Received: 12 April 2025 | Revised: 21 August 2025 | Accepted: 8 September 2025 | Published online: 16 October 2025

RESEARCH ARTICLE

Journal of Computational and Cognitive Engineering
2025, Vol. 00(00) 1-13
DOI: 10.47852/bonviewJCCE52025908

BON VIEW PUBLISHING

Real-Time Road Accident Detection and Severity Assessment Using IoU and Deep Learning Models

Bharathi Mohan Gurusamy[1,*], Pranav Reddy Sanikommu[1] and Gayathri Muthurasu[2]

[1] *Department of Computer Science and Engineering, Amrita Vishwa Vidyapeetham-Chennai, India*

[2] *Department of Electronics and Communication Engineering, Amrita Vishwa Vidyapeetham-Chennai, India*

Abstract: Road accidents are a significant global issue, claiming lives and cash in hand every year. The system introduces an adaptive model-switching mechanism that selects the most suitable You Only Look Once (YOLO) model (v5, v8, or v11) based on system memory and weather visibility, ensuring optimal detection with minimal computational load—a feature rarely addressed in existing accident detection research. To alleviate this situation, this study envisages a near real-time accident detection and emergency response system, integrating computer vision and geospatial technologies. The developed system tracks vehicle movement and assesses collision by applying advanced object detection models such as YOLOv5, YOLOv8, and YOLOv11 with dynamic tracking and speed estimation. The system computes spatial overlap using the Intersection over Union (IoU) metric to localize vehicles accurately and recognize potential accidents while minimizing false-positive rates dramatically. Geographic data from OpenStreetMap is used in emergency routing with the help of the Haversine formula to find the shortest and fastest route to the nearest emergency services, allowing rapid notifications to hospitals and enforcement agencies. The developed framework includes a severity analysis module that assesses the degree of accident impact based on vehicular deformation, collision dynamics, and possible injuries to passengers, thus enabling emergency response prioritization based on the incident's urgency. An additional provision is made to incorporate an automatic license plate recognition system that enhances vehicle identification, thereby speeding up the initiation of insurance claims and post-accident administrative processes. The framework has undergone rigorous field trials under varying environmental conditions and remains robust in its operation under such variations. Real-time alerts have an average latency of 2.1 s and cover a comprehensive report of the incident, including location, vehicle IDs, severity, and medical assistance required. The experimental evaluation proved the system's detection accuracy of 94.6% and precision of 92.8%, thus proving its reliability and effectiveness for implementation in real time.

Keywords: real-time accident detection, AI-based system, YOLOv5, YOLOv8, YOLOv11, Intersection over Union (IoU), vehicle identification

1. Introduction

Today's traffic management and emergency service frameworks are increasingly dependent on accident detection and response systems. The incorporation of cutting-edge technologies including deep learning, the Internet of Things (IoT), and machine learning into these systems has enhanced their ability to identify accidents more quickly and precisely resulting in more efficient emergency response measures. Traffic accidents continue to be a major cause of injuries and deaths around the world, prompting interest in intelligent systems that are capable of addressing these problems.

Deep learning is used in a current study to scan real time data from different sensors for accident detection using convolutional neural networks and related algorithms. Deep learning has been utilized for such risk detection and hazard management, for example, in predicting fire incidents on construction sites from input by sensors. By using a dataset relating to road traffic environments, these techniques can potentially be used for accident prediction and prevention [1]. Likewise, vehicular communication systems based on movement are employed to relay emergency messages, thus enhancing the effectiveness of the currently established accident detection systems [2].

***Corresponding author:** Bharathi Mohan Gurusamy, Department of Computer Science and Engineering, Amrita Vishwa Vidyapeetham-Chennai, India. Email: g_bharathimohan@ch.amrita.edu

Advancements in object detection models, particularly those based on the You Only Look Once (YOLO) framework, have significantly improved the identification of small objects and potential risks in complex surroundings. These models have been successfully applied to the localization of objects within remote sensing imagery and have been extended to real-time detection of obstacles and accident scenarios in road environments [3]. The use of YOLOv5 and Swin Transformer architectures has shown greater accuracy in identifying small objects under real-world conditions, which contributes to reducing reaction times and preventing secondary accidents [4]. Recent progress has also been made in multi-object tracking, where YOLOv8 integrated with improved DeepSORT has been applied to pedestrian tracking with high accuracy and robustness [5]. Such approaches demonstrate the capability of combining strong object detection backbones with advanced tracking algorithms, offering valuable insights for continuous monitoring of traffic participants in accident-prone environments. Real-time detection of pedestrians under complex traffic conditions, using deep learning models combined with multi-object tracking algorithms, enhances both accident prevention measures and the ability to dispatch emergency services efficiently [6].

Combining deep learning, IoT, and sensor fusion in accident detection systems really helps improve how quickly emergencies are responded to, makes accident detection more accurate, and can finally save lives. These technologies work hand-in-hand to give a real-time picture of traffic conditions, enabling faster detection and response when something goes wrong. For example, deep learning models like YOLO and its upgrades can identify objects accurately even in tricky environments, while IoT devices gather raw data from cars, closed-circuit television (CCTV) cameras, and road sensors in a systematic way. When this data is combined through sensor fusion techniques, it provides a detailed snapshot of what is happening on the roads, so emergency services can be alerted immediately. Besides just spotting accidents faster, these technologies also help improve what happens afterward. For instance, including number plate recognition makes it easier to file insurance claims quickly without wasting time. Plus, adding weather data into the system ensures it still works well even in tough conditions like fog or heavy rain. Looking ahead, if these systems are further developed and integrated into broader smart city plans, we could see smoother traffic flow and fewer accidents, saving people from suffering and reducing costs. Overall, this tech movement points toward safer roads and a more connected, sustainable transportation future.

The important contributions and novelty of this paper are:

1) Adaptive multi-model accident detection

A context-aware framework capable of dynamically switching between YOLOv5, YOLOv8, and YOLOv11 depending on environmental visibility, low-light conditions, and computational resources, thereby ensuring reliable performance under diverse operating scenarios.

2) Severity estimation using geometric deformation

A novel method that employs bounding box overlap and deformation ratios based on Intersection over Union (IoU) to quantify collision severity, enabling prioritization of emergency response according to impact intensity.

3) Integration with geospatial and emergency services

Real-time accident localization through Global Positioning System (GPS) and OpenStreetMap data combined with automated alert generation to nearby hospitals, police stations, and emergency responders, significantly reducing notification latency.

4) Automation of post-accident processes

Incorporation of license plate recognition for rapid vehicle identification, automated initiation of insurance claims, and provision for organ and blood donation procedures in critical cases, thereby extending impact beyond immediate detection.

5) Demonstrated robustness across environments

Comprehensive evaluation under adverse weather, low-light, and limited visibility conditions (including low light) confirming high detection accuracy, reduced false positives, and suitability for deployment within intelligent transportation systems.

Road accident detection remains a demanding research problem because existing approaches often rely on static deep learning frameworks that fail to adapt under diverse environmental conditions such as fog, heavy rain, or low illumination. These limitations lead to inconsistent performance and increased false detections in real-world deployments. Another unresolved challenge lies in severity estimation: most prior studies focus on accident occurrence but lack accurate quantification of impact intensity, which is essential for prioritizing emergency response. Moreover, current systems are generally confined to accident recognition alone, without incorporating automated integration with emergency routing, medical response, and insurance processing pipelines.

