Computers and Electronics in Agriculture 251 (2026) 111971

Contents lists available at ScienceDirect

Computers and Electronics in Agriculture

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

Original papers

Variational autoencoder enables unsupervised leaf diagnosis via hyperspectral imaging

Kangyu Ji [a,*], Zhiyuan Chen [a], Zhongzhong Niu [a], Alden Mo [a], Zhihao Qin [a], Katie DeCamp [b], Charles Wang [a], Weixuan Xu [a], Julie M. Young [c], Bryan G. Young [c], Jingqiu Chen [d], Anna O. Conrad [b], Jian Jin [a,*]

[a] *Department of Agricultural and Biological Engineering, Purdue University, West Lafayette 47907, IN, USA*
[b] *USDA Forest Service, Northern Research Station, Delaware 43015, OH, USA*
[c] *Department of Botany and Plant Pathology, Purdue University, West Lafayette 47907, IN, USA*
[d] *Biological Systems Engineering, College of Agriculture and Food Sciences, Florida A&M University, Tallahassee 32307, FL, USA*

ARTICLE INFO

Keywords:
Hyperspectral imaging
Variational autoencoder
Unsupervised learning
Plant phenotyping
Nitrogen deficiency
Plant disease diagnosis
Explainable artificial intelligence

ABSTRACT

Hyperspectral imaging provides rich spectral information for plant phenotyping, yet analyzing these high-dimensional data remains challenging due to limited labeled datasets and the reliance on manually extracted features, such as vegetation indices. Here, we present LeafVAE, a variational autoencoder framework that learns spectral signatures directly from hyperspectral leaf images without manual annotation. Our model encodes leaf spectra into a compact, two-dimensional latent space where pixels cluster by spectral similarity, creating quantitative leaf descriptors from signature composition that facilitate downstream phenotyping tasks. We show that the learned representations generalize robustly across genotypes, years, and imaging equipment on four different corn leaf datasets. We further demonstrate that LeafVAE outperforms traditional vegetation indices across multiple applications, including detecting nitrogen deficiency in corn, classifying herbicide stress in soybean, and identifying beech leaf disease using an affordable six-band multispectral camera. Pixel-level explainability, achieved by combining spectral signature composition with interpretable machine learning methods, identifies regions with abnormal spectral characteristics and highlights areas of the leaf that contribute to model predictions. By learning spectral features in an unsupervised manner while still allowing labels for interpretable diagnosis, LeafVAE provides a scalable foundation for automated plant phenotyping across diverse applications.

1. Introduction

Early detection of plant health issues, including nutrient deficiencies (Barbedo, 2019), herbicide stress (Henry et al., 2004), and diseases (Lee et al., 2023; Li et al., 2019), is essential for guiding field management decisions that increase productivity, optimize resource use, and support ecosystem sustainability (Kniss, 2017; Nosengo, 2003; Xu et al., 2024). Unmanned aerial vehicle (UAV)-based hyperspectral imaging has enabled large-scale, airborne field monitoring and high-throughput phenotyping. However, accurate plant diagnosis from airborne observations remains challenging due to unstable lighting conditions (e.g., varying solar angles and cloud cover), variation in canopy overlap and leaf orientation, and atmospheric interference (e.g., wind), all of which reduce data consistence and quality (Gano et al., 2024; Sun et al., 2021).

Handheld hyperspectral imaging provides a complementary approach by enabling close-range, high-resolution measurements (Rehman et al., 2020; Sarić et al., 2022). These systems are typically operated under more controlled acquisition conditions (e.g., fixed illumination and enclosure), allowing for consistent capture of spatially resolved reflectance/transmittance spectra. As a result, they enable the extraction of spatial-spectral features that reflect biochemical and structural variations at the single-leaf level (Zhang et al., 2026). Although limited in throughput compared to UAV-based platforms, handheld systems provide a cost-effective means of acquiring high-fidelity proximal spectral data, which supports the interpretation and validation of large-scale airborne observations.

Hyperspectral imaging generates high-dimensional data, acquiring hundreds of spectral bands per spatial pixel. Plant biochemical

* Corresponding authors.
E-mail addresses: axvcb1597382@gmail.com (K. Ji), jinjian@purdue.edu (J. Jin).

https://doi.org/10.1016/j.compag.2026.111971
Received 11 February 2026; Received in revised form 27 April 2026; Accepted 27 May 2026

compounds influence these spectra through both direct effects on light absorption by leaf pigments (e.g., chlorophyll, carotenoids, and anthocyanins) and indirect effects of leaf structural features that affect light scattering (Ollinger, 2011). This rich information introduces analytical challenges. Feature engineering (Rasti et al., 2020), the processes of identifying the most informative features for downstream modeling, become non-trivial (Gano et al., 2024). The redundancy and high intercorrelation among spectral bands can inflate computational costs and increase the risk of overfitting when sample sizes are limited (i.e., the curse of dimensionality (Rasti et al., 2020)), where the data collection is often constrained in plant studies by limited sample throughput, environmental variability, and biological heterogeneity. Therefore, models trained on a single dataset or species often generalize poorly to others due to biological variability (e.g., genotype and developmental stage of plants) and technical variability (e.g., imaging system configuration, sensor calibration, and acquisition conditions). This limits the practical deployment of supervised models across diverse genotype, environment, and management (G × E × M) conditions in plant science.

At the same time, there is also a trade-off between model interpretability and predictive performance. Early approaches relied on simple, physically interpretable representations, such as vegetation indices and linear models, which provide clear relationships between spectral features and plant physiological status (Bannari et al., 1995; Xue and Su, 2017). For example, the Normalized Difference Vegetation Index (NDVI) has been widely adopted due to its simplicity and direct interpretability (Rouse et al., 1974). On the other hand, with the increasing availability of high-resolution hyperspectral images, deep neural networks have been introduced to improve predictive accuracy in phenotyping (Gill et al., 2022). However, these models often operate as black boxes, providing limited insight into which spectral-spatial features drive predictions.

To address this limitation, recent research has focused on developing methods that improve model interpretability without sacrificing too much prediction performance for plant phenotyping (Danilevicz et al., 2025). These efforts include both inherently interpretable models, such as decision trees and attention-based architectures (Amirruddin et al., 2020; Aviles Toledo et al., 2024), and post hoc explainable artificial intelligence (XAI) techniques (Lapajne et al., 2025). Saliency-based methods, such as Gradient-weighted Class Activation Mapping (Grad-CAM) (Selvaraju et al., 2020), have been used to highlight important spatial regions contributing to predictions, but they do not explicitly quantify the contribution of individual spectral variables. Feature attribution methods such as Local Interpretable Model-agnostic Explanations (LIME) and SHapley Additive exPlanations (SHAP) (Lundberg and Lee, 2017; Ribeiro et al., 2016) decompose model predictions into contributions from individual inputs, enabling quantitative assessment of relationships between calculated features and crop health (Lapajne et al., 2025). These applications pave the way for approaches that fully leverage hyperspectral data while providing quantitatively interpretable spectral-spatial explanations.

In this work, we propose LeafVAE, a variational autoencoder (VAE) framework for hyperspectral leaf diagnosis (Fig. 1). The model learns spectral variations (i.e., subtle color differences) of various types of leaves by encoding each leaf spectrum, obtained from a spatial pixel in a hyperspectral image, into a compact two-dimensional (2D) latent embedding, from which the decoder reconstructs the original spectrum (Fig. 1**a**). This unsupervised learning approach does not require labeled data, as it learns intrinsic spectral patterns directly from the hyperspectral dataset. Thus, all pixels in the dataset can be encoded and projected into a common latent space, where they cluster according to spectral similarity (Fig. 1**b**). Each leaf can then be represented as a distribution of cluster compositions, the proportions of different spectral signatures (i.e., colors) within the leaf, which provide a simple and interpretable descriptor of leaf condition. These descriptors, combined with simple machine learning models (such as random forests) and explainable SHAP methods, can be linked to plant health status, enabling interpretable leaf diagnosis for downstream high-throughput phenotyping.

