ICPA
17th INTERNATIONAL CONFERENCE ON PRECISION AGRICULTURE
ConBAP 2026
CONGRESSO BRASILEIRO DE AGRICULTURA DE PRECISÃO E DIGITAL
13-16 July 2026 | Porto Alegre, Brazil

Enhancing Pest Detection Through Spectral Signature Extraction in Hyperspectral Data

André Oliveira Françani[a], João de Deus Ferreira e Silva[a], Kaique Micael Moreira de Oliveira[b], Lúcio André de Castro Jorge[c], Joaquim Cezar Felipe[a], Liang Zhao[a].

[a] Faculty of Philosophy, Sciences and Letters at Ribeirão Preto (FFCLRP), University of São Paulo, Ribeirão Preto, 14040-901, SP, Brazil.

[b] Institute of Mathematics and Computer Science (ICMC), University of São Paulo, São Carlos, 13566-590, SP, Brazil.

[c] National Research Center for Development of Agricultural Instrumentation, Brazilian Agricultural Research Agency, São Carlos, 13561-206, SP, Brazil.

A paper from the Proceedings of the

17th International Conference on Precision Agriculture and 11th Brazilian Congress on Precision and Digital Agriculture

13-15 July 2026

Porto Alegre, Brazil

Abstract.

Detecting insect infestation is essential for effective crop protection, particularly in large-scale systems. Caterpillars and stink bugs induce physiological and structural alterations in plant tissues that can be captured through hyperspectral reflectance sensing, which is a non-destructive technique that measures plant responses across hundreds of wavelengths. However, raw spectral signatures are characterized by high dimensionality, strong inter-band correlation, and they often exhibit baseline variability driven by crop-specific attributes, illumination conditions, and sensor effects. These factors can mask the subtle oscillatory patterns and absorption features associated with pest damage, limiting the discriminative power of conventional reflectance-based analyses. Motivated by the need for more robust spectral representations, this work investigates the use of multiple spectral transformations aimed at extracting informative oscillations and signature features from hyperspectral reflectance data. Instead of relying solely on raw spectra, we evaluate a set of preprocessing approaches designed to emphasize frequency components, local variations, and multi-scale patterns embedded in the signal. The evaluated transformations include first- and second-order derivatives (D1, D2), Discrete Fourier Transform (DFT), Discrete Wavelet Transform (DWT), Continuous Wavelet Transform (CWT), Discrete Cosine Transform (DCT), and Short-Time Fourier Transform (STFT). In this work, hyperspectral reflectance measurements were acquired over six days, collected from leaves of soybean, corn, and cotton plants cultivated under controlled experimental conditions. For each crop species, infestation trials were conducted independently for caterpillars and stink bugs, ensuring isolated and well-defined stress scenarios. The raw spectra were first filtered to remove bands affected by atmospheric

The authors are solely responsible for the content of this paper, which is not a refereed publication. Citation of this work should state that it is from the Proceedings of the 17th International Conference on Precision Agriculture and 11th Brazilian Congress on Precision and Digital Agriculture. EXAMPLE: Last Name, A. B. & Coauthor, C. D. (2026). Title of paper. In Proceedings of the 17th International Conference on Precision Agriculture and 11th Brazilian Congress on Precision and Digital Agriculture (unpaginated, online). Monticello, IL: International Society of Precision Agriculture and Brazilian Congress on Precision and Digital Agriculture.

water absorption and sensor-related instability. Later, the transformations were computed from the filtered reflectance curves to enhance local spectral variations and absorption features. All experiments were performed separately for each crop–pest combination, and the performance metrics were aggregated to evaluate the robustness of the extracted spectral features across different host species. Experimental results demonstrate that transformations capable of capturing localized or multi-scale spectral oscillations consistently enhance pest detection performance relative to raw reflectance. For caterpillar detection, CWT, DWT, STFT, D1, and D2 achieved average AUROC increases ranging from approximately +14% to +17%, with F1-score gains between +16% and +21%, the highest improvement being observed for CWT. In contrast, global frequency transforms such as DFT and DCT resulted in only marginal changes, remaining close to raw-spectrum performance. A similar trend was observed for stink bug detection, where CWT again yielded the largest gains. The CWT, DWT, D1, D2, and STFT transforms produced AUROC improvements between +14% and +17%, with F1 gains ranging from 14% to 21%. In contrast, the DFT and DCT had a minor impact compared to the raw performance. These findings indicate that feature representations emphasizing localized spectral oscillations and multi-scale signatures are more effective for pest discrimination than raw reflectance or purely global frequency decompositions. The performance gains are consistent across the classifiers, demonstrating that the improvements stem from enhanced spectral representation rather than model-specific tuning. Therefore, this work advances hyperspectral pest monitoring by highlighting the importance of transformation-based spectral feature extraction to improve sensitivity and robustness in agricultural pest diagnosis.

Keywords.
Spectral signature, Hyperspectral data, Pest detection, Feature extraction, Machine learning.

1. Introduction

Insect infestation is one of the main factors affecting crop development and productivity, especially in large-scale agricultural systems where early visual inspection is difficult and labor-intensive (Oerke 2006; Nansen *et al*. 2021). Pests such as caterpillars and stink bugs can induce physiological and structural changes in plant tissues before severe visual symptoms become evident (Nansen *et al*. 2021). Therefore, non-destructive sensing approaches capable of detecting early spectral responses to pest-induced stress are important for improving crop monitoring and supporting precision agriculture practices.

