Multi-Modal Sensor Platform for Low-Cost Detection of Landmines

Final Report

Team B:
Aaron Chong
Basel Alghanem
David Robinson
Edward Terry
Georgia Crowther

Advisor: Dimitrios Apostolopoulos

Date: 5/9/2018

Abstract

Landmines are a deadly remnant of war that continues to threaten the safety and economic stability of communities in the developing world. Humanitarian demining involves searching an area with a metal detector, then using a probe to confirm whether a detection is, in fact, a landmine. Robotic solutions to this problem exist, but their cost, complexity and bulk are a barrier to widespread adoption in the field. Additionally, most robotic solutions do not address the problem of high false-to-true-positive ratios.

This report outlines the progress of Clearfield Robotics (Figure 1) towards building a multi-modal sensor unit to be mounted on the Scorpion cart, developed by the National Robotics Engineering Centre (NREC), which has been previously deployed in Cambodia for unexploded ordnance (UXO) detection. In the first phase of searching, a metal detector identifies any buried potential threats. The unit sweeps back and forth until it detects a buried metallic object, which we have demonstrated to be able to locate a landmine target to within 10cm of its actual location. In the event of a detection, a metallic shaft probes the soil in a sequential pattern, classifying the shape of the potential hazard to confirm whether it is a landmine or just harmless shrapnel. We have demonstrated how our shape fitting algorithm calculates the geometry of a cylindrical buried object based on an array of subsurface contact points for targets buried outdoors. Finally, our localization system is used to record the total land coverage and where each landmine was found, to within 1m of its actual location.

Figure 1: Clearfield Robotics

Acronyms

BOM	Bill of Materials
CAD	Computer-Aided Design
CDR	Critical Design Review
CMAC	Cambodia Mine Action Centre
CMU	Carnegie Mellon University
DC	Direct Current
EKF	Extended Kalman Filter
FEA	Finite Element Analysis
FVE	Fall Validation Experiment
GPS	Global Positioning System
IMU	Inertial Measurement Unit
NGO	Non Governmental Organization
NREC	National Robotics Engineering Centre
PCB	Printed Circuit Board
PDB	Power Distribution Board
PR	Progress Review
RANSAC	Random Sampling Consensus
ROS	Robot Operating System
RTK	Real Time Kinematic
SVE	Spring Validation Experiment
TLA	Three-Letter Acronym
UI	User Interface
USB	Universal Serial Bus
UTM	Universal Transverse Mercator
UXO	Unexploded Ordnance

Table of Contents

1 Project Description

Landmines are a tragic remnant of war and are a prevalent issue in countries that have been affected by military conflicts. Many years after these conflicts have ended, landmines are still claiming lives. In 2015 alone, 6,461 landmine deaths were recorded [1]. Additionally, the presence of landmines prevents the use of agricultural land, which inhibits the economic growth of these countries [2]. Non-Governmental Organizations (NGOs) and local organizations are making an effort to remove landmines, but, due to the laborious nature of the process, progress has been slow. It is estimated that there are 110 million landmines (Figure 2) yet to be removed, with more placed each year than are removed [3].

The current process for demining varies but can include the use of several implements, including metal detectors, probes, bomb-sniffing dogs and rats, and mechanical flails. Given the limitations of particular sensing modalities, a range of complementary techniques are typically used to assure complete clearance. For example, metal detection alone yields a false to true positive ratio of up to 100:1, and every such signal must be investigated manually with a probe [4]. Autonomous robots using ground-penetrating radar have been proposed [5], but are infeasible for wide-spread use in the field because they are extremely costly.

In order for a landmine detection product to improve on current detection methods and achieve wide usage, it must consistently detect landmines, provide simple actionable feedback for landmine removal, be easy to use, and be inexpensive. The goal of Clearfield Robotics is to create a product that can increase the rate at which landmines are removed, and increase worker safety, while improving on current methods. Therefore, we propose an autonomous landmine detecting system that leverages human labor for mobility and sensor fusion of multiple cost-effective sensing modalities to consistently detect landmines and reject shrapnel and other false positives. It must interact well with its operator by providing clear instructions and warnings, and will provide greater worker safety by distancing the operator from a concentrated landmine blast in the case that a landmine is unintentionally detonated during detection. The result should be a surveyed field, with corresponding documentation, and physically marked potentially hazardous locations for later inspection.

Figure 2: Typical Landmines (Left: Anti-Tank Mine at Cambodia Landmine Museum. Right: Anti-Personnel Landmine [6]

2 Use Case

Andrew is a local resident of Krong Battambang province in Cambodia. He is employed as a land surveyor by the local branch of Cambodia Mine Action Centre (CMAC), a Government Organization responsible for the survey, clearance and release of land which is suspected to be contaminated with landmines from the Khmer Rouge conflict. Today he and his team will be surveying a patch of fertile land, as in Figure 3a, a task which involves scanning the ground for landmines, unearthing them and safely disposing of them, with the intent of allowing farmers to confidently use the land once more. CMAC has recently acquired a Clearfield Robotics Surveying Cart. Andrew and his colleague Brian lift the cart from off the back of their work truck, as in Figure 3b, and wheel it over to the edge of the field.

a) Cambodian Minefield *b) Unloading Process*

Figure 3: Use Case - Setup Process

Andrew switches on the system and the handlebar-mounted user interface (UI) displays his present location overlaid on a satellite image, along with the current system status and state, as in Figure 4a. With a few finger taps, he initializes the system and beings the sweeping sequence. The UI directs him to proceed pushing the cart through the field, as in Figure 4b.

a) System User Interface *b) Operational Procedure*

Figure 4: Use Case - Initialization and Operation

As Andrew pushes the machine along his first row, he notices that on his map display the land which he has covered is marked green. Suddenly, his attention is drawn by a warning on the tablet, as shown in Figure 5a. The system brakes engage and Andrew comes to a halt. The system pinpoints the target with the metal detector, and then investigates it with the probing unit. No landmine is detected, so all warnings clear and Andrew is directed to continue searching. He reaches the end of the lane and swings the machine around to intercept the new path. Shortly after, the warnings appear again, but this time the sensor signature confirms that a landmine has been found. The machine marks the landmine location with spray paint and marks the threat with a red dot on the map. Andrew reorients the cart and continues the survey, careful to avoid the marked location of the landmine.

a) Threat Detection Alert

b) Landmine Location and Dig Sheet

Figure 5: Use Case - Landmine Detection and Mapping

Once Andrew has investigated the entirety of the field, he presses the “End” button on the UI and is presented with a dig sheet that shows him the location of all the landmines in the searched area, as seen in Figure 5b. Andrew and his clearance team use the dig sheet to unearth the landmines and transport them elsewhere for destruction. By the end of the day, Andrew has surveyed the entire field and a record of his day has been transmitted to the local CMAC office. This record consists of a satellite image overlaid with his path and the number of landmines found.

3 System Requirements

Table 1 and Table 2, respectively, list the system's mandatory and desirable functional and non-functional requirements.

Table 1: Mandatory Requirements

Func.	M.F.1	Detects 4 out of 5 landmines in a grassy field buried to a maximum depth of 10cm
	M.F.2	Correctly classifies 3 out of 5 metallic, non-landmine-shaped objects as not landmines
	M.F.3	Detonates zero landmines during probing action
	M.F.4	Physically marks landmine locations to an accuracy of 10cm
	M.F.5	Warns user of potential hazards within 0.25s
Non-Func.	M.N.1	Costs less than $5000
	M.N.2	Able to be operated continuously for 1hr
	M.N.3	Surveys land at a rate of $3m^2$/min in an area known to be free of metallic objects
	M.N.4	Determines whether a detection is a true or false positive within 5 minutes

Table 2: Desirable Requirements

Func.	D.F.1	Records landmine locations on digital map to an accuracy of 1m
	D.F.2	Records area covered by system on digital map to an accuracy of 1m
Non-Func.	D.N.1	Presents a graphical representation of land area covered and system status in real-time
	D.N.2	Component assemblies weigh less than 50kg

4 Functional Architecture

The functional architecture in Figure 6 maps the major functions and the flow of data and energy between them based on the system requirements. To start the system, the gantry performs its initialization routine, and the UI marks the operator's starting position. When the operator is ready, they will press the "Start" button on their user interface and the gantry will begin the sweeping process. As the operator pushes the cart forwards, their location is updated and digitally recorded along with the area covered by the system.

As the cart moves, the metal detector will sweep the area to determine if there is a potential threat. If a potential threat is found, the operator will be directed to stop and the brakes will engage. The metal detector will pinpoint the exact position of the threat and then prompt the operator to begin probing. The probing unit will classify the shape of the detected object. If the target is determined to be a false positive, the operator will be instructed to continue moving the cart along its path. If the potential hazard is determined to be a true positive, its location will be recorded on a digital map of the area. The operator will be alerted of the hazard and will be given feedback to continue surveying, taking care to avoid the newly detected landmine.

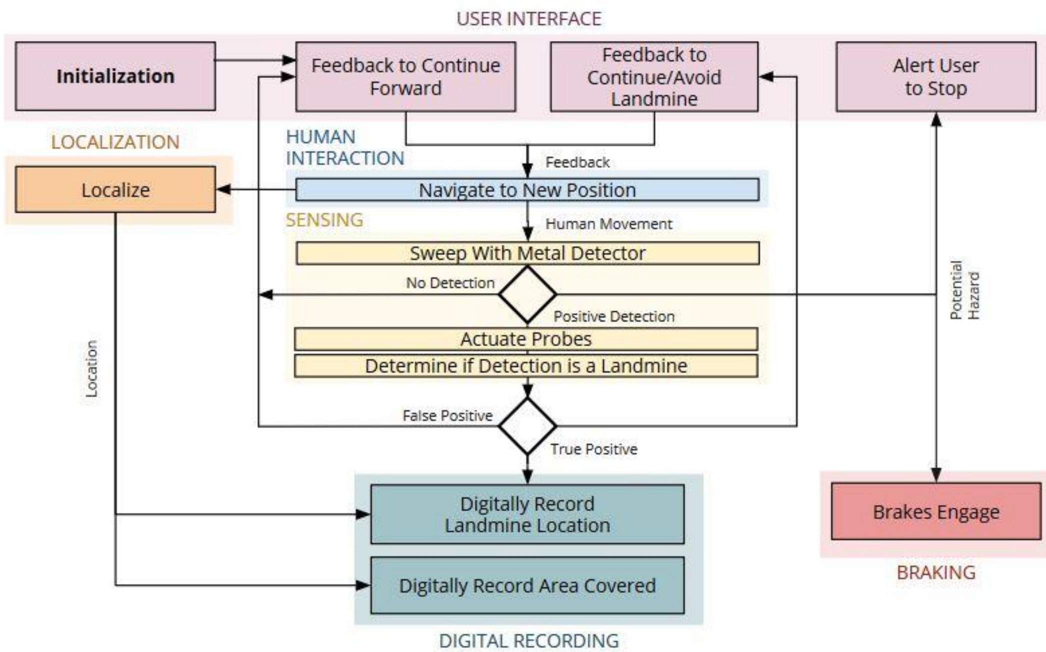

Figure 6: Functional Architecture

5 System-level Trade Studies

The trade study for each of three major components (the mechanical platform, the landmine sensing system, and the localization sensors) are displayed and discussed below. We ranked all attributes on a scale from 1 to 5, with 5 being the most favorable. Our original analyses have been replicated exactly, with the addition of a note on how we view each analysis in retrospect.

5.1 Mobility

Table 3: Mobility Trade Study

	Cost	Mass	Availability or Ease of Mfg.	Rugged-ness	Power Efficiency	Ease of sensor integration	Mobility	**Total**
	15%	15%	20%	15%	15%	10%	10%	**100%**
Cart	5	4	5	4	5	4	3	**4.40**
Drone	2	4	3	3	1	2	5	**2.80**
Autonomous Rover	1	2	3	3	2	4	3	**2.50**

The trade study for the system's mobility (Table 3) confirmed that a cart with a human operator is the most appropriate means of transporting the sensor unit around a contaminated field, because of the inherently low power consumption and low overhead nature of this approach. The mechanical platform's primary requirement is to enable the other systems, and a cart achieves that effectively without requiring a great deal of cost, manufacturing, or engineering effort. The other systems considered would not provide substantially more value over and above a human operator.

It is outside the scope of this project to factor in the potential for this system's deployment in the field. However, it is worth noting that technologies for humanitarian purposes should have the ability to be operated without foreign oversight. Manual control is chosen as a baseline architecture not just for scope management reasons but also because it would be the most viable option in the field, rather than an advanced autonomous platform, with the associated overhead of maintenance.

Looking back, this analysis was correct. Adding mobility would not only have been disastrous in expanding project workload, but also would not be feasible in the field due to cost, reliability, and power constraints.

