

ORIGINAL RESEARCH

Open access

Integrating Fall Armyworm Pressure into Echo State Network-Based Crop Yield Prediction

Samira Gharbi¹, Nabil Jebali², Ahmed Ben Ali^{1*}

Abstract

Crop production around the globe is increasingly at risk from invasive organisms, particularly the Fall Armyworm (FAW, *Spodoptera frugiperda*), which has caused severe deterioration of maize harvests across both Africa and Asia starting in 2016. To enable targeted pest control, an adaptive Echo State Network (ESN) was constructed that projects yearly maize output while factoring in FAW burden. A 15-year (2010–2024) monthly repository was built, integrating vegetation indicators from satellites, field-level meteorological measurements, soil nutrient readings, and FAW observation tallies. The intensity of FAW infestation is represented on a continuum from 0 to 100, derived by fusing monthly trap catch figures (40%) with larval concentration (60%). The ESN is initially calibrated on the entire body of environmental features to estimate crop performance. Next, isotonic regression is deployed to link infestation intensity with the ESN's positive residual errors, yielding a monotonic penalty mapping. This mapping captures the magnitude of output depression at varying infestation levels. When generating estimates, the inferred penalty is applied to the unadjusted ESN projection, modifying harvest forecasts to incorporate pest-driven losses without changing the foundational ESN configuration. In testing across held-out folds, the FAW-sensitive ESN yields an R^2 near 0.55 and cuts forecast discrepancies by as much as 67% compared to unadjusted reference models, faithfully reproducing documented harvest declines beyond 20% during major infestation episodes. The approach beats standard regression and deep neural architectures by comparable margins. It supports agriculturalists in directing countermeasures toward high-threat areas, lowering both chemical input volumes and operational expenditures. These outcomes emphasize its promise as a forward-looking alert mechanism for pinpointed actions that curtail agrochemical usage and improve input distribution. Ongoing field assessments will gauge its capacity to scale and its tangible efficacy within maize-growing zones affected by FAW.

Keywords Precision agriculture, Echo State Networks (ESN), Fall Armyworm, Pest management, Crop yield prediction

*Correspondence:

Ahmed Ben Ali
ahmed.benali@iresa.agrinet.tn

¹ Department of Pest Pressure Modeling and Yield Prediction, Faculty of Agriculture, University of Tunis, Tunis, Tunisia

² Department of Fall Armyworm Monitoring and ESN Forecasting, Faculty of Agriculture, University of Sfax, Sfax, Tunisia

Introduction

Farming systems across the world contend with formidable challenges arising from pests and pathogens, which are thought to drive annual harvest losses in the range of 20%–40% [1–3]. Such shrinkage threatens nutritional security and extracts a toll of hundreds of billions of dollars from the worldwide economy. Standing out among these hazards, the Fall Armyworm (FAW, *Spodoptera Frugiperda*) has surfaced in recent times as an exceptionally ruinous transboundary pest. After it first reached the African continent in 2016 [4], FAW extended swiftly into more than 50 nations spanning Africa and Asia, assaulting upward of 350 botanical varieties (with maize as a primary host) [5] and inflicting extensive destruction

[6]. Where outbreaks have been left untreated, maize harvest losses have in some cases topped 40%. The United Nations Food and Agriculture Organization (FAO) has labeled FAW as a foremost peril to African food systems [3], underlining how critical dependable observation and containment approaches have become. Long-standing routines centered on human field inspections and blanket chemical spraying are often inadequate, expensive, and environmentally harmful. A clear gap exists for precision farming instruments able to spot pest surges at an early stage and estimate their repercussions, thus permitting well-timed, site-specific responses while holding chemical dependence to a minimum [7, 8]. Because differing infestation severities call for customized management, the likelihood of severe monetary strain later on can be contained. Echo State Networks (ESNs) fit this purpose especially well owing to their strength in representing intricate, non-linear biological temporal patterns like pest population cycles. They also carry a far lighter training footprint relative to other thoroughly optimized machine learning pipelines. Evidence shows, for example, that catching pests while they are still in the larval phase demands a much smaller intervention footprint than tackling them at full maturity [9, 10]. In addition, substitute practices like habitat redesign can be brought to bear, cutting back the need for sweeping deployment of wide-acting insecticides [11].

Chemical measures, typified by uniform pesticide distribution, have traditionally dominated agricultural pest suppression. Yet these tactics routinely trigger profound ecological upset, encompassing chemical resistance build-up, the wiping out of useful species, and environmental fouling [12]. FAW throws up distinct management hurdles due to its compressed reproductive turnover, strong dispersal aptitude, and flexibility across diverse climatic regimes [13]. Fine-grained analyses of FAW life history pinpoint marked susceptibility during certain growth windows, most notably the larval instars, when well-timed action could drastically shrink infestation levels [14, 15]. Field evidence from nations like Zambia, Kenya, and South Africa lays bare the harsh toll of FAW episodes, where economic damages climb past millions of dollars each year through diminished harvests and swelling management outlays. A deep appreciation of FAW ecological dynamics is indispensable for engineering precision-based agricultural interventions and for pivoting away from environmentally corrosive chemical regimes toward pinpoint biological and ecological alternatives [16].

