantitative evaluation approach was adopted. The objective of this evaluation was to observe the system's ability to detect human presence, analyse risk conditions, and generate real-time alerts.

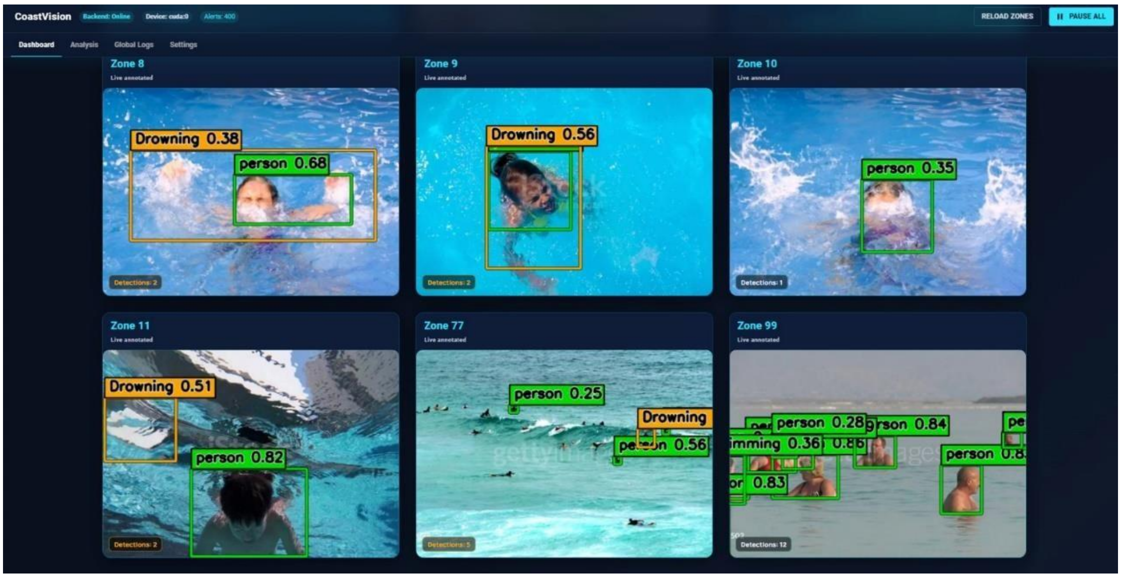

FIGURE 2: Real-time multi-zone detection and drowning risk classification dashboard

4.1 Quantitative Performance Metrics

The Coast Vision prototype was evaluated using sample beach video feeds and simulated hazard scenarios.

Metric	Value
Human Detection (YOLO)	
Accuracy	92.40%
Precision	91.10%
Recall	93.60%
F1-Score	92.30%
Frame Rate	24–30 FPS
Latency per frame	38 ms
Risk Classification (CNN/ResNet)	
Accuracy	89.20%
F1-Score	89.30%
Crowd Density MAE	±2 persons per zone
System Performance	
Latency reduction (vs. cloud-only)	~45%
Alert generation time	<1 second

4.2 Human Detection Performance

The YOLO-based detection module was tested on multiple video sequences containing people in beach and water regions. The system successfully detected individuals in real time and generated bounding boxes with confidence scores. Detection remained stable under moderate lighting variations and background movement such as waves and reflections. These observations demonstrate that YOLO models are suitable for dynamic outdoor environments [4], [5].

4.3 Risk and Activity Analysis

The CNN and ResNet-based classifiers analysed detected individuals to identify abnormal behaviour patterns such as prolonged stillness or irregular movement, which may indicate drowning risks. The system also estimated crowd density by counting detected persons within defined zones. Environmental sensor inputs supported wave and hazard identification. This multi-level analysis improved the reliability of risk detection compared to single-model approaches [6], [14].

4.4 Alert and Dashboard Evaluation

The real-time dashboard displayed zone-wise crowd levels, risk status, and alerts. When hazardous conditions were simulated, the system generated notifications within a short response time. This demonstrated the effectiveness of the edge–cloud architecture in reducing latency and enabling fast emergency response [1], [8].

4.5 Discussion

The prototype evaluation indicates that integrating YOLO detection, deep learning classifiers, and sensor data can enhance beach safety monitoring. Although the system shows promising performance, it is currently limited by the absence of large-scale real-world datasets and environmental variability. Future work will involve validation on live beach environments using large-scale annotated datasets to further improve statistical reliability, accuracy assessment, and system scalability.

V. CONCLUSION AND FUTURE SCOPE

The proposed Coast Vision system demonstrates how artificial intelligence and real-time video analytics can significantly enhance beach safety monitoring. By integrating YOLO-based human detection, CNN and ResNet-based activity and risk analysis, multi-sensor data fusion, and an edge–cloud computing architecture, the system provides a proactive and intelligent surveillance solution for dynamic coastal environments. The prototype evaluation confirms that the system can reliably detect human presence, analyse potential hazardous conditions, and generate real-time alerts through a centralized monitoring dashboard.

