Capstone Design Abstracts

Fall 2025

Air Environment Risk Indicator System (AERIS)

Team Members: Muhaiminul Ashraf, Cesar de Oleo, Aryash Jain, Adib Khan, Logan Schneider, Abhinav Yarlagadda

Advisor: Professor Charles DiMarzio

Abstract

The AERIS team has designed and built a system for wireless sensing of hazardous gases and particles, which can be deployed in many scenarios due to its small Size, Weight and Power (SWaP). This system can be utilized in multiple situations, ranging from ensuring the safety of first responders to polling the environment in a garage workshop. Realizing the existence of toxic yet untraceable gases such as carbon monoxide is crucial. Many sensing systems are often large and expensive, making it difficult to ensure safety inside an enclosed area. AERIS' design proposes interfacing with various sensors to alert everyone nearby of any immediate health risks as well as updating a web dashboard for long term data depiction.

To use AERIS, the switch on the enclosure is turned on, after which the system autonomously interfaces with 4 individual sensors (particulate matter (PM), toxic gases, CO/CO2, temperature/humidity) using 2 ESP32 microcontrollers. All sensors and components are housed within a portable mechanical enclosure. Two fans flanking either end of the enclosure create a consistent airflow for the PM sensor to effectively poll the area for any matter that's dangerous yet small enough to infiltrate the lungs. The enclosure is also fitted with LEDs to indicate proper operation, as well as a piezo buzzer to immediately indicate any hazards. The ESP32 uses a variety of communication protocols (I2C, analog) to efficiently field data from all the sensors. The data is then processed and converted into data packets that are sent to a Render API server over Wi-Fi. This API interfaces between the firmware and the frontend to ensure that the data being displayed is up to date and uncorrupted. The web dashboard then takes all the individual data and displays it in a graphical format, including markers for when concentrations of certain gases are too high. This wireless connection allows the device to operate autonomously in an empty room, eliminating the need to potentially risk human lives. Due to the compactness of the enclosure and the lack of a need for a wired connection, this device can also be attached to a bike to continuously sense hazardous pockets within a whole area, as long as a hotspot connection is secured.

For testing and validation, various stimuli such as Febreze and hand sanitizer were placed at varying distances from the device to test both operations as well as limits of operation. Testing was conducted in numerous locations, from the outdoors to inside a 3D printing enclosure, to ensure operationality in a variety of situations. The results showed that the device was able to not only detect anomalies but provide an accurate concentration per toxic gas. The web UI also updated rapidly and correctly indicated whether levels of any specific gas or particle size were too high. However, the sensors sometimes incorrectly flagged gas levels to be too high or too low. This is because the calibration process of these sensors need to occur periodically to ensure the best level of operation. The AERIS team believes that access to higher tech sensors would allow for more accurate and responsive results, which would unfortunately raise the price point of this device.

SkySense: A Multi-Sensor Drone Detection System

Team Members: Ishaan Desai, Luke Maroney, Lucas Confalonieri, Kentaro Kasai, Benjamin Wu, and Jack Wilmott

Advisor: Professor Masoud Salehi

Abstract

SkySense is a low-cost drone detection system designed to provide reliable airspace monitoring in controlled areas. Reports of drone violations in restricted areas like power plants, correctional facilities, and public venues are doubling every three years according to Dedrone Drone Public Safety Analytics. These reports may not even tell the whole story as drone technology develops to bypass traditional drone detection systems. SkySense integrates a 60 GHz mmWave radar, a high-resolution camera, and an NVIDIA Jetson Orin Nano AI computer to detect and identify aerial introducers. By incorporating multiple sensors, SkySense can detect more advanced drones in different atmospheric conditions better than conventional systems like CCTV. SkySense aims to address the growing concern around UAV malpractice by securing our borders from smugglers and protecting controlled areas like power plants.

The radar subsystem (TI IWR6843AOPEVM) controls the detection of intruders in the airspace. By generating a 3D point cloud containing the range, velocity, and heading of an object, the radar can locate intrusions while in poor light or weather conditions. The tracked position data is fed to an ESP32-controlled dual-axis gimbal which is powered by stepper motors, creating a semi-spherical dome of object detection. After detection, the camera subsystem (Arducam IMX477) uses the radar's coordinates to capture a high-resolution image used for classification. A YOLO-based computer vision model performs object detection on the image to classify the kind of intruder seen, with the main classes being drone, bird, and debris. The Jetson Nano fuses this data together to respond to the intruder accordingly, with drones being tracked with a predictive motion algorithm, and birds and debris being disregarded. The Nano also reports to an interface with an alert in the event of a drone incursion.

SkySense is designed with affordability and scalability as priorities, with a prototype costing under $1,000. All major subsystems were engineered to allow customization and hardware upgrades in the case of commercialization. Multiple domes could be assembled to secure a cluttered airspace such as a skyline, while improvements to the budget can afford much longer ranged detection. For example, according to TI's specifications, a maximum range of 300 meters is possible on Texas Instrument's AWR2944P radar for $733, while our current radar maxes out around 100 meters. SkySense demonstrates that a multi-sensor approach to drone detection can deliver reliable real-time detection without the restrictions of traditional security solutions.

PT Pal

Team Members: Temirlan Nurtayev, Vanessa Guan, Felipe De La Garza, Evan Li, Rachel Barume, and Lily Tsahalis

Advisor: Professor Sarah Ostadabbas

Abstract

Gross motor development is a critical component of early childhood growth, forming the foundation for mobility, coordination, cognitive development, and social-emotional wellbeing. However, an estimated one in six children in the United States experiences developmental delays, with nearly 10% demonstrating poor or very poor gross motor performance. Early intervention can substantially reduce long-term impacts, yet access to high-quality, continuous physical therapy remains limited due to resource constraints, socioeconomic barriers, and low adherence to traditional exercise programs.

