Cadernos de Pós-Graduação em Direito Político e Econômico
Published by: Centro de Estudos Acadêmicos Press
ISSN: 1678-2127
Volume 26 ***Issue 2,*** **2026**
Website: https://ceapress.org

From Rule-Based Systems to Retrieval-Augmented LLMs: A Comparative Review of AI Approaches for Maize Pest and Disease Management in Punjab

Rishabh Aryan [1*]
M.Tech (Artificial Intelligence and Data Science)
Department of Computer Science and Engineering,
Indian Institute of Information Technology, Bhagalpur, Bihar, India
Email: rishabh.250201011@iiitbh.ac.in
ORCID ID: https://orcid.org/0009-0004-7595-9440
Anju [2]
M.Sc. (Mathematics with Computer Science)
Department of Mathematics,
Maharishi Dayanand University, Rohtak, Haryana, India
Email: anjukaliraman1@gmail.com
ORCID ID: https://orcid.org/0009-0000-2248-9677
Satpal [3]
M.Sc. (Chemistry)
Department of Science,
Rayat Bahra University, Mohali, Punjab, India
Email: rajchouhansatpal@gmail.com
ORCID ID: https://orcid.org/0009-0001-1735-1873
Nishtha Kalia [4]
BCA
Department of Computer Science and Applications,
Mehr Chand Mahajan DAV College for Women,
Panjab University, Chandigarh, India
Email: nishtha.kalia.2411@gmail.com
ORCID ID: https://orcid.org/0009-0006-7522-1221
Dr. Sarwan Singh [5]
Scientist D and Joint Director
National Institute of Electronics and Information Technology, Ropar, Punjab, India
Email: sarwan@nielit.gov.in
ORCID ID: https://orcid.org/0000-0001-7062-2129
[1]Rishabh Aryan, [2]Anju, [3]Satpal, [4]Nishtha Kalia, and [5]Dr. Sarwan Singh contributed equally to this work and share first authorship.
* **Corresponding author:** Rishabh Aryan, M.Tech (Artificial Intelligence and Data Science), Department of Computer Science and Engineering, Indian Institute of Information Technology,

Bhagalpur, Bihar, India. Email: ID rishabh.250201011@iiitbh.ac.in. ORCID ID: https://orcid.org/0009-0004-7595-9440.

Submission:02.04.2026 Acceptance: 02.06.2026 Publication:07.09.2026

Abstract*:* Spring-sown maize in Punjab faces three recurring problems every season: shoot fly (Atherigona spp.) at the seedling stage, Maydis leaf blight (Bipolaris maydis) through the vegetative phase, and banded leaf and sheath blight (Rhizoctonia solani f. sp. sasakii) once the canopy closes. The Punjab Agricultural University (PAU) and the ICAR-Indian Institute of Maize Research (ICAR-IIMR), both based in Ludhiana, publish advisories for all three, and farmers in principle have access to this guidance every season. What has changed over the last four decades is not the advisories themselves but the computational tools built around them, moving from hand-coded rule-based expert systems in the 1980s and 1990s to classical machine learning on hand-crafted image features in the 2000s to convolutional neural networks and vision transformers from roughly 2015 onwards and now to retrieval-augmented large language models. Each of these four approaches has its own literature, and they rarely talk to each other: the deep-learning papers cite almost none of the expert-systems work, and the handful of agricultural RAG papers that exist barely reference either. The deep-learning literature also has a specific blind spot worth naming up front, which is that it leans heavily on the PlantVillage benchmark, a dataset of laboratory-photographed leaves that generalizes poorly once a model is pointed at an actual field and one that under-represents spring maize and the Punjab-specific disease complex in any case. This paper traces all four eras for one crop and one region rather than treating them as separate literatures, and it tries to hold itself to a stricter standard than is typical of a survey: every technical or agronomic claim is checked against at least two independent sources before it is stated. Where the practical, farmer-facing metrics that would actually distinguish these four approaches, hallucination rate, response latency on a cheap phone, and cost of running a village deployment simply have not been measured for Punjab maize by anyone, we say so and propose a protocol for measuring them rather than inventing numbers. The paper closes by naming, concretely, the gap in Punjab-specific field imagery and loss data that currently limits how far deep learning and RAG can be trusted here.
Keywords: Maize, Punjab, shoot fly, Maydis leaf blight, banded leaf and sheath blight, expert systems, machine learning, deep learning, convolutional neural networks, retrieval-augmented generation, large language models, precision agriculture.

I. INTRODUCTION

Maize is Punjab's third cereal after rice and wheat. The state's 2026 Package of Practices puts the crop's area at 86.1 thousand hectares and its production at 370.5 thousand tonnes for 2024 to 2025, an average yield of 17.41 quintals per acre [1]. A large share of that area is sown in spring, a window growers favour for its yield potential even though it comes at a cost: the crop establishes during the warm, dry weeks before the monsoon, which is exactly when shoot fly pressure is highest [2], [3]. Two more problems build through the Kharif season on top of that. Maydis leaf blight (MLB), caused by Bipolaris maydis, has been measured at 15 to 20 percent yield loss in an ordinary season and as high as 40 to 70 percent when conditions turn warm and humid, and a figure near 40 percent recurs often enough across the literature that it functions almost as a rule of thumb [4], [5]. Banded leaf and sheath blight (BLSB), caused by Rhizoctonia solani f. sp. sasakii, is soil-borne and spreads up the stalk once the canopy closes; North Indian trials put typical losses at 11 to 40 percent, and more than one paper reports losses approaching total crop failure under sufficiently humid, densely planted conditions [6], [7], [8].

Four distinct computational approaches have been proposed for diagnosing and managing problems like these, and they come out of four largely separate research communities. Rule-based expert systems encode an agronomist's if-then reasoning directly and predate the others by a wide margin. Classical machine learning came next, training statistical classifiers such as support vector machines on hand-engineered colour and texture descriptors pulled from leaf photographs. Deep learning replaced those hand-engineered features with convolutional neural networks and, more recently, vision transformers that learn their own filters end to end from labelled images. Most recently, retrieval-augmented large language models pair a conversational LLM with a retrieval layer over curated documents, aiming to answer a farmer's question in plain language while reducing the risk of the model simply making something up. Each paradigm has its own review literature [9] to [14], and CNN-based maize disease detection in particular is mature enough that Punjab-based groups, including one at Chitkara University, are already publishing in it [15]. Agricultural RAG systems are newer but growing quickly [16] to [20]. What none of this literature does is walk the full path from rules to RAG for a single crop in a single region, grounded in the advisories a Punjab extension officer would actually hand a farmer, and honest about where the evidence for Punjab spring maize specifically runs thin rather than borrowing numbers from other crops or countries to paper over the gap.

The remainder of the paper is organised as follows. Section II sets out the Punjab spring maize production context and the three target constraints, drawing on PAU and ICAR-IIMR figures together with seed-treatment and fungicide trials conducted at Ludhiana and comparable North Indian sites. Sections III through VI walk through the four computational eras in turn, each with a block diagram, a worked pseudocode algorithm, and the relevant equations. Section VII proposes an evaluation protocol for the practical metrics, hallucination rate, latency, and cost that no existing study has measured for this crop and region, rather than inventing figures for them. Section VIII works through a single farmer query end to end to show what grounding does and does not fix. Section IX lays out the Punjab-specific data gap in plain terms, Section X pulls the four eras

Cadernos de Pós-Graduação em Direito Político e Econômico
Published by: Centro de Estudos Acadêmicos Press
ISSN: 1678-2127
Volume 26 ***Issue 2,*** **2026**
Website: https://ceapress.org

together into a single comparison table, and Section XI closes with what we think the next steps actually are.

