

Article

Uncertainty-Aware Multimodal Fusion and Bayesian Decision-Making for DSS

Vesna Antoska Knights ^{1,*}, Marija Prchkovska ², Luka Krašnjak ³ and Jasenka Gajdoš Kljusurić ⁴¹ Faculty of Technology and Technical Sciences, University “St. Kliment Ohridski”—Bitola, 1400 Veles, North Macedonia² Faculty of Computer Science, Goce Delcev University—Stip, 2000 Štip, North Macedonia; maria.prchkovska@gmail.com³ Faculty of Electrical Engineering and Computing, University of Zagreb, Unska 3, 10000 Zagreb, Croatia; luka.krasnjak@fer.hr⁴ Faculty of Food Technology and Biotechnology, University of Zagreb, Pierottijeva 6, 10000 Zagreb, Croatia; jasenka.gajdos@pbf.unizg.hr

* Correspondence: vesna.knights@uklo.edu.mk

Abstract

Uncertainty-aware decision-making increasingly relies on multimodal sensing pipelines that must fuse correlated measurements, propagate uncertainty, and trigger reliable control actions. This study develops a unified mathematical framework for multimodal data fusion and Bayesian decision-making under uncertainty. The approach integrates adaptive Covariance Intersection (aCI) for correlation-robust sensor fusion, a Gaussian state-space backbone with Kalman filtering, heteroskedastic Bayesian regression with full posterior sampling via an affine-invariant MCMC sampler, and a Bayesian likelihood-ratio test (LRT) coupled to a risk-sensitive proportional-derivative (PD) control law. Theoretical guarantees are provided by bounding the state covariance under stability conditions, establishing convexity of the aCI weight optimization on the simplex, and deriving a Bayes-risk-optimal decision threshold for the LRT under symmetric Gaussian likelihoods. A proof-of-concept agro-environmental decision-support application is considered, where heterogeneous data streams (IoT soil sensors, meteorological stations, and drone-derived vegetation indices) are fused to generate early-warning alarms for crop stress and to adapt irrigation and fertilization inputs. The proposed pipeline reduces predictive variance and sharpens posterior credible intervals (up to 34% narrower 95% intervals and 44% lower NLL/Brier score under heteroskedastic modeling), while a Bayesian uncertainty-aware controller achieves 14.2% lower water usage and 35.5% fewer false stress alarms compared to a rule-based strategy. The framework is mathematically grounded yet domain-independent, providing a probabilistic pipeline that propagates uncertainty from raw multimodal data to operational control actions, and can be transferred beyond agriculture to robotics, signal processing, and environmental monitoring applications.



Received: 12 December 2025

Revised: 3 January 2026

Accepted: 7 January 2026

Published: 20 January 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

Keywords: multimodal data fusion; Bayesian inference; covariance intersection; uncertainty quantification; decision theory

1. Introduction

Uncertainty quantification and probabilistic modeling have become fundamental aspects of mathematical systems analysis, particularly in domains with data coming from multiple heterogeneous sources. Traditional deterministic formulations are often insufficient when correlations between measurement errors are unknown or when the uncertainty

in the data varies dynamically. Recent advances propose multimodal uncertainty quantification frameworks designed to address these challenges [1]. Applications across engineering and computational domains increasingly rely on multimodal sensing, parameter estimation, and robust decision-making under uncertainty [2]. In robotics, it has been demonstrated that symmetry-based modeling and nonlinear control contribute to improving stability and performance in uncertain environments [3]. Safety-oriented robotic systems further illustrate the importance of uncertainty-aware modeling for enhancing operational reliability and risk mitigation [4,5]. From a broader mathematical perspective, optimal-transport-based Bayesian inference methods provide rigorous tools for propagating uncertainty through complex dynamical systems [6]. Collectively, these studies highlight the growing need for unified Bayesian frameworks capable of fusing correlated multimodal observations while preserving mathematically principled uncertainty propagation.

Bayesian inference, multimodal data fusion, and decision theory have therefore been increasingly combined into unified probabilistic frameworks [6–9]. The state–space model provides a foundational structure for describing stochastic dynamical systems, where latent variables evolve over time while being observed indirectly through noisy measurements [10]. Classical estimation techniques, such as the Kalman filter, provide optimal linear–Gaussian solutions for such models [11]; however, their performance degrades when the correlation between sensors is unknown. To overcome this limitation, the Covariance Intersection (CI) method and its adaptive extension (aCI) have been proposed to fuse estimates conservatively without requiring explicit cross-correlation information [12–15]. These approaches rely on convex optimization over the space of positive-definite matrices, and guarantee bounded posterior covariance even under uncertainty [12–14].

Beyond state estimation, predictive modeling must also account for heteroskedasticity, i.e., data-dependent variance in the observation noise. Heteroskedastic Bayesian regression captures this behavior through probabilistic modeling of the conditional variance and has been widely adopted in recent stochastic modeling and machine learning research [16,17]. Parameter inference is typically achieved through Markov Chain Monte Carlo (MCMC) techniques, including affine-invariant ensemble samplers that ensure convergence in high-dimensional posterior spaces [18–20]. MCMC not only yields point estimates but also complete posterior distributions, providing interpretable uncertainty bounds essential for transparent decision-making.

Within the decision-theoretic layer, Bayesian Likelihood-Ratio Tests (LRT) translate posterior probabilities into optimal binary decisions by minimizing the Bayes risk, a criterion balancing false alarms and missed detections [21–24]. Integrating this with a risk-sensitive control law allows for dynamic systems to adapt their control inputs according to posterior mean and variance, providing stability even under uncertain operating conditions [25].

This paper presents a mathematically rigorous hybrid Bayesian framework that unifies the following:

- (i) adaptive covariance-based fusion;
- (ii) heteroskedastic Bayesian regression with full posterior sampling;
- (iii) Bayesian decision theory;
- (iv) risk-sensitive feedback control.

While demonstrated on an agricultural decision-support application, the proposed framework is domain-independent and applicable to any stochastic system requiring uncertainty-aware multimodal integration, including robotics, signal processing, and environmental modeling [26–28].

The contribution is a probabilistic fusion-to-decision pipeline that performs the following:

- (i) addresses unknown inter-sensor correlations via adaptive Covariance Intersection;

- (ii) models time-varying noise through heteroskedastic Bayesian regression;
- (iii) performs full Bayesian inference with emcee MCMC;
- (iv) optimizes decision thresholds by posterior predictive Bayes risk.

Unlike prior approaches that assume sensor independence, rely on homoskedastic noise, or decouple prediction from decision-making, the proposed framework propagates uncertainty consistently from raw multimodal data to operational actions.

Building on our previous work on ML-based crop recommendation, where six supervised classifiers achieved up to 99% accuracy when classifying 11 crop types from NPK, pH, temperature, humidity, and rainfall data [29], the present study shifts the focus from deterministic crop-type prediction to probabilistic stress assessment and uncertainty-aware irrigation control.

The main contributions of this work are as follows:

An end-to-end uncertainty-aware pipeline that propagates uncertainty from multimodal sensor fusion to probabilistic decision-making and closed-loop control.

The integration of adaptive Covariance Intersection with heteroskedastic Bayesian state-space modeling to ensure robustness under unknown inter-sensor correlations.

A Bayesian decision layer that explicitly accounts for class imbalance and asymmetric risk through Bayes-risk-optimal thresholding.

A closed-loop control strategy that leverages posterior predictive uncertainty rather than point estimates, enabling risk-sensitive actuation.

2. Materials and Methods

2.1. Multimodal Data and Latent State

The dataset contained 1100 field-level observations with 12 predictors and 4 heterogeneous sources: IoT soil sensors measuring macronutrients in soil (N, P, and K) and pH; meteorological stations recording temperature (T), humidity (H), and rainfall (Rf); and drone-borne remote-sensing imagery (RGB, multispectral, thermal) from which vegetation indices such as NDVI, SAVI, and canopy-temperature index (CTI) are derived.

The agro-climatic state vector at a discrete time step k is

$$x_k = [N, P, K, pH, T, H, Rf, NDVI, SAVI, CTI]_k^T \quad (1)$$

All state variables are assumed continuous and jointly distributed with finite second moments, i.e., $E[\|x_k\|^2] < \infty$.

2.2. State Dynamics (Forecasting Backbone)

Temporal evolution follows a stochastic linear state-space model:

$$x_{k+1} = F_k x_k + B_k u_k + \omega_k, \quad \omega_k \sim N(0, Q_k) \quad (2)$$

where, u_k denotes control or intervention inputs (e.g., irrigation/fertilization).

Matrices $F_k \in R^{n \times n}$ and $B_k \in R^{n \times r}$ describe deterministic transition and control dynamics.

For stability analysis, it is assumed that the spectral radius satisfies $\rho(F_k) < 1$ ensuring bounded propagation of the prior covariance.

Prediction (prior) moments are as follows:

$$x_{k|k-1} = F_k x_{k-1|k-1} + B_k u_{k-1}, \quad (3a)$$

$$P_{k|k-1} = F_k P_{k-1|k-1} F_k^T + Q_k \quad (3b)$$

Hypothesis 1. If $Q_k > 0$ and $\rho(F_k) < 1$ then $P_{k|k-1}$ is a uniformly bounded positive-definite sequence.

Proof. This follows directly from Lyapunov stability theory: the condition $\rho(F_k) < 1$ ensures the existence of a unique positive-definite solution to the discrete Lyapunov equation, which implies that $P_{k|k-1}$ converges to a bounded positive-definite matrix. $\square$

The linear–Gaussian state–space formulation in Equation (2) is adopted as an analytically tractable baseline. While agro–environmental processes may exhibit nonlinear and non-Gaussian behavior, the objective of this work is not to claim strict physical linearity, but to establish a mathematically grounded uncertainty-propagation pipeline in which conservative covariance bounds and decision-theoretic guarantees can be rigorously analyzed.

For nonlinear models, analogous boundedness results follow under standard assumptions of local Jacobian stability (Extended Kalman Filter (EKF)/Unscented Kalman Filter (UKF)) or bounded second moments of the state distribution (ensemble and particle-based filters), ensuring that posterior covariance remains finite and suitable for conservative fusion. In such settings, the forecasting and update steps in Equations (2)–(11) can be replaced by locally linearized or sampling-based estimators, yielding modality-specific posterior summaries in the form of mean and covariance estimates.