PTPal is an AI-powered pediatric physical therapy assistant designed to deliver accessible, engaging, and clinically informed gross motor skill intervention. Built on Google's BlazePose real-time pose estimation technology, PTPal provides accurate motion tracking, automated feedback, and progress monitoring, all processed locally to preserve user privacy. The system combines structured therapeutic exercises with a gamified user experience shown to improve motivation, adherence, and rehabilitation outcomes in children.Throughout development, the team completed a fully functional UI/UX system, including onboarding flows, webcam setup, exercise navigation, and feedback presentation, alongside a validated pose-estimation pipeline and structured exercise library. Real-time feedback cues, joint-angle-based error detection, and intuitive interface design ensure that therapy sessions remain clear, child-friendly, and supportive. The platform enables at-home guided therapy, complementing clinician-led sessions while expanding access for underserved populations.

By uniting evidence-based pediatric rehabilitation principles, engaging design, and modern computer vision, PTPal aims to improve gross motor performance, increase therapy adherence, and support long-term developmental outcomes. The system demonstrates the potential of AI-enhanced, gamified interventions to transform early childhood physical therapy and deliver meaningful impact for children and families.

Parkinson's Utensil

Team Members: Isabella Mor, Jordan Hoffman, Kyle Murrah, Fils Paul, Harrison Seeley, and Cameron Whiting

Advisor: Professor Bahram Shafai

Abstract

The Parkinson's Utensil team has developed an assistive eating device designed to reduce food spillage for individuals with hand tremors caused by Parkinson's disease, essential tremor, and aging. Today, conventional utensils lack stabilization properties for hand motions, while current active cancellation tools rely on expensive and inaccessible complex hardware. This project introduces an affordable and cost-effective solution using embedded sensing and a mechanical cover to secure food during movement. The device maintains the general form of a standard spoon, with all electronic components housed inside the enlarged handle so that it remains familiar and practical for everyday use.

The utensil consists of two main sections: a static spoon bowl and a rotating cover mounted on an axle driven by a compact servo motor. After the user collects food, the cover rotates over the bowl to prevent spills while the spoon is in motion. Once the spoon reaches the user's mouth, the cover retracts to allow the user to eat normally. Operation is fully automated and driven by real-time sensor feedback. A 9-axis IMU provides acceleration, angular velocity, and orientation data to an onboard ESP32 microcontroller. Using these measurements, the firmware identifies characteristic motion patterns and classifies the utensil's behavior into three states: scooping, in transport, and eating. The cover's position is adjusted based on these state transitions, allowing the system to respond naturally to the user's movements without requiring manual input.

When designing the hardware and its system, cost-effectiveness and robustness were at the forefront of all decision-making. The ESP32 microcontroller, low-cost IMU, micro servo motor, and lithium-ion battery were selected to balance performance with affordability, keeping total component cost relatively cheap but effective. The internal layout of the handle was designed to maximize durability, maintain water resistance, and ensure the device remains comfortable for users with reduced grip strength. Mechanical components were designed to withstand repeated movements and simple cleaning procedures.

Testing concentrated on validating sensor responsiveness, confirming accurate state transitions, and ensuring reliable and repeatable cover movement under typical eating motions. Prototype evaluations showed that the system can distinguish between scooping, lifting, and eating with consistency while maintaining smooth mechanical operation. Feedback from earlier trials helped refine sensor thresholds, movement timing, and mechanical tolerances. Overall, this project presents a practical, low-cost assistive utensil that enhances independence for people affected by tremors. By combining embedded sensing with a straightforward mechanical design, the system offers a spill-prevention solution that is simple to use, reliable, and accessible to a wide population of users.

RYDER

Team Members: Christopher Swanson, John Bergin, Taina Nieves, Steve Nketiah, Lane White, and Steven Lafferty

Advisor: Professor Charles Dimarzio

Abstract

The RYDER team has created a functional bicycle computer, capable of providing real-time biometric and positional data and providing higher levels of safety to bike users. This system is self-contained and compact, with minimally invasive integration for a user and bike system. A danger faced by many bicycle users is the possibility of significant injury without a system in place to easily call for help. The RYDER system provides easy access to emergency services and can automatically send location data to first responders. This system also provides group attendance checking to ensure groups stay together and are accounted for. The team also added functionality to provide useful data to a user beyond safety features to allow for improved fitness, recreation, or transportation.

The bike computer is a handlebar-mounted embedded system built around an ESP32-S3, enclosed with an integrated capacitive-touch TFT that provides all interaction and feedback. Its main PCB houses the processor, IMU, power circuitry, display interface, and an 6000 mAh battery chosen to provide long ride time without frequent charging. External sensors such as the wheel-mounted reed switch for speed, a haptic motor on the handlebars, and an oxygen-monitor sensor connect back to the main unit, allowing the computer to gather richer ride data beyond what can be measured in the enclosure. The firmware operates on a continuous state-machine loop that reads sensors, processes telemetry, updates the graphical display, and manages optional communication, keeping the system responsive even with many data sources. During use, the rider simply glances at or taps the touchscreen to view speed, heading, lean angle, GPS position, health metrics, and charging information. The design emphasizes modular sensing, low rider interaction, and clear real-time feedback, making it well-suited for cyclists who need consolidated information without distraction.

In order to ensure that our device provides reliable cellular connectivity, and ensures it maintains proper functionality, we allowed for rigorous testing during the entire project. This took the form of GPS, cellular connectivity, and battery life testing for our device, some of which still needs to be completed. The battery life testing is still to be done, but we have roughly estimated a total device charge time (for a full 6000 mAh battery) to be around 3-4 hours, which considering the estimated life time of 30-40 hours (depending on number of cellular messages, GPS connectivity, and screen brightness) is appropriate for the application. In order to test the GPS functionality, baseline testing data sets were collected in straight lines under various conditions including weather conditions, city/urban locations, and in open-sky and occluded conditions. These data sets also included other critical system sensor information such as the reed switch for speed, IMU data, and cellular connectivity information in order to test the signal strength. The hope is that analysis and calibration can be centered around these collected datasets, in order to further optimize the precision and accuracy of our product, providing trustworthy data.

