

Dynamic Sparse Partial Least Squares for Sensor Fault Detection and Diagnosis in Smart Agriculture

fouzi Harrou

King Abdullah University of Science and Technology

Muddu Madakyaru

`muddu.m@manipal.edu`

Manipal Academy of Higher Education

K. Ramakrishna Kini

Manipal Academy of Higher Education

Noreddine Ghaffour

King Abdullah University of Science and Technology

Ying Sun

King Abdullah University of Science and Technology

Article

Keywords: Dynamic sparse partial least squares, Sensor fault detection, Semi-supervised learning, Smart agriculture, Fault diagnosis

Posted Date: September 8th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10544232/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

Dynamic Sparse Partial Least Squares for Sensor Fault Detection and Diagnosis in Smart Agriculture

Fouzi Harrou^{1,2*}, Muddu Madakyaru^{3*}, K. Ramakrishna Kini³,
Noreddine Ghaffour², Ying Sun¹

¹Computer, Electrical and Mathematical Sciences and Engineering (CEMSE) Division, King Abdullah University of Science and Technology (KAUST), Thuwal, 23955-6900, Saudi Arabia.

^{2*}Environmental Science and Engineering Program, Biological and Environmental Science and Engineering (BESE) Division, King Abdullah University of Science and Technology (KAUST), Thuwal, 23955-6900, Saudi Arabia.

³Manipal Institute of Technology, Manipal Academy of Higher Education, Manipal, India.

*Corresponding author(s). E-mail(s): fouzi.harrou@kaust.edu.sa;
muddu.m@manipal.edu;

Contributing authors: ramakrishna.kini@manipal.edu;
noreddine.ghaffour@kaust.edu.sa; ying.sun@kaust.edu.sa;

Abstract

Reliable sensor measurements are critical for precision agriculture, where sensor faults can compromise monitoring and decision-making. This paper proposes a Dynamic Sparse Partial Least Squares (D-sPLS) framework for semi-supervised sensor fault detection and diagnosis. The proposed framework is trained exclusively using fault-free data, requiring neither labeled fault samples nor anomaly annotations. It combines dynamic feature augmentation with sparse latent-variable modeling to capture temporal dependencies while selecting the most informative predictors. Residuals generated under normal operating conditions are monitored using an Exponentially Weighted Moving Average (EWMA) statistic with a nonparametric Kernel Density Estimation (KDE)-based threshold. To enhance interpretability, an Individual Conditional Expectation (ICE)-based diagnostic layer is incorporated to explain detected anomalies and visualize fault severity. The framework is validated on two real agricultural datasets involving smart soil monitoring and soil quality assessment under multiple bias and

drift sensor fault scenarios. Experimental results are compared with representative latent-variable, regression-based, and statistical monitoring methods. The proposed D-sPLS framework consistently achieves high detection accuracy, short detection delays, and negligible false-alarm rates. Statistical significance is confirmed using bootstrap confidence intervals and the Wilcoxon signed-rank test, while computational analysis demonstrates its suitability for real-time monitoring. These results demonstrate that the proposed framework is a promising solution for accurate, robust, and interpretable sensor fault detection in smart agricultural systems.

Keywords: Dynamic sparse partial least squares, Sensor fault detection, Semi-supervised learning, Smart agriculture, Fault diagnosis.

1 Introduction

Smart agriculture increasingly relies on Internet-of-Things (IoT) technologies to continuously monitor soil and environmental conditions for precision irrigation, fertilization, and crop management [1–3]. Measurements collected from distributed sensing platforms, including soil moisture, air temperature, humidity, salinity, nutrient concentration, and pH, provide valuable information for improving agricultural productivity and resource efficiency [1, 3]. The development of soil sensing technologies has evolved from early studies on physicochemical soil properties, such as pH, salinity, nitrate, and organic matter, for site-specific crop management [4] to mobile and on-the-go sensing systems that enabled rapid and cost-effective soil characterization [5]. More recently, IoT-enabled smart farming platforms have facilitated dense sensor deployments capable of continuously collecting environmental and soil measurements in real time [6–8]. The reliability of these measurements is critical because faulty sensors or abnormal operating conditions may lead to inaccurate agronomic decisions, excessive water consumption, nutrient imbalance, and reduced crop yields [9, 10]. Despite significant advances in sensing infrastructure, most existing studies have focused on prediction, irrigation optimization, and communication architectures, whereas reliable multivariate fault detection for smart soil monitoring has received comparatively less attention [8, 11].

Detecting anomalies in agricultural systems presents unique challenges compared with traditional industrial processes. Environmental measurements are strongly coupled through weather conditions, irrigation practices, soil physics, and plant uptake mechanisms, leading to temporal dependencies, seasonal variability, measurement noise, missing observations, and non-Gaussian behaviors [12, 13]. Recent advances in machine learning have significantly improved anomaly detection capabilities by modeling complex nonlinear relationships among sensor measurements, enabling the identification of abnormal operating conditions that are difficult to capture using conventional model-based approaches [9, 14]. Supervised learning methods, including neural networks, support vector machines, decision trees, and ensemble learning techniques, have demonstrated promising detection performance in various monitoring

applications. However, these approaches require representative labeled fault data for training, which are often unavailable in practical smart agriculture systems because sensor faults occur infrequently, exhibit diverse characteristics, and are difficult to annotate. In practice, labeled fault samples are scarce, whereas normal operating data are abundant, making unsupervised and semi-supervised monitoring strategies particularly attractive for real-world deployments [15, 16]. By learning only the normal operating behavior, these approaches are naturally suited to detecting previously unseen or evolving faults while reducing the dependence on costly fault annotation. Robust fault detection is therefore essential to identify sensor bias, gradual drift, freezing behavior, communication failures, and abnormal environmental disturbances before they compromise decision-making processes [14, 17]. To address the strong correlations among environmental variables, latent-variable techniques such as Principal Component Analysis (PCA) and Partial Least Squares (PLS) have been widely adopted for multivariate process monitoring by projecting high-dimensional measurements onto a compact latent space that captures the dominant system behavior [18, 19]. However, conventional PCA and static PLS formulations generally ignore temporal dependencies and retain contributions from all measured variables, which may reduce interpretability and detection sensitivity under dynamic operating conditions [20–22]. These limitations are particularly critical in smart agriculture, where environmental responses evolve over time, and only a subset of variables may be relevant to specific agronomic or environmental phenomena.

Recent developments in sparse and dynamic learning provide promising directions for addressing these challenges. Sparse modeling introduces embedded feature selection mechanisms that improve interpretability and mitigate the influence of redundant or noisy measurements, while dynamic representations explicitly incorporate time-lagged information to capture temporal interactions among variables [23–25]. Integrating these two concepts within a unified latent-variable framework is particularly suitable for smart soil monitoring, where rainfall, temperature, humidity, and soil moisture exhibit delayed and strongly coupled behaviors [26, 27]. Moreover, residual-based statistical monitoring combined with nonparametric threshold estimation offers enhanced robustness under non-Gaussian and nonstationary operating conditions, enabling the detection of both abrupt and incipient anomalies [28–32].

Motivated by these considerations, this paper develops a Dynamic Sparse Partial Least Squares (D-sPLS)-based monitoring framework for reliable fault detection in smart agricultural soil sensors. The proposed approach integrates temporal embedding and sparse latent-variable modeling to characterize normal operating behavior using only healthy historical observations. Residual information generated by the D-sPLS model is subsequently monitored using Exponentially Weighted Moving Average (EWMA) statistics and Kernel Density Estimation (KDE)-based thresholds, allowing effective detection of subtle and progressive sensor abnormalities. The proposed methodology is designed to maintain high detection sensitivity while preserving interpretability and low false-alarm rates in realistic agricultural environments. To better position the proposed methodology within the existing literature, Table 1 summarizes the main advantages and limitations of representative multivariate monitoring approaches. The comparison highlights the need for a framework that simultaneously

captures temporal dependencies, performs embedded variable selection, and provides robust monitoring under non-Gaussian operating conditions.

Table 1: Comparison of existing multivariate monitoring approaches for smart agricultural soil sensing.

Method	Main Advantages	Limitations	Relevance to Smart Soil Monitoring
PCA-based monitoring [18, 20]	Simple implementation and effective dimensionality reduction for correlated variables.	Ignores temporal dependencies, retains all variables, and may struggle with non-Gaussian data.	Suitable for static monitoring but limited for dynamic agricultural environments.
PLS-based monitoring [18, 19]	Captures relationships between input and output variables and handles multicollinearity.	Conventional PLS lacks embedded variable selection and does not explicitly model temporal behavior.	Useful for correlated soil variables but less effective for evolving environmental conditions.
Dynamic PCA (DPCA) [21]	Introduces lagged information to represent process dynamics.	High-dimensional augmented spaces reduce interpretability and increase computational complexity.	Captures temporal effects but cannot identify the most informative sensors or lags.
Machine-learning approaches (OCSVM, Isolation Forest, kNN) [15, 16]	Flexible modeling capability and no strong distributional assumptions.	Limited interpretability and sensitivity to parameter tuning and data quality.	Effective for anomaly detection but provide little physical insight into soil processes.
Deep learning methods [12, 13]	Can learn complex nonlinear and temporal patterns automatically.	Require large datasets, high computational resources, and extensive hyperparameter tuning.	Challenging to deploy in data-scarce and resource-constrained agricultural applications.
Proposed D-sPLS-EWMA	Semi-supervised dynamic sparse monitoring with EWMA, KDE-based thresholds, and ICE-based diagnosis.	Requires lag and sparsity selection; currently addresses single-sensor faults.	Reliable, interpretable, and real-time monitoring for dynamic and non-Gaussian smart soil sensing systems.

The main contributions of this work are summarized as follows.

- A effective Dynamic Sparse Partial Least Squares (D-sPLS) framework is proposed for semi-supervised sensor fault detection in smart agriculture. The proposed model combines dynamic feature augmentation with sparse latent-variable modeling to capture temporal dependencies while identifying the most informative variables and time lags. The framework is trained exclusively using anomaly-free data, eliminating the need for labeled fault samples, which are typically scarce, expensive, or unavailable in agricultural monitoring. Consequently, the proposed approach is well-suited

for practical deployments where only normal operating data can be collected during system commissioning.

- A residual-based monitoring strategy integrating the Exponentially Weighted Moving Average (EWMA) statistic with nonparametric Kernel Density Estimation (KDE)-based threshold estimation is proposed to improve the detection of abrupt and incipient sensor faults. The EWMA statistic accumulates small persistent residual deviations, enhancing sensitivity to slowly evolving faults while attenuating random fluctuations. Meanwhile, the distribution-free KDE threshold eliminates the need for prior distributional assumptions, enabling reliable deployment on both Gaussian and non-Gaussian processes.
- The proposed framework is comprehensively validated on two real agricultural datasets involving smart soil monitoring and soil quality assessment. Its performance is compared with representative latent-variable, regression-based, machine-learning, and statistical-monitoring methods across multiple bias and drift fault scenarios. A comprehensive experimental analysis, including hyperparameter sensitivity, computational efficiency, detection delay, and ICE-based fault diagnosis, is complemented by 95% bootstrap confidence intervals and the Wilcoxon signed-rank test, demonstrating the accuracy, robustness, efficiency, and interpretability of the proposed framework.

The remainder of this paper is organized as follows. Section 2 presents the theoretical background and the proposed D-sPLS monitoring framework. Section 3 describes the experimental setup and discusses the results under different fault scenarios. Finally, Section 4 concludes the paper and outlines future research directions.

2 Proposed Methodology

This section presents the proposed semi-supervised D-sPLS framework for sensor monitoring and diagnosis. Trained exclusively on anomaly-free data, the framework comprises four stages: dynamic feature construction, sparse latent-variable modeling, residual-based EWMA monitoring with KDE threshold estimation, and ICE-based post-alarm diagnosis. It operates in two phases: offline model training and online fault detection and interpretation.

2.1 Dynamic Sparse PLS Model

Smart agricultural systems collect multivariate soil and environmental measurements that exhibit temporal dependence and cross-correlation due to weather conditions, irrigation practices, and soil–water interactions [2, 8]. For example, rainfall affects soil-moisture dynamics, whereas temperature and solar radiation influence evaporation and plant water uptake. Since these relationships evolve over time, static monitoring methods may provide an incomplete representation of normal operating behavior and reduced fault-detection sensitivity [20, 21, 33]. To address these limitations, the proposed Dynamic Sparse Partial Least Squares (D-sPLS) model combines temporal feature construction with embedded variable selection. Dynamic latent-variable modeling captures delayed process interactions, while sparse PLS retains the most

informative variables and temporal features and suppresses redundant information [18, 34]. The resulting representation is compact, interpretable, and suitable for modeling normal agricultural sensor behavior.

2.1.1 Dynamic Feature Construction

To accurately represent the temporal behavior of agricultural sensing systems, the proposed framework augments the current measurements with historical observations. This dynamic feature construction enables the model to capture delayed process responses and temporal correlations that cannot be represented by instantaneous measurements alone. Consequently, the resulting feature space provides a richer description of normal operating conditions and improves the predictive capability of the subsequent D-sPLS model. Let

$$\mathbf{x}_t = [x_{1,t}, x_{2,t}, \dots, x_{p,t}]^T \in \mathbb{R}^p \quad (1)$$

denote the multivariate measurement vector at time t , where p is the number of monitored variables. To incorporate temporal information, the current observation is combined with its previous L measurements to form the dynamic vector $\tilde{\mathbf{x}}_t = [\mathbf{x}_t^T, \mathbf{x}_{t-1}^T, \dots, \mathbf{x}_{t-L}^T]^T$. The lag order L determines the temporal memory of the model and is selected using process knowledge, autocorrelation analysis, or cross-validation. The corresponding dynamic data matrix is $\tilde{\mathbf{X}} = [\tilde{\mathbf{x}}_{L+1}^T, \tilde{\mathbf{x}}_{L+2}^T, \dots, \tilde{\mathbf{x}}_N^T]^T \in \mathbb{R}^{(N-L) \times p(L+1)}$, where N is the total number of observations. This representation preserves the sequential structure of the measurements and captures delayed responses, such as the effect of rainfall on soil moisture and the temporal evolution of soil physicochemical properties. It therefore provides a more representative description of normal operating conditions for subsequent sparse latent-variable modeling.