2. Methods

2.1. Dataset collection

2.1.1. Corn leaf dataset

Hyperspectral transmittance images of corn leaves were acquired using a custom-built handheld scanning device, LeafSpec, developed in

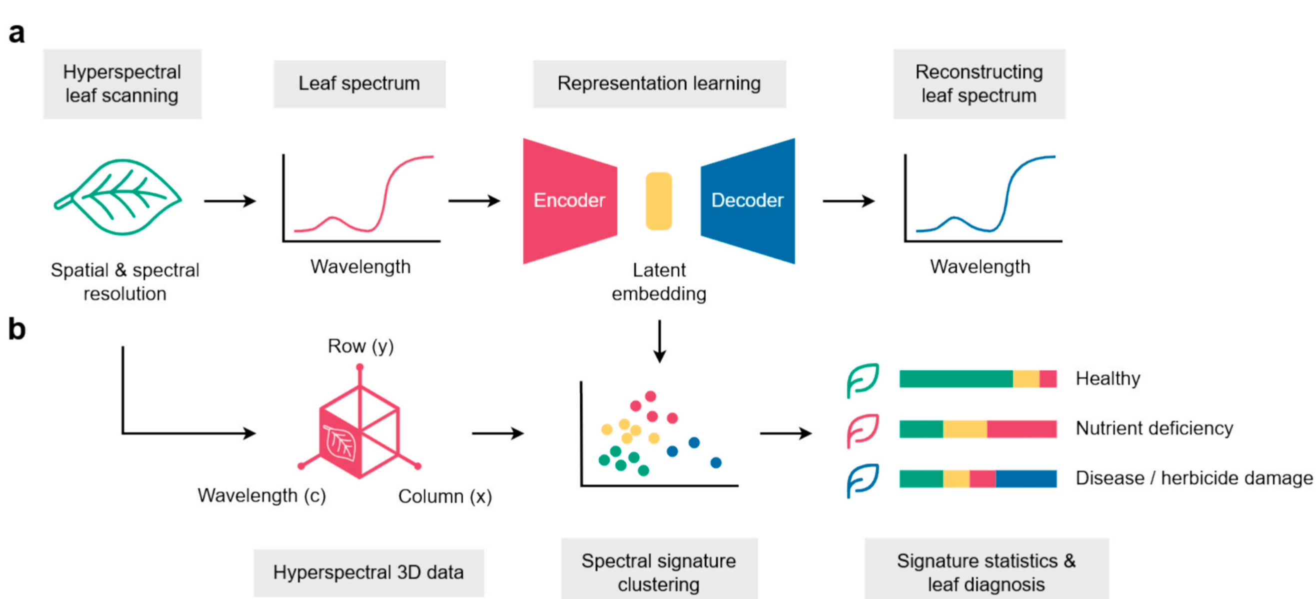

Fig. 1. Overview of LeafVAE for leaf diagnosis. a, Training LeafVAE using leaf spectra from hyperspectral leaf scanning data, where the encoder compresses spectra into low-dimensional latent embeddings and the decoder reconstructs the original leaf spectrum. The model learns latent representations of the leaf spectra through the reconstruction process. **b**, Analysis pipeline for leaf diagnosis. Hyperspectral data cubes are processed through the model by encoding wavelength information into latent space and assigning spectral signature clusters, enabling automated diagnosis of healthy leaves, nutrient deficiency, herbicide damage, and disease.

previous studies (Wang et al., 2020). The device employs a monochrome charge-coupled device (CCD) camera (BFLY-U3-05S2M-CS, Edmund Optics Inc., 66 dB dynamic range, 50 fps) for line-scan imaging, with one dimension representing the spatial x-axis (0.5 mm per pixel) and the other representing the spectral axis (0.59 nm per pixel), separated using a diffraction grating. A rotary encoder records the displacement of the leaf along the midrib direction (spatial y-axis) during scanning. The device uses an active halogen light source with a Teflon-coated window to diffuse illumination within sealed containers, minimizing the influence of external light. Spectral calibration was performed using a Krypton light source (KR1, Ocean Optics Inc., Largo, FL, USA). Before scanning, a dark image (light off) and a reference image (light on, no sample) were captured to calibrate intensity to percentage transmission.

Four corn leaf hyperspectral datasets were collected between 2022 and 2024. The first two were acquired one month apart (May and June 2022) under the same experimental conditions. Two corn genotypes (Beck's hybrid 5994V2P and Pioneer hybrid P1105AM) were grown under two nitrogen fertilizer treatments (224.2 kg/ha and 22.4 kg/ha) in the field north of the Indiana Corn and Soybean Innovation Center at Purdue University, USA. A total of 16 plots (2 genotypes × 2 nitrogen treatments × 4 duplicates), each containing 30 plants, were planted, and the experimental design and plot configuration for these trials are detailed in (Song et al., 2024). At the R1 growth stage, corn leaves were freshly cut from randomly selected plants between 9:00–11:30 a.m., then line-scanned across a wavelength range of 429 to 946 nm with a spectral interval of 1.4 nm.

The third dataset, consisting of one corn genotype (Pioneer hybrid P1108Q), was collected in a corn field of a local farmer in White County, IN, USA, in June 2023. Six graded nitrogen levels, corresponding to 0%, 40%, 60%, 80%, 100%, and 120% of the farmer's standard nitrogen rate (200 kg/ha), were used. A total of 24 plots (6 nitrogen treatments × 4 duplicates), each containing 150 plants, were planted. The plot configuration is detailed in Zhang et al., 2025. The leaf samples were cut at growth stage V8 and scanned between 8:00 a.m. and 12:00p.m. local time.

The fourth hyperspectral dataset was obtained from Beck's hybrid 5994V2P corn, grown in the same experimental field used for the first two collections in 2022. This trial included four nitrogen application rates (20, 40, 100, and 200 lbs N/acre). Sixteen plots were established (four nitrogen levels with four replicates), and each plot contained 14 plants. Leaves were sampled at the V10 growth stage, with fresh leaf segments taken from randomly selected plants during the morning between 8:00 a.m. and 12:00p.m.

The ground truth nitrogen content for each corn plant was chemically measured using the Dumas method (Nelson and Sommers, 1980). All available leaves of each plant were cut off at the collar and stored in a paper bag, then oven-dried for a week before being sealed and shipped to A&L Great Lakes Laboratories, Inc. (Fort Wayne, IN, USA) for total nitrogen analysis.

2.1.2. Soybean leaf herbicide dataset

Hyperspectral transmittance images of soybean leaves were acquired with a modified version of LeafSpec, LeafSpec-Dicot (Li et al., 2023), with a wavelength window from 443 to 937 nm. In this modified setup, leaves were placed flat inside the device light enclosure, and an internal motor drives the scanner to perform line-scan acquisition.

Soybean plants (Becks 296 L4, Liberty Link) were grown at the Lilly Greenhouses and Plant Growth Facility, Purdue University, each in a plastic pot containing a 2:1 mixture of potting soil and sand. The potting medium included peat moss, perlite, vermiculite, and limestone components. Fertilizer was applied to maintain healthy growth, and the insecticide Kontos (OHP, Inc., Mainland, PA, USA) was applied at the unifoliate stage to prevent insect damage. Two herbicides, Dicamba (XtendiMax, 2.9 lb ae/gal, Bayer AG) and 2,4-D choline (Enlist One, 3.8 lb ae/gal, Corteva), were applied at the V2 growth stage. The experiment followed a randomized block design with nine treatments, an untreated control and eight herbicide treatments (two herbicides at four application rates each), with 20 replicate plants per treatment. This design helps minimize variation in light and temperature across treatment groups within the greenhouse. Dicamba was applied at 0.56, 0.28, 0.14, and 0.0695 g/ha, corresponding to dilution ratios of 1:1,000, 1:2,000, 1:4,000, and 1:8,000, respectively. Similarly, 2,4-D choline was applied at 42.6, 21.3, 14.3, and 10.6 g/ha, corresponding to dilution ratios of 1:25, 1:50, 1:75, and 1:100.

2.1.3. Muscadine grape leaf disease dataset

Hyperspectral reflectance images of grape leaves were acquired using a commercial desktop hyperspectral imaging system (Via-Spec II Scanner Stage, Middleton Spectral Vision, WI, USA) equipped with a visible-near infrared spectral camera (MSV 101 VNIR, Middleton Spectral Vision, WI, USA) and a motorized stage with a scan speed of 10.16 mm/s. A dark image (light off) and a reference image (light on, no sample) were captured before scanning for spectral calibration.

Two perennial muscadine grape cultivars, Noble and Floriana, were sampled in May 2025 at the Florida A&M University Center for Viticulture and Small Fruit Research. The sampling design included both non-fungicide and fungicide treatment areas within the experimental vineyard. The non-fungicide treatment consisted of four Noble vines and five Floriana vines that received no fungicide application. These vines were selected in 2024 and marked with orange paint. The fungicide treatment consisted of three vines per cultivar located in the vineyard area managed under the vineyard fungicide application program.

For each selected vine, 10 symptomatic leaves and 5 asymptomatic leaves were collected by trained personnel at the FAMU research vineyard. Symptomatic leaves were selected based on visible foliar symptoms consistent with angular leaf spot or black rot, using the muscadine disease symptom references provided by the University of Georgia viticulture disease guide as the field reference (Brannen, 2019). The University of Georgia disease guide identifies black rot and angular leaf spot as primary southeastern muscadine diseases requiring fungicide management and describes angular leaf spot as a disease associated with Mycosphaerella angulata or Cercospora brachypus and black rot as a disease associated with Guignardia bidwellii f. muscadinii. Leaves were classified as symptomatic when they showed visible disease like lesions, including angular or irregular necrotic spots, brown to dark lesions, chlorotic areas surrounding lesions, or black rot like necrotic spots. Asymptomatic leaves were selected from the same sampled vines based on the absence of visible disease lesions at the time of collection. Because pathogen isolation or molecular confirmation was not performed, the disease labels in this dataset should be interpreted as symptom-based field labels rather than pathogen confirmed diagnoses.

In total, 15 vines were sampled, including 7 Noble vines and 8 Floriana vines. The final dataset included 225 leaves, consisting of 150 symptomatic leaves and 75 asymptomatic leaves. After collection, all leaves were placed in dry envelopes and shipped for hyperspectral imaging scanning. The fungicide treatment area received Captan and Myclobutanil on October 11, 2024, and Azoxystrobin and Captan on April 10, 2025, before the May 2025 sampling. The full seasonal fungicide management record also included Myclobutanil and Mancozeb on June 16, 2025, and Captan on September 25, 2025. Fungicide active ingredients, brand names, EPA registration numbers, and application rates are summarized in **Supplementary** Table 1.

