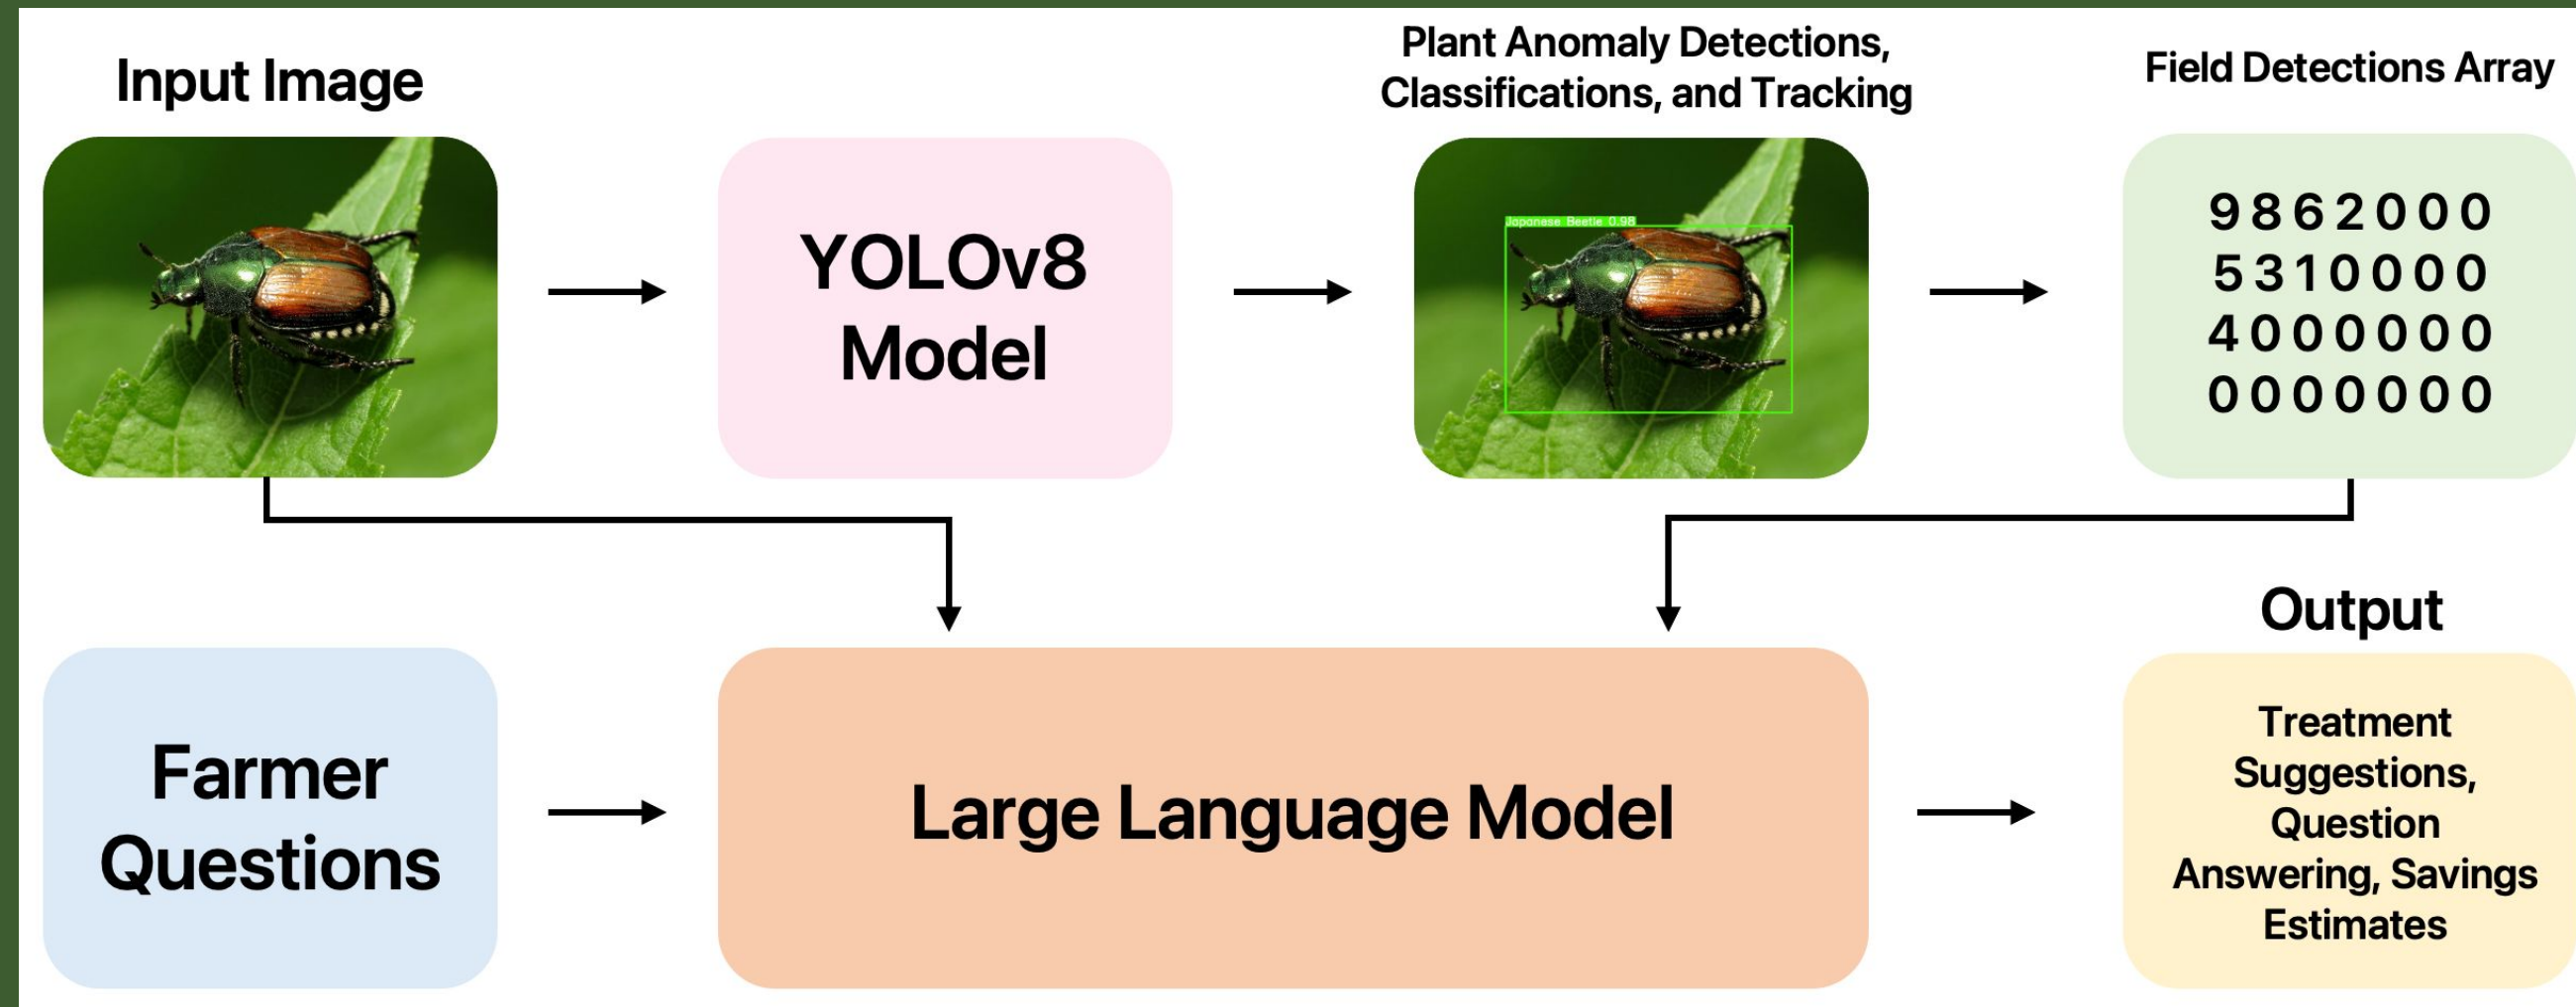


LeAF: *L*everaging *C*onvolutional *N*eural *N*etworks for *P*lant *A*nomaly *D*etection and *C*lassification for *F*armers with *L*arge *L*anguage *M*odels for *N*atural *L*anguage *I*nteraction



