

Letter

Watson on the Farm: Using Cloud-Based Artificial Intelligence to Identify Early Indicators of Water Stress

Daniel Freeman ¹, Shaurya Gupta ², D. Hudson Smith ³, Joe Mari Maja ^{4,*}, James Robbins ⁵, James S. Owen Jr. ⁶, Jose M. Peña ⁷ and Ana I. de Castro ⁸

¹ Department of Biochemistry and Genetics, Clemson University, Poole Agricultural Center, Clemson, SC 29634, USA; freemaf@g.clemson.edu

² Department of Electrical and Computer Engineering, Clemson University, Riggs Hall, Clemson, SC 29634, USA; shauryg@g.clemson.edu

³ Watt Family Innovation Center, Clemson University, 405 S Palmetto Blvd., Clemson, SC 29634, USA; dane2@clemson.edu

⁴ Department of Agricultural Science, Clemson University, 240 McAdams Hall, Clemson, SC 29634, USA

⁵ Division of Agriculture, University of Arkansas, Little Rock, AR 72204, USA; jrobbins@uaex.edu

⁶ School of Plant and Environmental Sciences, Virginia Tech, Hampton Roads Agricultural Research and Extension Center, Virginia Beach, VA 23455, USA; jsowen@vt.edu

⁷ Institute of Agricultural Sciences, CSIC, 28006 Madrid, Spain; jmpena@ica.csic.es

⁸ Institute for Sustainable Agriculture, CSIC, 14004 Cordoba, Spain; anadecastro@ias.csic.es

* Correspondence: jmaja@clemson.edu

Received: 27 September 2019; Accepted: 8 November 2019; Published: 13 November 2019



Abstract: As demand for freshwater increases while supply remains stagnant, the critical need for sustainable water use in agriculture has led the EPA Strategic Plan to call for new technologies that can optimize water allocation in real-time. This work assesses the use of cloud-based artificial intelligence to detect early indicators of water stress across six container-grown ornamental shrub species. Near-infrared images were previously collected with modified Canon and MAPIR Survey II cameras deployed via a small unmanned aircraft system (sUAS) at an altitude of 30 meters. Cropped images of plants in no, low-, and high-water stress conditions were split into four-fold cross-validation sets and used to train models through IBM Watson's Visual Recognition service. Despite constraints such as small sample size (36 plants, 150 images) and low image resolution (150 pixels by 150 pixels per plant), Watson generated models were able to detect indicators of stress after 48 hours of water deprivation with a significant to marginally significant degree of separation in four out of five species tested ($p < 0.10$). Two models were also able to detect indicators of water stress after only 24 hours, with models trained on images of as few as eight water-stressed Buddleia plants achieving an average area under the curve (AUC) of 0.9884 across four folds. Ease of pre-processing, minimal amount of training data required, and outsourced computation make cloud-based artificial intelligence services such as IBM Watson Visual Recognition an attractive tool for agriculture analytics. Cloud-based artificial intelligence can be combined with technologies such as sUAS and spectral imaging to help crop producers identify deficient irrigation strategies and intervene before crop value is diminished. When brought to scale, frameworks such as these can drive responsive irrigation systems that monitor crop status in real-time and maximize sustainable water use.

Keywords: sUAS; water stress; ornamental; container-grown; artificial intelligence; machine learning; deep learning; neural network; visual recognition; precision agriculture

1. Introduction

Freshwater is a finite resource that is required for the daily production of container crops to be used for food, ecosystem services, urban development, and other purposes. The United Nations Education, Scientific, and Cultural Organization (UNESCO) has indicated that the combined expansion of manufacturing, agriculture, and urban populations has created excessive strain on the existing fresh water supply and has called for more sustainable water management [1]. One opportunity to reduce water consumption lies in the development of intelligent irrigation systems that can optimize water use in real-time [2]. Crop producers routinely provide an excess of water to container-grown plants to mitigate plant stress and subsequent economic loss, resulting in inefficient use of agrichemicals, energy, and freshwater. Site-specific irrigation systems minimize these losses by using sensors to allocate water to plants as needed, improving crop production while minimizing operating costs [3]. Sensor-based irrigation is not a new concept [1,4–6]. Kim et al. [5] developed software for an in-field wireless sensor network (WSN) to implement site-specific irrigation management in greenhouse containers. Coates et al. [7] developed site-specific applications using soil water status data to control irrigation valves.

In 2017, the U.S. nursery industry had sales of \$5.9 billion and ornamental production accounted for 2.2 percent of all U.S. farms [8]. Plants grown in containers are the primary (73%) production method [9] and the majority (81%) of nursery production acreage is irrigated [10]. The largest production cost for nurseries is labor, which amounts to 39% of total costs [11], and labor shortages are linked to reduced production [12]. Adoption of appropriate technologies may offset increasing labor costs and labor shortages. Small unmanned aircraft systems (sUAS) have been suggested as an important tool in nursery production to help automate certain processes such as water resource management [13].