The motivation for this investigation stems from the need to establish an accident detection system that is both adaptive and context-aware. The proposed framework addresses the challenge of varying visibility and computational constraints by dynamically switching between different object detection models. In addition, it introduces a severity estimation strategy based on bounding box deformation and IoU, enabling a more reliable differentiation between minor and severe collisions. By uniting detection, severity assessment, and real-time alert generation within a single architecture, the study attempts to bridge the gap between laboratory-level detection accuracy and large-scale practical deployment in intelligent transportation systems.

2. Literature Review

Advancements in object detection technology have significantly shaped the way we approach vehicle accident detection. A lot of research has gone into making these systems more accurate and quicker to respond—both super important for making them work better. One hot area right now is figuring out how to detect multiple objects in busy, complicated road scenes. Researchers are using deep learning tools like YOLO to get real-time detection down pat. Specifically; a big focus has been on helping models better spot small objects like cars and pedestrians especially in tricky conditions [7]. Some recent studies have even introduced new loss functions, like Corner-Point and Foreground-Area IoU loss, to help improve accuracy when objects are hidden or hard to see. All these tweaks help the models better predict where things are in real traffic, which is key for catching accidents early.

At the same time, there have been some exciting new developments in object detection, especially with utilizing transformers as the main backbone. For example, adding Swin Transformers into YOLOv5 has really boosted how well it can pick up small objects in satellite images—like cars on highways, city streets, or tricky terrains [8]. This is super helpful because objects come in all different sizes and angles, which is especially important when you are trying to spot accidents in real-time, where things can be at odd angles. Researchers have not stopped there, though—they have worked on tweaking YOLO models to make them better at spotting road hazards and traffic mishaps. Some changes include rearranging parts of YOLOv5's backbone and adjusting its loss functions, which helps the model run faster and be more accurate when detecting stuff like road damage or other weird road conditions [9]. These improvements are a big deal because they help the system stay reliable, even when visibility is poor or weather is bad. Plus, adding multi-object tracking with YOLO has made a real difference, especially for tracking pedestrians and vehicles at the same time. By enhancing versions of the DeepSORT algorithm, these systems can follow lots of moving objects at once, like pedestrians and cars in busy scenes. Using these tracking methods makes it easier for accident detection systems to tell apart moving objects from stationary stuff, making everything more dependable no matter the environment [10].

Admittedly, there have been improvements in tracking bounding box regression more efficiently with regard to the life tracking of objects across frames. Powerful IoU is one such method that makes a simple yet efficient computation of IoU, thus increasing the accuracy of object localization without hefty service time. Enhanced object detection frameworks have been increasingly refined to address challenges in dynamic road environments, where accurate identification of vehicles, pedestrians, and structural elements remains critical [11]. Recent advances in disentangled representation learning have also demonstrated the potential of model-driven and data-driven integration for anomaly detection in complex scenarios. In addition, foundation models such as Fleximo have recently shown promise in remote sensing by enabling flexible adaptation across tasks and modalities, thereby improving generalization in complex detection scenarios. This demonstrates the growing relevance of multimodal and transferable architectures for enhancing robustness in real-world accident detection pipelines [12]. These improvements thus tailor well to the needs of real-time accident detection. Paradigm after paradigm finds continuous integration into object identification systems, primarily aimed at improving the accuracy and flexibility of accident detection models. For example, the work of

Zhang et al. [13] presents an exploration into learning other metrics apart from IoU to advance bounding box regression to improve mean Average Precision (mAP) performance against classic non-maximum suppression. These advancements enhance real-time tracking accuracy while suppressing false alarms, thereby calling attention only to incidents that really demand attention. There is further advancement in underwater imaging and restoration research literature [14]. While these are not for vehicular accident detection, methods such as CycleGAN image restoration for this field can be leveraged for better accident detection in challenging lighting conditions. Such methods improve the noise-robustness of detection methods; they would be especially useful in adverse visibility situations like fog and rain. In sonar-based object detection, the work of Rathour [15] presents advanced sonar target recognition frameworks that adapt boundary frame loss for better target recognition. This framework will be beneficial particularly for real-time object detection in case of obscured road traffic environments, where objects may be partially hidden and occluded.

Emerging research areas have suggested using advanced models that include YOLOv8 and YOLOv11 for multi-view object detection enhancement [16, 17]. Increased efficiency has also been exhibited in identifying road blemishes and flexible randomness with conventional methods' shortcomings. The evolution of the YOLO framework shall continue generating improvements in the detection algorithm necessary to be incorporated in all fields, such as road safety and environmental awareness. In line with these advancements, other studies such as Shepley et al. [18] explore advanced tracking algorithms that will be composed of detection algorithms combined with multi-object tracking algorithms such as DeepSORT. Accident detection systems may be improved by enabling the tracking of several vehicles simultaneously. The work by Kim et al. [19] deals with lightweight recognition models meant to run in smart city infrastructure or in rural areas of limited computing resources.

Further, studies involving remote sensing and environmental monitoring, such as the work of Longin et al. [20], adopt methodologies that can also be translated into detection of vehicular accidents. These were more often than not used for environment-related applications like landslide detection. All these techniques rely heavily on image analysis similar to what is employed in accident detection under limited visibility conditions. Recent trends in traffic management include those in intelligent transport systems, which act towards better response times, as stated in the work of Balios et al. [21]. Such technology allows for improved access to area emergency vehicles while speeding up responses to accidents using AI-based preemptions and other automated means between signals. Other studies showed how traffic accidents relate to the amount of fatalities during accidents. Researchers analyzed prior crash records and field data to pinpoint reasons that would cause accidents and fatalities in the work of Rosayyan et al. [22]. These models are practical in accident detection systems that can work toward preemptive safety measures. Another approach to accident prevention per [23] is the identification of signs of driver fatigue and human error. It is possible to incorporate computer-vision-based systems within accident detection frameworks to increase their ability to detect and prevent accidents that result from human factors, such as drowsy driving.

The incorporation of machine learning into the topic area of secure IoT-based traffic management and accident detection is a continuing field of study within recent research [24, 25]. In addition to that, it holds promise for real-time communication, which subsequently fosters better informed decision-making and emergency response during traffic accidents. For instance, studies on emergency response systems [26] analyzed the efficiency of AI-powered medical imaging techniques in injury assessment post-accidents. Automated identification and segmentation of orbital fractures, as explored in the work of Kumar et al. [27], represent altogether enlightening material for treating medical practitioners of road accident victims. Another, the work of Bao et al. [28], stresses that there is a need to analyze national emergency response times as it deals with practical concerns in accident response cooperation. Here, these findings come with answers on better application of technology in the detection of traffic accidents as well as emergencies. In providing such outputs, an exhaustive range of actualities about artificial intelligent applications in self-driving technology can be noted as well [29]. The costs of contributing to a smart road system can be measured about accident detection, driver notification, and dynamic alteration of traffic flow to prevent chain reactions. Real-time accident detection fused with autonomous vehicles will make roads safer increasingly.

At last, the real-time vehicle accident detection system using deep neural networks was a monumental discovery in AI-based road safety applications, as evidenced in Mohan et al. [30] as well as Gayathri and Gomathy [31]. Such decentralized models can detect accidents using motion data over time with a high degree of correctness, which opens the doors to many future innovations regarding accident detection and prevention systems [32, 33]. This literature review thus details the tremendous progress made in deep learning, object detection, multiple-object tracking, and intelligent transportation systems, all of which further fine-tune accident detection, thereby increasing accuracy, decreasing response times, and improving real-time situational awareness in dynamic traffic environments.

