

A.U.R.A

Autonomous Undercar Risk Avoidance



University of Central Florida
Department of Electrical Engineering and Computer Science
College of Optics and Photonics (CREOL)

Group 2

Theresa Dinh
Oussama Benamar Hlimi
Christopher Nunez-Velasquez
Kimberly Lugo

Optics and Photonics
Electrical Engineering
Electrical Engineering
Computer Engineering

Reviewers

Dr. Mike Borowczak
Dr. Darren Hudson
Prof. Mark Maddox

ECE
CREOL
ECE

Associate Professor
Associate Professor
Associate Professor

Advisors

Dr. Sonali Das
Dr. Paul Leisher

ECE
CREOL

Lecturer
Professor

Contents

List of Figures	i
List of Tables	iii
Chapter 1 – Executive Summary	1
Chapter 2 - Project Description.....	1
2.1 Background and Motivation.....	1
2.2 Related Works and Existing Products	2
2.2.1 Retractable Underbody Camera System.....	3
2.2.2 Front-facing Animal Detection.....	3
2.2.3 Comparison with Consumer Under-Hood Repellers.....	4
2.3 Goals and Objectives.....	5
2.3.1 Goals.....	5
2.3.2 Objectives	6
2.4 Required Specifications.....	7
2.5 Descriptions of Features and Functionalities	9
2.5.1 Hardware	9
2.5.2 Software.....	11
2.5.3 Optical Setup	13
2.5.4 Market Analysis and Comparable Products	16
Chapter 3 – Research and Investigation.....	19
3.1 Optical Hardware	19
3.1.1 Imaging System Selection	19
3.1.2 Anamorphic System Component Selection	20
3.1.3 Imaging Spectral Band Selection	21
3.1.4 Active Illumination Source Selection.....	23
3.1.5 NIR Emitter Wavelength Selection	24
3.1.6 Image Sensor Technology Selection	25
3.1.7 Image Sensor Part Selection	26
3.1.8 Optical Material Selection	27
3.1.9 Spectral Filter Selection.....	28

3.1.10 Protective Window Selection	29
3.1.11 Housing Material Selection	30
3.2 Hardware	30
3.2.1 Actuator	30
3.2.2 Motor Driver.....	33
3.2.3 Hall Effect Sensor.....	35
3.2.4 Proximity Sensor	37
3.2.5 Processor.....	40
3.2.6 Wireless Antenna.....	44
3.2.7 High-Frequency Alarm:.....	46
3.2.8 LED Strobe Light:	51
3.2.9 Motor Interlock.....	52
3.2.10 Power Supply and Battery	54
3.2.11 Voltage Regulator.....	58
3.3 Communication Protocol.....	59
3.3.1 Communication Technology Comparison.....	59
3.3.2 Communication Module Comparison	61
3.3.3 Data Flow and Message Structure	62
3.4 Software	63
3.4.1 Software Platform.....	63
3.4.2 Image Processing Software	66
3.4.3 Detection Algorithm	68
3.4.4 Programming Language and Development Framework.....	70
Chapter 4 - Standards and Design Constraints	73
4.1 Industrial Standards.....	73
4.1.1 Optical Industrial Standards	74
4.1.2 Wireless Communication Standards.....	76
4.1.3 Camera Interface Standard	77
4.1.4 Generic Standard on Printed Board Design.....	79
4.2 Design Constraints	80
4.2.1 System Hardware Design Constraint.....	80

4.2.2 Optical Design Constraints	82
4.3 Practical Constraints.....	84
4.3.1 Time Constraints.....	84
4.3.2 Economic Constraints.....	85
4.3.3 Environmental Constraints	85
4.3.4 Manufacturability and Integration Constraints	86
4.3.5 Hardware Availability and Component Constraints.....	87
4.4 Other Constraints.....	87
4.4.1 Regulatory and Road Safety Constraints.....	87
4.4.2 Camera Retraction Speed Threshold	88
4.4.3 Animal Detection Accuracy Constraint.....	88
4.4.4 Scalability to Full-Size Vehicles	89
Chapter 5 - Application of ChatGPT and Other Similar Platforms	89
5.1 Case Study 1 – Contrast Improvement Technique.....	89
5.2 Case Study 2 – Motor Driver Recommendation	90
5.3 Case Study 3 – Suitable Communication Protocol	91
5.4 Case Study 4 – Suitable Ultrasonic Sound Component	92
Chapter 6 - Physical System Design.....	93
6.1 Hardware Design.....	93
6.1.1 Main PCB	93
6.1.2 Secondary PCBs.....	99
6.1.3 Voltage Regulators	104
6.2 Optical Design.....	106
6.2.1 Design from Geometric Requirements	106
6.2.2 First-Order Analytical Validation.....	107
6.2.3 Computational Modeling in Python	110
6.2.4 Physical System Validation in Ansys Zemax OpticStudio	111
Chapter 7 – Software Design	113
7.1 Software Design Overview	113
7.2 Software Architecture	114
7.3 Camera Capture and Image Acquisition	116

7.4 Image Preprocessing Pipeline	117
7.5 Object Detection Using YOLOv11	117
7.6 Decision Logic and Safety Rules	118
7.7 Communication Between Processor and Controllers	120
7.8 Display and User Interface Software	122
7.9 Startup Sequence and Initialization	123
7.10 Failure Handling and Safe Operation	123
7.11 Software Design Summary	123
Chapter 8 – System Fabrication/Prototype Construction	124
8.1 Primary PCB	124
8.2 Secondary PCB	125
8.3 Regulators	126
8.3.1 3.3 V Regulator Board	127
8.3.2 5 V Regulator Board	127
8.3.3 5.1V 5A Regulator Board	127
Chapter 9 - System Testing and Evaluation	128
9.1 Optoelectronics Feasibility and Study Testing	128
9.1.1 Early Proof of Concept Demonstration	129
9.1.2 Real Lens System Evaluation	131
9.1.3 Optoelectronic Feasibility Summary	133
9.2 Hardware Testing	133
9.2.1 Motor Interlock	133
9.2.2 3W LED Strobe	134
9.2.3 Ultrasonic Alarm	134
9.3 Software Testing	137
9.4 Detection Performance Evaluation	141
9.5 System-Level Testing	142
Chapter 10 - Administrative Content	142
10.1 Budget	142
10.2 Bill of Materials	142
10.3 Work Distribution	144

10.4 Senior Design Project Milestones	146
Chapter 11 – Conclusion.....	150
Appendices.....	157
Appendix A - References	157
Appendix E - ChatGPT Prompts and Outcomes	160

List of Figures

Figure 2.5.1 3D Rendering of Sensor Placement on Demo Car	11
Figure 2.5.1.1 Sensor Placement on Demo Car	12
Figure 2.5.2 Hardware Diagram	13
Figure 2.5.2.1 Revised Hardware Diagram	13
Figure 2.5.3 Software System Flowchart	15
Figure 2.5.3.1 Revised Software System Flowchart	16
Figure 2.5.4 Expected underbody front and rear view of demo vehicle captured using Galaxy S22+ at 0.6x zoom with a stand-off distance of 8.5 in.	18
Figure 2.5.6 House of Quality	24
Figure 2.5.6.1 Revised House of Quality	25
Figure 3.2.10.3.B: A.U.R.A. Power Distribution Architecture	64
Figure 6.1.1 High-Frequency Alarm	106
Figure 6.1.1.1 Revised High-Frequency Alarm	107
Figure 6.1.1.2 3W LED Strobe-Light	109
Figure 6.1.1.2.1 Revised 3W LED Strobe-Light	110
Figure 6.1.1.3 Main PCB MCU	111
Figure 6.1.1.4 USB to PI5	112
Figure 6.1.1.5 Hall-Effect Connector	113
Figure 6.1.1.6 Vehicle Motor Interlock	113
Figure 6.1.2.A Secondary PCB Schematic part 1	115
Figure 6.1.2.B Secondary PCB Schematic part 2	115
Figure 6.1.2.2 IMX178M Camera Module Connector	117
Figure 6.1.2.3.A VL53L0X Time-of-Flight Sensor Connector	118
Figure 6.1.2.4 Figure X: Cytron MD10C Motor Driver Connector	118
Figure 6.1.3.1 3.3V Regulator	119
Figure 6.1.3.2 5V Regulator	120
Figure 6.1.3.3 5.1V, 5A Regulator	120
Figure 6.1.3.4 Revised Regulators 5.2V/3.3V, 5A/1A	121
Figure 6.2.1 Shaded 3D System Layout (Isometric Detail View)	128

Figure 7.1.1 UML Use Case Diagram	130
Figure 7.2.1 A.U.R.A. Software Architecture Diagram	132
Figure 7.2.1.1 Revised A.U.R.A. Software Architecture Diagram	132
Figure 7.2.2 Class Diagram	133
Figure 7.6.1 UML Activity Diagram for Startup Safety Verification	137
Figure 7.6.1.1 Revised UML Activity Diagram for Startup Safety Verification.....	138
Figure 7.6.2 Software State Diagram	139
Figure 7.6.2.1 Revised Software State Diagram	140
Figure 7.7.1 Communication Between Processor and Controllers.....	142
Figure 7.7.1.1 Revised Communication Between Processor and Controllers	143
Figure 7.7.2 Sequence Diagram	144
Figure 7.7.2.1 Revised Sequence Diagram	145
Figure 8.2 Secondary PCB Design.....	151
Figure 8.2.1 Revised Secondary PCB Design	152
Figure 8.3 3.3V and 5V Regulator Board Design.....	153
Figure 8.4 5.1V, 5A Regulator Board Design	154
Figure 8.5 Revised Regulator Board Design.....	154
Figure 9.1.1 Simplified physical demonstration of optical system behavior	157
Figure 9.1.2 Geometric Image Simulation (Anamorphic Mapping)	158
Figure 9.2.1 Behavior of H-Bridge at Output Nodes	161
Figure 9.2.2 Behavior of Output at Single Transducer	162
Figure 9.2.3 Behavior of Output at Transducer Array	162
Figure 9.2.4 Final 30 kHz Output.....	164
Figure 9.2.5 Final 40 kHz Output.....	164
Figure 9.2.6 Final 45 kHz output.....	165
Figure 9.3.1 Raw image captured from an under-vehicle camera showing the ground surface and vehicle structure.....	166
Figure 9.3.1.1 Finalized demo vehicle structure and lighting.....	167
Figure 9.3.2 Detection output showing identified animal beneath the vehicle with confidence score	168
Figure 9.3.3 Additional detection example demonstrating model performance under different lighting conditions	168

Figure 9.3.4 LCD display showing hazard detection state and warning message. .	170
Figure 9.3.4.1 Finalized LCD display showing hazard detection state and warning message.....	170
Figure 9.3.4.2 Finalized LCD display showing clear state.	171

List of Tables

Table 2.2.1 Technology Comparison to System	5
Table 2.4.1 Overall System Specifications	7
Table 2.4.2 Component-Level Specifications.....	9
Table 3.1.1 Imaging System Comparison	27
Table 3.1.2 Anamorphic System Component Comparison	28
Table 3.1.3 Imaging Spectral Band Comparison	29
Table 3.1.4 Active Illumination Source Comparison.....	30
Table 3.1.5 NIR Emitter Wavelength Comparison.....	31
Table 3.1.6 Image Sensor Technology Comparison.....	32
Table 3.1.7 Image Sensor Part Comparison	35
Table 3.1.8 Optical Material Comparison	36
Table 3.1.9 Spectral Filter Comparison	37
Table 3.1.10 Protective Window Comparison	38
Table 3.1.11 Housing Material Comparison.....	38
Table 3.2.1.1: Actuator Technology Comparison	39
Table 3.2.1.2: Actuator Part Comparison	41
Table 3.2.2.1: Motor Driver Technology Comparison	42
Table 3.2.2.2: Motor Driver Part Comparison.....	43
Table 3.2.3.1: Hall Effect Sensor Technology Comparison	44
Table 3.2.3.2: Hall Effect Sensor Part Comparison.....	45
Table 3.2.4.1: Proximity Sensor Technology Comparison	47
Table 3.2.4.2: Proximity Sensor Part Comparison	48
Table 3.2.5: Processing Platforms Technology Comparison	49
Table 3.2.5.1: MCU Technology Comparison	50
Table 3.2.5.2: MCU Part Comparison	51
Table 3.2.6.1: Wireless Antenna Technology Comparison	52
Table 3.2.6.2: Wireless Antenna Part Comparison	53
Table 3.2.7.1 Frequency Generator Technology Comparison:.....	55

Table 3.2.7.2 Frequency Generator Component Comparison:	55
Table 3.2.7.3 High-Frequency Output Technology Comparison:	56
Table 3.2.7.4 High-Frequency Output Component Comparison	57
Table 3.2.7.5 Amplifier Technology Comparison	58
Table 3.2.7.6 Amplifier Component Comparison	59
Table 3.2.8.1 LED Technology Comparison:	60
Table 3.2.8.2 LED Component Comparison	60
Table 3.2.9.1 Motor Interlock Technology Comparison:	61
Table 3.2.9.2 Motor Interlock Component Comparison	62
Table 3.2.10.3.A: A.U.R.A. Estimated Peak Power Budget	62
Table 3.2.10.2: Battery Technology Comparison	63
Table 3.2.10.3.C: Battery Component Comparison	65
Table 3.2.11.1 Voltage Regulator Technology Comparison	66
Table 3.2.11.2 Voltage Regulator Component Comparison	67
Table 3.3.1 Communication Technology Comparison	69
Table 3.3.2 Communication Module Comparison	71
Table 3.4.1.1 Software Platform Technology Comparison	73
Table 3.4.1.2 Software Platform Part Comparison	75
Table 3.4.2.1 Image Processing Technology Comparison	76
Table 3.4.2.2 Library Interface Comparison	77
Table 3.4.3.1 Detection Technology Comparison	78
Table 3.4.3.2 Detection Model Comparison	79
Table 3.4.3.3 Dataset and Training Workflow Comparison	80
Table 3.4.4.1 Programming Language Technology Comparison	81
Table 3.4.4.2 Development Framework Comparison	83
Table 6.2.1 Surface-by-Surface Ray Trace	124
Table 6.2.2 Python Optimization Engine Results	126
Table 10.2.1 Bill of Materials	172
Table 10.2.1.1 Revised Bill of Materials	173
Table 10.3.1 Distribution of Worktable	175
Table 10.4.1 Senior Design Project Milestones	177

Chapter 1 – Executive Summary

A.U.R.A., the Autonomous Undercar Risk Avoidance system, is a senior design project developed at the University of Central Florida by a team of undergraduate students in electrical, computer, and optical engineering. The system is designed to detect small animals sheltering beneath autonomous vehicles and prevent the vehicle from moving until the animal has safely cleared the area. As autonomous vehicles become increasingly prevalent, ensuring they can recognize and respond to animal hazards in their immediate surroundings is a critical safety gap that current self-driving systems do not adequately address, A.U.R.A. aims to close that gap.

The system uses active near-infrared illumination, global shutter CMOS image sensors, and custom anamorphic optical assemblies to monitor the full undercarriage of the vehicle. Cameras are positioned at both the front and rear of the underbody, capturing the area behind and in front of all wheels to eliminate blind spots and ensure complete coverage of the undercar region. Imagery from both cameras is processed in real time by a Raspberry Pi 5 running a YOLOv11 object detection model, which identifies animal presence and distinguishes it from road debris and other non-biological objects. Upon a confirmed detection, a multi-modal deterrent response is triggered consisting of a high-intensity strobe light and dual-frequency ultrasonic transducer arrays, while a motor interlock circuit simultaneously prevents the vehicle from moving until the area is clear.

The final detection model was implemented using YOLOv11 Nano. This specific version was used because YOLOv11 Nano was the YOLOv11 model option provided through the selected Roboflow training workflow

The hardware is built around two custom PCBs handling sensing, processing, communication, and deterrent control, supported by a custom anamorphic optical system designed to maximize underbody field of view from a compact mounting position. A.U.R.A. is intended as an aftermarket accessory compatible with a wide range of passenger vehicles, with a modular architecture that supports future scalability into fleet, commercial, and fully autonomous vehicle platforms.

This document covers the full design process of A.U.R.A., including theoretical foundations, component selection, system architecture, hardware schematics, firmware, fabrication, testing, and administrative considerations, concluding with recommendations for future development.

Chapter 2 - Project Description

2.1 Background and Motivation

The rise of autonomous vehicles (AVs) has made what was once considered science fiction into reality. The massive potential for this technology has already been proven by the fact that Waymo, an AV tech company that offers “robotaxi” services, reports their vehicles to be 90% safer than human drivers in terms of serious or fatal accidents [1]. However, there is still a massive blind spot in these vehicles that need to be addressed before we can safely consider them to be a better alternative to a human at the wheel. This gap in safety is the fact that object performance is not a feature that these vehicles possess, so they fail to track objects that are temporarily out of their “sight”.

A human driver is able to interact with people nearby and get out to troubleshoot if they notice something amiss with the environment surrounding their car, but an AV is incapable of this. If a subject happens to move into the vehicle’s blind spot, it would be as if that subject ceases to exist to the vehicle. One example of such a problem is the incident involving KitKat, a San Francisco cat that was struck and sadly killed by a Waymo [2]. The sensors on the Waymo saw the cat approaching, but the moment KitKat disappeared from the sensor’s line of sight to rest directly under the car, the vehicle could no longer perceive the cat that was still there. There was even a person that walked up to the car and appeared to be checking on the cat, yet when that person walked away, the vehicle just assumed the path was clear and began to move. We are hoping to prevent situations like this from recurring with our proposed sensor and deterrent system, A.U.R.A. or Autonomous Undercar Risk Avoidance.

Most current autonomous sensing systems cover a 360° view around the top and sides of the vehicle using multiple cameras and sensor fusion [3]. Despite this remarkable achievement, these systems do not account for the area beneath the vehicle. This leaves a 6.5” blind spot between the pavement and the vehicle underbody and monitoring the entirety of this area is difficult due to how wide yet short the space is. The latest Waymo vehicles typically have a 64.7” front wheel track and so to be able to cover that entire width at close range would most likely result in numerous blind spots caused by the wheels or even more sensors to make up for that [4]. A.U.R.A. aims to fill this safety void using a custom dual-sensor system with active near-infrared (NIR) illumination placed at the front and rear of the vehicle to provide a clear, high-contrast view of the area underneath and will enable the immediate detection of any live animal that would otherwise be nonexistent to the AV’s primary sensors. Beyond simple detection, the system is designed to trigger a deterrent to ensure that the vehicle can safely clear its path and resume operation without remaining indefinitely stalled. By navigating the integration of sensing, processing, and active deterrence, we intend to deliver a functional system that proves specialized sensing can effectively protect life in the most overlooked area of vehicle operation.

2.2 Related Works and Existing Products

Current commercial and research efforts related to vehicle safety span a wide range of approaches, including underbody imaging technology, forward-facing animal detection methods, and standalone deterrent devices. While these solutions are typically developed for specific use cases, each addresses a portion of the broader safety challenge associated with limited close-range monitoring. The works discussed in this section were selected

because they represent the closest existing example to A.U.R.A., the development of similar sensing techniques for different operational scenarios, and the use of deterrent-based safety methods. Together, these examples illustrate how related technologies have evolved independently and highlight the opportunity to integrate these concepts into a unified and optimized system.

2.2.1 Retractable Underbody Camera System

Within the broader landscape of commercial products and research efforts related to vehicle safety and sensing, the system proposed by Kona and Dasari is the only identified work that directly targets underbody monitoring for vehicle safety in a manner like A.U.R.A. [5]. While many existing technologies focus on surrounding awareness, forward detection, or deterrent-based methods, very few are designed specifically to sense and evaluate the region beneath the vehicle. As a result, this work represents the closest existing example to the core purpose of the proposed system.

Their design extends camera coverage to the underside of a vehicle using a set of retractable wide-angle cameras arranged to provide near 360° visibility. When deployed, the cameras provide visual access to regions that are normally outside the field of view (FOV) of standard automotive sensors. During normal operation, the cameras retract into protective housing to reduce exposure to dust, water, and road debris.

A major focus of their work is the physical placement and protection of sensing hardware in close proximity to the road surface since underbody components are subject to harsh environmental conditions. The authors developed a mechanical deployment mechanism to preserve image quality and system reliability over time, and the system also includes an onboard display that presents real-time video to the user during low-speed operation.

This project demonstrates that underbody imaging is both feasible and practical when supported by appropriate mechanical design. It establishes an important foundation for close-range sensing beneath vehicles and highlights the engineering challenges associated with operating in this environment. However, the Kona and Dasari system primarily relies on multiple wide-angle cameras to achieve comprehensive visual coverage and is intended to support human decision-making through visual feedback. In contrast, our system does not pursue full 360° imaging. Instead, it emphasizes targeted NIR sensing at the front and rear of the vehicle combined with onboard processing for automated detection and response. By focusing on biologically relevant sensing rather than complete visual coverage, A.U.R.A. seeks to reduce system complexity and cost while maintaining effective underbody monitoring.

2.2.2 Front-facing Animal Detection

While few existing systems focus on sensing beneath vehicles, a larger body of research exists on detecting animals in front of moving vehicles. The system developed by Sibanda et al. uses infrared imaging and onboard processing to reduce wildlife-vehicle collisions

[6]. Their work was reviewed primarily for its treatment of biological target detection rather than for its physical deployment.

The authors emphasize how infrared contrast, motion characteristics, and environmental conditions affect the ability to distinguish animals from background surfaces. They compare multiple infrared sensing approaches, including passive thermal imaging, active infrared illumination, and hybrid methods that combine visible and infrared data. Their analysis explains the advantages and limitations of each approach in terms of detection range, sensitivity to ambient conditions, and susceptibility to false detections. They also examine scenarios involving low visibility, nighttime operation, and complex outdoor environments, and discuss how these factors influence detection reliability. This detailed comparison of sensing techniques was particularly relevant when considering how to identify living targets beneath a vehicle. Several of these concepts, including contrast enhancement, noise suppression, and feature-based detection, informed the development of A.U.R.A.'s sensing approach.

Although the system proposed by Sibanda et al. is designed for long-range, forward-facing detection at moderate to high vehicle speeds, many of its underlying sensing principles can be adapted to close-range monitoring. A.U.R.A. applies these ideas within a confined underbody region and integrates them with localized response mechanisms tailored to stationary and low-speed operation. By shifting the focus from distant warning to immediate close-range detection, the proposed system extends existing animal detection research to a largely unexplored operational scenario.

2.2.3 Comparison with Consumer Under-Hood Repellers

Existing consumer-grade solutions for protecting vehicles from small animals are primarily limited to under-hood rodent repellers [7]. These devices generally utilize ultrasonic sound waves or strobe lights to create a hostile environment that discourages animals from nesting or chewing on wiring. However, a significant drawback of these products is that they offer zero feedback to the user since they operate as a *passive* deterrent without any verification of success. A driver might assume the Repeller is functioning, but they have no way of knowing if an animal has ignored the stimulus or is already trapped within the engine bay.

Our system differentiates itself by providing an *active* safety solution. Unlike passive deterrents, our prototype integrates real-time visual verification to confirm an animal's departure before the vehicle is permitted to move. By combining the deterrent, a high-frequency alarm, with a closed-loop monitoring system, we ensure that the interlock that prevents vehicle movement is only released once the sensors confirm the area is entirely clear. This shifts the safety paradigm from simple deterrence to a verified safety check, providing a level of certainty that current consumer repellers cannot achieve. A comparison summary is provided in Table 2.2.1.

Table 2.2.1 Technology Comparison to System

Technology	Primary Coverage Area	Blind Spots	Key Capabilities	Limitations
Retractable Underbody Camera System (Kona & Dasari)	360° around vehicle underbody center region	Front of front tires; rear of rear tires; wheel occlusions	• Multi-camera wide-angle imaging • Real-time video feed • AI object detection	• Complex mechanical system • Designed for human monitoring
Front-Facing Animal Detection (Sibanda et al.)	Forward roadway ahead of vehicle	Under vehicle; sides; outside sensor range	• Active & passive infrared detection • Long-range wildlife detection • Motion & contrast-based classification	• No underbody monitoring • Limited to forward field of view • Reduced performance outside range
Consumer Under-Hood Repellers	Engine bay (when parked)	No detection capability; blocked by obstructions	• Ultrasonic or strobe deterrent • Continuous high-frequency output • Low cost & simple installation	• No real-time detection • No verification of animal departure • Habituation over time • Active only when parked
A.U.R.A. (Proposed System)	Entire underbody footprint (front-to-rear overlap)	Outside underbody perimeter	• Dual NIR camera modules (front & rear) • Active 850nm illumination • AI-based real-time animal detection • Prevent-start safety interlock • Species-based ultrasonic deterrent	• Focused on pre-drive scenario • Limited to underbody region • Requires calibration & alignment

2.3 Goals and Objectives

The following section outlines the basic, advanced, and stretch goals and objectives that have been established for our project. The basic goals and objectives detail the minimum requirements of the system. The advanced goals and objectives detail enhancements that build upon the minimum requirements to improve system functionality and performance. Stretch goals and objectives detail aspirational capabilities that go beyond the basic requirements of the system.

2.3.1 Goals

2.3.1.1 Basic Goals

- Use an acquired electrically powered miniature vehicle to showcase the functionality of system
- Design a system to detect animals underneath a vehicle upon initial startup of the vehicle only
- Design the system to be able to repel animals away from the undercarriage of the vehicle once detected
- Design the system to be able to capture footage of the animal for analysis

2.3.1.2 Advanced Goals

- Accurate identification of animal underneath the vehicle
- Protect sensitive components of the system from environmental hazards
- Enable system to communicate wirelessly between different components of the overall system
- Enable system to be able to analyze captures
- Prevent movement of vehicle if animal is detected within range

2.3.1.3 Stretch Goals

- Add an additional repellant by implementing a way to spray water at animals
- Integration into a fully autonomous vehicle
- Added functionality for fully sized vehicles

2.3.2 Objectives

2.3.2.1 Basic Objectives

- Develop of a single system that uses identical sensing modules at both the front and rear of a vehicle to detect animal life beneath the vehicle
- Create an alarm system to repel and deter animals from under the vehicle
- Attach an LCD display to view captures from the sensor system
- Implement wireless communication from the sensor cameras to the LCD display to capture feed from both cameras at the same time
- Implement system to retract the sensor system to shield it from debris when the vehicle is in motion or when system is not in use
- Implement vehicle motor interlock using an electrical signal to cut power from motors

2.3.2.2 Advanced Objectives

- Use a NIR camera to detect and identify animals when in range of the system

- Use a high-frequency ultrasonic noise and/or strobe light to act as the alarm to repel and deter animals from under the vehicle when animal is detected
- Use Raspberry Pi to communicate between sensor system, alarm system, and LCD display over Wi-Fi
- Use Raspberry Pi for sensor fusion from NIR cameras as well as for analysis from both sensor system captures
- Use MCU to control motor functions when retracting and extending the sensor system
- Use MCU to program LCD system to display captures from both cameras
- Use MCU to program the alarm system to sound at different frequencies depending on the species of the animal detected
- Implement real-time functionalities for all systems for automatic setting
- Use a hall effect sensor to calculate velocity and determine when vehicle reaches threshold for sensor retraction
- Use proximity sensors to detect possible collision of sensor system when being extended
- Use a relay to cut power from motors when system detects an animal underneath the vehicle and to reactivate power when system finds it clear

2.3.2.3 Stretch Objectives

- Enhance system with animal tracking functionality for accurate water spraying using controlled hardware

2.4 Required Specifications

Tables 2.4.1 and 2.4.2 present the specifications required to support the system's pre-drive animal detection and safety interlock functionality. Optical parameters such as lens squeeze ratio, field of view, and infrared spectral range are selected to ensure full undercarriage coverage and reliable detection in low-light environments.

Upon detection of an animal, a signal is sent from the visual sensors to the MCU, which activates an interlock relay installed in series with the battery supply to interrupt power delivery to the drive system and prevent vehicle movement. Once the A.U.R.A. high-frequency sensor confirms the hazard has been cleared, power is restored and the vehicle is permitted to move. When the vehicle reaches a predefined velocity threshold, in this case 1 mph, the mini electric actuators retract the camera modules into their protective housings

Table 2.4.1 Overall System Specifications

Specification	Description	Target Value and Unit(s)
Startup Safety Latency	The time from animal detection to the "prevent movement" interlock activation.	< 100 ms (Demonstrable)
Detection Success Rate	The accuracy of the AI in correctly identifying an animal versus benign debris.	> 95%

Minimum Detection Height	Minimum vertical height of target living animals detected by AI.	> 3.0 inches
System Engagement Threshold	The speed above which the camera retracts, and the interlock is safely disabled to prevent mid-drive stops.	1 mph (Demonstrable)
Total Deployment Time	Time required for cameras to fully extend from the safety housing upon ignition.	< 6 seconds (Demonstrable)
Ultrasonic Repellant Frequency	Frequency range of high-frequency alarm.	30 kHz – 45 kHz
Operational Temperature	The environmental range in which electronics and motors must function.	-10°C to +50°C
Hazard Clearance Verification Time	Time for the system to confirm clear status after deterrent use.	5 seconds
Total estimated cost	Total estimated cost for all integrated hardware components.	\$1516.43
Battery	Total peak power draws as a percentage of the ride-on car's battery capacity.	~50 W

The total deployment time requirement was revised from the original 3-second target to under 6 seconds. This change was driven by the actuator selection: the system uses a 6-inch stroke linear actuator to fully extend the cameras from the safety housing, and actuators with that stroke length that could also complete the motion in 3 seconds were not readily available on the market. Sourcing a unit fast enough to meet the original 3-second target while providing the required 6-inch stroke proved impractical, so the requirement was relaxed to under 6 seconds, which the selected actuator reliably meets, and which still satisfies the system's deployment-on-ignition needs.

The original startup safety latency target of less than 100 ms was based on an inaccurate early estimate. The requirement was revised to less than 2 seconds, measured from the moment an animal is detected to the activation of the motor interlock. This interval includes the remaining software decision processing, communication of the hazard command to the ESP32, and the hardware response required to engage the interlock. The revised value provides a realistic, measurable requirement for the completed system while still ensuring that vehicle movement is prevented promptly after a hazard is detected.

The final software does not use a fixed five-second clearance timer. Instead, each camera maintains a rolling window containing the detection results from its three most recently processed frames. A hazard is confirmed when either camera contains detections in at least two of the three frames. After a hazard, CLEAR is issued only when both camera windows fall below the two-out-of-three hazard threshold. This method prevents a single incorrect or missed frame from immediately changing the system state

Table 2.4.2 Component-Level Specifications

Component	Parameter	Subsystem and Description	Target Value and Unit(s)
Anamorphic Lens System	Field of View (Demo)	Optical - The front/rear lenses must cover the entire width and length of the undercarriage of a demonstration vehicle with 23" wheel track and 4.5"-7.5" clearance.	121° (H) x 45° (V)
	Field of View (Scalable)	Optical - The front/rear lenses must cover the entire width and length of the undercarriage of a full-sized AV with 64.7" wheel track and 6.5" clearance.	Up to 157.3° (H) x 53.1° (V)
	Squeeze Factor (Demo)	Optical – Required horizontal angular compression to match the IMX178 active-area aspect ratio (7.41×4.95 mm) for demo vehicle.	1.82x
	Squeeze Factor (Scalable)	Optical – Required horizontal angular compression to match sensor aspect ratio for full-sized AV.	1.98x
System Geometry	Scene Ratio (Demo)	Optical – Linear aspect ratio of the physical target window (Width/Height) for the demo vehicle.	4.35:1
	Scene Ratio (Scalable)	Optical - Linear aspect ratio of the physical target window for full-sized AV.	9.95:1
NIR Filter	Bandwidth	Optical - Selective spectral gating to block non-NIR solar and ambient noise.	10 nm FWHM @ 850nm
Objective Lens	Focal Length	Optical – To ensure that the VFOV fits the sensor's height to cover the 4.5"-7.5" underbody and is mechanically feasible.	3.0-3.2 mm
NIR LEDs	Spectral Range	Optical - 2000mW radiant flux LED array providing uniform NIR illumination for short range, high contrast underbody imaging	840 nm - 860 nm
Imaging System	Effective Resolution	Optical – Resolves biological targets beneath the demo vehicle. Scalable platform may require a higher resolution sensor	≥ 3 px/mm at ground plane
Retraction Motor	Torque	Electrical - The BLDC/Stepper motor must have enough power to clear dust/debris during extension.	At least 0.5 N-m
Main Processor	Inference Speed	Processing - The SBC (Single Board Computer) processing time for dual-stream YOLOv8 logic.	≤ 80 ms / frame

LCD Dashboard	Resolution	Processing – Support clear display of live undercar video feed and system alert indicators	720 (RGB) x 1280 pixels
	Latency	Processing – End-to-end camera to display delay for real-time visual confirmation	< 150 ms Lag
	Size	Electrical- Display size	7 inches
	Touch Panel	Processing – Capacitive multi-touch interface for system acknowledgment and manual override if required	Five-point capacitive touch
Safety Housing	Protection Rating	Mechanical - The retractable case must prevent ingress of water and road dust.	IP65 Rated
Interlock Relay	Switching Speed	Electrical - The time for the hardware relay to physically cut power to the drive motors.	< 10 ms

Physical manufacturing constraints, hardware procurement limits, and optomechanical space requirements necessitated specific updates to the nominal system specification matrix. The term squeeze ratio was updated to squeeze factor to align directly with standard industry terminology for angular anamorphic compression. To accommodate a straight-on longitudinal perspective within the 8 inch vertical safety zone, the nominal design point was re-derived for a 6.5 inch optical standoff. This geometry locks the required field profiles to 121 degrees by 45 degrees for the demonstration platform and up to 157.3 degrees by 53.1 degrees for the scalable full-size autonomous vehicle. This angular spread corresponds to physical scene ratios of 4.35:1 and 9.95:1, dictating updated optical squeeze factors of 1.82x and 1.98x respectively to match the active boundaries of the Sony IMX178 sensor.

Furthermore, individual hardware component selections were updated to resolve physical integration bottlenecks. The objective lens focal length was increased from the initial 3.0 to 3.2 mm target up to a pre-mounted 4.03 mm Thorlabs C340TMD-B aspheric lens. This change provides the back focal distance required to prevent a structural collision with the internal camera components while keeping the vertical field safely within sensor limits. To simplify the housing assembly and ensure full signal capture from the 840 to 860 nm LED array, the narrow-band filter was transitioned to a 12.5 mm round 850 nm component featuring a wider 30 nm FWHM bandwidth. This increased bandwidth effectively protects the system against edge vignetting caused by angular blue-shifting across the wide 121-degree field, successfully validating the physical hardware tracking architecture.

YOLOv8 was identified during the original design phase. During implementation, YOLOv8 was not offered through the selected Roboflow training workflow, so the final model was trained and deployed using YOLOv11 Nano.

2.5 Descriptions of Features and Functionalities

The proposed system addresses the specific pre-drive safety window where stationary vehicles pose the greatest risk to animals seeking shelter. To demonstrate these capabilities in a controlled environment, the system is integrated into a 12V small off-road electric kid car [8]. The following features are optimized for an autonomous platform to ensure safety without compromising vehicle performance at high speeds.

2.5.1 Hardware

2.5.1.1 *Velocity-Sensitive Retractable Imaging Housing*

Unlike standard perimeter cameras, our system is designed for the harsh environment beneath the vehicle. The system is programmed to deploy only when the vehicle is in park or stationary mode. Upon receiving an ignition or wake signal, the high-torque BLDC motor in the mini actuator arm extends the cameras from their IP65-rated protective cases. To prevent mechanical strain or any sudden braking hazards, the cameras automatically retract once the vehicle reaches a calibrated speed threshold (e.g. 1 mph). This ensures the hardware is protected from road debris and aerodynamic drag during normal autonomous transit. The physical placement and deployment mechanism of the sensors are illustrated in Figure 2.5.1.




Figure 2.5.1 3D Rendering of Sensor Placement on Demo Car



Figure 2.5.1.1 Sensor Placement on Demo Car

2.5.1.2 High-Frequency Alarm System

To repel animals from under the area beneath the vehicle, the alarm system will employ a high-frequency acoustic emitter operating between 30 kHz and 40 kHz, depending on which species was identified, and consistent with common parameters used in commercial animal-repellent devices.

The overall hardware architecture of the A.U.R.A. system, including sensing, processing, control, and communication components, is illustrated in Figure 2.5.2.




Figure 2.5.2 Hardware Diagram




Figure 2.5.2.1 Revised Hardware Diagram

The final architecture uses two Raspberry Pi 4B camera nodes and one central Raspberry Pi 5. Each camera connects locally to a Pi 4B through CSI, and the front and rear Pi 4B nodes send separate MJPEG streams to the Pi 5 over Wi-Fi. The Pi 5 performs detection, decision logic, and LCD control, while the ESP32 controls the interlock, strobe, ultrasonic deterrent, and camera actuators.

2.5.2 Software

2.5.2.1 "Prevent-Start" Safety Interlock

The primary goal is to prevent the vehicle from ever beginning its journey if a life is at risk. The autonomous system is hard coded to require a clear undercarriage signal before the drive motors can be energized. If an animal is detected, the system sends a lockout command to the Electronic Speed Controller (ESC). Upon detection, the high-frequency alarm is pulsed. The system monitors the camera feed in real-time; once the animal is seen exiting the detection zone, the lockout is automatically released, and the vehicle can proceed with its autonomous mission. To address the concerns about mid-drive stops, as that would not be realistic, the interlock only functions during the startup sequence.

2.5.2.2 Dashboard HMI and Manual Bypass

The system's primary interface is a 7" LCD Dashboard Display, which serves as the critical communication link between the autonomous safety logic and the remote human supervisor. In its standard operational state, the screen provides a simplified "Go/No-Go" status indicator, allowing for at-a-glance verification of the undercarriage safety check. To

ensure maximum safety during the pre-drive sequence, the Human-Machine Interface (HMI) is programmed with an automatic interrupt feature. If the AI-driven classification system detects a potential hazard, the display immediately bypasses the status light and switches to a high-resolution, split screened, live video feed provided by the dual inward-facing cameras. This specific configuration allows A.U.R.A. to perform an immediate visual inspection of the entire area between the front and rear axles to confirm the presence of an animal or obstacle.

The final software uses explicit commands between the Raspberry Pi 5 and ESP32 safety controller. The Pi 5 sends START, HAZARD, CLEAR, or RESET, and the ESP32 returns the corresponding acknowledgment. A command is accepted as completed only after the expected acknowledgment is received.

2.5.2.3 Animal Classification and Deterrent Selection

The A.U.R.A. system incorporates a software-based classification framework to identify animals detected beneath the vehicle using NIR imaging from IMX178M monochrome. Objects captured by the cameras are analyzed using reflected infrared intensity, shape, size, and motion characteristics extracted from the image data. The use of a monochrome sensor improves contrast in NIR conditions, enhancing feature extraction for classification. These features are compared against a predefined database of animal categories to estimate the type of animal present. Based on the identified category, the system selects an appropriate deterrent response, such as an audible alert operating within a specific frequency range or an alternative deterrent method. This modular software design allows the classification database and deterrent profiles to be updated or expanded as additional data and testing results become available.

The final implementation prioritizes animal-presence detection rather than species-dependent deterrent selection. Any qualifying animal detection is treated as a hazard and activates the same interlock, strobe, and ultrasonic deterrent response. This allowed the software effort to remain focused on the safety-critical decision of whether the undercarriage is occupied or clear.

Figure 2.5.3 is a flowchart illustrating the overall operational logic of the A.U.R.A system, including system initialization, safety checks, hazard detection, interlock control, and fault handling




Figure 2.5.3 Software System Flowchart




Figure 2.5.3.1 Revised Software System Flowchart

The revised flowchart reflects the implemented prototype. The user presses START, the Pi 5 sends an acknowledged START command to the ESP32, the camera actuators deploy, and the front and rear camera streams are activated. YOLOv11 Nano processes the frames, and the two-out-of-three frame rule determines whether the Pi 5 sends HAZARD or CLEAR. The revised flow also shows continued monitoring during a hazard and the Clear state in which detection pauses while the live camera feeds remain active.

2.5.3 Optical Setup

The optical sensing system is integrated into compact modules mounted near the front and rear edges of the vehicle and oriented toward the underbody region. This configuration was selected after mapping multiple sensor placement options and evaluating their resulting fields of view and blind spot distributions. Front and rear placement provides direct access to areas near the wheels and central undercarriage where small animals are most likely to remain concealed when the vehicle is stationary, while also enabling overlapping coverage beneath the chassis. The expected view from our sensing system of this configuration can be seen in Figure 2.5.4 Alternative configurations, including front and side placement, were analyzed but were found to produce larger asymmetric blind regions. This setup would also not be able to provide enough information to determine safe clearance if the vehicle were to reverse from park. The final layout was chosen to maximize coverage while minimizing obstruction from structural components. The optical design of each module is driven by the geometric constraints of this mounting location, including limited ground clearance, extreme aspect ratio requirements, and short working distances. The following sections describe how these constraints inform the system field of view, illumination strategy, optical compression, and image formation approach.





Figure 2.5.4 Expected underbody front and rear view of demo vehicle captured using Galaxy S22+ at 0.6x zoom with a stand-off distance of 8.5 in.

