Machine Learning Techniques for the Detection of Powdery Mildew in Vineyards

Michael Acosta

Institution: California State Polytechnic University, Pomona

Date: Fall 2024

Committee Members: Dr. Amar Raheja, Dr. Subodh Bhandari, Dr. Yu Sun

Presentation Outline

1. Introduction
2. Methodology
3. Results
4. Conclusion

1. Introduction

Literature Survey Highlights

The survey breaks down as follows:

1. Powdery Mildew (PM)
2. Current Best Practices
3. Handheld Proximal Sensors
4. Drones and UAVs
5. Security Challenges

Powdery Mildew (PM)

- PM is caused by the Erysiphe necator fungus [1], affecting 7,600 plant species [2]; viticulture uses 19.5 kg/ha of fungicides yearly, contributing to pesticide-resistant strains [3].
- PM thrives in mild temperatures (ideal at 77°F)[2] and is challenging to detect early, especially under shaded canopies where growth is fostered.
- PM infection reduces fruit yield, causes berry splitting, and affects photosynthesis, sugar content, juice color, and wine quality [4].
- Current practices include disease-resistant grape species and pesticide use, but widespread application increases costs, resistance, and pollution [5].

Current Best Practices

- Traditional on-site visits and visual inspections are the primary methods for identifying pests in crops and vineyards but are labor-intensive and prone to error.
- Manual inspections face challenges such as bias, optical illusions, high labor costs, and inaccuracies leading to financial loss [6].
- Handheld sensors and lab equipment provide accurate data but increase time and cost.
- Precision agriculture offers scalable solutions to reduce labor reliance and enhance monitoring, addressing food production challenges.

Current Best Practices

- At least 10% of global food production is lost annually to plant diseases, worsening food insecurity for over 800 million people [7].
- Thermal imaging has been used to detect early-stage PM in crops like wheat [2].
- In grape cultivation, moisture management is critical to prevent mildew caused by environmental factors like temperature and humidity.
- Vineyard inspections and fungicides are employed to address moisture-induced issues.
- Growers use pest thresholds to guide intervention decisions and maintain grape quality.

Handheld Proximal Sensors

- Handheld farming devices, such as Chlorophyll Meters, NDVI Sensors, and Water Potential Meters, collect data to validate plant health.
- Some devices are field-usable, while others require lab testing, increasing cost and time.
- The SPAD 502 Chlorophyll Meter clamps onto leaves and measures chlorophyll nondestructively at two wavelengths.
- Chlorophyll measurements from the SPAD 502 allow quick nitrogen content calculation [8], speeding up nitrogen level assessment compared to full-plant testing.

Drones and UAVs

- Drones offer scalability, completing tasks in a fraction of the time compared to handheld devices, with the added benefit of consistent data collection.
- Weekly UAV flyovers establish a field baseline, enabling efficient progress tracking and predicting PM severity and grain yield.
- Satellite imagery provides structured data for broader applications like urban planning and climate change [9].

Security Challenges

- Precision agriculture faces vulnerabilities, including pranks and state-sponsored attacks, with an estimated $10 million in economic losses globally [10].
- Potential attack types include reverse engineering, MITM, spoofing, DoS, false data injection, and RF jamming [10].
- Data poisoning is a risk with computer vision and ML; federated learning can secure models by training them locally and updating only with consensus, while mitigating risks [10].

Introduction

- AI and ML in precision agriculture identify crops, weeds, and diseases via drones.
- Powdery mildew (PM) is a major challenge for vineyards.
- Early PM detection is critical; visible symptoms appear too late for action.
- Hyperspectral imaging aids detection but is computationally expensive.

Introduction

- YOLO model proposed for PM detection using conventional images.
- Model trained on annotated cell phone and drone images.
- Provides actionable insights and early warning for PM levels.
- Red Zinfandel grapes at Cal Poly Pomona are highly susceptible to PM.
- Harvest season for Red Zinfandel is typically October.