sUASs allow farmers to quickly survey large plots of land using aerial imagery. sUAS imagery has been used to detect diseases and weeds [14,15], predict cotton yield [16], measure the degree of stink bug aggregation [17], and identify water stress in ornamental plants [18]. Several thermal and spectral indices have been correlated to biophysical plant parameters based on sUAS imagery [19,20]. Analyses of sUAS imagery have been shown to be sensitive to time of day, cloud cover, light intensity, image pixel size, soil water buffering capacity, and atmospheric conditions at the canopy level [21,22]. Still, multispectral data collected with sUAS were shown to be more accurate than data collected using manned aircraft [23]. A variety of methodologies, including thermal and spectral imagery, have been used to assess water stress in conventional sustainable agriculture using sUAS [3]. Stagakis et al. [24] indicated that the high spatial and spectral resolution provided by sUAS-based imagery could be used to detect deficient irrigation strategies. Zovkoa et al. [25] reported difficulty measuring three levels of water stress of grape grown in soil; however, they were able to discern irrigated vs. non-irrigated plots via hyperspectral image analysis (409–988 nm and 950–2509 nm) when employing a support vector machine (SVM). de Castro et al. [18] successfully identified water-stressed and non-stressed containerized ornamental plants using two multispectral cameras aboard an sUAS, although the spectral separation was higher when information from the sensors was combined. Data being produced by de Castro and Zovkoa could be utilized as a roadmap for real-time, sustainable water management of specialty or container-grown crops using sUAS. Fulcher et al. [26] indicated that the adoption of sUAS to monitor crop water status will be useful in addressing the challenge of sustainable water use in container nurseries. Unlike conventional crops produced in soil systems, containerized soilless-based systems have low water buffering capacity, resulting in rapid physiological changes that may not be observed at the ground level visually, but can be monitored by reflected wavelengths captured by sUAS. To reduce size and cost, sUAS can collect and wirelessly transmit high-resolution image data to cloud providers that can perform analyses on offsite servers. Thus, the convergence of technologies—such as sUAS, Internet of Things (IoT), spectral imagery, and cloud-based computing—can be used to build intelligent irrigation systems that monitor crop status and optimize water allocation in real time.

In this study, images were analyzed with IBM Watson Visual Recognition, a cloud-hosted artificial intelligence service that allows users to train custom image classifiers using deep convolutional neural

networks (CNNs). Unlike linear algorithms, CNNs model complex non-linear relationships between the independent variables (pixels comprising the image) and the dependent variable (plant health) by transforming data through layers of increasingly abstract representation (Figure 1). The first layer is an array of pixel values from the original image; nodes in subsequent layers represent local features such as color, texture, and shape; deeper layers encode semantic information such as leaf or branch morphology. Individual nodes become optimized to represent different features of the image through an iterative learning process that rewards nodes that amplify aspects of the image that are useful for classification and suppresses those that do not [27]. The convolutional relationship from one layer to the next allows CNNs to model complex relationships between input variables, making it particularly useful for analyzing image data that cannot be understood by examining pixels in isolation. Given a set of images of stressed and non-stressed plants, for example, individual nodes in the network may become optimized to represent spectral indices that are sensitive to water stress. Those nodes can affect the outcome directly, or they can feed forward into higher-order features such as the specific location and pattern of discoloration within the plant. Spectral indices may combine with other plant features such as the unique structure of sagging branches or the distinct texture created by the shadows from drooping leaves. All of these features culminate in a single output node that returns a value from zero to one representing the confidence that a given image belongs to the desired class (i.e., water stress).

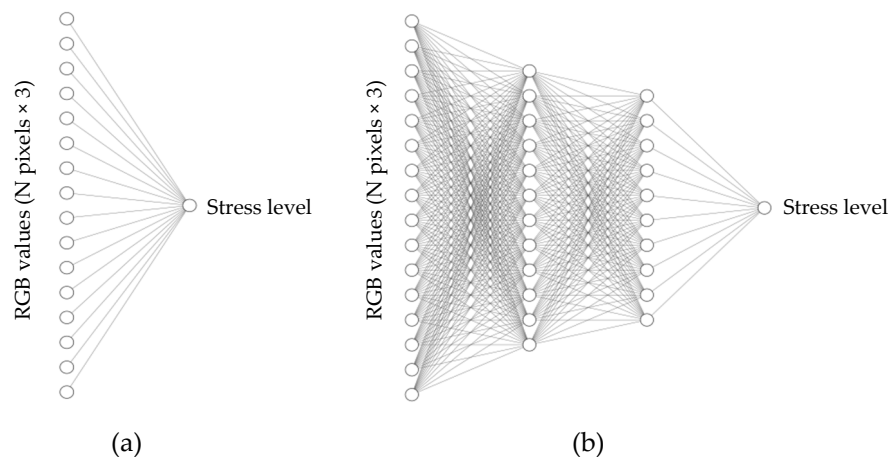


Figure 1. (a) Linear model in which each variable directly affects the outcome versus (b) a convolutional neural network (CNN) in which data is transformed through multiple layers.

