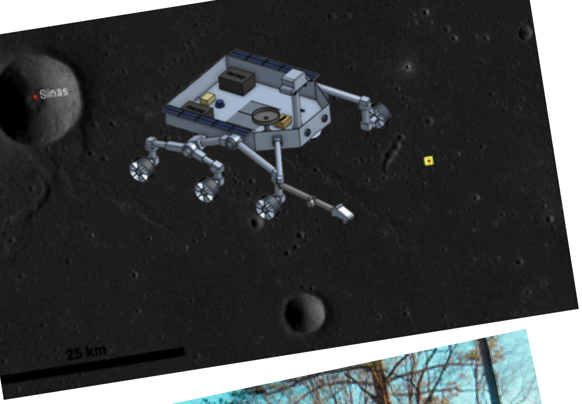




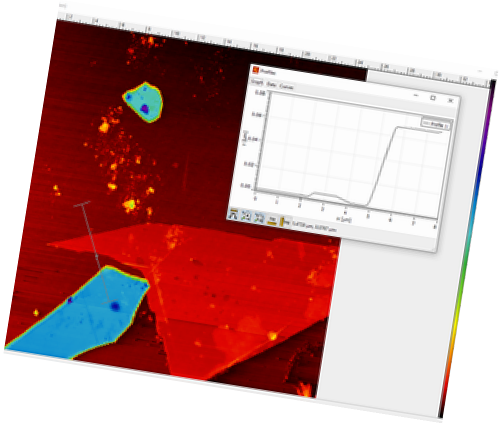


PROJECTS

Vedarsh Mishra

B.S. Applied Physics and CS , Boston College

Expected Graduation: May 2028

mishrave@bc.edu



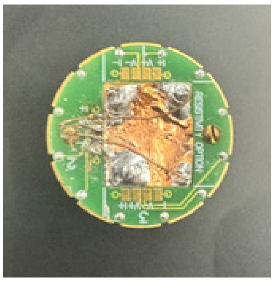
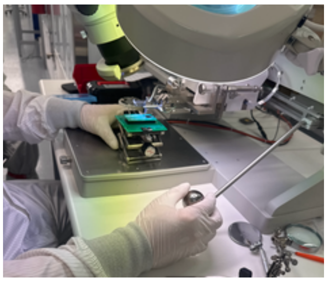

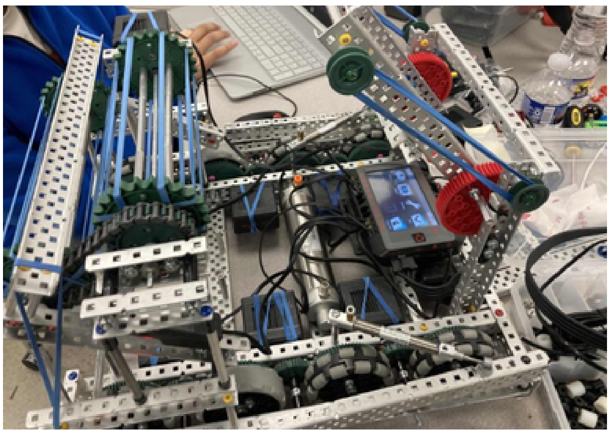
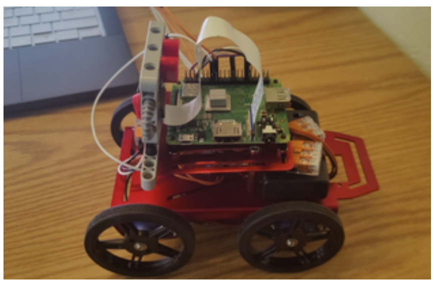

PROJECT INDEX

AUTONOMOUS VISION-GUIDED ROBOTICS CART

PERSON-FOLLOWING MOBILE ROBOT USING BEHAVIOR CLONING

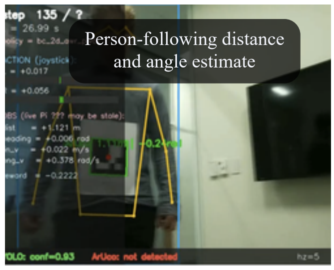

Role: Hardware Lead / Embedded Systems Contributor
Team: 3-person team, Spring 2026
Platform: Raspberry Pi 3B • Pi Camera 3 • 4× FS90R • GPIO/PWM • WAGO Power Distribution
Perception : YOLO-Pose • ArUco • Behavior Cloning • FastAPI/MJPEG

Overview: We built a vision-driven cart that **leads a person rather than following behind** them. A single backward-facing camera estimates the operator's distance and bearing, a compact behavior-cloned policy converts that state into left/right wheel commands, and the Raspberry Pi runs perception and inference at **5 Hz** on the edge. **My part** of the project was between software and hardware: getting the Pi, camera, PWM, servos, power distribution, and diagnostics to behave like one system instead of six separate parts. GITHUB

System Architecture:

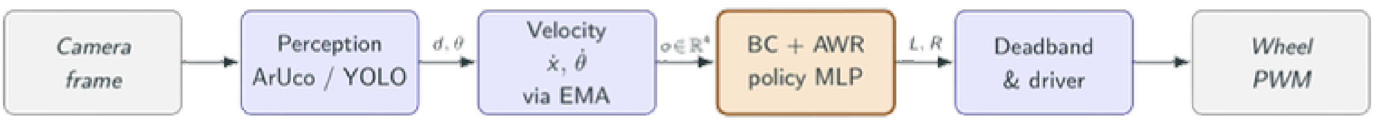

Only the behavior-cloning policy was learned; perception, velocity estimation, deadband logic, and wheel control were fixed geometric or rule-based stages.

Key Contributions:

- Developed the hardware-software interface across the Raspberry Pi, camera, GPIO/PWM control, servo wiring, and physical cart platform.
- Built the headless Pi and diagnostics/dashboard workflow used for camera streaming, control, testing, and data collection.
- Debugged critical hardware failures including servo brownouts, PWM jitter, overheating, camera issues, and repeated Pi failures.
- Supported final evaluation through hardware maintenance, camera replacement, and manual cart operation for data collection; produced the wiring and chassis technical figures.

Hardware:

Raspberry Pi 3B
1 GB RAM • no GPU
Pi Camera 3
Sony IMX708 • 640×480
4× FS90R servos
6 V • 130 RPM • 50 Hz PWM
GPIO
3.3 V logic
Power
5 V compute + dual ~6 V servo rails

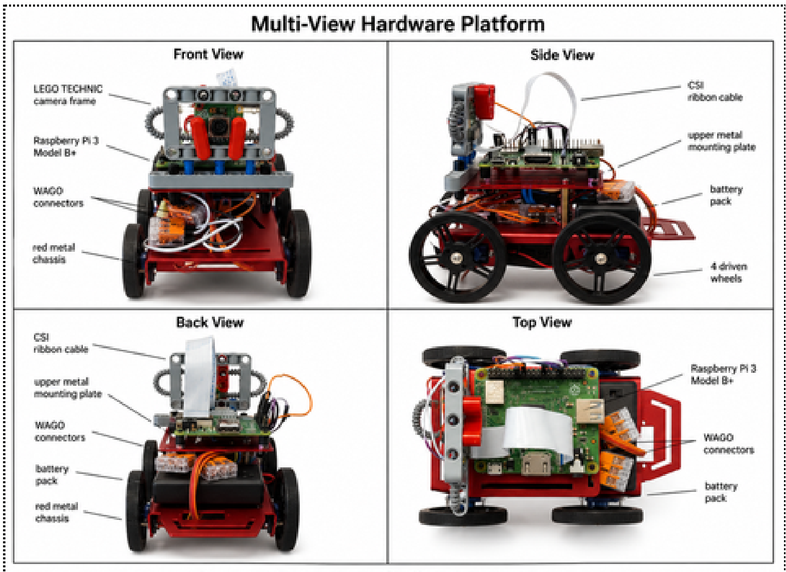

Orthographic-style reference views of the fifth-iteration cart prototype.

System Performance:

Calibration Accuracy	Tracking Reliability	Data Collected	Training Dataset	On-Device Inference
0.34 px	**93.5%**	**6,748 frames**	**2,592 samples**	**5 Hz**
ChArUco reprojection error	ArUco lock across data collection	449.5 s total	1,296 balanced rows + mirror augmentation	policy inference on Raspberry Pi

The fifth-generation platform integrated camera perception, GPIO/PWM actuation, browser diagnostics, and on-device policy inference, with stable operation after separating compute and servo power.