2.1.4. Beech leaf disease dataset

Multispectral transmittance images of beech leaves were acquired using an improved version of a previously developed scanning device, the Leaf Scanner (Zhang et al., 2019). The system scans an area of 12 × 9 cm in a line-scan configuration (width to length direction). Images were captured using an Arducam 8MP Sony IMX219 camera without an optical filter, across six transmittance bands (660, 560, 473, 595, 400, and 880 nm), with illumination diffused by a 1/8-inch Teflon light diffuser panel (McMaster-Carr Supply Company, IL, USA). Reference images (light on, no sample) were captured prior to scanning for spectral calibration.

American beech leaves were scanned at Johnson Woods State Nature Preserve in Marshallville, Ohio, where research on beech leaf disease is ongoing. Thirty symptomatic beech trees were surveyed in September 2025, and metadata were recorded for each sampled tree, including diameter at breast height (in cm), percent canopy cover, and percent healthy foliage. From each tree, two asymptomatic and two symptomatic leaves were collected and scanned.

The statistical details of all datasets are summarized in Table 1.

2.2. Data preprocessing

A reference image contains the lamp signal across image width (perpendicular to the line-scan direction) and wavelength, without any sample in place. The lamp signal was smoothed along the wavelength axis using a Savitzky-Golay filter (window length = 7, polynomial order = 3) to reduce noise. Leaf images were then calculated from the raw hyperspectral cube using the smoothed lamp signal and the dark background.

To reduce the influence of leaf thickness and overall intensity variation, each hyperspectral spectrum was normalized by its mean value in the 850–900 nm range. This near-infrared region is weakly affected by pigment absorption and primarily reflects leaf internal structure and scattering (Ollinger, 2011), making it a stable reference for intensity normalization. This helps the LeafVAE model focus on learning spectral shape representations rather than overall intensity variations. The full leaf area was segmented by thresholding selected wavelengths, and a binary leaf mask was generated to enable visual inspection and manual refinement of the segmentation results. The primary and secondary vein regions were segmented using an established method based on the local binary pattern algorithm (Song et al., 2024). After segmentation, the dataset was assembled into a dataloader, where masked regions were assigned NaN values so only the leaf area remained valid, with a cropped spectral range of 500–900 nm.

2.3. LeafVAE

To improve interpretability and generalizability across hyperspectral imaging data described in the Introduction, we design an interpretable learning pipeline that combines unsupervised representation learning (VAE), a simple predictive model (random forest), and post hoc explanation (SHAP).

Within this framework, we use a VAE model for dimensionality reduction. Unlike non-parametric embedding methods such as t-distributed Stochastic Neighbor Embedding (t-SNE) and Uniform Manifold Approximation and Projection (UMAP) (Van der Maaten and Hinton, 2008; McInnes et al., 2018), which require recomputation and may yield different results when re-run after introducing new data, a trained VAE provides a parametric mapping that can directly encode unseen spectra into a consistent latent space. Compared to Principal Component Analysis (PCA), which also generalizes to unseen data through a fixed linear projection, VAE further enforces a smooth latent distribution via variational regularization, making it more suitable for modeling continuous variation in leaf spectra. In addition, compared to standard or stacked autoencoders, its probabilistic formulation regularizes the latent space and reduces overfitting, which is beneficial for limited and noisy hyperspectral datasets.

This enables compression of high-dimensional spectral data into a compact latent space, followed by clustering to form approximately 20–25 representative spectral groups. These clustered features are then used as inputs for downstream prediction, allowing SHAP to identify the most influential spectral clusters associated with conditions such as low nitrogen or herbicide-induced damage, and to map their contributions back onto spatial leaf regions. This strategy reduces overfitting by compressing hundreds of spectral bands into a small set of interpretable features, while improving interpretability through direct visualization of condition-related spectral patterns mapped onto the spatial locations of leaves.

The LeafVAE employed a standard variational autoencoder architecture with fully connected encoder and decoder networks. The encoder projected the hyperspectral input spectra through a 400-unit hidden layer and a 200-unit latent feature layer, each with LeakyReLU activation ($\alpha = 0.2$). Its output was then passed through two separate

Table. 1
Summary of all experiments.

Dataset	Species	Topic	Treatments	Replications	Imaging equipment	Date
Corn leaf	Beck's hybrid 5994V2P, Pioneer hybrid P1105AM	Nitrogen fertilization	22.4 kg/ha, 224.2 kg/ha	4 duplicate plots, each with 30 plants	Handheld hyperspectral (LeafSpec), transmittance mode	May 2022
	Beck's hybrid 5994V2P, Pioneer hybrid P1105AM	Nitrogen fertilization	22.4 kg/ha, 224.2 kg/ha	4 duplicate plots, each with 30 plants	Handheld hyperspectral (LeafSpec), transmittance mode	June 2022
	Pioneer hybrid P1108Q	Nitrogen fertilization	0 kg/ha, 80 kg/ha, 120 kg/ha,160 kg/ha, 200 kg/ha, 240 kg/ha	4 duplicate plots, each with 150 plants	Handheld hyperspectral (LeafSpec), transmittance mode	June 2023
	Beck's hybrid 5994V2P	Nitrogen fertilization	22.4 kg/ha, 44.8 kg/ha, 112.1 kg/ha, 224.2 kg/ha	4 duplicate plots, each with 14 plants	Handheld hyperspectral (LeafSpec), transmittance mode	June 2024
Soybean leaf	Beck's 296 L4	Herbicide	Dicamba at 0.56, 0.28, 0.14, and 0.0695 g/ha; 2,4-D choline at 42.6, 21.3, 14.3, and 10.6 g/ha	20 duplicate plants	Handheld hyperspectral (LeafSpec-Dicot), transmittance mode	November 2020
Muscadine grape leaf	Noble muscadine, Floriana muscadine	Disease	With/without fungicide treatment	4 non-treated Novel veins, 5 non-treated Floriana veins, 3 treated Novel and 3 treated Floriana veins	Tabletop hyperspectral (Via-Spec II Scanner Stage), reflectance mode	May 2025
Beech leaf	American beech	Disease	N/A (surveyed)	30 symptomatic trees	Handheld 8-band multispectral (Leaf Scope), transmittance mode	September 2025

linear layers to independently predict the mean and log-variance of a two-dimensional latent embedding. The decoder mirrored the encoder, expanding the latent embedding through 200 and 400 hidden units with LeakyReLU before reconstructing the input via a sigmoid activation layer. The model was trained to minimize the sum of binary cross-entropy reconstruction loss and a Kullback-Leibler divergence regularizer.

The LeafVAE model was trained using stochastic gradient-based optimization with the Adam optimizer of learning rate = 0.0001. Each hyperspectral leaf image was flattened and reshaped into a 2D array (pixel × wavelength), and pixels containing NaN values were removed. Training was conducted over 1,000 epochs, with one randomly selected leaf image used per epoch. Within each epoch, spectra were processed in batches of 100 pixels, and model parameters were updated after each batch.

3. Results

We first collected hyperspectral transmittance images of corn leaves at the R1 growth stage, which marks the beginning of the reproductive phase, to evaluate the ability of LeafVAE to encode spectral information (Fig. 2). The images were calibrated to correct for lighting nonuniformity and camera noise using standard post-processing procedures, and leaf areas were segmented using thresholding at selected wavelengths (**Methods**). As the three-dimensional (3D) hyperspectral leaf image cannot be directly visualized, a representative corn leaf image is displayed using NDVI to summarize the spatial variation. Wavelength-dependent features are observed in slices of the transmittance maps at 550 nm (green), 620 nm (orange), and 720 nm (red) (Fig. 2**a**). To inspect spectral variation across the leaf, three regions of interest were selected including the midrib, center mesophyll, and side mesophyll, and their normalized transmittance spectra showed distinct signatures between 500 and 750 nm (Fig. 2**b**), which likely arise from variations in cellular organization, mesophyll density (i.e., leaf thickness), and photosynthetic pigment concentrations across the leaf.

After training LeafVAE on the corn leaf dataset, the model learned a compact 2D latent space that organizes spectral variability in a structured manner (Fig. 2, **c** and **d**). This is analogous to VAEs trained on the MNIST dataset, where handwritten digits are mapped into a low-dimensional latent manifold in which visually similar digits are placed in nearby regions, and each latent point can be decoded back into its corresponding image reconstruction (Li Deng, 2012). Similarly, in our case, LeafVAE embeds spectral signals into a continuous 2D latent space, where nearby coordinates correspond to similar spectral shapes. Each latent point represents a compressed encoding (x, y) that can be decoded back into the original spectrum (Fig. 2**c**).