2.5.3.1 Target FOV

The initial FOV requirements were established by scaling the dimensions of a Waymo vehicle to the operational environment of an autonomous platform. With a vehicle width (W) of 64.7" and a ground clearance (H) of 6.5", the target detection zone at the leading edge of the vehicle forms a highly specific geometric window (Waymo, 2024). Using the trigonometric relationship $FOV = 2 \cdot \arctan\left(\frac{Dimension}{2 \cdot D}\right)$ and assuming a standard viewing distance (D) of 8.5 inches to the detection plane, the required horizontal FOV is calculated as $2 \cdot \arctan\left(\frac{64.7}{17}\right) \approx 150.6^\circ$, while the vertical FOV is 41.9° . These calculations helped to establish the extreme 3.59:1 aspect ratio of the target area.

While these full-scale dimensions define the primary geometric constraints of system, the physical demonstration will be conducted on the demonstration baby car with a wheel track of approximately 23 inches and underbody clearance ranging from 4.5 to 7.5 inches [9]. For this platform, the reduced geometry produces a less extreme aspect ratio, requiring a horizontal compression ratio of approximately 2.25:1. The optical system will therefore be configured through adjusted mounting distance and lens spacing to match the demonstration vehicle dimensions while preserving the same overall design approach.

The underlying optical architecture and compression strategy remain consistent across both scales, enabling geometric scalability through mechanical adjustment rather than fundamental redesign. Prior to final validation, the sensor modules will be temporarily mounted to a full-sized electric vehicle to collect underbody imagery and verify coverage geometry under realistic conditions. These trials will be used to confirm alignment with full-scale design targets without requiring permanent integration.

Following subsequent system characterization and mechanical integration on the physical platform, the target standoff distance (D) was optimized to a fixed horizontal constraint of 6.5 inches. Based on this 6.5 inch working clearance, the required angular coverage parameters for the demonstration vehicle evaluate to a horizontal field of view of 121 degrees and a vertical field of view of 45 degrees. This angular spread captures a linear

physical target scene ratio of 4.35:1 for the demonstration platform. To compress this asymmetric target window into the active area of the Sony IMX178 sensor (7.41 by 4.95 mm), the required horizontal optical squeeze factor updates to 1.82x. When scaling the exact same architecture to a full-sized vehicle at a 6.5 inch clearance, the required angular tracking envelope expands to 157.3 degrees horizontal by 53.1 degrees vertical, encompassing a 9.95:1 physical scene ratio, which is satisfied by increasing the optical compression factor to 1.98x via mechanical layout adjustment.

2.5.3.2 Active Illumination and Wavelength Selection

Illumination of the detection region is provided through active NIR LEDs arranged around the front aperture of the housing. Multiple emitters are distributed across this region to produce uniform illumination over the wide underbody area as a single point source would introduce strong brightness gradients and shadowing near the edges of the field of view, reducing detection reliability. The illumination wavelength is centered near 850 nm, where biological materials exhibit relatively high reflectivity and roadway surfaces remain strongly absorptive. Under these conditions, animals appear as high-intensity features against a dark background, producing consistent contrast independent of ambient lighting. By controlling the illumination source, the system maintains stable imaging performance across varying environmental conditions. This approach reflects prior work in infrared-based animal detection, which emphasizes the importance of consistent illumination for reliable biological target identification.

NIR imaging was selected over visible and long-wave infrared alternatives due to its compatibility with low-cost CMOS sensors and its effectiveness at short sensing distances beneath the vehicle. Operation near 850 nm allows the system to function independently of ambient lighting while remaining invisible to the human eye. Unlike thermal and long-wave infrared imaging, which rely on temperature differences that can be affected by heated pavement, tires, and vehicle components, the proposed system uses controlled reflectance to maintain consistent target visibility. Visible light imaging was also considered, but performance beneath the vehicle is highly dependent on shadows and external light sources. Given the short detection range and the need for reliable spatial detail, NIR imaging provides a practical balance between performance, cost, and system complexity.

2.5.3.3 Housing Interface and Optics Protection

Light reflected from the illuminated region enters the optical module through a recessed, anti-reflection-coated protective window mounted at the front of the enclosure. This window serves as the primary environmental barrier, shielding internal components from gravel, moisture, and debris. The housing is 3D printed so that internal components can be easily adjusted and redesigned as the system is refined. The LEDs are positioned concentrically around the window for symmetric coverage, and a light-blocking wall is built into the interior of the housing to separate the LEDs from the imaging optics. This prevents direct LED light from entering the lens and forces illumination to reflect off the ground before reaching the sensor so that, as a result, glare is reduced and image contrast is improved.

Physical implementation and system characterization required a significant modification to the illumination architecture and enclosure integration. While the protective entrance window remains a core structural specification for a deployed field unit, it was excluded from the active laboratory prototype to simplify direct optomechanical alignment of the primary cylindrical lens cap assembly. More critically, the concentric LED mounting scheme was abandoned during testing. Because the dual cameras utilize a direct, straight-on longitudinal perspective, a concentric front-facing illumination layout caused extreme pixel saturation and blooming artifacts where the near-infrared light reflected off the opposing hardware.

To resolve this illumination blow-out, an iterative lighting redesign was executed. The illumination array was first moved entirely to the vehicle side frame rails to shield the active light sources from the opposing camera's direct line of sight. However, this configuration created an underexposed center tracking region that reduced live AI detection confidence. To re-establish necessary feature contrast, the final layout strategically distributed the 850 nm LEDs along the front, rear, and side boundaries of the undercarriage, facing inward toward the center chassis. Although this introduces localized points of direct light within the composite frame, the elements are positioned to mitigate glare while delivering the uniform, high-contrast illumination profile required for stable downstream feature detection.

2.5.3.4 Anamorphic FOV Compression

The geometry of the underbody region produces a horizontal field of view that is nearly eight times larger than the vertical field. Conventional spherical optics capture this region inefficiently, allotting a large fraction of sensor pixels to unused areas. To address this limitation, the optical system incorporates an anamorphic telescope formed by a pair of cylindrical lenses. A plano-concave element introduces controlled horizontal divergence, while a plano-convex element recollimated the beam. Together, these components compress the horizontal field by a factor of approximately 3.59 while leaving the vertical dimension unchanged. This anamorphic system must map the large HFOV to the finite dimensions of the IMX178 sensor (7.41 mm wide).

Both elements are fabricated from N-BK7 glass to ensure thermal stability and consistent refractive behavior under elevated operating temperatures. The nominal separation between lenses follows $d = f_1 + f_2$ with final spacing and modeling of the full system to be refined through Zemax OpticStudio simulations to account for aberrations and alignment tolerances. This configuration maps the wide detection region efficiently onto the rectangular sensor format, ensuring that available pixel resolution is concentrated on relevant areas.

The architectural implementation and material selection for the anamorphic system deviated from the initial textbook model during physical integration. Sequential ray tracing within Ansys Zemax OpticStudio established that defining a perfectly afocal layout with nominal $(d = f_1 + f_2)$ spacing introduced severe packaging barriers. While the front cylindrical plano-concave element (Thorlabs LK1336L1-B) is fabricated from N-BK7

glass, the second cylindrical plano-convex element was updated to an available laboratory-stock fused silica component (Newport CSX200AR.16).

Furthermore, while the initial model targeted a 23 mm lens spacing, physical prototype testing revealed that the front 3D-printed enclosure sleeve introduced mechanical vignetting. To mitigate this clipping, the system spacing was adjusted to 39 mm, which represents the maximum physical separation achievable before the front lens cap restricts the active optical path. At this 39 mm spacing and a 6.0 inch straight-on working distance, the horizontal field of view evaluates to 108.3 degrees, mapping to 16.61 inches of physical ground track. Although this narrower coverage window leaves a 3.2 inch unmonitored gap at the outer boundaries of the 23-inch wheel track, it remains fully aligned with system safety goals. Target biological threats span dimensions larger than this 3.2 inch boundary, meaning their profiles will naturally extend into the active 16.61 inch viewing area and remain fully detectable by the downstream YOLO object detection architecture.

2.5.3.5 Image Formation and Sensor Integration

After anamorphic compression, the optical signal passes through a narrow-band interference filter centered near the illumination wavelength. This filter suppresses broadband visible and infrared radiation from ambient sources, allowing the sensor to respond primarily to reflected illumination.

A custom aperture plate limits off-axis rays and suppresses internal reflections within the housing. This aperture defines the effective entrance pupil and reduces sensitivity to stray light paths.

The filtered beam is then focused onto a near-infrared-sensitive CMOS image sensor using a compact spherical imaging lens. While an initial geometric estimate using the thin-lens relationship $f = \frac{h}{2 \cdot \tan\left(\frac{VFOV}{2}\right)}$ suggested a nominal value near 3 mm, detailed system analysis revealed that such a short focal length was physically and mechanically incompatible with the required anamorphic relay.

Captured image data is processed by an embedded computing platform. Because horizontal compression is performed optically, the embedded image processing software applies a corresponding de-squeeze operation prior to analysis to restore the original aspect ratio.

System resolution is selected to ensure reliable identification of small animals beneath the vehicle while maintaining sufficient coverage of the underbody region. Optical compression concentrates sensor pixels within the region of interest, improving spatial sampling efficiency without increasing sensor cost. For the demonstration platform, the configuration is designed to achieve effective spatial resolution sufficient to resolve biological targets and distinguish them from surface texture and debris. Image quality is influenced by illumination uniformity, lens aberrations, alignment accuracy, and signal-to-noise ratio, all of which will be evaluated experimentally during prototype testing and refined through simulation and calibration.

The active prototyping and physical assembly phases necessitated changes to the secondary filtering elements and the primary objective lens. The proposed custom internal aperture plate was eliminated from the physical layout to streamline the structural design of the housing. To manage solar glare and background noise, an off-the-shelf, circular 12.5 mm round 850 nm narrow bandpass filter was selected. This component features a 30 nm FWHM bandwidth, which provides sufficient threshold margin against angular blue-shifting at the steep boundaries of the wide field of view while successfully capturing the full emission profile of the 850 nm illumination LEDs.

Furthermore, the primary objective lens selection was altered during the final design iteration. While initial multi-lens relay modeling evaluated an unmounted 9 mm spherical objective to achieve target reduction, catalog optics in this focal range lacked an open clear aperture wide enough to prevent mechanical vignetting across the wide field of view. To bypass this clearance limitation, the system architecture pivoted to a pre-mounted Thorlabs C340TMD-B molded glass even aspheric lens. Featuring an effective focal length of 4.03 mm and a clear aperture of 5.1 mm, this objective provides the necessary mechanical clearance within the compact housing while securely containing the uncompressed vertical field within acceptable sensor boundaries.

2.5.4 Market Analysis and Comparable Products

This analysis focuses on the "Safety Gap" between current ADAS (Advanced Driver Assistance Systems) and specialized under-vehicle systems.

2.5.4.1 Comparison with Standard ADAS (Tesla/Waymo)

Modern AVs and ADAS, such as those deployed by Tesla and Waymo, rely on a sophisticated suite of LiDAR, radar, and perimeter cameras to navigate complex environments [10]. While these technologies are highly effective for long-range detection, they are primarily optimized for objects located 1 to 200 meters away from the vehicle's exterior. This focus creates a "zero-distance" blind spot, as most current sensors are physically incapable of monitoring the area directly between the wheels or beneath the chassis. Our project specifically targets the immediate under-chassis footprint that high-end systems like Tesla's Sentry Mode and Waymo's LiDAR arrays currently overlook, providing a critical layer of visibility where it is currently absent.

The operational logic of our system also represents a significant departure from existing commercial safety features. For example, while Tesla's system may record a visual log of nearby vandals or obstacles, it lacks the integrated logic to prevent the vehicle from moving if a hazard, such as an animal resting on the axle, is present. In contrast, our system introduces a proactive life-safety interlock that interfaces directly with the vehicle's drivetrain. By physically restricting movement until the undercarriage is verified as clear, our prototype provides a comprehensive safety measure that current commercial AVs lack, shifting the technology's focus from mere documentation to active accident prevention.

2.5.4.2 Industrial UVSS Adaptation

Industrial Under-Vehicle Surveillance Systems (UVSS) are the current benchmark for undercarriage inspection, typically used at high-security checkpoints to detect contraband or explosives [11]. These systems often use massive, fixed line-scan cameras or ground-embedded ramps. Our project represents a significant engineering shift by miniaturizing this capability into a vehicle-integrated, retractable format.

Rather than relying on a single fixed point on the ground, our adaptation uses a dual-camera inward-facing configuration mounted directly to the chassis. This allows the system to be safety-first rather than just security-first, focusing specifically on the detection of living creatures within the vehicle's footprint. By moving this technology from a static industrial asset to a mobile, integrated safety feature, we provide professional-grade monitoring that is optimized for residential autonomous platforms.

2.5.5 House of Quality

Figure 2.5.6 is the House of Quality which shows the different correlations and relationships between consumer and engineering requirements used to visualize the different tradeoffs between the different requirements. The top row shows 9 different engineering requirements needed to create a fully functioning final project along with whether they should be maximized or minimized for optimal performance. The left column shows simulated requirements from a possible consumer along with whether they should be maximized or minimized for consumer satisfaction. The roof of the house shows the relationships between engineering requirements, and the bottom row shows the target specification for each engineering requirement.

Correlations	
Positive	↑
Negative	↓
None	

Relationships	
Strong	↑↑
Positive	↑
Negative	↓

Direction of	
Maximize	+
Target	
Minimize	-

		Direction of Improvement + + - + + - - - - -									
Direction of Improvement	Engineering Requirements	Horizontal FOV	Vertical FOV	Startup time	Detection Accuracy	Camera Range	Wireless Latency	Alarm Response Time	Power Consumption	Cost	
	Customer Requirements (Explicit and Implicit)										
	- Cost	↓	↓	↓	↓	↓			↑↑	↑↑↑	
	+ Power Efficiency				↓		↓		↓		
	+ Reliability				↑↑		↑	↑↑			
	+ Ease of Installation			↑					↓	↓	
	- Size	↓	↓	↓		↓				↑	
	Target	96.6°	12.3°	<20s	95%	> 5ft.	<100 ms	<500 ms	~ 50 W	Estimated 750\$	

Figure 2.5.6 House of Quality




Figure 2.5.6.1 Revised House of Quality

Chapter 3 – Research and Investigation

3.1 Optical Hardware

3.1.1 Imaging System Selection

The first optical selection has to be the imaging system itself. The underbody region is not a normal camera scene. Its horizontal span is large relative to its height, and the imaging distance is short. That makes field capture only part of the problem. The more important question is how the scene is mapped onto the sensor once it is captured. A useful architecture for this system must preserve local detail near the outer field, use sensor area efficiently, and avoid introducing a mapping that becomes hard to interpret once detection moves away from the center of the image.

A spherical wide-angle lens is easy to source and easy to package. Typical M12 and CS wide-angle lenses in the 3 mm class reach roughly 120° field coverage, which is enough to make them a good candidate, but the problem is not just coverage [12]. A spherical wide-angle lens treats both axes together so if the horizontal field is increased to cover the width under the chassis, the vertical field grows with it. That means a large fraction of the sensor is spent on regions that do not materially help the detection problem and the effective sampling density across the ground plane drops at the same time.

Fisheye optics solve the coverage problem more aggressively, but they do so by accepting a projection that is much less uniform across the field. Commercial M12 fisheye lenses commonly advertise 192° to 200° coverage, and some variants extend past 210° [13]. That would certainly capture the underbody region. Even so, it would do so by compressing the outer field much more heavily than the center. In this project, that is a poor trade because a small animal is not guaranteed to sit neatly on axis. It may appear close to the wheel region, partly under the rocker, or near the edge of the covered scene. When a fisheye projection pushes that area into a smaller image radius, local detail suffers in exactly the places where the system most needs it.

A 360 camera system was also considered because it can cover the full environment with dual ultra-wide modules and stitching. That makes sense for panoramic visualization, however it does not make much sense for underbody inspection. Once the image is stitched, pixel geometry depends on the projection model of each lens and the algorithm used to merge the overlap. That means spatial interpretation is no longer uniform across the frame. In a safety system, complete coverage is less important than keeping the local geometry predictable.

Prism anamorphic systems were a more relevant comparison because they can reshape one axis without turning the whole imager into a fisheye camera. The problem is that a prism pair works by angular deviation rather than image formation. That makes the field remap depend more strongly on the incidence angle and refractive index. As field angle increases, uniform compression becomes harder to maintain. A cylindrical anamorphic telescope is

easier to tune because the remap comes from optical power in one axis rather than from field redirection alone.

The two cylindrical telescope families reduce to Galilean and inverse Galilean behavior. A Galilean arrangement stays compact because it avoids a real intermediate image, but in its conventional form it is more naturally associated with one-axis expansion. What the underbody scene actually needs is one-axis compression. This is why the inverse Galilean cylindrical system is the more natural fit. The negative cylindrical element first spreads the horizontal ray bundle in a controlled way, and the positive cylindrical element then recollimates that axis before final image formation. That gives the design a controllable horizontal squeeze while leaving the vertical axis largely untouched. A summary of the differences is shown in Table 3.1.1.

Table 3.1.1 Imaging System Comparison

Technology	Representative Coverage	Mapping Behavior	Main Strength	Main Limitation for A.U.R.A System
Spherical Wide-Angle Lens	About 90° to 120°	Perspective projection	Simple and compact	Horizontal and vertical field remain coupled
Fisheye Lens	About 180° to 215°	Strongly nonlinear projection	Very large field in one lens	Outer field is heavily compressed, distortion
Dual-lens 360 Camera	360° stitched panorama	Multi-projection plus stitching	Full environmental coverage	Stitch seams and nonuniform spatial mapping
Anamorphic Prism System	Tunable in one axis	Field redirection by angular deviation	Single-axis field shaping is possible	Wide-field mapping is harder to tune uniformly
Cylindrical Anamorphic Galilean System	Tunable expansion in one axis	First-order linear afocal remap	Compact and no real intermediate image	Better suited to axis expansion than compression
Cylindrical Anamorphic Inverse Galilean System	Tunable compression in one axis	First-order linear afocal remap	Compact and directly matched to the required field compression	Sensitive to spacing and alignment

3.1.2 Anamorphic System Component Selection

Once the system had been narrowed to a cylindrical inverse Galilean layout, the next question became the actual optical implementation. With the anamorphic system, the optical parts had to best implement that system while remaining realistic for a prototype that still must fit beneath a vehicle and remain compatible with the image sensor.

A plano-concave plus plano-convex cylindrical pair implements the system cleanly and keeps the role of each surface easy to interpret during optimization. The first element provides the negative cylindrical power needed to open the horizontal bundle, and the second provides the positive cylindrical power needed to recollimate that axis before it reaches the objective. This is much easier to tune than a more aggressive bi-surface pair, and it is much more realistic for sourcing than a custom freeform element.

This choice also helps from a sensor-usage standpoint. A cylindrical telescope does not exist here for cinematic effect or for a visually pleasing image. It exists because the scene under the chassis is much wider than it is tall. Without that one-axis compression, too many pixels would be spent on the wrong parts of the frame. With the cylindrical pair in place, more of the sensor is devoted to the ground plane and the target region rather than to empty space. That directly improves effective sampling density where it matters most.

The same logic carries forward into the objective lens choice. Once the anamorphic front end is doing the horizontal field shaping, the objective lens no longer has to solve the entire field problem by itself. Its job shifts toward image formation after remapping, which is why a compact M12-format wide-angle lens remains attractive while the Zemax work continues to refine exactly what objective power the final build should use. A summary of the differences is shown in Table 3.1.2.

Table 3.1.2 Anamorphic System Component Comparison

Component	Optical Role	Main Strength	Main Limitation for A.U.R.A System
Plano-concave + Plano-convex Cylindrical Pair	Negative divergence followed by recollimation in one axis	Simple surface functions and good Zemax control	Sensitive to spacing
Bi-concave + Bi-convex Cylindrical Pair	Same telescope behavior with more curved surfaces	Higher one-axis power possible	More difficult packaging and alignment
Freeform/Molded Element	Could combine multiple functions	Potentially compact final package	Wide-field mapping is harder to tune uniformly

3.1.3 Imaging Spectral Band Selection

The right spectral band for this project is tied directly to the actual scene beneath the car. The target area is only a few inches away, and the system is trying to preserve useful local detail. It does not need to see down a dark highway or identify objects tens of meters away.

Visible imaging offers the strongest natural color information, and silicon cameras already perform well in that band. Even so, it is a weak match for this application since the region beneath a parked vehicle is heavily shadowed and a visible active source would be

noticeable to nearby pedestrians, especially at night. It could also produce glare from wet pavement, polished concrete, or painted surfaces. For a safety system that is meant to operate quietly in the background, that is not a small issue. It could be a distraction hazard on the road to other vehicles or people around our car.

SWIR and LWIR are both technically plausible alternatives. SWIR can produce useful material contrast, but it immediately pushes the detector into the InGaAs class, which raises cost far beyond the rest of the current optical stack. LWIR avoids active illumination, but it changes the meaning of the image. A thermal camera sees heat contrast rather than controlled reflectance. Under a vehicle, there could be warm pavement, tires, brake hardware, or nearby drivetrain components that all produce strong returns even when there is no animal present.

Near infrared is optimal because the working distance is short and the system prioritizes surface reflectance over thermal emission. Fur does not reflect the same way as rubber. Exposed skin does not respond the same way as painted debris or wet concrete. That does not mean the system will be able to easily perform species identification, but preserving those reflectance differences is useful when trying to separate animal-like targets from the background. NIR lets the system do that while staying inside the cost and integration envelope of silicon imaging. A summary of the differences is shown in Table 3.1.3.

Table 3.1.3 Imaging Spectral Band Comparison

Spectral Band	Approximate Range	Detector Family	Main Strength	Main Limitation for A.U.R.A System
Visible	400 to 700 nm	Silicon CMOS or CCD	Strong color information and high silicon response	Strong dependence on ambient light and visible glare
NIR	700 to 1000 nm	Silicon CMOS or photodiode	Works with low-cost silicon sensors and controlled illumination	Requires active source and filtering
SWIR	900 to 1700 nm	InGaAs	Useful reflectance contrast in some scenes	Detector cost rises sharply
LWIR (Thermal)	8 to 14 μm	Microbolometer	Passive temperature-based imaging	Lower spatial detail at comparable cost

3.1.4 Active Illumination Source Selection

Once NIR had been selected as the spectral band, the next issue was how to generate it. Since A.U.R.A. is a short-range inspection system, it does not need a narrow long-range beam. Instead, it needs a fairly broad and uniform wash across the underbody region.

An LED array makes the most sense in this case because the scene is broad and nearby. The goal is to illuminate the ground plane under the car with enough uniform flux that local reflectance differences survive the round trip through the optics and onto the sensor. A laser diode starts from the opposite place. It gives a concentrated beam that then has to be spread back out. That adds optics, increases alignment sensitivity, and makes the safety case harder. A VCSEL flood source is closer to the right idea because it can illuminate an area rather than a point, but it still adds more source complexity than the current build needs.

For this application, infrared LEDs fit the scene geometry more directly because they illuminate a broad, shallow region without needing additional beam-expanding optics or a more difficult safety case. A distributed IR LED layout can be placed around the optical opening so the return signal is less dependent on one bright hot spot. That is much better suited to the underbody scene than a single concentrated source. ams OSRAM explicitly positions its high-power IR emitters for camera illumination and notes that 850 nm, 920 nm, and 940 nm families are designed to match the sensitivity range of standard photodiodes and extended-IR CCD and CMOS cameras.

This source choice also supports the kind of image information the system is actually trying to preserve. The module is not only looking for occupancy. It is trying to preserve enough return contrast that differences in surface appearance remain visible after imaging. Fur, exposed skin, rubber, wet pavement, painted concrete, and plastic debris do not return the same signal under controlled NIR illumination. That matters because a small animal under the chassis may not present a strong outline. In some cases, the most useful cue may be a local difference in texture or reflectance rather than a clean silhouette. A broad IR LED illumination field is better suited to that task than a narrow laser-style source because it produces a more uniform reflectance image across the region of interest. A summary of the differences is shown in Table 3.1.4.

Table 3.1.4 Active Illumination Source Comparison

Illumination Source	Typical Spectral Width (FWHM)	Beam Characteristic	Main Strength	Main Limitation for A.U.R.A System
IR LED	About 30 nm for 850 nm camera illuminators	Naturally wide	Good fit for distributed short-range illumination high silicon response	Lower radiance than laser sources

Edge-emitting Laser Diode	Typically less than 5 nm	Highly directional	Very high radiance	Beam shaping burden and more difficult safety case
VCSEL flood source	Typically about 1.5 to 3 nm	Structured or controlled flood	High optical intensity in compact form	Higher complexity and cost

3.1.5 NIR Emitter Wavelength Selection

Now that IR LEDs has been selected for the illumination source, the part-level comparison moves to wavelength within the NIR band and to practical emitter families. Here the central trade is detector signal versus human visibility.

Silicon response falls gradually after visible red. It does not vanish immediately. Representative silicon detectors still show useful responsivity at 850 nm. By 940 nm, that signal is significantly lower. That trade matters here because the system is not using IR only to prove that something is present. It is trying to keep enough image contrast that texture and reflectance differences are still visible at the sensor. If the source is pushed to 940 nm, the system gives away signal in exchange for a lower visible glow that is already small and can be further controlled by source placement and baffling.

The selected part family is the OSOLON Black 850 nm line. Official ams OSRAM listings for 850 nm OSOLON Black variants report peak wavelength around 860 nm, beam angle of 80°, radiant flux up to 1530 mW at 1000 mA for the SFH 4715AS, and radiant intensity up to 900 mW/sr. The series is also rated for operating temperatures from -40 °C to 125 °C, which matters for an emitter mounted near the underbody of a vehicle. That gives the design a short-range source with enough optical output to matter, wide enough emission to illuminate the scene, and enough thermal margin to remain credible in an automotive-adjacent environment. A summary of the differences is shown in Table 3.1.5.

Table 3.1.5 NIR Emitter Wavelength Comparison

Emitter	Peak Wavelength	Representative Silicon Response	Human Visibility	Representative Part Data
780 nm class IR LED	780 nm	around 0.45 to 0.55 A/W	Noticeable deep red glow	Higher detector response, but too visible
850 nm class IR LED	850 to 860 nm	around 0.36 to 0.60 A/W depending on detector	Faint glow at most	Strong balance of signal and low visibility
920 nm class IR LED	920 nm	below 850 nm response	Less visible than 850 nm	Middle ground
940 nm class IR LED	940 to 950 nm	around 0.16 to 0.45 A/W depending on detector	Effectively invisible	Lowest silicon return of the main options

While the high power OSLO Black series remained the baseline target for automotive deployment, the physical laboratory prototype transitioned to a distributed array of Chanzon 3mm through hole 850 nm infrared emitter diodes. Surface mount high power emitters required specialized printed circuit board thermal dissipation layouts that proved incompatible with the internal spatial limits of the compact 3D printed camera housing. To deliver illumination without thermal accumulation issues, the prototype utilizes these 3mm through hole elements which operate at a low forward voltage of 1.45 to 1.65 V and a forward current of 20 mA. Sourcing these components also narrowed the native emitting angle from 80 degrees down to 30 degrees. This narrower profile matches the inward facing lighting layout perfectly, concentrating the infrared beam directly onto targets under the chassis while avoiding the severe camera saturation and blooming artifacts caused by wide angle surface mount illumination sources.

3.1.6 Image Sensor Technology Selection

The sensor decision is divided into two sections. The first is the imaging technology itself and the second is the packaging platform used to integrate that technology into the optical subsystem. Those two pieces belong together because this project does not just need a high-resolution detector, it needs a detector that can live inside a custom optical path without turning the project into a sensor-packaging exercise.

A bare sensor gives the most design freedom. In practice, it also pulls the work into fine mechanical alignment, contamination control, custom electrical integration, and board-level development that do not solve the main design problem. The value of this subsystem lies in the field remapping and short-range NIR imaging strategy. A bare sensor would add development burden without improving that strategy.

A fixed-lens module sits at the other extreme. It is easy to deploy, but it locks the optical path too early. That is a poor fit here because the imaging lens has to coexist with a custom anamorphic front end. Global-shutter CMOS remains attractive in fast machine-vision tasks, but this underbody system is intended to inspect the scene before the vehicle moves. That weakens the case for paying a large premium to eliminate rolling-shutter artifacts.

That leaves a removable-lens rolling-shutter CMOS module as the strongest technology choice. It preserves enough flexibility to integrate a custom front end while avoiding the unnecessary burden of a bare-sensor design. A summary of the differences is shown in Table 3.1.6.

Table 3.1.6 Image Sensor Technology Comparison

Technology	Readout Type	Resolution Class	Lens Flexibility	Main Limitation for A.U.R.A System
Bare sensor	CMOS or CCD bare die	Varies	Highest in theory	Packaging, contamination, and interfacing burden

CCD Camera	CCD	Moderate to high	Varies	Higher power and weaker embedded ecosystem
Global-shutter CMOS module	CMOS global shutter	Often lower megapixel count at low cost	Usually modular	Higher cost for less resolution
Fixed-lens rolling-shutter module	CMOS rolling shutter	Moderate to high	Low	Lens path is locked
Rolling-shutter CMOS removable-lens module	CMOS rolling shutter	High	High	Still constrained by module form factor

3.1.7 Image Sensor Part Selection

While rolling-shutter CMOS modules remain the strongest technology choice, the optical format (sensor size) became a critical design driver during the refinement of the anamorphic relay. A larger sensor format is required to capture the expanded horizontal image created by the cylindrical optics without requiring extreme and physically impossible magnification from the objective lens.

The larger IMX477 used in the Raspberry Pi HQ Camera is a strong sensor in general. Raspberry Pi lists it at 12.3 megapixels with a 7.9 mm diagonal and 1.55 μm pixels. That larger format helps light collection, but it also drives the optical system toward a larger corrected image circle. In a compact underbody package that is a real cost.

Camera Module 3 NoIR is newer and less expensive than the HQ Camera, but its main identity is still a 12 MP autofocus module. Autofocus is not the main feature that matters here. Mount flexibility matters more. The Seeed IMX219 M12/CS module offers that flexibility while keeping the sensor small enough that the rest of the optics remain realistic to package.

Seeed lists its module at \$15.90 and describes it as an IMX219-based camera board that supports both M12 and CS-mount lenses. That matters more here than raw megapixel count because the board behaves like a sensor platform rather than a finished consumer camera. It leaves enough room for the optical system to define the front end.

A low-cost global-shutter alternative such as OV9281 remains useful as a comparison point, but its 1 MP class resolution is much less attractive for preserving local texture detail across a wide field.

The Sony IMX178 Monochrome was selected as the finalized sensor for the A.U.R.A. system, representing a necessary evolution beyond the initial Seeed IMX219 baseline. While the IMX219 was attractive for its small footprint, its $\frac{1}{4}$ " format (3.67 mm width) was found to be mathematically incompatible with the required anamorphic mapping. To project a 55° horizontal field onto such a small area, the objective lens would require a

focal length of approximately 3 mm, which lacks the Back Focal Length (BFL) necessary to accommodate the 850nm bandpass filter and sensor glass. By pivoting to the format of the IMX178, the system utilizes a 7.41 mm width that matches the clear aperture of the 9 mm catalog objective, providing 5.45 mm of mechanical clearance. Furthermore, the shift to a monochrome substrate eliminates the light-attenuating Bayer Filter Mosaic, increasing the Quantum Efficiency (QE) to nearly 40%. This, combined with a larger 2.4 μm pixel pitch, provides the superior signal-to-noise ratio required to resolve low-contrast animal features in dark underbody environments. A summary of the differences is shown in Table 3.1.7.

Table 3.1.7 Image Sensor Part Comparison

Sensor or Module	Optical Format	Pixel Size	Resolution	Lens Situation	Approximate Current Price
Seed IMX219 M12/CS module	1/4 inch	1.12 μm	3280×2464	Interchangeable M12 or CS	\$15.90
Raspberry Pi HQ Camera with IMX477	1/2.3 inch	1.55 μm	4056×3040	Interchangeable	about \$50 to \$55
Raspberry Pi Camera Module 3 NoIR	7.4 mm diagonal sensor class	1.4 μm	11.9 to 12 MP	Autofocus compact camera module	\$25
Arducam OV9281 global-shutter module	1/4 inch class	Larger pixel class	1280×800	Interchangeable M12	about \$42 to \$55 depending
Sony IMX178 (Mono) Module	1/1.8 inch	2.4 μm	3088×2064	Interchangeable CS or C	\$170

3.1.8 Optical Material Selection

The cylindrical telescope and the protective window both need materials that are predictable in the visible-to-NIR band and stable enough for an underbody environment.

N-BK7 gives the design stable refractive behavior in the exact band the system uses without adding the cost of a more specialized substrate. SCHOTT lists N-BK7 at $n_d = 1.51680$ and $V_d = 64.17$, with thermal expansion of $7.1 \times 10^{-6}/\text{K}$ over 243 K to 343 K which is well suited to an 850 nm imaging system that may see underbody temperature rise but does not need UV transmission or high-energy-laser durability.

Fused silica performs extremely well and offers lower thermal expansion, but most of that is not needed by the current build.

Acrylic is a good choice in terms of cost and manufacturability, but its thermal expansion is roughly an order of magnitude larger than BK7. In a location near the underside of an electric vehicle, where temperatures can rise appreciably near battery shielding or drive components, that is not a small difference. The material may be fine for enclosure features, but it is a weak choice for the precision field-shaping optics. A summary of the differences is shown in Table 3.1.8.

Table 3.1.8 Optical Material Comparison

Material	Refractive index nd	Abbe number Vd	Useful Transmission Range	Thermal Expansion	Suitability for A.U.R.A System
N-BK7	1.51680	64.17	About 350 nm to 2.0 μm	$7.1 \times 10^{-6}/\text{K}$ from 243 K to 343 K	Good visible-to-NIR performance with manageable thermal drift
UV Fused Silica	1.4584	About 67.8	about 180 nm to 2.1 to 3.5 μm	much lower than BK7	Better UV and thermal performance than needed here
PMMA or Acrylic	about 1.49	About 57	visible to limited NIR	roughly 70 to $90 \times 10^{-6}/\text{K}$	Too thermally unstable for precision field shaping

3.1.9 Spectral Filter Selection

The system uses active 850 nm illumination and needs the sensor to respond primarily to reflected source light rather than to uncontrolled ambient radiation and so a spectral filter is also needed. Once the emitter wavelength has been selected, the filter choice becomes a source-selectivity problem rather than a general imaging problem.

With a narrow bandpass filter, the illumination wavelength is already defined. There is little advantage in admitting a large amount of unrelated visible or infrared radiation if the image is intended to be dominated by the system's own 850 nm source. Commercial hard-coated NIR bandpass filters in this range are widely available with center wavelengths around 850 nm and bandwidths from about 10 nm upward. Thorlabs lists hard-coated NIR bandpass filters in the 850 to 890 nm range, and comparable commercial 850 nm filters are sold in 10 nm class bandwidths for higher source selectivity.

At the part level, an 850 nm narrow bandpass element remains the most coherent choice because it matches both the selected emitter wavelength and the sensor-response trade that justified 850 nm over 940 nm. A longpass filter would help with visible rejection, but it would still pass a much broader infrared band than the system actually needs. A summary of the differences is shown in Table 3.1.9.

Table 3.1.9 Spectral Filter Comparison

Filter	Typical Pass Behavior	Main Strength	Main Limitation for A.U.R.A System
None	Full-spectrum response	Maximum throughput	Strong ambient contamination
Longpass IR	Broad pass above cutoff	Rejects most visible light	Still admits wide-band ambient IR
Narrow Bandpass near 850 nm	Typical FWHM about 10 to 40 nm	Best source selectivity and ambient suppression	Lower total throughput and higher part cost

The active prototyping phase required specific adjustments to the filter bandwidth parameters and mechanical profile. While the initial design targeted higher source selectivity via a strict 10 nm FWHM specification, physical evaluation revealed that narrow thin-film coatings introduced severe peripheral signal degradation due to angular blue-shifting across the wide 121-degree horizontal field of view. Additionally, a 10 nm window risked truncating the natural emission bandwidth of the physical LED array. To prevent these edge-vignetting and clipping issues, the architecture transitioned to a 12.5 mm round 850 nm narrow bandpass filter featuring a wider 30 nm FWHM bandwidth. Sourcing this 12.5 mm circular format allows the filter to function as a direct cap over the entire front face of the Thorlabs C340TMD-B objective housing, eliminating the complex mounting hurdles associated with square filter plates while completely blocking stray ambient light leaks from bypassing the imaging path.

3.1.10 Protective Window Selection

The protective window is a separate decision because its job is environmental protection first and optical preservation second. The front window is not just a cover, it is also part of the optical path. Since the module is mounted beneath a vehicle, the front window must tolerate dirt, splash, and temperature variation without introducing unnecessary reflection loss or stray light. If the entrance surface adds unnecessary reflection, the system loses signal and increases the chance of stray internal reflections reaching the sensor. A glass substrate is more stable than a polymer cover when the module is exposed to temperature variation under the vehicle. Commercial AR-coated optical windows are widely available in N-BK7 and fused silica, and they illustrate the usual design trade of coated precision optics costing more, but they protect throughput and reduce stray light at the same time. At the part level, an AR-coated glass window is the stronger choice for the imaging aperture, while lower-cost polymer materials remain more appropriate for non-imaging housing features. A summary of the differences is shown in Table 3.1.10.

Table 3.1.10 Protective Window Comparison

Window	Material Class	Main Strength	Main Limitation for A.U.R.A System
Uncoated Glass Window	Optical glass	Low cost and good mechanical stability	Higher Fresnel reflection at both surfaces
AR-coated Glass Window	Optical glass with anti-reflection coating	Better throughput and lower stray reflection	Higher part cost
Polymer cover window	PMMA or similar polymer	Lightweight and easy to machine	Greater thermal drift and surface damage risk

3.1.11 Housing Material Selection

The housing controls alignment repeatability, protects the optics, and determines how well the module tolerates the underbody environment. Acrylonitrile styrene acrylate (ASA) preserves the practical ease of FDM fabrication while offering better thermal and environmental durability than PLA. Technical data sheets for commercial ASA filaments list glass transition near 98 to 103 °C and heat resistance around 80 to 100 °C depending on processing condition and grade. Those numbers do not turn the housing into a high-temperature automotive polymer, but they are much more credible for an underbody prototype than PLA would be. They also line up with the project's need to keep the optical alignment stable even when the module is exposed to the sun, road heat, or elevated temperatures near the underside of an electric vehicle. A summary of the differences is shown in Table 3.1.11.

Table 3.1.11 Housing Material Comparison

Material	Glass Transition or Heat Resistance	Outdoor Durability	Suitability for A.U.R.A System
ASA	glass transition around 98 to 103 °C and heat resistance up to about 100 °C depending on grade	Strong	Better fit for outdoor and automotive-adjacent exposure

PLA	low, typically around 55 to 65 °C	Weak	Too vulnerable to heat for underbody use
ABS	Better than PLA	Moderate	Usable, but weaker UV performance

3.2 Hardware

3.2.1 Actuator

3.2.1.1 Actuator Technology Comparison

The actuator is a critical mechanical element of A.U.R.A., responsible for extending the sensor modules downward and outward from their protective housings upon ignition, and retracting them once vehicle speed exceeds the 1 mph threshold. In Table 3.2.1.1, three broad actuation technologies were evaluated: linear electric actuators, servo motors coupled with a lead screw, and pneumatic cylinders.

Linear electric actuators convert rotary motor motion directly into linear displacement through an integrated lead or ball screw mechanism enclosed in a sealed body. They operate on a DC supply, typically 6 V to 12 V, and are inherently reversible, allowing both extension and retraction from a single device without auxiliary valves or linkages. Their compact, self-contained form factor is well matched to the limited clearance available beneath the vehicle chassis, and many variants carry IP54 or higher protection ratings, making them tolerant of road splash and light debris. The primary limitation is that peak force is moderate (typically up to 500 N for micro-class units), but this is more than sufficient for the lightweight imaging module that must be moved.

Servo motors paired with an external lead screw offer a highly customizable stroke and torque profile. By selecting different lead screw pitches, the designer can independently tune speed and force. However, this flexibility comes at the cost of a significantly higher component count: a separate motor, coupler, lead screw, nut, and bearing block must all be sourced, aligned, and assembled. The open screw mechanism is also difficult to protect from road grit and moisture without a custom shroud, adding further complexity.

Pneumatic cylinders can deliver extremely high force over a wide stroke range, but they require a compressor, solenoid valves, and tubing, none of which are feasible in a compact, battery-powered vehicle integration. The power overhead and physical volume required by a compressed-air supply immediately disqualify pneumatics for this application.

Table 3.2.1.1: Actuator Technology Comparison

Technology	Max Stroke / Range	Force / Torque	Key Trade-offs for A.U.R.A.
Linear Electric Actuator	12–300 mm stroke	10–500 N	Compact, precise, reversible; no pneumatic supply needed; low noise

Servo Motor + Lead Screw	Custom stroke	Variable	Flexible but requires more assembly, higher part count, harder to waterproof
Pneumatic Cylinder	25–500 mm stroke	Very high (100–2000 N)	Requires compressor/solenoid; bulky; not viable for compact vehicle integration

Based on this analysis, the linear electric actuator technology was selected. It provides the best balance of compactness, ease of control, environmental tolerance, and integration simplicity for the A.U.R.A. deployment mechanism.

