Assessment of avocado water stress using multispectral imaging and deep learning: Toward scalable, cost-effective irrigation monitoring

by Estrada, J.S., Cetin, N., Sacilik, K., Ulu, B., Ulu, B. and Cheein, F.A.

Estrada, J.S., Cetin, N., Sacilik, K., Ulu, B., Ulu, B. and Cheein, F.A. (2026) 'Assessment of avocado water stress using multispectral imaging and deep learning: Toward scalable, cost-effective irrigation monitoring', *Expert Systems with Applications*, 326, article 132688.

Expert Systems With Applications 326 (2026) 132688

Contents lists available at ScienceDirect

Expert Systems With Applications

journal homepage: www.elsevier.com/locate/eswa

Assessment of avocado water stress using multispectral imaging and deep learning: Toward scalable, cost-effective irrigation monitoring

Juan Sebastián Estrada [a], Necati Cetin [b], Kamil Sacilik [b], Banu Ulu [c], Burak Ulu [d,e,f], Fernando Auat Cheein [g,*]

[a] Department of Electronic Engineering, Universidad Tecnica Federico Santa Maria, Av. España 1680, Valparaíso, Chile
[b] Faculty of Agriculture, Department of Agricultural Machinery and Technologies Engineering, Ankara University, Ankara, Turkey
[c] Faculty of Engineering, Architecture and Design, Department of Software Engineering, Kayseri University, Kayseri, Turkey
[d] Faculty of Engineering, Department of Mechatronics Engineering, Erciyes University, Kayseri, Turkey
[e] Erciyes University, ArGePark Research Building, Kayseri, Turkey
[f] Ulu Robotik Ltd., Technology Development Zone, Erciyes Teknopark, Kayseri, Turkey
[g] Department of Engineering, Harper Adams University, Newport, England, UK

ARTICLE INFO

Keywords:
Dehydration stage
Multi-spectral images
Hyperspectral reflectance values
Deep learning
Machine learning

ABSTRACT

The identification of leaf water content in avocados can help to improve the irrigation regimes, thus alleviating the great water consumption of the avocado production. The classification of leaves according to their dehydration stage can be used as an early assessment tool for leaf water content identification. This study focuses on classifying avocado leaves based on their dehydration stage, using two approaches: hyperspectral reflectance values with classic machine learning models, and multi-spectral images with deep convolutional neural networks. The classic machine learning algorithms trained include random forest, k-nearest neighbours, support vector machine, bagging, decision tree, reduced error pruning tree, simple classification and regression tree, Bayesian networks, logistic model tree, linear discriminant analysis, multinomial logistic regression, and multilayer perceptron. The deep learning approach employed convolutinal neural network, specifically ConvNeXt-Tiny, MobileNetV3-Large, and ResNet18. The models were evaluated using common metrics such as accuracy, precision, f1 score, receiving operating characteristic curves, area under curve and confusion matrices. Among the machine learning models, linear discriminant analysis achieved the highest accuracy of 0.98. Meanwhile, ResNet18 outperformed the other CNNs, achieving an accuracy of 0.99 on the validation dataset. These results highlight the effectiveness of the proposed methods for classifying avocado leaves based on their dehydration level.

1. Introduction

Fruits and vegetables are an essential part of a healthy diet because they are low in calories and provide a wide range of essential nutrients, including vitamins, minerals, and fiber. Rich in micronutrients and bioactive compounds, they play a critical role in disease prevention and overall health (Alemu, 2024; Guiné et al., 2024). The avocado (*Persea americana* Mill.), a tropical fruit native to Mexico and Central America, is renowned for its rich nutritional content and bioactive compounds (Santos et al., 2024). As a climacteric fruit, its respiratory and metabolic processes, combined with external factors, make it particularly susceptible to damage. Known globally for its high levels of carotenoids, antioxidants, and monounsaturated fatty acids, the avocado tree is evergreen and produces a distinctive pear-shaped fruit with peel, flesh, and seeds (Lieu et al., 2024). Often referred to as a "superfood," the avocado is celebrated for its nutritional and phytochemical properties, which are linked to numerous health benefits. Its oil, rich in monounsaturated fats, is especially noted for supporting cardiovascular health by reducing lipid adhesion and the formation of blood clots (Gonçalves et al., 2024).

Due to their nutritional benefits and versatile culinary applications, avocados are a significant agricultural product with strong global demand (Ramos-Aguilar et al., 2019). The fruit is cultivated in tropical and subtropical regions worldwide, with Mexico leading as the top producer (Menge et al., 2012). Global avocado production has been steadily increasing, reaching approximately 9 million tons in 2022. The top four avocado-producing countries are Mexico, which accounts for the largest share with 2.53 million tons, followed by Colombia with 1.09 million tons, Peru with 0.87 million tons, and the Dominican Republic with 0.74 million tons (Food and Agriculture Organization of the United Nations, 2024).

* Corresponding author.
E-mail address: f.auat@hw.ac.uk (F. Auat Cheein).

https://doi.org/10.1016/j.eswa.2026.132688
Received 2 May 2025; Received in revised form 1 October 2025; Accepted 28 April 2026
Available online 11 May 2026

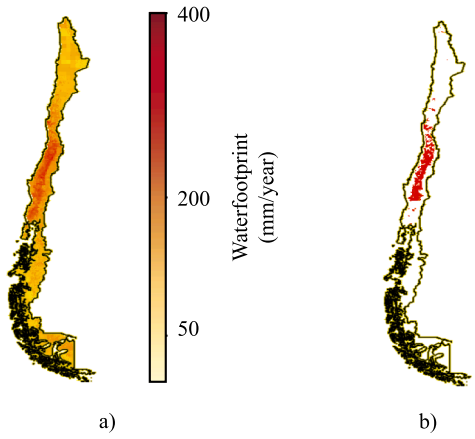

Fig. 1. a) Map of the water footprint of Chile, b) Map of the avocado water foot print in Chile.

The growing global market for avocados is associated with increasing water conflicts (Sommaruga & Eldridge, 2020), particularly in regions already experiencing droughts and water scarcity, such as Chile (Government of Chile and Ministry of Environment, 2016; UNCCD, 2016) and Mexico (Sommaruga & Eldridge, 2020).

In Chile, although it is not among the largest avocado-producing countries, its avocado water footprint is one of the highest among avocado-producing nations. For comparison, Chile's water footprint is slightly higher than Brazil's, despite Brazil producing twice as many avocados as Chile (Network, 2024; Sommaruga & Eldridge, 2020).

Fig. 1 shows the map of the current water footprint in Chile, and the main regions were avocados are grown (Network, 2024). It is clear that there is a relation between high water consumption particularly in regions of central Chile were avocados are grown. The use of water for avocado production has had significant effects, particularly in the Petorca province in the Valparaíso region, Chile, where most avocados are grown and the water cost to produce one fruit is of 320 liters (Lazzaris, 2024). This region faces severe water-related challenges, which are exacerbated by the prolonged drought affecting the area over the past decades (Madariaga et al., 2021). These water shortages not only impact avocado production but also create serious issues for the local population (Duran-Llacer et al., 2020; Schiappacasse et al., 2024).

To mitigate the negative impacts of avocado production on water management, especially in countries where water scarcity is a critical concern, it is essential to implement effective measures. These measures include optimizing avocado production processes (Sommaruga & Eldridge, 2020) and improving irrigation practices for the crop (Government of Chile and Ministry of Environment, 2016).

The detection of leaf water content offers a valuable reference for improving irrigation regimes, thereby helping to address water scarcity issues (Liu et al., 2018). While various methods exist for detecting leaf water content, traditional approaches often involve destructive sampling, are time- and resource-intensive, and lack scalability for large areas (Lawson & Hawkes, 1989; Liu et al., 2018). As a result, the development of non-destructive and rapid methods for detecting leaf water content has emerged as a significant focus of research.

One of the most common methods for detecting leaf water content is through hyperspectral information. Reflectance in the infrared region of the spectrum can provide insights into the chemical properties of vegetation (Wei et al., 2019). In particular, bands in the short infrared region (1000 to 2500 nm) are sensitive to changes in water content (Arevalo-Ramirez et al., 2020). As a result, there have been numerous efforts to predict leaf water content using vegetation indices derived from hyperspectral reflectance values across a wide range of species, including eucalyptus, pines (Villacrés et al., 2019), citrus tangerines (Dou et al., 2024), tea (Wei et al., 2019), tomato (Alordzinu et al., 2021), and potato (Liu et al., 2018), among others (Jin et al., 2023; Surase et al., 2019).

Although hyperspectral imaging (HS) is widely used to assess leaf water content, it has certain limitations. Many studies rely on handheld spectrometers to collect data (Dou et al., 2024; Liu et al., 2018; Surase et al., 2019; Villacrés et al., 2019), which necessitates manual sample collection and measurement; moreover, handheld spectrometers can only record data from a single point on the leaf (Estrada et al., 2023). In contrast, hyperspectral and multispectral (MS) cameras provide high spatial resolution and can be mounted on robotic platforms, enabling efficient surveys over large areas (Estrada et al., 2023). For instance, Liu et al. (2018) employed a hyperspectral camera to assess the water content of tea leaves, while Estrada et al. (2024) used multispectral imaging combined with deep learning to classify *Eucalyptus globulus* leaves by dehydration stage. Despite these practical constraints, HS sensing remains a standard for leaf-water assessment because of its fine spectral sampling of water-absorption features (approximately 970, 1200, 1450, 1940 nm) and a well-established body of water-sensitive indices (e.g., NDWI, MSI, NMDI, WI, DDI, SIWSI), which support continuous water-content regression and band-level interpretability.

At the same time, practical factors such as sensor cost, large data volumes (and the associated preprocessing/dimensionality-reduction), and sensitivity to illumination and calibration in field deployments motivate the investigation of few-band MS solutions for deployable monitoring. Although they offer less spectral detail and a narrower index repertoire than HS, MS cameras especially those including red-edge and near-infrared bands-provide lower costs, simpler acquisition/processing pipelines, faster operation, and easier integration on industrial lines, UAVs, and greenhouse systems. In our case, an MS + lightweight-CNN pipeline attains accuracy comparable to HS baselines for the discrete dehydration-stage classification that is the focus of this study, while being more readily deployable for irrigation monitoring; HS remains preferable when continuous, quantitative water-content estimation is required.

When comparing the performance of hyperspectral and multispectral data for estimating vegetation properties such as chlorophyll content, hyperspectral imagery demonstrates superior accuracy due to its higher number of spectral bands, which capture finer details in spectral reflectance. However, the improvements in accuracy are often marginal compared to models based on multi-spectral data (Lu et al., 2019). This indicates that multi-spectral cameras can also be effectively utilized to develop models for predicting vegetation traits with an acceptable degree of error, making them a viable and often more cost-efficient alternative.

Deep learning architectures, particularly convolutional neural networks (CNNs), have been widely applied in agriculture to detect stress factors in vegetation, such as pests and drought stress, using regular (RGB) and multi-spectral images (Estrada et al., 2024; Wang et al., 2024). For instance, models like ResNet-18 have been utilized to detect Hass avocados (Ramirez et al., 2021) and classify leaf diseases in rice plants (Ahad et al., 2023). Similarly, MobileNet has been successfully employed for biotic stress classification in coffee leaves (Tassis & Krohling, 2022) and for disease classification in olive leaves (Raouhi et al., 2022). Furthermore, Wu et al. (2023) applied the ConvNeXt model to classify soybean leaf diseases. In addition, Estrada et al. (2024) explored various deep learning approaches to classify *Eucalyptus globulus* leaves based on their water moisture content using both RGB and multispectral images.