II. PUNJAB SPRING MAIZE: PRODUCTION CONTEXT AND TARGET CONSTRAINTS

Maize in Punjab is sown across a fairly wide window, from the last week of May to the end of June for the irrigated kharif crop, and fertilisation follows a soil-test basis, with a default recommendation of 50 kg N, 24 kg P_2O_5 and 12 kg K_2O per acre for the long-duration hybrids PAU currently recommends, split across sowing, the knee-high stage, and pre-tasseling [1]. Three biotic constraints keep showing up, both in PAU and ICAR-IIMR advisories and in the peer-reviewed literature, and they form the running examples for the rest of this paper. Figure 1 turns the yield-loss figures from Table I into a chart, mainly so the reader can see at a glance not just which constraint is worse but how wide the reported range is for each one.

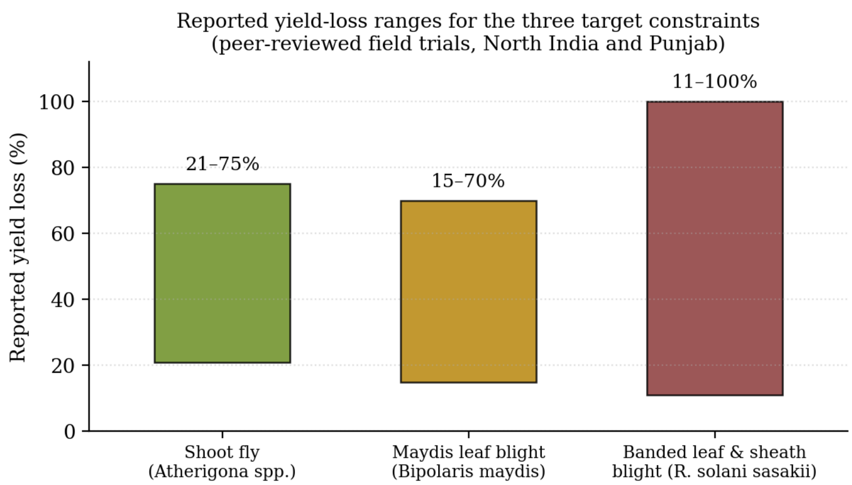

Figure 1. Reported yield-loss ranges for the three target constraints, drawn from the peer-reviewed field trials cited in Table I. Bars show the low-to-high range reported across studies, not a single controlled experiment.

A. Shoot Fly (Atherigona spp.)

Shoot fly hits spring-sown maize harder than kharif maize for a fairly simple reason: the crop is establishing exactly when fly activity peaks and natural predators have not yet built up, and reported losses reflect that, running up to 60% plant loss and 21% to 75% grain yield loss depending on the study [3], [23]. A multi-year trial at Ludhiana tested pre-sowing seed treatment with Gaucho (imidacloprid) 600 FS at 6 mL per kg seed against Atherigona naqvii in spring maize and found the treatment protected the crop through early establishment without hurting germination, even when seed was treated several days ahead of sowing [2]. That 6 mL per kg figure is not an outlier: a separate, independently written advisory for kharif maize across North India recommends the same rate for the same purpose and also lists phorate 10G at 5 kg per acre

Cadernos de Pós-Graduação em Direito Político e Econômico
Published by: Centro de Estudos Acadêmicos Press
ISSN: 1678-2127
Volume 26 ***Issue 2,*** **2026**
Website: https://ceapress.org

applied to the soil at sowing as an alternative where seed treatment is not practical [24]. Trials on the broader shoot-fly and stem-borer complex in cereals more generally point to thiamethoxam 30 FS at around 7 to 10 mL per kg of seed and fipronil at similar rates as comparably effective systemic options, and both tend to beat granular soil application once labour and cost are factored in [21], [22], [26]. PAU flags a related early-season pest in the same window, the spotted stem borer (Chilo partellus), strongly enough that the 2026 Package of Practices notes that timely chemical control against maize borer is especially important for early-sown crops [1]. The Package of Practices gives a directly confirmed rate for this pest in its fodder-maize section specifically: 40 mL of Coragen 18.5 SC (chlorantraniliprole) in 60 to 80 litres of water per acre, applied after uprooting and destroying attacked plants 2 to 3 weeks after sowing, with tricho-cards (50,000 Corcyra cephalonica eggs parasitised by Trichogramma chilonis, released twice) offered as a non-chemical alternative, alongside a strict 21-day withdrawal period before the fodder can be fed to animals [1]. A separate PAU-issued grower advisory gives a slightly different rate for the same pest, 30 mL in 60 litres of water per acre, without specifying whether it targets fodder or grain maize [52]; that gap between two PAU-linked sources is itself worth flagging rather than smoothing over, since it is exactly the kind of crop-context distinction, fodder versus grain maize, that a properly grounded system needs to preserve rather than average away.

B. Maydis Leaf Blight (Bipolaris maydis)

MLB shows up first as long, tan, cigar-shaped lesions along the leaf blade, and given enough warmth and humidity, it can strip the canopy before grain fill even finishes. A two-season trial run across Ludhiana, Karnal, and Delhi compared organic, chemical, and integrated disease-management modules, and the chemical module won at every site, cutting disease by 54.16% at Ludhiana and 52.92% at Karnal; at Delhi, a straightforward mancozeb 75WP spray at 2.5 g per litre of water outperformed everything else tested, reaching 64.29% control [5], [4]. That mancozeb figure holds up against an unrelated crop-production guide, which recommends mancozeb or zincate at 2 to 4 g per litre as soon as the disease appears, repeated every ten days if needed, sometimes paired with a propiconazole 25% EC spray at 0.1% around 35 and 50 days after sowing [27]. Propiconazole at that same 0.1% concentration also turns up, independently, as the single best-performing in vivo fungicide in two further efficacy trials that had nothing to do with the three-location study [28], [29], so the propiconazole recommendation in particular rests on three unrelated pieces of evidence rather than one. PAU's own hybrid recommendations track this pressure closely, and here the current Package of Practices can be quoted directly rather than through a secondary source: the 2026 Kharif edition describes NK 7328 as moderately resistant to Maydis leaf blight, PMH 17 as moderately resistant to both fall armyworm and Maydis leaf blight, and DKC 9144 as moderately resistant to Maydis leaf blight, charcoal rot, and fall armyworm together; two further hybrids, PMH 13 and ADV 9293, are listed as moderately resistant to Maydis leaf blight, charcoal rot, and maize stem borer as well [1].