We endured a significant amount of delays in shipping and import processing, which heavily delayed the bring-up time of our PCB. To counteract this, we purchased off-the-shelf sensors and devkits in order to develop the software while the hardware was delayed. Thankfully, once the parts came in, the PCB was mainly functional, with only minor tweaks to be made in order for full functionality, which was a tremendous relief especially considering the delays. For a future revision, a more accurate GPS and cellular module could be used, which would further improve the accuracy of our device. The same goes for the IMU in particular, and perhaps even the addition of a rudimentary sensor fusion algorithm could be applied with more time, such to fuse the GPS data with the IMU, furthering the positional accuracy.

VitalLink: A Wireless System for Continuous Patient Monitoring and Dynamic Triage in Emergency Departments

Team Members: Alejandro Baraona, Henry Butler, Edward DeBruyn, Andrew Hartley, Kevin Mei, Stephen Sodipo

Advisor: Professor Masoud Salehi

Abstract

The VitalLink team has designed and built a system for continuous wireless monitoring of patient vitals in emergency room waiting areas. The system addresses a critical gap in current hospital triage protocols, where patients often endure long wait times without medical oversight. During these periods, a patient's condition can deteriorate unnoticed, leading to critical emergencies that could have been preempted. Currently, hospital staff rely on initial triage data at check-in, which quickly becomes outdated as time passes. VitalLink's design proposes an additional safety layer using wireless wristbands to continuously track heart rate, blood oxygen levels, and skin temperature, ensuring that the most critical patients are prioritized dynamically.

To use VitalLink, a patient initiates the process at a self-service check-in kiosk, entering preliminary details which are synced via VitalLink's external API to the hospital's existing backend systems. Hospital staff are present to facilitate any additional information and to secure the wearable wristband to the patient. This device is built around a custom printed circuit board (PCB) featuring a Nordic nRF5340 microcontroller, which manages data collection from a suite of sensors, including a MAX30101 for heart rate and SpO2, and a MAX30208 for precise temperature readings. The wristband continuously packages this physiological data and transmits it via Bluetooth Low Energy (BLE) to a central base station. The base station, powered by a Raspberry Pi and housed in a custom laser-cut enclosure, acts as the system's hub. It runs a Node.js server that processes the incoming data packets and executes a custom triage algorithm based on US medical standards. This algorithm intelligently reorders patients based on the severity of their real-time vitals and wait time, displaying the prioritized list on a React-based staff dashboard. If a patient's vitals breach a safety threshold, the system triggers an immediate alert on the dashboard, prompting staff intervention.

For testing and validation, the team constructed fully functional wristband prototypes using 3d-printed enclosures and hand-assembled PCBs. The system successfully demonstrated a complete end-to-end workflow, validating the check-in procedure and wireless monitoring of live vitals using the VitalLink wristbands. To evaluate the system's scalability during peak hospital hours, the team utilized a "simulated band utility" to stress-test the backend infrastructure. This utility generated mock patient data reflecting specific medical emergencies, such as sepsis, heart attacks, and deteriorating vitals, confirming that the triage algorithm correctly prioritized these critical conditions in real-time. The primary limitation encountered was the low volume of physical wristbands available, which prevented full stress-testing of the Bluetooth network's bandwidth with a high density of real devices. Ultimately, the project successfully demonstrated VitalLink's potential as a crucial extra layer of safety, designed to preemptively identify at-risk patients and prevent fatalities in emergency waiting rooms.

Point-Based Autonomous Networked Drone Aiming System (PANDAS)

Team Members: Patrick Donnelly, Burke Havranek, James Lewis, Ryder Paulson, Shree Singhal, Simon Zhang

Advisor: Professor Sarah Ostadabbas

Abstract

Currently, drones used in the entertainment industry are limited by their battery capacities, which typically last between 15-50 minutes at a time. When capturing long events such as sports games, videographers must switch between multiple drones, leading to a loss of valuable aerial footage. To tackle this issue, the PANDAS team has designed an AI-powered turret aiming system to track and follow a drone in flight with a turret-mounted laser, using a drone-mounted photodetection system to simulate in-flight charging. With further development, we foresee the PANDAS system reducing the need for expensive grounding and battery replacement in such fields as entertainment and aerial videography.

The tracking system utilizes two computer vision models running on an Nvidia Jetson Orin Nano single-board computer (SBC), one for initial acquisition and one for tracking. Initial acquisition is accomplished using GroundingDINO, an open-sourced object detection model, which generates bounding boxes given a text prompt. These bounds are used to determine the coordinates of the photodetection system, which is then passed to CoTracker 3, a point tracking model from Meta, to track the photodetector. Point detection is a novel computer vision-based tracking technique used to dynamically follow specific feature points rather than entire objects. This technique is better suited to such applications as in-air drone charging over conventional bounding box techniques in that the latter is typically insufficient for tracking small objects, such as our photodetector, over time and space.

This integrated model continuously calculates an error delta relative to the center of the camera's perspective, to which the laser system is calibrated. This delta is then communicated to the turret in order to correct its aim in a manner agnostic to the geometry of its field of view. In particular, using an exponential moving average algorithm, the 3D-printed turret moves to minimize this error term using a two-axis motorized system, controlled using the same SBC as the tracking system. The photodetection system, implemented using a built-to-order PCB capable of measuring and recording the incidence of the turret-mounted laser, finally communicates the relative magnitude of our simulated in-air charging current, which is used to validate our design.