The adaptive Covariance Intersection (aCI) fusion layer (Section 2.4) operates directly on Gaussian posterior summaries and does not require knowledge of cross-covariances, thereby preserving conservative covariance bounds even when the underlying dynamics are nonlinear or non-Gaussian.

2.3. Observation Model and Stacking

Each modality $i = 1, \dots, M$ observes a linear projection with Gaussian noise:

$$z_k^{(i)} = H_k^{(i)} \mathbf{s}_k + v_k^{(i)}, \quad v_k^{(i)} \sim N(0, R_k^{(i)}) \quad (4)$$

where $z_k^{(i)}$ is the measurement from sensor i , $H_k^{(i)}$ is the measurement matrix for sensor i , $v_k^{(i)}$ represents measurement noise (i.e., uncertainty or random error in the soil sensor readings), and block diagonal matrix construction $R_k^{(i)} > 0$ is the noise covariance for modality i .

For fusion, measurements from M sensors are stacked as:

$$z_k = \begin{bmatrix} z_k^{(1)} \\ \vdots \\ z_k^{(M)} \end{bmatrix} \quad H_k = \begin{bmatrix} H_k^{(1)} \\ \vdots \\ H_k^{(M)} \end{bmatrix} \quad R_k = \text{blk diag}(R_k^{(1)}, \dots, R_k^{(M)}) \quad (5)$$

This enables the system to process heterogeneous sensor data within a unified fusion framework. In robustness experiments, measurement noise between heterogeneous sensors may exhibit non-zero cross-correlations. Let $v_k^{(i)}$ and $v_k^{(j)}$ denote measurement noises from modalities i and j . Their cross-sensor correlation coefficient is defined as

$$\rho_{ij} = \frac{\text{Cov}(v_k^{(i)}, v_k^{(j)})}{\sqrt{\text{Var}(v_k^{(i)}) \text{Var}(v_k^{(j)})}} \quad \rho_{ij} \in [-1, 1] \quad (6)$$

When two modalities are analyzed jointly, their joint noise covariance takes the form:

$$R_k^{(ij)} = \begin{bmatrix} \sigma_i^2 & \rho_{ij}\sigma_i\sigma_j \\ \rho_{ij}\sigma_i\sigma_j & \sigma_j^2 \end{bmatrix} \quad (7)$$

The scalar parameter ρ_{ij} quantifies the extent of dependence between sensor noise sources (e.g., soil sensors influenced by the same irrigation cycle, or drone indices affected by the same illumination changes). Note that the stacked covariance R_k in Equation (5) assumes block independence across modalities. The robustness experiments relax this assumption by injecting controlled cross-correlations ρ_{ij} into $R_k^{(ij)}$.

For Kalman fusion, standard practice assumes $\rho = 0$ (independence), while adaptive Covariance Intersection (aCI) performs correlation-agnostic fusion without requiring knowledge of ρ . Robustness analysis, therefore, varies $\rho \in \{0.0, 0.3, 0.6, 0.9\}$ to quantify how unknown correlations affect posterior uncertainty.

Kalman update:

$$S_k = H_k P_{k|k-1} H_k^T + R_k, \quad (8)$$

$$K_k = P_{k|k-1} H_k^T S_k^{-1}, \quad (9)$$

$$x_{k|k} = x_{k|k-1} + K_k (z_k - H_k x_{k|k-1}), \quad (10)$$

$$P_{k|k} = (I - K_k H_k) P_{k|k-1} \quad (11)$$

where S_k is innovation covariance, K_k is Kalman gain, $x_{k|k}$ posterior state estimate, $P_{k|k}$ posterior covariance.

Lemma 1. Given full-rank H_k , the $R_k > 0$ posterior covariance $P_{k|k}$ is positive-definite and $P_{k|k} \leq P_{k|k-1}$.

2.4. Adaptive Covariance Intersection (aCI) Under Unknown Correlations

When inter-sensor correlations are unknown, a correlation-agnostic fusion is achieved via adaptive Covariance Intersection (aCI).

Given modality-specific posteriors

$$\left\{ \hat{x}_{k|k}^{(i)} P_{k|k}^{(i)} \right\}_{i=1}^M \quad (12)$$

$$P_{aCI}^{-1} = \sum_{i=1}^M \omega_i (P_{k|k}^{(i)})^{-1} \quad (13)$$

$$x_{aCI} = P_{aCI} \sum_{i=1}^M \omega_i (P_{k|k}^{(i)})^{-1} x_{k|k}^{(i)} \quad (14)$$

$$\sum_{i=1}^M \omega_i = 1, \quad \omega_i > 0 \quad (15)$$

Weights ω are adapted each step by

$$\omega^* = \arg \min_{\omega \in \Delta^{M-1}} \det(P_{aCI}(\omega)) \quad (16)$$

which maximizes information under correlation uncertainty; (Δ^{M-1}) is the probability simplex.

Theorem 1. The optimization (16) is convex in ω because $\log \det P_{aCI} (\sum_{i=1}^M \omega_i A_i^{-1})$ is convex for positive-definite A_i . Hence, a unique global minimizer exists.

Proof. Let $A_i = P_{k|k}^{(i)} > 0$ and define

$$X(\omega) = \sum_{i=1}^M \omega_i A_i^{-1}, \quad \omega \in \Delta^{M-1} \quad (17)$$

The map $\omega \mapsto X(\omega)$ is affine on the simplex Δ^{M-1} , and $X(\omega)$ is positive definite for all $\omega > 0$. The aCI covariance is $P_{aCI}(\omega) = X(\omega)^{-1}$.

On the cone of positive-definite matrices, the function $f(X) = -\log \det X$ is convex. Thus, the composite function $\omega \mapsto f(X(\omega)) = -\log \det X(\omega)$ is convex in ω .

Moreover,

$$\det(P_{aCI}(\omega)) = \det(X(\omega)^{-1}) = \det(X(\omega))^{-1} \quad (18)$$

so minimizing $\det P_{aCI}(\omega)$ is equivalent to minimizing $-\log \det X(\omega)$, the optimization in Equation (8) is convex on Δ^{M-1} , which implies the existence of a unique global minimizer. $\square$

The convexity of the log-determinant objective and the existence of a unique minimizer follow from standard results in convex optimization [30], while covariance intersection under unknown correlations is well-established in multisensor fusion [31,32]. In practice, we solve Equation (16) on the simplex using a projected-gradient (or exponentiated-gradient) method with backtracking line search. Because the objective is convex and the feasible set is compact, initialization affects only convergence speed, not the final solution (up to numerical tolerance). Empirically, we repeat the optimization with multiple initializations (uniform $\omega_i \mapsto \frac{1}{M}$, random Dirichlet draws, or warm starts $\omega_k \leftarrow \omega_{k-1}$); all runs converge to the same ω^* and produce fused covariances ($\max |\Delta_\omega| < 10^{-3}$), including in higher-dimensional and strongly correlated settings.

2.5. Heteroskedastic Bayesian Regression (Response Layer)

Let $y_k \in \mathbb{R}$ denote a continuous response variable (e.g., crop yield index) predicted from the fused latent state: $x_k^* = \hat{x}_{aCI,k}$. In this study, y_k represents a normalized yield–stress index summarizing crop performance under agro-climatic stress conditions (higher values indicate increased crop stress).

The observation model is

$$y_k = \beta^T x_k^* + \epsilon_k, \quad \epsilon_k \sim N(0, \sigma_k^2) \quad (19)$$

with heteroskedastic variance parameterized as:

$$\log \sigma_k^2 = \gamma^T x_k^* + \eta_k, \quad \eta_k \sim N(0, \tau^2) \quad (20)$$

Priors are specified as weakly informative:

$$\beta \sim N(0, \tau^2), \quad \gamma \sim N(0, \tau^2, I), \quad \tau^2 \sim \Gamma(a_\tau, b_\tau) \quad (21)$$

where can be presented in vector form $\theta = (\beta_1, \dots, \beta_p, \gamma_1, \dots, \gamma_p, \log \tau)$

Remark 1. Equation (10) introduces log-variance regression, ensuring positivity of σ_k^2 and yielding an analytically tractable posterior.

In subsequent analyses, we report posterior summaries for the most influential regression coefficients, defined as the β - and γ -parameters associated with key agro-climatic drivers (temperature, humidity, CTI, and selected soil nutrients) that exhibit large posterior magnitudes and relatively tight credible intervals.

2.6. Log-likelihood, Log-Prior, and Log-Posterior

Given dataset $D = \{\hat{x}_k^*, y_k\}_{k=1}^n$ the conditional log-likelihood is:

$$\log p(y|\beta, \gamma, \tau, D) = -\frac{1}{2} \sum_{k=1}^n \left[\log(2\pi\sigma_k^2) + \frac{(y_k - \beta^T x_k^*)^2}{\sigma_k^2} \right] \quad \sigma_k^2 = \exp(\gamma^T x_k^*) + \tau^2 \quad (22)$$

The joint log-prior is:

$$\log p(\beta, \gamma, \tau) = -\frac{1}{2\tau_\beta^2} \|\beta\|_2^2 - \frac{1}{2\tau_\gamma^2} \|\gamma\|_2^2 + (a_\tau - 1)\log\tau^{-2} - b_\tau\tau^{-2} + C \quad (23)$$

Hence the unnormalized log-posterior becomes:

$$\log p(\beta, \gamma, \tau|D) = \log p(y|\beta, \gamma, \tau, D) + \log p(\beta, \gamma, \tau) \quad (24)$$

Proposition 1. Under bounded data and finite hyperparameters, the posterior of Equation (24) is log-concave in β conditional on (γ, τ) , ensuring unimodality of the conditional distribution.

2.7. MCMC Inference via the Affine-Invariant Ensemble Sampler

Posterior inference for (β, γ, τ) employs the affine-invariant ensemble MCMC algorithm [17,18].