Breakthroughs in Internet of Things (IoT) sensing and machine learning (ML) have unlocked fresh horizons for pest handling within precision farming paradigms. Intelligent detector webs (spanning automated pheromone attractant devices, airborne camera drones, and even electronic odor detectors tuned to pest signatures) can keep fields under uninterrupted watch and accumulate rich data caches on pest prevalence and ambient conditions [16, 17]. To illustrate, layered observation architectures that combine agro-environmental loggers and e-nose units have been put forward for prompt FAW identification. Such technologies permit mechanized data harvesting at elevated spatial granularity and revisit frequency, far outstripping what manual checking can achieve. Alongside these developments, artificial intelligence (AI) and predictive model frameworks are able to tap into these data streams to foresee pest behavior patterns. AI-driven forecasting structures have proven their worth in anticipating crop maladies and pest flare-ups by merging diverse input channels like orbital imagery, climatic archives, historical outbreak logs, and live field-sensor feeds. This melding of IoT and AI delivers real-time surveillance, advanced cautioning, and pinpoint risk quantification for pest surges, a genuinely transformative capability for safeguarding cultivated plants [7]. Receiving early notice enables agricultural producers to shift from responsive pest suppression toward preemptive stewardship – applying integrated pest management (IPM) tactics exclusively at infestation "hotspots" when circumstances warrant, replacing indiscriminate broadcast treatment [3]. Site-specific actions of this sort restrain superfluous pesticide discharge and the expenditures that go with it, in step with the aims of sustainable farming [1].

With respect to the model architecture at hand, AI and IoT figure in as follows: every detector feed and machine-vision output ultimately gets transformed into numerical descriptors that are fed straight into the ESN. Internet-connected instruments capture raw weather variables, edaphic readings, and trap interceptions at a temporal frequency finer than one day. A lightweight pre-processing step condenses or accumulates these signals into per-month values and dampens random fluctuations (one example being a Kalman filter applied to moth capture sequences [18]). Embedded or cloud-resident AI modules transcribe drone-acquired imagery into spectral vegetation measures matched to the identical monthly periodicity. The resulting fifteen-feature vector – climatic means, crop status metrics, and FAW descriptors – is min–max normalized and routed without modification into the ESN dynamical reservoir. No auxiliary inference or forecast

modeling takes place upstream of the ESN, preserving a clean demarcation among raw data capture, attribute construction, and yield forecasting operations.

The choice fell on Echo State Networks in preference to end-to-end trained recurrent designs such as LSTMs or GRUs on three practical grounds.

First, concerning sample size: the assembled 15-year chronicle contains only a few tens of full growing-cycle instances, a data-poor context where the millions of tunable weights inside deep RNN architectures commonly succumb to overfitting; by contrast, an ESN adjusts only its linear output stage, involving two to three orders of magnitude fewer parameters.

Second, regarding run-time economy: once reservoir meta-parameters are fixed, a forward pass consists of a lone matrix–vector product followed by a hyperbolic tangent per timestep; executing the full 12-month sequence on a basic low-power microcontroller takes less than 10 ms, which makes on-farm, instantaneous recalculation feasible without GPU hardware.

Third, concerning resilience when data distributions shift: because the randomized recurrent layer is immutable, additional sensor modalities (for instance, extra traps or meteorological stations) can be tacked on as fresh input channels with only modest retraining effort, and the model's restricted echo-state memory organically suppresses fleeting outlier spikes—a beneficial trait given the inherent noisiness of pest monitoring records. When taken as a whole, these properties make ESNs a markedly better match than bulkier RNN variants for a data-limited, edge-deployed, pest-cognizant yield-forecasting tool.

Within this framing, we concentrate on building a forward-looking analytical instrument for FAW within maize-dominant agricultural landscapes. Fall armyworm surveillance and harvest estimation form two sides of the same coin: prompt recognition of an infestation event can steer recalibrations of both output projections and intervention schedules. Our investigation confronts the problem of weaving pest demographic signals directly into the fabric of crop forecasting engines. The central offerings put forward in this manuscript are:

- i. The engineering of a bespoke adaptive Echo State Network calibrated expressly for maize output projection that folds in Fall Armyworm population fluctuations, harnessing the framework's strength at rendering elaborate temporal structures resident in agronomic data streams.
- ii. The devising of an unprecedented pest-contingent deduction scheme that overtly modulates harvest estimates as a function of FAW infestation magnitude, anchored in financial damage benchmarks and expert insight into pest-versus-yield linkages.
- iii. The merging of heterogeneous ecological and organismal data conduits (atmospheric observations, ground composition readings, distance-sensed greenery proxies, and organism tracking feeds) into an encompassing predictive shell that mirrors the layered interdependencies between growing conditions and biotic stress.
- iv. The tangible demonstration that model architectures cognizant of FAW presence deliver marked gains in yield estimation reliability under outbreak conditions, compressing forecast divergence by up to 67 % when set beside orthodox techniques.
- v. The delivery of an actionable pre-alert infrastructure for spatially precise pest governance, opening the door to swift and accurately aimed responses, thereby curtailing the need for manufactured insecticides and buttressing environmentally sensitive cultivation protocols.

We put forth a recurrent neural construct resting on Echo State Network (ESN) principles that digests biological and ecological variables sourced from disparate origins. Detector-derived FAW demographic indices serve as input signals for projecting harvest tonnage. The ESN blueprint earns its selection through computational thrift and a demonstrated facility

for rendering non-linear chronological behaviors [19, 20]—dynamics endemic to managed ecosystem trajectories. So as to encode the toll exacted by pest feeding, we augment the architecture with an original penalization apparatus that reweights yield projections contingent upon FAW infestation loads at instant t . Stated plainly, the system “internalizes” the inherent growth rhythm dictated by environmental covariates. At the point where a FAW population upswing is registered, a discounting term activates to translate that signal into estimated impairment. The fusion of statistical pattern extraction with theory-grounded pest consequence formalization aspires to sharpen both the precision and the operational relevance of seasonal output outlooks under pest duress.