Cadernos de Pós-Graduação em Direito Político e Econômico
Published by: Centro de Estudos Acadêmicos Press
ISSN: 1678-2127
Volume 26 ***Issue 2,*** **2026**
Website: https://ceapress.org

C. Banded Leaf and Sheath Blight (Rhizoctonia solani f. sp. sasakii)

BLSB starts at the base of the plant and works its way up the sheath once conditions turn warm, humid, and crowded, and North Indian trials generally put the resulting yield loss at 11% to 40%, though at least one paper reports losses approaching total loss when humidity is high enough for long enough [6], [7], [8]. A field trial in Karnataka found that treating seed with Pseudomonas fluorescens at 10 g per kg, combined with two propiconazole sprays at 0.1% timed to 30 and 40 days after sowing, produced both the highest grain yield and the lowest disease intensity of everything tested, a result a separate, unrelated efficacy screening more or less confirms by independently flagging propiconazole and related triazoles, alongside Trichoderma-based bioagents, as the strongest components of an integrated BLSB programme [8], [5]. A multi-year trial in Haryana adds useful timing detail: symptoms tend to first appear around the 31st to 32nd standard meteorological week, roughly early to mid-August, with severity peaking some six weeks after that, and spraying on a weather-informed schedule beat spraying on a fixed calendar [30]. A separate PAU-issued advisory for maize growers gives a third, independent figure for the same disease complex under the common name "leaf and sheath blight", recommending 100 mL of Amistar Top (azoxystrobin plus difenoconazole) in 200 litres of water per acre at disease appearance, repeated after 15 days if needed, alongside rather than in place of the propiconazole-based programme above [52].

D. Fall Armyworm as an Emerging Fourth Constraint

Fall armyworm (Spodoptera frugiperda) sits outside the three constraints this paper focuses on, but it is now established across most of India's major maize-growing states, Punjab included, and PAU has already started releasing hybrids with moderate tolerance to it, PMH 17 and DKC 9144 among them [1], [31], [32]. A PAU-issued grower advisory gives a specific spray protocol for it: Coragen 18.5 SC at 0.4 mL per litre, or Delegate 11.7 SC at 0.5 mL per litre, or Missile 5 SG at 0.4 g per litre, applied in 120 litres of water per acre on crops up to 20 days old with the nozzle aimed directly into the whorl [52]. The Coragen figure in that advisory is independently confirmed, mL for mL, by the current 2026 Package of Practices, which additionally specifies that on older crops, 20 to 40 days, the water volume should rise to 200 litres per acre with a correspondingly higher insecticide dose, and repeats the same 21-day fodder-withdrawal caution noted above for maize borer [1]. It is worth naming fall armyworm here for a practical reason: any retrieval-based knowledge base built for maize pest management in Punjab is going to need to tell fall armyworm damage apart from shoot-fly damage, because both attack the whorl of a young plant and look fairly similar to an untrained eye, but they call for different chemistries and different timing.

Table I. Reported yield-loss ranges and cross-checked management inputs for the three target constraints.

Cadernos de Pós-Graduação em Direito Político e Econômico
Published by: Centro de Estudos Acadêmicos Press
ISSN: 1678-2127
Volume 26 ***Issue 2,*** **2026**
Website: https://ceapress.org

Constraint	Causal agent	Reported yield-loss range	Cross-checked management input	Sources
Shoot fly	Atherigona spp.	Up to ~60% plant loss; 21 to 75% grain yield loss	Imidacloprid 600 FS, 6 mL/kg seed (confirmed by 2 independent sources)	[2], [3], [23], [24]
Maydis leaf blight	Bipolaris maydis	15 to 20% typical; up to 40 to 70% severe epidemic	Mancozeb 75WP, 2.5 to 4 g/L water, or propiconazole 0.1% (confirmed by 3 independent sources)	[4], [5], [27], [28], [29]
Banded leaf & sheath blight	Rhizoctonia solani f. sp. sasakii	11 to 40% typical; up to 100% extreme, humid conditions	Pseudomonas fluorescens seed treatment of 10 g/kg + propiconazole 0.1% spray at 30/40 DAS, or Amistar Top 100 mL/acre per PAU advisory (confirmed by 3 independent sources)	[6], [7], [8], [5], [52]

The Insect Pests and Diseases subsections of the grain-maize chapter, which would carry an official current-season shoot-fly, MLB, and BLSB dosage table, could not be retrieved from the 2026 Package of Practices despite repeated attempts; the agronomy, fertiliser, hybrid-resistance, weed-control, and fodder-maize plant protection sections were retrieved directly and are cited above. In their place, the figures reported here draw on peer-reviewed field trials conducted at PAU Ludhiana and ICAR-IIMR and on a PAU-issued grower advisory verified by direct retrieval; readers should confirm current product names and doses against the printed Package of Practices before use, since recommendations are revised each season and, as Section VIII notes, fodder-maize and grain-maize rates for the same pest are not always identical [1], [52].

III. ERA 1: RULE-BASED EXPERT SYSTEMS

The oldest computational tools for crop protection simply write down what a pathologist or entomologist already knows as production rules. A knowledge engineer sits with domain experts, turns their heuristics into an if-symptom-then-diagnosis-then-recommendation chain, and

Cadernos de Pós-Graduação em Direito Político e Econômico
Published by: Centro de Estudos Acadêmicos Press
ISSN: 1678-2127
Volume 26 ***Issue 2,*** **2026**
Website: https://ceapress.org

implements that chain in a rule engine such as CLIPS or EXSYS [9], [10], [33]. Reviewing this literature, most systems turn out to follow the same basic shape: a short symptom-elicitation interface feeding a backward- or forward-chaining inference engine, sometimes with a certainty-factor calculus bolted on to handle ambiguous symptoms rather than forcing a binary yes or no [34], [35]. More recent work has shifted toward ontology-based representations, the CropPestO ontology being one example, largely so the same rule base can be queried and extended across several crops without rewriting it in a new proprietary shell every time [36].

For the Punjab maize triad specifically, a rule-based system has an obvious appeal: PAU and ICAR-IIMR advisories are already written as structured if-then guidance, where crossing a given economic threshold level of dead hearts or leaf damage triggers a specific, named spray recommendation [1]. A minimal rule base can be built almost directly off the Package of Practices text without much reinterpretation, and that is both the paradigm's strongest selling point and its main limitation. The strength is that it cannot hallucinate, and every recommendation traces straight back to an authoritative source. The limitation is that it only ever knows the diagnoses its authors thought to include, cannot make sense of a farmer's own words without a separate natural-language layer bolted on top, and has to be hand-edited every time PAU revises a dosage or releases a hybrid with a new resistance rating.

Figure 2. Block diagram of a rule-based expert-system pipeline for maize pest and disease diagnosis.

Symptom Elicitation (structured Q&A / checklist)

↓

Working Memory (observed facts: crop stage, symptom set)

↓

Inference Engine (forward/backward chaining over PAU/ICAR rule base)

↓

Matched Diagnosis + ETL Check

↓

Recommendation (named product, dose, timing, from POP)

Cadernos de Pós-Graduação em Direito Político e Econômico
Published by: Centro de Estudos Acadêmicos Press
ISSN: 1678-2127
Volume 26 ***Issue 2,*** **2026**
Website: https://ceapress.org

When two independent symptom observations both point to the same diagnosis, most CLIPS-style systems combine their certainty factors rather than just taking the larger one so that corroborating evidence pushes confidence up without letting it exceed full certainty. Equation 1 gives that standard combination rule for two positive certainty factors, CF1 and CF2, supporting the same hypothesis [34].

$$CF_{combined} = CF_1 + CF_2\,(1 - CF_1) \quad (1)$$

Algorithm 1 sketches the core inference loop for a shoot-fly, MLB, and BLSB triage rule base of the kind implementable directly from Section II.

