A.U.R.A. (Autonomous Undercar Risk Avoidance): Underbody animal detection and deterrence for autonomous vehicles

Theresa Dinh, Oussama Benamar Hlimi, Christopher Nunez-Velasquez, Kimberly Lugo.
Dept. of Electrical and Computer Engineering and College of Optics and Photonics, University of Central Florida, Orlando, FL, 32816, USA

ABSTRACT

Autonomous vehicles cannot observe the region directly beneath the chassis, and documented incidents show they will resume motion with animals sheltering underneath. This paper presents A.U.R.A., a vehicle-mounted system that detects animals under a stationary vehicle and deters them before movement is permitted. Retractable camera modules with 850 nm near-infrared illumination and anamorphic optics feed a YOLOv11 model that achieved 91.7% mAP@50 and 92.1% recall. Upon detection, ultrasonic and strobe deterrents activate while a relay interlock holds the drive motors off until the underbody is verified clear. Hardware validation confirmed the interlock, strobe, actuation, and alarm subsystems.

Keywords: autonomous vehicles, embedded systems, near-infrared imaging, object detection, safety interlock, ultrasonic deterrent

1. INTRODUCTION

Autonomous vehicles (AVs) have progressed from research prototypes to commercially deployed robotaxi fleets, with operators such as Waymo reporting substantial safety improvements over human drivers in serious and fatal accident categories [1]. Despite this progress, current AV perception stacks share a common structural limitation: they are optimized for objects located around the perimeter of the vehicle, typically 1 to 200 m away, and are physically incapable of observing the region between the wheels and beneath the chassis [2]. On a typical passenger platform this creates a blind volume roughly 6.5 in. tall spanning the entire vehicle footprint, an area that small animals actively seek out for shelter and warmth.

The consequences of this blind spot are not hypothetical. In a widely reported 2025 incident, a cat known as KitKat was struck and killed by a Waymo vehicle in San Francisco after moving beneath the stationary car [3]. The vehicle's sensors tracked the animal while it approached, but the moment it passed under the chassis it effectively ceased to exist to the perception system, and the vehicle resumed motion. A human driver in the same situation could have been alerted by bystanders or stepped out to inspect the vehicle; an AV has no equivalent recourse. Because AV sensing suites do not implement object permanence for targets that leave their field of view, any solution must provide direct, persistent observation of the underbody region itself.

This paper presents A.U.R.A. (Autonomous Undercar Risk Avoidance), a senior design project developed at the University of Central Florida by an interdisciplinary team of electrical, computer, and photonic science and engineering students. A.U.R.A. is a self-contained aftermarket system that monitors the underbody of a vehicle during the pre-drive window, when a stationary vehicle poses the greatest risk to sheltering animals. The system detects animal presence using active near-infrared (NIR) imaging and machine-learning inference, prevents the vehicle from moving through a hardware motor interlock, and actively encourages the animal to vacate using species-aware ultrasonic and strobe deterrents. Only once the underbody is verified clear is the interlock released and the vehicle permitted to move. For demonstration, the full system is integrated onto a 12 V ride-on electric vehicle with a 23 in. wheel track and 4.5-to-7.5 in. of ground clearance, with the architecture designed to scale to full-size platforms.

The remainder of this paper describes the system requirements and specifications, the electrical and optical hardware design, the software architecture and detection pipeline, the fabrication of the custom printed circuit boards, and the testing and evaluation results obtained on the demonstration platform.

2. RELATED WORK

Existing technology addresses only fragments of this problem. Kona and Dasari demonstrated that retractable underbody cameras are mechanically feasible, but their system feeds a live view to a human rather than acting autonomously [4]. Sibanda et al. showed that active infrared illumination reliably distinguishes animals at driving speeds, informing

A.U.R.A.'s choice of NIR sensing [5]. Consumer under-hood repellers emit ultrasonic tones or strobes but are strictly passive, with no detection, no verification, and no drivetrain interaction [6]. A.U.R.A. combines all three threads into a closed-loop, verified safety check in which movement is physically inhibited until the underbody is confirmed clear.

3. SYSTEM OVERVIEW AND REQUIREMENTS

A.U.R.A. operates as a pre-drive safety gate. On ignition or wake, linear actuators extend two sealed camera modules from protective housings at the front and rear of the underbody, oriented inward so that their fields of view overlap and cover the full footprint between the axles. Active 850 nm NIR LED rings illuminate the scene, and each monochrome global-shutter camera streams imagery to an onboard single-board computer running a YOLOv11 object detection model. If an animal is detected, a lockout command holds a relay-based motor interlock open, cutting power to the drive motors, while the ultrasonic alarm and strobe deterrents are pulsed. Detection runs continuously until the animal is observed leaving the zone, at which point the interlock is released. Once the vehicle exceeds a 1 mph velocity threshold, measured by a Hall-effect sensor on the wheel hub, the cameras retract into their housings to protect the optics from road debris.