Index

AUTONOMOUS VISION-GUIDED ROBOTICS CART

REASONING & DESIGN TRADE STUDIES

Goal: Work toward a real-world application such as hands-free strollers or shopping carts for high-end malls and retail environments. V1 was intentionally kept simple as a low-cost indoor prototype to validate the core perception, control, and hardware stack before scaling.

Camera Geometry & Field-of-View

I needed the camera to reliably see a person at close indoor range, so I used simple geometry to check the field of view and think through camera placement before adding more sensing complexity.

Camera Calibration

Resolution: 640 × 480

$f_x = 493.67$ px

$f_y = 493.52$ px

$HFOV = 65.9°$

Reprojection Error = 0.34 px

Horizontal Field Coverage

$$W(d) = 2d \tan \frac{HFOV}{2}$$

At $d = 1$ m:

$$W(1) \approx 2(1)\tan(32.95°)$$

$$W(1) \approx 1.30 \text{ m}$$

At roughly 1 meter, the camera sees about 1.3 meters of horizontal width, which is enough to keep a person in frame during normal indoor motion.

Bearing from Image Position

$$\theta = \tan^{-1}\left(\frac{u - c_x}{f_x}\right)$$

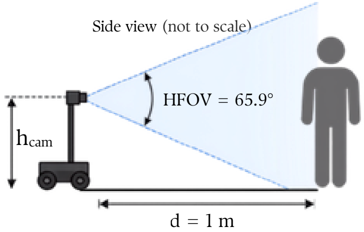

Camera Placement Trade

I wanted the camera high enough to reduce chassis occlusion, but low enough to keep the person large in frame at short range. For V1, the goal was not to find a perfect mounting height analytically, but to use the field-of-view geometry to keep the placement physically reasonable before testing.

Power Integrity & Thermal Considerations

During early tests, the Pi would brown out or reset when the servos started, turned, or reversed. I used basic circuit and thermal relationships to understand why.

Voltage Drop (Why brownouts happened)

$$V_P = V_S - I_{load} R_{path}$$

- Higher transient servo current increases voltage drop across the battery, wiring, connectors, and distribution path.
- This reduced the compute rail margin, causing brownouts,PWM jitter, and repeated Pi failures.

Thermal Margin (Ambient Temperature)

$$T_j \approx T_a + P\theta_{JA}$$

- The Pi also experienced overheating, so I considered how ambient temperature affects thermal margin.
- For a similar workload, higher ambient temperature directly increases device temperature.

Power Architecture Change

Before
Pi + 4 × servos
(shared supply path)
→ brownouts / instability

→

After
5 V compute rail
+
dual ~6 V servo rails
+
common ground (WAGO)
→ stable operation

Ambient T	Estimated T (TT = 45°C)
20°C	65°C
35°C	80°C

A system that is comfortable in a 20°C room can approach its thermal limit in a hotter environment without any change in software.

Sensor Trade Study.

I considered **camera, LiDAR, and radar** for perception. LiDAR would provide direct range through time-of-flight (at a 1 m target, the round trip is only **~6.7 ns**), while FMCW radar could provide radial velocity directly through Doppler shift. Neither addressed a V1 failure mode that justified the added hardware, power, integration, and learning time for this low-speed indoor prototype. I stayed **camera-only for V1** because it already provided person identification, bearing, and enough range information to close the control loop. If lighting, speed, or range accuracy became limiting, I would revisit LiDAR or radar.

What I Would Test Next

I see this prototype as a first step toward **hands-free strollers or shopping carts for malls and other controlled indoor spaces**, not as a finished product. Before adding more sensors or autonomy, I would instrument what we already built: servo current, rail voltage, Pi temperature, wheel speed, command-to-motion latency, and tracking loss versus distance and lighting. Those measurements would tell me what to change next: LiDAR for insufficient range accuracy, radar or sensor fusion for difficult lighting or motion, encoders/IMU for better state estimation, or improved power and thermal design if hardware reliability remains the limiting factor.

Index

MECHAEAGLES

[BUILDING COMPETITIVE ENGINEERING AT BOSTON COLLEGE]

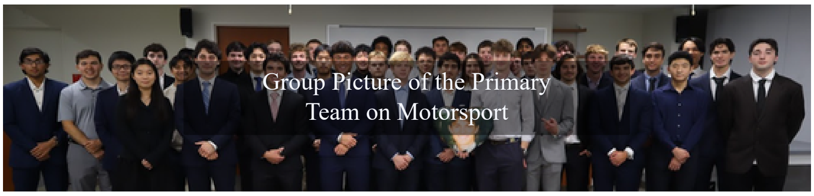
Group Picture of the Primary Team on Motorsport

Role: Founder & Systems Engineer
Timeline: Dec 2024–Present
Scope: Baja SAE • AeroEagles • Fabrication • Avionics • Engineering Infrastructure
Scale: 7–8 initial members → 160+ applicants → 60–70 selected engineers

Overview: I founded MechaEagles after realizing that the university did not have an established student pathway for projects like Baja SAE or high-power rocketry. Much of the early work was not glamorous engineering: working through university approval, securing workspace and welding access, establishing procurement and safety processes, and creating enough technical structure for teams to build hardware responsibly. From an initial 7–8 person group, MechaEagles grew into a 60–70 member engineering program supporting motorsport, aerospace, fabrication, and technical training.

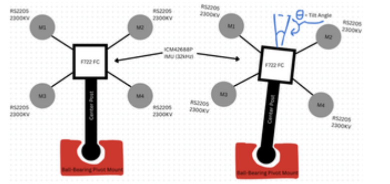

AeroEagles | From Concept to Hardware

AeroEagles began as a three-person effort inside MechaEagles. Our first major project explored a model-scale autonomous landing architecture and moved from low-fidelity modeling into a purpose-built pivot validation rig. Rather than jumping directly into a full drop or flight test, the rig isolated rotational dynamics so we could validate motor response and attitude-control behavior using real flight hardware while reducing the safety and facilities burden of an integrated test.

Hardware / Validation Platform

STM32F722 flight controller • ICM42688P IMU • 60 A 4-in-1 ESC • 4× RS2205 2300KV motors • 14.8 V 4S LiPo • Betaflight • 2 kHz IMU / RPM telemetry

I integrated sensing, power, actuation, telemetry, and bench-test interfaces into one hardware stack and helped connect the low-fidelity model to a repeatable physical test setup. Logged IMU/RPM telemetry at 2 kHz during a 0–70% throttle step and compared the measured motor response against our first-order model: $\tau = 0.28$ s measured vs. 0.25 s predicted, 12% error, within our <15% validation threshold. Full assumptions, plots, attitude-control results, and test limitations are in the linked report.

Outcome / Current Direction

MechaEagles now supports roughly 60–70 selected members, while AeroEagles has grown from the original three-person effort into a dedicated aerospace team of roughly eight members. The next stage is building toward larger-scale rocketry and the technical capability needed for competitions such as NASA Student Launch.

Technical Documentation: Final Report | Mathematical Model | Validation Results


High-power model rocket field test preparation

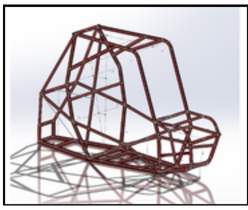

Ever to Excel student group recognition

Student-built motorsport shirt production


Welding preparation for chassis fabrication


Arduino / technical workshop session

Index

ARCADIA LUNAR ROVER PDR

NASA L'SPACE MISSION CONCEPT ACADEMY

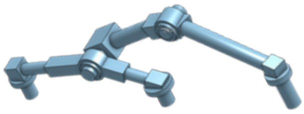

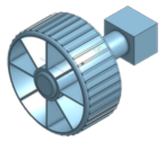

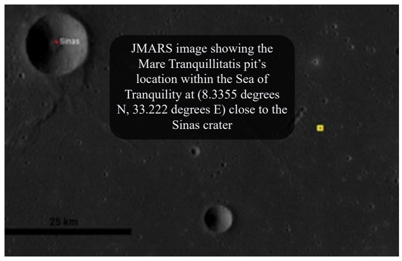

Role: Lead Systems / Project Manager + Mechanical Engineering | 2024
Team: 14-person team
Focus: Lunar Mobility • Mission Architecture • Requirements • Interfaces • Trade Studies • Verification

Mission Constraints:

- Mechanical mass: ≤ 180 kg
- Chassis: ≤ 70 kg
- Suspension: ≤ 30 kg
- Traverse speed: ≤ 0.15 mph
- Final spacecraft concept: 274.82 kg
- Stowed envelope: 2.0 × 1.5 × 1.5 m

Overview: I led the systems and mechanical development of a 6-wheel rocker-bogie rover architecture, translating mission science objectives into engineering requirements, subsystem interfaces, trade studies, and verification logic. My work focused most heavily on mobility architecture, bogie geometry, pivot placement, wheel concepts, packaging, and cross-subsystem constraints, while coordinating the team through formal NASA-style design reviews.