In this study, our goal is to classify avocado leaves according to their dehydration stage, using hyperspectral reflectance values collected with a spectrometer and multi-spectral images captured with a Micasense Red-Edge camera. For this purpose, we trained several machine learning models, including random forest (RF), k-nearest neighbor (kNN), support vector machine (SVM), bagging (BAG), decision tree (DT), reduced error pruning tree (REPTree), simple classification and regression tree (SimpleCART), Bayesian networks (BayesNet), logistic model

Table 1
Comparative table with the current literature.

References	Plants	Sensing Modality	Task	Learning Method	Best Results
Liu et al. (2018)	Potato	HS imaging	Leaf water content regression	MLR;PCA	$R^2 = 0.9507$
Surase et al. (2019)	Multi-species	HS spectroradiometer	Water content estimation	Statistical analysis	R: healthy 0.99 / diseased 0.76 / dry 0.97
Wei et al. (2019)	Tea	HS (VNIR line-scan)	Moisture content regression	PLSR; LS-SVR (RBF), LR	LS-SVR: $R^2 = 0.956$; RMSEP = 0.027
Alordzinu et al. (2021)	Tomato	HS spectrometer	Water-stress estimation	Correlation & linear regression	Yield stress: 0.976 (sandy) / 0.982 (silty)
Vera Ramirez et al. (2021)	Avocado	RGB imaging	Stage/quality classification	CNN (ResNet-18) with transfer learning	Acc 97.95
Raouhi et al. (2022)	Olive	RGB imaging	Disease classification	Multi-CNNs	Val-Acc 0.9843
Tassis & Krohling (2022)	Coffee	RGB imaging	Biotic stress + Severity level classification	Triplet Network (+kNN) & Prototypical Network; CNNs	Acc Leaf 96.03; Symp 96.72; Sev 93.25
Ahad et al. (2023)	Rice	RGB imaging	Disease classification	CNN comparison	Acc CNN and TL $\leq$ 97%; ensemble 97.6–98%
Jin et al. (2023)	Multi-species	HS spectrometer	Moisture content regression	Index computation + linear/ exponential regression	$R^2 = 0.90$ / NRMSE 20.6
Wu et al. (2023)	Soybean	RGB imaging	Disease classification	Multi-CNNs	Acc 85.4
Dou et al. (2024)	Citrus (tangerine)	HS spectrometer	Leaf water content regression	PLSR	$R^2 = 0.749$, RMSE = 0.028
Villacrés et al. (2019)	Eucalyptus globulus, Pinus radiata	HS spectrometer	Leaf water status regression	Linear regression per index	Eucalyptus: $R^2 \leq 0.96$; Pine: 0.90
Estrada et al. (2024)	Eucalyptus globulus	Multispectral + RGB imaging	Dehydration stage classification	Multi-CNNs	MS: Xception Acc 0.813; RGB: InceptionV3 0.730
The present study	Avocado	Multispectral + HS spectrometer	Leaf water stress classification	Multi-CNNs (3) + Multi ML models (12)	DL-MS: RN-18 0.99, MNv3 0.98; ML-HS: LDA 0.98

tree (LMT), linear discriminant analysis (LDA), multinomial logistic regression (Logistic), and multilayer perceptron (MLP) using hyperspectral reflectance values as input. Additionally, we employed three deep convolutional neural networks (ResNet-18, ConvNeXt-Tiny, and MobileNetV3) with multi-spectral images as input. The dataset consists of avocado leaves at different dehydration stages. The models were evaluated using common metrics, including accuracy, precision, recall, F1-score, receiver operating characteristic (ROC) curves, area under the curve (AUC), and confusion matrices, ensuring a comprehensive assessment of their classification effectiveness Table 1.

The work is organized as follows: Section 1 includes the introduction and a review of related work; Section 2 discusses the general system architecture, the creation of the dataset, the pre-processing of the images, and the selection of machine and deep learning models; Section 3 presents the results and evaluation of the implemented models; Section 4 contains the discussion of the results; and finally, Section 5 concludes the paper and outlines future work.

2. System architecture

For this study, we created a dataset consisting of 104 avocado leaves. Spectral measurements were collected across the wavelength range from 350 to 2500 nm, alongside multi-spectral images capturing the blue, green, red, near-infrared, and red-edge bands. The leaves underwent a dehydration process, resulting in five distinct sets of measurements: fresh leaves, dehydration stages 1 to 3, and completely dry leaves.

Twelve machine learning algorithms were trained to classify the spectral measurements into their respective dehydration stages. The algorithms used were RF, kNN, SVM, BAG, DT, REPTree, SimpleCART, BayesNet, LMT, LDA, multinomial logistic regression, MLP. These algorithms were trained using the spectral reflectance values.

Deep convolutional neural networks were trained to perform the classification of the leaves in their respective dehydration stage using as input the five bands of the multi-spectral images. The selected models were ResNet-18, ConvNeXt-Tiny and MobileNetV3. A general overview of the complete study is shown in Fig. 2.

Both, the deep learning and classic machine learning models, were assessed using common metrics for classification including accuracy, precision, recall, F1-score, ROC curves, AUC and confusion matrices.

2.1. Data creation and curation

The dataset for this work consists of 104 avocado leaves, collected from trees in the Valparaiso region, Chile. The leaves underwent a drying process to create five distinct dehydration stages resulting in a total of 520 samples for training and testing purposes. We have selected the number of samples based on other studies related to leaf water content, such as Villacrés et al. (2019) which used a total of 90 samples of *Pi-*

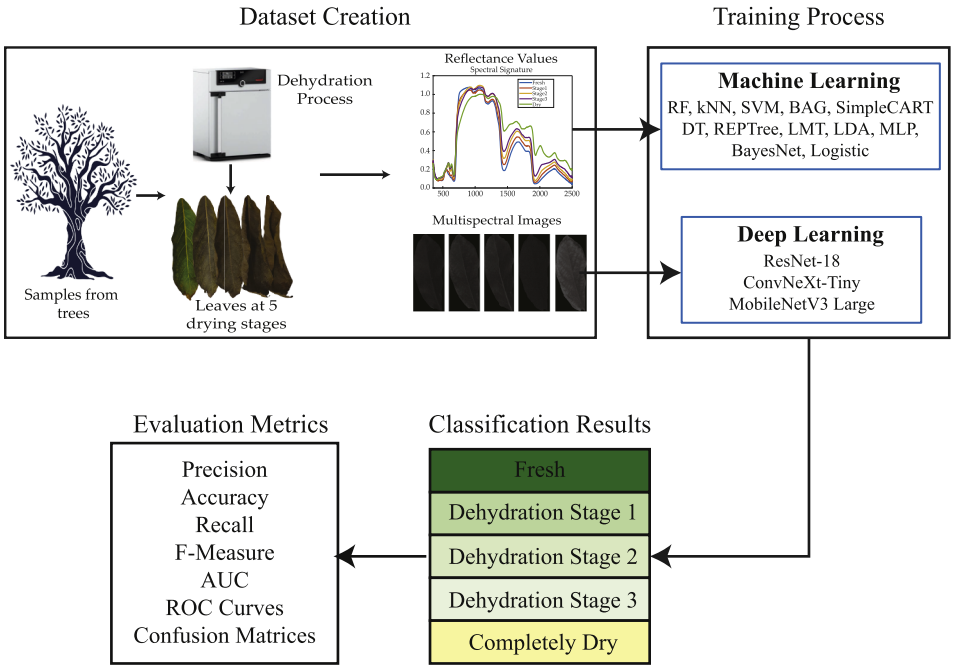

Fig. 2. General architecture implemented in this study.

nus radiata and *Eucalyptus Globulus* for training and testig purposes; and the study from Dou et al. (2024) which used 182 samples for traning their models. The data collection and dehydration process followed the guidelines presented by Villacrés et al. (2019).

The sample collection process took place during April and May of 2023. Individual branches were cut directly from the trees, and healthy leaves were selected through visual inspection. The leaves were placed in bags to preserve their moisture content and then transported to the laboratory for measurement and drying.

For measuring the spectral reflectance, an ASD TerraSpec Hi-Res spectrometer was used. The probe and light source were turned on for 15 min before the measuring process to pre-heath the instrument, following the vendor's instructions. The instrument was calibrated, and its parameters were optimized using a white Spectralon surface.

Once the instrument was calibrated, the measurement process began. Each sample was placed on a white surface (Spectralon), and the probe was positioned flush against the leaf. Three individual measurements were recorded at different locations on the leaf: one at the bottom, one at the middle section, and one at the top section. The final measurement used for the experiments was the average value of these three measurements. The spectrometer provides reflectance values in the visible and near-infrared (VNIR) region, ranging from 350 to 1000 nm, and in the short-wave infrared (SWIR) region, ranging from 1001 to 2500 nm.

The multi-spectral images were captured using a Micasense RedEdge M camera, mounted 45 cm above the leaves. Batches of four leaves were placed against a white background (white paper sheets). To ensure consistent lighting conditions, external light sources, such as windows and doors, were closed. The camera provides five spectral bands: blue (B), green (G), red (R), near infrared (NIR), and red-edge (RE).

To obtain hyperspectral reflectance values and multi-spectral images from leaves with varying leaf water content, a dehydration process was conducted. The leaves were placed in a Memmert UN30 oven for 15 min at a temperature of 65°C,a temperature chosen based on common practices for assessing the flammability of leaf contents (Villacrés et al., 2019). After 15 min, the leaves were removed from the oven, and a new set of measurements was taken. This process was repeated two more times, resulting in four dehydration stages: fresh leaves and dehydration stages 1 to 3. Finally, the leaves were placed in the oven for 24 h to remove any remaining traces of water, yielding the completely dry dehydration stage. The main difference of the dehydration stages is the time that they were dried in the oven, Table 2 showcases the dehydration stages and the total time that they were put in the oven Table 3 lists the key features of the equipment used to create the dataset.

Table 2
Description for the 5 different dehydration stages used in this work.

Dehydration Stage	Drying Time
Fresh Stage	0 min
Dehydration Stage 1	15 min
Dehydration Stage 2	30 min
Dehydration Stage 3	45 min
Completely Dehydrated Stage	24 h

As a result, a total of 520 spectral measurements were recorded (104 per dehydration stage), and 2600 images were collected, with 25 images per leaf, covering five bands (blue, green, red, near-infrared, and red-edge) at the five dehydration stages. Fig. 3 illustrates the spectral reflectances of all the samples across the different drying stages, along with the mean values. Fig. 4 displays images of the leaves at the five dehydration stages, and Fig. 5 presents the five multi-spectral bands of an avocado leaf at the fresh stage.

2.2. Image pre-processing and deep learning network architecture

The multi-spectral images underwent radiometric calibration. To achieve this, images of a calibrated panel were captured after each trial. The processing of the images was performed using the software provided by Micasense (https://github.com/micasense/imageprocessing). After radiometric calibration, the background of the images was removed, and individual leaves were segmented for subsequent analysis and use in training the deep learning models; for this purpose, we used simple thresholding using an automatic threshold value determined by Otsus's method (Otsu, 1979) then we visually inspected each leaf and, in cases where the segmentation was not satisfactory (i.e., parts of the leaf

Table 3
Instruments used to collect the hyperspectral reflectance values, multi-spectral images and to dehydrate the leaves.

