Final Report

Carnegie Mellon University

The Fire Blighters

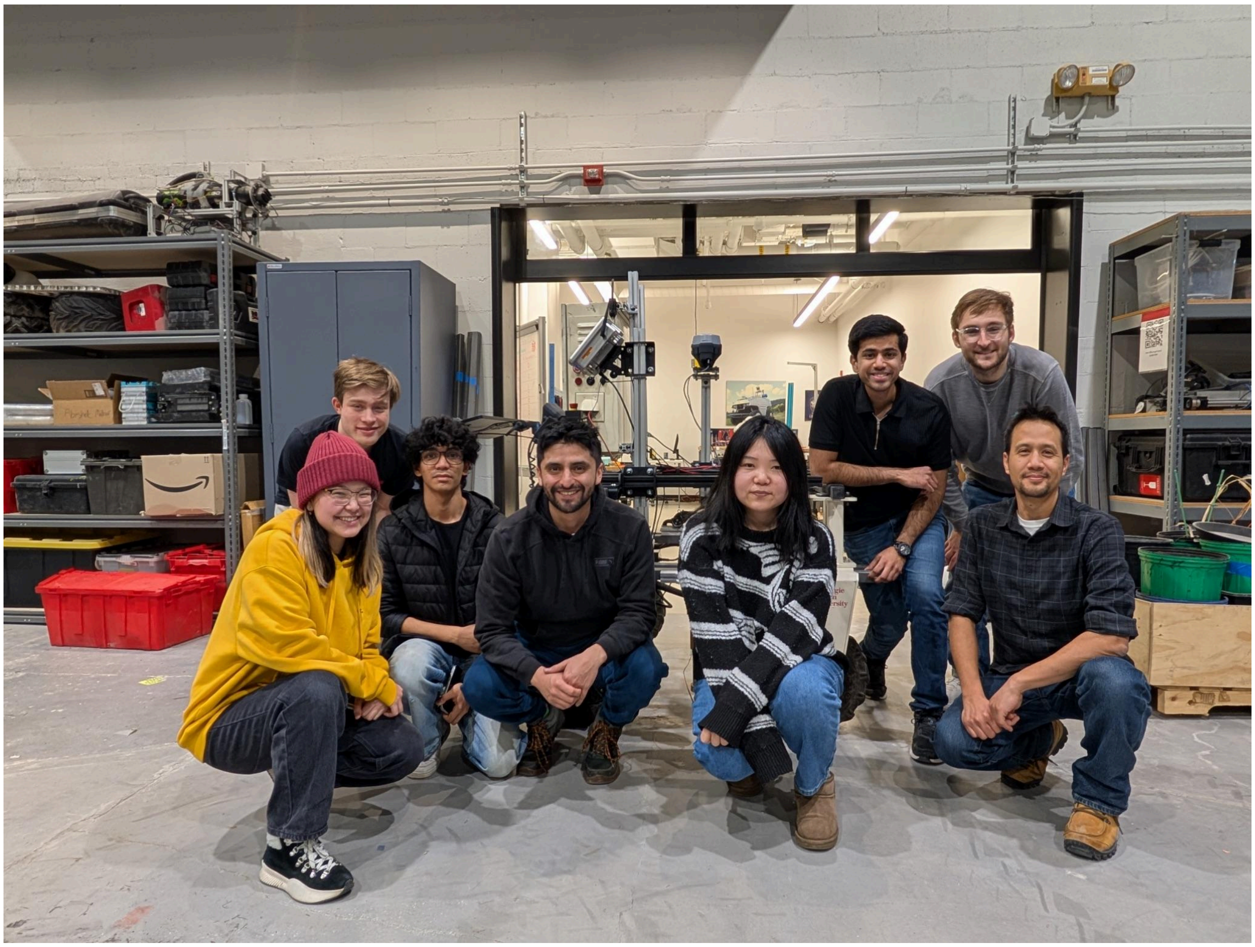

Competition Division: I

Team Advisors:

- Francisco Yandun, Senior Project Scientist - fyandun@andrew.cmu.edu
- Abhisesh Silwal, Systems Scientist - asilwal@andrew.cmu.edu

