Received: 4 June 2026 | Revised: 10 August 2026 | Accepted: 28 August 2026 | Published online: 16 September 2026

RESEARCH ARTICLE

Journal of Data Science and Intelligent Systems
2026, Vol. 00(00) 1–14
DOI: 10.47852/bonviewJDSIS620210658

BON VIEW PUBLISHING

Integrating Leaf-Image Deep Learning, Meteorological Risk-State Discovery, and Stochastic SICR Dynamics for Potato Early and Late Blight: A Data-Science and Intelligent-Systems Framework

Hongzhi Zhang[1,2], L. P. Varlamova[1,*] and Jianan Chen[3]

[1] Faculty of Applied Mathematics and Intellectual Technologies, National University of Uzbekistan named after Mirzo Ulugbek, Uzbekistan

[2] School of Mathematics and Statistics, Jining Normal University, China

[3] School of Physics and Electronic Information Engineering, Jining Normal University, China

Abstract: Potato early blight and late blight create substantial agronomic risk, yet leaf-image diagnosis, weather-driven warning, and transmission modeling are often treated as separate tasks. This paper develops an integrated data-science and intelligent-systems framework that combines convolutional neural network-based symptom recognition, meteorological environmental suitability modeling, AI-based weather-risk-state learning, and a stochastic susceptible-infected-carrier-recovered transmission mechanism. Leaf-image evidence is drawn from the Potato Disease Leaf Dataset. ResNet50 transfer learning provides a cross-source generalization baseline, whereas Xception fine-tuning provides the stronger image-classification branch. Weather-risk modeling uses daily temperature, relative humidity, and rainfall from Chahar Right Front Banner, Ulanqab, situated within an expanded three-growing-season record (2023–2025). Disease-specific environmental suitability indices are computed using Gaussian temperature, sigmoid humidity, and rainfall saturation responses; K-means clustering identifies recurring weather-risk states, and a cross-validated logistic surrogate classifier learns threshold-defined warning labels. Image performance is strongly protocol-dependent: ResNet50 reached 66.32% accuracy under a stricter cross-source test, whereas fine-tuned Xception exceeded 95% under the source split, so the two are not directly comparable. August 2025 showed markedly higher suitability than August 2024, and K-means recovered a 16-day high-risk state. Because the humidity-complete window is limited to two August periods (62 days), these warning-state results are demonstrative rather than broadly generalizable. The study contributes a transparent, modular, and reproducible early-warning framework while explicitly distinguishing surrogate warning-state learning from field-validated disease-incidence prediction.

Keywords: potato early blight, potato late blight, environmental suitability index, convolutional neural network, stochastic SICR model

1. Introduction

Potato production is exposed to biologically complex and environmentally sensitive disease pressures. Early blight, commonly associated with *Alternaria solani*, and late blight, associated with *Phytophthora infestans*, are especially important because their early foliar symptoms can be confused during field inspection and missed warning windows may delay management [1, 2]. In many production areas, growers still rely on visual scouting, local experience, and calendar-based management. These practices remain useful, but they are insufficient when environmental conditions change rapidly or when disease risk must be assessed across large areas.

The present study integrates four complementary analytical components: a ResNet-based potato leaf-classification baseline, an Xception-based progressive classification experiment, a stochastic susceptible-infected-carrier-recovered (SICR) transmission model, and an AI-augmented environmental suitability warning framework. Considered separately, each component has a narrower publication angle. A leaf-image analysis alone may be read as a conventional transfer-learning case study; a stochastic SICR analysis alone may remain too theoretical for a data-science journal; and an environmental suitability model alone may be

*Corresponding author: L. P. Varlamova, Faculty of Applied Mathematics and Intellectual Technologies, National University of Uzbekistan named after Mirzo Ulugbek, Uzbekistan. Email: vlp@bk.ru

regarded as a rule-based warning model with limited algorithmic depth. The stronger contribution is to integrate these components into a coherent data-science and intelligent-systems architecture.

This reframing is aligned with a journal scope that emphasizes data integration, knowledge extraction, AI-enabled analysis, and application-oriented intelligent systems. The manuscript addresses the following research question: how can leaf-image symptom recognition, meteorological risk-state learning, and stochastic transmission mechanisms be combined into a reproducible decision-support framework for potato early and late blight? The proposed answer is a modular pipeline. The image branch detects visible symptoms, the weather branch estimates pre-symptomatic environmental risk, the AI branch learns risk states and surrogate warning labels, and the stochastic SICR branch provides a mechanistic interpretation of how environmental risk may translate into transmission pressure.

The main contribution is integrative rather than a set of separate results. The framework (i) consolidates heterogeneous modalities, including potato leaf images and daily meteorological observations, within one intelligent-system architecture; (ii) explains why cross-source and random-split convolutional neural network (CNN) evaluations must not be interpreted as equivalent; (iii) upgrades the static suitability index into a data-science pipeline with entropy-weight robustness analysis, K-means clustering, and cross-validated logistic surrogate classification; and (iv) couples environmental suitability to a stochastic SICR transmission coefficient, retaining biological interpretability while allowing data-driven warning signals to enter the dynamical model.

The paper is deliberately conservative in its claims. The humidity-complete weather demonstration contains 62 daily observations from two August periods, situated within an expanded record of 459 daily temperature and rainfall observations spanning the May–September growing seasons of 2023–2025; the current dataset does not include field-confirmed incidence labels. Accordingly, the AI classifiers reported here are surrogate warning-state learners rather than validated disease-incidence predictors. This distinction is essential for scientific credibility. The framework is presented as a reproducible and extensible intelligent warning system that can be strengthened when multi-year field incidence, leaf wetness, cultivar resistance, fungicide records, remote-sensing data, and georeferenced disease observations become available.

2. Related Work and Positioning

2.1. Deep learning for image-based plant disease recognition

Deep learning has become a dominant approach for image-based plant disease recognition because CNNs can learn hierarchical visual representations from raw images [3–5]. Recent potato-specific studies combine deep learning with explainable AI and report high source-split accuracy and severity assessment for early blight, late blight, and healthy leaves [6–9]. ResNet addresses the degradation problem of very deep networks by using residual connections, making it possible to train deeper architectures effectively [10]. Xception replaces standard convolution with depthwise separable convolution and improves parameter efficiency while retaining strong feature-extraction capacity [11]. Both architectures are plausible candidates for potato leaf disease classification because early blight, late blight, and healthy leaves differ in texture, lesion geometry, color distribution, and disease-spot boundary characteristics. Broader reviews and efficient models confirm the maturity of image-based crop-disease recognition [12, 13], yet most results are obtained on visually homogeneous data, and systems that fuse leaf images with environmental signals in a genuine multimodal manner remain comparatively rare [14].

However, image-based plant disease recognition is highly sensitive to dataset construction. Models trained and tested on randomly partitioned images from a visually homogeneous dataset may report high accuracy yet fail under cross-region, cross-background, or cross-device deployment. Open plant-health image repositories and transfer-learning strategies have accelerated model development while also highlighting the importance of deployment-oriented validation with limited labeled data [15]. The ResNet component is valuable because it uses a development-test separation designed to introduce differences in region, lighting, background, and variety. The Xception component demonstrates how progressive optimization can improve performance under a structured train–validation–test split. The integrated manuscript therefore treats these two results as complementary evidence rather than as a direct one-to-one architecture benchmark.

2.2. Meteorological disease-risk modeling and environmental suitability

Early blight and late blight are not purely image-recognition problems. Before visible symptoms are severe, disease risk is shaped by temperature, relative humidity (RH), rainfall, canopy moisture, and pathogen biology. Weather-driven warning systems are valuable because they can operate before disease symptoms become visually obvious [16]. Machine-learning models built on temperature, humidity, and rainfall have been used to predict late blight risk level, disease incidence, and the timing of first appearance [17–19], and recent work explores generative pretrained and deep sequence models for late blight forecasting as well as deep learning on sequential environmental data [20, 21]. A suitability index is best understood as a feature-engineering layer: it transforms raw environmental observations into biologically interpretable risk features that can be combined with machine learning and decision-support rules.

The environmental suitability model used here contains three disease-specific response functions: a Gaussian temperature response, a sigmoid humidity response, and an exponential rainfall saturation response. This structure reflects the assumption that pathogen activity has an optimum temperature range, that infection risk increases sharply after humidity thresholds, and that rainfall has diminishing marginal contribution once sufficient wetness is present. A geometric aggregation is used because disease risk is constrained by the least favorable essential factor.

2.3. AI-based risk-state discovery and surrogate classification

For the supplied meteorological dataset, unsupervised and transparent supervised methods are more defensible than complex deep-learning sequence models. K-means clustering can discover recurring weather-risk regimes without requiring disease-incidence labels [22, 23]. Logistic regression can serve as a transparent surrogate classifier that learns threshold-defined warning labels from raw and engineered weather features [24]. The resulting metrics evaluate whether environmental suitability index (ESI)-defined warning logic is learnable from the meteorological feature space, not whether true disease incidence has been predicted.

