OXFORD JOURNAL
OF STUDENT
SCHOLARSHIP

Low-Cost RGB Imaging for Plant-Stress Detection in California: A Benchmark-to-Field Evaluation

Anik Koirala, Adit Shah & Ian Chan
anikoirala16@gmail.com, adit.shah0123@gmail.com, ian_chan@myabcusd.org

ABSTRACT

Insect pests and other stressors cost California agriculture billions of dollars each year, and early detection is one of the most effective defenses. AI-based detection systems that pair deep learning with multispectral or hyperspectral imagery report roughly 90–95% accuracy, but they require sensors costing $10,000 or more, large labeled datasets, and technical expertise that most small to medium-sized farms lack. This study asks whether a consumer RGB drone (about $400) combined with a training-free, color-based computer-vision algorithm can detect visible plant stress well enough to serve as a first step toward AI-based monitoring. The algorithm segments foliage in HSV color space, computes the fraction of foliage that is discolored, and labels a plant as stressed when that fraction reaches 0.20. On 2,012 plants annotated in drone video from four types of Southern California sites, it reached 70.1% accuracy but only 39.0% stress recall, missing 347 of 569 stressed plants. On all 54,305 images of the PlantVillage leaf benchmark, the same rule ranked diseased leaves above healthy ones well (AUC 0.814) but flagged only 17.6% of them, because most leaf lesions are dark rather than yellow; a classifier trained on RGB color-index features reached 96.7% accuracy on held-out PlantVillage images. The same algorithm's ability to separate stressed plants from healthy plants fell from AUC 0.814 on curated leaves to 0.645 in the field, showing that curated benchmarks overstate field readiness and that visible color alone misses subtle stress. The algorithm measures visible plant stress only; it does not detect insects or determine whether stress was caused by insect damage. We therefore utilized low-cost RGB imaging as a Phase-1 baseline that progresses towards CNN architecture (Phase 2) and that ultimately culminates into multispectral sensing with a proposed ResCNN–HDBSCAN model (Phase 3).

Keywords: *precision agriculture; plant stress detection; RGB imaging; unmanned aerial vehicles; benchmark-to-field gap*

INTRODUCTION

California's farms and ranches received $61.2 billion in cash receipts in 2024, and the state grows nearly half of the nation's vegetables and over three-quarters of its fruits and nuts [35]. At this scale, invasive insects carry undoubtly, large financial consequences. The Center for Invasive Species Research at UC Riverside mentions that, according to the California Department of Food and Agriculture (CDFA), the economic agricultural losses due to invasive species and insects exceed over $3 billion annually through crop damage, continuous treatment, and quarantine compliance [15].

The 2005 Diaprepes root weevil infestation in Orange and Los Angeles counties illustrates these costs. Jetter and Godfrey [1] found that treating established infestations averaged $609 per acre for citrus and avocado growers, raising costs by 21.61% for orange growers and 11.35% for avocado growers. The treatment program was cancelled for lack of funding in July 2008, and quarantine alone did not stop the spread; in Florida, movement of infested plants carried the weevil into 22 counties. Cases like this make early detection of insect damage a priority for California growers.

Conventional pest management relies on ground scouting and scheduled chemical treatment, which are labor-intensive and costly across large farmland, and early detection is difficult both operationally and economically [16][17]. Unmanned aerial vehicles (UAVs) offer wide-area coverage at lower cost, but many deployed systems collect imagery without integrated analysis, so the imagery is not turned into actionable detections [3]. AI-based systems close this gap in research settings (see Related Work), yet their cost and data requirements keep them out of reach for most small and medium operations.

This paper uses three terms with fixed meanings. Plant stress is any visible decline in plant health, such as yellowing, browning, or wilting, regardless of cause. Insect damage is plant stress caused specifically by insect feeding or infestation; it is one cause of plant stress among others, including drought, nutrient deficiency, and disease. Pest detection is identifying or counting the insects themselves. This study's algorithm only measures plant stress. It does not detect pests and does not determine whether a stressed plant has insect damage. Pest detection and insect-damage attribution are goals of the later phases described in the Discussion.

This study asks the following question: *Can a consumer RGB drone and a training-free color algorithm detect visible plant stress with meaningful accuracy under field conditions?* We hypothesized a field accuracy of 70–80%, and this value can be a practical benchmark for further advancements in AI modeling. The study makes the following contributions:

- A reproducible, training-free RGB pipeline that is used for detecting stress from analyzing the external conditions of the plants from consumer drone footage (Methods).
- An evaluation of the same algorithm on a curated leaf benchmark and on field imagery. A trained color-index classifier is also used to show how much curated datasets overstate field performance (Results and Discussion).
- A staged progression from low-cost RGB sensing and monitoring towards AI-integrated pest detection that involves integrating Phase-3 architecture (Methods and Discussion).

The Related Work section analyzes previous systems and identifies the gap to be insufficient evidence of using a low-cost RGB framework in real world field conditions, comparing that to curated datasets. The Methods describe the data, algorithm, and evaluation protocol, followed by the Results using data analysis, a Discussion of the results against the current literature and cost evaluations, and the Conclusion.

RELATED WORKS

Deep-Learning Pest and Disease Detection

Deep-learning systems report high accuracy for pest detection and disease identification. Gokeda and Yalavarthi's deep hybrid model for pest detection (DHMPD) reached a classification accuracy of 0.936 in experimental agricultural environments, using Internet of Things (IoT) connectivity so that models can act on data as it is being collected [2]. Khujamatov et al.'s Advanced Insect Detection Network (AIDN) uses multi-scale feature fusion to detect medium-to-large insects with bounding boxes, reaching accuracy up to 0.92 and outperforming YOLOv4, SSD, and R-CNN by a reported 10–15% [7]. AIDN weights the regions of each frame most likely to contain insects using spatial attention:

$$M_s(F) = \sigma(f^{7\times7}([\text{AvgPool}(F); \text{MaxPool}(F)])) \quad (1)$$

where F is the feature map, AvgPool and MaxPool are taken across channels, the semicolon denotes concatenation, $f^{7\times7}$ is a convolution with a 7×7 kernel, and σ is the sigmoid function, which scales each location's attention weight to between 0 and 1. The large kernel widens the receptive field, which helps with medium-to-large moving insects at some cost to fine detail [7].