Hyperspectral reflectance sensing is a promising tool for this purpose because it measures plant responses across narrow and contiguous wavelength bands, covering spectral regions associated with pigments, leaf structure, and water content (Knipling 1970; Blackburn 2007). These characteristics allow hyperspectral data to capture subtle changes caused by insect feeding activity (Nansen *et al*. 2021). However, hyperspectral signatures are also high-dimensional, highly correlated across neighboring bands, and affected by baseline variations related to crop species, illumination conditions, plant development, and sensor response (Zhu *et al*. 2017). As a result, using only the raw reflectance values as input features may not be sufficient to emphasize the most discriminative pest-related spectral patterns.

This limitation is illustrated in Fig. 1, which presents the average raw hyperspectral reflectance spectra for control and caterpillar-damaged samples across soybean, corn, and cotton. Each subplot corresponds to one crop, and the shaded regions represent one standard deviation around the mean reflectance curve. The strong overlap between control and damaged samples, together with the large within-class spectral variability, indicates that pest-induced changes may be partially masked in the raw reflectance domain. This observation motivates the use of spectral transformations capable of enhancing local variations, oscillatory patterns, and multi-scale structures that may be more discriminative for pest detection.

In this context, this work investigates the use of transformation-based spectral feature extraction for pest detection in hyperspectral data. Instead of relying only on raw spectra, several representations are evaluated, including first- and second-order derivatives (D1 and D2), Discrete Fourier Transform (DFT), Short-Time Fourier Transform (STFT), Discrete Wavelet Transform

(DWT), Continuous Wavelet Transform (CWT), and Discrete Cosine Transform (DCT). The transformed features are evaluated using Random Forest and XGBoost classifiers across different crop–pest detection scenarios involving soybean, corn, and cotton plants, as well as control and damaged plants by caterpillar and stink bug.

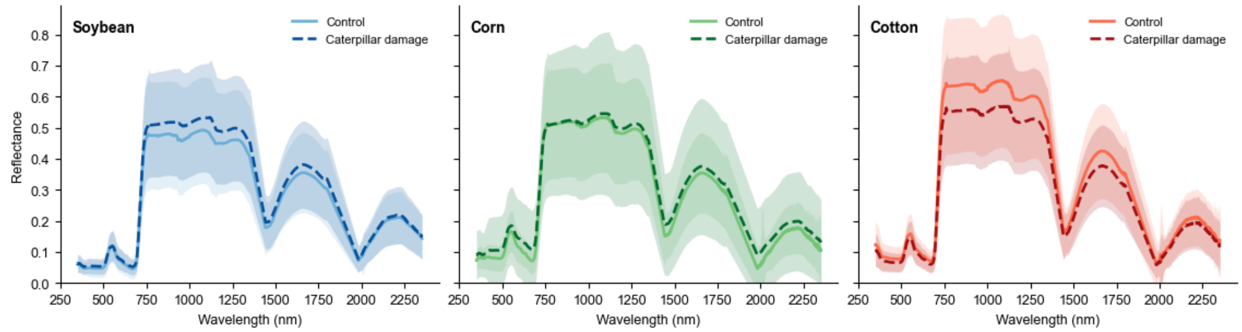

Fig. 1. Average raw hyperspectral reflectance spectra for control and caterpillar-damaged samples across soybean, corn, and cotton. Shaded regions indicate one standard deviation around the mean reflectance curve. The substantial overlap and variability between healthy and damaged spectra highlight the difficulty of relying only on raw reflectance values for pest detection and motivates the use of transformation-based spectral feature extraction.

The main contributions of this paper are as follows:

1. A systematic comparison of raw hyperspectral reflectance and multiple spectral transformations for pest detection.
2. An evaluation of local, global, and multi-scale spectral representations across different crop–pest combinations.
3. A classifier-independent analysis using Random Forest and XGBoost to verify whether performance gains are associated with the spectral transformations rather than a single machine learning model.
4. Empirical evidence that localized and multi-scale transformations, especially CWT, DWT, and STFT, are more effective than global frequency-domain representations for detecting pest-induced spectral changes in hyperspectral data.

The remainder of this paper is organized as follows. Section 2 describes the experimental setup, hyperspectral dataset, spectral transformations, classification models, and evaluation protocol. Section 3 presents and discusses the experimental results for caterpillar and stink bug detection, including the comparison between localized, multi-scale, and global transformations. Finally, Section 4 summarizes the main findings and presents directions for future work.

2. Materials and Methods

This section describes the experimental setup, hyperspectral dataset, spectral transformations, classification models, and evaluation protocol adopted in this work.

2.1 Experimental Setup

The experiments were conducted in collaboration with Embrapa Genetic Resources and Biotechnology, Brasília, DF, Brazil. Hyperspectral measurements were acquired from soybean (*Glycine max*), corn (*Zea mays*), and cotton (*Gossypium hirsutum*) plants cultivated under controlled greenhouse conditions. The crops were maintained under suitable environmental conditions to ensure homogeneous development prior to the infestation experiments.