Algorithm 1: Rule-Based Diagnosis and Recommendation
Input: observed_symptoms (set), crop_stage, dead_heart_pct, leaf_damage_pct
Output: diagnosis, recommendation

```
facts <- observed_symptoms U {crop_stage, dead_heart_pct, leaf_damage_pct}
agenda <- match(facts, RULE_BASE)      // pattern match against PAU/ICAR rules
if agenda is empty then
    return ("no confident match", "escalate to KVK / helpline")
rule <- highest_priority(agenda)       // e.g. ETL-triggered rules first
if rule.name == "shoot_fly" and dead_heart_pct >= rule.ETL then
    return ("shoot fly", rule.seed_treatment_or_spray)
else if rule.name == "maydis_leaf_blight" and lesion_pattern matches rule.signature then
    return ("Maydis leaf blight", rule.fungicide_schedule)
else if rule.name == "BLSB" and sheath banding matches rule.signature then
    return ("banded leaf and sheath blight", rule.fungicide_or_IDM_schedule)
return ("ambiguous", "request additional symptom or image evidence")
```

IV. ERA 2: CLASSICAL MACHINE LEARNING ON HAND-CRAFTED FEATURES

The second era replaces handwritten rules with a statistical classifier that is trained using numbers extracted from a leaf photograph, not hand-typed symptoms. Typically, these numbers are colour histograms or colour space statistics (either in HSV or Lab space); texture descriptors like the grey level co-occurrence matrix and local binary patterns; and shape descriptors like histograms of oriented gradients [13, 37], all fed into a classifier like a support vector machine, a random forest, k-nearest neighbours, or a small artificial neural network. These classical models get respectable but consistent results when compared to a CNN trained on the same problem, in the case of maize. A comparison of SVM, decision tree, KNN, and random forest against CNN was performed on a maize disease dataset with four classes (blight, common rust, grey leaf spot, and healthy) using the same features (combined GLCM, LBP, HSV, and Lab), and the CNN achieved the highest

Cadernos de Pós-Graduação em Direito Político e Econômico
Published by: Centro de Estudos Acadêmicos Press
ISSN: 1678-2127
Volume 26 ***Issue 2,*** **2026**
Website: https://ceapress.org

accuracy of 93.93%, while the classical models performed less well on the same features [38]; a previous study using Haralick texture features for plant-leaf disease recognition found approximately the same ordering of the models, with the CNN being the most successful [39]. Some do not compare the CNN features at all but instead use a light second-stage classical classifier on top of the CNN features extracted, such as DenseNet, EfficientNet, or ResNet with PCA and SVM, Random Forest, or ANN for fine-grained grading of severity of maize pest damages [40].

It is difficult to see many examples of this time period in particular for Punjab. Our search for published classical ML studies on spring maize collected in Punjab for shoot fly, MLB, or BLSB yielded no results; the only maize-specific classical ML studies are found in the literature from the other Indian states and the general plant-pathology benchmarks [13] to [40]. The same data gap that was mentioned in great detail in Section IX is responsible for that absence.

Figure 3. Classical machine-learning pipeline for leaf-image-based diagnosis.

Leaf Image Acquisition (smartphone / field camera)

↓

Pre-processing (background removal, colour normalisation)

↓

Hand-Crafted Feature Extraction (GLCM, LBP, HSV/L*a*b*, HOG)

↓

Classifier (SVM / Random Forest / KNN / Decision Tree)

↓

Predicted Disease Class + Confidence Score

This literature and the many papers on deep learning that follow it use the same few confusion-matrix metrics reported below in Equations 2-5, where TP, TN, FP, and FN refer to the counts of true positive, true negative, false positive, and false negative, respectively.

$$Accuracy = (TP + TN) / (TP + TN + FP + FN) \quad (2)$$

$$Precision = TP / (TP + FP) \quad (3)$$

Cadernos de Pós-Graduação em Direito Político e Econômico
Published by: Centro de Estudos Acadêmicos Press
ISSN: 1678-2127
Volume 26 ***Issue 2,*** **2026**
Website: https://ceapress.org

$Recall = TP / (TP + FN)$ (4)

$F1 = 2 \times (Precision \times Recall) / (Precision + Recall)$ (5)

In addition to field data on disease incidence (Equation 7), plant pathologists have also been monitoring disease severity in the field using the Percent Disease Index (Equation 6). Both are common field survey tests, and both have been reported multiple times in the MLB and BLSB trials referred to in Section II [41].

$PDI = [\ \Sigma(\text{numerical rating}) / (n_{plants} \times grade_{max})\ ] \times 100$ (6)

$\text{Disease incidence} = (n_{diseased} / n_{total}) \times 100$ (7)

Algorithm 2: Classical ML Classification Pipeline
Input: leaf_image I
Output: predicted_class, confidence

```
I_norm <- normalize_colour(segment_leaf(I))
f_colour <- extract_HSV_Lab_histograms(I_norm)
f_texture <- extract_GLCM_LBP(I_norm)
f_shape   <- extract_HOG(I_norm)
x <- concat(f_colour, f_texture, f_shape)
(y_hat, p) <- trained_classifier.predict_proba(x)   // e.g. SVM, RF, KNN
if p < CONFIDENCE_THRESHOLD then
    return ("uncertain", p) // fall back to rule-based ETL check or human review
return (y_hat, p)
```

V. ERA 3: DEEP LEARNING, CONVOLUTIONAL NETWORKS, AND VISION TRANSFORMERS

Deep learning eliminates the need for any handcrafted features and allows a convolutional neural network to extract its own hierarchy of visual filters directly from labelled images. The maize-specific literature here is extensive, with attention-augmented residual networks that achieve almost 98 percent accuracy on benchmark maize-leaf datasets [42]; transfer-learning comparisons of MobileNetV2, InceptionResNetV2, and custom-designed architectures on field imagery from four African countries [12]; hybrid CNN-Vision-Transformers that combine local attention (convolution) with global attention (self-attention) [15, 43]; and lightweight CNNs designed for offline operation on low-cost hardware [44, 45]. The only contribution from within Punjab was from Chitkara University, where they have released a hybrid CNN-ViT model for maize leaf disease classification and performed quite well in the PlantVillage benchmark [15]. That last is revealing, and it is not really that it does not concern Chitkara, for the researchers affiliated with Punjab are quite busy in this area, but, essentially without exception, the images they are training on are not even photographed in Punjab.

Cadernos de Pós-Graduação em Direito Político e Econômico
Published by: Centro de Estudos Acadêmicos Press
ISSN: 1678-2127
Volume 26 ***Issue 2,*** **2026**
Website: https://ceapress.org

The contradiction between the place of authorship and the images is the main issue that we have with this period, at least in our scope of interest. Still the benchmark in maize-CNN papers, PlantVillage is a massive dataset of 54,306 images across 14 crops and 38 disease classes, all of which are shot in a lab with a plain, uniform background [46, 47]. There are also several independent re-analyses that have since demonstrated that models trained on this dataset tend to focus on background and lighting rather than disease and that accuracy can drop precipitously when the same model is used to perform a field test on a photograph with natural clutter, occlusion, and uneven light; at least three of these re-evaluations report this phenomenon [47] to [49]. None of that is peculiar to Punjab; it shows up in every crop and every country, but it is worse here than elsewhere because, as far as the published literature goes, there is no set of public images of shoot-fly whorl damage or Maydis leaf blight lesions or BLSB banding that has been actually photographed in Punjab, separately, rather than as a whole. This is repeated in section IX.