By Aditya Sengupta - Redmond, WA, USA

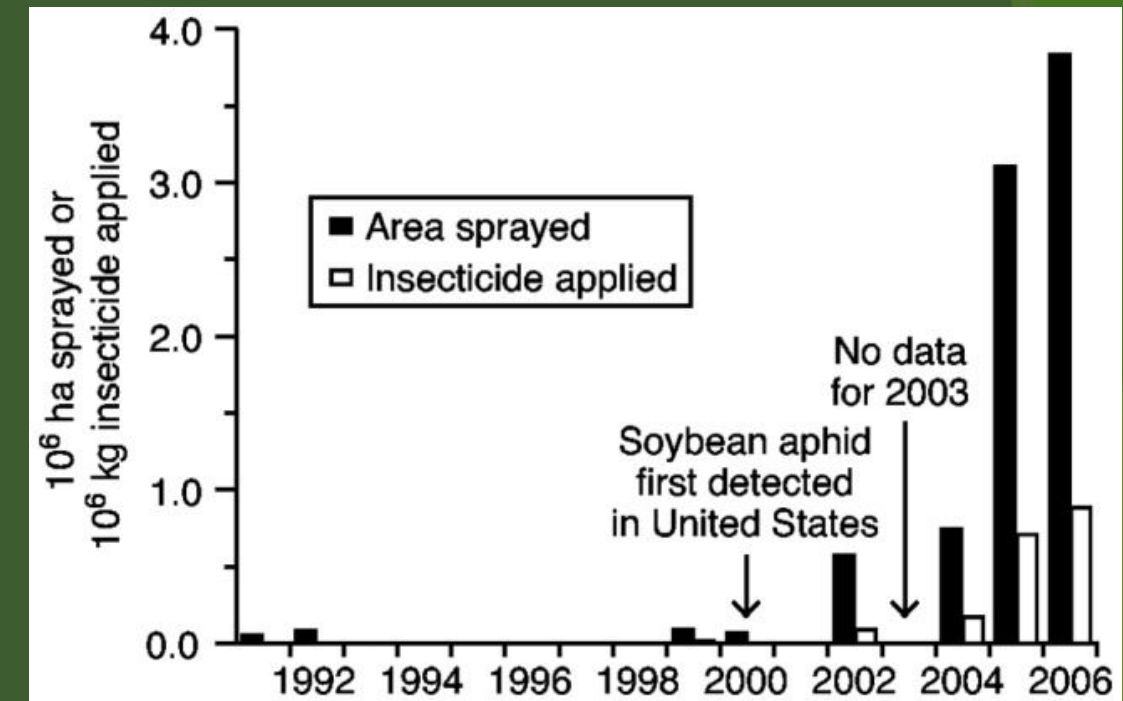
Motivation

- “You Pick Strawberries” at local farm
- Half-eaten and rotten strawberries
- Crops ravaged by pests
- Need for solution to help farmers deal with plant anomalies



Research Problem

- **Farmer pain point:** monitoring plant anomalies (pests, weeds, disease)
 - Cause losses **40%** of global crop production (**\$220 billion**) annually
- **Current Practices:** manual surveying and/or excessive and indiscriminate use of agricultural chemicals
- Environmental, Economic, and Ecosystem Impact of Chemicals
 - **\$60 billion** spent annually
 - **40 million kg** greenhouse gas emissions
 - **20,000+** new cases of cancer every year
 - **80+** million fish and bird deaths annually
- Need for low-cost and affordable automation to monitor plant anomalies and give focused treatment suggestions



Graph showing pesticide usage increase



Comparison with Existing Work

Dimension	Existing Work	LeAF
Solution Type	Classification Only	Classification & Bounding Box
Accuracy	30-40%	90%+
Extensibility	Only Specific Plants/Anomalies	All Plants/Anomalies
Model Size	20-100 million parameters	3 million parameters
End-to-End Solution?	No	Yes
Natural Language Interaction?	No	Fluent and Knowledgeable

Research Questions & Hypothesis

Question: Is it possible to automate plant anomaly detection for farmers, reduce chemical usage, and provide targeted treatment suggestions?

Hypothesis: Recent advances in AI/ML with deep learning based model architectures trained on agricultural datasets can provide effective solutions

Key Extensions

- **Generative AI:** Can natural language assistant help farmers for real-time monitoring and decision-making?
- **Deployability:** Can models be made lightweight for running on edge and mobile devices? Is the solution deployable with agricultural robots with cameras for data collection for analysis?