This distinction is important for peer review. In the absence of field-confirmed incidence data, using random forests, gradient boosting, or long short-term memory (LSTM) models to claim disease prediction would be methodologically weak. Such algorithms may be appropriate future extensions after incidence labels and multi-year observations are collected. In the current manuscript, AI is used only where the data support it: unsupervised risk-state discovery and surrogate learning of transparent warning labels.

2.4. Stochastic SICR modeling as a mechanistic layer

Stochastic epidemic models are useful when transmission is affected by random environmental disturbances and unobserved field variability [25–27]. The SICR structure divides the host population into susceptible, acutely infected, carrier, and recovered classes. For potato blight, this structure is biologically interpretable because infected plants may show acute symptoms, latent or carrier-like states may contribute to pathogen persistence, and management or plant resistance may move part of the population into a recovered or less susceptible state.

The stochastic SICR component provides theoretical properties, including positivity, stationary distribution, extinction, and persistence conditions. Its main limitation for a data-science journal is that it is weakly linked to observed data. The present framework addresses that limitation by allowing environmental suitability to drive the transmission coefficient. This turns the stochastic model into a mechanistic layer that can absorb warning signals from the data pipeline.

3. Materials and Methods

3.1. Overall framework

Figure 1 summarizes the compact integrated architecture. The framework contains two evidence streams. The first stream is symptom evidence derived from leaf images, which compares a ResNet50 transfer-learning branch and an Xception fine-tuning branch. The second stream is pre-symptomatic environmental risk evidence derived from daily meteorological observations. These two streams are connected to a stochastic SICR mechanism, which interprets high environmental suitability as a period of increased transmission pressure. The output is a daily decision-support signal containing risk level, risk-state cluster, classifier probability, and suggested scouting priority.

3.2. Leaf-image data and classification protocols

The image component uses the Potato Disease Leaf Dataset, which contains early blight, late blight, and healthy-leaf categories. In the ResNet component, the dataset is organized into a development set and an independent test set. The development set contains 7211 images, and the independent test set contains 793 images (363 early blight, 67 late blight, and 363 healthy (Table 1)). At evaluation time, 15 of these 793 test images could not be read or decoded and were therefore excluded from scoring, so all ResNet50 test metrics reported in this paper are computed on the remaining 778 evaluable images. The authors explicitly designed the test set to differ from the development set in region, lighting, background, device, and variety where possible. This design creates a more stringent generalization test than a random

Figure 1
Integrated data-science and intelligent-system architecture for potato blight early warning

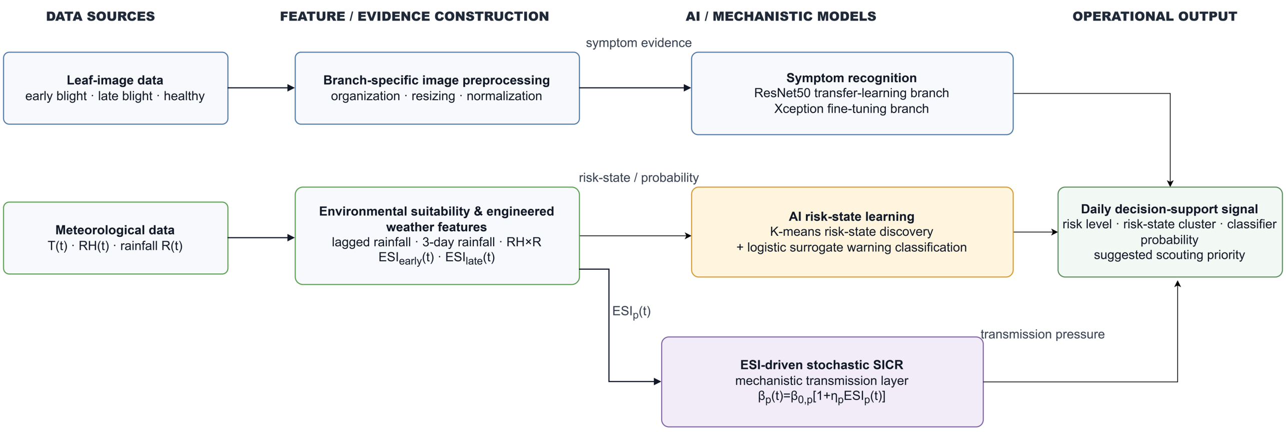

Table 1
Leaf-image data protocols consolidated from the ResNet and Xception analytical components

Dataset protocol	Early blight	Late blight	Healthy	Total	Purpose
ResNet development set	3615	2424	1172	7211	Model development and validation
ResNet independent test set	363	67	363	793	Cross-source/generalization-oriented evaluation
Xception training set	1989	1245	767	4001	Parameter learning
Xception validation set	994	622	383	1999	Hyperparameter monitoring
Xception test set	994	622	383	1999	Final split-based evaluation

split and explains why the ResNet result is treated as evidence of cross-source robustness rather than as a direct counterpart to the Xception split-based result.

The Xception component uses a normalized three-way train–validation–test split after preprocessing and folder standardization. The reported split contains 4001 training images, 1999 validation images, and 1999 test images. The difference between the ResNet and Xception data protocols is explicitly retained in this integrated paper. Results are therefore interpreted by protocol rather than used for an oversimplified model ranking.

3.3. CNN-based symptom recognition

The ResNet branch uses a pretrained ResNet50 encoder with frozen pretrained layers and trainable dense layers. Images are resized to 232 × 232 pixels and normalized using ImageNet statistics because the images are natural RGB field or leaf photographs rather than specialized medical or satellite images. This use of ImageNet-pretrained feature extractors follows the broader transfer-learning tradition in CNN-based visual recognition [15]. The reported environment includes Python 3.11, PyTorch 2.4, CUDA 12.4, and an NVIDIA GeForce RTX 4060 Ti GPU. The initial learning rate is 0.0001, with early stopping patience of 5.

The Xception branch follows a progressive experimental design. A two-convolution baseline CNN first establishes a performance floor. Data augmentation then increases sample diversity through random translation, scaling, rotation, brightness adjustment, and histogram equalization. Transfer learning freezes the pretrained Xception convolutional base and trains a new classifier. Fine-tuning then unfreezes the upper Xception block and uses a smaller learning rate to adapt the model to disease-specific features while limiting destructive parameter updates.

3.4. Meteorological data and environmental suitability index

The meteorological component uses daily weather records for Chahar Right Front Banner, Ulanqab City, Inner Mongolia, China. The humidity-complete demonstration covers August 2024 and August 2025, giving 62 daily records, each with mean temperature T, relative humidity RH, and rainfall R. To situate this window within a longer climatic context, an expanded record of 459 daily temperature and rainfall observations spanning the May–September growing seasons of 2023, 2024, and 2025 is also analyzed. RH, which is a required ESI input, was available only for the two August periods; therefore, the disease-specific ESI, clustering, and surrogate-classification analyses are reported for that RH-complete window and interpreted conservatively. This is not a large machine-learning dataset, but it is sufficient to demonstrate a reproducible data pipeline and to compare disease-risk pressure under clearly defined warning rules.

For each disease p, the temperature, humidity, and rainfall suitability components are computed as shown in Equations (1)–(3) and are aggregated into the disease-specific index in Equation (4). These component functions are central to the interpretability of the weather-risk branch. The disease-specific response parameters used in the demonstration are listed in Table 2.

$$G_p(T) = exp\left[-\left(T - \mu_p\right)^2 / \left(2\sigma_p^2\right)\right] \tag{1}$$

$$H_p(RH) = 1/\left\{1 + exp\left[-k_p\left(RH - h_p\right)\right]\right\} \tag{2}$$

$$P_p(R) = 1 - exp\left(-R/r_p\right) \tag{3}$$

$$ESI_p(t) = exp\left\{w_T \, ln\, G_p(t) + w_H \, ln\, H_p(t) + w_R \, ln\, P_p(t)\right\} \tag{4}$$

3.5. Entropy-weight robustness analysis

The default ESI uses equal weights because field incidence data are unavailable for parameter calibration. To test whether the main year-to-year conclusion depends on this assumption, entropy weighting is used as a robustness analysis. For each component j, normalized values are used to compute the entropy, diversity, and final weight of each component, as given in Equation (5):

$$e_j = -[1/\ln(n)]\,\Sigma_i\, p_{ij}\ln\left(p_{ij}\right),\ d_j = 1 - e_j,\quad w_j = d_j/\Sigma_j d_j \tag{5}$$

Entropy weighting is not treated as a biological calibration. It is used only to determine whether the conclusion that August 2025 was riskier than August 2024 remains stable when the weights are driven by the data distribution rather than set equally.