3.1 Key specifications

The governing system-level specifications are summarized in Table 1. The startup safety latency, defined as the time from animal detection to interlock activation, is specified at under 100 ms so that a detection can never race the drive motors. Detection accuracy is specified above 95% for animals with a minimum vertical height of 3.0 in., driven by automotive life-safety standards in which false-negative rates above a few percent are unacceptable. The ultrasonic deterrent must operate in the 30–45 kHz band, covering the aversive hearing ranges of common target species, and the complete integrated system was budgeted below $750 with an approximately 50 W peak power draw from the 12 V vehicle supply.

Table 1. Selected system-level specifications.

Specification	Target value
Startup safety latency (detection to interlock)	<100 ms
Detection success rate	>95%
Minimum detection height	>3.0 in.
System engagement threshold (retraction speed and interlock disabled)	1 mph
Total camera deployment time	<6 s
Ultrasonic repellent frequency	30-45 kHz
Hazard-clearance verification time	5 s
Operational temperature	-10 to +50 °C
Interlock relay switching speed	<10 ms
Total estimated cost	<$750

3.2 Revision of the actuation-time specification

One specification evolved significantly through prototyping and is worth highlighting as a design lesson. The camera deployment time, defined as the time required for the actuators to fully extend the camera modules from their protective housings upon ignition, was originally specified at under 3 s. This early target was set before an actuator had been physically evaluated and was based on the assumption that a short-stroke, high-speed unit would satisfy the mechanical envelope.

Hands-on prototyping proved otherwise. The team first acquired an 8 in. stroke linear actuator based on a preliminary estimate of the deployment distance, only to find that the total installed length of the unit, retracted body plus fully extended rod, exceeded 16 in. and was incompatible with the 4.5-to-7.5 in. underbody clearance of the demonstration vehicle. Re-evaluating the housing geometry identified a 6 in. stroke as the correct size, and within that class the units offering the required force, self-locking behavior, and IP54 environmental rating extend at a rate that cannot meet a 3 s full-stroke target. The specification was therefore formally revised to under 6 s. With the 6 in. stroke commercial linear actuator installed on the vehicle, the measured full extension time is approximately 5 s, comfortably meeting the revised specification while preserving the force margin, holding capability, and ingress protection that a faster, lighter-duty actuator would have sacrificed. The retraction timing analysis at the 1 mph threshold remains valid under the revised value: at 1 mph the vehicle covers approximately 17.6 in. per second, and the sub-100 ms threshold-to-command latency ensures retraction begins promptly and completes before the housing is at risk of ground contact at normal speeds.

4. ELECTRICAL HARDWARE DESIGN

The electrical architecture is split between a main PCB, which hosts the deterrent and interlock subsystems and the primary low-level microcontroller, and secondary PCBs, which handle camera interfacing, proximity sensing, and actuator control at each sensor module. All boards were designed in Fusion 360.

4.1 Actuation subsystem

Three actuation technologies were evaluated for deploying the camera modules: linear electric actuators, servo motors coupled to lead screws, and pneumatic cylinders. Pneumatics were disqualified immediately by the volume and power overhead of a compressed-air supply, and the servo plus 3D-printed linkage approach was rejected because mechanical play at each printed joint would prevent the repeatable optical alignment the NIR camera requires relative to the anamorphic lens assembly. A sealed linear electric actuator provides reversible motion, a compact self-contained form factor, and inherent tolerance to road splash.

The selected unit is a commercially available 12 V linear actuator with a 6 in. stroke, built-in end-of-travel limit switches, and a self-locking gear train that holds the deployed camera position without continuous current draw. Two units are used, one per sensor module. The actuator is driven by a Cytron MD10C motor driver, which accepts a 3.3 V logic-level PWM speed signal and a DIR direction signal from the microcontroller and switches the 12 V, high-current motor load. The MD10C's 10 A continuous rating provides ample margin over the actuator's stall current, and its sign-magnitude interface requires only two GPIO pins.

Two auxiliary sensors close the loop around actuation. A Hall-effect sensor mounted at the wheel hub counts magnet pulses per unit time to measure vehicle velocity, and the microcontroller firmware issues the retraction command within one interrupt service cycle of the 1 mph threshold being crossed. A VL53L0X time-of-flight proximity sensor mounted on the camera module detects potential obstructions in the deployment path during extension, allowing the firmware to halt actuation before a collision.

