Capstone Design Abstracts

Fall 2023

Autonomous Waste Collector (AWC)

Team Members: Ao Ruan, Aral Balsari, Ousmane Toure, Jianning Chen, Yide Song, Yizi Wang

Advisor: Professor Charles A. DiMarzio

Abstract

The Autonomous Waste Cleaner (AWC) team has designed a pet waste cleaning robot capable of operating on varying surfaces from grass, flat floors, and ramps. The AWC is characterized by its water-proof, lightweight, and durable design, coupled with an advanced navigation system to locate and pick up pet waste efficiently. This innovation was inspired by the need for a convenient and hygienic solution to pet waste management. With the global pet population exceeding one billion, a significant percentage of U.S. households owning pets, the challenges in pet waste cleanup are evident. These challenges include physical strain, difficulty in waste identification, and the risk of contamination. Market research underscores the rising demand for simplified waste management and disposal solutions, focusing particularly on resolving lawn care difficulties and reducing pet waste in public spaces. Notably, while automatic cat litter boxes address feline needs, an equivalent product for dogs has been lacking until now.

The AWC represents a significant technological advancement in robotic waste management systems. At its core, the AWC utilizes NAV2 for autonomous navigation, enabling it to skillfully maneuver to designated target points. To ensure optimal functionality and reliability, the AWC has been subjected to extensive testing in the Gazebo simulation environment. The testing process allowed for fine-tuning of the robot's operational mechanics, ensuring a high level of precision before introducing the software to physical components. The robot's mobility and functionality are enhanced by a set of four Mecanum wheels, enabling multidirectional movement, and powered by an 11.1V rechargeable battery. This design allows the AWC to efficiently navigate various terrains while carrying a load of up to 16kg and maintaining an operational time of approximately 1 hour. The integration of these advanced technologies positions the AWC as a pioneering solution in the realm of automated pet waste management.

For the object detection part, the team utilized both YOLOv8 and TensorFlow Lite to train the model. The model used are 3 fake excrements that are made of play dough and approximately 200 pictures were taken inside the ballroom with different backgrounds to ensure the maximum accuracy. The training process involved training using YOLOv8 Nano model and TensorFlow2 EfficientDet D0 512x512 model and the most accurate model is selected to be the one implemented in the raspberry pi. The graphical performance is enhanced by Google Coral USB accelerator. After testing on the webcam, the accuracy is able to reach 98% while the recall is able to reach 93%. The model can theoretically be improved by taking more pictures with more backgrounds, however, running the model will take more memories from the Raspberry Pi and slow down the responding speed of it.

SummitSignal

Team Members: Gabe Peralta, Cameron Gordon, Justin Kawakami, Sean Vumbaco, Ryan Lonkart, Elise Cooke

Advisor: Professor Masoud Salehi

Abstract

The SummitSignal team has developed a multifunctional SOS device and outdoor activity tracker, addressing the critical safety needs of outdoor adventures where unexpected and perilous situations may arise. In remote areas where cellular coverage is often lacking, adventurers face challenges seeking assistance from emergency services. Leveraging the GNSS network, specifically the Iridium satellite network, SummitSignal precisely determines the user's location and can send emergency messages to responders upon the user's request. This advanced solution guarantees global coverage, enabling seamless communication with emergency services in unconnected regions.

To utilize SummitSignal's multifunctional capabilities, users employ a device that serves dual purposes. It samples and records activity data, including step count, altitude, longitude, latitude, and speed from an array of I2C enabled sensors such as an inertial measurement unit and GNSS module. The recorded data is seamlessly tracked and visualized on a topographical map through a user-friendly web application. The device is centered around a fully integrated custom PCB, featuring a Teensy 4.1 MCU running custom firmware, a SDIO SD card port for data logging, multiple I2C and SPI sensors, and a RockBLOCK Iridium transceiver to communicate with the Iridium satellite network, ensuring global coverage. For data management, users upload their recorded data file to a website, where AWS Lambda, DynamoDB, and S3 functions handle computing and storage resources. The backend orchestrates the presentation of a comprehensive journey overview, enhancing the user's experience. In addition to its activity tracking capabilities, SummitSignal features an SOS function. Users can designate emergency contacts and provide trip details on the website. The physical device incorporates button-operating modes for various functions. A short press of the right button reveals the state of charge, while a long press restarts the device. The left button starts and stops logging with a short press and closes the current trip file with a long press. Simultaneously pressing both buttons triggers an SOS signal, sending a premade emergency message along with current GPS coordinates to the specified contacts.

For testing and validation, over 15 messages were sent using the SOS feature. The device's battery consumption was evaluated at various levels, and an AWS Lambda function was developed to process the logged data where it could then be visualized on the frontend. Operating flawlessly on battery power, the device operated for 40+ hours before being charged, on track to meet the designed 45 hour battery life. This robust performance is particularly crucial for trips lasting around one day, ensuring sustained functionality throughout the entire duration. After multiple enclosure iterations, the final version is over 45% smaller than the initial one, a notable achievement in line with our aim for portability and accessibility.