Other systems target specific settings. Zhang et al.'s pest density estimation framework analyzes small patches of drone imagery and pools them into density maps that estimate pest counts rather than species [6]. Juan et al.'s Qpest-RCNN counts tiny pests on stationary sticky traps with high accuracy [8], and Stojnić et al. reduced motion-related errors for small moving objects such as honeybees by training a CNN on synthetic video [10]. Yu et al.'s three-dimensional residual CNN (3D ResCNN) detected pine wilt disease from UAV hyperspectral imagery with an accuracy of 0.8106 while using only one-fifth of the training data [11], and Wu et al. improved region-based CNN segmentation of pests on crop leaves relative to Mask R-CNN [21]. Across these systems, a recurring tradeoff appears: YOLOv4 struggles with very small objects, AIDN's wide receptive field reduces fine detail, Qpest-RCNN is designed for stationary targets, and ResCNN keeps accuracy with limited data at the cost of high computation and training time [7][8][11]. All of these systems detect pests or specific diseases, which is a different task from the generic plant-stress detection evaluated here.

UAV Sensing Modalities

Sensor choice strongly affects what a UAV system can detect [25]. RGB cameras record visible light (about 380–780 nm); they are inexpensive, lightweight, and well suited to shape and texture, but visible color alone separates healthy and stressed plants poorly, because stress often changes color only after damage is advanced. Multispectral cameras add a few broader bands, typically including near-infrared, where the reflectance of stressed tissue changes before visible symptoms appear; they are sensitive to sunlight and weather, which adds noise. Hyperspectral cameras record many narrow bands and support early detection of decay: Nguyen et al. detected plant viral disease with confidence above 95% [4] using a SPECIM IQ sensor with 1.07 mm spatial resolution at 1 m [12], and identified individual sick fir trees in insect-infested forests [5]. Hyperspectral data, however, is costly to collect and computationally

demanding to process. Thermal infrared cameras (about 3–14 μm) detect abnormal canopy temperature linked to water stress and possible infestation, but require image fusion and geometric alignment with other sensors; Subramanian et al. report their use in pest management in South Indian rice fields [13]. Zhu et al. note that combining RGB, multispectral, hyperspectral, and thermal sensors captures the most complete crop information at low altitude, but the added payload and processing sharply limit flight time [14].

IoT and Closed-Loop Monitoring

IoT systems combine many sensors whose data is used to adjust or maintain a system's state [16]. Bhoi et al. transmitted UAV imagery of certain rice fields to a cloud-based machine-learning model that identified and located pests as well as automatically notifying the farmers [3]. Almalki et al. tested a low-cost IoT and UAV platform that measured temperature and humidity on a Tunisian farm [17], and Azfar et al. built an IoT pest-detection and response system for cotton that improved overall yields whilst limiting environmental degradation and impact [18]. These studies, together, suggest that AI-enabled IoT systems can ensure that monitoring that is being done is within a closed loop in which detections activate and enable automatic action and intervention.

Representation Learning and Density Clustering

Our proposed Phase-3 architecture (see Methods) draws on work that pairs learned image features with clustering. McInnes and Healy's accelerated HDBSCAN finds clusters of varying density and labels outliers as noise (−1), making it more practical than non-hierarchical DBSCAN for noisy spatial data, though it slows on very large datasets without GPU acceleration [24]. Hasana and Fitrianah reduced high-dimensional features before density clustering to map fire zones in Jakarta [23], and Yang et al. presented that variational-autoencoder embeddings improve clustering accuracy [22]. Zheng et al. combined a 3D attention convolutional autoencoder with a clustering layer for hyperspectral image analysis [26]. Residual CNNs have been used to map agriculture from satellite imagery [20] and could identify faults and abnormalities accurately in imbalanced solar-panel data [19], supporting their use as possible extracting tools of features when the actual labeled data is limited.

Accessibility Barriers and the Research Gap

The aforementioned systems share similar limitations that reduce their possibility of being adopted by small and medium farms. Multispectral drone systems typically cost between $10,000–15,000, while consumer RGB drones cost a few hundred dollars. Training a reliable deep-learning model requires a large labeled dataset for each of the crops analyzed and the variables being calculated, which often requires heavy GPU or cloud computing as well as advanced technical skills. Most of the reviewed studies also report accuracy under experimental or curated conditions. Few, among the studies we reviewed, test whether a low-cost, training-free RGB method can be used as a curated benchmark to real drone imagery. This study addresses this gap in the literature by measuring the same algorithm in both settings.

MATERIALS AND METHODS

Study Design

The study had two components. First, a structured literature review identified advances in UAV monitoring, AI, and pest detection relevant to Southern California. Second, a training-free RGB plant-stress detection algorithm was developed and evaluated on a public benchmark and on drone imagery collected by the authors; a classifier trained on RGB color-index features was also tested on the benchmark as a reference for how much signal color carries when a model can learn from labeled examples. Results were then compared with accuracies reported in the literature to assess the cost–capability tradeoff. The last Methods subsection describes a proposed, untested Phase-3 architecture.

Literature Review Protocol

Sources were searched on academic resources like EBSCO, ScienceDirect, and MDPI using the keywords "UAV," "agriculture," "CNN," "pest detection," "AI," "insect damage," "computer vision," "multispectral," "precision agriculture," and "deep learning," focusing on work published between 2010 and 2025, with earlier works included if deemed necessary. Screening followed the stages described in the PRISMA 2020 statement [9], which included duplicate removal, title and abstract screening, and full-text evaluation. Record counts for each of the stages were not preserved, so this can be seen as a more structured narrative review rather than a full systematic review. Inclusion required a target crop or insect, quantified UAV or AI performance, statistical analysis of the outcomes made, and for the article to be peer reviewed. Exclusion removed outdated technologies, smaller sample sizes, focused areas that were not similar to the environments of Southern California, and work unrelated to UAV or precision-agriculture monitoring. Twenty-seven peer-reviewed sources were retained and assessed for sample-size adequacy, similarity to the study setting, design quality, statistical rigor, and recency.

Equipment and Field Data Collection

Imagery was collected with a Potensic Atom 2, a manually controlled quadcopter under 249 g with an integrated 4K RGB camera recording at 30 frames per second, costing about $400 [34]. Flights took place in Southern California on July 6 and July 13, 2026, at four site types chosen to cover a range of species, growth stages, lighting, and management: a school garden, residential properties with fruit trees and ornamental plantings, a nursery, and public landscape areas with mixed vegetation. Flights were flown at 3–5 m altitude at varying times of day. Footage was recorded one continuous video stream rather than a fixed survey grid, creating a total of 27 clips. Figure 4 shows annotated frames of three types of potted plants.