While CNNs' layers allow networks to model complex nonlinear relationships that simpler algorithms might miss, they are also prone to overfitting. This occurs when CNNs learn patterns that are specific to the training set and do not generalize to the overall population. In one case study, for example, a model trained to predict a patient's age based on MRI images was found to have learned the shape of the head rather than the content of the scan itself [28]. The challenge of overfitting is compounded by CNNs' inherent 'black box' quality. Since information is passed through so many transformations, it is difficult to identify which input variables have the largest influence on the final outcome. While CNNs often must be trained with large datasets to overcome their tendency to overfit, transfer learning techniques allow fully trained networks to be repurposed for new classification tasks with much smaller datasets. A growing set of tools are also making it possible to introspect models to determine feature importance directly. Saliency heat maps, for example, can highlight regions of the image that are used for classification [29,30]. Overfitting can be tested with a cross-validation scheme in which models are trained with one set of images and then used to classify a new, previously unseen set of images. Performance metrics are based on how well the model's classification of unseen data matches a ground truth standard. A final limitation of CNNs is the significant amount of time and resources required to train them. To circumvent this, the computation may be outsourced to cloud

computing providers that train models on large servers and offer a suite of tools for hyperparameter tuning, transfer learning, and cross-validation [31].

Despite their inherent limitations, CNNs have become popular for image recognition tasks ranging from Facebook photo-tagging to self-driving cars [32,33]. In agriculture, CNNs have been used to predict wheat yield based on soil parameters, diagnose diseases with simple images of leaves, and detect nitrogen stress using hyperspectral imagery [34,35]. CNNs' ability to learn complex nonlinear features makes them particularly useful for analyzing image data in which individual pixels form larger features such as shape or texture. Extensive research has demonstrated that CNNs perform image classification tasks with higher accuracy than traditional machine vision algorithms [36].

In our study, a small set of aerial images were used to train custom image classification models to detect water stress in ornamental shrubs. The objective was to evaluate the ability of IBM Watson's Visual Recognition service to detect early indicators of plant stress. These experiments provide a strong rationale for the deployment of cloud-based artificial intelligence frameworks that use larger datasets to monitor crop status and maximize sustainable water use.

2. Materials and Methods

This research was conducted at the Hampton Roads Agricultural Research and Extension Center (Hampton Roads AREC-Virginia Tech), located in Virginia Beach, VA, USA (36.8919N, 76.1787W). Six plots with container-grown ornamental plants across two experimental areas were studied. Containers were established outdoors on gravel. The species and number of plants in each experimental plot are shown in Table 1.

Table 1. Species and number of plants in each experimental plot.

Scientific Name	Common Name	Height (cm)	Width \pm SD (cm)	Plants
<i>Buddleia</i> x 'Podaras #5' (BUD)	Buddleia Flutterby Grande [®] Peach Cobbler	82	109 \pm 12	34
<i>Cornus obliqua</i> 'Powell Gardens' (CO)	Red Rover [®] silky dogwood	43	55 \pm 6	42
<i>Hydrangea paniculata</i> 'Limelight' (HP)	Limelight panicle hydrangea	46	50 \pm 5	42
<i>Hydrangea quercifolia</i> 'Queen of Hearts' (HQ)	Queen of Hearts oakleaf hydrangea	31	49 \pm 7	21
<i>Physocarpus opulifolius</i> 'Seward' (PO)	Summer Wine [®] ninebark	61	100 \pm 23	36
<i>Spiraea japonica</i> 'Neon Flash' (SJ)	Neon Flash spirea	83	68 \pm 21	32

A subset of plants from each species was removed from the open-air nursery and transferred to a greenhouse where the plants experienced water stress due to the absence of overhead irrigation. High water stress (HWS) plants were transferred to the greenhouse on 8 Aug 2017 and low water stress (LWS) plants were transferred to the greenhouse on 9 Aug 2017. The plants were then returned to the open-air nursery on 10 Aug 2017 after non-stressed plants received overhead irrigation daily, including 10 Aug 2017. This process produced three levels of water stress for this experiment; high, low, and non-stressed (Table 2). At the time of flight, the soilless substrate of HWS plants contained ~19% less water (mL) than non-stress plants and soilless substrate of LWS plants contained ~13% less water (mL) than non-stress plants. There were no easily detectable visual symptoms of water stress in any of the treatment plants. After the data collection, all water-stressed plants were returned to normal irrigation on 10 August 2017 where they fully recovered and continued to grow. This strategy was part of a broader research program with the aim of studying the adaptation of ornamental species to stress conditions.

Table 2. Technical specifications of the sensors onboard the sUAS and the number of images taken at each field.

	Modified Canon	MAPIR Survey 2
Sensor resolution (pixels)	3246 × 2448	4608 × 3456
Focal length (mm)	28	23
Radiometric resolution (bit)	24	24
Image format	jpg	jpg
Image no. for Area 1	131	262
Image no. for Area 2	154	324
Wavebands	Red, Green, NIR	Red, Green, NIR
GSD (at 120 m)	4.05 cm	4.05 cm

2.1. Image Acquisition

Experimental plots were photographed with a quadcopter drone (DJI Inspire 2, DJI Science and Technology Co. Ltd., Shenzhen, China) (Figure 2a) mounted with two cameras: (1) a modified (Llewellyn Data Processing LLC, Carlstadt, NJ, USA) Canon (Canon, Tokyo, Japan) PowerShot ELPH 130 IS; and (2) a MAPIR Survey2 (MAPIR, Peau Productions, Inc., CA, USA) in Figure 2b. During each flight, the quadcopter took images using each camera at a height of 30 meters and a forward and side lap of 90% and 60%, respectively. The technical specifications of the two sensors are shown in Table 2. Figure 3 shows the data collected from both sensors.