Two insect pests with high economic relevance in Brazilian agriculture were evaluated: the fall armyworm caterpillar (*Spodoptera frugiperda*) and the green-belly stink bug (*Dichelops melacanthus*) (Barros *et al*. 2010; Overton *et al*. 2021; Sosa-Gómez *et al*. 2020). The insects were produced under controlled laboratory conditions and subsequently introduced into separate groups of plants to simulate pest infestation. Healthy plants maintained under identical cultivation conditions were used as control.

The infestation experiments were monitored over six acquisition days. During this period, hyperspectral measurements were collected repeatedly from both healthy and damaged plants, enabling the observation of spectral variations associated with pest-induced physiological stress.

Spectral measurements were acquired under natural illumination using an ASD FieldSpec 3 portable spectroradiometer (Analytical Spectral Devices Inc., Boulder, USA). Prior to data acquisition, the instrument was calibrated using a white reference panel. Measurements were performed with the sensor positioned approximately 50 cm above the plant canopy using an 8° field-of-view lens. The spectroradiometer operates within the 350–2460 nm spectral range, covering portions of the ultraviolet (UV), visible (VIS), near-infrared (NIR), and shortwave infrared (SWIR) regions.

After spectral acquisition, the reflectance values $\rho(\lambda)$ at wavelength λ were computed from the ratio between the radiance reflected by the target surface and the radiance measured from a calibrated white reference panel, according to

$$\rho(\lambda) = \kappa \frac{\phi_T(\lambda)}{\phi_R(\lambda)}, \tag{1}$$

where $\phi_T(\lambda)$ denotes the radiance reflected by the target plant surface, $\phi_R(\lambda)$ represents the radiance measured from the reference white panel, and κ is a calibration factor.

2.2 Hyperspectral Dataset

The resulting dataset contains 3,958 hyperspectral reflectance samples acquired from soybean, corn, and cotton plants subjected to different pest infestation conditions. Each sample corresponds to a single hyperspectral reflectance signature associated with a specific crop species and infestation label.

Table 1 summarizes the distribution of samples according to crop species and infestation condition. Soybean and corn experiments included both caterpillar and stink bug infestations, while cotton experiments were conducted exclusively with caterpillar infestations. In total, the dataset comprises 1,524 control samples, 1,428 caterpillar-damaged samples, and 1,006 stink bug-damaged samples.

To evaluate the generalization capability of the investigated spectral transformations, each crop–pest combination was treated as an independent binary classification problem. Consequently, five classification scenarios were considered throughout the experiments: soybean–caterpillar, soybean–stink bug, corn–caterpillar, corn–stink bug, and cotton–caterpillar.

Table 1. Distribution of hyperspectral samples according to crop species and infestation.

Crop	Control	Caterpillar	Stink Bug	Total
Soybean	485	428	656	1,569
Corn	474	505	350	1,329
Cotton	565	495	-	1,060
Total	1,524	1,428	1,006	3,958

The hyperspectral signatures capture detailed information regarding leaf biochemical and structural properties over the UV, VIS, NIR, and SWIR spectral regions (Blackburn 2007; Homolová *et al*. 2013). Since insect feeding alters pigment concentration, tissue structure, moisture content, and plant metabolism, these spectral measurements provide a rich source of information for detecting pest-induced stress and evaluating the effectiveness of different spectral representations (Nansen *et al*. 2021).

2.3 Spectral Transformations

Hyperspectral reflectance signatures contain information related to plant physiology, biochemical composition, and structural characteristics. However, pest-induced stress often manifests as subtle local variations and oscillatory patterns distributed across different wavelength scales, which may not be readily observable in the original reflectance domain. Although the raw spectrum preserves the measured reflectance magnitude at each wavelength, important discriminatory information can remain hidden due to spectral redundancy and baseline variations.

Motivated by the hypothesis that insect infestation alters not only the absolute reflectance values but also the underlying spectral signatures and oscillatory structures of the signal, this work investigates several spectral transformations capable of emphasizing local variations, frequency components, and multi-scale patterns. These transformations provide alternative representations of the same hyperspectral measurement, potentially revealing spectral characteristics associated with pest-induced physiological changes that are less evident in the original reflectance domain.

Let the signal $\mathbf{x} = [x_1, x_2, \cdots, x_N]$ denote a hyperspectral reflectance signature, where x_i represents the reflectance value measured at the i-th wavelength and N denotes the total number of spectral bands. The investigated spectral transformations can be viewed as operators that map an original reflectance signature into an alternative feature representation. Fig. 2 shows the feature extraction as a mapping formulation by a transformation $\mathcal{T}(\cdot)$ to obtain a feature vector $\mathbf{y}$.

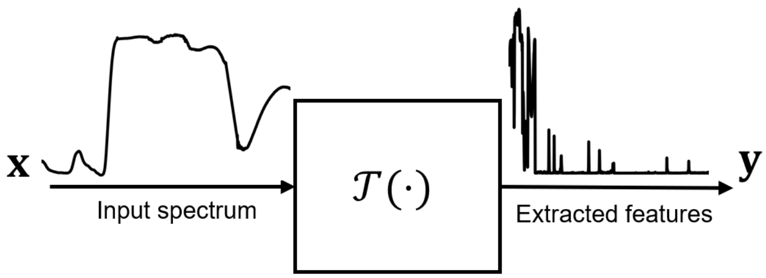

Fig. 2. Generic framework adopted for spectral feature extraction. A hyperspectral reflectance signature $\mathbf{x}$ is transformed by an operator $\mathcal{T}(\cdot)$, generating a new feature representation $\mathbf{y}$.