Instrument	Technical Specification
Multi-Spectral camera Micasense Red-Edge M	Bands: blue (B), green (G), red (R) red-edge (RE), near-infrared (NIR) Image resolution: 12 bit raw images Sensor resolution: 1280 × 960 pixels Ground samples distance at 120 m: 8 cm/pixel Horizontal field of view: 47.2° Capture Rate: 1 capture per second
Spectrometer ASD Terraspec 4 Hi-Res	Range: 350 to 2500 nm Resolution: 3 nm at VNIR (350 to 1000 nm) 6 nm at SWIR (1001 to 2500 nm) Reproducibility: 0.1 nm Accuracy: 0.5 nm
Oven Memmert UN30	Temperature: -5 to 100 °C with respect to room temperature Digital PID control

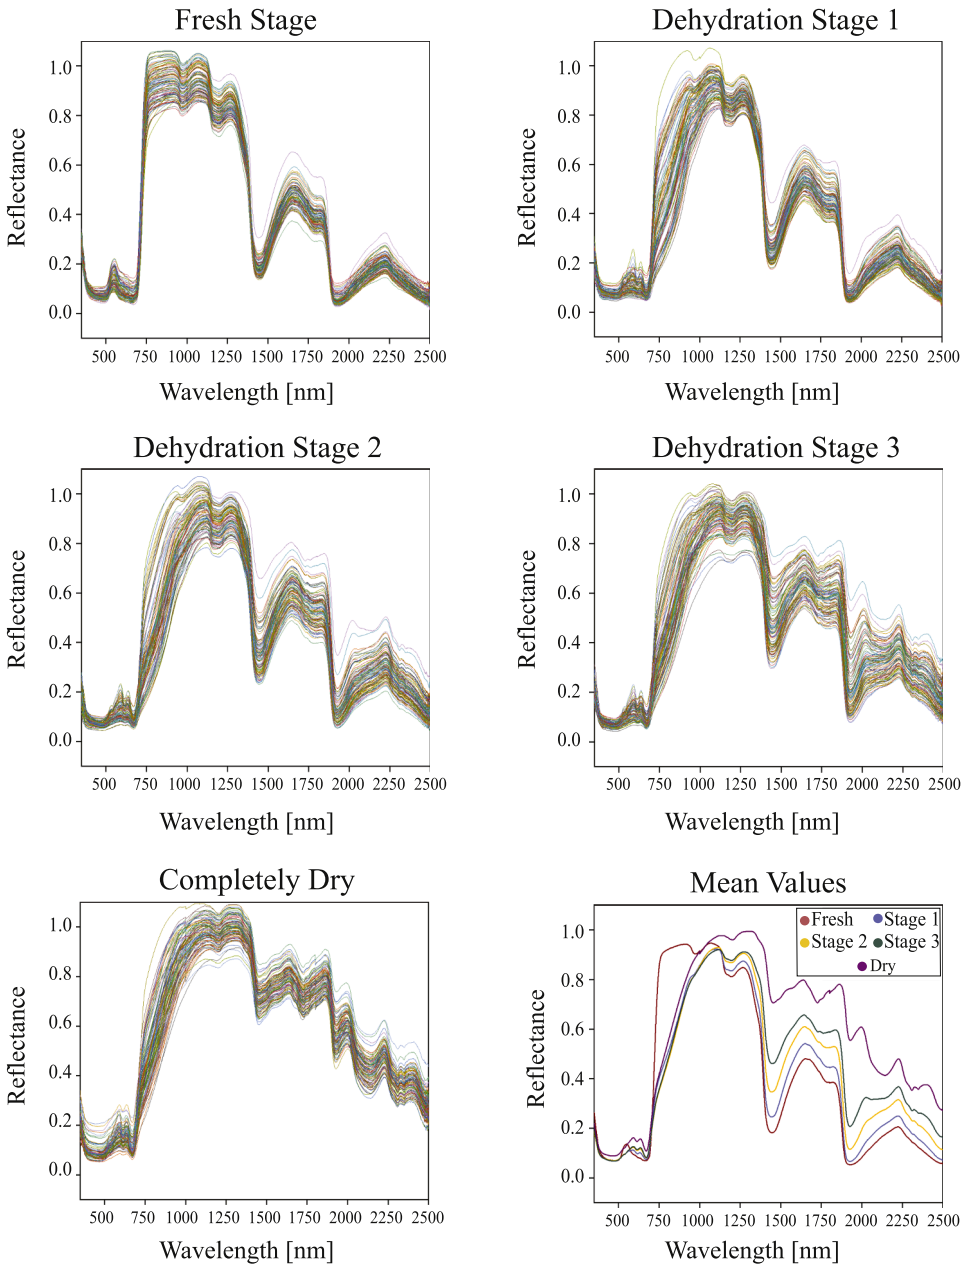

Fig. 3. Leaf reflectance at each drying stage for the avocado leaves. Figure in the bottom right corner shows the mean values for the drying stages.

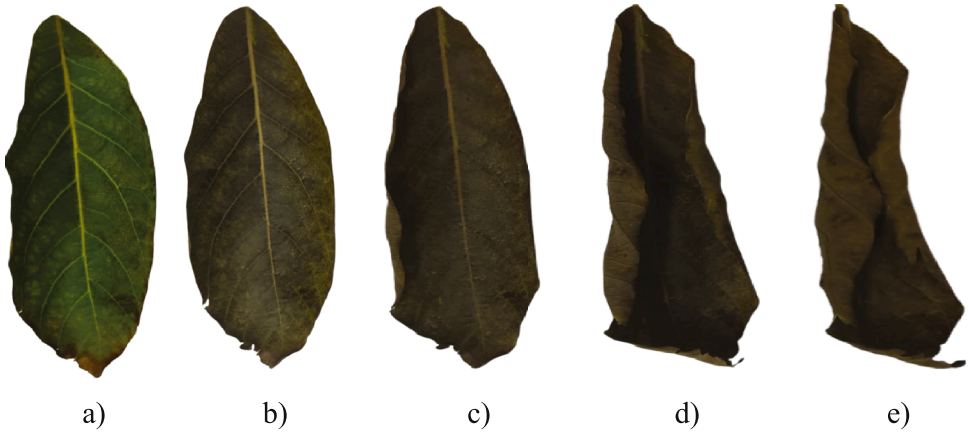

Fig. 4. Avocado leaves at the different drying stages: a) fresh stage, b) dehydration stage 1, c) dehydration stage 2, d) dehydration stage 3, and e) completely dehydrated stage.

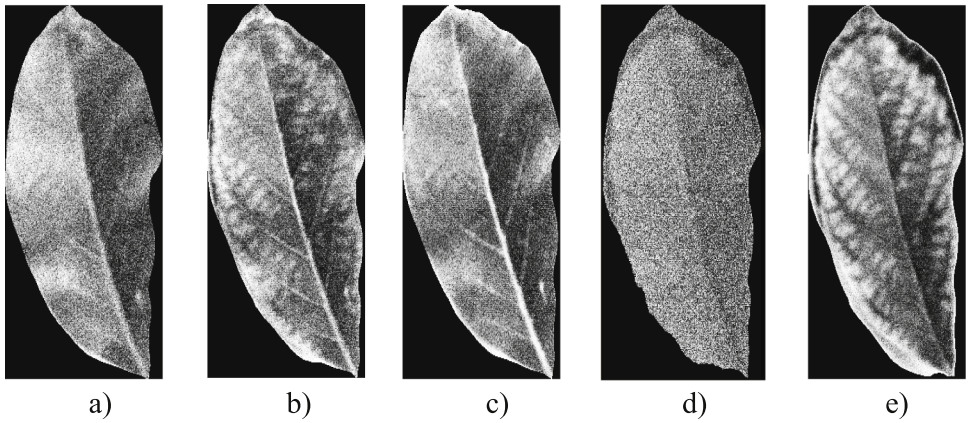

Fig. 5. Multi-spectral bands of an avocado leaf at fresh stage: a) blue, b) green, c) red, d)near infra-red, and e) red-edge.

were confused as background), manual segmentation of the background was conducted using Matlab's segmentation tool.

The dataset of images is divided into training, test, and validation sets in a 7:2:1 ratio, which is commonly used in deep learning applications. As a result, the data is randomly divided into 364 training samples, 104 test samples, and 52 validation samples. The study can be summarized in the following steps: normalizing the images, selecting the deep learning algorithms to be used, adjusting the input dimensions of the algorithms, determining the hyperparameters, training the algorithms, and obtaining the results.

The models are based on fastai algorithms and were developed on a laptop workstation featuring a 13th-generation i9 processor, an RTX 4060 graphics card, 32 GB of RAM, and the Windows 11 operating system. The code was compiled in a Jupyter notebook running on the Anaconda platform.

The images were normalized in both size and pixel values. First, all images were resized to a 401×210 format. As the images were in TIFF format with 16-bit depth, their pixel values ranged from 0 to 65,535. To normalize the pixel values to a range of 0 to 1, the images were divided by 65,535. Subsequently, the five spectral bands were concatenated, resulting in arrays of size $401 \times 210 \times 5$. These arrays were then saved as NumPy files for further processing.

The selection of deep learning algorithms was guided by their suitability for the dataset and task. The first model, ResNet-18, is a well-established and reliable architecture widely used in various image classification problems. Its fewer layers result in lower computational requirements and faster training times (He et al., 2015). Liu et al. (2022) combines classical convolutional neural networks (CNNs) with modern deep learning techniques. Its compact size reduces training time while maintaining sufficient complexity to achieve high accuracy for tasks such as the classification of avocado leaves. Lastly, MobileNetV3 Large (Qian et al., 2021) is optimized for devices with limited computational power, offering a balanced trade-off between model size and performance. This lightweight model enables efficient training and inference processes, making it suitable for potential deployment in mobile applications or field environments while maintaining high classification accuracy Table 5.

The third stage focuses on preparing the 5-channel files as input for the deep learning algorithms. Since the selected algorithms support three input channels (red, green, blue), their input layers were modified to accommodate five channels. This adjustment ensured compatibility with the data structure, which includes the additional bands. Following this modification, the process of determining the hyperparameters for each algorithm was initiated. The optimized hyperparameter values for the selected models are provided in Table 4.

As shown in Table 4, dropout and weight decay were employed to prevent overfitting. The learning rates were determined using the optimization function in the fastai library, with Adam selected as the most suitable optimizer for the dataset.

The next step is the training of the models. During the training phase, the batch size was set at 32 and the number of epochs at 15, to avoid overfitting. As a result of the training, classification is performed according to 5 classes. A schematic representation of all these stages is given in Fig. 6.