Team Members

Name	Email	Major	Current degree
Jack Nelson	jinelson@andrew.cmu.edu	Electrical and Computer Engineering	Undergraduate
Daniya Nussipbek	nussipbe@andrew.cmu.edu	Mechanical Engineering	Undergraduate
Sarthak Jain	sjain4@andrew.cmu.edu	Electrical and Computer Engineering	Undergraduate
Hayden Feddock	hfeddock@andrew.cmu.edu	Robotics	Graduate
Sandeep Sam Zachariah	sszachar@andrew.cmu.edu	Robotics	Graduate
Yi Wu	yiwu2@andrew.cmu.edu	Robotics	Graduate

1. Introduction

1.1. Summary

In this final report we present Erwin, a fully autonomous robotic inspection system equipped with custom stereo vision sensors for detecting and mapping fireblight in high density apple trees during dormant season.

Figure 1 Left. Erwin, a robotic inspection system that locates infected trees, marks them and creates a map of the ones affected with fireblight. Right. Common symptoms of the disease include cankers and bent branches (aka shepherd crooks). Fireblight is a bacterial disease that affects pome trees, causing millions of dollars in losses every year.

1.2. Ag Partner Overview

During the development of this project we had three Ag Partners:

- Dr. Srdjan Acimovic, plant pathologist from Virginia Tech. Dr. Acimovic helped us with domain expertise for labeling and training our computer vision model that detects the disease symptoms using a stereo camera.
- Mike Allridge, farm manager at Trax Farms. Mike showed us how fireblight affects his commercial orchards and provided key insights on current management practices and what would be the characteristics for a robot to be viable in practice.

- Dr. Kari Peter, plant pathologist from Penn State University. Dr. Peters gave us access to her research orchards located at Penn State Fruit Research and Extension Center, where we conducted all our field tests.

1.3. Problem Statement

Fire blight is caused by the bacterium *Erwinia amylovora*, a highly contagious pathogen that overwinters in bark lesions known as cankers and spreads rapidly via insects, rainwater, and wind; ultimately infecting blossoms and young shoots. Infected shoots develop a characteristic symptom called “shepherd’s crook,” while infected woody tissue leads to canker development and, in severe cases, rootstock blight that necessitates tree removal, causing significant losses. In the East coast of the United States, dormant-season pruning remains the most effective mitigation strategy, but it is costly, labor-intensive and must be performed at large scale. As labor availability declines and operational costs rise, growers often lack the capacity to perform routine scouting at the frequency and spatial resolution needed for timely intervention. Robotic monitoring systems offer a promising solution by enabling autonomous inspection and the construction of actionable maps of disease incidence. Such maps allow growers to precisely allocate resources---chemicals, labor, and equipment---while tracking disease trends across the orchard with higher temporal and spatial accuracy than manual scouting provides.

2. System Design

2.1. System Description

We divided our system into six subsystems: simulation, navigation, perception, manipulation, marking and user interface. All of the systems were orchestrated using mission control software, and were implemented in a ROS2 framework, which provides us with a set of libraries and tools to build the operating system of our robot.