Japanese Beetle 0.98



Pests

Dandelion 0.96



Weeds

Gray Leaf Spot 0.99



Plant Disease

Sample detections for different types of plant anomalies

LeAF Methodology

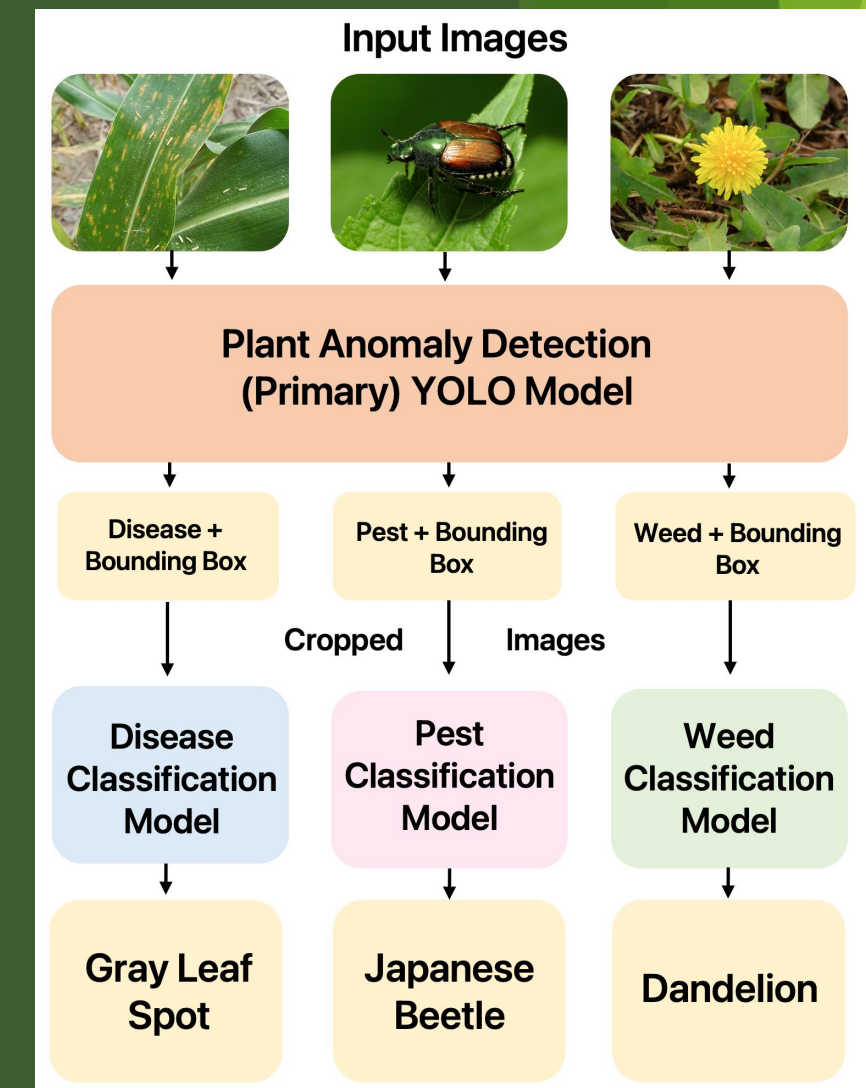
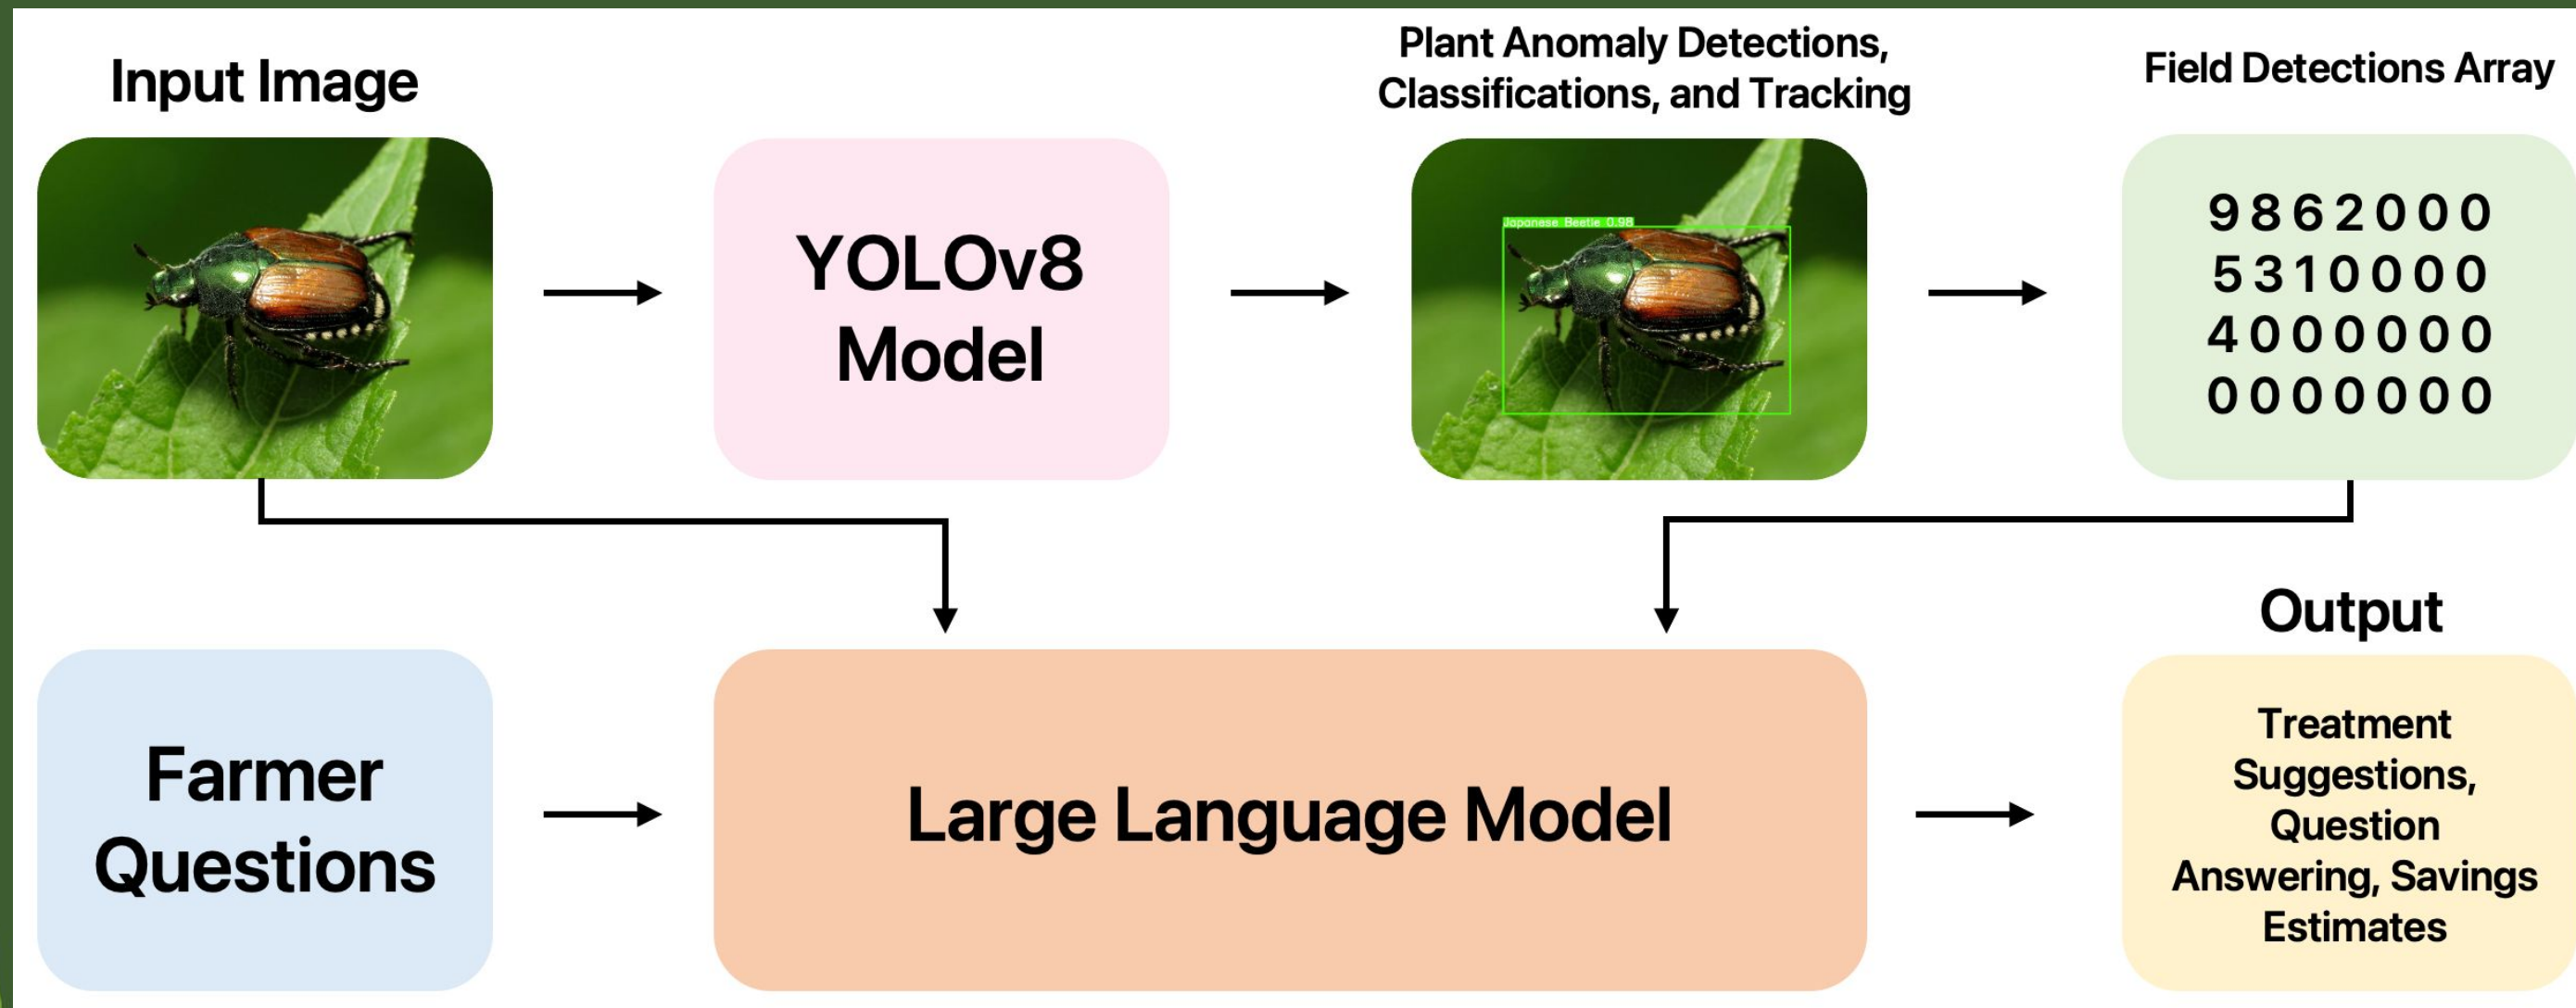