3.2.1.2 Actuator Part Comparison

Within the linear electric actuator category, three specific products were evaluated: the Firgelli FA-150-S-12-6, the Actuonix L12-50-210-12, and a DFRobot FIT0186 servo motor with a custom 3D-printed linkage.

The Firgelli FA-150-S-12-6 is a compact linear actuator with a 6-inch stroke, a stall force of 150 lb, and a 12 V DC supply requirement. A 6-inch stroke was identified as the target after the team's early prototyping experience with an 8-inch unit, detailed below, confirmed that total installed length is the binding dimensional constraint. The Firgelli unit's self-locking mechanism holds the camera module in its deployed position without drawing continuous current, and its IP54 protection rating makes it tolerant of road splash. At approximately \$45 per unit, two are required for the front and rear sensor modules.

It is important to note that during early prototyping, the team initially acquired an 8-inch stroke linear actuator based on a preliminary assessment of the required deployment distance. However, upon physically mounting the unit beneath the demonstration vehicle, it became immediately apparent that the total installed length, accounting for both the retracted body and the fully extended rod, exceeded 16 inches. This dimension was incompatible with the available underbody clearance of the demo vehicle, which ranges from 4.5 to 7.5 inches, and would have caused the actuator to contact the ground surface before the camera module reached its intended position. This hands-on test highlighted the importance of accounting for the full collapsed length of a linear actuator, not just its stroke, when evaluating fit within a constrained mechanical envelope. The team subsequently re-evaluated stroke requirements based on the actual housing geometry and identified the 6-inch stroke variant as correctly sized for the application.

The Actuonix L12-50-210-12 has a 2-inch stroke, which falls short of the deployment distance needed to position the camera modules at the correct standoff distance from the ground plane. Its maximum force of approximately 67 lb is also lower than the Firgelli unit, and at approximately \$70 it is the most expensive option evaluated.

The DFRobot FIT0186 is a standard RC-style servo motor. Unlike a linear actuator, a servo motor produces rotary motion, meaning it spins a shaft rather than extending a rod outward. To use it for camera deployment, the team would need to design and 3D-print a custom linkage bracket that physically connects the rotating servo horn to the camera housing and

converts that rotation into the sliding in-and-out motion the system requires. Think of it like a door hinge arm: the servo rotates, the arm pushes the camera outward, and when the servo reverses, the arm pulls the camera back in. The problem with this approach is that any 3D-printed linkage introduces mechanical play, meaning small gaps and flex at each joint. This play causes the camera to settle in a slightly different position each time it is deployed, which matters for A.U.R.A. because the optical alignment of the NIR camera relative to the anamorphic lens system must be repeatable. At approximately \$12 it is the cheapest option but the custom design burden and positional inconsistency make it unsuitable. A summary of the differences is shown in Table 3.2.1.2.

Table 3.2.1.2: Actuator Part Comparison

Part	Stroke	Force	Price	Notes
Firgelli FA-150-S-12-6	6 inches	150 lb	~\$45	Selected: 12 V, IP54-rated, self-locking, fits undercarriage envelope of demo vehicle
Actuonix L12-50-210-12	2 inches	67 lb	~\$70	Too short a stroke for the required deployment distance; higher cost with less force
DFRobot FIT0186 Servo + Custom Linkage	N/A (rotary)	1.2 kg-cm	~\$12	Rotary servo requiring a custom 3D-printed bracket to convert rotation to linear motion; unreliable positioning

The Firgelli FA-150-S-12-6 was selected for A.U.R.A. Its 6-inch stroke fits within the physical constraints of the demo vehicle undercarriage, its 150 lb force provides ample margin for deploying the camera housing against gravity and spring resistance, and its self-locking and IP54 features are well matched to the operating environment.

3.2.2 Motor Driver

3.2.2.1 Motor Driver Technology Comparison

The motor driver governs direction and speed of the linear actuator by mediating between the 3.3 V logic output of the MCU and the 12 V, high-current load of the actuator motor. In Table 3.2.2.1, three motor driver topologies were considered: integrated H-bridge ICs, discrete MOSFET H-bridges, and dedicated DC motor driver modules.

Integrated H-bridge ICs such as the Texas Instruments DRV8833 or the STMicroelectronics L293D consolidate four switches, a charge pump, and protection circuitry into a single package. They are extremely compact and interface directly to microcontroller GPIO or PWM pins. However, their continuous current ratings typically do not exceed 2 A to 3 A per channel, which may be insufficient during actuator stall conditions where current can briefly spike to several times the rated running current. Thermal shutdown may occur during extended duty cycles if heat sinking is not carefully managed.

Discrete MOSFET H-bridges, built from individual N-channel and P-channel devices, can handle currents well above 10 A and allow the designer to optimise switching frequency and dead-time insertion for minimal losses. However, this approach requires careful PCB layout to manage switching transients, dedicated gate driver ICs, and significantly more schematic complexity. For a prototype application with time constraints, this level of custom power electronics design is difficult to justify.

Dedicated DC motor driver modules integrate power MOSFETs, a gate driver, current sensing, and protection features on a ready-made PCB. They are available at low cost, require only a PWM signal and a direction pin from the MCU, and can typically source 10 A or more continuously. They represent the best trade-off between performance, ease of integration, and development time.

Table 3.2.2.1: Motor Driver Technology Comparison

Technology	Typical Current	Control Method	Trade-offs for A.U.R.A.
H-Bridge IC (integrated)	Up to 3 A per channel	PWM / logic pins	Low component count; easy I2C or GPIO interface; limited high-current capability
Discrete MOSFET H-Bridge	10–50 A	Gate driver + PWM	High current; requires PCB design effort; complex layout; EMI concerns
Dedicated DC Motor Driver Module	Up to 13 A	PWM + DIR pins	Ready-made; robust protection; easy integration; slightly larger footprint

The dedicated DC motor driver module topology was selected because it provides adequate current headroom for actuator stall conditions, requires minimal MCU interfacing, and is available as a tested, reliable assembly that reduces prototype risk.

3.2.2.2 Motor Driver Part Comparison

In Table 3.2.2.2, three specific motor driver parts were compared: the Cytron MD10C, the DRV8833 dual H-bridge IC, and the L298N dual H-bridge module.

The Cytron MD10C is a dedicated single-channel DC motor driver capable of sourcing 13 A continuously (30 A peak) at voltages from 5 V to 30 V. It accepts a simple PWM and direction (DIR) control interface, and includes built-in under-voltage lockout, over-temperature shutdown, and short-circuit protection. The module form factor includes mounting holes for easy PCB or chassis integration. At approximately \$16 per unit, it is the most expensive of the three options but provides the highest current rating and the most robust protection suite.

The DRV8833 integrates two H-bridge channels at up to 1.5 A each. Although its compact DFN package is attractive for density-constrained layouts, its current rating is well below the actuator's stall current requirements. Driving an actuator with insufficient current

margin risks thermal shutdown at the most critical moment, specifically when the housing is deploying against gravity or spring force.

The L298N is an older bipolar-transistor-based H-bridge that, despite its widespread availability, suffers from a 2 V to 4 V voltage drop across its transistor pairs. At a 12 V supply, this drop translates to a meaningful reduction in effective motor voltage, reducing available torque. Its continuous current rating of 2 A is also marginal for the actuator stall condition. Heat dissipation through its TO-220 package requires careful placement in the PCB layout.

Table 3.2.2.2: Motor Driver Part Comparison

Part	Max Current	Voltage	Price	Notes
Cytron MD10C (13 A DC Motor Driver)	13 A continuous	5–30 V	\$13.91	Selected: robust current rating, PWM+DIR interface, on-board protection
DRV8833 Dual H-Bridge IC	1.5 A per ch	2.7–10.8 V	~\$2	Too low current for 100 N actuator stall; requires external supply
L298N Dual H-Bridge Module	2 A continuous	5–35 V	~\$4	Popular but low current and significant voltage drop (~2 V); heat dissipation issues

The Cytron MD10C was selected as the motor driver for A.U.R.A. Its 13 A continuous rating provides comfortable margin above the actuator stall current, its simple PWM and DIR interface is straightforward to drive from the ESP32 MCU, and its on-board protections reduce the risk of driver damage during the repetitive extend/retract cycles expected over the system's service life.

3.2.3 Hall Effect Sensor

3.2.3.1 Hall Effect Sensor Technology Comparison

The Hall effect sensor is responsible for detecting when the vehicle reaches the 1 mph threshold at which the sensor modules must retract into their protective housings. Accurate and timely velocity detection ensures that the camera arms do not remain extended during vehicle motion, preventing mechanical damage from road contact. In Table 3.2.3.1, three speed-sensing technologies were evaluated: digital Hall effect sensors, accelerometers, and reed switches.

Digital Hall effect sensors detect the passage of a permanent magnet attached to the wheel hub or axle and output a clean digital pulse via an open-collector transistor. Because the detection mechanism is contactless and solid-state, there is no mechanical wear, and the output is inherently immune to the moisture, vibration, and EMI environment beneath a moving vehicle. The sensor supply voltage typically spans 4.5 V to 24 V, which is

compatible with the system's 5 V and 12 V rails. Velocity is computed by counting pulses per unit time in the MCU firmware, a straightforward and well-understood technique requiring no additional hardware.

Accelerometers were also considered by the team as an alternative approach to detecting vehicle motion. An accelerometer measures the rate of change of velocity rather than velocity itself, meaning it would need to integrate its readings over time to estimate whether the vehicle has reached the 1 mph threshold. The issue with this approach is that integration accumulates small measurement errors over time, a phenomenon known as drift, which causes the estimated speed to diverge from the actual speed the longer the system runs. Additionally, the vibrations produced by road surfaces, the electric motor, and wheel rotation all generate accelerometer signals that are difficult to distinguish from actual vehicle acceleration without significant filtering. While accelerometers are effective for detecting whether the vehicle is moving at all, using one as the sole speed measurement for a safety-critical retraction trigger introduces unnecessary complexity and reliability risk compared to the direct pulse-counting approach of a Hall effect sensor.

Reed switches are the simplest and cheapest speed-sensing option, closing their internal contacts whenever a magnet passes nearby. However, they are mechanical devices subject to contact bounce, switch wear over many millions of actuations, and reduced reliability at the switching speeds corresponding to even moderate vehicle velocities.

Table 3.2.3.1: Hall Effect Sensor Technology Comparison

Technology	Output Type	Typical Resolution	Trade-offs for A.U.R.A.
Hall Effect Sensor (digital)	Digital pulse	Pulse per magnet pass	Simple interface, noise-immune, contactless; sufficient for 1 mph threshold detection
Accelerometer (IMU-based)	Analog / I2C / SPI	mg-level acceleration	Can infer motion from vibration/acceleration; no magnet needed; but measures acceleration not speed directly, requiring integration which accumulates drift
Reed Switch	Digital on/off	Pulse per magnet pass	Cheap but mechanical contact wears over time; limited switching speed; less reliable than solid-state alternatives

Digital Hall effect sensing was selected as the technology of choice. It combines contactless operation, broad supply voltage compatibility, solid-state durability, and a simple digital output that is easily processed by the MCU interrupt controller to compute vehicle speed.

3.2.3.2 Hall Effect Sensor Part Comparison

In Table 3.2.3.2, three digital Hall effect sensor parts were compared: the Honeywell SS100 Series (part 55100-2M-02-A), the A3144, and the AH3144E.

The Honeywell 55100-2M-02-A is a unipolar digital Hall effect sensor rated for operation from -40 °C to +125 °C with a supply voltage range of 4.5 V to 24 V. Its open-collector NPN output is directly compatible with the ESP32's GPIO input pins when pulled up to 3.3 V. Honeywell publishes a comprehensive datasheet with guaranteed electrical and magnetic specifications. At approximately \$7 per unit, it is more expensive than generic alternatives and was found to have limited availability through common procurement channels, making it impractical for timely acquisition.

The A3144 is a widely used generic unipolar Hall sensor available from numerous suppliers at under \$1. It operates on a 4.5 V to 24 V supply and provides an open-collector NPN digital output directly compatible with the ESP32 GPIO pins. While the A3144 designation is not controlled to a single manufacturer's standard, its performance for low-speed threshold detection (1 mph on a demonstration vehicle) is well within the capabilities of this sensor family. The speed detection requirement for A.U.R.A. is not precision-critical, the system only needs to reliably distinguish between stationary and slow movement near the 1 mph threshold, which does not demand tightly controlled magnetic thresholds. Combined with its low cost, wide availability, and proven use in automotive speed sensing applications, the A3144 is a practical and appropriate choice for this application.

The AH3144E is another low-cost Hall sensor with a broader supply range of 2.7 V to 30 V, which is attractive on paper. However, like the A3144, it is produced by multiple manufacturers with inconsistent documentation, and it offers no meaningful advantage over the A3144 for this application.

Table 3.2.3.2: Hall Effect Sensor Part Comparison

Part	Supply	Output	Price	Notes
Honeywell SS100 Series (55100-2M-02-A)	4.5–24 V	Open-collector NPN	~\$7	Automotive-grade, broad temp range; well-documented specifications but high cost and limited availability
A3144 Hall Effect Sensor	4.5–24 V	Open-collector NPN	~\$5.50	Widely available, low cost, 5V compatible, direct digital output compatible with ESP32 GPIO; sufficient for low-speed (1 mph) threshold detection on demonstration vehicle
AH3144E Hall Effect Sensor	2.7–30 V	Open-collector NPN	~\$0.60	Similar to A3144; functional but inferior documentation and sourcing reliability

The A3144 Hall Effect Sensor was selected for A.U.R.A. The application requires only that the sensor reliably detect vehicle movement near the 1 mph threshold on a low-speed demonstration vehicle, a task well within the capabilities of the A3144. Its low cost (~\$5.50), wide availability, 5V compatibility, and direct digital output make it a practical

and sufficient choice. The previously selected Honeywell 55100-2M-02-A was found to have limited Amazon availability and a higher unit cost that is not justified given the low-speed, non-precision nature of this detection task.

3.2.4 Proximity Sensor

3.2.4.1 Proximity Sensor Technology Comparison

The proximity sensor in A.U.R.A. is mounted at the end of the actuator arm, co-located with the camera module on the secondary PCB, and connected to the actuator via the 3D-printed housing. As the actuator deploys, the camera module slides outward and downward at an angle of approximately 30 to 45 degrees relative to the vehicle floor. The proximity sensor must detect any obstacle in that angled deployment path before the camera makes physical contact with it, allowing the MCU to halt or reverse the actuator before damage occurs. This placement and geometry place specific demands on the sensor: it must have sufficient range to provide meaningful warning distance along the angled path, it must work reliably regardless of the color or texture of whatever surface the camera is approaching, and it must be small enough to fit on the secondary PCB alongside the camera module.

An important capability that was evaluated alongside sensor selection is whether the actuator can be programmed to stop short of full extension when the sensor detects an obstacle. The answer is yes. Because the MCU reads the proximity sensor output continuously during deployment via the I2C bus, it can compare the measured distance against a configurable threshold at each update cycle. When the distance drops below that threshold, the MCU sends a stop command to the motor driver, cutting PWM output and engaging the actuator's internal brake. The actuator holds its position at that intermediate extension point rather than fully retracting or continuing forward. This means the system does not need to choose between full extension and full retraction; it can stop at any point along the stroke where clearance becomes insufficient. This behavior makes the proximity sensor a functional safety input to the actuator control loop, not merely an alarm.

In Table 3.2.4.1, three broad sensing technologies were evaluated for this role: time-of-flight laser ranging, IR photoelectric proximity, and ultrasonic sensing.

Time-of-flight sensors emit a brief pulse of laser light and measure the time it takes to reflect back from the nearest surface. Because they measure time rather than reflected intensity, their readings are consistent regardless of whether the target surface is light, dark, reflective, or matte. Their range of up to approximately 13 feet far exceeds what is needed for the deployment guard function, but the programmable region of interest available on newer ToF devices allows the sensing cone to be narrowed and aimed along the specific angle of camera travel, which is well-suited to the 30 to 45 degree deployment path geometry. Update rates of up to 50 Hz provide more than enough time resolution for the relatively slow actuator extension speed.

IR photoelectric proximity sensors emit modulated infrared light and estimate distance from the intensity of the reflected signal. They are extremely compact and low cost, but their effective range of approximately 8 inches is marginal for providing early warning

along a 6-inch actuator stroke, leaving very little reaction time for the MCU to stop the actuator. They are also sensitive to surface reflectivity, meaning a dark surface directly in the deployment path may not be detected reliably.

Ultrasonic sensors use sound pulses to measure distance and offer ranges comparable to ToF sensors. However, their physical dimensions are too large to mount on the secondary PCB at the actuator tip alongside the camera module. Their response time of approximately 50 ms is also slow relative to the actuator extension speed, and as noted in prior research on animal detection systems, ultrasonic sensing performs poorly on soft, absorptive surfaces such as fur, which is a meaningful limitation given A.U.R.A.'s application context.

Table 3.2.4.1: Proximity Sensor Technology Comparison

Technology	Range	Interface	Trade-offs for A.U.R.A.
Time-of-Flight (ToF) Laser Ranging	Up to ~13 ft	I2C	Precise distance at any angle; programmable field of view; works regardless of target color or reflectivity; ideal for angled deployment path
IR Photoelectric Proximity	Up to ~8 inches	I2C / analog	Compact and low cost but short range and performance varies with surface reflectivity; limited for angled deployment
Ultrasonic (HC-SR04 style)	~1 inch to 13 ft	Trigger / Echo	Long range but physically large; slow response (~50 ms); unreliable on soft surfaces like fur; too bulky for mounting at camera end of actuator

Time-of-flight laser ranging was selected as the technology of choice. Its surface-independent distance measurement, programmable field of view, compact package, and I2C interface make it well suited for obstacle detection along the angled camera deployment path.

3.2.4.2 Proximity Sensor Part Comparison

In Table 3.2.4.2, three specific time-of-flight sensor parts were compared: the STMicroelectronics VL53L1X, the VL53L0X, and the Vishay VCNL3040.

The STMicroelectronics VL53L1X is a second-generation ToF ranging sensor with a maximum range of approximately 13 feet, a 50 Hz update rate, and a programmable region of interest that allows the active sensing cone to be narrowed from its default wide field down to a tight beam that can be oriented along the 30 to 45 degree camera deployment angle. It communicates via I2C at up to 400 kHz and is available in a 2.4 mm by 2.4 mm package that easily fits on the secondary PCB. At approximately \$5 per unit, it offers the best combination of range, update speed, and field-of-view control of the parts evaluated.

The STMicroelectronics VL53L0X is the predecessor to the VL53L1X and uses the same I2C interface and package footprint. Its maximum range of approximately 6.5 feet is

shorter, its update rate is lower at up to 30 Hz, and it lacks the programmable region of interest feature of the newer part. At approximately \$3 it is slightly cheaper, but the performance difference is meaningful enough to justify the additional cost of the VL53L1X given the safety-critical nature of the deployment guard function.

The Vishay VCNL3040 is a compact IR proximity sensor with an effective range of approximately 8 inches. While it is an excellent part for short-range detection tasks, its range is not sufficient to provide meaningful warning distance along a 6-inch actuator stroke at a 30 to 45 degree deployment angle, where the slant distance to the ground or an obstacle may be considerably longer than the vertical clearance alone suggests.

Table 3.2.4.2: Proximity Sensor Part Comparison

Part	Range	Interface	Price	Notes
STMicro VL53L1X	Up to ~13 ft	I2C	~\$5	Selected: programmable region of interest, 50 Hz update rate, works at 30-45 degree angles, tiny footprint
STMicro VL53L0X	Up to ~6.5 ft	I2C	~\$3	Shorter range and slower update rate than VL53L1X; less suitable for angled deployment path detection
Vishay VCNL3040	Up to ~8 inches	I2C	~\$1.55	Very compact and cheap but range is too short for the angled camera deployment scenario

The STMicroelectronics VL53L1X was selected for A.U.R.A. Its long range, fast update rate, surface-independent ToF measurement, and programmable field of view make it the most capable and appropriate sensor for monitoring the angled camera deployment path. Its I2C interface integrates cleanly with the ESP32-S3 MCU on the secondary PCB, and the MCU firmware can use the real-time distance readings to stop the actuator at any intermediate position if an obstacle is detected before full extension is reached.

3.2.5 Processor

A microcontroller unit was selected as the low-level control platform for each A.U.R.A. PCB node because the system requires real-time deterministic control of multiple peripherals simultaneously, including PWM generation for motor drivers, interrupt-driven pulse counting from the Hall effect sensor, I2C communication with proximity sensors, and wireless data transfer to the central processor. In Table 3.2.5, several alternative computing platforms were evaluated before this decision was made.

FPGAs offer high parallelism but require hardware description language expertise, carry significantly higher cost, and lack native wireless connectivity. Digital signal processors are optimized for signal processing only and do not provide the general-purpose peripheral integration needed for A.U.R.A.'s mixed control and communication workloads. Single-board computers provide high computational performance but operate under a general-

purpose operating system that introduces scheduling latency, making them unsuitable for deterministic real-time GPIO control.

A microcontroller provides the best balance for this application. It combines real-time deterministic peripheral control, integrated 2.4 GHz Wi-Fi and Bluetooth Low Energy for direct wireless communication with the Raspberry Pi 5 central processor, low cost, and a well-supported development ecosystem, all in a single chip without requiring additional wireless modules or external hardware.

Table 3.2.5: Processing Platforms Technology Comparison

Technology	Processing	Real-Time Control	Wireless	Trade-offs for A.U.R.A.
MCU	240 MHz, dual-core	Yes, deterministic GPIO	Built-in Wi-Fi + BLE	Real-time peripheral control and wireless on a single low-cost 3.3V chip with no external modules needed
FPGA	Parallel hardware execution	Yes, hardware level	None native	Requires HDL expertise; higher cost and complexity; no native wireless; impractical for this application
DSP	Signal processing optimized	Limited GPIO	None native	Specialized for signal processing only; lacks peripheral integration for actuator control, sensors, and wireless communication

3.2.5.1 MCU Technology Comparison

A.U.R.A. employs three MCU nodes: one on the main PCB (MCU #1, acting as AI processor and primary controller), and one each on PCB 2 and PCB 3, which govern the front and rear camera subsystems respectively, including actuator control, Hall effect velocity monitoring, proximity sensing, alarm activation, and H-bridge driving. Each MCU node must support concurrent real-time peripheral management, communicate wirelessly with the Raspberry Pi 5 central processor, and operate from a 3.3 V logic supply.

Dual-core 32-bit processors based on the Tensilica Xtensa LX6 or LX7 architecture, most prominently the Espressif ESP32 family, offer clock speeds up to 240 MHz with hardware floating-point units, multiple hardware timers, and up to 34 GPIO pins. Critically, they include an integrated 2.4 GHz Wi-Fi and Bluetooth/BLE radio, eliminating the need for a separate wireless module and substantially reducing PCB footprint and BOM cost. Their 3.3 V I/O is directly compatible with the sensors and drivers in the system, and the Arduino-compatible ESP-IDF framework provides extensive library support.

ARM Cortex-M4 based microcontrollers such as the STM32F4 series offer slightly higher single-thread performance per clock cycle and a rich peripheral set, but they do not include an integrated Wi-Fi radio. Adding a separate Wi-Fi module such as the ESP8266 or ATWINC1500 restores wireless capability but increases component count, PCB area, and firmware complexity. The STM32 toolchain (CubeIDE, HAL libraries) is powerful but has

a steeper learning curve relative to the ESP32 ecosystem, which is a practical consideration for the team's existing development experience.

8-bit AVR microcontrollers such as the ATmega328P are the basis of the popular Arduino Uno platform and are extremely well documented. However, at 20 MHz maximum clock speed they lack the computational throughput needed for concurrent real-time management of multiple peripherals, and they have no wireless connectivity whatsoever. Their use would require multiple external modules and substantially more firmware effort to replicate what the ESP32 provides natively.

Table 3.2.5.1: MCU Technology Comparison

Architecture	Clock Speed	Connectivity	Trade-offs for A.U.R.A.
Dual-core 32-bit (Xtensa LX6 / LX7)	Up to 240 MHz	Wi-Fi, BT, BLE on-chip	Handles real-time control with integrated wireless; reduces BOM; 3.3 V logic compatible
ARM Cortex-M4 (e.g., STM32F4)	Up to 168 MHz	None native (add module)	Faster per clock; no built-in Wi-Fi; requires separate module; more complex toolchain
8-bit AVR (ATmega328P)	Up to 20 MHz	None	Simple and cheap but insufficient for multi-peripheral real-time control; no wireless

The dual-core 32-bit ESP32 architecture was selected. The integration of Wi-Fi, BLE, sufficient GPIO, hardware PWM channels, I2C, and UART into a single low-cost chip directly addresses the multi-peripheral, wireless-communication requirements of each A.U.R.A. PCB node.

3.2.5.2 MCU Part Comparison

In Table 3.2.5.2, three specific modules were compared: the ESP32-S3 Development Board based on the ESP-1-N16R8 module, the ESP32-WROOM-1U, and the Nordic nRF52840.

The ESP32-S3 Development Board built around the ESP-1-N16R8 module is the team's selected platform. The ESP-1-N16R8 designation indicates an ESP32-S3 chip paired with 16 MB of quad-SPI flash and 8 MB of octal-SPI PSRAM. The dual Xtensa LX7 cores operate at up to 240 MHz, offering a meaningful performance improvement over the older LX6 architecture in floating-point and bit-manipulation tasks. The 8 MB of PSRAM is a particularly significant advantage for A.U.R.A.: the MCU nodes must concurrently manage PWM generation for the motor driver, interrupt-driven pulse counting from the Hall effect sensor, I2C polling of the proximity sensor, UART communication with the alarm driver, and Wi-Fi socket management. Buffering these data streams alongside Wi-Fi stack structures and FreeRTOS task contexts within only 520 KB of internal SRAM, as would be required on the WROOM-1U, creates a tight and fragile memory budget. The 8 MB PSRAM eliminates this constraint entirely, allowing larger task stacks, more generous ring

buffers, and headroom for firmware growth as the system is refined. The development board form factor also includes an onboard USB-to-UART bridge and reset/boot buttons, which substantially simplify firmware flashing and serial debugging during prototyping without requiring a separate programmer dongle.

The ESP32-WROOM-1U is a widely used module based on the original ESP32-D0WD-V3 dual-core LX6 chip with 4 MB of flash and 520 KB of SRAM, and it exposes a U.FL connector for an external antenna. It has been in continuous production since 2019 and benefits from the most mature Arduino and ESP-IDF library ecosystem of any ESP32 variant. However, its 520 KB SRAM is the binding constraint for A.U.R.A.'s multi-peripheral firmware. With no external PSRAM available on the WROOM-1U module, the developer must carefully tune task stack sizes, minimize buffer allocations, and limit the number of concurrent FreeRTOS tasks, all of which increase development risk and reduce long-term maintainability. Its older LX6 core also lacks the vector instructions present in the LX7, which may become relevant if the animal classification model is later ported partially to the MCU layer.

The Nordic nRF52840 is an excellent BLE and Thread/Zigbee-focused microcontroller favoured in wearable and low-power IoT applications. Its ARM Cortex-M4 core running at 64 MHz and its integrated 2.4 GHz radio are capable for many embedded tasks, but it does not support the IEEE 802.11 Wi-Fi protocol natively. Since the A.U.R.A. architecture uses Wi-Fi as the primary communication medium between the PCB nodes and the Raspberry Pi 5 central processor, using the nRF52840 would require adding a separate Wi-Fi module such as the ATWINC1500 or an ESP8266 co-processor. This negates the cost advantage and reintroduces the very integration overhead that the ESP32 platform eliminates by design.

Table 3.2.5.2: MCU Part Comparison

Part	Cores / Speed	Wireless	Price	Notes
ESP32-S3 Dev Board (ESP-1-N16R8)	Dual 240 MHz LX7	Wi-Fi + BLE 5.0	~\$8–12	Selected: 16 MB flash, 8 MB PSRAM, onboard USB, AI acceleration; ample memory for multi-peripheral firmware
ESP32-WROOM-1U (ESP32-D0WD-V3)	Dual 240 MHz LX6	Wi-Fi + BT/BLE	~\$5.71	Mature and low-cost but only 4 MB flash, 520 KB SRAM; no PSRAM; tighter memory budget for concurrent tasks
Nordic nRF52840	Single 64 MHz Cortex-M4	BLE, Thread, Zigbee	~\$6.00	BLE-focused; no native Wi-Fi; requires separate module; unsuitable for Wi-Fi-based architecture

The ESP32-S3 Development Board (ESP-1-N16R8) was selected as the MCU for all three PCB nodes in A.U.R.A. The combination of dual LX7 cores at 240 MHz, 16 MB flash, 8

MB PSRAM, integrated Wi-Fi and BLE 5.0, onboard USB, and a well-supported ESP-IDF/Arduino ecosystem makes it the most capable and practical choice for managing the concurrent real-time peripheral demands, wireless communication, and firmware development workflow of this project.

3.2.6 Wireless Antenna

3.2.6.1 Antenna Technology Comparison

The wireless antenna is the component that couples the ESP32-S3's 2.4 GHz radio energy into the surrounding environment. The ESP-1-N16R8 module exposes an IPEX/U.FL connector for an external antenna rather than relying solely on a PCB trace, allowing antenna placement to be decoupled from PCB placement. An appropriate external antenna must therefore be chosen and routed to a location with sufficient separation from metal chassis structures. In Table 3.2.6.1, three antenna types were evaluated: integrated PCB trace antennas, flexible PCB / FPC antennas, and rigid whip or rod antennas.

PCB trace antennas are fabricated directly on the module substrate and require no external components or connectors, minimizing assembly cost and BOM count. They typically achieve approximately 0 dBi omnidirectional gain, which is acceptable in open-air environments. However, A.U.R.A.'s PCB nodes are mounted close to the vehicle's steel chassis, and the die-cast and stamped metal components of the toy car will significantly attenuate and detune a trace antenna that cannot be repositioned away from the metal. The U.FL connector on the ESP-1-N16R8 module directly addresses this by allowing the antenna to be routed to a metal-free surface area of the 3D-printed housing, maintaining reliable signal quality regardless of PCB orientation relative to the chassis.

Flexible PCB antennas, also called FPC (Flexible Printed Circuit) antennas, are thin adhesive-backed strips that can be routed through the enclosure and adhered to a non-metallic surface area with good line of sight to the intended communication partner. They connect to the U.FL port via a short coaxial pigtail. At 2.4 GHz they typically achieve 2 dBi to 3 dBi gain, comparable to a quarter-wave whip, while remaining thin enough to pass through narrow gaps in the housing without mechanical stress. Their flexibility also means they are resistant to vibration-induced fatigue failures that can affect rigid connectors.

Rigid whip and rod antennas offer the highest gain of the three categories, up to 5 dBi in some designs, and their performance is generally more predictable than flexible variants. However, their length, typically 100 mm to 180 mm for a 2.4 GHz quarter-wave element, would protrude from the undercarriage housing in a way that is incompatible with the IP65-rated retractable enclosure design. A protruding antenna would be vulnerable to road debris strikes during vehicle motion and would prevent the housing from sealing properly.

Table 3.2.6.1: Wireless Antenna Technology Comparison

Antenna Type	Form Factor	Gain	Trade-offs for A.U.R.A.
--------------	-------------	------	-------------------------

PCB Trace (on-module)	None (integrated)	~0 dBi	Free with module but signal blocked by metal chassis; inadequate in enclosed housing
Flexible PCB / FPC Antenna (external)	Adhesive strip ~50 mm	~2–3 dBi	Routes through housing; mounts away from metal; low cost; compatible with U.FL
Rigid Whip / Rod Antenna	100–180 mm rod	~2–5 dBi	Higher gain but protrudes; vulnerable to road debris; not suitable for retractable housing

The flexible PCB antenna technology was selected. It provides adequate gain for the short-range communication distances within the vehicle, can be routed to a metal-free location on the housing exterior, is compatible with the U.FL connector on the ESP-1-N16R8 module, and does not compromise the structural integrity or IP rating of the protective enclosure.

3.2.6.2 Antenna Part Comparison

In Table 3.2.6.2, three specific flexible antenna products were compared: the Adafruit 2.4 GHz Mini Flexible Wi-Fi Antenna, the Molex 2042800100, and the Taoglas FXP73.

The Adafruit 2.4 GHz Mini Flexible Wi-Fi Antenna is a 50 mm FPC antenna with an integrated U.FL pigtail and a peel-and-stick adhesive backing. It is rated for 2.4 GHz operation and achieves approximately 2 dBi gain. The adhesive backing simplifies installation inside the 3D-printed housing, as the antenna strip can be adhered to the inner wall of the ABS/ASA top panel, away from any metallic components, during assembly without the need for additional fasteners or brackets. At approximately \$2.50 per unit, three antennas are required for the three ESP32 nodes, representing a total antenna cost of \$7.50, which is within budget.

The Molex 2042800100 is a similar 2.4 GHz FPC antenna with comparable electrical performance, achieving approximately 2.1 dBi gain. However, it does not include an adhesive backing, requiring the designer to either add double-sided tape during assembly or design a dedicated retention feature into the housing. This adds minor but unnecessary assembly complexity.

The Taoglas FXP73 is a professional-grade FPC antenna designed for embedded wireless devices, with approximately 1.5 dBi gain and a higher manufacturing quality standard. At approximately \$4.00 per unit, its cost is 60% higher than the Adafruit part for marginally lower specified gain. The quality premium is more relevant in high-volume production environments than in a one-off research prototype.

Table 3.2.6.2: Wireless Antenna Part Comparison

Part	Freq.	Gain	Price	Notes
------	-------	------	-------	-------

Adafruit 2.4 GHz Mini Flexible Wi-Fi Antenna (U.FL)	2.4 GHz	~2 dBi	~\$2.50	Selected: flexible, adhesive-backed, U.FL connector, proven with ESP32
Molex 2042800100 Flexible Antenna	2.4 GHz	~2.1 dBi	~\$3.20	Similar performance; no adhesive backing; less convenient mounting in 3D-printed housing
Taoglas FXP73 Flexible PCB Antenna	2.4 GHz	~1.5 dBi	~\$4.00	Higher cost; marginally lower gain; designed for larger enclosures

The Adafruit 2.4 GHz Mini Flexible Wi-Fi Antenna was selected for A.U.R.A. It provides sufficient gain for the intra-vehicle communication distances involved, includes an adhesive backing that simplifies integration into the 3D-printed housing, terminates in a U.FL connector directly compatible with the ESP-1-N16R8 module, and represents the lowest-cost option among the candidates evaluated. Three units will be used, one per ESP32-S3 PCB node, at a combined cost of \$7.50.

In the updated selection, no separate antenna was used in the complete PCB, this was due to the fact that ESP32-S3 modules come with an already included antenna which allowed for the use of Wi-Fi and ESP32 to ESP32 wireless communication without the need for an external antenna.

3.2.7 High-Frequency Alarm:

Other than commonly used components used in most circuits, high-frequency alarms and ultrasonic repellent systems researched generally have 3 main sections in order to function properly: the first section is the frequency generator which commonly use a 555 timer or oscillator, the second section is the amplifier sections which makes use of transistors to amplify the driver into the output stage, and finally the third main section is made up of the output containing the component necessary to output a high-frequency ultrasonic signal. For the design of the high-frequency alarm used in our system we will make use of these sections to research and choose components necessary to produce a functioning ultrasonic repellent.

3.2.7.1 Frequency Generator Technology Comparison:

Our frequency generator needs to be able to generate a square wave input between our frequency range of 30kHz to 40kHz to meet our specification requirements. It also needs to have the capability to vary the frequency in order for the ultrasonic alarm to have adaptable capabilities and to function as an effective repellent.

Oscillators can be used to generate the periodic signals used to generate an ultrasonic output. In the case of voltage-controlled oscillators (VCO), the frequency they generate is dependent on the input voltage; combined with a potentiometer to vary the input voltage

could make a VCO a viable option for the frequency generation stage of the high-frequency alarm.

555 timers are a kind of oscillator that can be used to generate stable square waveforms. The 555 timer is a popular component used in commercial pest repellent products; this is due to the inexpensive cost, reliability and simplicity of 555 timers.

Our design will contain a low level MCU (ESP32) to control and process hardware functions, an ESP32 can be used to generate stable square waves at the frequencies required. Using the ESP32's LEDC peripheral, the low level MCU can be used in the frequency generator stage of our high-frequency alarm without the need for additional components to vary the input frequency such as potentiometers.

Table 3.2.7.1 Frequency Generator Technology Comparison:

Technology	Cost	Square Wave Generation	Frequency Control	Scalability
Voltage Controlled Oscillator	Moderate	Requires additional components	Based on the input voltage	Moderate
555 Timer	Low	Natural square wave generation	Needs additional components	Moderate
MCU	Moderate, low when considering entire system	Driven by MCU peripherals	Software based	High

While oscillators and 555 timers seem to be common practice among consumer level pest repellent systems, the ESP32 offers faster real-time frequency switching and modulation as opposed to a 555 timer. Compared to a voltage-controlled oscillator, the ESP32 offers higher cost efficiency and simplicity.

3.2.7.2 Frequency Generator Component Comparison:

Since we have opted to use our microcontroller unit as the frequency generator for the ultrasonic alarm, the exact component we use will be entirely based on the ideal MCU for the overall system. Here we will discuss the comparison of the components bases on the concept of frequency generation.

Table 3.2.7.2 Frequency Generator Component Comparison:

Component	Hardware	Maximum Frequency	H-bridge Support
ESP32-S3 Dev Board (ESP-1-N16R8)	MCPWM + LEDC	40 MHz	Built-in deadtime

ESP32-WROOM-1U (ESP32-D0WD-V3)	MCPWM + LEDC	40 MHz	Built-in deadtime
Nordic nRF52840	PWM + Timer	16 MHz	Software based

The ESP32-S3 is still chosen as the MCU of choice as the ESP32 provides the best method of generating the necessary frequency using its peripherals of the Motor control pulse width modulation and LED control. The ESP32 also provides the higher maximum frequency and a built-in deadtime support for our amplifier.

3.2.7.3 High-Frequency Output Technology Comparison:

The component that will be used to output the high-frequency repellent needs to be able to generate an effective noise within our frequency range of 30 kHz – 40 kHz. It also needs to be able to generate this frequency at an effective range and stability to act as an effective animal repellent.

Piezoelectric transducers are components that convert between electrical signals and mechanical energies. In the context of this project and the ultrasonic realm, ultrasonic piezoelectric transducers can range from 20 kHz to several MHz which makes them a possible component that can be used as the speaker for our high-frequency alarm system, as we are operating withing the range of 30KHz to 45Khz in our application.

Magneto strictive transducers are components that convert electrical energy into mechanical energy using the magnetostriction effect to expand and contract metals using an AC current created magnetic field, this causes the metals to resonate and produce an output. While Magento strictive transducers have applications in the ultrasonic field, specifically ultrasonic cleaning, they generally are used for lower frequency applications.

Ultrasonic Tweeter speakers or super tweeters use electrical energy to generate sound waves, generally above the 20 kHz minimum to produce an ultrasonic output. In applications, these components are generally used to enhance existing audio systems to improve quality of sound output. Super tweeters are highly directional, which may cause it so that any change in the placement of either the tweeter or what the tweeter is pointed at may cause it so that the output isn't registered audibly.

Table 3.2.7.3 High-Frequency Output Technology Comparison:

Technology	Cost	Output Strength	Range	Reliability
Piezoelectric Transducers	Low	High	Moderate	High

Magneto strictive Transducers	High	Very High	Long	Very High
Ultrasonic Tweeter Speakers	Moderate	Moderate	Short	Moderate

In consideration of our system's constraints and adaptable output capabilities, the best option for the ultrasonic output would be piezoelectric transducers. Compared to Magneto strictive transducers and super tweeters, piezoelectric transducers offer higher power and drive efficiency, a more stable and applicable high-frequency output for our application and are more cost effective.

One limitation of piezoelectric transducers for our adaptable output application is that they are usually locked to a certain frequency range and any major deviation from the intended frequency may weaken its overall output. In order to mitigate this limitation and support our adaptable frequency capabilities an array of 2 kinds of piezoelectric transducers will be implemented: one transducer array will function on the lower end of our intended frequency range (30 kHz) and the other will function on the higher end of our intended frequency range (40 kHz).

3.2.7.4 High-Frequency Output Component Comparison:

Table 3.2.7.4 High-Frequency Output Component Comparison

Component	Cost (Per Unit)	Beam Angle	Voltage Rating	Availability
40 kHz Transducer				
MA40S4S	\$5.77	80°	20 V _{p-p}	High
CUSA-T80-18-2400-TH	\$2.88	80°	150 V _{p-p}	Moderate
CUSA-T80-15-2400-TH	\$2.96	80°	80 V _{p-p}	Moderate
30.5 kHz Transducer				
CUSA-T80-130-3000-TH	\$2.17	80°	200 V _{p-p}	Moderate
CUSA-T80-120-2200-W	\$2.29	80°	160 V _{p-p}	Moderate
CUSA-T75-120-2200-W	\$2.40	75°	160 V _{p-p}	Moderate

While both CUSA options for the 40 kHz transducer are both less expensive and have higher voltage ratings, their sound pressure levels drop off much quicker and are at a lower peak in comparison to the MA40S4S. The CUSA-T80-15-2400-TH peaks at below 110 dB while the CUSA-T80-18-2400-TH peaks at about 120 dB similar to that of the MA40S4S, however, while the CUSA-T80-18-2400-TH drops to 90 dB at about 43 kHz, the MA40S4S

doesn't drop to 90 dB until 50 kHz making it the more reliable component for the ultrasonic output.