Figure 4. Deep-learning (CNN / CNN-ViT) pipeline for maize leaf disease classification.

Field/Lab Image Dataset (e.g., PlantVillage, largely non-Punjab, lab-captured)

↓

Data Augmentation (rotation, flip, colour jitter, resize)

↓

CNN Backbone (local feature extraction, convolution + pooling)

↓

[optional] ViT Branch (global self-attention over patches)

↓

Fusion + Fully-Connected Classifier

↓

Predicted Disease Class + Severity + Confidence

Cadernos de Pós-Graduação em Direito Político e Econômico
Published by: Centro de Estudos Acadêmicos Press
ISSN: 1678-2127
Volume 26 ***Issue 2,*** **2026**
Website: https://ceapress.org

The core operation of the CNN backbone is the convolution operation shown in Equation 8 at one output feature map position (i, j) with an input image patch I, a learned kernel of m × n elements K, and a bias term b.

$$F(i, j) = \Sigma_m \Sigma_n I(i+m, j+n) \cdot K(m, n) + b \quad (8)$$

The categorical cross-entropy loss, Equation 9, is minimised during training, which is the difference between the true one-hot label, yc, and the predicted probability for class c, y-hat c.

$$L = -\Sigma_{c=1}^{C} y_c \log(\hat{y}c) \quad (9)$$

```
Algorithm 3: CNN (+/- ViT) Training and Inference
// Training
D <- augmented_training_set(images, labels)
for epoch in 1..N:
   for (x, y) in minibatches(D):
      f_cnn <- CNN_backbone(x)              // local texture/lesion features
      f_vit <- ViT_branch(patchify(x))       // optional global context
      f    <- concat(f_cnn, f_vit)
      y_hat <- classifier_head(f)
      loss <- cross_entropy(y_hat, y)        // Eq. 9
      backprop_and_update(loss)
// Inference
x <- preprocess(field_image)
y_hat, conf <- forward_pass(x)
if conf < THRESHOLD or domain_shift_detected(x) then
   flag("low confidence, possible lab/field domain shift")
return y_hat, conf
```

VI. ERA 4: RETRIEVAL-AUGMENTED LARGE LANGUAGE MODELS

You ask the generalist LLM to answer a general question about agriculture, and it answers it quite fluently, and that's the issue: it answers it fluently, but it's wrong; the literature calls it "hallucination" [16], [17], [50]. In most situations, this is not an important distinction, but in the context of dosages for pesticides, it can be costly as well as being unhelpful when the dosage is incorrect, as it can result in loss of crops and soil. The regular solution is retrieval-augmented generation, which is just to retrieve material from a trusted source, be it PAU and ICAR-IIMR advisories, peer-reviewed local trial reports, or Krishi Vigyan Kendra bulletins, and then give the model the instruction to answer only based on what it has retrieved, typically with explicit mention of the passage from which it is taking the information [18, 51, 52].

This has been rapidly picked up by agricultural-advisory research, but seldom with a maize crop from Punjab. A retrieval-only system can be configured in a Hindi-English bilingual fashion using

Cadernos de Pós-Graduação em Direito Político e Econômico
Published by: Centro de Estudos Acadêmicos Press
ISSN: 1678-2127
Volume 26 ***Issue 2,*** **2026**
Website: https://ceapress.org

data from the Kisan Call Centre and achieve near-perfect exact-match accuracy with essentially no hallucination for any closed-domain question, as this is a good upper bound for the retrieval system; it demonstrates what retrieval can achieve when the retrieval corpus actually contains the query space [18]. A region-aware RAG tested on 12 different agricultural subfields shows a 10-20% reduction in hallucination compared to the ungrounded RAG baseline and significantly higher scores from human reviewers, while a second comparison of RAG versus ungrounded across an independent agricultural domain shows essentially the same [19, 20]. As far as the image side is concerned, vision-language RAG systems, where a YOLO-based lesion detector is coupled to an LLM answer generator, have been proposed for coffee leaf disease, with the aim of maintaining the integrity of the treatment recommendations by feeding the LLM with the detected lesion and also a retrieved passage [47, 48]. The same architecture would be adopted in a maize shoot-fly and MLB, BLSB pipeline. A 5-crop RAG comparison implemented using the style of a PAU or ICAR package of practices conducted across 4 open LLMs (Mistral and Qwen2.5) shows that grounding results in better semantic faithfulness and source attribution than an ungrounded generation, though this study is not in the context of Punjab and was not tested against the actual PAU documents it was based on [52].

Figure 5. RAG-LLM architecture for a Punjab maize advisory assistant.

Farmer Query (free text or transcribed voice, Punjabi/Hindi/English)

↓

Retriever (embed query, vector search over PAU/ICAR-IIMR corpus)

↓

Top-k Retrieved Passages (POP dosage tables, trial reports, ETL rules)

↓

LLM Generator (prompted to answer only from retrieved text and cite source)

↓

Cadernos de Pós-Graduação em Direito Político e Econômico
Published by: Centro de Estudos Acadêmicos Press
ISSN: 1678-2127
Volume 26 ***Issue 2,*** **2026**
Website: https://ceapress.org

Grounded Answer + Citation (escalate to KVK if retrieval confidence is low)

The retriever simply returns the top-k cosine-similar passages, where the cosine similarity is between the query embedding q and the passage embedding d_i (Equation 10), and passes the context (top-k highest-scoring passages) to the generator.

$$sim(q, d_i) = (q \cdot d_i) / (\|q\| \times \|d_i\|) \quad (10)$$

```
Algorithm 4: RAG-Based Advisory Response
Input: farmer_query q, knowledge_base KB (PAU/ICAR-IIMR corpus, chunked)
Output: grounded_answer, citations

q_embed <- embed(translate_if_needed(q))
candidates <- vector_search(q_embed, KB, k=TOP_K)   // rank by Eq. 10
candidates <- rerank(candidates, q)
if max(similarity_score(candidates)) < RETRIEVAL_THRESHOLD then
    return ("insufficient local evidence, escalate to KVK/helpline", none)
context <- concat(candidates)
prompt <- SYSTEM("Answer only using CONTEXT. Cite passage IDs. If the
                 CONTEXT does not cover the question; say so.")
          + CONTEXT(context) + USER(q)
draft   <- LLM.generate(prompt)
verified <- check_claims_against_context(draft, context)
if verified.unsupported_claims is not empty then
    draft <- regenerate_or_strip(draft, verified)
return (draft, candidates.source_ids)
```

VII. PROPOSED EVALUATION FRAMEWORK: METHODOLOGY, NOT REPORTED RESULTS

One weakness runs through nearly all of the literature surveyed above: accuracy gets reported constantly, while the practical questions that would actually decide whether a Punjab farmer trusts and keeps using a system get asked rarely if at all. No published study puts rule-based, classical-ML, CNN, and RAG-LLM systems side by side on the same Punjab spring-maize query set, so this paper does not report comparative hallucination, latency, or cost numbers, because doing so would mean making them up. What Table II offers instead is a concrete, reproducible protocol for a follow-up study, including a future version of this paper once Punjab-specific field data actually exist per Section IX, could run to generate real numbers on a common basis, with the accuracy,

Cadernos de Pós-Graduação em Direito Político e Econômico
Published by: Centro de Estudos Acadêmicos Press
ISSN: 1678-2127
Volume 26 ***Issue 2,*** **2026**
Website: https://ceapress.org

precision, recall, and F1 definitions from Equations 2 to 5 forming the diagnostic-accuracy component.