Figure 2. (a) DJI Inspire 2 with (b) sensors installed underneath. Images on (b) were from www.canon.com (Canon Powershot) and www.mapir.camera (MAPIR).




Figure 3. Examples of false-color near-infrared images taken by (a) modified Canon and (b) MAPIR cameras.

2.2. Image Pre-Processing

Images were cropped using a custom, browser-based interface in LabelBox (Figure 4). Data were annotated by dragging a bounding box across each plant and labeling it as ‘high water stress’, ‘low

water stress', or 'no stress' according to the key provided in Figure 5. The GraphQL application programming interface (API) was used to pull the pixel coordinates of each bounding box onto a local computer so that individual plants could be cropped from the original aerial images. The resolution of the cropped images was approximately 150 by 150 pixels. The number of cropped images for each condition is shown in Table 3.

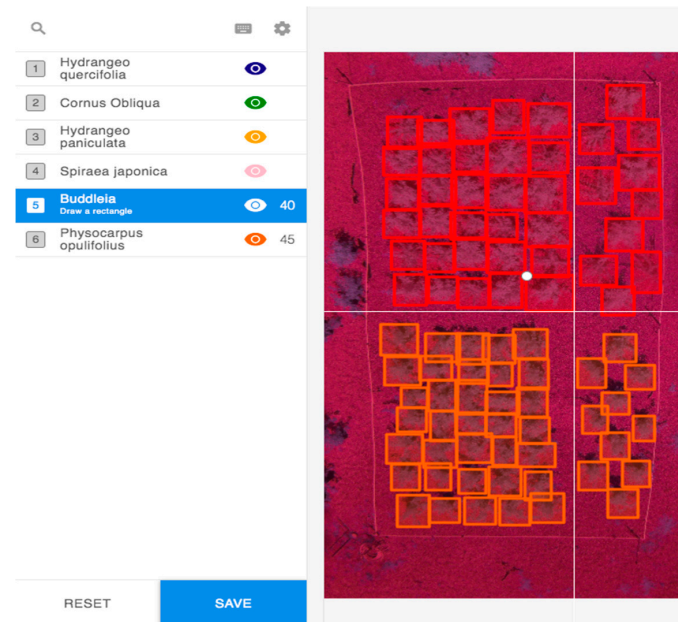


Figure 4. Custom LabelBox interface used to crop aerial images.



Figure 5. Example of the key used to identify plants according to species and condition. Stressed plants are labeled red (high water stress), yellow (low water stress), or blue (low phosphorus fertilizer, not used). Unmarked plants are non-stressed.

Since multiple photographs were taken of the same plots from different angles, cropped images of the same plants were grouped together so that they could be segregated into the training set or validation set as complete units. This procedure protected against overly optimistic performance estimates that would occur if photographs of the same plant appeared in both the training and validation datasets. For each species and treatment, the centers of each bounding box were calculated and normalized to a range of zero to one. Spatstat (<http://spatstat.org>), an open-source R package for analyzing point patterns, was then used to match plants from different aerial images based on the similarity of their pixel coordinates. For example, if there were eight plants in the HWS treatment of a certain species, all images of plants one through six would be used to train the model and all images

of plants seven and eight would be used for validation. This allowed us to make full use of the data during the training phase without artificially inflating performance metrics by validating models with images of the same plants they were trained with. The successful grouping was confirmed by visual inspection (Figure 6).

Table 3. Number of images for each camera, species, and treatment.

Species	Water Treatment	Plants	Modified Canon	MAPIR
<i>Buddleia</i> (BUD)	HWS	8	24	24
	LWS	8	24	24
	NS	18	72	72
<i>Cornus</i> (CO)	HWS	8	25	25
	LWS	8	25	25
	NS	28	85	85
<i>Hydrangea paniculata</i> (HP)	HWS	8	40	40
	LWS	8	40	40
	NS	30	150	150
<i>Hydrangea quercifolia</i> (HQ)	HWS	6	30	30
	LWS	6	30	30
	NS	25	125	125
<i>Physocarpus</i> (PO)	HWS	8	24	24
	LWS	6	18	18
	NS	31	93	93
<i>Spiraea</i> (SJ)	HWS	10	0	10
	LWS	10	0	10
	NS	36	0	36

HWS = high water stress; LWS = low water stress; and NS = no stress treatment.




Figure 6. Images of the same plants cropped from multiple aerial images (Img1–Img5) were matched based on their pixel coordinates so that unique plants could be grouped into either the training set or the validation set. Brackets show one of four train/test splits.

2.3. Model Training and Testing

Cropped images were used to train models with the Watson Visual Recognition API, a cloud-hosted artificial intelligence service provided by IBM that uses CNNs to build custom image classifiers. Here, models were trained to predict water stress status using red, green, and near-infrared pixel values of the cropped images. A Python script was used to access the service and transfer images from a local computer to a cloud server for model training and testing. For each species and camera, three-quarters of NS and HWS images were used to train a model that was then used to classify the remaining quarter. The API returned a prediction between zero and one for each validation image with zero indicating no stress and one indicating water stress (Table 4). This process was repeated four times so that a prediction could be made for each image in the dataset and compared to the ground truth.

Table 4. Predictions returned by the Watson Visual Recognition API (Score) are compared to the ground truth (Stress).