The impetus propelling this modeling endeavor is many-layered. Established statistical tools and modern algorithmic yield estimators alike exhibit four pivotal deficiencies. To begin with, the prevailing operational backbone remains exclusively climatological, which translates into methodical yield inflation on occasions when pestilence materializes. Secondly, deep neural schemes that indirectly absorb trap visitation frequencies seldom witness a sufficient corpus of outbreak episodes during fitting to reliably learn the authentic non-linear yield response curve. Thirdly, the vast majority of platforms disseminate a lone mid-cycle snapshot and possess no mechanism for revision when infestation intensity escalates during later phenological stages. Fourthly, neither paradigm generates a transparent bridge from pest load severity to projected harvest shrinkage—an interpretable conduit that is a prerequisite for setting integrated pest management decision triggers. A further consideration is that, with the ESN reservoir locked in place after a single round of hyperparameter tuning, each incoming observation bundle encounters nothing beyond a lone matrix–vector product plus a tanh evaluation [21]. Operational expense scales in proportion to reservoir width from one time increment to the next. It follows that, after the optimized hyperparameter set is stored locally, the model processes streaming inputs and emits a refreshed yield projection within milliseconds, affording bona fide in-field real-time nowcasting without any repeat training cycle or reliance on non-standard computation gear. These lacunae supply the rationale for the FAW-attentive Echo State Network advanced in this paper, wherein a lucid penalty term imbues expert domain awareness of pest depredation and empowers the estimate to shift in real time as supplementary infestation intelligence streams in.

To anchor our contribution within the extant body of knowledge, an exhaustive canvassing of modern predictive paradigms is first presented. The surveyed landscape reaches from deep convolutional and transformer-driven formulations to advanced reservoir computation constructs, and foregrounds the manner in which our Adaptive ESN, characterized by its overt integration of pest population kinetics, stands in relation to these comparator frameworks.

Survey of the Literature

Remote observation platforms together with algorithmic learning now furnish potent instruments for ecological surveillance, terrain-type discrimination, and harvest projection. Within this portion of the paper, we examine ten frameworks illustrative of contemporary progress. We begin by scrutinizing image-driven strategies, encompassing convolutional pipelines bolstered by shape-based descriptors and selective focus modules. Following that, we turn to blended spatio-temporal designs and reservoir computation formulations suited to time-varying estimation tasks. **Table 1** catalogs the data sources, methodologies, and principal outcomes associated with each model. The section closes by measuring these alternative techniques against our Adaptive ESN, a construct distinguished by its direct embedding of Fall Armyworm population behavior into crop output prognostics.

Table 1. Salient attributes within the FAW spread compilation, covering ecological (climatic/edaphic/vegetative) and organismal (pest abundance) indicators.

Feature group	Variables (Units)	Data sources
Climate conditions	Precipitation (mm), Air Temperature (°C), Relative Humidity (%)	In-field meteorological stations and archived climate datasets

Soil properties	Soil Water Content (%), Nutrient Composition (N-P-K levels)	IoT-based soil monitoring sensors and laboratory soil analysis
Crop condition	Near-Infrared (NIR) signal, reflectance percentage values	Proba-V satellite observations and NASA Earthdata platforms
Fall armyworm (FAW) infestation	Pheromone trap captures (moths/day, contributing 60% weight), Larval density (larvae/plant, contributing 40% weight); collectively defined as FAW Severity	IoT-enabled pheromone trapping systems and manual field scouting
Crop yield (target variable)	Final harvested grain yield (tons per hectare)	Direct harvest yield measurements collected at the plot level

Related works

A deep transfer-learning system that couples TrAdaBoost.R2, an adaptive boosting algorithm, with a Bidirectional LSTM network was put forward by Joshi *et al.* to estimate crop yields across disparate domains [22]. Their architecture ingests multi-source inputs (climate records and satellite time-series imagery) and leverages instance-weighting to contend with data-limited regions. Salient contributions consist of deploying BiLSTM as the sequential feature learner and adopting fine-tuning, domain-adversarial training, and boosted transfer to bolster generalization over different climatic zones. By harnessing knowledge from data-rich geographies for use in data-poor ones, this methodology yielded dependable winter wheat yield forecasts.

Comparison to our model: The BiLSTM transfer learning scheme targets cross-region domain alignment, whereas the Adaptive ESN is engineered to accommodate pest dynamics within time-series streams. Of particular note, pest infestation metrics and reservoir computing mechanisms are absent from Joshi *et al.*'s framework.

Farmonov *et al.* put forward HypsLiDNet [23], a deep learning pipeline for crop classification that fuses hyperspectral imagery with elevation information derived from LiDAR. The design employs a hybrid 3D-2D CNN layout. Spectral-spatial features are pulled from DLR DESIS hyperspectral inputs via 3D CNN layers, while parallel 2D CNN streams operate on LiDAR rasters. The LiDAR features undergo a novel spatial-spectral morphological attention procedure. Morphological transformations are applied to the LiDAR data foreground structural terrain and shape signals (e.g., field boundaries, plant height), after which an attention module selectively amplifies these structural signals in conjunction with spectral features. HypsLiDNet can resolve crops that exhibit similar spectral profiles but distinct physical architectures. HypsLiDNet surpassed conventional classifiers by > 14% accuracy and even edged out contemporary deep networks by roughly 1%–3% [23].