Plant-Stress Detection Algorithm

The algorithm classifies each plant from visible color alone, with no training and with all parameters fixed before evaluation (Figure 1). It was implemented in Python [29] using OpenCV [28] for image processing

and NumPy [33] for array operations. The algorithm classifies plants inside the pre-annotated bounding boxes (refer to Field Annotation) done by us; it does not attempt to locate plants on its own.

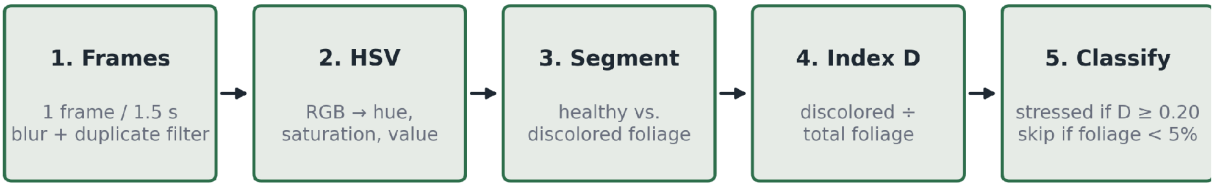


Figure 1: The five-step plant-stress detection pipeline. No step is trained on labeled data; all thresholds were set before evaluation.

Step 1: Frame extraction and quality control. Each video (~30 fps) was sampled at one frame every 1.5 s (~0.67 fps) with FFmpeg [30]. Frames that were motion-blurred were subsequently removed (low variance of the Laplacian) and/or if they were near duplicates of the previous kept frame (low mean absolute difference between 32×32 grayscale thumbnails). Remaining frames were downscaled and sized to 1,600 pixels. This created a total of 548 frames from the 27 clips, with an average similarity of 86.5% between consecutive frames.

Step 2: Color transform. Each plant region was converted from RGB to HSV color space, which separates hue from saturation and brightness and is less sensitive than raw RGB to changes in illumination between videos.

Step 3: Foliage segmentation. Healthy foliage was defined as hue 78–170° with saturation above 0.18. Discolored (chlorotic or necrotic) foliage was defined as hue 45–72° with saturation above 0.40 and value above 0.45 (Figure 2). The saturation and value limits on the discolored class exclude soil, mulch, and dead plant material, which share the yellow–tan hue range but are duller and darker. Grey hardscape, shadow, and most non-foliage flowers fall outside both classes, although bright yellow flower parts do not (see Error Patterns). Total foliage is the union of the two masks after a morphological opening to remove speckles. Hue is given in degrees (0–360°); in OpenCV's 8-bit HSV format, the equivalent values are half these numbers, and saturation and value are scaled to 0–255.

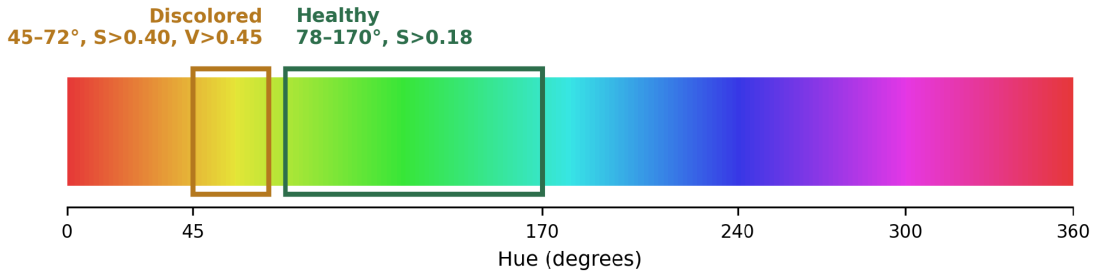


Figure 2: Hue ranges used to separate healthy and discolored foliage in Step 3. Pixels outside both boxes, or below the saturation and value limits, are treated as background.

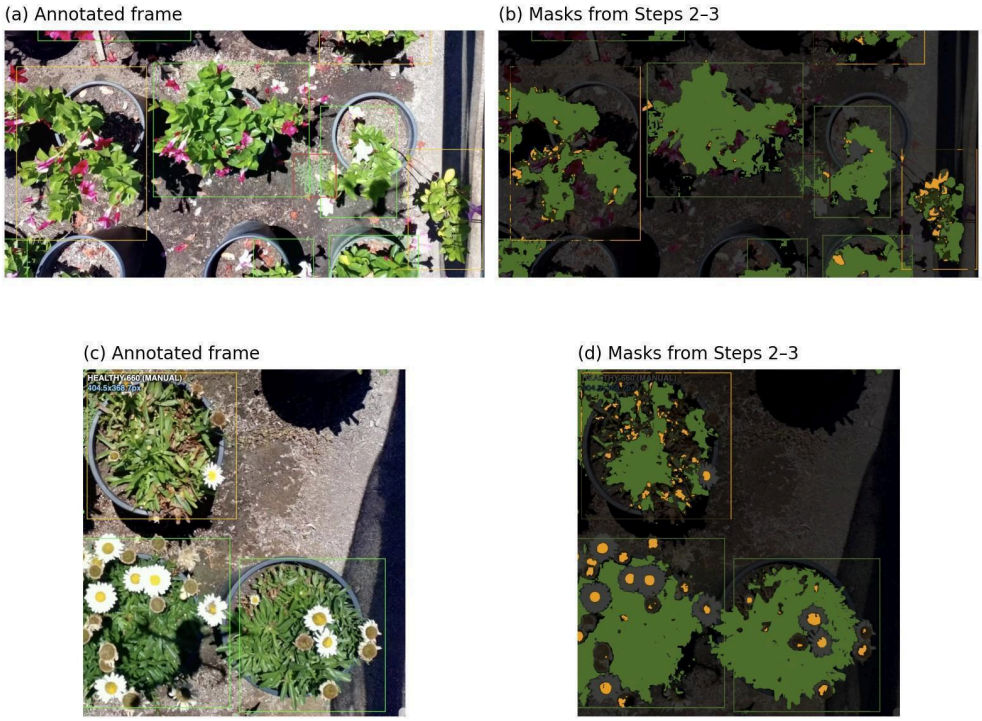


Figure 3: Steps 2–3 applied to two annotated field frames: green is healthy foliage and amber is discolored foliage; everything else is a darkened background. (a–b) Discolored parts of a stressed plant are marked in yellow. (c–d) The browning, dried leaves of the stressed plant at top left are dull and fall below the saturation and value limits, so they are mostly not flagged, while the bright yellow centers of the daisy-like flowers fall inside the discolored hue range. The segmentation masks were computed using specific thresholds and guidelines that were refined with the 3x3 kernel. Source: authors.