Row	Image	Score	Stress
1	IMG_7696_bud_lws_729.JPG	0.906	TRUE
2	IMG_7695_bud_lws_815.JPG	0.875	TRUE
3	IMG_7694_bud_lws_644.JPG	0.916	TRUE
4	IMG_7696_bud_lws_730.JPG	0.798	TRUE
5	IMG_7694_bud_lws_645.JPG	0.777	TRUE
6	IMG_7695_bud_lws_814.JPG	0.856	TRUE
7	IMG_7694_bud_ns_658.JPG	0.031	FALSE
8	IMG_7696_bud_ns_747.JPG	0.001	FALSE
...
23	IMG_7696_bud_ns_746.JPG	0.062	FALSE
24	IMG_7696_bud_ns_732.JPG	0.852	FALSE

2.4. Statistical Analysis

A receiver operating characteristic area under the curve (AUC) score was used to quantify the degree of separation between treatments for each species and camera. A one-sample *t*-test was used to compare the AUC scores returned by the four-fold validation sets to a hypothesized mean of 0.5, corresponding to random classification.

3. Results

Of the 11 combinations of species and camera used in this study, four produced models that were able to discriminate images of NS and HWS plants with a statistically significant degree of separation ($p < 0.05$): Canon and MAPIR images of *Buddleia*, Canon images of *Physocarpus opulifolius*, and MAPIR images of *Hydrangea paniculata* (Table 5). Of these four, models trained with MAPIR or Canon images of NS and HWS *Buddleia* were also able to discriminate NS and LWS plants with high separation (Table 6). Four datasets produced models with a marginally significant degree of separation ($0.05 < p < 0.10$): Canon and MAPIR images of *Hydrangea quercifolia*, Canon images of *Hydrangea paniculata*, and MAPIR images of *Physocarpus opulifolius* (Table 5). Images of *Spiraea japonica* were not tested because the HWS class in the training set did not meet the minimum of 10 images required by the Visual Recognition API. Overall, models trained with four of five species tested achieved marginal significance or better ($p < 0.10$) in one or both cameras (Figures 7 and 8).

Results were compared to a previous study by de Castro et al. [18] that described the same dataset by masking the background and comparing mean pixel values in stressed and non-stressed plants. The three wavelengths detected by each camera were delineated and differences between treatments were evaluated by performing an analysis of variance (ANOVA) significance by a Tukey honestly significant difference (HSD) range test. Experiments that demonstrated a significant difference in mean pixel value between water stress treatments in one or more wavelengths ($p < 0.05$) are highlighted

green in Table 5. Marginal significance is not shown because de Castro et al. [18] did not report specific *p*-values.

Table 5. Performance of models trained to classify HWS and NS images. Models achieving a statistically significant degree of separation (*p*-value < 0.05) are highlighted green and models achieving a marginal degree of separation (0.05 < *p*-value < 0.10) are highlighted yellow.

Species	Camera	Mean AUC	St. Dev	<i>P</i> -value	de Castro et al.
<i>Buddleia</i> (BUD)	Canon	0.9931	0.0046	1.10E−05	Green
	MAPIR	0.9907	0.0076	2.97E−05	
<i>Cornus</i> (CO)	Canon	0.463	0.0976	0.2635	Pink
	MAPIR	0.5094	0.1661	0.4602	
<i>Hydrangea</i> (HP)	Canon	0.6448	0.1381	0.0854	Pink
	MAPIR	0.7381	0.1036	0.0221	
<i>Hydrangea</i> (HQ)	Canon	0.6639	0.1284	0.0626	Yellow
	MAPIR	0.7113	0.1946	0.081	
<i>Physocarpus</i> (PO)	Canon	0.8177	0.1759	0.0344	Green
	MAPIR	0.5677	0.0503	0.0573	
<i>Spiraea</i> (SJ)	MAPIR	NA	NA	NA	Pink

Table 6. Models that achieved a statistically significant degree of separation on HWS images were also used to classify LWS images.

Species	Camera	Mean AUC	St. Dev.	<i>P</i> -value	de Castro et al.
<i>Buddleia</i> (BUD)	Canon	0.9884	0.0139	1.01E−04	Green
<i>Buddleia</i> (BUD)	MAPIR	0.9653	0.0306	5.40E−04	
<i>Hydrangea</i> (HP)	MAPIR	0.6144	0.1126	0.0897	Pink
<i>Physocarpus</i> (PO)	Canon	0.7167	0.1725	0.0643	




Figure 7. Performance of models trained to classify HWS and NS images.




Figure 8. Models that achieved a statistically significant degree of separation on HWS images were also used to classify LWS images.

4. Discussions

Unlike traditional machine vision models that require users to manually select features, CNNs have layers of neurons that allow them to automatically learn relevant features from data. CNNs improve with each training example by iteratively rewarding neurons that amplify aspects of the image that are important for discrimination and suppressing those that do not. For example, in traditional techniques, the background must be manually segmented prior to analysis. By contrast, CNNs can automatically ‘learn’ to ignore the background because it is not relevant to the classification task. Similarly, rather than manually delineating spectral indices thought to be correlated with plant health, networks can infer relevant transformation of the input color channels from data. Low level features inferred by the network feed into higher-order features such as the specific location or pattern of discoloration within the plant. Information from spectral indices may combine with other features such as the unique structure of sagging branches or the distinct texture created by the shadows from wilted leaves. Thus, CNNs can learn multiple features of the training images and are not limited by a priori hypotheses.