Comparison to our model: In line with ResMorCNN [24], HypsLiDNet exemplifies the cutting edge in utilizing spatial imagery (and LiDAR, in this instance) for classification tasks. By contrast, our adaptive ESN model does not operate directly on pixel-level spectral inputs; it instead derives yield estimates from temporally aggregated remote sensing indices and pest population readings. HypsLiDNet's attention-driven data fusion sets a benchmark for image classification, whereas our ESN is oriented toward learning temporal patterns. Our technique proves valuable when dynamic pest-crop feedback loops are paramount, yet it foregoes the inclusion of high-dimensional spatial descriptors. This divergence accentuates the mutually reinforcing roles of spatial CNN architectures and temporal reservoir models within precision agriculture.

Qiao *et al.* constructed a hybrid deep architecture termed the Spatial-Spectral-Temporal Neural Network (SSTNN), which unites 3D convolutional neural networks with recurrent layers for yield estimation [25]. An initial 3D-CNN isolates spatial and spectral features from sequences of multi-spectral satellite imagery, forwarding its output into stacked bidirectional LSTM cells to seize temporal dynamics. Training on multispectral MODIS time-series data allowed this model to assimilate joint spectral-temporal representations for predicting winter wheat and maize outputs across diverse Chinese

regions. The defining advance lies in the end-to-end fusion of spatial, spectral, and temporal cues by bridging a CNN for imagery with an RNN for sequential learning.

Comparison to our model: SSTNN operates on raw multi-temporal imagery through a deep CNN-LSTM conduit, unlike the Adaptive ESN, which rests on reservoir computing and curated inputs (e.g., indices and pest censuses). Moreover, pest population factors remain unaccounted for in Qiao *et al.*'s model, while our Adaptive ESN institutes a pest-driven penalty that modulates the reservoir state.

Marzvan *et al.* [26] showcase time-series remote sensing applied to ecological surveillance, specifically gauging the proliferation of the invasive water fern *Azolla filiculoides* within Iran's Anzali Lagoon. Drawing on three decades of Landsat records (1988–2018), they charted *Azolla* extent through spectral classification via the Spectral Angle Mapper algorithm and probed temporal trajectories against environmental drivers. Their investigation uncovered that *Azolla* infestation oscillated over decades rather than following a monotonic rise. Significantly, they established that *Azolla* spread correlated positively with lagoon water volume and rainfall, and negatively with temperature – humid years fostered *Azolla*, while hotter intervals triggered retreats. These findings, extracted from longitudinal satellite observations and field data, supported a risk appraisal: pinpointing years and zones of hazardous *Azolla* blooms that jeopardize native biodiversity. This case illustrates how established remote sensing procedures (SAM classification [27]) paired with time-series examination can steer environmental decision-making. Comparison to our model: Marzvan *et al.*'s inquiry [26] and our research both enlist time-series remote sensing to confront agricultural or environmental challenges. Their methodology, however, leans on a classical spectral classifier and statistical correlation, whereas we deploy a machine learning construct (ESN) to infer dynamic interdependencies. Their emphasis is retrospective charting of vegetation spread and its determinants; ours is prospective modeling of crop yield in the presence of pest dynamics. A shared thread is the weaving in of exogenous variables: Marzvan *et al.* [26] fold climate variables into the interpretation of *Azolla* trajectories, just as we explicitly embed pest populations into yield forecasting. Our ESN paradigm may be regarded as a more automated, learning-driven continuation of that investigative ethos, capturing nonlinear pest-induced crop responses analogous to capturing climate-driven weed proliferation in Marzvan *et al.*'s effort [26].

Felegari *et al.* delve into a machine-learning scheme for charting soil cadmium (Cd) contamination with the aid of the new Landsat-9 platform [28]. They amassed 100 field soil specimens of Cd concentration and aligned them with 20 multitemporal Landsat-9 scenes of the surveyed locale. The objective was to scrutinize an array of regression techniques (ranging from elementary linear to sophisticated machine learning) for faithfully inferring Cd levels from spectral readings. By invoking "experimental models," the authors indicate calibrations rooted in laboratory or field experiments that tie spectral indices or band reflectance to Cd concentrations. Their outcomes spotlighted the optimal model (judged by prediction error and R^2) for Cd mapping, facilitating the production of a contamination chart from the Landsat time series. This route delivered an economical proxy for tracking heavy metal pollution, with tolerable accuracy in light of the somewhat sparse ground truth sampling.

Comparison to our model: This investigation represents a further case of marrying remote sensing with machine learning to address an agricultural environmental concern (soil contamination). The procedure – trialing multiple model forms to pinpoint the strongest fit – is more empirically driven and oriented toward static spatial inference (mapping concentration at discrete locations) than our adaptive ESN, which is engineered for temporal forecasting. A pivotal distinction is that Felegari *et al.*'s selected model is plausibly a direct regression (perhaps linear or tree-based), which affords interpretability but may fall short in characterizing intricate dynamics. Our ESN, conversely, is purposefully chosen for rendering nonlinear temporal dynamics in crop–pest interplay. Furthermore, Felegari *et al.*'s framework demands ground specimens for calibration, a data-heavy prerequisite; our methodology can draw upon time-series archives (e.g., yield figures and pest counts) potentially more accessible through agricultural extension networks. In brief, both strategies mobilize remote sensing and learning algorithms, though our ESN framework stretches the paradigm into dynamic forecasting, whereas Felegari *et al.* operate within the bounds of spatial estimation.