2.1.2 Sparse Latent-Variable Modeling

Dynamic feature construction enriches the input space with temporal information, but it also increases dimensionality and may introduce redundant variables and lagged terms. In practice, only a subset of these features may be relevant for characterizing normal operating conditions. Sparse latent-variable modeling addresses this issue by combining dimensionality reduction with embedded feature selection, thereby improving interpretability and limiting noise propagation [20, 34].

PLS is well-suited to multivariate environmental monitoring because it can handle strongly correlated predictors and extracts latent variables that maximize the covariance between predictor and response spaces [18, 35]. Let

$$\mathbf{X} \in \mathbb{R}^{n \times m}, \quad \mathbf{Y} \in \mathbb{R}^{n \times q}, \quad (2)$$

where $\mathbf{X}$ denotes the dynamically augmented predictor matrix and $\mathbf{Y}$ contains the target measurements to be predicted. The corresponding latent scores are

$$\mathbf{t} = \mathbf{X}\mathbf{w}, \quad \mathbf{u} = \mathbf{Y}\mathbf{c}, \quad (3)$$




Fig. 1: Sparse PLS formulation used in the proposed D-sPLS framework.

where w and c are the projection vectors for X and Y , respectively. Classical PLS estimates these directions by solving

$$\begin{aligned} \max_{w, c} \quad & w^T X^T Y c, \\ \text{subject to} \quad & \|w\|_2 = 1, \quad \|c\|_2 = 1. \end{aligned} \quad (4)$$

Conventional PLS generally assigns nonzero weights to most predictors. This may reduce interpretability and retain irrelevant dynamic features, especially after temporal augmentation. To obtain a more compact representation, sparse PLS introduces an ℓ_1 penalty on the predictor weights [34]:

$$\begin{aligned} \max_{w, c} \quad & w^T X^T Y c - \lambda_s \|w\|_1, \\ \text{subject to} \quad & \|w\|_2 = 1, \quad \|c\|_2 = 1. \end{aligned} \quad (5)$$

where $\lambda_s > 0$ controls the sparsity level. Larger values of λ_s yield fewer active coefficients. Hence, weak variables and redundant lagged terms are suppressed, while the most informative dynamic predictors are retained.

Figure 1 illustrates this process. The augmented matrix X is projected onto a low-dimensional latent space while maximizing its covariance with Y . Sparsity is simultaneously imposed on w , so that the latent variables are constructed from a compact subset of sensors and temporal features.

For A retained components, the sparse score matrix is defined as $\mathbf{T} = [\mathbf{t}_1, \mathbf{t}_2, \dots, \mathbf{t}_A]$. The resulting D-sPLS decomposition is

$$\begin{cases} \mathbf{X} = \mathbf{T}\mathbf{P}^T + \mathbf{E}, \\ \mathbf{Y} = \mathbf{T}\mathbf{Q}^T + \mathbf{F}, \end{cases} \quad (6)$$

where $\mathbf{P}$ and $\mathbf{Q}$ are the loading matrices, and $\mathbf{E}$ and $\mathbf{F}$ are the residual matrices. This representation captures the dominant predictor–response structure while retaining only the most relevant dynamic features, providing a compact basis for prediction and residual-based fault monitoring.

2.1.3 Dynamic Prediction and Residual Generation

Once the sparse latent variables have been extracted, they are used to construct a predictive model that estimates the expected value of the monitored sensor from the remaining process variables. Under normal operating conditions, the predicted output should closely follow the measured sensor value, whereas sensor faults produce abnormal deviations between the prediction and the measurement. Consequently, the prediction residual provides an effective fault-sensitive indicator while remaining insensitive to normal process variations [18, 33]. For A retained latent components, the D-sPLS regression model is expressed as

$$\hat{\mathbf{Y}} = \mathbf{X}\mathbf{W}^*\mathbf{B}\mathbf{Q}^T, \quad (7)$$

where $\mathbf{W}^* = \mathbf{W}(\mathbf{P}^T\mathbf{W})^{-1}$ is the adjusted weight matrix, $\mathbf{B}$ is the diagonal regression coefficient matrix, and $\mathbf{Q}$ contains the response loadings. This formulation projects the dynamically augmented predictor matrix onto the sparse latent space and reconstructs the expected response using only the selected dynamic features. The prediction residual at sample k is then computed as

$$\mathbf{e}(k) = \mathbf{y}(k) - \hat{\mathbf{y}}(k), \quad (8)$$

where $\mathbf{y}(k)$ and $\hat{\mathbf{y}}(k)$ denote the measured and predicted sensor values, respectively. During fault-free operation, the residuals remain close to zero because the D-sPLS model accurately captures the normal multivariate relationships among the monitored variables. When a sensor fault occurs, the measured value becomes inconsistent with the learned process behavior, causing the residual magnitude to increase.

Figure 2 illustrates the proposed prediction stage. The dynamically augmented measurements are first projected onto the sparse latent space to estimate the expected sensor response. The resulting prediction residual is subsequently supplied to the residual-based monitoring stage described in the following subsection.

2.2 D-sPLS-Based EWMA Residual Monitoring

The D-sPLS model characterizes the normal relationship between the target sensor and the remaining measurements. A sensor fault disrupts this learned relationship




Fig. 2: Dynamic prediction and residual generation using the proposed D-sPLS model.

and produces a persistent prediction error. Therefore, the residual sequence is used as the primary fault-sensitive signal. For a scalar target sensor, the residual at time t is defined as

$$e_t = y_t - \hat{y}_t, \quad (9)$$

where y_t and $\hat{y}_t$ denote the measured and D-sPLS-predicted values, respectively. Under fault-free conditions, e_t remains close to its nominal distribution. Bias and drift faults produce sustained deviations that can be detected through sequential residual monitoring.

2.2.1 EWMA Residual Monitoring

Prediction residuals may be affected by measurement noise and natural process variability. Direct thresholding of individual residuals can therefore generate false alarms or fail to detect weak persistent faults. To improve monitoring sensitivity, the residuals are transformed into an EWMA statistic [36, 37]:

$$z_t = \lambda_{\text{EWMA}} e_t + (1 - \lambda_{\text{EWMA}}) z_{t-1}, \quad 0 < \lambda_{\text{EWMA}} \leq 1, \quad (10)$$

where z_t denotes the monitoring statistic and λ_{EWMA} is the forgetting factor. The initial value is set to the mean of the fault-free training residuals ($z_0 = \bar{e}_{\text{train}}$). The forgetting factor controls the trade-off between noise suppression and fault sensitivity. Smaller values assign greater weight to past residuals, making the EWMA statistic more sensitive to small persistent changes, whereas larger values respond more rapidly to abrupt faults but with less smoothing. In practice, values in the range of 0.2–0.3 are commonly recommended for detecting small to moderate process shifts [37]. Notably, $\lambda_{\text{EWMA}} = 1$ reduces the EWMA chart to the classical Shewhart chart. Consequently,

the EWMA statistic effectively suppresses random fluctuations while accumulating sustained residual deviations, making it particularly suitable for detecting both abrupt and incipient sensor faults.

2.2.2 KDE-Based Decision Threshold

Agricultural measurements and their residuals may be skewed, heavy-tailed, or multimodal because of weather variability, irrigation events, and heterogeneous soil conditions. Thresholds derived from Gaussian assumptions may therefore provide unreliable false-alarm control. To overcome this limitation, a nonparametric Kernel Density Estimation (KDE) approach is employed to estimate the distribution of the fault-free EWMA statistic directly from the training data. Let $\{z_i\}_{i=1}^{N_0}$ denote the EWMA statistics computed from the N_0 anomaly-free training samples. Their probability density is estimated as

$$\hat{f}(z) = \frac{1}{N_0 b} \sum_{i=1}^{N_0} K\left(\frac{z - z_i}{b}\right), \quad (11)$$

where $K(\cdot)$ denotes the kernel function and $b > 0$ is the bandwidth controlling the smoothness of the estimated density. In this work, a Gaussian kernel is employed because of its smoothness and widespread use in nonparametric statistical process monitoring. The bandwidth is selected automatically using Silverman's rule-of-thumb [38], providing a data-driven balance between estimation bias and variance. The corresponding cumulative distribution function is

$$\hat{F}(z) = \int_{-\infty}^z \hat{f}(u) du. \quad (12)$$

For a prescribed confidence level α , the monitoring threshold is determined as

$$h_\alpha = \hat{F}^{-1}(\alpha). \quad (13)$$

A fault alarm is generated whenever $z_t > h_\alpha$. Because the threshold is estimated directly from the empirical distribution of the fault-free residuals, the proposed KDE procedure does not require any prior distributional assumption. Consequently, it can be applied to both Gaussian and non-Gaussian monitoring data.

2.2.3 Offline Training and Online Detection

The proposed framework consists of an offline training phase and an online monitoring phase. During offline training, only anomaly-free observations are used. The variables are standardized using training-data statistics, dynamic features are constructed, and the D-sPLS model is estimated. The resulting residuals are transformed into EWMA statistics, from which the KDE-based threshold is obtained. During online operation, each incoming observation undergoes the same standardization and dynamic feature construction. The trained D-sPLS model predicts the expected target value,

after which the residual and EWMA statistic are updated. The sample is declared faulty when the monitoring statistic exceeds the learned threshold. The model parameters, normalization statistics, and decision threshold remain fixed during testing. Algorithm 1 summarizes the complete procedure.

Algorithm 1 Proposed D-sPLS Monitoring Framework

Input: Anomaly-free predictors $\mathbf{X}_{\text{train}}$, target measurements $\mathbf{Y}_{\text{train}}$, incoming observations, lag order L , number of components A , sparsity parameter λ_s , EWMA parameter λ_{EWMA} , and confidence level α . **Output:** Monitoring statistics and fault alarms.

Offline training

- 1: Estimate the normalization parameters from the anomaly-free training data.
- 2: Standardize $\mathbf{X}_{\text{train}}$ and $\mathbf{Y}_{\text{train}}$.
- 3: Construct the dynamically augmented predictor matrix $\tilde{\mathbf{X}}_{\text{train}}$ using lag order L .
- 4: Train D-sPLS using A latent components and sparsity parameter λ_s .
- 5: Predict the anomaly-free target values:

$$\hat{\mathbf{Y}}_{\text{train}} = f_{\text{D-sPLS}}(\tilde{\mathbf{X}}_{\text{train}}).$$

- 6: Compute the training residuals:

$$\mathbf{e}_{\text{train}} = \mathbf{Y}_{\text{train}} - \hat{\mathbf{Y}}_{\text{train}}.$$

- 7: Compute the training EWMA sequence:

$$z_t = \lambda_{\text{EWMA}} e_t + (1 - \lambda_{\text{EWMA}}) z_{t-1}.$$

- 8: Estimate the KDE-based threshold h_α from the fault-free EWMA statistics.

Online monitoring

- 9: **for** each incoming sample at time t **do**
- 10: Standardize the new measurements using the training parameters.
- 11: Construct the dynamic predictor vector $\tilde{\mathbf{x}}_t$.
- 12: Predict the expected target value:

$$\hat{y}_t = f_{\text{D-sPLS}}(\tilde{\mathbf{x}}_t).$$

- 13: Compute the prediction residual:

$$e_t = y_t - \hat{y}_t.$$

- 14: Update the EWMA statistic:

$$z_t = \lambda_{\text{EWMA}} e_t + (1 - \lambda_{\text{EWMA}}) z_{t-1}.$$

- 15: **if** $z_t > h_\alpha$ **then**
 - 16: Declare a sensor fault and activate the diagnostic layer.
 - 17: **else**
 - 18: Declare normal operation.
 - 19: **end if**
 - 20: **end for**
-

The monitoring performance was evaluated using the true positive rate (TPR), false positive rate (FPR), accuracy, precision, and F1-score, computed from the confusion matrix. Here, true positives (TP) and true negatives (TN) denote correctly

identified faulty and normal samples, respectively, whereas false positives (FP) and false negatives (FN) correspond to false alarms and missed detections.

$$\text{TPR} = \frac{TP}{TP + FN}, \quad \text{FPR} = \frac{FP}{FP + TN}, \quad (14)$$

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN}, \quad \text{Precision} = \frac{TP}{TP + FP}, \quad (15)$$

$$\text{F1} = \frac{2 \text{Precision} \times \text{TPR}}{\text{Precision} + \text{TPR}}. \quad (16)$$

2.3 ICE-Based Fault Diagnosis

The residual-monitoring stage determines whether the target sensor has departed from its expected behavior, but it does not explain the origin or severity of the alarm. To provide an interpretable post-alarm analysis, Individual Conditional Expectation (ICE) and Partial Dependence Plot (PDP) techniques are incorporated as a diagnostic layer [39, 40]. This layer is activated after fault detection and does not alter the monitoring decision.