Test Plot

- Vineyard near Cal Poly Pomona farm store is used as the test plot
- The targeted section consists of 14 rows of vines

2. Methodology

Research Goals

1. **Collect** and **process** imagery of vineyards.
2. **Train** ML models to detect PM.
3. **Evaluate** and **compare** the model's predictions.

Methodology

1. Data Collection and Processing
2. ML Training
3. Evaluate Model Performance

Handheld Camera Collection

- Images were captured at Cal Poly Pomona vineyard using an iPhone 11 camera from about 2.5 meters height (Handheld Camera Trained).
- Image resolution was 3024 by 4032 pixels, taken close to the drone flight on each collection day.
- Data collection occurred on Tuesdays around noon, from September to November 2024, at the end of the grape growth cycle.

Drone Collection

- The DJI Phantom 4 RTK has a 1-inch, 20-megapixel CMOS sensor, a gimbal with stabilization on tilt, roll, and yaw axes, and a max speed of 36 mph.
- It achieves a ground sample distance of 2.74 cm at 100 meters flight altitude, and the lens specs include FOV 84°, 8.8 mm/24 mm (35 mm equivalent), and an f/2.8 - f/11 aperture with auto focus from 1m to infinity.
- Imagery was collected at 5 meters altitude (lowest to avoid collisions with vines), saved with GPS data, under clear skies and low wind, based on preset flight waypoints, on Tuesdays from September to November 2024.

Data Collection - Drone

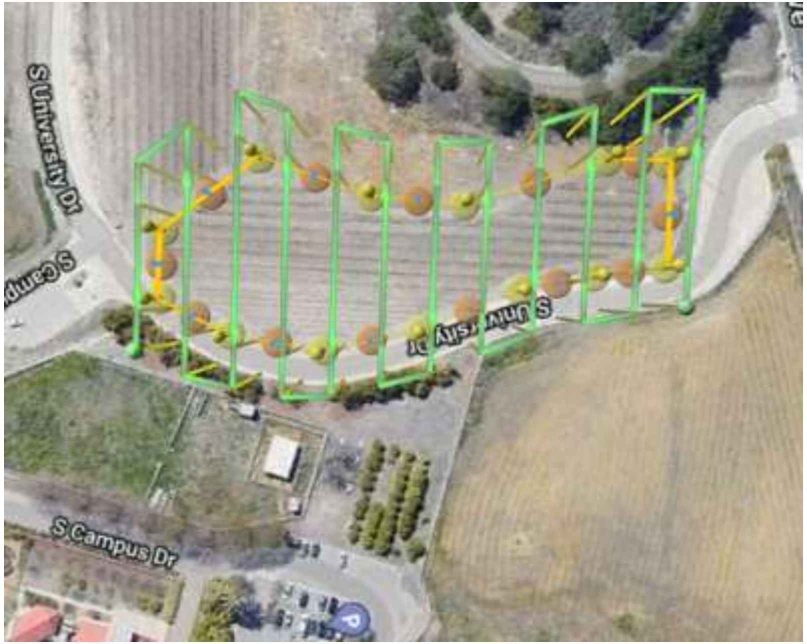

Pre-processing

- After the images are collected, manual pre-processing is necessary in order to select quality images.
- Each image should have a vine where a healthy or diseased detection is possible.
- Images should avoid being all dirt or miscellaneous objects.
- Both the classes should be balanced.

ML Training - YOLO

- The models will be trained using two methods.
- For both methods, the images and annotations will create a custom dataset that will be fed into the YOLO model. Each row of a text file represents one bounding box.
- Two YOLO models will be used in each combination of testing and training data: v8n and v10n. The n models prioritize lightweight, fast performance.

Object_Class	x_Center	y_Center	Width	Height
0	0.85	0.29	0.27	0.12

CVAT

- Computer Vision Annotation Tool
- Online annotation tool that generates annotations in a variety of formats.