DAMIT: Drone Attachment to Measure Ice Thickness

Team Members: Marcus Albertson, Stepan Kapreillian, Carrie Patterson, Alek Tunik, Diego Valdivia, Jack Zakian

Advisor: Professor John Kimani

Abstract

The absence of a safe and reliable method to measure ice thickness on frozen lakes and ponds has led to over 4,000 fatalities globally in the last 26 years. Global warming has compounded this issue, rendering traditional winter activities like ice fishing, skating, pond hockey, and snowmobiling, on historically safe bodies of water increasingly dangerous. Traditional methods for gauging ice thickness, such as manual exploration and core sample drilling, are unreliable as ice thickness can vary significantly between core sample locations and critically, these methods necessitate individuals risking their safety. As the uncertainties in temperature patterns continue to escalate with climate change, the necessity for a dependable and safe technique to assess ice safety and thickness becomes more urgent.

The purpose of this capstone project is to propose a solution for small scale, airborne measurements of ice thickness on frozen lakes and ponds. The proposed design couples a radar module attached beneath a drone with a mobile application to map the thickness of ice on dubious areas on frozen bodies of water. By measuring the time of radar reflections at the air-ice interface, and the ice-water interface, the ice thickness can be determined. The collected data will be displayed as a gradient over the body of water, enabling users to make informed decisions rather than blindly venturing onto the ice. This module is intended mainly for municipal recreation departments, ice fishermen, and frequent ice skaters.

The lightweight PLA 3-D printed module contains a radar sensor with a downward facing etched antenna, and a custom PCB, both powered by a 3.7V LiPo battery. On the PCB, there is a communications array controlled by a TI-series microcontroller, used to process the data from the radar module, as well as a modem which is used to transfer the data from the microcontroller to the server. The modem utilizes an external antenna connected to the housing to transfer the data to the server.

The software is distributed in three tiers. The firmware embedded in the microcontroller controls the GPR sensor and the modem, parses incoming requests and packages outgoing data. The highest layer—the user interface—is a react app containing a Google Maps instance to display data, menus, and information about the status of the device. The cloud layer lies in between the others and contains all the code necessary to communicate between the UI and microcontroller, including APIs to send commands and receive, parse, and store sensor and location data in the database.

Air Quality Unit for Interactive Logging and Observation (A.Q.U.I.L.O)

Team Members: Drew Balfour, Jack Duggan, Manas Mahajan, Ben Thacher, Eli Wiel, Erik Zilber

Advisor: Professor Bahram Shafai

Abstract

Numerous health issues such as respiratory and cardiovascular diseases, allergies, and increased risk of infections can be attributed to increased levels of airborne pollutants. These pollutants can originate from everyday household items such as gas stoves, disinfectants and cleaners, and paints and varnishes. To mitigate these health risks, the AQUILO team has designed a compact IoT device that tracks air quality data and enables users to view it in an application on their mobile devices. AQUILO implements a system of sensors that observe a variety of metrics that contribute to the overall air quality of a space, including temperature, humidity, particulate matter, CO2, CO, NO2, and O3.

To use AQUILO, the user must power on the device using either two 9V batteries or a 12V power supply. Once powered, the user can pair the device to their phone and connect it to their home network. Once operational, the device will collect measurements on a 5 minute interval. Multiple devices can be linked to one phone. The home screen of the application displays data for each device, including the name, the air quality index, and the state of charge. The user can generate a plot that can be customized to view any of the measured metrics for a specified time period.

For testing and validation, each of the selected sensors were verified to be working as intended prior to the assembly of the custom PCB. The sensors were tested at the Blue Hills reservation, the Weymouth Compressor Station, the Back Bay Commuter Rail Station, the Park Street Red Line Station, the Boston Common, the Massachusetts Turnpike, and a residential living space in Brighton. These testing sites were anticipated to encompass a wide range of overall air quality due to varying levels of nearby pollution sources. Once the PCB was assembled, its behavior was validated using a lab bench power supply and oscilloscope. An assembled PCB was then taken back to each of the testing locations to ensure that there was no discrepancy between the readings.

Hydrilla Hunter

Team Members: Lisa Byrne, Arjun Fulp, Jvalanti Prasad, Rohit Josyula, Colin McKissick, Darsh Jha

Advisor: Professor Charles DiMarzio

Abstract

Hydrilla is an invasive plant species that causes a myriad of issues for Connecticut waterbodies both environmentally and economically. The plant grows rapidly over the summer, growing up to an inch a day, and dies away in the winter. The fragmented plant particles created when the plant dies down in the winter create oxygen depletion zones that kill the local aquatic organisms. In the late summer, when the plant is most prolific, hydrilla can clog channels and get stuck in the motors of boats making it hard for recreational fishing and boating activities. The clogged boat motors can further spread hydrilla to new waterbodies as the plant is carried in the motor if not fully cleaned out. Therefore, it is integral for Connecticut scientists to clear waterbodies in hydrillas' early growth stages to minimize the spread.

Hydrilla Hunter is the integration of technology from a team of mechanical engineering students providing a semi-autonomous surface vessel (SV) and a team of electrical and computer engineering students providing a computer vision system, renewable energy source, and user interface for scientists to easily track the spread of hydrilla. The computer vision system utilizes a hyperspectral camera to collect spectral information about the scene below the SV. This information is run through a custom convolutional neural network to predict the detection of hydrilla in the captured scene. The GUI system automates the mapping process by dropping pins on a map where hydrilla is detected within about a 1.5-meter accuracy.