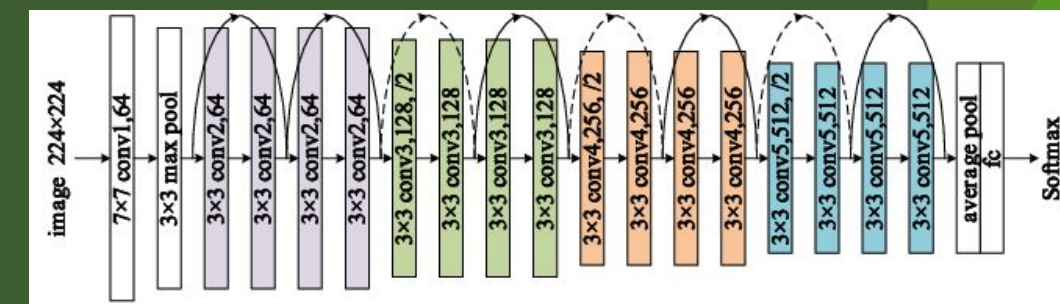
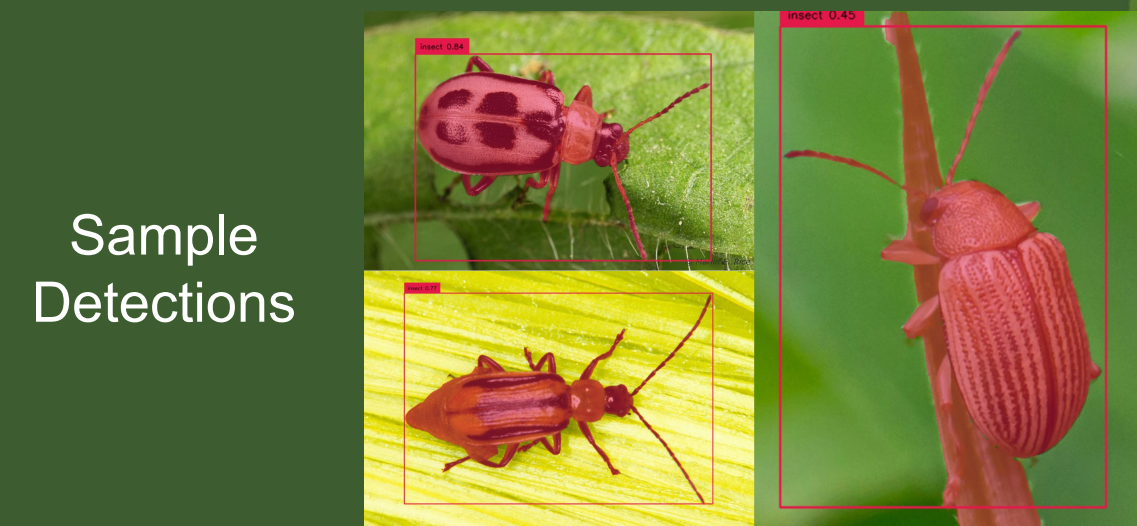
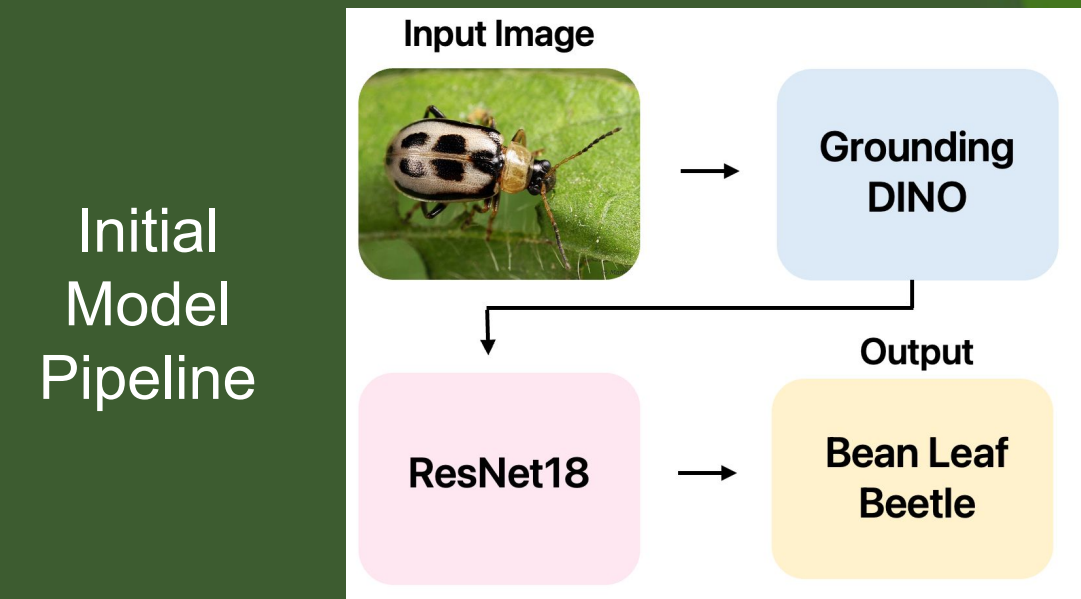
Image → Classification and Detection Model → Object Tracking → Field Array → Large Language Model → Output Treatment Suggestions

Framework Extensible to Multiple Plant Anomalies

Dataset & Model Pipeline



- 10 most harmful agricultural pests in US Midwest
- 10,000 images from iNaturalist (1,000 images per class)
- Data cleaning based on image quality, research grade, number of validations
- 80%-10%-10% Train-Validation-Test Split
- Pipeline: GroundingDINO for bounding box → cropped image → Classification model
- ResNet-18 classification model (most accurate compared to 5 other models)

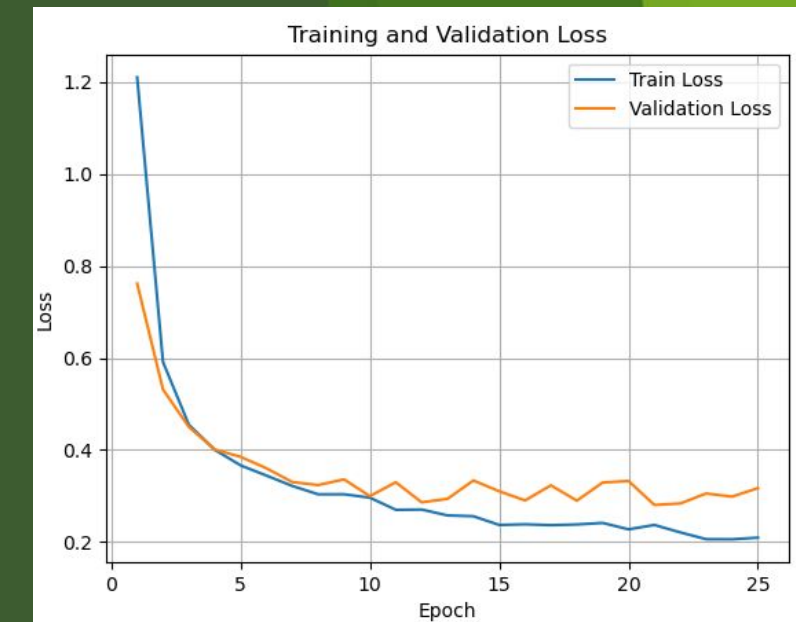


ResNet-18 Model Architecture

Results: Training & Evaluation

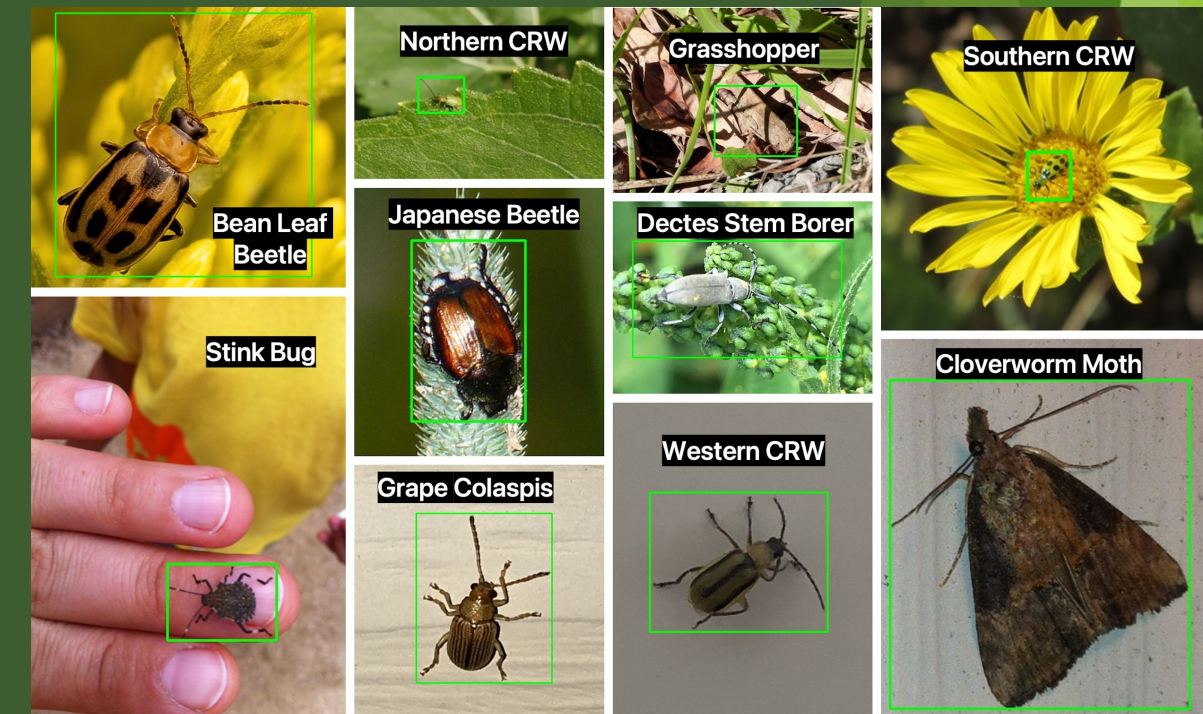
- Iterative training, testing, and tuning of hyperparameters: **accuracy increase from 65% to 91%**
- **Highly accurate** even on small and camouflaged pests
- Prediction Latency:
 - **3 seconds per image** on Nvidia T4 GPU
 - GroundingDINO taking majority of time
- Opportunity for a **lightweight model** to run on a robot with **low compute** and fast enough to keep up with **video feed**