To date, we have demonstrated the functionality of each individual subsystem under laboratory conditions. In order to test the integrated system in the field, we obtained access to the Expeditionary Cyber and Unmanned Autonomous Systems Shared Research Facility's drone cage at Northeastern University Innovation Campus at Burlington, enlisting the help of a pilot from Aerospace NU's drone racing team to test PANDAS in a realistic scenario. The team was able to verify that PANDAS can track low-velocity movement, while limitations imposed by the hardware's processing power prevent the system from accurately tracking faster moving objects.

Underwater Light Fidelity (Li-Fi) Modules

Team Members: Tyler Myers, Matthew Rock, Farhad Ibrahimzade, Lilian Miles, Annabel Hammerle, Bryan Pitts

Advisor: Professor Bahram Shafai

Abstract

The Underwater Li-Fi Module project examines the feasibility of short-range optical wireless communication (OWC) as an alternative to conventional underwater links. Tethered systems introduce entanglement hazards, restrict operational range, and represent single points of failure, while acoustic communication is fundamentally constrained by low bandwidth, high latency, and substantial signal distortion in aquatic environments. By exploiting the low attenuation of blue light in clear water, the system implements amplitude-shift keying (ASK) modulation at approximately 470 nm to enable a compact, low-power communication channel between submerged platforms, such as remotely operated and autonomous underwater vehicles (ROVs and AUVs). Each module serves as a half-duplex transceiver, alternating between transmit and receive modes, and is enclosed within an IP67 waterproof housing featuring an O-ring seal, a custom acrylic photodiode window, and a waterprooflight-emitting diode (LED) assembly.

The transmitter subsystem uses a 10 W blue LED driven through a high-efficiency buck converter capable of supporting high-frequency operation, while the receiver employs a photodiode front-end followed by an amplifier and comparator to generate a conditioned digital signal suitable for processing by an onboard ESP32 microcontroller. Practical throughput is currently limited due to LED rise-time constraints and the absence of optical collimation or filtering. The system is powered by a 7.4 V lithium-ion battery. With a total system consumption of roughly 5.1 W, the modules achieve an estimated continuous runtime of7.5 hours. The transmission protocol is built upon the W5500 Ethernet board to allow a drop-in replacement for a current tethered Ethernet control scheme. The protocol is compliant with IEEE 802.3 to ensure the safety and handling of all data being sent. The protocol's functions mirror a router with an incoming and outgoing buffer to handle packets. Due to the loss experienced underwater and the half-duplex nature of our system, a custom protocol for transmitting the data via light was developed with a custom preamble, fixed packet size, and a lightweight Cyclic Redundancy Check (CRC). The protocol closest mirrors a single-line transmission control protocol (TCP) connection by utilizing packet acknowledgments to handle loss.

To validate performance, testing was first conducted in an open air environment, then in a controlled freshwater laboratory tank environment. Initial results demonstrate reliable communication at distances of approximately half a meter, where enclosures are offset to direct each beam into its respective photodiode. Range is primarily constrained by beam divergence, LED switching characteristics, and photodiode sensitivity. While the current design establishes only a short-range proof of concept, it lays the foundation for a scalable underwater OWC architecture. Future work will prioritize integrating optical collimation, higher-speed emitters, enhanced receiver filtering, and improved mechanical alignment to extend the operational range, increase data throughput, and further position Li-Fi as a practical alternative to acoustic communication in short-range underwater applications.

GIANO: Guided Assistive Piano Learning Gloves

Team Members: Nikolas Varga, Rahul Sing, Skyla Marie Profita, Ajith George, Celine Habr, Nouriya Al Sumait

Advisor: Professor Charles Dimarzio

Abstract

The GIANO team has developed a prototype wearable glove system that enables accessible, low-cost piano learning without a physical keyboard, paid lessons, or dependence on visual feedback. Traditional piano education often requires expensive equipment, fees and access to lessons, and sight-based pattern recognition, which limits access for those without financial resources and low-vision and legally blind users. Music education supports motor coordination, memory, focus, and social-emotional development. GIANO offers a haptic feedback-driven, computer vision-assisted platform that delivers real-time instruction, feedback, and audio output across both free-play and guided learning modes.

At the center of the system is a Raspberry Pi 5 running computer vision algorithms that generate a virtual keyboard mask, track hand position, and manage hand positioning logic. It communicates with three Teensy 4.0 microcontrollers: two on glove-mounted PCBs and one dedicated to audio generation. Each glove integrates seven haptic motors (one per finger and two for hand direction) along with velostat pressure sensors and flex sensors that detect finger presses, movement, and articulation. By combining sensor data with camera-based tracking, the system provides finger-by-finger guidance in learning mode.

After a brief calibration, in learning mode, the Raspberry Pi issues data instruction sets containing data such as note, finger, and hand position goals. The glove-mounted Teensys activate the appropriate haptic motors through an I2C multiplexer and monitor sensors for user input. When a press is detected, they transmit the information to the Raspberry Pi, which verifies it against the expected note. All detected presses, in any mode, are also sent to the third Teensy, which produces audio using FM-based synthesis.

The prototype reliably detects single note presses, provides clear haptic cues, and generates accurate sound output. A key limitation appears during chord instruction: driving multiple motors simultaneously through the I2C multiplexer increases latency and reduces responsiveness.

Future work includes improving haptic control for multi-finger chords and creating a fully hands-free, visual-free setup process to enhance accessibility. Overall, GIANO demonstrates a portable and inclusive approach to music learning, lowering barriers for users of all abilities and backgrounds.

Pedal-n-Play: Exercise Bike Gamepad

Team Members: Grace Boltz, Evan Clifford, Neda Habboosh, Willow Spector, Hubert Chen, and Daniel Zhao

Advisor: Professor Masoud Salehi

Abstract

Smart exercise bikes are a rapidly growing market, offering a healthy, low-impact form of exercise. However, mainstream products are largely restricted to simple video playback, with more engaging smart bikes tagged at thousands of dollars. More affordable options have higher barriers of entry, requiring users to build the interaction system, and are limited in the modes of interaction that they can provide. In contrast, Pedal-n-Play is an affordable and highly interactive smart bike experience which addresses the lack of interactive products at accessible price points, with the potential for myriad options for users.

