Integrating Robotics, Computer Vision, and AI for Automated Stand Count and Vigor Assessment in Cornfields

Mr. Suseel Kumar Yerramsetti, University of Maryland Eastern Shore

Suseel Kumar Yerramsetti is currently a Ph.D. student in the Department of Food and Agricultural Sciences at the University of Maryland Eastern Shore (UMES). Mr. Yerramsetti received his bachelor's degree in Bioinformatics from Bharath University, Chennai, India, in 2011, and earned his master's degree in Computer Science from UMES in 2016. He brings an interdisciplinary background that bridges computing, engineering, and agricultural research. His doctoral research focuses on the application of artificial intelligence, robotics, remote sensing, and machine learning for precision agriculture. His work explores integrated, data-driven approaches to support crop monitoring, plant health assessment, yield prediction, and decision-support systems for sustainable crop management. In parallel with his academic research, Mr. Yerramsetti has more than ten years of professional experience in software engineering, application modernization, multi-cloud infrastructure, security governance, data engineering, Internet of Things (IoT) integrations, AI-driven analytics, and machine learning models. He currently works at IBM as a Solution Architect. He previously worked with The Walt Disney Company as a Senior Software Engineer and Cloud Engineer, and with Samsung as a Software Engineer. His broader interests lie at the intersection of artificial intelligence, cloud-based intelligent systems, predictive modeling, and agricultural production, with a focus on translating remote sensing, robotics, and machine learning advances into practical, accessible tools for farmers.

Dr. Abhijit Nagchaudhuri, University of Maryland Eastern Shore

Dr. Abhijit Nagchaudhuri is currently a tenured professor in the Department of Engineering and Aviation Sciences at the University of Maryland Eastern Shore(UMES). Dr. Nagchaudhuri received his baccalaureate degree from Jadavpur University (India) with honors in mechanical engineering in 1983. Thereafter, he worked in a multinational industry for a little over three years before joining Tulane University as a graduate student in the fall of 1987. He received a master's degree from Tulane University in 1989 and a doctoral degree from Duke University in 1992. He is a member of the American Society for Mechanical Engineers (ASME), the American Society for Engineering Education (ASEE), and, has also been involved with the American Society for Agricultural and Biological Engineers (ASABE). He is actively involved in teaching and research in the fields of (i) robotics and mechatronics, (ii) remote sensing and precision agriculture, and,(iii) biofuels and renewable energy. He has been involved with several extramurally funded projects as PI or Co-PI to support outreach, education, and research activities from various state, federal, and private agencies; served in leadership roles in professional societies; and has received several awards and certificates from his home institution as well as professional societies. He has published more than 90 refereed articles in journals, book chapters, and conference proceedings. He has been involved with several divisions of ASEE including the Minorities in Engineering Division (MIND), and served as the division chair of the Energy Conversion, Conservation, and Nuclear Engineering Division (ECCNED) in 2016. He is a member of the technical committee of IEEE/ASME Mechatronics and Embedded Systems Application(MESA) and has served as the division chair in 2022. He serves as a program committee member of the Maryland Space Grant Consortium.

Mr. Lance Ward, University of Maryland Eastern Shore

I am an undergraduate student researcher at the University of Maryland Eastern Shore currently working on a bachelor's degree in engineering with a concentration in computer engineering.

Ms. Divya vani Singamsetti, University of Maryland Eastern Shore

Divya vani Singamsetti is a Ph.D. student in the Department of Food and Agricultural Sciences at the University of Maryland Eastern Shore (UMES). She holds a Master's degree in Computer Science (2016) and a Bachelor's degree in Computer Science Engineering (2014) from Jawaharlal Nehru Technological University, Hyderabad (JNTUH), India. She has over five years of professional experience in software engineering, including application development and testing. Her research interests focus on artificial intelligence and robotics for agricultural applications. Her work centers on robotic crop monitoring and autonomous field navigation, supported by data-driven system design.

Mr. Jesu Raj Pandya, University of Maryland Eastern Shore

Currently a doctoral student in Food & Agriculture Sciences, University of Maryland Eastern Shore (UMES), MSc - Applied Computer Sciences at UMES, BSc - Electrical and Electronical Engineering (JNTU).