Initialization: Walkers are initialized around maximum a posteriori (MAP) estimates; it uses $N_w \in [100, 200]$ walkers, a burn-in phase $B \in [500, 2000]$, and a total of $S \in [3000, 10,000]$ iterations (with thinning as needed). In the experiments reported in Section 3, we used a fixed configuration with $N_w = 80$ walkers and $S = 3000$ iterations.

Outputs: The sampler produces posterior draws for (β, γ, τ) , corner plots of joint marginals, and standard convergence diagnostics, including the integrated autocorrelation time (IAT/IACT), the Gelman–Rubin $\hat{R}$ statistic (for split chains), and the effective sample size (ESS) for each parameter [33,34].

Posterior summaries reported in the Results are computed from the retained MCMC draws as the posterior median together with the 25–75% interquartile range (IQR).

Convergence diagnostics include the integrated autocorrelation time and Gelman–Rubin statistic [30,31].

The full Bayesian inference procedure using the affine-invariant ensemble sampler is summarized in Algorithm 1.

Algorithm 1. Bayesian Inference with Affine-Invariant Ensemble MCMC

Input:

Fused state inputs $\{x_k^*\}$ from aCI fusion (via Equations (12)–(16))
Observed outputs y_k

Steps:

1. Compute the conditional log-likelihood: $\log p(y|\beta, \gamma, \tau, D)$ by Equation (22)
2. Add log-priors: $\log p(\beta, \gamma, \tau)$ using Equation (23)
3. Form the unnormalized log-posterior: $\log p(\beta, \gamma, \tau|D)$ using Equation (24)
4. Run affine-invariant ensemble sampler;
 - initialize walkers near MAP
 - discard burn-in
 - thin chains if necessary.
5. Posterior predictive sampling:

For each retained draw l :

$$\sigma_k^{2l} = \exp\left(\gamma^{(l)T} x_k^*\right) + \tau^2$$

$$y_k^l \sim N\left(\beta^{(l)T}, x_k^*, \sigma_k^{2(l)}\right)$$

Output:

Posterior draws for (β, γ, τ)

Posterior predictive samples y_k^l

Convergence diagnostics: $IACT$, $\hat{R}$, ESS

2.8. Bayesian Likelihood-Ratio Test (Decision Layer)

For binary decisions (e.g., stress vs. safe), define hypotheses H_1 , H_0 . The Likelihood-Ratio Test (LRT) statistic is

$$\Lambda(z_k) = \frac{p(z_k|H_1)}{p(z_k|H_0)} \quad (25)$$

where H_1 denotes the hypothesis of a stressed condition (e.g., nutrient imbalance, insufficient irrigation), H_0 represents the safe condition, and η is a predefined threshold.

Using posterior predictive draws y_k (Section 2.7), compute the posterior probability of stress

$$\pi_k = Pr(H_1|D, z_k) = E[1\{\tilde{y}_k \in S\} | D] \quad (26)$$

Trigger an alarm if $\Lambda(z_k) > \eta$ or equivalently if $\pi_k > \pi^*$. The Bayes-risk-optimal threshold is

$$\eta^* = \frac{C_{10}(1 - \pi)}{C_{10}\pi} \quad \pi^* = \frac{C_{10}}{C_{10} + C_{01}} \quad (27)$$

With C_{10} (false-alarm cost), C_{01} (missed-detection cost), and prior $\pi = p(H_1)$. π^* is minimizing empirical Bayes risk over posterior predictive samples.

Theorem 2. *Given symmetric Gaussian likelihoods and linear loss, the threshold in Equation (27) minimizes the expected Bayes risk $R(\pi^*)$.*

Proof. Consider the Bayes risk

$$R(\delta) = C_{10}p(\delta(z_k) = H_1, H_0) + C_{10}p(\delta(z_k) = H_0, H_1)$$

where $\delta(z_k)$ is a deterministic decision rule. Under a Bayesian formulation with prior $\pi = p(H_1)$, the posterior odds satisfy

$$\frac{p(z_k|H_1)}{p(z_k|H_0)} = \frac{\pi}{1 - \pi} \Lambda(z_k), \quad (28)$$

where $\Lambda(z_k) = \frac{p(z_k|H_1)}{p(z_k|H_0)}$ is likelihood ratio. For liner loss with cost C_{10} (false alarm) and C_{01} (missed detection). It is optimal to decide H_1 whenever

$$C_{10}p(H_0|z_k) \leq C_{01}p(H_1|z_k) \quad (29)$$

which is equivalent to

$$\Lambda(z_k) \geq \eta^* = \frac{C_{10}(1 - \pi)}{C_{10}\pi} \quad (30)$$

and, in terms of posterior probability,

$$\pi_k = p(H_1|z_k) \quad (31)$$

$$\pi_k \geq \pi^* = \frac{C_{10}}{C_{10} + C_{01}} \quad (32)$$

Under standard Bayesian decision theory, any deviation from this threshold yields higher Bayes risk, making it the optimal decision rule under the assumed model. $\square$

2.9. Model-Averaging of Discriminative ML (Optional Late Fusion)

If multiple discriminative classifiers $\{f_m\}$ are available, Bayesian model-averaging combines their calibrated posteriors as

$$p_m(H_1|\cdot) = \sum_m \omega_m p_m(H_1|\cdot) \quad \sum_m \omega_m = 1 \quad (33)$$

with weights proportional to marginal likelihoods $\omega_m \propto p(D|M_m)$, where trained on x_k^* , combine their calibrated posteriors $p_m(H_1|\cdot)$ via Bayesian model-averaging.

Note: Bayesian model-averaging of discriminative classifiers is presented as an optional late-fusion extension of the framework. It was not used in the experiments reported in Section 3, where the primary pipeline relies on aCI-based fusion and heteroskedastic Bayesian regression.

2.10. Control Layer—Irrigation/Fertilization

Closed-loop adjustment uses a PD law informed by the posterior predictive mean $E[\tilde{y}_k]$ and uncertainty:

$$u_k = K_p(y^* - E[\tilde{y}_k]) + K_d \Delta E[\tilde{y}_k], \quad \Delta E[\tilde{y}_k] = E[\tilde{y}_k] - E[\tilde{y}_{k-1}] \quad (34)$$

Safety caps can incorporate predictive variance (risk-sensitive control).

Remark 2. Incorporating predictive variance in the control gain yields a risk-sensitive controller ensuring stability under stochastic perturbations.

2.11. Time Synchronization and Data Store

Since the data streams originated from distributed IoT devices, timestamps were synchronized using network protocols:

$$t_{global} = t_{local} + \Delta t \quad \Delta t \sim N(0, \sigma_t^2) \quad (35)$$

We store fused trajectories and posteriors

$$T_k = \{x_{0:k|0:k}, P_{0:k|0:k}, \Theta_{0:k}^{MCMC}, c_k\} \quad (36)$$

where $\Theta_{0:k}^{MCMC}$ are retained MCMC samples for continual learning.

2.12. Validation, Calibration, and Decision Utility

Model performance is evaluated using stratified K-fold cross-validation, where each metric is computed on every fold and then averaged:

$$\overline{M} = \frac{1}{K} \sum_{i=1}^K M^{(i)} \quad (37)$$

To assess predictive calibration and decision quality, three standard probabilistic metrics are reported.

Negative Log-Likelihood (NLL). For each test observation, with predictive mean $\mu_k = E[y_k]$ and variance σ_k^2 , the NLL is

$$NLL_k = \frac{1}{2} \left[\log(2\pi \sigma_k^2) + \frac{(y_k - \mu_k)^2}{\sigma_k^2} \right] \quad (38)$$

The final score is the average over all test samples:

$$NLL = \frac{1}{n} \sum_{k=1}^n NLL_k \quad (39)$$

Brier score. For binary stress detection, the predictive probability of stress $\hat{p}_k = P(H_k|x_k)$ is compared with the true label $y_k \in \{0, 1\}$

$$Brier = \frac{1}{n} \sum_{k=1}^n (\hat{p}_k - y_k)^2 \quad (40)$$

Credible interval width (CI width). The 95% posterior credible interval is computed as

$$CI = 1.96\sqrt{\sigma_k^2} \quad (41)$$

and the mean width is averaged across the test set.

Ablation studies assess the contribution of individual components by comparing the following: drone indices vs. no-drone inputs; Kalman fusion vs. adaptive Covariance Intersection; homoskedastic vs. heteroskedastic noise modeling; point-estimate (MAP) inference vs. full MCMC sampling; and fixed decision threshold vs. Bayes-risk-optimal threshold.

Performance degradation is quantified by:

$$\Delta NLL = NLL_{ablated} - NLL_{full} \text{ and } \Delta Brier = Brier_{ablated} - Brier_{fill} \quad (42)$$

Positive values indicate reduced predictive or probabilistic quality when a component is removed.

Additional Validation Measures: Further evaluation includes Accuracy, F1-score, PR-AUC, NLL, Brier score, Expected Calibration Error (ECE), reliability diagrams, and PIT histograms to assess predictive calibration and distributional correctness.

Robustness to correlation: Simulate cross-modal $\rho \in [0, 0.9]$; track coverage of 95% intervals and false-alarm rates at fixed recall.

Reliability diagrams and PIT histograms quantify the calibration quality of predictive distributions.

2.13. Hierarchical Structure of the Hybrid Bayesian Framework

Figure 1 presents the Bayesian multimodal fusion and decision architecture of the proposed uncertainty-aware Crop Decision-Support System (DSS). The framework integrates multi-source data (IoT soil sensors, laboratory assays, meteorological stations, and drone imagery) through adaptive covariance intersection (aCI) for robust fusion under unknown correlations.

The fused state vector enters the heteroskedastic Bayesian regression module, inferred via MCMC sampling (emcee), whose posterior distributions feed into the Bayesian likelihood-ratio test (LRT) for risk-sensitive decision-making and closed-loop PD control for irrigation and fertilization.

Table 1 outlines the hierarchical structure of the hybrid Bayesian framework. Each layer represents a mathematically well-defined operator acting on probability measures or state estimates.

Table 1. Structural classification—Hybrid Bayesian inference and decision algorithm with MCMC sampling.

