A SURVEY ON ACCIDENT DETECTION AND AUTOMATED EMERGENCY NOTIFICATION SYSTEM USING YOLO V8

Vipin Varghese, Annmaria Pauly, Nandana V A, Muhammed Nahshal A P

Vipin Varghese Student, Computer Science and Engineering, Holy Grace Academy of Engineering, Kerala, India
Ann Maria Pauly Student, Computer Science and Engineering, Holy Grace Academy of Engineering, Kerala, India
Nandana V A Student, Computer Science and Engineering, Holy Grace Academy of Engineering, Kerala, India
Muhammed Nahshal A P Student, Computer Science and Engineering, Holy Grace Academy of Engineering, Kerala, India
Ajith PJ Assistant Professor, Computer Science and Engineering, Holy Grace Academy of Engineering, Kerala, India
Sanam E Anto Head Of Department, Computer Science and Engineering, Holy Grace Academy of Engineering, Kerala, India

ABSTRACT

Road accidents and fire emergencies are considered two of the most critical emergencies concerning the safety and security of the general public. According to reports on the safety and security of the global transportation infrastructure, thousands of accidents take place on a daily basis, and many of these accidents lead to fatalities because of the lack of timely medical aid. Conventional emergency response systems usually rely on human reporting or communication with the emergency department. Such conventional systems face the major challenge of significant delays and the lack of accurate information regarding the accident location and type. Recent advances in the field of artificial intelligence, computer vision, and intelligent transportation systems offer new avenues for the development of automated accident detection systems. Modern surveillance systems, along with AI-powered video analysis, can effectively analyze the environment and detect unusual patterns such as accidents, fires, or individuals lying on the road. Accident Detection and Automated Emergency Notification System
The Accident Detection and Automated Emergency Notification System is a proposed solution that seeks to solve the problems that are inherent in a normal human-based accident reporting and communication to the emergency department. For instance, this proposed system would utilize a live video analysis and object detection models to detect accidents and communicate to the emergency department. In doing this, it reduces the amount of time it takes from when an accident occurs to when a response team arrives. Besides this, there is also a provision to report an accident and even check on the status of the response to that accident. Additionally, there is a provision to authenticate the rescue team and respond to emergencies appropriately.

Keyword : - *Accident Detection, Artificial Intelligence, Computer Vision, OpenCV, Deep Learning, Real-Time Monitoring, Emergency Response System, Intelligent Transportation System, IoT, Machine Learning, Road Safety, Image Processing, Automated Alert System, Surveillance Cameras, GPS Tracking, Smart City Infrastructure*

1. INTRODUCTION

In recent years, road safety has been one of the primary concerns because of the rising number of accidents on the roads worldwide. Some of the key factors causing accidents include overspeeding, distracted driving, poor road infrastructure, and human errors. Conventional approaches for reporting accidents involve communication, which sometimes results in delays in reporting the accident and providing timely relief to the accident victims. Considering the advancements in computer vision, AI, and deep learning, the possibility of developing an automated system for detecting accidents and hazardous occurrences has been achieved. This research aims to propose an Accident and Fire Detection System, which will enable the automated reporting of accidents and the sending of alerts with the location information of the accident to the accident rescue teams with the help of intelligent analysis of the accident scenarios and fire occurrences. The proposed system will be developed using the YOLO (You Only Look Once) object detection method, which will be used for the intelligent analysis of the accident scenarios and fire occurrences. The system will be developed using the Python programming language with the integration of computer vision. The proposed method aims to reduce emergency response time, improve road safety, and enhance smart city surveillance system infrastructure by offering an efficient detection tool. Moreover, it shows the potential of AI-based monitoring technologies in developing disaster management and emergency response systems in modern cities.

2. INFORMATION

Zhang, L., & Chen, Y. (2022). Real-time accident detection using deep convolutional neural networks, addressed one of the major issues in modern traffic management, which is to recognize road accidents as they occur without waiting for a human bystander to make a phone call. The authors' work presented a deep convolutional neural network model intended to analyze live feeds of video streams and automatically recognize accident scenes in real-time. What is significant in this work is the fact that it emphasized real-time response, as it is architected to recognize accident scenes as quickly as possible, as it is understood that in emergency response scenarios, every second counts and every second of delay can mean life and death. The authors' work is based on the known capability of deep CNNs to recognize hierarchical features of images, which allowed it to recognize accident scenes as they occur, as characterized by sudden deformation of vehicles, abnormal motion trajectories, and scatter of debris, which can happen at high speeds and can last only fractions of a second on screen. Experiments on benchmark traffic datasets demonstrated competitive accuracy figures and validated the model's potential for integration with the vast networks of CCTV cameras already deployed in urban environments worldwide. However, the authors candidly acknowledged that the model's hunger for large quantities of labeled training data poses a real barrier, since annotating accident footage is both time-consuming and emotionally taxing work. Additionally, the system's sensitivity to camera angle and mounting height means that a model trained on one intersection may not transfer cleanly to another, raising practical questions about how such systems would be maintained and recalibrated at scale across an entire city.