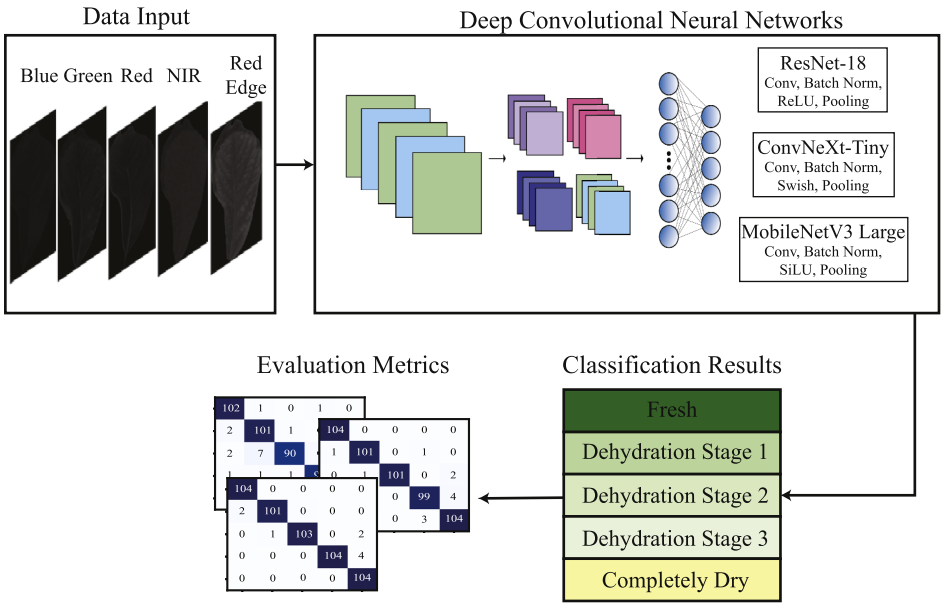

Fig. 6. Steps conducted for the training and evaluation of the deep learning models.

Table 4
Hyperparameters values for the deep learning algorithms.

Parameter	ResNet-18	ConvNeXt-Tiny	MobileNetV3
Learning Rate	0.01-0.001	0.0001-0.00001	0.001-0.0001
Batch Size	32	32	32
Weight Decay	0.01	0.01	0.01
Optimizer	Adam	Adam	Adam
Epochs	15	15	15
Dropout	0.5	0.5	0.5

2.3. Machine learning modelling

A total of 1,118,520 reflectance data points were obtained from avocado leaves across the five dehydration stages, spanning 2151 bands at 1 nm intervals within the spectrum range of 350 to 2500 nm. These data were used to train classical machine learning algorithms available in the WEKA software. The dataset comprised 104 leaves across five dehydration stages, resulting in 520 samples. The reflectance values in the VNIR (Visible and Near Infrared) and SWIR (Shortwave Infrared) regions were analyzed and visualized (see Fig. 3).

The raw data were directly used to train the models. The dataset was organized and structured specifically for each application and algorithm to ensure it was prepared appropriately for model training. For this study, 12 machine learning algorithms were used: RF, kNN, SVM, BAG, DT, REPTree, SimpleCART, BayesNet, LMT, LDA, Logistic, and MLP. For the data partitioning of the models, a k-fold cross-validation procedure was applied (Ropelewska et al., 2023). In this study, the k-fold value was set at 10. The important parameters for the models were as follows:

This work intentionally avoided the application of PCA or band selection in the hyperspectral (HS) pipeline. The aim was to establish a universal, reproducible upper-bound HS baseline for a fair comparison of accuracy and deployability with the few-band MS + CNN pipeline. Incorporating a selector would have rendered HS performance dependent on a particular selection technique and heightened the likelihood of optimistic bias, unless executed using leakage-free nested cross-validation. Notably, software band selection on hyperspectral (HS) systems does not diminish sensor-level complexity in field applications, whereas multispectral (MS) systems inherently minimize acquisition and processing demands. Consequently, we preserved the complete HS spectrum and managed model capacity using conventional regularization and hyperparameter optimization.

To evaluate model performance, the following metrics were used: precision, recall, F1-score, area under the curve (AUC), accuracy, confusion matrices, receiver operating characteristic (ROC) curves, learning rates from optimization, loss, and error rate criteria (Ropelewska et al., 2023).

2.4. Statistical significance analysis

To validate the statistical significance of performance differences between models, we conducted comprehensive statistical testing. McNemar's test is applied to compare classification performance between pairs of models using binary classification results (correct/incorrect predictions), while paired t-tests are used to compare continuous performance metrics (accuracy, precision, recall, F1-score) between models. The statistical tests are performed using Python (scipy.stats) with all pairwise comparisons between deep learning models. Results are reported with p-values and significance levels to ensure rigorous validation of our performance claims.

2.5. Robust performance metrics

To address potential class imbalance concerns and provide more comprehensive assessment, we calculated additional robust metrics including Balanced Accuracy, Matthews Correlation Coefficient (MCC), Cohen's Kappa, and both micro-averaged and macro-averaged Area Under Precision-Recall Curve (AUPRC). These metrics provide more robust evaluation compared to conventional accuracy, especially when dealing with imbalanced datasets or similar class distributions. The inclusion of these metrics ensures a comprehensive assessment that goes beyond simple accuracy measures and addresses potential limitations in conventional evaluation approaches.

2.6. Cross-Validation analysis

To ensure model generalization and address overfitting concerns, we conducted 5-fold cross-validation analysis for all deep learning models. The methodology involved calculating mean ± standard deviation for performance metrics including accuracy, precision, recall, and F1-score

Table 5
Hyperparameters of the machine learning classifiers.

Algorithms	Hyperparameters
RF	breakTiesRandomly: False; debug: False; maxDepth:0 numExecutionSlots: 1; numIterations: 100; seed: 1
k-NN	k:5; distanceWeighting: No distance Weighting; nearestNeigbourSearchAlgorithm: LinearNNSearch (distanceFunction: EuclideanDistance)
SVM	c: 1.0; calibrator: Logistic debug: False; epsilon: 1.0E-12 ; filterType: Normalize training data; kernel: PUK; randomSeed:1; toleranceParameter: 0.001
BAG	bagSizePercent: 100; calcOutOfBag: False; classifier: REPTree; numExecutionSlots:1; numIterations: 10; seed: 1
DT	binary splits: False collapseTree:True; confidenceFactor: 0.25; numFolds:5; reducedErrorPruning: False; seed: 1; useLaplace: False; useMDLcorrection: False
REPTree	initialCount: 0.0; maxDepth: -1; minVarianceProp: 0.0001; noPrunning: False; seed: 1; spreadInitialCount: False
SimpleCART	heuristic: True; minNumObj: 2.0; numFoldsPruning: 5; seed: 1; sizePer: 1.0; useOneSE: False; usePrune: True
BayesNet	estimator: SimpleEstimator; searchAlgorithm: K2; useADTree: False
LMT	fastRegression: True; minNumInstances: 15; numBoostingIterations: -1; splitOnResiduals: False; useAIC: False; weightTrimBeta:0.0
LDA	batchSize: 100; debug: False; ridge: 1.0E-8
Logistic	batchSize: 100; debug: False; maxIts: -1; ridge: 1.0E-6; useConjugateGradientDescent: False
MLP-1	hiddenLayers: a ((attribs + classes) / 2); momentum: 0.1; learningRate: 0.3; NominalToBinaryFilter: True; normalizeAttributes: True; trainingTime (epochs): 100; seed: 0; validationThreshold: 20; activation function: sigmoid
MLP-2	hiddenLayers: a ((attribs + classes) / 2); momentum: 0.1; learningRate: 0.3; NominalToBinaryFilter: True; normalizeAttributes: True; trainingTime (epochs): 300; seed: 0; validationThreshold: 20; activation function: sigmoid

across different data splits. This approach validates that reported performance is not due to overfitting and demonstrates model stability across different data partitions, providing confidence in the generalizability of our results.

3. Results

This study obtained hyperspectral images in five bands (R, G, B, Red Edge, and NIR) and spectral reflectance values from avocado leaves at different dehydration stages. The hyperspectral images were combined and classified using three deep-learning algorithms. In addition, 12 different machine-learning models were created for leaf classification using spectral reflectance values provided in the 350–2500 nm band range. To accurately sort the data into the different stages of dehydration and to test and compare the models' performance, a set of performance metrics were used. These included losses, accuracy, error rate, ROC curve, confusion matrix, precision, recall, F1-score, and AUC.

3.1. Spectral reflectance

The spectral signatures obtained at different dehydration stages and their means are shown in Fig. 3. The most notable changes were concentrated in the SWIR region, where the reflectance intensities changed significantly due to moisture loss in the leaves. The intensity and shape of the spectral reflectance in the SWIR region differed across the dehydration stages. The figure of the mean spectral reflectance shows that Fresh Leaves have the highest reflectance at approximately 750 nm and the lowest reflectance in the SWIR region. After Fresh Leaves, Dehydration Stage 1, Dehydration Stage 2, Dehydration Stage 3, and Dry Leaves had the highest reflectance values, especially in the SWIR region. Two valleys (1400–1500 nm and 1875–2000 nm) were observed for the five dehydration stages, and five peaks of different widths were observed. For all treatments, peaks with higher reflectance levels occurred at approximately 1000 and 1200 nm, whereas peaks with lower reflectance levels occurred at 2200 nm.

3.2. Deep learning training

Fig. 7 illustrates the progression of the loss function, accuracy, and error rate for the three deep learning models throughout the training phase. The horizontal axis of each graph represents the epoch number, while the vertical axis corresponds to the respective metric: loss, accuracy, or error rate.

The results demonstrate that no overfitting occurred up to the 14th iteration, thus the training steps were capped at 14. As expected, the loss and error rates consistently decreased with successive iterations, while accuracy steadily increased during training.

An adaptive learning rate strategy was employed to enhance the performance of the models. This approach was implemented using the functions provided by the fastai library. Fig. 8 shows the evolution of the learning rate for the three deep learning models during training.

3.3. Classification results

Accuracy and performance metrics values are shown in Figs. 9 and 10. ResNet-18 and MobileNet are the models with the best accuracy, reaching values of 0.990 and 0.980, respectively, which are deep learning strategies. For the classic machine learning models, LDA reached an accuracy of 0.975, comparable to the deep learning approaches. However, most of the models achieved accuracies over 0.79, except for the BayesNet model, which is the worst-performing model with an accuracy of 0.690.

