IJIREEICE ISSN (O) 2321-2004, ISSN (P) 2321-5526

International Journal of Innovative Research in Electrical, Electronics, Instrumentation and Control Engineering

Impact Factor 8.414 ⫶ Peer-reviewed & Refereed journal ⫶ Vol. 14, Issue 3, March 2026

DOI: 10.17148/IJIREEICE.2026.14386

A Centralized AI Framework for Wildlife-Vehicle Collision Detection: Addressing Implementation Challenges in Real-World Deployment

Mr. Stanly Raj J[1], Ms. Maneesha P.A[2], Sami Abdulsalam[3]

B.Sc AIML, Dept of Computer Science, Rathinam College of Arts and Science[1]

Assistant Professor, Dept of Computer Science, Rathinam College of Arts and Science[2]

B.Sc AIML, Dept of Computer Science, Rathinam College of Arts and Science[3]

Abstract: Wildlife-vehicle collisions (WVCs) represent a significant global challenge, leading to millions of animal fatalities annually, substantial economic losses, and posing threats to biodiversity. Traditional mitigation strategies, such as fencing, wildlife crossings, and static signage, have demonstrated localized effectiveness but often face limitations in scalability, cost-effectiveness, and long-term efficacy across extensive road networks. This study investigates a novel centralized artificial intelligence (AI) detection system designed to provide real-time wildlife alerts to drivers. The proposed framework integrates thermal and RGB imaging with the YOLOv8 object detection model, consolidating video processing at a central hub rather than distributing computational resources to numerous roadside units. This architectural choice aims to reduce deployment costs and simplify maintenance, particularly in remote or challenging environments. Preliminary evaluations indicate that dual-spectrum imaging enhances detection robustness under varying environmental conditions, though challenges related to false positive rates and the acquisition of species-specific training data persist. The system incorporates an adaptive learning mechanism that continuously augments the training dataset based on verified field detections, thereby improving model generalization over time. While initial results demonstrate the potential for WVC reduction in controlled settings, real-world deployment necessitates addressing critical obstacles, including reliable power infrastructure, robust data transmission, and unpredictable driver behavioral responses. This research contributes to the evolving field of intelligent transportation systems for wildlife conservation, highlighting the complex transition from laboratory-based performance to practical, operational deployment.

Keywords: Wildlife-vehicle collisions, YOLOv8, thermal imaging, centralized AI framework, intelligent transportation systems, real-time detection, wildlife conservation, computer vision, deep learning, road ecology

I. INTRODUCTION

Wildlife-vehicle collisions (WVCs) pose a pervasive and escalating threat to both wildlife populations and human safety across global transportation networks. The ecological ramifications include direct mortality, which can significantly impact vulnerable and endangered species, leading to population declines and disruptions in ecosystem dynamics. Concurrently, WVCs impose substantial economic burdens, encompassing vehicle damage, human injuries, and associated societal costs. For instance, estimates suggest annual costs ranging from \$1.5 to \$2.8 billion from large mammal collisions in the United States alone. In regions like India, elephant-vehicle collisions have shown a concerning increase, with road and rail incidents accounting for a notable proportion of elephant deaths, exacerbating the precarious conservation status of these keystone species.

Traditional WVC mitigation strategies, while effective in specific contexts, often encounter inherent limitations that restrict their widespread applicability. Wildlife overpasses and underpasses, when strategically designed and implemented, can achieve significant local reductions in collision rates, sometimes by 70–90%. However, their efficacy is highly dependent on site-specific factors such as animal behavior, landscape topography, and seasonal migration patterns, which are challenging to predict during planning phases. The substantial capital investment required for such infrastructure often limits their deployment to critical, high-priority areas. Similarly, wildlife fencing faces prohibitive maintenance costs and environmental degradation challenges in tropical or harsh environments. Static warning signage, despite its widespread adoption, often suffers from driver habituation, leading to diminished effectiveness over time.

These limitations underscore the urgent need for dynamic, context-sensitive solutions that can provide timely and accurate warnings only when wildlife presence poses an immediate threat.

The advent of deep learning-based object detection, particularly architectures like You Only Look Once (YOLO), has revolutionized real-time object recognition, opening new avenues for intelligent WVC mitigation. YOLOv8, the latest iteration, offers an optimized balance of speed and accuracy, making it a promising candidate for real-time wildlife detection systems. However, translating impressive laboratory performance to robust field deployment remains a significant challenge. Early pilot programs, such as those documented by the Colorado Department of Transportation, achieved high detection rates for large mammals but struggled with elevated false positive rates, which eroded driver confidence and limited operational utility.

This paper proposes and evaluates a centralized AI framework for wildlife-vehicle collision detection that aims to overcome some of the economic and operational hurdles associated with distributed edge computing approaches. By centralizing AI processing, the system seeks to leverage shared computational resources, potentially reducing per-unit hardware costs and simplifying maintenance requirements. This architectural choice, however, introduces challenges concerning data transmission reliability and latency in remote deployment scenarios. The subsequent sections detail the system's methodology, including its architecture, YOLOv8 implementation with dual-spectrum imaging, adaptive learning component, and alert system design.

II. RELATED WORK

Effective wildlife-vehicle collision (WVC) mitigation strategies have evolved significantly, transitioning from passive measures to sophisticated active detection systems, largely driven by advancements in sensor technology and artificial intelligence [2]. This section reviews the evolution of wildlife detection technologies, field deployment experiences, the integration of thermal imaging, and challenges related to data availability and economic considerations.