Step 4: Discoloration index. The discoloration index, *D*, is the portion of foliage pixels classified as discolored and was calculated for each plant:

$$D = N_{discolored} / (N_{healthy} + N_{discolored}) \quad (2)$$

Step 5: Classification. A plant was labeled stressed if $D \geq 0.20$ and healthy otherwise. This threshold was made clear before evaluations were made, and it was not changed after calculations of *D*. Plants whose box contained less than 5% foliage, usually because the box was mostly background or the plant was fully necrotic and indistinguishable from mulch, were marked unassessable and excluded.

Benchmark Protocol

Both methods were evaluated on the color version of the public PlantVillage dataset, which contains 54,305 photographs of single leaves against uniform backgrounds in 38 crop–condition classes [27]. The 12 classes named "healthy" (15,084 images) were labeled healthy and the 26 disease classes (39,221 images) were labeled stressed; no images were excluded, and images were used at their original 256×256 resolution.

- **Fixed rule (the field algorithm):** Steps 2–5 were applied to each whole image with the same thresholds, including $D \geq 0.20$. Since no training was involved, all 54,305 images served as the test data.
- **Trained color-index classifier (reference):** For each image, 50 color features were computed over non-background pixels (saturation above 0.15): mean and standard deviation of normalized red, green, and blue; mean and standard deviation of the excess-green index; mean VARI and GLI indices; mean saturation and value; *D* and foliage fraction; and a 36-bin hue histogram. A gradient-boosted tree classifier was trained on a stratified random 80% of images (43,444) and tested on the remaining 20% (10,861); a logistic regression on the same features was run for comparison.

Field Annotation

Two annotators labeled the 548 frames in CVAT [32], drawing a bounding box around each plant and assigning one of four categories: healthy, stressed, withering, or dead. For evaluation, stressed, withering, and dead were combined into a single stressed category. The annotators divided the frames by collection rather than double-labeling a shared subset, so inter-annotator agreement could not be computed. In total, 2,201 plant instances were labeled; 189 were unassessable under the 5% foliage rule, leaving 2,012 for evaluation (569 stressed, 1,443 healthy).

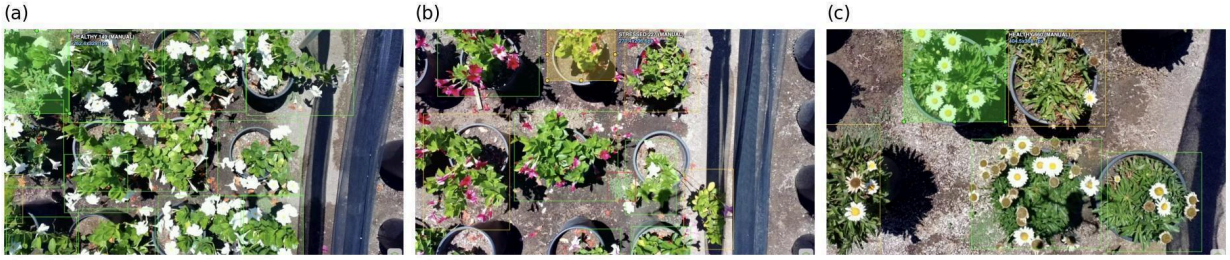


Figure 4: Field frames annotated in CVAT, showing three types of potted flowering plants. Green boxes are labeled healthy and amber boxes stressed; the red box in (b) marks a region annotated separately. The text overlay and shaded box in each frame are CVAT's label and highlight for the currently selected box, not part of the image. Source: authors.

Evaluation Metrics

Since only 28% of field plants were stressed, the overall accuracy can appear higher since some stressed plants may be missed. Therefore, we report accuracy alongside balanced accuracy (the mean of recall and specificity), precision, recall, specificity, and F1 for the stressed class, and the Matthews correlation coefficient (MCC), which stays informative under class imbalance. Recall is one of the most significant metrics as failing to take into account missed stressed plants can represent a costly error for farmers. Ninety-five percent confidence intervals utilized the Wilson method.

Proposed Phase-3 Architecture (Not Implemented)

For operations that later adopt multispectral sensors, we propose a density-adaptive mutual-reachability processing for encoded data (DAMPED) architecture that connects a 3D ResCNN with HDBSCAN (Figure 5). A 3D ResCNN encoder learns spectral and spatial features from multispectral frames, an approach that has kept accuracy with limited training data [11]. The resulting high-dimensional embeddings are reduced in dimension so that density structure can be found [22][23]. HDBSCAN then groups plants into clusters of similar condition and labels outliers as noise [24], following the encoder-plus-clustering design of Zheng et al. [26]. The output is a map of candidate stress hotspots that growers would confirm by ground scouting to determine whether insect damage or another cause is responsible. Because HDBSCAN is unsupervised, it could identify unusual patterns associated with under-documented pests without requiring labeled examples. Synthetic video of pests found in California could also be used to supplement the limited amount of real training data [10]. However, DAMPED has not been implemented or tested in this study, so these capabilities remain proposed rather than demonstrated results.



Figure 5: Proposed DAMPED architecture for Phase 3. Dashed outlines indicate components that are proposed and untested. Design based on [11], [24], and [26].

RESULTS

Benchmark Performance

Table 1 reports both methods. On PlantVillage, the fixed rule was almost never wrong when it flagged a leaf (precision 99.9%, specificity 100.0%, with 4 false alarms among 15,084 healthy leaves), but it flagged only 17.6% of diseased leaves, for an accuracy of 40.5%. Its discoloration index nonetheless ranked diseased leaves above healthy ones well (AUC 0.814): diseased leaves did tend to have more discolored foliage, but usually less than the 0.20 cutoff. The rule caught diseases that cause yellowing, such as grape leaf blight (58.2% flagged), peach bacterial spot (48.5%), strawberry leaf scorch (38.1%), and citrus greening (33.3%), but almost none that cause dark lesions, mottling, or stippling, including tomato target spot (0.0%), tomato mosaic virus (0.0%), apple black rot (0.0%), and spider-mite damage (0.8%). The trained color-index classifier reached 96.7% accuracy on held-out images (AUC 0.993), and the logistic regression reached 86.8% (AUC 0.926), confirming that color carries a strong stress signal on curated leaves when a model can learn which color patterns matter.