Mirhoseini *et al.* devise a deep learning apparatus for predicting crop yields from Earth observation feeds, constituting one of the newest frontiers in remote sensing-driven yield forecasting [29]. Their system, referenced here as ConvLSTM–ViT, weaves together three modules: a 3D-CNN, a ConvLSTM, and a Vision Transformer. The 3D-CNN commences by distilling spatial features from multispectral image sequences (interpreting spectral bands and, potentially, temporal frames as the 3D input). Those representations flow into a Convolutional LSTM, which seizes the temporal unfolding of crop development throughout the season (e.g., the progression of vegetation index trajectories). Following this, a ViT module administers global self-attention across the feature series, enabling the model to home in on the most salient spatio-temporal signatures for yield estimation. This composite architecture draws on the respective merits of CNNs (spatial hierarchies), RNNs (temporal memory), and Transformers (extended-range attention) within one end-to-end network. The authors trialed this structure on soybean yield prediction in a case study, harnessing multispectral satellite feeds and archival yield registers as training sources. ConvLSTM–ViT markedly surpassed baseline counterparts (including antecedent CNN+LSTM formulations such as DeepYield), posting considerably reduced prediction error. It drove the root-mean-square error down to roughly 6.3 (yield units) relative to ~11.8 for the runner-up, a gain of roughly 35%. This stands as a striking advancement, testifying to the model's facility in apprehending complex yield determinants.

Comparison to our model: The ConvLSTM–ViT construct bears direct relevance to our endeavor, as both grapple with yield prediction underpinned by remote sensing. The principal parting lies in the modeling philosophy: Mirhoseini *et al.* resort to a thoroughly trainable deep network encompassing millions of parameters, presupposing ample data volumes and a significant computational budget for training. We, in contrast, put forward an Adaptive Echo State Network, a reservoir computing variant wherein solely the output layer undergoes training. ESN was elected for its strength in rendering nonlinear dynamical systems with negligible training expense. This renders our design considerably leaner. Deep feature extractors are not trained; rather, we lean on the reservoir's immutable random projections and the adaptive scheme to absorb incoming information like pest incidence. Input features mark another divergence: ConvLSTM–ViT ingests dense multispectral image time-series and implicitly learns to surface vegetation signals, whereas our route explicitly operates on derived attributes (e.g., vegetation indices, pest population tallies). Incorporating Fall Armyworm dynamics supplies a novel perspective absent from ConvLSTM–ViT. While formidable, their model does not outwardly embed pest population data. Hence, our ESN-based formulation could be more responsive to abrupt pest outbreak shocks on yield (owing to the adaptive reservoir update), whereas ConvLSTM–ViT thrives at capturing typical weather and phenological patterns discernible through remote sensing. ConvLSTM–ViT epitomizes the state-of-the-art in deep-learning yield prediction, excelling on large corpora, while our scheme supplies a more lightweight, interpretable substitute that bakes in domain knowledge (pest influences) and may hold an advantage in data-constrained settings.

Khaki and Wang advanced a composite deep learning construct joining specialized CNN modules with a recurrent network to project crop yields [30]. In their blueprint, one CNN (W-CNN) extracts temporal patterns from sequential weather readings, while a second CNN (S-CNN) seizes spatial variability from soil profiles. The distilled features are relayed into an RNN tier that models year-over-year dependencies, indirectly accommodating genetic yield gains over successive seasons. Applied to multi-year environmental and management records across US Corn Belt crops, this CNN-RNN attained notable precision, with the RNN segment tracing long-run trajectories and the CNN constituents processing intra-seasonal weather and soil influences.

Comparison to our model: The CNN-RNN configuration deliberately disentangles weather, soil, and temporal trend constituents inside a deep neural network, whereas the Adaptive ESN furnishes a more frugal reservoir alternative. Crucially, Khaki and Wang's formulation omits pest outbreak intelligence; it zeroes in on climatic and edaphic factors. The Adaptive ESN, in opposition, dynamically recalibrates according to pest pressure (e.g., Fall Armyworm) by way of a novel reservoir penalty, confronting a dimension left out of the CNN-RNN scheme.

For the concurrent prediction of multiple crop yields, Khaki *et al.* [31] rolled out a deep CNN model christened YieldNet. This structure wields a shared "backbone" CNN feature extractor on remote sensing inputs and harnesses transfer learning to jointly estimate corn and soybean yields via shared learned representations. A novel multi-target loss formulation permits simultaneous training against two yield outputs, letting the model register interdependencies between

the yield patterns of different crops. YieldNet was conditioned on county-level satellite observations for U.S. corn and soybean, attaining an early-season prediction error below 9% of the mean yield for both commodities, outperforming single-crop counterparts.

Comparison to our model: YieldNet pursues multi-crop yield forecasting through a deep CNN and cross-crop weight-sharing, whereas the Adaptive ESN is purpose-built for a single crop (maize) augmented with pest-sensitive adaptive elements. Our ESN eschews multi-task learning; it instead implants a pest infestation penalty into the reservoir to lift single-crop yield fidelity during infestation episodes. YieldNet's forte resides in multi-crop data fusion and transfer learning, yet it lacks entomological or adaptive components.

Additional related investigations have centered on spatial analysis duties in remote sensing that bear less immediate relevance to yield forecasting. To illustrate, numerous deep architectures tackle crop classification or field segmentation – determining crop types or outlining crop parcels in imagery [32]. Super-resolution methods have likewise been pursued, enlisting deep generative networks to sharpen the spatial granularity of agricultural satellite scenes [33]. These avenues boost image definition and classification fidelity but engage with static spatial content. Yield prediction, by contrast, forms a temporal regression problem. Methods tailored for classification, segmentation, or image super-resolution do not directly contend with time-series yield forecasting and consequently sit outside the central ambit of crop yield modeling.