Layer	Algorithmic Principle	Role
Sensor fusion	Adaptive Covariance Intersection (aCI)	Robust fusion under unknown cross-correlations
State estimation	Bayesian/Kalman filter	Temporal tracking of agro-climatic state
Regression layer	Heteroskedastic Bayesian regression (inferred via emcee MCMC)	Models time-varying predictive uncertainty
Decision layer	Bayesian Likelihood-Ratio Test (LRT)	Converts posterior probabilities into early-warning alarms
Control layer	PD control/feedback law	Adaptive irrigation–fertilization control




Figure 1. Bayesian multimodal fusion, decision, and control framework integrating aCI, Kalman filtering, Bayesian likelihood-ratio testing, and PD feedback control for irrigation/fertilization.

The fusion layer performs convex optimization in the space of positive-definite matrices to obtain correlation-agnostic posteriors. The state-estimation layer executes recursive Bayesian updates for a stochastic state-space model. The regression layer introduces heteroskedastic variance modeling and posterior sampling via Markov Chain Monte Carlo. The decision layer applies a Bayes-risk-optimal likelihood-ratio test for binary hypotheses, and the control layer transforms posterior expectations and variances into real-time feedback actions. This decomposition formalizes the flow of uncertainty from data acquisition to decision and establishes a clear correspondence between algorithmic implementation and its underlying mathematical formulation.

Figure 2 presents the detailed probabilistic architecture underlying the multimodal Bayesian fusion and decision framework. The diagram highlights the flow from Kalman-based state estimation and adaptive covariance intersection, through heteroskedastic Bayesian regression inferred with MCMC, to Bayesian likelihood-ratio testing and closed-loop PD control. Each block corresponds to a well-defined probabilistic operator, making explicit how uncertainty propagates from sensor noise to posterior inference and operational decisions.




Figure 2. Probabilistic framework for uncertainty-aware multimodal fusion, Bayesian regression, and decision control using Kalman filtering, adaptive covariance intersection, and heteroskedastic MCMC inference. τ^* denotes a decision threshold used to trigger an alarm.

3. Results

The multimodal Bayesian framework was evaluated on a dataset consisting of 1100 samples and 12 agronomic, environmental, and spectral variables. Eleven crop types are included in the dataset: rice, maize, chickpea, kidney beans, lentil, pomegranate, watermelon, muskmelon, apple, cotton, and grapes, which are characteristic of Balkan and North Macedonia. These represent multiple agronomic categories (cereals, legumes, fruit crops, and industrial fiber crops), enabling the assessment of the proposed method under heterogeneous biological growth characteristics and stress responses. Their diversity strengthens the generalizability of the probabilistic decision model.

3.1. Performance of Multimodal Fusion

Simulation experiments were conducted to quantify the effect of unknown cross-sensor correlations on posterior uncertainty. A controlled linear–Gaussian system with two sensors was simulated, where the true measurement-noise covariance contained injected cross-correlations $\rho \in \{0.0, 0.3, 0.6, 0.9\}$. The naïve Kalman fusion baseline assumes independent sensor noise, while the adaptive Covariance Intersection (aCI) method fuses the single-sensor posteriors without requiring knowledge of cross-covariances. To quantitatively compare the uncertainty behavior of the Kalman fusion and the adaptive Covariance Intersection (aCI), the time-averaged posterior covariance trace and the associated 95% credible interval width are evaluated.

For a given filtering method (KF or aCI), the mean posterior covariance trace over a time horizon T is defined as:

$$\bar{P}_{KF} = \frac{1}{T} \sum_{k=1}^T \text{tr}(P_k^{KF}) \quad \bar{P}_{aCI} = \frac{1}{T} \sum_{k=1}^T \text{tr}(P_k^{aCI}) \quad (43)$$

where P_k^{KF} and P_k^{aCI} denote the posterior *state* covariance matrices at time step k for the Kalman filter and *aCI*, respectively.

Assuming a Gaussian posterior, the 95% credible-interval width, for a scalar state component with posterior variance $\sigma_{k,i}^2 = (P_k)_{ii}$, is

$$CI_{95\%} = 1.96\sqrt{\sigma_{k,i}^2} = 1.96\sqrt{(P_k)_{ii}} \quad (44)$$

In the experiments, we report the mean of these 95% credible-interval widths across time and state dimensions for each filtering method, providing a scalar summary of posterior uncertainty that is directly comparable between KF and aCI.

Because the Kalman filter assumes a diagonal measurement covariance R_{model} , the posterior covariance R_k depends solely on the assumed model and is therefore invariant to the injected correlation ρ . Only the empirical MSE reflects the effect of ρ , demonstrating that the filter becomes overconfident and increasingly biased under unmodeled cross-correlation.

The results in Table 2 confirm that the covariance-based metrics of the naïve Kalman filter remain insensitive to the injected correlation ρ , because the filter assumes a fixed diagonal measurement-noise model. For aCI, the reported covariance summaries were observed to be approximately constant in these experiments; however, aCI may change its conservative bound when the single-sensor posteriors diverge more strongly under correlation. Thus, the posterior covariance trace and credible-interval width remain effectively constant across all conditions. In contrast, the empirical mean-squared error (MSE) increases monotonically with ρ , reflecting degradation in estimation accuracy as the true sensor dependence grows. The naïve Kalman filter accumulates this error more rapidly, whereas aCI remains slightly more robust for high ρ (e.g., $\rho = 0.9$), consistent with its guaranteed conservative behavior under unknown cross-correlations. No divergence was observed, confirming the stability result stated in Theorem 1.

Table 2. Summary of the posterior covariance trace and 95% credible-interval width obtained under both fusion strategies.

ρ	Mean Posterior Covariance Trace (KF)	Mean Posterior Covariance Trace (aCI)	Relative Change	95% CI Width (KF)	95% CI Width (aCI)	MSE (KF)	MSE (aCI)
0.0	0.0218	0.0341	+56.4%	0.290	0.362	0.0218	0.0232
0.3	0.0218	0.0341	+56.4%	0.290	0.362	0.0254	0.0261
0.6	0.0218	0.0341	+56.4%	0.290	0.362	0.0288	0.0287
0.9	0.0218	0.0341	+56.4%	0.290	0.362	0.0325	0.0316

3.2. Effect of Cross-Sensor Correlation on Posterior Uncertainty Propagation

This subsection provides an interpretation of the results in Section 3.1 by linking injected cross-sensor correlation (ρ) to the practical risk of overconfident fusion under an independence assumption. These single-sensor posteriors are subsequently fused to obtain a global posterior used for prediction and decision-making. However, in real agro-environmental systems, sensor noises are often statistically dependent due to common environmental drivers and shared acquisition conditions. To study the impact of such dependencies, cross-sensor correlation was explicitly injected at the measurement-noise level through a controlled covariance structure parameterized by $\rho \in \{0.0, 0.3, 0.6, 0.9\}$. This

procedure allows for systematic evaluation of fusion robustness under increasing violation of the independence assumption.

Two fusion strategies were compared. The naïve Kalman filter (KF) assumes a diagonal measurement-noise covariance matrix and therefore treats all sensor measurements as statistically independent. In contrast, the adaptive Covariance Intersection (aCI) method fuses posterior distributions without requiring knowledge of cross-covariances, guaranteeing conservative uncertainty bounds even under unknown inter-sensor correlations. For both methods, posterior uncertainty was propagated over time and summarized by the mean trace of the posterior covariance and the corresponding 95% credible-interval width, as defined in Equations (42) and (43). This setup isolates the effect of unmodeled correlation on posterior uncertainty estimation and provides the theoretical basis for the comparative results reported in Section 3.1.

3.3. Posterior Inference of Heteroskedastic Bayesian Regression

Posterior inference was performed using an affine-invariant ensemble MCMC sampler with 80 walkers and 3000 iterations, resulting in 240,000 retained draws, of which N retained were kept after burn-in. Convergence diagnostics confirmed stable posterior exploration (Table 3).

Table 3. Convergence diagnostics for the heteroskedastic Bayesian regression posterior sampler.

Diagnostic	Value
Integrated autocorrelation time	34–52
Gelman–Rubin $\hat{R}$	1.02–1.05
Effective sample size	950–1600

Posterior summaries (median and interquartile range) for the most influential parameters are shown in Table 4.

Table 4. Key posterior parameters (median [25%, 75%]).

Parameter	Median	25%	75%
β_5 (temperature)	0.0835	0.0582	0.1099
β_6 (humidity)	−0.0320	−0.0397	−0.0241
β_{10} (CTI)	−15.365	−19.23	−11.35
γ_2 (phosphorus)	−0.933	−4.618	−0.095
γ_5 (temperature)	−1.682	−6.564	−0.021

Wide interquartile ranges observed for log-variance coefficients γ_2 associated with phosphorus indicate partial identifiability, which commonly arises when predictors are correlated or when variance heterogeneity is modest relative to observation noise. This behavior is expected in heteroskedastic regression models and does not compromise the proposed framework. Importantly, uncertainty in γ is fully propagated into the posterior predictive variance σ_k^2 and, consequently, into the stress probability π_k used for decision-making. In practice, larger uncertainty in γ leads to wider predictive intervals and more conservative decisions, rather than overconfident variance estimates. The operational reliability of the variance model is therefore assessed at the predictive level using calibration diagnostics and distributional metrics (NLL, ECE, PIT), which directly evaluate the adequacy of uncertainty quantification.

Temperature and CTI strongly govern the posterior mean of the yield–stress index, whereas the γ coefficients indicate differential contributions of environmental drivers to predictive variance, with strongest effects attributed to temperature and phosphorus.

3.4. Posterior Predictive Accuracy

Posterior predictive samples $\tilde{y}_k$ were generated from the retained MCMC chains. On the normalized yield scale, the model achieved the predictive results shown in Table 5.

Table 5. Posterior predictive performance metrics calculated from posterior draws.

Metric	Value
RMSE	0.43
Brier score	0.112
Negative log-likelihood (NLL)	0.285
Expected calibration error (ECE)	0.036

Reliability diagrams closely followed the perfect-calibration diagonal, and the PIT histogram was approximately uniform, indicating well-calibrated predictive distributions. The lower RMSE (0.43) reflects normalization; Section 3.9 reports out-of-sample (stratified 3-fold) performance on the original yield_index scale, which is expected to yield larger error and calibration metrics than the in-sample normalized diagnostics reported here.