Interest in robotics and automation in food prod

Integrating Robotics, Computer Vision, and AI for Automated Stand Count and Vigor Assessment in Cornfields

Abstract

Accurate stand counts and early vigor assessment are critical for precision agriculture decisions such as replanting, variable-rate input management, and yield forecasting. Manual counting and scoring are time-consuming, prone to observer bias, and difficult to scale across large fields. This study presents a hybrid computer-vision pipeline that combines OpenCV-based vegetation segmentation with YOLO (You Only Look Once) v-8 verification to detect and count corn plants under variable field conditions. Data were collected using an AgileX Scout Mini ground robot equipped with an Intel RealSense D435i RGB-D camera and a ROS2 data capture workflow. The robot records fixed-length video segments of approximately 17 ft 5 in (5.33 m), representing 1/1000th of an acre at 30-inch row spacing, across two management zones (Green and Orange) that differ in agronomic conditions. Captured video segments are processed through a multi-stage workflow: candidate vegetation regions are segmented using OpenCV, YOLOv8 confirms plant instances, and temporal tracking enforces persistence to reduce false positives and prevent duplicate counting. In addition to stand counts, the method computes a composite vigor score using plant size, stem/morphological structure, and color-based metrics. The proposed approach achieved an average stand-count accuracy of 96.67% compared with seed maps. Results indicate a density–vigor trade-off between zones: the green zone showed higher plant density (30,330 plants/acre) with slightly lower mean vigor (56.61), while the orange zone had lower density (26,750 plants/acre) with slightly higher vigor (58.05). Overall, this hybrid OpenCV–YOLOv8 pipeline demonstrates a practical path for field-deployable stand counting and vigor scoring that can integrate into precision agriculture workflows for decision support.

1.0 Introduction

The ability to precisely quantify and collect stand counts and assess vegetation, such as vigor data, in the early stages of crop growth is a cornerstone of precision farming. These details are critical for making agronomic decisions, including whether to replant, adjust fertilization, or modify the watering system as per data. Traditionally, stand counts have been performed using the "1/1000th-acre methodology"[1]. Manual counting in farming, especially plant counts, row length, and extrapolating to per-acre quantitative calculations, has been proven successful in principle, but this methodology is labor-intensive, prone to sampling bias, and difficult to apply on a larger scale.

Simultaneously, vigor assessment, such as evaluating plant size, color, density, and structural robustness, has typically been a subjective visual exercise conducted by trained agronomists. The variability inherent in human assessment introduces inconsistencies, particularly across large teams or extended timeframes.

High-throughput phenotyping has advanced rapidly through imaging, sensors, and automated analysis pipelines. Image-based phenotyping enables scalable measurement of plant traits and supports breeding and agronomic decision-making [2-5]. Broader phenomics efforts emphasize the need to connect sensors to actionable agronomic knowledge [6, 7]. In parallel, deep-learning object detectors (e.g., YOLO-family models [8]) have been applied to agricultural scenes, but field conditions such as illumination changes, occlusion, and background clutter still create reliability challenges. Reviews of agricultural computer vision literature highlight that practical systems often benefit from combining traditional computer vision steps (e.g., segmentation and filtering) with deep learning for robust verification [9-11].

2.0 Objective

Existing automated plant stand counting systems mostly rely on two different approaches: deep learning techniques or traditional computer vision, such as OpenCV. Deep learning approaches require large, annotated datasets and intensive computing resources for training, whereas computer vision-only systems face issues of reduced robustness in variable lighting, obstruction, or growth stage conditions. This study addresses the need for a hybrid approach that leverages the interpretability and efficiency of traditional CV for the initial candidate systematic recognition. YOLOv8 was employed for consistent, reliable, and systematic confirmation under different field conditions, and temporal tracking was integrated to ensure persistence-based tracking and confirmation of plants. Furthermore, the system produces not only counts but also detailed vigor metrics of a plant.