4.2 Deterrent subsystems

The primary deterrent is a high-frequency ultrasonic alarm implemented as two nearly identical driver chains, one tuned to 40 kHz transducers and the other to 30.5 kHz transducers, so that the output frequency can be matched to the aversive range of the species detected. Each chain consists of three stages: square-wave frequency generation from the microcontroller's LEDC peripheral, a pair of IRS2104 half-bridge gate drivers with bootstrap diode and capacitor networks per the datasheet reference design, and an N-channel MOSFET H-bridge amplifier with a 470 µF and 100 nF capacitor used to filter and decouple. The H-bridge drives an array of five piezoelectric transducers connected in parallel, converting the 12 V supply into a differential drive of approximately 24 V peak-to-peak across the array. Transducers are placed along all edges of the main PCB, plus a downward-facing pair, to maximize acoustic coverage of the underbody volume.

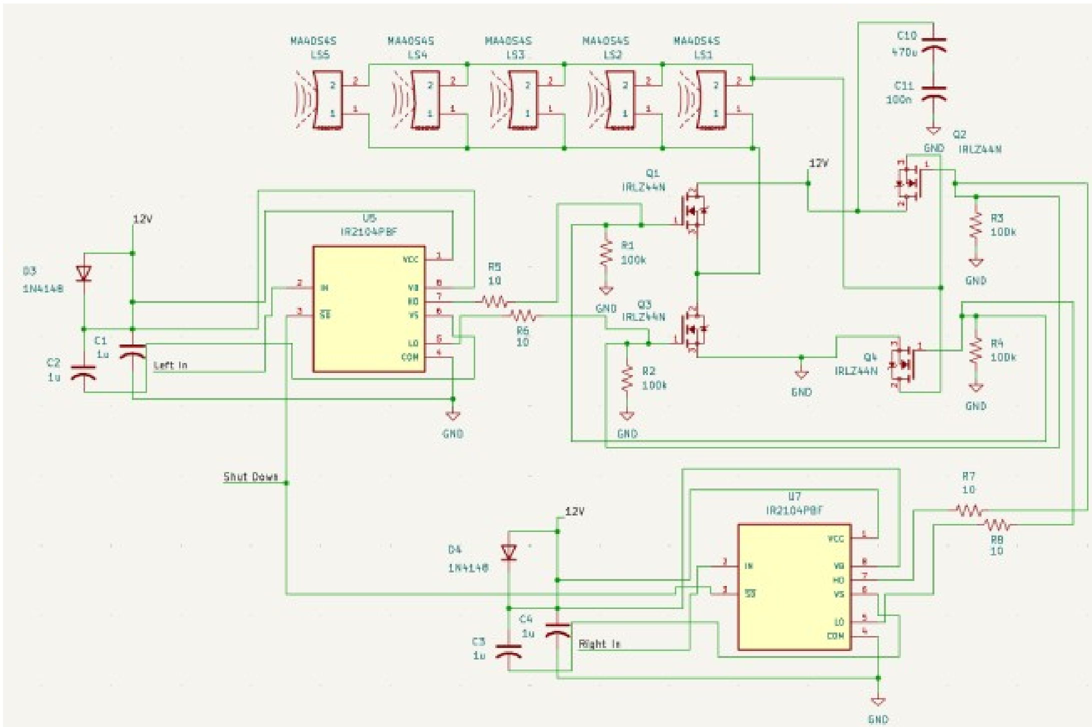

Figure 1. Ultrasonic alarm driver chain: MCU frequency generation, IR2104 half-bridge drivers, N-channel H-bridge, and five-transducer array.

The secondary deterrent is a strobe light array from three lines of five indicator LEDs connected in parallel with 100 Ω current-limiting resistors. A logic-level MOSFET configured as a low-side switch acts as the current source, with a 10 kΩ gate pulldown, and the microcontroller controls pulse timing to produce a disorienting strobe pattern.

4.3 Motor interlock

The motor interlock is the subsystem that converts a detection into physical prevention of movement. A JD1912 automotive SPST relay is installed in series between the demonstration vehicle's 12 V battery and its 10 A fuse. The relay coil is driven by a 2N2222A NPN transistor whose base is connected to a microcontroller GPIO through a series resistor, with a 1N4007 flyback diode parallel with the coil to protect the logic from inductive spikes generated by the dual-motor drive circuit. The interlock is intentionally scoped to the startup sequence only; once the vehicle is in motion the interlock is disabled to avoid unrealistic and hazardous mid-drive stops.