Thus, a hyperspectral leaf image can be represented as a density heatmap in the latent space, where each point corresponds to an individual pixel and density indicates the population of pixels with particular spectral signatures (Fig. 2**d**). To assess the biological interpretability of the latent space, we decoded spectra at each latent coordinate and calculated vegetation indices from the reconstructed spectra. The resulting maps showed smooth variations in NDVI (Fig. 2**e**), photochemical reflectance index (PRI, Fig. 2**f**, (Gamon et al., 1997)), and anthocyanin reflectance index (ARI, Fig. 2**g**, (Gitelson et al., 2007)) across the latent space. We note that these indices are commonly used as indirect spectral proxies for chlorophyll content, xanthophyll cycle activity, and anthocyanin accumulation, whereas direct quantification of the biochemical concentrations would require dedicated chemical measurements.

While the 2D latent embedding provides a smooth and continuous representation of spectral variation, its complexity still poses challenges for direct interpretation or for use in downstream machine learning models. To convert this continuous landscape into discrete and comparable units, we applied unsupervised clustering to the latent embeddings from the corn leaf dataset, generating a set of representative

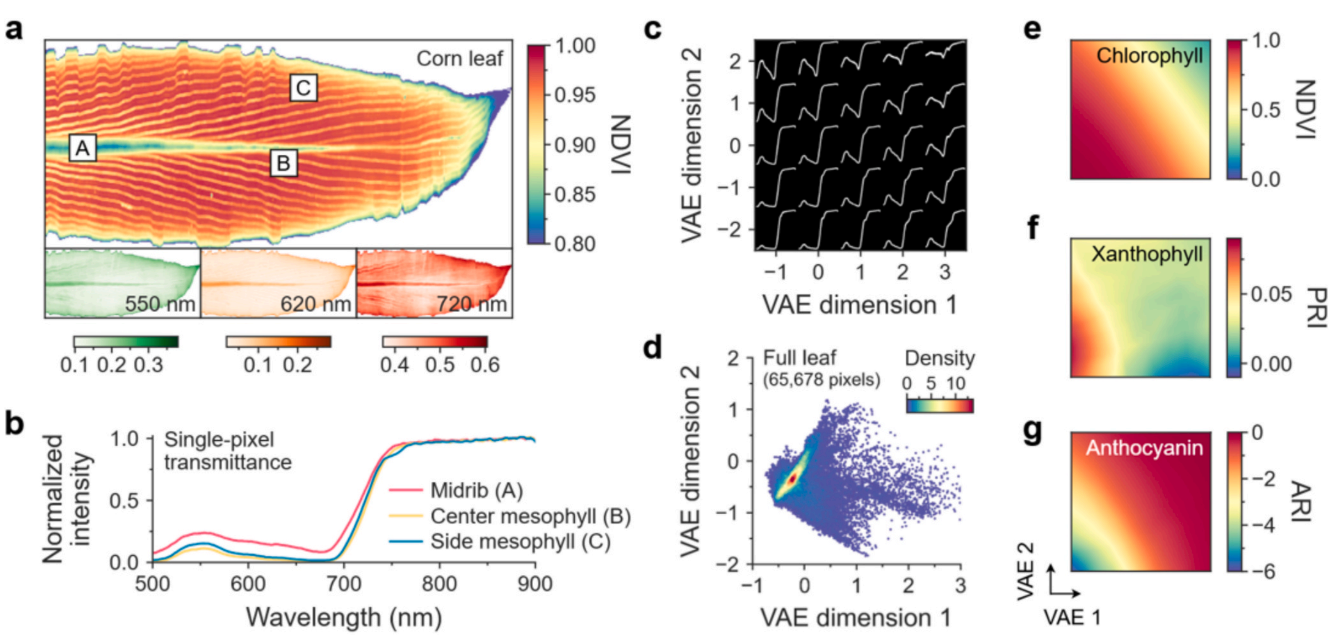

Fig. 2. The latent embedding encodes spectral information of corn leaves. a, Top: a typical NDVI image of a top-collar corn leaf. Bottom: wavelength-specific transmittance images. Regions of interest are marked as A, B, and C. **b**, Single-pixel transmittance spectra from the selected regions in **a**, showing distinct spectral signatures for (A) midrib, (B) center mesophyll, and (C) side mesophyll. The transmittance spectra are normalized to mitigate intensity variations caused by differences in leaf thickness across regions. **c**, Two-dimensional latent space representation learned by LeafVAE, where each (x, y) point corresponds to a decoded hyperspectral spectrum. **d**, Visualization of the entire leaf spectrum in latent space. **e-g**, Latent space distribution of photosynthetic pigments: (**e**) chlorophyll (NDVI), (**f**) xanthophyll (photochemical reflectance index, PRI), and (**g**) anthocyanin (anthocyanin reflectance index, ARI). Here, we compute vegetation indices directly from the measured transmittance spectra by substituting them into the standard index formulations, without converting between modalities. This is presented as a demonstration to illustrate relative spectral variations, and the resulting values are not expected to exactly match those obtained from standard reflectance-based measurements.

spectral signatures that can be easily tracked and analyzed across leaves (Fig. 3). We first visualized all spectral signatures as a heatmap showing their spectral profiles across the hyperspectral wavelength window. Each row represents a distinct spectral signature, with colors indicating the percentile rank at each wavelength relative to the full dataset spectrum (Fig. 3**a**, left). The pixel count distribution (Fig. 3**a**, right) revealed that certain signatures were highly abundant across the dataset, while others were relatively rare, reflecting the natural heterogeneity of the leaf.

The number of spectral signature clusters was determined by first calculating silhouette scores across a range of k values (Fig. 3**b**). The analysis revealed several local maxima (k = 6, 15, and 23), indicating that the spectral data possess a hierarchical structure with meaningful groupings at multiple levels of granularity. While the silhouette score favored lower k values, we selected a slightly higher number of clusters (k = 25) to capture finer spectral variations and improve the generalizability to new datasets with potentially greater spectral diversity. Two least populated signatures were arbitrarily removed using a conservative threshold (<0.01% of total pixels) to exclude likely noise signals, yielding a final set of 23 representative clusters. This removal had minimal impact on downstream tasks (**Supplementary** Fig. 1).

Visualization of the spectral embeddings in latent space revealed clear, well-separated spectral signature clusters in the training dataset (Fig. 3**d**). To evaluate the generalization of the identified clusters, we fixed the positions of cluster centroids and used them as reference points for projecting new datasets into the same latent space. Spectral data from independent validation corn leaf datasets, including same-year data under identical experimental conditions (Fig. 3**e**), cross-year data with different genotypes and field environments (Fig. 3**f**), and cross-year data collected using different equipment and imaging protocols (Fig. 3**g**), were consistently aligned with the predefined cluster centroids (dataset details in **Methods**). We then assessed the reconstruction fidelity of LeafVAE across these datasets by calculating the absolute reconstruction error in spectra. The reconstruction error remained consistently low across all wavelengths for the training dataset (Fig. 3**h**) and generalized well to the validation datasets (Fig. 3**i–k**), with mean errors below 2% across the spectrum. We note that changes in the imaging device can shift the captured spectrum (likely due to calibration and instrumental differences), which explains the higher reconstruction error observed in Fig. 3**k**. These results demonstrate that the spectral signatures can generalize across new datasets with varying experimental conditions.

We then investigated the diagnostic utility of these spectral signatures for detecting plant physiological stress. We started with nitrogen deficiency, a major factor affecting corn yield. We collected corn leaves from well-fertilized (healthy) and low-nitrogen (low-N) plots for comparison (Fig. 4). Visualization of spectral signature composition revealed consistent differences between healthy and low-N leaves across all four datasets (**Supplementary** Fig. 2). The spatial signature maps showed that low-N leaves displayed more heterogeneous distributions of spectral signatures compared to healthy controls (Fig. 4**a, b**).

To measure these differences and find the spectral signatures that best indicate nitrogen status, we trained a random forest regression model using spectral signature compositions as input features to predict leaf nitrogen content (**Methods**). This downstream modeling step enables interpretable analysis through Shapley additive explanations (SHAP) (Lundberg and Lee, 2017), a method that quantifies the contribution of each feature to individual predictions. Analysis of pixel percentages across signatures confirmed that certain signatures were more prevalent in each treatment, and SHAP values revealed which signatures contributed positively or negatively to nitrogen predictions (Fig. 4**c, d**). When comparing the SHAP values of these two leaves, the prediction contribution from spectral signatures C5 and C14 showed notable changes, from 0.05% and 0.08% to −0.03% and −0.04%, respectively. Pixel-level nitrogen prediction maps were generated by assigning SHAP values to spectral signatures at each pixel (Fig. 4**e, f**). By plotting the prediction maps across the dataset, we found that low-N signals (regions with low SHAP value allocation) varied with leaf position on the plant (**Supplementary** Fig. 3), consistent with source-sink nitrogen remobilization dynamics (Tegeder and Masclaux-Daubresse, 2018). Lower leaves showed nitrogen deficiency patterns near the leaf base, whereas top-collar leaves exhibited deficiency signatures along the midrib. Such