Comparing the other metrics shown in Fig. 10; still the ResNet18 was the best performing model achieving values of 0.99 in the precision, recall and F1-score metrics. The second best-performing model is MobileNet large, which achieved values 0.98 for all the metrics; also it is important to note that both of these models reached an AUC of

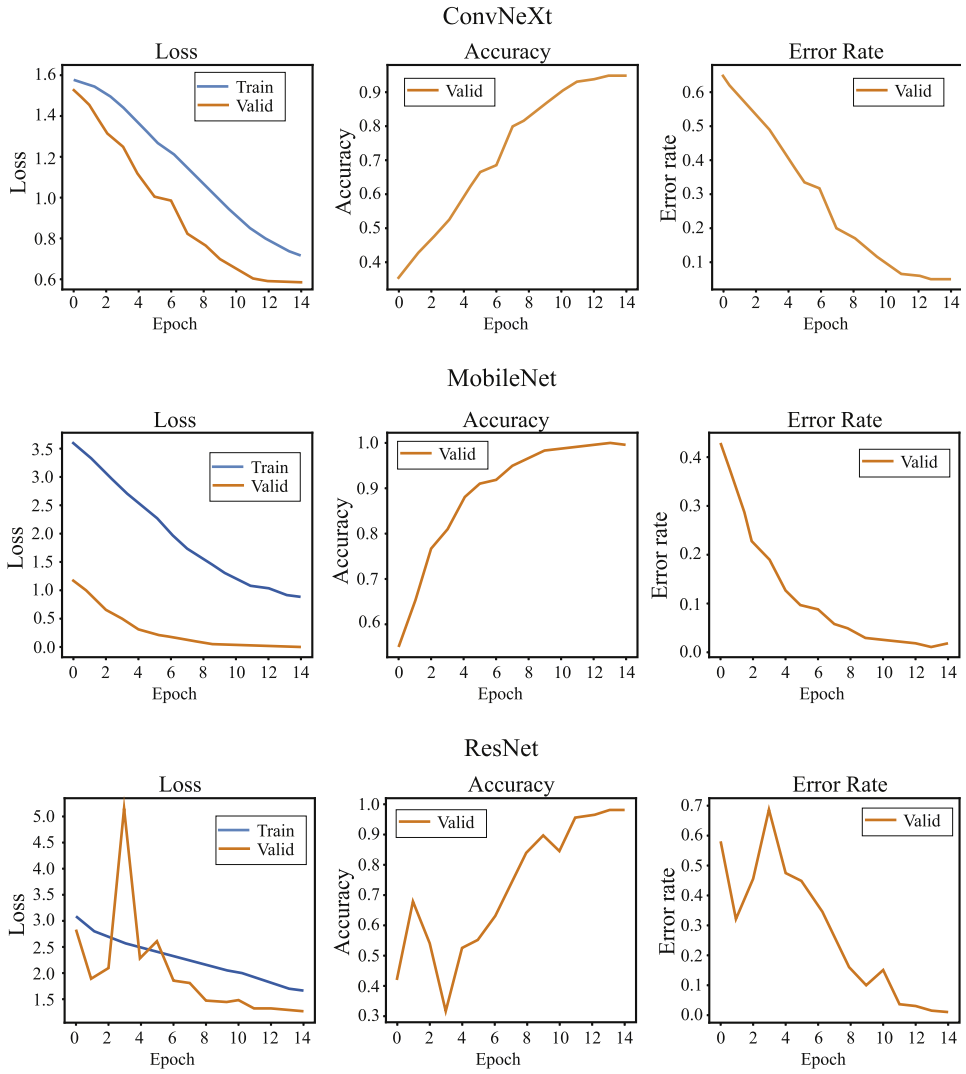

Fig. 7. Performance metrics including the loss function, accuracy and error rate for the deep learning models across the training stage.

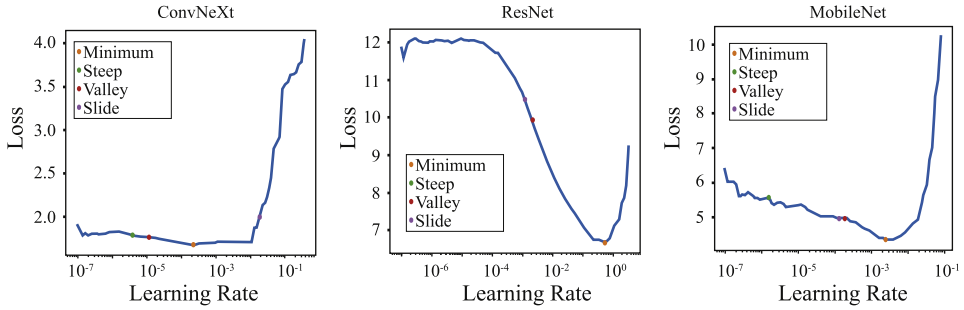

Fig. 8. Evolution of the learning rate during the training stage for the deep learning models.

1.000. The worst performing deep learning approach was the ConvNeXt-tiny model, with metrics' values of 0.94 for precision, recall and F1-score.

The classic machine learning models generally performed worse than the deep learning approaches, with MLP-2 being the best-performing model, achieving 0.990, 0.981, and 0.986 for precision, recall, and F1-score, respectively. The next best-performing machine learning model was LDA, with a precision of 0.975, recall of 0.975, and F1-score of 0.975. In contrast, the worst-performing machine learning model was BayesNet, which achieved 0.687, 0.690, and 0.688 for precision, recall, and F1-score, respectively. Following BayesNet, the next model with the lowest metrics was SimpleCART, which achieved values of 0.799, 0.802, and 0.800 for precision, recall, and F1-score, respectively. The other machine learning models achieved metrics above 0.800, with SVR, LMT, Logistic, and MLP1 reaching values over 0.9 in all the metrics. This demonstrates that these models are also capable of correctly classifying the leaves in their respective dehydration stages using the hyperspectral information.

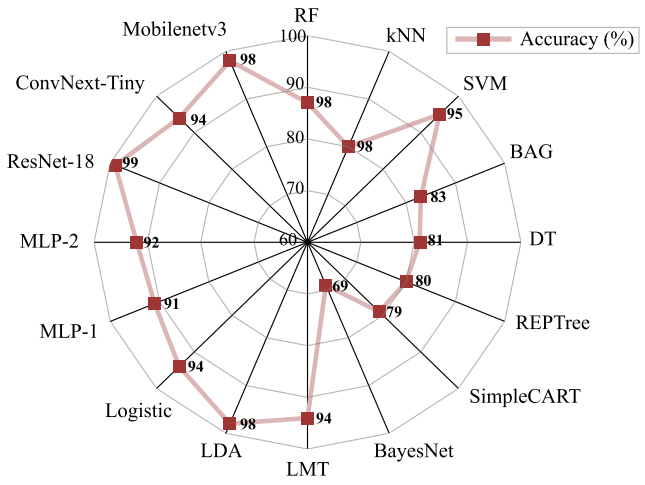

Fig. 9. Accuracy results for the machine and deep learning models.

Confusion matrices were used to further analyze the classification results achieved by the models. Figs. 11 and 12 show the confusion matrices for all the models. The confusion matrix for the ResNet-18 model, shows that only 3 samples were misclassified, 2 were classified as fresh leaves but their true class was the first dehydration stage; and 1 sample was classified as a stage 1 leaf, but the dehydration stage 2 was the true class for the sample. The confusion matrix for the MobileNetV3 model shows more classification errors in the dehydration stages 1 through 3, but correctly classifying the leaves in the fresh and completely dry stage. Finally the ConvNeXt model again presents more classification errors between dehydration stages 2 and 3, correctly classifying 90 and 92 samples respectively.

The confusion matrices for the machine learning models again show that the majority of classification errors occur in dehydration stages 1 through 3. The LDA model correctly classifies all the samples in the fresh and completely dry stages, with 4, 4, and 5 misclassified labels in stages 1 through 3, respectively. In general, all the models correctly classify most of the samples in the fresh and completely dry stages, with more than 99 samples correctly classified in the fresh stage.

The model with the worst performance, according to the confusion matrices, is the BayesNet model. While most of the samples in the fresh and dry stages are correctly classified, performance in dehydration stages 1 through 3 significantly deteriorates, particularly in dehydration stage 2, where only 44 samples out of 104 were predicted with their true label. A similar behavior is observed in the kNN, DT, REPTree, and SimpleCART models, where most classification errors occur in stage 2 of dehydration, with samples mostly misclassified as stage 1 or stage 3 leaves. However, these models still correctly classify most of the samples in the fresh and completely dry stages.

To address overfitting concerns, we present detailed learning curves for all three deep learning models (ResNet-18, MobileNetV3, ConvNeXt-Tiny) showing training vs validation performance across 15 epochs. For ResNet-18, validation loss (1.21) was consistently lower than training loss (1.55), indicating no overfitting. MobileNetV3 showed a similar pattern with validation performance matching or exceeding training performance. ConvNeXt-Tiny demonstrated stable convergence with minimal overfitting indicators. These learning curves demonstrate that models achieve stable convergence without overfitting, validating the reported high performance metrics and addressing concerns about potential overfitting in small datasets (Table 6).

Additional robust metrics provide comprehensive assessment beyond conventional accuracy measures. ResNet-18 achieved balanced accuracy of 0.980, Matthews Correlation Coefficient of 0.973, Cohen's Kappa of 0.973, and both micro-averaged and macro-averaged AUPRC of 0.980. ConvNeXt-Tiny demonstrated perfect performance across all robust metrics with values of 1.000. MobileNetV3 showed similar performance to ResNet-18 with balanced accuracy of 0.980, MCC of 0.973, Cohen's Kappa of 0.973, and AUPRC values of 0.980. These robust metrics confirm that conventional accuracy results are not misleading and provide comprehensive performance assessment that addresses potential concerns about class imbalance or evaluation bias (Table 7).

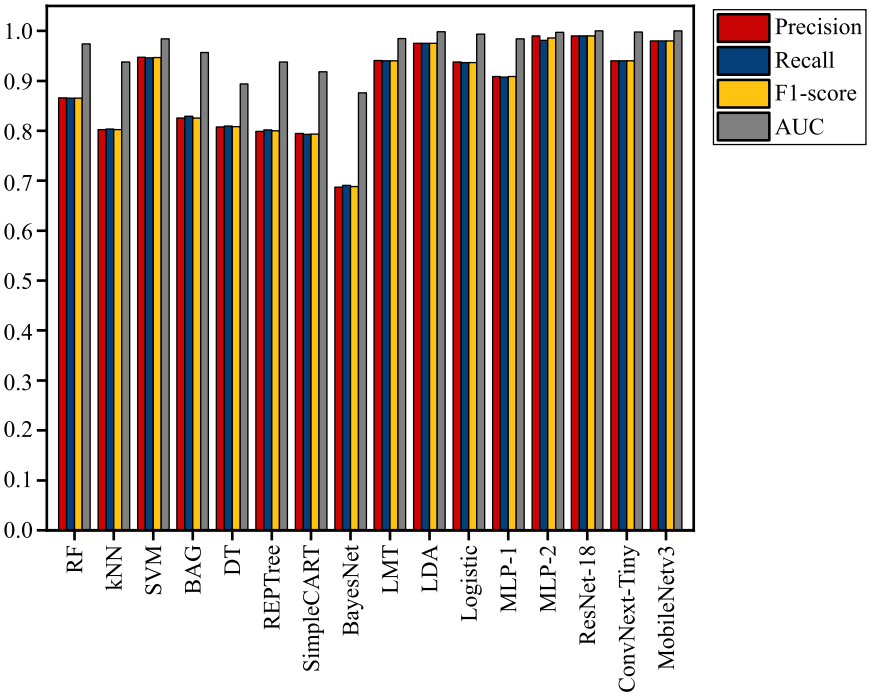

Fig. 10. Performance metrics for the machine and deep learning models, including precision, recall, F1-score and AUC.