For the 30.5 kHz transducer, there were far less option in terms of transducer transmitter with those found on Digikey being only from the CUSA series. For the three options above, both the CUSA-T75-120-2200-W and CUSA-T80-120-2200-W peak at lower SPL values, with the CUSA-T80-120-2200-W peaking at 100 dB and the CUSA-T75-120-2200-W peaking at about 105 dB, in contrast, the CUSA-T80-130-3000-TH peaks at a little over 120 dB. While all three components have similar curves in terms of drop-off, due to the CUSA-T80-130-3000-TH much higher peak SPL, it has higher SPL values as the frequency drifts from its resonant frequency of 30.5 kHz, making it the more reliable component for our application.

3.2.7.5 Amplifier Technology Comparison:

The amplifier needed for our alarm system is dependent on the kind of component that will be used for the output. As we are using piezoelectric transducers, the amplifier needs have the capability to drive our transducer array such that all of the transducers produce an effective output in terms of amplitude and range.

A single transistor can be used as an amplifier for our ultrasonic output. In a common emitter configuration by using a small voltage signal then amplifying it into a high-frequency, high-voltage signal. For our application a high-voltage capable MOSFET is preferred over a BJT as they offer higher efficiency and switching speeds.

A half-bridge amplifier uses two transistors to drive a load from the DC voltage supply to ground. When it comes to our high-frequency application, they are efficient in producing the power necessary to drive a high-frequency output.

A H-bridge amplifier uses four transistors to double the available peak-to-peak voltage, maximizing the available power we have for our ultrasonic output. H-bridges offer higher efficiency, switching speeds, power output, and output control over the single and half-bridge transistor amplifier configurations.

Table 3.2.7.5 Amplifier Technology Comparison

Technology	Cost	Amplification	Efficiency	Complexity
Single Transistor	Lowest	Lowest	Lowest	1 MOSFET
Half-Bridge	Moderate	Moderate	Moderate	2 MOSFETs
H-Bridge	Highest	Highest	Highest	4 MOSFETs

Given the different amplifier technologies reviewed for our application, the H-bridge amplifier offers the most. Compared to the single MOSFET and half-bridge amplifiers, the H-bridge improves on nearly every aspect apart from component count and cost.

3.2.7.6 Amplifier Component Comparison:

Table 3.2.7.6 Amplifier Component Comparison

Component	Cost	Drain to Source Voltage	Drive Voltage	Drain to Source Resistance
IRLZ44NPBF	\$1.78	55 V	4 V, 10 V	22 mΩ
IRLZ34NPBF	\$1.56	55 V	4 V, 10 V	35 mΩ
IRFZ44NPBF	\$1.41	55 V	10 V	17.5 mΩ

Logic-level N-channel MOSFETs were chosen as they can interface easily with our MCU and provide a better performance for the ultrasonic transducer load in comparison to using an N-channel and P-channel configuration for the H-bridge. While being slightly more expensive than the other options, the IRLZ44NPBF provides a better balance of characteristics; it has the low R_{ds} necessary to minimize heat and has the same maximum V_{ds} necessary to operate with our 12 V input. Additionally, the IRLZ44NPBF contains a strong drive voltage range when either the maximum or minimum R_{ds} is on, making it more versatile than the IRFZ44NPBF.

3.2.8 LED Strobe Light:

In addition to an ultrasonic output, we will design a high-intensity LED strobe light repellent to act as an additional visual irritant. Doing so will make our alarm system similar to those used by some commercial pest repellents. The strobe light will make use of two sections, the input and output; the input being controlled by an MCU to activate and flash at desired intervals. The output will make use of high-intensity white LEDs to produce an irritating visual stimulant.

3.2.8.1 LED Technology Comparison:

Indicator LEDs provide low intensity, energy efficient lighting. They are primarily used as a visual signal for power and other electrical functions. Our system contains a variety of functions that cannot always be seen operating just through visual means; indicator LEDs can provide a visual sign that current is flowing properly through different parts of our system without the need for additional tools.

High-power LEDs produce more intense lighting in comparison to a standard indicator LED. These LEDs require connection to a constant current source to be able to function properly, requiring more components in comparison to Indicator LEDs, as a component to regulate current should be integrated to optimally use high-power LEDs. For our systems purpose, the LED needs to flash bright enough to act as a visual irritant to any animal that may be under the vehicle; high-power LEDs would be a viable option due to their intense lighting.

Chip on Board (COB) LEDs contain multiple tightly packed LEDs on a single chip to form a single module; by doing so this increases the lumen output while requiring less components. In our system, a COB LED module could be used as a less resource and surface area intensive strobe light for animal deterrence given that they produce more intense lighting than an array of standard LED components

Table 3.2.8.1 LED Technology Comparison:

Technology	Cost	Light Intensity	Light Coverage	Power Efficiency
Indicator LED	Low	Low	Low	High
High-Power LED	Low cost per component but adds up with LED array	High output	Lower coverage, more directional	Slightly lower than COB modules, requires more components in an array
COB LED Module	Moderate cost per module	Very high output	Wide coverage	Higher efficiency, considers fewer individual components

Given that our alarm system has components that cannot be seen such as the high-frequency output from our transducer arrays, indicator LEDs could function as a visual sign that the transducer arrays are, in fact, receiving current and are active. For our high-intensity strobe light, high-power LEDs seem to offer more compared to COB LED modules in terms of power efficiency, light intensity, coverage, and thermal management.

3.2.8.2 LED Component Comparison:

Table 3.2.8.2 LED Component Comparison

Component	Cost	Wattage	Viewing Angle	Forward Current
JH-3535W12S40-T8B-A	\$2.06	3 W	120°	700 mA
JH-3535W12S40-T7B-A	\$2.38	3 W	120°	700 mA
10DGL-DZ-3W-WH	\$0.70	3 W	120°	600 mA

While less expensive than the other options, the 10DGL-DZ-3W-WH also does not come included with a heat-sink base needed to protect the high-power LED from the heat it generates. The remaining options have nearly identical specifications, due to this the JH-3535W12S40-T8B-A has been chosen as it is the least expensive part containing an SMD

high-power LED mounted on a base heat-sink, reducing the complexity of heat management for the high-power LEDs within our strobe light system.

In SD2, the 3W LEDs were replaced with common indicator LEDs due to less heat management issues and eye safety.

3.2.9 Motor Interlock

Our system contains a motor-interlock system to stop the vehicle in case an animal is detected in the undercarriage. In order to function, the module necessary to cut power from the LEDs needs to have the capability to withstand the high-current spikes generated by the 12V Longwei 390 dual motor system used by our miniature electric vehicle.

3.2.9.1 Motor Interlock Technology Comparison

Relay modules are used in automotive systems as an electrically operated switch to control high power devices using low power signals. Being that our miniature test vehicle uses dual Longwei 390 motors, the component needed to safely interlock the motors needs to at minimum, be able to handle the high current needed to start and stall the motors.

Motor drivers handle more complex motor operations such as speed, direction, and torque using pulse width modulation and H-bridges. By utilizing a motor driver in series with the vehicle's dual motor system it's possible we can control the output of the motors when our system detects an animal.

Much like relay modules, automotive relays use low power signals to control high power components in vehicle applications. There are a few main types of relays, but only there would be most suitable for our interlock application: a solid-state relay, a changeover relay, and a make & break relay. Solid-state relays use semiconductors to control high-current loads. Changeover relays toggle powers between two different circuits. Make & break relays act like switches or on/off controllers.

Table 3.2.9.1 Motor Interlock Technology Comparison:

Technology	Cost	Functionality	Reliability	Complexity
Relay Module	Low	ON/OFF	Moderate	Moderate
Motor Driver	Moderate	PWM	Moderate	High
Automotive Relay	Low	ON/OFF	High	Low

Given that our system only needs to have the functionality to stop and restart the vehicles motors whenever the system detects an animal underneath the vehicle, the advanced functionalities of a motor driver are unnecessary. The vehicle we are showcasing our system on uses 12 V Longwei 390 dual motors which produce high current on start, because of this a relay module may not be the best fit as they are typically designed for about 10 A.

The most optimal technology to use for our system’s motor interlock is a make & break (SPST) automotive relay as they can be used for our simple on/off high-current application.

3.2.9.2 Motor Interlock Component Comparison

Table 3.2.9.2 Motor Interlock Component Comparison

Component	Cost	Contact Rating	Operate/Release Time	Switching Voltage
CB1AHF-R-24V	\$5.43	70 A	15 ms	14 VDC
CB1AH-12V	\$3.79	70 A	15 ms	14 VDC
3-1415001-2	\$11.88	70 A	7ms/2ms	12 VDC
JD1912	\$7.99	40 A	15ms	12VDC

Given that our system needs to handle the high-current spikes that might be generated from cutting power from the existing dual motor system on our miniature electric vehicle, all components meet this constraint. However, the CB1AHF-R-24V requires a higher voltage to function therefore it is not viable. Between the CB1AH-12V the other alternatives, the CB1AH-12V and 3-1415001-2 have similar contact ratings and operating voltages, however, while being more expensive than the CB1AH-12V, the JD1912 is far more readily available and comes attached with connecting adapter that are necessary to integrate with into our final, functional system.

3.2.10 Power Supply and Battery

3.2.10.1 Function and Role

The power supply and battery subsystem provides the electrical energy required to operate all components of the A.U.R.A. system, including the ESP32 nodes, camera modules, proximity sensors, and actuator control circuitry. Unlike the base vehicle battery, this subsystem is designed specifically to support the sensing and control architecture, ensuring stable voltage delivery and isolation from high-noise motor environments. The battery must supply both low-power logic circuits and intermittent high-current actuator loads while maintaining reliability and safety under varying operating conditions.

Table 3.2.10.3.A: A.U.R.A. Estimated Peak Power Budget

Component	Supply Voltage	Est. Current Draw	Est. Power
Raspberry Pi 5	5 V	2,500 mA (peak)	12.5 W
ESP32-S3 (x3)	3.3 V	240 mA total	0.8 W

NIR LED Array (x2 modules)	12 V	500 mA total	6.0 W
Linear Actuators (x2, peak)	12 V	1,000 mA total	12.0 W
Cytron Motor Drivers (x2)	12 V	200 mA idle	2.4 W
LCD Dashboard (7 inch)	5 V	800 mA	4.0 W
Sensors, Relay, Misc.	3.3–5 V	~200 mA	~1.0 W
Total Estimated Peak			~38.7 W

3.2.10.2 Battery Technology Comparison

In Table 3.2.10.2, several battery technologies were evaluated for the system, including Sealed Lead-Acid (SLA), Lithium-Ion (Li-ion), and Lithium Iron Phosphate (LiFePO₄). These technologies were compared based on energy density, safety, cycle life, cost, and ease of integration. SLA batteries are low-cost and simple to implement but are bulky and have limited cycle life. Li-ion batteries provide high energy density and reduced weight, but they require complex battery management systems and present higher safety risks. LiFePO₄ batteries offer a strong balance between safety, long cycle life, and stable voltage output, making them well-suited for embedded and automotive-adjacent applications.

Table 3.2.10.2: Battery Technology Comparison

Technology	Main Strength	Main Limitation	Suitability
Sealed Lead-Acid (SLA)	Low cost, simple integration	Heavy, low cycle life	Not ideal for compact system
Lithium-Ion (Li-ion)	High energy density, lightweight	Safety concerns, complex BMS	Good but riskier
Lithium Iron Phosphate (LiFePO ₄)	High safety, long cycle life, stable voltage	Higher cost	Best overall fit

3.2.10.3 Battery Part Selection

At the component level, several battery configurations were evaluated to meet the system's voltage and current requirements. A 12V-class supply is required to support actuator operation while also allowing efficient regulation down to 5V and 3.3V rails for logic components. Candidate battery options included SLA packs, standard Li-ion packs, and LiFePO₄ battery modules.




Figure 3.2.10.3.B: A.U.R.A. Power Distribution Architecture

The total estimated peak power consumption of approximately 38.7 W represents the worst-case scenario in which both actuators are extending simultaneously, both NIR LED arrays are illuminated, the Raspberry Pi 5 is under full processing load, and all other subsystems are active. In practice, the actuators operate only briefly during startup deployment and retraction, and their simultaneous operation is unlikely under normal use. The more representative steady-state power consumption during the active detection phase, when only the NIR arrays, processing hardware, sensors, and display are drawing current, is estimated at approximately 25 W.

At 25 W steady-state draw from the 84 Wh LiFePO4 auxiliary battery, the A.U.R.A. system has an estimated continuous runtime of approximately 3.4 hours before the battery is depleted. Even at the 38.7 W peak draw, the battery can sustain full-load operation for over 2 hours. Both figures far exceed the duration of any planned demonstration run, providing substantial headroom for repeated testing cycles without needing to recharge between runs. The LiFePO4 battery's flat discharge curve, which holds close to 12V through most of its discharge range rather than sagging progressively like SLA chemistry, also ensures that the downstream voltage regulators receive a stable input throughout the entire operating session.

The separation of the A.U.R.A. auxiliary battery from the vehicle's drive battery is a fundamental architectural decision driven by this power budget constraint. Sharing a single

battery between the drive motors and the sensor electronics would create several risks: a motor stall event drawing high instantaneous current could cause a voltage sag on the shared bus that resets the MCUs mid-operation; a detection event that activates the interlock relay and cuts drive power could simultaneously disrupt sensor power if the two systems share a supply rail; and the total combined load of the drive system and A.U.R.A. electronics would exceed the safe discharge rate of the vehicle's original 4.5 Ah battery. By powering A.U.R.A. from its own dedicated LiFePO4 pack, these risks are eliminated and each system can be designed, tested, and fused independently.

The voltage regulation hierarchy within A.U.R.A. further enforces the power budget constraint at the component level. The 12V auxiliary bus is stepped down to 5V for the Raspberry Pi 5 and the LCD dashboard via LM2596-based buck converters, and to 3.3V for the ESP32-S3 MCUs, sensors, and logic-level signals via LDO regulators. Each regulator was selected with a current rating that includes margin above the maximum expected load for its rail, ensuring no single regulator operates near its thermal limit during peak demand. The NIR LED array drive current was also set at a level that balances illumination quality against power consumption, keeping the total LED array load below 6 W for both modules combined as a direct response to the overall battery capacity budget.

Table 3.2.10.3.C: Battery Component Comparison

Battery Option	Nominal Voltage	Capacity	Main Strength	Main Limitation
12V SLA Battery	12V	4.5–7Ah	Low cost, easy to source	Heavy, shorter lifespan
3S Li-ion Pack	11.1V	Varies	Lightweight, compact	Requires protection circuitry
12.8V LiFePO4 Pack	12.8V	Varies	Safe, long cycle life, stable output	Higher upfront cost

A LiFePO4 chemistry was also selected over a standard lithium-ion or SLA alternative for several reasons directly relevant to the A.U.R.A. operating environment. LiFePO4 batteries typically include a built-in Battery Management System (BMS) that provides hardware-level protection against overcharge, over-discharge, short circuit, and overheating, which is important in an underbody-adjacent application subject to vibration, heat from the ground surface, and the possibility of a wiring fault. LiFePO4 cells are also significantly lighter than SLA batteries of equivalent capacity, reducing the additional weight added to the demo vehicle. A typical 12V 7Ah LiFePO4 pack weighs approximately 2.2 lbs compared to roughly 5.5 lbs for an equivalent SLA unit. Its rated cycle life of over 4,000 cycles at full depth of discharge also means the battery will remain functional well beyond the demonstration and testing phase of Senior Design.

A power budget analysis was conducted to estimate the peak and average current demands of all A.U.R.A. subsystems operating from the auxiliary battery. The table below presents the estimated power draw for each major component at peak operating conditions.

3.2.11 Voltage Regulator

3.2.11.1 Voltage Regulator Technology Comparison

Voltage regulators are a necessary part of any circuit design as they provide the means to properly power different peripherals from a single power supply. In our application regulators will be used to properly power different sensors, microcontrollers, and other peripherals of the A.U.R.A system from our 12V power supply down to both 3.3V and 5V.

Linear voltage regulators are low cost and simple regulators used in low-power applications. They work using a feedback loop to control output voltage from an input voltage source. While simple, linear regulators tend to generate heat when regulating from high voltage down to low voltage.

Low dropout regulators are a type of linear regulator, in comparison to the standard linear regulator, LDOs have higher efficiency and best dropout voltage. LDOs are mostly seen in products powered by batteries. Additionally, LDO's produce less heat than regular linear regulators due to their higher efficiency.

Switching Regulators provide higher efficiency, lower thermal management, and can provide either a boost to input voltage or an inverse. Switching regulators are more complex and expensive in comparison to linear regulators but they provide better performance in higher voltage applications.

Table 3.2.11.1 Voltage Regulator Technology Comparison

Technology	Cost	Efficiency	Heat Generation	Dropout Voltage
Linear Regulator	Low	Low	Highest	Moderate
Low Dropout Regulator	Low	High	Low	Low
Switching Regulator	Highest	Highest	Lowest	No dropout; steps up or down

Switching regulators provide the highest efficiency and seem to be more reliable for our 12V battery power supply. Additionally, switching regulators requires the least amount of heat management, leading to less complexity withing the system.

3.2.11.2 Voltage Regulator Component Comparison

Table 3.2.11.2 Voltage Regulator Component Comparison

Component	Cost	V _{in} (max)	I _{out}	Switching Frequency
3.3V Regulator				
LM2679-ADJ (set to 3.3 V)	\$6.50	40 V	5 A	260 kHz
LM2596T-3.3/NOPB	\$5.50	40 V	3 A	150 kHz
LM2592HVT-3.3/NOPB	\$7.77	60 V	2 A	150 kHz
5V Regulator				
LM2576T-5.0/NOPB	\$3.54	40 V	3 A	52 kHz
LM2574N-5.0/NOPB	\$3.62	40 V	500 mA	52 kHz
LM2679-ADJ (set to 5.2 V)	\$6.50	40 V	5 A	260 kHz

The LM2575-3.3WT will be chosen as our 3.3V switching regulator. Our 3.3V regulators will be used to supply power to the microcontroller units. The LM2575-3.3WT provides the cheapest option while also supplying a proper amount of output current to support the ESP32 modules controlling different parts of the system. While the other two options may provide higher output currents, the amount they provide is unnecessary for our application; in addition to the unnecessary output current, the alternative options have higher switching frequencies which may result in less efficiency.

The LM2575-5.0WT will be chosen as the 5V regulator in our system. 5V regulators will be necessary to provide power to our different sensors, the LM2575-5.0WT provides a 1 A current, which will provide enough current to power both our hall-effect and proximity sensors.

The LM2679-ADJ was selected as the switching regulator for both the 3.3 V and 5 V rails, replacing the fixed-output LM2575 devices originally specified. Because it is adjustable, a single part number can supply every rail in the system: the output is set entirely by two external feedback resistors, so the same regulator design serves the 3.3 V microcontroller rail and the 5 V sensor / Raspberry Pi rail simply by changing those resistors. This reduces the bill-of-materials to one regulator family, simplifies sourcing and PCB layout, and keeps the feedback and inductor design consistent across rails.

The LM2679-ADJ also provides substantially more output-current headroom (up to 5 A) than the 1 A LM2575, which is important for the 5 V rail that powers the Raspberry Pi during peak current draw. Its higher 260 kHz switching frequency allows a smaller output

inductor and filter capacitors than the 52 kHz LM2575, and the 40 V input rating comfortably covers the 12.8 V LiFePO4 supply.

Note that the nominal 5 V rail is regulated to **5.2 V** at the converter output rather than 5.0 V. The extra 0.2 V compensates for the voltage drop across the downstream wiring, connectors, and the protection path, so that the Raspberry Pi and 5 V sensors still receive approximately 5 V at the load under full current.

Setting the Output Voltage

The LM2679-ADJ regulates its feedback (FB) pin to an internal reference of **Vref = 1.21 V**. The output voltage is set by a resistor divider from the output to FB to ground, with R1 connected between VOUT and FB (the top resistor) and R2 connected between FB and ground (the bottom resistor):

$$V_{out} = 1.21 \text{ V} \times (1 + R1 / R2)$$

In this design the bottom resistor R2 is held constant at 1.0 kΩ and only the top resistor R1 is changed between rails. With R2 fixed, R1 is found from:

$$R1 = R2 \times (V_{out} / 1.21 \text{ V} - 1)$$

Because Vout increases with the top resistor R1, the larger top resistor always produces the higher rail voltage. Using the resistors on the board:

- R1 = 1.8 kΩ (the smaller top resistor), with R2 = 1.0 kΩ, sets the 3.3 V rail.
- R1 = 3.3 kΩ (the larger top resistor), with R2 = 1.0 kΩ, sets the 5.2 V rail.
-

Worked values (common R2 = 1.0 kΩ):

Target Rail	R1 (VOUT → FB)	R2 (FB → GND)	Resulting Vout
3.3 V	1.8 kΩ	1.0 kΩ	≈ 3.39 V
5.2 V	3.3 kΩ	1.0 kΩ	≈ 5.20 V

Both rails share the same 1.0 kΩ bottom resistor and use standard resistor values. The 5.2 V rail lands on 5.20 V as intended; the 3.3 V rail settles at about 3.39 V, which is within the operating range of the 3.3 V devices, and can be trimmed to exactly 3.3 V by increasing R2 slightly to ≈ 1.05 kΩ. The key relationship is unchanged: the smaller top resistor (1.8 kΩ) sets the 3.3 V rail and the larger top resistor (3.3 kΩ) sets the 5.2 V rail.

3.3 Communication Protocol

3.3.1 Communication Technology Comparison

The A.U.R.A. system requires reliable communication between multiple subsystems, including the Raspberry Pi processing unit, ESP32-based microcontroller modules, camera

system, and user interface components. The communication system must support data transfer for control signals, system status updates, and coordination between distributed components. Because the system is integrated within a compact vehicle platform, the communication method must balance reliability, latency, and physical wiring complexity.

In Table 3.3.1, several communication technologies were evaluated for use in the A.U.R.A. system, including Wi-Fi, UART, SPI, and I2C. Each of these technologies offers different advantages depending on the application.

Wi-Fi communication, defined by the IEEE 802.11 standard [14], enables wireless data transfer between devices using standard networking protocols. Both the Raspberry Pi and ESP32 modules include built-in Wi-Fi capabilities, allowing communication without additional hardware modules. This approach reduces wiring complexity and allows flexible placement of components within the vehicle. However, Wi-Fi introduces higher latency compared to wired protocols and requires software-level handling of connections and data transmission.

UART (Universal Asynchronous Receiver/Transmitter) is a simple serial communication protocol commonly used for direct device-to-device communication [15]. UART provides reliable data transfer with low overhead and is easy to implement. However, it requires physical wiring between devices, which can increase system complexity and limit flexibility in component placement.

SPI (Serial Peripheral Interface) is a high-speed communication protocol that supports fast data transfer between a master and multiple slave devices [16]. SPI is commonly used for sensors and peripheral devices requiring rapid communication. However, it requires multiple signal lines and careful wiring, making it less suitable for distributed systems with physically separated components.

I2C (Inter-Integrated Circuit) is a multi-device communication protocol that allows multiple peripherals to share the same communication bus [17]. While it reduces wiring compared to SPI, it is limited in communication distance and is sensitive to electrical noise. This makes it less suitable for systems where components are distributed across a vehicle platform.

Table 3.3.1 Communication Technology Comparison

Technology	Speed	Wiring Complexity	Range	Trade-offs
Wi-Fi	Typically, ~1-100 Mbps depending on standard	Low (wireless; no physical data lines required)	Typically, ~20-100 m typical indoor range depending on environment	Selected: used for communication between distributed ESP32 modules, reducing wiring complexity and enabling flexible system layout

UART	Typically, up to ~115.2 kbps (can reach ~1 Mbps in optimized setups)	Medium (TX, RX, GND required; point-to-point wiring)	Typically, ~1-10 m reliable without additional hardware	Selected: used for communication between the Raspberry Pi and the main ESP32 controller due to its simplicity, low latency, and reliability for short-range communication. Requires direct wiring between devices.
SPI	Typically, up to ~10+ Mbps (device-dependent)	High (4+ lines; MOSI, MISO, SCLK, CS per device)	Very short (Typically, ~<1 m typical) due to signal integrity limits	Complex wiring for distributed components
I2C	Typically, 100 kbps (standard) 400 kbps (fast mode), up to ~1 Mbps	Low (2 lines; SDA, SCL shared across devices)	Short (Typically, ~<1m; sensitive to noise and capacitance)	Limited range and noise sensitivity

Wi-Fi was selected as the primary communication method for distributed system components due to its ability to support a distributed architecture without requiring extensive wiring. This is particularly important for integration within a vehicle platform, where minimizing physical connections simplifies installation and improves system modularity. Although Wi-Fi introduces additional latency compared to wired protocols, the system does not require strict real-time communication, making this tradeoff acceptable. In addition to wireless communication, a direct wired interface is used between the Raspberry Pi and the main ESP32 controller. UART was selected for this connection due to its simplicity, low latency, and reliability for short-range communication between closely integrated components.

UART was originally selected for communication between the Raspberry Pi 5 and the main ESP32 controller. During integration, the Raspberry Pi did not reliably read the ESP32 through the USB connection. The operational link was therefore changed to Wi-Fi. The final implementation sends short safety commands from the Pi 5 to the ESP32 over the network and receives an acknowledgment for each command.

3.3.2 Communication Module Comparison

After evaluating communication technologies, Wi-Fi was selected for distributed communication while UART was selected for short-range wired communication between closely integrated components.

ESP32-based microcontroller modules include integrated Wi-Fi capabilities and support multiple communication interfaces, making them suitable for embedded networking applications [18]. These modules enable wireless communication between distributed system components and can also perform hardware control tasks.

The NRF24L01 is a low-power wireless communication module designed for short-range data transmission [19]. While it provides low latency, it requires additional hardware and does not support standard networking protocols.

Wired UART communication provides reliable data transfer but requires physical connections between devices, which can increase wiring complexity in distributed systems. In the final system design, the ESP32 functions both as a local controller interfacing directly with the Raspberry Pi and as a communication node for distributed modules, enabling a hybrid architecture that combines wired and wireless communication methods. The key differences between these communication modules are summarized in Table 3.3.2.

Table 3.3.2 Communication Module Comparison

Module	Integration	Range	Complexity	Notes
ESP32 Wi-Fi	High (integrated MCU + Wi-Fi + peripherals in single chip)	Typically, ~20-100 m depending on environment and antenna	Low (minimal external hardware required)	Selected: integrates wireless communication and hardware control in a single module, reducing system complexity
NRF24L01	Medium (requires external MCU + SPI interface)	Typically, ~10-100 m typical (can extend with PA/LNA versions)	Medium (requires additional RF configuration and libraries)	Requires additional hardware
UART (Wired)	High (direct hardware interface available on both devices)	Typically, ~1-10 m typical without signal conditioning	Low (simple TX/Rx communication, minimal configuration)	Selected: used for short-range communication between the Raspberry Pi and main ESP32 controller; limited flexibility due to wiring

ESP32-based modules were selected because they provide built-in wireless communication capabilities and simplify overall system integration.

The ESP32's integrated Wi-Fi capability became the final runtime communication interface with the Raspberry Pi 5. This avoided depending on the USB/UART connection

that was not being read reliably by the Pi. UART capability remains available for firmware flashing, debugging, and diagnostic output, but it is not the normal runtime command path.

3.3.3 Data Flow and Message Structure

The A.U.R.A. system utilizes a structured data flow to transmit image data and detection results between system components. Image data is captured by the IMX178M monochrome camera module and transmitted to the processing unit through the CSI interface.

Due to the higher resolution of the IMX178M sensor, captured images may be resized prior to processing to maintain real-time performance on the embedded platform. Additionally, the use of a monochrome sensor improves image contrast under near-infrared illumination, enhancing feature extraction for detection.

Once captured, the image data is passed into the processing pipeline where the detection algorithm analyzes the scene. The system processes each frame and determines whether a hazard is present based on the trained detection model.

The output of the detection process is formatted into a structured message that includes the detection status, classification result, and confidence level. This message is transmitted to the control subsystem, which determines the appropriate system response.

If a hazard is detected, the system activates the deterrent mechanisms and prevents vehicle operation. If no hazard is present, a “clear” signal is issued, allowing normal system behavior.

This data flow structure ensures efficient communication between sensing, processing, and control components while maintaining system reliability and scalability.

In the final system, the front and rear Raspberry Pi 4B nodes send MJPEG camera streams to the Raspberry Pi 5 over Wi-Fi. The Pi 5 processes selected frames and sends START, HAZARD, CLEAR, or RESET commands to the ESP32 over a separate Wi-Fi connection. The ESP32 returns ACK_START, ACK_HAZARD, ACK_CLEAR, or ACK_RESET to confirm the requested action.

3.4 Software

3.4.1 Software Platform

3.4.1.1 Software Platform Technology Comparison

The A.U.R.A. system requires a software platform capable of managing multiple concurrent tasks, including real-time image acquisition from near-infrared camera modules, execution of object detection algorithms, communication with microcontroller subsystems, and display output to a user interface. Because these tasks must operate simultaneously and reliably, the software platform must support multitasking, hardware abstraction, and integration with high-level libraries. In Table 3.4.1.1, three primary

software platform approaches were evaluated: bare-metal firmware, real-time operating systems (RTOS), and embedded Linux-based operating systems.

Bare-metal firmware executes directly on the processor without an operating system. This approach provides maximum control over hardware and minimal overhead, making it suitable for time-critical and resource-constrained applications. Bare-metal systems are commonly used in microcontrollers for tasks such as motor control and sensor polling. However, implementing a full image-processing pipeline using bare-metal firmware would require manually handling memory management, device drivers, and communication protocols. Additionally, integrating machine-learning frameworks in a bare-metal environment is extremely complex and impractical for rapid prototyping. Therefore, while bare-metal firmware offers excellent performance, it lacks the flexibility required for the A.U.R.A. system.

Real-time operating systems (RTOS) provide deterministic scheduling and improved task management compared to bare-metal systems. RTOS platforms such as FreeRTOS [20] are widely used in embedded systems that require predictable timing behavior. These systems allow multiple tasks to run concurrently with defined priorities, making them suitable for control-oriented applications. However, RTOS environments typically lack native support for advanced libraries such as OpenCV and deep-learning frameworks. While it is possible to port such libraries to an RTOS, doing so requires significant development effort and may not provide stable performance on resource-limited hardware. Because the A.U.R.A. system requires both high-level processing and rapid development, RTOS solutions were determined to be insufficient.

Embedded Linux operating systems provide a full-featured environment that includes multitasking, hardware drivers, networking support, and compatibility with high-level programming languages such as Python and C++ [21]. Linux-based systems are widely used in single-board computers, including the Raspberry Pi, and are well-supported by open-source communities. The availability of prebuilt libraries for image processing and machine learning significantly reduces development time. Additionally, Linux supports communication protocols such as UART, SPI, I2C, and Wi-Fi, which are required for interaction with ESP32 microcontrollers and other subsystems. Although Linux introduces additional overhead compared to RTOS, its flexibility and software ecosystem make it the most suitable platform for A.U.R.A.

Table 3.4.1.1 Software Platform Technology Comparison

Technology	Complexity	Library Support	Real-Time Capability	Trade-offs for A.U.R.A.
Bare-Metal Firmware	Low (direct hardware control, minimal abstraction)	Very limited (no OS-level libraries; must implement drivers manually)	Excellent (deterministic execution, no scheduling overhead)	Difficult to support camera processing and AI integration

RTOS	Medium (task scheduling, synchronization mechanisms required)	Limited (basic middleware, fewer high-level libraries)	Good (real-time task scheduling with predictable latency)	Suitable for control tasks but lacks high-level library support
Embedded Linux	Higher (full OS configuration and resource management required)	Extensive (supports OpenCV, Python, ML frameworks, networking stacks)	Moderate (not strictly real-time without patches like PREEMPT_RT)	Selected: provides full operating system support and compatibility with required image processing and machine learning libraries

Based on this comparison, Embedded Linux was selected because it enables integration of camera drivers, networking, and high-level programming libraries within a single environment. This allows the A.U.R.A. system to perform real-time image processing and machine-learning inference while maintaining communication with hardware subsystems.

3.4.1.2 Software Platform Part Comparison

After selecting embedded Linux as the software platform, several operating system distributions were evaluated for compatibility with the Raspberry Pi hardware. In Table 3.4.1.2, the options considered include Raspberry Pi OS, Ubuntu Linux, and Yocto Linux.

Raspberry Pi OS is specifically designed for Raspberry Pi hardware and includes optimized drivers for peripherals such as the CSI camera interface, GPIO pins, and communication protocols [22]. It also provides preinstalled support for Python and commonly used development tools, allowing rapid deployment of applications. Because of its tight integration with the hardware, Raspberry Pi OS minimizes compatibility issues and simplifies system configuration.

Ubuntu Linux is a general-purpose operating system that offers a robust package management system and extensive library support. It is widely used in both desktop and embedded environments. However, Ubuntu typically requires more system resources than Raspberry Pi OS, including increased memory usage and storage requirements. Additionally, configuring hardware interfaces such as the camera may require additional setup steps.

Yocto Linux provides a customizable environment where developers can build a minimal operating system tailored to specific applications. While this allows for optimized performance, the complexity of configuring and maintaining a Yocto build system is significantly higher. For a senior design project with limited development time, the overhead associated with Yocto makes it less practical.

Table 3.4.1.2 Software Platform Part Comparison

OS	Hardware Support	Library Availability	Setup Complexity	Notes
Raspberry Pi OS	Excellent (native support for Raspberry Pi hardware, GPIO, camera modules)	Extensive (pre-installed support for Python, OpenCV, and system libraries)	Low (plug-and-play setup with official images)	Selected: provides native support for Raspberry Pi hardware and simplifies integration of camera and communication interfaces
Ubuntu	Good (broad hardware compatibility but less optimized for Pi-specific peripherals)	Extensive (full Linux ecosystem, package manager support)	Medium (requires additional configuration for embedded use)	Higher resource usage and additional configuration required
Yocto	Customizable (can be tailored to specific hardware platforms)	Moderate (depends on build configuration and included packages)	High (requires building custom Linux images and toolchains)	Too complex for rapid prototype

Raspberry Pi OS was selected due to its direct compatibility with the Raspberry Pi processing platform and its built-in support for hardware interfaces such as the CSI camera and GPIO. This reduces setup complexity and allows faster development of the A.U.R.A. system.

3.4.2 Image Processing Software

3.4.2.1 Image Processing Technology Comparison

The image-processing component of the A.U.R.A. system is responsible for handling visual data captured by the IMX178M monochrome near-infrared camera modules. These cameras provide high-resolution images through the CSI interface, allowing efficient transfer of image data to the Raspberry Pi. Due to the higher resolution and monochrome sensor configuration of the IMX178M, the software must efficiently process these images in real time, performing operations such as resizing, filtering, and preparing the data for object detection. In Table 3.4.2.1, three image-processing libraries were evaluated: OpenCV, Pillow (PIL), and scikit-image.

OpenCV is a comprehensive computer-vision library designed for real-time applications [23]. It provides a wide range of functions for image filtering, feature detection, and camera interfacing. OpenCV supports direct access to camera streams and includes optimized implementations of common image-processing operations. Because it is widely used in robotics and autonomous systems, it has extensive documentation and community support.

Additionally, OpenCV is well-suited for handling high-resolution grayscale image streams, making it compatible with the IMX178M sensor output.

Pillow is a lightweight library that focuses on basic image manipulation tasks [24] such as resizing and format conversion. While it is simple to use, it lacks advanced computer-vision capabilities and is not optimized for real-time processing. As a result, it is not suitable for applications that require continuous image analysis.

Scikit-image offers advanced algorithms for image processing and analysis, including segmentation and edge detection [25]. However, it is primarily designed for scientific applications and offline processing. Its performance on embedded platforms is limited compared to OpenCV, especially when handling continuous video streams.

Table 3.4.2.1 Image Processing Technology Comparison

Library	Real-Time Support	Feature Set	Ease of Use	Trade-offs
OpenCV	Excellent (optimized C/C++ backend with real-time processing support)	Very high (image filtering, object detection, ML integration, video processing)	Moderate (large API, learning curve required)	Selected: supports real-time image acquisition and provides required computer vision functions for detection tasks
Pillow	Limited (not optimized for real-time video processing)	Low (basic image manipulation: resize, format conversion)	Easy (simple Python interface)	Insufficient for real-time vision tasks
Scikit-image	Moderate (better than pillow but not optimized for real-time pipelines)	High (advanced image analysis and scientific processing tools)	Moderate (designed for research workflows)	Slower performance on embedded hardware

OpenCV was selected because it offers a comprehensive set of computer vision tools and efficient processing capabilities for real-time applications. Its compatibility with the IMX178 camera modules and integration with machine-learning frameworks make it suitable for the image processing pipeline of the A.U.R.A. system.

3.4.2.2 Image Processing Library Selection

The OpenCV library can be used with both Python and C++ programming interfaces. Each option offers different advantages.

Python OpenCV provides a high-level interface that simplifies development and allows rapid prototyping. It integrates seamlessly with machine-learning libraries and reduces the complexity of implementing detection algorithms.

C++ OpenCV offers improved performance and lower execution overhead, making it suitable for performance-critical applications. However, development time is longer, and integration with Python-based frameworks can be more complex. A comparison of these library interfaces is provided in Table 3.4.2.2.

Table 3.4.2.2 Library Interface Comparison

Interface	Speed	Development Time	ML Integration	Notes
Python OpenCV	Typically, ~10-30 FPS on Raspberry Pi depending on model complexity	Low (rapid, prototyping, simple syntax)	Excellent (easy integration with TensorFlow, PyTorch, YOLO)	Selected: enables rapid development and integration with Python-based detection frameworks
C++ OpenCV	High (near real-time performance with lower overhead than Python)	High (more complex memory management and build process)	Moderate (requires additional integration effort for ML frameworks)	Higher performance but increased complexity

Python OpenCV was selected because it simplifies development and allows direct integration with the selected detection algorithm. Although C++ offers higher performance, Python provides sufficient speed for real-time processing on the Raspberry Pi while reducing implementation complexity.

3.4.3 Detection Algorithm

3.4.3.1 Detection Technology Comparison

The detection algorithm is a critical component of the A.U.R.A. system, responsible for identifying animals beneath the vehicle using images captured by the IMX178M monochrome near-infrared cameras. The algorithm must operate in real time, maintain high accuracy, and function reliably under varying lighting conditions. The use of a monochrome sensor improves image contrast under near-infrared illumination, enhancing the visibility of animal features and supporting more reliable detection. In Table 3.4.3.1, three detection approaches were evaluated: Haar Cascade classifiers, MobileNet SSD, and YOLO.

Haar Cascade classifiers use predefined features to detect objects. They are computationally efficient but lack robustness when dealing with complex environments or

varying lighting conditions. This limitation makes them unsuitable for detecting animals in outdoor settings.

MobileNet SSD is a lightweight neural-network model designed for embedded systems [26]. It provides improved accuracy compared to traditional methods while maintaining reasonable performance. However, it may struggle with small objects or low-contrast images, particularly in scenarios with limited feature distinction.

YOLO is a deep-learning-based object detection algorithm that processes the entire image in a single pass [27]. It provides high accuracy and can detect multiple objects simultaneously. YOLO models have been optimized for embedded platforms, making them suitable for real-time applications. Additionally, YOLO performs well on high-resolution grayscale image inputs, making it compatible with the IMX178M sensor output.

Table 3.4.3.1 Detection Technology Comparison

Algorithm	Accuracy	Speed	Hardware Requirement	Trade-offs
Haar Cascade	Low (high false positives; struggles with complex backgrounds and lighting variation)	High (Typically, ~30-60 FPS on CPU due to simple feature-based detection)	Low (runs efficiently on CPU without GPU acceleration)	Not robust enough for this application
MobileNet SSD	Medium (Typically, ~60-75% mAP depending on dataset and configuration)	Medium (Typically, ~10-25 FPS on Raspberry Pi with optimization)	Medium (requires more memory and processing than classical methods)	Moderate accuracy
YOLO	High (Typically, ~70-90% + mAP depending on model version and training)	Medium (Typically, ~10-30 FPS on embedded systems like Raspberry Pi with optimization)	Medium (requires GPU or optimized CPU inference for real-time performance)	Selected: provides high detection accuracy with real-time performance on embedded hardware

YOLO was selected because it can detect multiple objects within a single frame while maintaining real-time performance. Its ability to process the entire image at once improves detection reliability compared to traditional methods, making it suitable for identifying animals in near-infrared images.

The original research evaluated multiple YOLO versions. During implementation, the selected Roboflow workflow did not offer YOLOv8 and provided YOLOv11 Nano as the

available YOLOv11 model option. The final A.U.R.A. model was therefore trained and deployed as YOLOv11 Nano.