3.6. AI risk-state discovery and surrogate warning classification

K-means clustering is applied to standardized feature vectors containing temperature, humidity, rainfall, three-day rainfall, early blight ESI, and late blight ESI. Three clusters are used because the operational decision objective is to separate low-, moderate-, and high-risk weather states. This choice was checked against internal cluster-validity indices, namely, the silhouette coefficient, the Davies–Bouldin index, and the Calinski–Harabasz index, together with the within-cluster inertia (elbow) across 2–6 clusters. Although a two-cluster partition maximizes the silhouette score, the three-cluster solution is retained because it maps directly onto the three operational warning tiers while still isolating a clearly separated high-risk regime. Cluster labels are ordered after fitting by the mean of the early and late blight ESI values so that the labels are interpretable.

Table 2
Disease-specific environmental suitability parameters used in the demonstration

Disease	Temperature response	Humidity response	Rainfall response	Interpretation
Early blight	μ = 25 °C; σ = 5 °C	h = 80%; k = 0.20	r = 10 mm	Warmer optimum and moderate rainfall saturation
Late blight	μ = 18 °C; σ = 4 °C	h = 90%; k = 0.25	r = 15 mm	Cooler optimum and stronger high-humidity dependence

A logistic regression surrogate classifier is then trained to learn ESI-defined warning labels from raw and engineered meteorological inputs. Predictors include T, RH, R, previous-day rainfall, three-day rainfall, RH x R, and temperature deviation from each disease-specific optimum. Labels are produced by ESI thresholds rather than by field incidence. Class imbalance is handled through inverse-frequency weights, and five-fold stratified cross-validation is used. Metrics include accuracy, balanced accuracy, precision, recall, F1 score, and the area under the receiver operating characteristic curve (AUC).

3.7. ESI-driven stochastic SICR mechanism

The stochastic SICR subsystem partitions the host population into susceptible plants S(t), acutely infected plants $I(t)$, carrier plants $C(t)$, and recovered plants $R(t)$. The disease-transmission coefficient is coupled to environmental suitability as given in Equation (6):

$$\beta_p(t) = \beta_{0,p}\left[1 + \eta_p ESI_p(t)\right] \tag{6}$$

The stochastic subsystem can be written in reduced form as Equations (7)–(10):

$$\mathrm{d}S = \left[\Lambda - \beta_p(t)\,S(I + \theta C) - \mu S\right]\mathrm{d}t + \sigma_1 S\mathrm{d}B_1(t) \tag{7}$$

$$\mathrm{d}I = \left[\beta_p(t)\,S(I + \theta C) - (\alpha + \gamma_I + \mu)\,I\right]\mathrm{d}t + \sigma_2 I\mathrm{d}B_2(t) \tag{8}$$

$$\mathrm{d}C = \left[\alpha I - (\gamma_C + \mu)\,C\right]\mathrm{d}t + \sigma_3 C\mathrm{d}B_3(t) \tag{9}$$

$$\mathrm{d}R = \left[\gamma_I I + \gamma_C C - \mu R\right]\mathrm{d}t \tag{10}$$

When $\eta_p = 0$ or ESI is constant, the model reduces to the constant-coefficient stochastic SICR structure. In this manuscript, SICR is used as the mechanistic coupling layer; calibrated Monte Carlo incidence prediction is not claimed because field incidence labels and independently confirmed parameter values are not available in the current dataset.

4. Results

4.1. Image-classification results

The image-classification results show both the value and the limitation of transfer learning. Under the stricter ResNet50 protocol, which emphasizes cross-source generalization, 15 of the 793 assembled independent-test images could not be read or decoded and were excluded, leaving 778 evaluable images. The model correctly classified 516 of these 778 images, giving an independent-test accuracy of 66.32% (516/778), with macro recall of 0.739 and macro F1 score of 0.594. Reporting recall and F1 alongside accuracy is important because, under class imbalance and domain shift, accuracy alone would overstate performance; the protocol-aware metrics are consolidated in Table 3. Under the Xception protocol, the baseline CNN achieved 77.9% test accuracy and showed clear overfitting. Adding data augmentation improved accuracy to 82.2% and reduced the training–validation gap. Xception transfer learning raised test accuracy above 91%, and fine-tuning raised it above 95% (Table 4). Because the two branches use different datasets and evaluation protocols, their headline accuracies are reported side by side only with explicit caveats and are not interpreted as a controlled architecture comparison; a like-for-like ResNet50-versus-Xception benchmark under one shared split and metric set is identified as necessary future work (Section 6). Figure 2 presents these protocol-aware results in two separate panels.

These results should not be interpreted as proving that Xception is universally superior to ResNet50. The stricter ResNet50 test was designed to contain distribution shifts, whereas the Xception result comes from a different split and optimization protocol. The scientifically safer conclusion is that potato disease recognition is highly sensitive to data split design and deployment-domain mismatch. For an intelligent agricultural system, cross-source evaluation should be retained as a mandatory validation step before field deployment.

4.2. Environmental suitability patterns

The equal-weight ESI results show a sharp contrast between August 2024 and August 2025. August 2024 had only two rainy days and total rainfall of 20.27 mm. August 2025 had 27 rainy

Table 3
Cross-source generalization result (ResNet50 branch)

Model/strategy	Protocol	Reported metric	Interpretation
ResNet50 transfer learning	Development set + stricter independent test	Accuracy = 66.32%; recall = 0.739; F1 = 0.594	Cross-source generalization remains difficult

Note: Of the 793 assembled independent-test images, 15 could not be read or decoded and were excluded; the metrics above are computed on the 778 evaluable images (accuracy = 66.32% = 516/778).

Table 4
Source-split progressive results (Xception branch)

Model/strategy	Protocol	Reported metric	Interpretation
Basic CNN	Xception source split	Accuracy = 77.9%	Baseline with obvious overfitting
CNN + augmentation	Xception source split	Accuracy = 82.2%	Improved generalization but still limited
Xception transfer learning	Xception source split	Accuracy > 91%	Pretrained features substantially improve recognition
Xception fine-tuning	Xception source split	Accuracy > 95%	Best reported split-based performance

Figure 2
Reported image-classification performance under two different protocols, shown as separate panels: (a) ResNet50 cross-source generalization test and (b) Xception source-split progressive results. Protocol differences must be considered when interpreting performance (empirical figure derived from the reported results)

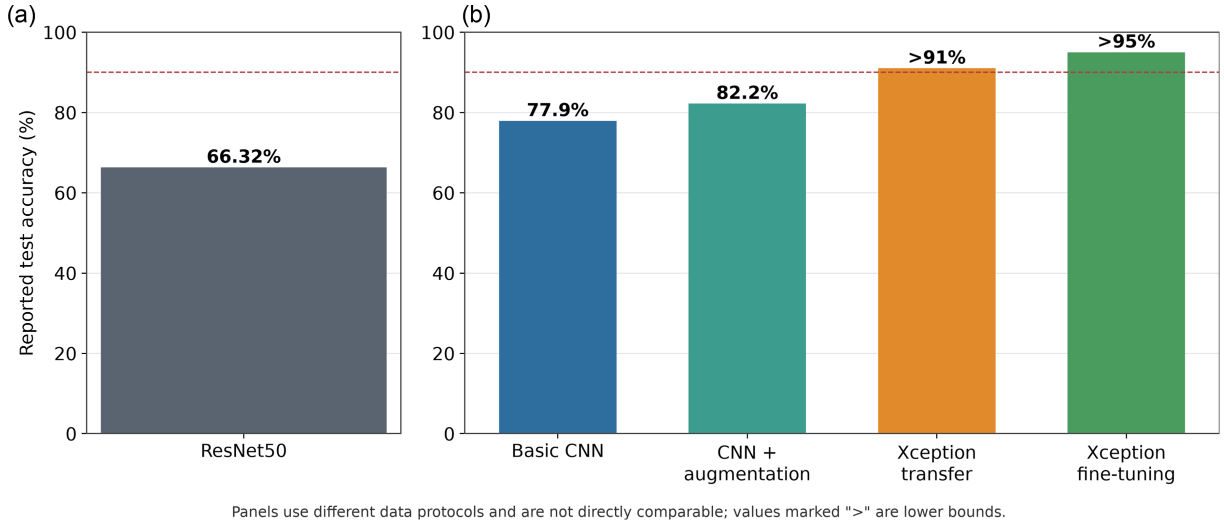

days and total rainfall of 575.12 mm. For early blight, cumulative ESI increased from 0.9295 in August 2024 to 14.0575 in August 2025. For late blight, cumulative ESI increased from 0.9633 to 11.3069. High-risk days increased from 0 to 8 for early blight and from 1 to 7 for late blight (Table 5 and Figure 3).

4.2.1. Three-growing-season meteorological context

The expanded record confirms that the elevated 2025 suitability reflects an anomalously wet growing season rather than an isolated August. Growing-season (May–September) rainfall reached 2273.8 mm in 2025, compared with 350.2 mm in 2023 and 186.9 mm in 2024, and 2025 recorded 27 heavy-rain days (at least 25 mm) versus two and one, respectively, while the growing-season mean temperature remained close to 18 °C in all three years (Table 6 and Figure 4). Because RH was recorded only in August, the disease-specific ESI analysis remains restricted to the two August periods; nevertheless, the broader temperature and rainfall record shows that the moisture-driven risk contrast between 2024 and 2025 is robust within a substantially larger sample.