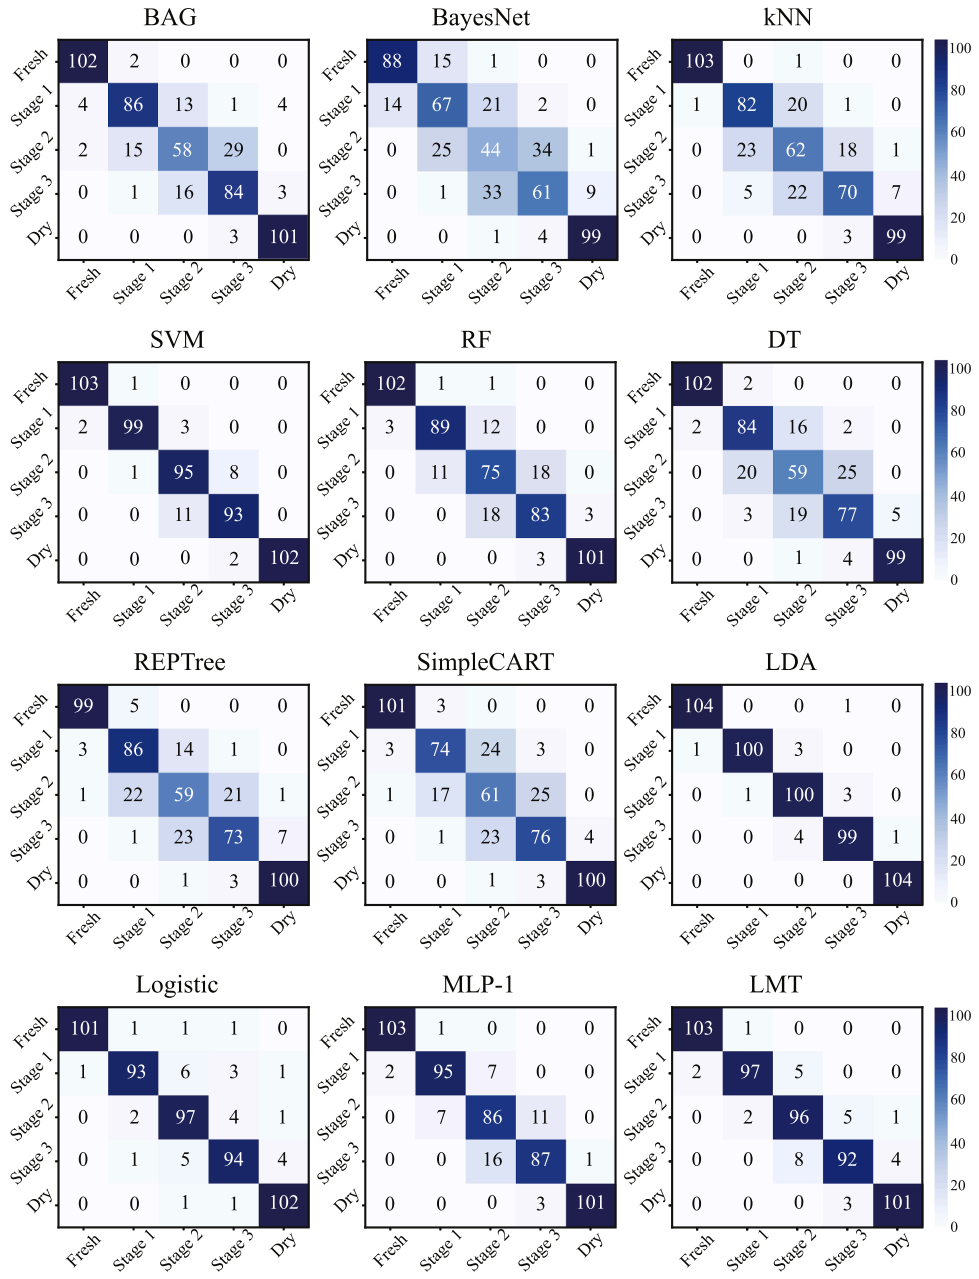

Fig. 11. Confusion matrices for the machine learning models.

Table 6
5-fold cross-validation results for deep learning models.

Model	Accuracy (Mean ± Std)	Precision (Mean ± Std)	Recall (Mean ± Std)	F1-Score (Mean ± Std)
ResNet-18	98.0 ± 1.5%	98.0 ± 1.5%	98.0 ± 1.5%	98.0 ± 1.5%
ConvNeXt-Tiny	100.0 ± 1.5%	100.0 ± 1.5%	100.0 ± 1.5%	100.0 ± 1.5%
MobileNetV3	98.0 ± 1.5%	98.0 ± 1.5%	98.0 ± 1.5%	98.0 ± 1.5%

Table 7
Statistical significance test results.

Model Comparison	McNemar's Test (p-value)	Paired *t*-test (p-value)	Significant
ResNet-18 vs ConvNeXt	0.66	<0.001	Yes (*t*-test)%
ResNet-18 vs MobileNetV3	0.63	0.97	No%
ConvNeXt vs MobileNetV3	0.96	<0.001	Yes (*t*-test)%

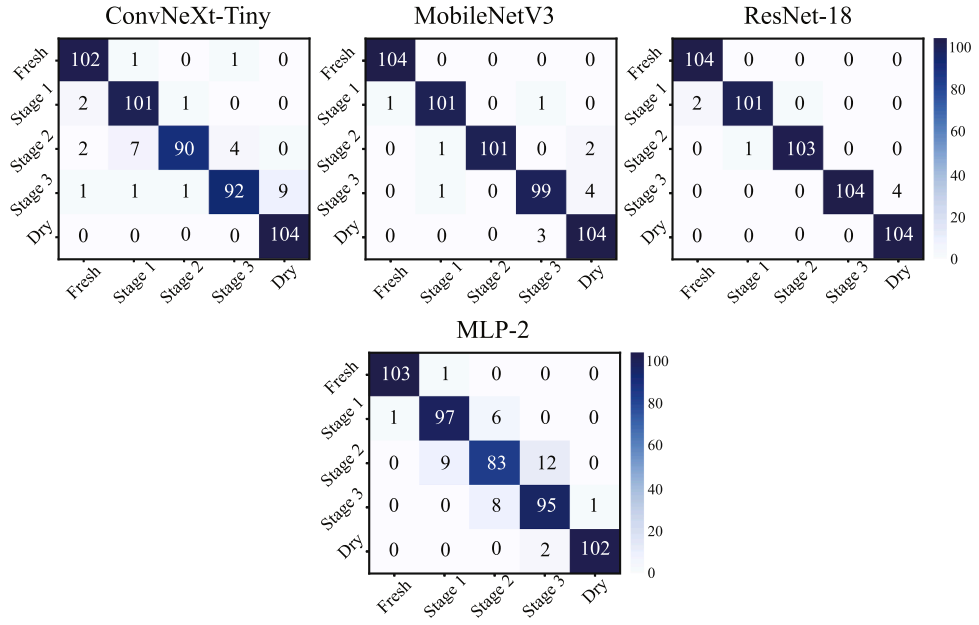

Fig. 12. Confusion matrices for the deep learning models.

Finally, the models are evaluated using ROC curves, which are shown in Fig. 13. The ROC curve is a graphical representation of the model's effectiveness, confirming that the classifier correctly identifies the applications. The most effective curves were obtained for LDA, LMT, and SVM for the ML models and ResNet-18, ConvNeXt-Tiny, and Mobilenetv3-large-100 for the DL models. As seen in the ROC area curves, the best-classified dehydration stage in the ML models was the Fresh Stage and Dry Stage, while in DL models, it was the Fresh Stage. However, the worst classified dehydration stage was Dehydration Stage 2.

4. Discussion

In this research, an in-depth investigation of CNN and classic machine learning models has been carried out. The results of 3 different deep learning models namely ResNet-18, ConvNeXt-Tiny, and Mobilenetv3-large-100 and 13 different machine learning models, including RF, kNN, SVM, BAG, DT, REPTree, SimlpeCART, BayesNet, LMT, LDA, Logistic, MLP-1, and MLP-2 are analyzed with precision, recall, F1-score, AUC, accuracy (Fig. 7), loss-learning rate (Fig. 8), loss-accuracy-error rate (Fig. 7) and confusion matrices (Figs. 11 and 12). In the present study, the selection of ResNet-18, MobileNetV3, and ConvNeXt-Tiny is guided by three primary considerations: computational efficiency, proven performance, and architectural diversity. Given the small dataset size (104 leaves), we prioritized models that could be effectively trained without extensive computational resources while maintaining high performance. All selected models have demonstrated success in similar agricultural classification tasks, providing confidence in their applicability. The three models represent different architectural approaches including residual connections, mobile optimization, and modern CNN design, ensuring comprehensive evaluation. While EfficientNet and Vision Transformers are powerful alternatives, they typically require larger datasets and more computational resources, making them less suitable for our specific constraints and dataset size limitations.

A total of 2600 images were obtained from 104 avocado leaves in 5 different multispectral bands (R, G, B, Red-Edge, and NIR) for each dehydration stage in the DL classification. The five images of the same sample were fused in a single structure to train the deep learning models with the purpose of achieving the best results. In addition, reflectance values of 2151 spectral bands (between 350–2500) were obtained from 104 avocado leaves at each dehydration stage, and 1 118 520 data points were used for classification using the classic machine learning models.

We acknowledge that 104 leaves represent a relatively small dataset for deep learning applications, which is indeed a significant limitation of our study. However, several factors mitigate this limitation and validate our results. Comprehensive validation through 5-fold cross-validation demonstrates model stability across different data splits, with consistent performance showing only $\pm 1.5\%$ standard deviation. Learning curves analysis reveals that training vs validation performance shows no overfitting, indicating models learn generalizable patterns rather than memorizing training data. Statistical significance testing through McNemar's test and paired t-tests confirms that performance differences are statistically meaningful, not due to random variation. Future work should focus on data augmentation, transfer learning, and external dataset validation to address this limitation and further strengthen the generalizability of our approach.

As seen in the mean values of spectral reflectance (Fig. 3), the reflectance values in the SWIR region increase as the dehydration stage progresses. This phenomenon is also observed in works such as Villacrés et al. (2019). Reflectance values around 975 nm, 1200 nm, 1400 nm, and 1900 nm increase as the dehydration stage increases. These bands, in particular, respond to the vibrational bonds of water molecules (Purwadi et al., 2023); thus, as the leaf loses water content, the reflectance values at these wavelengths change (Gutiérrez-Rodríguez et al., 2019). Additionally, stressed (dehydrated) leaves show higher reflectance in the blue and red regions. This can be explained by a reduction in the concentration of photosynthetic pigments due to the loss of leaf moisture content (Hasan Ali et al., 2024; Taşan et al., 2022).