5.2 Landmine Sensing

Table 4: Landmine Sensing Trade Study

	Cost	Sensing accuracy	Sensing efficiency	Ease of sig. proc.	Availability	Tech. maturity	Mass	**Total**
	15%	25%	15%	20%	10%	10%	5%	**100%**
Probe	4	3	2	4	3	5	4	**3.45**
Metal Detector	3	2	4	3	5	5	3	**3.30**
GPR	1	4	5	2	4	4	2	**3.20**
Chemical Sensing	2	2	4	5	1	3	5	**3.05**

The trade study (Table 4) has revealed that two complementary sensing modalities for detecting landmines are a metal detector, for its non-contact sensing ability which contributes to an efficient sensing approach, and probes, for their low-cost contact sensing ability. This combination of sensors reflects the approach that deminers use in the field, where a metal detector is used to detect potential landmines, and probes are used to determine whether a positive signal from a metal detector is true or false. The trade study showed what is already known in the field: that no single sensor can satisfy all requirements.

In retrospect, this study had the right conclusions, but we did not accurately estimate the importance of cost. We designed a system that cost under $5,000, whereas Ground Penetrating Radar (GPR) used in the field costs two orders of magnitude greater. A better analysis at the time would have set us more firmly in the direction we ultimately went. Additionally, we felt that we improperly rated the maturity of probe technology. Though probes themselves are a simple concept, we had to develop the technology surrounding automating the probing process entirely.

5.3 Localization

Table 5: Localization Trade Study

	Cost	Sensing Accuracy	Ease of signal processing	Technology Maturity	**Total**
	15%	35%	30%	20%	**100%**
Wheel encoders	5	3	5	5	**4.30**
RTK-GPS	2	4	5	4	**4.00**
GPS	5	2	5	5	**3.95**
IMU	5	2	4	5	**3.65**
Lidar	1	5	3	4	**3.60**
Stereo odometry	2	2	2	3	**2.20**

The trade study for localization (Table 5) has indicated that the most suitable technologies for determining the position of the system in a local context is a fusion of wheel odometry and an

Inertial Measurement Unit (IMU). For global accuracy the most effective means of sensing are Global Positioning System (GPS) modules that allow either Real-Time-Kinematics (RTK) or Precise-Point-Positioning (PPP). The only tradeoff between these two GPS technologies would be accuracy and cost, with the more expensive RTK providing centimeter level accuracy, and PPP providing decimeter level accuracy. With cost requirements, PPP is the ideal choice. These sensors together enable acceptable localization accuracy and each can be acquired and integrated with a reasonable cost and engineering effort.

Looking back, this trade study was somewhat inaccurate in estimating the importance of differences in GPS accuracy. We ultimately had to reduce our localization accuracy requirements because PPP GPS was not sufficiently accurate and we struggled with GPS drift. Using several localization sensors (GPS, IMU, encoders) and fusing the information was the final approach.

6 Cyberphysical Architecture

6.1 Overview

The cyberphysical architecture in Figure 7 maps the flow of data and energy between components and subsystems based on the results of our trade studies. The notable aspects are that:

- The vehicle is manually pushed
- The localization is provided by a synthesis of GPS, wheel encoder, and IMU data
- A touchscreen display is used for user interaction
- The probe force sensing is provided by a compression load cell
- A combination of off-the-shelf Direct Current (DC) motors, servos and stepper motors are used for actuation
- An off-the-shelf metal detector is used

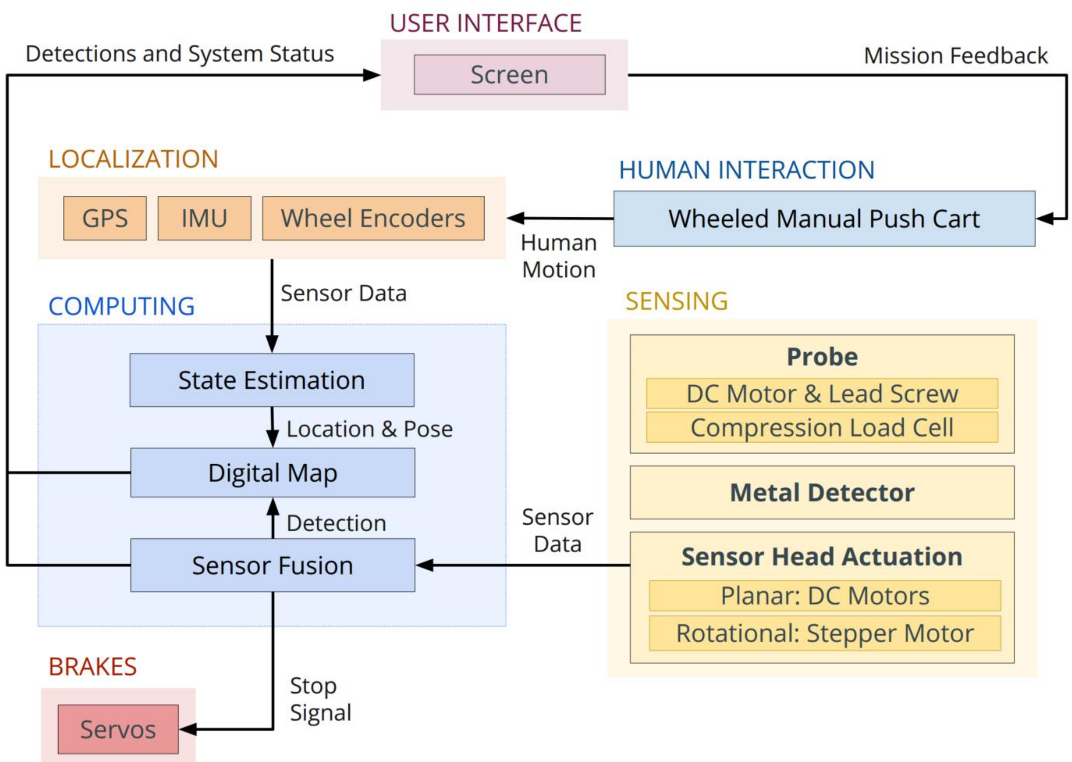

Figure 7: Cyberphysical Architecture

6.2 Connections

The majority of the subsystems are physically integrated through Universal Serial Bus (USB), connected to an industrial-grade USB hub, which relays all serial information to our computing unit through a single high-speed USB connector. Figure 8 illustrates a high-level overview of how the subsystems are integrated physically.

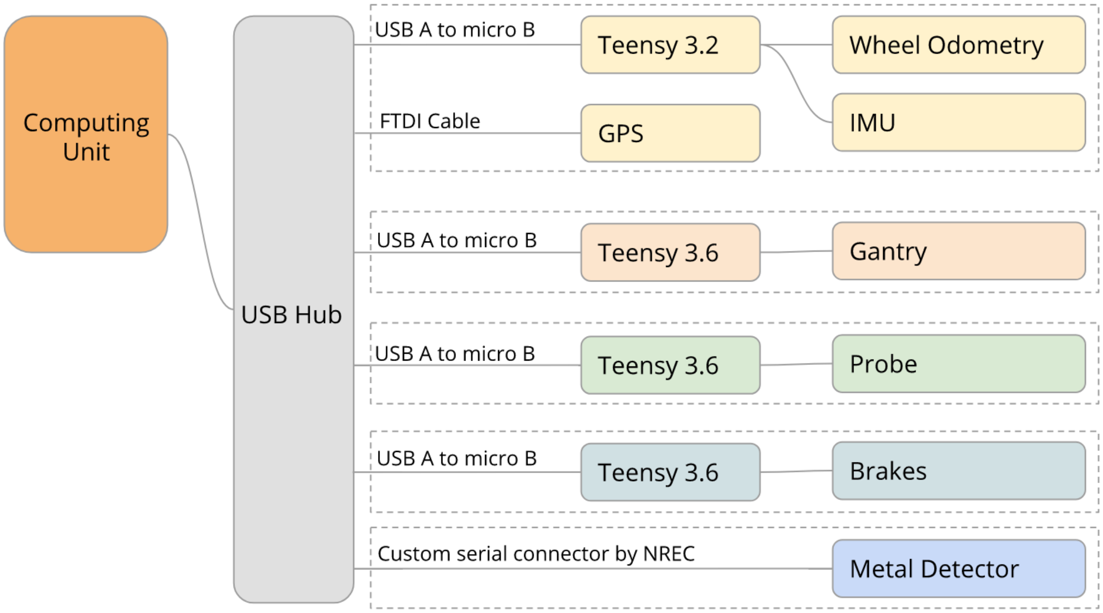

Figure 8: Subsystem Connections

All our subsystems are powered through our power distribution board using a 12V lead-acid battery. Figure 37 in Section 11.1.2 shows how power is distributed to the subsystems and actuators.

6.3 Software

The overall system runs on Robot Operating System (ROS), which provides a high level of modularity and allows for ease of collaboration and integration over multiple subsystems. It utilizes the robust serial packages in ROS to communicate with multiple serial devices and microcontrollers, while our computing unit controls the flow of the system with a comprehensive state machine.

Table 6 lists out the different packages, libraries and algorithms used in each of the subsystems. The specific uses of these components will be further elaborated and discussed in the relevant system description subsections.

Table 6: Subsystem software packages and algorithms

Gantry	Rosserial, Rapidly-Exploring Random Trees (RRT), Proportional-Integral-Derivative Control (PID)
Metal Detector	Custom Python Serial Driver
Probing	Rosserial, Random Sampling Consensus (RANSAC), Hough Fit with K-means Clustering
Localization	Rosserial, Nmea_serial_driver, Robot_localization, Extended Kalman Filter (EKF), Navsat_transform
Computing Unit	ROS
User Interface	Mapviz, rqt_gui
Brakes	Rosserial

7 System Description & Evaluation

7.1 Vehicle Platform

Figure 9 shows the complete system assembly. The vehicle is known as the Scorpion, which has been already used for UXO clearance tasks in Cambodia, and whose modular construction makes it easy to transport. The gantry and all sensor electronics are located at the head of the platform to distance the operator from potential threats as much as possible. Power electronics and the system battery are located near the base of the platform to give the operator easy access to the power and emergency switches. Localization and computing electronics are mounted towards the center of the cart to separate sensitive electronics from electrical noise emitted by the DC motor driver. Finally, a touch screen is mounted on the cart handlebars so the user can easily see and respond to instructions, as well as review the coverage area map.

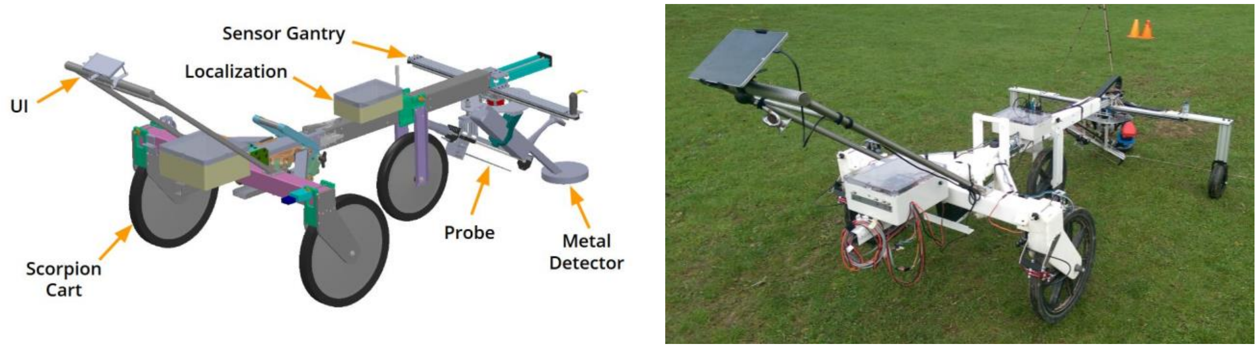

Figure 9: Final Design of the System, Modeled (Left) and in Reality (Right)

7.2 Gantry Subsystem

7.2.1 Current System

The gantry provides a structure to enable three degrees of freedom for the probe assembly and metal detector (Figure 10). Its main functions are to sweep the metal detector as the platform traverses the minefield, and then to position the probe assembly to investigate targets. Support wheels were added to the ends of the gantry in the Spring semester to increase stability under load.

In sweeping mode, the gantry moves the metal detector back and forth at a user-specified speed. At its maximum speed, the metal detector can cover just over $3m^2$ per minute. Once a metal target is detected, the gantry reports the target's position with respect to the platform and updates that position during the metal detector pinpointing process.

In probing mode, the gantry sends the probe assembly to the necessary X, Y, and θ positions as commanded by the main planner. The gantry can position itself with an accuracy of a millimeter in both the X and Y directions, and to less than a degree in the θ direction.