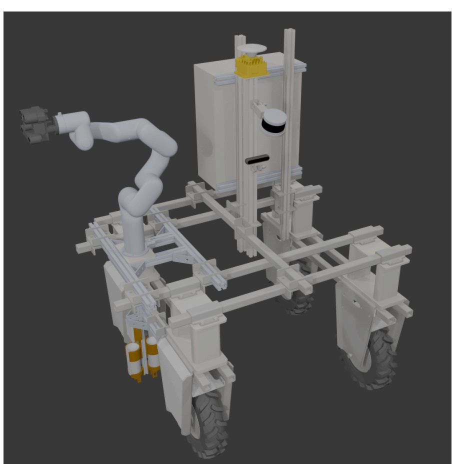

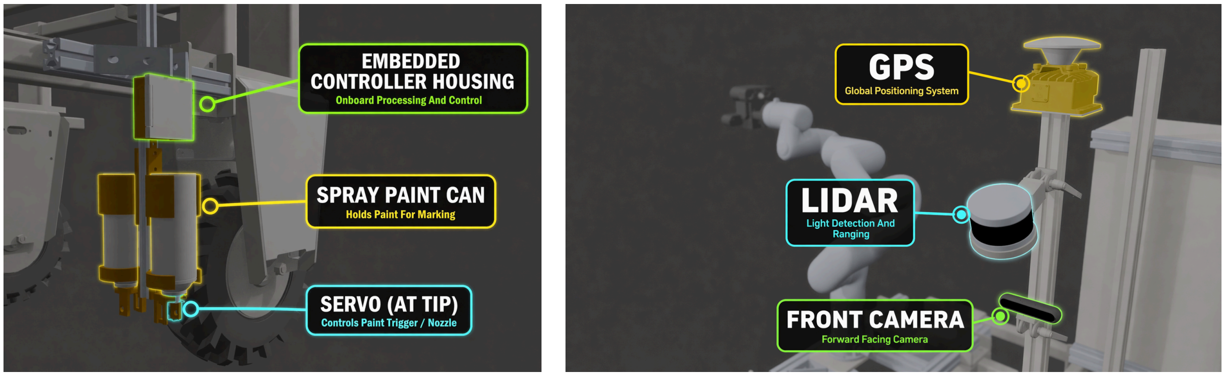

Fig 2. Erwin and its main components. We used the Amiga mobile base, a 6 Dof xArm manipulator and several sensors for navigation and for mapping fireblight. The marking tool uses a servo motor and a generic spray paint can.

2.1.1. Simulation

In order to test our autonomous navigation stack, we created a realistic looking simulation environment using the Isaac Sim from Nvidia. The simulator allowed us to develop and test the autonomous navigation pipeline that included GPS waypoint following, obstacle avoidance and safety stops. This component was led by Yi Wu.

Fig 3. Simulated path of Erwin in Isaac Sim, a realistic simulation environment for robotics developed by Nvidia.

2.1.2. Navigation

Our navigation subsystem consists of localization, path following and obstacle avoidance algorithms, relying on Real-Time Kinematic (RTK) GNSS technology for precise positioning. This subsystem takes two inputs: the path to follow within the orchard, and the locations of the trees, where the manipulation system will scan the trees looking for symptoms of fireblight. In the case a person stands in front of the robot, Erwin's front-facing camera will detect them and will stop its operation. Regular operation only resumes after the person is not detected anymore. This component was led by Yi Wu.

2.1.3. Perception

The perception pipeline transforms stereo camera images into a confidence-aware semantic 3D representation of the tree. This subsystem consists of four stages: (1) flash-illuminated stereo image acquisition, (2) instance-level disease detection, (3) real-time depth estimation, and (4) semantic point cloud construction. This modular structure allows perception components to be independently swapped while maintaining a consistent 3D representation for downstream planning. This component was led by Hayden Feddock.

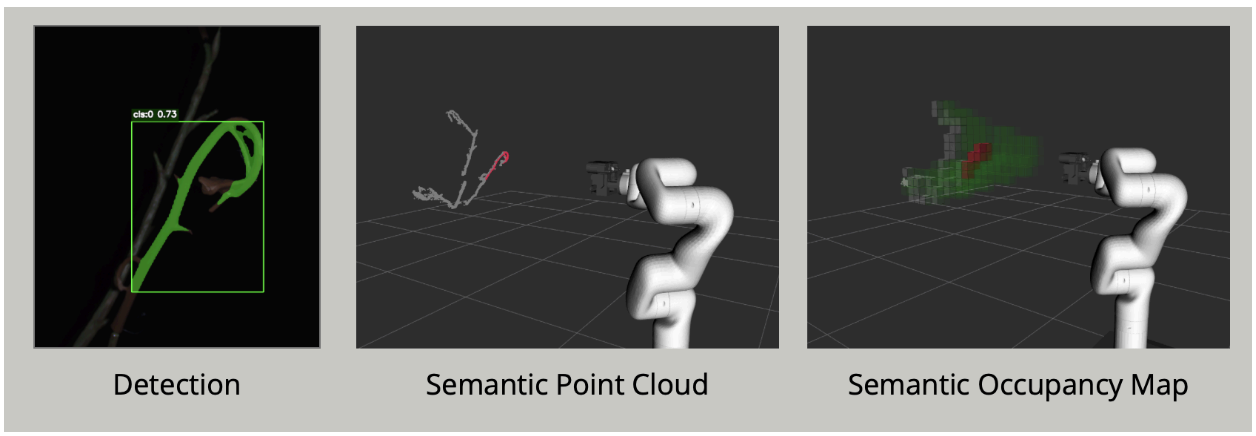

Fig 4. Perception pipeline. A neural network detects pixels corresponding to fireblight symptoms. Later, a stereo camera projects the detections to three dimensions (semantic point cloud). To operate in real time, we compress the 3D view of the tree to a 3D grid representation, which requires less memory (semantic occupancy map). Cubes are then colored red if either a canker or shepherd's crook was observed in the original image.