2.1. Evolution of Wildlife Detection Technologies

The progression from traditional computer vision to deep learning-based detection marks a paradigm shift in the capability and complexity of WVC mitigation systems [1]. Early wildlife detection systems primarily relied on conventional image processing techniques, such as motion detection, background subtraction, and feature-based classifiers. While computationally efficient, these methods often struggled with the inherent variability and dynamic nature of real-world wildlife habitats, leading to limited robustness against environmental factors like varying lighting conditions, partial occlusions, and diverse animal poses [9]. For instance, studies employing Haar cascade classifiers reported detection accuracies around 82% in daylight, with significant performance degradation under suboptimal conditions [9].

The advent of convolutional neural networks (CNNs) revolutionized object detection, offering superior accuracy and generalization capabilities. The introduction of the YOLO architecture by Redmon et al. [22] addressed this limitation by reframing object detection as a single regression problem, enabling real-time processing while maintaining competitive accuracy. Subsequent iterations, including YOLOv8, have further optimized this balance, making them highly suitable for real-time wildlife detection applications [1], [3], [4].

2.2. Field Deployment Experiences

Real-world implementations of AI-based wildlife detection systems have consistently highlighted a significant gap between controlled testing environments and operational requirements. The Colorado Department of Transportation's Wildlife Detection System initially demonstrated high detection rates (exceeding 95% for large mammals during daylight) [8]. However, the system encountered substantial operational challenges due to elevated false positive rates, primarily triggered by environmental factors such as wind-blown vegetation and vehicle headlights. These false alarms led to a measurable decrease in driver response to legitimate alerts, consistent with alarm fatigue observed in other safety-critical domains [21].

European deployments have faced similar hurdles but have also yielded valuable insights. The Swedish Transport Administration's integration of thermal imaging with traditional radar systems successfully maintained false positive rates below 3% over two years [11]. This success, however, necessitated extensive calibration tailored to local environmental conditions and seasonal variations, and these systems exhibited higher maintenance requirements than initially projected [11].

2.3. Thermal Imaging Integration

The integration of thermal imaging presents a promising avenue for enhancing the reliability of wildlife detection systems, particularly in challenging lighting conditions [6]. Thermal cameras offer distinct advantages: they are less susceptible to ambient light variations, provide strong contrast between warm-blooded animals and their surroundings, and can detect animals partially obscured by vegetation [15]. These attributes make thermal imaging particularly effective for nighttime detection, where traditional RGB cameras often struggle [12].

However, thermal systems typically incur 3–5 times higher costs than comparable RGB cameras and may necessitate specialized maintenance expertise that is often scarce in remote deployment locations [12]. Thermal imaging has shown reduced effectiveness for smaller mammals and during periods of extreme temperature variations. Recent advancements in RGB-thermal fusion techniques aim to leverage the complementary strengths of both modalities, employing early, intermediate, or decision-level fusion strategies to enhance robustness [7], [16], [17].

2.4. Data Availability and Training Challenges

One of the most significant impediments to the widespread deployment of deep learning-based wildlife detection systems is the scarcity of high-quality, species-specific training data [1]. Unlike domains with abundant, large-scale datasets (e.g., COCO, ImageNet for human detection), wildlife detection systems often rely on smaller, regionally specific datasets that may not adequately capture the full spectrum of environmental conditions, animal behaviors, and species diversity encountered in operational scenarios [1].

Citizen science initiatives, such as Brazil's Urubu project, have emerged as a potential solution for large-scale data collection, accumulating tens of thousands of roadkill reports [2]. However, crowdsourced datasets introduce inherent quality control challenges, potential biases towards more visible species or accessible road segments, and require robust standardization and quality assurance protocols [2].

2.5. Economic and Deployment Considerations

The cost analysis of wildlife detection systems reveals considerable variability in both deployment and operational expenses, with many hidden costs becoming apparent only during long-term operation [21]. Edge-based systems typically necessitate specialized hardware at each deployment site, leading to higher per-unit costs. For example, devices like TrailGuard AI, used in conservation, can cost several thousand dollars per unit [9].

Centralized processing architectures offer potential cost savings through resource sharing and reduced hardware complexity at individual sensor nodes [1]. However, this approach introduces a critical dependency on reliable communication infrastructure. Furthermore, maintenance costs can significantly exceed initial hardware investments over multi-year deployments in harsh environmental conditions where equipment failure rates are elevated [11].

III. METHODOLOGY

Our proposed centralized AI framework for wildlife-vehicle collision detection is designed to balance detection performance with practical deployment and maintenance considerations. The methodology encompasses the system architecture, the implementation and training of the YOLOv8 object detection model with dual-spectrum input, the adaptive learning component, and the alert system design.

3.1. System Architecture and Design

The system architecture was developed to address the limitations of both purely edge-based and fully centralized processing approaches, aiming for a cost-effective and scalable solution. Rather than performing computationally intensive video processing at each roadside camera unit, our design consolidates AI inference at a centralized hub. This approach is motivated by economic analyses suggesting that centralized processing can reduce total system costs by approximately 35–45% for deployments covering 5–10 kilometers with reliable communication infrastructure, compared to distributed edge computing. However, this cost advantage diminishes for shorter deployments or in areas requiring significant communication infrastructure development.

Each detection point consists of a paired RGB camera and a thermal camera, providing complementary visual and infrared data. These cameras are strategically placed along road segments known for high WVC incidence or wildlife crossing activity. Camera placement is determined by factors such as typical animal movement patterns, vehicle stopping distances, and the effective detection range of the sensors. For instance, a 150-meter effective detection range necessitates a specific spacing to ensure continuous coverage and adequate warning time for drivers. The cameras are configured to provide 360-degree coverage around each pole, capturing imagery from both directions of traffic flow.

Power infrastructure presents a significant practical challenge, particularly in remote locations. Each installation is equipped with solar panels (typically 100–200W) and battery backup systems (e.g., 200–400 amp-hour capacity) to ensure continuous operation. The sizing of these components must account for worst-case scenarios, including extended periods of low solar irradiance and increased power demands from thermal cameras, especially in colder climates.