Metric	**Operational definition**	**How to measure it**
Diagnostic accuracy	Accuracy, precision, recall, and F1 (Eq. 2 to 5) of the system's top diagnosis against expert-adjudicated ground truth	The blind expert panel (PAU entomology/pathology) adjudicates a held-out query set per class.
Grounding / hallucination rate	Proportion of factual claims in a generated answer, dosage, timing, and product name that are not traceable to a cited PAU/ICAR/peer-reviewed source	Manual or LLM-assisted claim extraction plus citation check against the retrieval corpus, per [18] to [20]
Response latency	Wall-clock time from query submission to final answer, measured on a representative low-end Android handset over a 2G/3G rural connection	On-device or edge-server timing logs; report median and 95th percentile.
Source citeability / farmer trust	Whether the answer names a specific, checkable source, a POP page or a trial citation that a Krishi Vigyan Kendra officer could verify	Presence or absence of a verifiable citation, plus a structured farmer/extension-officer trust survey (Likert scale)
Deployment cost	One-time and recurring costs per village-level deployment, device, connectivity, model hosting or API cost, and KVK staff time	Itemised cost model per deployment mode: offline rule engine vs. cloud LLM API vs. on-device quantised model

Table II is deliberately a measurement protocol, not a results table, and it is worth being explicit about why: any of the four eras could score well or badly on any given row. A well-maintained rule base cannot hallucinate by construction, but its coverage is limited to whatever its authors anticipated. A RAG-LLM can feel instant on a good connection and become unusable the moment the connection is not, which a benchmark run entirely on a lab network would never catch. A CNN is fast and works fine offline, but it gives no account of why it reached a diagnosis, and that silence affects how much a farmer trusts it independently of how often it happens to be right.

Cadernos de Pós-Graduação em Direito Político e Econômico
Published by: Centro de Estudos Acadêmicos Press
ISSN: 1678-2127
Volume 26 ***Issue 2,*** **2026**
Website: https://ceapress.org

VIII. GROUNDING CASE STUDY: CHECKING A RECOMMENDATION AGAINST CROSS-CHECKED EVIDENCE

It helps to illustrate with concrete examples. If a farmer types, 'My spring maize plants are drooping in the middle and not growing,' what might they be experiencing? What should I treat the seed with the next time? This is more or less the same description as shoot-fly and dead-heart. This question might be answered with a semi-correct answer, such as the name of a registered product for maize in another country or an incorrect rate, or an answer that addresses a different issue altogether, as reported by the widespread agricultural RAG literature [16], [17], [50]. A RAG system that is based on the corpus created in Section II. Instead, I would present the Ludhiana seed-treatment trial and a separate North Indian kharif-maize advisory, which both agree that 6 mL/600 FS imidacloprid (Gaucho) is an effective pre-treatment of maize seed before sowing, having no negative effect on germination in three seasons of testing, and that 600 FS thiamethoxam is similarly effective, as reported in related seed-treatment trials [21; 25]. Please consult the list of products currently on the market and the ratings for the newest PAU Package of Practices before purchasing a product, as there are recommendations that vary from season to season [1].

That last sentence is hedging no more; it is working. This 6 mL per kg figure is not taken from this season's Package of Practices text itself (which we could not completely retrieve in its entirety), but rather from peer-reviewed trials, and a well-founded system should state this as such and not subtly slip in a research-trial figure as if it were this year's official number. Published evidence that supports a claim and PAU's current advisory about the claim are two different things; it's this distinction that a citation-based RAG design maintains, whereas an ungrounded LLM will tend to collapse the claim into a single overconfident statement.

A second, more acute version of this problem arose when putting together Table I for this paper. Two sources for PAU-linked rates of Coragen 18.5 SC for maize stem borer exist, and both are different; one from the Package of Practices' fodder-maize section (where maize is cut green at 50 to 60 days) and one from a more general grower advisory (no specification of fodder or grain use). The difference is likely due to the different crop contexts of the two, rather than to error; the rate of 40 mL per acre for the fodder-maize section of the Package of Practices, where maize is cut green at 50 to 60 days and has an 18-day pre-harvest withdrawal caution, whereas the general grower advisory (no specification of use for fodder or grain) has a 21-day pre-harvest withdrawal caution. A retrieval system which used these interchangeably, returning the highest-ranking passage for a query, would pass a grain maize farmer a fodder-crop rate, or the other way around, without an enquirer realising that there is a mismatch. The fix is not a better dosage number but rather a retrieval design that removes the notion of ‘crop use metadata’ from a soft ranking signal to a hard filter and removes the notion of ‘grain’ vs ‘fodder’ as a soft ranking signal from a hard filter, and that surfaces the stated crop use of the source in the answer itself rather than dropping it off during the retrieval process.

Cadernos de Pós-Graduação em Direito Político e Econômico
Published by: Centro de Estudos Acadêmicos Press
ISSN: 1678-2127
Volume 26 ***Issue 2,*** **2026**
Website: https://ceapress.org

IX. THE PUNJAB SPRING-MAIZE DATA GAP

The generalisation problem documented for PlantVillage in Section V is not crop specific [46]–[49] but represents three distinct challenges in this specific example which is the topic of this paper. The first is scale and diversity: Public plant-disease data is massively privileged towards lab-captured images (covering a narrow range of crops and diseases), and alternatives that are field-orientated, like PlantDoc [43] [44], are an order of magnitude smaller than PlantVillage. The second reason is geography; obviously, none of the maize-CNN studies included in Section V utilised data from the province of Punjab. None of the field trials conducted in Uganda, Tanzania, Ghana and Namibia [12] or the general disease-severity trials in the north of India [5, 8, 30] describes the combination of soil, cultivar and microclimate that affects the expression of these symptoms in Punjab. The third is class specificity; to the best of our knowledge, there is no public dataset that distinguishes between damage from whorl-feeding pests, including fall armyworm, and damage from spring maize shoot-fly whorl or dead heart, and class specificity is important because these two call for different chemistries and may appear near identical to an untrained eye. The straightforward answer is that most of the publicly available CNN efforts for maize disease are based on PlantVillage or have images from the field that come from areas outside Punjab, and we were unable to find a publicly available, labelled image set of maize images in the field collected from Punjab that contain both shoot fly, maydis leaf blight, and banded leaf and sheath blight. However, this is not a reason to not use deep learning or RAG for Punjab maize; it is a specific and solvable problem. The trial plots are already available for research at PAU and ICAR-IIMR Ludhiana, as the disease incidences are already known, and the experts are available for labelling. They are already conducting the seasonal field trials as mentioned in Section II, and thus the necessary ground truth is available. Combining this infrastructure with a well-defined protocol for the capture of images would bridge a reasonable part of this gap at an incremental expense that is likely modest.

X. COMPARATIVE SYNTHESIS ACROSS THE FOUR ERAS

Table III collects the arguments of the qualitative trade-offs presented in Sections III through VI. Each entry is a description extracted from the cited literature and from the realisation of what each paradigm architecture makes possible or impossible and why: because, as stated in Section VII, nobody has yet run all four paradigms head-to-head against the same set of queries for Punjab.