Models tested in this study demonstrated significant variation in their ability to identify water stress in different species. Models trained on *Buddleia* achieved near-perfect separation while those trained on *Cornus* approximated random classification. Such variation is consistent with previous literature showing differences in morphological and physiological responses to water stress across genera, species, and even cultivar. In Michigan, Warsaw et al. [37] tracked daily water use and water use efficiency of 24 temperate ornamental taxa from 2006 and 2008. Daily water use varied from 12 to 24 mm per container and daily water use efficiency (increase in growth index per total liters applied) varied from 0.16 to 0.31. Of the similar taxa used, *Buddleia davidii* ‘Guinevere’ (24 mm per container) had the greatest water use followed by *Spirea japonica* ‘Flaming Mound’ (18 mm per container), *Hydrangea paniculata* ‘Unique’ (14 mm per container), and *Cornus sericea* ‘Farrow’ (12 mm per container) with estimated crop coefficients (KC) of 6.8, 5.0, 3.6, and 3.4, respectively. Low-water tolerant taxa such as *Cornus* may simply not have been demonstrating symptoms of water stress when they were photographed. Models that achieved moderate performance were likely provided with too few examples to distinguish patterns relevant to the classification task from those specific to the training data, causing them to generalize poorly to new data during the testing phase. Such overfitting

bias can be overcome by training models with a larger and more diverse set of training images. Varying the location, weather, and growing period in which images are taken, for example, can force models to learn features that generalize to all conditions. Future studies can also use images of plants with multiple degrees of water stress to train regression models that return a value along a numeric scale rather than a stressed or not-stressed binary.

While CNNs' complicated nature prevents us from knowing what features are driving the model, insight can be gained from the conditions in which classifiers succeed or fail. For example, classifiers trained by pooling images of all species had significantly lower performance than classifiers trained with images of just one species despite having a considerably larger training set. This suggests that symptoms of water stress differ from one species to the next. Subsequent studies can identify what features are driving the model by iteratively removing them from the image. For example, one experiment could train models with individual R, G, or near-infrared channels to determine if certain spectral indices are more sensitive to water stress than others. Another experiment could crop a rectangle circumscribed to the plant in order to see if plant shape or other peripheral features aid the classifier. Features that significantly reduce performance when removed may represent biologically relevant phenotypes that are worthy of further study.

5. Conclusions

Our findings confirm that the IBM Watson Visual Recognition service can be used to identify early indicators of water stress in ornamental shrubs despite constraints such as small sample size, low image resolution, and lack of clear visual differences. Watson-generated models were able to detect indicators of stress after 48 hours of water deprivation with a significant to marginally significant degree of separation in four out of five species tested ($p < 0.10$). Models trained on images of *Buddleia* achieved near-perfect separation after only 24 hours with a max AUC of 0.9884. Furthermore, unlike traditional algorithms that require users to manually select plant parameters believed to correlate with health status, CNNs were able to automatically infer relevant features from the training data and combine multiple types of visual information. Despite this, not all models were successful. Failure of models trained on images of *Cornus* was consistent with previous literature, suggesting higher water stress tolerance in *Cornus* compared to the other species tested. Because all plants were grown in the same experimental area, authors cannot be certain that these models will generalize well to new situations.

Future studies can focus on improving model accuracy and generalizability by increasing the number of training examples and varying the conditions in which images are taken. Fully trained networks can also be introspected to give biological backing to the most predictive features. Other studies can expand the application of this workflow by testing data collected with different sensors and on different species. These experiments provide a valuable case study for the use of CNNs to monitor plant health. Brought to scale, artificial intelligence frameworks such as these can drive responsive irrigation systems that monitor plant status in real time and maximize sustainable water use.

Author Contributions: Conceptualization, A.I.d.C., J.M.P., J.M.M., J.R., and J.S.O.J.; Methodology, D.F., S.G., and D.H.S.; Software, D.F. and S.G.; Investigation, A.I.d.C., J.M.P., J.M.M., J.R., and J.S.O.J.; Resources, J.S.O.J. and J.M.M.; Validation, A.I.d.C., and J.M.P.; Formal Analysis, D.H.S., D.F., and S.G.; Writing—original draft preparation, D.F.; Writing—review and editing, D.F., S.G., D.H.S., J.M.M., J.R., and J.S.O.J.; Supervision, J.M.M.; Project Administration, J.M.M.; Funding Acquisition, J.M.M. and J.R.

Funding: This work was partially supported by a grant from the J. Frank Schmidt Family Charitable Foundation and is based on work supported by NIFA/USDA under project numbers SC-1700540, SC-1700543 and 2014-51181-22372 (USDA-SCRI Clean Water3). Research of Drs. Peña and de Castro was financed by the “Salvador de Madariaga” for Visiting Researchers in Foreign Centers Program (Spanish MICINN funds) and the Intramural-CSIC Project (ref. 201940E074), respectively.