Model Training Curve



Tuning Hyperparameters

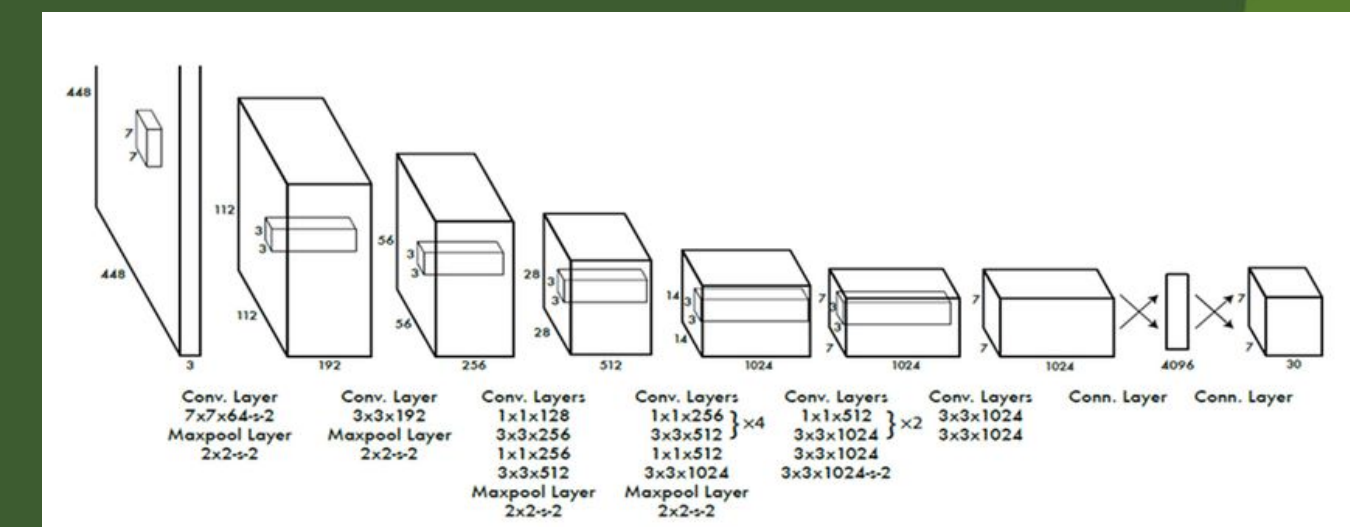
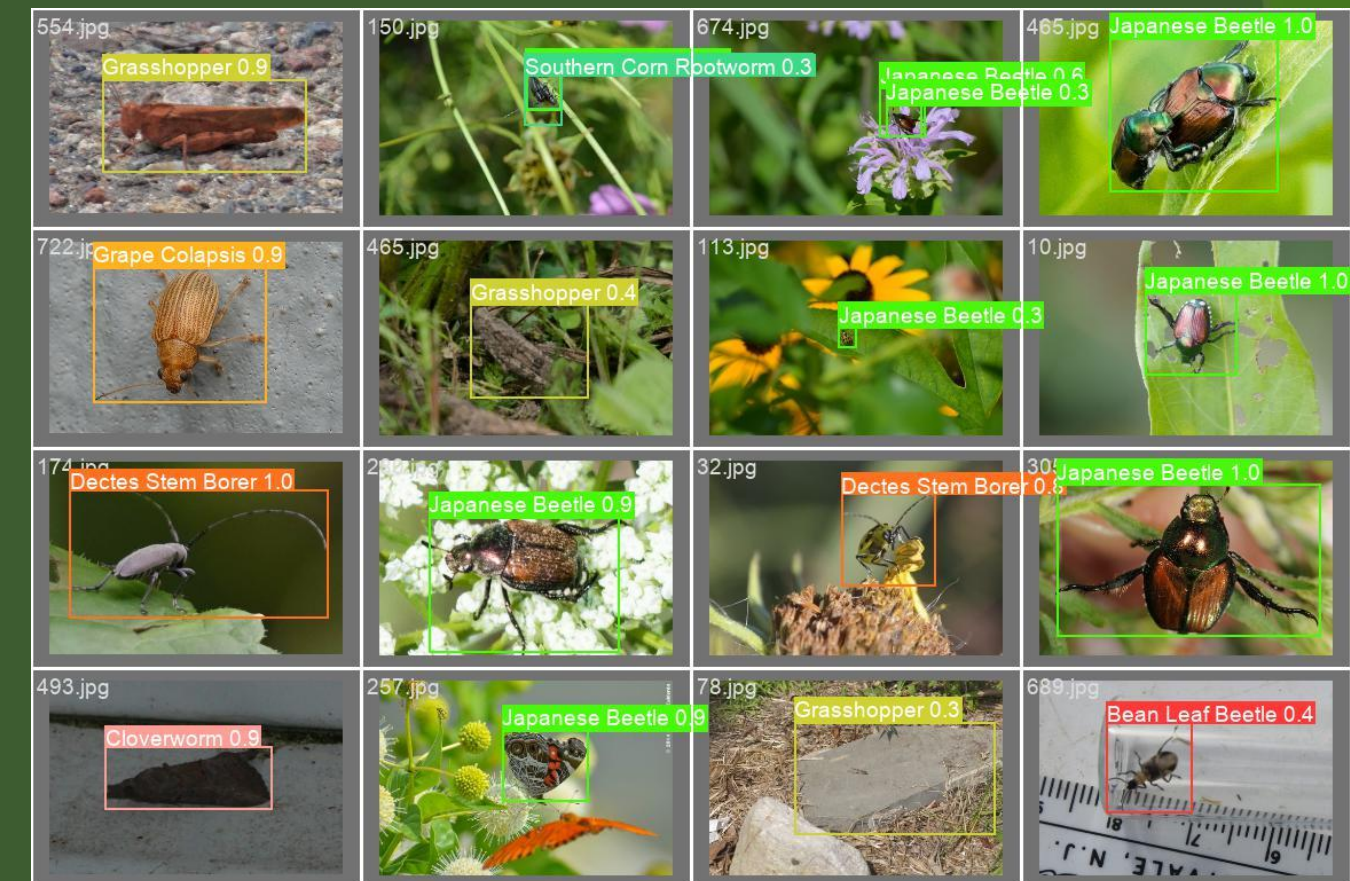
Tuning Technique	Validation Accuracy
Original Accuracy	65.71%
Dataset Filtering and Augmentation	81.29%
Adam Optimizer	88.97%
Learning Rate 0.001	90.96%
Batch Size 32	91.29%
Train 10 Epochs (Last Layer)	91.43%