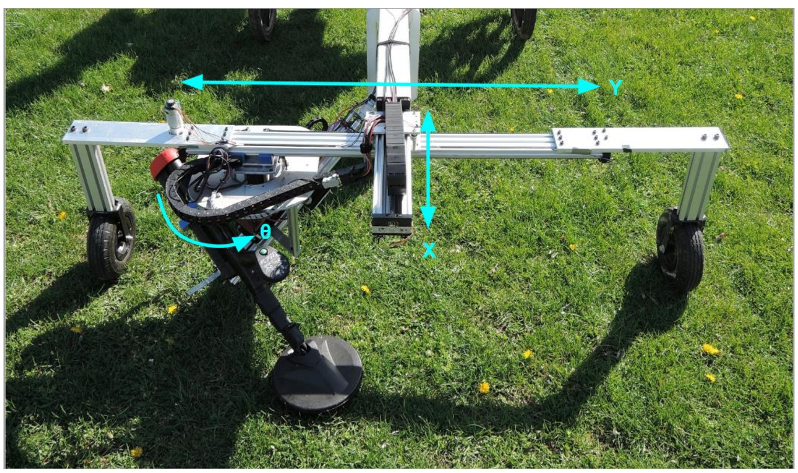

Figure 10: Gantry Degrees of Freedom

The gantry has four main software states: Waiting, Initializing, Sweeping, and Positioning. Its first state stops all of the motors. The second moves the gantry to its X and Y limits to calibrate the DC motor encoders. The third state sweeps the metal detector at a commanded speed. The fourth state accepts positioning commands for all three degrees of freedom. If a limit switch is hit at any time in the Sweeping or Positioning states, the gantry will immediately return to the Waiting state. ROS messages communicate states, positions, and speeds to and from the main system.

Additionally, to comply with the weight limit and mobility of the system, the gantry was designed to easily disconnect from the platform for storage and transit. The entire assembly mounts as a simple press-fit, with some alignment screws. All electrical connections are sided, locking and unique, so connecting the gantry is a plug-and-play operation. Finally, the gantry can connect to either the Scorpion platform for use in the field, or directly to the main power electronics box (the E-Core), for bench testing or demonstrations.

7.2.2 Modeling, Analysis and Testing

The gantry was designed and modeled in SolidWorks, as shown in Figure 11. In order to ensure the proper fit of components, ease of assembly, and completeness of the Bill of Materials (BOM), the Computer-Aided Design (CAD) model includes all fasteners, belts, connectors, and other mechanical components.

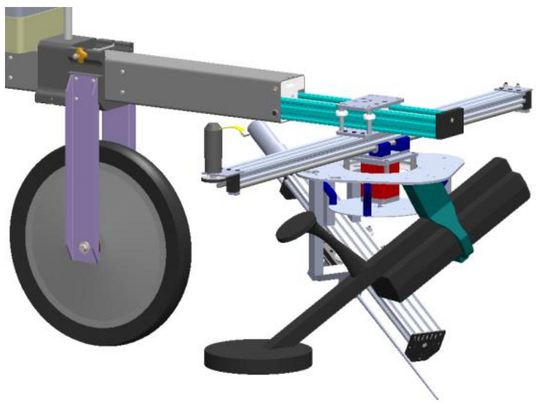

Figure 11: Complete CAD Model of the Gantry

Finite Element Analysis (FEA) was performed to some mechanical parts to assure that they could withstand the load of both sensors, as well as the force from probing into the ground at their fully cantilevered positions (Figure 12). While the analysis of this particular failure mode proved to be accurate, we neglected to consider the strength of mechanical connections on the mid-section of the Scorpion frame, as well as the cumulative effects of tolerance mismatches along the frame and gantry itself, affecting the stability of the gantry when in motion. To correct this, we added wheels on either side of the gantry to prevent defections or excessive forces on the Scorpion assembly.

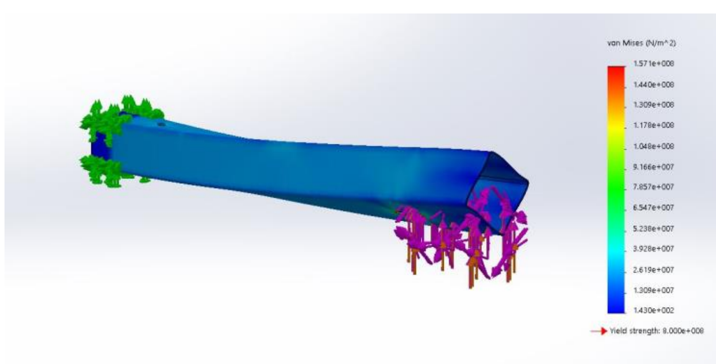

Figure 12: FEA Simulation of Gantry Mounting Forces on Carbon Fiber Scorpion Frame

7.2.3 SVE Performance Evaluation

The pertinent requirements for the gantry are:

- Surveys land at a rate of 3 m^2/min in an area known to be free of metallic objects (M.N.3).
- Individual component assemblies weigh less than 50kg (D.N.2).

We demonstrated the first requirement by presenting a video in which the metal detector covered 3.15m^2 in 60 seconds. To verify complete ground coverage, markers were laid out at intervals of one metal detector head diameter and the cart operator passed over all of them. The weight of the Scorpion alone is 15kg and the gantry 41kg, meeting the second requirement.

7.2.4 Strengths and Weaknesses

As mentioned previously, the mid-section connection clamps on the Scorpion platform were not robust enough under the full load of the gantry in motion. Whenever the metal detector entered sweeping mode, these clamps would often vibrate open, causing the vehicle to split in half. To remedy this, we added two wheels on either side of the gantry for added support. We also replaced the top clasp with a structural aluminum plate. Each solution is presented in Figure 13. Given additional time and resources, we would have remedied this through a more application-specific platform design.

Figure 13: Faulty Clasps (left) and Support Wheels (right)

In each of the gantry's degrees of freedom, there is some level of backlash or slack that adds to system wobble or detracts from accuracy. Much of this could be easily remedied with slightly more expensive precision extrusion for our X- and Y-axes, and an industrial-grade turntable for the rotational axis. Despite these weaknesses, the gantry proved to be robust and reliable even under almost constant use during demonstration preparation.

7.3 Metal Detector Subsystem

7.3.1 Subsystem Description

7.3.1.1 Mechanical Design

The metal detector is mounted on the gantry using a 3D-printed bracket that attaches to the rotational plate of the gantry and firmly clamps the metal detector (Figure 14). A cable leads from the metal detector to the Scorpion's on-board processor for computing.

Figure 14: Metal Detector Mounted on Gantry

7.3.1.2 Software Design

The metal detector software architecture is described in Figure 15. It consists of several modular software packages: an acquisition node that interfaces with the metal detector, a filtering node that combines the data from the metal detector and the gantry, and a data analysis node that takes in the combined data, analyzes that data, and uses the conclusion of the analysis to generate commands for the gantry to sweep or to move to specific positions in order to pinpoint. A gantry command library interprets those commands and relays them to the gantry.

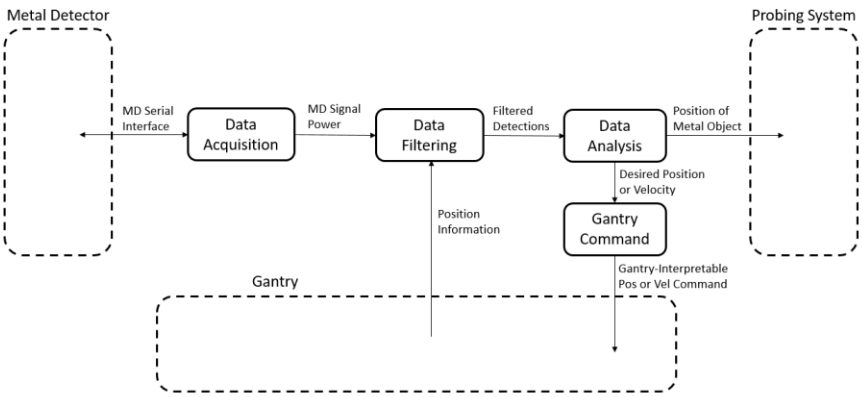

Figure 15: Metal Detector Subsystem Software Architecture

Once a sufficiently high detection has been found during sweeping, the pinpointing begins. The analysis package uses a recursive two-axis ascent algorithm to pinpoint a mine. The algorithm collects data while moving a fixed distance in each direction along the X-axis. It then finds the maximum signal location along that axis. It moves to that maximum signal location, then repeats the procedure along the Y-axis. This procedure is repeated in each axis, and the final maximum signal location is determined to be the metal object center.

7.3.2 Modeling, Analysis and Testing

As with every piece of software in our system, the metal detector software was fully testable in simulation. The important inputs and outputs of the software subsystem are metal detector signal and gantry position information in and gantry commands out. We collected the metal detector signature of our targets (Figure 16) and used that data to create an accurate metal detector simulation of a mine. A model of the gantry was also built to simulate the real gantry. In combination, these allowed the metal detector software to be fully tested in simulation.

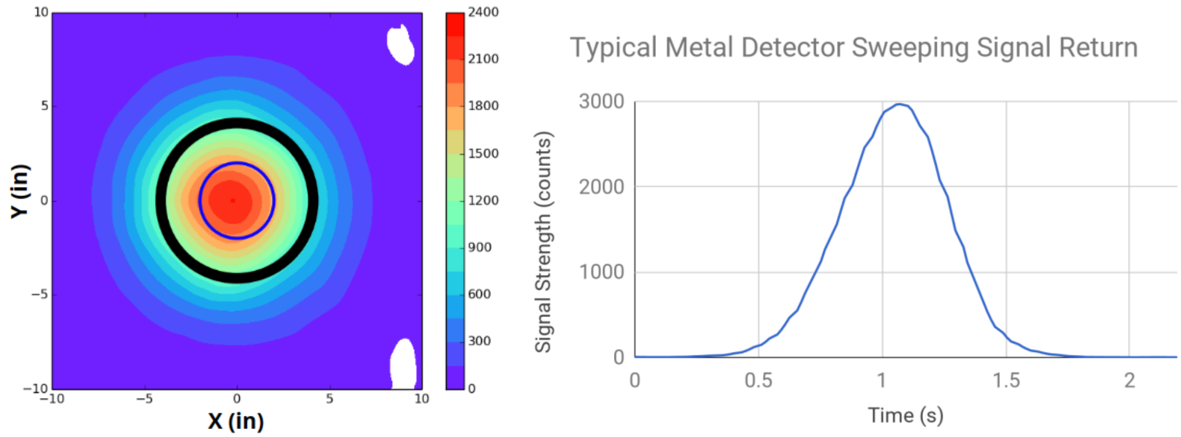

Figure 16: Metal Detector Signals over a Target (Left: Static. Right: At 1m/s sweeping speed)

7.3.3 SVE Performance Evaluation

The pertinent requirement for the metal detector subsystem is: "Physically marks landmine locations to an accuracy of 10cm" (M.F.4). This requirement is indeed a system-level requirement, but as our scheme estimates the final location of the object using the results from the metal detector pinpointing sequence, it is presented in this section. Figure 17 shows a histogram of the differences between the predicted and actual centers of landmine targets, based on pre-SVE trials, verifying that the requirement is met.

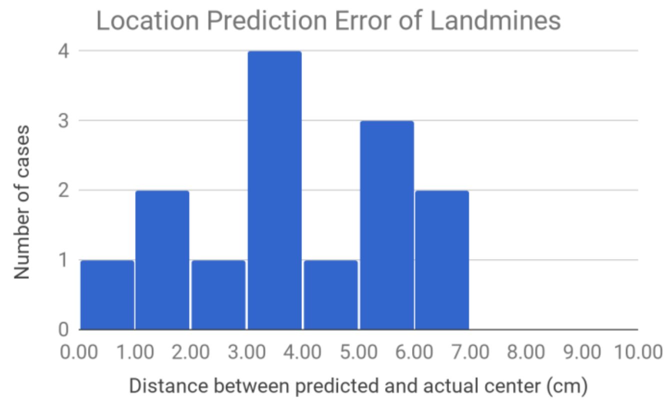

Figure 17: Location Prediction Error of Landmines

7.3.4 Strengths and Weaknesses

As mentioned, the metal detector subsystem fulfilled its requirements in localizing metal targets. Additionally, the subsystem would always find the same metal object center when it first detected the metal object at the same position. This consistency made it easier to test and tune the pinpointing. Additionally, the system was relatively robust and reliable, suffering few hardware or software failures once it was assembled.

There were two weaknesses of the metal detector subsystem. The first was inherent in the challenge of metal detecting for localization. The metal detector cannot tell the difference between a small object with a strong signature and a large object with a weaker signature. This shortfall meant that the subsystem could not detect the edges or size of an object, only the center. Another weakness was that the pinpointed center location depended on the rotational symmetry of an object. The pinpointed center is at the max metal detector signal strength, which did not always match the actual center of an object.

In the first Spring Validation Experiment (SVE) demonstration, the metal detector was inconsistent in finding the correct center location of metal targets. We adjusted the pinpointing algorithm, and the updated metal detector subsystem fulfilled its requirements and accomplished its goal of accurate target localization. Because of the limited travel available on the gantry X-axis, we had to make the first detection of a metal target in its back half to pinpoint it correctly. However, when we did that, it correctly localized the object each time.

7.4 Probe Subsystem

7.4.1 Subsystem Description

The probe's high-level function is to emulate the process by which humanitarian deminers use a probe to determine the shape of target which the metal detector has sensed in order to decide whether the target is, in fact, a mine which is worth removing. The unit uses the full range of motion of the gantry to generate contact points around the circumference of the mine.

7.4.1.1 Mechanical Design

Figure 18 shows the final mechanical design of the probe unit, where the system uses a linear rail with a lead screw and DC motor to actuate the probe. During probe insertion, the shaft compresses against the load cell, generating a signal proportional to the force at the tip. The unit is mounted at 30° to the ground to emulate the probe approach angle which actual deminers use.