Let $f(\mathbf{x})$ denote the trained D-sPLS prediction model and let x_j be a predictor selected for diagnostic analysis. For observation i , the ICE curve is defined as

$$\text{ICE}_{i,j}(v) = f(x_{i,1}, \dots, x_{i,j-1}, v, x_{i,j+1}, \dots, x_{i,p}), \quad (17)$$

where the selected predictor is varied over a grid of values v , while the remaining predictors are fixed at their observed values. The PDP is obtained by averaging the ICE curves:

$$\text{PDP}_j(v) = \frac{1}{N_0} \sum_{i=1}^{N_0} \text{ICE}_{i,j}(v). \quad (18)$$

The fault-free ICE curves define the expected predictor–response region learned from anomaly-free data. During monitoring, the measured target values are superimposed on this region. Measurements that remain inside the ICE envelope are consistent with the learned multivariate behavior, whereas systematic departures indicate an abnormal sensor response. This diagnostic representation provides information that is not available from the binary alarm alone. An abrupt displacement beyond the normal envelope is indicative of a bias-type fault, whereas a progressive departure reflects gradual sensor drift. The number and magnitude of outlying measurements also provide an intuitive indication of fault severity. Consequently, the ICE/PDP layer helps operators verify the alarm, distinguish abnormal sensor behavior from legitimate process changes, and prioritize recalibration or maintenance actions.

3 Results and Discussion

This section evaluates the proposed D-sPLS framework using two real agricultural datasets for smart soil sensor monitoring and soil quality assessment. Performance is assessed under multiple bias and drift fault scenarios and compared with representative

latent-variable, regression-based, kernel-based, neural-network, and statistical monitoring methods. For a fair comparison, all prediction-based methods employ the same EWMA residual monitoring and KDE-based thresholding strategy. The benchmark methods are summarized in Table 2.

Table 2: Summary of the benchmark methods considered for comparison.

Method	Category	Main characteristic
D-sPLS-EWMA	Proposed latent-variable model	Dynamic sparse PLS with EWMA residual monitoring and KDE thresholding.
SparsePLS-EWMA	Sparse latent-variable model	Sparse Partial Least Squares with embedded variable selection [34].
PLS-EWMA	Latent-variable model	Classical Partial Least Squares for correlated multivariate regression [35].
ElasticNet-EWMA	Regularized regression	Linear regression using combined ℓ_1 and ℓ_2 regularization [41].
DecisionTree-EWMA	Tree-based regression	Regression tree based on recursive partitioning [42].
ELM-EWMA	Neural network	Single-hidden-layer Extreme Learning Machine with analytical training [43].
SVR	Kernel regression	Support Vector Regression using structural risk minimization [44].
GPR	Probabilistic regression	Bayesian Gaussian Process Regression with uncertainty estimation [45].
KRR-EWMA	Kernel regression	Kernel Ridge Regression using regularized least squares [46, 47].
PCA- T^2	Statistical monitoring	Hotelling’s T^2 statistic in the PCA latent space [48].
PCA-SPE	Statistical monitoring	Squared Prediction Error for PCA residual monitoring [48].

3.1 Application 1: Smart Soil Sensor Fault Detection

3.1.1 Dataset Description and Exploratory Analysis

The first case study considers a real-world soil sensing dataset collected in Guayaquil, Ecuador [49]. The sensing platform continuously monitors environmental and soil conditions relevant to precision agriculture, providing a realistic benchmark for evaluating multivariate sensor fault detection methods. Figure 3 illustrates the experimental sensing platform. The sensing platform consists of an Arduino Uno connected to four low-cost sensors measuring air humidity (AH), air temperature (AT), soil moisture (SM), light intensity (LI), and rain occurrence (RO). Measurements were acquired every 30 s over three consecutive days, yielding 2,864 multivariate observations.

The combination of temporal dynamics and cross-correlations among the monitored variables makes this dataset suitable for evaluating dynamic sensor fault detection.




Fig. 3: Experimental soil sensing platform used in Application 1.

Figure 4 presents the autocorrelation functions (ACFs) of the anomaly-free sensor measurements over a 4 h lag. Air humidity, air temperature, soil moisture, and light intensity exhibit clear temporal dependence. Soil moisture shows the slowest decay, reflecting its gradual dynamics. In contrast, rain occurrence remains weakly autocorrelated because it is binary and sporadic. These results confirm the temporal structure of the continuous variables and justify the dynamic formulation of the proposed D-sPLS model.

To further characterize the relationships among the monitored variables, a dependency analysis was performed using the anomaly-free training data. Figure 5(a) presents the Spearman correlation matrix, showing a strong negative correlation between air humidity and air temperature ($\rho = -0.70$, $p < 0.001$). Soil moisture also exhibits statistically significant associations with air humidity, air temperature, and light intensity. The corresponding p -value matrix in Figure 5(b) confirms that most pairwise relationships are statistically significant ($p < 0.05$). Figure 5(c) shows the normalized mutual information (NMI) matrix, which captures both linear and nonlinear dependencies. The highest NMI values are observed between soil moisture and air humidity (0.40) and between soil moisture and air temperature (0.26), indicating nonlinear interactions not fully captured by correlation analysis. Overall, the complementary correlation and NMI analyses confirm the complex dependency structure of the sensor measurements, supporting the use of the proposed D-sPLS framework for multivariate sensor fault detection.

Figure 6 ranks the predictor variables according to their association with soil moisture using the composite score $S_j = \frac{1}{2} (|\rho_{s,j}| + \text{NMI}_j)$. Air humidity is identified as the most informative predictor (0.33), followed by air temperature and light intensity






Fig. 4: ACFs of the anomaly-free sensor measurements over a 4 h lag.

(0.17), indicating that atmospheric conditions largely govern soil moisture dynamics. Rain occurrence exhibits the lowest association score (0.08), primarily due to its binary, sporadic nature despite its physical relevance. These results highlight the unequal predictive contributions of the measured variables and further motivate the sparse variable selection mechanism employed by the proposed D-sPLS framework.



(a) Spearman correlation coefficients.

(b) Corresponding p -values.



(c) Normalized mutual information.

Fig. 5: Dependency analysis of the anomaly-free sensor measurements: (a) Spearman correlation coefficients, (b) corresponding p -values, and (c) normalized mutual information (NMI).

3.1.2 Model Development and Hyperparameter Setting

The proposed monitoring framework employs a semi-supervised strategy, using only anomaly-free observations to develop the model. The available fault-free data were partitioned into training (80%) and testing (20%) subsets. The training set, comprising approximately 1000 healthy observations, was used for model training, hyperparameter




Fig. 6: Association ranking of the predictor variables with respect to soil moisture based on the composite association score.

optimization, and threshold estimation, whereas the testing set was reserved exclusively for fault injection and performance evaluation. All variables were standardized using the mean and standard deviation computed from the training data. Hyperparameters were selected using five-fold cross-validation by minimizing the fault-free prediction RMSE, thereby ensuring accurate residual generation for subsequent monitoring. For the proposed D-sPLS model, the candidate feature pool consisted of the current measurements, lagged variables, first differences, moving averages, and interaction terms. Sparse optimization retained 25 informative predictors and 20 latent components. Conventional SparsePLS and PLS retained four latent components, while the remaining benchmark models were tuned using their standard cross-validation procedures. Under fault-free conditions, the proposed D-sPLS model achieved an RMSE of 76.46 and an R^2 of 0.882, demonstrating an accurate characterization of the normal process behavior. The selected hyperparameter settings are summarized in Table 3. During online monitoring, all prediction-based methods employed the same EWMA residual monitoring scheme with forgetting factor $\lambda_{\text{EWMA}} = 0.15$ and a KDE-based decision threshold estimated from the fault-free training residuals at a confidence level of 97%.

3.1.3 Fault Detection Results

The fault detection capability of the proposed D-sPLS–EWMA framework is evaluated under two representative sensor fault scenarios: abrupt bias faults and gradual drift faults, which emulate the most common sensor degradations in smart agriculture. Performance is assessed using TPR, FPR, accuracy, precision, F1-score, and detection delay.

Table 3: Main hyperparameter settings of the investigated prediction models.

Method	Main hyperparameter settings
D-sPLS	20 latent components; 25 retained predictors after sparse optimization; retention fraction = 0.50; two residual-correction stages with weights 1 and 0.5.
SparsePLS	4 latent components; 4 retained predictors.
PLS	4 latent components.
ElasticNet	Mixing parameter $\alpha = 0.5$; regularization parameter $\lambda = 1.2852 \times 10^{-4}$.
DecisionTree	Minimum leaf size = 5; maximum number of splits = 50.
SVR	Gaussian kernel with automatic kernel-scale selection.
GPR	ARD squared-exponential kernel; subset-of-regressors approximation; fully independent conditional prediction.
KRR	RBF kernel; kernel width $\sigma = 1.9191$; ridge parameter = 0.01.
PCA	Four principal components used for both T^2 and SPE monitoring.

Scenario 1: Bias Fault Detection

Sensor bias is one of the most common faults in smart soil sensing systems, arising from sensor aging, calibration errors, fouling, or environmental influences. Such faults introduce a constant offset into the measured soil-moisture signal, potentially leading to inaccurate irrigation decisions and inefficient water management. During testing, bias faults were injected into the soil-moisture sensor from sample $k_f = 550$, while the remaining sensors were assumed to operate normally and served as predictors. The faulty measurement was generated as

$$y_f(k) = y(k) + \frac{b}{100}y(k), \quad k \geq k_f,$$

where $y(k)$ and $y_f(k)$ denote the fault-free and faulty soil-moisture measurements, respectively, and b is the bias magnitude. Five bias levels (10%, 20%, 30%, and 50%) were considered to assess the robustness of the monitoring methods under increasing fault severity.

Table 4 summarizes the monitoring performance under the different soil-moisture sensor bias levels. The reported values correspond to the mean of 300 bootstrap resampling runs, while the associated uncertainty is expressed as the half-width of the 95% bootstrap confidence interval, providing a statistical assessment of the robustness and repeatability of each monitoring method.

The proposed D-sPLS consistently provides the best overall monitoring performance across all bias levels. Even for the smallest bias (10%), it achieves a TPR of 0.95 ± 0.02 with zero false alarms, resulting in the highest F1-score (0.97 ± 0.01) among the methods that maintain a zero false-alarm rate. Although GPR and PCA-SPE attain comparable or slightly higher detection rates, they consistently produce nonzero

Table 4: Performance under different soil-moisture sensor bias levels (mean $\pm$ 95% bootstrap CI half-width).