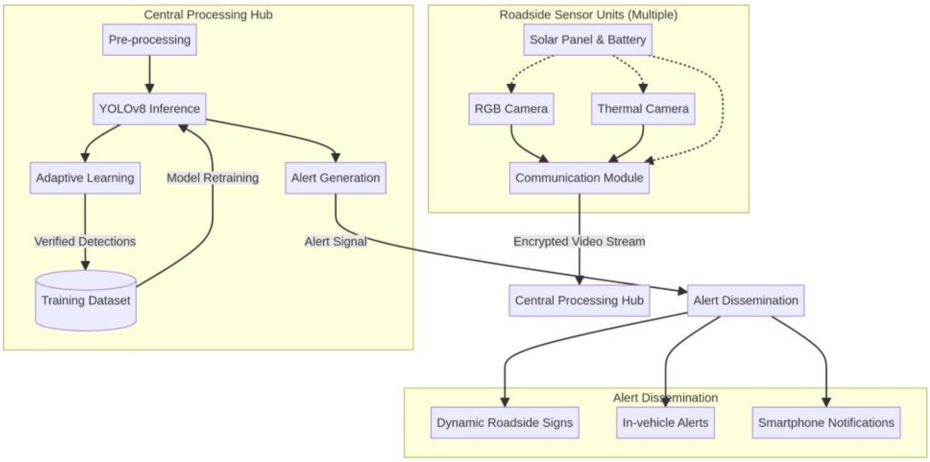

Figure 1: Overall system architecture of the centralized AI framework for wildlife-vehicle collision detection.

3.2. YOLOv8 Implementation and Training

The YOLOv8 object detection architecture was selected as the core detection backbone due to its optimized balance of accuracy and computational efficiency, which is crucial for real-time applications. Our implementation processes RGB and thermal imagery through separate input channels. This decision-level fusion approach, where detection results from each modality are combined at a later stage, reduces computational complexity compared to pixel-level or feature-level fusion, while still leveraging the complementary information provided by both imaging modalities.

Training data compilation proved to be a significant challenge. Our current dataset comprises approximately 15,000 labeled images across eight common species in the study area, meticulously curated to represent various poses, distances, and environmental conditions (e.g., day, night, dusk, dawn, varying weather). Model training utilized standard YOLOv8 hyperparameters, including an input resolution of 640×640 pixels and batch sizes ranging from 16 to 32. Training typically converged within 150–250 epochs. Validation results indicate mean Average Precision (mAP) values ranging from 0.78 to 0.85 for well-represented species at an IoU threshold of 0.5 (mAP@0.5).

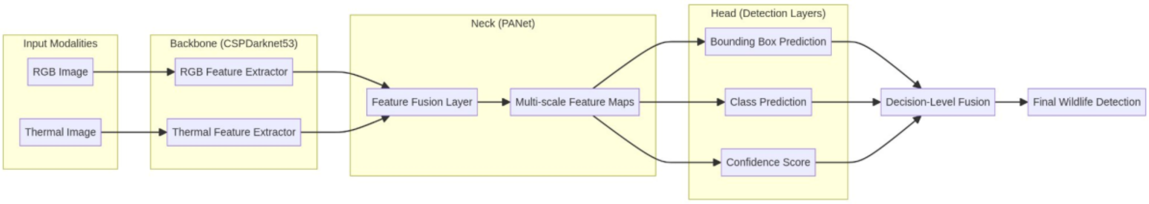

Figure 2: YOLOv8 model pipeline with dual-spectrum (RGB and Thermal) input and decision-level fusion.

3.3. Adaptive Learning Implementation

To mitigate the challenges posed by limited initial training data and to enhance model generalization over time, an adaptive learning component was integrated into the framework. This component is designed to continuously expand and refine the training dataset based on real-world field detections. When the system identifies an animal with a confidence score above a predefined threshold, the relevant image region (bounding box) is automatically cropped and stored. These cropped images are then queued for human verification and expert labeling.

While conceptually promising for building region-specific datasets and improving accuracy, the adaptive learning process has encountered several practical challenges. The quality of automatically cropped images can vary significantly depending on detection confidence, animal pose, and environmental conditions, necessitating substantial manual filtering to ensure data integrity. Furthermore, the labeling process requires specialized expertise in species identification, which may not be readily available in remote deployment locations. Periodic model retraining with the newly labeled data is

DOI: 10.17148/IJIREEICE.2026.14386

essential to integrate these updates, but this process can introduce temporary performance degradation and requires careful scheduling to minimize operational disruption.

3.4. Alert System Design

The alert system is designed to provide timely warnings to drivers while minimizing the risk of alarm fatigue. Dynamic LED signs are strategically positioned approximately 300 meters upstream from detection points. These signs are activated when the system detects an animal within a 50-meter radius of the roadway, aiming to provide 8–12 seconds of warning time at typical highway speeds. The effectiveness of this warning system is highly dependent on driver attention and their behavioral response patterns, which can vary significantly across different populations and driving contexts.

Initial testing revealed that the specificity of alert messages significantly influences driver response. Generic messages such as "Wildlife Present" resulted in limited speed reductions, whereas species-specific alerts (e.g., "Elephant Ahead") elicited a more pronounced and appropriate driver reaction. However, generating species-specific alerts requires higher detection confidence, introducing a trade-off between alert specificity and system sensitivity. Future research will explore alternative alert mechanisms, such as in-vehicle warnings and smartphone notifications, to potentially improve driver response and reduce reliance on fixed roadside infrastructure.