Sample Detections using custom-trained ResNet

Distilling for Mobile Devices

Sample Detections using Custom-Trained YOLO

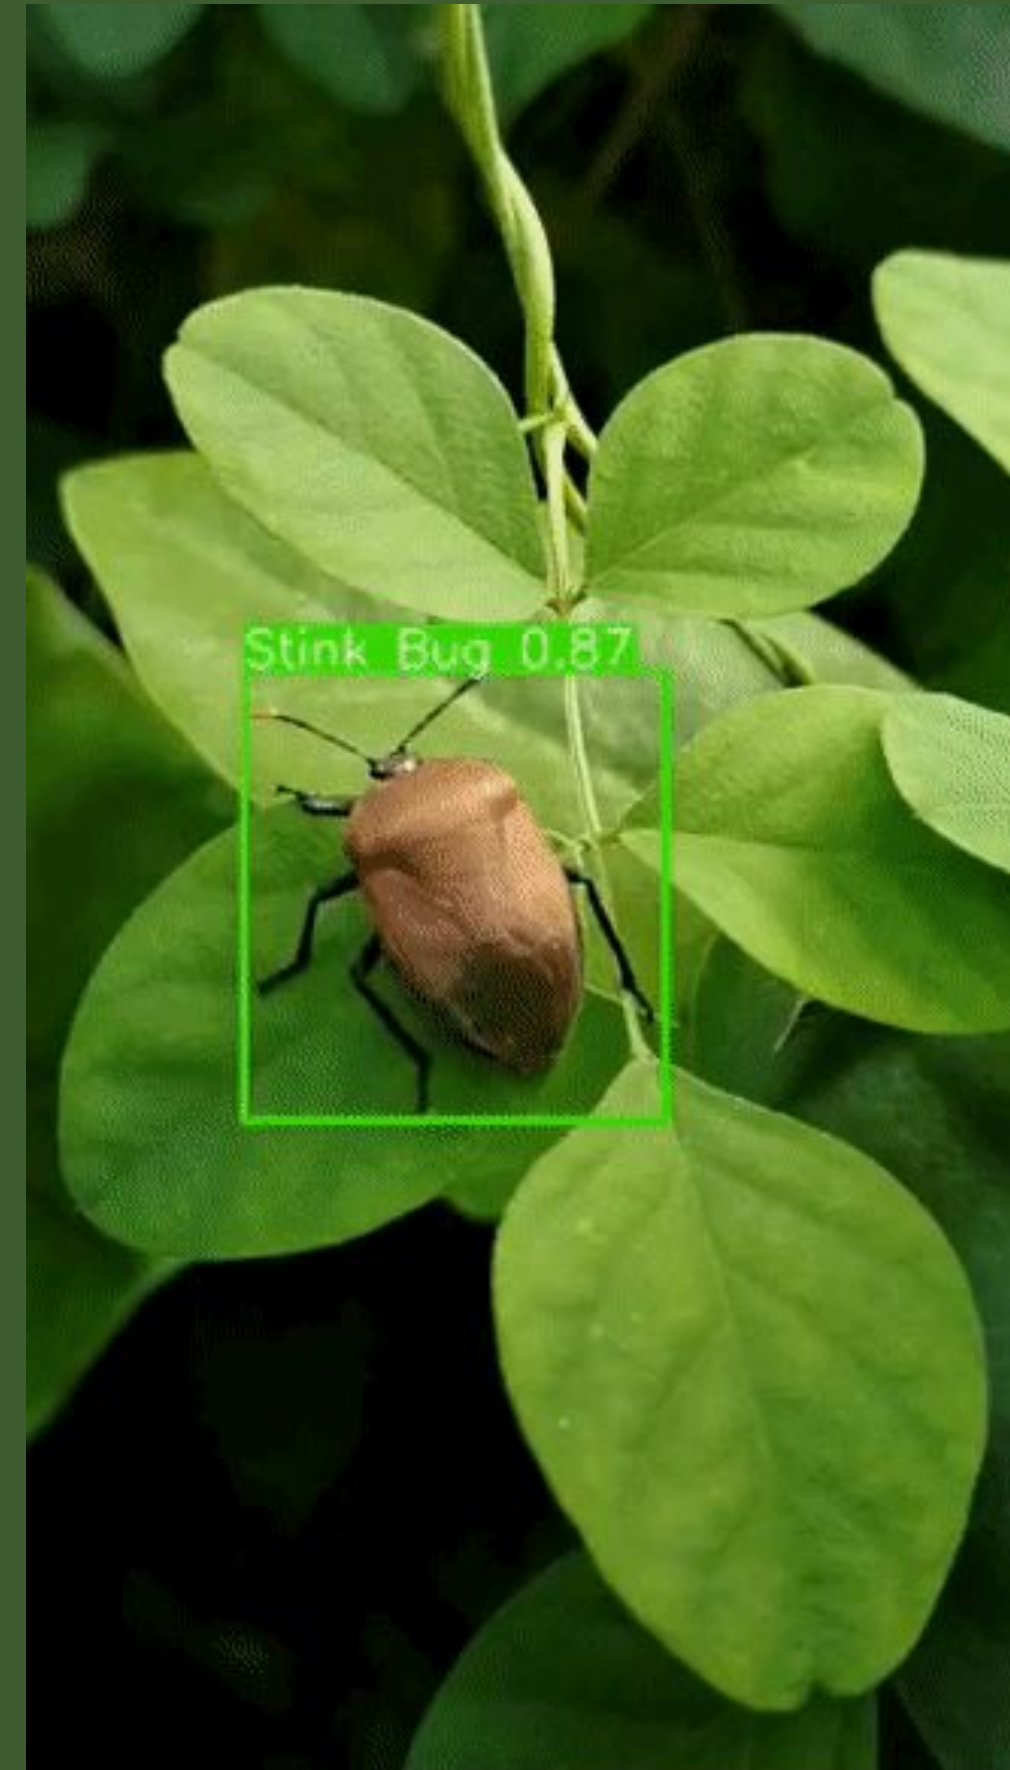
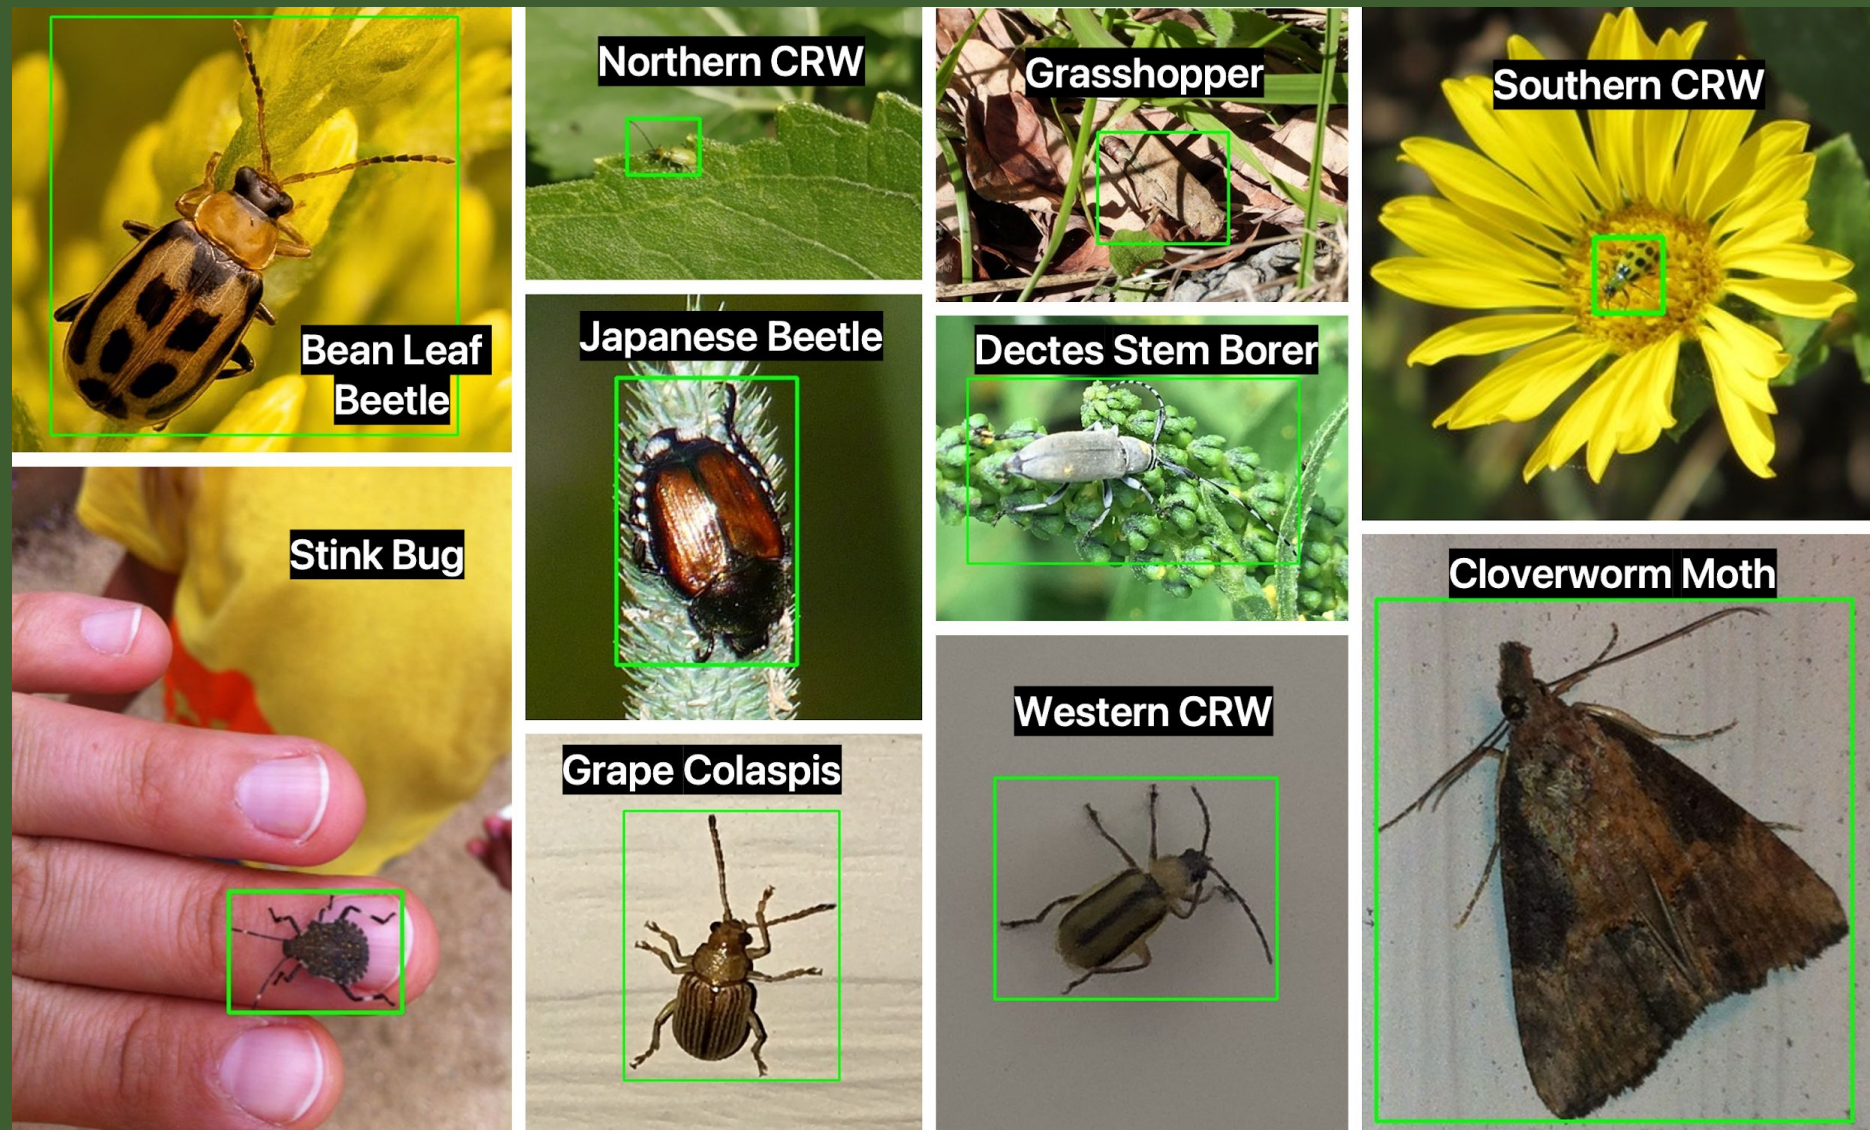