Field Performance

On drone imagery (Table 1), across 2,012 plants, accuracy was 70.1% (95% CI 68–72%), within the hypothesized 70–80% range, but this figure was inflated by the larger number of healthy plants. Balanced accuracy was 60.7%, precision 46.5%, and recall only 39.0% (95% CI 35–43%). MCC was 0.226 (Cohen's κ was similar at 0.224), indicating weak agreement with the human labels beyond chance. Using the discoloration index D as a continuous score, the area under the ROC curve (AUC) was 0.645, so even the best choice of threshold would separate stress from healthy plants only modestly.

Metric	PlantVillage: fixed rule	Field: fixed rule	PlantVillage: trained classifier
Accuracy	40.5% (95% CI 40–41)	70.1% (95% CI 68–72)	96.7%
Balanced accuracy	58.8%	60.7%	95.8%
Precision (stressed)	99.9%	46.5%	97.5%

Metric	PlantVillage: fixed rule	Field: fixed rule	PlantVillage: trained classifier
Recall (stressed)	17.6% (95% CI 17–18)	39.0% (95% CI 35–43)	98.0%
Specificity	100.0%	82.3%	93.5%
F1 (stressed)	29.9%	42.4%	97.8%
MCC	0.236	0.226	0.919
AUC	0.814	0.645	0.993
Instances evaluated	54,305 images	2,012 plants	10,861 test images

Table 1: Performance of the fixed-rule RGB algorithm on the PlantVillage benchmark and on field imagery, and of a gradient-boosted classifier trained on RGB color-index features (PlantVillage held-out 20%).

The confusion matrix below (Figure 6) indicates the predictions of the algorithm in comparison to human labels. The algorithm correctly flagged 222 stressed plants and missed 347, or 61% of all stressed plants. It correctly predicted 1,188 healthy plants (82.3%), but it also reported back 255 false positives. The calculations found that fewer than half of its stress flags were correct. Figure 7 compares the same rule on this PlantVillage benchmark and in the field.

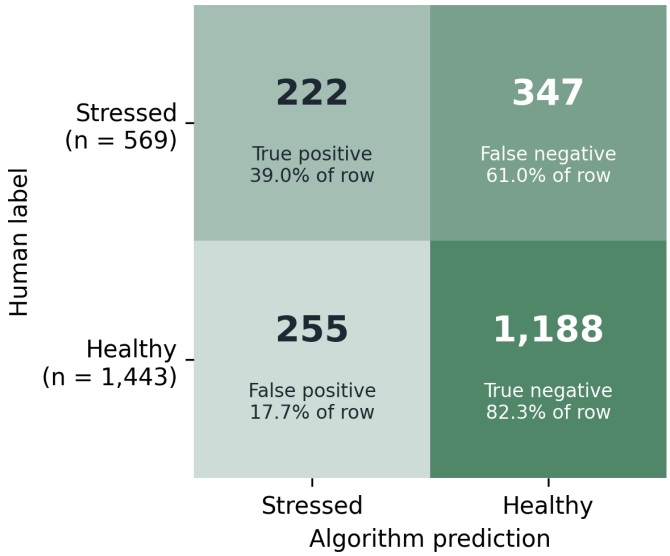


Figure 6: Field confusion matrix for 2,012 assessed plants. Shading shows the proportion of plants in each classification category. Counts were reconstructed from the reported recall, precision, and 347 false negatives and they reproduce every reported metric, including both confidence intervals.

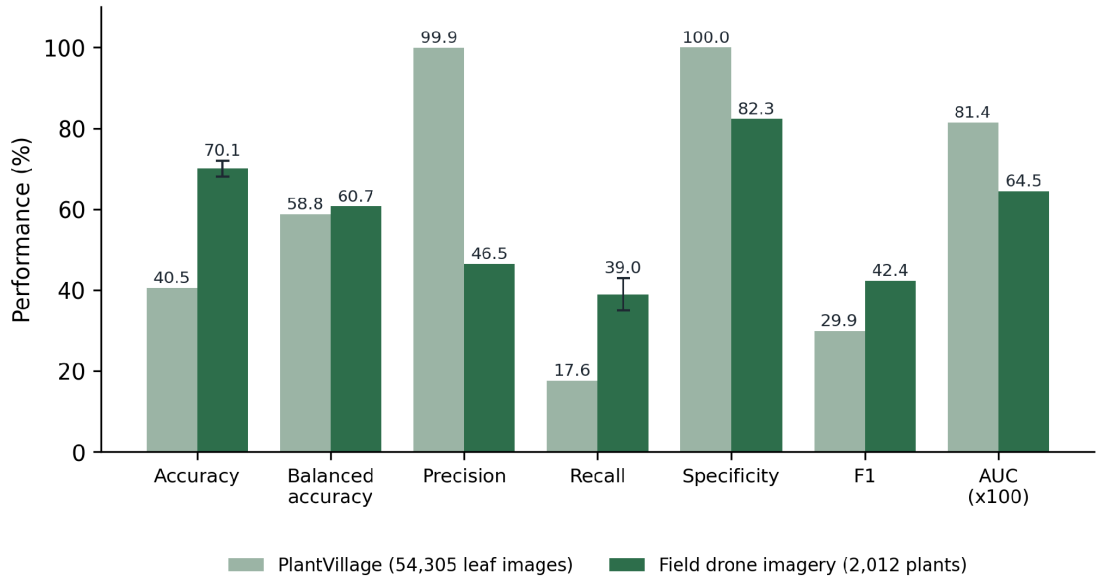


Figure 7: The same fixed rule on PlantVillage and on field imagery. Error bars show 95% Wilson confidence intervals for field accuracy and recall. AUC is multiplied by 100 to determine the performance percentage. Source: data collected by authors and [27].

Error Patterns

The algorithm reliably flagged plants with obvious yellowing but missed early or subtle stress. The example frames in Figure 3 show two mechanisms behind these errors. First, browning or dried foliage is dull and dark, so the saturation and value limits that exclude soil and mulch also exclude it, and such plants are missed. Second, bright yellow flower parts, such as the centers of daisy-like flowers, can fall within the discolored hue range and cause a healthy plant to be incorrectly flagged as stressed. The benchmark shows the same pattern: the rule caught diseases that cause yellowing but almost none that cause dark lesions or stippling, including spider-mite damage.