4.4 Processing and communication

An ESP32-S3 serves as the primary low-level controller on the main PCB, generating the alarm frequencies, driving the strobe and interlock, and reading the Hall-effect sensor. It connects to the onboard computer, a Raspberry Pi 5, over a native USB serial link using the ESP32-S3's built-in USB peripheral on GPIO 19 and 20. Each of the two camera modules is hosted by a dedicated Raspberry Pi 4 Model B, which interfaces with its IMX178M sensor over the CSI ribbon connection and streams the video feed wirelessly to the central Pi 5. The Raspberry Pi 5 performs all image processing and machine-learning inference on the received streams and presents system status on a 7 in. LCD touchscreen dashboard. Communication between the main controller and the distributed sensor-module controllers is likewise carried over Wi-Fi, which eliminates long wiring runs through the chassis; the local Pi-to-ESP32 link uses wired UART for deterministic low-latency command delivery. Every command message carries an identifier and status code, and each controller returns an acknowledgment so that lost or malformed messages leave the system in its safe state.

4.5 Power architecture

The entire system operates from the vehicle's 12 V supply, with an estimated 50 W peak budget. Regulation is deliberately modular: each rail is implemented as a separate plug-in regulator board on a 4-pin header so that a failed rail can be isolated and replaced independently. Every board is built around the LM2679-ADJ, a 5 A adjustable step-down switching regulator whose output voltage is programmed by a feedback resistor divider. One board is set to 3.3 V for the microcontrollers and logic-level sensors, one to 5 V for the 5 V peripherals including the motor driver logic, and one to 5.1 V for the Raspberry Pi 5 and Pi 4 camera nodes, whose combined current demand sits comfortably within the LM2679's 5 A capability. Using a single adjustable part across every rail means one board design, one bill of materials, and interchangeable spares. Ground plane pours on both copper layers of every board reduce impedance and electromagnetic interference.

Table 2. Estimated peak power budget.

Component	Supply voltage	Estimated current draw
Raspberry Pi 5	5 V	2,500 mA (peak)
Raspberry Pi 4 (×2)	5.1 V	3,000 mA total
ESP32-S3 (×3)	3.3 V	240 mA total
NIR LED array (×2 modules)	12 V	500 mA total
Linear actuators (×2, peak)	12 V	1,000 mA total
Cytron motor drivers (×2)	12 V	200 mA idle
LCD dashboard (7 in.)	5 V	800 mA
Sensors, relay, misc.	3.3-5 V	~200 mA
Total estimated peak		~54.0 W

5. OPTICAL DESIGN

The underbody sensing geometry is the central optical challenge of the project. While the physical mounting bracket constraint sits at a 6 in. horizontal standoff, angling the dual-camera configuration extends the active optical path length to an effective 6.5 in. optical standoff distance. Based on this geometry, scaling the system to a full-size AV with a 64.7 in. wheel track and 6.5 in. clearance yields a required horizontal field of view of approximately 157.3 degrees against a vertical field of 53.1 degrees. On the demonstration vehicle with a 23 in. wheel track and 4.5-to-7.5 in. clearance, the geometry relaxes to 121 degrees horizontal by 44.3 degrees vertical. Because the scene remains far wider than it is tall, a conventional spherical lens imaging this region would waste a large fraction of sensor pixels on an empty vertical field.

5.1 Active NIR illumination

Illumination is provided by rings of 850 nm NIR LEDs distributed concentrically around each camera aperture, producing uniform coverage that a single point source could not achieve. The 850 nm band was selected because biological materials such as fur and skin exhibit relatively high reflectivity there while asphalt and painted surfaces remain strongly absorptive, so animals appear as high-intensity features against a dark background regardless of ambient lighting. NIR also remains invisible to the human eye and, unlike long-wave thermal imaging, is immune to confounding heat signatures from pavement, tires, and drivetrain components. Compatibility with inexpensive silicon CMOS sensors was the final deciding factor over thermal alternatives.

5.2 Anamorphic field compression

To map the established wide fields efficiently, the optical train incorporates an anamorphic telescope formed by an N-BK7 cylindrical plano-concave lens (Thorlabs LK1336L1-B) and a fused-silica cylindrical plano-convex lens (Newport CSX200AR.16). This configuration provides a 1.82× horizontal squeeze to compress the demonstration vehicle's 121-degree field into the native sensor aspect ratio; the same architecture scales up to a 1.98× squeeze for full-size vehicle profiles via independent mechanical rail adjustments.

The compressed beam passes through a 12.5 mm round 850 nm narrow bandpass filter (30 nm FWHM) for ambient solar rejection before entering a pre-mounted Thorlabs C340TMD-B aspheric objective (4.03 mm focal length). A 30 nm FWHM bandwidth was selected over a narrower 10 nm baseline to prevent angular blue-shifting at the steep boundaries of the wide field and to capture the full emission profile of the 850 nm illumination LEDs. While first-principles calculated an ideal 3.0 to 3.2 mm objective focal length to bound the vertical field, micro-scale optics in that range lacked the back focal distance to clear the internal mirror assembly. The 4.03 mm objective resolves this mechanical collision constraint while maintaining vertical margins, and its 12.5 mm circular filter cap spans the housing without introducing geometric vignetting.