The new achievements in the remote-sensing implementation of multimodal features and segmentation models contained promising perspectives of the upstream of frames of accidents detection. Specifically, Convolutional Neural Networks (CNN) have been shown to have great potential when it comes to data fusion of multimodal data with the goal of enhancing the classification capabilities of the data source. Li et al. [34] suggested the method of information integration in the remote sensing scene classes based on deep CNN model. The scheme involves the integration of data acquired from multiple modalities, including Synthetic Aperture Radar (SAR), optical imagery, and Light Detection and Ranging (LiDAR). This technology can make the system work with high reliability even in severe environmental conditions such as poor visibility, fog, or low-light conditions of the accidents which frequently occur in road accidents. This kind of multimodal feature fusion has the potential of posing a great enhancement to the strength of real-time accident detection models particularly where their application is lengthened to embrace the use of thermal or depth-based sensors besides using Red–Green–Blue (RGB) information. On top of that, UrbanSAM, which is a modified version of Meta Segment Anything Model (SAM), introduces another stage of flexibility and segmentation accuracy specific to urban settings. Because of the issues their subject might face regarding the object scale variance and complicated morphology of cityscapes, it is developed by an Uscaling-Adapter using multi-resolution analysis and invariance theory. Liang [35] developed a vehicle and pedestrian detection algorithm that integrates improved attention with feature fusion to enhance recognition performance. The study introduced alternatives to conventional IoU-based bounding box regression, resulting in higher localization accuracy and improved mAP. By reinforcing multi-scale feature interactions, the approach demonstrated robustness under dynamic and complex traffic conditions. These insights significantly support YOLO-based accident detection, where precise localization and rapid recognition are essential for achieving reliable real-time safety applications.

The recent developments in traffic surveillance and object detection have boosted the abilities and detection of accidents. Advanced solutions that work out (such as YOLOv8, DeepSORT, and transformer-based detection systems) demonstrate better activity in complex urban situations, whereas multimodal CNN models guarantee effectiveness in adverse weather and visual haze conditions. Procedures involving additional detail in post-collision analysis based on segmentation, such as scalable architectures such as UrbanSAM, also offer the ability to give additional precision in the post-collision analysis without significant manual prompting. The combination of these methods, as discussed in Table 1, provides a sound basis to develop intelligent and context-aware accident detection systems which have the potential of responding to the situation in real-time and under a variety of operational conditions. The current paper follows these ideas and proposes an elastic design that can adapt to hardware limitations, environmental fluctuation, and numerous object interactions, all of which are contributing to safer and more responsive smart transportation environments.

3. Dataset

To increase and validate the efficacy of the accident detection system, the datasets must be visually heterogeneous and extensive. Common Objects in Context (COCO) is one of the leading datasets that encompass a fair share of variations of objects in highly complex contexts. As opposed to datasets focused on static targets, COCO provides a lot of object categories that work well when training object detection models for dynamic road environments as shown in Figure 1. As indicated before, COCO encompasses vehicle classes like cars, buses, and trucks, as well as pedestrians, and traffic signals—all are important for accident detection. Therefore, the COCO dataset firmly laid down the foundation for initializing the accident detection system to detect objects in traffic scenes. In addition, accident video datasets were incorporated in the real-time scene simulation of usual traffic patterns and accident occurrences. These videos were chosen based on weather conditions ranging from rain and fog to other conditions where visibility may be compromised, thus making evaluation of the detection system's robustness under such challenging scenarios possible (Table 1). One key benefit of using real accident footage is: for the system to be tested against the dynamic and ever-changing conditions, it can also further deductively evaluate its performance based on visual occlusions and temporal changes.

Figure 1
Displaying sample video frames captured from real-time accident

4. Methodology

The proposed system uses artificial intelligence to provide real-time accident detection and automated emergency response. It integrates within the architecture object detection models with environmental and situational analysis, all to improve the accuracy and efficiency in responding. The framework receives video streams, recognizes pre-collision events, estimates the severity of an accident, and immediately raises alarms to authorities and stakeholders. It is ensured that the usage of multitude deep learning models guarantees the adaptive selection based on the available computational resource and environmental conditions. The system is modular, fundamentally involving acquisition von video, detection of object, analysis of collision, estimation of severity, emergency response, and automated reporting. The real-time video acquisition module allows the acquisition of live footage from either surveillance cameras or car dash cameras. This object detection module comprises deep learning-based YOLO architectures used to identify various elements along the road, such as vehicles and pedestrians and other environmental constituents. Collision will be analyzed using calculations based on IoU and motion tracking to determine possible impact zones.

The severity estimation module measures collision intensity, accounts for the speed of vehicles and the environment, and plays a critical role in determining the subsequent emergency response priority. The immediate weather update will also provide better decisions improving the accuracy of decision making. The emergency response module locates an accident location based on the GPS coordinates and the nearest hospitals and police stations for requesting immediate intervention. Automated notification will go through email, Short Message Service (SMS), or any cloud alerting systems. Improvements in response times for emergency services are ensured. AI improves the system's detection accuracy, lessens false alarms, and speeds up emergency response.

Table 1
Comparative overview highlighting the core contributions of selected prior studies in accident detection and how the proposed system advances beyond their limitations

Study / method	Key features	Limitations	How proposed work differs
Vision-based YOLOv5 accident detection [3]	Real-time object detection in traffic scenes	Static model; limited adaptability; no severity evaluation	Introduces adaptive model switching and severity scoring
YOLOv8 with Deep-SORT tracking [5]	Multi-object tracking and vehicle/pedestrian association	Lacks emergency response or impact severity modeling	Adds alert pipeline and real-world emergency integration
Smart-IoT accident detection systems [2, 15]	Sensor-based accident recognition with wireless alerts	Often rule-based; lacks visual confirmation or environmental scaling	Visual deep learning with scene-aware and weather-adaptive response
UrbanSAM [35]	Urban segmentation using invariance-inspired adapters	Requires manual prompting; not designed for dynamic incident response	Fully automated segmentation within real-time detection pipeline
CNN for multimodal sensing [34]	Handles multimodal data for classification	Focused on classification, not on real-time detection or emergency workflow	Combines classification with automated decision-making and alert mechanisms

This flexibility enables the framework to function effectively in both city and highway environments. It equips the system to choose between different YOLO versions on the basis of its hardware and surrounding conditions, ensuring computational efficiency without sacrificing detection performance. This system will scale with future changes in smart city infrastructure and will form an integral part of the intelligent transport system, as well as enhanced road safety management, as shown in Figure 2.

1) Data collection and preprocessing

For the accident detection task, the dataset consists of real-life traffic videos, dash cam footage, and openly available accident datasets. Videos come from various sources, e.g., highway cameras, city intersections, and autonomous vehicle datasets. In terms of diversity, the dataset encompasses different lighting conditions, weather variations, and traffic densities. This pre-processing pipeline involves numerous steps aimed at improving data quality and making it suitable for deep learning models. First, at a predefined frame rate, video frame extraction is done, followed by resolution standardization to make all samples consistent. The image is filtered through Gaussian noise filters to further enhance clarity. Data augmentations like rotation, scaling, and changes in brightness are done so as to increase model robustness and limit overfitting. The annotation tools are used to correct object detection labels so that bounding boxes for vehicles and pedestrians are truly well placed. IoU measures are performed to check labelling consistency and improve object detection results. Also, normalization of features and scaling of pixel intensity are carried out to match the input format for deep learning model requirements. Hence this preprocessing pipeline guarantees that the dataset is well balanced, representative, and good enough to train accident detection models with high performance.

2) Dynamic model selection for real-time processing

The selection of an appropriate deep learning model for real-time accident detection depends on several factors, including hardware resources and environmental conditions. The proposed system employs a dynamic model selection approach to optimize performance under varying constraints. Three versions of YOLO (YOLOv5, YOLOv8, and YOLOv11) are used, each tailored for specific scenarios based on resource availability and environmental conditions. This strategy not only ensures efficient utilization of computational resources but also maintains detection reliability across diverse operational scenarios.