Bias (%)	Method	TPR	FPR	Accuracy	Precision	F1
10	D-sPLS	0.95 $\pm$ 0.02	0.00 $\pm$ 0.00	0.98 $\pm$ 0.01	1.00 $\pm$ 0.00	0.97 $\pm$ 0.01
	SparsePLS	0.64 $\pm$ 0.04	0.00 $\pm$ 0.00	0.84 $\pm$ 0.02	1.00 $\pm$ 0.00	0.78 $\pm$ 0.03
	PLS	0.62 $\pm$ 0.05	0.00 $\pm$ 0.00	0.83 $\pm$ 0.02	1.00 $\pm$ 0.00	0.76 $\pm$ 0.03
	ElasticNet	0.64 $\pm$ 0.04	0.02 $\pm$ 0.01	0.83 $\pm$ 0.02	0.97 $\pm$ 0.02	0.77 $\pm$ 0.03
	DecisionTree	0.56 $\pm$ 0.04	0.00 $\pm$ 0.00	0.80 $\pm$ 0.02	1.00 $\pm$ 0.00	0.72 $\pm$ 0.03
	SVR	0.31 $\pm$ 0.04	0.03 $\pm$ 0.01	0.68 $\pm$ 0.03	0.90 $\pm$ 0.05	0.47 $\pm$ 0.05
	GPR	0.96 $\pm$ 0.02	0.04 $\pm$ 0.01	0.96 $\pm$ 0.01	0.96 $\pm$ 0.02	0.96 $\pm$ 0.01
	KRR	0.42 $\pm$ 0.05	0.06 $\pm$ 0.02	0.70 $\pm$ 0.03	0.85 $\pm$ 0.04	0.56 $\pm$ 0.04
	PCA- T^2	0.46 $\pm$ 0.04	0.06 $\pm$ 0.02	0.73 $\pm$ 0.03	0.87 $\pm$ 0.04	0.60 $\pm$ 0.04
20	PCA-SPE	0.70 $\pm$ 0.04	0.02 $\pm$ 0.01	0.85 $\pm$ 0.02	0.96 $\pm$ 0.02	0.81 $\pm$ 0.03
	D-sPLS	0.99 $\pm$ 0.01	0.00 $\pm$ 0.00	0.99 $\pm$ 0.00	1.00 $\pm$ 0.00	0.99 $\pm$ 0.00
	SparsePLS	0.88 $\pm$ 0.03	0.00 $\pm$ 0.00	0.94 $\pm$ 0.01	1.00 $\pm$ 0.00	0.94 $\pm$ 0.02
	PLS	0.86 $\pm$ 0.03	0.00 $\pm$ 0.00	0.94 $\pm$ 0.02	1.00 $\pm$ 0.00	0.93 $\pm$ 0.02
	ElasticNet	0.88 $\pm$ 0.03	0.00 $\pm$ 0.00	0.94 $\pm$ 0.01	1.00 $\pm$ 0.00	0.94 $\pm$ 0.02
	DecisionTree	0.66 $\pm$ 0.04	0.00 $\pm$ 0.00	0.85 $\pm$ 0.02	1.00 $\pm$ 0.00	0.79 $\pm$ 0.03
	SVR	0.45 $\pm$ 0.05	0.00 $\pm$ 0.00	0.75 $\pm$ 0.02	1.00 $\pm$ 0.00	0.62 $\pm$ 0.04
	GPR	0.99 $\pm$ 0.01	0.04 $\pm$ 0.02	0.98 $\pm$ 0.01	0.96 $\pm$ 0.02	0.97 $\pm$ 0.01
	KRR	0.56 $\pm$ 0.05	0.06 $\pm$ 0.02	0.77 $\pm$ 0.03	0.88 $\pm$ 0.04	0.68 $\pm$ 0.04
30	PCA- T^2	0.50 $\pm$ 0.04	0.07 $\pm$ 0.02	0.74 $\pm$ 0.03	0.86 $\pm$ 0.05	0.63 $\pm$ 0.04
	PCA-SPE	1.00 $\pm$ 0.00	0.04 $\pm$ 0.02	0.98 $\pm$ 0.01	0.95 $\pm$ 0.02	0.98 $\pm$ 0.01
	D-sPLS	0.99 $\pm$ 0.01	0.00 $\pm$ 0.00	1.00 $\pm$ 0.00	1.00 $\pm$ 0.00	1.00 $\pm$ 0.00
	SparsePLS	0.97 $\pm$ 0.01	0.00 $\pm$ 0.00	0.99 $\pm$ 0.01	1.00 $\pm$ 0.00	0.98 $\pm$ 0.01
	PLS	0.95 $\pm$ 0.02	0.00 $\pm$ 0.00	0.98 $\pm$ 0.01	1.00 $\pm$ 0.00	0.97 $\pm$ 0.01
	ElasticNet	0.98 $\pm$ 0.01	0.00 $\pm$ 0.00	0.99 $\pm$ 0.01	1.00 $\pm$ 0.00	0.99 $\pm$ 0.01
	DecisionTree	0.68 $\pm$ 0.04	0.00 $\pm$ 0.00	0.86 $\pm$ 0.02	1.00 $\pm$ 0.00	0.81 $\pm$ 0.03
	SVR	0.36 $\pm$ 0.05	0.00 $\pm$ 0.00	0.71 $\pm$ 0.03	1.00 $\pm$ 0.00	0.53 $\pm$ 0.05
	GPR	0.99 $\pm$ 0.01	0.04 $\pm$ 0.02	0.97 $\pm$ 0.01	0.95 $\pm$ 0.02	0.97 $\pm$ 0.01
50	KRR	0.65 $\pm$ 0.04	0.06 $\pm$ 0.02	0.81 $\pm$ 0.02	0.90 $\pm$ 0.03	0.76 $\pm$ 0.03
	PCA- T^2	0.50 $\pm$ 0.05	0.06 $\pm$ 0.02	0.74 $\pm$ 0.03	0.88 $\pm$ 0.04	0.63 $\pm$ 0.05
	PCA-SPE	1.00 $\pm$ 0.00	0.04 $\pm$ 0.02	0.98 $\pm$ 0.01	0.95 $\pm$ 0.02	0.98 $\pm$ 0.01
	D-sPLS	0.99 $\pm$ 0.01	0.00 $\pm$ 0.00	1.00 $\pm$ 0.00	1.00 $\pm$ 0.00	1.00 $\pm$ 0.00
	SparsePLS	0.99 $\pm$ 0.01	0.00 $\pm$ 0.00	1.00 $\pm$ 0.00	1.00 $\pm$ 0.00	1.00 $\pm$ 0.00
	PLS	0.99 $\pm$ 0.01	0.00 $\pm$ 0.00	1.00 $\pm$ 0.00	1.00 $\pm$ 0.00	1.00 $\pm$ 0.00
	ElasticNet	0.99 $\pm$ 0.01	0.00 $\pm$ 0.00	1.00 $\pm$ 0.00	1.00 $\pm$ 0.00	1.00 $\pm$ 0.00
	DecisionTree	0.74 $\pm$ 0.04	0.00 $\pm$ 0.00	0.88 $\pm$ 0.02	1.00 $\pm$ 0.00	0.85 $\pm$ 0.02
	SVR	0.99 $\pm$ 0.01	0.01 $\pm$ 0.01	0.99 $\pm$ 0.01	0.99 $\pm$ 0.01	0.99 $\pm$ 0.01
	GPR	1.00 $\pm$ 0.01	0.04 $\pm$ 0.02	0.97 $\pm$ 0.01	0.95 $\pm$ 0.02	0.97 $\pm$ 0.01
	KRR	0.99 $\pm$ 0.01	0.13 $\pm$ 0.03	0.92 $\pm$ 0.02	0.86 $\pm$ 0.03	0.92 $\pm$ 0.02
	PCA- T^2	0.57 $\pm$ 0.05	0.06 $\pm$ 0.02	0.77 $\pm$ 0.03	0.89 $\pm$ 0.03	0.69 $\pm$ 0.04
	PCA-SPE	0.99 $\pm$ 0.01	0.00 $\pm$ 0.00	0.99 $\pm$ 0.00	1.00 $\pm$ 0.01	0.99 $\pm$ 0.01

false-alarm rates (2–4%), making them less suitable for continuous online monitoring where false alarms are costly. As the bias magnitude increases, D-sPLS rapidly approaches perfect detection while preserving zero false alarms, achieving nearly perfect accuracy, precision, and F1-score for bias levels of 30% and above. SparsePLS, PLS, and ElasticNet also exhibit progressively improved detection performance as the bias magnitude increases, approaching the performance of D-sPLS for large bias faults. However, their sensitivity is noticeably lower for small faults, particularly at the 10% bias level, where early detection is most challenging. DecisionTree consistently provides

lower detection rates across all bias levels, indicating limited capability to capture the multivariate relationships underlying the sensor dynamics. Among the kernel-based approaches, SVR performs poorly for small and moderate bias faults but achieves near-perfect detection only for the largest bias level (50%). In contrast, KRR exhibits moderate detection performance together with relatively high false-alarm rates, reaching an FPR of 0.13 ± 0.03 at the 50% bias level. PCA- T^2 remains the least sensitive method throughout all bias levels, whereas PCA-SPE achieves excellent detection rates but at the expense of persistent false alarms. The narrow bootstrap confidence intervals demonstrate that the proposed D-sPLS provides highly stable performance under repeated resampling, indicating strong statistical robustness. Its superior performance can be attributed to the combination of dynamic feature augmentation and sparse variable selection, which captures temporal dependencies while retaining only the most informative predictors. Consequently, the generated residuals are more discriminative, enabling accurate bias-fault detection with consistently low false-alarm rates across all fault severities.

The Friedman omnibus test [50] was applied to compare the F1-scores of all monitoring methods across the seven soil-moisture bias levels. The test revealed a significant difference among the methods ($\chi^2(9) = 46.454$, $p < 0.001$), with Kendall's coefficient of concordance ($W = 0.737$) indicating a large overall effect. Based on the mean ranks, D-sPLS achieved the best overall performance, followed by PCA-SPE, SparsePLS, PLS, and ElasticNet. Consequently, pairwise Wilcoxon signed-rank tests with Holm correction were performed to compare D-sPLS against each competing method. As summarized in Table 5, D-sPLS consistently outperformed all competing methods, with the largest F1-score improvements observed over PCA- T^2 , SVR, and KRR. Although the Holm-adjusted p -values exceeded 0.05 because of the limited number of paired bias levels, all paired F1-score differences were positive and accompanied by moderate-to-large rank-biserial effect sizes, providing strong evidence of the practical superiority and consistent performance of the proposed D-sPLS framework.

Table 5: Wilcoxon signed-rank test for D-sPLS under soil-moisture bias faults.

Compared model	Mean F1	Mean Δ F1	Median Δ F1	W	p -value	Holm p -value	Effect size
GPR	0.970	0.023	0.025	28.0	0.008	0.070	1.000
PCA-SPE	0.962	0.031	0.011	25.0	0.039	0.117	0.786
SparsePLS	0.955	0.038	0.001	25.0	0.039	0.117	0.786
ElasticNet	0.952	0.040	0.006	28.0	0.008	0.070	1.000
PLS	0.950	0.043	0.003	25.0	0.039	0.117	0.786
DecisionTree	0.854	0.139	0.176	27.0	0.016	0.070	0.929
KRR	0.808	0.184	0.104	28.0	0.008	0.070	1.000
SVR	0.787	0.206	0.074	28.0	0.008	0.070	1.000
PCA- T^2	0.681	0.312	0.331	28.0	0.008	0.070	1.000

Scenario 2: Sensor Drift Fault

The second scenario evaluates the proposed framework under gradual sensor drift, which commonly results from sensor aging, calibration degradation, or environmental

exposure. Unlike abrupt bias faults, drift evolves progressively, making early detection more challenging. Drift faults were injected into the soil-moisture sensor from sample $k_f = 550$, while the remaining sensors were assumed to operate normally and served as predictors. The faulty measurement was generated as

$$y_f(k) = \begin{cases} y(k), & k < k_f, \\ y(k) + \frac{s}{100}y(k)\frac{k - k_f}{100}, & k \geq k_f, \end{cases}$$

where $y(k)$ and $y_f(k)$ denote the fault-free and faulty measurements, respectively, and s is the drift slope. Four drift severities (5%, 10%, 15%, and 20% per 100 samples) were considered to assess the sensitivity of the monitoring methods to gradually developing faults.

Table 6 summarizes the monitoring performance under different drift levels, reported as the mean of 300 bootstrap resampling runs with 95% bootstrap confidence intervals. The proposed D-sPLS consistently achieves the best overall performance across all drift severities. For the smallest drift (5%), it attains the highest detection rate (TPR = 0.73) with zero false alarms, yielding an F1-score of 0.85. As drift severity increases, detection performance improves steadily, reaching a TPR of 0.92 and an F1-score of 0.96 at the 15% and 20% drift levels, while maintaining zero false alarms. SparsePLS and PLS provide the closest competing performance but consistently remain below D-sPLS. GPR and PCA-SPE also achieve competitive results with zero false alarms, whereas KRR produces higher false-alarm rates despite relatively high detection rates. ElasticNet, DecisionTree, and SVR show lower sensitivity, particularly for small drifts, while PCA- T^2 exhibits the weakest overall performance. The detection-delay results further highlight the trade-off between detection speed and monitoring reliability. As the drift severity increases, the detection delay of D-sPLS decreases from 133 samples at the 5% drift level to 39 samples at the 15% and 20% drift levels. Although GPR, KRR, DecisionTree, and PCA-based methods detect faults earlier, their shorter delays are generally accompanied by lower detection accuracy or higher false-alarm rates. The narrow bootstrap confidence intervals demonstrate the robustness of D-sPLS under repeated resampling. Overall, the combination of dynamic feature augmentation and sparse variable selection enables D-sPLS to achieve a favorable balance among early detection, high sensitivity, and zero false alarms, making it well-suited for continuous soil monitoring applications.

The Friedman ranking confirms the overall superiority of D-sPLS under soil-moisture drift faults, achieving the best mean rank (1.000) and the highest mean F1-score (0.927) across the four drift slopes. PCA-SPE, GPR, SparsePLS, and PLS formed the next best group of methods, whereas KRR, DecisionTree, SVR, PCA- T^2 , and ElasticNet showed lower and less consistent performance. Consistent with these results, the pairwise Wilcoxon signed-rank tests (Table 7) show positive mean and median $\Delta F1$ values for all comparisons, indicating that D-sPLS consistently outperformed every competing method. The largest improvements were observed over ElasticNet, PCA- T^2 , SVR, and DecisionTree, while the smallest margins were against PCA-SPE, GPR, SparsePLS, and PLS. Although the Holm-adjusted p -values were

Table 6: Performance under different soil-moisture sensor drift levels (mean $\pm$ 95% bootstrap CI half-width).