A Sony IMX178M monochrome CMOS sensor captures the image, maximizing NIR quantum efficiency and contrast across the scene. The image-processing software applies a digital de-squeeze prior to analysis to restore the true aspect ratio.

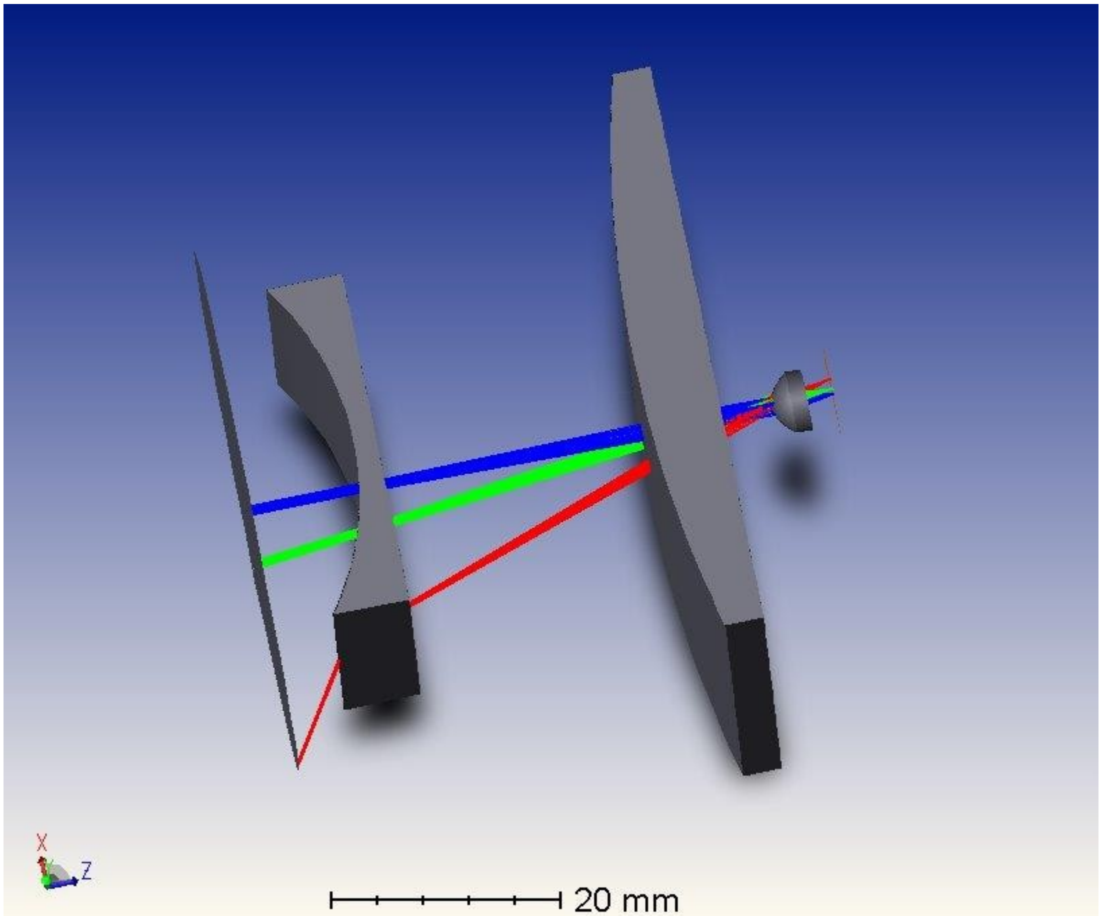

Figure 2. 3D lens layout modeled in Ansys Zemax OpticStudio demonstrating horizontal beam compression across the undercar tracking assembly.

6. SOFTWARE DESIGN

6.1 Layered architecture

The software follows a layered architecture that strictly separates perception, decision making, communication, and hardware control. The perception layer on the Raspberry Pi captures and preprocesses frames and runs inference; the decision layer converts detections into system states; the communication layer exchanges validated messages with the microcontrollers; and the control layer, running on the ESP32 controllers, is the only code permitted to touch safety-critical hardware signals. This structure guarantees that a fault in the detection code cannot directly energize or de-energize the interlock, and it allows each module to be developed and tested independently.