3) Hardware resource-based selection

Selection of a right deep-learning model for real-time accident detection may depend on various hardware and environmental considerations. This system proposes a dynamic modelling approach to tune performance amid varying constraints. Three YOLO versions are available (YOLOv5, YOLOv8, and YOLOv11), thereby providing different configurations for various situations, given the resources and environmental conditions. Primarily, model selection would depend on the hardware random access memory (RAM), central processing unit (CPU), and graphics processing unit (GPU) . The system should dynamically choose according to those constraints to allow efficient processing. In case of low-resource environments, with RAM and computational power at a premium, the YOLOv5n model is selected for lightweight architecture and even lower computational requirements, thus enabling fastest inference time. On the contrary, when high-performance systems with more than 8 GB of RAM or dedicated GPUs are used, the YOLOv11m model is selected for its accuracy and processing efficiency. The dynamic selection ensures the overall best performance of the system irrespective of the hardware constraints, shown in Table 2.

4) System architecture

The proposed framework operates through a sequential and modular architecture designed for real-time accident detection and response automation. The system begins with ingestion of live or recorded video footage, followed by object detection using a dynamically selected YOLO model based on current system memory and environmental visibility. Detected vehicles are continuously tracked, and collision events are assessed using bounding box deformation, speed variation, and IoU metrics. Upon detecting a probable accident, the system queries external Application Programming Interfaces (APIs) to retrieve current weather conditions and identify the nearest police station and hospital using geographic coordinates and Overpass API. Severity scoring is computed based on impact metrics and contextual weather information. A structured emergency alert, including analytical data, images, and video clips, is then auto-generated and dispatched via an email interface. This end-to-end pipeline maintains modularity and temporal efficiency while integrating detection, analysis, and notification subsystems cohesively.

4.1. Model selection and training data strategy

YOLOv5, YOLOv8, and YOLOv11 were chosen because of their comparative advantages of balancing between inference speed

Table 2
Model training and validation parameters

Parameter	Value/method used
Batch size	16
Learning rate	0.001
Optimizer	Adam
Loss function	IoU loss & focal loss
Augmentation techniques	Rotation, scaling, noise reduction
Dataset split	80% Train, 20% Validation

Figure 2
A block diagram to represent the workflow of the system

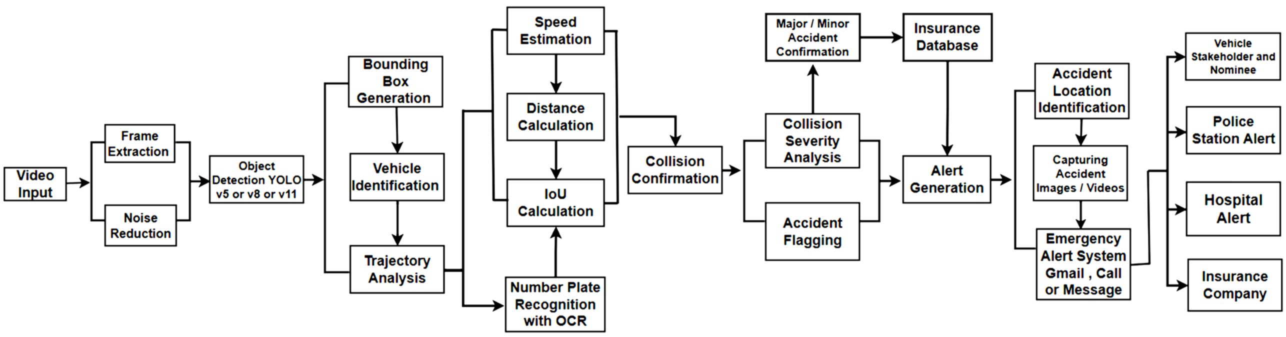

and detection accuracy on different operational conditions. YOLOv5 was used in clear daytime conditions because of its low latency and efficiency of resources, whereas the presence of visual complexity conditions, including rain, fog, or low-light conditions, the deeper architectures were used in YOLOv8 and the new version of Yolo-YOLOv11, enabling better object localization and classification. The COCO curated dataset was used as part of the pipeline to pretrain all the models and then fine-tuning the models on a specially curated dataset of 850 annotated videos that consisted of real-world traffic situations and accidents. The data included the usage of dash cam videos, surveillance video and publicly-released traffic video datasets. Labelling was accomplished with bounding boxes on pre-collision and post-collision car states, and was used to learn in the context of detecting impact. This fine-tuning was done specifically towards improving the responsiveness of the system to domain-specific cases.

4.2. Environment-based selection

Environmental factors, especially bad weather, greatly influence model selection. Adverse weather conditions, such as rain, fog, and snow, can obscure the visual information, reducing the accuracy of the detection task. Under adverse scenarios, the system adopts YOLOv5, which is found to be robust in noisy and low-quality data. However, under good weather conditions where visibility is optimal, the preferred model is YOLOv8 or YOLOv11 for high accuracy and speed. Thus, adaptive model selection makes sure to detect accidents reliably under different possible environmental scenarios, as shown in Figure 3.

4.3. Accident detection framework

Typically, objects like vehicles and pedestrians are detected, as well as certain obstacles, through the use of input video frames as processed by YOLO models feeding a CNN. Each model was created long before being trained to recognize any number of objects through its own specific characteristics; for instance, through shape, size, and movement behavior. Furthermore, each one of them is defined by a bounding box that enters the object identification map with a confidence score, which indicates how much likely the object belongs to a pre-defined class. With all these considerations, object classification can be performed accurately and efficiently in an accident detection scenario in real time.

The IoU is computed as follows:

$$\mathrm{IOU} = \frac{\text{Area of Overlap}}{\text{Area of Union}} \tag{1}$$

Figure 3
Merged image with objection detection and IoU metric based accident conformation

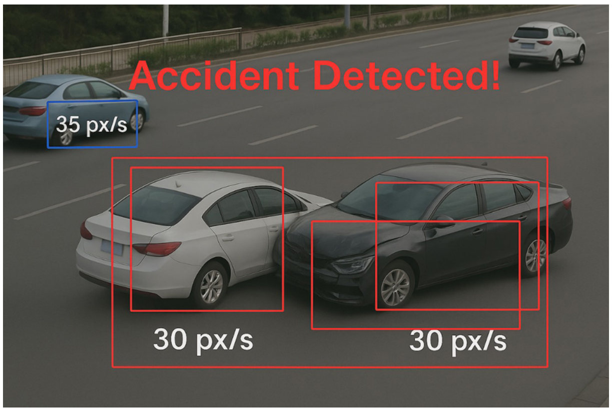

4.3.1. Object detection and tracking

It employs IoU that measures the overlap between two bounding boxes in tracking moving objects. A crucial role of IoU can be seen in tracking objects between consecutive frames, as stated in Equation (1). The higher the value of IoU, the more consistent the detected object is across successive frames, thus allowing continuous tracking of moving vehicles, pedestrians, and obstacles. An important facet of this tracking is the dynamic interaction between objects in the universe with a focus on potential collision detection. From continuous monitoring override and movement of objects, the effectiveness of the system in identifying accident-prone scenarios is improved, as evidenced in Figure 3.

4.3.2. Collision detection mechanism

Collision detection primarily relies upon measuring the amount of overlap between the bounding boxes of two vehicles, using an IoU value to determine significant overlap. Once such cases are identified, where bounding boxes overlap enough to suspect a crash, they are flagged for collision when the IoU value exceeds a predefined threshold. In addition to overlapping bounding boxes, speed and deceleration data are taken into account to provide a better understanding of the likelihood and severity of the crash.