This system benefits the Connecticut Agricultural Experiment Station (CAES) as it allows the scientists to survey water bodies with more ease and in turn have more time back to address the problem while the hydrilla is still in its early growth stages.

E-Mobility Charging

Team Members: Alexander Zabib, Ali Hussein, Ezekiel Gunnison, Grady Kirsch, Michael Montanaro, Valentina Ribeiro

Advisor: Professor Masoud Salehi

Abstract

Our team has designed a solar powered charging station for e-mobility that provides accessible charging to university students whose institutions place restrictions on indoor charging. The rise and popularity of modes of e-mobility like electric bikes and scooters has led to cases of battery-related fires on campuses; pressuring administrations for issuing bans on e-mobility. This design proposes an outdoors charging station, employing a charge controller and inverter, whose power delivery is regulated by an app.

The system contains a 3x2 solar photovoltaic array arrangement using a total of 6 panels that is pole-mounted at a height of 7 feet atop a concrete-anchored base. The electronic components are housed in a NEMA 3R-rated enclosure, and the system is grounded with an external earthing design utilizing 3 parallel steel electrodes. The enclosure contains circuit subsystems consisting of a charge controller and inverter. The charge controller connects the solar panel to the internal battery of the system. The purpose of the converter is to regulate the voltage produced by the solar panel to match the specifications of the internal battery, while employing a maximum power point tracking (MPPT) algorithm to maximize the power generated by the panels. The controller responsible for this regulation is capable of temperature measurement of the internal battery and input/output voltage sensing. The inverter is designed to draw from the battery and output 120VAC, providing up to 300W to the e-mobility. It employs an ESP32 utilizing a Sinusoidal Pulse Width Modulation (SPWM) to generate the output sine wave, as well as facilitate seamless communication between the app and the physical charging system. An MQTT broker setup using Eclipse Mosquitto is employed, ensuring efficient real-time data exchange. Hosted on an EC2 instance, the MQTT broker establishes a scalable and reliable connection, laying the foundation for potential future integrations in smart charging solutions. This innovative approach enhances the charging system's functionality, contributing to a more efficient and controlled energy transfer process. This control scheme underscores a streamlined approach to mobile app development, emphasizing security, user experience, and advanced communication protocols, complemented by a novel technique of SPWM for optimal charging system performance.

To model the scalability of the microgrid, a simulation was performed using Simulink and Simscape electrical toolbox. The model shows the main subsystems, simulates Boston irradiance data for the panels, and includes an additional interconnection with the national grid. The system aims to model a scaled-up version of our design, capable of a larger load that supports up to ten E-bikes, sourcing a solar farm array and large battery storage, and employing a bidirectional, grid-connected inverter designed that complies with national grid standards regarding voltage and frequency thresholds.

Testing and validation was performed on the charge controller and inverter. The charger controller had its MPPT abilities tested on a solar panel. The inverter was tested for its ability to generate 120VAC for a range of voltages that simulate the charge profile of the battery. The powering of the load was initiated by the app to emulate a user charging their e-mobility at a station, as the goal of the design to provide a reliable and accessible charging method for students that overcomes indoor charging restrictions.

SPEAK: An Innovative Wireless Chording Peripheral for Real-Time Vocal Communication

Team Members: Conner Ellis, Dan Pitsch, John Rodrigues, Thomas Colao, Kontantin Vasilyev, Andre Danahy

Advisor: Professor John Kimani

Abstract

Aphonia, the loss of the ability to speak, can result from various conditions, including diseases, trauma, or neurological issues. The SPEAK (Single Peripheral for Ease of Accessibility via Keyer) peripheral addresses the communication challenges faced by aphonic individuals, providing a wireless, one-handed chording keyer that enables real-time vocal communication. SPEAK represents a significant advancement in assistive technology for individuals with aphonia, combining the efficiency of chorded input with a unique wireless, portable design and integrated text-to-speech capabilities to enhance communication options and improve the quality of life for those with speech impairments.
SPEAK consists of a portable and ergonomic chording keyer, which users can operate with one hand, and a hip-mounted computer for processing. This setup allows users to communicate vocally in real-time by keying inputs single-handedly, keeping the non-dominant hand free for other tasks. The keyer leverages three 5-way switches for chording, along with an additional switch on the thumb for operation. It employs a Raspberry Pi Pico to record and stream keyer input wirelessly to a Raspberry Pi 4, the platform for the HMC (Hip-mounted computer). The HMC converts the key signals into "chords", or arrays of character strings, which are then transformed into sentences or phrases depending on the chord map being utilized. These are processed by PIPER, a text-to-speech program, facilitating real-time speech communication.

Chording, defined by multiple simultaneous key presses or inputs, allows for faster input compared to traditional keyboards, reducing keystrokes and input time. Customizable chord maps add personalization and ease of use. Unlike existing two-handed or desktop-oriented chorded keyboards, SPEAK is designed for one-handed, portable use, making it a more practical solution for aphonic individuals. Its customizable chord maps and ergonomic design make it accessible and user-friendly, especially for the elderly and those with limited mobility.