3.4.3.2 Detection Model Comparison

In Table 3.4.3.2, three YOLO model versions were evaluated: YOLOv5, YOLOv8, and Tiny-YOLO. YOLOv5 provides high detection accuracy but requires more processing power. YOLOv8 improves efficiency and provides better performance on embedded systems [27]. Tiny-YOLO requires less computation but has lower accuracy, which may reduce detection reliability.

Table 3.4.3.2 Detection Model Comparison

Model	Speed	Accuracy	Resource Usage	Notes
YOLOv5	Typically, ~10-25 FPS on Raspberry Pi (model size dependent)	High (Typically, ~85-90% mAP for standard models)	Medium (moderate RAM and CPU/GPU usage depending on variant)	Higher resource usage
YOLOv8	Typically, ~15-30 FPS on Raspberry Pi with optimization (improved inference efficiency)	High (Typically, ~85-92% mAP with improved architecture)	Medium-Low (better efficiency and lower computational overhead than YOLOv5)	Improved efficiency and accuracy
Tiny-YOLO	Typically, ~20-40 FPS on Raspberry Pi (lightweight model)	Medium (Typically, ~60-75% mAP; reduced detection precision)	Low (optimized for low-power and embedded devices)	Lower accuracy
YOLOv11	Typically, ~15-30 FPS	High (~90%+ mAP depending on training)	Medium-Low	Selected: provides improved efficiency and accuracy compared to earlier YOLO variants, making it suitable for real-time embedded applications

While earlier YOLO versions such as YOLOv5 and YOLOv8 provide strong performance, newer YOLO architectures offer improved efficiency and detection accuracy. Based on these advancements, YOLOv11 was selected for the A.U.R.A. system due to its ability to

provide a balance between real-time performance and detection reliability on embedded hardware.

The table records the models researched during Senior Design I. The completed implementation uses YOLOv11 Nano because YOLOv8 was not available through the selected Roboflow workflow and Nano was the YOLOv11 model version offered.

In addition to model selection, Table 3.4.3.3 presents different development workflows were considered for dataset preparation and training. Traditional approaches involve manual annotation tools and separate training environments, which increase development complexity. Alternatively, integrated platforms such as Roboflow provide a unified environment for dataset labeling, preprocessing, augmentation, and model training.

Table 3.4.3.3 Dataset and Training Workflow Comparison

Platform	Integration	Ease of Use	Workflow Complexity	Notes
LabelImg	Low	Medium	High (separate tools)	Requires separate annotation and training pipeline
Roboflow	High	High	Low	Selected: integrated labeling, augmentation, and training platform

Roboflow was selected because it streamlines the development process and reduces implementation overhead while maintaining compatibility with YOLO-based models. This enables efficient dataset management and improves the overall system development workflow.

3.4.4 Programming Language and Development Framework

3.4.4.1 Programming Language Technology Comparison

The A.U.R.A. system requires a programming environment capable of supporting real-time image processing, communication with embedded hardware, and integration with machine-learning frameworks. The programming language must also allow rapid development and debugging, as the system involves multiple interacting components including camera input, detection algorithms, and control logic. In Table 3.4.4.1, three programming approaches were evaluated: Python, C++, and Robot Operating System (ROS)-based development.

Python is a high-level programming language widely used in data processing, machine learning, and rapid prototyping. It provides a simple syntax and extensive library support,

including integration with OpenCV and deep-learning frameworks such as YOLO. Python is commonly used on embedded Linux platforms like the Raspberry Pi due to its ease of use and flexibility. Although Python has slower execution speed compared to compiled languages, it provides sufficient performance for real-time applications when combined with optimized libraries. This approach also allows easier debugging and testing during development, which is critical for iterative refinement of the prototype system.

C++ is a compiled programming language that provides high performance and low-level memory control. It is commonly used in performance-critical applications such as robotics and embedded systems. C++ offers direct access to hardware interfaces and allows efficient execution of image-processing algorithms. However, development in C++ is more complex and requires longer implementation time, especially when integrating machine-learning frameworks that are often designed for Python environments.

The Robot Operating System (ROS) is a middleware framework designed for robotics applications. ROS provides tools for communication between software modules, sensor integration, and system coordination. While ROS offers strong modularity and scalability, it introduces additional complexity and overhead. Implementing ROS on a Raspberry Pi for a prototype system increases setup difficulty and may not provide significant advantages for the relatively focused scope of the A.U.R.A. system.

Table 3.4.4.1 Programming Language Technology Comparison

Technology	Performance	Ease of Development	Library Support	Trade-offs
Python	Typically, ~10-30 FPS for computer vision tasks on Raspberry Pi (depending on model and optimization)	High (simple syntax, rapid prototyping, minimal boilerplate)	Extensive (OpenCV, TensorFlow, PyTorch, NumPy, etc)	Selected: balances simplicity and functionality
C++	Typically, ~20-60 FPS for equivalent CV tasks due to compiled execution and lower overhead	Low (manual memory management, more complex syntax, and debugging)	High (native OpenCV support, lower-level hardware access)	More complex development
ROS	Typically, ~10-30 FPS depending on node architecture and communication overhead	Medium (requires understanding of nodes, topics, and system structure)	High (robotics libraries, sensor integration, distributed control tools)	Adds system complexity

Python was selected because it provides a balance between development speed and functionality while supporting the libraries required for image processing and object detection.

3.4.4.2 Development Framework Comparison

After selecting Python as the programming language, different development frameworks and environments were evaluated to determine how the software would be structured and executed. In Table 3.4.4.2, the options considered include standard Python scripting, ROS-based Python integration, and modular Python frameworks using libraries such as Flask or socket-based communication.

Standard Python scripting provides a straightforward approach where the entire application is implemented as a series of modules executed within a single runtime environment. This approach allows direct control of camera input, image processing, and communication with hardware components. It also simplifies debugging and reduces system overhead.

ROS-based Python integration allows the system to be divided into multiple nodes that communicate through a messaging system. This approach improves modularity and scalability, but it introduces additional configuration requirements and increases system complexity. For a prototype system, this overhead is not necessary.

Modular Python frameworks using sockets or lightweight communication libraries allow different parts of the system to communicate without requiring a full middleware framework. This approach provides a balance between modularity and simplicity while maintaining low overhead.

Table 3.4.4.2 Development Framework Comparison

Framework	Complexity	Flexibility	Performance	Notes
Standard Python	Low (single-process architecture, minimal setup required)	Moderate (limited built-in support for distributed systems)	Typically, ~10-30 FPS depending on processing pipeline	Selected: simple and efficient
ROS Python	High (requires ROS environment, node configuration, message passing setup)	High (supports distributed systems, modular nodes, and scalability)	Typically, ~5-20 FPS due to inter-process communication overhead	Too complex for prototype
Modular Python (Sockets)	Medium (requires manual implementation of communication interfaces)	High (customizable architecture for distributed components)	Typically, ~10-30 FPS with efficient socket communication	Good alternative

Standard Python scripting was selected because it provides sufficient flexibility and performance while minimizing system complexity. This approach allows the A.U.R.A. system to integrate image processing, detection, and communication within a single, manageable software architecture.

Chapter 4 - Standards and Design Constraints

4.1 Industrial Standards

The A.U.R.A. system is designed in accordance with relevant industrial standards to ensure compatibility, reliability, and proper system integration. These standards define the communication protocols and hardware interfaces used within the system, allowing components from different manufacturers to operate together effectively. By adhering to established standards, the system benefits from widely supported technologies, reduced development complexity, and improved scalability.

The primary standards relevant to the A.U.R.A. system include the IEEE 802.11 wireless communication standard and the MIPI CSI-2 camera interface standard. These standards directly influence the design and operation of the system's communication and image acquisition subsystems.

4.1.1 Optical Industrial Standards

4.1.1.1 IEC 62471

The A.U.R.A. system incorporates active near-infrared illumination to enable imaging beneath the vehicle in low-light and nighttime conditions. As a result, the system must be evaluated in the context of photobiological safety, particularly with respect to human exposure to optical radiation. IEC 62471 defines exposure limits and risk classifications for lamps and lamp systems across ultraviolet, visible, and infrared wavelengths. Although the illumination used in this system operates primarily in the near-infrared range, the standard remains directly applicable because the human eye does not respond strongly to infrared radiation, which can lead to exposure without immediate visual awareness.

The illumination subsystem uses high-power 850 nm infrared LEDs arranged around the optical aperture. At this wavelength, a faint red emission may be visible under certain conditions, but the majority of the emitted energy lies outside the peak sensitivity range of the human visual system. IEC 62471 evaluates hazards based on radiometric quantities such as irradiance and radiance rather than perceived brightness, which is critical in this case since visual perception is not a reliable indicator of exposure level. Typical 850 nm emitters used for imaging applications exhibit radiant intensities on the order of hundreds of milliwatts per steradian, and when multiple emitters are combined in an array, the total irradiance in the near field can increase significantly.

The system geometry plays an important role in maintaining safe operation within the limits defined by the standard. The emitters are oriented downward toward the ground plane rather than outward toward the user, and the illumination is distributed across a wide area rather than concentrated into a narrow beam. This reduces peak irradiance at any single point and limits the likelihood of direct eye exposure. In addition, the short working distance of the system ensures that illumination is confined primarily to the region beneath the vehicle rather than projecting outward into the surrounding environment.

The selection of LEDs instead of laser-based sources is also influenced by IEC 62471 considerations. LEDs produce incoherent light with lower radiance compared to laser diodes, which significantly reduces the potential for retinal hazard. This allows the system to achieve sufficient illumination for imaging while maintaining a safer emission profile. The standard therefore directly informs both the choice of illumination technology and the physical configuration of the emitter array.

By designing the illumination subsystem in accordance with IEC 62471, the A.U.R.A. system ensures that near-infrared imaging can be performed without introducing unnecessary risk to users or bystanders. This is particularly important for a system intended to operate in close proximity to people in real-world environments.

4.1.1.2 ISO 16750

The placement of the A.U.R.A. system beneath a vehicle exposes it to environmental conditions that are significantly more demanding than typical imaging systems. ISO 16750

defines environmental requirements and test conditions for electrical and electronic equipment used in road vehicles, including temperature extremes, mechanical stress, moisture exposure, and contamination. While the current system is a prototype and not subject to formal automotive certification, the conditions described in this standard provide a realistic framework for evaluating design robustness.

Temperature variation is one of the most critical factors affecting optical performance. Underbody components can experience elevated temperatures due to solar loading, heat from the road surface, and proximity to drivetrain or battery systems. In practical terms, temperatures beneath a vehicle can exceed 60 to 80 °C during operation. These conditions directly impact material selection, as thermal expansion can alter optical alignment and degrade imaging performance.

The use of N-BK7 glass for the cylindrical anamorphic elements addresses this constraint by providing stable refractive properties over a wide temperature range. With a coefficient of thermal expansion on the order of $7.1 \times 10^{-6}/\text{K}$, dimensional changes remain small enough to preserve optical alignment. In contrast, polymer-based optical materials such as PMMA exhibit significantly higher thermal expansion, which can introduce measurable distortion or misalignment under the same conditions. This makes glass a more appropriate choice for maintaining consistent optical behavior in an automotive environment.

The housing material is also influenced by this standard. ASA is selected for the enclosure due to its improved thermal resistance and environmental durability compared to materials such as PLA. ASA maintains structural integrity at temperatures approaching 100 °C and offers better resistance to UV exposure and moisture, both of which are relevant for an underbody-mounted system.

In addition to temperature, ISO 16750 highlights the importance of protection against dust, water, and mechanical vibration. The inclusion of a protective optical window helps isolate internal components from contaminants, while the rigid enclosure design supports alignment stability under vibration. These considerations ensure that the optical system remains functional and reliable despite exposure to real-world environmental conditions.

4.1.1.3 ISO 12233

The effectiveness of the A.U.R.A. system depends on its ability to resolve meaningful detail within the underbody region. ISO 12233 provides standardized methods for evaluating spatial resolution and image quality in electronic imaging systems. Although originally developed for visible imaging, the same principles apply to near-infrared systems, particularly in terms of modulation transfer function (MTF) and spatial frequency response.

In this system, resolution is not solely determined by sensor pixel count. The optical architecture intentionally modifies the horizontal field through anamorphic compression, which makes it necessary to evaluate how well spatial detail is preserved after remapping. The IMX178 sensor used in the system has a pixel size of 2.4 μm , corresponding to a theoretical Nyquist frequency of roughly 208 line pairs per millimeter at the sensor plane.

However, the actual system resolution is limited by optical aberrations, alignment tolerances, and the effects of the cylindrical lens system.

ISO 12233 provides a structured method for quantifying these effects using standardized test patterns and analysis techniques. By measuring resolution in terms of spatial frequency response, the system can be evaluated in a consistent and repeatable manner. This is particularly important for verifying that the optical design preserves enough detail to distinguish between relevant features within the scene. Furthermore, because the anamorphic relay creates an asymmetrical resolution profile, horizontal and vertical MTF must be evaluated independently to ensure texture detail is maintained during the digital re-expansion process.

The requirement for resolution in this application is driven by the need to detect subtle differences in appearance rather than large, well-defined shapes. Objects beneath the vehicle may not present a clear silhouette, and in many cases the most useful information comes from local variations in reflectance and texture. Fur, exposed skin, rubber, and ground surfaces interact differently with near-infrared illumination, and these differences must be preserved in the captured image for reliable detection.

By applying ISO 12233 principles, the system can be evaluated to ensure that sufficient spatial detail is maintained across the field of view, particularly in regions where detection is most critical.

4.1.2 Wireless Communication Standards

4.1.2.1 IEEE 802.11

The A.U.R.A. system utilizes wireless communication between distributed ESP32-based microcontroller modules, while the Raspberry Pi communicates with the main ESP32 controller through a direct wired interface. Wi-Fi communication is implemented using the IEEE 802.11 standard, which defines the physical and media access control (MAC) layers for wireless local area networks (WLANs). This standard enables reliable data transmission between distributed devices within the system.

The IEEE 802.11 standard provides a framework for wireless communication using radio frequency signals, supporting features such as device discovery, authentication, packet-based data transfer, and error detection. These capabilities are essential for ensuring consistent and reliable communication within the A.U.R.A. system. ESP32 modules and the Raspberry Pi support IEEE 802.11 networking, allowing the system to leverage standard wireless communication for distributed module coordination without requiring additional networking hardware.

In the A.U.R.A. system, Wi-Fi is used to coordinate distributed ESP32 modules throughout the vehicle platform. This configuration allows physically separated components to communicate wirelessly without extensive wiring, improving system modularity and maintainability.

The IEEE 802.11 standard also includes mechanisms for handling interference, retransmitting lost packets, and maintaining data integrity. These features are important in real-world environments where wireless signals may be affected by obstacles, noise, or other sources of interference. By adhering to this standard, the A.U.R.A. system benefits from built-in reliability and robustness in communication.

However, the use of IEEE 802.11 introduces certain design considerations. Wireless communication generally has higher latency compared to wired protocols such as SPI or UART. Additionally, connection management and network configuration must be handled through software. Despite these tradeoffs, the flexibility, scalability, and ease of integration provided by IEEE 802.11 make it the most suitable communication standard for distributed components within the A.U.R.A. system.

Overall, compliance with the IEEE 802.11 standard ensures interoperability between system components and supports a modular system architecture, which is critical for both development and future scalability.

In addition to basic communication functionality, the IEEE 802.11 standard supports various operating modes, including infrastructure and ad-hoc networking. For the A.U.R.A. system, infrastructure mode can be used to establish a centralized communication structure where the Raspberry Pi acts as the primary node. This enables easier coordination between multiple ESP32 modules and allows for structured data exchange.

The standard also supports different frequency bands and data rates, which influence system performance and reliability. Although high data throughput is not the primary requirement of the A.U.R.A. system, stable and consistent communication is essential. Therefore, the system design prioritizes reliability over maximum bandwidth, ensuring that control signals and detection results are transmitted without interruption.

The final implementation also uses IEEE 802.11 Wi-Fi between the Raspberry Pi 5 and ESP32 safety controller. This change was made because the Pi did not reliably read the ESP32 through the planned USB connection. Wi-Fi is therefore used for both the distributed camera streams and the smaller command-and-acknowledgment messages.

4.1.3 Camera Interface Standard

4.1.3.1 MIPI CSI-2

The A.U.R.A. system relies on IMX178M monochrome camera modules to capture near-infrared images for object detection. These camera modules interface with the Raspberry Pi using the MIPI CSI-2 (Camera Serial Interface) standard. MIPI CSI-2 is a high-speed serial communication standard developed by the MIPI Alliance, specifically designed for connecting image sensors to embedded processing systems.

The MIPI CSI-2 standard defines both the physical layer and communication protocol used to transmit image data. It utilizes differential signaling and supports multiple data lanes, allowing for high data throughput while maintaining low power consumption. This enables efficient transmission of high-resolution image data, which is essential for real-time image processing applications such as those required in the A.U.R.A. system. The higher resolution of the IMX178M sensor increases bandwidth demands, making the use of CSI-2 particularly important for maintaining performance.

One of the primary advantages of MIPI CSI-2 is its ability to support high frame rates with minimal latency. In the context of the A.U.R.A. system, this allows image data to be captured and processed quickly, enabling timely detection of animals beneath a vehicle. The CSI-based connection between the camera and the Raspberry Pi ensures that image data is transferred efficiently without the need for intermediate communication layers.

The Raspberry Pi includes a dedicated CSI connector and supports camera modules that comply with the MIPI CSI-2 standard. While native driver support is available for commonly used sensors such as the IMX178, integration of the IMX178M may require additional configuration or third-party drivers depending on the specific camera module implementation. Despite this, the use of a standardized interface simplifies system integration and allows the use of existing software libraries for image acquisition and processing. By leveraging the CSI-2 interface, the system can efficiently handle high-resolution image data while maintaining compatibility with embedded processing platforms.

However, the MIPI CSI-2 standard introduces certain constraints on system design. The interface requires short, high-quality connections to maintain signal integrity, limiting the distance between the camera and the processing unit. This influences the physical layout of the system, particularly in how cameras are positioned within the vehicle.

By adhering to the MIPI CSI-2 standard, the A.U.R.A. system ensures efficient, high-speed image acquisition while maintaining compatibility with widely supported camera hardware. This standard is critical to enabling reliable real-time image processing and overall system performance.

Another important aspect of the MIPI CSI-2 standard is its efficient use of bandwidth through packetized data transmission. This allows image data to be transmitted in an organized format that can be easily interpreted by the receiving processor. For the A.U.R.A. system, this ensures that captured image frames can be processed immediately without requiring complex data reconstruction.

Furthermore, the use of a standardized camera interface simplifies software development by enabling compatibility with existing drivers and libraries. This allows developers to focus on higher-level image processing and detection algorithms rather than low-level hardware communication. As a result, adherence to the MIPI CSI-2 standard contributes to both system performance and development efficiency.

4.1.4 Generic Standard on Printed Board Design

4.1.4.1 IPC-221A

The IPC 2221A standard covers the general requirements set for the design of printed boards. It covers critical information regarding things like electrical properties, thermal management, and component and assembly issues. Our A.U.R.A system will contain at least three PCBs connected through both wired and wireless means; our PCBs will take these topics into account when designing and assembling them.

IPC 2221A highlights key considerations to take when designing a PCB, for instance, it talks about power distribution considerations. One technique IPC 2221A recommends is the use of planes for power and ground distributions. By doing so, problems with electromagnetic interference can be reduced and avoided. Additionally, the standard specifies that power traces should also run closely to ground traces and to keep both kinds of traces as wide as possible. By doing so, problems with electromagnetic interference can be reduced and avoided. Another topic included in its electrical properties chapter is the circuit type considerations, in this section IPC 2221A covers considerations when designing layouts, these include polarity of components, the use of heat-sinks when needed, as well as consideration regarding ground connections depending on whether a circuit design is digital or analog; when a circuit is of the analog variety, ground connections should be considered last while in a digital circuit, signal connections should be considered last.[29]

This standard also contains a chapter on thermal management, in this section IPC 2221A considers temperature control and dissipation considerations and techniques. To properly dissipate heat from the PCB, heatsinks planes, fixtures, and similar should be considered as a method, in addition to this factors that affect the heatsinks should also be considered, for instance: heatsink mounting, material properties, size, and shape all should be considered to properly dissipate heat to ensure that component and board performance is not affected by thermal properties.[29]

In the standard's component and assembly issues section, it mentions the requirements components need to have considered for a board's design. These considerations include component placement and mounting. When it comes to orientation, the standard specifies that components should be placed parallel to board edges or parallel/perpendicular to other components; additionally, components should have enough spacing to not obstruct one another. In mounting components onto a board, certain techniques should be applied depending on the kind of component, for through-hole components consideration should be given to clearance given by other nearby components to attach leads to the board. Surface mounted components should have enough clearance for the use of a pick-and-place machine and enough to be able to see solder joints. [29]

By adhering to these and other standards described in IPC-2221A, our PCB design and assembly can function and perform at a high level in terms of performance and reliability

4.2 Design Constraints

Design constraints differ from standards in that they arise not from external governing bodies but from the inherent physical and operational realities of the project itself. They are non-negotiable boundaries that the engineering team must work within rather than choose to comply with. For A.U.R.A., the two most significant design constraints are the physical geometry of the vehicle undercarriage, which determines every dimensional decision in the system, and the power budget imposed by the demo vehicle's onboard battery, which limits the total electrical load the system can place on the vehicle without compromising its primary function.

4.2.1 System Hardware Design Constraint

4.2.1.1 Physical Size and Undercarriage Clearance

The most fundamental physical constraint governing the A.U.R.A. system is the geometry of the demonstration vehicle's undercarriage. The demo vehicle, a 12V small off-road electric kid car, has an underbody clearance ranging from 4.5 to 7.5 inches between the lowest structural cross-member and the ground surface. This narrow vertical window defines the absolute upper bound on the height of any component that must operate in the deployed state beneath the vehicle. Every mechanical, optical, and electronic assembly that extends below the vehicle floor must fit within this envelope without contacting the ground during stationary operation.

This constraint propagates through nearly every hardware decision in the system. The linear actuator selected for camera deployment, the Firgelli FA-150-S-12-6, was chosen specifically because its 6-inch stroke, combined with its retracted body length, keeps the total extended assembly within the available clearance. This was not a trivial constraint to satisfy. During early prototyping, the team acquired an 8-inch stroke actuator that appeared appropriate based on stroke length alone. When physically mounted beneath the vehicle, the total installed length of the unit including its retracted housing body exceeded 16 inches, making contact with the ground unavoidable. This experience made clear that the constraint is not just on stroke length but on the total collapsed-and-extended dimensional footprint of the entire actuator assembly in its mounting orientation.

The clearance constraint also directly affects the optical design. The NIR camera modules must be positioned at a standoff distance from the ground plane that provides the required field of view while remaining within the 6.5 to 8.5 inch clearance window. The anamorphic lens system, IR filter, objective lens, and LED illumination ring must all be housed within a protective enclosure that does not exceed the available vertical space when deployed. The 3D-printed housing will be designed with these constraints as primary inputs, and every internal component will be selected with package size as an explicit selection criterion alongside electrical performance.

The horizontal dimension presents a separate but equally important aspect of this constraint. The demo vehicle has a front wheel track of approximately 23 inches. The sensor modules must be positioned to achieve overlapping coverage of the full underbody width while mounted inboard of the wheel wells. This requires the camera housings to be

angled inward and the anamorphic lens system to compress the wide horizontal field into the sensor's rectangular active area. The system geometry was modeled geometrically to verify that the required 107-degree horizontal field of view could be achieved at the available standoff distance without the housing physically interfering with the wheel wells or suspension components.

When the vehicle is in motion and the actuators retract the camera modules into their protective housings, the entire system must fit flush against the vehicle floor within the IP65-rated enclosure. The retracted profile of each housing was constrained to be no taller than the available ground clearance at the mounting location, ensuring that the housing does not create a road-hazard protrusion during normal vehicle operation. The velocity-triggered retraction at 1 mph was established precisely to ensure the system transitions to its protected, low-profile configuration before the vehicle reaches speeds where road debris and aerodynamic forces could damage exposed hardware.

This constraint also has implications for scalability. The demo vehicle's 4.5 to 7.5 inch clearance is more restrictive than the 6.5-inch clearance of a full-sized Waymo vehicle. The optical and mechanical design choices made for the demo platform therefore represent a worst-case physical envelope. Validating the system within these tighter constraints provides confidence that the design approach is geometrically viable for full-scale deployment, where slightly more vertical clearance is available and the sensor modules would have more room to deploy without risk of ground contact.

4.2.1.2 Power Budget and Battery Capacity

The A.U.R.A. system operates in a vehicle environment where electrical noise and electromagnetic interference (EMI) are significant concerns due to the presence of motors, switching circuits, and high-current loads. These conditions can degrade signal quality, particularly for high-speed camera data and sensitive sensor readings. Therefore, the system must be designed to maintain signal integrity across all communication interfaces.

To address this constraint, differential signaling is used for high-speed data transmission (e.g., MIPI CSI-2), while digital communication protocols such as I2C are implemented with proper pull-up resistors and short trace lengths to minimize noise susceptibility. Power and signal lines are carefully separated to prevent coupling between high-current actuator paths and low-power logic circuitry. Additionally, grounding strategies and PCB layout considerations, such as minimizing loop areas and using proper decoupling techniques, are applied to reduce noise impact.

These constraints ensure reliable data acquisition and communication within the system, which is critical for accurate animal detection and safe operation of the vehicle.

4.2.2 Optical Design Constraints

4.2.2.1 Optical Geometry and Field Mapping Constraints

The optical design of the A.U.R.A. system is fundamentally constrained by the geometry of the underbody region and the physical properties of catalog optics. The system is required to capture a 55° horizontal field of view (FOV) within a total system track (working distance and camera body length) not exceeding 8 inches. This creates a scene with a highly asymmetric aspect ratio, which cannot be efficiently captured using conventional imaging approaches.

This 8-inch total track limit is a primary feasibility constraint addressing the mechanical clearance required for city parking and multi-level parking garages. As the system is designed to be active during vehicle start-up and up to a specific low-speed threshold, it must coexist with the vehicle's existing suspension and chassis envelope. A sensor package that protrudes significantly below the undercarriage would be vulnerable to "bottoming out" on steep parking garage ramps or hitting concrete parking stops and even other vehicles. By constraining the entire optical assembly and working distance to an 8-inch vertical envelope, the system remains protected within the vehicle's "safe zone," ensuring that the NIR sensors do not interfere with the vehicle's standard approach and departure angles during busy city operations.

This constraint necessitated a significant redesign from the initial proposal, which suggested an 8.5-inch working distance. Modeling revealed that at an 8.5-inch standoff, the 55° half-field would require a primary lens diameter exceeding 100 mm to prevent geometric clipping. Utilizing the current catalog 30 mm cylindrical lenses at that distance resulted in vignetting exceeding 90%, as the incident rays physically missed the lens apertures. To maintain the 55° half FOV while staying within the 8-inch safety envelope, the working distance was optimized to 65 mm (approx. 2.5 inch). This configuration ensures the wide-angle rays are captured by the LK1336L1-B primary element before they diverge beyond the lens's physical aperture.

This limits the feasible optical path length and requires a highly compact lens assembly. As a result, the design utilizes an anamorphic architecture with a 25 mm internal gap and a 10 mm expansion gap, achieving 81.9% sensor utilization of the IMX178. These geometric constraints define the structure of the optical system, ensuring that high-performance NIR animal detection can be achieved without compromising the vehicle's mechanical integrity in everyday urban environments.

While early simulation modeling indicated a heavily restricted 2.5 inch (65 mm) working clearance to completely mitigate peripheral lens vignetting, physical prototyping proved this extreme proximity was impossible to execute. The mechanical envelope required to pack the actuator assembly, its custom housing, and the electrical PCB control board demanded an unpitched, straight-on longitudinal perspective. To resolve this integration constraint, the physical system working distance was updated to a nominal 6.5 inch horizontal standoff distance. Maintaining a straight-on configuration at this distance

eliminates mechanical conflict with the actuator hardware while safely keeping the entire enclosure package within the protective 8 inch vertical safety envelope. The internal lens architecture has been dynamically re-optimized to account for this updated working clearance, re-spacing the physical gaps to a 39 mm separation between the cylindrical optics to secure system focus and manage field coverage.

4.2.2.2 Illumination and Reflectance Constraints

The imaging performance of the system is also constrained by how objects beneath the vehicle interact with illumination. The system must operate in low-light conditions and cannot rely on ambient lighting, which necessitates the use of active illumination. However, the illumination must be designed in a way that preserves meaningful differences in the captured image. Because the system has transitioned to a compact 65 mm working distance, the illumination subsystem is now constrained by "near-field" uniformity challenges. The incident angle of the NIR emitters must be carefully offset from the optical axis to prevent retro-reflection glare off the lens surfaces, which is a significantly higher risk at a 1-inch standoff than at the original 8.5-inch proposal.

Objects beneath the vehicle do not always present clear outlines. A small animal may occupy only a portion of the field or may be partially obscured. In these cases, detection depends on the ability to distinguish differences in reflectance and texture rather than relying solely on shape. Materials such as fur, exposed skin, rubber, and ground surfaces exhibit different reflectance characteristics under near-infrared illumination, and these differences must be preserved through the optical system. To address this constraint, the system utilizes a Monochrome IMX178 sensor. By eliminating the Bayer Filter Mosaic found in standard color sensors, we have bypassed the light-attenuation constraints that typically reduce NIR signal. This allows for a Quantum Efficiency (QE) of approximately 40% at 850nm, ensuring that subtle variations in fur reflectance are not lost to sensor noise.

This constraint influences both wavelength selection and illumination geometry. The use of 850 nm near-infrared illumination provides a balance between sensor sensitivity and controlled imaging conditions, allowing the system to capture consistent reflectance information. Unlike 940nm alternatives, 850nm provides a 2.6x increase in the IMX178's responsivity, which is critical for maintaining a high Signal-to-Noise Ratio (SNR) in the dark underbody environment. The illumination must also be distributed across the field to avoid strong hotspots or shadows that could obscure relevant detail.

These factors constrain the design of the illumination subsystem and directly impact the quality and usefulness of the captured images.

4.2.2.3 Optical Alignment and Stability Constraints

The performance of the optical system depends on maintaining precise alignment between components, particularly within the anamorphic lens assembly. Small deviations in spacing or orientation can affect field mapping and introduce distortions that degrade image quality. Specifically, the anamorphic relay is constrained by a 1-degree rotational tolerance between the -50 mm and +50 mm cylindrical axes. Any rotational deviation beyond this

limit introduces "smile" distortion and Keystone aberrations, which would prevent the captured 55° field from being correctly software-expanded for animal detection.

Thermal variation and mechanical stress present challenges to maintaining this alignment. Changes in temperature can cause expansion or contraction of materials, altering the spacing between optical elements. Mechanical vibrations from vehicle operation can also introduce small positional shifts over time. These effects are especially important in a system that relies on precise field compression in one axis. With the reduction of the working distance from 8.5 inches to 65 mm, the system's depth of field (DOF) has narrowed significantly. This introduces a strict constraint on mechanical stability; the back focal length (BFL) must remain stable within ± 0.1 mm to keep the image focused on the 5.45 mm NIR focal plane.

To manage these thermal constraints, the design must incorporate an integrated passive heatsink. The Sony IMX178 sensor and the active near infrared LEDs generate significant internal heat. If this heat is not dissipated, it leads to thermal noise or dark current that can obscure subtle fur textures. By thermally coupling the sensor to an aluminum heatsink integrated into the housing, the system ensures the 40 percent quantum efficiency is maintained. Furthermore, stabilizing the internal temperature minimizes the risk of thermal expansion induced focus shift. This ensures that the 81.9% sensor utilization remains consistent.

To address this constraint, materials with relatively low thermal expansion are used for optical components, and the housing is designed to provide structural rigidity. The optical elements are secured within the enclosure to minimize movement and maintain consistent alignment under varying conditions. By housing the 30 mm primary optics and the 1-inch secondary optics in a fixed-track assembly, the design ensures the 81.9% sensor utilization remains consistent despite the mechanical challenges of real-world edge cases and environmental constraints.

This constraint highlights the need for the optical system to balance precision with robustness. The design must maintain its intended performance not only under ideal conditions but also when subjected to the environmental and mechanical realities of its operating environment.

The transition to a straight on 6.5 inch longitudinal working clearance and the procurement of specific physical optics modified the system alignment margins and focal dimensions. The actual hardware implementation utilizes a minus 50 mm focal length primary lens and a positive 200 mm focal length secondary cylindrical lens instead of the initial balanced matching pair. Moving the assembly out to a 6.5 inch standoff significantly relaxed the extreme depth of field constraints that were previously imposed by a tight 65 mm proximity. This extra focal depth expanded the physical alignment tolerances, allowing the system to achieve sharp target focus along the ground track.

To manage these mechanical tolerances in the laboratory, the fixed track configuration was replaced by a custom 3D printed telescoping housing. This custom enclosure utilizes

internal structural alignment lips and specialized slot barriers that are precisely molded to the profiles of the cylindrical lenses and the pre-mounted 4.03 mm Thorlabs C340TMD-B aspheric objective barrel. These physical features completely lock the optics into the exact 39 mm internal air gap required to maintain crisp focus. While initial testing of the telescoping sleeves revealed a minor optical centration offset due to 3D printing tolerances, this spatial error was successfully resolved through manual bracket adjustments. The rigid 3D printed assembly provides stable alignment and focal preservation for prototype validation under active tracking conditions.

4.3 Practical Constraints

The design and implementation of the A.U.R.A. system are influenced by several practical constraints that impact system performance, development strategy, and overall feasibility. These constraints include limitations related to time, cost, environmental conditions, and system integration. Addressing these constraints is essential to ensure that the system remains achievable within the scope of a senior design project while still meeting functional requirements. These practical constraints play a critical role in guiding system design decisions, ensuring that the A.U.R.A. system remains feasible, reliable, and effective within real-world limitations.

4.3.1 Time Constraints

The development of the A.U.R.A. system is constrained by a limited project timeline defined by the academic semester. This restricts the ability to design and implement highly customized hardware or develop complex software solutions from scratch. As a result, the system design prioritizes the use of existing platforms, libraries, and development tools.

The use of a Raspberry Pi running a Linux-based operating system allows rapid development by providing access to pre-built drivers and software frameworks. Similarly, the use of ESP32-based microcontroller modules reduces development time by offering integrated wireless communication and support for standard programming environments.

Time constraints also affect testing and validation. Because the system must be developed and demonstrated within a fixed timeframe, testing is performed incrementally throughout development rather than relying on extended validation phases. This requires the system to be modular, allowing individual components such as communication, image processing, and control systems to be tested independently.

To accommodate these limitations, the system design emphasizes simplicity, modularity, and the use of well-documented technologies. This approach reduces development risk and ensures that a functional prototype can be completed within the available time. These constraints also limit the ability to explore alternative architectures or perform extensive optimization, reinforcing the need for a practical and efficient design approach.

4.3.2 Economic Constraints

The A.U.R.A. system must be developed within a limited budget, which influences the selection of hardware components and system architecture. Cost considerations restrict the use of specialized or high-performance hardware such as dedicated GPUs or industrial-grade imaging systems.

Instead, the system utilizes cost-effective components such as the Raspberry Pi and ESP32 modules, which provide sufficient performance for real-time processing while remaining affordable. Budget limitations also restrict the use of redundant components and advanced sensing technologies, requiring the system to rely on a minimal but effective set of hardware. Additionally, the IMX178M camera modules were selected to provide improved imaging performance, particularly in low-light and near-infrared conditions, while remaining more cost-effective than industrial-grade imaging systems. Although integration of the IMX178M may require additional configuration compared to standard camera modules, it offers a balance between performance and cost suitable for the system requirements.

Economic constraints also affect scalability and redundancy. The system is designed as a prototype, meaning that cost optimization is prioritized over redundancy and fault tolerance. For example, the system does not include backup processing units or redundant communication channels, as these would significantly increase overall cost.

These constraints require careful selection of components to balance performance and affordability. By leveraging widely available and well-supported hardware, the system achieves its functional goals while remaining within budget limitations.

4.3.3 Environmental Constraints

The A.U.R.A. system operates in an environment that presents several challenges, particularly due to its placement beneath a vehicle. Environmental conditions such as low lighting, dust, temperature variations, and mechanical vibrations can impact system performance. These conditions require the system to maintain consistent performance despite external variability, which introduces additional challenges for image processing and detection accuracy.

Lighting conditions are a primary concern, as the system must detect animals in low-light or nighttime environments. To address this, the system uses near-infrared imaging, allowing detection even when visible light is limited. However, variations in ambient lighting and reflective surfaces can still affect image quality and detection accuracy.

The physical environment also introduces challenges related to dust and debris, which may accumulate on camera lenses and degrade image clarity. Additionally, the system must tolerate temperature fluctuations that could affect electronic components and system stability.

Mechanical vibrations from vehicle movement can impact both hardware connections and image stability. To mitigate these effects, the system design must ensure secure mounting of components and stable camera placement.

These environmental constraints influence both hardware and software design decisions, requiring the system to be robust and adaptable to varying operating conditions.

4.3.4 Manufacturability and Integration Constraints

The A.U.R.A. system must be designed for practical assembly and integration within a vehicle platform. This introduces constraints related to component placement, wiring, and overall system layout.

Because the system includes multiple distributed components, minimizing wiring complexity is essential. The use of wireless communication reduces the need for extensive cabling, simplifying installation and improving maintainability. Components such as the Raspberry Pi and ESP32 modules must be positioned in a way that allows efficient communication while maintaining accessibility for testing and maintenance.

The use of standardized interfaces, such as the CSI connection for camera modules, simplifies integration by ensuring compatibility between hardware components. This reduces the need for custom hardware design and allows the system to be assembled using commercially available parts.

Manufacturability constraints also influence enclosure design and physical mounting. Components must be securely mounted to withstand environmental conditions while remaining accessible for adjustments and troubleshooting.

By prioritizing modular design and standard interfaces, the system can be assembled, modified, and tested more efficiently, which is essential for a prototype system developed within limited time and resource constraints. This also allows individual components to be replaced or upgraded without requiring a complete system redesign. The use of a hybrid communication architecture, combining UART for local communication and Wi-Fi for distributed modules, further simplifies system integration while reducing wiring complexity.

The final system uses Wi-Fi rather than UART for the Pi 5-to-ESP32 runtime link. Removing the required USB/UART connection simplified the physical integration after the Pi was unable to read that connection reliably. The camera-stream and safety-command links remain logically separate even though both use Wi-Fi.

4.3.5 Hardware Availability and Component Constraints

The A.U.R.A. system is also constrained by the availability of hardware components and the selection of specific development platforms. Components such as the Raspberry Pi, ESP32 modules, and camera sensors must be readily obtainable within the project timeline.

Delays in procurement or limited availability of newer hardware, such as advanced ESP32 variants, can impact development progress and require design adjustments.

In particular, the selection of microcontroller platforms introduces constraints related to compatibility and supported interfaces. Not all microcontrollers support high-bandwidth data interfaces or advanced peripheral connectivity, which can limit system design options. This requires careful consideration when selecting hardware components to ensure they meet system requirements while remaining accessible.

Additionally, reliance on commercially available components means that the system must adapt to the limitations of existing hardware rather than using fully customized solutions. While this reduces development complexity, it also constrains the level of optimization that can be achieved.

To address these challenges, the system design prioritizes widely available and well-supported components. This reduces the risk of delays and ensures that replacement parts can be obtained if necessary during development and testing.

4.4 Other Constraints

In addition to the two primary design constraints discussed above, several other constraints shape the boundaries of acceptable system behavior. These constraints are each addressed briefly in the following subsections.

4.4.1 Regulatory and Road Safety Constraints

Any system integrated into or mounted beneath a vehicle that operates on public or shared roadways is subject to regulatory oversight from transportation and vehicle safety authorities. In the United States, the National Highway Traffic Safety Administration (NHTSA) governs vehicle equipment standards under the Federal Motor Vehicle Safety Standards (FMVSS). While A.U.R.A. is currently demonstrated on a low-speed electric toy vehicle that does not operate on public roads and is therefore not subject to FMVSS compliance in its prototype form, the design must acknowledge these regulations as binding constraints for any future commercial or road-legal deployment.

The most directly relevant regulatory concern is that no component of A.U.R.A. may create a new road hazard. A protruding sensor housing that could catch on road surfaces, strike pedestrians, or generate debris in the event of a mechanical failure would violate the spirit of road safety regulations regardless of the vehicle's speed class. This is the regulatory basis for the IP65-rated retractable housing design: the sensor modules must be fully contained within the vehicle's body envelope whenever the vehicle is in motion, eliminating any protrusion that could contact the road or endanger other road users. The 1 mph retraction threshold was set conservatively to ensure compliance with this principle well before the vehicle reaches any speed at which a protruding object would pose a meaningful hazard.

4.4.2 Camera Retraction Speed Threshold

The 1 mph vehicle speed threshold at which the camera modules must retract is a system-level constraint that bridges the physical, mechanical, and software domains. It was established by considering three competing factors: the time required for the actuator to fully retract, the minimum safe ground clearance of the deployed housing at low vehicle speeds, and the detection logic requirement that the system must complete a full underbody scan before the vehicle begins moving.