Key Contribution:

- Owned rocker-bogie geometry, pivot placement, wheel concepts, clearance, and mechanical interfaces.
- Managed system constraints including ≤180 kg total subsystem mass, ≤70 kg chassis, ≤30 kg suspension, ≥25 kg science payload, and ≤0.15 mph traverse speed.
- Converted science objectives into requirements, interface definitions, traceability, and verification criteria.
- Led the team through MCR, SRR, MDR and culminating in a 160+ page Preliminary Design Review

Architecture Baseline

- *6-wheel rocker-bogie:* passive articulation to maintain wheel contact and improve stability over uneven terrain.
- *2000-series Al chassis*: selected for aerospace heritage, manufacturability, cost, and exposed-environment robustness.
- *NiTi superelastic tires:* compliant spring-wheel concept chosen for recoverable deformation, traction, and mass reduction.
- *LiDAR + stereo + IMU:* autonomous terrain mapping, hazard detection, attitude sensing, and path planning.
- *Science interfaces:* robotic arm, spectrometer, thermal/radiation instruments, GPR, cameras, lighting, and power/thermal accommodation.

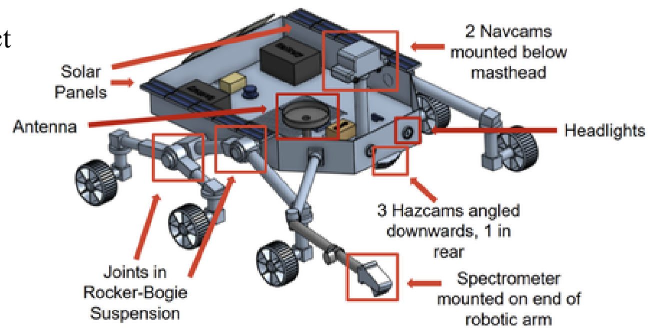

Mobility Trade Study

I treated mobility as a trade between traction, sinkage, obstacle clearance, mass, wheel compliance, and packaging. I compared rigid vs. compliant wheel concepts, tread geometry, rocker-bogie pivot placement, and regolith interaction before converging on the rover configuration used in the PDR.

Ground pressure:
increasing total wheel-ground contact area reduces pressure on loose regolith, helping limit sinkage.

$$P_g = \frac{W}{A}$$

Wheel torque:
a larger wheel can improve obstacle capability, but for the sametractive force it also increases required axle torque.

$$\tau_{wheel} = F_t r$$

Outcome:
Delivered a 160+ page Preliminary Design Review covering mission requirements, ConOps, subsystem architecture, interfaces, science instrumentation, mass/power budgets, FMEA/risk, schedule, cost, and verification planning

Index

ARCADIA LUNAR ROVER EXAMPLE WORKS

NASA L'SPACE MISSION CONCEPT ACADEMY

Science Traceability Matrix

Science Goals	Science Objectives	Science Measurement Requirements		Instrument Performance Requirements		Predicted Instrument Performance	Instrument	Mission Requirements
		Physical Parameters	Observables					
The moon is one of the few places in the Solar System capable of preserving centuries of geological transformations including pits, caves, and lava tubes. Do habitable conditions exist within these features and do they contain resources for in-situ use for future explorers? *(Lunar Exploration Analysis Group (LEAG) Habitation: HAB-SAT Report, Priority Objectives)*	Characterize the thermal environment of Mare Tranquillitatis Pit (MTP)	Identify temperature fluctuations throughout mission lifespan on the lunar surface and on MTP floor	Observe stable floor and wall temperatures ~ 290 Kelvin	Temperature Sensitivity:	150-300 K	6–35 μm & 14.5-15.5 μm	Thermal IR Sensor aboard Mars Environmental Dynamics Analyzer (MEDA)	Rover will land within a 5km radius of Mare Tranqulitatis Pit
				Spatial Resolution:	< 10 sq. m	~ 10 sq. m		
				Measurement stability:	± 1 K	< 1 μm		
	Assess the structural integrity of the walls and overhangs of MTP	Conduct nondestructive testing to identify signs of structural instability through imaging	Image cracks, fractures, deformations, and debris piles near MTP entryway, walls, and floor.	Spectral Range:	Visible	400-700 nm	Close - Up Imager (CLUPI)	
				Resolution	< 2.5 mm/pixel	79 μm/pixel at 10 cm		Vehicle shall maintain at most a 100 cm distance from imaging target
				Field of View:	< 100°	14 +/- 2°		
	Investigate and characterize potential cave networks and terrain in and around MTP, respectively	Identify hidden cave network(s) and other features using ground penetrating radar and map their terrain	Collect radar signatures of MTP subsurface	Spatial Resolution:	< 1 sq. m	< 0.5 sq m	Under-Rover Ultra-Wide Band Non-Contact Ground Penetrating Radar Antenna	
				Depth:	≤ 50 m	> 5 m		
				Central Frequency:	~ 100 MHz	~1000 MHz		
				Frequency Range:	> 100 Mhz	120 MHz to 2000 MHz		
	Quantify the radiation environment in and around MTP to understand its effect on humans and instrumentation	Identify radiation dose rate and total ionizing dose in areas surrounding MTP and within MTP.	Collect telemetry on total ionizing dose (TID) rate, displacement damage dose (DDD) rate, and cosmic ray flux	Energy Coverage:	100 keV - 10 MeV	5 MeV - 100 MeV	Radiation Assessment Detector (RAD)	Vehicle shall collect temperature samples daily for thermal mapping
				Exposure Time:	< 1 hr	< 15 min		
				Sensitivity:	< 10^4 erg/s	~ 100 keV		
	Map the distribution of volatiles and hydrogen concentrations found in and around MTP	Identify abundances of H, CH4, SiO2, N2, NH3	Collect relative abundance of H, CH4, SiO2, N2, NH3 in MTP regolith	Infrared Wavelength Range:	780 nm - 4 μm	1–25 keV	Alpha Particle X-Ray Spectrometer (APXS)	Vehicle shall identify terrain characteristics and ease of access into MTP
				Optical/UV Range:	100 - 700 nm			
				Resolution:	~1-2 nm	~ 200 eV		

FMEA Chart

Function	Failure Mode	Effects	Sev	Cause	Occ	Prevention	Det	RPN	Actions
Power	Solar power reduction in efficiency	Greatly reduces power available to the spacecraft's systems.	7	Dust accumulation over the solar panels from lunar regolith.	7	Special dust repellent coating, or using electrostatic devices to prevent dust particles from building up.	2	98	Dust needs to be cleaned off (repelled) using electrostatic devices.
	Lithium Ion battery capacity reduction	Increases risk of current power not being sufficient when solar panels cannot be used. Lifespan of battery can also be reduced due to cold temperatures. Lithium plating can also occur which is dangerous to the components.	7	Cold temperatures on lunar surface.	6	Proper temperature control such as heating during cold conditions as well as proper insulation, proper UV protection.	3	126	Battery needs to be warmed in case of cold temperature. If lithium plating occurs, the only solution is to rely on a secondary battery.
	PMAD fault	Loss of power regulation causing damage to components and loss of data. Also affects efficiency	8	Temperature changes	2	Proper temperature control such as heating during cold conditions as well as proper insulation, proper UV protection.	5	80	Identify the cause of the issue, but needs a redundant system to rely upon in case of failure.
CDH	Computer Failure	Vehicle unable to operate.	9	Dust, radiation, temperature changes, power fluctuations, memory failure.	4	Proper temperature control such as heating during cold conditions as well as proper insulation, proper UV protection. Dust protection is also needed.	2	72	A backup computer is needed in case the primary one fails. Once the system is stable on the secondary computer, the issue can be troubleshooted and solved.
	Data Failure	Data potentially lost, extremely negative to the science objectives of the mission with the loss of collected data.	9	Radiation, temperature changes, debris, collisions, software glitch.	4	Data needs to be regularly backed up and system must be protected from environmental conditions.	2	72	Troubleshoot system and see what the error is and if the system can be repaired.
	Telecommunications failure	Unable to transmit data back to Earth which is detrimental to the mission's science objectives. Furthermore, can leave the whereabouts of the vehicle unknown, and unable to troubleshoot other problems.	8	Radiation, interference from other celestial bodies, data failure	3	Needs protection from radiation. Frequencies must be chosen specifically to reduce the impact of other bodies in space from affecting transmissions.	4	96	Issue needs to be diagnosed and repaired.
Thermal	Loss of thermal protection	Damage of sensitive components and even destruction of spacecraft.	10	Degradation of passive insulation over time due to lunar temperature cycling.	2	Needs a protective layer to prevent damage from physical/environmental damage. Cannot be subject to excessive compression.	4	80	Needs redundant layer of thermal insulation in case of breach.
	Radiator fault	Increased energy consumption, loss of efficiency, reduction in heat	10	Dust accumulation over the radiator panels. The dust acts as a layer of insulation trapping	2	Needs protection from dust.	3	60	Dust needs to be cleaned from panels. Can also have a redundant