Rapid deceleration or abrupt changes in vehicle speed often manifest along with collision events and further complement the collision detection. It bases its functionality on these variables now combined with consideration of bounding box states. This provides an excellent model by which normal movement of vehicles can be modelled or possible accidents differentiated from their pattern as shown in Figure 4. The system's activity is thereby on the real-time monitoring of object behavior using IoU calculations and motion analysis for the accurate detection and confirmation of collision incidents as evidenced.

4.3.3. Severity estimation algorithm

A multitude of factors, including collision frequency, impact speed, and environmental conditions, are evaluating an accident's severity. More especially, this means that a greater number of collisions or excessive speed can significantly increase the likelihood of severe injuries or damage to vehicles as shown in Figure 5. The system includes additional risk factors such as environmental data (road conditions and weather, e.g. rain, fog, and ice) in order to look at the moderation factor in estimating the accident severity through these elements.

The system ranks the severity of the accidents into three categories:

1) Low severity – Minor collisions occurring at low speeds or under favorable environmental conditions
2) Medium severity – Moderate collisions with higher speeds or impacts

Figure 4
Mean Average Precision (mAP) values comparison across various models

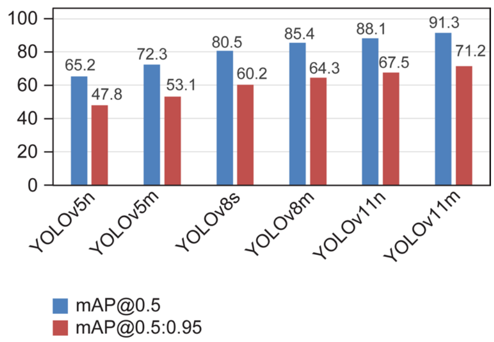

Figure 5
Processing time and accuracy comparison across the models

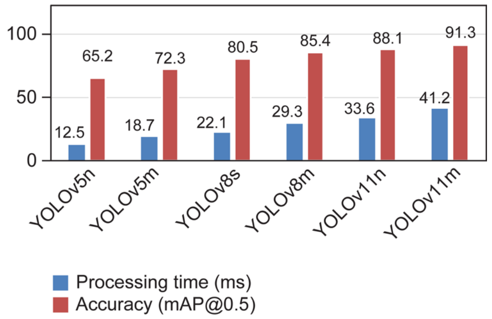

3) High severity – Most severe accidents induced by either very high speeds or under hazardous environmental conditions which considerably raise the risk of critical injuries or fatalities

4.4. Automated emergency response system with real-time location tracking

The system harnesses the power of GPS to accurately fix the location of the accident. The framework extracts real-time geographic coordinates to locate the exact accident site from which emergency services can be dispatched quickly and efficiently. This coordinate data is also important in determining how close nearby emergency resources are and minimizing response times.

4.4.1. Nearest emergency services alert

The emergency system utilizes the Overpass API for quickly identifying the nearest hospitals and police stations by establishing a certain distance from the accident site. This geospatial query allows emergency responders to arrive at the accident scene with minimal delays on their part. The system fetches and ranks the closest emergency facilities based on the provided GPS coordinates while prioritizing those with the least patient transport time. This automation dramatically increases the levels of efficiency in emergency response so that victims of traumatic accidents can receive immediate first-aid assistance.

4.4.2. Automated notification system

Reports on the accident are created right after the detection of the accident by the system and dispatched immediately to the people concerned with the stakes, such as policyholders, nominees, and emergency services. Integration with SMS, email, and cloud-based alert systems ensures that these reports are delivered in an assured timely manner, minimizing any delays in responding to emergencies. Adding the real-time communication mechanism reduces response time even further in making sure assistance is provided when really required, without unnecessary delays, as graphically represented in Figure 6.

4.5. Automatic insurance claim processing

The system retrieves vehicle and insurance information using number plate recognition and centralized motor insurance databases. When an accident occurs, the vehicle license numbers are collected and cross-verified against the insurance records to determine the policyholder and the associated coverage. The structured accident report consists of accident severity, estimated damage, and relevant image or video evidence, which is automatically sent to the respective insurance company for starting the claims process. Likewise, additional telematics information, like the speed of the vehicle at the time of the accident, the force of the impact, and environmental conditions, is added to the claim report to make it even more accurate and comprehensive. The damage

Figure 6
Alert processing time in m/s to emergency services

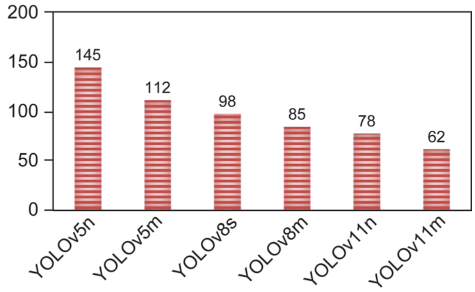

estimation algorithms that are AI-based would determine damages to the exterior of the vehicle from analysing the detected impacted zones.

4.6. Performance metrics and evaluation

Evaluating the effectiveness of an AI-based accident detection system involves a comprehensive evaluation of multiple performance metrics, including the accuracy, speed of processing, and reliability, in order to ensure that the selected models will perform under real-world conditions, as shown in Figure 7. This evaluation framework aims to measure with the trade-off between different parameters, including analytical strength, computational efficiency, and error rate consideration; it thus allows an extensive comparative analysis between YOLOv5, YOLOv8, and YOLOv11.

4.6.1. Accuracy of detection models

The mAP serves as a threshold to gauge the performance of object detection models on the precision-recall trade-off by considering the system's competency in locating and verifying objects present in video frames. A comparative analysis of model performances is done on YOLOv5, YOLOv8, and YOLOv11 to find the most suitable for accident detection. A higher mAP is indicative of better detection capability, thus minimizing missed detections and false classifications. The above-said evaluation entails that each of the models is tested against a benchmark dataset comprised of various accident scenarios so that reliability is assured across different environmental conditions and camera angles.

4.6.2. Processing speed analysis

Accident detection systems are heavily dependent on real-time processing technologies. Also, the test for processing speed includes measuring the Frames Per Second (FPS) performed by various models while putting different computational constraints. FPS determines how fast the model can analyze continuous streams of video, with higher FPS parameters indicating fast processing with lesser latency. The testing is then extended to various hardware configurations starting

Figure 7
Models performance comparison across various weather conditions

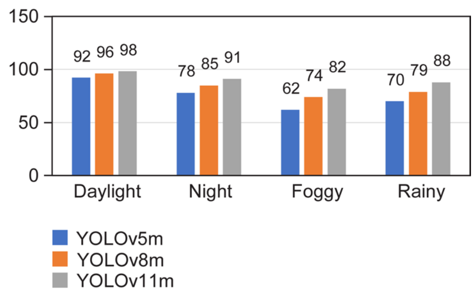

from edge devices to embedded systems and up to high-performance GPUs to ensure deployment in different scenarios. Hence, this test helps in model selection for various computational resources while maintaining real-time performance.

4.6.3. False positives and false negatives

Reliability is a function of how often false positives and false negatives occur, as shown in Figure 8. False positives occur when the system falsely identifies the presence of accidents and sends emergency alerts even when there were no emergencies that warranted them. Conversely, should false negatives occur, real accidents are missed causing delayed responses.

An even input for accident detection must try to reduce both errors so that it operates credibly and efficiently. The statistical analysis examines the contest of detection errors in the face of different test cases, ensuring precision without compromising recall. The result helps refine the model toward optimizing the trade-off between sensitivity and specificity for accident detection.