Figure 18: Probe Assembly Mounted on Gantry

7.4.1.2 Low-Level Software Design

The nature of the classification task performed by the probes is binary; it seeks to answer the question of whether the sensed subsurface shape of an object is similar to a landmine. This approach is based on current demining practices in which the regular spatial intervals of probe insertion represent the deminer's prior knowledge of the geometry of the expected target. If the deminer does not feel a continuous sequence of subsurface contact points then the suspected target can be safely dismissed as shrapnel, and not worth further investigation for disposal.

As depicted in the visual description of the state machine in Figure 19, the first stage of operation is to initialize the force sensor. In order to effectively determine whether the probing system has contacted a subsurface object or not, it is necessary to determine a force derivative threshold. This was achieved by adding a calibration state to the system and then empirically determining an appropriate force derivative threshold. Initial results with this method were successful for determining whether the probes have come into contact with any subsurface object. After the initialization is complete, the probing can begin. For an estimated target location, a sequence of probes is generated, corresponding to an array of gantry positions. After the gantry position has been sent and it has reached its commanded position, the “Probe” state is commanded. Two conditions can trigger the “Stop” state, in which the carriage returns to its home position. The first is if the force gradient threshold is exceeded and the second is if the end stop is triggered without any contact. The microcontroller managing the probe constantly sends back the carriage position, which is used as the basis for calculating the coordinates of the contact point.

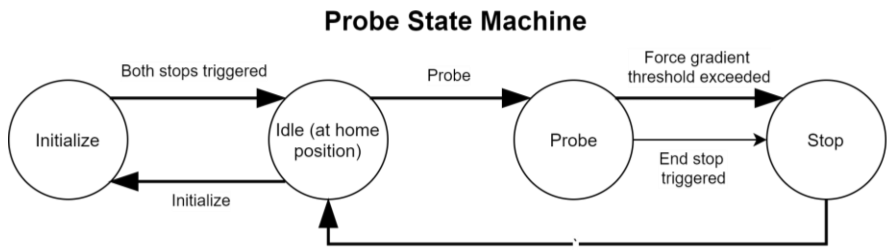

Figure 19: Probe State Machine

To better emulate the manner in which a human performs a probing activity, a feature which slows down the probing velocity as the depth of insertion and reaction force increases was implemented. This process also serves the crucial function of reducing the opportunity for the probe to displace or disturb the buried object when contacting it. A model for this motion was derived empirically, as in the function below, where the velocity of the linear actuator is a function of the probe reaction force. The constant α is computed after calibration, using the maximum insertion force value. Section 11.2.3 shows the linear velocity response of the system over a range of probe reaction force values, for a calibrated maximum probe reaction force of 10kgf.

$$Velocity_{Linear\ Actuator} = 1/(Force_{Probe} \times \alpha) + 1)$$

7.4.1.3 Probe Sequence Planning

The approach to planning the probe contact locations was modelled upon the current strategy used by deminers, which is to approach the suspected landmine location in small increments as to not detonate the target, as illustrated in Figure 20. The planned probe location is based on the desired probe tip position at full extension, and homogeneous transformations which are parameterized by the system’s coordinate origins are used to determine the associated gantry location for every desired probe tip location. In our implementation of this incremental ‘stepping’ approach, we empirically tuned the stepping safety factor to optimize performance in our test field. Once the probe has made contact with the object, a range of 11 feasible sample angles are generated and executed, as in Figure 20.

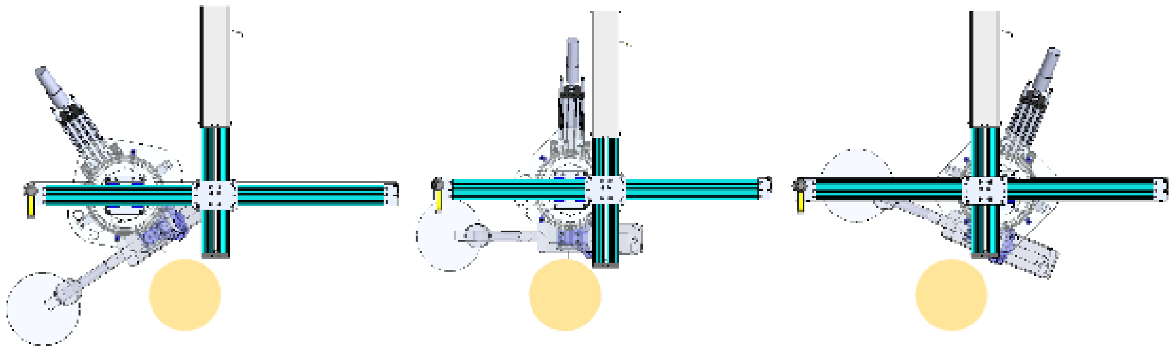

Figure 20: Range of Approach Angles to Detect the Landmine

The inclusion of the gantry support wheels imposed a set of motion constraints which needed to be accounted for while planning the motion of the gantry. Figure 21 shows the worst-case transition condition, where a simple rotation of the sensor head would cause the metal detector to contact the support wheel. An RRT path planner with path shortening was implemented to find the sequence of gantry motions to avoid collision. Figure 22 shows the three motion commands required to avoid the constraints in the blue boxes. While there were few times during testing that this complex dynamic sequence was required, the protections that it provided were essential to avoiding hardware damage.

Figure 21: Worst-Case Gantry Transition Condition (Left: Start. Right: End)

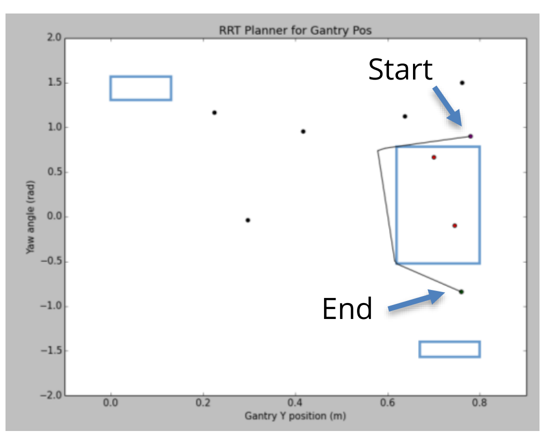

Figure 22: RRT Planner Solution for Avoiding Obstructions

7.4.1.4 Landmine Classification

The team is operating under the assumption that the landmine shape is cylindrical, oriented such that the face of the cylinder is parallel with the ground plane and that the diameter is of a known size. This is a fair assumption in most places where particular landmine types are associated with particular conflicts. As such, our approach for determining whether or not the object is a landmine or not is based on the quality of a fitted circle to the ground-plane projection of the contact points. The initial technique which was implemented was a least square fit, but we found that this was sensitive to outliers and could easily diverge from the real solution. The final approach which we employed is a Hough transform, which provided stable results on simulated and real data. The implementation is described for a set of probe contact points in Figure 23a. In Figure 23b, the Hough circles of the known landmine diameter are drawn and the intersections between each circle are found and plotted. In Figure 23c, the intersection points are thresholded to a bounding box to remove outliers, and then K-means clustering is performed to find the approximate center of the most significant cluster. It was determined empirically that finding two K-means points and then selecting the one with the greatest number of 'closest intersection points' was most suitable for further removing noisy intersections. Figure 23d finally shows the fit error process, which is the mean Euclidean distance between the points and the newly fitted circle, normalized by the number of contact points.

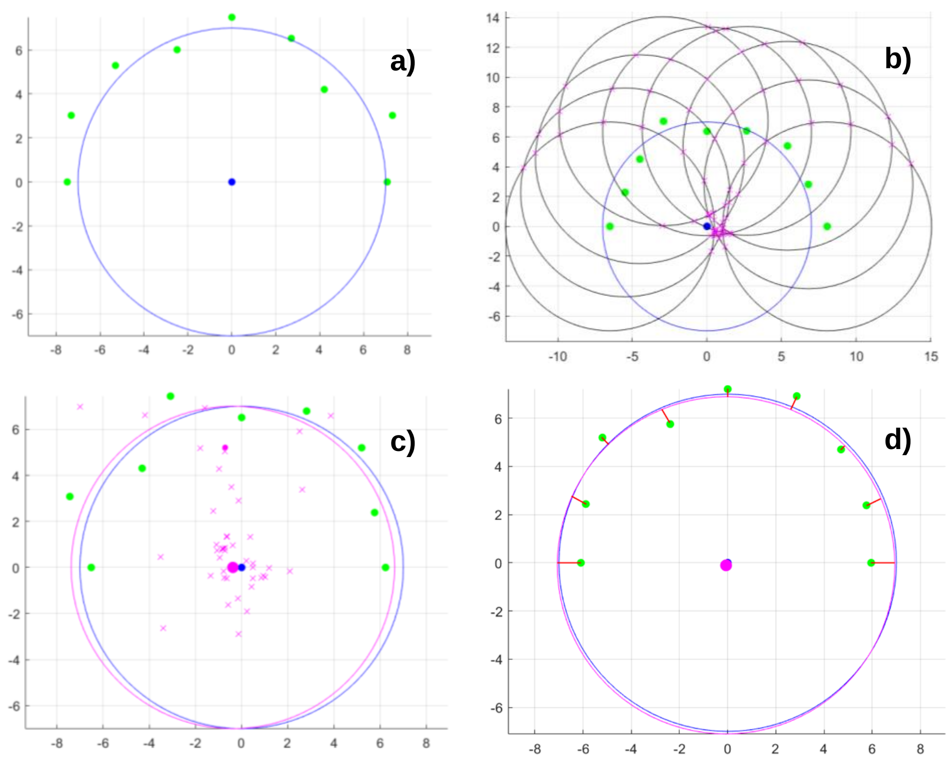

Figure 23: Hough Transform Circle Fit

Once the contact points are gathered, RANSAC is used to randomly select five points, generate a circle fit using the Hough transform, and then evaluate the number of inliers based on the threshold of the Euclidean error metric depicted in Figure 23d. This is repeated for a fixed set of random attempts and then the fit which has the highest number of inliers is selected. The classification criteria is to use a threshold of 0.5 for the landmine metric (described below), where a set of points with a landmine metric greater than 0.5 are classified as landmines. The value for the threshold was determined empirically from the limited data which was collected from the operational system.

$$Landmine\ metric\ = \#RANSAC\ inliers/\#probe\ attempts$$

7.4.2 Modeling, Analysis and Testing

The team converged on a selection of the probe as a ⅛" steel rod, which produced the lowest mean insertion force of each surveyed option, based on results in Section 11.2.1. The selection of this probe served to relax the force delivery requirements of the linear actuation system and the reaction force for which the mounting platform must repeatedly withstand.

The most effective approach for learning about the probe subsystem was to test it in real soil conditions. Figure 24 compares the force and its gradient for five insertion cycles in potting soil with and without contact with a buried object. The force gradually increases with probing depth and when there is an obstruction, increases drastically. The force gradient, normalized by insertion speed, is used to determine whether contact has been made. It is clear in the case of Figure 24 that a threshold of 5kgf/s for the normalized force gradient will suffice to determine whether the probe has contacted an object or not.

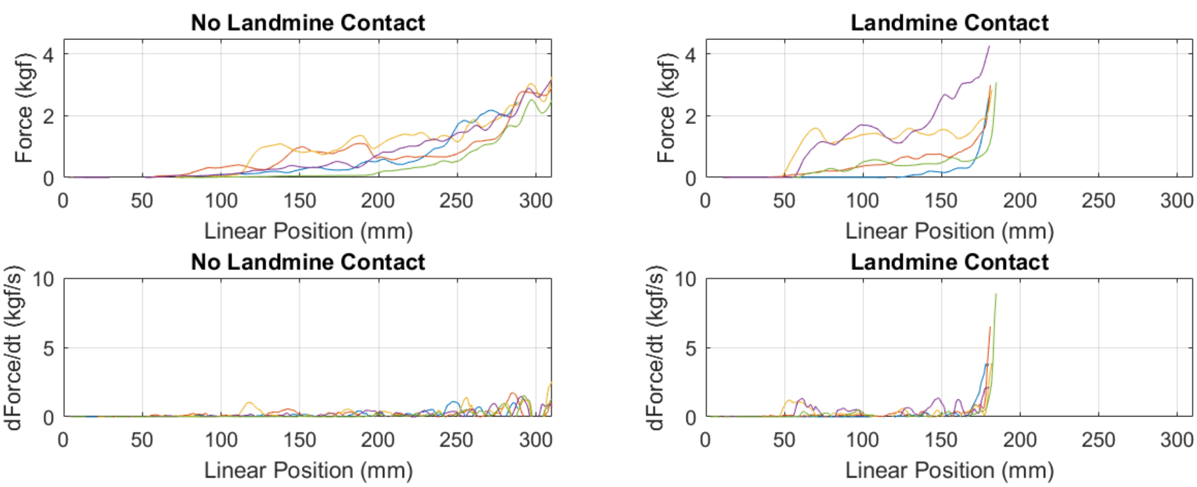

Figure 24: Probe Insertion and Contact Force Profiles

A simulation of the probing procedure was developed using the information that was known about the real-world constraints of the system. This was particularly useful for developing the planning strategy for the landmine probing procedure, and ensuring that it was in relatively ready state for testing on the real system after it had been rigorously tested in simulation. The simulation was also developed in a modular fashion such that simulated probe node could be easily replaced with the real microcontroller.