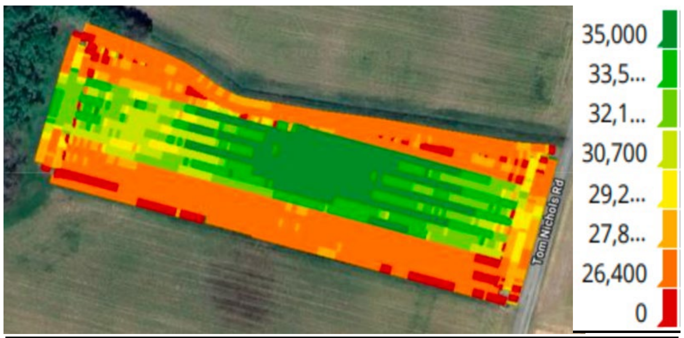

Figure 1: Green and Orange Seeding Zones

The primary objective of this study was to initiate the development a farm-deployable plant and vegetation tracking system by integrating a conceptual systemic structure combining the OpenCV system with YOLOv8 object recognition. Temporal tracking was implemented to eliminate false positives and duplicate counting of plants. A composite vigor scoring system structure incorporating morphological, structural, and color-based features of a plant was created. Finally, the study aimed to evaluate operational effectiveness under two different farming conditions (Green and Orange zones with seeding variability, see Figure 1) and determine vegetation density and vigor.

3.0 Materials and Methods

Field data were collected using an AgileX Scout Mini ground robot equipped with an Intel RealSense D435i RGB-D camera (Figure 2) . The D435i provides RGB and depth streams (depth range approximately 0.2–10 m) and includes an IMU to support orientation tracking during motion. The camera can capture high frame rates (up to 90 fps), enabling stable sampling under field movement and wind. The robot ran ROS2 Foxy on an onboard computer. A custom-built ROS2 node was utilized to obtain RGB video frames, depth frames, and the odometry of the wheels. Odometry-based triggering was used to automatically record fixed-length video segments (17 ft. 5 inch). Each segment was labeled using a standardized naming convention (plot ID, segment number, timestamp) to support repeatable analysis and comparison across sampling locations. Two cultivation zone conditions were targeted: the Green Zone (G) with a higher seeding rate and the Orange Zone(O) with a lower seeding rate (Figure 1).

Figure 2: AgileX Scout Mini with Intel RealSense D435i

Seven random segments were collected from these two zones, including G segments rs_005, rs_011, and rs_012, and O segments rs_019, rs_020, rs_021, and rs_023. Several challenges were faced during data collection. These included lighting variations owing to cloud movement and time-of-day differences which affect color-based segmentation; plants with wider leaves that overlap in later growth stages mimicking stems causing partial obstruction; weed interference from small weeds with broad leaves occasionally mimicking crop signatures, leading to false positives; and wind-induced motion which blurred some frames and temporarily obscured stems with leaves, as well as camera vibrations from uneven terrain causing positional noise. The data processing pipeline was designed as a multilayered approach by integrating both traditional computer vision (OpenCV) methods and deep learning (YOLOv8) to achieve highly precise detection performance under different field conditions (Figure 3). This multilayered method also provides the flexibility to adjust the parameters for different growth stages and field conditions while maintaining robustness against obstruction, lighting changes, and weed interference. Each collected video segment of approximately 17 ft 5 in (1/1000th acre) was processed through a sequence of operations. First, RGB frames were extracted from the recorded video at 15–30 fps, and depth frames were stored in parallel for potential use in later 3D reconstruction. Next, vegetation and stem segmentation using OpenCV involved HSV-based masking to identify vegetation pixels and morphological filtering to emphasize plant stems and reduce noise. This was followed by plant detection using a trained YOLOv8 model, with detections cross-verified against the OpenCV-based stem candidates. Temporal tracking was then implemented to ensure persistent plant tracking across consecutive frames to avoid duplicate counting. Vigor metrics were computed using per-plant measurements (size, stem, and color) via pixel-based calculations to determine a composite vigor score. Finally, plant counts and vigor summaries were exported to CSV files for further analysis and validation, and annotated videos were created for visual verification. For the identification process, each frame was