Dimension	Rule-Based (Era 1)	Classical ML (Era 2)	Deep Learning (Era 3)	RAG-LLM (Era 4)
Input	Structured symptom checklist	Leaf image, hand-crafted features	Raw leaf image, end-to-end	Free-text/voice query, optionally

Cadernos de Pós-Graduação em Direito Político e Econômico
Published by: Centro de Estudos Acadêmicos Press
ISSN: 1678-2127
Volume 26 ***Issue 2,*** **2026**
Website: https://ceapress.org

Dimension	Rule-Based (Era 1)	Classical ML (Era 2)	Deep Learning (Era 3)	RAG-LLM (Era 4)
				with image via detector
Coverage of novel/ambiguous cases	Low, limited to anticipated rules [9], [33]	Moderate, bounded by feature design	High for classes seen in training; degrades under domain shift [47] to [49]	High in principle, bounded by retrieval-corpus coverage [18]
Traceability to an authoritative source	Very high; rule maps directly to POP text.	Low, opaque feature weighting	Low, opaque unless explainable—AI added [12]	High if citation-constrained [18] to [20], [51]
Offline / low-connectivity operation	Yes, trivially.	Yes, lightweight	Yes, if the model is quantised/on-device [44], [45]	Difficult; typically needs a network call to an LLM API
Maintenance burden as guidance changes	Manual rule edits each POP revision.	Retraining on new labelled data	Retraining/fine-tuning on new labelled data	Corpus update only; architecturally the lowest-friction to keep current
Representative Punjab-linked evidence	General Indian rule-based crop systems [10], [35]	None identified specifically for Punjab maize	Chitkara University CNN-ViT, PlantVillage-trained [15]	None identified specifically for Punjab maize; general Indian agri-RAG exists [18]

XI. CONCLUSION AND FUTURE WORK

It has aimed to keep the latter thread on the actual agronomy of spring-sown maize in Punjab, in the context of shoot fly, Maydis leaf blight, and banded leaf and sheath blight, all of which have been supported by PAU and ICAR-IIMR recommendations, as well as peer-reviewed local or regional field trials [1] to [8] and [21] to [30]. Whenever PAU guidance is traditionally rule-based, and it is more important to be able to track the rules than to be flexible, this rule-based approach still gets its due. Classical ML and deep learning provide image-based diagnoses, but there is very little literature on maize images from non-Punjab or laboratory imagery that was found here; furthermore, the literature from plant pathology generally agreed that this type of model fails badly as soon as the real field is encountered [46] to [49]. The advantage of RAG-LLMs for the farmer is that they can simply be asked a question in natural language, and if the generation is controlled to keep it within the sources cited, then hallucination can be significantly reduced compared to an ungrounded LLM [18]-[20]. They just need a stable connection and a retrieval corpus that has not yet been systematised for the specific case of Punjab spring maize.

There are three things that follow fairly directly from the gaps which this paper has attempted to name honestly. One is a coordinated effort for image collection for the Punjab spring maize crop, ideally utilising the existing multi-location field trials of PAU and ICAR-IIMR instead of using a new one to develop a labelled set of images of shoot fly, MLB and BLSB under field conditions. Another is to test the evaluation framework proposed in Section VII for real: a head-to-head test of a rule-based system, a CNN and a RAG-LLM system on the same held-out set of queries and images from Punjab, so that the qualitative comparison given in Table III can be later replaced by numbers rather than words. The third is "architectural", not empirical, meaning that it is a hybrid system, with an on-device CNN or a rule-based triage layer, that would be used for offline, low-connectivity first-pass diagnosis followed by an escalation to a RAG-LLM only when connectivity is available and confidence is really low, marrying the offline reliability of the first three eras with the conversational range of the fourth. As for the three, this last one, namely some working combination, rather than any one paradigm presented alone, appears to be the most realistic short-term option to achieve a maize advisory system that would be adopted and used by maize farmers in Punjab.

XII. AUTHOR CONTRIBUTIONS

[1]Rishabh Aryan, [2]Anju, [3]Satpal, [4]Nishtha Kalia, and [5]Dr. Sarwan Singh are co-first authors and contributed equally to the conceptualisation, methodology, and execution of this study.

XIII. FUNDING

This work was supported by the National Institute of Electronics and Information Technology (NIELIT), Ropar, Punjab, India, which also provided research facilities and infrastructure. The funder had no role in study design, data collection, analysis, interpretation, or the decision to submit the manuscript.

XIV. CONFLICT OF INTEREST

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

XV. DECLARATION OF GENERATIVE AI AND AI-ASSISTED TECHNOLOGIES IN THE WRITING PROCESS

Cadernos de Pós-Graduação em Direito Político e Econômico
Published by: Centro de Estudos Acadêmicos Press
ISSN: 1678-2127
Volume 26 ***Issue 2,*** **2026**
Website: https://ceapress.org

During the preparation of this manuscript, the author(s) utilised Claude, ChatGPT and Gemini to refine the phrasing, check English grammar, assist in structuring the text, and perform text and literature summarisation. Following the application of these technologies, the author(s) thoroughly reviewed, verified and edited all generated narrative outputs to ensure scientific accuracy and alignment with the original experimental data. The author(s) maintain full responsibility for the ultimate integrity, objectivity and originality of the published work.

XVI. REFERENCES

[1] Punjab Agricultural University, Package of Practices for Crops of Punjab: Kharif 2026, vol. 43, Ludhiana, PAU, 2026.

[2] Efficacy of pre-treated seed with Gaucho (imidacloprid) 600 FS against shoot fly, Atherigona naqvii Steyskal, in spring sown maize, field trial conducted at Ludhiana, 2010 to 2012, ResearchGate Publication 290334455.

[3] Shoot fly complex, a significant threat to cereals and their management, Journal of Entomology and Zoology Studies, vol. 12, no. 3, 2024.

[4] Maydis leaf blight of maize, update on status, sustainable management and genetic architecture of its resistance, ScienceDirect, Crop Protection, 2022.

[5] Comparative evaluation of management modules against Maydis leaf blight disease in maize (Zea mays), European Journal of Plant Pathology, 2023.

[6] Comparative Management Studies on Banded Leaf and Sheath Blight of Maize Caused by Rhizoctonia solani f. sp. sasakii, Applied and Environmental Soil Science, pubs.sciepub.com, 2020.

[7] Alleviation of banded leaf and sheath blight disease incidence in maize by bacterial volatile organic compounds and molecular docking of targeted inhibitors in Rhizoctonia solani, Frontiers in Plant Science, vol. 14, 2023.

[8] Efficacy of Newer Molecules, Bioagents and Botanicals against Maydis Leaf Blight and Banded Leaf and Sheath Blight of Maize, Punjab Pathology Journal, ppjonline.org, reporting up to 40% (MLB) and 100% (BLSB) yield loss.

[9] Expert Systems Applied to Plant Disease Diagnosis, Survey and Critical View, ResearchGate, 2016.

[10] Applications of Expert Systems for Agricultural Crop Disease Diagnosis, A Review, ResearchGate, 2017.

[11] Deep learning and content-based filtering techniques for improving plant disease identification and treatment recommendations, a comprehensive review, PMC, 2024.