To develop an understanding of the performance of the classification system, 30 tests were carried out on the test field, where 20 were performed on landmine targets and 10 were conducted on non-landmine targets under the same conditions as the SVE. Table 7 describes these results and associated statistics. Figure 25 depicts a small subset of the visualized data during this testing, where the red spheres indicate contact point locations, and the red circle depicts the circle fit. The variance in the reported contact points can be attributed to deflection of the system as it makes contact with a landmine, in addition to the variation in soil conditions.

Table 7: Probing System Test Results

	Landmine	**Non-Landmine**
Targets Surveyed (#)	20	10
Classification Accuracy (%)	100%	90%
Mean Attempted Probes (#)	8.0	10.4
Mean Inlier Contact Points (#)	6.7	2.0
Mean Landmine Metric (%)	83.8%	19.2%
Mean Contact Point Error (mm)	6.7	-

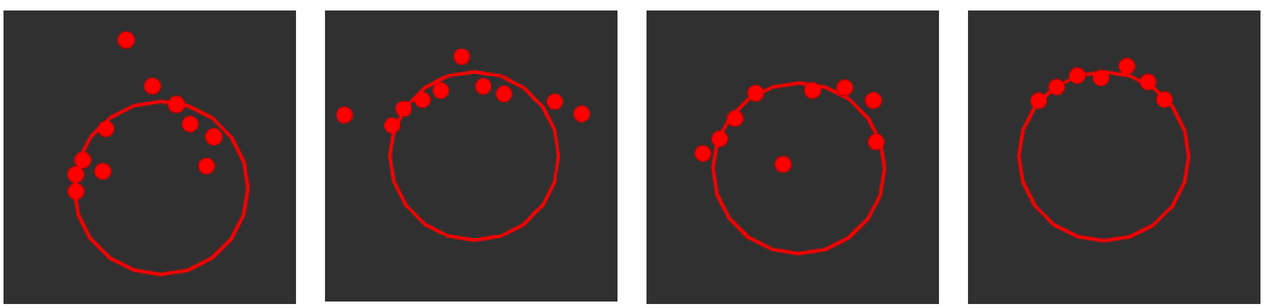

Figure 25: Contact Point Visualization

7.4.3 SVE Performance Evaluation

The pertinent requirements for the probing subsystem are:

- Detects 4 out of 5 landmines in a grassy field buried to a maximum depth of 10cm (M.F.1) (Figure 26)
- Correctly classifies 3 out of 5 metallic non-landmine-shaped objects as not landmines (M.F.2) (Figure 26)
- Detonates zero landmines during probing action (M.F.3)

- Determines whether a detection is a true or false positive within 5 minutes (M.N.4)

Each requirement was met in the SVE, where 2/2 true landmines were classified as such, and 2/2 non-landmines were classified as such, without the probe tip contacting the upper surface of the target. The average time for a typical probing sequence was approximately 2 minutes, satisfying our probe time requirement.

Figure 26: Target Types (Left: Landmine. Right: Non-Landmine)

7.4.4 Strengths and Weaknesses

A key strength of this subsystem was the effective use of off-the-shelf components for the mechanical and electrical design. This enabled us to quickly design and assemble the system, and to easily swap out components if they became worn or damaged. Eventually, if this system were to be deployed in remote locations, the use of accessible components would be key for any sustainable maintenance of the unit.

The probe subsystem exhibited vast performance variations in different soil conditions and was sensitive to inconsistent patches of soil or areas with high concentrations of rocks and hard objects. This was somewhat compensated for by the probing procedure adjustments from calibration and the relatively high number of probe samples in an attempt to effectively increase the signal-to-noise ratio from the noisy samples which resulted from premature contact with rocks or similar objects. The probe could also transmit large forces to the gantry, depending on the hardness of the soil which is being probed, which limited the range of operating conditions that we could act in. In order to mitigate this system weakness, we reduced the effective soil hardness by softening the ground with water before probing, a process which is actually employed in the field by deminers. Finally, another weakness of the probe itself is that it is subject to deformation after repeated use in hard soil. This aspect of the design does not affect our results significantly because we were able to quickly replace a damaged or deformed probe because of our modular design.

7.6 Localization Subsystem

7.6.1 Subsystem Description

The localization subsystem allows the system to record and display on the User Interface the location of detected landmines as well as the path traversed by the platform. It comprises a GPS capable of Precise Point Positioning (Ublox NEO7P), an IMU (BNO055) and a custom magnetic wheel encoding system (AS5045), as shown in Figure 27.

The absolute position from the GPS, heading and gyroscopic data from the IMU as well as the linear velocity of the wheels are used to estimate the global position of the platform using an Extended Kalman Filter and the navsat_transform node from the robot_localization ROS package, which uses an omnidirectional motion model. The location of the classified target is estimated using the transform between the estimated location relative to the platform and Universal Transverse Mercator (UTM) origin.

As we were not allowed to permanently modify the Scorpion, we designed and built a non-intrusive method of measuring wheel velocity using an extra set of smaller wheels tensioned against the main wheels of the platform (Figure 27).

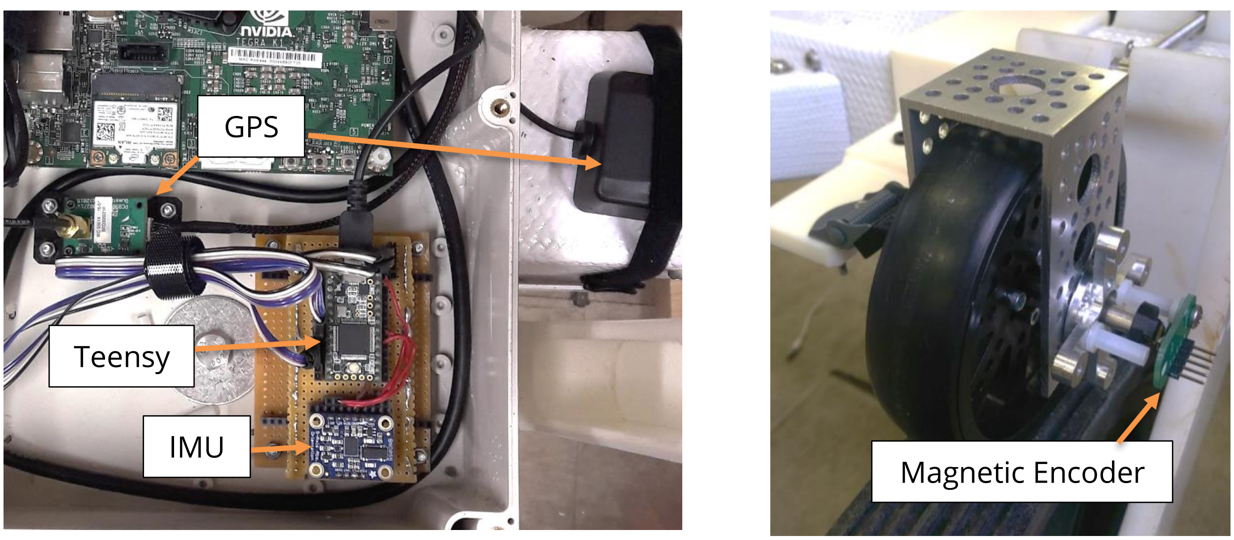

Figure 27: Localization Sensors (Left: GPS (NEO7) and IMU (BNO055). Right: Custom Wheel Encoding System (AS5045))

7.6.2 Modeling, Analysis, Testing

The combination of packages used in the localization stack to achieve our requirements was verified during the Fall Validation Experiment (FVE). This was a combination derived from a Dual EKF, which feeds the filtered odometry from the EKF into the global state estimator, and uses the output of that as another sensor input into the EKF for the next time step. Figure 42 in Section 11.3 shows the integration of the different packages.

We conducted two tests in the NREC carpark to verify the vehicle's target localization accuracy and global path traversal accuracy (Figure 28). We used the visible parking space lines as the ground truth and pushed the cart at the standard field traversal speed. The closely overlapping circuits in Figure 28 show the repeatability of localization to a relative and absolute distance of less than 1m.

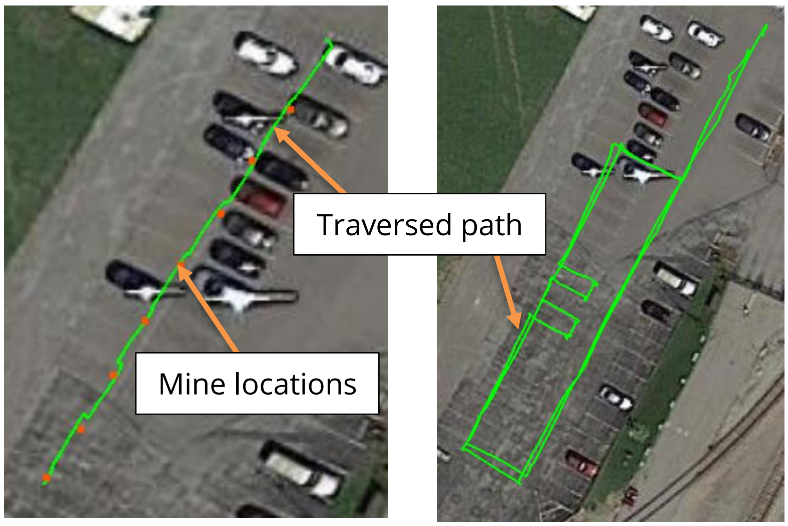

Figure 28: Vehicle Track (Left: Target Localization Test. Right: Path Traversal Test)

7.6.3 SVE Performance Evaluation

The pertinent requirements for the localization subsystem are:

- Records landmine locations on digital map to an accuracy of 1m (D.F.1)
- Records area covered by system on digital map to an accuracy of 1m (D.F.2)

For the final demo, we presented a fully integrated system, which was able to show the path traversed by the platform, as well as the classified targets. We also conducted a trial showing performance over longer periods of time, as shown in Figure 29.


Figure 29: Sample Inspection Sequence Showing the Area Covered and Mine Locations

7.6.4 Strengths and Weaknesses

The strength of our localization subsystem are its modular ROS-based architecture, which makes it easy to reconfigure to our localization sensor suite and its low cost of under $200. The drawback of using low-cost components is that our global localization suffers in several cases. Specifically, when the platform is stationary during the probing sequence, the GPS would report a drift in the estimated position. To rectify this, the input from the GPS was ignored for the period that the brakes were engaged. Despite this, we believe the low cost of this subsystem, motivated by our specific use case, outweighs the drawbacks in localization accuracy.

7.7 Electronics and Computing Subsystem

7.7.1 Subsystem Description

In an attempt to reduce the potential for electrical connections to become a point of error or confusion in our system, all electronics were mounted in two separate electronics boxes. The E-Core houses power electronics (including our power distribution board and a 12V to 24V DC-DC converter), motor drivers (for the stepper and DC motors) and gantry control computing unit (a Teensy 3.6 on a custom circuit board). Our localization electronics box (known as the Localization Core) contained our GPS, IMU, braking Teensy and USB hub (Figure 30). These two sets of electronics were separated from each other to prevent potential interference between high-power electronics and low-power computing.

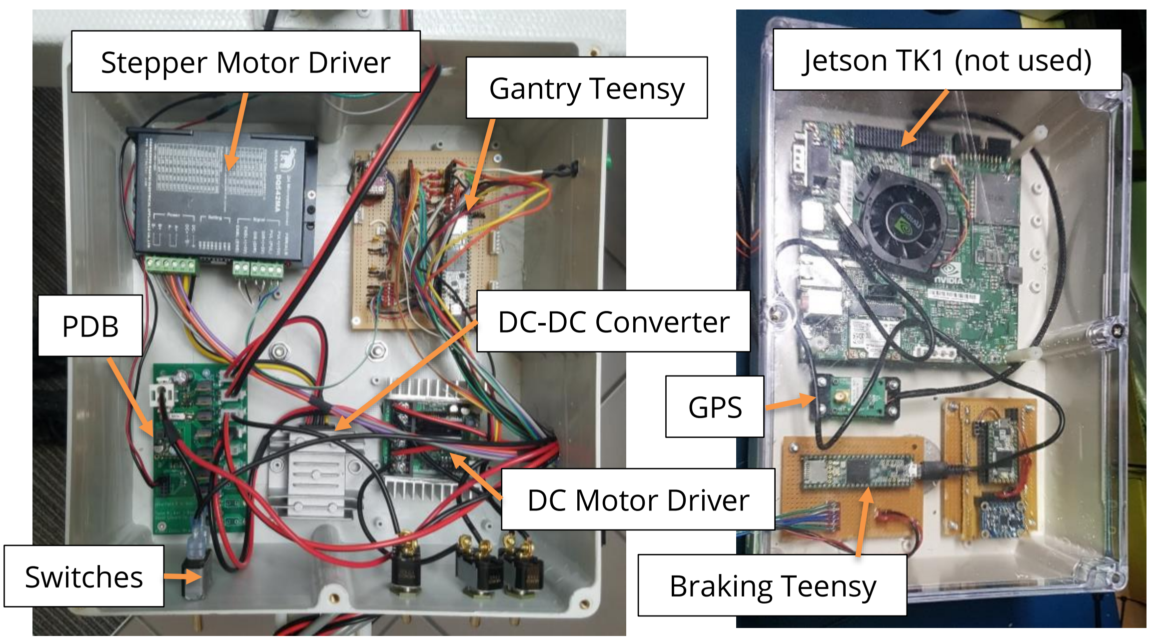

Figure 30: E-Core (Left) and Localization Core (Right)

With safety in mind, all power to the system can be turned off via an emergency switch, placed at an easily accessible distance from the operator. Additionally, all motors were equipped with individual power switches, allowing us to test subsystems independently. Our power distribution board was designed with fuses on all outputs and inputs, in addition to reverse and over-voltage protection. The power distribution board (PDB) supplies a regulated 12V to all motors, the UI and other peripherals and 5V to the braking system (Figure 31).