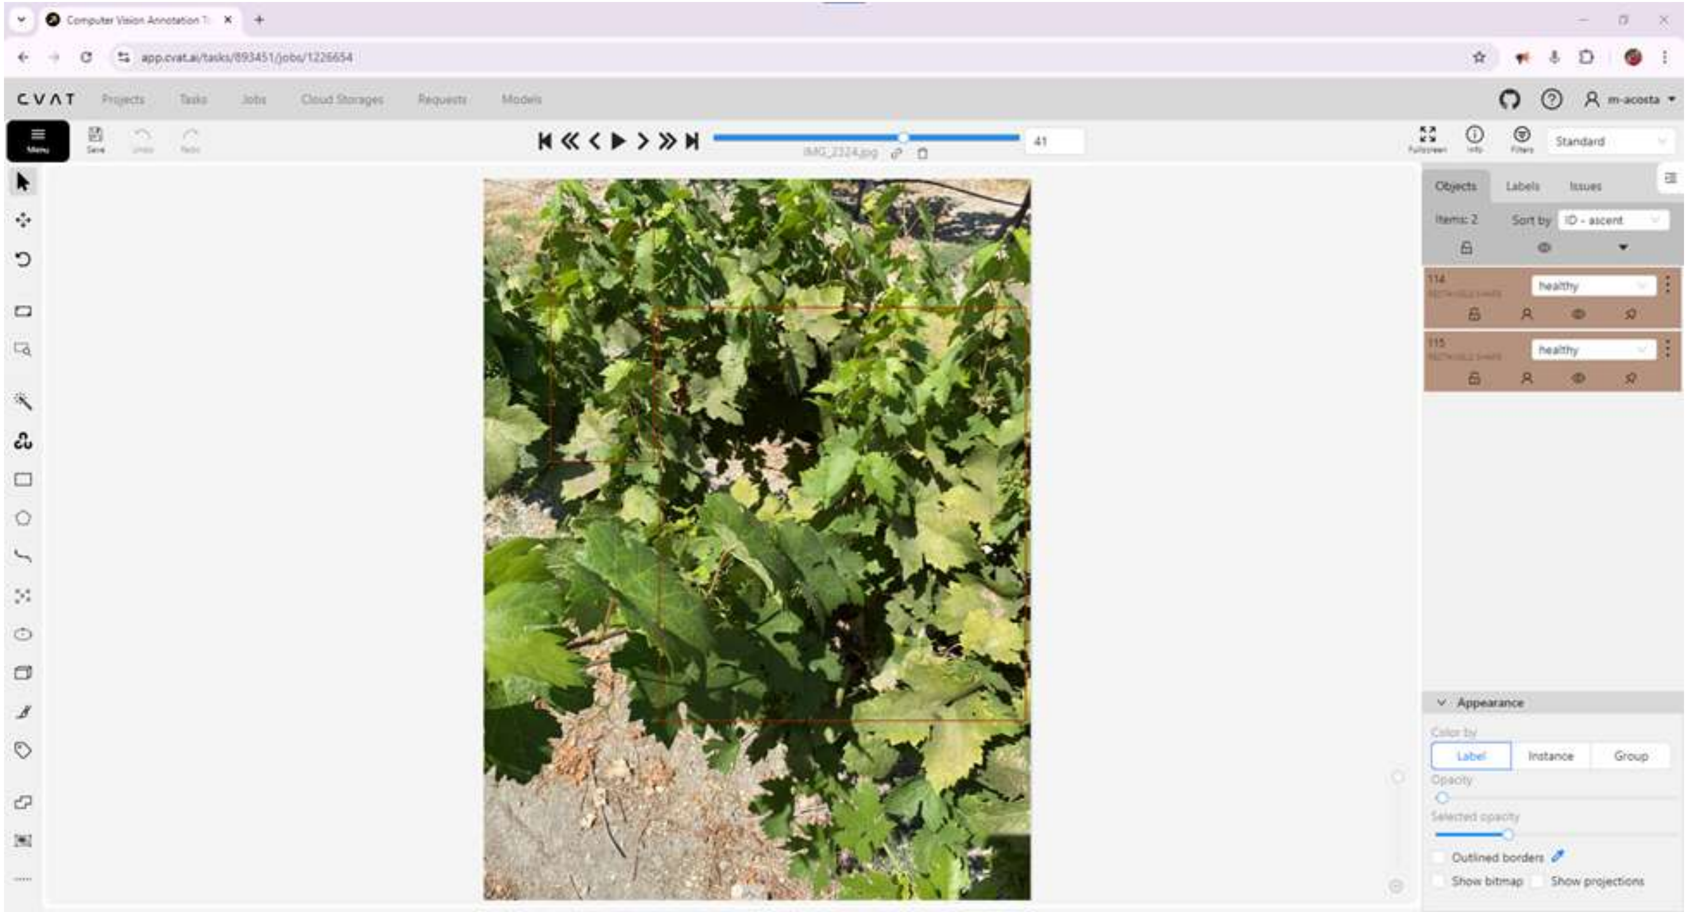
Computer Vision Annotation T
app.cvat.ai/tasks/893451/jobs/1226654
CVAT
Projects
Tasks
Jobs
Cloud Storages
Requests
Models
m-acosta
41
Standard
Objects
Labels
Issues
Items: 2
Sort by
ID - ascent
114
healthy
115
healthy
Appearance
Label
Instance
Group
Outlined borders
Show bitmap
Show projections