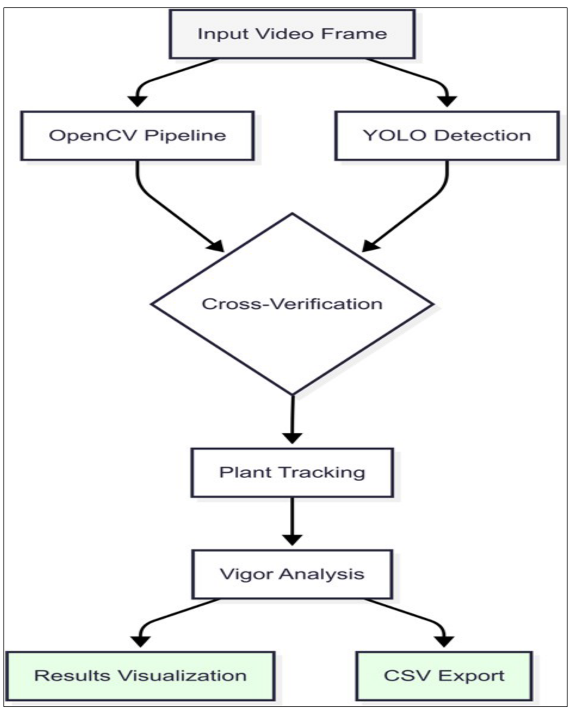

Figure 3: Processing Pipeline (Computer Vison-Deep learning 'CV-DL')

converted from RGB to the HSV color space, which distinguishes chromatic content from intensity and provides robustness against changing lighting conditions. Vegetation masking parameters were optimized for early to mid-growth stages with a Hue range of 32–88, Saturation of 28–255, and Value of 28–255. Morphological operations included closing with a 7x5 kernel to fill small gaps and opening with a 5x5 kernel to remove isolated noise pixels. These steps produced a binary vegetation mask capturing leaf area while excluding background pixels. As corn plants have a strong vertical stem structure, a vertical morphological kernel (15x3) was applied to enhance stem features. Stem candidate filtering used specific criteria: a minimum height of at least 48 pixels, a height-to-width ratio of at least 2.0, and a stem-pixel density threshold of at least 0.3 (ratio of stem pixels to bounding box area). These parameters helped rule out low-lying weeds and horizontal leaf clusters, focusing only on the upright growth form of corn stems (see Figure 4). The YOLOv8 (You Only Look Once, Version 8 [8]) object detection model was integrated as a secondary validation step for the OpenCV detections. YOLOv8n (nano) was used for edge deployments on devices like the Scout Mini with limited computing power, while YOLOv8s (small) was used for more accurate detections on a GPU-enabled laptop. Training and fine-tuning utilized base weights from (**C**ommon **O**bjects in **Co**ntext) COCO-pretrained models [12], transfer learning with a