Having surveyed a broad body of literature on cutting-edge prediction frameworks, the subsequent section elaborates the methodology for embedding FAW parameters into the ESN. We additionally lay out experimental findings measured against baseline models and reflect on the ramifications for precision pest management and sustainable agriculture.

Materials and Methods

We assembled a multivariate collection of observations spanning the crop production environment together with pest multiplication indicators. **Table 1** presents a summary of the data streams and variables incorporated. The ecological descriptors encompass conventional agro-climatic readings and vegetation indicators, whereas the biological descriptors record FAW infestation magnitude across time. Information was gathered from a trial maize site outfitted with IoT devices and FAW surveillance traps over several cropping cycles:

- **Climatic readings:** Uninterrupted registrations of precipitation (mm), ambient temperature (°C), atmospheric moisture (%), and ground wetness (%) were logged, given that these elements shape both plant development and pest reproductive rhythms. Archived climatological sequences (e.g., dry periods or extended rainfall events) were additionally folded in, as such patterns frequently herald infestation surges.
- **Satellite-derived vegetation proxies:** Canopy condition markers such as Near-infrared Spectroscopy (NIR) were collected through recurring orbital overpasses. These indicators function as surrogates for plant robustness and strain. Elevated NIR reflectance typically aligns with denser, more vigorous canopies, thereby serving as a direct signal of thriving crops [34].
- **Agronomic and edaphic descriptors:** Supplementary situational variables covered crop developmental phase, sowing date, and ground characteristics. Growth stage records (for instance, vegetative versus tasseling) carry weight because FAW herbivory at initial growth stages can exert a different yield penalty than damage incurred later in the cycle. Edaphic information (type and nutrient status) supplies the foundational production ceiling and can temper pest consequences (vigorous stands may endure herbivory more effectively).
- **FAW infestation measurements:** The principal biological signal is the FAW abundance trajectory, sourced from pheromone lure tallies and in-field inspection logs. Automated lure stations deployed across the trial site furnished daily enumerations of captured adult moths, acting as an advance cue for larval colonization. We also incorporated per-day field counts of larval concentration (number of FAW juveniles per plant) and the fraction of plants displaying fresh

herbivory symptoms, subsequently condensed into monthly bins. Together, these readings capture pest load at instant t . Within our compilation, FAW counts displayed episodic surges mirroring generational turnovers and migratory waves. Crucially, we register these demographic chronicles against harvest outcomes so that the model can absorb the mapping between covariates and output. A point worth underscoring is that FAW incidence records necessitated imputation owing to insufficient locally accessible measurements. The data repository primarily draws upon the FAO FAW compilation [35].

Gaps in monthly FAW variables were supplied from the FAO Global Fall Armyworm Monitoring and Early Warning System (FAMEWS) repository [36]. For every incomplete season, we began by pairing it with the geographically closest FAMEWS administrative unit (latitude band within $\pm 1.5^\circ$). The mean weekly moth lure catch for that district during the matching agronomic year was then determined and rescaled to our plot level via the ratio of the longest uninterrupted trap log at the site to the FAMEWS mean (a single multiplicative constant per location). The rescaled weekly sequence was binned into monthly totals and linearly interpolated wherever gaps of ≤ 2 consecutive weeks existed. This pipeline injects ecologically plausible cyclic behavior but likely underrepresents severe localized eruptions; consequently, every assessment partition retains at least one season possessing fully recorded trap logs, obliging the ESN to generalize past the imputed entries and ensuring any introduced distortion is captured within the cross-validation discrepancy. **Table 1** enumerates the dataset constituents, their measurement units, and provenance.

All chronologically organized series were condensed to a monthly cadence ($\Delta t = 1$ month) for modeling purposes; meteorological variables were averaged or totaled as fitting, while pest enumerations were summed over each month. We collapsed every input stream onto a monthly time base for three reasons:

Data completeness: climatic loggers and automated lure stations yielded daily observations, yet archival FAW enumerations preceding 2018 existed solely as monthly agricultural extension summaries; retaining a finer temporal mesh would have created systematic lacunae.

Signal fidelity relative to noise: FAW moth interceptions oscillate sharply with wind conditions and lunar phase; a 30-day averaging window suppresses this stochasticity while conserving the generational surges (roughly 30–35 days) that govern crop injury.

Empirical verification: pilot trials employing weekly resolution (52×11 input nodes) shrunk the fitting corpus to fewer than 350 snapshots and drove the baseline ESN into overfitting, delivering no precision benefit relative to coarser monthly binning. We therefore regard monthly summarization as the optimum balance between biological verisimilitude and model resilience. The consolidated dataset encompassed numerous cropping cycles, furnishing a spectrum of conditions from negligible pest load to calamitous epidemics.