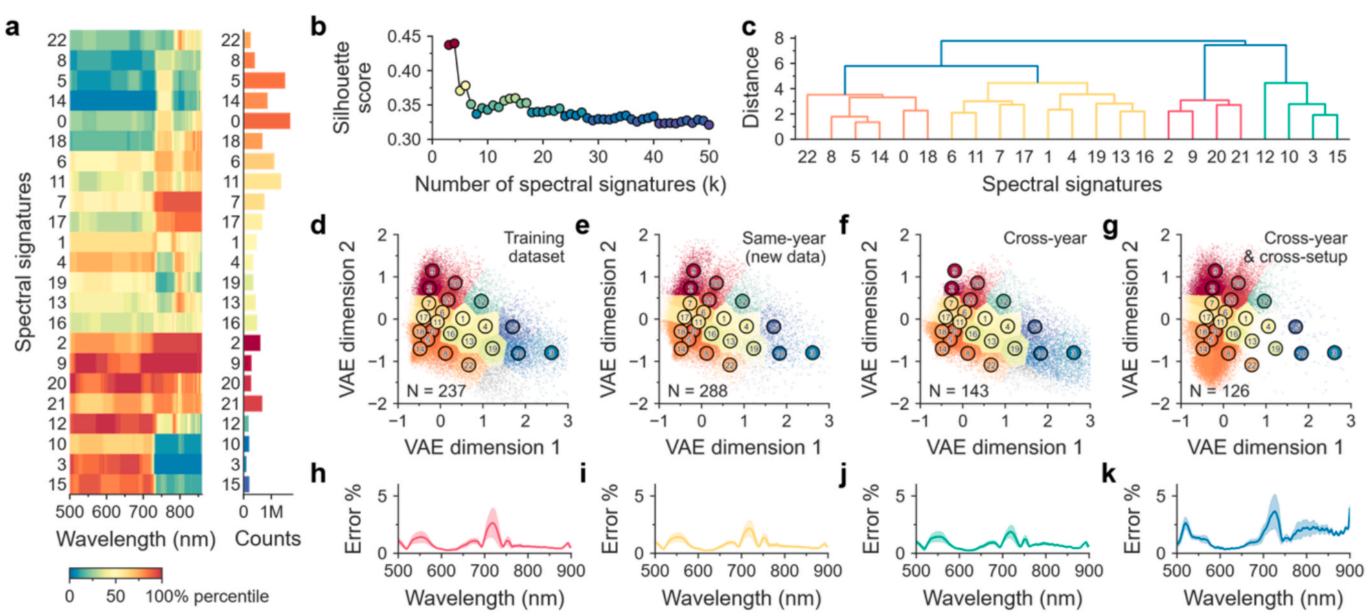

Fig. 3. Unsupervised clustering of spectral signatures of corn leaves. a, Left: Heatmap showing the spectral signatures across the hyperspectral wavelength window (860–900 nm used for normalization). Colors indicate the percentile of each spectral signature at a given channel relative to the spectrum of the corn leaf dataset. Right: Total number of pixels in each cluster across the dataset. **b**, Silhouette score analysis for determining the optimal number of spectral signature clusters (k). **c**, Hierarchical clustering dendrogram of the 23 identified spectral signatures. **d-g**, Visualization of the full dataset spectrum in latent space: (**d**) training dataset, (**e**) same-year validation dataset, (**f**) cross-year dataset, and (**g**) cross-year dataset with different equipment and imaging/calibration protocol. **h-k**, Absolute reconstruction error of the LeafVAE encoder-decoder on normalized transmittance spectra from the dataset corresponding to **d-g**. Shaded regions represent the standard error across leaves.

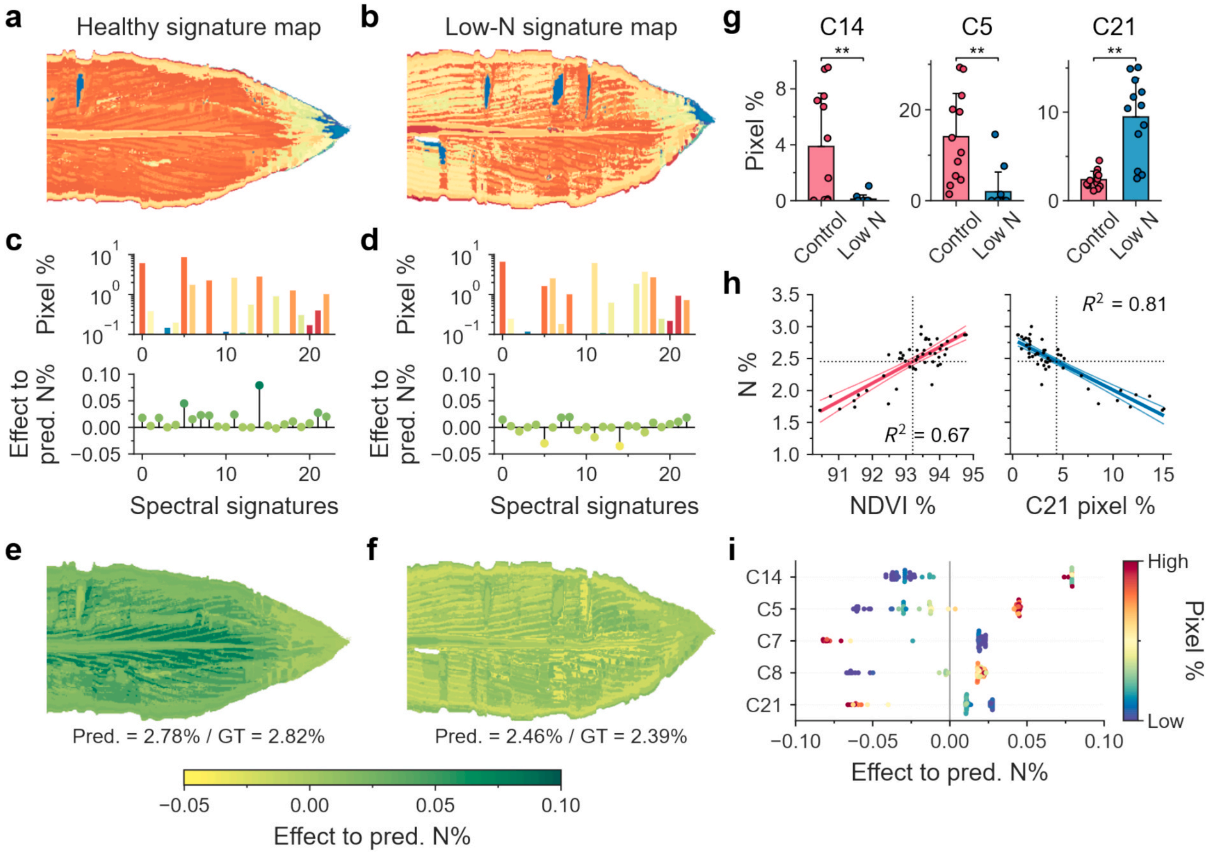

Fig. 4. Spectral signature analysis enables visualization of nitrogen deficiency in corn leaves. a, b, Spatial distribution of spectral signatures for representative (**a**) healthy and (**b**) nitrogen-deficient (low-N) corn leaves. **c, d**, Top: pixel percentage distribution across 23 spectral signatures for leaves in **a** and **b**. Bottom: SHAP values of pixel percentages on predicted corn nitrogen content. A random forest model was trained to extract SHAP values. **e, f**, Nitrogen prediction maps for healthy and low-N leaves in **a** and **b**, visualized with SHAP values from each spectral signature. **g**, Pixel percentages of key spectral signatures in control versus low-N leaves. **$P < 0.01$ (two-sided Mann-Whitney U test, N = 12, false discovery rate corrected via Benjamini–Hochberg adjustment). **h**, Correlation between corn nitrogen content and NDVI/spectral signature pixel percentage. **i**, Relationship between pixel percentage and nitrogen prediction contribution of the top five important signature clusters.

patterns were consistently observed in both genotypes (**Supplementary** Fig. 3 and 4).

To identify key diagnostic spectral signatures, we first performed a model-free statistical analysis. Among the 23 spectral signatures, 19 showed significant differences between healthy and low-N leaves. Three clusters (C14, C5, and C21) exhibited the most pronounced differences between the two treatments (Fig. 4**g, P** < 0.01, two-sided Mann-Whitney U test, N = 12), with C14 and C5 increased in healthy leaves and C21 in nitrogen-deficient leaves. These signatures also showed the strongest correlations with chemically measured plant nitrogen content. Among them, C21 achieved the best predictive performance, with its pixel percentage alone ($R^2 = 0.81$) outperforming NDVI ($R^2 = 0.67$) (Fig. 4**h**). The top five features identified by the random forest model and the SHAP importance ranking were consistent with the statistical findings (Fig. 4**i**).

Beyond diagnosing nutrient deficiencies, we tested whether LeafVAE could distinguish chemical stresses induced by herbicides (Fig. 5). Herbicide damage is a major challenge in modern agriculture, as misapplication can greatly reduce yields and harm nearby fields. We applied LeafVAE to hyperspectral reflectance images of soybean leaves exposed to two synthetic auxin herbicides with similar modes of action, 3,6-dichloro-2-methoxybenzoic acid (dicamba) and 2,4-dichlorophenoxyacetic acid (2,4-D), along with untreated controls (Fig. 5**a, b**; dataset collection described in **Methods**). Although the soybean images were captured using a tabletop system with lower spatial resolution than the handheld device used for corn, their spectral signatures showed spectral patterns consistent with those observed in corn (Fig. 3**a**), suggesting that transfer learning could be effective across species.