Kumar, A., & Bhatia, S. (2023). Smart emergency response through IoT and machine learning integration. proposed a solution to emergency response in terms of systems integration, where not only is the detection of emergencies not sufficient but it is essential that it is done in such a manner that it is linked with how quickly it is converted into an emergency response. They propose a framework in which there is a seamless integration of IoT devices placed in the environment and ML classifiers running in the backend servers to create an end-to-end framework that not only senses emergencies but is intelligent enough to send out notifications without requiring human intervention in between. This is necessary because of the reality of emergency response in terms of human intervention, where witnesses may or may not be present at the scene and may or may not know if what they have seen is actually worthy of calling for help. By eliminating human intervention altogether, Kumar and Bhatia's framework addresses the issue at its core. The ML component is not only limited to triggering but is intelligent enough to distinguish between actual emergencies and sensor noise. Despite these strengths, the framework carries a significant vulnerability that the authors themselves highlight: it assumes a stable and reasonably fast network connection between the sensor nodes and the backend processing layer. In rural areas, disaster zones, or locations with degraded infrastructure — precisely the environments where emergency response is most needed and most delayed — this assumption may not hold, and the system could go silent at the worst possible moment.

Li, D., & Zhao, P. (2021). Automatic fire detection from surveillance videos using hybrid CNN models. the researchers decided to concentrate on one of the most visually distinctive and devastatingly fast-moving emergencies a surveillance system must deal with: fire. Unlike vehicle accidents, which leave wreckage behind and which can be detected even minutes after the accident, fire is moving and growing rapidly, with the possibility of covering the entire view of the camera in just minutes. The researchers designed a hybrid CNN model that processes the frames of the video feed and learns to identify the early signs of fire, which include the flicker, the color gradient, and the smoke rising before the flames. The hybrid model, which combines the feature maps of the convolutional branches, was particularly designed with the aim of detecting small, nascent fires before they become dangerous, which single models fail to do. The work demonstrated reliable detection rates in controlled test environments and proposed practical integration pathways with existing CCTV infrastructure, making it accessible to organizations that cannot afford to replace their current camera hardware. The honest limitation the authors raise is that real-world surveillance backgrounds are messy — car headlights, neon signs, and sunlight reflections can mimic the color profile of fire, and dense foliage or structural elements partially obscuring a flame source can cause the model to miss detections that a trained human observer would catch immediately.

Prakash, R., & Thomas, N. (2023). AI-driven accident detection and alert systems using OpenCV. Built their system with a strong emphasis on accessibility and practical deployability, choosing OpenCV as the foundation for their video preprocessing pipeline precisely because it is widely available, well-documented, and requires no specialized hardware. Their deep learning classifier sits downstream of this preprocessing layer, receiving cleaned and normalized frame data and producing accident or no-accident decisions at speeds compatible with live monitoring. What distinguishes this work is its attention to the full pipeline rather than just the model — the authors carefully address questions of frame sampling rate, resolution trade-offs, and alert delivery latency, treating the system as a product to be deployed rather than just an experiment to be published. The system was validated against a curated dataset of real road footage from multiple camera installations and reported strong precision scores that suggest practical viability. The authors were forthright about the two environmental conditions that most severely degrade performance: night-time scenes, where low ambient light strips away the visual contrast the model relies on, and heavily occluded scenarios where a large vehicle blocks the camera's view of the accident in the critical moment of impact. Both of these are not edge cases — they are everyday realities in road surveillance — making their resolution a clear priority for follow-on research.

Singh, J., & Mehta, V. (2022). Edge computing approach for intelligent transportation safety. made a compelling argument that the conventional architecture of sending all video data to a centralized cloud for processing is fundamentally mismatched with the time-critical nature of accident detection. Every millisecond spent transmitting data to a distant server and waiting for a response is a millisecond during which an injured person is not receiving help. Their proposed solution repositions the intelligence at the edge — deploying lightweight, optimized inference models directly on roadside computing nodes that sit physically close to the cameras they serve. This architectural change not only significantly cuts latency but also eases bandwidth requirements and provides more robustness since the edge node can function even if it is temporarily disconnected from its central cloud node. The authors gave detailed benchmarks showing how well their approach compared to a cloud-only approach. They demonstrated quantified and significant performance benefits in terms of end-to-end response times. The issue that the authors acknowledged is one that everyone who has ever had to manage hardware infrastructure will understand at once: it is difficult to keep dozens or hundreds of edge nodes updated and running in the outdoors in various states of bad weather and with potential for vandalism and power supply issues.