Ahead of model construction, we rescaled descriptors to a common interval and undertook feature crafting activities, including shifting pest variables temporally (to permit yield repercussions to surface after colonization) and deriving additive weather summaries (e.g., cumulative seasonal precipitation up to a given month). These synthetic inputs aided the model in seizing lagged relationships, such as the influence of preceding moisture deficits on pest vulnerability. The verified harvest tonnage for each season/parcel was registered at crop maturity. This yields the predictive endpoint for the ESN formulation. It merits mention that yield is construed here as an end-of-cycle consequence molded by temporally varying inputs across the season. Accordingly, the estimation challenge is structured as a time-series-to-scalar regression, wherein monthly bundles of predictors (climate, NIR, FAW enumerations, etc.) translate into a single seasonal output figure. We partitioned the compilation into a fitting subset (roughly 70% of cropping cycles) and a held-out testing subset (30%). Care was taken that the testing fold contained at least one season marked by a significant FAW episode, enabling the model's appraisal under most adverse pest circumstances.

Echo State Network model adaptation

The central estimation engine rests upon an Echo State Network (ESN), a variant of recurrent neural architecture situated within the reservoir computing family. The choice of ESN was motivated by its facility for representing non-linear dynamical systems while demanding only a light training footprint. In contrast to an exhaustively optimized Deep Neural Network (DNN), an ESN deploys an immutable, expansive "reservoir" of recurrent processing nodes wired through random linkages, leaving exclusively the readout coefficients to be fitted [37]. This arrangement renders ESNs computationally inexpensive and comparatively resistant to overfitting—a valuable trait given our somewhat constrained observational pool. ESNs have demonstrated proficiency in projecting volatile temporal sequences across disparate domains, and agriculture exhibits equivalently intricate chronological motifs (atmospheric oscillations [38], pest demographic surges and collapses [37]).

Within our instantiation, the ESN's input interface receives the multivariate descriptor bundle at each monthly step (every **Table 1** field apart from yield). The reservoir dimensionality is held constant at 11 processing units interlinked through sparse (~1%–2%) recurrent pathways. Reservoir scale and sparseness were established through methodical hyperparameter exploration, with trials spanning 5 to 200 units and 0.5% to 5% inter-unit linkage density. The lean specification of 11 units at roughly 1%–2% connectivity emerged as the optimum from a randomized grid sweep that maximized cross-validated performance. This compact reservoir supplies an optimal trade-off between computational frugality and estimation fidelity for our agricultural chronosequences. Although more capacious reservoirs (50–200 units) underwent evaluation, they exhibited a propensity to overfit on our restricted dataset, most notably under severe epidemic scenarios. The adopted architecture furnishes adequate remembrance capacity to embed seasonal temporal signatures while generalizing robustly across fluctuating pest intensities. Input projection weights W_{in} are sampled from $U(-1,1)$ with no extra scaling modifier. Recurrent coupling weights W are set uniformly at random and subsequently adjusted to uphold the echo state criterion. The spectral ceiling ρ is calibrated through random scouting and repeatedly settles near $\rho \sim 1.03$. We execute a grid-based randomized traverse across $\rho \in [0.1, 1.5]$ (incremented by 0.02) alongside ridge-regression penalization $\alpha \in [10^{-10}, 10^{-4}]$ (log-uniform spacing), grading each combination through leave-one-year-out cross-validation and picking the pairing that attains the maximal average R^2 on excluded seasons. The linear readout stage is subsequently resolved analytically over the complete reservoir state matrix, incorporating FAW-informed target attenuation via the fitted isotonic penalty mapping. This rigorous, cross-validation-steered calibration guarantees that the ESN reservoir faithfully encodes season-spanning dynamics under both minimal and elevated infestation regimes.

Readout weight estimation proceeds through ridge regression (L2-penalized least squares) to forestall overfitting in light of the reservoir's substantial dimensionality. This ridge procedure tacks on a modest penalty against inflated weight magnitudes (L2 norm) to foster generalizability. The fitted ESN thereby furnishes a non-linear transformation from the chronological succession of input variables over a cropping cycle to the estimated harvest. In conceptual terms, the reservoir operates as a non-linear digest of the entire seasonal trajectory (precipitation sequences, pest colonizations, etc.), while the readout layer acquires the capacity to pair particular reservoir activation configurations with ultimate yield realizations. **Figure 1** underneath displays a schematic of the constructed model.




Figure 1. Penalty-aware ESN flowchart.

To subject model behavior to stringent assessment while extracting maximal value from our temporally confined repository (spanning 15 years), we adopted a leave-one-year-out cross-validation regimen [39]. For every designated target year appearing in the outcomes summary, the model underwent fitting on records from the other 14 annual cycles and was subsequently appraised on the excluded year. This operation cycled through all 15 iterations, yielding projections for each observational year while guaranteeing that no held-out instance contaminated the fitting phase. Crucially, the pest-contingent deduction mapping was likewise shaped solely on the training splits, obstructing any seepage of test-period intelligence into the construct.

The ESN is built around an immutable, randomly wired recurrent reservoir that converts input chronologies (ecological descriptors and FAW variables) into an expanded dynamical representation. Exclusively, the readout tier undergoes fitting (via linear least squares) to link the terminal reservoir configuration at season closure to the ultimate harvest estimate. Within this tailored formulation, a pest-responsive penalty mapping $f(P(t))$ is engaged to recalibrate the yield projection whenever elevated FAW abundance P registers at instant t .

The assorted detectors and remotely sensed feeds arriving at each step t (climatic, edaphic, vegetation proxy, FAW enumeration) are routed into the reservoir, which revises its internal activation profile. Upon digesting the full time-

sequence up to cycle termination (harvest point), the concluding reservoir state is tapped by the readout layer to emit the projected tonnage. Model fitting proceeded under supervision, drawing upon recorded harvests from antecedent seasons to determine the optimal readout coefficient array (through L2-penalized least squares). During forward inference on a fresh cropping cycle, the ESN advances stepwise from sowing through the present month t , and can issue an estimated yield at any juncture. This projection can be refreshed each month as supplementary readings stream in, effectively delivering a rolling output outlook whose fidelity tightens as the season unfolds. Nevertheless, an unmodified ESN construct, even when supplied pest demographics among its inputs, may not autonomously internalize the full magnitude of pest-attributable yield erosion—most notably if severe epidemic episodes are sparse across the conditioning corpus. To deliberately steer the formulation toward reflecting pest injury, we embedded a yield attenuation apparatus detailed in the following subsection.