[12] A survey on using deep learning techniques for plant disease diagnosis and recommendations for development of appropriate tools, ScienceDirect, 2022.

Cadernos de Pós-Graduação em Direito Político e Econômico
Published by: Centro de Estudos Acadêmicos Press
ISSN: 1678-2127
Volume 26 ***Issue 2,*** **2026**
Website: https://ceapress.org

[13] Comparative Analysis of CNN, SVM, Decision Tree, Random Forest, and KNN for Maize Leaf Disease Detection Using Color and Texture Feature Extraction, Jurnal Teknik Informatika (JUTIF), 2025.
[14] Critical analysis of advanced machine learning and deep learning models for crop disease detection, challenges, innovations, and future directions, ScienceDirect, 2026.
[15] G. Shandilya, S. Gupta, H. G. Mohamed, S. Bharany, A. U. Rehman, and S. Hussen, Enhanced Maize Leaf Disease Detection and Classification Using an Integrated CNN-ViT Model, Food Science & Nutrition, 2025, doi:10.1002/fsn3.70513.
[16] International Journal for Multidisciplinary Research (IJFMR), AgriSaathi, LLMs, multimodal AI and RAG for agricultural advisory, IJFMR, 2025.
[17] Fine-Tuning and Evaluating Conversational AI for Agricultural Advisory, arXiv:2603.03294, 2026.
[18] Empowering farmers with artificial intelligence, a retrieval-augmented generation based large language model advisory framework (Kisan-BiQAS), Journal of Agricultural Engineering, 2026.
[19] AgriRegion, Region-Aware Retrieval for High-Fidelity Agricultural Advice, arXiv:2512.10114, 2025.
[20] Advancing site-specific disease and pest management in precision agriculture, from reasoning-driven foundation models to adaptive, feedback-based learning, arXiv:2510.24650, 2025.
[21] Control efficacy of different seed dressing insecticides against sorghum/maize shoot fly, Atherigona spp., ResearchGate Publication 323726565.
[22] Assessment of efficacy of thiamethoxam 30% FS as seed treatment against early-stage pests of maize, International Journal of Chemical Studies, biochemjournal.com, 2026.
[23] Evaluation of fungicides and biocontrol agents for suppression of banded leaf and sheath blight of maize (Zea mays), S. I. Harlapur and L. S. Rajput, ICAR AGRIS record 64748097425ec3c088f8cf84, 2015.
[24] Management of Shoot Fly and Stem Borers on Maize, P. Meti, A. G. Sreenivas, K. Prakash, M. L. Jat, Venkateshalu, A. Prabhuraj, and N. Manjunath.
[25] Maize Kharif, Kharif Maize Crop information, apnikheti.com, independent general advisory corroborating imidacloprid 6 mL/kg seed treatment for shoot fly.
[26] Bioefficacy of different insecticides against shoot fly (Atherigona spp.) in kharif sorghum and maize, thepharmajournal.com, 2022.
[27] Crop Production Guide, Maize (Zea mays L.), Tamil Nadu Agricultural Department, tnagriculture.in, mancozeb/zineb 2 to 4 g/L and propiconazole 0.1% recommendation for Maydis leaf blight.

[28] Insights into maydis leaf blight resistance in maize, a comprehensive genome-wide association study evaluated at PAU Ludhiana and PAU Gurdaspur, PMC, 2024.
[29] Management of maydis leaf blight, in-vitro and in-vivo fungicide screening identifying propiconazole (0.1%) as most effective, ResearchGate Publication 363535414.
[30] Impact of weather parameters and integrated disease management practices on banded leaf and sheath blight and grain yield of maize, European Journal of Plant Pathology, 2025.
[31] Sustainable Management of Invasive Fall Armyworm, Spodoptera frugiperda, Agriculture (MDPI), vol. 12, no. 9, p. 2150, 2022.
[32] Insecticide susceptibility vis-a-vis molecular variations in geographical populations of fall armyworm, Spodoptera frugiperda, in India, PMC, 2022.
[33] Developing an Expert System for Plant Disease Diagnosis, CLIPS-based, Journal of Artificial Intelligence, 2008.
[34] J. Durkin, Expert Systems, Design and Development, and related CLIPS-based plant-disease expert system work and certainty-factor combination rule, Scialert jai.2008.78.85.
[35] An expert system for diagnosing plant diseases using certainty factor and backward chaining based on android, IOP Conference Series, 2020.
[36] CropPestO, an ontology-based expert system for crop pests and diseases recognition, Electronics (MDPI), vol. 10, no. 8, p. 905, 2021.
[37] Plant Leaf Disease Recognition Using Random Forest, KNN, SVM and CNN, ResearchGate/Academia.edu.
[38] A Hybrid Machine Learning Approach for Detecting and Assessing Zyginidia pullula Damage in Maize Leaves, Applied Sciences (MDPI), vol. 15, no. 10, p. 5432, 2025.
[39] Enhanced residual-attention deep neural network (MaizeNet) for disease classification in maize leaf images, Scientific Reports, 2025.
[40] Deep learning models for enhanced in-field maize leaf disease diagnosis, ScienceDirect, 2025, field data from Uganda, Tanzania, Ghana, Namibia.
[41] Mobile based deep CNN model for maize leaf disease detection and classification, Plant Methods, 2025.
[42] A mobile-based system for maize plant leaf disease detection and classification using deep learning, PMC, 2023.
[43] N. Noyan, Uncovering bias in the PlantVillage dataset, arXiv, 2022.
[44] A Unified Self-Supervised Framework for Plant Disease Detection on Laboratory and In-Field Images, Electronics (MDPI), vol. 14, no. 17, p. 3410, 2025.
[45] Plant disease classification in the wild using vision transformers and mixture of experts, Frontiers in Plant Science, 2025, reporting the in-distribution versus field accuracy gap.
[46] AgriPath, A Systematic Exploration of Architectural Trade-offs for Crop Disease Classification, arXiv:2603.13354, 2026.

[47] Overcoming LLM Challenges using RAG-Driven Precision in Coffee Leaf Disease Remediation, arXiv:2405.01310, 2024.
[48] Vision Meets Language, A RAG-Augmented YOLOv8 Framework for Coffee Disease Diagnosis and Farmer Assistance, arXiv:2505.21544, 2025.
[49] Artificial Intelligence-Driven Plant Disease Detection and Diagnosis, A Comprehensive Review of Deep Learning Approaches, Multimodal Sensing Technologies, and Future Perspectives in Precision Agriculture, PMC, 2026.
[50] Wheeler formula for Per cent Disease Index and standard plant-pathology field-severity scoring conventions, B. E. J. Wheeler, An Introduction to Plant Diseases, 1969, as applied in An Integrated Approach to Evaluate the Efficacy of Different Fungicides, Botanicals, Bioagents and Their Combinations against Maydis Leaf Blight, aatcc.peerjournals.net, 2025.
[51] ICAR-Indian Institute of Maize Research, PAU Campus Ludhiana, Punjab, Maize Production for Food, Feed and Fodder, ResearchGate Publication 380898025, 2024.
[52] PAU Advisory for Maize Crop, posted by PAU, Ludhiana, Punjab, 2020-09-17, Apni Kheti Expert Advisory, https://www.apnikheti.com/en/hr/expert-advisory-details/pau-advisory-for-maize-crop/3475.