6.2 Detection pipeline

Frames from the front and rear cameras are captured over the CSI interface by their Raspberry Pi 4 host nodes, streamed wirelessly to the Pi 5, and preprocessed with OpenCV: resizing to the model's fixed input resolution, normalization, and contrast and brightness correction tuned for the NIR illumination conditions. Detection is performed by a YOLOv11 model [7]; object detection was chosen over simple classification because the system must localize the animal within the underbody region, not merely confirm its presence. The nano variant of the model was selected as the best balance of speed and accuracy on the embedded processor, which has no dedicated GPU. Because no public dataset exists for under-vehicle

NIR imagery, a custom dataset was annotated, augmented, and trained through the Roboflow platform, with training performed offline and the resulting weights deployed to the runtime.

6.3 Decision logic and fail-safe behavior

The decision module enforces the safety contract. A detection above the confidence threshold places the system in a hazard state: the alarm is pulsed, the strobe fires, and the interlock remains engaged. To suppress false triggers from noise or transient shadows, the logic requires multiple consecutive confirming frames before entering the hazard state, and it requires a sustained period with no detections before declaring the area clear. Critically, every failure mode collapses to the safe state: a disconnected camera, a crashed detection process, a communication timeout, or malformed data all leave the interlock engaged. The vehicle can only move when the detection system affirmatively confirms clearance, so no software error can silently disable the safety function.

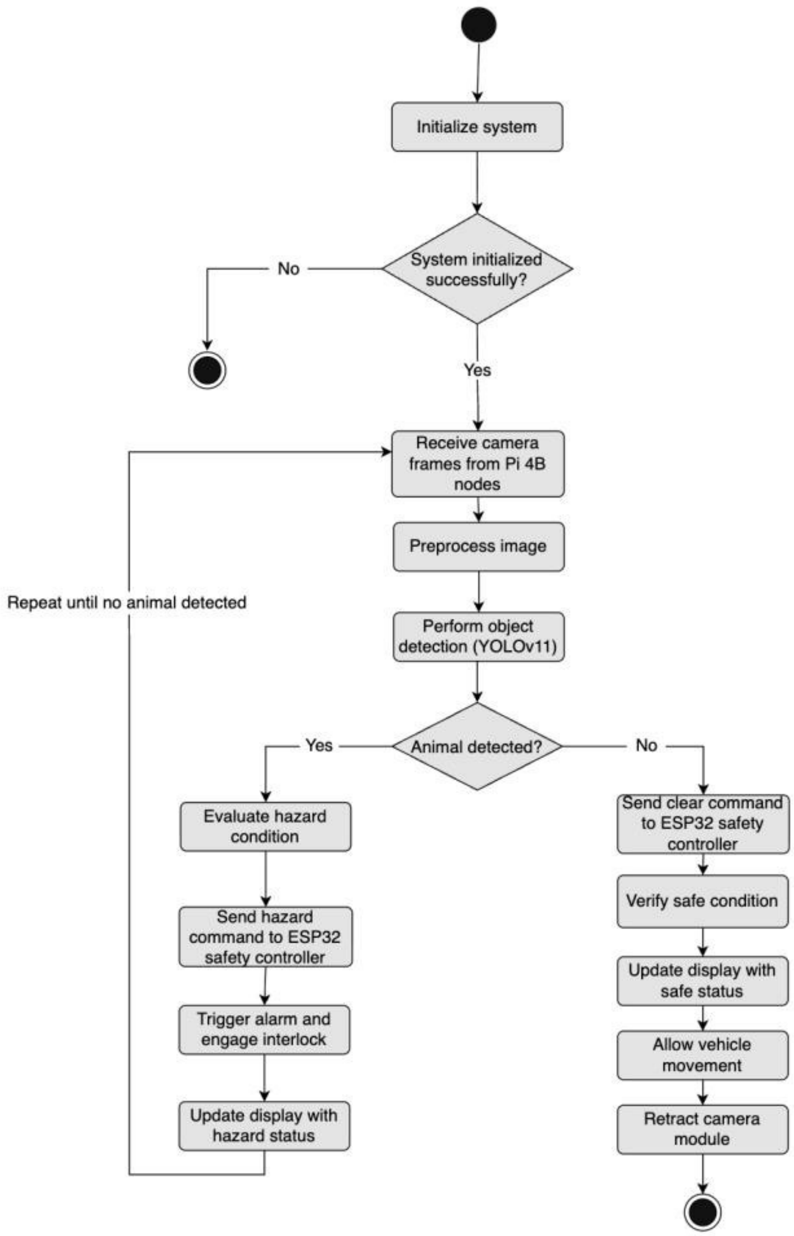

Figure 3. Software safety-state flow showing initialization, YOLOv11 detection, hazard response, clearance verification, and fault handling.

6.4 Human-machine interface

A 7 in. LCD dashboard provides the human supervision link. In normal operation, it displays a simple go/no-go status indicator. When a hazard is detected, the display automatically interrupts to a split-screen live view from the front and rear cameras, allowing immediate visual confirmation of the detection, and presents a warning instructing the operator to check beneath the vehicle. Debug information such as detection confidence and communication status is available during development.