The product consists of a trainer stand and attachable handlebar controllers, which interacts with a custom API. The trainer stand uses a pair of brushed DC motors to provide up to 500W of continuous braking power, with the ability to adjust electronically to lower power settings. The use of a trainer stand rather than a traditional stationary bike allows users to mount any bike, reducing the cost for consumers by allowing users to use their own bike. The Pedal-n-Play's compatibility with any bike also provides an immersive gaming workout experience without compromising the comfort of using the bike that the user always rides with.

Through the use of the custom API, Pedal-n-Play is extendable to any game that could be made to utilize its features. The Pedal-n-Play is a fully compliant Gamepad class Human Interface Device, enabling the control of any video game with the creation of drivers. Users can also choose to map the various buttons, joysticks, as well as the reading on handlebars and bike speed, to map Pedal-n-Play inputs to keyboard inputs for their favorite games. Finally, developers can use the Pedal-n-Play API to create a game which is not only fully compatible with the Pedal-n-Play from inception, but can take advantage of some of the device's unique features, such as Programmatic Resistance Control, a feature previously reserved for the most expensive standing bikes on the market.

To demonstrate the potential of our product, the team developed News Cycle, a custom infinite runner video game with in-game actions for moving by steering the bike, aiming with a joystick, and using various buttons on a controller. News Cycle draws from classic arcade and mobile games like Paperboy and Temple Run. Our software suite emphasizes mainstream appeal, open market development practices, and affordability.

ARC²: Augmented Reality Cycling Against Rearview Collisions

Team Members: Avery Corbin, Jack Gladowsky, Jude Percy-Allen, Dominick Villamor, and Jingpeng Zhang

Advisor: Professor Sarah Ostadabbas

Abstract

The ARC team has designed and built a cycling safety system enabling cyclists to be aware of oncoming threats on the road such as overtaking vehicles through an augmented reality heads-up display (AR HUD), eliminating the need for them to take their attention away from the road ahead. Municipal data and policy direction indicates cycling is an underserved and quickly growing area of public safety in the city of Boston, and national data points to the disproportionate risk of injury in cyclist-automobile traffic incidents, especially in cases of overtaking and close passes. ARC takes inspiration from driver-assistance systems developed for cars to give cyclists the agency and confidence to avoid threats on the road behind them where the current infrastructure fails to. ARC's design proposes the fusion of rearview radar and computer vision object detection to detect and characterize vehicular traffic before it becomes a collision risk for the cyclist. This information is relayed to the cyclist via an AR HUD, providing them with full front and back awareness.

The ARC system utilizes a fusion of radar and computer vision to detect traffic behind the cyclist and display it to them using augmented reality. The system utilizes a radar module and a webcam to gather information about the traffic behind the cyclist. This information is fed into a Jetson Orin Nano onboard the bike which handles processing the radar and computer vision detections. The processing algorithm computes the angle of each detection from the +Y-axis of our sensor suite mounted to the bike. Using these angles we can determine which radar detections correspond to the webcam computer vision detections, resulting in us having a proper position estimate for each rear detection. The Jetson then sends these resulting hazard detections through a websocket to our Magic Leap 2 Augmented Reality (AR) headset. The AR headset has a custom UI to display the hazard detections with the class and position of the object, while not blocking the cyclist's view of the road ahead. This enables the cyclist to see detections behind them without ever having to turn their head.

For testing and validation, members of team ARC rode around a bike equipped with our radar and webcam, along with another webcam facing the front of the bike. The sensor data from these rides was recorded allowing the team to analyze the recordings and hazard detections to verify that our algorithms successfully detected and classified hazards. The team found that the ARC system does alert cyclists to hazards behind them that they otherwise would not be aware of. However, we determined that although our system does detect hazards, cyclists should still check over their shoulders and behind them while using the ARC system to verify hazards.

InductiSense: Self-Powered Sensor Network

Team Members: Druhi Bhargava, Adityaa Chandna, Leonardo Chen, Vasishta Malisetty, Haley Pflasterer, Owen Richmond

Advisor: Professor Bahram Shafai

Abstract

The InductiSense team has designed and built a self-powered network of sensors that is highly power efficient and can be adapted and used in several different applications, including monitoring environmental factors, agricultural and ranching tracking, and power line monitoring. The sensor network is designed to last several years without needing replacement and can be recharged using solar power or a custom wireless clamp that can inductively harvest power from current flowing through power lines. Rural and remote areas often lack scalable, continuous monitoring for environmental or agricultural factors or power line monitoring, and existing sensing solutions are expensive. InductiSense provides quantitative insights into these factors while reducing maintenance costs, all at a much reduced cost, making it a viable product.

To power InductiSense, the team designed two cases: one containing a DC solar panel and the other containing a split-core current transformer (referred to as clamp) with a matching network and rectifier. Both cases produce sufficient power at a low voltage, which is converted to a higher voltage to supply the microcontroller with sufficient power and charge an energy storage device (either a battery or a supercapacitor) during device sleep. The microprocessor used is a Nordic nRF52840 (capable of Bluetooth Low Energy or BLE) on a RAK4631 module which also contains an SX1262 LoRa (long range) transceiver module to enable low-power communication. The RAK module has universal connectors which allow for a variety of sensors to be attached depending on the application. The microprocessor is optimized for power consumption and is in a sleep mode unless it is actively sampling sensor data. Every ten minutes, it activates the LoRa transceiver and sends out a data packet with the sensor readings on the network. Each LoRa transceiver then relays the message until it reaches the endpoint. Data can be transmitted to a phone using BLE from the endpoint, where it is uploaded to a server and can be displayed on an app or a website.