ML Training - Handheld Camera

- There are 156 total collected images, all YOLO annotated in CVAT.
- The images are split 80/10/10 for training, validation, and test.
- There are 165 healthy instances and 200 diseased instances in the training and validation datasets.
- Healthy plants are large sections of the plant highlighted with a bounding box, avoiding extraneous objects such as the dirt below or metal support wires.
- Diseased labels will be much smaller and target areas where PM is the worst.


HEALTHY 1 (MANUAL)

```yaml
augmentations:
  hsv_h: 0.02              # Hue augmentation value, allowing hue shifts of ±2%
  hsv_s: 0.8               # Saturation augmentation value, allowing saturation shifts of ±80%
  hsv_v: 0.5               # Brightness (value) augmentation, allowing brightness shifts of ±50%
  degrees: 45.0            # Allow rotation by ±45 degrees to account for different orientations
  translate: 0.2           # Translate image by up to 20% of its width/height
  scale: 0.5-1.5           # Scale the image by 50% to 150% of its original size
  shear: 0.1               # Small shear transformations
  perspective: 0.001       # Small perspective shift to simulate different camera angles
  flipud: 0.5              # Enable vertical flips with a probability of 50%
  fliplr: 0.5              # Enable horizontal flips with a probability of 50%
  mosaic: 1.0              # Enable mosaic augmentation (combining 4 images into 1)
  mixup: 0.0               # Disable mixup augmentation
  copy_paste: 0.0          # Disable copy-paste augmentation
  erasing: 0.5             # Random erasing probability of 50%
  auto_augment: randaugment  # Use RandAugment for automatic transformations
  crop_fraction: 0.8       # Crop up to 80% of the image to simulate different viewpoints
  # Additional settings:
  blur: 0.3                # Optional: Apply a slight blur to simulate reduced resolution
  noise: 0.2               # Optional: Add Gaussian noise to simulate sensor noise
```

Any questions so far?