In Fig. 2, depending on the selected transformation, $\mathbf{y}$ may emphasize local spectral variations, oscillatory behavior, frequency-domain characteristics, or multi-scale. Also, if no transformation is applied to the input spectrum, it corresponds to the raw reflectance spectrum, i.e., the original hyperspectral signature. The following subsections describe each investigated transformation and its corresponding mathematical formulation.

2.3.1 First and Second Derivative

The first- and second-order derivatives are commonly used in hyperspectral analysis to enhance absorption features and emphasize local spectral variations. The first derivative (D1) captures the rate of change of the reflectance signal with respect to wavelength, while the second derivative (D2) measures the variation of the first derivative, that is, the spectral curvature. Both derivatives are defined from mathematics by:

$$\mathbf{y}_{\mathrm{D1}} = \frac{d\mathbf{x}}{d\lambda}, \quad \mathbf{y}_{\mathrm{D2}} = \frac{d^2\mathbf{x}}{d\lambda^2}, \tag{2}$$

where $\mathbf{y}_{\mathrm{D1}}$ and $\mathbf{y}_{\mathrm{D2}}$ correspond to the first and second-order derivatives, respectively.

By emphasizing local spectral variations rather than absolute reflectance values, D1 and D2 can reveal subtle modifications in the spectral signature with pest-induced stress. In this work, both derivatives were estimated using the *Savitzky-Golay* filter (Savitzky and Golay 1964) implemented in the *SciPy* library (Virtanen *et al*. 2020).

2.3.2 Fourier Transform

The Fourier Transform decomposes a signal into a set of sinusoidal components with different frequencies, providing a frequency-domain representation of the original spectrum (Oppenheim and Schafer 2010). In the context of hyperspectral analysis, this transformation can reveal global oscillatory patterns that may not be directly observable in the reflectance domain.

Given a discrete spectral signature $\mathbf{x}$, the Discrete Fourier Transform (DFT) is defined as

$$\mathrm{X}(k) = \sum_{n=0}^{N-1} \mathrm{x}(n) e^{-j2\pi\frac{kn}{N}}, \tag{3}$$

where $\mathrm{x}(n)$ denotes the reflectance value at the n-th spectral band, N is the total number of bands, k is the frequency index, and $j = \sqrt{-1}$. The transformed signal $\mathrm{X}(k)$ represents the contribution of each frequency component to the original spectrum.

In this work, the DFT was computed using the Fast Fourier Transform (FFT) (Nussbaumer 1981), which is an efficient algorithm with low computational complexity. The feature representation was constructed from the magnitude spectrum,

$$\mathrm{y}_{\mathrm{DFT}}(k) = |\mathrm{X}(\mathrm{k})|, \tag{4}$$

which describes the intensity of the frequency components ignoring their phase.

By transforming the spectral signature into the frequency domain, the DFT captures the global frequency content of the spectrum, emphasizing oscillatory patterns distributed across the entire wavelength range rather than local variations occurring within specific spectral regions.

2.3.3 Short-Time Fourier Transform

The Short-Time Fourier Transform (STFT) was used to extract localized frequency information (Gröchenig 2001) from the hyperspectral signature. While the Fourier Transform represents the spectrum in terms of global frequency components, it does not indicate where these components occur along the wavelength λ axis. This limitation is relevant for hyperspectral pest detection because insect-induced changes may appear in specific absorption regions rather than across the entire spectrum.

To preserve this local wavelength information, the STFT applies the Fourier Transform over short overlapping windows of the spectral signal, and it is defined as

$$\mathrm{S}(m,k) = \sum_{n=0}^{L-1} \mathrm{x}(n+mH)\, w(n) e^{-j2\pi\frac{kn}{L}}, \tag{5}$$

where $S(m,k)$ is the STFT coefficient at window position m and frequency index k, L is the window length, H is the hop size between consecutive windows, and $w(n)$ is the window function. In this work, the STFT was computed using window length $L = 64$, hop size $H = 16$, and a Hann window function. Then, for each window, the power spectrum was computed as

$$P(m,k) = |S(m,k)|^2. \tag{6}$$

After that, the feature vector constructed from the STFT was obtained by averaging the power coefficients across all window positions and applying a logarithmic transformation:

$$\mathrm{y}_{\mathrm{STFT}}(k) = \ln\left(1 + \frac{1}{M}\sum_{m=0}^{M-1} P(m,k)\right), \tag{7}$$

where M is the number of spectral windows.

Therefore, unlike the DFT, which summarizes oscillatory behavior over the entire spectral range, the STFT preserves information about both the frequency content and its approximate location in the wavelength domain. This makes it more suitable for detecting localized spectral oscillations and absorption-related variations associated with pest-induced stress.

2.3.4 Discrete Wavelet Transform

The Discrete Wavelet Transform (DWT) decomposes the spectral signature into components associated with different frequency ranges and resolution levels (Mallat 1989). The DWT also provides a localized multi-resolution representation, making it suitable for capturing spectral oscillations and absorption-related variations occurring at different wavelength scales.