At 1 mph, the vehicle travels approximately 17.6 inches per second. Given the total deployment time specification of under 3 seconds for full extension, the retraction must be completed in a comparable time to ensure the housing is fully stowed before the vehicle covers enough ground to encounter a surface irregularity that could strike the deployed housing. The Hall effect sensor on the wheel hub detects the 1 mph threshold by counting magnet pulses per unit time, and the MCU firmware issues a retraction command to the motor driver within one interrupt service cycle of the threshold being crossed. The system-level latency from threshold detection to motor driver activation is specified at under 100 ms, ensuring the retraction begins promptly and the full retraction is completed well before any risk of ground contact at normal vehicle speeds.

4.4.3 Animal Detection Accuracy Constraint

The A.U.R.A. system specification requires an animal detection accuracy of greater than 95 percent, meaning the AI classification model must correctly identify a living animal beneath the vehicle in at least 95 out of every 100 detection events. This figure was established by reference to automotive safety standards for object detection systems, where false negative rates above a few percent are considered unacceptable in life-safety applications. A false negative in the A.U.R.A. context is a missed detection: the system clears the interlock and permits vehicle movement when an animal is still present beneath the chassis, which is the primary failure mode the system is designed to prevent.

Achieving greater than 95 percent accuracy under real underbody conditions is a meaningful engineering challenge. The NIR imaging environment beneath the vehicle is characterized by non-uniform illumination, variable ground surface reflectivity, wheel occlusion zones, and potential interference from road debris that may superficially resemble animal shapes in infrared. The YOLOv8-based classification model running on the Raspberry Pi 5 must be trained on a representative dataset of underbody infrared images covering a range of animal species, sizes, and positions within the detection zone. The 95 percent accuracy constraint therefore drives not only the choice of inference hardware but also the scope and diversity of the training dataset, the illumination design, and the image resolution of the sensor system.

The final accuracy evaluation was performed using the YOLOv11 Nano model trained through Roboflow. YOLOv8 remained part of the original design research, but it was not the final deployed model because it was not offered through the selected training workflow.

4.4.4 Scalability to Full-Size Vehicles

Although A.U.R.A. is demonstrated on a small-scale electric vehicle, the system is designed with scalability to full-sized autonomous vehicles as an explicit architectural constraint. This means that no design decision may be made that is inherently incompatible with deployment on a full-scale platform, even if it is expedient for the prototype. The target full-scale reference vehicle is the Waymo autonomous robotaxi, which has a front wheel track of 64.7 inches and a minimum ground clearance of 6.5 inches.

The scalability constraint manifests most concretely in the optical design. The anamorphic lens system was designed with two sets of specifications: one for the demo vehicle's 23-inch wheel track and one for the full-scale 64.7-inch track. The underlying optical architecture, a plano-concave and plano-convex cylindrical lens pair forming an inverse Galilean telescope, is the same in both cases. Scaling from demo to full size requires adjusting the lens spacing and the squeeze ratio of the anamorphic pair rather than redesigning the system from scratch. This geometric scalability was a deliberate constraint imposed on the optical design from the outset, and it is reflected in the dual specification table in Chapter 2, which presents both demo-scale and full-scale target values side by side. The system's mechanical and electronic architecture is similarly modular, with the intent that the same PCB designs, firmware, and software pipeline can be carried forward to a full-scale implementation with mechanical adaptation only.

Chapter 5 - Application of ChatGPT and Other Similar Platforms

5.1 Case Study 1 – Contrast Improvement Technique

Please refer to Appendix E, [\[1a\]](#), [\[1b\]](#), and [\[1c\]](#) for the Case Study prompt and the returned responses from Google AI and ChatGPT.

In this case study, the goal of the prompt was to determine which image processing techniques would improve contrast between organic materials, such as fur or skin, and ground surfaces under 850 nm near-infrared illumination [\[1a\]](#). This is directly relevant to the A.U.R.A. system, where detection does not rely solely on object shape but also on differences in reflectance and texture between animals and the ground. Because the system operates in a constrained lighting environment using active NIR illumination, the usefulness of the response depends on how well each model accounts for both physical reflectance behavior and practical implementation.

What I found interesting when comparing the two responses is that they approached the problem from completely different angles, even though the prompt was the same. Google AI [\[1b\]](#) focused much more on reflectance and intensity differences at 850 nm, which is technically correct. It clearly understands that organic materials tend to appear brighter than ground surfaces and leaned heavily into methods like CLAHE, histogram normalization, and thresholding. That part was helpful, but it also started suggesting things

like multi-spectral fusion and intrinsic image decomposition, which do not really apply to my system since I am only working with a single NIR band. It felt like the answer was strong in theory, but not really filtered for what I can actually implement nor catered to A.U.R.A. as a system.

ChatGPT [\[1c\]](#), on the other hand, went in a direction that I honestly found more useful for this specific project. Instead of just focusing on brightness differences, it emphasized contrast in texture. Under a vehicle, an animal is not always going to stand out as a bright object. It might be partially hidden or take up a small portion of the frame, so relying only on intensity is not enough. The idea that local texture and micro-contrast, like fur patterns versus the smoother surface of asphalt, could be a stronger indicator was something that stood out to me since that was also a topic covered in OSE4830 Imaging and Display. It also helped that ChatGPT organized everything into a pipeline which made it easier to understand how I would implement it instead of just picking random techniques.

At the same time, both responses had similar limitations. They both included methods that do not fully apply to my setup and neither one really accounted for system specific constraints like sensor noise from the IMX219, uneven illumination from the LED array, or the fact that everything has to run on limited processing power. That was probably the biggest thought I had. That even when the answers seem correct, they do not automatically match the reality of the system unless you guide them very specifically. This can sometimes be very tedious and require a lot of patience if I were to continue the conversation because you tend to have to repeat certain instructions sometimes when working with these LLMs.

I would say Google AI gave a stronger theoretical explanation, while ChatGPT was more helpful in thinking through how I would actually approach the problem. Neither one gave a complete answer on its own but comparing them did give some useful insight into how I could improve our design.

5.2 Case Study 2 – Motor Driver Recommendation

Please refer to Appendix E, [\[2a\]](#), [\[2b\]](#), and [\[2c\]](#) for the Case Study 2 prompt and outcomes from Gemini and ChatGPT.

After identifying that a motor driver was necessary to operate the linear actuator, the next step was determining which specific motor driver would be best suited for the A.U.R.A. system. The prompt used for this case study, shown in Appendix E [\[2a\]](#), asked both platforms to recommend a motor driver capable of controlling a 12V DC linear actuator from a 3.3V microcontroller, with enough current capacity to handle stall conditions.

Gemini [\[2b\]](#) provided a well-structured response that not only listed candidate motor drivers but explained the reasoning behind each recommendation. It correctly identified that the primary requirements for this application are bidirectional control through an H-bridge topology, a continuous current rating sufficient to handle actuator stall, and logic-level compatibility with a 3.3V microcontroller. Gemini specifically recommended the Cytron MD10C, citing its 13 A continuous current rating, its PWM and direction pin interface, and its built-in protections against over-current and short circuit faults. It also noted that many lower-cost alternatives such as the L298N module, while popular in

hobbyist projects, suffer from a significant internal voltage drop that reduces the effective voltage delivered to the motor and increases heat dissipation, making them unsuitable for applications requiring reliable force output near the rated load.

ChatGPT [\[2c\]](#), when given the same prompt, also recommended the Cytron MD10C and agreed with Gemini's assessment of the L298N's limitations. However, ChatGPT's response went further in explaining the underlying physics of why the L298N's bipolar transistor architecture produces a higher voltage drop compared to the MOSFET-based design of the MD10C. This additional depth was useful for understanding the technical trade-off rather than simply accepting a recommendation at face value. ChatGPT also flagged that for higher-voltage or higher-current applications beyond A.U.R.A.'s scope, discrete MOSFET H-bridge designs would be preferred, which was relevant context for the technology comparison section of Chapter 3.

An interesting difference between the two platforms emerged around source reliability. ChatGPT's citations in this case leaned toward electronics hobbyist blogs and community forum posts, which tend to reflect practical real-world experience but may not represent the most rigorous technical sources. Gemini, by contrast, surfaced references from Cytron's official product documentation and engineering review sites, which provided more authoritative specifications. This difference in sourcing behavior reinforced the importance of verifying any AI-generated component recommendation against the official manufacturer datasheet before committing to a selection.

Overall, this case study demonstrated that AI platforms are most valuable in the motor driver research context when used to quickly identify the right component category and narrow down a shortlist of candidates, rather than as a replacement for datasheet-level verification. Both Gemini and ChatGPT performed well in this role, and their agreement on the Cytron MD10C provided a useful starting point that was subsequently confirmed through independent review of the manufacturer's documentation.

5.3 Case Study 3 – Suitable Communication Protocol

Please refer to Appendix E [\[3a\]](#) and [\[3b\]](#) for the Case Study prompt and outcome from ChatGPT.

In this case study, the goal of the prompt was to determine the most appropriate communication method between the Raspberry Pi and ESP32 microcontrollers in the A.U.R.A. system. This is directly relevant to the project because reliable communication is required for transferring detection results, control commands, and system status between processing and control units.

ChatGPT provided a comparison of several communication protocols, including UART, I2C, SPI, and Wi-Fi. It explained that UART is well-suited for direct, short-range communication due to its simplicity and low latency, while Wi-Fi is more appropriate for distributed systems where multiple sensor nodes must communicate wirelessly. This was

helpful in understanding the trade-offs between wired and wireless communication in embedded systems.

What I found useful about the response was that it introduced the concept of bidirectional communication, where commands and status updates can be exchanged between devices. This helped clarify how the Raspberry Pi would issue control signals and how the ESP32 controllers would respond with system feedback. It also provided a general framework for structuring communication within the system.

However, the response did not fully account for the specific architecture of the A.U.R.A. system. While it mentioned that a hybrid approach may be considered depending on system requirements, it did not explicitly determine how such an approach should be implemented for this project. Additional reasoning was required to determine that UART should be used between the Raspberry Pi and the main ESP32 controller, while Wi-Fi should be used for communication with front and rear sensor nodes.

A limitation of the response was that it did not consider hardware constraints such as wiring complexity, latency requirements for safety-critical signals, or the reliability of wireless communication in different environments. These factors had to be evaluated separately to ensure the system would operate safely and consistently.

Overall, ChatGPT provided a strong starting point and helped guide the communication design process. However, the final implementation required refinement and adaptation to match the specific requirements of the system. This case study demonstrates that while ChatGPT is useful for exploring design options, its responses must be critically evaluated before being applied to real-world engineering problems.

5.4 Case Study 4 – Suitable Ultrasonic Sound Component

Please refer to Appendix E [\[4a\]](#), [\[4b\]](#), and [\[4c\]](#) for the Case Study prompt and outcome from ChatGPT and Google Gemini.

When it comes to starting out on a project or research, ChatGPT can be quite a useful tool. It gives a solid foundation on where and how to begin looking for information. Regarding my prompt, ChatGPT not only gave me the kinds of components I should be looking for but also a variety of available market products that might work for my ultrasonic application. In addition to presenting a variety of possible suitable components, it also weighed the pros and cons of each and reasoned as to what each component could be used for and why it was inferior to the standard component of piezoelectric transducers.

In addition to the comparison of different components, ChatGPT also provided more information as to the properties that would serve well in my application of an ultrasonic repellent. For instance, it presented the frequency range quality a component must have to function as a viable ultrasonic deterrent, in addition to this it also presented that a continuous fixed output would become ineffective over time.

ChatGPT also provided the basic architecture used to drive piezoelectric transducers to output ultrasonic noise, with the architecture of a frequency generator, and an amplifying stage to drive the output which is similar to those found in commercial products and is the architecture our system will attempt to use. Additionally, ChatGPT provided a variety of sources.

While ChatGPT did provide a solid foundation as what should be looked for in the design of the ultrasonic alarm system, the information it provided still needed to be fact checked for accuracy.

In comparison, Google Gemini provided little comparison to other possible components and instead opted to compare piezoelectric transducers and piezo tweeters. Unlike ChatGPT, Gemini did not weigh the pros and cons of each and instead detailed each in their functionality. Similarly to ChatGPT, Gemini also provided the basic architecture needed to use a piezoelectric transducer, including a signal generator, an amplifier, and an output. One more difference between ChatGPT and Gemini is that Gemini provided a video of commercial ultrasonic pest repellers. While both platforms provided a foundation and starting point for continued research, ChatGPT was more detailed in its comparisons while Gemini was more straightforward.

Chapter 6 - Physical System Design

6.1 Hardware Design

6.1.1 Main PCB

The main PCB contains our primary low-level MCU connected to all the peripherals attached to the main board including a USB port used for a serial UART connection to our on-board computer using the ESP32-S3's native USB port. This system also contains our strobe-light deterrent configured with a MOSFET being used as a constant current source for the LEDs connected in parallel. It also contains our motor interlock system that is connected to our miniature electrical vehicle, with the interlock being configured with an SPST relay driven by a NPN BJT and a feedback diode parallel to the relay. The two lead hall-effect sensor is connected to the MCU through a female pin connector. Finally, two nearly identical ultrasonic output systems are also connected to this system, each containing an array of five piezoelectric transducers, an N-channel H-bridge connected to two half-bridge drivers.

Figure 6.1.1 High-Frequency Alarm




Figure 6.1.1.1 Revised High-Frequency Alarm

The high-frequency alarm consists of three stages, the frequency generation from the MCU connected to high and low side inputs of the half-bridge drivers, the H-bridge amplifier, and the five-component transducer transmitter array. Two nearly identical systems will be connected, one using our 40kHz transducers and the other using our 30.5 kHz transducers

For the updated schematic in SD2, the half-bridge drivers were changed to IR2104 drivers which have a built in deadtime to prevent shoot through. Additionally, extra connections were made between the VS pin of the half-bridge driver to the high and low MOSFET junction. Finally, the 4 feedback diodes were removed from the circuit as they proved to be unnecessary, instead two capacitors were placed in series to the 12V input rail to filter the signal.

Half-Bridge Driver:

Each half bridge contains the following pin connection: Vcc, HIN, LIN, COM, Vb, HO, Vs, LO. Following the datasheet recommendation and typical connections for the IRS2003 half-bridge driver, a bootstrap diode was placed between the connection of Vcc and Vb and an additional bootstrap capacitor was placed between Vb and Vs. In addition, to the bootstrap capacitor another capacitor was placed between Vcc and ground. The purpose of the bootstrap circuit is to supply voltage to the high-side MOSFETs. Each driver is connected our MCU through the HIN pin and connected to each half of the H-bridge amplifier through HO and LO pins and series resistors.

H-bridge:

The H-bridge amplifier circuit to be used to drive our piezoelectric transducer array uses four N-Channel MOSFETs. All MOSFETs are grounded with a high ohm resistor at the gate, high side MOSFETs are connected at drain, low side MOSFETs are connected at source. High side and low side MOSFETs are connected through high side source and low side drain. A $12V_{cc}$ supply is connected between high side gate connections, ground is connected between low side source connections. Left side high MOSFET is connected to right side low MOSFET by the gate; the same connection is shared by the left high side and right side low MOSFETs. Feedback diodes are placed between top and low side MOSFETs to protect them from voltage spikes. Driver circuits are connected to half-bridge connections at high and low side gates.

Transducer Array:

Transducers are arranged in a parallel connection to the output of the N-channel H-bridge amplifier.

Strobe-Light:




Figure 6.1.1.2 3W LED Strobe-Light




Figure 6.1.1.2.1 Revised 3W LED Strobe-Light

The strobe-light will consist of a MOSFET behaving as a constant current source connected to an array of high-power LEDs arranged in two parallel, series pairs of LEDs; pulses and timing intervals are being controlled by the MCU. A 10k pull down resistor was connected to the output from the MCU and a resistor towards the gate of the MOSFET. The source of the MOSFET is connected to ground and the drain is connected to the cathode side of the array and a 12V input connected to the anode side.

For the updated schematic in SD2, the 3W LEDs were replaced with simple indicator LEDs. A MOSFET was still used to drive the LED array; however instead, 2 parallel lines of 3 LEDs, 3 parallel lines of 5 LEDs were used to make up for the drop in intensity. The reason for the change was that 3W LEDs needed a heat sink to properly function. This may have caused damage to the PCB if they ran hot. Additionally, 3W LEDs could prove to be not eye safe, so indicator LEDs seemed to be the better overall fit.

MCU:




Figure 6.1.1.3 Main PCB MCU

USB port is connected to MCU through the ESP32-S3's native USB peripheral using GPIO pins 19 and 20 connected to the port's D- and D+ connections. MCU also contains a 10k resistor and 1u capacitor connected to the enable pin, in-line with the datasheets recommended set-up. Additional connections to this MCU also include left and right side connections to 40 kHz and 30.5 kHz resonant alarm drivers, relay control connection, strobe light control, and output from the hall-effect sensor

The MCUs pin layouts were changed for SD2, instead of having 4 GPIO pins for each ultrasonic alarm, 3 for each were used for the IN and SD pins of the drivers. Additionally, capacitors and buttons were placed for flashing the MCU using BOOT and RESET. Finally, a CP2102 USB-UART bridge chip was used to program the MCU via the transmit and receive pins of the MCU.

USB Connection to On-Board Computer:




Figure 6.1.1.4 USB to PI5

Our on-board computer will be connected to our MCU via a USB connection, with the raspberry PI also being connected to our LCD screen through a direct wired connection from the raspberry PI. The USB port is connected to the MCU via the D+ and D- pins.

In SD2, this connection was not used as when tested, the Raspberry PI5 would not recognize the MCU even after through troubleshooting, instead we opted for a Wi-Fi connection between MCU and Raspberry PI5.

Hall-Effect Sensor:



Figure 6.1.1.5 Hall-Effect Connector

The hall-effect sensor contains three wire leads connected to a female pin connector; one connected to ground, one to a GPIO for the sensor's output, and the other to a 5V input.

Motor Interlock:




Figure 6.1.1.6 Vehicle Motor Interlock

The motor interlock subsystem consists of an NPN transistor connected to the MCU as the base through a series resistor, the transistors emitter is connected to ground, and its collector is connected to a feedback diode and the coil of the SPST automotive relay. The SPST relay will be connected between the 12 V vehicle battery and its 10 A fuse, the feedback diode will be connected in parallel to the SPST relay coil to protect components from large current spikes from the dual motor circuit.

In SD2 this schematic did not change, slight modifications were made or the integration into our vehicle, instead of being connected between the vehicle's battery and fuse, it was connected after the fuse and before the control box of the vehicle.

6.1.2 Secondary PCBs




Figure 6.1.2.A Secondary PCB Schematic part 1

The first part of the secondary PCB schematic shows the USB-C connector, I2C pull-up resistors, ESP_EN pull-up, decoupling capacitor, camera module FPC connector (CAM_MODULE), and the left side of the ESP32-P4 module. The USB-C connector includes 5.1k Ω CC resistors for device mode enumeration and carries USB_DP and USB_DM signals to the ESP32-P4. The camera module connects to the ESP32-P4 via MIPI CSI-2 differential pairs (MIPI_CLK, MIPI_L0, MIPI_L1), with SDA and SCL lines also routed through the connector for camera register configuration. The two LM2575 external regulator modules are also shown, providing the 5V and 3.3V power rails to the system.




Figure 6.1.2.B Secondary PCB Schematic part 2

The second part of the secondary PCB schematic shows the right side of the ESP32-P4 module alongside the CP2102N USB-to-UART bridge, VL53L0X time-of-flight sensor, and Cytron MD10C motor driver. The CP2102N is wired with a VBUS resistor divider, bypass decoupling capacitors, and RST pull-up resistor, with its UART lines connected to the ESP32-P4 for programming and debugging. The VL53L0X sensor is connected via I2C on GPIO0 and GPIO1, with XSHUT on GPIO2 for hardware standby control. The Cytron motor driver receives PWM and direction signals from GPIO3 and GPIO4 respectively.

6.1.2.1 High-Speed Vision (CSI Interface)




Figure 6.1.2.2 IMX178M Camera Module Connector

The IMX178M camera module connector is a 30-pin FPC connector (Molex 524353071, 0.5mm pitch) used to interface the IMX178M camera sensor with the ESP32-P4. The connector carries multiple power pins including 5V on pins 2-5 and 3.3V on pin 7 for camera logic and analog supply. The MIPI CSI-2 differential pairs are routed through the connector with MIPI_L0 on pins 9-10, MIPI_L1 on pins 12-13, and the clock lane MIPI_CK on pins 15-16, each pair surrounded by GND pins for signal integrity. I2C communication lines SCL and SDA are present on pins 24 and 25 respectively, used for camera register configuration via the VEYE register map. Pins 17-22 and 26-29 are left unconnected as they are not required for this application.

6.1.2.3 Proximity Sensing (I2C Bus)




Figure 6.1.2.3.A VL53L0X Time-of-Flight Sensor Connector

The VL53L0X time-of-flight sensor is connected via a 6-pin header. The sensor is powered by the 3.3V rail on pin 1 with GND on pin 2. I2C communication is handled through SCL and SDA on pins 3 and 4 respectively, shared with the I2C bus used for camera register configuration. XSHUT on pin 5 is connected to GPIO2 on the ESP32-P4, allowing the firmware to place the sensor in hardware standby mode when not in use. Pin 6 is left unconnected as the GPIO1 interrupt pin of the VL53L0X is not utilized in this design.

6.1.2.4 Actuation and Motor Control Interface



Figure 6.1.2.4 Figure X: Cytron MD10C Motor Driver Connector

The Cytron MD10C motor driver is connected via a 3-pin header carrying DIR on pin 1, PWM on pin 2, and GND on pin 3. The DIR signal is driven by GPIO4 on the ESP32-P4 and controls the rotational direction of the motor, while the PWM signal from GPIO3 controls the motor speed through pulse width modulation. The MD10C is capable of driving DC motors up to 10A, making it suitable for driving the linear actuator used in this

design. The ACTUATORFRONT and ACTUATORREAR headers, along with the FROM_CYTRON pass-throughs, manage the high-current side of the system. While the MCU operates at 3.3V, these headers are designed to handle the 12V rail required to drive the electric actuators. This power isolation is a critical feature, ensuring that the high-current draw of the motors does not cause voltage drops that would crash the MCU logic.

6.1.3 Voltage Regulators

Voltage regulators will be needed to appropriately power the different components and systems within A.U.R.A.




Figure 6.1.3.1 3.3V Regulator

The schematic for the 3.3V regulator was drawn from the LM2575-3.3WT datasheet's application diagram. These regulators will be used to sufficiently power our microcontroller units from our 12V power supply.




Figure 6.1.3.2 5V Regulator

Similarly to that of 3.3V regulators, the schematic for the 5V regulator was drawn from the reference design in the same datasheet as both regulator ICs use the same datasheet. This 5V regulator will be used to power the different sensors in our system.




Figure 6.1.3.3 5.1V, 5A Regulator

The 5.1V, 5A regulator will be used in our system to power the Raspberry PI5 on-board computer on the main PCB. It uses an AP64501 synchronous converter IC to program the

appropriate voltage output. It uses a voltage divider to program the necessary output, for our purpose R1 will be set to 118k Ω and R2 will be set to 22.1k Ω . The design of the regulator was obtained from the typical application figure in the AP64501 datasheet.




Figure 6.1.3.4 Revised Regulators 5.2V/3.3V, 5A/1A

The 5 V and 3.3 V regulators were revised into the single LM2679-based design shown in Figure 6.1.3.4, which replaces both the earlier LM2575-5.0WT 5 V regulator (Figure 6.1.3.2) and the AP64501 5.1 V, 5 A regulator (Figure 6.1.3.3). This change was made on the recommendation of Dr. Weeks, who advised using the LM2679 regulators. The LM2679-ADJ is an adjustable 5 A step-down switching regulator, so a single regulator design is able to supply both rails, a 5.2 V, 5 A rail for the Raspberry Pi 5 and a 3.3 V, 1 A rail for the microcontrollers, with only the feedback resistors changed between them.

The output voltage of each LM2679-ADJ is programmed with a feedback resistor divider running from the output to the FB pin to ground. The regulator holds the FB pin at its internal reference of 1.21 V, so the output is given by $V_{out} = 1.21 \text{ V} \times (1 + R1/R2)$, where R1 is the top resistor (VOUT to FB) and R2 is the bottom resistor (FB to GND). The bottom resistor is fixed at $R2 = 1.0 \text{ k}\Omega$, and the desired rail voltage is then set by choosing the top resistor from $R1 = R2 \times (V_{out} / 1.21 \text{ V} - 1)$. Because the output rises with the top resistor, a larger R1 produces a higher rail voltage: $R1 = 3.3 \text{ k}\Omega$ sets the 5.2 V rail, while $R1 = 1.8 \text{ k}\Omega$ sets the 3.3 V rail.

The 5 V rail is deliberately programmed slightly high, to 5.2 V, so that after the small voltage drop across the downstream wiring and connectors the Raspberry Pi 5 still receives approximately 5 V at its input under full load. Using one adjustable part for both rails also reduces the bill-of-materials to a single regulator family and keeps the feedback and filter design consistent across the two outputs.

6.2 Optical Design

6.2.1 Design from Geometric Requirements

The optical subsystem was developed directly from the geometric and imaging constraints defining the underbody sensing region and the physical dimensions of the selected image sensor. These constraints produce a highly asymmetric field of view where the horizontal extent of the scene is significantly larger than the vertical extent. As a result, a conventional imaging approach that treats both axes equally would lead to an inefficient use of sensor area and reduced effective resolution across the region of interest. Instead, the optical system must selectively compress the horizontal field while preserving vertical information, allowing the scene to be mapped efficiently onto the sensor.

The required mapping between object space and image space was previously defined through geometric relationships involving field of view and working distance. A critical parameter in this mapping is the Squeeze Ratio (Demo), which was established at 100:1 to match the IMX178 active area. In the initial calculations, this ratio was associated with a lateral magnification (M) of approximately 0.49x based on an 8.5 inch working distance and a 3 mm lens. However, these relationships alone do not guarantee that a physically realizable optical system can achieve the desired result. The purpose of this design stage is therefore to translate these geometric targets into a sequential optical model and evaluate whether the selected architecture can physically propagate rays, compress the field, and form an image within the sensor bounds.

During the initial modeling phase, it became necessary to stray from the original geometric proposal of an 8.5 inch working distance. While an 8.5-inch standoff provided a comfortable mechanical envelope, a quantitative analysis of the Lagrange Invariant proved that the 55° horizontal field of view would require a primary lens diameter exceeding 100 mm to prevent total ray clipping at that distance. Because the project was constrained to use 30 mm catalog optics for reasons of cost and mechanical packaging, the 8.5 inch working distance resulted in over 90 percent vignetting in the simulation.

To resolve these conflicts, a custom Python based computational model was developed to iterate through potential system geometries. Using the Python raytracing library, we were able to perform high speed paraxial sweeps of various working distances and lens spacings. This computational approach allowed us to identify a short-standoff configuration with a 65 mm working distance as the optimal pivot point. By reducing the working distance, the light from the vehicle underbody can enter the 30 mm lens aperture before the ray bundle diverges beyond the physical bounds of the glass, enabling the system to maintain high light collection efficiency within the now 8-inch total track limit.

Once the Python model provided a validated paraxial solution that achieved 81.9 percent sensor utilization, Zemax OpticStudio was used to construct and evaluate the real lens implementation. This two-stage modeling approach allows for direct visualization of ray propagation, surface interactions, and image formation behavior. The design process followed a progressive approach, beginning with the Python paraxial validation and

advancing toward a complete Zemax system implementation. At each stage, the system was evaluated for ray propagation, vignetting, image scaling, and mechanical feasibility, ensuring that both optical and physical constraints were addressed simultaneously.

6.2.2 First-Order Analytical Validation

To provide the technical depth required for the final design, the system was validated through a series of analytical derivations. This stage moves beyond initial geometric estimates to establish the specific focal lengths and positional requirements of the optical train through a formal paraxial ray trace.

6.2.2.1 Objective Focal Length Recalculation

A critical requirement for system feasibility involved a formal recalculation of the objective focal length based on the required system magnification. In the preliminary design phase, a magnification of 0.49x was targeted based on an 8.5 inch working distance. However, as the design shifted to a 65 mm working distance to satisfy aperture constraints, the target magnification was re-derived to match the specific dimensions of the Sony IMX178 sensor and the 92.83 mm half-object height required for dual-camera overlap.

The required lateral magnification (M) is defined as the ratio of the target image height on the sensor (h_{img}) to the half-object height in the scene (h_{obj}), $M = \frac{h_{img}}{h_{obj}}$.

The active area of the IMX178 sensor is 7.41 mm wide, giving a half-sensor width (h_{img}) of 3.705 mm. To map the required 92.83 mm half-object scene onto this area, the required magnification is calculated as follows, $M = \frac{h_{img}}{h_{obj}} = \frac{3.705}{92.83} \approx 0.0327 \times$.

This 0.0327x magnification represents the finalized paraxial requirement necessary to achieve the intended 2.04:1 horizontal squeeze at a short standoff. While the initial calculations used a 0.49x target, that value was incompatible with a 65 mm working distance and the physical diameters of the catalog optics. Using the paraxial relationship between focal length f and magnification, we get $f = \frac{M \cdot D}{1+M} = \frac{0.0327 \cdot 65 \text{ mm}}{1+0.0327} \approx 2.06 \text{ mm}$.

Although the paraxial math suggests a 2.06 mm lens would satisfy the 0.0327x mapping, a lens of this focal length possesses a back focal length of less than 1.5 mm, which was determined to be mechanically infeasible for housing the 850 nm interference filter. By evaluating the system as a multi-lens relay, the 9 mm objective was identified as the optimal selection. When placed at the calculated expansion gap in the anamorphic relay, the 9 mm lens achieves the target 0.0327x reduction while providing a stable 5.45 mm back focal length. This recalculation proves that the shift in magnification was a necessary adjustment to meet the original 2.04:1 squeeze ratio within the physical limits of the hardware.

6.2.2.2 ABCD Transfer Matrix Derivation

The system was validated using the ABCD Transfer Matrix method to track the ray height (y) and angle (θ) through the anamorphic relay. This method is essential for proving the imaging condition, which ensures the rays converge to a sharp focus on the sensor. The fundamental equation for the matrix transformation is as follows:

$$\begin{bmatrix} y_{out} \\ \theta_{out} \end{bmatrix} = \begin{bmatrix} A & B \\ C & D \end{bmatrix} \begin{bmatrix} y_{in} \\ \theta_{in} \end{bmatrix}$$

The transfer matrix shown above provides a complete mathematical description of the horizontal optical path. It should be noted that in the matrix above, the variable D is a mathematical element of the transfer matrix and is distinct from the system working distance D used in our geometric calculations. To solve for the final system state, the individual matrices for each component are multiplied in reverse order, starting from the sensor and moving back toward the object. In this derivation, the total system matrix is the product of the translation through the working distance D , the refraction through the first cylindrical lens, the internal spacing between the cylinders, the second cylindrical lens, the expansion gap, and finally the 9 mm objective lens.

The most critical result of this matrix multiplication is the B -element, which represents the system focus. By setting the B -element to zero, we determined the exact sensor placement required to form a sharp image. This analytical approach revealed that the 9 mm objective lens provides a 5.45 mm back focal length, which is a significant improvement over the initial 3 mm estimate. This distance is a vital mechanical constraint as it provides the physical clearance necessary for the 850 nm near infrared interference filter and the sensor protective glass.

Furthermore, this matrix model allows for the calculation of the System Magnification, which is found in the A -element of the final product. The derivation confirms that the high-power cylindrical relay, when combined with the 9 mm objective, produces the required 0.0327x reduction. This ensures that the 110 degree full field of view is compressed to fit within the horizontal bounds of the IMX178 sensor, achieving the target 81.9 percent utilization without the geometric clipping that would occur at shorter working distances. This mathematical verification establishes the paraxial foundation required to begin the computational optimization phase.

6.2.2.3 Paraxial Ray Trace

To verify that the ray bundle remains within the clear aperture limits of the catalog optics, a formal paraxial ray trace was performed at every optical interface. This analysis ensures that the incident rays from the front focal plane do not exceed the physical semi-diameter of the lens substrates. The chief ray at an incident angle of 55° , representing the edge of the 110° total field of view, was tracked to determine the maximum ray height, y , at each stage of the relay.

At the optimized working distance, D , of 65 mm, the targeted half-object height, h_{obj} , is established by the field requirement. For the A.U.R.A. system to resolve a 110° full field of view at this standoff, the system must capture an object height of approximately 92.8 mm. This represents the "worst-case" geometric requirement that ensures the 30 mm catalog optics can handle the extreme angles required for underbody imaging without internal vignetting.

The ray height as it hits the first plano surface of the LK1336L1-B lens is calculated using the tangent of the 55° half-field angle and the 65 mm working distance:

$$y_1 = y_{obj} - (D \cdot \tan(55^\circ))$$

$$y_1 = 92.8 \text{ mm} - (65 \text{ mm} \cdot 1.428) \approx 0.0 \text{ mm (on-axis interaction)}$$

This result represents an on-axis interaction, indicating that the chief ray from the field edge intersects the optical axis at the first refractive interface. This is a critical technical achievement, as it proves the 65 mm standoff ensures the wide-angle light bundle is successfully captured within the 15 mm clear semi-diameter of the primary lens. By ensuring the rays enter the system at this optimized height, the design maintains a robust engineering buffer, preventing the mechanical vignetting that would occur at shorter working distances.

Table 6.2.1 Surface-by-Surface Ray Trace

Surface #	Component	Description	Ray Height (y), in mm	Ray Angle (θ), in $^\circ$	Clear Aperture, in mm
0	Object	Front Focus Plane	92.83	55.00	N/A
1	LK1336L1-B	Front Surface (Radius)	7.10	55.0	15.00
2	LK1336L1-B	Back Surface (Plano)	6.50	36.10	15.00
3	Air Gap	Space to Lens 2	10.00	36.10	N/A
4	LJ1909L1-B	Front Surface (Plano)	10.00	-24.80	7.50
5	LJ1909L1-B	Back Surface (Radius)	8.10	-14.20	7.50
6	Air Gap	Space to Objective	4.80	-14.20	N/A
7	EO #65-549	Objective Front	4.80	-41.50	4.50
8	EO #65-549	Objective Back	4.80	-38.20	4.50

9	Image	IMX178 Sensor	3.04	0.00	3.705
---	-------	------------------	------	------	-------

The data in Table 6.2.2.3 utilizes a thick-lens surface-by-surface analysis to account for the physical center thickness of the N-BK7 and achromatic substrates. Unlike a simplified thin-lens model, tracking both the front and back interfaces of each lens is necessary to verify the total optical track length and ensure the ray bundle remains within the clear aperture for the entire duration of the glass transit. A critical distinction in this paraxial trace is the variation of ray height within the glass substrates of different components. In the Radius-Plano and Plano-Radius cylindrical elements (Surfaces 1 through 5), the ray height changes significantly because the refractive power is applied at the curved interface. This initiates refraction immediately upon entry in the primary stage and provides a controlled exit in the secondary stage, ensuring the rays remain on a converging trajectory as they travel through the system. Specifically, the LK1336L1-B cylinder reduces the ray angle from 55.00° to 36.10°, allowing the ray height to reach a manageable 10.00 mm as it enters the secondary optics.

Conversely, in the 9 mm objective lens (Surfaces 7 and 8), the ray height is modeled as constant at 4.80 mm. This is a standard paraxial treatment for a compact achromatic doublet where the refraction is captured at the principal plane. Because the cylindrical relay has already "squeezed" the beam, the rays enter the objective at a much shallower trajectory than at the system entry. The final transition from 4.80 mm to the image height of 3.04 mm occurs across the 5.45 mm back focal length, where the ray angle of -38.20 degrees directs the bundle to the sensor. This achieves the target 81.9 percent sensor utilization, providing a robust engineering buffer for the IMX178 sensor. This surface-by-surface analysis proves that the Radius-Plano / Plano-Radius architecture and the 65 mm working distance D are the mathematically required solutions to enable high-angle imaging with standard catalog components.

6.2.3 Computational Modeling in Python

Once the analytical framework was established through manual ray tracing, a custom computational model was developed in Python to identify the physical aperture limits of the catalog optics. The Python environment functioned as a numerical engine to evaluate thousands of potential lens spacings against the 15 mm semi-diameter of the Thorlabs cylindrical elements.

The primary discovery of the numerical validation was the identification of the geometric aperture limit. The simulation revealed that at a 65.0 mm working distance, the 55-degree chief rays from the field edge are successfully directed toward the optical axis. As shown in the final output of the A.U.R.A. Numerical Validation Engine, the ray height at the primary lens entry is reduced to 0.00 mm, representing an on-axis intersection. This result provides a maximum 15.00 mm aperture safety margin, ensuring the wide-angle light bundle is captured with 100 percent efficiency and providing a robust engineering buffer that protects the system from mechanical vignetting.

By centering the simulation on this 65.0 mm standoff, the Python engine verified the final image height required for the IMX178 sensor. The model confirmed that the 110-degree scene is successfully compressed to a 3.04 mm image height, achieving a finalized 81.9 percent sensor utilization. This satisfies the project requirement for high spatial data density while maintaining a 0.66 mm safety buffer on each side of the active pixel area. This computational step locked the axial coordinates for the system before moving into high-fidelity physical validation.

Table 6.2.2 Python Optimization Engine Results

Parameter	Optimized Value	Status
System Working Distance (D)	65.0 mm	Locked
Field Entry Ray Height (y_1)	0.00 mm	On-Axis
Aperture Safety Margin	15.00 mm	Validated
Final Image Height (y_{img})	3.04 mm	Target Met
Total Sensor Utilization	81.9 %	Finalized

6.2.4 Physical System Validation in Ansys Zemax OpticStudio

The final stage of the optical design process utilized Ansys Zemax OpticStudio to provide a high-fidelity physical validation of the A.U.R.A. system. While the previous computational stages established the paraxial framework, this physical simulation accounts for the real-world mechanical constraints of the selected Thorlabs and Edmund Optics catalog components. By implementing the specific center thicknesses, radii of curvature, and spectral dispersion of the glass substrates at 850 nm, the model provides a definitive proof of concept for the hardware assembly.

To reconcile the paraxial framework with real-world constraints, each element was modeled as a sequential refractive surface using the N-BK7 glass catalog at a target wavelength of 850 nm. A critical step in this implementation was the orientation of the cylindrical radii to maximize refractive efficiency; the LK1336L1-B was oriented in a Radius-Plano configuration, while the high-power LJ1909L1-B utilized a Plano-Radius orientation. To simulate physical aperture bottlenecks, the Aperture Stop was positioned at the entrance of the 9 mm achromatic objective lens. By utilizing a Float-by-Stop Size configuration, the engine was forced to back-map ray bundles through the relay, ensuring the resulting layout respects the physical clear apertures of the hardware.




Figure 6.2.1 Shaded 3D System Layout (Isometric Detail View)

As illustrated in Figure 6.2.1, the implementation confirms the mechanical feasibility of the 110-degree capture. To maximize technical clarity, the layout utilizes an isometric perspective cropped to the final 10 mm of the working distance, focusing on the primary refractive interfaces. The simulation verified that at the optimized 65.0 mm working distance, the 55° chief rays are successfully captured within the 15 mm clear semi-diameter of the primary lens, providing a 7.9 mm engineering buffer against mounting misalignments.

A critical observation in the 3D trace is the presence of marginal stray rays and the close-proximity interaction at the objective entrance. While the extreme field edges graze the physical aperture of the 9 mm lens, this tight convergence is an intentional result of the high-power anamorphic squeeze. The rays seen bypassing the sensor substrate represent a controlled geometric overfill, ensuring that the active pixel area remains fully saturated with the near-infrared signal without edge-vignetting. Furthermore, the physical validation accounts for spectral dispersion, confirming that while the objective possesses a theoretical 5.45 mm focal length, a finalized 2.0 mm mechanical air gap is required to achieve optimal focus. This configuration provides the necessary clearance for the 850 nm interference filter, resulting in a build-ready specification that serves as the foundation for the subsequent housing design.

The final sequential layout in Ansys Zemax OpticStudio was completely reconstructed and re-optimized to capture the true hardware updates introduced during the laboratory prototyping phase. The updated model substitutes the secondary cylindrical lens with a premium fused silica element (Newport CSX200AR.16) and replaces the baseline 9 mm lens with the pre-mounted Thorlabs C340TMD-B aspheric objective. The mechanical air gaps within the telescoping enclosure were adjusted to a 39 mm separation between the cylindrical pair, matching the 6.5 inch straight-on longitudinal standoff distance.

As displayed in the updated 3D shaded layout (Figure 6.2.2), the sequential ray trace verifies clean, unvignetted light transmission for the full 121-degree horizontal field of view. The corresponding matrix spot diagram (Figure 6.2.3) tracks aberrations across the fields, revealing a geometric spot radius of $40.2\text{ }\mu\text{m}$ on-axis, which expands to a maximum of $82.8\text{ }\mu\text{m}$ at the extreme 55-degree horizontal field margin. While the spot diagram demonstrates controlled vertical astigmatism at the 25-degree vertical field margin, the overall resolution boundaries cleanly satisfy the geometric mapping constraints required to successfully compress the physical underbody track onto the active pixel boundaries of the Sony IMX178M monochrome sensor without induced edge truncation.




Figure 6.2.2 Updated Ansys Zemax OpticStudio 3D shaded lens layout of the final system configuration, demonstrating unvignetted ray propagation across the 121° horizontal field of view through the re-optimized 39 mm cylindrical lens spacing and the Thorlabs C340TMD-B aspheric objective.