Figure 31: Power Distribution Board

7.7.2 Modeling, Analysis and Testing

Routing diagrams were produced for each electronics box in anticipation of the complexity of wire routing and harnessing. This allowed us to plan and prepare harnesses and electronics boxes long before they were needed, and to plan wire routing in a structured way. See Figure 37 in Section 11.3 for the wiring diagrams.

7.7.3 SVE Performance Evaluation

The pertinent requirement for the electronics and computing subsystem is: “Able to be operated continuously for 1hr” (M.N.2). To verify this, a simulation of a typical mission complete with simulated targets was set up indoors and the system was run for the full 1hr period.

7.7.4 Strengths and Weaknesses

The effort that was invested into building the electronics boxes was justified, as it lead to almost no issues with poor-quality or incorrect electrical connections. We found that it was easy to control the power to individual systems and that the flow of data and energy was easy to trace. In one instance, the PDB was accidentally connected directly to a benchtop power supply rather than a battery, causing a TVS diode to burn up. Despite this, all the peripheral electronics were protected.

7.8 User Interface Subsystem

7.8.1 Subsystem Description

The user interface subsystem consists of a touch screen interface as well as graphical applications displaying path coverage, system status, user commands and allowing user input. Figure 32 shows the control UI.

Figure 32: Screenshot of the UI

7.8.2 Modeling, Analysis and Testing

The UI was subject to informal user testing during its development which led to incremental improvements to its appearance.

7.8.3 SVE Performance Evaluation

The pertinent requirement for the user interface is: "Presents a graphical representation of land area covered and system status in real-time" (D.N.1), which was met by demonstrating it during the SVE.

7.8.4 Strong/Weak Points

The user interface subsystem was subject to competing design interests. On the one hand, we intended to present a model of a deployment-ready high-level interface where the details of the system's operation would be hidden from the user. On the other hand, some of the more advanced reporting and data visualization functions in favor of ease of debugging in the field. In light of this, the subsystem's strength was the UI's simplicity and its main weakness was deploying the ROS mapviz plugin, which included a complex and fragile set of package dependencies.

7.9 Braking Subsystem

7.9.1 Subsystem Description

The braking subsystem consists of a pair of bicycle caliper brakes mounted on custom brackets to the Scorpion platform (Figure 33). Each brake is actuated by a single servo, and remains activated for the duration of the period between an initial metal detector signal and the completion of the inspection sequence. The servos are driven by a Teensy microcontroller.

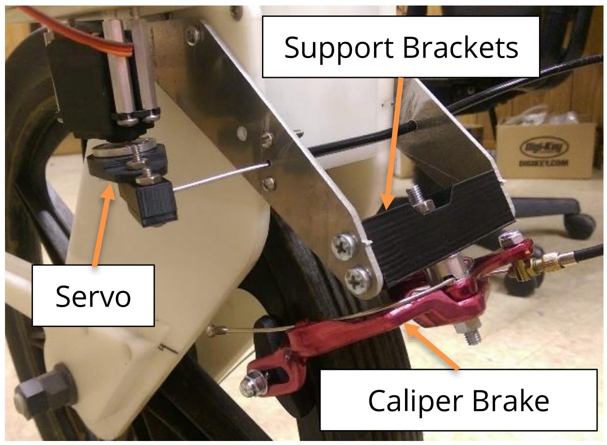

Figure 33: Braking Subsystem

7.9.2 Modeling, Analysis and Testing

The rigging of the brake system was time-consuming and intricate, involving the simultaneous adjustment of springs, cable tension and servo rotation. The use of servos with their fine gradations in rotation angle helped with on-board tuning which would have taken far longer to model accurately.

7.9.3 SVE Performance Evaluation

The pertinent requirement for the brakes is "Warns user of potential hazards within 0.25s" (M.F.5). The time between the command signal being sent and the brakes engaging was captured on video and the sequence took less than 0.25s, meeting the requirement.

7.9.4 Strengths and Weaknesses

The brakes were included in the system architecture as a means of preventing the operator from pushing the vehicle into danger. The strengths of the subsystem included the use of off-the-shelf components and its elegant spatial integration. Its main weakness was the servos being undersized for the level of pressure needed to slow the wheel effectively. In the future, we would use a design which offered a greater mechanical advantage in actuation.

8 Project Management

8.1 Schedule

The final project schedule for the Fall and Spring semester, in Section 11.5, was adhered to by the team. The development of the sensing modalities was limited by lead times associated with the acquisition of the Minelab F3 Metal Detector, and the lack of ability to gather actual data from the final system. This delayed many of the sensing process design decisions to much later in the semester, while the assembly and testing of the gantry actuation system was much more time consuming than expected. The development of this subsystem was on the critical path to full system integration and testing, and as a consequence, the validation of the performance of the gantry with the probe and metal detector was significantly delayed in our schedule. The software subsystem development was scheduled in tandem with the hardware build and testing, where a simulation was developed to mimic the anticipated hardware behavior. Using the ROS framework, the codebase was designed with modularity as a key objective, enabling the simulated software components to be substituted for the microcontrollers on-board the system, which eased the process of integrating the software and hardware. This scheduling decision contributed significantly to the success of the project, and was able to effectively compensate for the delays associated with the gantry integration.

The team failed to correctly estimate the time taken for full system integration of mechanical, electrical and hardware subsystems, and account for the large number of deliverables which were required from other technical courses at the stage of the semester where the integration task was planned for. As such, our soft deadlines were overshot by approximately two weeks, and there was significant pressure on the team to address integration bugs mere days before the demo. Finally, the use of Google Sheets to plan our schedule had added benefits of collaboration and viewing privileges for all team members, but at the expense of a using a less flexible tool for scheduling. Upon reflection, there was limited collaboration that actually occurred with project management, and it would perhaps have been more effective to use a tool which does not possess free collaboration features, but is designed for the specific application of Gantt chart and Work Breakdown Structure tracking.

8.2 Budget

An overview chart of our finances is shown in Figure 34 and a table of major spending categories in Table 8. This shows that we have spent and planned to spend most of our funds. Our budgeting process was very flexible. One member of the team was tasked with periodically updating expected spending in the several categories and confirming purchase requests made by other members. Because our expected spending never exceeded our total project budget, budgetary constraints never had a major impact on our project. Our final metal detector ($1550) was not included because it was purchased on our behalf by our advisor. Our full ordered parts list is available here. In spending less than $5000, the requirement "Costs less than $5000" (M.N.1) was met.

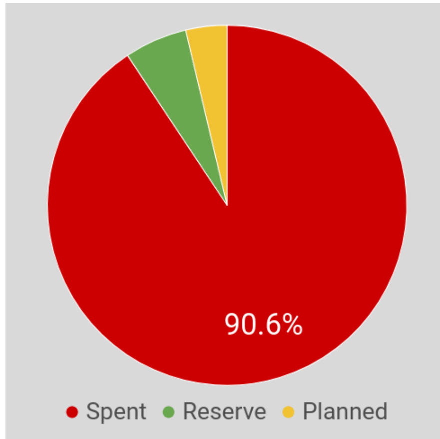

Figure 34: Final Budget Status

Table 8: Budget Spending Categories

Subsystem	Category	Total cost	Spent
Power	Power distribution electronics	$900.00	94%
Gantry	Gantry materials	$1,500.00	100%
Landmine sensing	Force Sensor	$73.00	100%
	Actuators	$400.00	96%
	Probes	$220.00	100%
	Metal detector	$650.00	96%
Pose sensing	Wheel encoders	$105.00	96%
	GPS	$120.00	97%
	IMU	$40.00	87%
User interface	Touch Screen	$120.00	88%
Processor	Processor	$400.00	95%
Mobility	Misc mechanical	$100.00	72%
Testing	Outdoor Testing Equipment	$90.00	90%
Total		$4,718.00	$4,532.46

8.3 Risk Management

Figure 35 shows two risk matrices. The left hand matrix shows the estimate of risks established at the beginning of the project, and the right hand matrix shows the final status after mitigation measures. Table 10 in Section 11.4 details the risks in the matrices. The overall risk profile decreased over the course of the project.

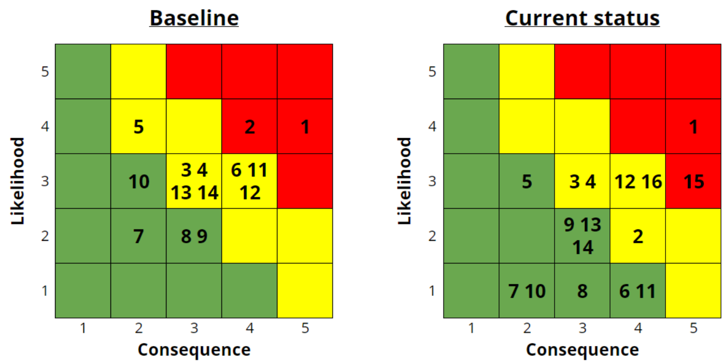

Figure 35: Risk Matrices (Left: Baseline. Right: Final)

In the early stages, the risk profile centered around the "known unknowns" of the viability of our proposed sensor suite and whether our best-guess performance benchmarks could be met. The nature of the most pressing risks evolved over the course of the project. The most persistent risks related to the uncertainty about the performance of the metal detector (#2) and probes (#1).

We were mostly satisfied with our mitigation for this category of risks, which involved presenting a scaled-down version of each sensing modality in the Fall. This bolstered our confidence in the sensitivity and reliability of the metal detector (which we upgraded to a higher-end model in the Spring). We were able to validate the concept of inferring the shape of a buried cylindrical object with a probe using an indoor test soil pit. We were cognizant of having to effectively start from scratch with hardware in the Spring but felt that it was an appropriate trade-off to be able to validate the concept while not being able to transfer the physical hardware. In light of this, we identified that the gantry development would lie on the critical path in the Spring.

One of the lessons we learned from the Fall semester was our over-reliance on physical access to our subscale gantry for testing. To avoid repeating this mistake, we set about developing a high-fidelity simulation which allowed us to test algorithms and subsystem interfaces long before the physical hardware was available, and then integrate them incrementally for hardware-in-the-loop tests. This was a crucial part of our success, which allowed development to continue alongside a series of unexpected mechanical design challenges faced with the gantry.

We did not apply this structured risk management template to our regular team operations, but the documentation we prepared for each design milestone was a helpful guide for how to allocate team resources. Instead, we shaped our work around 2-week sprints between milestones, tackling the major risks for each as they arose.

9 Conclusions

9.1 Lessons Learned

This experience has been very instructive for each team member, and resulted in the following lessons learned:

1. While our system requirements clearly defined our high-level performance benchmarks, the fact that we did not document our subsystem requirements with the same rigor was a setback. For example, we underestimated the magnitude of the gantry deflection under load, using the assumption that the entire mounting structure would be rigid. This discovery - and rapid reaction only four weeks before SVE - may have been avoided if we had set this as a subsystem requirement. In fairness, it is unlikely that we could have predicted this analytically, given that it was caused by a tolerance build up across multiple interfacing components.
2. The time required for hardware and software integration must not be underestimated, especially for subsystems conceived independently. Furthermore, as each subsystem's owner is present at the time of integration, this means that delays due to debugging have a multiplying effect, reducing productivity. The best remedy for this is to establish unit testing criteria well before integration. Additionally, the designer of the individual subsystem should define a series of incremental tests to verify basic capabilities as a foundation to more advanced capabilities. One of the great successes in Spring and strongest improvements over Fall was making early agreements about communications protocols and state transition logic and developing our simulator to allow modular hardware-in-the-loop testing.
3. The disciplined approach we took to cable routing and electrical system integration was an improvement over the Fall. We recognized the risk that a casual attitude towards this activity would likely have brought testing to a halt at the most crucial time, and opted to invest many long, tedious hours into ensuring its reliability.
4. In many instances, the work allocated to a particular team member could not viably be delegated to others, for either lack of skill, a specialized knowledge or single-person physical access. On the software front we made effective use of code branching and a modular architecture to parallelize development, but the burden for hardware development or testing activities often fell to the same one or two people. In the future we would possibly allocate work packages to pairs of people.