Figure 3 presents the posterior predictive calibration diagnostics for the proposed Bayesian model.




Figure 3. Posterior predictive calibration diagnostics: (a) PIT histogram for the normalized yield_index, showing near-uniform distribution with mild underdispersion, consistent with well-calibrated predictive uncertainty; (b) reliability diagram comparing predicted stress probabilities with empirical frequencies, demonstrating close adherence to the perfect-calibration diagonal.

Figure 3a shows the Probability Integral Transform (PIT) histogram for the normalized yield_index. The histogram is approximately uniform (mean = 0.50, std = 0.31), indicating that the posterior predictive distribution is, in general, well calibrated. A mild U-shaped pattern is visible, with slightly increased mass in the 0–0.2 and 0.8–1.0 intervals, suggesting marginal underdispersion, meaning that predictive intervals are slightly narrower than the true empirical variability. Despite this, the PIT values remain well within acceptable limits for heteroskedastic Bayesian regression models. Combined with the low Expected Calibration Error (ECE = 0.036), these results confirm good overall probabilistic calibration.

Figure 3b provides a reliability diagram (calibration curve) for the posterior stress probabilities π_k . The calibration curve lies close to the diagonal, demonstrating that the predicted probabilities closely match the empirical frequency of stress across most probability bins. Deviations from perfect calibration are small, which is consistent with the

low ECE and further supports that the proposed heteroskedastic Bayesian model achieves well-calibrated probabilistic predictions.

3.5. Bayesian Stress Decision Layer

Building on the posterior predictive distributions obtained in Section 3.3, a probabilistic decision layer was constructed to translate continuous yield–stress predictions into binary stress alarms suitable for agro-environmental decision-making.

Stress was defined using a yield-index threshold on the original (non-normalized) scale:

$$S = \{y_k < y_{crit}\}, \quad y_{crit} = 6.5$$

For each sample k the posterior stress probability was estimated from the posterior predictive draws:

$$\pi_k = Pr(y_k \in S | D_k) \approx \frac{1}{N_{draw}} \sum_{j=1}^{N_{draw}} I\{\tilde{y}_k^{(j)} < y_{crit}\}$$

where $\tilde{y}_k^{(j)}$ denotes the j -th posterior predictive sample, and $I\{\cdot\}$ is the indicator function.

The dataset exhibits moderate class imbalance, with 762 SAFE and 338 STRESS samples (STRESS ratio ≈ 0.307), making accuracy alone an unreliable metric for threshold selection.

Under a symmetric loss assumption, where false alarms and missed detections have equal costs ($C_{10} = C_{01}$), the Bayes-optimal decision threshold becomes:

$$\pi^* = \frac{C_{10}}{C_{10} + C_{01}} = 0.5$$

A binary stress decision is then obtained as:

$$\hat{z}_k = \begin{cases} 1, & \text{if } \pi_k \geq \pi^* \\ 0, & \text{otherwise,} \end{cases}$$

where $\hat{z}_k = 1$ denotes STRESS and $\hat{z}_k = 0$ denotes SAFE. Using $\pi^* = 0.5$ yields a balanced baseline classifier that detects both SAFE and STRESS samples. However, due to overlapping posterior stress probabilities and class imbalance, this theoretically optimal threshold does not maximize task-specific performance metrics, such as the F_1 -score.

To analyze the sensitivity of stress detection to the decision threshold, the threshold $\tau \in (0,1)$ was varied, and the corresponding accuracy and F_1 -score were evaluated.

Maximizing accuracy leads to a high threshold ($\tau \approx 0.94$), which trivially predicts all samples as SAFE. While this operating point achieves an accuracy of approximately 0.69, it yields an F_1 -score of zero and completely suppresses stress detection, rendering it unsuitable for risk-sensitive agricultural applications.

In contrast, maximizing the F_1 -score yields a low threshold ($\tau^* \approx 0.04$), corresponding to a safety-oriented operating point. This threshold substantially increases recall for the STRESS class at the expense of increased false alarms, reflecting a cost-asymmetric decision rule consistent with agro-climatic risk management, where missed stress events are considerably more costly than false alarms.

Table 6 summarizes stress classification performance at representative operating points, clearly illustrating how different optimization criteria lead to fundamentally different decision behaviors.

Table 6. Stress classification performance at representative decision thresholds.

Threshold Strategy	Threshold τ	Accuracy	F ₁ -Score (STRESS)	Interpretation
Symmetric Bayes	0.50	moderate	moderate	Balanced theoretical baseline
Accuracy-optimal	≈ 0.94	≈ 0.69	0.00	Trivial all-SAFE classifier
F ₁ -optimal	≈ 0.04	≈ 0.31	≈ 0.47	Safety-oriented, high recall

To illustrate the effect of threshold selection, confusion matrices are reported for two representative operating points: Figure 4a presents the confusion matrix obtained using the symmetric Bayes threshold $\tau = 0.5$, serving as a neutral baseline; Figure 4b presents the confusion matrix obtained using the F₁-optimal threshold $\tau^* \approx 0.04$, highlighting the substantial increase in STRESS recall at the cost of increased false alarms.




Figure 4. Confusion matrix of stress classification based on posterior stress probabilities. Darker shades indicate higher sample counts, while lighter shades indicate lower counts. (a) Symmetric Bayes decision threshold $\tau = 0.5$, yielding a balanced baseline classifier under equal error costs; (b) Safety- or F₁-optimal threshold $\tau^* \approx 0.04$, where τ^* denotes the threshold selected to optimize the F1 score, prioritizing stress detection and achieving high recall for the STRESS class at the expense of increased false alarms.

These results demonstrate that threshold selection critically influences the operational behavior of the stress decision layer. While the symmetric Bayes threshold provides a theoretically sound baseline under equal error costs, practical agro-climatic decision-making typically involves asymmetric risk, where missed stress events (e.g., under-irrigation) are substantially more costly than false alarms. Consequently, optimal decision performance cannot be defined solely in terms of accuracy, but must account for application-specific risk preferences and cost asymmetry. The proposed probabilistic framework enables transparent exploration of this trade-off and supports application-specific threshold selection based on risk tolerance rather than accuracy alone.

To further elucidate the behavior of the Bayesian stress decision layer, the empirical posterior stress probability distributions π_k for SAFE and STRESS samples are shown in Figure 5.

The two distributions exhibit substantial overlap across a wide range of probability values, indicating that stress and non-stress conditions cannot be cleanly separated by a single probability threshold. Importantly, this overlap reflects intrinsic uncertainty and class imbalance in the data, rather than a lack of probabilistic calibration. This overlap directly explains the degeneracy observed at the symmetric Bayes threshold ($\tau = 0.5$), where most samples fall on the SAFE side of the decision boundary, leading to suppressed STRESS recall despite well-calibrated posterior probabilities.




Figure 5. Posterior stress probability distributions for SAFE and STRESS samples. The dashed vertical line indicates the symmetric Bayes decision threshold $\tau = 0.5$, while the dotted vertical line denotes the F1-optimal threshold $\tau^* = 0.04$ where τ^* is selected to optimize the F1 score, prioritizing stress detection.

Conversely, the F1-optimal threshold $\tau^* \approx 0.04$ shifts the decision boundary into a low-probability region, intentionally prioritizing sensitivity to stress events at the cost of increased false alarms. This operating point corresponds to a cost-sensitive decision rule that favors recall over precision, which is appropriate in risk-averse agro-environmental applications.

Figure 5 demonstrates that the trade-off between accuracy and recall is not a modeling artifact but an intrinsic consequence of posterior uncertainty and class overlap.

This behavior underscores the necessity of probabilistic decision-making and cost-aware threshold selection in agro-environmental applications, where missed stress events are substantially more costly than conservative interventions. Notably, decision performance is governed at the decision layer, while the underlying Bayesian model remains well-calibrated and uncertainty-aware.

Figure 6 presents accuracy and F1-score as a function of the decision threshold τ . Maximizing accuracy leads to a trivial all-SAFE classifier, while maximizing the F1-score yields a safety-oriented operating point suitable for agro-environmental decision-making.




Figure 6. Accuracy and F1-score as a function of the decision threshold τ .

The strong overlap between posterior stress probability distributions for SAFE and STRESS samples prevents the existence of a single threshold that simultaneously maximizes accuracy and recall. This limitation is inherent to the uncertainty structure of the problem and is exacerbated by class imbalance.

Accordingly, the proposed Bayesian framework ensures calibrated probabilities and supports flexible, cost-aware decision rules, rather than relying on fixed deterministic thresholds or post hoc accuracy optimization.

3.6. Heteroskedastic Noise Modeling

Table 7 compares two noise models: homoskedastic, where the observation noise has constant variance across all conditions; and heteroskedastic, where the variance changes as a function of agro-climatic drivers (temperature, humidity, CTI, etc.), as defined by γ in Equation (20).

Table 7. Utility of heteroskedastic modeling.

Component Removed	NLL ↓	Brier ↓	CI Width ↓
Homoscedastic	0.517	0.201	2.41
Heteroskedastic	0.285	0.112	1.76

↓ indicates that lower values correspond to better performance.

To evaluate comparative performance, three probabilistic metrics are reported:

- NLL (Negative Log-Likelihood): Assesses the overall quality of the predictive distribution—lower values indicate better-calibrated uncertainty.
- Brier score: Measures the accuracy of probabilistic stress predictions—lower is better.
- CI width: Average width of the 95% credible interval—narrower intervals indicate reduced predictive uncertainty.

Table 7 demonstrates a clear advantage of heteroskedastic modeling: a 34–44% reduction in predictive uncertainty (narrower credible intervals) and a $\approx 45\%$ improvement in probabilistic decision quality (substantially lower NLL and Brier score).

These results confirm that allowing for the noise variance to depend on environmental stressors (via γ) yields better-calibrated, sharper, and more informative predictive distributions compared with a homoskedastic assumption.

3.7. Ablation Study

MCMC inference and drone vegetation indices are the two largest contributors to uncertainty reduction.