] Rahman, M., & Alam, S. (2021). A machine learning-based framework for automatic road incident detection using CCTV feeds designed their road incident detection framework with modularity as a guiding principle, recognizing that no single machine learning algorithm performs best across all road environments and all types of incidents. Their pipeline combines handcrafted visual features — motion vectors, scene change metrics, and vehicle density estimates — with ensemble classifiers that can be swapped or retrained depending on the deployment context. This flexibility is genuinely valuable in practice, where a system deployed at a busy highway interchange must behave differently from one monitoring a narrow urban alley. The framework was evaluated on publicly available traffic datasets and demonstrated competitive detection accuracy, with the modular design allowing the authors to isolate which components contributed most to performance. The limitation that the authors and the broader research community have grappled with is the tension between handcrafted features and deep learning — manual feature engineering is interpretable and lightweight but requires domain expertise to design well, while deep learning automates feature discovery at the cost of computational weight and reduced interpretability. As deployment environments diversify, the need to re-engineer features for each new context becomes a scaling bottleneck.

] Patel, K., & Das, P. (2023). Deep learning-enabled accident severity classification for smart cities. asked a question that most accident detection papers do not address: once you know an accident has happened, how bad is it? The answer to that question should determine whether one ambulance is dispatched or three, whether the nearest hospital is simply notified or placed on full trauma alert, and how traffic management responds to prevent secondary collisions. Their deep learning model was trained to classify detected accidents based on severity levels such as minor, moderate, and severe using visual features such as the degree of deformation of vehicles, number of vehicles involved, and presence of pedestrians in the accident zone. The output of this model was integrated into a simulated smart city dispatch system and shown to have positive impacts on efficiency in scenarios where multiple accidents occur, preventing scenarios where multiple resources are deployed to a minor accident while another severe accident is under-resourced. The actual challenge the authors of this paper had to deal with is quite honest and shows they are not trying to oversell their model: severe accidents are rare and therefore underrepresented in data sets, which means that a model trained on this data is overconfident in its classification of the majority class.

Reddy, S., & Nair, G. (2020). Smart emergency notification via GPS and GSM modules in IoT-based accident systems. took a fundamentally different philosophical approach from the vision-based papers surrounding them — instead of trying to see an accident happen, they chose to feel it. Their system uses accelerometers and gyroscopes mounted inside vehicles to detect the sudden, violent changes in motion that characterize a collision, treating the vehicle itself as a sensing platform rather than relying on external infrastructure. When an impact event is detected, the system cross-references onboard GPS data to obtain precise coordinates and immediately dispatches an automated SMS alert to emergency services via a GSM module, all without requiring the driver — who may be unconscious — to take any action. This is a deeply human-centered design insight: the scenarios where automated reporting is most critical are precisely the scenarios where manual reporting is least likely to happen. The system was field-tested and demonstrated reliable detection of significant collisions. The practical limitation that field testing revealed is the difficulty of calibrating impact detection thresholds that are sensitive enough to catch real accidents without being triggered by the aggressive braking, pothole impacts, and rough road surfaces that are part of normal driving, particularly in regions with poor road infrastructure.

Wang, T., & Liu, C. (2024). Vision-based human fall and accident detection in outdoor environments. undertook a noble and ambitious idea of unified detection, in which a deep learning model is utilized to identify two different classes of outdoor emergencies, namely human falls and vehicle accidents, from a common feed of a surveillance camera. This is a practically important problem, as it is expensive to install and maintain a series of cameras to detect emergencies. Having a camera that can detect multiple classes of emergencies is significantly more useful and valuable than a series of cameras that can detect a single class of emergencies individually. Moreover, this is a diverse range of scenes that are outdoors and include a variety of lighting and seasons, making this a realistic evaluation of this model compared to many of the papers in this area. Multitasking allowed this model to learn a common representation of scenes, such as a sudden stop in motion that is characteristic of a human falling and a car accident. The authors found that while the model performed strongly under the controlled conditions of their test set, performance degraded noticeably in scenes with extreme weather, such as heavy rain or fog, where the visual cues the model relies upon are physically obscured.