Drift (%)	Method	TPR	FPR	Accuracy	Precision	F1
5	D-sPLS	0.73 $\pm$ 0.04	0.00 $\pm$ 0.00	0.87 $\pm$ 0.02	1.00 $\pm$ 0.00	0.85 $\pm$ 0.03
	SparsePLS	0.62 $\pm$ 0.04	0.00 $\pm$ 0.00	0.81 $\pm$ 0.02	1.00 $\pm$ 0.00	0.77 $\pm$ 0.03
	PLS	0.61 $\pm$ 0.04	0.00 $\pm$ 0.00	0.81 $\pm$ 0.03	1.00 $\pm$ 0.00	0.76 $\pm$ 0.03
	ElasticNet	0.35 $\pm$ 0.04	0.00 $\pm$ 0.00	0.68 $\pm$ 0.03	1.00 $\pm$ 0.00	0.52 $\pm$ 0.04
	DecisionTree	0.49 $\pm$ 0.04	0.00 $\pm$ 0.00	0.75 $\pm$ 0.03	1.00 $\pm$ 0.00	0.66 $\pm$ 0.04
	SVR	0.50 $\pm$ 0.04	0.04 $\pm$ 0.02	0.73 $\pm$ 0.03	0.93 $\pm$ 0.03	0.65 $\pm$ 0.03
	GPR	0.68 $\pm$ 0.04	0.00 $\pm$ 0.00	0.84 $\pm$ 0.02	1.00 $\pm$ 0.00	0.81 $\pm$ 0.03
	KRR	0.67 $\pm$ 0.04	0.07 $\pm$ 0.02	0.80 $\pm$ 0.03	0.91 $\pm$ 0.03	0.77 $\pm$ 0.03
	PCA- T^2	0.59 $\pm$ 0.04	0.08 $\pm$ 0.02	0.76 $\pm$ 0.03	0.88 $\pm$ 0.04	0.71 $\pm$ 0.03
	PCA-SPE	0.60 $\pm$ 0.04	0.00 $\pm$ 0.00	0.80 $\pm$ 0.03	1.00 $\pm$ 0.00	0.75 $\pm$ 0.03
10	D-sPLS	0.89 $\pm$ 0.03	0.00 $\pm$ 0.00	0.94 $\pm$ 0.01	1.00 $\pm$ 0.00	0.94 $\pm$ 0.01
	SparsePLS	0.82 $\pm$ 0.03	0.00 $\pm$ 0.00	0.91 $\pm$ 0.02	1.00 $\pm$ 0.00	0.90 $\pm$ 0.02
	PLS	0.81 $\pm$ 0.03	0.00 $\pm$ 0.00	0.91 $\pm$ 0.02	1.00 $\pm$ 0.00	0.90 $\pm$ 0.02
	ElasticNet	0.58 $\pm$ 0.04	0.00 $\pm$ 0.00	0.79 $\pm$ 0.02	1.00 $\pm$ 0.00	0.73 $\pm$ 0.03
	DecisionTree	0.62 $\pm$ 0.04	0.00 $\pm$ 0.00	0.81 $\pm$ 0.02	1.00 $\pm$ 0.00	0.77 $\pm$ 0.03
	SVR	0.73 $\pm$ 0.04	0.01 $\pm$ 0.01	0.86 $\pm$ 0.02	0.99 $\pm$ 0.01	0.84 $\pm$ 0.03
	GPR	0.79 $\pm$ 0.03	0.00 $\pm$ 0.00	0.90 $\pm$ 0.02	1.00 $\pm$ 0.00	0.88 $\pm$ 0.02
	KRR	0.80 $\pm$ 0.03	0.07 $\pm$ 0.02	0.87 $\pm$ 0.02	0.92 $\pm$ 0.02	0.86 $\pm$ 0.02
	PCA- T^2	0.65 $\pm$ 0.04	0.09 $\pm$ 0.03	0.78 $\pm$ 0.03	0.88 $\pm$ 0.03	0.75 $\pm$ 0.03
	PCA-SPE	0.86 $\pm$ 0.03	0.00 $\pm$ 0.00	0.93 $\pm$ 0.02	1.00 $\pm$ 0.00	0.93 $\pm$ 0.02
15	D-sPLS	0.92 $\pm$ 0.02	0.00 $\pm$ 0.00	0.96 $\pm$ 0.01	1.00 $\pm$ 0.00	0.96 $\pm$ 0.01
	SparsePLS	0.89 $\pm$ 0.03	0.01 $\pm$ 0.01	0.94 $\pm$ 0.01	0.99 $\pm$ 0.01	0.94 $\pm$ 0.02
	PLS	0.87 $\pm$ 0.03	0.00 $\pm$ 0.00	0.94 $\pm$ 0.02	1.00 $\pm$ 0.00	0.93 $\pm$ 0.02
	ElasticNet	0.72 $\pm$ 0.04	0.00 $\pm$ 0.00	0.86 $\pm$ 0.02	1.00 $\pm$ 0.00	0.84 $\pm$ 0.02
	DecisionTree	0.84 $\pm$ 0.03	0.00 $\pm$ 0.00	0.92 $\pm$ 0.02	1.00 $\pm$ 0.00	0.92 $\pm$ 0.02
	SVR	0.78 $\pm$ 0.03	0.00 $\pm$ 0.00	0.89 $\pm$ 0.02	1.00 $\pm$ 0.00	0.87 $\pm$ 0.02
	GPR	0.87 $\pm$ 0.03	0.00 $\pm$ 0.00	0.93 $\pm$ 0.02	1.00 $\pm$ 0.00	0.93 $\pm$ 0.03
	KRR	0.85 $\pm$ 0.03	0.06 $\pm$ 0.02	0.89 $\pm$ 0.02	0.93 $\pm$ 0.02	0.89 $\pm$ 0.02
	PCA- T^2	0.71 $\pm$ 0.04	0.09 $\pm$ 0.03	0.81 $\pm$ 0.02	0.89 $\pm$ 0.03	0.79 $\pm$ 0.03
	PCA-SPE	0.90 $\pm$ 0.03	0.00 $\pm$ 0.00	0.95 $\pm$ 0.01	1.00 $\pm$ 0.00	0.95 $\pm$ 0.02
20	D-sPLS	0.92 $\pm$ 0.03	0.00 $\pm$ 0.00	0.96 $\pm$ 0.01	1.00 $\pm$ 0.00	0.96 $\pm$ 0.01
	SparsePLS	0.89 $\pm$ 0.03	0.00 $\pm$ 0.00	0.94 $\pm$ 0.01	1.00 $\pm$ 0.00	0.94 $\pm$ 0.01
	PLS	0.89 $\pm$ 0.03	0.00 $\pm$ 0.00	0.94 $\pm$ 0.01	1.00 $\pm$ 0.00	0.94 $\pm$ 0.02
	ElasticNet	0.78 $\pm$ 0.04	0.00 $\pm$ 0.00	0.89 $\pm$ 0.02	1.00 $\pm$ 0.00	0.88 $\pm$ 0.02
	DecisionTree	0.86 $\pm$ 0.03	0.00 $\pm$ 0.00	0.93 $\pm$ 0.02	1.00 $\pm$ 0.00	0.93 $\pm$ 0.02
	SVR	0.82 $\pm$ 0.03	0.00 $\pm$ 0.00	0.91 $\pm$ 0.02	1.00 $\pm$ 0.00	0.90 $\pm$ 0.02
	GPR	0.88 $\pm$ 0.03	0.00 $\pm$ 0.00	0.94 $\pm$ 0.02	1.00 $\pm$ 0.00	0.94 $\pm$ 0.02
	KRR	0.88 $\pm$ 0.03	0.06 $\pm$ 0.02	0.91 $\pm$ 0.02	0.93 $\pm$ 0.02	0.91 $\pm$ 0.02
	PCA- T^2	0.72 $\pm$ 0.04	0.08 $\pm$ 0.02	0.82 $\pm$ 0.02	0.90 $\pm$ 0.03	0.80 $\pm$ 0.03
	PCA-SPE	0.89 $\pm$ 0.03	0.00 $\pm$ 0.00	0.95 $\pm$ 0.01	1.00 $\pm$ 0.00	0.94 $\pm$ 0.02

not significant because of the limited number of paired drift levels, all comparisons showed large rank-biserial effect sizes. This indicates that the performance improvements are practically meaningful and supports D-sPLS as the most robust monitoring method for gradual soil-moisture drift faults.

Overall, the first case study confirms the effectiveness of D-sPLS for reliable soil-moisture sensor monitoring, delivering superior fault detection performance with low false-alarm rates across diverse fault scenarios.

Table 7: Wilcoxon signed-rank comparison of F1-score between D-sPLS and competing methods under soil-moisture drift faults.

Compared model	Mean F1	Mean $\Delta F1$	Median $\Delta F1$	W	p -value	Holm p -value	Effect size
PCA-SPE	0.891	0.036	0.016	10.000	0.063	0.563	1.000
GPR	0.889	0.037	0.034	10.000	0.063	0.563	1.000
SparsePLS	0.886	0.040	0.031	10.000	0.063	0.563	1.000
PLS	0.882	0.045	0.037	10.000	0.063	0.563	1.000
KRR	0.854	0.072	0.075	10.000	0.063	0.563	1.000
DecisionTree	0.818	0.109	0.108	10.000	0.063	0.563	1.000
SVR	0.818	0.109	0.092	10.000	0.063	0.563	1.000
PCA-T2	0.761	0.166	0.167	10.000	0.063	0.563	1.000
ElasticNet	0.743	0.184	0.163	10.000	0.063	0.563	1.000

3.2 Application 2: Soil Quality Monitoring

3.2.1 Dataset Description and Exploratory Data Analysis

The second case study employs a multivariate soil quality dataset collected from agricultural fields in northern Greece between 2015 and 2019 by the Soil and Water Research Institute (SWRI) [51]. The dataset contains measurements from 781 sampling locations covering approximately 270 km² and was developed to support soil fertility assessment and precision agriculture. Sixteen physicochemical soil properties, including soil texture, nutrient concentrations, and chemical characteristics, were measured at each location. The monitored variables are listed in Table 8. Zinc (Zn) is selected as the response variable because it is an essential micronutrient and an important indicator of soil fertility, while the remaining 15 variables are used as predictors for model development and fault detection.

Table 8: Monitored soil physicochemical variables.

Var.	Description	Unit
S	Sand	%
C	Clay	%
Si	Silt	%
pH	Soil acidity	—
EC	Electrical conductivity	mS/cm
OM	Organic matter	ppm
CaCO ₃	Calcium carbonate	ppm
NO ₃ -N	Nitrate nitrogen	ppm
P	Phosphorus	ppm
K	Potassium	ppm
Mg	Magnesium	ppm
Fe	Iron	ppm
Zn	Zinc	ppm
Mn	Manganese	ppm
Cu	Copper	ppm
B	Boron	ppm

Figure 7 shows the distributions of the fault-free training data for the monitored soil variables. The variables exhibit heterogeneous statistical characteristics. Soil texture components have relatively broad distributions, whereas most chemical properties (e.g., EC, NO₃-N, P, K, Mg, Fe, Zn, Mn, Cu, and B) are positively skewed with long right tails. In contrast, pH is concentrated within a narrow range. These heterogeneous scales and non-Gaussian distributions motivate the use of the proposed D-sPLS framework, which employs nonparametric KDE-based threshold estimation.






Fig. 7: Distributions of the fault-free soil variables.

Figure 8 summarizes the dependency structure of the fault-free soil variables. Figure 8(a) shows the Spearman correlation coefficients, revealing several strong monotonic relationships. Sand is strongly negatively correlated with clay ($\rho = -0.64$), whereas Fe exhibits strong negative correlations with CaCO₃ ($\rho = -0.73$) and pH ($\rho = -0.72$), and a strong positive correlation with Mn ($\rho = 0.76$). Figure 8(b) shows that most moderate and strong correlations are highly significant ($p < 0.001$). Figure 8(c) presents the normalized mutual information (NMI), highlighting additional nonlinear dependencies, particularly among pH, Fe, Mn, and CaCO₃. Overall, these results demonstrate strong linear and nonlinear dependencies among the soil variables, supporting the use of the proposed D-sPLS framework to exploit their correlation structure while reducing redundancy through sparse variable selection.


Figure 9 ranks the predictor variables according to their association with Zn concentration using a composite score based on the absolute Spearman correlation and



(a) Spearman correlation coefficients.



(b) Corresponding p -values.



(c) Normalized mutual information.

Fig. 8: Dependency analysis of the fault-free soil variables: (a) Spearman correlation coefficients, (b) corresponding p -values, and (c) normalized mutual information.

normalized mutual information. Cu and P exhibit the strongest associations, indicating that they are the most informative predictors for Zn estimation. Moderate

associations are observed for Mn, Clay, EC, CaCO_3 , B, and Fe, whereas the remaining variables show weaker individual associations. Nevertheless, all variables are retained because weak individual relationships may provide complementary information in the proposed D-sPLS framework.




Fig. 9: Association ranking of the predictor variables with Zn concentration based on a composite association score.

The prediction and monitoring models were developed using only the fault-free training subset, the first 400 samples, while the last 380 samples were kept for independent testing. Hyperparameters were selected during the training stage using cross-validation or fixed model-specific settings. PLS used one latent component selected from a maximum of ten candidates with five-fold cross-validation, while SparsePLS retained eight predictors using five-fold cross-validation. Elastic Net was trained with a mixing parameter $\alpha = 0.5$, a selected regularization parameter $\lambda_{\text{EN}} = 0.10055$, and five-fold cross-validation. The ELM model used 80 hidden neurons with a hyperbolic tangent activation function and ridge regularization of 10^{-3} . SVR employed a Gaussian kernel with automatic kernel-scale selection. GPR used an ARD squared exponential kernel with sparse regression fitting and fully independent conditional prediction. KRR was implemented with an RBF kernel, kernel width $\sigma = 4.5947$, and ridge regularization of 0.01. For residual-based monitoring, the prediction residuals were transformed into EWMA charting statistics using a forgetting factor of $\eta_{\text{EWMA}} = 0.15$, and the alarm thresholds were estimated from the fault-free training data using a KDE-based thresholding strategy.

3.2.2 Scenario 1: Sensor Bias Fault Detection

In the first scenario, abrupt bias faults were introduced into the Zn sensor measurements to evaluate the sensitivity of the proposed monitoring framework to additive sensor offsets. The prediction models were trained using the fault-free data and remained unchanged during testing. The last 380 samples were used as the independent test set, with the bias fault injected from sample 150 onward. Four representative bias levels (15%, 20%, 30%, and 50%) were considered. For each bias level, the prediction residuals were monitored using the common EWMA detector with KDE-based thresholding. Detection performance was evaluated using TPR, FPR, accuracy, precision, F1-score, and detection delay. To assess statistical robustness, 95% confidence intervals were estimated from 300 bootstrap resampling runs.

Tables 9 summarize the monitoring performance under different soil-moisture sensor bias magnitudes. At the lowest bias level (15%), the proposed D-sPLS-EWMA achieves the highest detection performance (TPR = 0.89) while maintaining zero false alarms, yielding the best F1-score (0.94). SparsePLS-EWMA, PLS-EWMA, and PCA- T^2 also provide relatively good detection performance, whereas ElasticNet-EWMA, ELM-EWMA, GPR, KRR-EWMA, and PCA-SPE exhibit considerably lower sensitivity. This behavior is expected because small bias faults generate residual variations comparable to the normal prediction uncertainty, making early detection inherently challenging. As the bias magnitude increases, the proposed framework consistently improves, reaching a TPR of 0.95 and an F1-score of 0.98 for the 50% bias while preserving a zero false-alarm rate. In contrast, SparsePLS-EWMA and conventional PLS show only marginal improvement, indicating that static latent-variable models cannot fully exploit the underlying temporal information. Kernel- and neural-network-based methods (GPR, KRR-EWMA, and ELM-EWMA) become more effective for larger faults because the induced residuals are more easily separated from normal process variability. However, their performance remains below that of the proposed approach. The superior performance of D-sPLS-EWMA can be attributed to the combination of dynamic feature augmentation and sparse latent-variable modeling. By incorporating temporal information while retaining only the most informative dynamic predictors, the proposed model generates more accurate fault-free predictions and more discriminative residuals. Consequently, the EWMA detector can identify sensor bias faults with high sensitivity without increasing false alarms. Overall, the results demonstrate that exploiting both temporal dynamics and sparse variable selection leads to more reliable and robust monitoring than conventional latent-variable, regression-based, and statistical approaches.