SPEAK represents a significant advancement in assistive technology, offering a novel approach to real-time vocal communication for individuals with aphonia. By combining the efficiency of chorded input with wireless, portable design, and integrated text-to-speech capabilities, SPEAK offers a unique solution that promotes independence and provides a novel solution for real-time vocal communication. Its user-friendly design, emphasizing customization and ergonomic considerations, makes it particularly accessible to the elderly and those with limited mobility. This project not only enhances communication options for aphonic individuals but also contributes meaningfully to the broader field of assistive technologies.

Project Limelight: Modern Live Lightshow Depiction for Arbitrary Audio

Team Members: Isaac Kaper, Abdulla Alshehhi, Bradley Perritt, Nelson Barnes, and Smit Patel

Advisor: Professor Bahram Shafai

Abstract

In the era of technological evolution, the "Project Limelight" stands as a testament to the innovative integration of audio-visual experiences. Led by an adept team under the guidance of Professor Bahram Shafai, this project revolutionizes the way we perceive lightshows. It utilizes a state-of-the-art Nvidia Jetson running Ubuntu to create a dynamic lightshow, synchronized with any given audio input.

The process begins with audio input, either through Bluetooth or an auxiliary connection, feeding into the Jetson. A crucial aspect of the project is the 3-second audio buffer, which provides the necessary context for the audio algorithm. This innovative algorithm, a core component of the system, interprets the audio and orchestrates a visually stunning lightshow. The sophistication of the system is further highlighted by its ability to simultaneously release audio to the speaker while dispatching light commands to a microcontroller.

The Raspberry Pi Pico, connected to w2812b individually addressable lights, acts as the executor of these light commands, bringing the algorithm's vision to life. The result is a mesmerizing display of lights that dance in harmony with the audio, capturing the essence of the sound in a visual spectacle.

There are rudimentary LED music visualizers that work in real-time but only track large, sudden spikes in volume such as bass kicks. There are also music analysis programs that can do very complex analysis given the data for an entire song. What we are trying to do is bridge the gap between these two methods. We take the audio signal of a song and buffer it for a few seconds so that we can perform analysis on the music while looking slightly ahead into the future. This lets us create the effect of a real-time music visualizer while also having some of the benefits of more nuanced signal analysis techniques.

This project not only exemplifies the seamless fusion of audio and visual elements but also sets a new standard in live entertainment technology. It opens new avenues for exploring the relationship between sound and sight, promising an immersive experience for audiences. The ingenuity of Project Limelight lies in its ability to adapt to any audio input, making it a versatile tool for various applications, from concerts to private events.

In conclusion, Project Limelight is a groundbreaking venture that redefines the boundaries of audio-visual synchronization. Its innovative approach, technical prowess, and adaptability mark a significant milestone in the field of live entertainment technology. The project stands as a beacon of innovation, showcasing the potential of combining art and technology to create extraordinary experiences.

Garden Robot for Overseeing Water Treatment in Hydroponics (GROWTH)

Team Members: Tasmiha Amir, Susan Badanai, Rida Jawed, Tracy Qiu, Adam Reed, David Savoie

Advisor: Professor Charles DiMarzio

Abstract

The GROWTH team has designed and implemented a system to automate the regulation of light and water conditions in a hydroponics farming operation. While solutions exist which may target either consumer or industrial applications, GROWTH's unique design presents a more flexible and scalable solution. The GROWTH solution uses two linear rail axes to intelligently move its sensors between completely isolated water circulation systems, enabling it to regulate the nutrient and pH levels differently for each isolated environment and allowing for easy expansion and flexibility as needed.

GROWTH provides users with real-time insights into the water conditions of distinct plants within the system, automating adjustments to cater to their specific needs. The horizontal rail allows the system to move between several plots, detect fiducials using computer vision to accurately locate and identify the plants in each plot. The system contains water conductivity, pH, and temperature sensors which are attached to a 3D-printed dispense head that is moved by a vertical rail. To stabilize the water conditions of each hydroponic set up, the dispense head is also equipped with four peristaltic pumps that automatically dispense nutrients, and acid/base solutions as necessary. Grow lights controlled via relays provide the appropriate duration of direct light exposure for each individual plant. A custom PCB with an RP2040 microcontroller was developed to carry out control of the functions detailed above. Custom firmware written in C controls the hardware and sends and receives serial data from a Raspberry Pi. The Raspberry Pi hosts a Python Flask web server connected to a GUI using a React client and MongoDB database. Users can create their own account and login to view their data. Through the GUI, the user can specify the number of plots present in their hydroponic setup, and the GUI displays the plots' data (type of plant, last calibration time, last run time, temperature, pH, conductivity) in an accessible format. All the plant data is stored in the database.

To test the design, basil, scallion, and lettuce plants were placed in the system. A calibration routine uses computer vision to detect the relative positions of the plots and a measurement routine takes sensor readings of each plot and actuates pumps to dispense the proper solutions to correct water quality. Using these processes, appropriate water conditions are maintained with little user intervention. By using only one set of sensors and pumps, GROWTH can also keep costs down when scaled to much larger numbers of plots.