We mapped SHAP values from spectral clusters onto the full leaf to visualize how the model distinguishes each herbicide class (Fig. 5**c**). Interestingly, the stress patterns for 2,4-D appeared mainly along the leaf edges, likely due to leaf deformation under stress that alters the reflected spectra. We then compared LeafVAE against conventional vegetation indices using 5-fold cross-validation with the same random forest classifier architecture (Fig. 5**d**). While NDVI achieved 63.5% accuracy and PRI and ARI performed lower (33.5% and 53.1%, respectively), LeafVAE reached 74.2% accuracy. The improved performance can be attributed to learning data-driven spectral representations from the full spectrum, which allows it to capture subtle physiological responses associated with different herbicide modes of action. We note that a key advantage of the LeafVAE embedding, compared to a single vegetation index, is that its spectral encoding and clustered representation enables the use of SHAP to map model attributions back onto spatial leaf regions (Fig. 5**c**). Visual inspection of spectral signature components revealed noticeable differences between different herbicide dosages (**Supplementary** Fig. 5), and LeafVAE achieved higher accuracy than other methods for herbicide dose classification for both Dicamba and 2,4-D (**Supplementary** Fig. 6).

We also applied the same LeafVAE pipeline to hyperspectral images of muscadine grape leaves to diagnose fungal diseases (**Supplementary** Fig. 7). Across two muscadine grape cultivars, Floriana and Noble, we observed distinct spectral responses associated with disease symptoms, including black rot, angular leaf spot, and combined symptoms. Noble grape leaves exhibited more subtle spectral changes under disease stress,

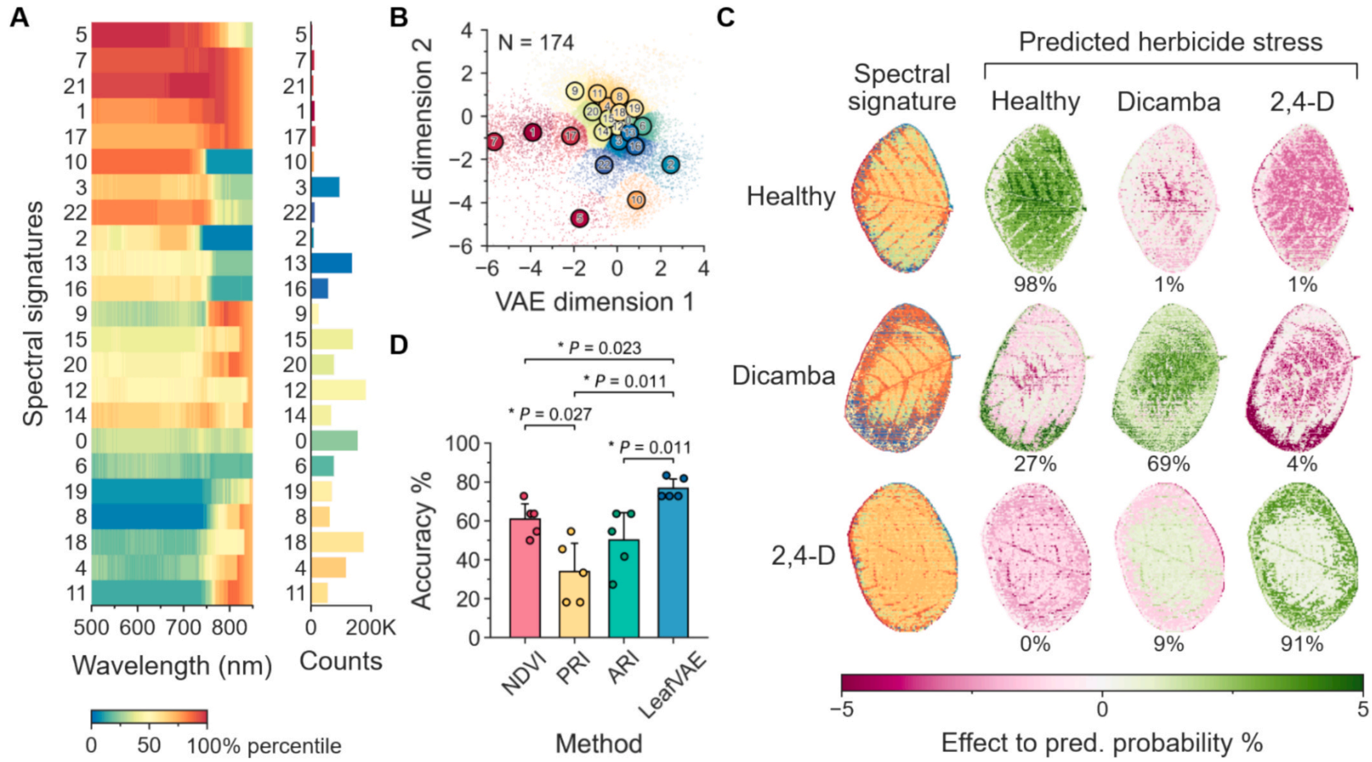

Fig. 5. Herbicide stress assessment of soybean leaves. a, Left: Heatmap showing the spectral signatures across the hyperspectral wavelength window. Colors indicate the percentile of each spectral signature at a given channel relative to the spectrum of the soybean dataset. Right: Total number of pixels in each cluster across the dataset. **b**, Visualization of the full dataset spectrum in the latent space of LeafVAE. **c**, Visualization of spectral signatures and their predicted health status (healthy, dicamba, or 2,4-D herbicide stress). Colors represent Shapley Additive Explanations (SHAP) values, indicating the influence of spectral features on the predicted probability of the health state: green contributes to a positive prediction, purple to a negative prediction. **d**, Classification accuracy for herbicide stress classification of untreated, dicamba-treated (0.56 g ae/ha), and 2,4-D-treated (42.6 g ae/ha) soybean leaves at 7 days after herbicide spray treatment. A random forest model was trained and evaluated using five-fold cross-validation for each method.

suggesting a greater inherent robustness compared to Floriana.

A key practical consideration for field deployment is whether LeafVAE can perform effectively with affordable multispectral cameras, which have lower spectral resolution than conventional hyperspectral imaging systems. To test this, we applied LeafVAE to 6-band multispectral images of American beech leaves affected by beech leaf disease (BLD), an emerging and rapidly spreading disease caused by a foliar nematode (Ewing et al., 2019) (Fig. 6, dataset details in **Methods**). The spectral heatmap visualizations showed clear differentiation in spectral patterns across these channels (Fig. 6**a** and **Supplementary** Fig. 8). Note we did not calibrate the lighting across the images to mimic a low-cost setup that can be deployed in the field without a technician, so variations in the reflectance spectra may reflect illumination differences rather than biological differences. For each infected tree surveyed, we sampled two asymptomatic leaves and two leaves with symptoms of BLD. We visualized the composition of spectral signatures (left heatmap) alongside three tree-level characteristics recorded by field experts: percentage of canopy cover (%CC), diameter at breast height (DBH in cm), and the percentage of healthy leaves (converted to a 1–4 infection severity level) (Fig. 6**b**). The spectral signatures differed between healthy and diseased leaves, with BLD-affected samples exhibiting higher proportions of specific clusters, such as C20 (shown in deep blue of the bar plot).

We next generated spatial prediction maps to visualize how spectral signatures contributed to BLD classification (Fig. 6**c**). Note that the leaf tip and base had uneven lighting conditions in band 6, resulting in minimal contributions from these regions. In a moderately affected leaf, 38% of signatures were classified as healthy and concentrated in less-damaged regions, while 62% corresponded to BLD signatures in symptomatic areas. A severely diseased leaf showed only 6% healthy signatures, with 94% identified as BLD-affected across nearly the entire leaf surface. These results suggest that our method can both classify and potentially quantify disease severity at the individual leaf level.

4. Conclusions

In this work, we developed LeafVAE, a variational autoencoder framework for unsupervised hyperspectral leaf analysis that learns spectral signatures without manual annotation. These signatures generalize across genotypes, years, and imaging equipment, and can be extracted from both conventional hyperspectral imaging systems and affordable six-band multispectral cameras. When coupled with explainable AI methods, our workflow enables pixel-level visualization of physiological stress patterns, providing spatially resolved diagnostics for plant phenotyping.