4.3. Threshold sensitivity and entropy-weight robustness

Threshold sensitivity supports the robustness of the year-to-year contrast. At the moderate threshold of 0.40, August 2025 has 17 early blight days and 11 late blight days, compared with one day for each disease in 2024. At the strict threshold of 0.80, August 2025 still has four early blight days and three late blight days, whereas August 2024 has none (Table 7).

Entropy weighting assigns the largest weight to rainfall suitability because rainfall varies most sharply across the 62-day dataset. For early blight, weights are 0.0419 for temperature, 0.1911 for humidity, and 0.7670 for rainfall. For late blight, weights are 0.0156 for temperature, 0.3465 for humidity, and 0.6380 for rainfall. These weights do not overturn the main conclusion: August 2025 remains much riskier than August 2024 for both diseases (Table 8).

4.4. AI-based risk-state discovery and surrogate warning classification

K-means clustering separates the 62 daily observations into interpretable weather-risk states. The high-risk weather state contains 16 days and is characterized by high humidity, repeated rainfall, and accumulated wetness. Its mean RH is 89.81%, mean daily rainfall is 35.31 mm, mean three-day rainfall is 72.08 mm, mean early blight ESI is 0.7083, and mean late blight ESI is 0.6369. This cluster contains the main threshold-defined high-risk days for both diseases, indicating that unsupervised clustering can recover a biologically meaningful risk regime without observed disease labels (Table 9).

Figure 5 visualizes the K-means cluster profile and clarifies why the 16-day high-risk state is epidemiologically meaningful: it combines high humidity, intense rainfall, accumulated wetness,

Table 5
Equal-weight ESI summary for potato early and late blight

Year	Disease	Mean ESI	Max ESI	Date of max	Cumulative ESI	High-risk days	Moderate/high days	Rainy days	Rainfall (mm)
2024	Early blight	0.0300	0.6062	2024-08-25	0.9295	0	1	2	20.27
2024	Late blight	0.0311	0.7982	2024-08-25	0.9633	1	1	2	20.27
2025	Early blight	0.4535	0.9213	2025-08-04	14.0575	8	17	27	575.12
2025	Late blight	0.3647	0.8909	2025-08-15	11.3069	7	11	27	575.12

Figure 3
Monthly environmental suitability burden and threshold-defined high-risk day counts for August 2024 and August 2025 (empirical figure derived from the ESI analysis)

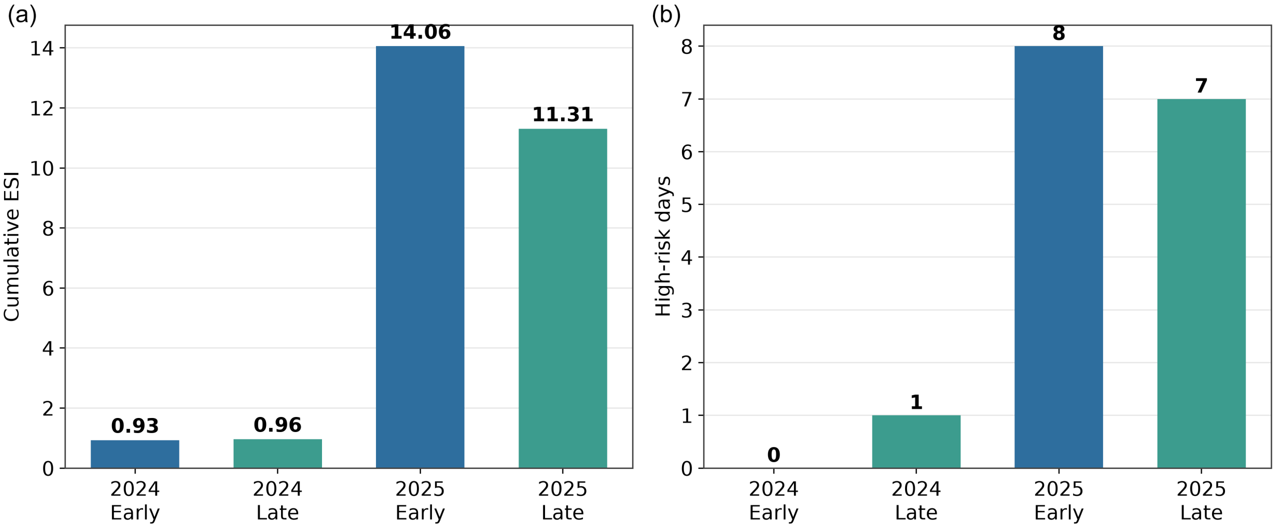

Table 6
Three-growing-season (May–September 2023–2025) temperature and rainfall summary for the study location

Year	Days	Mean T (C)	Total rainfall (mm)	Rainy days	Heavy-rain days (>= 25 mm)	Max daily rainfall (mm)
2023	153	18.00	350.24	41	2	88.74
2024	153	18.15	186.90	17	1	42.75
2025	153	17.99	2273.75	96	27	185.03

Figure 4
Three-growing-season meteorological context: (a) monthly growing-season rainfall for 2023–2025 and (b) May–September rainfall totals by year (empirical figure derived from the expanded meteorological record)

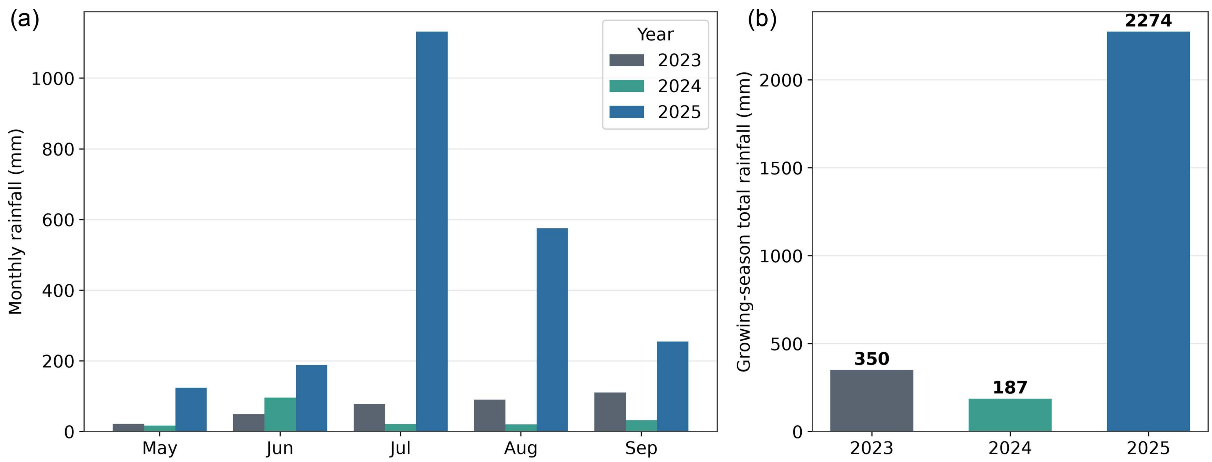

Table 7
Threshold sensitivity of high/moderate-risk day counts

Threshold	2024 early	2024 late	2025 early	2025 late
0.40	1	1	17	11
0.50	1	1	16	10
0.60	1	1	12	8
0.70	0	1	8	7
0.80	0	0	4	3

Table 8
Entropy-weight robustness results

Disease	Year	wT	wH	wR	Max entropy ESI	Cumulative entropy ESI	High-risk days	Moderate/high days
Early	2024	0.0419	0.1911	0.7670	0.8273	0.9946	1	1
Early	2025	0.0419	0.1911	0.7670	0.9763	13.9723	11	16
Late	2024	0.0156	0.3465	0.6380	0.7257	0.8027	1	1
Late	2025	0.0156	0.3465	0.6380	0.9189	10.0596	6	11

Table 9
Interpreted K-means weather-risk states

Risk state	Days	Mean RH (%)	Mean rainfall (mm)	Mean 3-day rainfall (mm)	Mean early ESI	Mean late ESI	Interpretation
Low-risk weather state	18	67.17	0.26	18.50	0.0469	0.0172	Low moisture pressure and low suitability
Moderate-risk weather state	28	81.57	0.92	10.44	0.1003	0.0632	Humid conditions with limited rainfall intensity
High-risk weather state	16	89.81	35.31	72.08	0.7083	0.6369	Sustained wetness and high blight suitability

Figure 5
AI-discovered weather-risk states and their epidemiological interpretation (empirical figure derived from the K-means analysis)