Anusha, P., & Raj, M. (2022). Automated rescue alert system using AI-powered surveillance cameras. designed their AI-powered surveillance system with a specific deployment context in mind — smart campuses and institutional facilities where conventional manual security monitoring is staffed inconsistently and where delayed emergency response to an accident or fire could have severe consequences for students, staff, or visitors. The system monitors camera feeds continuously, applies a trained detection model to identify accident and fire events, and routes alerts directly to emergency service channels with enough contextual information — location, event type, timestamp — for responders to act immediately without needing to ask clarifying questions. The system was validated in a real campus deployment and demonstrated low false positive rates during daytime operation, building the kind of operational trust that is essential before institutions will rely on automated systems for genuine emergency notifications. The deployment revealed two practical lessons: first, that blind spots in camera coverage create dangerous gaps in detection coverage that no amount of algorithmic sophistication can compensate for; and second, that under high concurrent load — such as during a campus-wide event — processing latency can increase enough to meaningfully delay alerts, suggesting the need for load-aware resource scaling.

Yadav, P., & Sharma, H. (2021). Integration of cloud and IoT for real-time emergency management.. designed their emergency management architecture with the acknowledgment that "the potential of IoT sensor data is only realized if it is processed intelligently in large quantities and that cloud computing offers the required computational capabilities without requiring each deployment site to purchase server hardware." They have designed a system in which sensor nodes like cameras, smoke detectors, and vibration sensors communicate with cloud-based ML services that use anomaly detection algorithms and emergency response workflows. Certainly, one of the strengths of the design is its scalability. "To add a new monitoring site in the system, it is only required to add sensor nodes and register them with the cloud service." There was verification of the design with reduced incident-to-response times compared to manual reporting methods. The vulnerability that the authors acknowledge is both obvious and serious: when network connectivity between the sensors and the cloud is lost — due to infrastructure damage in the very emergency being detected, for instance — the system loses its decision-making capability entirely, making redundancy planning and local fallback logic an essential feature for any production deployment.

Farooq, S., & Khan, Z. (2020). Accident detection using image processing and sensor fusion techniques.. argued that the redundancy and complementary nature of multiple sensor modalities are not luxuries, but necessities for robust accident detection in realistic conditions. In their work, they fused video information from camera sensors and acceleration and rotation information from inertial sensors, using the richness of spatial information provided by visual data and the direct measurement of physical impact force, which is not available through camera sensors.

The fusion algorithm is also designed to be simple and lightweight for sensor node implementation, yet rich enough to capture the relevant correlation between the two sensor modalities. A camera sensor may detect a sudden change in traffic pattern, indicating an accident, but cannot confirm whether there was physical impact; on the other hand, an inertial sensor may detect a jarring impact but cannot confirm whether there was a collision or just a speed bump. Experiments on a collected dataset of staged collisions demonstrated improved precision and recall compared to single-modality baselines. The engineering challenge the authors highlight is the synchronization of heterogeneous data streams — even small timing offsets between sensor readings can degrade fusion performance, and maintaining synchronization reliably over extended deployment periods requires careful system design.

Borse, R., & Patil, A. (2023). AI-based vehicle crash detection with GPS-enabled alert systems.addressed a problem that is both technically interesting and humanly urgent — knowing that an accident has occurred is useful, but knowing exactly where it occurred is what enables responders to actually reach the scene. Their system pairs a trained vehicle crash detection model with real-time GPS tracking, ensuring that when a crash event is identified, the precise coordinates of the incident are immediately attached to the outgoing alert. This location-awareness transforms the alert from a vague notification into an actionable dispatch order. The system was evaluated across multiple vehicle classes and reported consistent detection performance for moderate to severe crashes, with the automatic alert generation demonstrably reducing the time required for responders to locate the incident compared to manual caller-based reporting. The limitation the authors encountered is a familiar one in GPS-reliant systems: signal quality degrades in urban canyons flanked by tall buildings and disappears entirely in tunnels, which are, ironically, among the most dangerous locations for high-speed vehicle collisions.

Chen, J., & Luo, K. (2024). Dynamic route optimization for emergency vehicles using real-time traffic data. took a new approach to focus on the minutes after detection, rather than detection itself – the time it takes for an emergency vehicle to travel from its location to the scene of the emergency. Their dynamic route optimization algorithm continually updates its knowledge of the quickest route through a stream of live data from connected infrastructure and adapts to congestion, closed roads, and even the subsequent traffic problems that may have been caused by the emergency. The system has been benchmarked on a simulated city road network and has shown substantial reductions in estimated time of arrival compared to static routing based on pre-computed shortest routes. Such reductions have very real human significance – with cardiac arrest, brain damage, and loss of blood all requiring an ambulance to arrive within three minutes of the difference between life and death. The authors discussed practical integration pathways with existing traffic management center software and identified two performance bottlenecks: the computational overhead of continuous re-optimization under peak congestion conditions, and the latency inherent in receiving traffic updates from third-party data feeds that may themselves be running on refresh cycles of tens of seconds.