2.1.4. Manipulation

The manipulation strategy leverages the robotic arm to systematically capture images of the entire tree. Using a 3D grid representation generated by the perception pipeline, the onboard computer plans collision-free trajectories around the canopy, ensuring complete and efficient visual coverage. Our system tries to maintain a trade off between scan speed and accuracy. This component was led by Hayden Feddock.

2.1.5. Marking

Once Erwin detects several symptoms in a tree, it marks them by spraying the ground with a non-toxic paint. This feature was developed with the final user in mind. Out in the field, a farmer or manager does not need to worry about scouting the trees in detail, they just need to spot the marks on the ground to quickly identify and verify the ones infected with fireblight. This component was led by Daniya Nussipbek and Sarthak Jain.

Fig 5. Marking system. Erwin sprays non-toxic paint on the ground to mark infected trees.

2.1.6. User Interface

All the subsystems described before are orchestrated using a behaviour tree. This algorithm ensures the sequential execution of each one at the right time. It also handles errors and unexpected behaviors. At each point of the operation, we also provide a visual interface so the user can have visual feedback of how everything is operating. This includes raw images from the cameras and other sensors, the robot's telemetry and more importantly a geo-referenced map of the location of the robot and the detections of the fireblight. The systems integration was led by Sandeep Zachariah and the UI development was led by Jack Nelson.

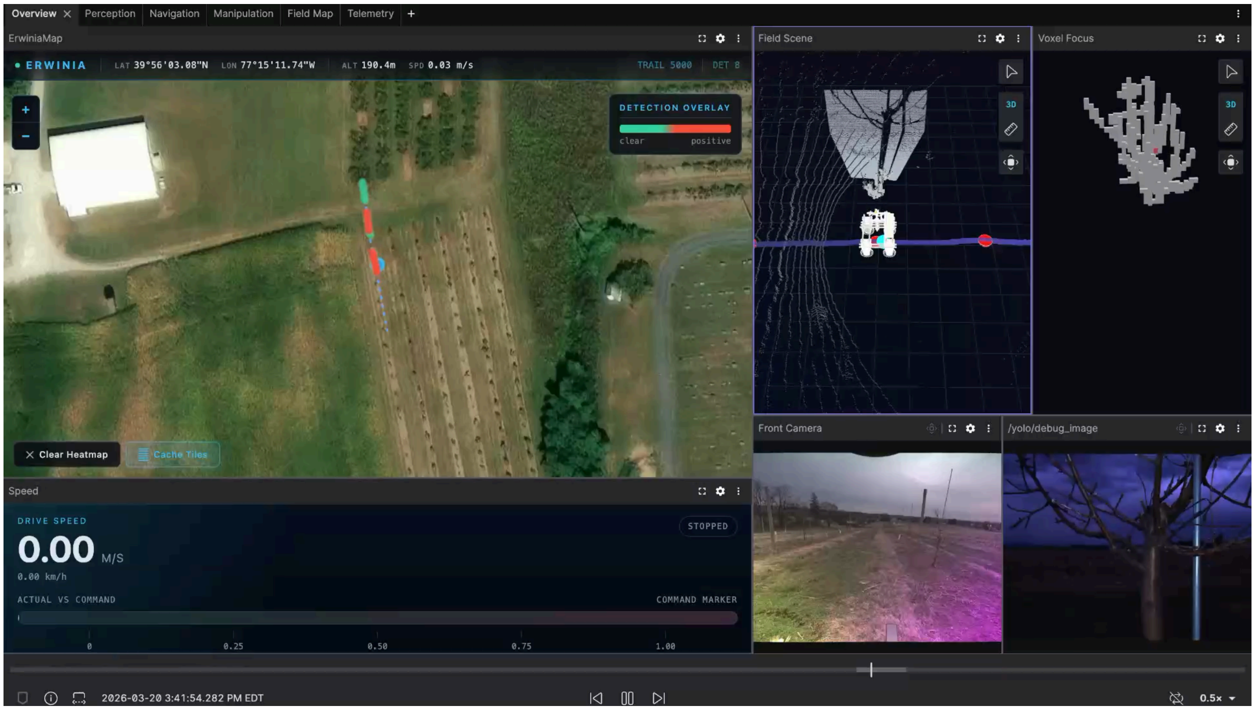

Figure 6. User Interface. The main interface displays a geo-referenced map of the orchard, where infected trees are highlighted in red and overlaid on a satellite basemap. The UI also includes the robot's trajectory and stopping points, providing context on coverage and inspection locations. In addition, a 3D reconstruction of each scanned tree is shown, with detected symptoms clearly marked in red. The interface further integrates live feeds from both the navigation and flash cameras, enabling users to inspect raw imagery alongside the processed outputs.