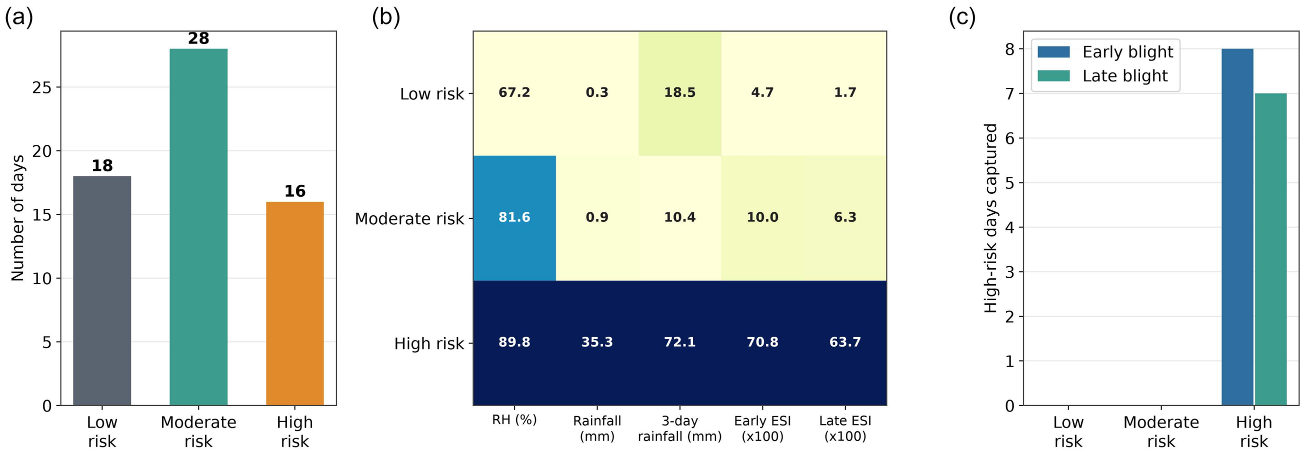

The high-risk cluster combines high humidity, intense rainfall, accumulated wetness and high suitability; labels are ordered by mean ESI after fitting.

Note: Cluster labels are ordered by mean ESI after fitting; high-risk day counts are threshold-defined warning labels rather than field-confirmed incidence counts.

and high suitability scores for both early and late blight. The internal cluster-validity indices supporting the choice of three clusters are reported in Table 10.

The logistic surrogate classifier achieves high cross-validated AUC values when learning ESI-defined warning labels. For early blight, AUC is 0.9848 at the 0.40 threshold and 0.9560 at the 0.70 threshold. For late blight, AUC is 0.9900 at the 0.40 threshold and 0.9907 at the 0.70 threshold. These results demonstrate feature consistency and learnability of the warning rules, but they should not be interpreted as independent biological validation because the labels are produced from ESI thresholds rather than field observations (Table 11).

4.5. Connection to stochastic SICR threshold dynamics

The stochastic SICR layer gives the framework a mechanistic interpretation. High ESI days and high-risk clusters can be treated as periods in which the time-varying transmission coefficient $\beta_p(t)$ increases. This coupling links environmental forcing to the transmission coefficient and separates deterministic state transitions from stochastic perturbations. If field incidence data are later collected, the coupling strength can be calibrated to quantify how weather-risk states amplify transmission. Monte Carlo simulation can then estimate peak infection, extinction frequency, time

Table 10
Internal cluster-validity indices for K-means solutions with 2–6 clusters

Number of clusters (k)	Silhouette	Davies–Bouldin	Calinski–Harabasz	Within-cluster inertia
2	0.446	0.983	44.59	213.41
3	0.277	1.220	32.99	175.61
4	0.294	1.208	30.75	143.59
5	0.270	1.113	28.68	123.47
6	0.276	1.053	27.18	108.55

Note: The three-cluster solution is retained for operational interpretability (low-, moderate-, and high-risk tiers); the within-cluster inertia shows diminishing returns beyond three clusters, and the high-risk regime remains clearly separated.

Table 11
Cross-validated logistic surrogate classification of ESI-defined warning labels

Disease	ESI threshold	Positive days	Accuracy	Balanced accuracy	Precision	Recall	F1	AUC
Early	0.40	18	0.9355	0.8889	1.0000	0.7778	0.8750	0.9848
Early	0.70	8	0.9032	0.8912	0.5833	0.8750	0.7000	0.9560
Late	0.40	12	0.9677	0.9800	0.8571	1.0000	0.9231	0.9900
Late	0.70	8	0.9677	0.9282	0.8750	0.8750	0.8750	0.9907

to peak, and the probability of sustained persistence. The current manuscript does not report calibrated incidence simulations because the required incidence labels and parameter-confirmation dataset are not available. Instead, it provides a transparent computational pathway from daily meteorological observations to stochastic transmission analysis.

5. Discussion

The integrated results demonstrate why the study is stronger as a data-science and intelligent-systems paper than as four separate analytical components. The CNN branch shows that symptom recognition is feasible but strongly dependent on the validation protocol. The meteorological branch shows that daily weather data can be transformed into interpretable disease-specific risk features. The AI branch adds algorithmic structure through clustering and surrogate classification. The stochastic SICR branch links warning signals to disease-transmission mechanisms. Figures 6, 7, 8, 9, 10, 11, and 12 are therefore not decorative additions; they connect the method, result, and deployment layers into a single intelligible manuscript narrative.

The most important methodological lesson is that evaluation design must match deployment claims. A model with over 95% split-based accuracy may still perform poorly if deployed in different regions, varieties, backgrounds, or imaging devices. Conversely, a lower cross-source accuracy may be more informative for field robustness because it exposes domain shift. Therefore, the manuscript should not overstate CNN performance. It should present the Xception branch as a promising symptom classifier and the ResNet branch as evidence that true field generalization remains challenging.

For weather-risk learning, the main substantive finding is the sharp difference between August 2024 and August 2025. The 2025 period is characterized by frequent rainfall, high humidity, and accumulated moisture, which jointly produce sustained early and late blight suitability. Cumulative ESI is more actionable than an isolated maximum value because disease management often depends on sustained risk pressure rather than a single

Figure 6
Disease-specific environmental suitability response functions for early blight and late blight. The curves are schematic visualizations of the ESI component functions used to improve interpretability of the meteorological modeling layer

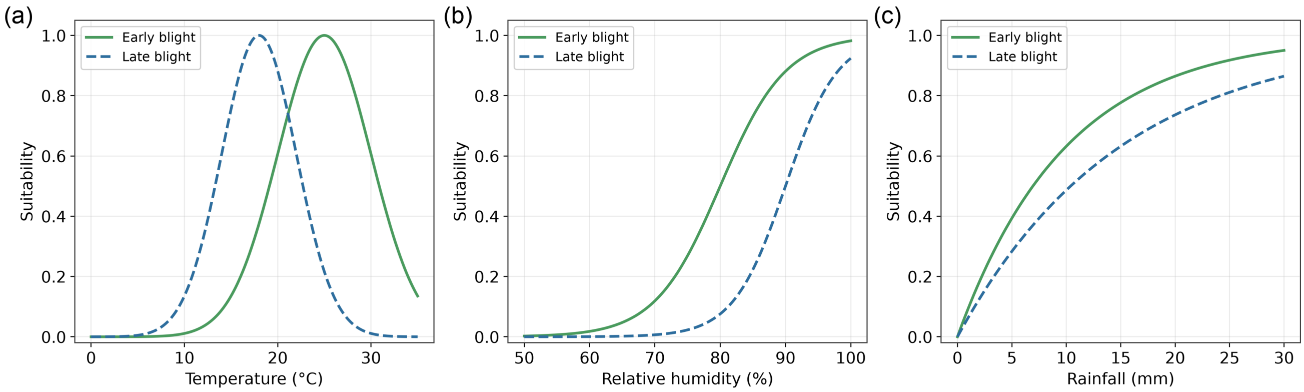

Figure 7
Integrated multi-source research framework linking heterogeneous inputs, analytical modules, and actionable outputs for potato blight early warning

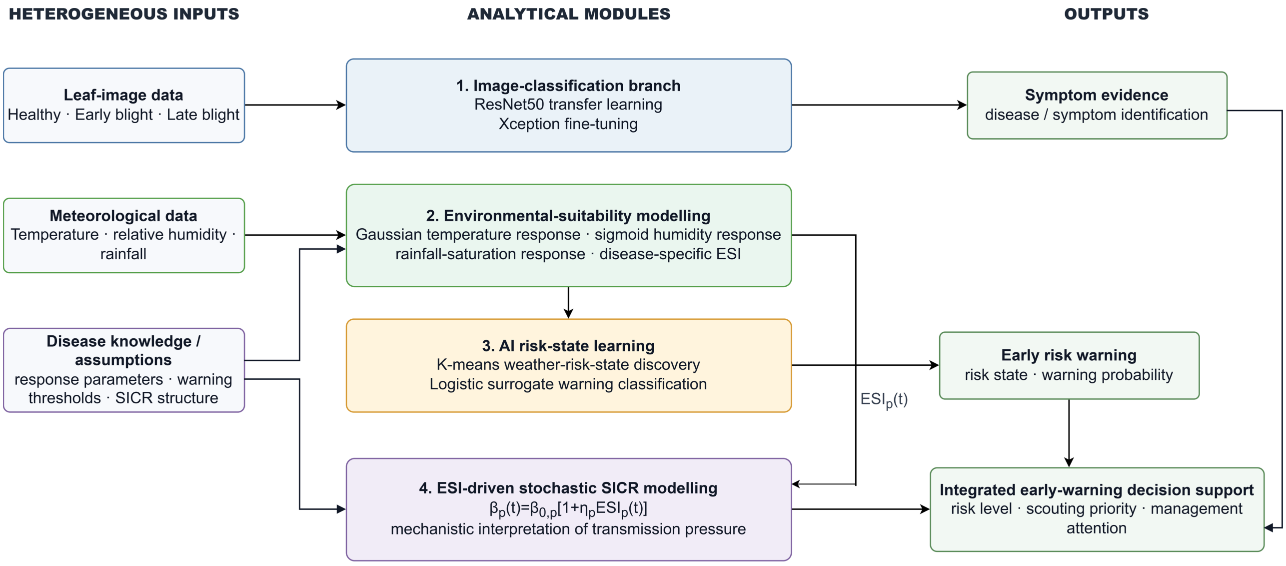

Figure 8
Experimental pipeline for potato leaf disease classification, showing the ResNet50 cross-source branch and the Xception source-split branch as two separate tracks that converge on a protocol-aware interpretation