] George, A., & Kannan, T. (2021). Computer vision-based fire and smoke detection using YOLOv5 architecture, turned to one of the most successful and widely adopted object detection architectures in computer vision — YOLOv5 — and adapted it specifically for the task of detecting fire and smoke in surveillance footage. The choice of YOLO-family architecture was deliberate and practically motivated: its single-pass detection design delivers real-time inference speeds that make it viable for live monitoring, whereas two-stage detection approaches, while potentially more accurate, introduce latency that is unacceptable in fire detection scenarios where speed of identification is paramount. The authors assembled a custom training dataset from surveillance camera footage of real fire incidents captured across diverse environments and fine-tuned the model's detection heads to respond to the visual signatures of fire and smoke specifically. Strong detection performance was reported for well-developed fire events with clear visual presence in the camera frame. The performance gap the authors identified for small, early-stage flames reflects a fundamental challenge in fire detection: the very moment when intervention is most effective — before the fire spreads — is the moment when the visual evidence is smallest and most easily confused with non-fire phenomena like orange lighting or reflections.

Dinesh, R., & Shankar, P. designed their emergency detection platform to solve the scaling problem that confronts any organization trying to monitor a large number of camera feeds simultaneously: centralized cloud processing of high-resolution video streams from dozens or hundreds of cameras simultaneously is bandwidth-intensive and expensive, while fully distributed on-device processing using current embedded hardware may not meet accuracy requirements. Their hybrid architecture resolves this tension by assigning time-critical detection decisions to lightweight models running locally at the camera level, reserving cloud resources for the computationally heavier tasks of aggregating events across locations, identifying spatial patterns, generating reports, and supporting human operator dashboards. The system was demonstrated in a multi-camera smart city pilot and showed both the operational benefits of the hybrid approach and its practical complexity. The authors surfaced two challenges that any large-scale video surveillance deployment must confront: the bandwidth demands of continuous video streaming even at reduced quality levels, and the growing societal discomfort with the privacy implications of always-on AI-analyzed video surveillance in public spaces.

Ahmed, N., & Hussain, R. (2020). Machine learning approaches to accident prediction and prevention. took a longer view than most papers in this survey, asking not just how to respond to accidents after they happen but whether it is possible to predict and prevent them before they occur. Their approach applies a range of machine learning techniques to historical road accident databases, learning statistical associations between accident occurrence and factors such as road geometry, weather conditions, time of day, traffic density, and historical incident frequency at specific locations. The resulting predictive models can identify high-risk zones and high-risk time windows, providing actionable intelligence for infrastructure investment decisions, traffic management interventions, and targeted enforcement. The work offers something that real-time detection systems cannot: the ability to act before the accident rather than after. The limitation the authors acknowledge is equally clear: historical patterns do not always predict future events, particularly in rapidly changing urban environments, and the retrospective nature of the approach means it cannot directly inform live emergency response to an ongoing incident.

] Zhang, Y., & Xu, L. (2022). Hybrid AI model for recognizing road hazards in real-time video feeds.. understood that there are road hazard events that involve a vehicle slowly moving towards a barrier, a pedestrian slowly moving into a fast-moving lane of traffic, a driver's behavior becoming increasingly erratic, and so forth. These are not static events and therefore cannot be determined by a single frame of a video. They are, in a sense, a series of events that unfold over time, and in order to recognize them, there must be a way to reason about what is happening over time. This is where the hybrid model shines, using CNNs to recognize the spatial features of a single frame of a video, but then using recurrent networks to recognize patterns over time. This is what ultimately helps to differentiate between a real hazardous situation and a false alarm caused by a sudden change in lighting, a sudden movement of the camera, or a large vehicle driving into the frame of view. The system was shown to perform well in terms of detection for the classes of temporal hazards that it was trained to recognize. The practical barrier the authors identify is the computational weight of the combined CNN-RNN architecture: while powerful, it demands more processing resources than simpler frame-by-frame classifiers, making real-time deployment on resource-constrained embedded hardware a significant engineering challenge.