Table 8 reports the ablation results used to quantify the contribution of each component to predictive performance and decision reliability. Removing the aCI fusion mechanism led to higher uncertainty and degraded calibration ($\Delta\text{NLL} = +0.19$), confirming that robust aggregation of heterogeneous sensors is essential when cross-correlation is present among soil, weather, and drone modalities. Excluding the drone vegetation indices (NDVI, SAVI, CTI) produced an even larger deterioration ($\Delta\text{NLL} = +0.25$), demonstrating that canopy-level spectral measurements capture variance that cannot be explained by ground sensors alone.

Table 8. Ablation study results showing the effect of removing key model components on uncertainty and decision quality.

Component Removed	$\Delta\text{NLL} \uparrow$	$\Delta\text{Brier} \uparrow$
No aCI fusion	+0.19	+0.13
No drone indices (NDVI, SAVI, CTI)	+0.25	+0.21
No MCMC (MAP only)	+0.31	+0.18
Fixed threshold $\pi = 0.5$	0.14	+0.11

↑ indicates that lower values correspond to worse performance.

Replacing full MCMC inference with a MAP estimator resulted in the largest increase in NLL (+0.31), indicating that accurate uncertainty quantification—not only point predictions—is crucial for reliable stress probability estimation. Finally, enforcing a fixed

decision threshold ($\tau = 0.5$) reduced robustness across environmental conditions, primarily degrading decision-level metrics (e.g., F1/recall) under class imbalance, even when the underlying probabilistic predictions remained unchanged.

The ablation study highlights that the combination of MCMC uncertainty estimation, drone-based vegetation indices, and robust aCI fusion provides the greatest improvement in both calibration and decision quality.

3.8. Control Performance

The posterior uncertainty was propagated into a proportional–derivative (PD) control law:

$$u_k = K_p(y^* - E[y_k]) + K_d\Delta E[y_k] \quad (45)$$

where y^* is the target yield index, $E[y_k]$ denotes the posterior predictive mean at day k , and $\Delta E[y_k] = E[y_k] - E[y_{k-1}]$. The proportional and derivative gains were empirically tuned to ensure closed-loop stability, smooth irrigation actuation, and realistic soil–plant dynamics. The final values used in all simulations were $K_p = 0.85$, and $K_d = 0.12$. These values were selected by grid search over stabilizing ranges to minimize water usage while avoiding oscillatory moisture dynamics. The same gains were used for both controllers to ensure a fair comparison.

To quantitatively evaluate the effect of incorporating posterior predictive uncertainty into the irrigation controller, a stylized 90-day soil–plant dynamic simulation was performed. The Bayesian proportional–derivative controller was compared to a conventional rule-based strategy across 200 Monte Carlo realizations with stochastic evapotranspiration, sensor noise, and soil-moisture variability.

True soil-moisture dynamics:

$$M_{k+1} = M_k + I_k - ET_k + \omega_k \quad (46)$$

where I_k is irrigation applied by the controller, ET_k is evapotranspiration (randomly varying daily), and $\omega_k \sim N(0, \sigma_\omega^2)$ process noise.

Sensor model (noisy measurements):

$$\tilde{M}_k = M_k + v_k, \quad v_k \sim N(0, \sigma_v^2) \quad (47)$$

which generates occasional false stress signals.

Stress was defined as $stress_k = (y_k < y_{crit})$ with $y_{crit} = 6.5$. In this stylized 90-day simulation, soil moisture M_k acts as a proxy driver that influences the yield–stress index y_k through the predictive model; therefore, irrigation affects stress probability indirectly via the dynamics of M_k .

Two controllers were tested

Rule-based: irrigates when measured moisture < threshold

Bayesian PD controller: irrigates only when $P(M_k < M_{crit}|D_k) > \tau$ where τ is the decision threshold selected according to the Bayesian stress decision layer (Section 3.5).

Irrigation was scaled using the PD control law in Equation (41), while Equation (43) specifies the noisy measurement model used for posterior state updates.

For each simulation run, the following are recorded:

Total water usage:

$$W = \sum_{k=1}^{90} I_k$$

False stress alarms: Days when the controller irrigated due to noise, even though

$$(M_k \geq M_{crit})$$

After running 200 stochastic simulations, the following are computed:
Water usage:

$$W_{rule} = \frac{1}{200} \sum_{k=1}^{200} W_{rule}^{(i)} = 238.65 \quad W_{bayes} = \frac{1}{200} \sum_{k=1}^{200} W_{bayes}^{(i)} = 204.69$$

False stress alarms:

$$FA_{rule} = \frac{1}{200} \sum_{k=1}^{200} FA_{rule}^{(i)} = 0.78 \quad FA_{bayes} = \frac{1}{200} \sum_{k=1}^{200} FA_{bayes}^{(i)} = 0.50$$

The performance metrics (204.69 vs. 238.65 units of water; 0.50 vs. 0.78 false-alarm days) represent the ensemble means across all 200 simulations, providing statistically stable estimates of long-term behavior. Each simulation yields slightly different outcomes due to climatic and sensor uncertainty, and the reported averages thus quantify systematic, rather than incidental, performance differences.

The Bayesian controller demonstrated substantial gains in both resource efficiency and decision reliability:

- 14.2% reduction in total water usage (204.69 vs. 238.65 units), achieved by acting only when the posterior predictive distribution indicated a credible risk of plant stress;
- 35.5% fewer false stress alarms (0.50 vs. 0.78 days), confirming that calibrated uncertainty estimates prevent unnecessary irrigation events;
- More stable control inputs, because posterior uncertainty regularizes abrupt changes in the expected yield-index trajectory $E[y_k]$.

Compared to rule-based irrigation, the results report 14.2% lower water usage and 35.5% fewer false stress alarms.

Figure 7 compares the relative irrigation performance of the proposed Bayesian uncertainty-aware controller and the conventional rule-based strategy, normalized to the rule-based baseline. The Bayesian controller achieves approximately 14% lower water consumption and 35% fewer false stress alarms on average. Error bars denote 95% confidence intervals obtained from 200 Monte Carlo simulations. The results demonstrate that integrating posterior predictive uncertainty into the control law leads to both improved water-use efficiency and enhanced decision reliability.




Figure 7. Relative control performance of Bayesian and rule-based irrigation strategies (mean $\pm$ 95% CI).

3.9. Cross-Validated Decision Performance

A stratified 3-fold Bayesian cross-validation was conducted following the protocol in Section 2.12. The full dataset ($N = 1100$) was partitioned into three folds while preserving the

stress proportion ($\approx 30.7\%$, defined as $\text{yield_index} < 6.5$). For each fold, the heteroskedastic Bayesian regression model was re-trained on the training subset, and posterior predictive distributions for the validation subset were obtained using MCMC sampling. Table 9 reports performance across folds using RMSE for the continuous yield_index, Gaussian negative log-likelihood (NLL), Brier score for stress probability, accuracy and F1-score for binary stress classification, and expected calibration error (ECE).

Table 9. Three-fold Bayesian cross-validation.

Metric	Mean $\pm$ SD
RMSE	15.08 ± 1.78
NLL	3.01 ± 0.31
Brier score	0.318 ± 0.035
Accuracy	0.69 ± 0.065
F1 score	0.333 ± 0.057
ECE	0.254 ± 0.074

The model achieved an RMSE of 15.08 ± 1.78 , indicating moderate reconstruction accuracy of the yield index. The NLL (3.01 ± 0.31) and Brier score (0.32 ± 0.04) show that the model produces informative probabilistic predictions. Stress-detection accuracy averaged 0.69 ± 0.065 , with an F1-score of 0.33 ± 0.057 , reflecting reasonable discrimination under pronounced class imbalance.

Calibration analysis yielded an ECE of 0.25 ± 0.07 , indicating that predicted stress probabilities are not perfectly calibrated but still retain meaningful uncertainty structure. Notably, folds with lower Brier scores tended to show lower ECE, suggesting internal consistency between probabilistic accuracy and calibration.

Overall, these results confirm that the heteroskedastic Bayesian model provides coherent and practically useful uncertainty estimates that can be directly exploited in the risk-sensitive decision and control framework described in Sections 2.8 and 2.10. Given the class imbalance ($\approx 30\%$ stressed crops), an accuracy near 51% would be expected from a naïve classifier, whereas the observed F1-score of 0.33 demonstrates non-trivial stress-detection capability.

4. Discussion

The results of this study demonstrate the importance of explicitly modeling uncertainty in agro-environmental decision-making systems, particularly when predictions are directly coupled with control actions, such as irrigation. Unlike many existing smart-agriculture pipelines that rely on deterministic regression or classification models and do not propagate uncertainty into the control loop, the proposed framework integrates heteroskedastic Bayesian inference supported by advanced MCMC techniques [7–9,18–20], robust multimodal sensor fusion via adaptive Covariance Intersection grounded in optimal covariance-bounding theory [14,25,35,36], and a risk-sensitive stress decision layer formulated in accordance with Bayesian decision theory [22]. This holistic probabilistic formulation allows for not only point prediction of crop stress but also calibrated quantification of confidence, which is essential for reliable actuation under stochastic environmental and dynamic-system conditions, consistent with principles of uncertainty-aware control developed in related nonlinear systems [3–5,37].

The uncertainty fusion analysis confirmed that unmodeled cross-sensor correlations lead to systematic overconfidence in the naïve Kalman filter, as evidenced by the approximately invariant posterior covariance trace despite increasing injected correlation ρ . Such behavior reflects the well-known inconsistency of Kalman-based fusion when indepen-

dence assumptions are violated [12,13]. In contrast, the adaptive Covariance Intersection method maintained conservative and stable uncertainty bounds across all correlation levels, while the empirical MSE increased monotonically with ρ . These findings validate the robustness guarantees of CI-type fusion methods [14,38,39] and highlight their importance for real-world agro-ecological sensing systems, where soil, weather, and UAV-derived measurements are often mutually dependent [25,26,40,41].