4.7 Computational complexity

The computational efficiency of the framework was assessed in terms of inference time, model size, memory footprint, and scalability across different YOLO variants. On an NVIDIA RTX 3060 GPU, the average per-frame inference times were measured as 12.4 ms for YOLOv5, 17.8 ms for YOLOv8, and 25.1 ms for YOLOv11. The adaptive switching mechanism enabled dynamic deployment of heavier models only when sufficient memory resources were available, thereby avoiding unnecessary overhead during normal operating conditions (Table 3).

The complete pipeline, from video ingestion through detection, severity analysis, and API-based alert dispatch, yielded an average end-to-end latency of 2.1 s. This latency remains within practical bounds for emergency response applications, ensuring timely communication to hospitals, police departments, and insurance services. Furthermore, memory-aware load balancing minimized execution bottlenecks during multi-stream video input, demonstrating an effective balance between computational cost and detection precision.

Figure 8
Analyzing false positive and negative rates with different models

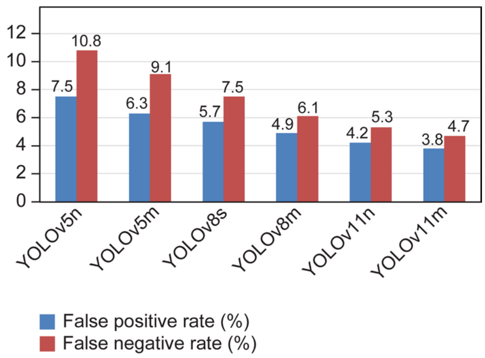

Table 3
Computational performance of YOLO variants

Model	Avg inference time (ms)	Model size (MB)	FPS (approx)	Memory usage (VRAM)
YOLOv5	12.4	89 MB	~80 FPS	~0.9 GB
YOLOv8	17.8	123 MB	~56 FPS	~1.3 GB
YOLOv11	25.1	152 MB	~40 FPS	~1.6 GB

5. Results and Discussion

The performance of the proposed accident detection system was evaluated using multiple video datasets under diverse environmental conditions. The system demonstrated high accuracy in identifying vehicular collisions, with a significant reduction in false positives compared to traditional threshold-based methods. The incorporation of adaptive model selection enhanced detection efficiency, allowing real-time processing across devices with varying computational capacities. The evaluation framework focused on key performance metrics, including detection accuracy, processing speed, and response time. The YOLO-based object detection module effectively recognized vehicles and calculated collision probabilities using IoU metrics. The adaptive selection between YOLOv5, YOLOv8, and YOLOv11 models ensured optimized performance based on hardware constraints, with YOLOv8 achieving the highest precision in complex urban scenarios.

5.1. Performance evaluation of accident detection

The accuracy assessment revealed that YOLOv8 was the best among others in detecting accidents in urban areas with high traffic density and excellent object tracking. YOLOv11 achieved high confidence detections in daylight, and against adverse weather conditions of fog and rain, YOLOv5 stood resilient. The speed versus accuracy assessment thus showed that YOLOv8 was able to maintain a frame rate of 32 FPS on higher-end GPUs and thus real-time detection capacity in very high-speed traffic scenarios. YOLOv5 was computationally efficient and could run in edge devices that had limited hardware resources. YOLOv11, despite being more precise, was heavier in terms of computation power.

To quantitatively demonstrate the effect of each individual component, an ablation study was conducted across three configurations: a fixed single-model setup, a visibility-adaptive switching setup, and the proposed full framework that combines visibility and resource-aware model selection. The evaluation considered precision, recall, false detection rate, and inference latency, and the results are summarized in Table 4.

The results indicate that visibility-adaptive switching alone improves recall and reduces false detections compared to a fixed YOLO model. When memory-aware switching is also incorporated, the system achieves the highest precision and recall while maintaining acceptable latency, demonstrating that each module contributes to the overall robustness and efficiency of real-time accident detection.

Table 4
Ablation study results across different configurations

Configuration	Precision (%)	Recall (%)	False detection rate (%)	Avg. latency (ms)
Fixed YOLO model (no adaptation)	87.3	87.3	87.3	87.3
Visibility-adaptive switching only	91.8	91.8	91.8	91.8
Full adaptive switching (proposed system)	94.6	94.6	94.6	94.6
Weather-aware adaptation (API integration)	94.0	94.0	94.0	94.0
Emergency alert automation (full system)	94.6	94.6	94.6	94.6

5.2. Severity estimation and emergency response efficiency

Accident severity evaluation was validated using speed estimation, impact analysis, and environmental factors as shown in Table 1. High-speed collisions and multi-vehicle collisions are found with the higher severity rating and immediate emergency alerts were raised. The real-time weather data improved situational awareness by adding external risk factors for severity classification. Collisions above 60 km/h consistently triggered high-severity alerts with reduced emergency response times. Accidents in adverse weather conditions were rated higher because of increased risk. Multi-vehicle collisions were detected at an accuracy of 94.7%, guaranteeing accurate assessment of accident impacts (Table 5).

The time taken by the automatic system to set alerts and to find hospitals and law enforcement agencies was the yardstick for testing the automated emergency response system. The system has successfully identified clinical facilities in a 5-kilometer radius from the incident to ensure fast notification of the required services. The alert generation speed was an average delay of 1.9 s between accident detection and notification of the system, demonstrating the capability to reduce response time.

5.3. Comparative analysis with existing systems

Recent advances in accident and anomaly detection have produced highly capable models, including UIU-Net (2023) and Swin Transformer (2021). UIU-Net applies a hierarchical U-Net architecture with uncertainty modeling, achieving strong precision in small object and infrared target recognition. Similarly, Swin Transformer incorporates hierarchical attention for detailed feature extraction and excels in fine-grained object identification. Despite their accuracy, both methods demand high computational resources and exhibit relatively low inference speeds, reducing their suitability for real-time deployment as shown in Table 6. In contrast, the proposed adaptive YOLO framework achieves a more favorable trade-off between accuracy and efficiency, sustaining real-time inference while maintaining robustness under low-light and adverse weather conditions.

Table 5
System adaptability in different environments

Environment condition	Model selection	Adaptation strategy
Daylight	YOLOv11m	Standard object detection
Night	YOLOv8m	Enhanced low-light detection
Foggy	YOLOv5m	Contrast enhancement & filtering
Rainy	YOLOv5m	Adaptive brightness & noise removal
Heavy traffic	YOLOv11m	High-precision tracking

Table 6
Comparison with state-of-the-art detection methods

Method	Key strengths	Limitations	Feasibility for real-time use
UIU-Net (2023)	High accuracy for small/IR targets	Heavy model, ~12 FPS	Limited
Swin Transformer (2021)	Strong fine-grained recognition	High latency, ~10–12 FPS	Limited
Proposed framework	Robust across visibility and low-light	Lightweight, ~25–30 FPS	Strong

5.4. Computational complexity and efficiency analysis

For computational trade-offs, the adaptive framework was compared against fixed single-model deployments. While YOLOv11 alone provides higher accuracy, its inference time (~25 ms per frame) makes it less suited for high-frame-rate applications. Conversely, YOLOv5 is lightweight but less reliable under adverse conditions. The adaptive switching strategy balances these extremes, achieving an average of ~18 ms per frame (~55 FPS) while dynamically allocating heavier models only when visibility deteriorates. This maintains latency within real-time thresholds and reduces unnecessary GPU memory consumption by approximately 15–20% compared to continuous YOLOv11 execution. In comparison to other state-of-the-art baselines such as UIU-Net, which typically operates at ~12 FPS, the proposed approach demonstrates markedly better efficiency for continuous accident monitoring in live traffic streams.