DISCUSSION

The Benchmark-to-Field Gap

The difference between the benchmark and field results is clearer when comparing the algorithm's ability to rank plants by stress rather than comparing raw accuracy. The accuracy values are not directly comparable because 72% of PlantVillage images are diseased, whereas only 28% of the field plants were classified as stressed. As a result, a rule that flags relatively few plants can perform poorly on the benchmark while producing a higher accuracy in the field. AUC does not depend on class balance or on the cutoff, and it fell from 0.814 on curated leaves to 0.645 in the field for the same algorithm. Meanwhile, a classifier trained on color features reached 96.7% on curated leaves. Curated leaf datasets

therefore substantially overstate how well color-based methods will perform on real drone imagery. Two factors likely contribute to this difference. First, the tasks differ: PlantVillage shows one leaf against a plain background, while field boxes contain whole plants with shadows, mixed backgrounds, flowers, and overlapping neighbors. Second, RGB responds only to visible discoloration, while much stress, such as early decline, wilting, dark lesions, or insect feeding, produces little of the bright yellowing the rule detects; on PlantVillage, it flagged only 0.8% of leaves with spider-mite damage. Overall, the results show that RGB rule can provide a useful starting point for detecting visible stress, but its performance is limited when plant stress is subtle or does not produce a strong color change

Comparison with Published Systems

The table below (Table 2) compares the RGB results against alternative systems discussed in the Related Work section. On curated leaves, the trained color-index classifier (96.7%) matches or exceeds the accuracies reported for DHMPD (0.936) [2], AIDN (up to 0.92) [7], and 3D ResCNN (0.8106) [11], while the fixed rule does not (40.5%). Neither comparison is like-for-like. Those systems detect pests or specific diseases, often use richer sensors, and are evaluated on different data; the RGB algorithm detects generic visible stress. In the field, RGB accuracy trials reported AI-based multispectral systems (about 88–96%) by roughly 18–26 percentage points, and its recall is still lower. This is expected: multispectral and hyperspectral sensors measure near-infrared reflectance, which changes in stressed tissue before visible symptoms appear. Low-cost RGB therefore approaches published accuracy only for coarse, visibly symptomatic distinctions under favorable conditions.

Approach	Sensor	Task	Reported accuracy
This work (color threshold, no training)	Consumer RGB (~$400)	Visible plant stress (binary)	Fixed rule: 40.5% benchmark, 70.1% field; trained classifier: 96.7% benchmark
DHMPD, deep hybrid CNN + IoT [2]	RGB	Pest detection	0.936
AIDN, multi-scale CNN [7]	RGB (UAV)	Insect detection	up to 0.92
Qpest-RCNN [8]	RGB	Pest counting on sticky traps	High (not quantified)
3D ResCNN [11]	Hyperspectral (UAV)	Pine wilt disease	0.8106 (one-fifth of data)
Hyperspectral + deep learning [4][5]	Hyperspectral	Plant disease / sick trees	>95% confidence

Table 2: Indicative comparison with published systems. Accuracies come from different datasets, tasks, and sensors and are not head-to-head comparisons.

Cost–Capability Tradeoff

The RGB system costs about 4% of a comparable multispectral system (Table 3). Its value lies not in matching those systems but in offering an affordable entry point: for an operation with no monitoring at all, a tool that reliably flags advanced, visibly symptomatic stress is a meaningful first capability, provided growers understand that it misses most subtle stress and plan scouting around that limitation.

System	Field accuracy	Hardware cost	Practical accessibility
RGB (this work)	70.1% (39.0% stress recall)	~$300–500	Small and medium operations
Multispectral + AI	~88–96% (reported [2][4][7][11])	$10,000+	Large operations
Satellite imagery	Not measured here	Subscription; multi-day latency	Resolution and latency can be limiting

Table 3: Cost–capability comparison. RGB figures are measured in this study; other figures are drawn from the literature and are indicative.

A Staged Pathway from RGB to AI-Based Pest Monitoring

The tradeoff points to a staged pathway rather than a choice between cheap RGB and expensive AI (Figure 8).

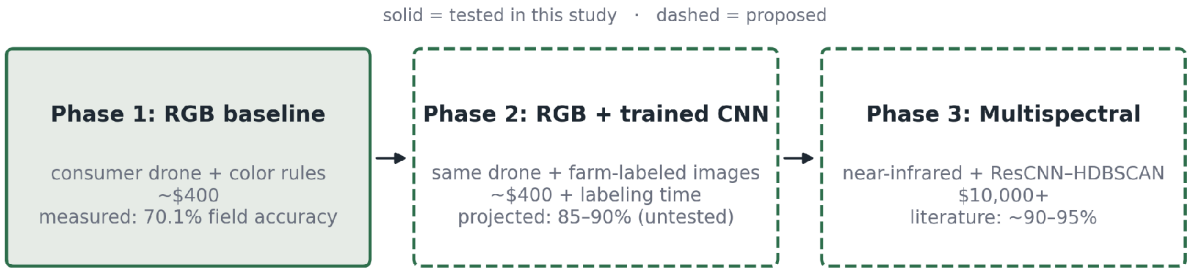


Figure 8: Proposed pathway from low-cost RGB to AI-based pest monitoring. Only Phase 1 was tested; Phase 2 accuracy is a projection and Phase 3 accuracy comes from the literature.

Phase 1: RGB baseline. A consumer drone and color rules establish feasibility for about $400. This study measured 70.1% accuracy and 39.0% recall: visible color captures overt discoloration but misses subtle stress, and it cannot distinguish insect damage from other causes.

Phase 2: RGB with a trained CNN (proposed). The same drone, paired with a modest set of farm-collected labeled images used to fine-tune a compact detector such as a YOLO-family model [31], could learn crop- and damage-specific visual patterns that color rules cannot capture. This phase directly

targets the recall deficit and is the first step toward identifying insect damage specifically. The trained color-index classifier achieved 96.7 accuracy on curated leaves, compared with 40.5% for the fixed rule. This difference suggests that using labeled examples allowed the classifier to better distinguish the color patterns associated with plant stress. Fine-tuned systems tested in field conditions have reported 85-90%, but these results still need to be validated utilizing data collected from farms.