The results produced by the deep learning models: ResNet-18 (Accuracy rate: 99) and Mobilenetv3-large-100 (Accuracy rate: 98) were the best performing models using as input multi-spectral images. In contrast, the LDA model was the best classic machine learning model, using as input the 2151 reflectance values of the spectrum. As stated in (Lu et al., 2019), using hyperspectral information as input, generally produces better results when predicting vegetation chemical traits; however, in cases such as drought stress or in our case severe water stress (completely dehydrated stage), leaves undergo physical changes such as rolled edges (Wang et al., 2024) and a reduction in size (see Fig. 4)., such features cannot be detected using only reflectance values from a single point in the leaf. In such cases, the use of convolutional neural networks allows

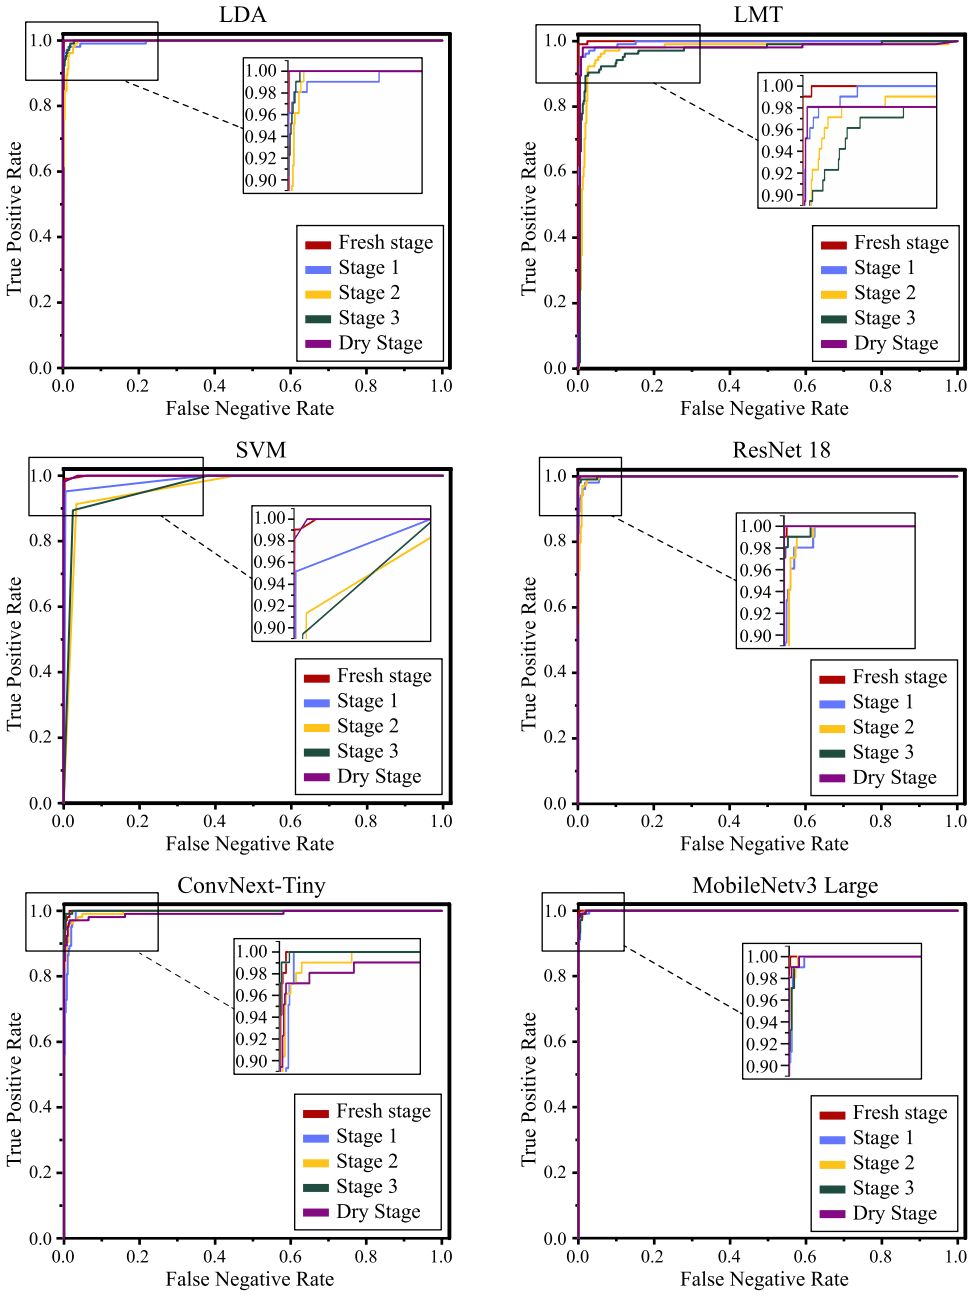

Fig. 13. ROC curves for the most effective machine learning models, LDA, LMT and SVM; and for the three deep learning models: ResNet-18, ConvNext-Tiny and MobileNetv3 Large.

for the detection of physical changes in the leaves, as well as changes in the multi-spectral reflectance values, leading to better classification results.

Although we maintained the all-bands HS baseline for the sake of fairness and reproducibility, we recognize that leakage-free dimensionality reduction and band selection could produce more compact models. Future research will assess PCA and embedded/filter selectors through nested cross-validation, as well as index-driven HS models, to measure potential improvements while maintaining generalizability.

In general, our work produced better results than (Estrada et al., 2023), which achieved accuracies of 0.8 using the Xception model. In this study, the deep learning architectures were different, and the use of strategies such as varying the learning rate and hyperparameter optimization may be contributing factors to the improved performance in a similar task. Other factors that can explain the variation in results include the different plant species and the differences in reflectance values between the species.

A further examination of the models using the confusion matrices shows that fresh and completely dehydrated leaves are the classes with the best results in all of the models, while the dehydration stages from 1 through 3 tend to show the most classification mistakes, particularly in some of the ML models (BayesNet, SimpleCART, REPTree, DT, and kNN). This finding is consistent with the results presented by Estrada et al. (2023), and it can be attributed to the fact that the drying process may be uneven inside the oven, and there is not a clear distinction between these stages of dehydration among the samples (Arevalo-Ramirez et al., 2020).

The ROC curves obtained for the proposed models for the classification of avocado leaves provide more information than simple error rates because the distribution of true positive and false positive rates

can be traced. The obtained distributions were even close to the corner points, indicating that the models performed well. This is further supported by the AUC values, which demonstrate the high performance of the proposed models as reflected in the ROC curves.

Although both the ResNet-18 and MobileNetV3-large-100 models produced the best results, the MobileNetV3 model offers an efficient structure and generally requires less data for fine-tuning. Additionally, it can be run on lower-end hardware setups, such as mobile phones (Howard, 2017). In contrast, the ResNet-18 model has increased robustness by capturing a wider range of features (He et al., 2016; Szegedy et al., 2016). The high performance of our models is consistent with other classification tasks in vegetation, such as the detection of avocado diseases (Mir et al., 2024; Mishra et al., 2022) and the classification of soybean leaf diseases (Wu et al., 2023).

The proposed method can be adopted for separating different leaves according to moisture levels and can provide valuable information about plant health/conditions using multi-spectral images, in contrast to hyperspectral reflectance values, which is the common approach when assessing water content status. However, similar methods require further studies for different plant species and humidity levels. When comparing the study presented by Estrada et al. (2024) with our results, it is important to note that different vegetation species may exhibit distinct reflectance values and may respond differently to severe water stress conditions.

In this work, we successfully demonstrated that combining deep learning with input from a multi-spectral camera (five spectral bands) can achieve better performance than using full spectral information (reflectance across 350-2500 nm) when classifying leaves according to their dehydration stage. This finding is particularly relevant for real-time and in-field applications, where, as noted by Ram et al. (2024), it is often preferable to rely on sensors that capture only specific regions of the spectrum, such as multi-spectral cameras, rather than complex and non-portable devices like spectrometers. An additional advantage of cameras is that they can be easily integrated into robotic platforms, enabling large-scale monitoring of crops and vegetation (Estrada et al., 2023). With our results, we showed that assessing leaf water status is feasible using only multi-spectral imagery; however, to ensure robust performance under the more variable conditions encountered in the field, recalibration or retraining of the deep learning models would be necessary.

5. Conclusions

In this study, deep learning and machine learning models were developed for the dehydration stage classification of avocado leaves using both multispectral images and spectral reflectance, and the effectiveness of the models was investigated. First, images from five different bands were combined for multispectral-assisted DL classification, and a single image was obtained. For spectral reflectance-assisted ML classification, spectral reflectance values obtained from 2151 bands were used as the raw data. Subsequently, a set of performance criteria was visualized to determine the proposed models, which were analyzed in detail. In summary, the ResNet-18 and Mobilenetv3-large-100 deep learning models achieved 0.99 and 0.98 accuracy rates, respectively, while the LDA classical model achieved 0.98 accuracy rates, supported by other performance metrics, confusion matrices, and ROC curves. However, the BayesNet model performed worse than most other models.

In addition, the lowest spectral reflectance values were found in the Dry Stages and were particularly pronounced in the SWIR region. These results are promising for both image- and reflectance-based separation of leaves at different dehydration levels. The important contribution of this research is that it helps detect the moisture levels of plants at an early stage so that necessary precautions can be taken and crop losses can be prevented.

In this study we successfully predicted the dehydration stage of avocado leaves, using a multi-spectral camera, a cheaper and more portable sensor than a hand-held spectrometer. We were able to achieve superior results when compared to approaches using the full spectrum obtained from a spectrometer, thus demonstrating the capabilities of multispectral cameras for the assessment of water content in leaves, and their potential in future irrigation tasks.

Our work is a first step towards an affordable and scalable methodology for irrigation monitoring and future work includes the testing of the deep learning models in field conditions and using images of larger portions of trees from the canopy, instead of individual leaves, from avocados and from different crops. For this purpose, the use of robotic platforms such as UAV's are considered. Furthermore, the exploration of a larger dataset involving different crops and with a finer differentiation among the dehydration stages, as well as using more robust models, can improve the results obtained in this research and expand the scope of the present work.

CRediT authorship contribution statement

J.S. Estrada: Conceptualization, Software, Data collection, Data curation, Validation, Data analysis, Writing; **N. Cetin:** Software, Data validation, Data analysis; **K. Sacilik:** Software, Validation, Writing, Data analysis; **Banu Ulu:** Software, Data analysis, Writing; **Burak Ulu:** Software, Data analysis, Writing; **F. Auat Cheein:** Conceptualization, Data analysis, Validation, Writing, Editing, Reviewing.

Data availability

Data will be made available on request.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

Ahad, M. T., Li, Y., Song, B., & Bhuiyan, T. (2023). Comparison of CNN-based deep learning architectures for rice diseases classification. *Artificial Intelligence in Agriculture*, *9*, 22–35.

Alemu, T. T. (2024). Nutritional contribution of fruit and vegetable for human health: A review. *International Journal Health Policy Plann*, *3*(1), 1–9.

Alordzinu, K. E., Li, J., Lan, Y., Appiah, S. A., Alaa, A. L. A., Wang, H., Liao, J., Sam-Amoah, L. K., & Qiao, S. (2021). Ground-based hyperspectral remote sensing for estimating water stress in tomato growth in sandy loam and silty loam soils. *Sensors*, *21*(17), 5705. https://doi.org/10.3390/s21175705

Arevalo-Ramirez, T., Villacrés, J., Fuentes, A., Reszka, P., & Auat Cheein, F. A. (2020). Moisture content estimation of pinus radiata and eucalyptus globulus from reconstructed leaf reflectance in the SWIR region. *Biosystems Engineering*, *193*, 187-205. https://doi.org/10.1016/j.biosystemseng.2020.03.004.

Dou, S., Zhang, W., Deng, Y., Zhang, C., Mei, Z., Yan, J., & Li, M. (2024). Comparison of citrus leaf water content estimations based on the continuous wavelet transform and fractional derivative methods. *Horticulturae*, *10*(2), 177. https://doi.org/10.3390/horticulturae10020177.

Duran-Llacer, I., Munizaga, J., Arumí, J., Ruybal, C., Aguayo, M., Sáez-Carrillo, K., Arriagada, L., & Rojas, O. (2020). Lessons to be learned: Groundwater depletion in chile's ligua and petorca watersheds through an interdisciplinary approach. *Water*, *12*(9), 2446. https://doi.org/10.3390/w12092446.