9.2 Future Work

In taking on this project, the team chose to pursue the relatively unexplored method of automated subsurface probing. In doing so we believe that we have made a unique and novel contribution to this area of research. We opted early in the design process to prefer a fully integrated field-ready system over a highly refined standalone sensor unit without an operational context. This led us to spread our efforts across multiple domains, perhaps at the expense of a more thorough treatment of the probing subsystem.

In light of this, we hope that someone might be able to carry on our work in the domain of subsurface object classification. As we discovered, the complexity of the hardware poses a very high barrier to entry to even get to the point of being able to collect data in the field, but an indoor test rig with more realistic - in other words, very tough - soil may be a suitable replacement. While the method for classifying mines was deterministic and robust, it was based on the few observations which were made during testing and may not generalize well to other soil conditions and mine types, which was a weakness of the system. In order to generalize to more conditions, classification should be performed in a data-driven method. This kind of method was not used because the time-cost of obtaining data was very high and as such, and an insufficient amount was collected to develop a data-driven classification method.

If we were to evolve this into a business venture, the next logical step would be to travel to a landmine-affected area and embed ourselves in a demining organization. Such exposure would help us to validate our design assumptions and to better understand our place in the workflow of humanitarian demining. Overall, we found a global design optimum subject to the constraints of using off-the-shelf parts and retrofitting an existing vehicle. A revised design would likely step back up the design chain to put gantry stability as the main focus, and make the sensor head more tolerant of a wider variety of terrain conditions.

10 References

[1] - Landmine Monitor 2016. (n.d.). Retrieved September 29, 2017, from http://the-monitor.org/en-gb/reports/2016/landmine-monitor-2016/casualties-and-victim-assistance.aspx#ftn2

[2] - Lin, E. (2016). How War Changes Land. Retrieved September 29, 2017, from http://seareg.org/wp-content/uploads/2016/05/2016_LinSEAREGpaper.pdf

[3] - Care. (2013, October 25). Facts About Landmines. Retrieved September 29, 2017, from http://www.care.org/emergencies/facts-about-landmines

[4] - Bruschini, C., Gros, B., Guerne, F., Pièce, P., & Carmona, O. (1998). Ground penetrating radar and imaging metal detector for antipersonnel mine detection. Journal of Applied Geophysics,40(1-3), 59-71.

[5] - Trevelyan, J., et al. (2016). 58.2.1. In Springer Handbook of Robotics. Cham: Springer International Publishing.

[6] - (n.d.). Retrieved September 29, 2017, from http://pics.siampedia.org/Minefields/Minefields-010.jpg

11 Appendices

11.1 PCB Schematics

11.1.1 Magnetic Encoder PCB

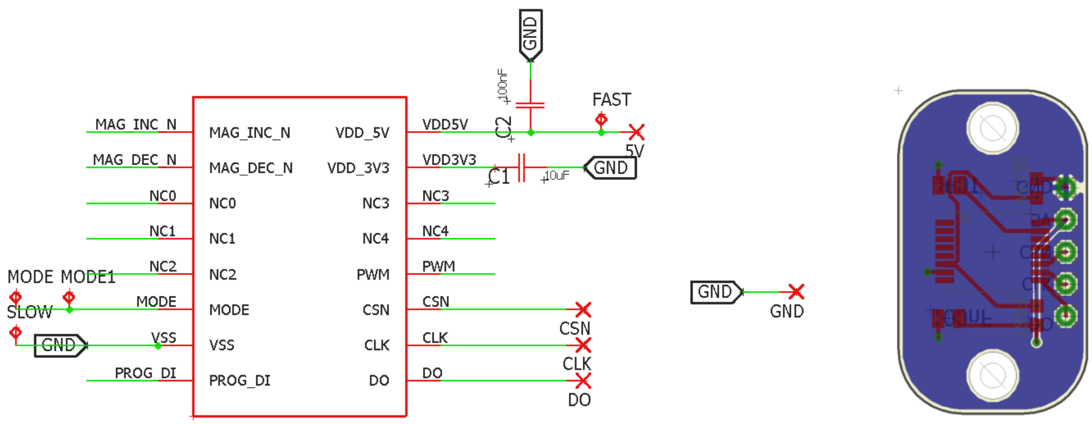

Figure 36: Magnetic Encoder Schematic (Left) and Board (Right)

11.1.2 Power Distribution PCB

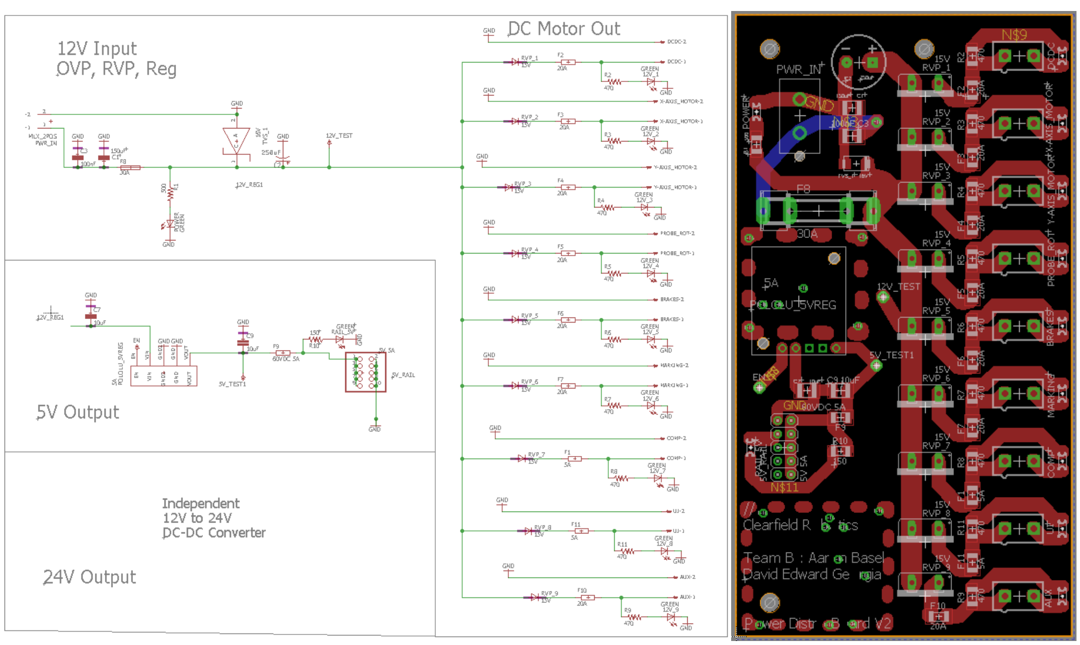

Figure 37: Power Distribution Schematic (Left) and Board (Right)

11.1.3 Wiring Diagrams

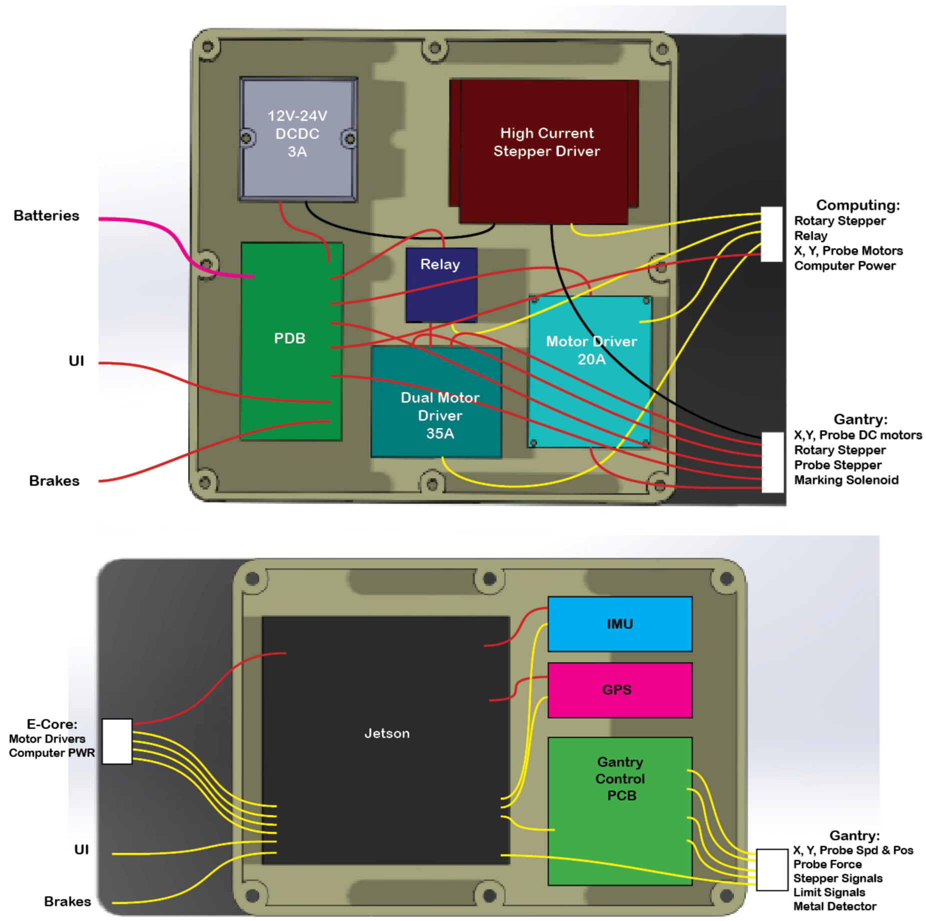

Figure 38: Initial Wiring Diagrams for E-Core (Top) and Localization Core (Bottom)

11.2 Probes

11.2.1 Probe Force Assessment

Table 9: Maximum forces for probe designs

Probe Design	4-8mm Tapered Shaft	3/16” Shaft	1/8” Shaft
Max Force (kgf)	35.15	20.44	10.71