MARS CTX CRATER DETECTION

GEOSPATIAL ML ▪ DATA PIPELINE ▪ CONTROLLED DETECTOR VALIDATION

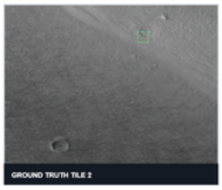

Role: Research / ML Engineering Contributor
Timeline: May 2026 – Aug. 2026
Tools: Mars CTX • Geospatial Data • Faster R-CNN • Evaluation • GIS Output

Overview: I worked on a Mars crater-detection pipeline using CTX orbital imagery and crater-catalog data. What began as reconstruction of an existing notebook evolved into a more controlled benchmark: source imagery was tiled, crater catalogs were projected into image coordinates, geographic leakage was explicitly blocked, and the detector was evaluated with one-to-one matching instead of raw candidate counts. The emphasis became less about "training a model" and more about making sure the result was actually traceable and defensible.

Data & Geospatial Pipeline

The CTX region was cut into overlapping **1024 × 1024 tiles** with **128 px** overlap and an **896 px** step. Catalog crater centers and diameter measurements were projected through raster metadata into tile-local bounding boxes, with crater size derived from trusted meters-per-pixel information rather than screenshot scale.

Dataset Snapshot

- **2,809** total generated tiles
- **1,683** catalog craters in benchmark
- **1,908 train / 318 validation / 371 test**
- **1,561 negative tiles**
- **75–930 m** crater-diameter range
- 0 cross-split overlapping tile pairs
- 0 crater IDs shared across splits

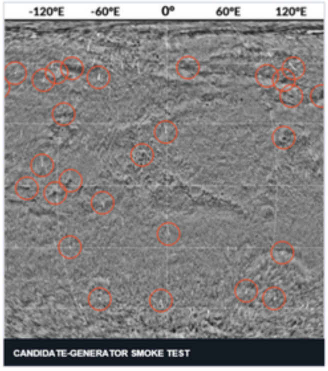

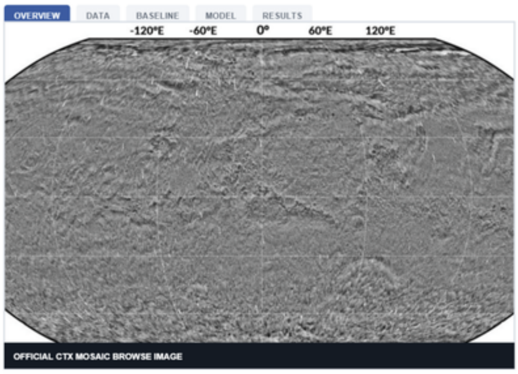

Key Contributions

- Built catalog audits, affine label generation, geographic splits, map-level deduplication, and traceable run manifests.
- Implemented a deterministic **geographic-block split** to prevent nearby overlapping tiles or repeated crater IDs from leaking across train/validation/test.
- Designed and ran a controlled Faster R-CNN comparison with validation-only model/threshold selection and a single locked test evaluation.
- Published failure cases alongside headline metrics rather than treating candidate volume as detector accuracy.

Model Validation

A deliberately weak YOLO11n run was first used as a pipeline smoke test. It successfully exercised sourcing, tiling, training, inference, evaluation, and GIS deduplication, but produced **0 AP50** and more than **53,000 post-deduplication false positives**. The failure was useful because it showed the data path worked while the detector did not.

That motivated a controlled Faster R-CNN experiment.

Locked Faster R-CNN Decision

- Candidate A, epoch 11
- Best validation AP50: 0.292
- Locked score threshold: 0.8
- IoU threshold: 0.5
- 14 automated pipeline tests passed
- Test set evaluated once after selection was frozen

Final Locked Test

At the frozen test settings:

- 6 TP
- 10 FP
- 8 FN
- Precision: 0.375
- Recall: 0.429
- F1: 0.400
- AP50: 0.263
- Median matched IoU: 0.663

Takeaway

The model found real craters, but the limiting factor was still **recall and false-positive burden**, especially in low-contrast or terrain-confused regions. The next step was expanding geographically separated training data, preserving hard negatives, and repeating the experiment with a new untouched test set.

Index

CASSINI JUPITER LIGHTNING

PLANETARY IMAGE ANALYSIS • CANDIDATE REDUCTION • VALIDATION • GEOMETRY

Role: Research / Data Analysis Contributor
Timeline: June 2026 – Aug. 2026
Tools: Cassini ISS • Jupiter Nightside • Python • SNR Detection • Candidate Review

Overview: Find bright pixels. Jupiter's nightside images contain real lightning candidates alongside **hot pixels, cosmic rays, streaks, compression artifacts, limb leakage, and cloud structure**. I built a reproducible search-and-review workflow that reduced a large number of bright regions into traceable candidates while keeping a strict distinction between a candidate and confirmed lightning.

What I Built

The detector used **Python / NumPy / Pillow**, local-background SNR, connected bright regions, shape measurements, temporal linking, and artifact flags. The workflow produced repeatable date-window runs with per-image provenance, candidate CSVs, contact sheets, review plans, and geometry outputs.

Search Reduction

The detector processed: 196,233→12,611→106

196,233 raw bright regions
↓ artifact rejection + prioritization
12,611 review candidates
↓ controlled first-pass selection
106 review rows

A total of 183,622 regions were rejected or artifact-flagged before the review packet was built.

Locked Parameters

- Background blur radius: **18 px**
- Detection SNR threshold: **7**
- Review SNR threshold: **8**
- Minimum review area: **3 px**
- Maximum component area: **500 px**
- Track displacement: **≤85 px**
- Maximum temporal gap: **12 min**

Detector Logic

The pipeline measured, for each connected region:

- frame ID
- x/y position
- peak / mean SNR
- integrated SNR
- area
- sharpness
- elongation
- artifact flags
- timestamp
- provenance

Temporal linking allowed persistence and motion to be reviewed separately from one-frame brightness.

Published-Match Validation

Before treating unmatched detections seriously, I validated the detector against published Cassini lightning coordinates.
Result: 6 / 6 published marks recovered

The detector matched all six within the 8 px validation radius, with offsets between: $1.54 \text{ px} \le \Delta r \le 5.56 \text{ px}$

Geometry / Localization

ISIS geometry processing produced candidate coordinates for:
42 / 106 requested rows
The remaining 64 rays produced no surface intersection, with zero projection failures. The coordinate outputs still required validation before being treated as scientific results.