For design, testing, and validation, several clamps were prototyped and tested in a laboratory setting. Power line deployment of these clamps would require regulatory approval, which was not possible within the project timeline. Preliminary tests look extremely promising, and the proof-of-concept product is functional based on this testing. The InductiSense team believes that with appropriate safety validation, this product can be implemented in a real-world setting. Software was continuously tested with the maximum range of LoRa transmission tested via movement throughout Boston and the system tuned for optimal transmission. The app and website was developed in conjunction and progressed from development with sample values to creating a fully functional pipeline from LoRa transmitted packets to BLE to endpoints. Data is uploaded onto a web server for easy retrieval and access.

IR ASL Glove

Team Members: Norah Bumrah, Camille Ferrell, Lauren Hutchison, David Noble, Clara Small

Advisor: Professor Charles DiMarzio

Abstract

The IR ALS Glove team has designed and built a low-cost, portable system capable of recognizing letters from the American Sign Language (ASL) using a combination of inertial measurement units (IMUs) and infrared (IR) optical tracking. This glove is intended to support the Deaf and Hard of Hearing community by providing real-time translations in environments where interpreters may be unavailable, e.g. hospitals, emergency situations, or public speaking events. Current ASL recognition technologies are often expensive, require strictly controlled conditions, or depend on two-handed input, making them unsuitable for daily mobile use. The IR ASL Glove's design addresses these limitations by enabling single-handed ASL recognition with compact, embedded electronics.

To operate, the glove integrates five ICM-20948 IMUs placed on each fingertip. All sensors communicate through a single TCA9548A I^2C multiplexer, which routes data to an ESP32-C3 microcontroller responsible for synchronizing fingertip IR LED timing and transmitting the data to a computer via Bluetooth Low Energy. A modified Logitech C922 webcam (with the IR filter removed) captures the illuminated fingertip markers, enabling IR blob detection. On the computer, IMU-derived orientation is fused with camera-based fingertip position measurements to mitigate IMU drift and form a reliable estimate of finger posture. These reconstructed finger positions are then compared against a calibrated dictionary of ASL alphabet configurations to determine the signed letter, which is displayed on a user interface.

System development proceeded through IMU verification, camera modification and testing, and integration of the sensor-fusion pipeline. During testing and validation, the team confirmed reliable detection of illuminated fingertips via blob-detection software, successful Bluetooth transmission of IMU data, and effective fusion of both data streams for gesture reconstruction. Challenges such as sensor drift, camera occlusions, and user-dependent finger geometry were addressed through calibration procedures, timing synchronization, and iterative filter tuning. The prototype demonstrated functional single-hand ASL alphabet recognition and established a foundation for future expansion toward full sign-language translation.

Autonomous Robotic Collection Bots (ARC)

Team Members: Jackson Knapp, Ojasvi Mukkoti, Ananya Kumar, Derek Jackson, Kate Rober, and Carlos Rodriguez

Advisor: Professor Masoud Salehi

Abstract

Over the course of the last 6 months, the ARC team has designed and built a semi-autonomous robot for trash collecting. Over 2 billion tons of waste is generated each year, and improper disposal of trash and recyclables poses serious environmental, health, and economic challenges. There are small-scale in-home solutions, such as Roombas and large industrial sorting systems for industrial use, but there remains a significant gap for a system designed specifically for everyday indoor public environments. Thus, the ARC team presents a semi-autonomous robotic solution that merges the gap between manual cleanup and fully automated waste management in everyday indoor public spaces.

To use ARC, the user first navigates the robot using simple arrow key commands, allowing for intuitive manual control in indoor spaces. Once the integrated object-detection neural network identifies a trash or recycling item through real-time camera input, the system transitions into autonomous mode. During this phase, the robot performs self-centralization to optimally align itself with the object, ensuring the robotic arm is positioned for accurate pickup. A ROS-based pipeline, run on a Jetson Orin Nano, processes the camera data and coordinates the overall system execution, directing the robotic arm to grasp the detected object while an integrated vacuum system assists in collecting the item. The neural network continuously classifies specific waste categories, which determine the motion planning parameters of the robotic arm, including the required shoulder and elbow angles for successful grasping. To maintain simplicity and efficiency, the arm is designed with two degrees of freedom – shoulder and elbow- enabling reliable pickup while minimizing mechanical complexity. The robotic base and arm structure are fully 3D printed, allowing for rapid prototyping and cost-effective scalability. This hybrid control structure allows for efficient human-robot collaboration, combining user guidance with intelligent automation to ensure accurate, safe, and effective waste and object collection.

Testing and validation were conducted using multiple methods to evaluate the performance of each subsystem. The neural network was tested in the capstone lab using various objects, demonstrating generally consistent accuracy, though ambiguous or irregular items occasionally caused misclassification. Mobility testing focused on user-controlled navigation, with the ARC bot operated through direct connections to a monitor, keyboard, and mouse to imitate system setup and control. This demonstrated moderate success in maneuverability and responsiveness within an indoor environment. Robotic arm testing evaluated its programmed movements for object retrieval based on neural network input; while the arm operated as intended, the vacuum pickup system repeatedly failed due to mechanical fragility, with several instances of structural breakage. These results indicate that a more durable and simplified pickup mechanism would significantly improve overall system reliability in future iterations. If given more time, a more reliable localization tool such as vSLAM would have been able to increase the autonomous abilities of the robot, allowing for less manual control from the user.