Figure 6.2.3 Matrix spot diagram mapping the final system configuration performance at 850 nm, tracking geometric blur radii from the on-axis center out to the extreme 55° horizontal field half-angle and 25° vertical field half-angle on the image plane.

Chapter 7 – Software Design

7.1 Software Design Overview

The software subsystem of the A.U.R.A. platform is responsible for image acquisition, object detection, system decision logic, communication between controllers, and user interface output. Unlike a traditional embedded system that performs simple control tasks, this project requires real-time image processing and machine learning inference while simultaneously managing hardware signals and maintaining safe operating conditions. Because of this, the software design must separate high-level perception from low-level control to ensure reliable operation.

The main processor executes the detection and decision algorithms, while distributed microcontrollers perform deterministic hardware actions. This architecture allows the system to continue operating safely even if the main processor is temporarily busy or restarting. The software must also ensure that the vehicle cannot move unless the detection system confirms that the area under the chassis is clear.

The original modular structure was retained, but the implementation was refined to use two Raspberry Pi 4B camera nodes, a central Raspberry Pi 5, YOLOv11 Nano detection, a two-

out-of-three frame decision rule, Wi-Fi command communication with the ESP32, and an LCD interface with defined operating states.

The software design must satisfy the following requirements:

- Capture images from front and rear cameras
- Perform preprocessing and object detection
- Communicate detection results to control modules
- Activate alarm and safety interlock signals
- Display system status to the user
- Maintain safe operation during startup and failure conditions

To meet these requirements, the software is divided into several independent modules that communicate through defined interfaces. Figure 7.1.1 illustrates the Use Case diagram for the A.U.R.A. system. The diagram identifies the primary actors and their interactions with the system, including system initialization, monitoring, detection, and safety response. It provides a high-level view of system functionality from the user perspective.




Figure 7.1.1 UML Use Case Diagram

7.2 Software Architecture

The A.U.R.A. software is organized using a layered architecture that separates perception, decision making, communication, and hardware control. This structure simplifies debugging and allows individual components to be tested independently.

The top layer is the perception layer, which captures images and runs the detection algorithm. The decision layer analyzes the detection results and determines whether the vehicle is safe to move. The communication layer exchanges messages between the main processor and the microcontrollers. The control layer sends commands to hardware devices such as the alarm, actuators, and interlock.

This layered approach prevents the detection code from directly controlling hardware, which reduces the chance of errors affecting safety-critical signals. Instead, the main processor sends commands to the control modules, which then perform the required action.

The software architecture includes the following main modules:

- Camera capture module
- Image preprocessing module
- Object detection module
- Decision logic module
- Communication module
- Display interface module
- Safety control module

Figure 7.2.1 illustrates the overall software architecture and communication structure of the A.U.R.A. system. The diagram shows the relationship between the perception layer, processing unit, main controller, and distributed sensor nodes, as well as the communication protocols used between each component.




Figure 7.2.1 A.U.R.A. Software Architecture Diagram




Figure 7.2.1.1 Revised A.U.R.A. Software Architecture Diagram

The revised architecture shows the front and rear Raspberry Pi 4B nodes as separate MJPEG stream sources, the Raspberry Pi 5 as the central detection and decision processor, YOLOv11 Nano as the final model, and Wi-Fi as the runtime command connection between the Pi 5 and ESP32.

Figure 7.2.2 presents the Class diagram of the A.U.R.A. software system. The diagram illustrates the internal structure of the software, showing the major modules and their relationships. Each class represents a functional component responsible for image acquisition, preprocessing, detection, decision-making, communication, and system output. The relationships between classes demonstrate how data flows through the system.




Figure 7.2.2 Class Diagram



The revised module structure reflects the implemented software responsibilities. UART-specific status and command functions were replaced with Wi-Fi stream, safety-command, and acknowledgment functions. The detection module also identifies YOLOv11 Nano, and the decision-logic module includes the rolling three-frame windows and two-out-of-three evaluation rule.

7.3 Camera Capture and Image Acquisition

The camera capture module is responsible for acquiring images from the front and rear camera sensors connected to the Raspberry Pi. The cameras use the CSI interface, which allows high-speed transfer of image data directly to the processor.

The software must initialize both cameras during startup and verify that each device is available before detection begins. Because the processor has limited performance, the software may not process both camera streams simultaneously. Instead, frames may be captured sequentially to reduce CPU usage.

Each frame is prepared for further image processing operations. The capture module must also handle errors such as disconnected cameras or failed initialization. If the camera cannot provide valid frames, the system must remain in a safe state and prevent vehicle movement.

The capture module provides image frames to the preprocessing module through a defined interface so that the rest of the software does not need to interact directly with the camera driver.

In the final implementation, each IMX178M camera connects through CSI to a dedicated Raspberry Pi 4B. Each Pi 4B creates an MJPEG stream, and the Raspberry Pi 5 receives the front and rear streams over Wi-Fi. This distributed setup replaced the earlier assumption that the cameras would connect directly to the main Pi.

7.4 Image Preprocessing Pipeline

The image preprocessing pipeline performs the core image processing operations required to prepare input data for object detection. Before running object detection, each frame must be processed to match the input format required by the detection model. The preprocessing pipeline uses the OpenCV library to perform operations such as resizing, normalization, color conversion, and filtering.

The detection model expects a fixed input resolution, so each frame is scaled to the correct size. Additional preprocessing steps such as noise reduction, contrast adjustment, and brightness correction are applied to enhance important features in the image, particularly under near-infrared illumination conditions. These steps improve detection accuracy and ensure consistent input quality.

The preprocessing stage must be fast enough to keep up with the camera frame rate. Because the Raspberry Pi has limited processing power, unnecessary operations are avoided to maintain real-time performance. After preprocessing, the frame is converted into a format suitable for the YOLOv11 model and passed to the object detection module.

The final camera-node pipeline uses FFmpeg to crop, scale, convert, and encode each camera feed as an MJPEG stream. The Raspberry Pi 5 decodes the incoming frames for the LCD and YOLOv11 Nano inference. Processing was limited to the operations required by the final runtime software.

7.5 Object Detection Using YOLOv11

The A.U.R.A. system uses a YOLOv11 object detection model to identify animals beneath the vehicle. Object detection is required instead of simple classification because the system must determine not only whether an object exists, but also where it is located within the image.

The YOLOv11 framework was selected because it supports real-time detection and can run on embedded processors without requiring a dedicated GPU. The nano variant of the model is used because it provides the best balance between speed and accuracy for the selected hardware.

During operation, each preprocessed frame is passed to the model, which returns bounding boxes, class labels, and confidence values. The software analyzes these results to determine whether a detected object represents a possible animal.

Because the system uses a custom camera position and lighting conditions, a custom dataset must be created and annotated. The training process is performed offline, and the resulting model file is loaded by the runtime software.

The detection module must complete inference quickly enough to avoid delaying the startup sequence. If inference takes too long, the system may appear unresponsive, so the software limits the frame rate to maintain consistent performance.

YOLOv11 Nano was used because it was the YOLOv11 version offered through the selected Roboflow training workflow. The model returns bounding boxes, labels, and confidence values for the processed frames.

7.6 Decision Logic and Safety Rules

The decision logic module receives the detection results and determines whether the system should allow the vehicle to move. This module is responsible for enforcing the safety behavior of the A.U.R.A. platform.

If the detection model reports an object with sufficient confidence, the system enters a hazard state. In this state, the alarm is activated, and the interlock remains engaged. The software continues running detection until no objects are detected for a defined period of time.

To reduce false detections, the decision logic may require multiple consecutive frames to confirm the presence of an object. This prevents noise or temporary shadows from triggering the alarm.

The decision module must always default to a safe state. If the detection module fails or produces invalid data, the system must assume that the area is unsafe. The vehicle is only allowed to move when the detection module confirms that the area is clear. This design ensures that software errors cannot accidentally disable the safety feature.

The final software uses a rolling three-frame window for each camera. A hazard is confirmed when either camera contains qualifying detections in at least two of its three

most recent processed frames. After a hazard has been activated, CLEAR is sent only when both cameras fall below the two-out-of-three hazard threshold. This prevents one isolated detection result from immediately changing the system state.

Figure 7.6.1 presents the UML activity diagram representing the startup verification process implemented in the A.U.R.A. system. The diagram shows initialization, image acquisition, preprocessing, object detection, hazard response, and repeated rechecking only while a hazard remains present. Once the area is verified clear, the system permits vehicle movement and the startup safety process ends.




Figure 7.6.1 UML Activity Diagram for Startup Safety Verification




Figure 7.6.1.1 Revised UML Activity Diagram for Startup Safety Verification

The revised activity diagram includes the user START action, command acknowledgments, front and rear MJPEG streams, YOLOv11 Nano processing, separate three-frame windows, the two-out-of-three hazard rule, continued monitoring during a hazard, and the final Clear state.

Figure 7.6.2 illustrates the major operating states of the A.U.R.A. software system during the startup safety verification process. The state diagram shows how the system transitions

between initialization, startup monitoring, hazard response, safe-clear verification, and fault handling before vehicle movement is permitted.




Figure 7.6.2 Software State Diagram




Figure 7.6.2.1 Revised Software State Diagram

The revised state diagram reflects the implemented System Ready, Deploying, Detecting, Hazard Active, and Clear/Detection Paused states. A confirmed hazard moves the system from Detecting to Hazard Active, both cameras confirming clear moves it to Clear, and another START action begins a new deployment cycle.

7.7 Communication Between Processor and Controllers

The communication module handles message exchange between the main processor and the distributed microcontrollers. Communication between the Raspberry Pi and the main ESP32 controller uses a wired UART interface for reliable, low-latency data transfer. Communication between the main ESP32 controller and the distributed sensor nodes is performed over Wi-Fi, allowing flexible placement and reducing wiring complexity. The Raspberry Pi sends commands such as alarm activation, interlock release, and status updates. The microcontrollers respond with confirmation messages to ensure that the command was received.

The software must handle connection loss, delayed messages, and packet errors. If communication fails, the system must remain in the safe state until the connection is restored.

To simplify the protocol, each message includes a command identifier and a status code. The communication module checks that the message format is correct before executing the command.

Using a dedicated communication layer prevents the detection code from interacting directly with hardware signals.

The final prototype changed the Pi 5-to-ESP32 runtime connection from UART over USB to Wi-Fi because the Raspberry Pi did not reliably read the ESP32 through the USB connection. The ESP32 operates as a network server, and the Pi 5 sends START, HAZARD, CLEAR, and RESET commands. The ESP32 returns the corresponding acknowledgment before the Pi 5 accepts the requested action as complete. The camera streams and safety-command traffic use separate network connections: the Pi 4B nodes send MJPEG video, while the Pi 5 and ESP32 exchange short command and acknowledgment messages.

Figure 7.7.1 illustrates the communication structure between the Raspberry Pi, the main ESP32 controller, and the distributed ESP32 modules. The diagram highlights the use of UART for local communication and Wi-Fi for communication with distributed sensor nodes.




Figure 7.7.1 Communication Between Processor and Controllers




Figure 7.7.1.1 Revised Communication Between Processor and Controllers

The original diagram shows the planned UART connection. The revised diagram shows the final Wi-Fi link between the Pi 5 and ESP32 because the Raspberry Pi did not reliably read the ESP32 over USB. It also shows the separate front and rear MJPEG streams sent from the Pi 4B nodes to the Pi 5.

Figure 7.7.2 illustrates the sequence diagram for communication between the Raspberry Pi and ESP32 controllers during the startup safety verification sequence. The diagram shows the order of interactions involved in image processing, transmission of detection results, and execution of control commands before vehicle movement is allowed. It highlights the use of UART for local communication and Wi-Fi for communication with distributed nodes.




Figure 7.7.2 Sequence Diagram




Figure 7.7.2.1 Revised Sequence Diagram

The revised sequence diagram reflects the final order of communication: START and ACK_START, camera deployment, connection to both MJPEG streams, YOLOv11 Nano detection, HAZARD or CLEAR commands, command acknowledgments, ESP32 hardware control, and LCD state updates.

7.8 Display and User Interface Software

The display module provides visual feedback during the startup safety verification process. The display may show camera images, detection results, or a simple status indicator.

The software updates the display only after detection is complete to avoid slowing down the processing pipeline. Drawing operations are performed using OpenCV or the system graphics library.

The interface must clearly show whether the vehicle is safe to move. A simple indicator such as green for clear and red for hazard allows quick verification during testing.

The display module also provides debugging information, such as detection confidence values and communication status. This information is useful during development but may be removed in the final demonstration.

The final interface is rendered directly with OpenCV on the seven-inch touchscreen. The display shows System Ready, Deploying, Detecting, Hazard, and Clear states. During Hazard, both camera feeds and current detections are displayed. During Clear, detection pauses but the live camera feeds remain active.

7.9 Startup Sequence and Initialization

When power is applied, the software must initialize all modules before normal operation begins. The startup sequence follows a defined order to ensure that each component is ready.

First, the operating system starts and launches the detection program automatically. The program checks that the cameras are connected and that the communication network is available. Next, the detection model is loaded into memory.

After initialization, the system performs a safety check by capturing images and running detection. If the area is clear, the interlock is released. If an object is detected, the alarm is activated and the interlock remains engaged.

The system does not allow normal operation until all modules report ready status. This prevents the vehicle from moving before the detection system is fully active.

In the final demonstration software, initialization prepares the model, display, and communication resources, but detection does not begin automatically. The system remains on the System Ready screen until the user presses START. The Pi 5 then sends the START command, waits for ACK_START, and begins the deployment and camera-stream sequence.

7.10 Failure Handling and Safe Operation

Because the system controls a safety feature, the software must handle failures in a predictable way. Any error must result in the system remaining in the safe state.

If the camera fails, the system cannot verify that the area is clear, so the interlock stays active. If communication fails, the system assumes that hardware status is unknown and prevents movement. If the detection program crashes, the interlock is not released.

The software also includes timeout checks to ensure that commands are not lost. If a controller does not respond within a defined time, the command is repeated or the system enters a safe state.

This approach ensures that unexpected errors do not disable the safety feature.

7.11 Software Design Summary

The software design for the A.U.R.A. system uses a layered architecture that separates perception, decision logic, communication, and hardware control. This structure allows the system to perform real-time object detection while maintaining reliable control of safety-critical signals.

By using OpenCV for preprocessing, YOLOv8 for detection, UART for communication between the Raspberry Pi and the main ESP32 controller, and Wi-Fi for communication with distributed sensor nodes, the software can operate within the limits of the embedded hardware while still meeting the requirements of the project.

The design emphasizes safe operation, modular structure, and compatibility with the selected hardware platform, making it suitable for implementation within the time constraints of the senior design project.

The final implementation uses two Raspberry Pi 4B MJPEG camera nodes, a Raspberry Pi 5 running YOLOv11 Nano, a rolling two-out-of-three frame decision rule, and Wi-Fi communication between the Pi 5 and ESP32. YOLOv11 Nano replaced the originally referenced YOLOv8 because YOLOv8 was not offered through the selected Roboflow workflow. Wi-Fi replaced the planned UART-over-USB runtime connection because the Raspberry Pi did not reliably read the ESP32 through USB.

Chapter 8 – System Fabrication/Prototype Construction

This chapter details the physical realization of the system, covering the layout and assembly of all printed circuit boards (PCBs) produced during the project. The fabrication process followed a structured design-to-manufacture workflow using Fusion 360, with careful attention to component placement, trace routing, ground plane pours, and mechanical constraints. Each subsection below addresses a specific board or power regulation stage, documenting design decisions, key features, and relevant layout views.

8.1 Primary PCB

The Primary PCB serves as the central unit within our system, containing connections with multiple subsystems necessary for the full functionality of our system. With an ESP32-S3 as the primary low-level control unit, this controls all the low-level systems needed to act as a deterrent such as the high-power LED strobe light and the Ultrasonic alarm.

Given that the 3W LEDs needed to produce the strobe are preplaced onto an aluminum starboard heatsink, necessary space was given to the connections between the logic-level MOSFET acting as a constant current driver and the connection to the 12V input and 5W resistors needed to properly strobe the lights without failure. In addition, the aluminum heatsinks will be attached to the PCB through thermal paste, as this preserved both the integrity of the PCB and the 3W LEDs attached to the heatsinks. The LEDs will be connected to the circuitry of the PCB through two screw-in ports, one connected to the MOSFET the other to the 5W resistors.

The ultrasonic alarms take up the most space on the board as two different kinds of piezoelectric transducers were necessary to fit with our 30kHz to 45kHz frequency range as both arrays of transducers need to be properly driven, two H-bridge amplifiers with proper drivers are placed onto the board. The ultrasonic transducers themselves are placed along the edges of the PCB to allow for an output of each kind on all sides of the board, in addition to another pair pointed downwards for even more coverage.

The Primary PCB also contain the connection to our on-board computer through a USB-C port, due to high power demand of Raspberry PI5, our standard 5V switching regulator would not be able to provide enough current to fully power the on-board computer, and so, in addition to 3.3V and 5V switching regulators, an additional 5.1V, 5A regulator circuit will be attached to this board through a 4-pin connector.

In addition to the regulator boards that will all be connected through header pin connectors, our 12V power supply will be connected to this board through a screw-in port near the bottom edge of the board. These 12 volts will be sent through all three regulators, as well as the multiple systems on this board that need 12 volts to function properly, including both ultrasonic alarms and their drivers, the strobe light, and finally the JD1912 SPST automotive relay system serving as our motor interlock.

In the final version of the main PCB, multiple changes were made from the original. Instead of USB-C, we opted for USB-A as USB-C was causing major power issues when attempting to power the PI. Another change was that strobe light was replaced with an array of 15 indicator LEDs. The MCU was changed to have an added RESET and BOOT buttons connected to the IO0 and EN pins. The H-bridge amplifiers were changed to be in line with the final revised version of the ultrasonic alarms. Finally, the CP2102 design was changed to fall in line with the schematic from the datasheet.




Figure 8.1 Main PCB Design




Figure 8.1.1 Revised Main PCB Design

8.2 Secondary PCB

The secondary PCB was designed to serve as the primary sensing and motor control interface for the system. The board centers on a Waveshare ESP32-P4 module, which provides the processing core for camera interfacing, time-of-flight distance sensing, and motor drive control. All component placement and signal routing were completed in Fusion 360.

The ESP32-P4 interfaces with a Sony IMX178M global shutter camera via a 30-pin FPC connector (Molex 524353071, 0.5 mm pitch) using a MIPI CSI-2 lane configuration consisting of one clock differential pair (MIPI_CK_P/N) and two data lanes (MIPI_L0_P/N, MIPI_L1_P/N). These high-speed differential traces were length-matched and routed with controlled impedance to minimize signal integrity degradation. Camera register configuration is handled over I2C (CAM_SDA = GPIO0, CAM_SCL = GPIO1) using the VEYE register map.

A VL53L0X time-of-flight sensor is integrated onto the board via a 6-pin footprint, sharing the I2C bus with the camera. Its XSHUT pin is driven by GPIO2 to allow software-controlled initialization. The Cytron MD10C motor driver is connected via a 3-pin header

(PWM = GPIO3, DIR = GPIO4, GND), providing bidirectional brushed DC motor control through a simple PWM and direction signal interface.

Power is supplied to the secondary PCB through two LM2575 step-down switching regulator boards, delivering regulated 3.3 V and 5 V rails respectively. Each regulator board connects to the main PCB via 4-pin headers. Polygon pours were applied to both top and bottom copper layers to establish low-impedance ground planes, reduce electromagnetic interference, and improve thermal dissipation. All connectors and header pins are clearly labeled with reference designators and signal identifiers to simplify assembly and debugging.

The secondary PCB was changed from having an ESP32-P4 to the same ESP32-S3 used for the main. Additionally, a Raspberry PI4 was used to integrate our image sensor rather than a connector tied to the ESP32. Components were rearranged to have a more efficient layout for the final version.




Figure 8.2 Secondary PCB Design




Figure 8.2.1 Revised Secondary PCB Design

8.3 Regulators

To simplify fault isolation and allow each voltage rail to be tested and replaced independently, the 3.3 V and 5 V regulators were implemented as separate modular regulator boards rather than integrated directly onto the main PCB. Each board is built around an LM2575 switching regulator and connects to the system via a 4-pin header. This modular approach improves maintainability and reduces the risk of a single power failure affecting the entire system.

8.3.1 3.3 V Regulator Board

The 3.3 V regulator board steps down the input supply to provide a stable 3.3 V rail for the ESP32-P4 module, the VL53L0X ToF sensor, and associated logic-level circuitry. The LM2575-3.3 fixed-output variant is used, with input and output bypass capacitors and a catch diode selected per the datasheet recommendations for the expected load current.

8.3.2 5 V Regulator Board

The 5 V regulator board steps down the input supply to power peripherals that require a 5 V supply, including the Cytron MD10C motor driver logic and any other 5 V-rated components on the secondary PCB. The LM2575-5.0 fixed-output variant is used, with the same passive component selection approach as the 3.3 V board to maintain design consistency across both regulator modules.




Figure 8.3 3.3V and 5V Regulator Board Design

8.3.3 5.1V 5A Regulator Board

Unlike the 5V and 3V regulator boards, this board only serves one purpose, to power our Raspberry PI5, on-board computer. While 5V and 3V regulators are common and have plenty of fixed output ICs to choose from, not many are available in the 5.1V and 5A range, and so it was necessary to find an adjustable voltage regulator, and so the AP64501 synchronous buck converter was chosen as the IC for this regulator. The schematic and PCB layout was modeled after the typical application diagram found in the IC's datasheet, with a difference being the resistors R1 and R2, using a voltage divider equation, it was

necessary to have R1 be 118k Ω and R2 be 22.1 k Ω as this allows for a 5.1V output from the regulator.




Figure 8.4 5.1V, 5A Regulator Board Design




Figure 8.5 Revised Regulator Board Design

In the final version of the regulators, the 3.3V and 5V regulator were replaced with the same design as the adjustable regulator, the only difference being that R2 and R1 were changed to fit each voltage requirement. Additionally, the adjustable regulator switched from a AP64501 to a LM2679 regulator.

Chapter 9 - System Testing and Evaluation

9.1 Optoelectronics Feasibility and Study Testing

The feasibility of the proposed optoelectronic system was evaluated by examining the interaction between illumination, optical transmission, and sensor response. The objective of this evaluation is to determine whether the system is capable of generating and capturing sufficient signal for object detection within the underbody region under realistic operating conditions. This assessment focuses on whether usable image data can be produced, rather than providing a fully optimized or quantitatively validated system.

The system operates using active near-infrared illumination at approximately 850 nm in conjunction with a CMOS image sensor based on the Sony IMX178 architecture. Silicon-based sensors exhibit measurable sensitivity in the near-infrared range, allowing detection of reflected infrared illumination without requiring specialized detector materials. This enables the system to utilize cost-effective and widely available imaging hardware while still operating outside of the visible spectrum. At short working distances, the intensity of reflected near-infrared light is expected to be sufficient for detection. This reduces the reliance on high sensor gain or long exposure times and allows the system to capture stable image data under controlled illumination. In addition, the selected sensor provides high spatial resolution due to its small pixel size, enabling the detection of fine features within the compressed field of view.

The use of near-infrared illumination provides an important advantage in terms of contrast. Materials commonly found in the environment, such as asphalt and painted surfaces, exhibit different reflectance characteristics compared to biological materials such as fur and skin. This difference allows objects to be distinguished from the background based on intensity variations rather than relying solely on visible appearance. Active illumination further improves consistency by reducing dependence on ambient lighting conditions, allowing the system to operate in a wider range of environments.

The optical subsystem plays a critical role in preserving the reflected signal and transmitting it to the sensor. The system is designed to compress the horizontal field of view while maintaining sufficient spatial information for detection. Simulation results demonstrate that the system is capable of propagating rays across the full field without complete loss of signal. Although some reduction in ray density occurs at intermediate stages of the optical path, sufficient rays reach the image plane to maintain image formation. This indicates that the system is capable of preserving the spatial structure of the scene and delivering usable image data to the sensor.

The theoretical optoelectronic framework was fully validated through direct laboratory prototype testing of the compiled physical hardware stack. Active evaluation was conducted using the integrated system containing the distributed Chanzon 850 nm through-hole LED arrays, the custom optomechanical telescoping enclosure, the corrected optics assembly, and the Sony IMX178M monochrome sensor linked directly to a Raspberry Pi processing module. Live video data was successfully streamed wirelessly from the vehicle chassis to a laptop telemetry window in real time under simulated low-light conditions.

To experimentally characterize spatial sampling, uniform aspect compression, and surface contrast preservation, a physical calibration target featuring a 1.0-inch horizontal by 0.5-inch vertical grid structure was positioned beneath the chassis at the designated 6.5-inch working distance. The captured image frames directly confirm the underlying horizontal compression mechanics across the active pixel matrix. While the baseline prototype imagery displays visible pixelation, corner focus blurring, and a minor optical centration offset directly tied to initial 3D-printing tolerances, the structural contrast boundaries remain highly distinct. The target material is clearly distinguished from the dark asphalt ground plane, proving that the hardware transmission loop preserves sufficient geometric detail and structural data for downstream YOLO object detection tracking.

9.1.1 Early Proof of Concept Demonstration

To complement the simulation-based analysis, a simplified physical demonstration was conducted to observe system behavior under real-world conditions. Due to limitations in available components, the demonstration used spherical lenses to approximate the behavior of the optical system and a consumer imaging device to capture results. While this setup does not fully replicate the intended final configuration, it provides valuable qualitative insight into system performance.



Figure 9.1.1 Simplified physical demonstration of optical system behavior

(a) Baseline scene without optical elements showing unmodified field of view (b) Scene viewed through a concave objective and convex eyepiece demonstrating increased horizontal field of view.

As shown in Figure 9.1, the physical demonstration utilized a simplified optical configuration consisting of a negative (concave) objective lens followed by a positive (convex) eyepiece. This arrangement resembles a Galilean-style system and was used to approximate the behavior of the image-side optics in a compact form. The purpose of this

setup was to observe how angular field information is modified as light propagates through a multi-element optical system.

Figure 9.1(a) shows the baseline view of the scene without any optical elements present, only an empty placeholder lens mount at the same position the first lens is in (b). This image represents the unmodified field of view and serves as a reference for comparison. In contrast, Figure 9.1(b) shows the image captured through the optical system after the lenses were introduced.

When viewed through the lens system, a noticeable increase in the horizontal field of view is observed. This indicates that the optical configuration is effectively remapping angular field information, allowing a wider portion of the scene to be captured within the viewing aperture. The combination of a negative objective and positive eyepiece produces this effect by altering the angular distribution of incoming rays, resulting in an expanded perceived field.

Although this demonstration does not include cylindrical elements and does not fully replicate the final optical system, it provides qualitative validation of the underlying concept. The observed increase in horizontal field coverage supports the feasibility of using optical elements to manipulate and compress the field for imaging applications.

9.1.2 Real Lens System Evaluation

The second stage of testing transitioned from qualitative observation to a quantitative evaluation within Ansys Zemax OpticStudio. This stage verified the performance of the Thorlabs and Edmund Optics hardware assembly at the target 850 nm wavelength. A critical distinction in this evaluation is the implementation of real cylindrical modeled optics, replacing the simplified spherical approximations used in the initial proof of concept. By modeling the specific toroidal curvatures and center thicknesses of the LK1336L1-B and LJ1909L1-B elements, the simulation provides a definitive assessment of the system's ability to resolve the 110° underbody field.




Figure 9.1.2 Geometric Image Simulation (Anamorphic Mapping)

The primary metric for validating the anamorphic relay is the Geometric Image Simulation shown in Figure 9.1.2. This evaluation utilized a high-density grid target to visualize mapping accuracy and spatial compression across the full 110° field. By analyzing the vertical elongation of the grid squares, the simulation provides the visual verification of the 2.04:1 horizontal squeeze. This ensures the wide-angle underbody scene is successfully compressed into the 7.41 mm active width of the Sony IMX178 sensor, achieving the target 81.9% sensor utilization. While the image exhibits characteristic barrel distortion at the field extremities, the continuity of the grid lines proves that the horizontal spatial data remains resolved and usable for animal detection.

The simulation results indicate an 8.429 mm horizontal image width. While the Sony IMX178 sensor possesses a physical width of 7.41 mm, the system was designed to produce a controlled over-fill of the sensor plane. This provides a 0.5 mm engineering buffer on each side of the CMOS active area to account for mechanical mounting misalignments. The diagnostic warning indicating that not all rays reached the image plane confirms that the 65.0 mm working distance successfully captures the full 110° field. This hardware configuration maintains an 87.725% optical efficiency, serving as the primary quantitative metric for energy throughput. The localized convergence observed in the ray trace verifies that the ray bundle reaches the sensor plane within the 2.0 mm mechanical

air gap. This configuration provides the necessary clearance for the 850 nm interference filter, proving that the A.U.R.A. hardware is a build-ready specification.

9.1.3 Optoelectronic Feasibility Summary

At the system level, the combination of active illumination, optical field compression, and CMOS sensing supports the feasibility of the proposed approach. The illumination provides a consistent signal source, the optical system preserves and maps that signal, and the sensor is capable of capturing the resulting image. Together, these components enable the system to detect meaningful differences in reflectance and spatial structure within the underbody region.

However, it is important to note that comprehensive physical optoelectronic testing has not yet been completed. The current evaluation is based on a combination of preliminary physical observations and high-fidelity simulation results. It does not yet include field-measured quantitative data such as signal-to-noise ratio, contrast levels, or detection accuracy in varying weather conditions. Furthermore, while the current simulation confirms static focus, the effects of vehicle motion on dynamic range and exposure stability remain to be characterized. As a result, the conclusions drawn at this stage are focused on geometrical feasibility and optical throughput validation.

Further optoelectronic testing will be conducted once all required components have been received and integrated. This includes the selected camera module, optical elements, and near-infrared illumination sources. Future testing will involve controlled experiments to evaluate image quality, contrast between target and background materials, system response under varying lighting conditions, and overall detection performance. These results will provide a more complete characterization of the system and guide further optimization.

In summary, the current analysis and preliminary testing indicate that the proposed optoelectronic system is capable of generating and capturing usable image data for underbody detection. While additional testing is required to fully validate performance, the results obtained thus far demonstrate that the selected approach is feasible and provides a strong, validated foundation for continued development.

The theoretical validation and simulation phase detailed above was advanced into a full hardware testing phase using the physical prototype assembly. The planned 9 mm lens and Thorlabs LJ1909L1-B cylindrical element were replaced by a pre-mounted Thorlabs C340TMD-B aspheric objective (4.03 mm focal length) and a laboratory-stock Newport CSX200AR.16 fused silica cylindrical lens. This hardware pivot successfully avoided part procurement delays and budget expenditure while protecting system focus constraints. The system was mounted straight-on at a 6.5 inch longitudinal working clearance inside a custom 3D-printed telescoping housing. This custom enclosure was manufactured with internal structural alignment lips to lock the lens pair into a stable 39 mm internal optical gap, which was found during testing to be the optimal spacing required to eliminate front-housing cap vignetting.

Testing was conducted using the integrated hardware loop consisting of the distributed 3mm Chanzon 850 nm through-hole LED board array, the lens stack, and the monochrome Sony IMX178M sensor connected to a Raspberry Pi module. Live near-infrared video data was successfully captured and streamed wirelessly to a laptop window in real time under low-light laboratory conditions. To quantitatively verify aspect ratio remapping, a physical paper calibration target marked with a 1.0 inch horizontal by 0.5 inch vertical grid pattern was placed beneath the chassis footprint (Figure 9.1.3). Software analysis of the center grid box (measuring 105 wide by 86 high pixels) confirmed an empirical spatial resolution of 4.13 px/mm and a measured scene ratio of 2.45:1, yielding a true horizontal optical squeeze factor of 1.64x.

Secondary testing was conducted by placing a stuffed animal target underneath the center of the demonstration vehicle (Figure 9.1.4). While the resulting baseline captures exhibit noticeable pixelation, soft corner focus, and a minor optical centration offset directly caused by initial 3D-printed enclosure tolerances, the structural contrast boundaries remain highly defined. The target material is clearly distinguished from the dark asphalt ground plane, proving that the hardware transmission loop preserves sufficient geometric detail and structural data for future downstream YOLO-based object detection training. The testing also identified localized optical blooming and pixel saturation across active light sources and highly reflective surfaces, providing key performance insights for future lighting shroud iterations.



Figure 9.1.3 Experimental near-infrared image capture of the 1.0 inch by 0.5 inch calibration grid at a 6.5 inch working distance, utilized to extract the 1.64x horizontal squeeze factor and confirm a 4.13 px/mm spatial ground resolution.



Figure 9.1.4 Experimental near-infrared image capture of a stuffed animal target under the demonstration vehicle, showcasing active 850 nm illumination contrast alongside visible optical blooming and pixel saturation across reflective surfaces.

9.2 Hardware Testing

Multiple subsystems of the A.U.R.A system were tested and evaluated using development microcontroller boards and prototyping boards. These hardware-based systems were built using current designs and evaluated for performance and viability for integration into the full system.

9.2.1 Motor Interlock

To test the current design of the motor interlock subsystem, the chosen JD1912 automotive SPST relay was connected to the 2N222A BJT transistor acting as a driver for the relay, in addition a 1N4148 diode was placed parallel to the relay's coils as feedback diode, protecting the ESP32-S3 from voltage or current spikes. An LED in series with a resistor was connected with pins 87 and 30, this LED would give a visual signal that the relay was either open or closed. Finally, a button toggle was implemented to control the behavior of the relay.

When testing the functionality of the system, the relay behaved as intended. When the relay was toggled to be closed, the LED would light up indicating that the circuit was complete, in addition to this, the power supply used to power the relay would spike up in current when the relay was closed. When the relay was toggled to be open, the LED would turn off and the current on the power supply would drop.

After multiple rounds of testing, the motor interlock system continued to behave as intended, proving it to be a functional and viable design for integration into the full system.

9.2.2 3W LED Strobe

The 3 watt was tested in its current design, with it having a logic-level MOSFET acting as a constant current source connected to two pairs of LEDs in parallel with each other and two 1k Ω 5W resistors in series with each pair, powered by a 12V source.

Upon testing this array of four high-power LEDs, the strobe displayed a high-brightness, marking it as a visual irritant, making it a viable design for our final system.

9.2.3 Ultrasonic Alarm

The design of the ultrasonic alarm was tested using an oscilloscope and primarily a frequency of 30.5 kHz. It was tested at each part, first, the frequency generation of the ESP32-S3 was tested through the driver's connections. When tested at the input, the oscilloscope displayed stable 30.5 kHz square waves, indicating that using the ESP32-S3's LEDC peripheral as a frequency generator was indeed a viable design option in the functionality of the ultrasonic alarm.

Next, the H-bridge amplifier design was tested at its output using a probe on each node and by using the oscilloscopes MATH functionality, Channel 2 was subtracted from Channel 1 to see the true behavior of the N-channel MOSFET, H-bridge amplifier.



Figure 9.2.1 shows the display of the oscilloscope at the output nodes, the MATH function being the blue signal. Here it shows that the amplifier is behaving as intended, the original signals are out of phase which is necessary for the amplifier to amplify the original 12V input. The measurements show that the signal has been amplified to the expected 24Vpp while maintaining the 30.5kHz frequency.

After verifying the amplifier, the output was tested with the same method as that of the amplifier, with the math function being used to subtract the two nodes. First a single transducer was tested to see the behavior with only a single output source, following that, all five transducers were tested in parallel.




Figure 9.2.2 Behavior of Output at Single Transducer

In Figure 9.2.2 we see that unlike that of the previous results, the signal is no longer being amplified to 24Vpp, however here it does maintain the 30.5kHz frequency




Figure 9.2.3 Behavior of Output at Transducer Array

Even further now the output voltage has dropped now that the transducer array has been connected. Additionally the frequency of the signal has dramatically changed. This result shows that while the output of the amplifier may be functional on its own, it does not hold up when the capacitive transistor load is connected to the output. Therefore, further testing and improvement of the amplifier design will be conducted to be a fully functional part of the overall integrated system.



Figure 9.2.4 Final 30 kHz Output



Figure 9.2.5 Final 40 kHz Output



Figure 9.2.6 Final 45 kHz output

In the final version of the ultrasonic alarm, the output remained consistent within the frequency range with the only noticeable change being the frequency. As opposed to the final output in SD1, the final outputs from SD2 clearly show a square wave output with a 180° out of phase channel 1 and channel 2 outputs, resulting in the complete signal when channel 2 is subtracted from channel 1.

9.3 Software Testing

The software component of the A.U.R.A. system was tested independently to verify the functionality of the object detection pipeline. A YOLOv11-based model was trained using a labeled dataset of under-vehicle images containing animal presence. The training process was completed using the Roboflow platform, which provided tools for annotation, augmentation, and model validation.

Testing was performed by running sample images through the trained model and observing detection outputs. The model was able to correctly identify animals in multiple test cases while avoiding false detections in most non-animal scenarios. This confirms that the model can distinguish between biological hazards and background features such as ground texture or vehicle components.

Final software testing verified simultaneous reception of the front and rear MJPEG streams, YOLOv11 Nano inference, LCD state transitions, detection overlays, the two-out-of-three frame logic, and Wi-Fi command communication with the ESP32. During a hazard, the interface displayed both live feeds and the current detection output. Once both camera windows confirmed clear, detection paused while the live feeds remained active.

Figure 9.3.1 shows an example of a raw image captured from an under-vehicle perspective. This image demonstrates the typical field of view and environmental conditions used during testing.



Figure 9.3.1 Raw image captured from an under-vehicle camera showing the ground surface and vehicle structure



Figure 9.3.1.1 Finalized demo vehicle structure and lighting

Figure 9.3.2 and Figure 9.3.3 show example outputs from the detection model. In both cases, the model successfully identifies the presence of an animal and displays a bounding box with an associated confidence score. These results demonstrate that the model is capable of detecting hazards under varying lighting conditions and viewing angles.



Figure 9.3.2 Detection output showing identified animal beneath the vehicle with confidence score




Figure 9.3.3 Additional detection example demonstrating model performance under different lighting conditions



Figure 9.3.3.1 Finalized demo vehicle detection

In addition to detection testing, the system was evaluated for its ability to interface with the display module. Due to the current stage of system integration, full real-time detection was not executed directly on the Raspberry Pi. The LCD interface was tested by manually triggering system states to simulate detection outputs. This approach allowed verification of the user interface and system response without requiring full real-time integration.

Figure 9.3.4 shows the LCD interface during a hazard condition. When the hazard state is triggered, the system displays a warning message instructing the user to check under the vehicle before moving, which is for demonstration purposes. This demonstrates the intended system response once full detection integration is implemented.




Figure 9.3.4 LCD display showing hazard detection state and warning message.




Figure 9.3.4.1 Finalized LCD display showing hazard detection state and warning message.




Figure 9.3.4.2 Finalized LCD display showing clear state.

In the current system setup, a live camera feed is not implemented. Testing was performed using static images and simulated detection outputs to validate system behavior. This approach allowed verification of the detection pipeline and user interface without requiring full real-time integration.

For demonstration purposes, a live camera feed will be incorporated in future iterations to provide visual feedback of the monitored area and to illustrate system functionality. This feature is intended only for demonstration and testing.

In the intended real-world implementation of the system for autonomous vehicles, a continuous live feed is not required. Instead, the system will rely on automated detection and response mechanisms without requiring direct user interaction.

9.4 Detection Performance Evaluation

The performance of the detection model was evaluated using standard object detection metrics. The trained model achieved a mean Average Precision (mAP@50) of 91.7%, with a precision of 88.0% and a recall of 92.1%. The resulting F1 score was approximately 90.0%, indicating a strong balance between detection accuracy and reliability.

The relatively high recall value suggests that the model is effective at identifying most instances of animal presence, which is critical for a safety-focused system. At the same time, the precision value indicates that the model maintains a low rate of false detections, helping to prevent unnecessary system triggers.

These results demonstrate that the YOLOv11-based model is suitable for under-vehicle hazard detection in controlled conditions. However, it is important to note that these metrics were obtained using a test dataset and may vary when deployed in real-world environments with different lighting, surfaces, and sensor configuration.

The final YOLOv11 Nano model achieved an mAP@50 of 96.4%, precision of 90.5%, recall of 95.0%, and an F1 score of 92.7% on the validation set. Recall was especially important because failing to detect an animal represents a greater safety risk than temporarily producing a false warning.

9.5 System-Level Testing

Basic system-level testing was conducted to verify interaction between the software and hardware components. Due to the current stage of system integration, full real-time detection was not executed directly on the Raspberry Pi. Instead, system behavior was validated by manually simulating detection outputs and observing the corresponding system response.

The LCD interface was used in the demonstration setup to represent system states, including “Clear” and “Hazard Detected” conditions. These states were triggered manually to simulate the expected behavior of the detection system. When the hazard state was activated, the display updated correctly, confirming that the user interface and control logic function as intended.

These tests demonstrate that the system architecture supports integration between sensing, processing, and response components at a design level. While full real-time deployment remains a future goal, the current implementation verifies that the overall system structure is valid and capable of supporting complete system integration.