Outcome

- processed 221 Cassini images
- covered 9 date windows
- reduced 196,233 bright regions into traceable review products
- recovered 6/6 published lightning marks
- generated geometry for 42/106 prioritized rows
- produced review batches, temporal summaries, provenance records, and training gates

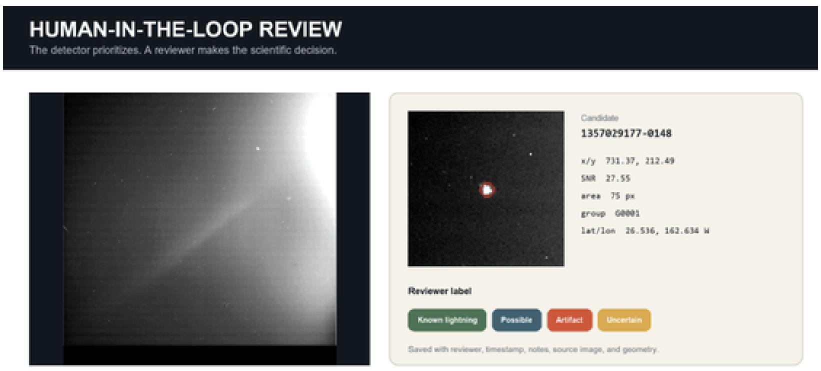

Takeaway

The project was not ready for a learned model yet because no reviewed positive/negative training set existed. The next step was to complete blinded review, validate the 42 geometry intersections, diagnose the 64 misses, review temporal tracks against spacecraft motion, freeze a labeled benchmark, and only then compare a learned detector.

SAFEBYTE

WEARABLE FOOD-SAFETY SCANNER FOR DIETARY RESTRICTIONS

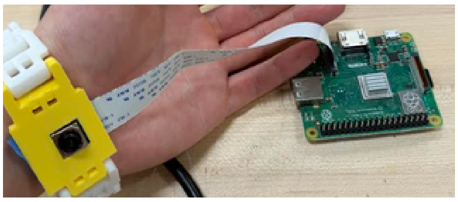

Role: Software + Systems Integration
Timeline: 24-Hour Hackathon
Result: 1st Place
Tools: Raspberry Pi • Camera • OCR • Barcode Recognition • Nutrition / Allergen Data • Browser UI

Goal: Build a wearable food-safety assistant that could quickly identify food products or menu items and flag dietary risks using a single Raspberry Pi camera pipeline.

Overview: I integrated the Raspberry Pi camera, OCR, barcode recognition, nutrition/allergen data, and browser UI into an end-to-end wearable scanner. The system supported three identification paths: **barcode scanning for packaged foods, OCR for menus and labels, and visual recognition for cases where text or barcodes were unavailable.** The output was then compared against the user's dietary profile to flag potential allergens or restricted ingredients. [Demo Vid]

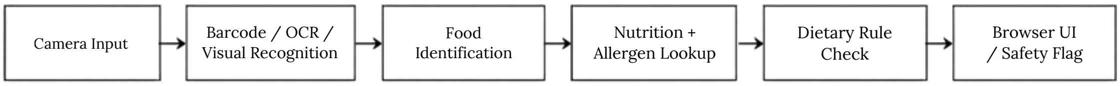

Key Contribution:

- Integrated the Raspberry Pi camera, barcode recognition, OCR, nutrition/allergen data, and browser interface into one working hardware/software pipeline.
- Built the camera-to-interface workflow used for image capture, recognition, safety checks, and real-time feedback.
- Combined multiple identification paths so the system could fall back between barcode, text, and visual input depending on what information was available.
- Helped move the project from separate components into a functional end-to-end prototype within 24 hours.

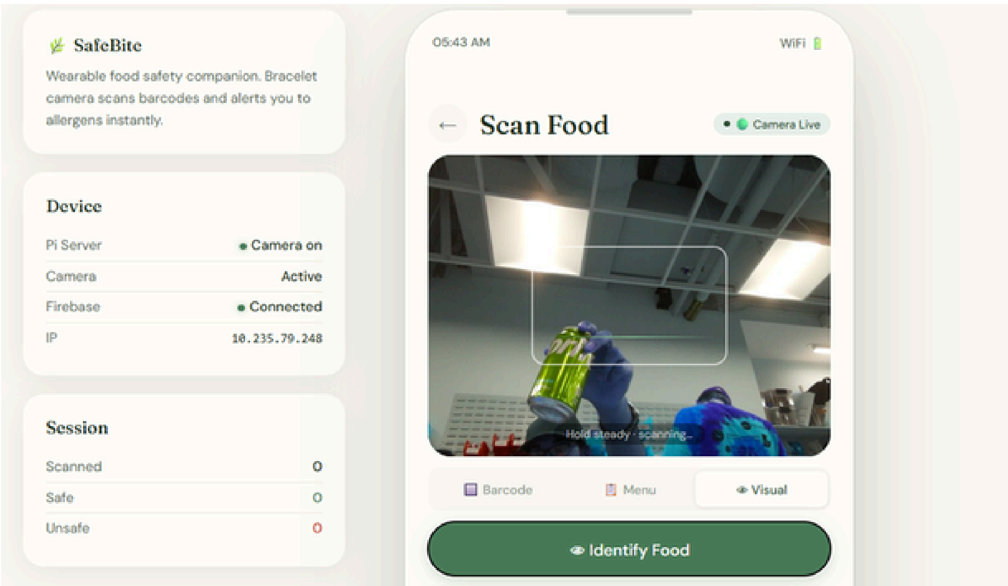

Engineering Tradeoff:

The main constraint was time. The raw camera stream was unstable and network-heavy, so instead of building a more complex distributed pipeline, I kept preprocessing and recognition close to the Raspberry Pi and only sent the information needed by the browser UI. That reduced unnecessary data transfer and made the demo more reliable under hackathon conditions.

What Worked

The final prototype could capture camera input, identify food through multiple recognition paths, query nutrition/allergen data, and return a user-facing safety result through the browser interface.

VEX ROBOTICS COMPETITION

RIVETRAIN • MECHANISM DESIGN • AUTONOMOUS

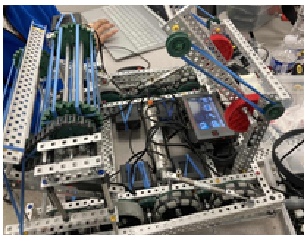

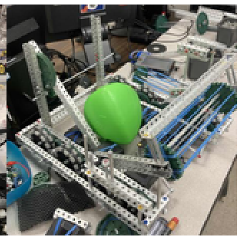

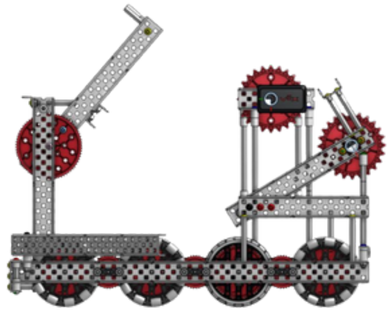

Role: Co-Captain 1356
Timeline: 2021–2024
Focus: Autodesk Inventor • VEX V5 • Pneumatics • Encoder/PID Autonomous

Overview: I co-led the mechanical and autonomous development of our competition robots across drivetrain, intake, shooter, wings, elevation, and match-autonomous systems. The design process was iterative rather than CAD-first: we moved from **problem definition → concepts → weighted decision matrix → CAD/build → subsystem testing → competition feedback**, then rebuilt around what actually failed on the field.

Drivetrain Design

The 2023–24 robot converged on a **360 RPM tank drive** using **600 RPM cartridges**, **3:4 external gearing**, and **3.25-in wheels,** later moving to an **all-omni configuration** when driver testing showed the prior traction setup constrained turning.

$$T_{out} = niT_m$$

$$w_{out} = \frac{w_m}{i}$$

$$v = w_{out} r$$

The trade was straightforward: gearing changes wheel torque and speed together, while wheel radius changes ground force and linear speed. The final configuration prioritized the maneuverability that better matched competition driving rather than maximizing traction in isolation.

Mechanism Stack: 11 W / 600 RPM intake • 42 RPM slapper • deployable wings • 70 PSI pneumatic elevation

Failure → Redesign: One recurring failure mode was drive stalling and gear/axle problems under high load. The fix was not “use a stronger motor.” We changed bearing support, unsupported shaft spans, gearing, and structural load paths to reduce deflection and improve reliability.

Pneumatic Elevation: The elevation system operated at 70 PSI. Ideal cylinder force begins with: F = PA but the useful mechanism force depended just as much on linkage geometry, friction, moment arm, and center-of-gravity placement. The initial design missed the speed target. After shifting mass forward and iterating the mechanism, five measured elevation trials were: 0.8, 0.4, 0.3, 0.3, 0.5 s
That reduced the final mechanism cycle to roughly 0.3–0.8 s.