The heteroskedastic Bayesian regression results further demonstrate that allowing for predictive variance to depend on environmental stressors (such as temperature, moisture imbalance, and CTI) yields sharper and better-calibrated posterior distributions than homoskedastic models. The 34% reduction in credible-interval width and the substantial improvement in NLL and Brier score indicate that heteroskedastic modeling captures meaningful structure in both the mean and variance of the yield–stress process. Similar conclusions regarding the superiority of heteroskedastic uncertainty models in environmental prediction tasks have been reported in recent probabilistic and sensor-fusion studies [3,6,22,39].

From a decision-theoretic perspective, the Bayesian stress decision layer illustrates the practical limitations of symmetric-loss Bayes rules under class imbalance. While the theoretical threshold $\pi^* = 0.5$ is optimal only when misclassification costs are equal, agro-climatic decision-making is inherently asymmetric, since missed stress events (e.g., under-irrigation) are considerably more costly than false alarms. By optimizing the decision threshold using the F_1 -score, the practically optimal operating point shifted to $\tau^* \approx 0.04$, resolving the degeneracy of the symmetric-cost classifier and significantly improving recall for the stress class. Similar threshold-adaptation strategies have been advocated in recent agricultural risk-detection and stress-monitoring systems [19,21,24,40,41].

The closed-loop irrigation control experiments further confirm the operational value of uncertainty-aware decision-making. By propagating posterior predictive uncertainty into a proportional–derivative control law, the Bayesian controller achieved a 14.2% reduction in total water usage and a 35.5% reduction in false stress alarms compared with a conventional rule-based strategy. These gains are fully consistent with recent IoT- and ML-enabled smart irrigation systems, which report that closed-loop control combined with probabilistic prediction can substantially reduce water consumption while maintaining yield stability [19–21]. However, unlike most existing smart-irrigation schemes that rely on deterministic thresholds or point forecasts [19,20], the proposed controller explicitly integrates posterior uncertainty into the control logic, thereby reducing actuation driven by sensor noise and short-term fluctuations.

In previous work on deterministic crop recommendation, Knights et al. [29] developed a supervised learning pipeline combining exploratory data analysis, mathematical modeling, and six machine learning classifiers (Decision Tree, Random Forest, Gaussian Naïve Bayes, Logistic Regression, XGBoost, and SVM) to recommend the most suitable crop based on soil nutrients and agro-environmental conditions. Using a dataset with 1100 samples and 11 crop types, ensemble and probabilistic models achieved classification accuracies above 99%, confirming that well-structured agro-environmental features can support highly reliable deterministic crop-type prediction [25,27,28].

However, the problem addressed in the present study is fundamentally different. Instead of multi-class crop-type classification under relatively balanced classes, we address binary stress detection on a continuous yield_index under class imbalance ($\approx 30\%$ stressed cases), heteroskedastic noise, and explicit uncertainty propagation into irrigation control. Consequently, overall accuracy is not directly comparable between the two regimes: near-perfect deterministic accuracies reflect a simpler classification setting with point predictions,

whereas the accuracy of 0.69 and F1-score of 0.33 obtained here arise from a substantially more demanding, risk-sensitive probabilistic decision-making task.

Taken together, the two lines of work are complementary: deterministic ML models offer excellent crop-selection accuracy [9,14,25], whereas the Bayesian framework proposed here enables calibrated stress-risk assessment and robust control under uncertainty in dynamic agro-environmental conditions.

The proposed pipeline is modular by design. Components that are largely domain-agnostic include adaptive Covariance Intersection (aCI) fusion, which operates on Gaussian posterior summaries without requiring cross-covariance knowledge; Bayesian posterior sampling and uncertainty propagation; Bayes-risk-based decision thresholding; and probabilistic calibration and evaluation metrics. Domain-specific elements are the definition of the latent state vector x_k , the measurement operators $H_k^{(i)}$ and noise models $R_k^{(i)}$, the event or stress set S used to define posterior event probabilities π_k , and the choice of control law (a PD controller in this study).

To apply the framework to robotic navigation, x_k may represent pose and velocity, measurements may originate from LiDAR or vision sensors, and S may encode collision or localization-failure risk. In environmental monitoring, x_k may represent pollutant concentrations or hydrological states, and S may define regulatory exceedance events. In all cases, the fusion-to-decision structure of the framework remains unchanged.

From a broader perspective, these findings support the growing role of Bayesian learning and uncertainty-aware control in smart agriculture, particularly in IoT-enabled systems where sensing reliability, environmental variability, and economic risk must be jointly addressed [15–21]. The proposed framework provides a principled mathematical bridge between probabilistic inference and real-time actuation, enabling sustainable water management under uncertain climatic conditions and complementing recent UAV–IoT–ML integration efforts in precision farming [17,18,20].

Future research will extend the current decision-support framework beyond multimodal Bayesian inference toward full cyber–physical autonomy.

First, strengthening the IoT network and system security will be essential to ensure reliable data acquisition and to protect the DSS against vulnerabilities, intrusion attempts, and manipulation of sensor streams, in line with previously proposed methodologies for IoT threat detection and prevention [42].

Second, the framework will be advanced toward integration with robotic and mechatronic platforms, enabling autonomous monitoring, precision spraying, soil sampling, and stress mitigation using mobile and anthropomimetic robots. Prior work on safety-oriented robot operation [43] and balance-stability control in anthropomimetic robots [44] provides a strong foundation for incorporating robotic actuation into the DSS architecture. Additional advances in agricultural robotics, particularly in autonomous ground vehicles and UAV-enabled monitoring [45], will further support this integration.

Ultimately, the DSS is envisioned to evolve into a secure, uncertainty-aware cyber–physical ecosystem in which Bayesian reasoning, secure IoT infrastructure, and autonomous robotics jointly support sustainable and resilient agricultural management.

5. Conclusions

This study presented a unified Bayesian framework for multimodal agro-environmental decision support that integrates robust sensor fusion, heteroskedastic Bayesian regression, and a risk-sensitive stress–decision layer within a closed-loop irrigation controller. The results show that explicitly modeling and propagating uncertainty—from sensor measurements to predictive inference and control actions—substantially improves reliability compared with deterministic baselines. Adaptive Covariance Intersection maintained

stable uncertainty bounds under unknown cross-sensor correlations, while heteroskedastic modeling produced sharper and better-calibrated predictive distributions. The Bayesian decision layer effectively addressed class imbalance through threshold optimization, and the uncertainty-aware controller achieved meaningful improvements in water efficiency and reduction of false stress alarms.

Collectively, these findings demonstrate that uncertainty-aware Bayesian reasoning is essential for robust smart-agriculture decision systems operating under environmental variability and sensor noise. Future work will extend the framework toward secure IoT network integration and autonomous robotic platforms, paving the way for a fully cyber–physical, uncertainty-aware ecosystem for sustainable agricultural management.

Author Contributions: Conceptualization, V.A.K.; methodology, V.A.K.; software, V.A.K., M.P., L.K. and J.G.K.; investigation, V.A.K. and J.G.K.; resources, V.A.K.; writing—original draft preparation, M.P. and L.K.; writing—review and editing, V.A.K., J.G.K. and L.K.; visualization, M.P. and L.K.; supervision, V.A.K. All authors have read and agreed to the published version of the manuscript.

Funding: This research is based on the methodology developed within the project “*Mathematical Models in Machine Learning for Smart Agriculture and High-Value Food*”, funded by the Ministry of Education and Science of the Republic of North Macedonia under the Call for Financing Scientific Research Projects at Public Universities and Scientific Institutes for 2025–2026 (Contract No. 15-6171/22, dated 7 August 2025). No project funds were received at the time of publication. The article processing charge (APC) was covered by the authors.

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding author.

Acknowledgments: The authors acknowledge the institutional framework of the project “*Mathematical Models in Machine Learning for Smart Agriculture and High-Value Food*”.

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

Abbreviations

The following abbreviations are used in this manuscript:

aCI	Adaptive covariance intersection
AICT	Integrated autocorrelation time
CI	Covariance intersection
CTI	Canopy temperature index
DSS	Decision support system
ECE	Expected calibration error
ESS	Effective sample size
ET	Evapotranspiration
F1	F-measure (harmonic mean of precision and recall)
GP	Gaussian process
IoT	Internet of Things
KF	Kalman filter
LRT	Likelihood-ratio test
MAP	Maximum a posteriori
MCMC	Markov Chain Monte Carlo
ML	Machine learning
MSE	Mean-squared error
NLL	Negative log-likelihood
PD	Proportional–derivative (controller)
RMSE	Root-mean-squared error
SLR	Stress likelihood ratio
EKF/UKF	Extended Kalman filter/unscented Kalman filter