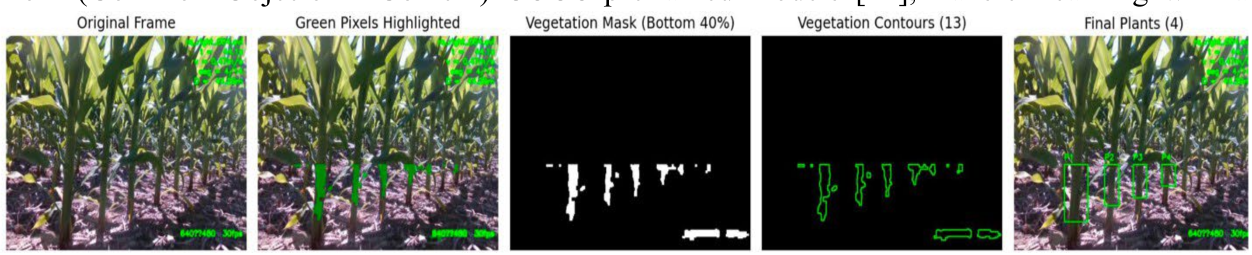

Figure 4: Visualization and Stem Detection Pipeline

custom dataset of 2,000 annotated corn plant images, and specific inference parameters including a confidence threshold of 0.25 and an Intersection over Union (IoU) threshold of 0.45, with a class filter for "plant" only. The final detection set for each frame was determined by comparing OpenCV-based stem detections and YOLOv8 detections using an Intersection-over-Union (IoU) threshold of 0.5. If both methods detected the same plant (IoU $\geq$0.5), it was accepted as cross-verified. If YOLOv8 detected a plant not found by OpenCV, it was included only if its confidence score exceeded 0.5. If OpenCV found a plant not confirmed by YOLOv8, it was included only if its stem density exceeded 0.4. This condition-based approach minimized both false positives and false negatives. Field testing revealed that parameters needed adjustment according to the corn plant growth stage [13] for accurate results. For the early stage (V2–V4), a lower stem density threshold (0.25) and wider HSV hue range were used. For the mid stage (V6–V8), a stricter stem height requirement ($\geq$ 60 px) and adjusted morphological kernels were implemented due to higher leaf overlap. In the late stage (Pre-Tassel), reliance on YOLOv8 cross-verification increased due to obstruction, and the role of stem detection was reduced in favor of whole plant bounding boxes. The primary challenge in plant detection from videos is to avoid duplicate counting as the camera moves along the row. Inconsistent or false detections caused by obstructions, lighting changes, or model errors can lead to incorrect counts. To resolve these issues, a temporal tracking module was implemented that creates and maintains plant identities across consecutive frames of the video. Tracking objectives included creating a unique ID for each plant, distinguishing between new and existing plants, handling obstructions without deleting IDs prematurely, and feeding per-plant temporal data into the vigor scoring module. The tracking algorithm operates after the hybrid detection stage. Detection extraction involves using bounding boxes and center coordinates.

Spatial association uses a distance-based matching algorithm with a threshold of 45 pixels. Overlap verification requires at least 30% IoU. ID assignment follows the match results, and a persistence rule requires a plant to exist in at least seven consecutive frames to be confirmed. Lost frame tolerance allows a plant ID to remain active even if it disappears for up to 10 frames. For each tracked plant ID, the system logs bounding box coordinates per frame, pixel height/width, vegetation area, color statistics, stem indicator, and per-frame vigor components. These time-series features support robust averaging and segment-level summaries. The temporal tracking approach reduces duplicate counts, improves robustness to environmental changes, and stabilizes vigor estimation by averaging measurements across frames. The vigor scoring component converts raw per-plant data into a quantitative measure of plant health on a 0–100 scale. This allows for easy comparison of plant health across conditions. Vigor scoring combines plant size, structure, and canopy color quality into one objective, repeatable index. The composite vigor score was calculated as a weighted sum of five component scores:

$$\text{Vigor Score} = w_{!}\,S_{\text{size}} + w_{"}S_{\text{height}} + w_{\#}\,S_{\text{color}} + w_{\$}S_{\text{density}} + w_{\%}\,S_{\text{stem}} \qquad (1)$$

Where:

- S_{size} = Normalized size score from plants canopy area (pixels²)
- S_{height} = Normalized height score from plants height (pixels)
- S_{color} = Color health score based on hue, saturation and color uniformity
- S_{density} = Canopy density score reflecting leaf fullness
- S_{stem} = Stem robustness score based on stem presence and thickness

The weights were selected based on empirical trials and agronomic input, and component scores were calculated and then averaged across frames. Vigor components are interpreted with the growth stage in mind, and final scores are mapped to qualitative categories for interpretation.