Recorded Team Results:
174 driver skills • 154 autonomous skills • 83% qualification win rate • Design Award 2× VEX Worlds

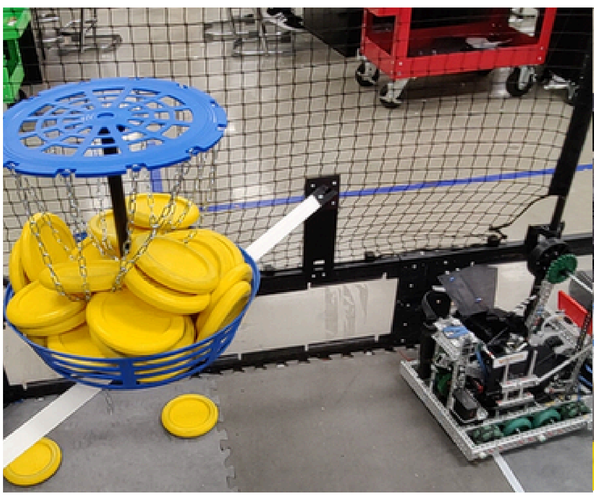

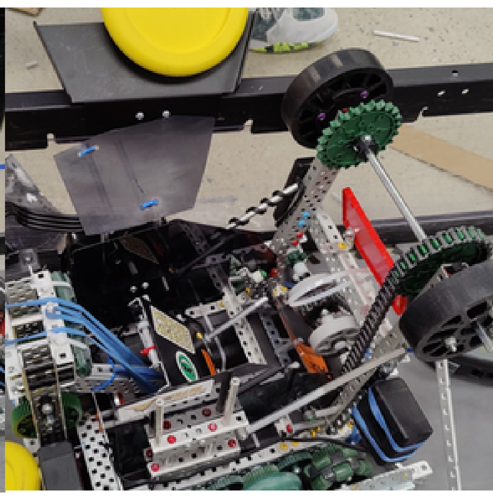

[Index]

GRAPHENE TRANSFER + CHARACTERIZATION

SAMPLE PREPARATION, RAMAN IDENTIFICATION, AND AFM ANALYSIS

Role: Experimental Physics / Materials Processing Contributor
Timeline: Spring 2026
Focus: Process control, sample preparation, material identification, surface/thickness validation
Tools: Graphite, CVD Graphene, Tape Exfoliation, Wet Transfer, Optical Microscopy, Raman Spectroscopy, AFM

Overview: This project focused on preparing and characterizing graphene samples through a materials-processing workflow. Rather than treating graphene as only a physics concept, the goal was to understand the engineering difficulty of producing usable samples: transfer quality, contamination, tearing, folding, flake variability, and thickness nonuniformity. The workflow connected sample preparation with measurement validation. Graphene was prepared through exfoliation and/or wet-transfer pathways, inspected under optical microscopy, and then characterized using Raman spectroscopy and AFM to evaluate material identity, structure, and local thickness.

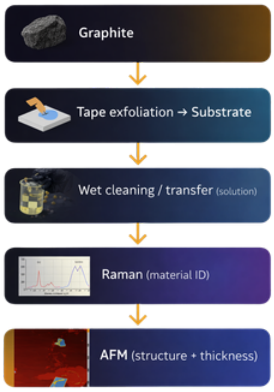

Key Contribution:

- Prepared graphene samples through exfoliation and transfer-based workflows
- Used optical microscopy to inspect flake size, shape, contrast, and sample quality
- Helped evaluate whether transfer steps preserved usable graphene regions
- Interpreted Raman spectroscopy as a material-identification and quality-check tool
- Used AFM data to evaluate surface structure, thickness variation, and local features
- Compared how different preparation paths affected consistency and measurement

Challenges:

- Graphene transfer is highly sensitive to handling and process conditions. Small changes in transfer timing, substrate contact, cleaning, adhesion, or drying can create defects that are not obvious until characterization.

Outcome:

- Built hands-on experience with nanoscale materials processing and characterization workflows. The project developed practical familiarity with graphene transfer, optical inspection, Raman-based material identification, AFM validation, and the process sensitivity involved in advanced materials research.

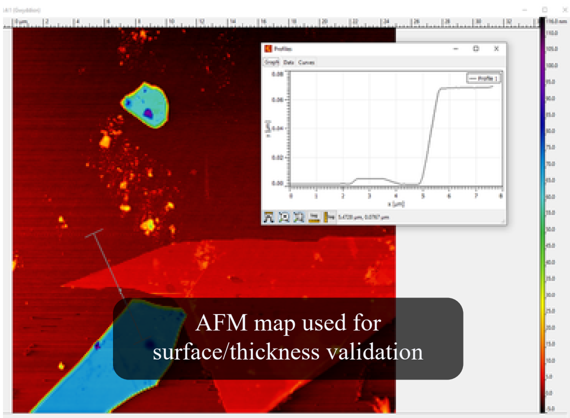

AFM map used for surface/thickness validation

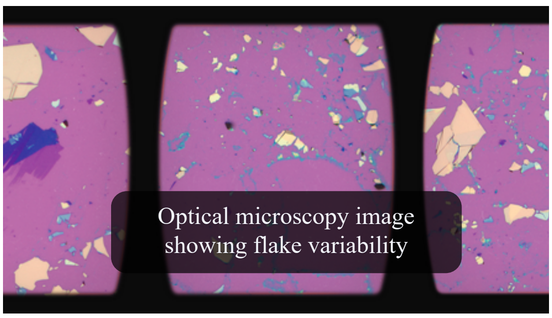

Optical microscopy image showing flake variability

Lab preparation and sample-handling setup

Index

HALL EFFECT MEASUREMENT IN COPPER

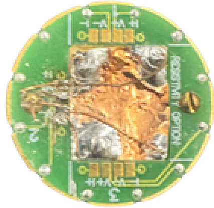

MAGNETOTRANSPORT TEST SETUP, CARRIER DENSITY EXTRACTION, AND SENSOR VALIDATION

Role: Experimental Physics / Measurement Systems Contributor
Timeline: Spring 2026
Focus: Sensor measurement, data acquisition, material-property extraction, uncertainty-aware analysis
Tools: Copper Sample, Magnetic Field Sweep, Constant Current Source, Voltage/Resistance Measurement, Linear Fit, Temperature Comparison

Overview: This project measured the Hall effect in copper as a precision measurement and data-analysis workflow. The experiment applied a controlled current through a copper sample while sweeping the magnetic field from negative to positive values, then measured the transverse response to extract material properties. The focus was measurement reliability. Hall signals can be small and sensitive to sample geometry, contact alignment, field direction, temperature, and noise. The project required treating the lab setup as a calibrated measurement system, not just a formula-based physics exercise.

Key Contribution:

- Measured transverse resistance during controlled magnetic-field sweeps
- Used copper sample dimensions to support Hall coefficient calculation
- Extracted Hall slope through linear fitting of resistance vs. magnetic field
- Used slope direction to identify dominant carrier behavior
- Compared measurements across temperature conditions to evaluate consistency
- Connected the Hall measurement method to magnetic sensors, current detection, and material characterization

Challenges:

- Hall measurements are sensitive to experimental imperfections. Small voltage signals can be affected by contact placement, sample alignment, temperature drift, noise, and geometry assumptions.

Outcome:

- Developed experience with magnetotransport testing, sensor-style measurement, linear regression, carrier-density extraction, and uncertainty-aware data analysis. This project strengthened my ability to move from raw physical measurement to validated engineering quantities.