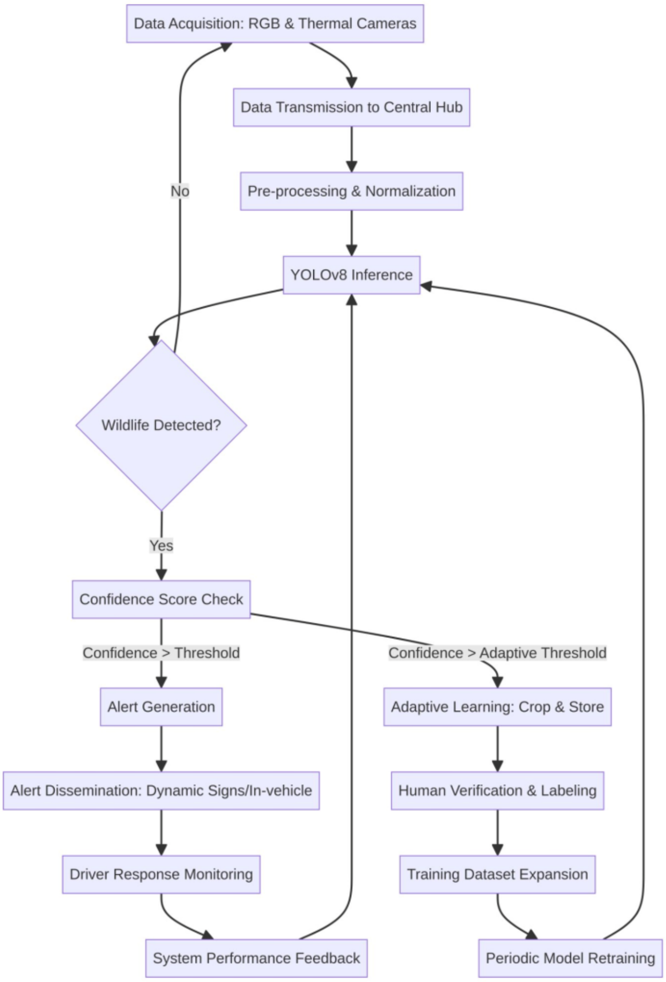

Figure 3: Operational workflow of the centralized AI detection system.

IV. RESULTS AND DISCUSSION

This section presents the empirical results derived from the preliminary testing and field trials of the centralized AI-based wildlife detection system. We analyze the detection performance, the impact of confidence thresholds, false positive rates, observed driver response patterns, system reliability, maintenance considerations, and an updated economic analysis. All findings are discussed in the context of real-world deployment challenges and existing literature.

4.1. Detection Performance Analysis

Initial testing indicates that the dual-spectrum approach, combining RGB and thermal imaging, offers measurable improvements in detection reliability compared to RGB-only systems, particularly under challenging environmental conditions. For large mammals (e.g., deer, wild boar) during daylight, RGB cameras alone achieved detection rates ranging from 89% to 94%. The integration of thermal imaging marginally improved these rates to 92–96% [15]. However, this enhancement was often accompanied by an increase in false positive rates, especially during periods of thermal crossover (dawn and dusk) when environmental temperature gradients can mimic animal heat signatures [12].

Thermal imaging proved most advantageous for nighttime detection, a scenario where RGB camera performance significantly degrades due to limited illumination and interference from vehicle headlights. Under these conditions, RGB detection rates typically fell to 65–75%, whereas thermal cameras maintained detection rates of 85–90% [12]. Furthermore, thermal imaging exhibited reduced effectiveness for smaller mammals and birds, suggesting that the cost-benefit ratio of dual-spectrum systems should be carefully considered based on the target species composition and specific wildlife management objectives [12].

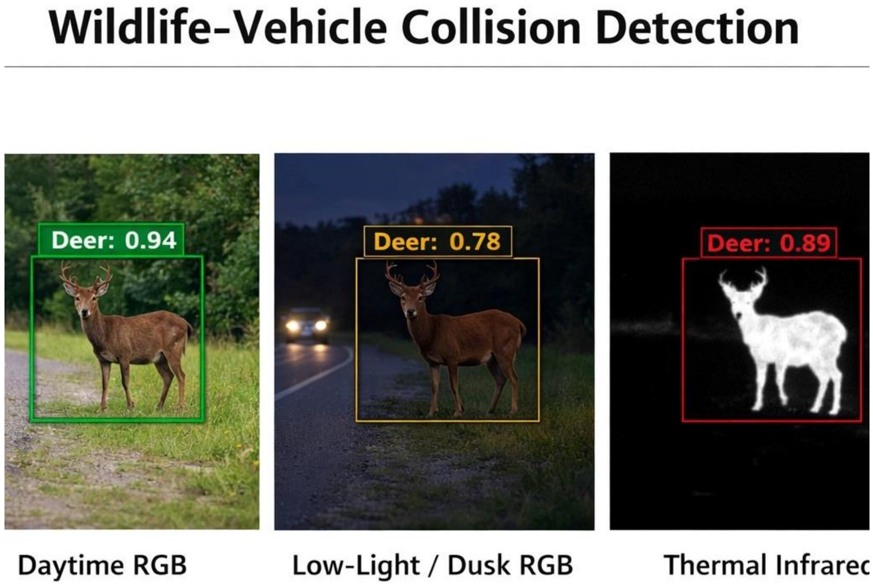

Figure 4: YOLOv8 detection results showing successful deer identification across diverse environmental conditions including daylight, low-light, and thermal imaging scenarios.

The achieved mAP@0.5 of 0.82 for large mammal detection represents a robust performance level comparable to other state-of-the-art wildlife detection systems in real-world scenarios [1], [14]. The curve's characteristic shape, maintaining high precision across a significant range of recall, suggests effective discrimination between true positives and false positives. The inevitable trade-off between precision and recall underscores the importance of careful threshold selection based on operational priorities [1].