The DWT computes approximation and detail coefficients through successive low-pass and high-pass filtering operations, followed by downsampling. At decomposition level j, these coefficients can be expressed as:

$$a_j(k) = \sum_n h(n-2k)a_{j-1}(n), \tag{8}$$

$$d_j(k) = \sum_n g(n-2k)a_{j-1}(n), \tag{9}$$

where $a_j(k)$ and $d_j(k)$ are the approximation and detail coefficients at level j, respectively, $h(\cdot)$ is the low-pass filter, $g(\cdot)$ is the high-pass filter, and $a_0 = \mathbf{x}$. In this work, the DWT was computed using the Daubechies-4 wavelet with symmetric boundary extension (Daubechies 1992). Furthermore, the maximum valid decomposition level was used, and both approximation and detail coefficients were retained to build the feature vector.

The final DWT feature vector was obtained by concatenating the coefficients from all decomposition levels $\mathbf{y}_{\mathrm{DWT}} = [a_J,\ d_J, d_{J-1}, \cdots, d_1]$, where J is the maximum decomposition level. This representation preserves multi-scale spectral information and can emphasize localized oscillations in plant reflectance.

2.3.5 Continuous Wavelet Transform

The Continuous Wavelet Transform (CWT) can also provide a localized multi-scale representation of the spectral signature by comparing the signal with shifted and scaled versions of a mother wavelet (Grossmann 1990). In contrast to the DWT, which decomposes the signal into discrete resolution levels, the CWT evaluates the spectral response over a continuous set of scales, allowing a detailed characterization of localized oscillations and absorption-related patterns.

The CWT coefficient at scale a and spectral position b is defined as:

$$W_{\mathrm{x}}(a,b) = \frac{1}{\sqrt{a}} \sum_{n=0}^{N-1} \mathrm{x}(n)\psi^{*\left(\frac{n-b}{a}\right)}, \tag{10}$$

where $W_{\mathrm{x}}(a,b)$ is the wavelet coefficient, $\psi(\cdot)$ is the mother wavelet, $\psi^*(\cdot)$ its complex conjugate, and N is the number of spectral bands. In this work, the CWT was computed using the Morlet wavelet over 32 scales.

Finally, the feature representation was obtained by computing the average energy of the wavelet coefficients at each scale:

$$\mathrm{y}_{\mathrm{CWT}}(a) = \frac{1}{N} \sum_{b=0}^{N-1} |W_{\mathrm{x}}(a,b)|^2 . \tag{11}$$

Therefore, the final CWT feature vector is given by $\mathbf{y}_{\mathrm{CWT}} = [\mathrm{y}_{\mathrm{CWT}}(a_1), \mathrm{y}_{\mathrm{CWT}}(a_2), \cdots, \mathrm{y}_{\mathrm{CWT}}(a_A)]$, where A is the number of evaluated scales. This representation summarizes the multi-scale energy distribution of the spectral signal and can also capture localized spectral variations.

2.3.6 Discrete Cosine Transform

As an alternative global spectral representation, the Discrete Cosine Transform (DCT) expresses the reflectance signature as a combination of cosine basis functions (Oppenheim and Schafer 2010). This transformation is useful for describing the overall oscillatory structure of the signal using real-valued coefficients.

The DCT coefficient at frequency index k is defined as:

$$\mathrm{C}(k) = \alpha(k) \sum_{n=0}^{N-1} \mathrm{x}(n) \cos\left[\frac{\pi}{N}\left(n + \frac{1}{2}\right) k\right], \tag{12}$$

where $\alpha(k)$ is the normalization term given by:

$$\alpha(k) = \begin{cases} \sqrt{\frac{1}{N}}, & k = 0, \\ \sqrt{\frac{2}{N}}, & k > 0. \end{cases} \tag{13}$$

Then, the final DCT-based feature representation was simply obtained from the DCT coefficients, that is, $\mathrm{y}_{\mathrm{DCT}}(k) = \mathrm{C}(k)$.

Although the DCT can compactly represent broad oscillatory behavior in the hyperspectral signature, it does not explicitly preserve the wavelength location of localized absorption features. Therefore, similarly to the DFT, it mainly describes global spectral frequency content rather than local variations.

2.4 Classification Models

As part of the evaluation process, two machine learning classifiers were used to evaluate the discriminative capability of each spectral representation: Random Forest (RF) and Extreme Gradient Boosting (XGBoost). These models were selected once they are widely used for classification tasks and require limited feature engineering.

Random Forest (Breiman 2001) is an ensemble learning method based on multiple decision trees trained using bootstrap samples of the data and random subsets of features. The final prediction is obtained by aggregating the predictions of the individual trees. This approach is robust to noise, handles high-dimensional feature spaces, and reduces overfitting compared with a single decision tree. However, RF may require a large number of trees to capture complex decisions.

The other selected classifier, XGBoost (Chen and Guestrin 2016), is a gradient boosting algorithm that builds an ensemble of decision trees sequentially, where each new tree is trained to correct the errors of the previous trees. It is effective for structured datasets and includes regularization terms that help control overfitting. XGBoost often provides strong predictive performance, but it is more sensitive to hyperparameter choices than RF.

In this work, both classifiers were applied independently to each spectral representation, including raw spectra and all transformed feature sets. By evaluating both RF and XGBoost, the experiments aimed to reduce the dependence of the analysis on a single classifier. Thus, when performance gains were observed for both models, they were interpreted as evidence that the improvements were primarily related to the spectral transformation rather than to model-specific behavior.