YOLO Model Architecture

- Large model pipeline → smaller model
- Classification and detection in one model
- YOLOv8 Nano
 - Smaller model (0.17% of original pipeline)
 - Runs on low compute devices
- Results
 - Same accuracy with 600x size reduction
 - 10 msec for prediction per image (300x latency reduction)
 - Can run live on robot on video feed

Results: Evaluation in Challenging Conditions

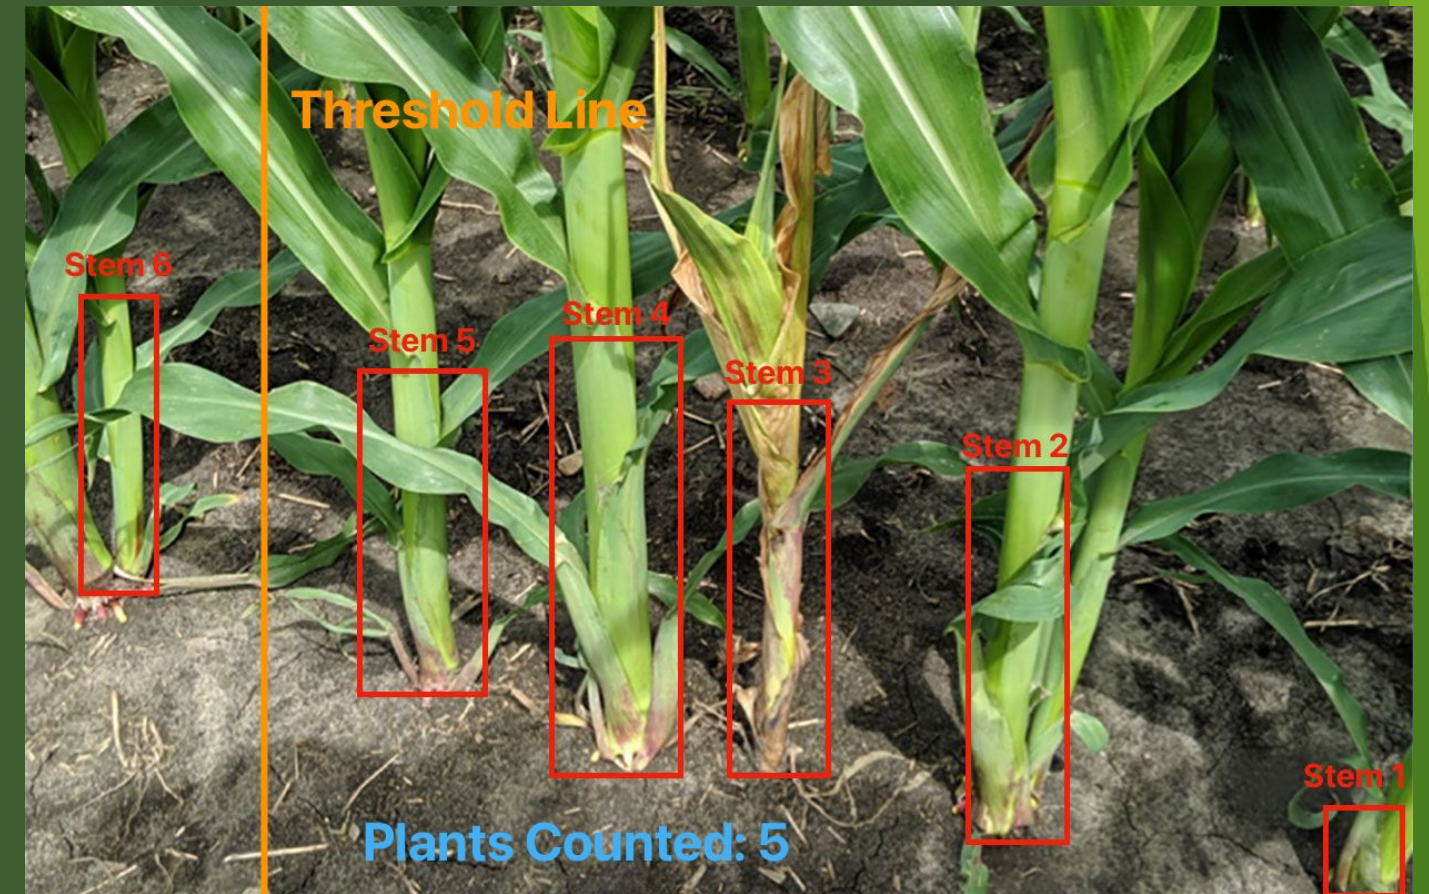
- Similar Background Color: Bean Leaf Beetle, Southern CRW
- Small and Blending with Background: Northern CRW
- Extreme Camouflage: Grasshopper
- Different pest orientations and viewing angles