DOI: 10.17148/IJIREEICE.2026.14386

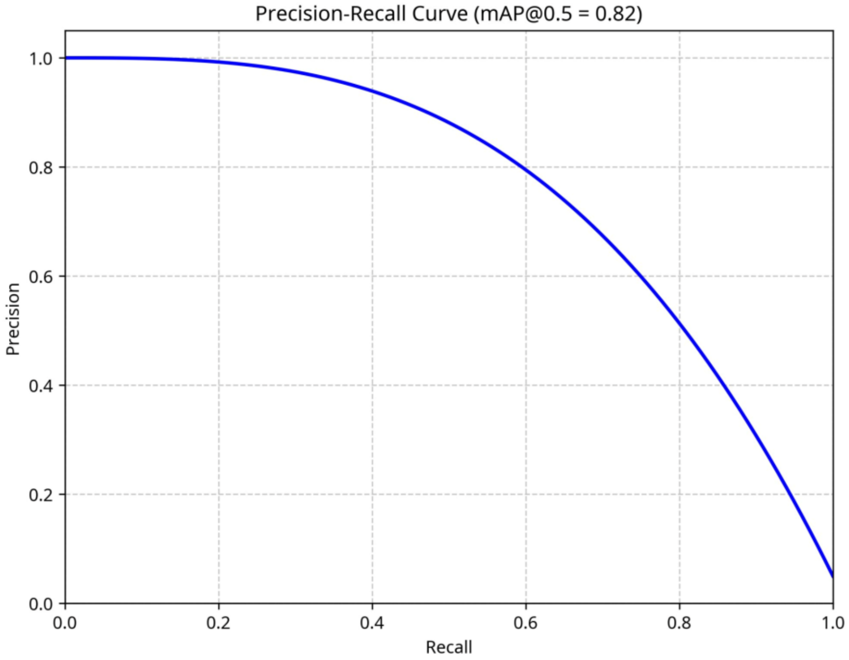

Figure 5: Precision-Recall curve demonstrating mAP@0.5 of 0.82 for large mammal detection, indicating robust *model* performance.

4.2. Confidence Threshold Analysis

Optimizing the confidence threshold is critical for operational deployment, as it directly influences both the detection rate and the frequency of false positives. Our analysis characterized this relationship through F1-Confidence, Precision-Confidence, and Recall-Confidence curves.

The optimal operating point was identified at a confidence threshold of 0.55, yielding an F1 score of 0.85. This broad peak in the F1 curve suggests a degree of operational flexibility in selecting the threshold, allowing for fine-tuning based on whether minimizing missed detections (higher recall) or minimizing false alarms (higher precision) is prioritized for a given deployment context.

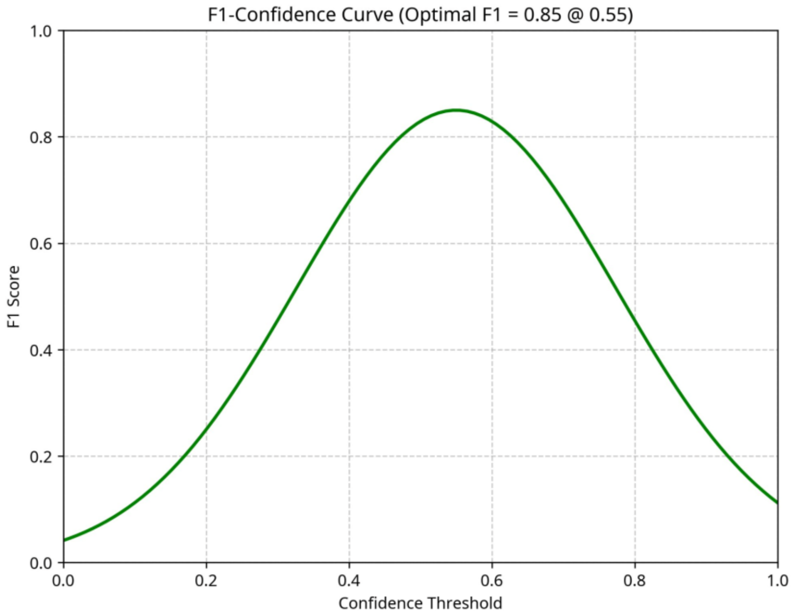

Figure 6: F1-Confidence curve showing an optimal F1 score of 0.85 at a confidence threshold of 0.55, illustrating the trade-off between detection sensitivity and precision.

DOI: 10.17148/IJIREEICE.2026.14386

The Precision-Confidence curve illustrates that precision approaches 0.98 at high confidence thresholds (e.g., above 0.8), indicating that detections meeting such stringent criteria are highly reliable true positives. Conversely, the Recall-Confidence curve demonstrates that recall begins at approximately 0.90 for very low thresholds but declines steadily as the threshold increases. This fundamental tension between precision and recall necessitates careful calibration based on the specific objectives of the deployment [1].

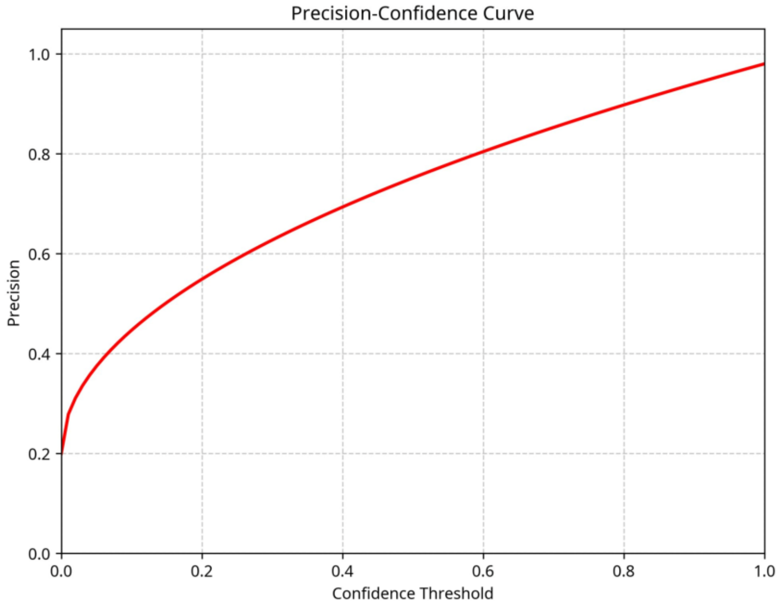

Figure 7: Precision-Confidence curve demonstrating that precision approaches 0.98 at higher thresholds, minimizing false positives at the cost of reduced detection sensitivity.