2.2. Operation

Erwin's operation is relatively simple. The user needs to provide the path the robot is going to follow and the location of the trees to be scanned. This can be achieved by different ways but our standard deployment includes the following steps:

1. First, a drone flight is conducted to generate a high-resolution (centimeter-level) orthomosaic map of the orchard in TIF format, compatible with standard GIS tools.
2. Next, using GIS software, the user defines the path for Erwin and selects the trees to be scanned.
3. This plan is then uploaded to the robot's onboard computer, enabling the system to begin autonomous operation.
4. During execution, the user can monitor progress through the user interface developed in Foxglove. The interface provides real-time feedback, including raw camera feeds, detected disease symptoms, robot telemetry, and—most importantly—a georeferenced map highlighting the locations of infected trees (see Figure 6).
5. In addition to digital outputs, Erwin physically marks the ground near detected infected trees, enabling quick in situ verification.

6. The final georeferenced map serves as the primary output of each operation, offering growers and managers a practical tool to assess disease distribution and plan targeted mitigation strategies.

2.3. Performance Factors

There are several factors that affect Erwin's performance. In terms of navigation, sloped terrains make the Amiga mobile base to slip and keep track of the path. Another condition that affects the performance of the navigation system is the conditions of the ground. Wet soil can reduce the traction of the wheels and cause excessive slip which causes the robot to lose track of the desired path. Our approach for mitigating these issues was to test in flat dry conditions. In the future we plan to implement a better control system to minimize the impact of these issues in the robot.

For the computer vision model and imaging system, wind is the primary factor affecting performance. The system assumes a static scene while scanning the tree, detecting disease symptoms, and reconstructing the 3D model. When wind causes the tree to move at any time during this process, this assumption is violated, leading to misalignment and reduced accuracy in the scan pattern. The mitigation measure in this case is to make sure the robot operates when wind speeds are under a specific threshold.

Another factor affecting Erwin—though to a lesser extent—is lighting. Our imaging system incorporates a flash to reduce sensitivity to typical ambient illumination changes. However, under extreme lighting variations, the performance of the disease detection model can degrade. We plan to mitigate this issue by collecting more training images that cover the extreme cases.

2.4. Safety Features

To reduce the risks that Erwin, as any other autonomous machine can have,, we included several features that will halt the operation of the mobile base or the robotic arm. Table I summarizes the risk and the safety features we included in our system.

RISK	LIKELIHOOD (High, Moderate, Low)	IMPACT (High, Moderate, Low)	MITIGATION STEPS Safety Features
People in front of the robot while navigating	**Moderate**	**High**	➔ Use the camera to detect people and stop the movement. Only resume if the person is not detected anymore.
Navigation controller fails	**Low**	**High**	➔ Use a hierarchical control scheme. Any signal of the autonomy

			controller will be overridden if a joystick (controlled by a human) is given. ➔ Use the Amiga e-stop
Navigation planner fails	**Low**	**Medium**	➔ The controller will only move the robot if a valid planting plan is given.
Free rolling on a slope	**Moderate**	**High**	➔ We included a hand brake that will prevent free rolling when the motors are unpowered.
Loss of GPS signal	**Moderate**	**Moderate**	➔ Do not operate until the GPS signal is recovered.
Robotic manipulator operation fails	**Moderate**	**High**	➔ Limit arm workspace. ➔ Make a prior model of the tree so the arm does not collide with it. ➔ An e-stop button will halt the operation of the arm.
Operating an autonomous robot	**Low**	**Low**	➔ Proper training to students and staff

2.5. Cost

Item	Cost Market	Cost for our team	Notes
Amiga	$12,990.00	$0.00	The robot is part of the lab
Onboard Computer	$5,000.00	$0.00	The computer was already mounted on the robot
Robotic Manipulator	$9,000.00	$0.00	The robot is part of the lab
Detection camera system	$800.00	$0.00	This sensor was part of the lab
Stereo Camera (front facing)	$300.00	$0.00	This sensor was part of the lab

LiDAR	$4,600.00	$0.00	This sensor was part of the lab
RTK GPS	$3,000.00	$0.00	This sensor was part of the lab
Misc components	$500.00	$500.00	
	$36,190.00	**$500.00**	

3. Design Evaluation

We evaluated three key aspects of Erwin's operation: individual symptom detection, tree marking accuracy, and operational efficiency.

