Towards an AI-Driven Cyber Physical System for Closed-Loop Control of Plant Diseases

Abhisesh Silwal[1], Xuemei M. Zhang[2], Thomas Hadlock[3], Jacob Neice[3], Shadab Haque[4], Adwait Kaundanya[4], Chang Lu[3], Boris A. Vinatzer[2], George Kantor[1], Song Li[2]

[1]Robotic Institute, Carnegie Mellon University, Pittsburg, PA, 15213
[2]School of Plant and Environmental Sciences
[3]Department of Chemical Engineering
[4]Department of Electronic and Computer Engineering
Virginia Tech, Blacksburg, VA 24061 USA
asilwal@andrew.cmu.edu, songli@vt.edu

Abstract

Plant diseases are a major biosecurity threat to food production and the bio-energy industry. Early detection and control of plant diseases can improve producers' profitability and reduce environmental impacts from chemical inputs. We proposed to develop a cyber-physical system with three major components: an AI-driven imaging system for early stress detection, an autonomous robotic system to collect plant samples, and a sequencing pipeline to detect molecular signatures of pathogens for disease confirmation. This system is envisioned to control a detected disease by removing or pruning infected plants. This manuscript describes the major milestones achieved by this CPS project and provides a future perspective on disease control automation in agriculture.

Introduction

Plant diseases pose an increasing threat to the nation's food supply and biosecurity in a changing environment. Recent failures to prevent plant disease emergence and spread in the United States have resulted in major economic losses (20-40% of total yield) for growers (Savary et al. 2019, 2012).

Conventionally, detecting infected plants requires inspecting plants manually, culturing bacteria and fungi from infected materials, and confirming infective agents using molecular biology approaches, such as polymerase chain reaction (PCR) followed by sequencing. These procedures are time-consuming and, thus, are often too slow to achieve real-time disease control in production. In recent years, the technical advancements in the application of AI in computer vision, robotics, and genomics have been changing the landscape of crop protection in agriculture. For example, machine learning methods have been developed to detect unhealthy plants from images taken by aerial or ground-based robotic platforms (Pilli et al. 2015; Chin, Catal, and Kassahun 2023) or to detect pathogens from meta-genomic samples (Johnson, Vinatzer, and Li 2023; Espindola et al. 2015).

One central challenge for preventing accidental pathogen dissemination and disease outbreaks is that many plant diseases are difficult to detect early (Trippa et al. 2024), and infected (possibly asymptomatic) plants can spread pathogens undetected when being shipped from one location to another. Once an emerging disease takes a foothold in a new environment, eradicating such disease is extremely challenging. For emerging pathogenic species, the process of detecting these emerging threats usually takes months, and the methods for developing novel detection assays, such as strain-specific antibodies or PCR assay, take even longer, thus giving the pathogen time to spread further and establish in additional areas. To solve all these challenges, we have been developing a closed-loop control system including three main components (Figure 1): (1) detecting early disease symptoms with advanced imaging technologies and AI, (2) sampling plant tissues automatically with robotics, and (3) analyzing plant metagenomes with a microfluidic device for library preparation followed by nanopore sequencing and classifying meta-genome samples with machine learning. These steps ensure accurate and rapid detection of a plant disease, which can guide plant pruning to remove infected individuals from a plant production system and prevent disease from spreading. Our system is designed to focus on greenhouse environments with high-density plant production of horticulture crops such as tomatoes.

Methods and Results

Detecting Plant Stress Phenotypes With Hyper-Spectral Imaging and Foundation AI Model

To achieve early stress detection, we performed two sets of experiments (biotic and abiotic stresses) using tomato seedlings as our model system (Figure 1.1). For biotic stress, we used a bacterial leaf spot pathogen, Xanthomonas perforans, to inoculate healthy plants and collected hyperspectral images of tomato leaves. For abiotic stress, we performed drought treatment of tomato seedlings and collected side-view images with a Canon DSLR. In both experiments, the goals were to train machine learning models to detect signatures of plants under stress. In the case of biotic stress, hyperspectral images were collected, and machine learning models were trained using individual leaf and whole plant images. In the case of abiotic stress, the main goal is to automatically detect wilting phenotyping, which is a common symptom of plants infected by pathogens. We tested advanced foundation models, including the Segment Any-