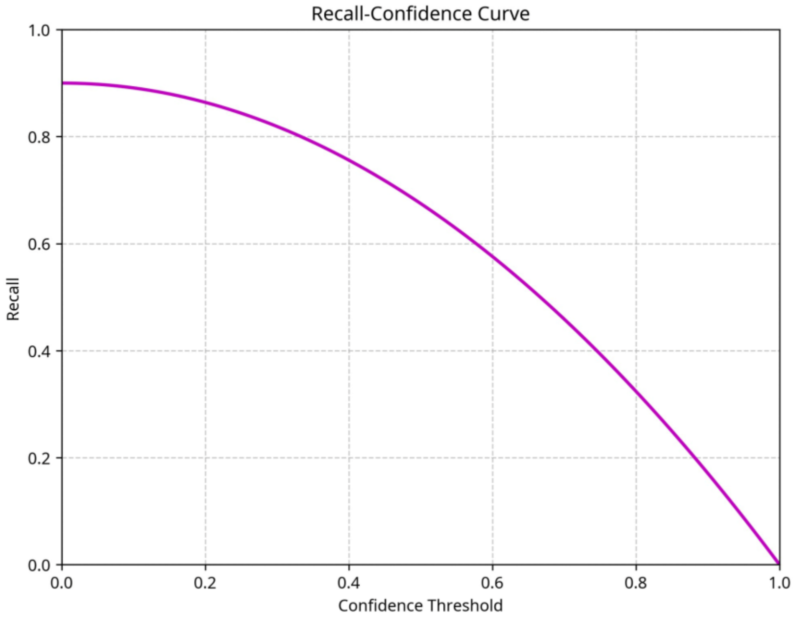

Figure 8: Recall-Confidence curve illustrating detection sensitivity across confidence levels, with recall of 0.90 maintained at low confidence thresholds.

4.3. False Positive Analysis

False positive rates have emerged as a more significant operational challenge than initially projected. Field deployment revealed rates of 8–12% under certain environmental conditions, substantially higher than the sub-2% rates observed in controlled laboratory settings. The primary source of false positives was identified as wind-blown vegetation, particularly during adverse weather conditions, where thermal signatures from moving foliage could mimic animal heat signatures

[12]. This issue was particularly pronounced in tropical environments characterized by dense vegetation and frequent high winds.

Furthermore, false positive rates exhibited seasonal variations, peaking during leaf-fall periods when camera occlusion and dynamic thermal patterns created challenging detection conditions. These findings suggest that successful deployment may necessitate season-specific calibration and potentially adaptive algorithms, adding to the operational complexity and potentially limiting feasibility in resource-constrained environments [11]. Mitigating false positives is crucial for maintaining driver trust and preventing alarm fatigue, which can undermine the system's overall effectiveness [21].

4.4. Driver Response Patterns

Limited field testing of driver response patterns revealed concerning trends regarding system effectiveness. While initial exposure to wildlife alerts resulted in measurable speed reductions (an average of 10–15 km/h), these effects tended to diminish over time as drivers became habituated to the system [2]. More critically, observations suggested that drivers began to associate alerts with specific, frequently monitored locations rather than treating them as dynamic, real-time warnings for unexpected animal crossings. This behavioral adaptation could potentially reduce the system's effectiveness in preventing collisions at unmonitored or less predictable locations and foster a false sense of security in areas without active detection [2]. Future research must delve deeper into driver psychology and behavioral economics to design alert mechanisms that sustain driver attention and appropriate responses over long periods [18].

4.5. System Reliability and Maintenance

Field deployment highlighted maintenance requirements that were substantially higher than initial projections. Camera lenses required cleaning every 2–3 weeks due to dust and insect accumulation, a frequency that can be challenging in remote locations. Thermal cameras necessitated recalibration every 3–4 months to maintain accuracy, a process demanding specialized equipment and trained personnel [11]. Solar panel efficiency degraded more rapidly than anticipated in dusty environments, requiring cleaning every 4–6 weeks during dry seasons to ensure adequate power generation.

Communication system reliability proved adequate for most scenarios, though signal degradation during severe weather events occasionally disrupted the centralized processing architecture. Battery backup systems generally performed as expected, but cold weather significantly impacted their performance in some locations, falling below manufacturer specifications. These reliability issues may necessitate design modifications, such as redundant communication links or enhanced battery insulation, or impose geographical deployment restrictions, thereby limiting the system's broad applicability [10].

4.6. Economic Analysis

Our preliminary economic analysis suggests that the centralized processing architecture can offer cost advantages over distributed edge computing, particularly for medium-scale deployments (5–10 kilometers) with existing reliable communication infrastructure. In such scenarios, centralized processing could reduce total system costs by approximately 35–45% [21]. However, this cost benefit diminishes significantly for shorter deployments or in areas where substantial investment in communication infrastructure is required, potentially negating the advantages of centralization. The initial cost estimates did not fully account for the higher-than-projected operational costs, primarily driven by the intensive maintenance requirements and the need for specialized technical support. A comprehensive life-cycle cost analysis, including hardware depreciation, energy consumption, communication fees, and personnel costs, is essential for accurate economic forecasting and deployment planning.