Smart Hydroponic Countertop Garden

Team Members: Devon Holdgate, Sam Sheehan, Priya Khadilkar, George Hagopian, Joseph Pak

Advisor: Professor Masoud Salehi

Abstract

The Smart Hydroponic Countertop Garden Team has designed and built a system for urban consumers to affordably and sleekly integrate hydroponic food production into their own home. Such a system allows users to grow their own microgreens and small vegetables in the comfort of their home, without requiring soil or other messy and often expensive gardening essentials. The intelligent control, automation, and alert system allows users to monitor all status information in an easy to understand way. In response to most stimuli detected by onboard sensors, automatic responses are deployed to ensure the garden maintains optimal conditions without any input from the end user. The purpose of this design is to create an affordable and extremely low maintenance entry into the concept of urban hydroponic gardens, as raising awareness of this type of farming is critical to developing the support necessary for social change that would incentivize the adoption of this highly efficient technology.

Our countertop garden uses a Raspberry Pi as the main controller of the system. It monitors the garden through a camera and various other sensors (such as humidity, temperature, and water level). Using the Pi, the water pump, humidifier, fans, and lights can all be automatically controlled, without need for user intervention. To allow users to interact and obtain feedback on the system, we created an iOS app that receives live sensor data through a Flask API endpoint on our Pi. An image recognition API capable of identifying health problems in plants is also utilized through a camera on our Pi. It suggests optimal remedies, which the system may automatically implement, or alternatively, it will alert the user on the app with detailed corrective instructions. As it stands the only necessary user interactions include sewing seeds, refilling water, and harvesting, notwithstanding emergency health issues that may be the result of pests, sun burning, or other causes of premature plant death.

To test the system, we went through two trial growth phases (time-limited) with both the software (user app, database, server) and firmware (state machine, sensor control) fully integrated into the hardware. We used one trial to monitor how well the system acted under normal conditions, with an uninterrupted growth period, and the second trial to test how the system responded to special cases (burnt leaves, dirty water, too much sun, too hot, dead plants) by manually causing them. Using these two trials we were able to calibrate the firmware's sensor controls and troubleshoot errors in user alerts and connectivity issues.

EMAK: Emergency Medical Assessment Kit

Team Members: Siddharth Simon, Michael Delaney, Nicholas Balog, Casey Holden, and Sean Persaud
Advisor: Professor John Kimani

Abstract

The Emergency Medical Assessment Kit (EMAK) is a system aimed at instructing members of the public on how to aid in medical emergencies before professionals arrive. The system integrates an interactive Graphical User Interface (GUI) on a Raspberry Pi 4 connected to a touchscreen tablet and essential medical devices, including a glucometer, pulse oximeter, and blood pressure monitor. Housed in a portable case with user-friendly instructions and a rechargeable power source with approximately 5 hours of battery life, the EMAK aims to give efficient and effective responses reliably in critical situations.

The EMAK runs a Python application on its Raspberry Pi 4 that includes the GUI, serving as the primary interface between the user and the system. Upon starting the EMAK, the user is prompted with a screen informing them that they are protected by the Massachusetts Good Samaritan Law before providing diagnostics and treatment solutions. While this message would be configured depending on the state, a user's informed involvement will be covered by Good Samaritan Laws that are largely present in all 50 states. After receiving user consent, the software employs a decision tree structure, informed by discrete user inputs and vital sign information, guiding user actions prior to EMS arrival. This guidance can range from instructions and a tempo for performing CPR to directing the diagnostics and stabilization of a stroke patient. In instances where the user is aware of the situation's cause, there is a quick access panel that provides efficient navigation to specific procedures related to the incident. Users can also end the instructions at any time and get a breakdown of vital information when professionals do arrive. The data logic, collection, and communication are performed using a model-view-controller architecture and a device adapter interface to communicate with the external diagnostic tools.

The EMAK's medical devices include a glucometer, a pulse oximeter, and a blood pressure monitor. These hardware components are essential for assessing critical vital signs, enabling users to assess and respond to emergencies more effectively. The glucometer reads blood glucose levels, allowing swift identification and management of diabetic emergencies. The pulse oximeter connects to an Arduino that calculates blood oxygen levels and measurement confidence based on the pulse oximeter's raw data. This data is then transferred to the Raspberry Pi 4 through its GPIO pins. Additionally, the blood pressure monitor has the calculated patient blood pressure read from its EEPROM chip to the GPIO pins on the Raspberry Pi 4 using I^2C protocol.

By including these vital readings, the EMAK interprets and displays real-time patient information, improving its ability to suggest informed directions for the user to follow. The kit is physically designed to be robust, portable, and accessible, ensuring that non-professionals can efficiently and confidently rely on the kit in emergencies. In doing so, the EMAK holds the potential to help save countless lives by improving pre-hospital patient outcomes and giving members of the public the confidence to be good Samaritans.

Sound Floats: A Low-Cost Acoustic Release

Team B3 Members: Isabelle Brandicourt. Max Brotman, Charlie Daigle, Luis Soto, Joshua Cheng, Evan Tran

Advisor: Professor Bahram Shafai

Abstract

The Sound Floats team has designed a low-cost acoustic release device for oceanographic device management. Acoustic release mechanisms use sound waves to communicate between the ocean surface and the seafloor, allowing users to retrieve their deployed devices without the use of invasive rope and buoy tethering systems. Researchers use these types of tools to gain critical insight into areas that are difficult to monitor long term, providing data that is vital in this age of severe climate change. A major issue with existing acoustic release mechanisms is cost – generally between $5,000 and $10,000. While there is a significant cost reduction from the industry standard releases to the research releases, there are consequential compromises in device range and depth ratings.