7. SYSTEM FABRICATION

All boards were laid out in Fusion 360 following a structured design-to-manufacture workflow with attention to component placement, controlled trace routing, and ground plane pours on both copper layers.

The main PCB integrates the ESP32-S3 controller, both ultrasonic driver chains, the strobe driver, the interlock relay circuit, the Hall-effect sensor connector, the USB link to the Raspberry Pi, and the 12 V power entry through a screw terminal. The transducers occupy the board edges plus a downward-facing pair for omnidirectional coverage, and the 3 W LED star boards are thermally coupled to the PCB with thermal paste and connected through screw terminals so the heatsinks can be serviced without rework. Power trace widths are graded by rail: 100 mil for 12 V, 50 mil for 5 V, and 33 mil for 3.3 V and signals.

The secondary PCB at each sensor module hosts the camera interface, the VL53L0X time-of-flight sensor on a shared I2C bus [8] with software-controlled shutdown, and the 3-pin PWM/DIR header for the Cytron MD10C actuator driver. The high-current 12 V actuator path is routed through dedicated headers isolated from the 3.3 V logic domain so that motor current draw cannot brown out the controller. High-speed camera differential pairs were length-matched and routed with controlled impedance. The modular regulator boards plug into both PCBs on labeled 4-pin headers.

The camera and optics stack is enclosed in a 3D-printed housing with a recessed, anti-reflection-coated protective window as the environmental barrier. The NIR LEDs are positioned concentrically around the window, and an internal light-blocking wall separates the emitters from the imaging optics so that illumination must reflect off the ground before reaching the sensor, suppressing glare. Printing the housing in-house allowed rapid iteration of internal spacing as the optical design was refined.

8. TESTING AND EVALUATION

8.1 Optical validation

The optical architecture was validated via sequential ray tracing in Ansys Zemax OpticStudio followed by physical prototype evaluation. Zemax simulations utilizing exact cylindrical lens parameters at 850 nm established a baseline 1.82× horizontal squeeze across a 121° field of view.

Physical hardware validation was conducted using a custom-designed, 3D-printed optomechanical telescoping enclosure that houses a Raspberry Pi computing module and a wireless live-stream telemetry pipeline. This enclosure incorporates precision internal alignment features to constrain the cylindrical pair and the aspheric objective lens at the determined axial spacing required for system focus. At a 6.5 in. working distance, a physical calibration target with 1.0 in. horizontal and 0.5 in. vertical grid line spacing was imaged under active 850 nm near-infrared illumination (Figure 4).

While the resulting baseline capture exhibits pixelation, peripheral focus degradation, and a minor optical centration offset characteristic of additive manufacturing tolerances, the image successfully provides functional verification of uniform aspect ratio compression across the active field window. These alignment and centration offsets have since been resolved through iterative mechanical mounting bracket adjustments. Although the raw resolution limits of the initial prototype are apparent, the captured spatial contrast and structural geometry are anticipated to meet the threshold requirements for downstream YOLO-based object detection models.

Figure 4. Laboratory near-infrared image capture of a physical calibration grid (1.0 in. horizontal by 0.5 in. vertical spacing) utilized to verify uniform horizontal compression and spatial resolution across the active field.

8.2 Hardware subsystem testing

The motor interlock was breadboarded with the JD1912 relay, 2N2222A driver, and 1N4148 flyback diode, with an indicator LED across the switched contacts. Across repeated toggle cycles the relay opened and closed the load path exactly as commanded, with the supply current confirming contact state, validating the design for integration. The strobe circuit was tested in its final configuration and produced the high-brightness output required to act as a visual irritant.

The ultrasonic alarm was characterized stage by stage on an oscilloscope at 30.5 kHz. The ESP32-S3 LEDC peripheral produced stable square waves at the driver inputs, and differential measurement across the H-bridge output nodes confirmed correct out-of-phase switching with the expected 24 V peak-to-peak amplification at the target frequency. The circuit was tested by applying two oscilloscope probes to each pin of the 5th array transducer in the array then using the math function the channels were subtracted and the behavioral signal was able to be observed. Driving the full array of transducers resulted in the necessary 30.5 kHz at over double the input of 12 Vpp. This is a stark improvement to issues that arrived during the initial testing and design of the amplifier circuit that contained four feedback diodes instead of the filtering capacitors where the output would fail to show all three characteristics necessary, including the square wave output, the amplified peak-to-peak voltage, and the necessary frequency. While the signal observed below appears to be noisy, this is most likely due to the switching of the MOSFETs within the H-bridge amplifier.