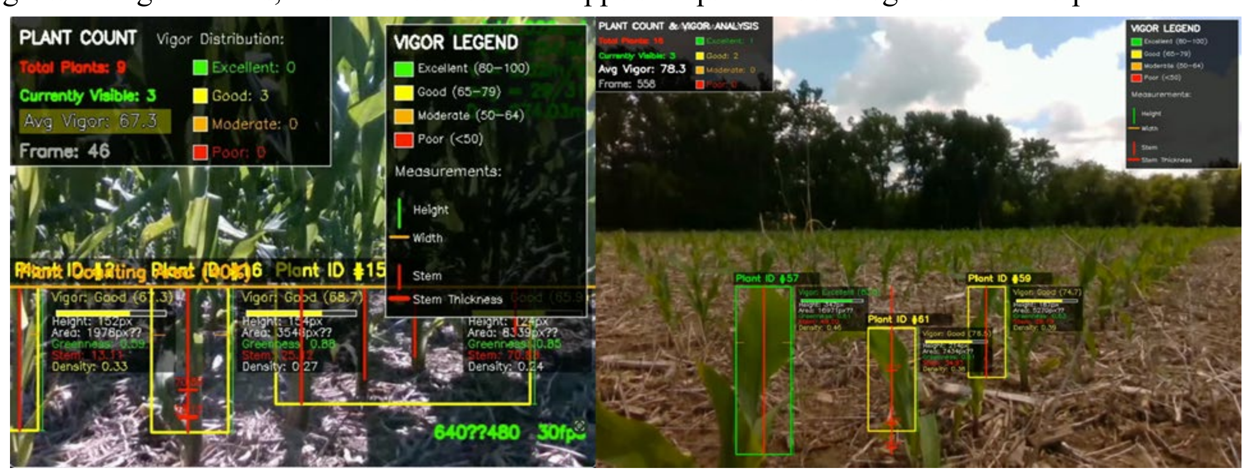

Figure 5: Annotated Video Frame Showing Detected Plants with Vigor-based Color Coding.

For each 17 ft. 5 in. segment, the system generates a summary CSV file with the average vigor score, a detailed CSV file with all plant data, and an annotated video with color-coded bounding boxes. In corn crops, stand counts are traditionally converted to plants per acre using the 1/1000th acre measurement [11]. An acre is 43,560 square feet, so 1/1000th is 43.56 square feet. For 30-

inch row spacing, the corresponding row length is approximately 17.42 ft, or 17 ft 5 inch. By collecting video segments of this length, each segment directly represents 1/1000th of an acre. For each segment, confirmed plant IDs from the temporal tracking module are counted once. Segment vigor is the mean of all per-plant vigor scores within that segment. Plants per acre are calculated as plants per segment multiplied by 1000. Zone-level aggregation averages these metrics across all segments for comparison between zones. This links density to vigor, supporting management decisions. Vigor components are interpreted with growth stage in mind, because early stages emphasize canopy size and color, while later stages rely more on structure and stem robustness. Final vigor scores are mapped to qualitative categories for interpretation. Figure 5 illustrates a sample annotated video frame showing detected plants with vigor based color coding.

4.0 Results and Discussion

The following results were obtained from seven 1/1000th-acre segments collected across the two cultivation zone conditions. The G zone condition produced approximately 13.4% higher plant density than the O zone, while the O zone condition produced plants with approximately 2.5% higher average vigor scores. The measured plant counts matched the seed application rates with an average accuracy of 96.67%. Data analysis revealed an inverse relationship between plant density and average vigor score, illustrating a density–vigor trade-off between the zones (as shown in Tables 1 and 2 below)

Table 1: Per-Segment Plant Counts, Plants/ Acre, And Average Vigor Scores

Sample ID	Zone	Plants (Row Count)	Plants/acre	Avg. Vigor Score
rs_005	G	29	29,000	55.49
rs_011	G	31	31,000	57.21
rs_012	G	31	31,000	57.14
rs_019	O	26	26,000	57.63
rs_020	O	27	27,000	58.12
rs_021	O	26	26,000	57.75
rs_023	O	28	28,000	58.72

Table 2: Mean of Zone Plants Counts/Segment, Plants/ Acre, Vigor Scores

Zone	Mean Plants per Segment	Plants/acre	Mean Vigor Score
G	30.33	30,330	56.61
O	26.75	26,750	58.05

The results of this study provide both agronomic insights and technical validation for the hybrid computer vision – deep learning (CV-DL) pipeline deployed in operational corn field conditions. The clear density-vigor trade-off is consistent with increased competition for resources at higher populations. The hybrid pipeline achieved high stand count accuracy, with cross-verification and temporal tracking improving reliability. The system's strengths include its hybrid workflow, adaptability, real-time potential, and comprehensive output. Limitations include lighting variability, dense canopy obstruction, weed interference, computational load, and the need for stage-specific parameter tuning.

Although the system demonstrated strong overall performance, several limitations and challenges remain. Variability in lighting conditions, particularly sudden changes in sunlight, can alter color values and negatively affect vegetation masking as well as YOLO detection confidence scores. As the crop progresses into later growth stages, dense canopy formation and overlapping leaves can obstruct stem visibility, reducing the accuracy of stem-based detection and increasing reliance on the YOLO model. In addition, small weeds that share similar color and structural characteristics with young corn plants occasionally pass the detection filters, leading to false positives. The computational demands of running YOLOv8 in real time also present a challenge, as effective deployment typically requires GPU acceleration and necessitates lightweight model variants for edge-based systems. Furthermore, achieving optimal accuracy currently requires manual parameter tuning across early, mid, and late growth stages, highlighting the need for an automated, growth-stage-aware adaptation mechanism to improve usability.