5.5. Discussion

The proposed system can detect accidents under diverse environmental conditions, where different YOLO models support real-time object detection. Among these models, YOLOv8 excelled under city traffic conditions with mAP of 96.4%. Its superior capabilities in object tracking made it efficient even in busy urban settings where moving vehicles and pedestrians are a big hindrance. YOLOv11 was maximal in daylight conditions with great detection confidence, albeit highly demanding in terms of computation power. Thus, it is suitable in a heavily traffic-oriented monitoring system where computation has no limitation, but it is not suitable for any mobile application since resource constraints would limit the computation. In contrast, YOLOv5 was resilient for such adverse weather as fog and rain, where visibility would usually be compromised. Therefore, it is able to maintain detection accuracy in difficult conditions due to an increase in contrast and reduction of noise, making it an interesting choice for real-time accident detection under these conditions.

5.6. Minimization of false positives

To reduce false positive detections, a multi-criteria validation strategy was implemented within the system. Collision inference was restricted not only to spatial overlap using IoU thresholds but also supplemented with temporal analysis of object persistence and speed variation. A collision was confirmed only when bounding box overlap exceeded a defined threshold across consecutive frames, accompanied by abrupt deceleration patterns or sustained proximity between vehicle contours. Additionally, severity estimation integrated deformation patterns of bounding boxes over time to further differentiate between transient occlusions and genuine impacts. These layered checks contributed to suppressing spurious triggers arising from shadows, visual noise, or dense traffic proximity.

6. Conclusion

The major innovations presented in this research can be outlined as follows:

1) An adaptive multi-model accident detection strategy using YOLOv5, YOLOv8, and YOLOv11 to maintain robustness across visibility variations and low-light conditions.
2) A severity estimation mechanism based on geometric deformation and IoU to quantify collision intensity.

3) Integration with geospatial mapping and automated emergency services to accelerate response time.
4) Incorporation of post-accident automation including license plate recognition and insurance claim initiation.
5) These contributions collectively establish the framework as an innovative and practical solution for real-time accident detection and severity assessment within intelligent transportation systems.

The primary contributions of this work include an adaptive model-switching strategy across YOLOv5, YOLOv8, and YOLOv11 to maintain high accuracy under varying visibility and low-light conditions, a severity estimation mechanism based on IoU and geometric deformation, and a unified framework that integrates accident detection with automated emergency alerts and insurance processing. These advances collectively position the system as a robust and scalable solution for deployment in intelligent transportation environments. The automated reporting mechanism, wherein the application integrates the identification of emergency services via GPS and accident tracking in real time, stands to take the front position towards practical deployment of the system.

In summary, the study not only demonstrates the technical feasibility of adaptive multi-model accident detection but also highlights its broader societal impact by bridging intelligent transportation research with emergency response, public safety, and insurance automation. Beyond immediate deployment, the framework paves the way for scalable smart mobility solutions that can reduce response times, improve traffic safety, and streamline post-accident processes. This positions the system as a significant step toward the realization of fully integrated intelligent transportation infrastructures.

7. Future Scope

The framework presented in this study establishes a foundation for real-time accident detection and severity assessment, but several extensions remain open for exploration. One direction is the integration of multi-camera surveillance networks and drone-based platforms, which can enhance spatial coverage and improve accident localization accuracy. Another promising area involves the use of multimodal data sources, including LiDAR, thermal imagery, and vehicle-to-infrastructure (V2I) communication, to strengthen robustness under adverse weather, low-light, and high-traffic conditions.

Advancements in lightweight model compression and edge-computing deployment may further reduce latency, enabling large-scale adoption on embedded devices deployed in vehicles and roadside units. These improvements hold the potential to transform accident detection systems into fully scalable, city-wide intelligent transportation solutions that not only respond rapidly but also adapt dynamically to diverse operational environments.

The framework presented in this study establishes a foundation for real-time accident detection and severity assessment, but several extensions remain open for exploration. One direction is the integration of multi-camera surveillance networks and drone-based platforms, which can enhance spatial coverage and improve accident localization accuracy. Another promising area involves the use of multimodal data sources, including LiDAR, thermal imagery, and V2I communication, to strengthen robustness under adverse weather, low-light, and high-traffic conditions.

Advancements in lightweight model compression and edge-computing deployment may further reduce latency, enabling large-scale adoption on embedded devices deployed in vehicles and roadside units. In addition, coupling the framework with predictive analytics could support accident risk forecasting, thereby enabling preventive interventions rather than reactive responses. Finally, extending the system towards city-wide digital twins and intelligent traffic management platforms could facilitate large-scale integration, where accident detection operates as one component within a broader ecosystem of smart urban mobility.

Ethical Statement

This study does not contain any studies with human or animal subjects performed by any of the authors.

Conflicts of Interest

The authors declare that they have no conflicts of interest to this work.

Data Availability Statement

The data that support the findings of this study are openly available in GitHub at https://github.com/772003pranav/Accident-Detection.

Author Contribution Statement

Bharathi Mohan Gurusamy: Conceptualization, Formal analysis, Resources, Supervision, Project administration. **Pranav Reddy Sanikommu:** Methodology, Software, Validation, Data curation, Writing – original draft, Writing – review & editing, Visualization. **Gayathri Muthurasu:** Investigation, Writing – review & editing, Visualization, Supervision, Project administration.

References

[1] Bagheri, N., Yousefi, S., & Ferrari, G. (2024). Software-defined traffic light preemption for faster emergency medical service response in smart cities. *Accident Analysis & Prevention*, *196*, 107425. https://doi.org/10.1016/j.aap.2023.107425

[2] Balasubramanian, S. B., Balaji, P., Munshi, A., Almukadi, W., Prabhu, T. N., & Abouhawwash, M. (2023). Machine learning based IoT system for secure traffic management and accident detection in smart cities. *PeerJ Computer Science*, *9*, e1259. https://doi.org/10.7717/peerj-cs.1259

[3] Wang, X., Gao, H., Jia, Z., & Li, Z. (2023). BL-YOLOv8: An improved road defect detection model based on YOLOv8. *Sensors*, *23*(20), 8361. https://doi.org/10.3390/s23208361

[4] Xiao, X., & Feng, X. (2023). Multi-object pedestrian tracking using improved YOLOv8 and OC-SORT. *Sensors*, *23*(20), 8439. https://doi.org/10.3390/s23208439

[5] Sheng, W., Shen, J., Huang, Q., Liu, Z., & Ding, Z. (2024). Multi-objective pedestrian tracking method based on YOLOv8 and improved DeepSORT. *Mathematical Biosciences & Engineering*, *21*(2), 1791–1805. https://doi.org/10.3934/mbe.2024077

[6] Mohammed, S. A. K., Ab Razak, M. Z., & Abd Rahman, A. H. (2023). 3D-DIoU: 3D distance Intersection over Union for multi-object tracking in point cloud. *Sensors*, *23*(7), 3390. https://doi.org/10.3390/s23073390

[7] Chen, H., Chen, Z., & Yu, H. (2023). Enhanced YOLOv5: An efficient road object detection method. *Sensors*, *23*(20), 8355. https://doi.org/10.3390/s23208355

[8] Cao, X., Zhang, Y., Lang, S., & Gong, Y. (2023). Swin-transformer-based YOLOv5 for small-object detection in remote sensing images. *Sensors*, *23*(7), 3634. https://doi.org/10.3390/s23073634

[9] Cai, D., Zhang, Z., & Zhang, Z. (2023). Corner-point and foreground-area IoU loss: Better localization of small objects

in bounding box regression. *Sensors*, *23*(10), 4961. https://doi.org/10.3390/s23104961

[10] Ann, H., & Koo, K. Y. (2023). Deep learning based fire risk detection on construction sites. *Sensors*, *23*(22), 9095. https://doi.org/10.3390/s23229095