The first focuses on the segmentation and localization of disease symptoms using the flash stereo camera. To assess this, we employed standard precision and recall metrics to evaluate the performance of the YOLOv26-seg model. The model achieved a **precision** of 0.872 and a **recall** of 0.871], which we considered suitable for field deployment.

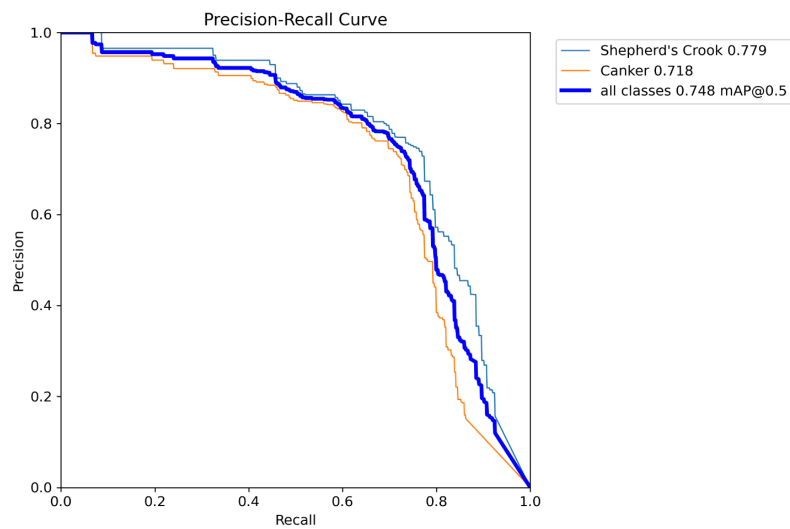

Figure 7. Precision-recall curve for the computer vision model used for segmentation of fireblight symptoms.

The second metric evaluates the field performance of the marking operation. During field tests, we scanned 10 trees, of which 7 were infected with fire blight. The robot successfully navigated to each tree, stopped, and marked 6 of the 7 infected trees. In the missed case, the camera system failed to detect visible symptoms, and the tree was therefore not marked. Overall, this resulted in an **accuracy** of 85.71% for infected trees (out of 10 infected trees, 9 were correctly marked) and 100% **specificity** for healthy trees, as none of the uninfected trees were incorrectly marked.

In terms of operational speed, Erwin currently takes approximately 3-4 minutes to scan a single tree, depending on canopy size and structure. In our field tests, the system was able to scan, localize, and map 10 trees in about 35 minutes, which includes perception and mapping overhead. While this timing is modest in its current prototype stage, it is important to note that Erwin is designed for continuous operation, enabling day-and-night deployment effectively increasing total daily coverage compared to manual scouting. Additionally, we expect substantial improvements in speed and efficiency in future iterations through better motion planning, sensing optimization, and hardware integration.

4. Design History

The challenge of fire blight and the need for more effective management strategies were introduced to our team through our agricultural partner, Srdjan Acimovic. As a plant pathologist whose research focuses on fire blight, Prof. Acimovic expressed interest in using technology to generate geo-referenced maps of disease incidence. After evaluating existing approaches—such as aerial monitoring with drones and traditional manual scouting—Dr. Yandun proposed an alternative solution based on a mobile ground robot equipped with a vision system capable of detecting symptoms, generating maps, and marking infected trees.

Building on this concept, we developed an initial prototype using a Husky mobile base paired with an xArm6 manipulator. Early experiments demonstrated that the robotic arm enabled a robust and precise scanning strategy; however, they also revealed limitations in coverage and platform robustness for real-world deployment. To better understand practical requirements, we conducted field visits and engaged directly with growers, presenting the concept of a robotic system for detecting and mapping infections. These interactions were well received and provided valuable feedback that informed key design considerations for subsequent iterations of the system.