When compared with existing approaches, the proposed hybrid pipeline offers several advantages. Relative to deep learning–only methods, it reduces dependence on large, labeled datasets while maintaining interpretability through explicit segmentation and rule-based verification. In contrast to traditional OpenCV-only approaches, the inclusion of YOLO-based verification improves robustness under real-world field variability that often violates strict color and shape assumptions. Compared with manual 1/1000th-acre plant counts, the system minimizes human bias, enables broader spatial sampling in significantly less time, and provides additional plant health and structural metrics beyond conventional population estimates.

Building on these findings, future improvements will focus on increasing robustness, scalability, and autonomy. Planned enhancements include adaptive parameter adjustment to automatically tune HSV thresholds, stem height filters, and YOLO confidence levels in response to changing lighting conditions and crop growth stages. Incorporating 3D plant structure analysis using RealSense depth data would enable more accurate estimation of plant height, volume, and biomass. Expanding the YOLO model to support multi-class detection would allow simultaneous identification of weeds, pest damage, and disease symptoms alongside crop plants. Finally, optimizing the pipeline for edge computing through deployment of lightweight YOLO variants on embedded GPU platforms would support fully autonomous, real-time field operation.

5.0 Conclusion and Educational Implications

This work demonstrates how contemporary artificial intelligence tools—specifically modern object detection frameworks initialized with large-scale pretrained models—can be meaningfully integrated into engineering and STEM education and research while addressing real-world challenges in agriculture and autonomous systems. Beyond the technical performance of the detection pipeline, the project highlights an instructional approach that bridges theory, data, and application, offering participants an authentic learning experience grounded in current practices.

From an educational perspective, the use of COCO-pretrained YOLOv8 models provides a practical entry point for students to engage with state-of-the-art computer vision without the prohibitive data and computational requirements associated with training deep networks from scratch. By leveraging transfer learning, students are able to focus on higher-level concepts such as dataset curation, model adaptation, threshold selection, validation, and deployment trade-offs. This shift in emphasis—from algorithmic derivation alone to system-level reasoning—aligns well with ASEE's emphasis on preparing students for modern, interdisciplinary roles.

The project also reinforces the value of contextual learning. Framing object detection within the domain of crop monitoring, phenological assessment, or field robotics and instrumentation allows students to connect abstract concepts in machine learning to tangible outcomes such as plant growth stage identification, stress detection, or autonomous navigation. This domain grounding enhances student motivation and supports deeper conceptual understanding, particularly for learners who benefit from applied, problem-driven instruction. Importantly, the approach is flexible and can be adapted to other application areas, making it suitable for integration across multiple engineering and/or STEM courses, laboratories, or capstone experiences.

In addition, the work supports inclusive and STEM education. The use of open-source tools, publicly available pretrained models, and modest hardware requirements lowers barriers to participation, enabling implementation at minority-serving institutions and resource-constrained programs. Students with diverse academic backgrounds—including those with limited prior exposure to artificial intelligence—can meaningfully contribute to the project through data annotation, experimental design, performance analysis, and system integration. Such distributed engagement fosters teamwork, builds confidence, and supports equitable participation in emerging technology areas.

From a workforce development standpoint, the skills emphasized in this work—model evaluation, threshold tuning, domain adaptation, and ethical use of pretrained systems—are directly transferable to industry and research settings. Students gain experience not only in using advanced AI tools, but also in critically assessing their limitations, an increasingly important competency as pretrained and foundation models become ubiquitous across STEM disciplines. Exposure to these considerations prepares graduates to be thoughtful practitioners who can responsibly deploy AI-enabled systems in societally relevant contexts.

By situating AI-based perception within agricultural and environmental applications, the project illustrates how engineering solutions can contribute to sustainability, food security, and data-driven decision making. Such framing aligns strongly with ASEE's mission to advance engineering and STEM education that is socially responsive, interdisciplinary, and forward-looking.