ML Training - Drone

- There are 300 total collected images, all YOLO annotated in CVAT.
- The images are split 80/10/10 for training, validation, and test.
- There are 868 healthy instances and 868 diseased instances in the training and validation datasets.
- Healthy plants are small bounding boxes (that **match the diseased ones**).
- Diseased labels are small and target areas where PM is the worst.

```yaml
augmentations:
  hsv_h: 0.02             # Allows hue shifts of ±2%, useful for natural color variations in leaves
  hsv_s: 0.8              # Allows saturation shifts of ±80%, covering a range of leaf color saturations
  hsv_v: 0.5              # Allows brightness shifts of ±50%, simulating various lighting conditions
  degrees: 10.0           # Rotation by ±10 degrees for slight orientation adjustments
  translate: 0.1          # Translates image by up to 10% to mimic shifts in the leaf’s position within frame
  scale: 0.8-1.2          # Scales the image between 80% and 120% to simulate small zoom-ins/outs
  shear: 0.05             # Small shear transformations to adjust for slight distortions
  perspective: 0.002      # Perspective shift for subtle changes in camera angle
  flipud: 0.2             # Enables vertical flips with a probability of 20%
  fliplr: 0.5             # Enables horizontal flips with a probability of 50%
  mosaic: 0.5             # Mosaic with a probability of 50%, combining four images for better context
  mixup: 0.0              # Disable mixup augmentation (not required for this context)
  copy_paste: 0.0         # Disable copy-paste augmentation
  erasing: 0.3            # Random erasing with a probability of 30% to simulate occlusion
  auto_augment: randaugment  # Applies RandAugment for automatic varied transformations
  crop_fraction: 0.8      # Crop up to 80% of the image to highlight different parts of the scene
  blur: 0.2               # Slight blur with 20% probability to mimic slight focus issues
  noise: 0.2              # Adds Gaussian noise with 20% probability to simulate sensor noise
```

Evaluate Model Performance

- Images from handheld cameras and drones are analyzed using YOLO versions, with predictions made as bounding boxes, each image likely having multiple predictions.
- GPS location is extracted from EXIF data and mapped to pixels, with the center of the bounding box representing the GPS location.
- Predictions are classified as healthy or diseased, with manual checks by experts to verify accuracy; mismatches are labeled false positives, while correct matches are true positives.

Metrics

Recall (True Positive Rate): $\frac{TP}{TP+FN}$

Precision: $\frac{TP}{TP+FP}$

F1 Score (F Measure): $\frac{2\times Recall\times Precision}{Recall+Precision}$

Confidence:

The confidence value ranges from 0 to 1 where higher values mean the model is more confident in the prediction.

Computational Time:

Calculate as the average run time of evaluation for each image during the test phase (run on an Intel Core(TM) i7-4770 3.40GHz CPU).

3. Results

Results

Model and Train/Test Combination			Precision	Recall	Average Confidence	F1
YOLOv8n	Handheld Trained	Handheld Tested	1	0.24	0.36	0.39
		Drone Tested	*no possible predictions*			
	Drone Trained	Handheld Tested	0.62	0.64	0.36	0.63
		Drone Tested	1	0.48	0.33	0.65
YOLOv10n	Handheld Trained	Handheld Tested	0.79	0.76	0.75	0.77
		Drone Tested	1	0.009	0.43	0.018
	Drone Trained	Handheld Tested	0.4	0.72	0.34	0.51
		Drone Tested	0.94	0.64	0.35	0.76

Comparison of Models

- There are 8 combinations of model architecture, training dataset, and test dataset.
- The focus is solely on the diseased class. The model is not evaluated on its predictions for the healthy class.
- Once the predictions are made, they are reviewed by an expert (ground truth).

Comparison of Models

- Each image in the test set is annotated in advance in order to capture the entire area that is infected with PM.
- This shows all of the positive cases of PM that are present in the vineyard.
- Recall is more valuable, as identifying as much PM in the vineyard as possible is the goal. If this means more false positives, that is acceptable.

Comparison of Models

- YOLO models performed better at detecting powdery mildew when training and test images were the same (e.g., handheld vs. handheld), with reduced success when images varied (e.g., handheld vs. drone).
- The best YOLO model achieved 79% precision, 76% recall, and 77% F1 score using YOLOv10n with handheld images for both training and testing.
- The model performed well with heavily diseased cases of PM but struggled with overexposed leaves and areas of dirt; factors like bounding box size and image quantity affected results, suggesting room for improvement.