Phase 3: Multispectral sensing with DAMPED (proposed). For operations able to invest further, multispectral imagery and the DAMPED architecture (see Methods) add near-infrared and spatial-density information for detecting stress before it is visible, with comparable systems reporting about 90–95%. Confirming that a flagged hotspot is insect damage would still require ground scouting or a dedicated pest-detection model.

In this pathway, the RGB serves as an initial step rather than a replacement for the later methods. Its low cost allows the basic approach to be tested before moving to more advanced systems. At the same time, its measured limitations, particularly its low recall in field conditions, indicate where the baseline is insufficient and provide a basis for moving to later phases.

Limitations

- Setting: Imagery came from a small number of ornamental and mixed-vegetation sites rather than commercial monoculture farms, and from a single drone operator.
- Ground truth: Labels were visual judgments by the authors, not trained agronomists or instrument measurements such as chlorophyll or moisture readings, and no subset was double-labeled, so inter-annotator reliability is unknown.
- Non-independent samples: With 86.5% similarity between consecutive frames, the same plant can appear in several frames. Confidence intervals that treat each instance as independent are therefore likely too narrow; the released evaluation code also computes intervals that resample whole video clips.
- Record-keeping: Screening counts for the literature review, the exact cutoffs of the frame filters, and written criteria for each annotation category were not recorded, and results were not broken down by site type or time of day.
- Benchmark differences: The trained classifier was tested only on PlantVillage; it was not applied to the field imagery, so its benchmark accuracy shows the strength of the color signal on curated leaves, not expected field performance.
- Human-drawn boxes: The algorithm classified plants inside human-drawn boxes. A fully automatic system would also need to locate plants, which would likely lower performance.
- No cause attribution: The algorithm detects visible stress, not pests, and cannot say whether stress comes from insects, water, nutrients, or disease.
- Environment: Wind, rain, and fog limit UAV data collection [25].

CONCLUSION

This study evaluated whether a $400 consumer RGB drone and a training-free color algorithm can detect visible plant stress well enough to serve as a starting point for AI-based monitoring. The algorithm reached 70.1% accuracy but only 39.0% stress recall on real drone imagery, missing most subtly stressed plants, and its ability to separate stressed from healthy plants fell from AUC 0.814 on a curated leaf benchmark to 0.645 in the field. A classifier trained on the same kind of color features reached 96.7% on curated leaves, showing that curated benchmarks overstate field readiness. The results support the use of a low-cost RGB imaging as a Phase-1 foundation for plant-stress monitoring, particularly for detecting more obvious signs of stress. However, the approach is not sufficient for identifying pests or recognizing insect damage. The results therefore justify moving to trained CNN detectors and, ultimately, multispectral sensing through the proposed DAMPED architecture and IoT-based monitoring system. Future research should validate each phase across multiple commercial farms, crops, and growing seasons. Validation should also include trained-agronomist verification, double-labeled annotations, and comparison of RGB and multispectral imagery collected from the same plot.

FUNDING

This research was funded by Drone Slone Squadron (DSS), a private student-led project focused on advancing the adoption of precision-agriculture adoption. No other external funding or support was received to conduct this study.

CONFLICT OF INTEREST

The authors declare no commercial or financial relationships that could be construed as a potential conflict of interest. The study compares RGB and AI-based multispectral approaches and does not endorse any specific commercial product or vendor

DATA AVAILABILITY

Raw drone footage, processed frames, annotations, and the evaluation code (field evaluation, fixed-rule benchmark, and color-index classifier) are available from the corresponding author on request.

REFERENCES

[1] Jetter, K.M.; Godfrey, K. Diaprepes root weevil, a new California pest, will raise costs for pest control and trigger quarantines. *California Agriculture* 2009, 63, 121–126. DOI:10.3733/ca.v063n03p121

[2] Gokeda, V.; Yalavarthi, R. Deep Hybrid Model for Pest Detection: IoT-UAV-Based Smart Agriculture System. *Journal of Phytopathology* 2024, 172(5), e13381. DOI:10.1111/jph.13381

[3] Bhoi, S.K.; Jena, K.K.; Panda, S.K.; Long, H.V.; Kumar, R.; Subbulakshmi, P.; Jebreen, H.B. An Internet of Things assisted Unmanned Aerial Vehicle based artificial intelligence model for rice pest detection. *Microprocessors and Microsystems* 2021, 80, 103607. DOI:10.1016/j.micpro.2020.103607

[4] Nguyen, C.; Sagan, V.; Maimaitiyiming, M.; Maimaitijiang, M.; Bhadra, S.; Kwasniewski, M.T. Early Detection of Plant Viral Disease Using Hyperspectral Imaging and Deep Learning. *Sensors* 2021, 21, 742. DOI:10.3390/s21030742

[5] Nguyen, H.T.; Lopez Caceres, M.L.; Moritake, K.; Kentsch, S.; Shu, H.; Diez, Y. Individual Sick Fir Tree (Abies mariesii) Identification in Insect Infested Forests by Means of UAV Images and Deep Learning. *Remote Sensing* 2021, 13, 260. DOI:10.3390/rs13020260

[6] Zhang, Y.; Chen, J.; Zeng, X.; Chen, R.; Chen, L.; Xiao, S.; Song, Y. A Multimodal UAV-IoT Sensing Framework for Intelligent Pest Density Estimation in Smart Agricultural Systems. *Sensors* 2026, 26, 2877. DOI:10.3390/s26092877

[7] Khujamatov, H.; Muksimova, S.; Abdullaev, M.; Cho, J.; Jeon, H.-S. Advanced Insect Detection Network for UAV-Based Biodiversity Monitoring. *Remote Sensing* 2025, 17, 962. DOI:10.3390/rs17060962

[8] Juan, Y.; Ke, Z.; Chen, Z.; Zhong, D.; Chen, W.; Yin, L. Rapid density estimation of tiny pests from sticky traps using Qpest RCNN in conjunction with UWB-UAV-based IoT framework. *Neural Computing and Applications* 2024, 36, 9779–9803. DOI:10.1007/s00521-023-09230-4

[9] Page, M.J.; McKenzie, J.E.; Bossuyt, P.M.; Boutron, I.; Hoffmann, T.C.; Mulrow, C.D.; et al. The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. *BMJ* 2021, 372, n71. DOI:10.1136/bmj.n71

[10] Stojnić, V.; Risojević, V.; Muštra, M.; Jovanović, V.; Filipi, J.; Kezić, N.; Babić, Z.; Cao, L. A Method for Detection of Small Moving Objects in UAV Videos. *Remote Sensing* 2021, 13, 653. DOI:10.3390/rs13040653