6.0 Acknowledgments

The authors would like to acknowledge the USDA NIFA Evans Allen Grant, "Contemporary and Smart Farming Technologies to Advance Learning Discovery and Engagement", NASA MSTAR Grant # 80NSSC23M0206 titled "DREAM: Developing Robotic Explorations with AgroBots and MoonBots", the UMES efforts of the 'Climate Smart Farming' integral to "The 1890 Foundation Center of Excellence for Emerging IoT Technologies and Smart Farming " and the "AIRSPACES (Autonomous Instrumented Robotic Sensory Platforms to Advance Creativity and Engage Students)" project supported by Maryland Space Grant Consortium.

Bibliography

[1] Iowa State University Extension and Outreach, "Estimating Corn Plant Populations," *Integrated Crop Management*, 2021. [Online]. Available: https://crops.extension.iastate.edu. Accessed: Dec. 17, 2025.

[2]N. Fahlgren, M. A. Gehan, and I. Baxter, "Lights, camera, action: High-throughput plant phenotyping is ready for a close-up," *Current Opinion in Plant Biology*, vol. 24, pp. 93–99, 2015.

[3]J. L. Araus and J. E. Cairns, "Field high-throughput phenotyping: The new crop breeding frontier," *Trends in Plant Science*, vol. 19, no. 1, pp. 52–61, 2014.

[4]R. Khanna, L. Schmid, A. Walter, J. Nieto, R. Siegwart, and F. Liebisch, "A spatio-temporal spectral framework for plant stress phenotyping," *Plant Methods*, vol. 15, no. 1, pp. 1–18, 2019. [Online]. Available: https://plantmethods.biomedcentral.com/articles/10.1186/s13007-019-0398-8. Accessed: Dec. 17, 2025.

[5]S. Madec, F. Baret, B. De Solan, S. Thomas, D. Dutartre, S. Jezequel, *et al.*, "High-throughput phenotyping of plant height: Comparing unmanned aerial vehicles and ground LiDAR estimates," *Frontiers in Plant Science*, vol. 8, Art. no. 2002, 2017, https://doi.org/10.3389/fpls.2017.02002.

[6]F. Tardieu, L. Cabrera-Bosquet, T. Pridmore, and M. Bennett, "Plant phenomics, from sensors to knowledge," *Current Biology*, vol. 27, no. 15, pp. R770–R783, 2017.

[7]W. Yang *et al.*, "Crop phenomics and high-throughput phenotyping: Past decades, current challenges, and future perspectives," *Molecular Plant*, vol. 13, no. 2, pp. 187–214, 2020.

[8]G. Jocher, A. Chaurasia, and J. Qiu, "YOLO by Ultralytics," ver. 8.0.0, computer software, 2023. [Online]. Available: https://github.com/ultralytics/ultralytics. Accessed: Dec. 17, 2025.

[9]H. Tian, T. Wang, Y. Liu, X. Qiao, and Y. Li, "Computer vision technology in agricultural automation—A review," *Information Processing in Agriculture*, vol. 7, no. 1, pp. 1–19, 2020, https://doi.org/10.1016/j.inpa.2019.09.006.

[10]W. Guo, B. Zheng, A. B. Potgieter, J. Diot, K. Watanabe, K. Noshita, *et al.*, "Aerial imagery analysis – quantifying appearance and number of sorghum heads for applications in breeding and agronomy," *Frontiers in Plant Science*, vol. 9, Art. no. 1544, 2018, https://doi.org/10.3389/fpls.2018.01544.

[11] D. Wu, S. Lv, M. Jiang, and H. Song, "Using channel pruning based YOLOv4 deep learning algorithm for the real-time and accurate detection of apple flowers in natural

environments," *Computers and Electronics in Agriculture*, vol. 178, Art. no. 105742, 2020, https://doi.org/10.1016/j.compag.2020.105742.

[12] T.-Y. Lin, et al, "Microsoft COCO: Common objects in context," Proceedings of European Conference on Computer Vision (ECCV), 2014, pp. 740–755.

[13] S.W., Ritchie, et al, (1993). *How a corn plant develops* (PM 1666). Iowa State University Extension.