Field Mapping

- YOLOv8 Nano model runs on the robot camera feed in the farm
- **Object tracking algorithm** to count pests
- Plant stem detection to cluster pests
- **2D field matrix** representing plant-specific pest counts
 - Provides **field map** to help farmers identify **clusters** of pest infestations
 - Plant-by-plant management enables early detection

Plant Stem Detections & Counting



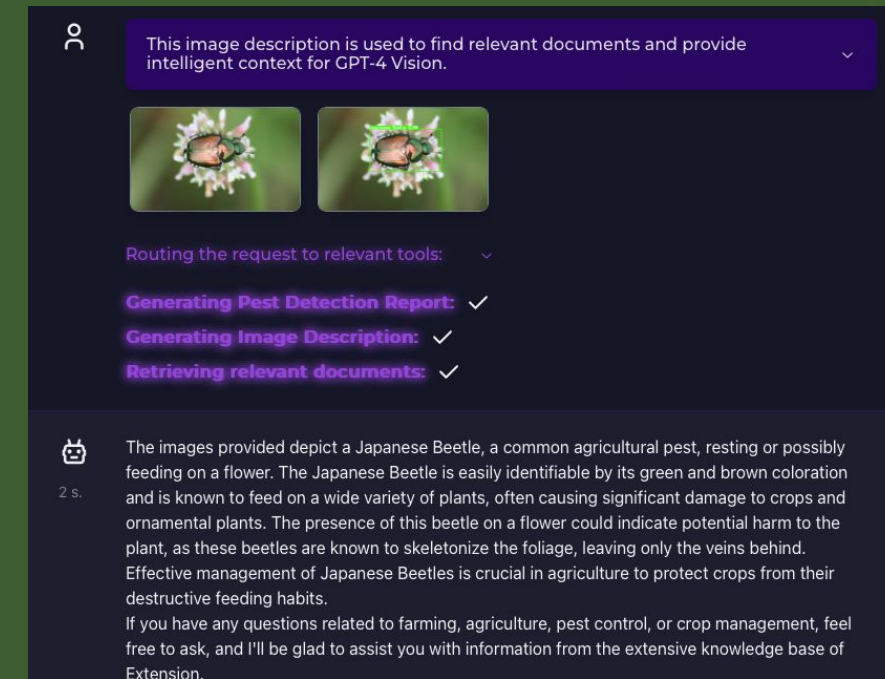
```
0 1 4 1 0 0
0 1 9 9 0 0
0 2 7 5 0 0
0 0 4 0 0 0
```

Sample Pest Field Array

Large Language Model & Web Tool

- Fine-tuned Google Gemini Multimodal Model
 - Pest images
 - Field array description
 - Farmer questions
 - Example answers
- Model provides **analysis of infestations** on field based on the map
 - Treatment suggestions
 - Farmer can ask **follow-up questions**
- Web tool to provide personalized estimates
 - **Financial, chemical, and emissions savings**


Natural Language Interaction



Farmer Input	LeAF Pesticide Analysis
Number of rows of plants 25 - +	Pesticide Usage Metrics
Number of plants per row 45 - +	Total litres of pesticide used before LeAF: 1383.75 liters
Pesticide usage per plant (litres) 1.23 - +	Total litres of pesticide used with LeAF: 167.28 liters
Number of plants treated with LeAF 136 - +	Total litres pesticides saved: 1216.47 liters
Cost of pesticide per litre (\$) 5.36 - +	Amount of money saved: \$6520.28
	Amount of carbon emissions saved: 3041.18 kg
	<input type="button" value="Generate LLM Response"/>

Web Estimate Tool

LeAF Natural Language Interaction with Farmer: Input & Output

INPUT	pest class	INPUT	field	INPUT	question	INPUT	image:	OUTPUT	answer
1	Japanese Beetle	1 9 9 0 0	3 8 4 0 0	0 3 1 0 0	0 0 0 0 0	0 0 0 0 0		<p>In this field, there is a cluster of Japanese Beetles concentrated in the northwest corner, extending 3 rows south and 3 rows west. Other areas of the field have little to no Japanese Beetle presence.</p> <p>Combat Japanese Beetle damage with targeted measures. Apply neonicotinoids or pyrethroids early for effective insecticide control. Enhance resilience by adopting cultural practices like crop rotation and debris removal, disrupting the beetles' life cycle to minimize future infestations. Regular monitoring ensures timely intervention, safeguarding crops from defoliation and yield loss.</p>	

BRANCH Robot

- Commercial robots cost \$5,000+
- Custom-made **BRANCH Robot** (Budget-friendly Robot for Agricultural Nonintrusive Crop Photography)
 - Accessible to farmers at **less than \$500**
 - Manual and autonomous surveying
- BRANCH robot design
 - **Rugged robot drivetrain** with tripod and phone
 - ARM CPU based brain with sensors for manual and autonomous driving on the farm



Custom-made
BRANCH Robot

Deployment in Farms

- Ongoing deployment in Washington State farms using BRANCH
 - **Nelson Farms** – Farmington, WA
 - **Dancing Crow Farm** – Carnation, WA (Organic)
- Partnering with SnoValley Tilth
 - Nonprofit organization of **80+** smallholder farmers
 - **1,000 acres** in Carnation, WA
- Meeting with farmers **weekly** for deployment
- Working towards grant proposal for **WSDA** and **USDA** funding



Explaining LeAF to farmer at Dancing Crow Farm in Carnation, WA

Discussion: LeAF Impact

- **Economic impact**
 - Up to **60% reduction** of pesticides → Billions of dollars saved
- **Environmental impact**
 - Millions of kgs of greenhouse gas emissions saved, soil health preserved
- **Ecological impact**
 - Reduced food contamination and fish and bird deaths

Limitations & Future Direction

- **Handling unseen pests** - 'unknown' class and expanded model
- **Dense/rough terrain farms** - expand to Aerial/Drone imagery
- **Predict pest infestations** - data from Soil Moisture/Humidity sensors
- Accepted at **top IEEE conference** on Artificial Intelligence

Conclusion

LeAF revolutionizes agriculture by minimizing costs, reducing chemical usage, and maximizing crop health with precise anomaly detection and treatment suggestions leading to significant savings for farmers and a positive environmental impact.

Acknowledgements

- Partnership with local farmers for deployment and feedback
- Organizers for creating research opportunities for high schoolers
- Judges and audience for your time and interest

Works Cited (Complete List in Paper)

He, K., Zhang, X., Ren, S., & Sun, J. 2015, December 10. Deep Residual Learning for Image Recognition.

iNaturalist: <https://www.inaturalist.org/>

Liu, S., Zeng, Z., Ren, T., Li, F., Zhang, H., Yang, J., ... Zhang, L. 2023, March 20. Grounding DINO: Marrying DINO with Grounded Pre-Training for Open-Set Object Detection.

Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., ... Sutskever, I. 2021, February 26. Learning Transferable Visual Models From Natural Language.

Redmon, J., Divvala, S., Girshick, R., & Farhadi, A. 2016, May 9. You Only Look Once: Unified, Real Time Object Detection.

Ritchie, H., Rosado, P., & Roser, M. 2022, December 2. Environmental Impacts of Food Production. Tudi, M., Ruan, H. D., Wang, L., Lyu, J., Sadler, R., Connell, D., Chu, C., & Phung, D. T. 2021, February 18.

Agriculture Development, Pesticide Application and Its Impact on the Environment. United Nations. Feeding the World Sustainably. <https://www.un.org/en/chronicle/article/feeding-world-sustainably> University, S. 2020, October 7.

Why Laughing Gas is a Growing Climate Problem. World Bank Group. 2019, April 26. For Up to 800 Million Rural Poor, a Strong World Bank Commitment to Agriculture.

THANK YOU!

Questions?

adityasngpta@gmail.com

+1 425-229-4439