References

- Obayemi, A.; Nguyen, K.A. Uncertainty Quantification of Multimodal Models. In *Advances and Trends in Artificial Intelligence. Theory and Applications, Proceedings of the IEA/AIE 2025, Kitakyushu, Japan, 1–4 July 2025*; Fujita, H., Watanobe, Y., Ali, M., Wang, Y., Eds.; Lecture Notes in Computer Science; Springer: Berlin/Heidelberg, Germany, 2026; Volume 15706. [\[CrossRef\]](#)
- Schmerling, E. Multimodal Modeling and Uncertainty Quantification for Robot Planning and Decision Making. Doctoral Dissertation, Stanford University, Stanford, CA, USA, 2019. Available online: <https://stacks.stanford.edu/file/druid:xx848nq8857/SchmerlingPhD-augmented.pdf> (accessed on 8 December 2025).
- Knights, V.A.; Petrovska, O.; Kljusurić, J.G. Symmetry and Asymmetry in Dynamic Modeling and Nonlinear Control of a Mobile Robot. *Symmetry* **2025**, *17*, 1488. [\[CrossRef\]](#)
- Knights, V.A.; Petrovska, O.; Kljusurić, J.G. Nonlinear Dynamics and Machine Learning for Robotic Control Systems in IoT Applications. *Future Internet* **2024**, *16*, 435. [\[CrossRef\]](#)
- Knights, V.; Petrovska, O. Dynamic Modeling and Simulation of Mobile Robot Under Disturbances and Obstacles in an Environment. *J. Appl. Math. Comput.* **2024**, *8*, 59–67. [\[CrossRef\]](#)
- El Moselhy, T.A.; Marzouk, Y.M. Bayesian Inference with Optimal Maps. *J. Comput. Phys.* **2012**, *231*, 7815–7850. [\[CrossRef\]](#)
- Robert, C.P.; Casella, G. *Monte Carlo Statistical Methods*, 2nd ed.; Springer: Berlin, Germany, 2004. [\[CrossRef\]](#)
- Stuart, A.M. Inverse Problems: A Bayesian Perspective. *Acta Numer.* **2010**, *19*, 451–559. [\[CrossRef\]](#)
- Girolami, M.; Calderhead, B. Riemann Manifold Langevin and Hamiltonian Monte Carlo Methods. *J. R. Stat. Soc. B* **2011**, *73*, 123–214. [\[CrossRef\]](#)
- Kalman, R.E. A New Approach to Linear Filtering and Prediction Problems. *ASME J. Basic Eng.* **1960**, *82*, 35–45. [\[CrossRef\]](#)
- Anderson, B.D.O.; Moore, J.B. *Optimal Filtering: In Prentice-Hall Information and System Science Series*, 9th ed.; Kailath, T., Ed.; Prentice Hall: Englewood Cliffs, NJ, USA, 1979; Available online: <http://www.dii.unimore.it/~lbiagiotti/MaterialeKalmanFilter/AndersonMoore2005.pdf> (accessed on 8 December 2025).
- Uhlmann, J.K. Covariance Consistency Methods for Fault-Tolerant Distributed Data Fusion. *Inf. Fusion* **2003**, *4*, 201–215. [\[CrossRef\]](#)
- Julier, S.J.; Uhlmann, J.K. A Non-Divergent Estimation Algorithm in the Presence of Unknown Correlations. In Proceedings of the American Control Conference, Albuquerque, NM, USA, 4–6 June 1997; Volume 4, pp. 2369–2373. [\[CrossRef\]](#)
- Reinhardt, M.; Noack, B.; Arambel, P.O.; Hanebeck, U.D. Minimum Covariance Bounds for the Fusion under Unknown Correlations. *IEEE Signal Process. Lett.* **2015**, *22*, 1210–1214. [\[CrossRef\]](#)
- Rasmussen, C.E.; Williams, C.K.I. *Gaussian Processes for Machine Learning*; MIT Press: Cambridge, MA, USA, 2006.
- Dick, J.; Goda, T.; Murata, H. Toeplitz Monte Carlo. *Stat. Comput.* **2021**, *31*, 1121–1135. [\[CrossRef\]](#)
- Sivia, D.S.; Skilling, J. *Data Analysis: A Bayesian Tutorial*, 2nd ed.; Oxford University Press: Oxford, UK, 2006. Available online: <https://global.oup.com/academic/product/data-analysis-9780198568322> (accessed on 11 December 2025).
- Goodman, J.; Weare, J. Ensemble Samplers with Affine Invariance. *Commun. Appl. Math. Comput. Sci.* **2010**, *5*, 65–80. [\[CrossRef\]](#)
- Foreman-Mackey, D.; Hogg, D.W.; Lang, D.; Goodman, J. Emcee: The MCMC Hammer. *Publ. Astron. Soc. Pac.* **2013**, *125*, 306–312. [\[CrossRef\]](#)
- Knights, V.; Petrovska, O.; Prchkowska, M. Mathematical Analysis of Stochastic Transition Matrices in Markov Models Using Monte Carlo and MCMC Simulation. *J. Nat. Sci. Math. UT* **2024**, *10*, 100–110. [\[CrossRef\]](#)
- Berger, J.O. *Statistical Decision Theory and Bayesian Analysis*, 2nd ed.; Springer: New York, NY, USA, 1985.
- Raiffa, H.; Schlaifer, R. Applied Statistical Decision Theory. Available online: <https://gwern.net/doc/statistics/decision/1961-raiffa-appliedstatisticaldecisiontheory.pdf> (accessed on 11 December 2025).
- Cover, T.M.; Thomas, J.A. *Elements of Information Theory*, 2nd ed.; Wiley: Hoboken, NJ, USA, 1991; Available online: https://cs-114.org/wp-content/uploads/2015/01/Elements_of_Information_Theory_Elements.pdf (accessed on 11 December 2025).
- Whittle, P. *Risk-Sensitive Optimal Control*; Wiley: Chichester, UK, 1990. Available online: <https://scispace.com/pdf/risk-sensitive-optimal-control-3fpaelp8n4.pdf> (accessed on 11 December 2025).
- Russell, J.; Bergmann, J.H.M.; Nagaraja, V.H. Towards Dynamic Multi-Modal Intent Sensing Using Probabilistic Sensor Networks. *Sensors* **2022**, *22*, 2603. [\[CrossRef\]](#) [\[PubMed\]](#)
- Piechocki, R.J.; Wang, X.; Bocus, M.J. Multimodal Sensor Fusion in the Latent Representation Space. *Sci. Rep.* **2023**, *13*, 2005. [\[CrossRef\]](#)
- Mondal, M.; Alam, T.; Ahmed, S.; Rahman, M. A Cross-Comparative Review of Machine Learning for Plant Disease Detection. *Artif. Intell. Agric.* **2024**, *12*, 127–151. [\[CrossRef\]](#)
- Lin, R.; Hu, H. Adapt and explore: Multimodal mixup for representation learning. *Inf. Fusion* **2024**, *105*, 102216. [\[CrossRef\]](#)
- Knights, V.; Petrovska, O.; Prchkowska, M. Machine Learning and Mathematical Modeling in Agricultural Development. In *AI and Digital Transformation: Opportunities, Challenges, and Emerging Threats in Technology, Business, and Security, Proceedings of the ICITBT 2025, Tirana, Albania, 29–30 May 2025*; Dhoska, K., Spaho, E., Eds.; Communications in Computer and Information Science; Springer: Cham, Switzerland, 2026; Volume 2669. [\[CrossRef\]](#)

30. Boyd, S.; Vandenberghe, L. *Convex Optimization*; Cambridge University Press: Cambridge, UK, 2004. Available online: https://web.stanford.edu/~boyd/cvxbook/bv_cvxbook.pdf (accessed on 2 January 2026).
31. Julier, S.J.; Uhlmann, J.K. General decentralized data fusion with covariance intersection. In *Handbook of Multisensor Data Fusion*; Hall, D.L., Llinas, J., Eds.; CRC Press: Boca Raton, FL, USA, 2001; pp. 319–343. Available online: <https://dsp-book.narod.ru/HMDF/2379ch12.pdf> (accessed on 11 December 2025).
32. Bar-Shalom, Y.; Li, X.R.; Kirubarajan, T. *Estimation with Applications to Tracking and Navigation: Theory Algorithms and Software*; John Wiley & Sons: New York, NY, USA, 2002. [CrossRef]
33. Geyer, C.J. Practical Markov Chain Monte Carlo. *Stat. Sci.* **1992**, *7*, 473–511. [CrossRef]
34. Gelman, A.; Rubin, D.B. Inference from Iterative Simulation Using Multiple Sequences. *Stat. Sci.* **1992**, *7*, 457–472. [CrossRef]
35. Wang, X.; Sun, S.; Liu, Y.; Liu, Z. A Fast Covariance Union Algorithm for Inconsistent Sensor Data Fusion. *IEEE Access* **2021**, *9*, 143941–143949. [CrossRef]
36. Mulla, D.J. Twenty Five Years of Remote Sensing in Precision Agriculture: Key Advances and Remaining Knowledge Gaps. *Biosyst. Eng.* **2013**, *114*, 358–371. [CrossRef]
37. Liang, W.; Zhang, Y.; Chong, A.; Cochran Hameen, E.; Loftness, V. Exploring Gaussian Process Regression for Indoor Environmental Quality: Spatiotemporal Thermal and Air Quality Modeling with Mobile Sensing. *Build. Environ.* **2025**, *281*, 113143. [CrossRef]
38. Dubois, D.; Fargier, H.; Prade, H. Multiple-Sources Information Fusion—A Practical Inconsistency-Tolerant Approach. In Proceedings of the 8th International Conference on Information Processing and Management of Uncertainty in Knowledge-Based Systems (IPMU 2000), Madrid, Spain, 3–7 July 2000; pp. 1047–1054. Available online: <https://hal.science/hal-03405306/document> (accessed on 11 December 2025).
39. Nandgude, N.; Singh, T.P.; Nandgude, S.; Tiwari, M. Drought Prediction: A Comprehensive Review of Different Drought Prediction Models and Adopted Technologies. *Sustainability* **2023**, *15*, 11684. [CrossRef]
40. Wang, S.; Yu, Z.; Aorigele; Zhang, W. Study on the modeling method of sunflower seed particles based on the discrete element method. *Comput. Electron. Agric.* **2022**, *198*, 107012. [CrossRef]
41. Kim, S.; Kim, S.; Hwang, S.; Lee, H.; Kwak, J.; Song, J.-H.; Jun, S.-M.; Kang, M.-S. Impact assessment of water-level management on water quality in an estuary reservoir using a watershed-reservoir linkage model. *Agric. Water Manag.* **2023**, *280*, 108234. [CrossRef]
42. Antoska Knights, V.; Gacovski, Z. *Methods for Detection and Prevention of Vulnerabilities in the IoT (Internet of Things) Systems [Internet]. Internet of Things—New Insights*; IntechOpen: London, UK, 2024. [CrossRef]
43. Knights, V.A.; Stankovski, M.; Nusev, S.; Temeljkovski, D.; Petrovska, O. Robots for Safety and Health at Work. *Mech. Eng. Sci. J.* **2015**, *33*, 275–279. Available online: <https://www.mesj.ukim.edu.mk/journals/issue/download/17/9> (accessed on 15 November 2025).
44. Antoska, V.; Jovanović, K.; Petrović, V.M.; Bašćarević, N.; Stankovski, M. Balance Analysis of the Mobile Anthropomorphic Robot Under Disturbances—ZMP Approach. *Int. J. Adv. Robot. Syst.* **2013**, *10*, 206. [CrossRef]
45. Shi, X.; Han, W.; Zhao, T.; Tang, J. Decision Support System for Variable Rate Irrigation Based on UAV Multispectral Remote Sensing. *Sensors* **2019**, *19*, 2880. [CrossRef]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.