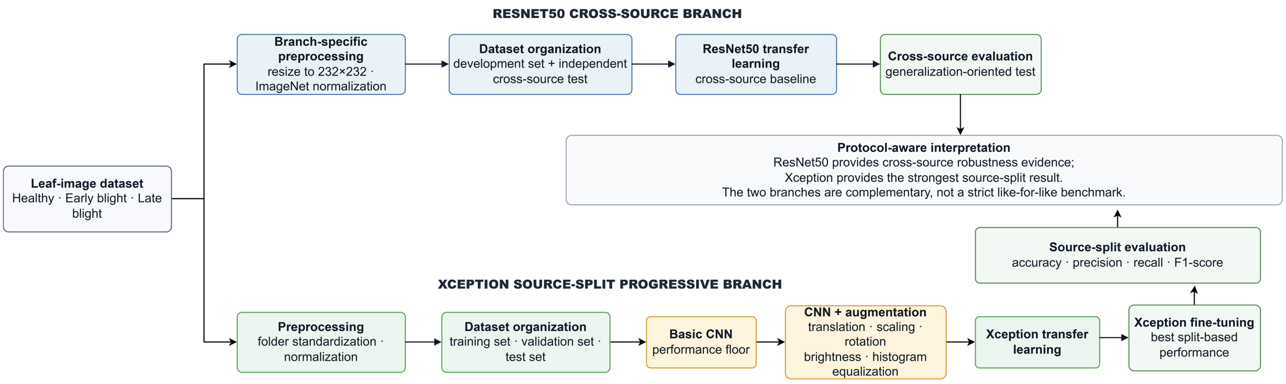

Figure 9
Comparative model performance across progressive learning strategies. The chart is used as an integrative visual summary rather than as a strict architecture benchmark because the underlying protocols differ across image-classification branches

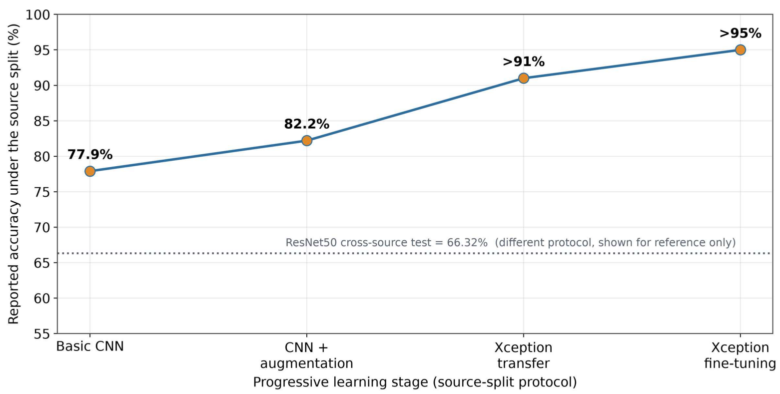

Conceptual synthesis of the source-split accuracy trajectory (basic CNN → augmentation → transfer learning → fine-tuning). It is not a controlled architecture benchmark; the ResNet50 value comes from a stricter, non-comparable protocol.

Figure 10
Weather-driven risk-state learning and warning logic, showing the pathway from meteorological inputs to risk-state outputs and recommended actions

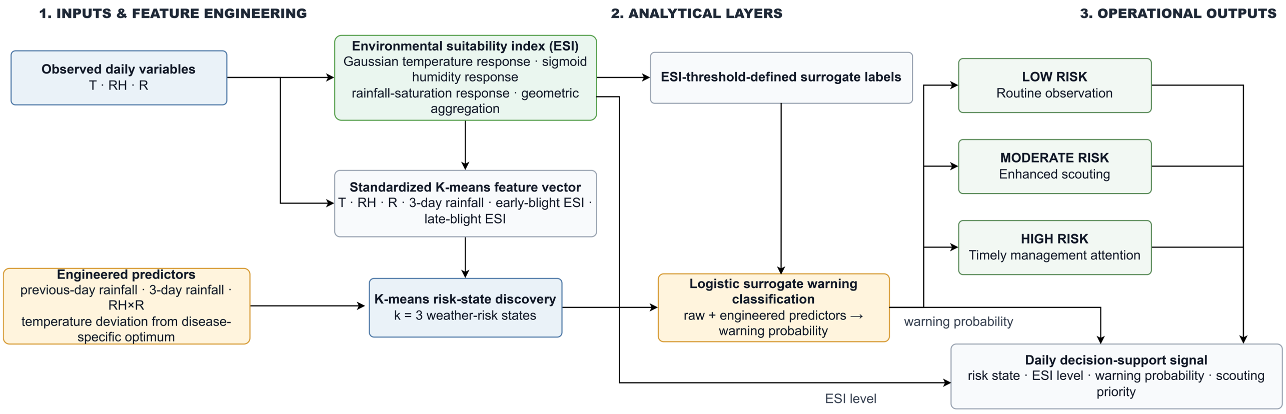

Figure 11
ESI-driven stochastic SICR transmission mechanism. Environmental suitability influences the time-varying transmission coefficient, while stochastic perturbations represent uncertainty

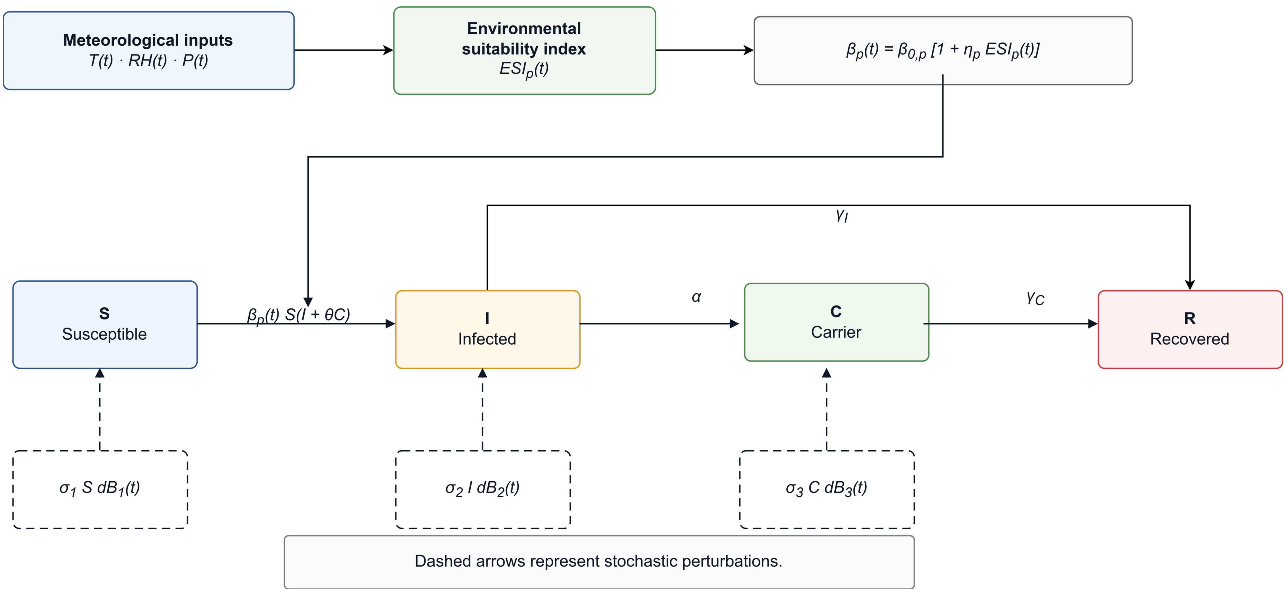

daily spike. The high-risk K-means cluster supports this interpretation by grouping days with persistent wetness and high suitability.