Copper sample and measurement geometry

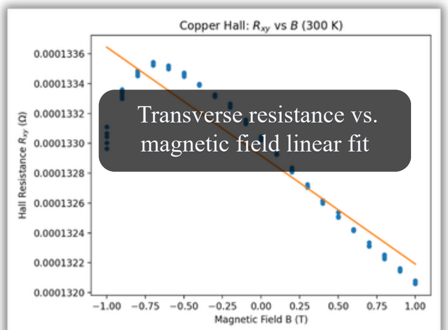

Transverse resistance vs. magnetic field linear fit

Lab preparation and sample-handling setup

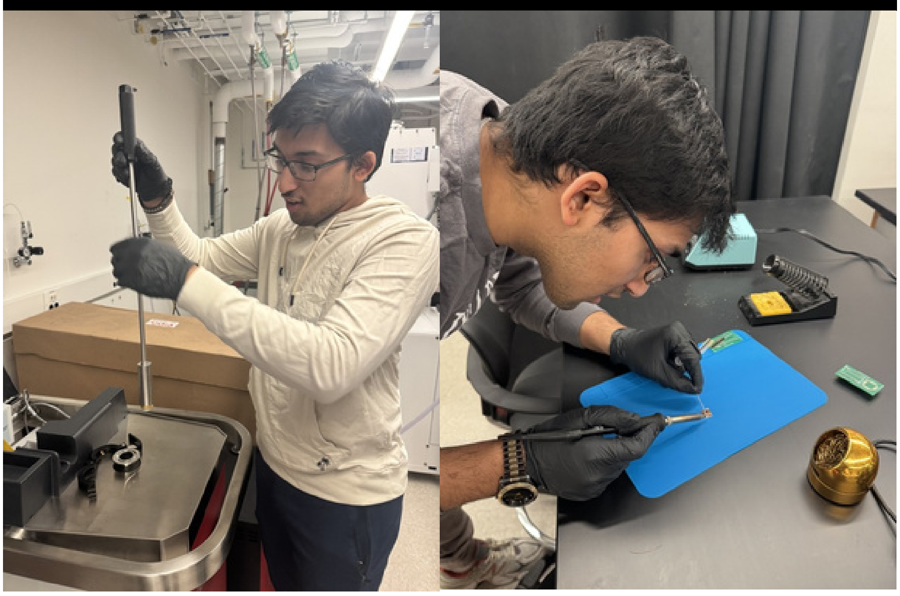

Index

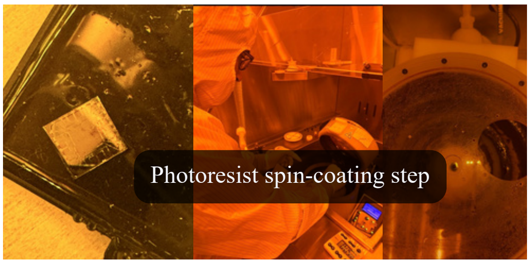
Photoresist spin-coating step

MICROFABRICATION PROCESS WORKFLOW

PHOTOLITHOGRAPHY, UMLA EXPOSURE, DEVELOPMENT, AND WIRE BONDING

Role: Experimental Physics / Measurement Systems Contributor
Timeline: Spring 2026
Tools: Spin Coater, Photoresist, Hot Plate, uMLA, Developer Bath, Optical Microscope, Wire Bonder
Focus: Semiconductor process sequencing, chip-scale handling, patterning, inspection, electrical interfacing

Overview: This project introduced a semiconductor-style fabrication workflow where process sequence, timing, cleanliness, exposure, and inspection determine whether a chip-scale pattern or device interface succeeds. The workflow moved from surface preparation to photoresist coating, hot plate baking, uMLA exposure, development, microscope inspection, and wire bonding. The engineering focus was process discipline: each step had to be performed carefully because small deviations in resist thickness, bake time, exposure, development, or handling could affect final pattern quality and electrical connectivity.

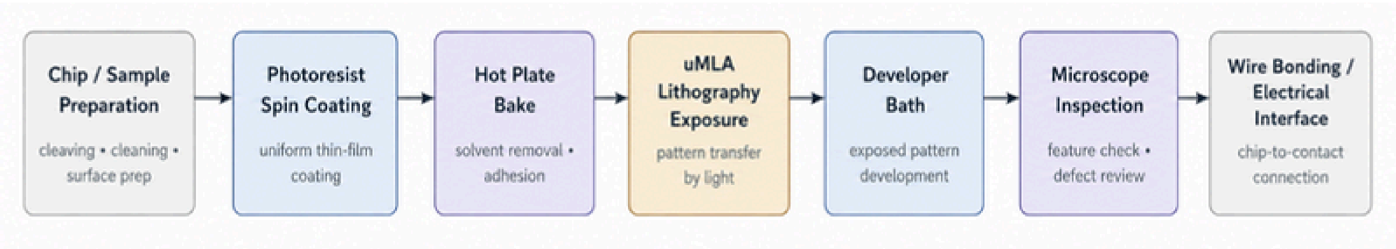

Key Contribution:

- Prepared chip/sample surfaces for lithography processing
- Applied photoresist using spin-coating procedures
- Followed thermal bake steps to stabilize resist layers
- Used uMLA exposure to pattern selected regions
- Developed exposed regions and inspected results under microscopy
- Practiced wire bonding between chip/device interfaces

Challenges:

- Microfabrication is failure-sensitive because each step affects every step after it. A small issue in resist coating, exposure, development, or handling can propagate into poor pattern definition or failed electrical connection.

Outcome:

- Built exposure to semiconductor fabrication methods, lithography process flow, chip-scale inspection, and electrical-interface preparation. The project strengthened my understanding of how materials processing, device patterning, and electrical connection come together in experimental hardware.

Wire bonding process

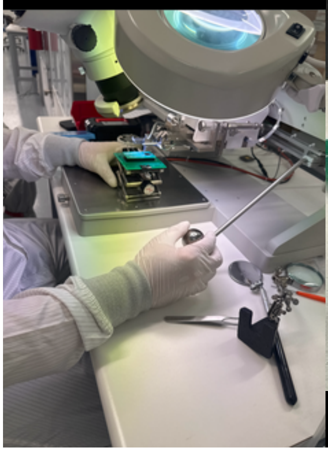

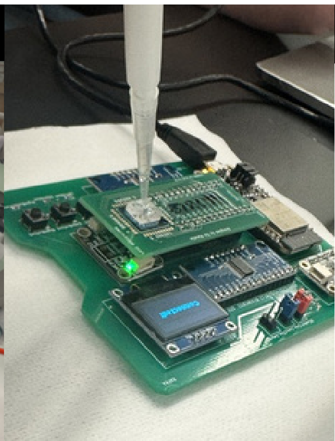
Final chip/sample interface

The graph shows resistance vs gate voltage for all 20 graphene devices.

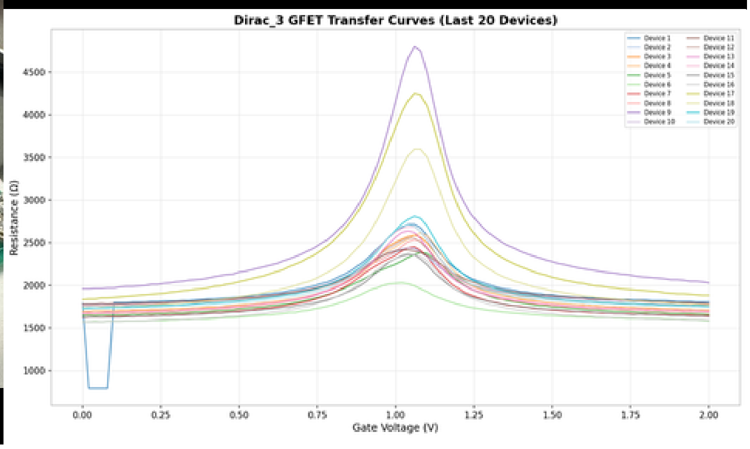

Index

FABRICATION + RESIDENTIAL INFRASTRUCTURE

HANDS-ON CONSTRUCTION, UTILITIES, AND BUILD SEQUENCING

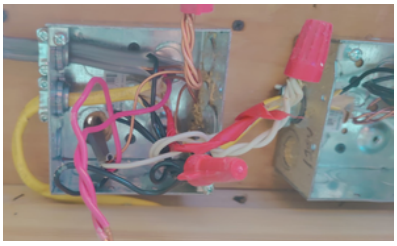

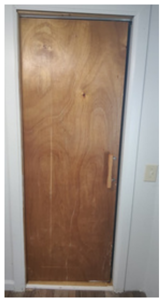

Role: Builder / Assistant Fabricator
Timeline: Early High School (During Covid)
Focus: Construction • Utility Routing Preparation • Tile Work • Sliding Doors • Gazebo Build

Overview: Assisted with hands-on residential fabrication and infrastructure projects, including shed build-outs, utility-routing preparation, tile installation, sliding-door assembly, and gazebo construction. Jumping into this type of work early on, especially during the pandemic to help out and earn income meant learning how to adapt quickly to demanding, physical field environments where you don't have heavy machinery to do the lifting for you. Working closely on a tiny, three-person crew forced me to grasp project sequencing and resourcefulness right on the spot. These projects gave me early experience with tools, measurement, physical constraints, and the reality that every build has to survive weather, alignment, material limits, and human error.