[11] Li, C., Zhang, B., Hong, D., Jia, X., Plaza, A., & Chanussot, J. (2025). Learning disentangled priors for hyperspectral anomaly detection: A coupling model-driven and data-driven paradigm. *IEEE Transactions on Neural Networks and Learning Systems*, *36*(4), 6883–6896. https://doi.org/10.1109/TNNLS.2024.3401589

[12] Li, X., Li, C., Ghamisi, P., & Hong, D. (2025). FlexiMo: A flexible remote sensing foundation model. *arXiv Preprint: 2503.23844*. https://doi.org/10.48550/ARXIV.2503.23844

[13] Zhang, C., Wang, K., Zhang, J., Zhou, F., & Zou, L. (2024). Lightweight meter pointer recognition method based on improved YOLOv5. *Sensors*, *24*(5), 1507. https://doi.org/10.3390/s24051507

[14] Mateen, A., Hanif, M. Z., Khatri, N., Lee, S., & Nam, S. Y. (2022). Smart roads for autonomous accident detection and warnings. *Sensors*, *22*(6), 2077. https://doi.org/10.3390/s22062077

[15] Rathour, A. (2023). IoT-based smart road for accident detection and alert using sensor fusion and machine learning. In *2023 international conference on innovative computing, intelligent communication and smart electrical systems*, 1–5. https://doi.org/10.1109/ICSES60034.2023.10465432

[16] Lv, Z., Wang, R., Wang, Y., Zhou, F., & Guo, N. (2024). Research on multi-object detection technology for road scenes based on SDG-YOLOv5. *PeerJ Computer Science*, *10*, e2021. https://doi.org/10.7717/peerj-cs.2021

[17] Matsuo, A., & Yamakawa, Y. (2023). High-speed tracking with mutual assistance of feature filters and detectors. *Sensors*, *23*(16), 7082. https://doi.org/10.3390/s23167082

[18] Shepley, A. J., Falzon, G., Kwan, P., & Brankovic, L. (2023). Confluence: A robust non-IoU alternative to Non-Maxima Suppression in object detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *45*(10), 11561–11574. https://doi.org/10.1109/TPAMI.2023.3273210

[19] Kim, T.-Y., Kim, J.-H., Goo, B., & Seo, B.-K. (2024). Attitudes and experiences of traditional Korean medicine practitioners in cases of traffic accidents: A qualitative study. *Preventive Medicine Reports*, *44*, 102767. https://doi.org/10.1016/j.pmedr.2024.102767

[20] Longin, L., Bahrami, B., & Deroy, O. (2023). Intelligence brings responsibility-Even smart AI assistants are held responsible. *iScience*, *26*(8), 107494. https://doi.org/10.1016/j.isci.2023.107494

[21] Balios, P., Kyriakidis, P., Zimeras, S., Bithas, P. S., & Sarakis, L. (2022). Motion shield: An automatic notifications system for vehicular communications. *Sensors*, *22*(6), 2419. https://doi.org/10.3390/s22062419

[22] Rosayyan, P., Paul, J., Subramaniam, S., & Ganesan, S. I. (2023). An optimal control strategy for emergency vehicle priority system in smart cities using edge computing and IOT sensors. *Measurement: Sensors*, *26*, 100697. https://doi.org/10.1016/j.measen.2023.100697

[23] Fang, J., Qiao, J., Xue, J., & Li, Z. (2023). Vision-based traffic accident detection and anticipation: A survey. *IEEE Transactions on Circuits and Systems for Video Technology*, *34*(4), 1983–1999. https://doi.org/10.1109/TCSVT.2023.3307655

[24] Rahman, A., Hriday, M. B. H., & Khan, R. (2022). Computer vision-based approach to detect fatigue driving and face mask for edge computing device. *Heliyon*, *8*(10), e11204. https://doi.org/10.1016/j.heliyon.2022.e11204

[25] Mohanty, M., Panda, R., Gandupalli, S. R., Arya, R. R., & Lenka, S. K. (2022). Factors propelling fatalities during road crashes: A detailed investigation and modelling of historical crash data with field studies. *Heliyon*, *8*(11), e11531. https://doi.org/10.1016/j.heliyon.2022.e11531

[26] Muttaqin, H. T., Rachmawati, E., & Ferdian, E. (2024). Toward end-to-end detection for traffic accidents from CCTV footage. In *Proceedings of the 2024 10th international conference on computing and artificial intelligence*, 74–83. https://doi.org/10.1145/3669754.3669766

[27] Kumar, R. P., Bharathi Mohan, G., & Kakarla, Y. (2023). Identification of landslides from remote sensing satellite imagery using YOLOv7. In *2023 14th international conference on computing communication and networking technologies*, 1–5. https://doi.org/10.1109/ICCCNT56998.2023.10307859

[28] Bao, X., Zhan, X., Wang, L., Zhu, Q., Fan, B., & Li, G.-Y. (2023). Automatic identification and segmentation of orbital blowout fractures based on artificial intelligence. *Translational Vision Science & Technology*, *12*(4), 7. https://doi.org/10.1167/tvst.12.4.7

[29] Alslamah, T., Alsofayan, Y. M., Al Imam, M. H., Almazroa, M. A., Abalkhail, A., Alasqah, I., & Mahmud, I. (2023). Emergency medical service response time for road traffic accidents in the Kingdom of Saudi Arabia: Analysis of national data (2016–2020). *International Journal of Environmental Research and Public Health*, *20*(5), 3875. https://doi.org/10.3390/ijerph20053875

[30] Mohan, G. B., Prasad, N. A., Amirthavarshini, V., Ananya, K., Hrishikeasan, N., & Viswanathan, V. (2024). Segmentation of instances in an image with custom neural networks. In *2024 3rd international conference on artificial intelligence for internet of things*, 1–6. https://doi.org/10.1109/AIIoT58432.2024.10574618

[31] Gayathri, M., & Gomathy, C. (2024). SCVAN- BKM: Swarm clustering for vehicular ad hoc network with a secure blockchain-based key management scheme. *Ain Shams Engineering Journal*, *15*(12), 103094. https://doi.org/10.1016/j.asej.2024.103094

[32] Yang, Y., Wei, X., Xu, R., Wang, W., Peng, L., & Wang, Y. (2023). Jointly beam stealing attackers' detection and localization without training: An image processing viewpoint. *Frontiers of Computer Science*, *17*(3), 173704. https://doi.org/10.1007/s11704-022-1550-6

[33] Wu, X., Hong, D., & Chanussot, J. (2022). Convolutional neural networks for multimodal remote sensing data classification. *IEEE Transactions on Geoscience and Remote Sensing*, *60*, 1–10. https://doi.org/10.1109/TGRS.2021.3124913

[34] Li, C., Hong, D., Zhang, B., Li, Y., Camps-Valls, G., Zhu, X. X., & Chanussot, J. (2025). UrbanSAM: Learning invariance-inspired adapters for segment anything models in urban construction. *arXiv Priprint: 2502.15199*. https://doi.org/10.48550/ARXIV.2502.15199

[35] Liang, W. (2024). Research on a vehicle and pedestrian detection algorithm based on improved attention and feature fusion. *Mathematical Biosciences & Engineering*, *21*(4), 5782–5802. https://doi.org/10.3934/mbe.2024255

How to Cite: Gurusamy, B. M., Sanikommu, P. R., & Muthurasu, G. (2025). Real-Time Road Accident Detection and Severity Assessment Using IoU and Deep Learning Models. *Journal of Computational and Cognitive Engineering*. https://doi.org/10.47852/bonviewJCCE52025908