To complement the descriptive performance metrics and verify whether the observed differences among monitoring methods are statistically meaningful, nonparametric statistical tests were applied to the F1-score results. The Friedman omnibus test was first used to examine whether the F1-score differences among the monitoring methods were statistically significant across the four Zn bias-fault magnitudes. As reported in Table ??, the test rejects the null hypothesis of equal performance, with $\chi_F^2 = 33.253$, $df = 10$, and $p = 2.469 \times 10^{-4}$. Kendall's coefficient $W = 0.831$ indicates a strong agreement in the method rankings across the considered fault levels. The

Table 9: Performance comparison under different bias fault levels. Results are reported as mean $\pm$ half-width of the 95% bootstrap confidence interval.

Bias (%)	Method	TPR	FPR	Accuracy	Precision	F1
15	D-sPLS-EWMA	0.89 $\pm$ 0.04	0.00 $\pm$ 0.00	0.93 $\pm$ 0.02	1.00 $\pm$ 0.00	0.94 $\pm$ 0.02
	SparsePLS-EWMA	0.86 $\pm$ 0.04	0.01 $\pm$ 0.01	0.91 $\pm$ 0.03	1.00 $\pm$ 0.01	0.92 $\pm$ 0.03
	PLS-EWMA	0.78 $\pm$ 0.06	0.01 $\pm$ 0.02	0.86 $\pm$ 0.04	0.99 $\pm$ 0.01	0.87 $\pm$ 0.04
	ElasticNet-EWMA	0.17 $\pm$ 0.05	0.00 $\pm$ 0.00	0.50 $\pm$ 0.05	1.00 $\pm$ 0.00	0.30 $\pm$ 0.07
	DecisionTree-EWMA	0.78 $\pm$ 0.06	0.09 $\pm$ 0.05	0.83 $\pm$ 0.04	0.93 $\pm$ 0.04	0.85 $\pm$ 0.04
	ELM-EWMA	0.31 $\pm$ 0.06	0.00 $\pm$ 0.00	0.58 $\pm$ 0.05	1.00 $\pm$ 0.00	0.48 $\pm$ 0.07
	SVR	0.78 $\pm$ 0.05	0.02 $\pm$ 0.03	0.86 $\pm$ 0.03	0.98 $\pm$ 0.02	0.87 $\pm$ 0.04
	GPR	0.18 $\pm$ 0.05	0.00 $\pm$ 0.00	0.50 $\pm$ 0.05	1.00 $\pm$ 0.00	0.30 $\pm$ 0.07
	KRR-EWMA	0.17 $\pm$ 0.05	0.00 $\pm$ 0.00	0.50 $\pm$ 0.05	1.00 $\pm$ 0.00	0.30 $\pm$ 0.07
	PCA- T^2	0.80 $\pm$ 0.05	0.00 $\pm$ 0.00	0.88 $\pm$ 0.03	1.00 $\pm$ 0.00	0.89 $\pm$ 0.03
20	PCA-SPE	0.44 $\pm$ 0.06	0.00 $\pm$ 0.00	0.66 $\pm$ 0.05	1.00 $\pm$ 0.00	0.61 $\pm$ 0.06
	D-sPLS-EWMA	0.89 $\pm$ 0.04	0.00 $\pm$ 0.00	0.93 $\pm$ 0.02	1.00 $\pm$ 0.00	0.94 $\pm$ 0.02
	SparsePLS-EWMA	0.86 $\pm$ 0.05	0.01 $\pm$ 0.01	0.91 $\pm$ 0.03	0.99 $\pm$ 0.01	0.92 $\pm$ 0.03
	PLS-EWMA	0.78 $\pm$ 0.05	0.01 $\pm$ 0.02	0.86 $\pm$ 0.03	0.99 $\pm$ 0.01	0.87 $\pm$ 0.03
	ElasticNet-EWMA	0.18 $\pm$ 0.04	0.00 $\pm$ 0.00	0.50 $\pm$ 0.05	1.00 $\pm$ 0.00	0.30 $\pm$ 0.06
	DecisionTree-EWMA	0.78 $\pm$ 0.05	0.09 $\pm$ 0.05	0.83 $\pm$ 0.04	0.93 $\pm$ 0.04	0.85 $\pm$ 0.03
	ELM-EWMA	0.32 $\pm$ 0.06	0.00 $\pm$ 0.00	0.58 $\pm$ 0.05	1.00 $\pm$ 0.00	0.48 $\pm$ 0.07
	SVR	0.78 $\pm$ 0.05	0.02 $\pm$ 0.02	0.86 $\pm$ 0.03	0.98 $\pm$ 0.02	0.87 $\pm$ 0.03
	GPR	0.39 $\pm$ 0.06	0.00 $\pm$ 0.00	0.63 $\pm$ 0.05	1.00 $\pm$ 0.00	0.56 $\pm$ 0.06
	KRR-EWMA	0.18 $\pm$ 0.05	0.00 $\pm$ 0.00	0.50 $\pm$ 0.05	1.00 $\pm$ 0.00	0.30 $\pm$ 0.07
30	PCA- T^2	0.80 $\pm$ 0.05	0.00 $\pm$ 0.00	0.88 $\pm$ 0.03	1.00 $\pm$ 0.00	0.89 $\pm$ 0.03
	PCA-SPE	0.44 $\pm$ 0.06	0.00 $\pm$ 0.00	0.66 $\pm$ 0.04	1.00 $\pm$ 0.00	0.61 $\pm$ 0.06
	D-sPLS-EWMA	0.91 $\pm$ 0.03	0.00 $\pm$ 0.00	0.94 $\pm$ 0.02	1.00 $\pm$ 0.00	0.95 $\pm$ 0.02
	SparsePLS-EWMA	0.85 $\pm$ 0.04	0.01 $\pm$ 0.01	0.91 $\pm$ 0.03	0.99 $\pm$ 0.01	0.92 $\pm$ 0.02
	PLS-EWMA	0.78 $\pm$ 0.05	0.01 $\pm$ 0.02	0.86 $\pm$ 0.03	0.99 $\pm$ 0.01	0.87 $\pm$ 0.03
	ElasticNet-EWMA	0.41 $\pm$ 0.07	0.00 $\pm$ 0.00	0.64 $\pm$ 0.05	1.00 $\pm$ 0.00	0.58 $\pm$ 0.07
	DecisionTree-EWMA	0.79 $\pm$ 0.05	0.09 $\pm$ 0.05	0.83 $\pm$ 0.04	0.93 $\pm$ 0.04	0.85 $\pm$ 0.04
	ELM-EWMA	0.32 $\pm$ 0.06	0.00 $\pm$ 0.00	0.59 $\pm$ 0.05	1.00 $\pm$ 0.00	0.49 $\pm$ 0.07
	SVR	0.78 $\pm$ 0.05	0.02 $\pm$ 0.02	0.86 $\pm$ 0.04	0.98 $\pm$ 0.02	0.87 $\pm$ 0.04
	GPR	0.41 $\pm$ 0.06	0.00 $\pm$ 0.00	0.64 $\pm$ 0.05	1.00 $\pm$ 0.00	0.58 $\pm$ 0.06
50	KRR-EWMA	0.41 $\pm$ 0.06	0.00 $\pm$ 0.00	0.64 $\pm$ 0.05	1.00 $\pm$ 0.00	0.58 $\pm$ 0.06
	PCA- T^2	0.80 $\pm$ 0.05	0.00 $\pm$ 0.00	0.88 $\pm$ 0.03	1.00 $\pm$ 0.00	0.89 $\pm$ 0.03
	PCA-SPE	0.44 $\pm$ 0.06	0.00 $\pm$ 0.00	0.66 $\pm$ 0.05	1.00 $\pm$ 0.00	0.61 $\pm$ 0.06
	D-sPLS-EWMA	0.95 $\pm$ 0.03	0.00 $\pm$ 0.00	0.97 $\pm$ 0.02	1.00 $\pm$ 0.00	0.98 $\pm$ 0.01
	SparsePLS-EWMA	0.78 $\pm$ 0.06	0.01 $\pm$ 0.01	0.86 $\pm$ 0.04	0.99 $\pm$ 0.01	0.87 $\pm$ 0.04
	PLS-EWMA	0.79 $\pm$ 0.05	0.01 $\pm$ 0.02	0.87 $\pm$ 0.03	0.99 $\pm$ 0.01	0.88 $\pm$ 0.03
	ElasticNet-EWMA	0.71 $\pm$ 0.06	0.00 $\pm$ 0.00	0.83 $\pm$ 0.03	1.00 $\pm$ 0.00	0.83 $\pm$ 0.04
	DecisionTree-EWMA	0.80 $\pm$ 0.05	0.09 $\pm$ 0.05	0.84 $\pm$ 0.03	0.93 $\pm$ 0.03	0.86 $\pm$ 0.03
	ELM-EWMA	0.77 $\pm$ 0.05	0.00 $\pm$ 0.00	0.86 $\pm$ 0.03	1.00 $\pm$ 0.00	0.87 $\pm$ 0.03
	SVR	0.79 $\pm$ 0.04	0.02 $\pm$ 0.02	0.87 $\pm$ 0.03	0.98 $\pm$ 0.02	0.88 $\pm$ 0.03
	GPR	0.78 $\pm$ 0.05	0.00 $\pm$ 0.00	0.87 $\pm$ 0.03	1.00 $\pm$ 0.00	0.88 $\pm$ 0.03
	KRR-EWMA	0.78 $\pm$ 0.06	0.00 $\pm$ 0.00	0.87 $\pm$ 0.04	1.00 $\pm$ 0.00	0.88 $\pm$ 0.04
	PCA- T^2	0.80 $\pm$ 0.05	0.00 $\pm$ 0.00	0.88 $\pm$ 0.03	1.00 $\pm$ 0.00	0.89 $\pm$ 0.03
	PCA-SPE	0.44 $\pm$ 0.06	0.00 $\pm$ 0.00	0.66 $\pm$ 0.05	1.00 $\pm$ 0.00	0.61 $\pm$ 0.05

mean-rank analysis further confirms that D-sPLS-EWMA provides the best overall performance, achieving the lowest mean rank of 1.000 and the highest mean F1-score of 0.9528.

Pairwise statistical comparisons were further carried out to identify which methods were significantly outperformed by the proposed approach. For this purpose, paired Wilcoxon signed-rank tests were applied between D-sPLS-EWMA and each competing method using the paired bootstrap F1-score samples, and Holm correction was used to control the family-wise error rate due to multiple comparisons. Here, $\Delta F1$ denotes the

paired F1-score difference between D-sPLS-EWMA and the compared method, while the effect size corresponds to the rank-biserial effect. As reported in Table 10, D-sPLS-EWMA provides positive mean F1-score differences against all benchmark methods. The improvement is largest against ElasticNet-EWMA, KRR-EWMA, ELM-EWMA, GPR, and PCA-SPE, with mean F1 gains ranging from 0.3414 to 0.4499. More importantly, D-sPLS-EWMA also remains superior to the strongest competing methods, including SparsePLS-EWMA, PCA- T^2 , PLS-EWMA, and SVR, with mean F1 gains of 0.0435, 0.0661, 0.0791, and 0.0809, respectively. All reported p -values are Holm-adjusted and are lower than 0.001, confirming that these improvements are statistically significant. In addition, the rank-biserial effect sizes are close to their maximum, ranging from 0.96 to 1.00, indicating that the superiority of D-sPLS-EWMA is not due to isolated bootstrap samples but is consistently observed across resampling runs.

Table 10: Paired Wilcoxon signed-rank comparison between D-sPLS-EWMA and competing methods for the Zn bias-fault scenario.

Compared model	Mean F1	Mean $\Delta F1$	W	Holm p	Effect
ElasticNet-EWMA	0.5027	0.4499	720600	< 0.001	1.00
KRR-EWMA	0.5134	0.4392	720600	< 0.001	1.00
ELM-EWMA	0.5788	0.3738	720600	< 0.001	1.00
GPR	0.5799	0.3726	720600	< 0.001	1.00
PCA-SPE	0.6112	0.3414	720600	< 0.001	1.00
DecisionTree-EWMA	0.8531	0.0995	720600	< 0.001	1.00
SVR	0.8717	0.0809	720600	< 0.001	1.00
PLS-EWMA	0.8735	0.0791	720600	< 0.001	1.00
PCA- T^2	0.8865	0.0661	720600	< 0.001	1.00
SparsePLS-EWMA	0.9091	0.0435	703198	< 0.001	0.96

3.2.3 Scenario 2: Drift Sensor Fault Detection

To further evaluate its robustness, the proposed framework is assessed under gradual sensor drift faults. Drift was introduced into the Zn sensor measurements to emulate progressive sensor degradation. The prediction models were trained using the fault-free data and remained unchanged during testing. The last 380 samples were used as the independent test set, with the drift injected from sample 150 onward. Three drift severities (10%, 15%, and 20% of the Zn operating range per 100 samples) were considered. Detection performance was evaluated using TPR, FPR, accuracy, precision, F1-score, and detection delay. Statistical robustness was assessed using 300 bootstrap resampling runs to estimate the corresponding 95% confidence intervals.

Table 11 summarizes the monitoring performance under different Zn sensor drift severities. Overall, D-sPLS-EWMA provides the best balance between detection sensitivity and false-alarm rate. At the 10% drift level, it achieves an F1-score of 0.88 with zero false alarms, outperforming SparsePLS-EWMA and the remaining competing methods. As drift severity increases, the proposed method consistently maintains the best overall performance, achieving an F1-score of 0.92 at the 20% drift level,

with perfect precision and zero false alarms. SparsePLS-EWMA remains the closest competitor, whereas ELM-EWMA, GPR, and KRR-EWMA improve as the drift becomes more pronounced. In contrast, PCA-SPE shows limited sensitivity, with nearly unchanged performance across all drift levels. The narrow bootstrap confidence intervals further confirm the stability and robustness of the proposed framework.

Table 11: Performance under Zn sensor drift faults (mean $\pm$ 95% bootstrap CI half-width).