Autonomous Mine Detection Vehicle (AMDV)

Team Members: Edwin Yang, Ronan Hinshaw, Akshaj Sirineni, Luke Sanders, Tarika Selvaraj

Advisor: Professor Sarah Ostadabbas

Abstract

The AMDV team designed and built an RC car capable of detecting underground landmines across various terrains by utilizing a DIY metal detector and live cameras with obstacle detection algorithms. Fleets of AMDVs can be deployed across the world by citizens and militaries alike in low-cost, no-personnel mine detection and clearing campaigns. With almost 6,000 casualties from landmines annually (APOPO), these campaigns are pertinent to the safety of the residents in the 58 global landmine-inflicted countries where more than 2100 of these casualties are children. Current landmine detection solutions including mine detection rats, mine cleaning tanks, and manual human detection are expensive, impractical, and incredibly dangerous. The AMDV's design proposes a combination of multi-sensor integration, object detection algorithms, remote laptop assisted operation, and an interactive digital dashboard to mark the precise location of landmines in real time as they are detected.

For operation of The AMDV system, the vehicle integrates several sensing and computing subsystems into a compact and lightweight platform designed specifically for uneven and unpredictable terrain. The chassis is built on a modified RC truck with reinforced suspension springs and custom 3D-printed housing that protects the onboard electronics during operation. A tailor made metal-detection coil driven by an Arduino Uno produces an oscillating electromagnetic field and detects changes in the frequency of oscillation to identify subterranean metallic objects. This allows the vehicle to detect hazards without any direct contact. The AMDV uses stereo cameras connected to an NVIDIA Jetson Orin Nano Super Developer Kit to run real-time object detection algorithms on a live video feed. The Jetson relays video data with inertial and positional data to assist in navigation decisions to the digital dashboard, while the Arduino independently handles steering and throttle control to ensure responsive maneuverability. The AMDV can be programmed to run in a predetermined path over a selected geofence, or driven manually via controls accessible on the digital dashboard. The digital dash communicates with the Arduino's metal-detection circuit and continuously updates a satellite map of the search area with locations of potential landmines. This integrated approach allows The AMDV to autonomously explore, identify potential landmines, and relay clear visual indicators of object detection events to the operator.

In the process of testing The AMDVs onboard systems, a variety of tests were conducted in various areas, including open fields, indoors and in more rugged terrain, all areas in which The AMDV could prove to be life saving. Through these tests, it was determined that The AMDV functioned best when operating in locations with a clear view of the sky, as the onboard GPS module faced interference when used indoors. However, the onboard stereo cameras provide key information which allow users to track and validate the rovers location as a backup for the GPS. It is believed that the inclusion of a magnetometer would benefit the rover as The AMDV's heading could be tracked, providing even greater navigational detail for The AMDV's operator.

FirePod: Wildfire Detection System

Team Members: Cameron Brown, Charlie Amante, Charlie Lange, Paarth Soni, Trevor Stempien, William Fox

Advisor: Professor Bahram Shafai

Abstract

Wildfires continue to escalate in frequency and severity, emphasizing the urgent need for distributed early-detection systems capable of identifying hazardous environmental conditions before flames spread uncontrollably. Current detection infrastructures, such as satellites, lookout towers, and aerial surveillance often suffer from long update intervals, limited coverage in remote terrain, and obscuration from vegetation or weather conditions. To address these critical limitations, the FirePod team has designed a compact, multi-sensor device intended to serve as a scalable, ground-level wildfire detection system.

FirePod incorporates a variety of environmental sensing components, including a gas and VOC sensor (ENS160), an environmental sensor measuring pressure, humidity, and air quality (BME688), and a precision digital temperature sensor (TMP117). These sensors communicate over an I²C bus with an STM32 microcontroller, which coordinates data acquisition and prepares readings for wireless transmission. A u-blox MAX-M10S GPS module provides accurate positional data to geotag sensor alerts, while a Semtech SX1262 LoRa transceiver enables long-range communication to remote monitoring stations. The system is powered through a micro USB and solar-compatible battery management circuit featuring a dedicated charging IC, a buck-boost converter for stable power delivery, and low-dropout regulators (LDO)to ensure consistent operation under varying load conditions. The buck-boost IC manages a 3.3 V rail for most components while the LDO manages a 1.8 V line dedicated for the ENS160. All circuitry is implemented on a custom four-layer PCB designed in Altium, with careful attention to power plane layout and accessible test points to support hardware debugging and subsystem validation.

At the current project stage, the team has completed full PCB fabrication, component soldering, and hardware assembly of the FirePod device. Initial electrical validation has confirmed correct power sequencing and successful continuity across major subsystems. The next phase will involve comprehensive subsystem testing such as verifying sensor accuracy, GPS acquisition performance, LoRa communication range, and microcontroller firmware behavior followed by integrated system simulation under controlled environmental conditions. Future work will focus on optimizing RF performance, strengthening environmental robustness, and preparing the device for field testing. FirePod lays the foundation for an accessible and distributed wildfire monitoring network, offering a promising step toward faster detection and improved response capabilities in fire-prone regions.

Runway Incursion Monitoring and Alerting System

Team Members: Maxwell Boulerice, Maggie Heaney, Tyler Johnson, Andrew Krumpter, Jonah Lefkoff, Jessica Taddeo

Advisor: Professor Charles DiMarzio

Abstract

Runway incursions at uncontrolled airports remain a significant contributor to general aviation accidents. Without an air traffic controller to coordinate ground operations, aircraft and ground vehicles rely on self-reporting and visual awareness to avoid entering active runways at the wrong time. These environments are especially prone to human-factor challenges such as miscommunication, limited visibility, and high workload. To address this gap, our capstone team is developing runway incursion monitoring and alerting system, a low-cost deployable device designed to detect aircraft on final approach and warn pilots preparing for takeoff before a hazardous situation occurs.

The system combines three sensing modalities to ensure robust detection: a commercial 77 GHz automotive radar for tracking approaching aircraft, dual high-resolution Sony IMX477 cameras paired with an aircraft-detection machine-learning model, and an ADS-B receiver that captures position broadcasts from cooperative aircraft along with basic weather information. These data streams are merged through a sensor-fusion pipeline that estimates an approaching aircraft's motion vector in real time. If the system predicts that an aircraft on final approach is close enough to pose a risk to an aircraft entering the runway, it activates high-visibility indicator lights to signal to pilots to hold short.