Train/Test with Different Dataset

- 4 of the models are training with different images than the test set.
- These performed much worse that matching image sets.
- The initial approach was to explore this method.
- The models just don’t generalize very well with the change in resolution.

Handheld Camera: Test Data

- The test set is comprised of 15 images across the weeks of collection.
- This is 10% of the total images from the handheld camera.
- There are 25 instances of PM present within these images.

Bounding Box Collisions

Distant False Positives

Drone: Test Data

- The test set is comprised of 30 images across the weeks of collection.
- This is 10% of the total images from the handheld camera.
- There are 109 instances of PM present within these images.

Image Name	GPS location	Health	Confidence	Ground Truth	Classification Outcome
0903_DJI_0074.JPG	34.0494329722222, -117.820467111111	Diseased	0.48	Diseased	True Positive
0903_DJI_0074.JPG	34.0494329722222, -117.820467111111	Diseased	0.36	Diseased	True Positive
0903_DJI_0074.JPG	34.0494329722222, -117.820467111111	Diseased	0.27	Diseased	True Positive
0903_DJI_0077.JPG	34.0493429166667, -117.820498	Diseased	0.33	Diseased	True Positive
0903_DJI_0085.JPG	34.0491014722222, -117.8205835	Diseased	0.54	Diseased	True Positive
0903_DJI_0085.JPG	34.0491014722222, -117.8205835	Diseased	0.37	Diseased	True Positive
0903_DJI_0085.JPG	34.0491014722222, -117.8205835	Diseased	0.27	Diseased	True Positive
0903_DJI_0085.JPG	34.0491014722222, -117.8205835	Diseased	0.27	Diseased	True Positive
1001_DJI_0017.JPG	34.049141388888884, -117.820421111111	Diseased	0.47	Diseased	True Positive
1001_DJI_0017.JPG	34.049141388888884, -117.820421111111	Diseased	0.38	Diseased	True Positive
1001_DJI_0017.JPG	34.049141388888884, -117.820421111111	Diseased	0.33	Diseased	True Positive
1001_DJI_0017.JPG	34.049141388888884, -117.820421111111	Diseased	0.27	Diseased	True Positive
1001_DJI_0035.JPG	34.0490250277777, -117.820511027777	Diseased	0.32	Healthy	False Positive

Confidence

- The low confidence values could come from a lack of diversity in the dataset or size.
- Another issue could be the small size of PM compared to the resolution.
- It needs to be explored in the future why the precision is so high, but the confidence values are not.
- None of the predictions were excluded based on confidence because recall was the priority.

True Positive Prediction

False Positive Prediction

Comparison

- Drone-trained models have lower prediction confidence, but still perform well, while handheld camera-trained models sometimes fail to make predictions due to fewer PM instances in the dataset.
- Precision generally outperforms recall, indicating high-quality predictions with fewer false positives, but the system aims to prioritize recall to detect more PM.
- Drone-based training is more promising than handheld cameras, as it provides faster, reproducible, and more easily automated image collection, enhancing the effectiveness of ML in precision agriculture.
- The drone training also scored high with F1 across the board when compared to handheld training.

4. Conclusion

Discussion

- Images collected later in the growth cycle are more valuable as PM becomes more apparent. Early-stage detection is harder, especially with images taken from above.
- PM thrives on the underside of vines, and an AR app with a YOLO model could help identify diseased leaves using a phone camera.

Discussion

- False positives, such as dirt resembling PM, are an issue; increasing bounding box size could reduce this problem.
- The number of training epochs affects performance, and the ideal value needs further exploration to avoid overfitting/underfitting.
- The study used YOLO nano architecture for real-time processing, trading off precision and recall for lightweight size and speed suitable for drones. Much higher scores could be achieved with the performance models.