11.2.2 Load Cell Mechanical Design

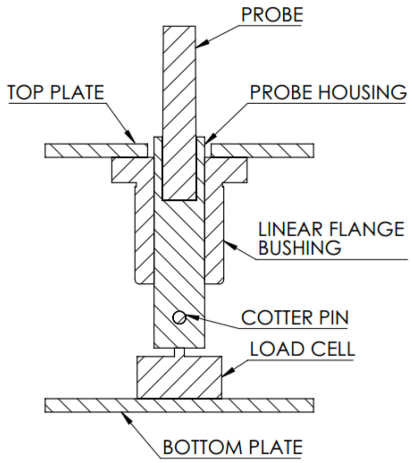

Figure 39: Load Cell Mechanical Design

11.2.3 Probe Linear Velocity and Probe Reaction Force Relationship

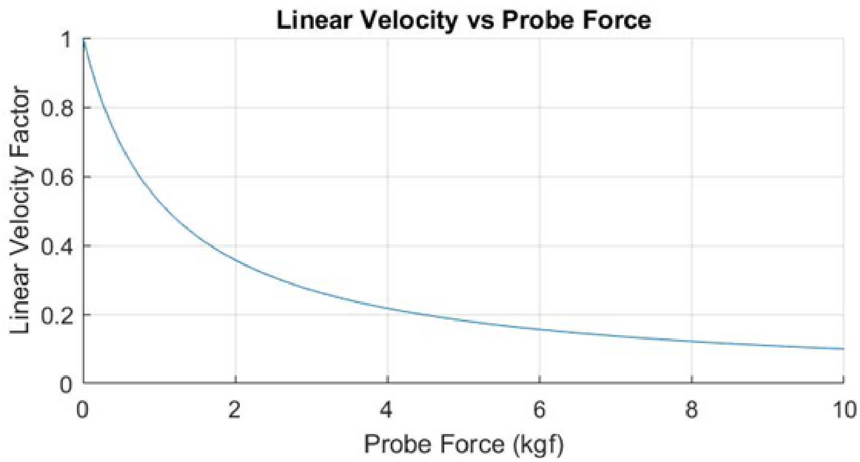

Figure 40: Probe Linear Velocity and Probe Reaction Force Relationship

11.2.4 Probe Reaction Force Results

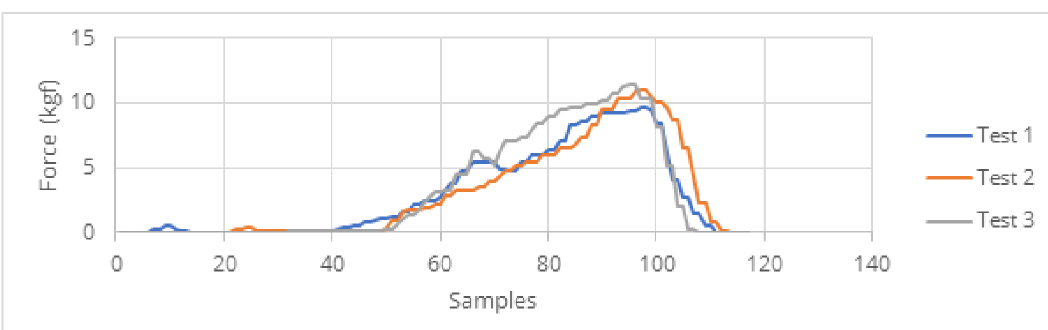

Figure 41: Force Profile for 1/8" Probe

11.3 Localization

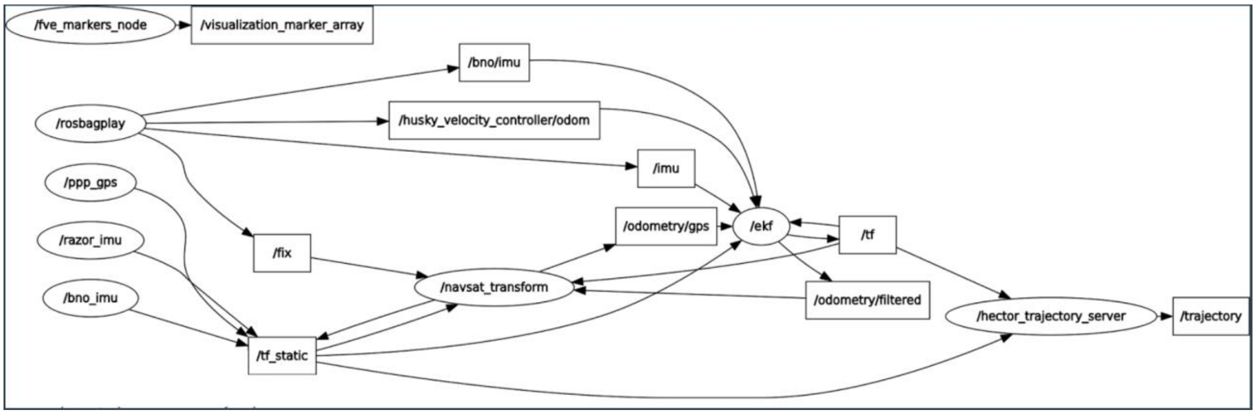

Figure 42: Rosnode graph of localization subsystem

11.4 Risk Management

Table 10: Risk Management Table

ID	Type	Req't	Title	L	C	Possible Consequences	Mitigation strategies
1	Technical	MF1-3, MF5	Low performance of custom mechanical probe design	3	5	Loss of one of two sensing modalities.	Devote proportionately more man-hours in early stages to design and testing. Plan to complete 2 full iterations by FVE.
2	Technical	MF1	Interference between metal detector and electronics	2	4	Metal detector produces false positives.	Conduct subsystem tests with components in close proximity to characterize behavior. Change components to plastic.
3	Technical	All	Critical sensors break	2	3	Unable to proceed with testing or integration.	Identify high-impact or high lead-time sensors. Buy duplicates of fragile electronics <$100. Allow only certain individuals to work with critical components.
4	Schedule	All	Variability in outdoor testing conditions	3	3	Unable to meet testing milestones.	Incorporate a 2 week buffer for outdoor tests and be ready to seize fair weather opportunities as they arise. Become familiar with the specific test environment.
5	Schedule	N/A	High overhead of travel time to NREC	2	2	Loss of productivity.	Limit traveling group to key personnel
6	Resources	N/A	Dependent on NREC resources beyond our control	3	4	Delays in testing.	Confirm all engagements involving NREC personnel or equipment with 1 week lead time. Get written agreement for the terms of use of equipment, facilities and personnel.
7	Resources	N/A	Difficulty in acquiring simulated targets	1	2	Reduced credibility.	If NREC supply is insufficient, contact military training supply companies
8	Resources	MN1	Exceed budget	4	3	Unable to buy spares. Necessity to make design decisions with an artificial cost constraint.	Break down total budget into monthly limits
9	Resources	N/A	Inefficiency with build materials being split between RI and NREC	1	3	Parts get lost in transit. Don't have required parts at either place.	Conduct tests with small-scale and portable equipment at RI and tests with large-scale, bulky equipment at NREC

10	Resources	N/A	Access to areas to plant simulated targets in.	2	2	Unable to perform developmental and validation tests.	Establish relationship with facilities department at CMU and NREC.
11	Resources	N/A	Metal detector acquisition	1	4	Loss of one of two sensing modalities.	Build relationships with multiple vendors.
12	Technical	All	Spatial integration of subsystems	3	4	Unplanned effort for vehicle integration.	Make wireframe integration models early on
13	Technical	DF1, DF2	Bad conditions for GPS	3	3	Mine location mapped will not be accurate location on ground. Unable to show accurate path sweeped and spots missed.	Gather data on good day, post process and visualize later
14	Technical	DF1, DF2	GPS not accurate enough	3	3	Mine location mapped will not be accurate location on ground. Unable to show accurate path sweeped and spots missed.	Get GPS RTK or loosen accuracy requirement
15	Technical	All	Delays in gantry availability prevent testing	3	5	Gantry availability and reliability is on the critical path for outdoor testing.	Ensure modularity and integrity of software packages to reduce delays in physical integration
16	Schedule	MF4, MF5	Marking and braking subsystems not ready in time	3	4	Design of small subsystems not yet started	Complete design and order parts by the end of Spring Break

11.5 Schedule

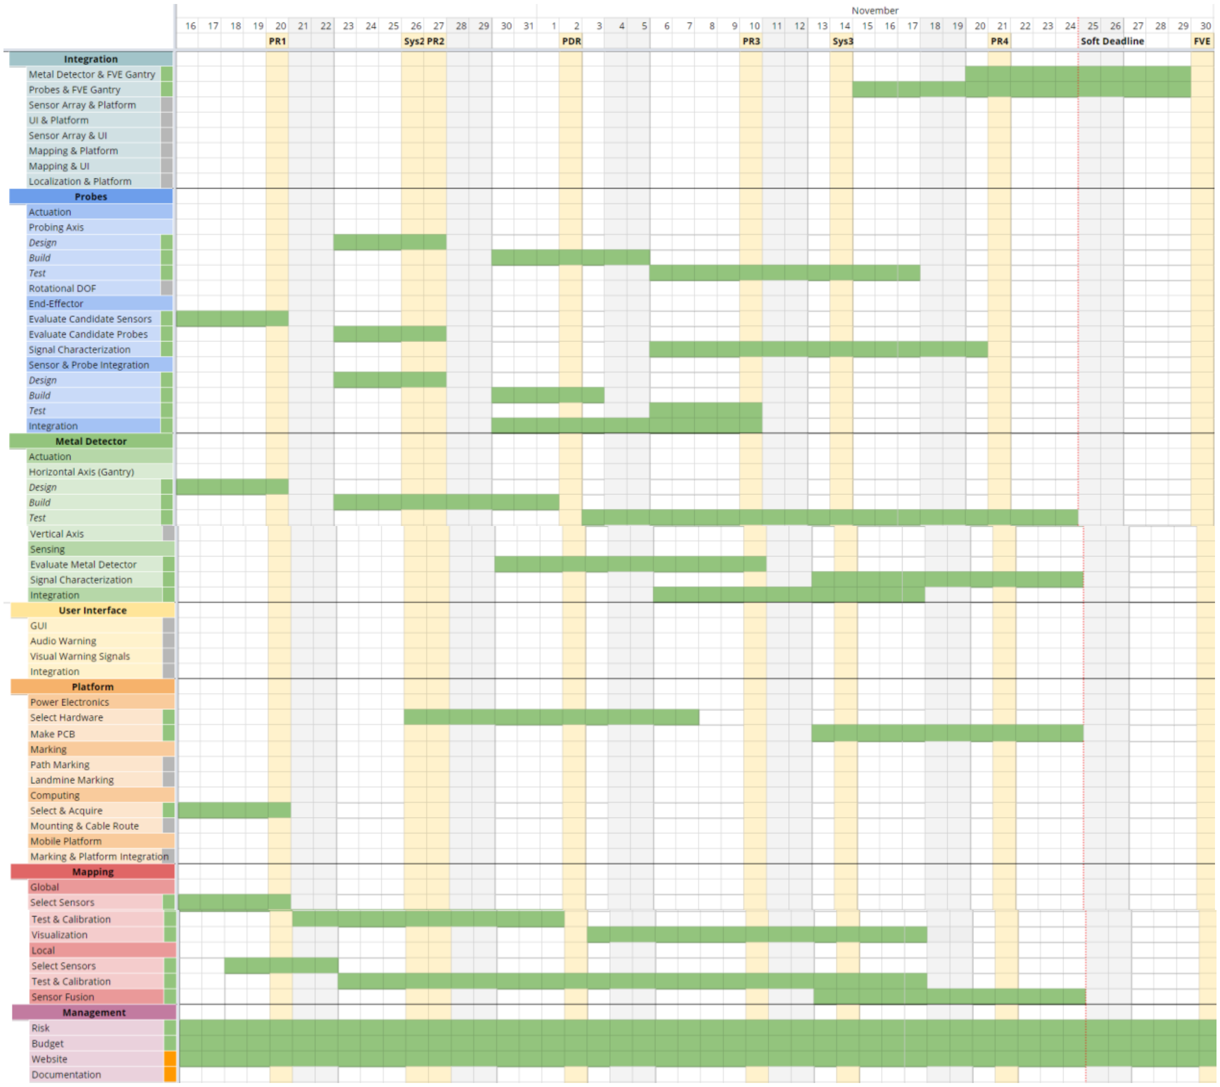

Figure 43: Fall Semester Schedule

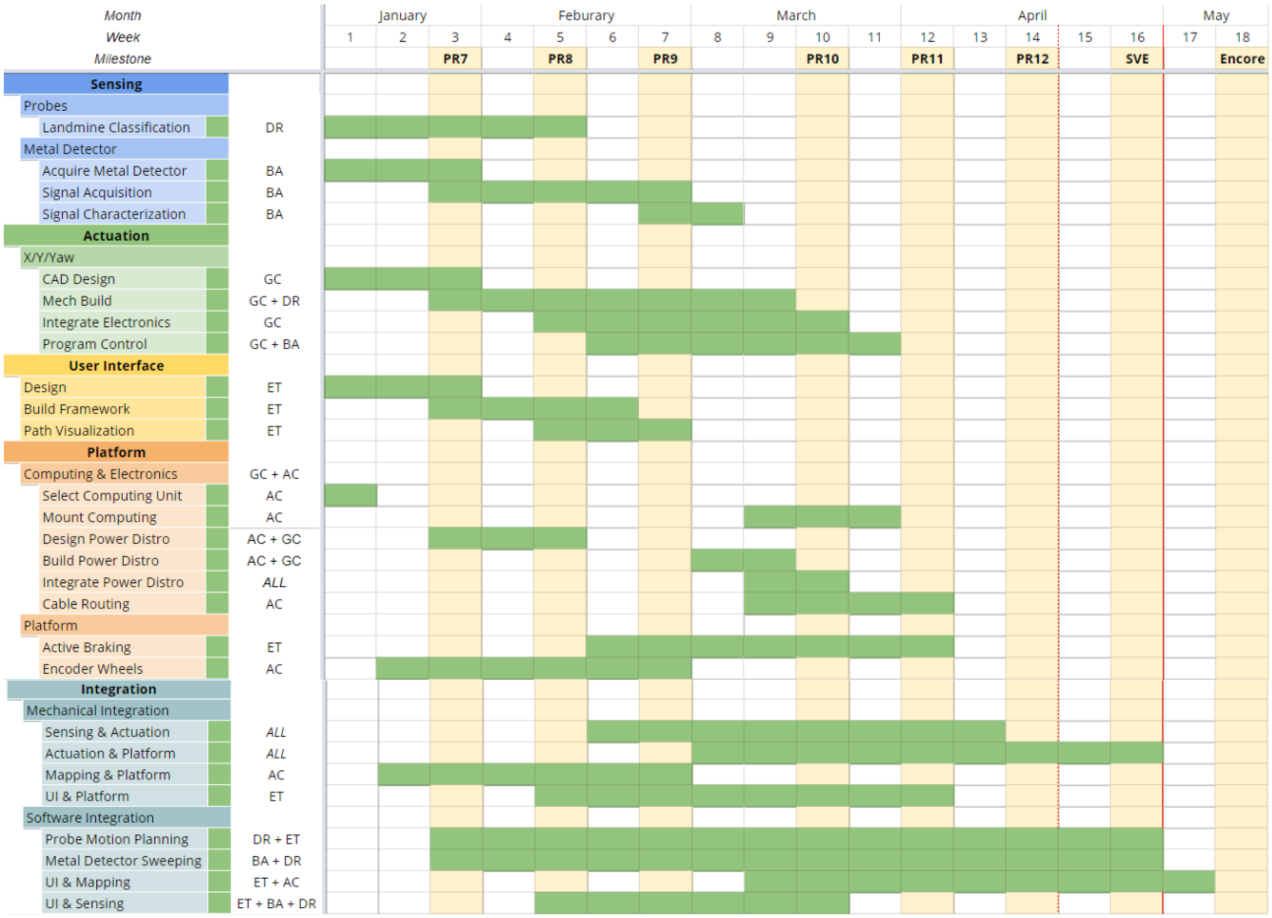

Figure 44: Spring Semester Schedule