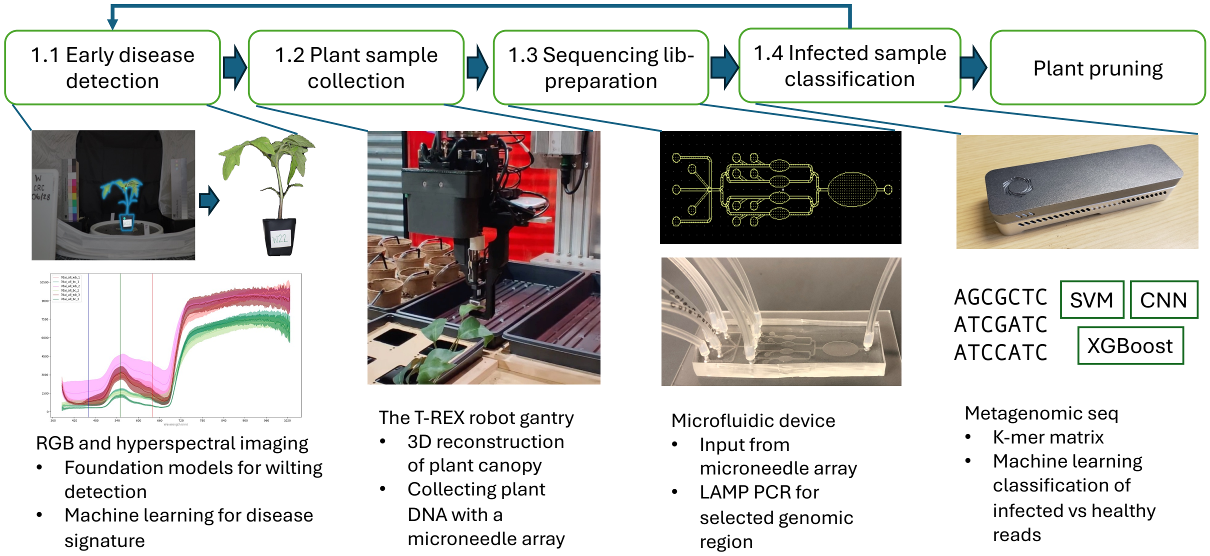

Figure 1: A CPS for early plant disease detection and control. (1) image analysis with AI to detect stressed plant phenotype. (2) T-REX robotic platform for plant 3D modeling and automatic leaf collection. (3) microfluidic sample preparation, and (4) machine learning for classifying infected plants using nanopore reads.

thing (Kirillov et al. 2023) and the Grounding DINO (Liu et al. 2023) models, to automatically recognize the plants in the images and measure plant heights to detect wilted plants.

For the biotic stress and hyperspectral analysis, details of the hyperspectral analysis and machine learning model are described elsewhere (Zhang, Vinatzer, and Li 2023). In brief, we collected spectral imaging data (with PIKA-L, 400-1000 nm and PIKL-NIR, 900-1700 nm cameras from Resonon Inc.) from healthy and inoculated leaves. Seven machine learning and non-ML models were tested to classify spectral signatures at the whole leaf and the pixel levels. We found that linear discriminative analysis best separates infected leaves from healthy leaves and achieved the highest F1 score of 0.76. However, we found that leaves at the top of the canopy showed different spectral reflectance from the leaves at the lower canopy. A model that understands the 3D structure of plants will be helpful to interpret results from hyperspectral data collected from a whole plant.

For the abiotic stress analysis, we used water-stressed plants as our model system. The GroundingDINO model was used to recognize the tomato plant using a text prompt (Figure 2). The pixels for a whole plant were segmented using SAM followed by a simple script to measure the plant height. We found a significant shift in the height distribution between healthy and stressed plants.

Automating Plant Sample Collection With the T-REX Robot System

The goal of the T-REX (The Robot for extracting leaf samples) is to autonomously scout plants in a greenhouse and use imaging-based techniques to sense and extract leaf samples. We have completed the development of a custom gantry robot system to fully automate the laborious task of acquiring temporal images of plants, completed the design

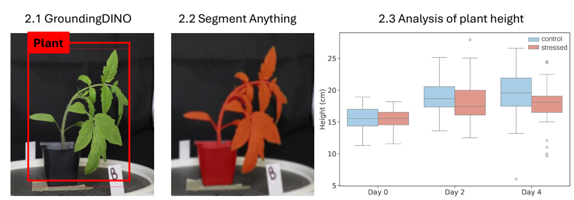

Figure 2: Image analysis pipeline to detect stressed plants. (2.1) Grounding DINO was used with text prompt "plant" to detect an individual plant. (2.2) The Segment Anything model was used to segment plants from the background. (2.3) Boxplot shows a significant reduction of plant height at day 4.

and development of a novel end-effector to extract leaf samples physically, and the AI-based vision algorithm to understand plant semantics to guide the robot to extract leaf samples at the most optimal locations (Figure 1.2).