The logistic surrogate classifier is useful but must be interpreted carefully. Because the target labels are generated from ESI thresholds, high AUC values show that the engineered meteorological features encode the warning logic. They do not prove that observed disease incidence has been predicted. This conservative interpretation is preferable for peer review because it prevents overstating the evidence. In a future field-validated version, the same feature set can be evaluated against observed incidence or severity labels.

The stochastic SICR layer is also positioned conservatively. The mathematical structure provides theoretical support for positivity, persistence, and extinction reasoning, but calibrated stochastic epidemic forecasting requires incidence data and confirmed parameter values. Coupling $\beta_p(t)$ to ESI is therefore presented as a methodological bridge rather than as a completed incidence-prediction result. This is more defensible than fabricating calibrated transmission outcomes from insufficient evidence.

The practical implication is that potato disease warning should not rely on a single data modality. Leaf images are useful once symptoms are visible; weather-risk models are useful before symptoms become severe; stochastic transmission models help interpret whether short-term risk may become persistent; and intelligent-system architecture helps organize these components into operational decision support. Figure 12 translates this implication into a deployable architecture in which field sensing, data platforms, analytics engines and user-facing outputs form a modular warning system. Agricultural extension services can use such a framework to prioritize scouting, identify high-risk weather windows and prepare targeted management recommendations according to local regulations and agronomic practice.

Figure 12
Practical deployment architecture for intelligent potato disease early warning, illustrating a modular pathway from sensing and data platforms to analytics and user-facing outputs

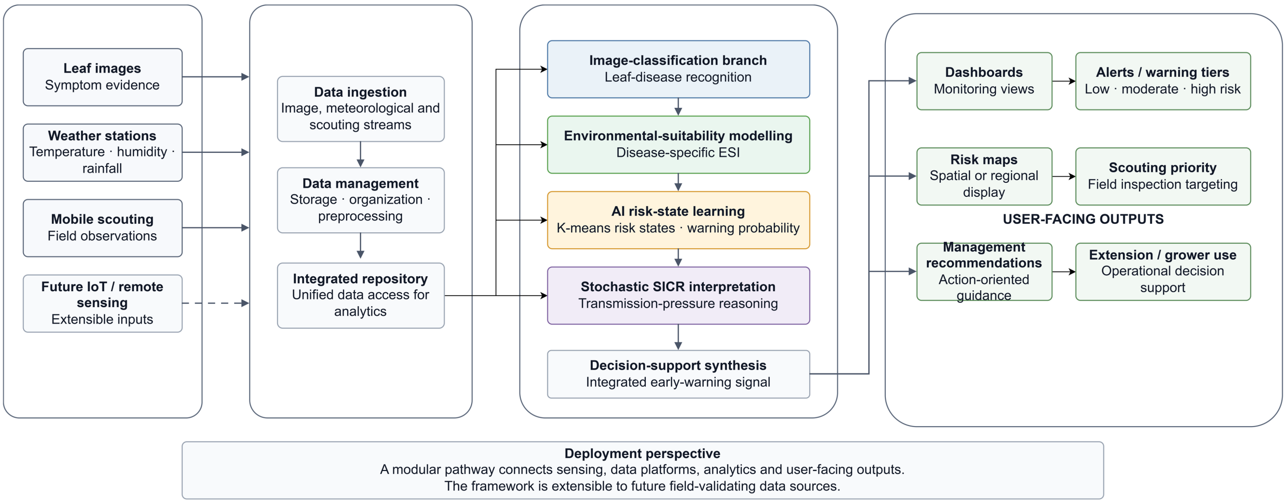

5.1. Visualizing disease-specific environmental responses

Figure 6 complements Section 3.4 by making the environmental suitability layer more transparent. It visualizes the response functions assumed for temperature, RH, and rainfall. The figure clarifies why the same meteorological sequence can produce different early- and late-blight risk profiles: early blight is represented as warmer-temperature responsive, whereas late blight is represented as more strongly dependent on cooler temperatures, high humidity, and sustained moisture. This improves the interpretability of the ESI equations before the empirical risk-burden results are interpreted.

5.2. Integrating heterogeneous evidence streams

Figure 7 presents the expanded conceptual framework for the whole paper. It shows that the study is not merely an image-classification article and not merely a stochastic-modeling article; its contribution is the integration of heterogeneous data sources, intelligent analysis, and mechanistic interpretation. Figure 8 narrows the focus to the leaf-image branch and explains how dataset construction, preprocessing, splitting, modeling, and evaluation are connected. This pairing prevents the classification results from appearing isolated from the wider early-warning system.

5.3. Progressive performance synthesis

Figure 9 summarizes the progressive learning logic—from a basic CNN, to augmentation, to transfer learning, and finally to fine-tuning—without claiming a strict controlled benchmark across all architectures. This is consistent with the caution expressed in Section 4.1: the ResNet and Xception results originate from different validation designs. The figure therefore strengthens readability while preserving methodological honesty.

5.4. Warning logic and mechanistic interpretation

Figure 10 translates the weather-risk results into an operational warning logic. It shows how temperature, RH, rainfall, lagged rainfall, and rolling rainfall can be converted into ESI values, risk-state learning outputs, and tiered actions. Figure 11 links this warning output to transmission dynamics: environmental suitability modulates $\beta_p(t)$, while stochastic perturbations represent field uncertainty. This pairing connects data-driven warning logic with epidemiological reasoning, which strengthens the manuscript's positioning as an intelligent-systems paper rather than a purely descriptive weather analysis.

5.5. Deployment implications and translational value

Figure 12 extends the manuscript from methodological demonstration to future deployment. It shows how leaf images, weather stations, mobile scouting, and future Internet of Things (IoT) or remote-sensing inputs can be organized through a data platform and analytics engine into dashboards, alerts, risk maps, and management recommendations. This deployment view is important for JDSIS positioning because it emphasizes modularity, scalability, and practical decision-support relevance. Figures 6, 7, 8, 9, 10, 11, and 12 function as synthesis figures that improve coherence and readability by connecting equations, pipelines, risk logic, and deployment architecture.

6. Limitations and Future Work

Several limitations must be addressed before field deployment:

Limited humidity-complete window. Although an expanded three-growing-season record of 459 daily temperature and rainfall observations is analyzed, the humidity-complete window that supports the full ESI, clustering, and surrogate-classification analyses is limited to 62 daily observations from two August periods; the warning-state results are therefore demonstrative rather than broadly generalizable, and multi-year, multi-location humidity records are required for regional generalization.

Surrogate rather than field-confirmed labels. The weather-risk classifier uses ESI-derived labels rather than field-confirmed incidence labels and should be described as a surrogate warning classifier, not a disease-incidence prediction model.

No like-for-like image benchmark. The ResNet and Xception image results are based on different datasets and data splits and should not be interpreted as a strict architecture benchmark; a like-for-like comparison under one shared split and a genuine multimodal fusion of leaf images with pre-symptomatic weather risk are priorities for future work.

Uncalibrated stochastic component. The stochastic SICR formulae and threshold conditions require additional parameter confirmation before calibrated incidence simulation.

Scope of the deposited data. The leaf-image dataset, the compiled meteorological records, the derived suitability indices, and the analysis code are openly deposited, and the redistribution conditions of the original sources were verified before deposition; the deposit does not contain field-confirmed disease-incidence observations because none were available for this study.

Schematic versus empirical figures. Figures 6, 7, 8, 9, 10, 11, and 12 are synthesis figures used to strengthen interpretability and presentation rather than independently validated field-deployment outputs; the empirical figures (Figures 2, 3, 4, and 5) are clearly distinguished from these schematic figures in the captions.

Future work should collect multi-year and multi-location field incidence data, disease severity scores, leaf wetness, soil moisture, cultivar resistance, planting density, irrigation, fungicide records, remote-sensing features, and georeferenced image data. Once these labels are available, the surrogate classifier can be replaced by true supervised disease prediction, and models such as random forest, gradient boosting, support vector machines, temporal convolutional networks, LSTM networks, and graph-based spatial models can be compared legitimately. The image branch can also be extended into a multimodal fusion model that combines visible symptoms with pre-symptomatic weather risk.

7. Conclusion

This manuscript integrates four previously separate research components into a coherent data-science and intelligent-systems framework for potato early and late blight. The framework combines leaf-image symptom recognition, meteorological environmental suitability modeling, entropy-weight robustness analysis, K-means weather-risk-state discovery, logistic surrogate warning classification, and stochastic SICR transmission logic. The image results show that fine-tuned Xception can exceed 95% accuracy under the source split, whereas ResNet50 cross-source testing exposes the difficulty of real-world generalization. The weather results show that August 2025 had substantially higher cumulative early and late blight suitability than August 2024. The AI results demonstrate that unsupervised clustering can identify a high-risk weather regime and that and that a logistic surrogate classifier can learn ESI-defined warning labels from meteorological features with high AUC. The synthesis figures strengthen the manuscript by linking equations, pipelines, risk logic, and deployment architecture to the written argument. Overall, the manuscript contributes a modular and reproducible early-warning architecture while maintaining a clear boundary between demonstrated results, model assumptions, and future field-validated prediction.