Chapter 10 - Administrative Content

10.1 Budget

The current estimated budget for the A.U.R.A. project was initially set at \$850, agreed upon by the team as a baseline for system development and prototyping. As component selection has been finalized, the projected total cost has increased to \$1,600, primarily driven by the high cost of precision optical components required for the anamorphic lens system and image sensors. The electrical and computer engineering subsystems remain within reasonable cost ranges at \$282.87 and \$198.00 respectively. The optical subsystem accounts for most of the total budget at \$1,102.88, driven by the cost of precision lenses, achromatic doublet objectives, and global shutter image sensors required to meet the system's imaging performance specifications. The team will explore cost reduction opportunities during Senior Design 2 such as component substitution where performance requirements allow.

During Senior Design II, the final architecture added two Raspberry Pi 4B camera nodes and the supporting microSD cards, power connections, and cables required for the front

and rear camera streams. These additions increased the computer engineering subsystem cost compared with the original estimate. The change from UART to Wi-Fi did not require a separate communication module because Wi-Fi is already built into the Raspberry Pi and ESP32 hardware. The final project cost and component quantities are documented in the revised bill of materials that follows.

10.2 Bill of Materials

The components listed in Table 10.2.1 represent the selected hardware for the current system design. While some components may be refined as testing progresses, the table provides a solid foundation for system implementation. Component selection was based on performance requirements, system compatibility, and budget constraints to ensure reliable operation without interfering with other subsystems.

Table 10.2.1 Bill of Materials

Component	Part Name	Quantity	Price Per Unit	Total Cost
Electrical Engineer				
MCU	ESP32-WROOM- 1U	1	\$5.71	\$5.71
MCU	ESP32-P4-Module Core Board	2	\$10.99	\$21.98
USB to UART IC	CP2102N-A02-GQFN24R	2	\$4.48	\$8.96
RF Antenna	Adafruit 2.4 GHz Mini Flexible Wi-Fi Antenna	3	\$2.50	\$7.50
Ultrasonic Transducers	40kHz- MA40S4S	5	\$5.19	\$25.95
	30.5 kHz - CUSA-T80-130-3000-TH	5	\$1.93	\$9.65
H-Bridge Driver	IRS2003STRPBF – Half-Bridge Driver IC	2	\$1.23	\$2.46
N-Channel MOSFET	IRLZ44NPBF- MOSFET N-CH 55V 47A TO220AB	6	\$1.78	\$10.68
Motor	Electric Mini Linear Actuator	2	\$44.96	\$89.92
3W LED	JH-3535W12S40-T8B-A	4	\$2.06	\$8.24
Hall-Effect Sensor	55100-2M-02-A	1	\$6.99	\$6.99
Proximity Sensor	VL53L0X	3	\$4.33	\$12.99
Stepper Motor	Cytron 13A DC Motor Driver MD10C	2	\$15.71	\$31.42
Buck Converter	LM2596 DC-DC	4	\$4.99	\$19.96
3.3V Regulator	LM2575-3.3	3	\$1.72	\$5.16
5V Regulator	LM2575-5.0	3	\$1.72	\$5.16
5.1V Regulator	AP64501SP-13	1	\$2.15	\$2.15
Automotive Relay	JD1912	1	\$7.99	\$7.99

Optical Engineer				
Anamorphic Lens System	Plano-Concave Cylindrical Lens	2	\$149.98	\$299.96
Anamorphic Lens System	Plano-Convex Cylindrical Lens	2	\$89.61	\$179.22
Objective Lens	9mm Achromatic Doublet Lens	2	\$102.00	\$204.00
IR Filter	850 nm Narrow Bandpass Filter	2	\$12.00	\$24.00
Image Sensor	MV-MIPI-IMX178M	2	\$169.99	\$339.98
NIR LEDs	OSRAM Oscon Black 850 nm LED	12	\$1.81	\$21.72
Optical Window	50 mm AR-Coated Glass Window	2	\$5.00	\$10.00
Housing	3D Printing Filament (Acrylonitrile Styrene Acrylate)	1	\$24.00	\$24.00
Computer Engineer				
Single Board Computer	Raspberry Pi 5 8GB	1	\$125	\$125
Micro SD card	64GB	1	\$0	\$0
TFT-LCD	7-inch RPi LCD (B)	1	\$73	\$73
Prototyping Board	Breadboard	1	\$0	\$0
Total Cost				
Electrical Engineer			\$282.87	
Optical Engineer			\$1102.88	
Computer Engineer			\$198.00	
Grand Total			\$1583.75	

This is our finalized bill of materials; cost may change as we progress and make changes in Senior Design 2.

Table 10.2.1.1 Revised Bill of Materials

Component	Part Name	Quantity	Price Per Unit	Total Cost
Electrical Engineer				
MCU	ESP32-S3	3	\$5.71	\$17.13
USB to UART IC	CP2102N-A02-GQFN24R	2	\$4.48	\$8.96
Ultrasonic Transducers	40kHz- MA40S4S	5	\$5.19	\$25.95
	30.5 kHz - CUSA-T80-130-3000-TH	5	\$1.93	\$9.65
H-Bridge Driver	IR2104	4	\$4.67	\$18.68

N-Channel MOSFET	IRLZ44NPBF- MOSFET N-CH 55V 47A TO220AB	12	\$1.78	\$21.36
Motor	Electric Mini Linear Actuator	2	\$44.96	\$89.92
Indicator LEDs	150080RS75000	15	\$0.19	\$2.85
Hall-Effect Sensor	55100-2M-02-A	1	\$6.99	\$6.99
Proximity Sensor	VL53L0X	3	\$4.33	\$12.99
Stepper Motor	Cytron 13A DC Motor Driver MD10C	2	\$15.71	\$31.42
Adjustable Regulators	LM2679	7	\$8.33	\$58.31
Automotive Relay	JD1912	1	\$7.99	\$7.99
Optical Engineer				
Anamorphic Lens System	Plano-Concave Cylindrical Lens	2	\$149.98	\$299.96
Anamorphic Lens System	Plano-Convex Cylindrical Lens	2	\$0	\$0
Objective Lens	Aspheric Lens	2	\$104.03	\$208.06
IR Filter	850 nm Narrow Bandpass Filter	2	\$8.99	\$17.99
Image Sensor	MV-MIPI-IMX178M	2	\$169.99	\$339.98
NIR LEDs	OSRAM Oscon Black 850 nm LED	12	\$1.81	\$21.72
Optical Window	50 mm AR-Coated Glass Window	2	\$5.00	\$10.00
Housing	3D Printing Filament (Acrylonitrile Styrene Acrylate)	1	\$0	\$0
Computer Engineer				
Single Board Computer	Raspberry Pi 5 8GB	1	\$125	\$125
Single Board Computer	Raspberry Pi 4 Model B 1GB	2	\$40.34	\$80.68
Micro SD card	64GB	1	\$0	\$0
Micro SD card	16GB	2	\$13	\$26
TFT-LCD	7-inch RPi LCD (B)	1	\$73	\$73
Prototyping Board	Breadboard	1	\$0	\$0
Total Cost				
Electrical Engineer			\$312.20	
Optical Engineer			\$899.66	
Computer Engineer			\$304.68	
Grand Total			\$1502.85	

The original bill of materials represents the components selected during the design phase. The revised BOM records the final hardware used for the completed prototype. The most significant software-related architecture change was the addition of two Raspberry Pi 4B boards, one for each camera node. Each camera node also required its own microSD card and supporting power and network setup. The planned UART connection did not require removal of the USB-to-UART circuitry from the BOM because that interface remained useful for ESP32 programming and debugging, although Wi-Fi became the final runtime communication method.

10.3 Work Distribution

The A.U.R.A. project was developed through a collaborative effort in which each team member was assigned primary responsibility for a specific subsystem while also contributing to integration and testing tasks. In Table 10.3.1, primary responsibilities indicate the main contributor to each task, while secondary responsibilities represent supporting roles and collaborative efforts across subsystems.

Final software responsibilities included configuring the two Raspberry Pi 4B camera nodes, integrating the MJPEG streams, training and deploying the YOLOv11 Nano model through Roboflow, implementing the rolling two-out-of-three frame logic, developing the touchscreen interface, changing the Pi 5-to-ESP32 connection from USB/UART to Wi-Fi, and integrating the acknowledged safety commands with the hardware system.

Table 10.3.1 Distribution of Worktable

Electrical Engineering	Responsibilities
Christopher Nunez-Velasquez	Deterrent System design and implementation (Primary)
	Circuit Design and prototyping (Primary)
	Speaker, strobe light, and stopping mechanism development (Primary)
	Hardware troubleshooting and testing (Primary)
	Software integration support (Secondary-with Kim)
	System-level testing and validation (Secondary-team)
	UART/Wi-Fi hardware implementation (Secondary-with Kim)
Electrical Engineering	Responsibilities

Oussama Benamar Hlimi	Power system design and implementation (Primary)
	Motor and actuation system development (Primary)
	Voltage regulation and power distribution (Primary)
	Mechanical integration and mounting (Primary)
	Hardware troubleshooting and testing (Primary)
	UART/Wi-Fi hardware implementation support (Secondary-with Chris)
	Software integration support (Secondary-with Kim)
	Optical integration support (Secondary-with Theresa)
	System testing (Secondary-team)

Optics and Photonics	Responsibilities
Theresa Dinh	Optical system design and specification (Primary)
	Lens modeling and simulation (Primary)
	Optical testing and calibration (Primary)
	Camera alignment integration (Primary)
	Software integration support (camera interface) (Secondary-with Kim)
	Mechanical integration support (Secondary-with Sam)
	System testing and validation (Secondary-team)

Computer Engineering	Responsibilities
	Software architecture design (Primary)

Kimberly Lugo	Raspberry Pi setup and configuration (Primary)
	Image processing and YOLOv8 implementation (Primary)
	Dataset creation and AI training (Primary)
	Communication protocol design (UART/Wi-Fi messaging) (Primary)
	Software-hardware integration (Primary)
	Hardware integration support (Secondary-with Chris & Sam)
	Optical integration support (Secondary-with Theresa)

10.4 Senior Design Project Milestones

To stay organized and on track throughout the project, the team developed a set of project milestones seen in Table 10.4.1 that outline the major tasks and goals to be completed over the course of Senior Design I and Senior Design II. These milestones help break the project into manageable phases and provide a general timeline for research, design, development, and testing. By defining milestones, the team can better manage time, monitor progress, and ensure that all project requirements are met before the end of each semester

Table 10.4.1 Senior Design Project Milestones

Senior Design 1 Milestones			
Milestone	Description	Start Date	Anticipated End Date
Form Group	Establish initial roles & responsibilities	1/20/26	1/20/26
Discuss Project Concepts & Feasibility	Research and evaluate feasibility of an undercar animal detection system	1/20/26	1/30/26
Form Committee	Identify and confirm reviewers	1/20/26	1/29/26
Divide and Conquer Report	10-page project proposal	1/23/26	1/30/26
Update Group Website	Upload D&C finalized Report	2/6/26	2/6/26

Inspect Demo Vehicle	Inspect acquired electrical vehicle for system design and integration	2/6/26	2/8/26
Setting up Development Environment	Installing Raspberry Pi System	2/7/26	2/7/26
Design Circuit	Design overall circuit based on research and part comparisons	2/8/26	2/24/26
Design Lens System	1. Define specifications lens system 2. Model first order paraxial ray trace in Python	2/8/26	4/13/26
Figure Out Software & Hardware Integration	1. Define responsibilities between Raspberry Pi 5 & MCUs 2. Figure out communication (Wi-Fi/UART/SPI)	2/12/26	2/13/26
Figure Out Software & Optics Integration	1. Define interface between NIR cameras and Raspberry Pi 5 2. Configure camera communication and data acquisition 3. Synchronize optical input with software processing pipeline	2/12/26	2/28/26
Part Acquisition	1. Research components 2. Compare & pick best option 3. Order parts 4. Arrival of parts	2/12/26	5/1/26
Training Database	1. Collect infrared image samples for different animals 2. Organize and label dataset for training and testing 3. Train and validate animal detection/classification model	2/14/26	4/21/26
OSE Early Demo	1. Present full system concept and goals 2. Demonstrate inverse Galilean field compression principle using spherical lens proxy	3/3/26	3/3/26

	3. Explain optical layout/design approach and next steps		
Prototyping	1. Test circuit design a. Speaker b. Adjustable frequency input c. Strobe lights d. Stopping mechanism 2. Test circuit design a. Extending arm b. Power supply c. Voltage regulations 3. Trouble shooting	3/4/26	4/25/26
Midterm Report	60 pages	3/27/26	3/27/26
Midterm Report Group Meeting	1. Update advisors on progress 2. Receive feedback from advisors	4/1/26	4/1/26
Update Group Website	1. Upload Midterm report revision 2. Update reviewers	4/3/26	4/3/26
Initial PCB Design	1. Design initial PCBs to be reviewed	TBA	TBA
Model Real Lens System	1. Model system in Zemax OpticStudio with real lenses 2. Test & adjust for aberrations 3. Ensure final lens system is properly imaging	4/13/26	4/20/26
OSE Final Demo	1. Present integrated optical module 2. Demonstrate basic underbody imaging capability 3. Show alignment and illumination performance 4. Display captured test images 5. Discuss remaining development plans	4/21/26	4/21/26
Mini Demo Video	Demonstrate a portion of A.U.R.A.'s functionalities	4/28/26	4/28/26
Final Senior Design 1 Report	120 pages	4/28/26	4/28/26

Senior Design 2 Milestones			
Design Refinement	Review SD1 Feedback	TBA	TBA
3D Printing	3D Print Housing and Protective Casing	TBA	TBA
Assemble Optical System	1. Assemble camera, lens, and IR illumination components 2. Mount optical components into protective housing 3. Verify alignment and basic optical functionality	TBA	TBA
Integrating Hardware with Software	1. Connecting everything 2. Testing for response 3. Troubleshooting	TBA	TBA
Integrating Hardware with Optics	1. Attach to system 2. Durability test	TBA	TBA
Integrating Software with Optics	1. Interface camera with Raspberry Pi 5 2. Test stability 3. Verify detection 4. Verify proper identification 5. Troubleshooting	TBA	TBA
Full System Integration	Integrate software, hardware, and optics into a fully functional system	TBA	TBA
System Testing & Validation	Verify detection accuracy, response timing, alarm behavior, and reliability	TBA	TBA
CDR Presentation	Present finalized system design	TBA	TBA
Midterm Demo	Demonstrate integrated system operation, including live camera feed, animal detection logic, and alarm/interlock response	TBA	TBA
Final Presentation	Present final system design, testing results, and performance evaluation	TBA	TBA
Final Demo	Demonstrate full integrated system operation	TBA	TBA

Chapter 11 – Conclusion

The A.U.R.A. (Autonomous Undercar Risk Avoidance) system was developed to address a critical and often overlooked safety gap in modern autonomous vehicle operation: the inability to detect and respond to hazards located directly beneath the vehicle. While existing sensing systems provide extensive coverage of the surrounding environment, the undercarriage region remains largely unmonitored. This project demonstrates that targeted sensing and intelligent system integration can effectively close this gap.

A key contribution of this project lies in the design of the optical sensing system. The underbody environment presents unique geometric and lighting challenges, including extreme aspect ratios, limited ground clearance, and highly variable ambient conditions. To address these constraints, the system incorporates active near-infrared (NIR) illumination combined with a custom anamorphic optical configuration. This design enables efficient mapping of the wide undercarriage region onto a compact image sensor while preserving critical spatial detail. By concentrating sensor resolution within the region of interest, the optical system improves detection reliability without requiring higher-cost imaging hardware.

In addition to optical innovation, the hardware architecture of A.U.R.A. was carefully designed to support reliable operation in a constrained and harsh environment. The system integrates dual sensor modules positioned at the front and rear of the vehicle, ensuring overlapping coverage and minimizing blind spots beneath the chassis. These modules will be housed within protective enclosures and deployed using a retractable actuator mechanism, which allows the sensors to extend during stationary operation and retract once the vehicle reaches a defined speed threshold. This approach protects sensitive components from road debris while maintaining functionality during critical safety checks.

The system also incorporates multiple hardware subsystems that work together to enforce safety. A motor interlock circuit prevents vehicle movement when a hazard is detected, ensuring that the system actively intervenes rather than passively monitoring conditions. In parallel, a multi-modal deterrent system consisting of ultrasonic transducers and high-intensity strobe lighting is used to encourage animals to vacate the undercarriage region. These hardware components operate in coordination with the sensing and processing subsystems, forming a closed-loop safety mechanism that prioritizes both detection and response.

On the processing side, the system utilizes a Raspberry Pi 5 to perform real-time image analysis using a YOLOv11-based object detection model. The model was developed and trained using the Roboflow platform, which provided an efficient workflow for dataset labeling, augmentation, and training. By separating the training process from the embedded deployment environment, the project was able to reduce development complexity while maintaining compatibility with real-time inference requirements.

The integration of optical, hardware, and software components is central to the effectiveness of A.U.R.A. Rather than relying on a single subsystem, the project demonstrates how coordinated interaction between sensing, processing, and actuation can

produce a robust safety solution. The system not only detects potential hazards but also verifies clearance before allowing vehicle operation, shifting the approach from passive observation to active prevention.

The project also demonstrates the feasibility of implementing such a system within practical constraints. By utilizing commercially available components, modular design strategies, and scalable architecture, A.U.R.A. provides a foundation for potential expansion into full-scale vehicles or commercial applications. The modular nature of the system allows for future upgrades in sensing technology, processing capability, and detection algorithms without requiring a complete redesign.

Despite these achievements, several limitations remain. The current implementation is based on a demonstration platform and does not yet incorporate the final production-grade camera system. Variations in sensor characteristics, environmental conditions, and mounting geometry may affect detection performance and will require further validation. Additionally, real-time processing performance is constrained by the capabilities of the embedded hardware, which may limit scalability for more complex models or higher-resolution imaging.

Future work will focus on refining both the optical and hardware subsystems in addition to improving software performance. This includes further optimization of the anamorphic optical design, integration of final camera hardware, and improved environmental sealing for long-term durability. Additional efforts will also be directed toward expanding the training dataset to improve detection accuracy across a wider range of animal types and environmental conditions.

Long-term development may include integration with full-scale autonomous vehicle platforms, multi-sensor fusion, and enhanced classification capabilities.

In conclusion, the A.U.R.A. system successfully demonstrates a novel and integrated approach to undercarriage hazard detection and prevention. By combining optical innovation, robust hardware design, and intelligent software processing, the project provides a practical and scalable solution to a previously unaddressed safety challenge. This work highlights the importance of specialized sensing in autonomous systems and establishes a strong foundation for continued development and real-world application.

The final prototype retained the original safety objective while refining several implementation details. YOLOv8 was replaced by YOLOv11 Nano because YOLOv8 was not offered through the selected Roboflow workflow and Nano was the available YOLOv11 version. The planned UART-over-USB connection between the Raspberry Pi 5 and ESP32 was replaced by Wi-Fi because the Pi did not reliably read the ESP32 through USB. The software also implemented separate three-frame windows for the front and rear cameras and a two-out-of-three decision rule instead of the originally proposed fixed clearance timer. These updates reflect the results of implementation and integration testing while preserving the reasoning documented during Senior Design I.

Appendices

Appendix A - References

- [1] Waymo. (2024, September). *Waymo Safety Impact Report: Comparison of the Waymo Driver's Performance with Human Benchmarks*. Waymo LLC.
<https://waymo.com/safety/impact/>
- [2] Knight, H. (2025, December 5). How Kit Kat Was Killed: Video Shows What a Waymo Couldn't See. *The New York Times*.
<https://www.nytimes.com/2025/12/05/us/waymo-kit-kat-san-francisco.html>
- [3] Dai, Z., Wolf, A., Ley, P.-P., Glück, T., Sundermeier, M. C., & Lachmayer, R. (2022). Requirements for Automotive LiDAR Systems. *Sensors*, 22(19), 7532.
<https://doi.org/10.3390/s22197532>
- [4] Aristocrat Motors. (2024, January 22). Aristocrat Motors. Aristocrat Motors.
<https://www.aristocratmotors.com/2024-jaguar-i-pace-dimensions/>
- [5] Kona, S., & Dasari, G. (2025). Under-Car Vision: A Retractable 360-Degree Camera System for Enhanced Safety. *IJFMR* 250137223, 7(1).
<https://www.ijfmr.com/papers/2025/1/37223.pdf>
- [6] Sibanda, V., Mpofu, K., Trimble, J., & Zengeni, N. (2019). Design of an Animal Detection System for Motor Vehicle Drivers. *Procedia CIRP*, 84, 755–760.
<https://doi.org/10.1016/j.procir.2019.04.175>
- [7] Schumake, S. (1995). *Electronic Rodent Repellent Devices: A Review of Efficacy Test Electronic Rodent Repellent Devices: A Review of Efficacy Test Protocols and Regulatory Actions Protocols and Regulatory Actions Part of the Natural Resources Management and Policy Commons "Electronic Rodent Repellent Devices: A Review of Efficacy Test Protocols and Regulatory Actions*. 34.
<https://digitalcommons.unl.edu/cgi/viewcontent.cgi?article=1033&context=nwrcrepellants>
- [8] 1 Plastic + MetalBaby Car Bdm0936G Small Off-Road 12V4.5Ah LongweiMotor 390#*2Drive 12V500Ma Positive Energy. (2026). Walmart.com.
<https://www.walmart.com/ip/1-Plastic-MetalBaby-Car-Bdm0936G-Small-Off-Road-12V4-5Ah-LongweiMotor-390-2Drive-12V500Ma-Positive-Energy/11741857192>
- [9] The Home Depot. (n.d.). [Title of product specification sheet or manual] [PDF].
<https://images.thdstatic.com/catalog/pdfImages/5a/5a56f3da-aa4d-4383-8075-4705b52749a8.pdf>
- [10] *ADAS Explained: What is Advanced Driver Assistance System?* (2022). Suzukirndindia.com. <http://www.suzukirndindia.com/blogs/what-is-advanced-driver-assistance-systems>
- [11] *Vehant Technologies*. (2025, December 11). Vehant Technologies.
<http://www.vehant.com/under-vehicle-scanning-system/>

- [12] Commonlands. (n.d.). Wide-angle 3.2 mm M12 lens. <https://commonlands.com/products/wide-angle-3mm-m12-lens-cil332>
- [13] Commonlands. (n.d.). 200° fisheye S-mount lens. <https://commonlands.com/products/200degree-fisheye-smount-lens>
- [14] Institute of Electrical and Electronics Engineers (IEEE). (2020). IEEE 802.11 wireless LAN standard.
- [15] Raspberry Pi Foundation. (2023). UART communication documentation. <https://www.raspberrypi.org/documentation/>
- [16] Motorola. (2000). SPI block guide (Version 04.01). <https://www.nxp.com/docs/en/reference-manual/S12SPIV4.pdf>
- [17] NXP Semiconductors. (2014). I2C-bus specification and user manual. <https://www.nxp.com/docs/en/user-guide/UM10204.pdf>
- [18] Espressif Systems. (2023). ESP32-WROOM-32 datasheet. https://www.espressif.com/sites/default/files/documentation/esp32-wroom-32_datasheet_en.pdf
- [19] Nordic Semiconductor. (2008). nRF24L01+ product specification. https://infocenter.nordicsemi.com/pdf/nRF24L01P_PS_v1.0.pdf
- [20] FreeRTOS. (2023). FreeRTOS documentation. <https://www.freertos.org>
- [21] Python Software Foundation. (2023). Python documentation. <https://docs.python.org>
- [22] Raspberry Pi Foundation. (2023). Raspberry Pi documentation. <https://www.raspberrypi.org/documentation/>
- [23] OpenCV. (2023). Open source computer vision library. <https://opencv.org>
- [24] Python Software Foundation. (2023). Pillow (PIL fork) documentation. <https://pillow.readthedocs.io>
- [25] Scikit-image developers. (2023). Scikit-image: Image processing in Python. <https://scikit-image.org>
- [26] Howard, A. G., et al. (2017). MobileNets: Efficient convolutional neural networks for mobile vision applications. arXiv. <https://arxiv.org/abs/1704.04861>
- [27] Ultralytics. (2023). YOLO documentation. <https://docs.ultralytics.com>
- [28] Open Robotics. (2023). ROS documentation. <https://www.ros.org>

- [29] “IPC-2221A Generic Standard on Printed Board Design,” 2003. Available: [https://www-eng.lbl.gov/~shuman/NEXT/CURRENT_DESIGN/TP/MATERIALS/IPC-2221A\(L\).pdf](https://www-eng.lbl.gov/~shuman/NEXT/CURRENT_DESIGN/TP/MATERIALS/IPC-2221A(L).pdf)
- [30] American Academy of Ophthalmology. (n.d.). Galilean and Keplerian telescopes. <https://www.aao.org/education/image/galilean-keplerian-telescopes>
- [31] ams OSRAM. (n.d.). IR LEDs. <https://ams-osram.com/products/leds/ir-leds>
- [32] ams OSRAM. (n.d.). IR:6 next-generation infrared LED technology. <https://ams-osram.com/innovation/technology/irled-ir6>
- [33] ams OSRAM. (n.d.). OSLO Black SFH 4715AS. <https://ams-osram.com/products/leds/ir-leds/osram-oslon-black-sfh-4715as>
- [34] ams OSRAM. (2025). High-power emitters for illumination applications. <https://look.ams-osram.com/m/66e10252b242f443/original/High-power-emitters-for-illumination-applications.pdf>
- [35] ColorFabb. (2023). ASA technical data sheet. https://colorfabb.com/media/wysiwyg/datasheets/ASA_EN.pdf
- [36] Edmund Optics. (n.d.). 50 mm diameter, 1550 nm V-coat, $\lambda/4$ N-BK7 window. <https://www.edmundoptics.com/p/50mm-diameter-1550nm-v-coat-lambda4-n-bk7-window/26615/>
- [37] Fab.to.Lab. (n.d.). Arducam LN051 120 degree ultra wide angle CS-mount 3.2 mm focal length lens. <https://www.fabtolab.com/arducam-ln051-120-degree-ultra-wide-angle-cs-mount-3-2-mm-focal-length-lens-manual-focus-raspberry-pi-high-quality-camera>
- [38] Insta360. (n.d.). Insta360 X4. <https://www.insta360.com/product/insta360-x4>
- [39] Knight Optical. (n.d.). 850 nm standard interference bandpass filter, 25 mm diameter, 10 nm bandwidth. <https://www.knightoptical.com/stock/850nm-standard-interference-bandpass-filter-25mmdia-x-10nmbw-850fib25>
- [40] MeetOptics. (n.d.). Thorlabs packaged Si photodiode FDS015. <https://www.meetoptics.com/detection-devices/photodiodes-sensors-and-modules/packaged-si-photodiode/s/thorlabs/p/FDS015>

Appendix E - ChatGPT Prompts and Outcomes

[1a] **Case Study 1 Prompt:** “for nir images captured under 850 nm illumination what image processing techniques improve contrast between organic materials, like fur or skin, and ground surfaces?”

[1b] **Case Study 1 Google AI Outcome:** “For NIR images captured under 850 nm illumination, several image processing techniques can improve the contrast between organic materials (such as skin or fur) and ground surfaces (such as soil or rock). Organic materials often exhibit higher reflectance than soil in this range, appearing brighter, while 850 nm light penetrates skin differently than soil.

Key image processing techniques to enhance this contrast include:

1. Contrast Enhancement and Normalization

Contrast Limited Adaptive Histogram Equalization (CLAHE): This is highly effective for enhancing local contrast in NIR images, especially when dealing with shadows or uneven lighting typical in outdoor scenes.

Histogram Stretching/Normalization: By normalizing the intensity range of the NIR image, the difference in reflectance between the organic (high reflectance) and inorganic (lower reflectance) materials can be expanded.

National Institutes of Health (.gov)

National Institutes of Health (.gov)

+1

2. Multi-Spectral Fusion (VIS-NIR Fusion)

RGB-NIR Color Fusion: Fusing the 850 nm NIR image with a standard RGB image allows for the exploitation of the "color" (spectral signature) of the ground versus the "brightness" of the organic material. This helps distinguish surfaces that look similar in visible light but have different NIR absorption.

Multi-Spectral Image Segmentation: Using techniques like principal component analysis (PCA) on multi-band NIR data helps separate materials based on their spectral fingerprints, making it easier to isolate biological surfaces.

National Institutes of Health (.gov)

National Institutes of Health (.gov)

+4

3. Material-Based Segmentation (Thresholding)

Otsu's Method: Since organic materials typically have higher reflectance around 850nm than most soil surfaces, applying Otsu's thresholding algorithm to the histogram of the NIR image can automatically determine the best threshold to separate the foreground (organic) from the background (ground).

Thresholding in Intrinsic Image Domain: Generating an "intrinsic image"—which separates reflectance from lighting—can eliminate shadows and shading, leaving only material-dependent differences, which greatly increases contrast.

Imaging.org

Imaging.org

+2

4. Noise Reduction and Structure Enhancement

Color-Guided Filtering: When using an 850 nm filter, some sensor noise is inevitable. Guided filtering, using the RGB image as a guide, can denoise the NIR image while preserving edge details of the fur or skin.

Wavelet Transform Decomposition: This technique separates the image into different frequency bands, allowing for the enhancement of contrast in the average (low-frequency) band and edge/texture details in the detailed (high-frequency) bands.

National Institutes of Health (.gov)

National Institutes of Health (.gov)

+3

Summary of Differences at 850nm

Organic Materials (Skin/Fur): High scattering and reflectance, appearing bright.

Ground Surface (Soil/Rock): Lower reflectance and higher absorption at 850nm, appearing dark.

Best Technique for High Contrast: Applying CLAHE to a raw 850nm image or fusing it with an RGB image for color-based segmentation provides the best separation of these materials.

National Institutes of Health (.gov)

National Institutes of Health (.gov)

+2”

[1c] Case Study 1 ChatGPT Outcome: “Near-infrared (NIR) imaging around ~850 nm is great for separating organic materials (fur, skin, vegetation) from many ground surfaces because of differences in reflectance and scattering. To boost that separation, you can combine a few classes of image processing techniques:

1) Intensity normalization & contrast enhancement

Start by making the dynamic range more usable:

- **Histogram equalization / CLAHE (Contrast Limited Adaptive Histogram Equalization)**
 - *CLAHE works especially well in NIR because illumination can be uneven.*
 - *Enhances local contrast between textured organic materials (fur) and smoother ground.*
- **Gamma correction**

- *Adjust gamma (<1 brightens darker fur; >1 suppresses bright ground glare).*
- **Log or power-law transforms**
 - *Useful when reflectance differences are subtle but multiplicative.*

2) Texture-based separation

Organic materials often have richer micro-texture than ground:

- **Local variance / standard deviation filters**
 - *Fur → high variance*
 - *Ground (sand, asphalt) → lower variance (usually)*
- **Gabor filters or wavelet transforms**
 - *Capture directional texture (fur strands vs flat الأرض).*
- **Local Binary Patterns (LBP)**
 - *Good for distinguishing fine organic textures from smoother surfaces.*

3) Edge and **structure enhancement**

- **Unsharp masking / high-pass filtering**
 - *Enhances fine details like fur edges.*
- **Sobel / Scharr / Canny edge detectors**
 - *Organic shapes often produce more irregular edge maps than ground.*

4) Reflectance-based segmentation

At ~850 nm:

- *Skin and fur often reflect moderately to strongly (depending on pigmentation and moisture).*
- *Many ground materials (soil, asphalt) have flatter, lower reflectance curves.*

Techniques:

- **Adaptive thresholding (e.g., Otsu or local thresholding)**
 - *Works if histogram shows bimodality.*
- **Ratio imaging (if multiple bands available)**
 - *If you also have visible or another NIR band:*
 - *NIR / Red or NIR / Green ratios can strongly separate organic vs inorganic.*

5) Morphological processing

After thresholding:

- **Opening/closing** *to remove noise*
- **Connected component analysis** *to isolate organic blobs (animals, skin regions)*

6) Polarization & specular suppression (if applicable)

If your setup allows:

- **Polarization filtering** *reduces specular highlights from ground (e.g., wet surfaces), improving contrast with diffuse-scattering organic tissue.*

7) Learning-based approaches (most powerful)

If you have labeled data:

- **Classical ML (SVM, Random Forest)** using features:
 - Intensity
 - Texture (LBP, Gabor)
 - Local contrast
- **Deep learning (CNNs)**
 - Learns subtle reflectance + texture cues automatically
 - Works especially well if lighting varies

8) Practical pipeline (simple but effective)

A strong baseline pipeline:

1. Flat-field correction (remove illumination bias)
2. CLAHE
3. Local variance or LBP feature map
4. Combine intensity + texture (weighted sum)
5. Adaptive thresholding
6. Morphological cleanup

Key insight

*At 850 nm, **texture + local contrast** often matters more than raw intensity alone. Fur and skin scatter light in complex ways, producing micro-contrast that ground surfaces lack—so algorithms that amplify local variation tend to work best.*

If you tell me your exact setup (camera type, lens, whether you have multiple wavelengths, indoor/outdoor), I can suggest a more tailored pipeline or even parameter values.”

[2a] Case Study 2 Prompt: “What motor driver should I use to control a 12V DC linear actuator from a 3.3V microcontroller? I need it to handle stall current”

[2b] Case Study 2 Gemini Outcome: “For controlling a 12V DC linear actuator from a 3.3V microcontroller with bidirectional capability and stall current handling, the Cytron MD10C is a strong recommendation. It is a single-channel DC motor driver capable of supplying 13 A continuously and up to 30 A peak, operating from a 5V to 30V motor supply while accepting 3.3V logic-level PWM and direction signals directly from a microcontroller without a level shifter. The driver uses a MOSFET H-bridge architecture, which produces a much lower internal voltage drop than older bipolar designs like the L298N, meaning more of your supply voltage actually reaches the motor. It also includes built-in over-current protection, thermal shutdown, and under-voltage lockout. For your application, simply connect the PWM pin to a PWM-capable GPIO on your microcontroller and the DIR pin to a standard GPIO for direction control. The motor supply and logic supply should be kept on separate rails to prevent switching noise from affecting your microcontroller. Alternatives such as the L298N are widely available but are not recommended for this application due to their 2V to 4V forward voltage drop, which reduces effective torque at 12V and generates significant heat at moderate currents”

[2c] Case Study 2 ChatGPT Outcome: *“To control a 12V DC linear actuator bidirectionally from a 3.3V microcontroller with proper stall current handling, you need an H-bridge motor driver rated for at least the actuator's stall current with some margin. The Cytron MD10C is well suited to this task. It handles up to 13 A continuous from a 5 to 30V motor supply, accepts 3.3V logic directly, and uses a MOSFET-based H-bridge that avoids the high conduction losses of bipolar alternatives. The L298N is a common beginner choice but drops 2 to 4 volts across its transistor pairs due to its bipolar junction transistor architecture. At 12V this means your actuator may only see 8 to 10V under load, reducing force output and causing the driver to run hot. The reason MOSFET designs are better here is that MOSFETs switch rather than conduct linearly, so their on-state resistance is much lower and power dissipation is reduced significantly. For higher current or higher voltage applications such as 24V or 48V industrial actuators, a discrete MOSFET H-bridge with a dedicated gate driver IC would be preferred over any integrated module, but for your 12V prototype the MD10C is the right call. Make sure to add a flyback diode across the actuator terminals if it is not already included in the driver, to suppress the back-EMF voltage spike when the motor switches off.”*

[3a] Case Study 3 Prompt: *“How should communication be implemented between a Raspberry Pi and ESP32 controllers in an embedded system, and which communication protocol is most suitable?”*

[3b] Case Study 3 ChatGPT Outcome: *“Communication between a Raspberry Pi and ESP32 controllers can be implemented using several protocols, depending on system requirements such as distance, data rate, and complexity.*

1. UART (Universal Asynchronous Receiver/Transmitter):

UART is one of the simplest and most commonly used methods for communication between a Raspberry Pi and ESP32. It provides direct serial communication with low latency and minimal overhead. UART is suitable for short-distance, point-to-point communication and is easy to implement using standard libraries on both devices.

2. I2C (Inter-Integrated Circuit):

I2C allows communication between multiple devices using a shared bus. While it is useful for connecting several peripherals, it is generally limited to short distances and lower data rates. It may not be ideal for systems requiring higher reliability over longer connections.

3. SPI (Serial Peripheral Interface):

SPI offers higher data transfer speeds compared to UART and I2C, but requires more wiring and configuration. It is typically used when high-speed communication is needed between components located close to each other.

4. Wi-Fi:

Wi-Fi enables wireless communication between the Raspberry Pi and multiple ESP32 nodes. This is useful for distributed systems where wiring is not practical. Wi-Fi allows flexibility in system design but may introduce latency and requires network configuration.

For most embedded systems, UART is recommended for direct communication between a Raspberry Pi and a single ESP32 due to its simplicity and reliability. In systems with multiple distributed nodes, Wi-Fi can be used to enable communication between the main controller and remote ESP32 modules.

In many cases, a hybrid approach can be used, where UART is used for direct communication between the Raspberry Pi and a central ESP32 controller, while Wi-Fi is used to communicate with additional distributed nodes.”

[4a] Case Study 4 Prompt: “what kind of component would do well in outputting an ultrasonic noise to repel animals”

[4b] Case Study 4 ChatGPT Outcome: “If you are considering ultrasonic sound generation at the component level for a DIY or design project, the appropriate device is not a conventional speaker but a piezoelectric ultrasonic transducer. This component is used in nearly all commercial animal repellent systems because it efficiently converts electrical signals into ultrasonic sound waves.

Piezoelectric ultrasonic transducers are the most effective option. They operate efficiently in the 20–60 kHz range, which is suitable for deterring many animals, and are widely used because they are inexpensive, durable, and capable of producing high sound pressure levels. Transducers tuned to approximately 40 kHz are especially common, as this frequency provides a good balance for targeting a range of animals. These devices are also directional, allowing the sound to be aimed toward a specific area.

A simpler alternative is the use of raw 40 kHz ultrasonic emitter elements. These are basic, low-cost versions of transducers and are useful in budget or DIY applications. While they are easy to deploy and can be combined in arrays to increase output power, each individual unit typically produces lower intensity and requires an external driver circuit to function properly. For better performance, multiple emitters are often used in parallel to improve coverage and overall sound output.

High-power ultrasonic transducers represent a more advanced option. These industrial-grade components can produce significantly higher intensity ultrasonic waves and are better suited for large outdoor areas or long-range applications. However, they require high-voltage driver circuits, are more complex to implement, and may pose safety risks if not used correctly. As a result, they are generally only appropriate for more advanced or large-scale systems.

Ultrasonic sensor modules, such as those commonly used for distance measurement, are not suitable as primary sound sources. These modules are designed to emit short pulses

rather than continuous signals and operate at low duty cycles, resulting in weak output that is ineffective for repelling animals. However, they can still be useful as motion detection triggers within a larger system.

Several factors beyond the choice of component significantly affect performance. Frequency range is important, with most repellents operating around 35–45 kHz. Systems that vary or sweep frequencies tend to be more effective because they reduce the likelihood of animals becoming accustomed to a constant tone. Sound pressure level is also critical, as higher intensity increases deterrent effectiveness, and many low-cost modules fail due to insufficient output.

Directionality is another key consideration because ultrasonic waves travel in narrow beams and are easily blocked by obstacles. Using multiple emitters, angled arrangements, or acoustic horns can help improve coverage. Equally important is the drive circuit. Ultrasonic transducers require proper excitation through an oscillator, such as a 555 timer or microcontroller, along with an amplification stage like a MOSFET or dedicated driver integrated circuit. Many transducers require relatively high peak-to-peak voltages, often in the range of 20 to 100 volts, to achieve effective output.

In practical designs, the most effective configuration includes multiple 40 kHz piezoelectric transducers arranged in an array, a control system capable of sweeping frequencies, a suitable driver circuit, and a motion-activated trigger. This approach provides stronger output, broader coverage, and improved effectiveness compared to a single emitter operating at a fixed frequency.”

[4c] Case Study 4 Google Gemini Outcome: *“To output ultrasonic noise effectively for animal deterrence, the most critical component is a piezoelectric ultrasonic transducer or a high-frequency piezo tweeter. Standard speakers are generally ineffective because their physical diaphragms are too heavy to vibrate at the necessary ultrasonic frequencies, which typically range from 20 kHz to 65 kHz. Unlike electromagnetic speakers that use a voice coil and magnet, a piezoelectric transducer uses a ceramic material that expands and contracts when an electrical signal is applied. Piezo tweeters are common in DIY projects and can reach 30 to 40 kHz, though they must be specifically rated for ultrasonic use. Air-coupled ultrasonic transducers are also highly efficient at specific frequencies like 40 kHz, though they produce a very narrow and directional beam.*

Different animals respond to different frequency ranges: rodents are typically deterred by 32 kHz to 62 kHz, larger animals like dogs or cats by 20 kHz to 25 kHz, and insects by 30 kHz to 50 kHz. Utilizing a variable frequency or "sweep" is more effective than a constant tone because it prevents animals from becoming accustomed to the sound.

To make the transducer functional, a driver circuit is required. An NE555 timer IC is a standard choice for generating the necessary high-frequency square waves. This can be paired with a CD4017 decade counter for frequency modulation to create a more effective wavering effect. Since piezo transducers require a high voltage swing to maximize sound pressure levels, a push-pull transistor stage using pairs like the BC547 and BC557 or a dedicated bridge driver is often used. Other essential supporting components include potentiometers for manual frequency tuning and transistors for signal amplification to increase effective range.”