Plants grown in a greenhouse environment with modern techniques such as hydroponic systems usually have flat beds as a base for plants to grow, and growth lights are usually hung from above. The design of the robot system also emulates a similar setting with a workspace of 3-meter by 1.5-meter flat test bed with timed growth lights. This gantry-like robot is an open-look kinematic chain robot arm with three prismatic and three revolute joints (3P3R) to achieve six degrees of freedom pose. The T-REX is equipped with a custom-built end-effector and an active light camera to image the plants (Silwal et al. 2021). The robot uses digital servo motors as actuators and runs on ROS2 (Macenski et al. 2022). Currently, the robot traverses to each known location of the plants and acquires images of plants using six pre-

Figure 3: Top view image of the plants with grasp point per leaf (left), SDF of the workspace highlighting free vs occupied (center), optimal leaf and grasp point (right).

programmed poses that later are used for 3D modeling and image-based AI analytics.

Novel end-effector: The design of the end-effector takes into account that plants/leaves are deformable, and highly prone to damage. To minimize damage to surrounding leaves as well as to unintentionally push the plants away, the motion planner approaches the plant from the top and tries to grasp the leaf by clamping the leaf from its top and bottom surfaces. The tip of the end-effector houses a microneedle array that is pushed against the leaf surface to destructively extract DNA samples.

Microneedle dispensing mechanism: The section of the end-effector that houses the microneedle array is a thin rectangular plastic to which the microneedle arrays are glued. The section where the microneedle array sticks has laser-etched tabs that break away when an external force is applied (Figure 1.2). The dispensing system consists of several rectangular microneedle housings and automatically reloads the tip of the end-effector for end-to-end sampling without human intervention.

Novel leaf grasping algorithm: To guide the robot to physically grasp a leaf requires a semantic understanding of the plant's location and its parts, such as leaves, with respect to adjacent plants in the robot's workspace. To enable this capability, we leverage both 2D spatial and 3D information. The 2D information contains semantic segmentation of instances of leaves. Using deep learning algorithms (https://github.com/ultralytics/ultralytics) and the 3D information is extracted from the merged point cloud data using multi-view stereo registration (Choi, Zhou, and Koltun 2015). The pipeline first uses instance segmentation and 3D information to generate proposals of grasping locations for each individual leaf using leaf centrality and flatness as the decision variables. This part of the pipeline then automatically generates large training samples for a heatmap-based regression model to identify optimal grasping locations (Figure 3). However, as plants have indeterminate growth, some of the leaves are very hard to reach. The next step includes selecting the most "easy" leaf to grasp to increase the chances of successful sampling. This capability uses signed distance field (SDF) to represent robot's workspace (free vs occupied space) and pareto front to select leaves that are maximally away from the surrounding leaves and minimally away from free un-occluded regions of the workspace.

Microfluidic-Based Nanopore Sample Preparation and Reference-Free Machine Learning-Based Disease Detection

In this project, a microfluidic device was designed and fabricated using basic modules of our existing library platform (Murphy et al. 2020), with modifications to enable nanopore library preparation and automation of the process using a portable liquid handling system (Figure 1.3). We have successfully tested this device, which can produce multi-region LAMP amplified libraries from low sample inputs generated by microneedle extraction. Significant protocol optimization was performed to adapt this device to low-input field-extracted samples.

Another major component of this project is to develop machine learning algorithms that can classify infected plant samples from healthy samples. The details of this machine learning approach were published in (Johnson, Vinatzer, and Li 2023). In brief, the machine learning model was developed to classify nanopore reads sequenced from infected and healthy leaf samples. Multiple machine learning models were trained to classify reads based on the k-mer frequency. The model was tested for two plant diseases, including the tomato leaf spot disease. Several machine learning methods achieved higher than 0.90 accuracy in testing data. A convolutional neural network (CNN) model performed best in accuracy, while a random forest model achieved the fastest speed. Another key aspect of this model is that the model was evaluated using only 5000 reads per sample. A typical Nanopore sequencing run can generate 4 to 10 million reads, and the cost is approximately $2000, resulting in $1-2.5 per 5000 reads, supporting a low-cost way to detect plant pathogens using nanopore sequencing.

Discussion and Future Directions

The work in this manuscript lays the foundation for further developing closed-loop disease control to reduce pathogen spread within densely cultivated greenhouses. We envision a complete system including a robotic platform that will constantly scout large greenhouses and sample plants that may already be infected. Samples will be subject to metagenomic sequencing for automatic pathogen identification. To close the loop, the results from image-based and sequencing-based disease detection will be combined with temperature and humidity data from environmental sensors to guide the implementation of control strategies. In summary, the next version of this CPS would combine real-time risk estimation during disease outbreaks with targeted preventative strategies to mitigate and possibly stop the spread of pathogens in greenhouses. In the long run, this system can be expanded to trace pathogens through the supply chain by barcoding each plant and comparing pathogen signatures in the field or the market to those from production greenhouses.