Drift (%)	Method	TPR	FPR	Accuracy	Precision	F1
10	D-sPLS-EWMA	0.79 $\pm$ 0.06	0.00 $\pm$ 0.00	0.87 $\pm$ 0.03	1.00 $\pm$ 0.00	0.88 $\pm$ 0.03
	SparsePLS-EWMA	0.78 $\pm$ 0.05	0.00 $\pm$ 0.00	0.87 $\pm$ 0.03	1.00 $\pm$ 0.00	0.88 $\pm$ 0.03
	PLS-EWMA	0.78 $\pm$ 0.06	0.01 $\pm$ 0.02	0.86 $\pm$ 0.04	0.99 $\pm$ 0.01	0.87 $\pm$ 0.04
	ElasticNet-EWMA	0.71 $\pm$ 0.06	0.00 $\pm$ 0.00	0.83 $\pm$ 0.04	1.00 $\pm$ 0.00	0.83 $\pm$ 0.04
	DecisionTree-EWMA	0.80 $\pm$ 0.05	0.09 $\pm$ 0.05	0.84 $\pm$ 0.03	0.93 $\pm$ 0.04	0.86 $\pm$ 0.04
	ELM-EWMA	0.58 $\pm$ 0.06	0.00 $\pm$ 0.00	0.74 $\pm$ 0.04	1.00 $\pm$ 0.00	0.73 $\pm$ 0.05
	SVR	0.79 $\pm$ 0.05	0.02 $\pm$ 0.03	0.86 $\pm$ 0.03	0.98 $\pm$ 0.02	0.88 $\pm$ 0.03
	GPR	0.72 $\pm$ 0.05	0.00 $\pm$ 0.00	0.83 $\pm$ 0.04	1.00 $\pm$ 0.00	0.84 $\pm$ 0.04
	KRR-EWMA	0.72 $\pm$ 0.06	0.00 $\pm$ 0.00	0.83 $\pm$ 0.04	1.00 $\pm$ 0.00	0.84 $\pm$ 0.04
	PCA- T^2	0.81 $\pm$ 0.05	0.03 $\pm$ 0.03	0.87 $\pm$ 0.03	0.98 $\pm$ 0.02	0.88 $\pm$ 0.03
	PCA-SPE	0.49 $\pm$ 0.06	0.00 $\pm$ 0.00	0.69 $\pm$ 0.05	1.00 $\pm$ 0.00	0.66 $\pm$ 0.06
15	D-sPLS-EWMA	0.84 $\pm$ 0.05	0.00 $\pm$ 0.00	0.90 $\pm$ 0.03	1.00 $\pm$ 0.00	0.91 $\pm$ 0.03
	SparsePLS-EWMA	0.79 $\pm$ 0.05	0.00 $\pm$ 0.00	0.87 $\pm$ 0.03	1.00 $\pm$ 0.00	0.88 $\pm$ 0.03
	PLS-EWMA	0.79 $\pm$ 0.05	0.01 $\pm$ 0.02	0.87 $\pm$ 0.03	0.99 $\pm$ 0.01	0.88 $\pm$ 0.03
	ElasticNet-EWMA	0.74 $\pm$ 0.06	0.00 $\pm$ 0.00	0.84 $\pm$ 0.04	1.00 $\pm$ 0.00	0.85 $\pm$ 0.04
	DecisionTree-EWMA	0.81 $\pm$ 0.05	0.09 $\pm$ 0.05	0.85 $\pm$ 0.04	0.93 $\pm$ 0.04	0.87 $\pm$ 0.03
	ELM-EWMA	0.78 $\pm$ 0.05	0.00 $\pm$ 0.00	0.87 $\pm$ 0.03	1.00 $\pm$ 0.00	0.88 $\pm$ 0.03
	SVR	0.81 $\pm$ 0.05	0.02 $\pm$ 0.02	0.87 $\pm$ 0.03	0.98 $\pm$ 0.02	0.89 $\pm$ 0.03
	GPR	0.77 $\pm$ 0.06	0.00 $\pm$ 0.00	0.86 $\pm$ 0.04	1.00 $\pm$ 0.00	0.87 $\pm$ 0.04
	KRR-EWMA	0.78 $\pm$ 0.05	0.00 $\pm$ 0.00	0.87 $\pm$ 0.03	1.00 $\pm$ 0.00	0.88 $\pm$ 0.03
	PCA- T^2	0.81 $\pm$ 0.05	0.03 $\pm$ 0.02	0.87 $\pm$ 0.03	0.98 $\pm$ 0.02	0.88 $\pm$ 0.03
	PCA-SPE	0.49 $\pm$ 0.07	0.00 $\pm$ 0.00	0.69 $\pm$ 0.04	1.00 $\pm$ 0.00	0.66 $\pm$ 0.06
20	D-sPLS-EWMA	0.84 $\pm$ 0.05	0.00 $\pm$ 0.00	0.91 $\pm$ 0.03	1.00 $\pm$ 0.00	0.92 $\pm$ 0.03
	SparsePLS-EWMA	0.84 $\pm$ 0.04	0.00 $\pm$ 0.00	0.91 $\pm$ 0.03	1.00 $\pm$ 0.00	0.92 $\pm$ 0.02
	PLS-EWMA	0.81 $\pm$ 0.05	0.01 $\pm$ 0.02	0.88 $\pm$ 0.03	0.99 $\pm$ 0.01	0.89 $\pm$ 0.03
	ElasticNet-EWMA	0.77 $\pm$ 0.05	0.00 $\pm$ 0.00	0.86 $\pm$ 0.03	1.00 $\pm$ 0.00	0.87 $\pm$ 0.03
	DecisionTree-EWMA	0.84 $\pm$ 0.05	0.09 $\pm$ 0.05	0.87 $\pm$ 0.03	0.93 $\pm$ 0.04	0.89 $\pm$ 0.03
	ELM-EWMA	0.79 $\pm$ 0.05	0.00 $\pm$ 0.00	0.87 $\pm$ 0.03	1.00 $\pm$ 0.00	0.88 $\pm$ 0.03
	SVR	0.84 $\pm$ 0.05	0.02 $\pm$ 0.02	0.89 $\pm$ 0.03	0.98 $\pm$ 0.02	0.91 $\pm$ 0.03
	GPR	0.80 $\pm$ 0.05	0.00 $\pm$ 0.00	0.88 $\pm$ 0.03	1.00 $\pm$ 0.00	0.89 $\pm$ 0.03
	KRR-EWMA	0.80 $\pm$ 0.06	0.00 $\pm$ 0.00	0.88 $\pm$ 0.04	1.00 $\pm$ 0.00	0.89 $\pm$ 0.04
	PCA- T^2	0.82 $\pm$ 0.05	0.03 $\pm$ 0.02	0.88 $\pm$ 0.03	0.98 $\pm$ 0.02	0.89 $\pm$ 0.03
	PCA-SPE	0.49 $\pm$ 0.06	0.00 $\pm$ 0.00	0.69 $\pm$ 0.04	1.00 $\pm$ 0.00	0.66 $\pm$ 0.05

The detection delay results further demonstrate the effectiveness of the proposed framework. D-sPLS-EWMA detects the drift after 49, 38, and 36 samples for the 10%, 15%, and 20% drift levels, respectively, indicating faster detection as the drift becomes more pronounced. At the lowest drift level, PCA- T^2 and DecisionTree-EWMA achieve slightly shorter delays (45 and 47 samples, respectively), but both

exhibit nonzero false-alarm rates. SparsePLS-EWMA provides comparable performance, particularly for the 20% drift where it also detects the fault after 36 samples. In contrast, ElasticNet-EWMA, ELM-EWMA, GPR, KRR-EWMA, and especially PCA-SPE require substantially longer detection times, with PCA-SPE consistently exceeding 117 samples. Overall, D-sPLS-EWMA offers one of the shortest detection delays while simultaneously maintaining zero false alarms, providing a favorable compromise between rapid fault detection and monitoring reliability.

Table 12 reports the paired Wilcoxon signed-rank test used to assess whether D-sPLS-EWMA provides statistically superior F1-scores compared with the competing monitoring strategies under the drift-fault scenarios. The test was applied to paired bootstrap F1-score samples, with D-sPLS-EWMA used as the reference method. The reported p -values were adjusted using the Holm correction to account for multiple comparisons.

The results show that D-sPLS-EWMA achieves the highest mean F1-score, equal to 0.9020, and outperforms all compared methods with positive mean differences. The largest improvements are observed against PCA-SPE, ELM-EWMA, ElasticNet-EWMA, and GPR, with mean F1 gains of 0.2451, 0.0730, 0.0494, and 0.0389, respectively. The Holm-adjusted p -values are far below 0.05 for all comparisons, indicating that the superiority of D-sPLS-EWMA is statistically significant after multiple-comparison correction. The rank-biserial effect sizes are also positive in all cases, confirming that the improvement is not only statistically significant but also consistent across the paired bootstrap samples. Although the margin over SparsePLS-EWMA is smaller, with a mean F1 difference of 0.0098, the corrected p -value remains significant, supporting the added value of the dynamic sparse formulation for drift-fault monitoring.

Table 12: Wilcoxon signed-rank comparison of D-sPLS-EWMA against the competing methods under drift-fault scenarios.

Compared method	Mean F1 D-sPLS	Mean F1 method	Mean diff.	W stat.	Holm p -value	Effect
PCA-SPE	0.9020	0.6569	0.2451	405450.0	< 0.001	1.0000
ELM-EWMA	0.9020	0.8290	0.0730	403090.0	< 0.001	0.9884
ElasticNet-EWMA	0.9020	0.8526	0.0494	404317.0	< 0.001	0.9944
GPR	0.9020	0.8630	0.0389	399839.0	< 0.001	0.9767
KRR-EWMA	0.9020	0.8672	0.0348	396352.0	< 0.001	0.9595
DecisionTree-EWMA	0.9020	0.8713	0.0307	388674.0	< 0.001	0.9172
PLS-EWMA	0.9020	0.8796	0.0223	361244.0	< 0.001	0.7859
PCA- T^2	0.9020	0.8858	0.0162	331316.0	< 0.001	0.6379
SVR	0.9020	0.8887	0.0133	316116.0	< 0.001	0.5593
SparsePLS-EWMA	0.9020	0.8922	0.0098	282619.0	< 0.001	0.3972

3.3 ICE-Based Fault Diagnosis

Although the proposed D-sPLS framework effectively detects sensor faults through residual monitoring, it does not explain the cause of an alarm. To improve interpretability, an ICE-based diagnostic layer is applied after fault detection [39, 40]. Unlike the EWMA-KDE detector, which provides a binary decision, the ICE analysis visualizes abnormal sensor behavior and helps assess fault severity. Since similar observations were obtained for both applications, the ICE analysis is presented only for the second case study. The analysis is performed with respect to Cu because it exhibits the strongest association with Zn (Fig. 9). ICE curves are generated by varying the Cu concentration while keeping the remaining predictors fixed. The resulting ICE envelope represents the expected Zn response under normal operating conditions, while the corresponding PDP summarizes the global Cu–Zn relationship. During monitoring, measured Zn values are superimposed on the learned envelope. Samples inside the envelope are considered consistent with normal operation, whereas those outside indicate abnormal sensor behavior. Figure 10 illustrates the ICE/PDP diagnosis under Zn sensor bias faults. During fault-free operation, most measurements remain within the ICE envelope. After introducing a 15% bias, many Zn measurements exceed the upper envelope, and this separation increases with larger bias levels. Under a 50% bias, nearly all measurements fall outside the expected region, indicating severe sensor degradation. These results show that the ICE representation not only confirms the detected fault but also visualizes its severity through the increasing deviation from the learned Cu–Zn relationship.




Fig. 10: ICE/PDP-based diagnosis of Zn sensor bias faults. The shaded region represents the 5th–99th percentile ICE envelope.

Figure 11 illustrates the ICE/PDP diagnosis under Zn sensor drift faults. During fault-free operation, most samples remain within the ICE envelope, confirming consistency with the learned Cu–Zn relationship. As the drift severity increases from 10% to 20% per 100 samples, an increasing number of Zn measurements move beyond the upper ICE boundary. Unlike the abrupt shift caused by bias faults, drift produces a gradual departure from the expected predictor–response relationship. This progressive migration outside the ICE envelope provides a clear visual indication of sensor degradation. Thus, the ICE/PDP analysis complements the EWMA detector by confirming the detected fault and revealing its progression and severity.




Fig. 11: ICE/PDP-based diagnosis of Zn sensor drift faults. The shaded region represents the 5th–99th percentile ICE envelope. Green and red markers denote measurements inside and outside the expected response region, respectively.

Overall, the proposed ICE-based diagnostic layer complements the D-sPLS monitoring framework by providing an interpretable explanation of detected faults. It helps distinguish abrupt bias from gradual drift, visualize fault severity, and support timely maintenance in smart soil monitoring systems.

3.4 Computational Efficiency

Computational efficiency is an important consideration for real-time deployment in smart agricultural sensing systems. The computational cost of each monitoring strategy was evaluated during the testing phase, in which only the online monitoring step (prediction, monitoring statistic computation, and decision-making) was timed. Training time was excluded because model development is performed offline. All experiments

were conducted on a laptop equipped with an Intel Core i9-9880H CPU (2.30 GHz), 16 GB RAM, running Windows 10 Education (64-bit), without GPU acceleration. The reported execution times correspond to the average of 200 independent runs, and the error bars represent the half-width of the 95% confidence interval.




Fig. 12: Average testing execution time of the investigated monitoring methods. Error bars denote the 95% confidence interval.