Acknowledgments: The authors would like to thank Julie Brindley and Ana Birnbaum for their support and assistance in this project. Special thanks to IBM for supporting our research with access to their artificial intelligence services.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Maja, J.M.; Robbins, J. Controlling irrigation in a container nursery using IoT. *AIMS Agric. Food* **2018**, *3*, 205–215. [[CrossRef](#)]
2. Cohen, Y.; Alchanatis, V.; Meron, M.; Saranga, S.; Tsipris, J. Estimation of leaf water potential by thermal imagery and spatial analysis. *J. Exp. Bot.* **2005**, *56*, 1843–1852. [[CrossRef](#)] [[PubMed](#)]
3. Gago, J.; Douthe, C.; Coopman, R.E.; Gallego, P.P.; Ribas-Carbo, M.; Flexas, J.; Escalona, J.; Medrano, H. UAVs challenge to assess water stress for sustainable agriculture. *Agric. Water Manag.* **2015**, *153*, 9–19. [[CrossRef](#)]
4. Jacobson, B.K.; Jones, P.H.; Jones, J.W.; Paramore, J.A. Real-time greenhouse monitoring and control with an expert system. *Comput. Electron. Agric.* **1989**, *3*, 273–285. [[CrossRef](#)]
5. Kim, Y.; Evans, A.R.G. Software design for wireless sensor-based site-specific irrigation. *Comput. Electron. Agric.* **2009**, *66*, 159–165. [[CrossRef](#)]
6. Stone, K.C.; Smajstrla, A.G.; Zazueta, F.S. Microcomputer-based data acquisition system for continuous soilwater potential measurements. *Soil Crop Sci. Soc. Fla. Proc.* **1985**, *44*, 49–53.
7. Coates, R.W.; Delwiche, J.M.; Brown, P.H. Design of a system for individual microsprinkler control. *Trans. ASABE* **2006**, *49*, 1963–1970. [[CrossRef](#)]
8. USDA; National Agricultural Service. *2017 Census of Agriculture, United States Summary and State Data; Geographic Area Series, Part 51*. Publ AC-17-A-51; USDA: Washington, DC, USA; National Agricultural Service: Washington, DC, USA, 2019; Volume1, p. 820.
9. Vilsack, T.; Reilly, A.J.T. *2012 Census of Agriculture. Census of Horticultural Specialties (2014); Special Studies. Part 3*. AC-12-SS-3; 2015; Volume 3, p. 422. Available online: https://www.nass.usda.gov/Publications/AgCensus/2012/Online_Resources/Farm_and_Ranch_Irrigation_Survey/fris13.pdf (accessed on 8 November 2019).
10. Vilsack, T.; Reilly, A.J.T. *2012 Census of Agriculture. Farm and Ranch Irrigation Survey (2013); Special Studies. Part 1*. AC-12-SS-1; USDA NASS; Volume 3, pp. 84–195. Available online: https://www.nass.usda.gov/Publications/AgCensus/2012/Online_Resources/Farm_and_Ranch_Irrigation_Survey/fris13.pdf (accessed on 8 November 2019).
11. USDA. USDA Agricultural Resource Management Survey (ARMS) 2016. Available online: <https://www.ers.usda.gov/topics/farm-economy/farm-labor/> (accessed on 16 September 2019).
12. Posadas, B.C.; Knight, P.R.; Coker, C.H.; Coker, R.Y.; Langlois, S.A. Hiring Preferences of Nurseries and Greenhouses in Southern United States. *Hort Technol.* **2014**, *24*, 101–117.
13. Robbins, J.A. Small unmanned aircraft systems (sUAS): An emerging technology for horticulture. *Hortic. Rev.* **2018**, *45*, 33–71.
14. De Castro, A.I.; Ehsani, R.; Ploetz, R.; Crane, J.H.; Abdulridha, J. Optimum spectral and geometric parameters for early detection of laurel wilt disease in avocado. *Remote Sens. Environ.* **2015**, *171*, 33–44. [[CrossRef](#)]
15. De Castro, A.I.; Torres-Sánchez, J.; Peña, J.M.; Jiménez-Brenes, F.M.; Csillik, O.; López-Granados, F. An Automatic Random Forest-OBIA Algorithm for Early Weed Mapping between and within Crop Rows Using UAV Imagery. *Remote Sens.* **2018**, *10*, 285. [[CrossRef](#)]
16. Maja, J.M.; Campbell, T.; Camargo-Neto, J.; Astillo, P. Predicting cotton yield of small field plots in a cotton breeding program using UAV imagery data. In Proceedings of the Autonomous Air and Ground Sensing Systems for Agricultural Optimization and Phenotyping, Baltimore, MD, USA, 17 May 2016; p. 98660C. [[CrossRef](#)]
17. Reay-Jones, F.; Greene, J.; Maja, J.M.; Bauer, P. Using remote sensing to improve the management of stink bugs in cotton in South Carolina. In Proceedings of the XXV International Congress of Entomology, Orlando, FL, USA, 26 September 2016.
18. De Castro, A.; Maja, J.M.; Owen, J.; Robbins, J.; Peña, J.M. 2018. Experimental approach to detect water stress in ornamental plants using sUAS-imagery. In Proceedings of the Autonomous Air and Ground Sensing Systems for Agricultural Optimization and Phenotyping, SPIE Commercial + Scientific Sensing and Imaging, Orlando, FL, USA, 21 May 2018. [[CrossRef](#)]
19. Salami, E.; Barrado, C.; Pastor, E. UAV flight experiments applied to remote sensing of vegetated areas. *Remote Sens.* **2014**, *6*, 11051–11081. [[CrossRef](#)]