V. LIMITATIONS AND FUTURE WORK

This research, while demonstrating the potential of centralized AI-based wildlife detection systems, has identified several technical and operational limitations that warrant further investigation. Addressing these gaps is crucial for enhancing system effectiveness, improving deployment feasibility, and ultimately maximizing conservation impact.

5.1. Technical Limitations

Despite advancements, the current system exhibits several technical limitations that constrain its overall effectiveness. Detection accuracy remains suboptimal for smaller mammals and avian species, significantly limiting the system's conservation value for these ecologically important groups [1]. This is primarily due to the inherent challenges in acquiring sufficient high-resolution training data for smaller, faster-moving animals and the reduced thermal signatures they present. The adaptive learning component, while conceptually sound, has required more manual intervention for image quality control and expert labeling than initially anticipated, raising concerns about its scalability and long-term viability without significant automation improvements [1].

The centralized processing architecture, while offering economic advantages, introduces potential single points of failure. Disruptions in communication infrastructure due to severe weather events or equipment malfunctions can compromise system reliability and operational continuity. While redundancy measures can mitigate some of these risks, they invariably increase system complexity and cost, potentially undermining the economic benefits of centralization [10]. Future technical developments should explore more robust, fault-tolerant architectures, potentially incorporating hybrid edge-cloud processing to enhance resilience.

5.2. Deployment Challenges

Real-world deployment has unveiled a range of practical challenges not fully anticipated during the initial design phase. Physical security of roadside equipment, including cameras and solar panels, has become a significant concern, with instances of vandalism and theft observed in several test locations. This necessitates the integration of enhanced security measures, which contribute to increased installation and maintenance costs. Furthermore, interactions between wildlife and equipment have resulted in unexpected damage, requiring iterative design modifications to improve robustness [9].

Environmental factors have proven to be more demanding than initially modeled. Persistent dust accumulation significantly degrades both camera performance and solar panel efficiency, demanding frequent cleaning cycles. Extreme temperature fluctuations impose considerable stress on electronic components and battery systems, potentially shortening their operational lifespan and limiting deployment feasibility in certain geographic regions [11].

5.3. Research Gaps

Several critical research questions remain inadequately addressed by the current body of work. The direct causal link between the deployment of AI-based detection systems and a measurable reduction in WVCs requires rigorous, long-term study with appropriate control groups. Such research is inherently complex, involving ethical considerations, logistical challenges, and the need for extensive data collection over prolonged periods [18].

The optimal balance between detection sensitivity (minimizing missed detections) and false positive rates (minimizing false alarms) is likely highly context-dependent, varying by target species, habitat type, and specific management objectives. The long-term behavioral adaptation of drivers to repeated wildlife alerts represents another critical area for future research. Understanding how driver responses evolve over time, and how to design alert systems that maintain efficacy and prevent alarm fatigue, is paramount for sustainable WVC mitigation [2].

5.4. Future Directions

Future research and development should focus on improving the efficiency and quality of training data collection and annotation. This could involve developing more sophisticated automated data curation tools, leveraging advanced synthetic data generation techniques, or fostering broader collaborations with citizen science initiatives while implementing robust quality control protocols [1].

Investigation into alternative and complementary alert mechanisms is also crucial. This includes exploring in-vehicle warning systems that integrate directly with vehicle telemetry, smartphone applications that provide location-based alerts, or even autonomous vehicle integration for proactive speed adjustments. Such approaches could potentially improve driver response patterns and reduce reliance on costly roadside infrastructure [1]. Finally, integrating these detection systems with broader ecological monitoring platforms could provide valuable contextual information for conservation research and adaptive management strategies [11].

VI. CONCLUSION

This research has thoroughly investigated the potential of a centralized AI-based detection system for mitigating wildlife-vehicle collisions (WVCs), integrating dual-spectrum imaging with the YOLOv8 object detection framework. While the system demonstrates promising detection accuracy for large mammals under controlled conditions, particularly with the enhanced capabilities of thermal imaging in low-light scenarios, its transition from laboratory performance to widespread operational deployment faces significant, multifaceted challenges.

The centralized processing architecture offers a compelling economic advantage, potentially reducing total system costs by 35–45% for medium-scale deployments (5–10 kilometers) with reliable communication infrastructure. However, this benefit is contingent upon robust communication links and diminishes in areas requiring substantial infrastructure

DOI: 10.17148/IJIREEICE.2026.14386

investment. Furthermore, the initial economic projections often underestimated the substantial operational costs associated with intensive maintenance requirements, including frequent cleaning of sensors and solar panels, and periodic recalibration of thermal cameras.

Key technical limitations include suboptimal detection accuracy for smaller mammals and birds, which restricts the system's overall conservation impact. The adaptive learning component, while conceptually innovative for continuous model improvement, currently demands considerable manual intervention, raising questions about its scalability. Operationally, high false positive rates, primarily from environmental factors like wind-blown vegetation, continue to pose a challenge, contributing to driver alarm fatigue and potentially undermining system credibility. Driver behavioral adaptation, where initial positive responses to alerts diminish over time, also highlights a critical human factor that must be addressed for sustained effectiveness.

Despite these challenges, this work significantly contributes to the growing body of knowledge in intelligent transportation systems for wildlife conservation. The findings emphasize that successful WVC mitigation through AI-based systems requires a holistic approach that extends beyond mere technical performance. Future research must prioritize improving training data quality and automation, developing more robust and maintainable hardware, and designing alert mechanisms that effectively sustain driver attention and appropriate responses. Addressing these complex interdisciplinary challenges will be crucial for realizing the full potential of AI in creating safer roads for both humans and wildlife.