Both classifiers were implemented using standard machine learning libraries. The RF model was employed from the *scikit-learn* library, using 150 trees with maximum depth of 6. The traditional XGBoost Python library was used for the XGBoost classifier, using 150 estimators and maximum depth equals to 4, subsampling ratio of 0.8, and the binary logistic objective function.

2.5 Evaluation Protocol

Each crop–pest combination was evaluated as an independent binary classification problem, in which pest-damaged samples were assigned to the positive class and control samples were assigned to the negative class. The same evaluation protocol was applied to all spectral representations, including raw spectra and the transformed feature vectors obtained from D1, D2, DFT, STFT, DWT, CWT, and DCT.

For each experiment, the dataset was partitioned into training (70%) and test (30%) subsets using stratified random sampling, preserving the proportion of control and damaged samples in both subsets. In all experiments, the random seed was fixed to ensure reproducibility. With that, the same data split was used for all preprocessing methods and classifiers within each crop–pest scenario.

The classification performance was evaluated using precision, recall, $F1$-score, and area under the receiver operating characteristic curve (AUROC) (Bishop 2006). These metrics were computed from the confusion matrix, using true positives (T_P), true negatives (T_N), false positives (F_P), and false negatives (F_N).

The precision (P) measures the proportion of samples predicted as damaged that were correctly classified, while the recall (R) measures the proportion of damaged samples that were correctly detected. They are defined as:

$$P = \frac{T_P}{T_P + F_P}, \quad R = \frac{T_P}{T_P + F_N}. \tag{14}$$

The $F1$-score is the harmonic mean between precision and recall:

$$F1 = 2\frac{P \cdot R}{P + R} = \frac{2T_P}{2T_P + F_P + F_N}. \tag{15}$$

The AUROC was used as a threshold-independent metric to evaluate the ability of each model to distinguish between control and pest-damaged samples. It is computed as the area under the ROC curve, which relates the true positive rate (T_{PR}) to the false positive rate (F_{PR}),

$$T_{PR} = \frac{T_P}{T_P + F_N}, \quad T_{PR} = \frac{F_P}{F_P + T_N}, \tag{16}$$

across different classification thresholds. AUROC values close to 1 indicate strong class separability, whereas values close to 0.5 indicate performance similar to random classification.

3. Experimental Results and Discussion

The experimental results were analyzed separately for caterpillar and stink bug detection. For each pest, the performance values were averaged across the corresponding crop experiments, classifying between pest-damaged and control samples. Furthermore, performance gains for each transformation were computed as the difference relative to the raw spectra, averaged over the two classifiers.

3.1 Caterpillar Detection

Initially for the caterpillar, experiments were conducted specifically for this pest over the two classifiers and all transformations. Table 2 presents the average results over the crops for caterpillar detection using Random Forest and XGBoost. Note that the raw spectra provided a baseline performance, with AUROC values of 80.2% for RF and 82.7% for XGBoost.

As seen in Table 2, the best results were obtained with transformations that preserve local or multi-scale spectral information. CWT achieved the highest performance for both classifiers, reaching 97.1% AUROC and 89.9% $F1$-score with RF, and 99.4% AUROC and 96.3% $F1$-score with XGBoost. Besides, DWT and STFT also achieved strong results, with AUROC values above 95% for RF and above 98% for XGBoost. The derivative-based methods, D1 and D2, also substantially improved the results compared with the raw baseline, confirming the importance of local spectral transitions and absorption-related variations for caterpillar detection.

Table 2. Average caterpillar detection performance across crop species for Random Forest and XGBoost. Best values are highlighted in bold.

Model	Transform	P	R	$F1$	AUROC
RF	Raw	71.4	70.3	70.0	80.2
	D1	86.7	85.6	85.6	94.2
	D2	86.3	84.8	84.8	93.6
	DFT	71.8	71.2	71.1	79.8
	DWT	89.0	88.0	88.1	95.7
	CWT	**90.7**	**89.8**	**89.9**	**97.1**
	DCT	73.9	73.3	73.2	80.6
	STFT	89.2	88.7	88.7	96.4
XGBoost	Raw	74.5	74.0	73.9	82.7
	D1	92.3	91.9	91.9	97.9
	D2	92.7	92.2	92.3	98.2
	DFT	75.3	75.1	75.1	83.6
	DWT	93.8	93.6	93.7	98.8
	CWT	**96.4**	**96.3**	**96.3**	**99.4**
	DCT	75.4	75.0	74.9	83.4
	STFT	93.6	93.4	93.4	98.5

For a better analysis of the performance, the gains of each transformation relative to raw spectra are summarized in Table 3. CWT produced the highest average improvement, with gains of +20.6 in precision, +20.9 in recall, +21.1 in $F1$-score, and +16.8 in AUROC. DWT and STFT also showed high gains, with $F1$-score improvements of +18.9 and +19.1, respectively. In contrast, DFT and DCT produced only marginal improvements, remaining close to the raw baseline. This suggests that transformations that capture global frequency representations were less effective than localized or multi-scale transformations for detecting caterpillar-induced spectral changes.

Table 3. Average caterpillar detection gain relative to raw spectra. Gains were computed as percentage-point differences and averaged over Random Forest and XGBoost. The highest gains are highlighted in bold.