We propose a relatively low-cost acoustic release mechanism ($1,000 - $2,000) with comparable functionality to those currently on the market but specified to bioacoustic research needs. Our system comprises a deck box for human interfacing, a set of transducers for acoustic communication, and a release mechanism that attaches to the oceanographic sensors of choice. The deck box implements a Raspberry Pi to run the user interface, read the status of all devices connected, document deployment, send the retrieval signal, and document collection. The hydrophones operate within the 20-35 kHz range as optimal to complement bioacoustics monitoring needs and can both send and receive signals via the frequency shift keying (FSK) modulation scheme. The release mechanism operates from a low power mode to enable longer deployments (~1 month). Once a wake-up signal is registered, it can send its status and wait for the release command. Once the release command is registered, there is a motor-actuated release that unscrews the entire system from an anchor plate. Flotation rings make the system buoyant once the anchor is released, and the entire system can float to the ocean surface for retrieval.

Our system differs from existing acoustic releases by communicating over a longer range than typical research modems but costing significantly less than industry modems as a result of decreased bitrate and battery life. This will enable more ecologically minded entities to purchase a higher quantity of release mechanisms and deploy their own large-scale arrays of oceanographic equipment.

MoCAP: Method of Capturing Anatomical Position

Team Members: Katherine Cloutier, Christian Cole, Matthew Low, Noah Mitchell, Zachary Spiegel

Advisor: Professor Chuck DiMarzio

Abstract

The MoCAP team has designed and built a battery-powered, wearable motion tracking device which can be used for a wide variety of applications, including the medical field (physical therapy and rehabilitation), motion capture (video games and film), and virtual reality gaming. Currently, the market for motion capture solutions varies by price, accuracy, technological limitations, and applications, with high-quality tracking systems relying on laser base stations that are expensive and require open spaces for a line-of-sight between the trackers and base stations. Many of the current fully wireless solutions use lower-cost inertial measurement units (IMUs) which are susceptible to large drift errors. Our objective was to create a high-precision IMU-based motion capture solution that addresses a gap in the market, filling the needs of users who are after the freedom offered by IMU-based capture solutions and the quality that has previously been exclusive to laser base station tracking.

MoCAP consists of three major elements: the self-contained tracking devices, the host computer, and a 3D visualization of the tracking results. The custom designed and manufactured printed circuit board (PCB) contains all the necessary circuitry to capture real-time motion data. There are four major systems on these custom boards: the power system, microcontroller (MCU), IMU, and wireless module. The power system manages the module's lithium-ion battery charging and voltage regulation for all components. A central piece of the system is the MCU, which handles all the communication protocols with the IMU and battery system, as well as all the front-end processing to transform the raw data before transmission to the host computer. To capture the real-time angle (gyroscope) and velocity (accelerometer) of the tracking device, an IMU is used, which contains the necessary sensors to capture this information. In the final stage of the data path, the wireless module takes transformed data from the MCU and transmits over user datagram protocol (UDP) to the host computer. A UDP server is run on the host, and all tracking devices are connected wirelessly. Each arriving UDP packet contains a unique identifier to distinguish which tracker the data has come from. With this information, the real-time data can be applied to the visualization to provide an accurate representation of the object being tracked.

For system functionality, a total of six trackers were assembled and validated as functional. In the project's current state, these trackers can be used for human motion tracking, where two trackers are placed on the legs, two are placed on the arms, one is placed on the chest, and one is placed on the head. The results we have obtained demonstrate that with precise IMUs that have low drift error, wireless motion tracking is feasible without the need for bulky base stations. Given the current design, there are a few areas where improvements can be made to the overall system functionality. Primarily, there are some limitations with the speed of motion, as the gyroscope dynamic range is not ideal for this type of application. However, given the time, budget, and constraints of the project, we were successful in exploring the feasibility of high precision IMU-based motion tracking solutions.

The Dojo: Intelligent Arcade Controller

Team Members: Nalin Storm, Viraj Lawande, Tyler Jackson, Yiqi Dou, Kevin Fu, Femi Lamptey

Advisor: Professor Masoud Salehi

Abstract

The Dojo team has designed and built an arcade controller and software package for the game Street Fighter 6 that provides the player with feedback on combo execution. Players can string together sequences of button presses with specific timings to perform extended actions in the game, referred to as "combos." The game features a training mode which allows players to practice combos outside of a competitive match environment, however there is no way to indicate a desired combo to practice, search for combos from other players, or receive direct feedback on combo execution. The Dojo design proposes a device with the form factor of a traditional arcade controller that can download a sequence of button presses, or "combo code," from a video analyzed by a machine learning (ML) model and provide players feedback on their inputs while practicing a saved combo.