[11] Yu, R.; Luo, Y.; Li, H.; Yang, L.; Huang, H.; Yu, L.; Ren, L. Three-Dimensional Convolutional Neural Network Model for Early Detection of Pine Wilt Disease Using UAV-Based Hyperspectral Images. *Remote Sensing* 2021, 13, 4065. DOI:10.3390/rs13204065

[12] Behmann, J.; Acebron, K.; Emin, D.; Bennertz, S.; Matsubara, S.; Thomas, S.; Bohnenkamp, D.; Kuska, M.T.; Jussila, J.; Salo, H.; et al. Specim IQ: Evaluation of a New, Miniaturized Handheld Hyperspectral Camera and Its Application for Plant Phenotyping and Disease Detection. *Sensors* 2018, 18, 441. DOI:10.3390/s18020441

[13] Subramanian, K.S.; Pazhanivelan, S.; Srinivasan, G.; Santhi, R.; Sathiah, N. Drones in Insect Pest Management. *Frontiers in Agronomy* 2021, 3, 640885. DOI:10.3389/fagro.2021.640885

[14] Zhu, H.; Lin, C.; Liu, G.; Wang, D.; Qin, S.; Li, A.; Xu, J.; He, Y. Intelligent agriculture: deep learning in UAV-based remote sensing imagery for crop diseases and pests detection. *Frontiers in Plant Science* 2024, 15, 1435016. DOI:10.3389/fpls.2024.1435016

[15] California Department of Food and Agriculture (CDFA). The Battle Against Invasive Pests and Diseases. https://www.cdfa.ca.gov/invasives/

[16] Yasaei, R.; Hernandez, F.; Al Faruque, M.A. IoT-CAD: Context-aware adaptive anomaly detection in IoT systems through sensor association. In *2020 IEEE/ACM International Conference on Computer Aided Design (ICCAD)*, 2020, pp. 1–9.

[17] Almalki, F.A.; Soufiene, B.O.; Alsamhi, S.H.; Sakli, H. A Low-Cost Platform for Environmental Smart Farming Monitoring System Based on IoT and UAVs. *Sustainability* 2021, 13(11), 5908. DOI:10.3390/su13115908

[18] Azfar, S.; Nadeem, A.; Ahsan, K.; Mehmood, A.; Almoamari, H.; Alqahtany, S.S. IoT-Based Cotton Plant Pest Detection and Smart-Response System. *Applied Sciences* 2023, 13, 1851. DOI:10.3390/app13031851

[19] Bhati, N.; Tripathy, M. Generative residual framework for imbalanced solar PV fault classification and identification. *Solar Energy* 2026, 114678. DOI:10.1016/j.solener.2026.114678

[20] Xie, B.; Zhang, H.K.; Xue, J. Deep Convolutional Neural Network for Mapping Smallholder Agriculture Using High Spatial Resolution Satellite Image. *Sensors* 2019, 19, 2398. DOI:10.3390/s19102398

[21] Wu, X.; Liu, Y.; Xing, M.; Yang, C.; Hong, S. Image Segmentation for Pest Detection of Crop Leaves by Improvement of Regional Convolutional Neural Network. *Scientific Reports* 2024, 14, 24160. DOI:10.1038/s41598-024-75391-4

[22] Yang, L.; Fan, W.; Bouguila, N. Deep Clustering Analysis via Dual Variational Autoencoder With Spherical Latent Embeddings. *IEEE Transactions on Neural Networks and Learning Systems* 2023, 34, 6303–6312. DOI:10.1109/TNNLS.2021.3135460

[23] Hasana, S.; Fitrianah, D. A Study on Enhanced Spatial Clustering Using Ensemble DBSCAN and UMAP to Map Fire Zones in Greater Jakarta, Indonesia. *Jurnal Riset Informatika* 2023, 5, 409–418. DOI:10.34288/jri.v5i3.234

[24] McInnes, L.; Healy, J. Accelerated Hierarchical Density Based Clustering. In *Proceedings of the 2017 IEEE International Conference on Data Mining Workshops (ICDMW)*, New Orleans, LA, USA, 2017; pp. 33–42. DOI:10.48550/arXiv.1705.07321

[25] Mashori, A.S.; Li, F.; Aman, M.; Zhang, W.; Jia, S.; Ali, A.; Jabeen, N.; Hao, W.; Yan, Y. Remote Sensing through UAVs for Precision Agriculture: Applications, Technical Foundations, Current Barriers, and Future Opportunities. *Smart Agricultural Technology* 2026. DOI:10.1016/j.atech.2026.102074

[26] Zheng, Z.; Zhang, S.; Song, H.; Yan, Q. Deep Clustering Using 3D Attention Convolutional Autoencoder for Hyperspectral Image Analysis. *Scientific Reports* 2024, 14, 4209. DOI:10.1038/s41598-024-54547-2

[27] Hughes, D.P.; Salathé, M. An open access repository of images on plant health to enable the development of mobile disease diagnostics. *arXiv* 2015, arXiv:1511.08060.

[28] Bradski, G.; Kaehler, A. *Learning OpenCV: Computer Vision in C++ with the OpenCV Library*; O'Reilly Media, 2008.

[29] Van Rossum, G.; Drake, F.L. *The Python Language Reference Manual*; Network Theory Ltd., 2009.

[30] FFmpeg Project. FFmpeg multimedia framework, 2026. https://ffmpeg.org/

[31] Jocher, G.; Chaurasia, A.; Qiu, Z. YOLO by Ultralytics, 2023. https://github.com/ultralytics/ultralytics

[32] Boesch, G. CVAT: Computer Vision Annotation Tool, 2019. https://github.com/opencv/cvat

[33] Harris, C.R.; et al. Array programming with NumPy. *Nature* 2020, 585, 357–362. DOI:10.1038/s41586-020-2649-2

[34] Potensic. Atom 2 Specifications and User Manual, 2026. https://www.potensic.com/

[35] California Department of Food and Agriculture (CDFA). California Agricultural Production Statistics. https://www.cdfa.ca.gov/Statistics/ (accessed 27 September 2026).

ABOUT THE AUTHORS

Anik Koirala, Adit Shah, and Ian Chan are seniors at Gretchen Whitney High School in Cerritos, California, and members of Drone Slone Squadron, a student-led project on using drones for agricultural assistance.