To bring this vision into reality, our current progress has addressed several fundamental challenges in the CPS science, which include: (1) The T-REX robot uses advanced machine learning algorithms to understand plant architecture and to control the movement of a robotic manipulator to achieve precise robot-plant interaction. (2) In-field metagenomic sampling and microfluidics have been developed to include semi-automated DNA library preparation, sequencing, and data analytics using a portable nanopore sequencing device and machine learning. (3) Real-time sensing with machine learning-powered hyperspectral image analysis and foundation AI models enables rapid phenotyping and presymptomatic detection of stressed plants. However, integrating these methods in a closed-loop control system with environmental sensing capability and disease models in outdoor agriculture is more challenging due to constantly changing environmental conditions.

Acknowledgments

This project is supported by a USDA CPS grant 2021-67021-34037 to AS and GK at CMU, and SL, BV and CL at Virginia Tech. The metagenomic analysis and machine learning is supported USDA FACT-CIN 2021-67021-34343 to SL, and BV labs at Virginia Tech.

References

Chin, R.; Catal, C.; and Kassahun, A. 2023. Plant disease detection using drones in precision agriculture. *Precision Agriculture*, 24(5): 1663–1682.

Choi, S.; Zhou, Q.-Y.; and Koltun, V. 2015. Robust reconstruction of indoor scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5556–5565.

Espindola, A.; Schneider, W.; Hoyt, P. R.; Marek, S. M.; and Garzon, C. 2015. A new approach for detecting fungal and oomycete plant pathogens in next generation sequencing metagenome data utilising electronic probes. *International journal of data mining and bioinformatics*, 12(2): 115–128.

Johnson, M. A.; Vinatzer, B. A.; and Li, S. 2023. Reference-free plant disease detection using machine learning and long-read metagenomic sequencing. *Applied and Environmental Microbiology*, 89(6): e00260–23.

Kirillov, A.; Mintun, E.; Ravi, N.; Mao, H.; Rolland, C.; Gustafson, L.; Xiao, T.; Whitehead, S.; Berg, A. C.; Lo, W.-Y.; et al. 2023. Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4015–4026.

Liu, S.; Zeng, Z.; Ren, T.; Li, F.; Zhang, H.; Yang, J.; Li, C.; Yang, J.; Su, H.; Zhu, J.; et al. 2023. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. *arXiv preprint arXiv:2303.05499*.

Macenski, S.; Foote, T.; Gerkey, B.; Lalancette, C.; and Woodall, W. 2022. Robot operating system 2: Design, architecture, and uses in the wild. *Science robotics*, 7(66): eabm6074.

Murphy, T. W.; Hsieh, Y.-P.; Zhu, B.; Naler, L. B.; and Lu, C. 2020. Microfluidic platform for Next-Generation Sequencing library preparation with low-input samples. *Analytical chemistry*, 92(3): 2519–2526.

Pilli, S. K.; Nallathambi, B.; George, S. J.; and Diwanji, V. 2015. eAGROBOT—A robot for early crop disease detection using image processing. In *2015 2nd international conference on electronics and communication systems (ICECS)*, 1684–1689. IEEE.

Savary, S.; Ficke, A.; Aubertot, J.-N.; and Hollier, C. 2012. Crop losses due to diseases and their implications for global food production losses and food security. *Food security*, 4(4): 519–537.

Savary, S.; Willocquet, L.; Pethybridge, S. J.; Esker, P.; McRoberts, N.; and Nelson, A. 2019. The global burden of pathogens and pests on major food crops. *Nature ecology & evolution*, 3(3): 430–439.

Silwal, A.; Parhar, T.; Yandun, F.; Baweja, H.; and Kantor, G. 2021. A robust illumination-invariant camera system for agricultural applications. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 3292–3298. IEEE.

Trippa, D.; Scalenghe, R.; Basso, M. F.; Panno, S.; Davino, S.; Morone, C.; Giovino, A.; Oufensou, S.; Luchi, N.; Yousefi, S.; et al. 2024. Next-generation methods for early disease detection in crops. *Pest management science*, 80(2): 245–261.

Zhang, X.; Vinatzer, B. A.; and Li, S. 2023. Hyperspectral Imaging Analysis for the Early Detection of Tomato Bacterial Leaf Spot Disease. *Research Square Preprint*.