To use the Dojo controller, a user will upload a video clip of the desired combo to the ML model which will analyze the video frames and output the sequence of moves. The model was trained on thousands of annotated images of the character Ryu, with each image tied to one of several moves he can execute. Once analyzed, the combo code is sent to a Firebase server which the Dojo controller can then download and save locally on the Raspberry Pi inside. Using the attached touchscreen at the top of the controller, the user can then select any saved combo to begin practice. The user has the option to select a demonstration mode which will illuminate the buttons in the order they must be pressed and at the ideal time. Once the user has selected a combo, they can begin to practice as normal. The buttons on the dojo controller are connected to a breakout board which splits the signal between a Brook Universal Fighting Board responsible for passing the inputs to the game as a standard controller would, and the Raspberry Pi which handles the input processing for evaluating player inputs and the touchscreen component. Once an attempt has been completed, the Raspberry Pi parses the inputs into the game specific format, and these inputs are then compared against the saved combo using a lookup table of frame data. Since the game runs at a locked 60 frames per second, or 60Hz, the widely available information about each character's move durations and properties in terms of in-game frames can be converted to time and referenced against the time-stamped player inputs. Finally, the display can return information about each player input, such as whether a button was pressed X number of frames early or late for it to properly create the combo or if the correct button was pressed. An internal battery connected to the Raspberry Pi also allows the device to be used when not explicitly connected to a game console, such that users may continue to practice saved combos while away from the game.

For testing and validation, sequences of known button inputs were sent to the controller while connected to the game to verify that the sequences displayed on the game screen match the results of the analyzed inputs on the Raspberry Pi. The results indicated that despite the inherent variability of player input, the controller is reliably able to detect the intended game output, and the ML model was consistently able to detect the correct moves with a high degree of accuracy.

Distant Duels

Team Members: Lulwa Alghanim, Sarah Moyer, Kai Johnson, Katie Nikolov, Aidan Di Salvo, and Alejandro Manuel Campa

Advisor: Professor John Kimani

Abstract

Distant Duels blends traditional, face-to-face chess with online gameplay. It serves as a bridge between the physical and virtual chess worlds by integrating with LIChess, an online chess platform. With Distant Duels, your moves are transmitted to LIChess, creating a synchronized experience. As your opponent responds, the chessboard illuminates the coordinates of the moved piece and its final destination.

Made of high-quality MDF imported from Germany, the board was designed, first in CAD, in compliance with the tournament standard chess board size of 21"x21" in outer frame perimeter and a 16"x16" tile surface. The tile surface consists of opaque white and translucent black 1/8" thick cast acrylic sheets that have been laser cut to 2"x2" squares per the United States Chess Federation (USCF). Fastened to the board is the timer block that houses the arcade pushbuttons and LCDs. The components on the timer block allow the player to set game-play settings and complete moves.

The housing component, located in the heart of the board, consists of four 3D-printed quadrants designed for the precisely angled positioning of RFID readers, LED strips, and wiring. These quadrants can be lifted for facilitated access to wire hubs and for making sectional adjustments.

Embedded into the chess pieces are RFID tags (antenna and chip) extracted from key fob casings adhered to the bottom of the pieces and secured with fabric. These RFID tags within the pieces are registered by the RFID readers found below the checkered acrylic. When a move is made, the RFID readers will register the RFID tag found within the piece placed. When the player hits their button on the timer block this move will be uploaded to LIChess. Then when an opponent makes a move the coordinate points, such as A1 or F5, are sent from LIChess to the Raspberry Pi. The Pi lights up LEDs that correspond to that chess cell. The RFID readers register the piece placed through the RFID tags and verify that the piece was placed in the correct coordinate. If not the board lights up red and will not allow the player to proceed until the piece is in the correct location.

Object Detection Visual Assistant (ODVA)

Team Members: Rose Hensley, Mei Lin Mulvey, Aditi Purandare, Samantha Wong, Min Jung You, Jeffrey Zhang

Advisor: Bahram Shafai

Abstract

The ODVA team has created a compact, wearable device capable of detecting, analyzing, and describing the user's surroundings to empower individuals with visual impairments. The applications of the ODVA include many everyday tasks like reading labels on products, distinguishing between tactilely-identical items, and navigating one's environment. For these common tasks, someone who is blind or visually impaired often must rely on the assistance of a sighted person. Instead, the ODVA enables a user to increase independence in their daily life. These benefits are possible for individuals with varied levels and types of visual impairments due to the flexible nature of the ODVA responding according to the specific requests of each user.

Appearing as a hat and sling bag, the ODVA is portable and discreet. When the user desires the assistance of the ODVA, they must press a button and speak their inquiry out loud. Upon the button press, the camera captures an image of what is in front of the user and the microphone captures the audio request for the duration of the button press. This audio input goes through natural language processing (NLP) run on the Jetson Orin Nano to generate the request in text. This request along with the captured image are run through a combination of Azure and GPT-4 to identify necessary information for the user's question, such as present object types, labels, and a probabilistic map of the locations of the objects. The image analysis is carried out in Azure, and the text description generated is processed by GPT. The localization is carried out by a combination of bounding boxes and NLP. This location data is further processed by an ODVA-specific algorithm to more precisely determine the location of present objects in relation to one another. Finally, an audio response is generated by another ODVA-specific algorithm to appropriately answer the user's question from the extracted information.