20. Shahbazi, M.; Jérôme, T.; Ménard, P. Recent applications of unmanned aerial imagery in natural resource management. *GISci. Remote Sens. J.* **2014**, *51*, 1548–1603. [[CrossRef](#)]
21. Jackson, R.D. Canopy temperature and crop water stress. In *Advances in Irrigation*; Hillel, D., Ed.; Academic Press: New York, NY, USA, 1982; Volume 1, pp. 43–80.
22. Bellvert, J.; Zarco-Tejada, P.J.; Girona, J.; Fereres, E. Mapping crop water stress index in a ‘Pinot-noir’ vineyard: Comparing ground measurements with thermal remote sensing imagery from an unmanned aerial vehicle. *Precis. Agric* **2014**, *15*, 361. [[CrossRef](#)]
23. Garcia Ruiz, F.; Sankaran, S.; Maja, J.M.; Lee, W.S.; Rasmussen, J.; Ehsani, R. Comparison of two aerial imaging platforms for identification of Huanglongbing-infected citrus trees. *Comput. Electron. Agric.* **2013**, *91*, 106–115. [[CrossRef](#)]
24. Stagakis, S.; Gonzalez-Dugo, V.; Cid, P.; Guillen-Climent, M.L.; Zarco-Tejada, P.J. Monitoring water stress and fruit quality in an orange orchard under regulated deficit irrigation using narrow-band structural and physiological remote sensing indices. *ISPRS J. Photogramm. Remote Sens.* **2012**, *71*, 47–61. [[CrossRef](#)]
25. Zovkoa, M.; Žibrat, U.; Knapič, M.; Kovačić, M.B.; Romić, D. Hyperspectral remote sensing of grapevine drought stress. *Precis. Agric* **2019**, *20*, 335–347. [[CrossRef](#)]
26. Fulcher, A.; LeBude, A.V.; Owen, J.S., Jr.; White, S.A.; Beeson, R.C. The Next Ten Years: Strategic Vision of Water Resources for Nursery Producers. *HortTechnology* **2016**, *26*, 121–132. [[CrossRef](#)]
27. LeCun, Y.; Bengio, Y.; Hinton, G. Deep learning. *Nature* **2015**, *521*, 436–444. [[CrossRef](#)]
28. Herent, P.; Jegoue, S.; Wainrib, G.; Clozel, T. Brain age prediction of healthy subjects on anatomic MRI with deep learning: going beyond with an “explainable AI” mindset. *bioRxiv* **2018**, 413302. [[CrossRef](#)]
29. Zhou, B.; Khosla, A.; Lapedriza, A.; Aude, O.; Torralba, A. Learning deep features for discriminative localization. *arXiv* **2015**, arXiv:1512.04150.
30. Zhang, Q.; Zhu, S. Visual interpretability for deep learning: a survey. *Front. Inf. Technol. Electron. Eng.* **2018**, *19*, 27–39. [[CrossRef](#)]
31. Awad, M.; Khanna, R. *Efficient Learning Machines*; Springer Link Apress: Berkeley, CA, USA, 2015; pp. 167–184. ISBN 978-1-4302-5989-3.
32. Ahmad, S.; Abdou, M.; Perot, E.; Yogamani, S. Deep reinforcement learning framework for autonomous driving. *arXiv* **2017**, arXiv:1704.02532.
33. Taigman, Y.; Yang, M.; Ranzato, M.; Wolf, L. DeepFace: closing the gap to human-level performance in face verification. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2014), Columbus, OH, USA, 24–27 June 2014; pp. 1701–1708.
34. Knyazikhin, Y.; Schull, M.A.; Stenberg, P.; Möttus, M.; Rautiainen, M.; Yang, Y.; Marshak, A.; Carmona, P.L.; Kaufmann, R.K.; Lewis, P. Hyperspectral remote sensing of foliar nitrogen content. *Proc. Natl. Acad. Sci. USA* **2013**, *110*, E185–E192. [[CrossRef](#)]
35. Gómez-Casero, M.T.; López-Granados, F.; Peña-Barragán, J.M.; Jurado-Expósito, M.; García Torres, L.; Fernández-Escobar, R. Assessing Nitrogen and Potassium Deficiencies in Olive Orchards through Discriminant Analysis of Hyperspectral Data. *J. Am. Soc. Hort. Sci.* **2007**, *132*, 611–618. [[CrossRef](#)]
36. Alom, M.Z.; Taha, T.M.; Yakopcic, C.; Westberg, S.; Sidike, P.; Nasrin, M.S.; Van Essen, B.C.; Awwal, A.A.S.; Asari, V.K. The history began from alexNet: A comprehensive survey on deep learning approaches. *arXiv* **2018**, arXiv:1803.01164.
37. Warsaw, A.L.; Fernandez, R.T.; Cregg, B.M.; Andersen, J.A. Water conservation, growth, water use efficiency of container grown woody ornamentals irrigated based on daily water use. *HortScience* **2009**, *44*, 1308–1318. [[CrossRef](#)]