Nirmal, J., & Varma, S. (2021). IoT-based smart rescue management with automated alert prioritization built their rescue management platform around a deceptively simple but operationally critical insight: not all emergency alerts deserve the same urgency of response, and a system that treats them all equally actually serves none of them well. When a dispatch center receives simultaneous alerts — a minor fender-bender, a pedestrian fall, and a multi-vehicle highway collision all flagged within seconds of each other — the ability to instantly rank these by severity and route each to the appropriate response resource is the difference between an efficient emergency service and a chaotic one. Their platform assigns dynamic priority scores to incoming alerts based on sensor-derived severity indicators, historical incident data for the location, and contextual factors such as time of day and proximity to hospitals or rescue stations. The prioritization engine feeds a real-time dispatch dashboard that allows operators to see all active incidents ranked by urgency and act on the most critical first. The honest limitation is that the quality of prioritization is entirely dependent on the quality and completeness of the sensor data feeding into it — a poorly calibrated sensor or a partially obscured camera at the incident site can result in a severe incident being underscored and under-resourced.

Bhattacharya, P., & Roy, D. (2023). Deep learning and IoT convergence for smart city emergency response systems envisioned their work as part of a larger smart city initiative, which is "the concept of developing smarter and more efficient cities using intelligent and connected infrastructure, which can provide more human and more humane responses to crises and emergency situations." The work of Bhattacharya and Roy is an architectural approach that "puts deep learning inference capability inside IoT devices deployed throughout the city, allowing each device to make local decisions and send structured event data to a central emergency coordination dashboard, rather than sending raw video data." This approach significantly reduces the data bandwidth requirements for city-wide monitoring and provides the "big picture" situational awareness at the central emergency coordination dashboard, while still allowing for rapid escalation of emergency incidents, which was shown through simulation. The authors were candid about two long-term operational challenges: the difficulty of ensuring that the dozens of different IoT device types from different manufacturers that populate a real city can all communicate effectively with a central platform, and the gradual drift in model performance that occurs as city conditions change over months and years without corresponding model updates.

Bhattacharya, P., & Roy, D. (2023). Deep learning and IoT convergence for smart city emergency response systems focused their work specifically within the context of the highway environment, recognizing that highway accidents constitute a particularly critical emergency scenario that warrants its own specialized and purpose-built detection capabilities. This is because of the speed at which vehicles operate on the highway, which makes the potential consequences of accident impact much more severe and critical. There is also the issue of emergency vehicle access and the potential for secondary accidents caused by in-approaching traffic. Their approach was designed around high-speed video feeds with sufficient resolution to discern meaningful visual detail from moving vehicles and a detection algorithm specifically designed to recognize visual patterns of highway incidents: the distinctive "spray" of debris kicked up by high-velocity impacts, the abnormal rotation of a vehicle in roll, and the sudden stoppage of traffic flow in an otherwise free-flowing lane. Notification of highway authorities in under a sub-second was achieved and validated in their work. They also touched upon potential integration with variable message signs for warning in-approaching drivers of potential danger, a capability with direct and immediate life-saving potential for preventing secondary accidents that often compound and contribute to the severity of highway accidents.

Menon, C .and Joseph, M. In 2024, someone suggested using an AI system to automatically spot accidents and send alerts right away on highways. The idea was to help emergency services get to the scene faster. They use a bunch of different tools—like AI, cameras that watch traffic, smart sensors, GPS, and cloud communication—to keep an eye on how cars are moving and what's happening with traffic all the time. The system looks at things like if a car quickly slows down, moves strangely, or stops unexpectedly in the road. This helps it figure out if an accident might be happening. When the system spots an accident and confirms it with different pieces of information, it automatically tells emergency services, hospitals, and traffic people where it happened, when, and how bad it is. This helps them get there quicker and work together better. This study shows that when you mix AI with systems that watch things instantly, it can really help cut down the time it takes to report accidents. This makes roads much safer for everyone. But the people who wrote this paper say the system needs good surveillance setups, accurate information, and a steady internet connection to work well. This could make it not perform as well in places where technology isn't very good.

Gupta, A., & Arora, P built their system with an explicit commitment to cost accessibility, recognizing that the majority of vehicles on roads worldwide — particularly in developing regions where road accident mortality rates are highest — will never be equipped with sophisticated onboard computing platforms. Their embedded accident detection system uses commodity accelerometer and gyroscope components, a standard GPS module, and a GSM communication chip — components that are inexpensive, widely available, and proven in mass-market consumer electronics. When the inertial sensors detect a crash event, the GPS module provides coordinates and the GSM module dispatches an automated emergency message, all without requiring any action from the driver who may be severely injured or unconscious. The entire system was designed to operate on a simple microcontroller platform accessible to low-budget manufacturing. The fundamental design trade-off the authors navigated is between detection sensitivity and false positive rate: a threshold sensitive enough to catch every real crash will inevitably also trigger on the aggressive braking, rough terrain, and road surface variations that characterize driving conditions in many regions where this system is most needed.