To operate reliably at small airports with limited infrastructure, the system is designed as a self-contained field unit. The unit includes built-in power redundancy to ensure continuous operation during outages or unexpected failures. A Jetson Orin Nano module provides on-device machine-vision processing, and custom PCBs handle power distribution, high-speed camera interfaces, and external I/O. The enclosure integrates cameras, radar, antennas, and indicators into a weather-resistant housing suitable for long-term outdoor use. The power subsystem supports both AC input and battery operation, enabling flexible placement near runway entry points without requiring permanent installation.

Throughout development, the team collected field data at a local uncontrolled airport. These recordings were replayed through the system to validate radar and camera performance, strengthen model training, and evaluate sensor-fusion accuracy. The resulting prototype demonstrates that a multi-sensor device can reliably detect approaching aircraft and provide clear, timely alerts to prevent runway-entry accidents at airports that lack traditional air-traffic control resources.

AquaSentri (Rescue Drone and Wearable System)

Team Members: Abhishek Narang, Anthony Ngumah, Rohi Ahn, Ryder Robbins, Steven Lam

Advisor: Professor Masoud Salehi

Abstract

AquaSentri is a prototype rescue system designed to reduce time-to-floatation in small-vessel overboard emergencies using a wearable-triggered autonomous drone. Many recreational boat rental environments lack dedicated lookout personnel, and operators may not immediately notice when a passenger falls into the water, especially in choppy or low-visibility conditions. While traditional life jackets provide buoyancy, they do not provide real-time location updates or automated rescue support. AquaSentri addresses this gap by integrating a low-cost wearable transmitter with an autonomous drone capable of navigating to the user's last known position and delivering a flotation device.

AquaSentri is composed of a two-part system: a wearable beacon and an autonomous drone. To use the product, the user attaches a wrist-wearable waterproof transmitter that continually broadcasts its GPS location to the drone station using LoRa communication, a long-range wireless link. During this period of consistent location updates from the wearable to the drone, the drone will remain idle and keep a log of the user's position. If these transmissions become intermittent or stop entirely, such as when a person falls overboard and submerges under, the drone automatically launches and flies to the last reported GPS coordinate. Once at the waypoint, the drone will begin its search and use its thermal sensor to visually identify the user within the surrounding area. After the user is located, the drone releases a flotation device directly above their position and returns back to its station.

During development, several limitations were identified. LoRa signals attenuate quickly underwater, resulting in loss of live tracking when the wearable is submerged. The thermal sensor used provides limited resolution and struggles with thermal contrast over water, restricting performance in challenging lighting or wave conditions. Additionally, drift between the last known GPS fix and the user's true location can grow rapidly due to wind, waves, and current. Finally, drone flight time remains constrained by battery capacity, and environmental conditions such as wind and spray may impact search accuracy and stability. For testing and validation, the system will measure LoRa communication range in open field conditions, GPS accuracy of the wearable, drone payload drop precision, thermal camera performance in controlled water scenarios, and drone battery endurance under expected mission loads. Autonomous testing was limited due to drone license difficulty with FAA shutdown.

Despite these constraints, it is believed that AquaSentri provides a feasible proof-of-concept for wearable-triggered, drone-assisted rescue within bounded maritime environments, particularly for boat rental operators responsible for passenger safety.

RFID Tracking System

Team Members: Scott Abramson, Harrison Broad, Matthew Friedman, Ephram Jacquin, Christopher Lam, Arnav Mishra, Brandon Walters

Advisor: Professor Bahram Shafai

Abstract

The RFID Tracking System Team designed a project that implements a modern-day cloud-connected tracking system that aims to improve accuracy, reliability and security to various items that tend to shift in their given environment. This system has a wide variety of environments in which it would prove to be an asset. Many organizations such as hospitals, airports, manufacturing facilities, and even retail stores can benefit from this system as it can help reduce costs due to misplaced or unlabeled items. Our system hopes to modernize tracking by integrating an automated data collection service through a wireless connection and accessible through a web application. At its core, the system uses a Raspberry Pi microcontroller connected to an RFID scanner to collect RFID tags. This data is then collected to a cloud database which is then accessible to the web application to display real time data.

The design uses Amazon Web Services (AWS) to process and store information reliably. Specifically, the design uses the AWS lambda function along with the Amazon Relational Database Service (RDS). The Raspberry Pi sends HTTP requests containing the RFID tag data to an API endpoint where a lambda function processes the requests and updates the database. There is a structured relational database schema that links organizations, users, scans sessions, RFID reads, and trackers. These tables make the database organized and allow for efficient querying. Having the RFID hardware scan real-world items which are then stored and processed through the cloud will allow for scalability and security. The system web application presents the data in an organized format to allow users to view scanner activity and track tagged items. The application communicates with the backend through the same API gateway and lambda functions used by the hardware to retrieve data. The web application is modular because it can always be updated for future features or provide types of data.

The project development was divided into two groups, the hardware team and the Software team. The hardware team is responsible for the microcontroller setup and integrating the RFID scanner. The software team is responsible for setting up the database through AWS, implementing lambda logic, along with setting up the web application. Once the two groups finish the setup, the groups come together to connect the hardware with the software with scans being formatted into JSON objects and sent to the cloud. Testing was conducted at each level whether it was testing the software for any bugs, misconceptions of the database schema, or testing the hardware for scans and sent JSON objects. Through the implementation of the Raspberry Pi, RFID scanner, and cloud platforms, the group is able to connect RFID tag reads to the homemade database and demonstrates how a modern-day tracking system can be built to be reliable and scalable. The result is a modern solution to inventory accuracy that helps reduce loss and enhance operational efficiency throughout many industries.