Figure 5. Ultrasonic alarm characterization: amplified 12 V output at 30.5 kHz.

The actuation subsystem was verified against the revised deployment specification described in Section 3.2. With the 6 in. stroke actuator driven by the MD10C at the full 12 V rail, full extension of the camera module completes in approximately 5 s, meeting the under 6 s requirement, and the self-locking gear train holds the deployed position with zero holding current.

8.3 Detection performance

The trained YOLOv11 model was evaluated on a held-out test set using standard object detection metrics. The model achieved a mean average precision (mAP@50) of 91.7%, a precision of 88.0%, and a recall of 92.1%, for an F1 score of approximately 90.0%. The high recall is the most important figure for a life-safety application, since a missed animal is the failure mode the system exists to prevent, while the precision indicates a low rate of nuisance triggers. In preliminary qualitative testing on under-vehicle imagery, the model correctly localized animals across varied lighting conditions and viewing angles while rejecting ground texture and vehicle structure. These results represent preliminary model performance prior to full end-to-end vehicle integration.

Table 3. YOLOv11 detection performance.

Metric	Result
Mean average precision (mAP@50)	91.7%
Precision	88.0%
Recall	92.1%
F1 score	~90.0%

Figure 6. YOLOv11 detection outputs on under-vehicle imagery showing localized animal detections with confidence scores.

8.4 System-level testing

System-level behavior was validated by exercising the full state machine with simulated detection outputs. Triggering the hazard state correctly interrupted the LCD to the warning display, held the interlock, and returned to the clear state on command, confirming that the interfaces between the perception, decision, communication, and control layers behave as designed. Full real-time end-to-end integration on the vehicle, with the final camera hardware in the optical housings, is the remaining step ahead of the final demonstration.

9. CONCLUSION

A.U.R.A. demonstrates that the underbody blind spot of autonomous vehicles, a documented cause of animal fatalities, can be closed with a targeted, practical, and affordable sensing system. The project's central contributions are an anamorphic NIR optical design that maps an extreme-aspect-ratio scene efficiently onto a low-cost CMOS sensor, a fail-safe layered software architecture in which every fault leaves the vehicle immobilized, and a closed-loop deterrent-and-verify response that goes beyond the passive repellers available to consumers today. Validated subsystem performance includes 91.7% mAP@50 detection with 92.1% recall, a functional hardware interlock, a 24 Vpp ultrasonic drive stage, and camera deployment in approximately 5 s against the revised under 6 s specification.

The evolution of the actuation specification from 3 s to under 6 s illustrates a broader lesson from the project: specifications written before hardware is in hand must be revisited against physical constraints, in this case the collapsed length of a linear actuator rather than its stroke alone. Future work includes an improved ultrasonic output stage capable of driving the full transducer array, integration of the final camera hardware into the optical housings, expansion of the NIR training dataset across additional species and surfaces, and scaling trials on a full-size vehicle. The modular architecture, built entirely from commercially available components, provides a direct path to those extensions without redesign.

ACKNOWLEDGEMENTS

We extend our sincere thanks to the professors who generously volunteered their time to review this project. We are also grateful to the Department of Electrical and Computer Engineering and the College of Optics and Photonics for the laboratory facilities, equipment, and resources that made the design, fabrication, and testing of this system possible.

REFERENCES

[1] Waymo, “Waymo Safety Impact Report: Comparison of the Waymo Driver’s Performance with Human Benchmarks,” Waymo LLC, Sep. 2024. [Online]. Available: https://waymo.com/safety/impact/
[2] Z. Dai, A. Wolf, P.-P. Ley, T. Glück, M. C. Sundermeier, and R. Lachmayer, “Requirements for Automotive LiDAR Systems,” Sensors, vol. 22, no. 19, p. 7532, 2022.
[3] H. Knight, “How Kit Kat Was Killed: Video Shows What a Waymo Couldn’t See,” The New York Times, Dec. 5, 2025.
[4] S. Kona and G. Dasari, “Under-Car Vision: A Retractable 360-Degree Camera System for Enhanced Safety,” IJFMR, vol. 7, no. 1, 2025.
[5] V. Sibanda, K. Mpofu, J. Trimble, and N. Zengeni, “Design of an Animal Detection System for Motor Vehicle Drivers,” Procedia CIRP, vol. 84, pp. 755–760, 2019.
[6] S. Schumake, “Electronic Rodent Repellent Devices: A Review of Efficacy Test Protocols and Regulatory Actions,” National Wildlife Research Center Repellents Conference, 1995.
[7] Ultralytics, “YOLO Documentation,” 2023. [Online]. Available: https://docs.ultralytics.com
[8] NXP Semiconductors, “I2C-Bus Specification and User Manual,” UM10204, 2014.