Hassan, F., & Rahimi, M. built their accident detection system on a principled rejection of single-modality approaches, arguing that the real world presents too many failure scenarios for any one sensing technology to handle reliably on its own. Cameras go blind in the dark and in fog. Microphones are defeated by wind noise and traffic drone. Vibration sensors cannot tell a collision from a construction site. But a system that listens for the distinctive acoustic signature of an impact, sees the visual disruption in the traffic scene, and simultaneously feels the physical shock through vibration sensors is far harder to fool and far more robust to the failure of any individual component. Their AI fusion model was trained to combine these three streams of evidence into a unified detection decision, with the model learning to weight each modality's contribution according to its reliability in the current environmental conditions. The system was evaluated on a custom multimodal dataset and outperformed all single-modality baselines across all tested conditions. The authors are realistic about the deployment implications: instrumented monitoring nodes carrying three sensor types cost more to manufacture, install, and maintain than camera-only systems, and this cost multiplier must be justified by the reliability gains in the specific contexts where deployment is being considered.

Mathew, D., & Thomas, A. placed their research within the dispatch center, the human nerve center of emergency response, and asked how AI could make the people within this center more effective, rather than replacing them. The AI system for decision support monitors the flow of emergency messages, uses its trained model to evaluate the urgency of each message, and displays to the human dispatchers a ranked list of active emergencies, annotated so that they can focus their attention and resource allocation where the emergencies are most urgently in need of their attention. It was trained using historical dispatch records, which contain the expertise of experienced human dispatchers, and tested against human baselines for decision-making. The design philosophy is explicitly collaborative: the AI ranks and recommends, but the human operator makes the final call, preserving accountability while reducing cognitive overload. The limitation the authors acknowledge is the data dependency: the model's prioritization quality is only as good as the historical records it was trained on, and those records reflect the patterns of the specific city and time period from which they were drawn — a model trained on data from one city may perform poorly when deployed in another with different road layouts, incident types, and response conventions.

Johnson, R., & Lee, H. offered perhaps the most ambitious vision of all: an autonomous response pipeline where AI detection, GPS locationing, and cloud-based coordination work so well together that human involvement becomes not only desirable but optional from detection through response. The rationale is not only laudable but also urgent: humans introduce delays, make mistakes when tired, and aren't available at 3 AM to staff a dispatch center. A system that eliminates the need to wait for a human to notice and respond has the potential to save lives in exactly the situations where the existing system is most likely to fail. The authors evaluated their architecture through large-scale simulation and demonstrated substantial reductions in incident-to-response time compared to human-operated baselines, alongside discussion of integration pathways with existing emergency service infrastructure that would make phased real-world deployment feasible. Yet the authors were admirably clear-eyed about the profound challenges that stand between this vision and responsible production deployment — the system must remain functional when network connections fail during the very infrastructure-damaging emergencies it is designed to handle; the ethical and legal frameworks for autonomous systems making life-and-death dispatch decisions without human oversight do not yet exist in most jurisdictions; and the risk of model performance degradation over time, as road conditions, urban layouts, and vehicle types evolve in ways the training data did not anticipate, demands continuous monitoring and retraining pipelines that are themselves complex systems to maintain. These are not merely technical problems — they are questions about how much trust society is prepared to place in automated systems when the stakes are human lives.

3. CONCLUSION

This research aims to examine the process of developing an Accident and Fire Detection System that makes use of artificial intelligence, computer vision, and geolocation technology to improve emergency response mechanisms. The proposed Accident and Fire Detection System makes use of the YOLO (You Only Look Once) object detection algorithm to detect accident scenarios and fire incidents through live feed images. Current research studies and developments in deep learning and computer vision technology allow the system to analyze live feed images from surveillance cameras and detect abnormal conditions such as accidents, fire outbreaks, and hazardous conditions with high accuracy and speed. Current research studies prove that the proposed system can achieve efficient performance using real-time object detection algorithms, such as the YOLO algorithm, for detecting abnormal conditions such as accidents and fire outbreaks. However, there are challenges in handling different lighting conditions and traffic conditions while achieving low latency for real-time applications. The proposed Accident and Fire Detection System can instantly send notifications to accident rescue teams whenever there is an accident or fire detected using the YOLO algorithm. This helps reduce the response time for emergency teams and enables them to reach the location more quickly.