Testing and validation was done with the developed fully functional system running on a Jetson Orin Nano powered by a 12V lithium-ion rechargeable battery. Testing focused on the scenario of distinguishing the type and location of tactilely-identical objects. Specifically, three soda cans of identical size were placed together on a table for the test case. Various questions were asked, such as "What are the types of soda?" and "Where is the Coke?". The results indicated that the methods of speech processing, object detection, localization, and speech synthesis were satisfactory. The ODVA team believes that present testing confirms the system's proficiency in reading product labels, identifying objects, and determining relative object location; future testing is required to determine the system's capacity for navigational assistance.

TouchTyper

Team Members: Nathaniel Dorson, Aidan Dougherty, Timothy Frieden, Nicholas Gale, Christopher Lackowski, Steven Scangas

Advisor: Professor Charles DiMarzio

Abstract

The TouchTyper team has built gloves which track finger movements and allow translation of the wearer's actions to keystrokes on an imaginary keyboard. The gloves constantly stream sensor data to a host device, where the data are run through algorithms to track all the finger positions, and these positions are mapped onto a virtual keyboard, allowing the wearer to type without the use of an actual keyboard. This set of gloves is more portable than a normal physical keyboard, as well as more discreet. In theory the gloves could be used with varied devices such as smartphones, tablets, laptops, and desktops. The core piece of technology – the finger tracking – also can be imagined to have further applications such as for sign language translation, playing MIDI instruments, or controlling robots, to name a few.

To use the TouchTyper device, the user wears the gloves, which are connected via wires or paired over Bluetooth to the desired computer or other smart device. The host device runs a program to track the inertial measurement units (IMUs) on each fingertip. Acceleration, gyro, and magnetometer data from every finger are transmitted from the gloves to the host device at around 30Hz, allowing the rotation to be tracked and allowing linear acceleration to be twice integrated to determine the position offset of each fingertip. To counteract drift and account for sensor noise, the TouchTyper team implemented a number of filters and heuristics that remove constants of integration and keep the system in a known state at all times. With proper tracking, the keypress detection routine outputs detected keypresses, and also feeds the keypress detections back into the position estimation routine in order to increase confidence in the tracked position with every keystroke.

For testing and demonstration, the TouchTyper team set up a virtual 3D scene containing a keyboard model and sensor models for each tracked finger, and had the sensor models track the physical sensor positions. In this way, it was possible to visually assess whether the sensors were being tracked correctly and were hitting the correct keys. A single set of gloves was created to show the proof of concept of finger tracking and the application thereof for keyboard input. Cost restraints dictated that an affordable IMU be selected for the project, but even more success could be achieved with higher grade sensor components, since sensor inconsistency caused the most issues when trying to achieve accurate tracking and keypress detection. When prices for such sensors drop in the future, these sorts of input devices may be more viable.

Piano Pedal Accessibility Device (PPAD)

Team Members: Luke Wisner, Roberto Chadwick, Lucas Laya Marina, Hanchen Liu, Josh Moffat, Tristan Siu

Advisor: Professor Charles DiMarzio

Abstract

The PPAD team has created a simple device and iPhone app to simulate pressing pedals in a digital piano, using only the movements of your head. This concept has many applications to help people with disabilities that would otherwise prevent them from being able to use the piano pedals as they play. Jonathan Goodwin is one such person, being paralyzed from the waist down. Jonathan is a family friend of a member of our group, so we wanted to deliver a product to him that would once again make playing his piano fully possible. With Jonathan's suggestions and guidance, we were able to hone in on the most important aspects of a system like this. Using what he told us, we landed on the user tilting their head in different directions to activate the pedals.

The electrical engineering team members created a custom PCB with an esp32 chip embedded into it as a means of communicating with the app over Bluetooth. The firmware on the chip receives a data packet sent by the app and then activates the pins that correspond to the correct pedals. The board has three built in 1/4-inch jacks to allow the app to send signals for a one or three-pedal piano. Given that three-pedal pianos have proprietary cables, the user would have to create their own custom cable to connect their piano to the jacks, but this functionality is still supported through the app and firmware. The firmware also supports an easy connection through Bluetooth, with unique UUIDs so there is no conflict between multiple devices in the same home.

The software engineering team members designed an iPhone app using SwiftUI and Apple's Vision Library. The Vision Library has functions to designate facial features into eyes, nose, and mouth, so from there we wrote code that can determine the head position. Another important aspect of the app is the settings page that allows the user to select between one and three pedals as well as adjusting individual pedal settings. With only one pedal, the user can choose the tilt direction to activate the device. With three pedals this is more restricted because only certain pedals can be used in tandem with others, however, every pedal can have its angle range adjusted so the user can choose what range feels best to them. In addition, there are options to change the activation to be inverse or held if the song requires the pedal to be held down for a long amount of time. These settings are all processed when the head tracking runs, and then a data packet is sent over Bluetooth containing which pedals have been activated by the user tilting their head past their set boundary. After improving the code to add debouncing and lowered latency, the processes run quickly to be more natural for the user. The design is meant to be simple and intuitive, so that Jonathan and anyone in the future can easily learn it and gain the ability to use their piano fully again.