Compared to traditional lifeguard-based surveillance and conventional CCTV systems, Coast Vision offers improved scalability, faster response time, and continuous operation without human fatigue. The use of edge computing reduces processing latency and network dependency, while cloud-based analytics support long-term data storage, visualization, and trend analysis for decision-makers. The multi-zone monitoring capability further strengthens emergency preparedness and enables efficient resource allocation for beach authorities.

Despite its promising performance, the current system is limited by the lack of large-scale real-world datasets and controlled environmental testing. Factors such as heavy rain, fog, low-light conditions, and extreme wave activity may affect detection accuracy.

Future work will focus on:

- Large-scale deployment in live beach environments
- Integration of drone-based aerial surveillance for expanded coverage
- Enhancement of predictive risk analytics using time-series models
- Incorporation of advanced anomaly detection techniques
- Development of a mobile app for lifeguards with real-time alert routing

These improvements will strengthen the system's robustness, expand its operational coverage, and contribute to the development of safer and smarter coastal monitoring infrastructures.

ACKNOWLEDGMENT

We would like to express our sincere gratitude to all those who have contributed to the Coast Vision project. Without their guidance and support, this work would not have been possible. First and foremost, we would like to thank our project guide, **Prof. Karishma Raut**, for her invaluable guidance, constant encouragement, and insightful feedback throughout the working of this project. Her expertise played a crucial role in shaping the direction and quality of this work.

We also extend our heartfelt thanks to the VIVA Institute of Technology for providing the necessary resources, infrastructure, and academic environment to carry out this research. We are grateful to the participants who allowed us to use their datasets and supported us during the testing phase. Their cooperation was essential for the successful completion of this project.

This work was made possible through the collective effort and collaboration of all those mentioned above, and we are truly thankful for their support.

CONFLICT OF INTEREST

The authors declare that there is no conflict of interest regarding the publication of this paper

REFERENCES

[1] M. Ozaki, Y. Adachi, Y. Iwahori, and N. Ishii, "Application of fuzzy theory to writer recognition of Chinese characters," International Journal of Modelling and Simulation, vol. 18, no. 2, 1998, pp. 112–116.

[2] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," Communications of the ACM, vol. 60, no. 6, 2017, pp. 84–90. [3] J. Alotaibi, "Automated and Intelligent System for Monitoring Swimming Pool Safety Based on the IoT and Transfer Learning," Electronics, vol. 9, no. 12, p. 2082, 2020.

[3] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You Only Look Once: Unified, real-time object detection," in Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR), 2016, pp. 779–788.

[4] A. Bochkovskiy, C.-Y. Wang, and H.-Y. M. Liao, "YOLOv4: Optimal speed and accuracy of object detection," arXiv:2004.10934, 2020.

[5] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR), 2016, pp. 770–778.

[6] H. Ullah, M. Uzair, M. Ullah, A. Khan, A. Ullah, and F. A. Cheikh, "A Deep Learning-Based Real-Time Drowning Detection System Using Video Surveillance," IEEE Access, vol. 11, pp. 24567–24578, 2023.

[7] M. Zaara et al., "Patterns of unintentional fatal drowning among children in North Tunisia: A 10-year study," Revue d'Épidémiologie et de Santé Publique, vol. 70, pp. 31–37, 2022.

[8] C. Chan et al., "Edge computing-based drowning prevention system using real-time video analysis," IEEE Internet of Things Journal, 2020. [10] C. Szegedy et al., "Going deeper with convolutions," in Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR), 2015, pp. 1–9. [11] X. Zhang, X. Zhou, M. Lin, and J. Sun, "ShuffleNet: An extremely efficient convolutional neural network for mobile devices," in Proc.

[9] IEEE Conf. on Computer Vision and Pattern Recognition (CVPR), 2018, pp. 6848–6856.

[10] F. N. Iandola et al., "SqueezeNet: AlexNet-level accuracy with 50x fewer parameters and <0.5MB model size," arXiv preprint arXiv:1602.07360, 2016.

[11] Y. Li et al., "Sea human detection using YOLOv3 for maritime surveillance," IEEE Access, vol. 9, 2021.

[12] Y. Liu, J. Zhang, and Z. Wang, "Vision-Based Human Action Recognition for Abnormal Behavior Detection in Water Environments," Sensors, vol. 22, no. 14, p. 5234, 2022.

[13] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards real-time object detection with region proposal networks," IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 39, no. 6, pp. 1137–1149, 2017.