Conclusion

- AI and ML provide cost-effective solutions for detecting PM in vineyards using lightweight YOLO models with handheld cameras and drones.
- YOLOv10n, trained on handheld images, achieved 79% precision, 76% recall, and 77% F1 score; drone models showed strong large-scale monitoring potential.
- Challenges include false positives (due to dirt) and variability in image quality, with potential improvements from better image diversity and collection angles.
- Drone-based data collection offers superior scalability, making it ideal for precision agriculture and crop health management.

Future Work

- The ultimate goal of this work at Cal Poly Pomona is to have drones and UGVs working maintain a field of crops.
- This is useful for weed removal, pesticide application, etc.
- The drone will fly and identify the areas that contain weeds or pests and communicate that information to the UGV. The UGV will then go in with that general location and use its targeted sensors to correct the problem.
- There will be a trained YOLO model on a drone equipped with an NVIDIA Jetson processor to detect in real time.
- Also, work with hyperspectral imagery will also continue within this project.

Acknowledgements

- The project was supported by the ARI Grant Number 23-04-113.

References

[1] Knauer, Uwe, et al. "Improved Classification Accuracy of Powdery Mildew Infection Levels of Wine Grapes by Spatial-Spectral Analysis of Hyperspectral Images." *Plant Methods*, vol. 13, no. 1, 2017, https://doi.org/10.1186/s13007-017-0198-y. Accessed 14 Apr. 2023.
[2] Bayoumi, Tarek Y., and A. A. Abdullah. "Application of thermal imaging sensor to early detect powdery mildew disease in wheat." *The Journal of Middle East and North Africa Sciences*, vol. 2, no. 2, 2, 2016, pp. 56–63, https://doi.org/10.12816/0032657.
[3] Khan, Imran Haider, et al. "Early detection of powdery mildew disease and accurate quantification of its severity using hyperspectral images in wheat." *Remote Sensing*, vol. 13, no. 18, 2021, p. 3612, https://doi.org/10.3390/rs13183612.
[4] Kunova, Andrea, et al. "Grapevine powdery mildew: Fungicides for its management and advances in molecular detection of markers associated with resistance." *Microorganisms*, vol. 9, no. 7, 20 July 2021, p. 1541, https://doi.org/10.3390/microorganisms9071541.
[5] Khan, Imran Haider, et al. "Early detection of powdery mildew disease and accurate quantification of its severity using hyperspectral images in wheat." *Remote Sensing*, vol. 13, no. 18, 2021, p. 3612, https://doi.org/10.3390/rs13183612.

References

[6] Neupane, Krishna, and Fulya Baysal-Gurel. "Automatic identification and monitoring of plant diseases using unmanned aerial vehicles: A Review." *Remote Sensing*, vol. 13, no. 19, 2021, p. 3841, https://doi.org/10.3390/rs13193841.
[7] Strange, Richard N., and Peter R. Scott. "Plant disease: A threat to Global Food Security." *Annual Review of Phytopathology*, vol. 43, no. 1, 1 Sept. 2005, pp. 83–116, https://doi.org/10.1146/annurev.phyto.43.113004.133839.
[8] Yang, Yang, et al. "Nitrogen concentration and dry-matter accumulation in maize crop: Assessing maize nitrogen status with an allometric function and a chlorophyll meter." *Communications in Soil Science and Plant Analysis*, vol. 43, no. 11, 2012, pp. 1563–1575, https://doi.org/10.1080/00103624.2012.675393.
[9] Demir, Ilke, et al. "DeepGlobe 2018: A Challenge to parse the Earth through satellite images." *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, June 2018, https://doi.org/10.1109/cvprw.2018.00031.

References

[10] Balaji, Shree Ram, et al. “Cybersecurity challenges and solutions in IOT-based Precision Farming Systems.” *2023 IEEE 14th Annual Ubiquitous Computing, Electronics & Mobile Communication Conference (UEMCON)*, 12 Oct. 2023, pp. 237–246, https://doi.org/10.1109/uemcon59035.2023.10316154.

Q & A