Transform	P	R	$F1$	AUROC
D1	+16.6	+16.5	+16.8	+14.6
D2	+16.5	+16.3	+16.6	+14.5
DFT	+0.6	+0.9	+1.1	+0.2
DWT	+18.5	+18.6	+18.9	+15.8
CWT	**+20.6**	**+20.9**	**+21.1**	**+16.8**
DCT	+1.7	+1.9	+2.1	+0.5
STFT	+18.4	+18.9	+19.1	+16.0

3.2 Stink Bug Detection

Similar to the caterpillar experiments, Table 4 shows the average results for stink bug detection. It can be seen from Table 4 that the transformed representations substantially improved the detection performance, especially those based on local or multi-scale spectral information. CWT again achieved the best overall performance, reaching 99.1% AUROC and 93.0% $F1$-score with RF, and 99.8% AUROC and 96.9% $F1$-score with XGBoost. Furthermore, D1, D2, DWT, and STFT also produced strong results, generally exceeding 95% AUROC. These results indicate that stink bug damage also generates spectral modifications that are better represented when local variations and multi-scale oscillations are emphasized by the transformations.

Following the same, now for the stink bug, the average gains relative to raw spectra are presented in Table 5. CWT again produced the largest improvement, with gains of +20.4 in precision, +20.7 in recall, +20.8 in $F1$-score, and +16.9 in AUROC. D1, D2, DWT, and STFT also produced consistent improvements, with AUROC gains ranging from +14.1 to +15.4. In contrast, DFT and DCT produced small gains variations for all metrics relative to raw experiments, showing sometimes negative gains. This confirms that global transforms like DFT and DCT did not capture the most relevant spectral changes associated with stink bug infestation.

Table 4. Average stink bug detection performance across crop species for Random Forest and XGBoost. Best metric values are highlighted in bold.

Model	Transform	P	R	$F1$	AUROC
RF	Raw	73.6	72.2	72.5	81.7
	D1	89.5	87.1	87.8	95.4
	D2	87.9	85.9	86.6	96.4
	DFT	70.9	68.2	68.5	78.4
	DWT	89.3	84.2	85.3	97.3
	CWT	**94.3**	**92.4**	**93.0**	**99.1**
	DCT	75.7	70.3	70.2	80.4
	STFT	87.7	83.0	84.0	95.6
XGBoost	Raw	77.1	75.5	75.9	83.5
	D1	91.7	91.7	91.7	97.9
	D2	93.7	94.0	93.8	98.5
	DFT	77.3	75.5	76.0	84.5
	DWT	92.8	92.3	92.5	98.6
	CWT	**97.2**	**96.7**	**96.9**	**99.8**
	DCT	77.9	76.0	76.5	84.5
	STFT	92.9	91.9	92.3	98.3

Table 5. Average stink bug detection gain relative to raw spectra. Gains were computed as percentage-point differences and averaged over Random Forest and XGBoost. Highest gains are highlighted in bold.

Transform	P	R	$F1$	AUROC
D1	+15.2	+15.6	+15.6	+14.1
D2	+15.4	+16.1	+16.0	+14.9
DFT	-1.2	-2.0	-2.0	-1.1
DWT	+15.7	+14.4	+14.7	+15.4
CWT	**+20.4**	**+20.7**	**+20.8**	**+16.9**
DCT	+1.4	-0.7	-0.9	-0.1
STFT	+14.9	+13.6	+13.9	+14.4

3.3 Closing Remarks

To summarize the results presented in Tables 2 to 5, the experimental results indicate a clear difference between localized or multi-scale transformations and global spectral transformations. In other words, localized and multi-scale transformations were more effective than global transformations for hyperspectral pest detection. D1, D2, STFT, DWT, and CWT preserve local spectral transitions, windowed oscillations, or multi-scale structures, while DFT and DCT mainly describe the global frequency content of the entire spectrum. This distinction explains why DFT and DCT produced only marginal or negative gains, whereas CWT, DWT, and STFT consistently improved performance for both pests. Considering the average gains across precision, recall, $F1$-score, and AUROC, the three best transformations were, in order, **CWT**, **DWT**, and **STFT**, indicating that pest-induced stress is better represented by localized spectral signatures than by global frequency patterns.

4. Conclusion

Raw hyperspectral reflectance signatures contain relevant information for pest detection. However, their high dimensionality, spectral redundancy, and baseline variability can mask subtle patterns associated with insect-induced stress. For this reason, extracting more informative spectral representations is an important step for improving hyperspectral pest diagnosis. In this work, experimental results show that transformations designed to emphasize local variations and multi-scale oscillatory structures are more effective than using raw reflectance directly or relying only on global frequency representations.

The best overall transformation was CWT, which produced the highest gains relative to the raw spectra for both pests. Specifically, CWT improved the $F1$-score by +21.1 and +20.8 percentage points, and the AUROC by +16.8 and +16.9 percentage points, for caterpillar and stink bug detection, respectively. DWT and STFT were the next most effective methods, also showing consistent improvements across the evaluated metrics. Therefore, based on the experiments, the three most effective transformations for hyperspectral pest detection were, in order, CWT, DWT, and STFT.