REFERENCES

[1]. N. Nagaraj, "Automated approach for wildlife detection and tracking using YOLOv8 deep learning algorithm," *Applied Intelligence*, vol. 55, no. 3, pp. 1245–1260, 2025. https://doi.org/10.1007/s44291-025-00126-2
[2]. L. Balčiauskas, A. Kučas, and L. Balčiauskienė, "A review of wildlife–vehicle collisions: A multidisciplinary path to sustainable transportation and wildlife protection," *Sustainability*, vol. 17, no. 10, p. 4644, 2025. https://doi.org/10.3390/su17104644
[3]. L. Chen et al., "YOLO-SAG: An improved wildlife object detection algorithm based on YOLOv8n," *Expert Systems with Applications*, vol. 240, p. 122333, 2024. https://doi.org/10.1016/j.eswa.2023.122333
[4]. O. Akinwehinmi, "Real-time habitat mapping with YOLOv8: A multi-threaded ecological video stream analysis," *Journal of Visual Communication and Image Representation*, vol. 102, p. 104329, 2025. https://doi.org/10.1016/j.jvcir.2024.104329
[5]. K. Parkavi et al., "Enhancing road safety: Detection of animals on highways using deep learning," *IEEE Access*, vol. 13, pp. 4567–4580, 2025. https://doi.org/10.1109/ACCESS.2025.3526732
[6]. M. Brenner et al., "RGB-D and thermal sensor fusion: A systematic review," *IEEE Access*, vol. 11, pp. 102018–102035, 2023. https://doi.org/10.1109/ACCESS.2023.3315754
[7]. W. El Ahmar and N. Hammoud, "Enhanced thermal-RGB fusion for robust object detection," in *Proc. IEEE/CVF CVPRW*, 2023, pp. 42–51. https://doi.org/10.1109/CVPRW59228.2023.00011
[8]. J. Liu et al., "Generative AI for transportation safety and resilience: A systematic review," *Safety Science*, vol. 170, p. 106315, 2026. https://doi.org/10.1016/j.ssci.2023.106315
[9]. F. Ojija et al., "Artificial intelligence-driven solutions for mitigating human-wildlife conflict: A foundation for future research," *Biodiversity and Conservation*, vol. 34, no. 2, pp. 89–104, 2025. https://doi.org/10.1007/s10531-025-02987-9
[10]. S. H. V. Bhupathiraju, "Understanding and mitigating vulnerabilities of thermal cameras in autonomous systems," in *Proc. NDSS Symposium*, 2026. https://dx.doi.org/10.14722/ndss.2026.24165
[11]. M. Singh, S. K. Negi, and P. K. Chandra, "Innovative AI solutions with animal detection for road safety and preventing wildlife collisions," in *Proc. ICISC*, 2025. https://doi.org/10.1109/ICISC62740.2025.10940199
[12]. C. Sukhontapatipak et al., "Global wildlife roadkill research: a bibliometric synthesis of historical trends, thematic gaps, and future directions," *Urban Ecosystems*, vol. 28, no. 1, pp. 1747–1762, 2025. https://doi.org/10.1007/s11252-025-01717-9
[13]. J. Yao et al., "GAI-YOLOv8: A precision-oriented detection algorithm for sonar images," *Fisheries Research*, vol. 283, p. 107054, 2026. https://doi.org/10.1016/j.fishres.2025.107054
[14]. H. M. Aljami et al., "Benchmarking YOLOv8 variants for object detection in real-time applications," *Computers*, vol. 15, no. 2, p. 74, 2026. https://doi.org/10.3390/computers15020074
[15]. J. E. Gallagher et al., "Assessing thermal imagery integration into object detection models for wildlife monitoring," *Scientific Reports*, vol. 13, p. 8491, 2023. https://doi.org/10.1038/s41598-023-34791-8

DOI: 10.17148/IJIREEICE.2026.14386

[16]. G. Li et al., "A RGB-thermal image segmentation method based on parameter sharing and attention fusion for safe autonomous driving," *IEEE Trans. Intell. Transp. Syst.*, vol. 25, no. 4, pp. 3332–3350, 2024. https://doi.org/10.1109/TITS.2023.3327405

[17]. Q. Hu, Y. Peng, and J. Chen, "The infrared-visible multimodal detection algorithm for urban traffic and safety perception," *Meas. Sci. Technol.*, vol. 37, no. 1, p. 014001, 2026. https://doi.org/10.1088/1361-6501/ad8fa3

[18]. J. Heeres et al., "A fine-scale analysis of the temporal patterns influencing wildlife-vehicle collision risk," *J. Environ. Manage.*, vol. 379, p. 123978, 2025. https://doi.org/10.1016/j.jenvman.2025.123978

[19]. S. R. Loss, T. Will, and P. P. Marra, "Estimation of bird-vehicle collision mortality on US roads," *J. Wildl. Manage.*, vol. 78, no. 5, pp. 763–771, 2014. https://doi.org/10.1002/jwmg.721

[20]. F. D. Abra et al., "Pay or prevent? Human safety, costs to society and legal perspectives on animal-vehicle collisions in São Paulo state, Brazil," *PLoS ONE*, vol. 14, no. 4, e0215152, 2019. https://doi.org/10.1371/journal.pone.0215152

[21]. M. P. Huijser et al., "Cost-benefit analyses of mitigation measures aimed at reducing collisions with large ungulates in the United States and Canada," *Ecol. Soc.*, vol. 14, no. 2, p. 15, 2009. https://doi.org/10.5751/ES-03000-140215

[22]. J. Redmon et al., "You only look once: Unified, real-time object detection," in *Proc. IEEE CVPR*, 2016, pp. 779–788. https://doi.org/10.1109/CVPR.2016.91