Estrada, J. S., Fuentes, A., Reszka, P., & Cheein, F. A. (2023). Machine learning assisted remote forestry health assessment: A comprehensive state of the art review. *Frontiers in Plant Science*, *14*. https://doi.org/10.3389/fpls.2023.1139232.

Estrada, J. S., Zanartu, M., Demarco, R., Fuentes, A., & Cheein, F. A. (2024). Water content classification on leaves based on multi-spectral imagery and machine learning techniques for wildfire prevention. In *2024 IEEE international conference on industrial technology (ICIT)* (pp. 1–7). https://doi.org/10.1109/ICIT58233.2024.10540771

Food and Agriculture Organization of the United Nations (2024). World Food and Agriculture - Statistical Yearbook 2024 FAO, addressRome, Italy, https://openknowledge.fao.org/handle/20.500.14283/cd2971en, Accessed: May 7 2026.

Gonçalves, D., Gouveia, C. S. S., Ferreira, M. J., Ganança, J. F. T., Pinto, D. C. G., & de Carvalho, M. A. A. P. (2024). Comparative analysis of antioxidant and fatty acid composition in avocado (persea americana mill.) fruits: Exploring regional and commercial varieties. *Food Chemistry*, *442*, 138403.

Government of Chile, Ministry of Environment (2016). Chile's second biennial update report on climate change. Ministry of Environment, Chile, Santiago, Chile. Biennial Update Report (BUR) to the UNFCCC, Accessed: May 7 2026.

https://unfccc.int/files/national_reports/non-annex_i_parties/biennial_update_reports/application/pdf/bur2_chile_english2017.pdf
Guiné, R., De Guiné E, F. S., Teixeira-Lemos, E., Lima, M. J., & C. Goncalves, J. (2024). Consumption of fruits and vegetables in two european countries: Results from a survey in france and portuga. *AIMS Agriculture and Food*, *9*(3), 767–788.
Gutiérrez-Rodríguez, B. J., Argüello-Tovar, J. O., & García-Navarrete, O. L. (2019). Use of VIS-NIR-SWIR spectroscopy for the prediction of water status in soybean plants in the colombian piedmont plains. *Dyna*, *86*(210), 125–130.
Hasan Ali, İ., Ropelewska, E., & Çetin, N. (2024). Using spectral vegetation indices and machine learning models for predicting the yield of sugar beet (beta vulgaris l.) under different irrigation treatments. *Computers and Electronics in Agriculture*, *221*, 109019.
He, K., Zhang, X., Ren, S., & Sun, J. (2015). Deep residual learning for image recognition. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, (pp. 770–778). https://api.semanticscholar.org/CorpusID:206594692.
He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 770–778).
Howard, A. G. (2017). MobileNets: Efficient convolutional neural networks for mobile vision applications. arxiv preprint arXiv:1704.04861.
Jin, J., Wang, Q., & Song, G. (2023). Exploring low order fractional derivative spectra indices for estimating leaf fuel moisture content across a variety of plant species. *International Journal of Remote Sensing*, *44*(7), 2342-2358. http://dx.doi.org/10.1080/01431161.2023.2201384.
Lawson, B. D., & Hawkes, B. (1989). Field evaluation of moisture content model for medium-sized logging slash. In *Proceedings of the 10th conference on fire and forest meteotology* (pp. 247–257).
Lazzaris, S. (2024). AVOCADO: The cost of production. https://www.foodunfolded.com/article/avocado-the-cost-of-production.
Lieu, M. D., Phuong, T. V., Nguyen, T. T. B., Dang, T. K. T., & Nguyen, T. H. (2024). A review of preservation approaches for extending avocado fruit shelf-life. *Journal of Agriculture and Food Research*, *16*, 101102. https://doi.org/10.1016/j.jafr.2024.101102
Liu, N., Wu, L., Chen, L., Sun, H., Dong, Q., & Wu, J. (2018). Spectral characteristics analysis and water content detection of potato plants leaves. *IFAC-PapersOnLine*, *51*(17), 541-546. http://dx.doi.org/10.1016/j.ifacol.2018.08.152.
Liu, Z., Mao, H., Wu, C.-Y., Feichtenhofer, C., Darrell, T., & Xie, S. (2022). A convnet for the 2020s. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 11966–11976. https://doi.org/10.1109/CVPR52688.2022.01167
Lu, B., He, Y., & Dao, P. D. (2019). Comparing the performance of multispectral and hyperspectral images for estimating vegetation properties. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, *12*(6), 1784–1797. https://doi.org/10.1109/JSTARS.2019.2910558
Madariaga, A., Maillet, A., & Rozas, J. (2021). Multilevel business power in environmental politics: the avocado boom and water scarcity in chile. *Environmental Politics*, *30*(7), 1174-1195. https://doi.org/10.1080/09644016.2021.1892981.
Menge, J. A., Douhan, G. W., McKee, B., Pond, E., Bender, G. S., & Faber, B. (2012). Three new avocado rootstock cultivars tolerant to phytophthora root rot:'zentmyer','uzi', and 'steddom'. *HortScience*, *47*(8), 1191–1194.
Mir, T. A., Gupta, S., Chauhan, R., Singh, M., Banerjee, D., & Kumar, B. V. (2024). Enhanced multiclassification of avocado leaf diseases: CNN and random forest integration. In *2024 3rd international conference for innovation in technology (INOCON)* (pp. 1–6). IEEE.
Mishra, S., Ayane, T. H., Ellappan, V., Rathee, D. S., & Kalla, H. (2022). Avocado fruit disease detection and classification using modified SCA–PSO algorithm-based mobilenetv2 convolutional neural network. *Iran Journal of Computer Science*, *5*(4), 345–358.
Network, W. F. (2024). Water footprint assessment tool. https://www.waterfootprintassessmenttool.org/.
Otsu, N. (1979). A threshold selection method from gray-level histograms. *IEEE Transactions on Systems, Man, and Cybernetics*, *9*(1), 62–66. https://doi.org/10.1109/TSMC.1979.4310076
Purwadi, I., Erskine, P. D., & van der Ent, A. (2023). Reflectance spectroscopy as a promising tool for 'sensing' metals in hyperaccumulator plants. *Planta*, *258*(2). https://doi.org/10.1007/s00425-023-04167-3.
Qian, S., Ning, C., & Hu, Y. (2021). Mobilenetv3 for image classification. In *2021 IEEE 2nd international conference on big data, artificial intelligence and internet of things engineering (ICBAIE)* (pp. 490–497). https://doi.org/10.1109/ICBAIE52039.2021.9389905
Ram, B. G., Oduor, P., Igathinathane, C., Howatt, K., & Sun, X. (2024). A systematic review of hyperspectral imaging in precision agriculture: Analysis of its current state and future prospects. *Computers and Electronics in Agriculture*, *222*, 109037. http://dx.doi.org/10.1016/j.compag.2024.109037.
Ramirez, O. J. V., de la Cruz, J. E. C., & Machaca, W. A. M. (2021). Agroindustrial plant for the classification of hass avocados in real-time with resnet-18 architecture. In *2021 5th international conference on robotics and automation sciences (ICRAS)* (pp. 206–210). IEEE.
Ramos-Aguilar, A. L., Ornelas-Paz, J., Tapia-Vargas, L. M., Ruiz-Cruz, S., Gardea-Béjar, A. A., Yahia, E. M., de Jesús Ornelas-Paz, J., Pérez-Martínez, J. D., Rios-Velasco, C., & Ibarra-Junquera, V. (2019). The importance of the bioactive compounds of avocado fruit (persea americana mill) on human health. *Biotecnia*, *21*(3), 154–162.
Raouhi, E. M., Lachgar, M., Hrimech, H., Kartit, A. et al. (2022). Optimization techniques in deep convolutional neuronal networks applied to olive diseases classification. *Artificial Intelligence in Agriculture*, *6*, 77–89.
Ropelewska, E., Sabanci, K., Aslan, M. F., & Çetin, N. (2023). Rapid detection of changes in image textures of carrots caused by freeze-drying using image processing techniques and machine learning algorithms. *Sustainability*, *15*(8), 7011.
Santos, N. C., Almeida, R. L. J., Monteiro, S. S., de Alcântara Silva, V. M., de Lima, T. L. B., Saraiva, M. M. T., Santos, R. M. S., Martins, A. N. A., Paiva, Y. F., de Lima, J. S. et al. (2024). Ultrasound and microwaves reduce stress in probiotics during avocado drying: Impact on mass transfer and cell viability. *Food Bioscience*, *61*, 104655.
Schiappacasse, I., Segura, P., & Rozas, J. (2024). Social conflicts over the use of water resources in chile: The role of social movements and business power. *Oxford Development Studies*, *52*(4), 381-395. https://doi.org/10.1080/13600818.2024.2365721.
Sommaruga, R., & Eldridge, H. M. (2020). Avocado production: Water footprint and socio-economic implications. *EuroChoices*, *20*(2), 48-53. https://doi.org/10.1111/1746-692x.12289.
Surase, R. R., Kale, K. V., Varpe, A. B., Vibhute, A. D., Gite, H. R., Solankar, M. M., Gaikwad, S., & Nalawade, D. B. (2019). Estimation of water contents from vegetation using hyperspectral indices. In G. Panda, S. C. Satapathy, B. Biswal, & R. Bansal (Eds.), *Microelectronics, electromagnetics and telecommunications* (pp. 247–255). Singapore: Springer Singapore.
Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., & Wojna, Z. (2016). Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 2818–2826).
Taşan, S., Cemek, B., Taşan, M., & Cantürk, A. (2022). Estimation of eggplant yield with machine learning methods using spectral vegetation indices. *Computers and electronics in agriculture*, *202*, 107367.
Tassis, L. M., & Krohling, R. A. (2022). Few-shot learning for biotic stress classification of coffee leaves. *Artificial Intelligence in Agriculture*, *6*, 55–67.
UNCCD (2016). Update on the risk of desertification, land degradation and drought in chile. United Nations Convention to Combat Desertification, https://www.unccd.int/sites/default/files/inline-files/Chile_1.pdf, Accessed: May 7 2026
Villacrés, J., Arevalo-Ramirez, T., Fuentes, A., Reszka, P., & Cheein, F. A. (2019). Foliar moisture content from the spectral signature for wildfire risk assessments in valparaíso-chile. *Sensors*, *19*(24), 5475. https://doi.org/10.3390/s19245475. https://doi.org/10.3390/s19245475
Wang, Y., Jing, X., Gao, Y., Han, X., Zhao, C., & Pan, W. (2024). Leaf rolling detection in maize under complex environments using an improved deep learning method. *Plant Molecular Biology*, *114*(5). https://doi.org/10.1007/s11103-024-01491-4.
Wei, Y., Wu, F., Xu, J., Sha, J., Zhao, Z., He, Y., & Li, X. (2019). Visual detection of the moisture content of tea leaves with hyperspectral imaging technology. *Journal of Food Engineering*, *248*, 89-96. https://doi.org/10.1016/j.jfoodeng.2019.01.004.
Wu, Q., Ma, X., Liu, H., Bi, C., Yu, H., Liang, M., Zhang, J., Li, Q., Tang, Y., & Ye, G. (2023). A classification method for soybean leaf diseases based on an improved convnext model. *Scientific Reports*, *13*(1), 19141.