These results indicate that both caterpillar and stink bug detection benefit from spectral representations that preserve localized and multi-scale information. The consistency of the gains across RF and XGBoost also suggests that the improvements were mainly associated with the enhanced spectral representations rather than with model-specific behavior. In contrast, DFT and DCT produced only marginal gains for caterpillar and stink bug detection, indicating that global frequency-domain representations were less effective for capturing pest-related spectral changes in this dataset.

The findings highlight the importance of transformation-based spectral feature extraction for improving the robustness of hyperspectral pest monitoring. In particular, localized and multi-scale transformations were more suitable for representing pest-induced stress, likely because the relevant spectral changes occur in specific wavelength regions rather than uniformly across the entire spectrum.

Future work should evaluate the selected transformations under field conditions, with larger datasets and additional crop–pest combinations. Further studies should also investigate temporal modeling strategies to explicitly represent the evolution of pest-induced spectral changes over time. With larger data availability, deep learning approaches, including recurrent neural networks (Elman 1990; Hochreiter and Schmidhuber 1997) and Transformer-based architectures (Vaswani *et al*. 2017), could be explored for automatic spectral-temporal feature learning.

Acknowledgments

This work was supported in part by the São Paulo Research Foundation (FAPESP) under Grants #2019/07665-4, #2024/15430-5, #2025/10840-3, and #2025/08971-2, by the Brazilian National Council for Scientific and Technological Development (CNPq) under Grants #406417/2022-9, #403307/2025-2 and #306788/2023-3, and by the National Natural Science Foundation of China (NSFC) under Grant #W2412112.

References

Barros, E. M., Torres, J. B., Ruberson, J. R., & Oliveira, M. D. (2010). Development of Spodoptera frugiperda on different hosts and damage to reproductive structures in cotton. *Entomologia Experimentalis et Applicata*, *137*(3), 237-245.

Breiman, L. (2001). Random forests. *Machine Learning*, *45*(1), 5-32.

Blackburn, G. A. (2007). Hyperspectral remote sensing of plant pigments. *Journal of experimental botany*, *58*(4), 855-867.

Bishop, C. M. (2006). *Pattern Recognition and Machine Learning*. New York, NY, USA: Springer.

Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 785–794). New York, NY, USA: ACM.

Daubechies, I. (1992). *Ten lectures on wavelets*. Society for industrial and applied mathematics.

Elman, J. L. (1990). Finding structure in time. *Cognitive science*, *14*(2), 179-211.

Gröchenig, K. (2001). *Foundations of time-frequency analysis* (Vol. 359). Boston: Birkhäuser.

Grossmann, A., Kronland-Martinet, R., & Morlet, J. (1990). Reading and understanding continuous wavelet transforms. In *Wavelets: Time-Frequency Methods and Phase Space Proceedings of the International Conference, Marseille, France, December 14–18, 1987* (pp. 2-20). Berlin, Heidelberg: Springer Berlin Heidelberg.

Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural computation*, *9*(8), 1735-1780.

Homolová, Lucie, et al. "Review of optical-based remote sensing for plant trait mapping." *Ecological Complexity* 15 (2013): 1-16.

Knipling, E. B. (1970). Physical and physiological basis for the reflectance of visible and near-infrared radiation from vegetation. *Remote sensing of environment*, *1*(3), 155-159.

Mallat, S. G. (1989). A theory for multiresolution signal decomposition: the wavelet representation. *IEEE transactions on pattern analysis and machine intelligence*, *11*(7), 674-693.

Nansen, C., Murdock, M., Purington, R., & Marshall, S. (2021). Early infestations by arthropod pests induce unique changes in plant compositional traits and leaf reflectance. *Pest Management Science*, *77*(11), 5158-5169.

Nussbaumer, H. J. (1981). The fast Fourier transform. In *Fast Fourier transform and convolution algorithms* (pp. 80-111). Berlin, Heidelberg: Springer Berlin Heidelberg.

Oerke, E. C. (2006). Crop losses to pests. *The Journal of Agricultural Science*, *144*(1), 31-43.

Oppenheim, A.V. and Schafer, R.W. (2010) Discrete-Time Signal Processing. Pearson Higher Education, Inc., Upper Saddle River.

Overton, K., Maino, J. L., Day, R., Umina, P. A., Bett, B., Carnovale, D., ... & Reynolds, O. L. (2021). Global crop impacts, yield losses and action thresholds for fall armyworm (Spodoptera frugiperda): A review. *Crop Protection*, *145*, 105641.

Savitzky, A., & Golay, M. J. (1964). Smoothing and differentiation of data by simplified least squares procedures. *Analytical chemistry*, *36*(8), 1627-1639.

Sosa-Gómez, D. R., Corrêa-Ferreira, B. S., Kraemer, B., Pasini, A., Husch, P. E., Delfino Vieira, C. E., ... & Negrao Lopes, I. O. (2020). Prevalence, damage, management and insecticide resistance of stink bug populations (Hemiptera: Pentatomidae) in commodity crops. *Agricultural and Forest entomology*, *22*(2), 99-118.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems*, *30*.

Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., ... & Van Mulbregt, P. (2020). SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nature methods*, *17*(3), 261-272.

Zhu, H., Chu, B., Zhang, C., Liu, F., Jiang, L., & He, Y. (2017). Hyperspectral imaging for presymptomatic detection of tobacco disease with successive projections algorithm and machine-learning classifiers. *Scientific reports*, *7*(1), 4125.