Funding Support

This work was supported by the Jining Normal University Scientific Research Project: Convolutional Neural Network-based Potato Leaf Disease Identification Method (jsky2024001).

Conflicts of Interest

The authors declare that they have no conflicts of interest to this work.

Data Availability Statement

The data that support the findings of this study are openly available on Zenodo at https://doi.org/10.5281/zenodo.22105735. This deposit comprises the leaf-image dataset that supports the image-classification results (the Potato Disease Leaf Dataset, with early blight, late blight, and healthy-leaf categories), the daily meteorological records for Chahar Right Front Banner (Ulanqab City, Inner Mongolia) used in the environmental suitability analyses, the derived ESIs, and the analysis code. The redistribution conditions of the original meteorological sources were verified prior to deposition.

Author Contribution Statement

Hongzhi Zhang: Conceptualization, Methodology, Investigation, Resources, Data curation, Writing – original draft, Writing – review & editing, Project administration, Funding acquisition. **L. P. Varlamova:** Validation, Writing – review & editing, Supervision. **Jianan Chen:** Software, Formal analysis, Data curation, Visualization.

Use of AI-Assisted Tools

The authors used AI-assisted tools only to a minimal extent, limited to light language editing (grammar and phrasing consistency) and formatting of the submission file. These tools were not used for the study design, data collection or processing, statistical analysis, generation of results, or the interpretation and conclusions, all of which were carried out and verified by the authors. The authors have reviewed the entire manuscript and take full responsibility for its content.

References

[1] Ristaino, J. B., Anderson, P. K., Bebber, D. P., Brauman, K. A., Cunniffe, N. J., Fedoroff, N. V., . . . , & Wei, Q. (2021). The persistent threat of emerging plant disease pandemics to global food security. *Proceedings of the National Academy of Sciences*, *118*(23), e2022239118. https://doi.org/10.1073/pnas.2022239118

[2] Kumar, D., & Mukhopadhyay, R. (2025). Climate change and plant pathogens: Understanding dynamics, risks and mitigation strategies. *Plant Pathology*, *74*(1), 59–68. https://doi.org/10.1111/ppa.14033

[3] Liu, J., & Wang, X. (2021). Plant diseases and pests detection based on deep learning: A review. *Plant Methods*, *17*(1), 22. https://doi.org/10.1186/s13007-021-00722-9

[4] Shafik, W., Tufail, A., Namoun, A., de Silva, L. C., & Apong, R. A. A. H. M. (2023). A systematic literature review on plant disease detection: Motivations, classification techniques, datasets, challenges, and future trends. *IEEE Access*, *11*, 59174–59203. https://doi.org/10.1109/ACCESS.2023.3284760

[5] Shoaib, M., Sadeghi-Niaraki, A., Ali, F., Hussain, I., & Khalid, S. (2025). Leveraging deep learning for plant disease and pest detection: A comprehensive review and future directions. *Frontiers in Plant Science*, *16*, 1538163. https://doi.org/10.3389/fpls.2025.1538163

[6] Alhammad, S. M., Khafaga, D. S., El-Hady, W. M., Samy, F. M., & Hosny, K. M. (2025). Deep learning and explainable AI for classification of potato leaf diseases. *Frontiers in Artificial Intelligence*, *7*, 1449329. https://doi.org/10.3389/frai.2024.1449329
[7] Mulugeta, A. T., Jifara, W., Bogale, E., Desiyo, T., & Mokonnen, A. (2025). Early detection and classification of potato leaf disease using convolutional neural networks. *Applied Computational Intelligence and Soft Computing*, *2025*(1), 7614841. https://doi.org/10.1155/acis/7614841
[8] Dame, T. A., Adera, G. B., & Girmaw, D. W. (2025). Deep learning-based potato leaf disease classification and severity assessment. *Discover Applied Sciences*, *7*(7), 649. https://doi.org/10.1007/s42452-025-07240-6
[9] Erlin, Fuadi, I., Putri, R. N., Nasien, D., Gusrianty, & Oktarina, D. (2024) Deep learning approaches for potato leaf disease detection: Evaluating the efficacy of convolutional neural network architectures. *Revue d'Intelligence Artificielle*, *38*(2), 717–727. https://doi.org/10.18280/ria.380236
[10] He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition*, 770–778. https://doi.org/10.1109/CVPR.2016.90
[11] Chollet, F. (2017). Xception: Deep learning with depthwise separable convolutions. In *2017 IEEE Conference on Computer Vision and Pattern Recognition*, 1800–1807. https://doi.org/10.1109/CVPR.2017.195
[12] Dolatabadian, A., Neik, T. X., Danilevicz, M. F., Upadhyaya, S. R., Batley, J., & Edwards, D. (2025). Image-based crop disease detection using machine learning. *Plant Pathology*, *74*(1), 18–38. https://doi.org/10.1111/ppa.14006
[13] Wang, X., & Liu, J. (2024). An efficient deep learning model for tomato disease detection. *Plant Methods*, *20*(1), 61. https://doi.org/10.1186/s13007-024-01188-1
[14] Albahli, S. (2025). AgriFusionNet: A lightweight deep learning model for multisource plant disease diagnosis. *Agriculture*, *15*(14), 1523. https://doi.org/10.3390/agriculture15141523
[15] Xu, M., Yoon, S., Jeong, Y., & Park, D. S. (2022). Transfer learning for versatile plant disease recognition with limited data. *Frontiers in Plant Science*, *13*, 1010981. https://doi.org/10.3389/fpls.2022.1010981
[16] Bagchi, P., Sawicka, B., Stamenkovic, Z., Markovic, D., & Bhattacharjee, D. (2024). Potato late blight outbreak: A study on advanced classification models based on meteorological data. *Sensors*, *24*(23), 7864. https://doi.org/10.3390/s24237864
[17] Meno, L., Escuredo, O., Abuley, I. K., & Seijo, M. C. (2023). Predicting daily aerobiological risk level of potato late blight using C5.0 and random forest algorithms under field conditions. *Sensors*, *23*(8), 3818. https://doi.org/10.3390/s23083818
[18] Vaidheki, M., Gupta, D. S., Basak, P., Debnath, M. K., Hembram, S., & Ajith, S. (2023). Prediction of potato late blight disease incidence based on weather variables using statistical and machine learning models: A case study from West Bengal. *Journal of Agrometeorology*, *25*(4), 583–588. https://doi.org/10.54386/jam.v25i4.2272
[19] Joshi, D., Krishnan, P., Vashisth, A., Kundu, M., Rani, A., & Bag, T. K. (2025). AI-based machine learning and multiple linear regression approach to simulate the effect of weather on the crop age at first appearance of potato late blight (*Phytophthora infestans* (Mont.) de Bary) disease. *Potato Research*, *68*(2), 1825–1848. https://doi.org/10.1007/s11540-024-09795-0
[20] Yuan, P., Xia, Y., Dong, M., He, C., Liu, D., Tan, Y., & Dong, S. (2026). PLB-GPT: Potato late blight prediction with generative pretrained transformer and optimizing. *Mathematics*, *14*(2), 225. https://doi.org/10.3390/math14020225
[21] Lee, S., & Yun, C. M. (2023). A deep learning model for predicting risks of crop pests and diseases from sequential environmental data. *Plant Methods*, *19*(1), 145. https://doi.org/10.1186/s13007-023-01122-x
[22] MacQueen, J. B. (1967). Some methods for classification and analysis of multivariate observations. In *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*, *1*, 281–297.
[23] Gallardo-Romero, D. J., Apolo-Apolo, O. E., Martinez-Guanter, J., & Perez-Ruiz, M. (2023). Multilayer data and artificial intelligence for the delineation of homogeneous management zones in maize cultivation. *Remote Sensing*, *15*(12), 3131. https://doi.org/10.3390/rs15123131
[24] Cox, D. R. (1958). The regression analysis of binary sequences. *Journal of the Royal Statistical Society: Series B (Methodological)*, *20*(2), 215–232. https://doi.org/10.1111/j.2517-6161.1958.tb00292.x
[25] Mao, X. (2007). *Stochastic differential equations and applications* (2nd ed.). UK: Horwood Publishing.
[26] Chen, T., & Li, Z. (2022). Dynamics and stationary distribution of a stochastic SIRS epidemic model with a general incidence and immunity. *Boundary Value Problems*, *2022*(1), 85. https://doi.org/10.1186/s13661-022-01668-0
[27] Li, R., & Guo, X. (2024). Dynamics of a stochastic SEIR epidemic model with vertical transmission and standard incidence. *Mathematics*, *12*(3), 359. https://doi.org/10.3390/math12030359

How to Cite: Zhang, H., Varlamova, L. P., & Chen, J. (2026). Integrating Leaf-Image Deep Learning, Meteorological Risk-State Discovery, and Stochastic SICR Dynamics for Potato Early and Late Blight: A Data-Science and Intelligent-Systems Framework. *Journal of Data Science and Intelligent Systems*. https://doi.org/10.47852/bonviewJDSIS620210658