However, a major challenge in plant phenotyping is data scarcity, which makes it difficult to obtain large, well-annotated datasets covering the diversity of G × E × M conditions. This issue is particularly relevant for developing new genotypes, testing novel herbicides, or studying emerging diseases, where expert labeling is costly and time-consuming. While LeafVAE helps address part of this challenge through unsupervised spectral feature learning, unsupervised learning alone is not sufficient. Several methods could help reduce data scarcity in future plant phenotyping studies. Transfer learning can reduce data needs by pre-training models on large, diverse, or simulated datasets across species and then fine-tuning them for specific tasks. This approach takes advantage of spectral and physiological features that are conserved across plant systems (Zhang et al., 2021). Data augmentation methods, such as interpolation along stress gradients, wavelength-specific perturbations, or simulated lighting conditions, could expand

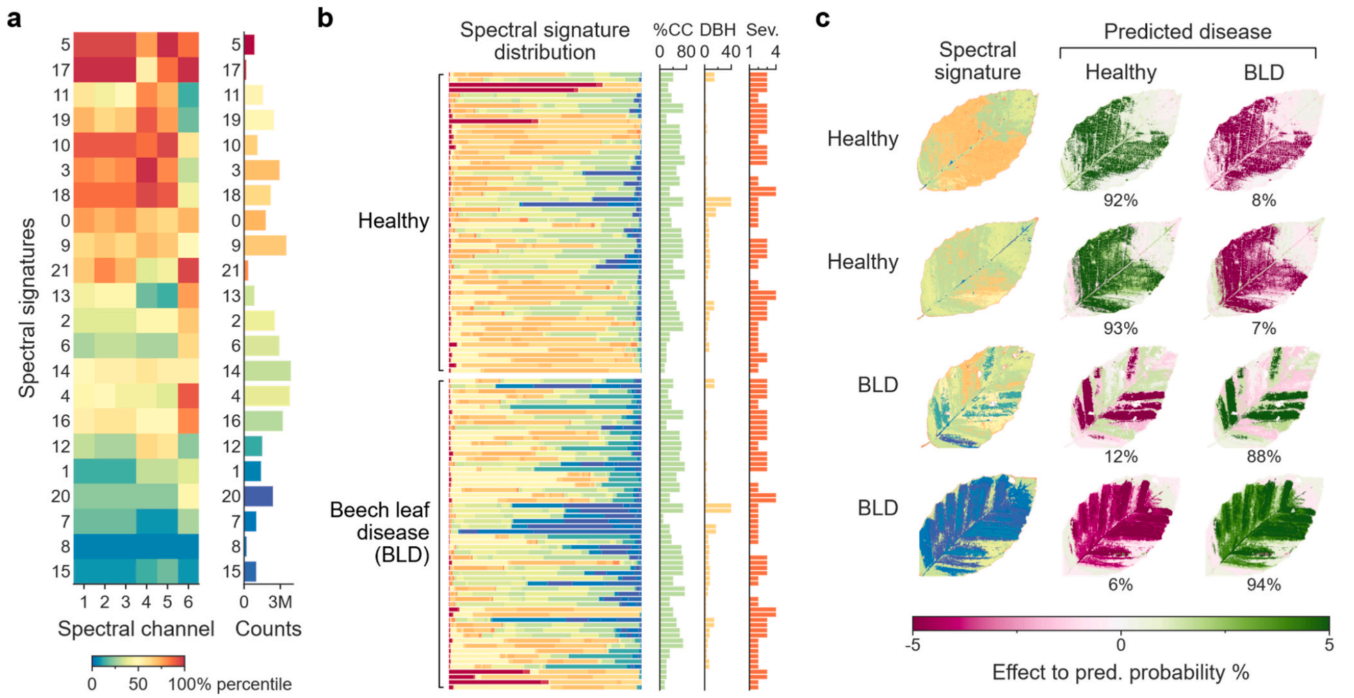

Fig. 6. Detecting beech leaf disease (BLD) on American beech using an affordable 6-band multispectral camera. a, Left: Heatmap showing spectral signatures across 6-band wavelength channels (660, 560, 473, 595, 400, and 880 nm, respectively). Colors indicate the percentile of each spectral signature at a given channel relative to the spectrum of the American beech leaf dataset. Right: Total number of pixels in each cluster across the dataset. **b**, Spectral signature distribution of individual leaves separated by asymptomatic and BLD-symptomatic samples. Adjacent bars indicate tree-level characteristics recorded by experts: %CC, DBH (in cm), and the infection severity level. Each row corresponds to a single leaf sample. **c**, Visualization of spectral signatures (left) and their predicted health status (middle: healthy, right: BLD). Colors represent SHAP values, showing the influence of spectral features on the predicted probability of each health state: green indicates a positive contribution, while purple indicates a negative contribution.

training datasets, but it is important to preserve realistic relationships between pigment composition and spectral signals. Finally, creating open-source, standardized benchmark datasets in plant phenotyping, similar to how open datasets in biomedical research enable the rapid development of novel and robust AI diagnostic tools, would support careful evaluation of models and accelerate progress across the field.

The two-stage design of LeafVAE with explainable AI introduces several limitations. In this framework, spectral signatures are first learned in an unsupervised manner from hyperspectral data, and then a separate classifier or regressor (a random forest in our demonstrations) is trained using the signature composition percentages as input features. The interpretability of predictions depends on post-hoc explainability methods, such as SHAP feature values applied to the downstream model, rather than emerging directly from the latent representation itself. The need for labeled data to train the second-stage model also limits the applicability of the framework to new research where no labeled samples are available.

Achieving fully unsupervised, end-to-end diagnosis, a “one-for-all” model capable of detecting arbitrary stress conditions without any labeled training data, remains an open challenge. Several promising directions could help move toward this goal. First, while LeafVAE currently operates on pixel-level spectra, integrating spatial context (i.e., spatial encoding) through convolutional or graph-based representations could improve the identification of localized or patterned stress responses. Second, multimodal learning frameworks that combine spectral information with complementary data sources (such as soil properties, weather records, or management metadata) could capture the broader environmental and physiological context. Third, implementing hierarchical or multi-resolution clustering within an unsupervised architecture could automatically determine the appropriate number of spectral clusters, reducing reliance on manual k-means selection. Such an approach would better capture both broad stress categories and fine-grained spectral variations, enabling consistent interpretation across genotypes and species within a single, unified model.

Despite these limitations, LeafVAE represents a step toward scalable and interpretable plant phenotyping. By demonstrating consistent performance across multiple species (corn, soybean, grape, and beech), stress types (nutrient deficiency, herbicide injury, and disease), and imaging modalities (hyperspectral and multispectral), the framework provides a foundation for data-driven approaches in agriculture and plant science. Looking ahead, integrating the strategies outlined above could advance the phenotyping technology toward fully autonomous monitoring systems capable of diagnosing plant health directly in field environments.

CRediT authorship contribution statement

Kangyu Ji: Writing – original draft, Visualization, Software, Methodology, Formal analysis, Data curation, Conceptualization. **Zhiyuan Chen:** Writing – original draft, Resources, Investigation, Data curation. **Zhongzhong Niu:** Resources, Investigation, Data curation. **Alden Mo:** Resources, Investigation. **Zhihao Qin:** Resources, Investigation. **Katie DeCamp:** Investigation. **Charles Wang:** Investigation. **Weixuan Xu:** Investigation. **Julie M. Young:** Resources, Investigation. **Bryan G. Young:** Resources, Investigation. **Jingqiu Chen:** Writing – review & editing, Resources, Investigation, Data curation. **Anna O. Conrad:** Writing – review & editing, Resources, Investigation. **Jian Jin:** Writing – review & editing, Supervision, Project administration, Conceptualization.

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.

Acknowledgments

There is no direct funding source for this study. The beech leaf dataset was collected with support from the USDA Forest Service. Special thanks to the Ohio Department of Natural Resources for allowing access to Johnson Woods State Nature Preserve. JC would like to acknowledge the 1890 Institution Teaching, Research and Extension Capacity Building Grants Program, Award 2024-38821-42099, from the U.S. Department of Agriculture's National Institute of Food and Agriculture. The authors gratefully acknowledge Dr. Violeta M. Tsolova, Professor and Director of the Center for Viticulture and Small Fruit Research at Florida A&M University, for her support with vineyard access, field sampling, and vineyard management information.

Appendix A. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.compag.2026.111971.

Data availability

Data will be made available on request.

References

Amirruddin, A.D., Muharam, F.M., Ismail, M.H., Ismail, M.F., Tan, N.P., Karam, D.S., 2020. Hyperspectral remote sensing for assessment of chlorophyll sufficiency levels in mature oil palm (*Elaeis guineensis*) based on frond numbers: Analysis of decision tree and random forest. Comput. Electron. Agric. 169, 105221. https://doi.org/10.1016/j.compag.2020.105221.

Aviles Toledo, C., Crawford, M.M., Tuinstra, M.R., 2024. Integrating multi-modal remote sensing, deep learning, and attention mechanisms for yield prediction in plant breeding experiments. Front. Plant Sci. 15, 1408047. https://doi.org/10.3389/fpls.2024.1408047.

Bannari, A., Morin, D., Bonn, F., Huete, A.R., 1995. A review of vegetation indices. Remote Sens. Rev. 13, 95–120. https://doi.org/10.1080/02757259509532298.

Barbedo, J.G.A., 2019. Detection of nutrition deficiencies in plants using proximal images and machine learning: a review. Comput. Electron. Agric. 162, 482–492. https://doi.org/10.1016/j.compag.2019.04.035.

Brannen, P.M., 2019. The State of the Art in Muscadine Disease Management [WWW Document]. URL https://viticulture.uga.edu/files/2019/01/2019-Muscadine-Diseases-Presentation.pdf.

Danilevicz, M.F., Upadhyaya, S.R., Batley, J., Bennamoun, M., Bayer, P.E., Edwards, D., 2025. Understanding plant phenotypes in crop breeding through explainable AI. Plant Biotechnol. J. 23, 4200–4213. https://doi.org/10.1111/pbi.70208.

Ewing, C.J., Hausman, C.E., Pogacnik, J., Slot, J., Bonello, P., 2019. Beech leaf disease: an emerging forest epidemic. For. Pathol. 49, e12488. https://doi.org/10.1111/efp.12488.