Figure 12 shows that all methods require less than 0.31 ms per monitoring cycle, confirming their suitability for real-time soil sensor monitoring. PCA- T^2 (0.105 ms) and PCA-SPE (0.109 ms) are the fastest methods, followed by ELM-EWMA (0.187 ms), ElasticNet-EWMA (0.188 ms), KRR-EWMA (0.195 ms), and PLS-EWMA (0.196 ms). SparsePLS-EWMA (0.208 ms), GPR (0.216 ms), SVR (0.246 ms), and DecisionTree-EWMA (0.256 ms) require slightly longer execution times. The proposed D-sPLS-EWMA has the highest average execution time (0.300 ms) because it incorporates dynamic feature construction and sparse latent-variable modeling. However, the execution time remains well below one millisecond, making the additional computational cost negligible for practical smart soil monitoring systems. These results demonstrate that D-sPLS-EWMA achieves superior monitoring performance while maintaining good computational efficiency.

4 Conclusion

This paper presented a semi-supervised D-sPLS-EWMA framework for sensor fault detection in smart agricultural monitoring systems. The proposed approach combines dynamic sparse partial least squares modeling with EWMA residual monitoring and KDE-based threshold estimation to capture temporal dependencies while providing robust, data-driven decision limits. The framework was validated using two agricultural case studies involving soil moisture and soil quality monitoring under

multiple sensor-bias and sensor-drift fault scenarios. The experimental results demonstrated that D-sPLS-EWMA consistently outperformed representative latent-variable, regression-based, and statistical monitoring methods, particularly for incipient faults, while maintaining low false-alarm rates. Bootstrap analysis confirmed the robustness of the proposed framework. The ICE-based diagnostic layer improved fault interpretability, while the computational cost remained below 0.31 ms per monitoring cycle, enabling real-time deployment.

The proposed framework provides a flexible foundation for intelligent sensor monitoring in precision agriculture. Future research will focus on adaptive online model updating, detection and isolation of multiple concurrent sensor faults, validation on larger and more diverse agricultural sensing platforms, and integration with edge-enabled IoT systems for real-time deployment. Extending the framework with uncertainty quantification and probabilistic decision-making also represents a promising direction for further improving monitoring reliability and supporting next-generation smart agricultural monitoring systems.

Conflict-of-Interest Statement

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.

Data Availability

The datasets used in this study are publicly available from the sources cited in the manuscript.

Application 1 data:

{<https://iee-dataport.org/documents/weather-monitoring-station-farms-and-agriculture> [49]}

Application 2 data:

{<https://data.mendeley.com/datasets/r7tjn68rmw/1>[51]}

Funding Declaration

No funding was recieved for conducting this study.

References

- [1] Mansoor, S., Iqbal, S., Popescu, S.M., Kim, S.L., Chung, Y.S., Baek, J.-H.: Integration of smart sensors and iot in precision agriculture: trends, challenges and future prospectives. *Frontiers in Plant Science* **16**, 1587869 (2025)
- [2] Sophocleous, M., Karkotis, A., Papasavva, A., Goldberger, M., Vassiliou, L., Ros-Lis, J.V., Shacham-Diamand, Y., Georgiou, J.: A stand-alone, in situ, soil quality sensing system for precision agriculture. *IEEE Transactions on AgriFood Electronics* **2**(1), 43–50 (2024)

- [3] Song, W., Li, X., Zhang, J., Huang, J., Wang, J., Feng, W., Guo, Y., Ren, L.: Soil sensors in smart agriculture: Multi-type monitoring technologies and ecological development pathways. *Agriculture* **16**(3), 359 (2026)
- [4] Hummel, J., Gaultney, L., Sudduth, K.: Soil property sensing for site-specific crop management. *Computers and Electronics in Agriculture* **14**(2-3), 121–136 (1996)
- [5] Adamchuk, V.I., Hummel, J.W., Morgan, M.T., Upadhyaya, S.K.: On-the-go soil sensors for precision agriculture. *Computers and electronics in agriculture* **44**(1), 71–91 (2004)
- [6] Soussi, A., Zero, E., Sacile, R., Trincherro, D., Fossa, M.: Smart sensors and smart data for precision agriculture: a review. *Sensors* **24**(8), 2647 (2024)
- [7] Zhang, X., Feng, G., Sun, X.: Advanced technologies of soil moisture monitoring in precision agriculture: A review. *Journal of Agriculture and Food Research* **18**, 101473 (2024)
- [8] Elashmawy, R., Uysal, I.: Precision agriculture using soil sensor driven machine learning for smart strawberry production. *Sensors* **23**(4), 2247 (2023)
- [9] Wang, Z., Zhang, L., Guo, J., Huang, Y., Ping, L., Zheng, W., Zhangzhong, L.: Lgca-net: A time-series anomaly detection method for multiple types of agricultural sensors. *Information Processing in Agriculture* (2026)
- [10] Sreeram, R., Krishna, S.A., Kumar, A.S., Remya, S., Cho, Y.Y.: Soil moisture monitoring technologies in smart agriculture: A comprehensive review. *Farming System*, 100189 (2025)
- [11] Catalano, C., Paiano, L., Calabrese, F., Cataldo, M., Mancarella, L., Tommasi, F.: Anomaly detection in smart agriculture systems. *Computers in Industry* **143**, 103750 (2022)
- [12] Deforce, B., Baesens, B., Diels, J., Asensio, E.S.: Self-supervised anomaly detection of rogue soil moisture sensors. *arXiv preprint arXiv:2305.05495* (2023)
- [13] Bandaru, L., Irigireddy, B.C., Davis, B., *et al.*: Deepqc: A deep learning system for automatic quality control of in-situ soil moisture sensor time series data. *Smart Agricultural Technology* **8**, 100514 (2024)
- [14] Choe, H.-O., Lee, M.-H.: Artificial intelligence-based anomaly detection technology for equipment condition monitoring in smart farms. *Applied Sciences* **15**(23), 12843 (2025)
- [15] Chalapathy, R., Chawla, S.: Deep learning for anomaly detection: A survey. *arXiv preprint arXiv:1901.03407* (2019)
- [16] Gkountakos, K., Ioannidis, K., Demestichas, K., Vrochidis, S., Kompatsiaris, I.:

- A comprehensive review of deep learning-based anomaly detection methods for precision agriculture. *IEEE Access* **12**, 197715–197733 (2024)
- [17] Catalano, C., Paiano, L., Calabrese, F., Cataldo, M., Mancarella, L., Tommasi, F.: Anomaly detection in smart agriculture systems. *Computers in Industry* **143**, 103750 (2022)
 - [18] Harrou, F., Sun, Y., Hering, A.S., Madakyaru, M., *et al.*: Statistical Process Monitoring Using Advanced Data-driven and Deep Learning Approaches: Theory and Practical Applications. Elsevier, Cambridge, MA (2020)
 - [19] Park, Y.-J., Fan, S.-K.S., Hsu, C.-Y.: A review on fault detection and process diagnostics in industrial processes. *Processes* **8**(9), 1123 (2020)
 - [20] Melo, A., Câmara, M.M., Pinto, J.C.: Data-driven process monitoring and fault diagnosis: A comprehensive survey. *Processes* **12**(2), 251 (2024)
 - [21] Zhou, Y., Wang, K., Zhang, Y., Liang, D., Jia, L.: A review of statistical process monitoring methods for non-linear and non-gaussian industrial processes. *The Canadian Journal of Chemical Engineering* **103**(7), 3092–3119 (2025)
 - [22] Harrou, F., Dorbane, A., Taghezouit, B., Sun, Y.: Deep Anomaly Detection in Photovoltaic Systems: A Comprehensive Review and Avenues for Enhancement. Zenodo (2026). <https://doi.org/10.5281/zenodo.19193218> . <https://doi.org/10.5281/zenodo.19193218>
 - [23] Chun, H., Keleş, S.: Sparse partial least squares regression for simultaneous dimension reduction and variable selection. *Journal of the Royal Statistical Society Series B: Statistical Methodology* **72**(1), 3–25 (2010)
 - [24] Dong, Y., Qin, S.J.: Dynamic latent variable analytics for process operations and control. *Computers & Chemical Engineering* **114**, 69–80 (2018)
 - [25] Qin, S.J., Dong, Y., Zhu, Q., Wang, J., Liu, Q.: Bridging systems theory and data science: A unifying review of dynamic latent variable analytics and process monitoring. *Annual Reviews in Control* **50**, 29–48 (2020)
 - [26] Li, G., Qin, S.J., Zhou, D.: A new method of dynamic latent-variable modeling for process monitoring. *IEEE Transactions on Industrial Electronics* **61**(11), 6438–6445 (2014)
 - [27] Fernandes, N.C., Rato, T.J., Reis, M.S.: Modeling in the observable or latent space? a comparison of dynamic latent variable based monitoring methods for sensor fault detection. *Chemometrics and Intelligent Laboratory Systems* **231**, 104684 (2022)
 - [28] Kini, K.R., Harrou, F., Madakyaru, M., Sun, Y.: Enhancing fault detection

- in wastewater treatment plants: a multi-scale principal component analysis approach with the kantorovich distance. *Environmental Science: Water Research & Technology* **11**(5), 1215–1232 (2025)
- [29] Harrou, F., Kini, K.R., Madakyaru, M., Sun, Y.: Sensor fault detection and diagnosis in photovoltaic systems using hellinger distance and individual conditional expectation analysis. *Solar Energy* **298**, 113633 (2025)
 - [30] Tang, Y., Ma, P., Li, L., Liu, X., Liu, Y., Wang, Q.: Incipient fault detection based on double kullback–leibler divergence kldattw improved by a self-attention mechanism. *Reliability Engineering & System Safety* **264**, 111247 (2025)
 - [31] Xu, C., Huang, D., Yu, G., Liu, Y.: Distributed process monitoring of the large-scale system using spatio-temporal-causality and wasserstein-distance-based canonical variate analysis. *Journal of Process Control* **146**, 103367 (2025)
 - [32] Li, T., Han, Y., Xu, W., Geng, Z.: Novel adaptive fault detection method based on kernel entropy component analysis integrating moving window of dissimilarity for nonlinear dynamic processes. *Journal of Process Control* **125**, 1–18 (2023)
 - [33] Fan, W., Zhu, Q., Ren, S., Zhang, L., Si, F.: Dynamic probabilistic predictable feature analysis for multivariate temporal process monitoring. *IEEE Transactions on Control Systems Technology* **30**(6), 2573–2584 (2022)
 - [34] Chun, H., Keleş, S.: Sparse partial least squares regression for simultaneous dimension reduction and variable selection. *Journal of the Royal Statistical Society Series B: Statistical Methodology* **72**(1), 3–25 (2010)
 - [35] Rosipal, R., Krämer, N.: Overview and recent advances in partial least squares. In: *International Statistical and Optimization Perspectives Workshop” Subspace, Latent Structure and Feature Selection”*, pp. 34–51 (2005). Springer
 - [36] Harrou, F., Nounou, M.N.: Monitoring linear antenna arrays using an exponentially weighted moving average-based fault detection scheme. *Systems Science & Control Engineering: An Open Access Journal* **2**(1), 433–443 (2014)
 - [37] Montgomery, D.C.: *Introduction to Statistical Quality Control*, 8th edn. John Wiley & Sons, Hoboken, NJ (2020)
 - [38] Silverman, B.W.: *Density Estimation for Statistics and Data Analysis*. Routledge, London (2018)
 - [39] Goldstein, A., Kapelner, A., Bleich, J., Pitkin, E.: Peeking inside the black box: Visualizing statistical learning with plots of individual conditional expectation. *journal of Computational and Graphical Statistics* **24**(1), 44–65 (2015)

- [40] Harrou, F., Kini, K.R., Madakyaru, M., Sun, Y.: Sensor fault detection and diagnosis in photovoltaic systems using hellinger distance and individual conditional expectation analysis. *Solar Energy* **298**, 113633 (2025)
- [41] Zou, H., Hastie, T.: Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society Series B: Statistical Methodology* **67**(2), 301–320 (2005)
- [42] Breiman, L., Friedman, J., Olshen, R.A., Stone, C.J.: *Classification and Regression Trees*. Chapman and Hall/CRC, New York (2017)
- [43] Huang, G.-B., Zhu, Q.-Y., Siew, C.-K.: Extreme learning machine: theory and applications. *Neurocomputing* **70**(1-3), 489–501 (2006)
- [44] Wu, J., Wang, Y.-G.: A working likelihood approach to support vector regression with a data-driven insensitivity parameter. *International Journal of Machine Learning and Cybernetics* **14**(3), 929–945 (2023)
- [45] Kanagawa, M., Hennig, P., Sejdinovic, D., Sriperumbudur, B.K.: Gaussian processes and kernel methods: A review on connections and equivalences. *arXiv preprint arXiv:1807.02582* (2018)
- [46] Wang, W., Jing, B.-Y.: Gaussian process regression: Optimality, robustness, and relationship with kernel ridge regression. *Journal of Machine Learning Research* **23**(193), 1–67 (2022)
- [47] Vu, K., Snyder, J.C., Li, L., Rupp, M., Chen, B.F., Khelif, T., Müller, K.-R., Burke, K.: Understanding kernel ridge regression: Common behaviors from simple functions to density functionals. *International Journal of Quantum Chemistry* **115**(16), 1115–1128 (2015)
- [48] Joe Qin, S.: Statistical process monitoring: basics and beyond. *Journal of Chemometrics: A Journal of the Chemometrics Society* **17**(8-9), 480–502 (2003)
- [49] Juan, C., Steven, S., Víctor, A., Rebeca, E.: Weather Monitoring Station For Farms And Agriculture. <https://doi.org/10.21227/mdfs-ya42> . <https://ieee-dataport.org/documents/weather-monitoring-station-farms-and-agriculture>
- [50] Demšar, J.: Statistical comparisons of classifiers over multiple data sets. *Journal of Machine learning research* **7**(Jan), 1–30 (2006)
- [51] Tziachris, P.: Soil Data Grevena. <https://doi.org/10.17632/r7tjn68rmw.1> . <https://data.mendeley.com/datasets/r7tjn68rmw/1>