4. REFERENCES

[1] Zhang, L., & Chen, Y. (2022). Real-time accident detection using deep convolutional neural networks. IEEE Transactions on Intelligent Transportation Systems, 23(9), 15012–15021.

[2] Kumar, A., & Bhatia, S. (2023). Smart emergency response through IoT and machine learning integration. International Journal of Computer Applications, 182(3), 45–52.

[3] Li, D., & Zhao, P. (2021). Automatic fire detection from surveillance videos using hybrid CNN models. Computers & Electrical Engineering, 95, 107433.

[4] Prakash, R., & Thomas, N. (2023). AI-driven accident detection and alert systems using OpenCV. International Journal of Emerging Trends in Engineering Research, 11(4), 280–287.

[5] Singh, J., & Mehta, V. (2022). Edge computing approach for intelligent transportation safety. IEEE Access, 10, 78234–78245.

[6] Rahman, M., & Alam, S. (2021). A machine learning-based framework for automatic road incident detection using CCTV feeds. Sensors, 21(8), 2739.

[7] Patel, K., & Das, P. (2023). Deep learning-enabled accident severity classification for smart cities. Journal of Artificial Intelligence and Data Mining, 11(1), 31–39.

[8] Reddy, S., & Nair, G. (2020). Smart emergency notification via GPS and GSM modules in IoT-based accident systems. International Journal of Advanced Research in Computer Engineering & Technology, 9(6), 612–618.

[9] Wang, T., & Liu, C. (2024). Vision-based human fall and accident detection in outdoor environments. Pattern Recognition Letters, 181, 1–9.

[10] Anusha, P., & Raj, M. (2022). Automated rescue alert system using AI-powered surveillance cameras. International Conference on Smart Computing and Systems (ICSC), 210–216.

[11] Yadav, P., & Sharma, H. (2021). Integration of cloud and IoT for real-time emergency management. Journal of Network and Computer Applications, 191, 103176.

[12] Farooq, S., & Khan, Z. (2020). Accident detection using image processing and sensor fusion techniques. International Journal of Computer Vision and Signal Processing, 10(2), 48–55.

[13] Borse, R., & Patil, A. (2023). AI-based vehicle crash detection with GPS-enabled alert systems. International Journal of Intelligent Systems and Applications, 15(2), 102–110.

[14] Chen, J., & Luo, K. (2024). Dynamic route optimization for emergency vehicles using real-time traffic data. IEEE Intelligent Transportation Systems Magazine, 16(1), 54–62.

[15] George, A., & Kannan, T. (2021). Computer vision-based fire and smoke detection using YOLOv5 architecture. Applied Artificial Intelligence, 35(12), 1198–1209.
[16] Dinesh, R., & Shankar, P. (2023). Smart surveillance for emergency detection using AI and cloud analytics. International Journal of Smart Sensors and Ad Hoc Networks, 13(3), 89–96.

[17] Ahmed, N., & Hussain, R. (2020). Machine learning approaches to accident prediction and prevention. Transportation Research Part C: Emerging Technologies, 118, 102715.

[18] Zhang, Y., & Xu, L. (2022). Hybrid AI model for recognizing road hazards in real-time video feeds. Neural Computing and Applications, 34(22), 20147–20158.

[19] Nirmal, J., & Varma, S. (2021). IoT-based smart rescue management with automated alert prioritization. International Journal of Internet of Things, 10(4), 102–109.

[20] Bhattacharya, P., & Roy, D. (2023). Deep learning and IoT convergence for smart city emergency response systems. Future Internet, 15(7), 189.

[21] Menon, C., & Joseph, M. (2024). Design of AI-enabled accident detection and real-time notification system for highways. International Journal of Innovative Research in Computer and Communication Engineering, 12(255–63.

[22] Gupta, A., & Arora, P. (2020). Smart accident detection and ambulance dispatch using GSM and GPS modules. Procedia Computer Science, 172, 619–626.

[23] Hassan, F., & Rahimi, M. (2023). Multimodal sensing for accident detection using AI fusion models. IEEE Sensors Journal, 23(12), 15345–15352.

[24] Mathew, D., & Thomas, A. (2021). AI-based prioritization for emergency response optimization. International Conference on Data Science and Communication (ICONDSC), 312–319.

[25] Johnson, R., & Lee, H. (2024). Towards autonomous emergency response: integration of AI, GPS, and cloud systems. Journal of Intelligent Systems Research, 18(5), 210–219.