Key Contribution:

- Hand-dug trenches to lay down water supply lines and septic pipes across the property
- Set up full bathroom fixtures, including toilets, showers, and the drainage pipe system under supervision
- Helped pull and route heavy-gauge electrical lines from the main house panel out to multiple detached sheds.
- Mixed concrete for the foundation, handled the structural wood framing, and completed the full timber build for a large 28' x 35' outdoor gazebo.
- Mounted and calibrated track-guided sliding door mechanisms, balancing tolerances so they would slide smoothly without catching.

Challenges:

- Operating as the youngest member of a three-person crew alongside my dad and a 60-year-old uncle during the pandemic forced me to adapt instantly to heavy, un-mechanized workloads. Working with them meant keeping up with decades of grit, hand-digging deep septic and water line trenches, and mixing concrete by hand without complaining. The real test came after a heavy rainfall revealed that compressed edge spacing on our custom gazebo was causing water to pool and back up instead of shedding safely off the roofline. To fix this, I had to analyze the structural runoff pattern, calculate corrective panel angles, and retrofit a custom gutter system, which taught me how to turn real-world field failures into functional design adjustments.

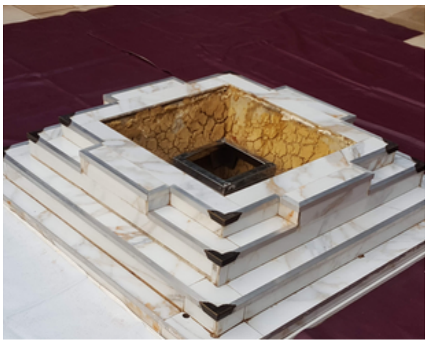

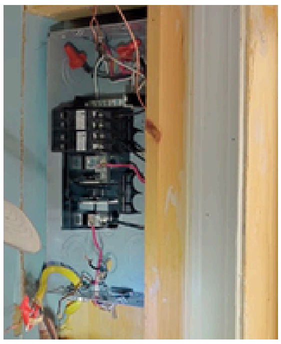

Index

MANHUNT WEB APP

[*REAL-TIME LOCATION-BASED OUTDOOR GAME PLATFORM*]

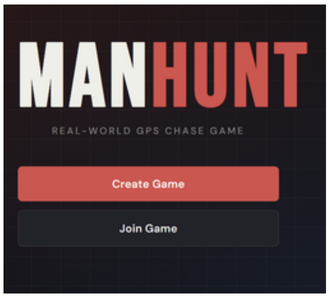

Role: Full-Stack Developer
Timeline: Spring 2026
Tools: Web App • Maps • GPS Location • Team Codes • Live Game State • Boundary Drawing

Overview: Built Manhunt, a web-based multiplayer location game created after a ski trip was unexpectedly canceled and my friends needed a new way to turn a snowy 40-acre property into a playable outdoor arena. The app lets a host create a game, draw a playable boundary on a map, generate a join code, assign players as hunters or runners, and control location-reveal intervals during gameplay.

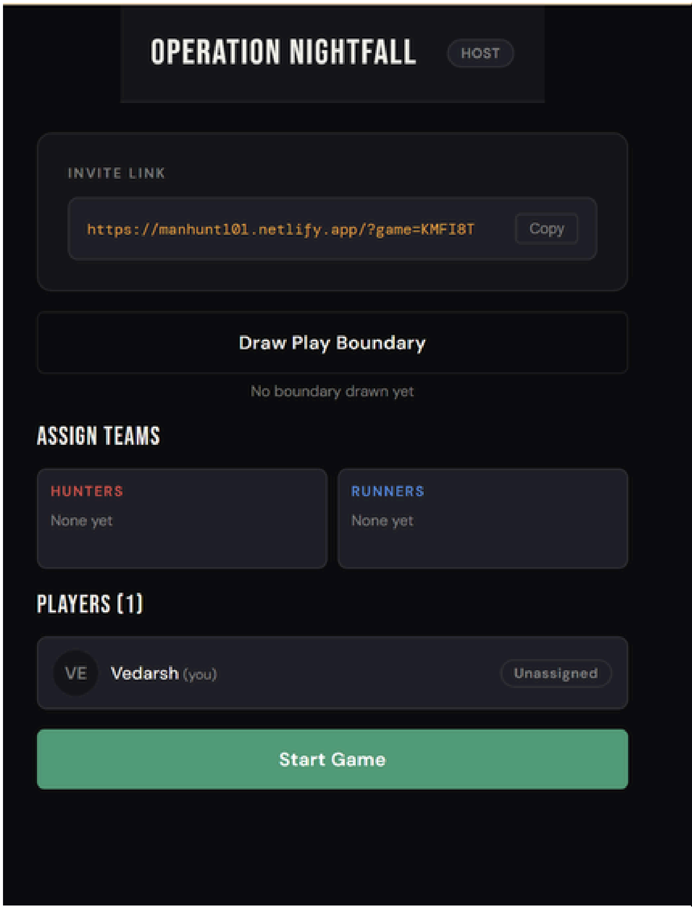

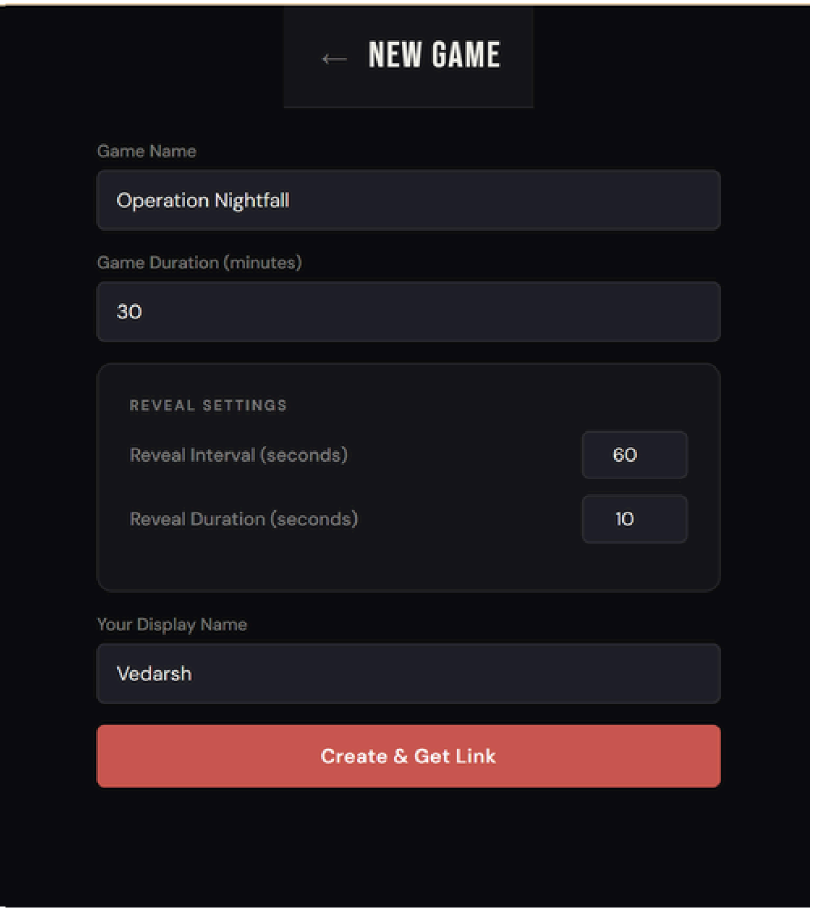

Key Features:

- Game creation with shareable join links and team codes
- Map-based boundary drawing for custom outdoor play
- Player role assignment for hunters and runners
- Configurable location-reveal intervals
- Live game state for group outdoor gameplay
- Mobile-friendly interface for field use

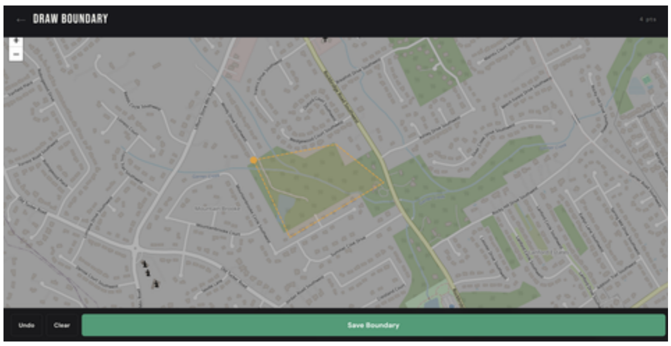

Outcome:

- Built a functional location-based game platform and learned how to turn a social idea into a working interactive system with maps, GPS logic, user flows, and real-time play constraints.

Index