In fact, one of the main challenges of the project was obtaining access to test sites with the right conditions for deployment (i.e., having infected trees with visible symptoms). We conducted initial visits to Soergel and Trax Farms in Pennsylvania; however, these orchards were well managed, and little to no fire blight infection was present, limiting their usefulness for validation. We finally established a collaboration with Dr. Kari Peters at the Penn State Fruit Research and Extension Center, who maintains dedicated research plots with actively infected trees. This partnership was critical, as it provided consistent access to representative disease conditions necessary for testing and evaluating Erwin in real conditions.

5. Impacts

5.1. Summary

This project presents Erwin, an autonomous mobile robot designed to detect, mark, and generate geo-referenced maps of fire blight incidence in apple orchards. Field evaluations demonstrate that the system performs reliably as an initial prototype and proof of concept under real-world conditions. The key strengths of Erwin lie in its ability to accurately detect disease symptoms in the outdoors, coupled with a user interface that provides real-time feedback during operation. In particular, the generation of geo-referenced maps highlighting infected trees represents a powerful decision-support tool, enabling growers to plan targeted mitigation strategies and respond in a timely manner.

Feedback from our agricultural partners has been positive regarding Erwin's practical potential; however, they emphasized that improving scanning speed will be critical for adoption at scale. Addressing this limitation will be a central focus of our future development, as increased throughput will directly enhance the system's operational viability in commercial orchard settings.

5.2. Commercial Potential

Fire blight management in apple orchards relies heavily on frequent scouting and timely intervention, yet current approaches are both labor-intensive and economically inefficient. Preventive measures such as antibiotic or biological sprays typically cost on the order of $50–150 per acre annually, with additional labor required for monitoring bloom stages and assessing infection risk University of Maine Cooperative Extension, spray cost estimates). However, the true cost burden emerges during outbreaks, when missed or late detections allow the disease to spread rapidly. In these situations, scouting itself becomes significantly more expensive, as growers must conduct repeated, high-intensity inspections to identify new infections, often alongside extensive pruning and sanitation efforts. These combined activities can add $25,000–$75,000 in labor costs per season at the farm level, driven in large part by the need for continuous manual scouting to track disease progression (Cornell Cooperative Extension, fire blight management reports). Despite this investment, human scouting remains inconsistent and constrained by time and labor availability, increasing the risk of undetected infections.

These issues create a suitable market opportunity where Erwin can be an attractive solution. We see our robot as a **prototype** that addresses these challenges by enabling autonomous, repeatable field monitoring and generating actionable maps to support targeted interventions. Given our development and field tests have been apple orchards, our primary market can be the apple growing industry in the U.S. East Coast, where infected trees are pruned during dormant season.

We acknowledge that Erwin's current cost may be prohibitive for many growers; however, we view it as a prototype whose construction, form factor, and components can be further optimized to reduce cost and improve scalability. In addition, we envision a service-based deployment model, where Erwin—or a more mature version of the system—can be offered as an on-demand service, allowing growers to seasonally map their orchards without the need for large upfront investments. This approach lowers the barrier to adoption while preserving the benefits of high-frequency, data-driven disease monitoring.

In the future, we envision a venture where we can be a successful provider of robotic sensing solutions for specialty crop agriculture, with systems deployed across several orchard types and providing growers with continuous, high-resolution insight into their crop health. The robot will consist of a modular perception and mapping system that can integrate with existing or emerging agricultural robotic platforms, utilizing crop-specific detection models that can be deployed per-application.

5.3. Social and Environmental

The deployment of robotic systems like Erwin for detecting and mapping fire blight has important societal and environmental implications. By enabling early and precise identification of infections, growers can be supported for applying targeted interventions, reducing the need for blanket chemical applications and therefore lowering environmental impact on soil, water, and surrounding ecosystems. More efficient disease management can also help to preserve orchard productivity, contributing to food security and economic stability for growers, particularly in regions where specialty crops like apples are a major source of income.

From a societal perspective, automation can help address ongoing agricultural labor shortages and reduce the physical burden of repetitive scouting tasks, while allowing skilled labor to focus on higher-value decision-making. Additionally, the generation of high-resolution, geo-referenced disease maps promotes more data-driven and sustainable farming practices, aligning with broader goals of precision agriculture and responsible resource management.