Gamon, J.A., Serrano, L., Surfus, J.S., 1997. The photochemical reflectance index: an optical indicator of photosynthetic radiation use efficiency across species, functional types, and nutrient levels. Oecologia 112, 492–501. https://doi.org/10.1007/s004420050337.

Gano, B., Bhadra, S., Vilbig, J.M., Ahmed, N., Sagan, V., Shakoor, N., 2024. Drone-based imaging sensors, techniques, and applications in plant phenotyping for crop breeding: a comprehensive review. Plant Phenome J. 7, e20100. https://doi.org/10.1002/ppj2.20100.

Gill, T., Gill, S.K., Saini, D.K., Chopra, Y., De Koff, J.P., Sandhu, K.S., 2022. A comprehensive review of high throughput phenotyping and machine learning for plant stress phenotyping. Phenomics 2, 156–183. https://doi.org/10.1007/s43657-022-00048-z.

Gitelson, A.A., Merzlyak, M.N., Chivkunova, O.B., 2007. Optical properties and nondestructive estimation of anthocyanin content in plant leaves¶. Photochem. Photobiol. 74, 38–45. https://doi.org/10.1562/0031-8655(2001)0740038OPANEO2.0.CO2.

Henry, W.B., Shaw, D.R., Reddy, K.R., Bruce, L.M., Tamhankar, H.D., 2004. Remote sensing to detect herbicide drift on crops. Weed Technol. 18, 358–368. https://doi.org/10.1614/WT-03-098.

Kniss, A.R., 2017. Long-term trends in the intensity and relative toxicity of herbicide use. Nat. Commun. 8, 14865. https://doi.org/10.1038/ncomms14865.

Lapajne, J., Vončina, A., Vojnović, A., Donša, D., Dolničar, P., Žibrat, U., 2025. Field-scale UAV-based multispectral phenomics: Leveraging machine learning, explainable AI, and hybrid feature engineering for enhancements in potato phenotyping. Comput. Electron. Agric. 229, 109746. https://doi.org/10.1016/j.compag.2024.109746.

Lee, G., Hossain, O., Jamalzadegan, S., Liu, Y., Wang, H., Saville, A.C., Shymanovich, T., Paul, R., Rotenberg, D., Whitfield, A.E., Ristaino, J.B., Zhu, Y., Wei, Q., 2023. Abaxial leaf surface-mounted multimodal wearable sensor for continuous plant physiology monitoring. Sci. Adv. 9, eade2232. https://doi.org/10.1126/sciadv.ade2232.

Deng, L.i., 2012. The MNIST database of handwritten digit images for machine learning research [best of the web]. IEEE Signal Process. Mag. 29, 141–142. https://doi.org/10.1109/MSP.2012.2211477.

Li, X., Chen, Z., Wang, J., Jin, J., 2023. LeafSpec-dicot: an accurate and portable hyperspectral imaging device for dicot leaves. Sensors 23, 3687. https://doi.org/10.3390/s23073687.

Li, Z., Paul, R., Ba Tis, T., Saville, A.C., Hansel, J.C., Yu, T., Ristaino, J.B., Wei, Q., 2019. Non-invasive plant disease diagnostics enabled by smartphone-based fingerprinting of leaf volatiles. Nat. Plants 5, 856–866. https://doi.org/10.1038/s41477-019-0476-y.

Lundberg, S., Lee, S.-I., 2017. A Unified Approach to Interpreting Model Predictions. Doi: 10.48550/ARXIV.1705.07874.

McInnes, L., Healy, J., Melville, J., 2018. Umap: Uniform manifold approximation and projection for dimension reduction. ArXiv Prepr. ArXiv180203426.

Nelson, D.W., Sommers, L.E., 1980. Total nitrogen analysis of soil and plant tissues. J. AOAC Int. 63, 770–778. https://doi.org/10.1093/jaoac/63.4.770.

Nosengo, N., 2003. Fertilized to death. Nature 425, 894–895. https://doi.org/10.1038/425894a.

Ollinger, S.V., 2011. Sources of variability in canopy reflectance and the convergent properties of plants. New Phytol. 189, 375–394. https://doi.org/10.1111/j.1469-8137.2010.03536.x.

Rasti, B., Hong, D., Hang, R., Ghamisi, P., Kang, X., Chanussot, J., Benediktsson, J.A., 2020. Feature extraction for hyperspectral imagery: the evolution from shallow to deep: overview and toolbox. IEEE Geosci. Remote Sens. Mag. 8, 60–88. https://doi.org/10.1109/MGRS.2020.2979764.

Rehman, T.U., Ma, D., Wang, L., Zhang, L., Jin, J., 2020. Predictive spectral analysis using an end-to-end deep model from hyperspectral images for high-throughput plant phenotyping. Comput. Electron. Agric. 177, 105713. https://doi.org/10.1016/j.compag.2020.105713.

Ribeiro, M.T., Singh, S., Guestrin, C., 2016. "Why Should I Trust You?": Explaining the Predictions of Any Classifier, in: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. Presented at the KDD '16: The 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, ACM, San Francisco California USA, pp. 1135–1144. Doi: 10.1145/2939672.2939778.

Rouse Jr, J.W., Haas, R.H., Deering, D.W., Schell, J.A., Harlan, J.C., 1974. Monitoring the vernal advancement and retrogradation (green wave effect) of natural vegetation.

Sarić, R., Nguyen, V.D., Burge, T., Berkowitz, O., Trtílek, M., Whelan, J., Lewsey, M.G., Čustović, E., 2022. Applications of hyperspectral imaging in plant phenotyping. Trends Plant Sci. 27, 301–315. https://doi.org/10.1016/j.tplants.2021.12.003.

Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D., 2020. Grad-CAM: visual explanations from deep networks via gradient-based localization. Int. J. Comput. vis. 128, 336–359. https://doi.org/10.1007/s11263-019-01228-7.

Song, Z., Wei, X., Zhang, J., Chen, Z., Jin, J., 2024. Spatial-spectral feature mining in hyperspectral corn leaf venation structure and its application in nitrogen content estimation. Comput. Electron. Agric. 227, 109495. https://doi.org/10.1016/j.compag.2024.109495.

Sun, Z., Wang, X., Wang, Z., Yang, L., Xie, Y., Huang, Y., 2021. UAVs as remote sensing platforms in plant ecology: review of applications and challenges. J. Plant Ecol. 14, 1003–1023. https://doi.org/10.1093/jpe/rtab089.

Tegeder, M., Masclaux-Daubresse, C., 2018. Source and sink mechanisms of nitrogen transport and use. New Phytol. 217, 35–53. https://doi.org/10.1111/nph.14876.

Van der Maaten, L., Hinton, G., 2008. Visualizing data using t-SNE. J. Mach. Learn. Res. 9.

Wang, L., Jin, J., Song, Z., Wang, J., Zhang, L., Rehman, T.U., Ma, D., Carpenter, N.R., Tuinstra, M.R., 2020. LeafSpec: an accurate and portable hyperspectral corn leaf imager. Comput. Electron. Agric. 169, 105209. https://doi.org/10.1016/j.compag.2019.105209.

Xu, P., Li, G., Zheng, Y.i., Fung, J.C.H., Chen, A., Zeng, Z., Shen, H., Hu, M., Mao, J., Zheng, Y., Cui, X., Guo, Z., Chen, Y., Feng, L., He, S., Zhang, X., Lau, A.K.H., Tao, S., Houlton, B.Z., 2024. Fertilizer management for global ammonia emission reduction. Nature 626, 792–798. https://doi.org/10.1038/s41586-024-07020-z.

Xue, J., Su, B., 2017. Significant remote sensing vegetation indices: a review of developments and applications. J. Sens. 2017, 1–17. https://doi.org/10.1155/2017/1353691.

Zhang, J., Wei, X., Song, Z., Ampong, K., Beltrame, A., Zhao, T., Penn, C.J., Jin, J., 2026. Hyperspectral image-based leaf-level spatial and spectral feature mining for phosphorus deficiency symptom differentiation in corn plants at early vegetative stage. Comput. Electron. Agric. 241, 111245. https://doi.org/10.1016/j.compag.2025.111245.

Zhang, J., Wei, X., Song, Z., Chen, Z., Jin, J., 2025. A high-precision spatial and spectral imaging solution for accurate corn nitrogen content level prediction at early vegetative growth stages. Comput. Electron. Agric. 230, 109940. https://doi.org/10.1016/j.compag.2025.109940.

Zhang, L., Wang, L., Wang, J., Song, Z., Rehman, T.U., Bureetes, T., Ma, D., Chen, Z., Neeno, S., Jin, J., 2019. Leaf Scanner: a portable and low-cost multispectral corn leaf scanning device for precise phenotyping. Comput. Electron. Agric. 167, 105069. https://doi.org/10.1016/j.compag.2019.105069.

Zhang, Y., Hui, J., Qin, Q., Sun, Y., Zhang, T., Sun, H., Li, M., 2021. Transfer-learning-based approach for leaf chlorophyll content estimation of winter wheat from hyperspectral data. Remote Sens. Environ. 267, 112724. https://doi.org/10.1016/j.rse.2021.112724.