Penalty function for pest impact

To fuse discipline-specific understanding of pest herbivory into the predictive engine [40, 41], a penalty mapping was engineered that modulates the ESN's raw yield estimate as a function of FAW population intensity. The underlying concept is to replicate the harvest contraction occasioned by a given infestation load, in effect discounting the forecast commensurately with epidemic severity. In operation, the ESN first issues a foundational yield figure Y_{ESN} after assimilating all predictors. A penalty term $f(P)$ is then specified that amplifies alongside the pest abundance $P(t)$ (for instance, lure interceptions or larval density at instant t), and this term is employed to produce a corrected yield Y_{adj} according to the expression:

$$Y_{adj}(t) = Y_{ESN}(t) [1 - f(P(t))] \quad (1)$$

Here, $f(P)$ is calibrated so that $0 \leq f(P) < 1$ throughout the empirical pest range, ensuring the correction consistently shrinks the estimate (a zero-pest scenario yields $f = 0$ and $Y_{adj} = Y_{ESN}$). The parametric contour of $f(P)$ was guided by economic injury threshold constructs drawn from entomological scholarship, estimated at roughly 34% for Sub-Saharan African settings [42–45]. Put differently, established threshold concentration values of FAW—beyond which yield diminution becomes pronounced—are exploited. Building upon advisory benchmarks and published experimental evidence, a demarcation threshold is recognized (e.g., X moths per trap weekly, or $Y\%$ of crop units colonized) after which output loss accelerates. The shape of $f(P)$ is constructed to stay modest when P remains beneath this boundary and to swell more steeply once P breaches it. In lieu of imposing a piecewise-linear formulation manually, $f(P)$ is inferred through isotonic regression applied to the fitting seasons within each cross-validation partition. Once the ESN is fitted, the "realized penalty" is calculated for every conditioning year:

$$f_{obs} = \max \left(0, 1 - \frac{Y_{true}}{Y_{ESN}} \right) \quad (2)$$

The resultant function is confined to the interval $[0,1]$. An isotonic regression curve is subsequently fitted:

$$P \rightarrow f_{iso}(P)$$

This enforces monotonic growth (so amplified FAW pressure can never correspond to a diminished penalty) and confines output scores within $[0,1]$. To forestall any possibility of information seepage or memorization, both the ESN construct and the isotonic regression penalty mapping were shaped exclusively on the conditioning folds, never on evaluation data. Inside each partition of our leave-one-year-out cycling, the ESN underwent fitting on the designated training years, realized penalties were computed across those same conditioning years, and the isotonic regression was then tuned to those penalty figures. The complete assembly (ESN plus penalty module) was then appraised on the entirely sequestered test year. This protocol certifies that no intelligence from the test season infiltrated either the ESN coefficient array or the penalty mapping parameters, upholding a clean separation between training and assessment sources. This empirically derived trajectory encodes the degree to which the ESN overshoots actual harvests under diverse infestation intensities,

all while averting information leakage: the excluded season's FAW magnitude is never consulted during curve construction. At estimation time, $f_{iso}(P_{test})$ is evaluated for the unseen year's peak infestation, and the adjustment is computed.

$$Y_{adj,test} = Y_{ESN,test} [1 - f_{iso}(P_{test})] \quad (3)$$

Figure 2 offers a graphical depiction of the FAW penalty mapping.




Figure 2. FAW penalty visualization.

In substance, this corrective mechanism compels the model to assimilate the pattern of yield attrition that tracks pest intensity. When fielded, the system functions across two distinct regimes.

Forecasting mode: while the season is underway, the ESN continuously supplies an updated yield outlook, and the penalty is imposed using the instantaneous pest level $P(t)$ at each time step t . This arrangement permits the yield forecast to drop dynamically upon detecting an abrupt pest upsurge, thereby furnishing an advance alert of likely harvest impairment.

Scenario analysis mode: synthetic "counterfactual" explorations become possible by inserting alternate pest burden trajectories $P(t)$ (for illustration, a conjectural severe epidemic against a subdued one) through the penalty apparatus to gauge how ultimate tonnage would respond. This capacity can steer risk quantification and preemptive planning for pest containment.

We have erected a reservoir computation framework that marries statistically derived learning with statistically derived correction. The ESN absorbs intricate multivariate associations and chronological regularities from heterogeneous data streams, capturing how climatic, edaphic, and crop vitality signals interrelate with yield. Superimposed on this, the penalty apparatus infuses specialist awareness of pest herbivory, guaranteeing that the model's output estimates remain faithful to established pest-effect linkages even under circumstances potentially underrepresented in the fitting archive. This flexible formulation is anticipated to deliver more resilient projections, particularly when confronting severe FAW epidemic conditions.

Results and Discussion

The adapted ESN formulation was assessed against the held-out data partition, centering on its estimation fidelity for end-of-cycle harvest under a spectrum of pest burden intensities. Its behaviour was benchmarked against a set of reference constructs: a baseline multiple linear regression drawing exclusively on climatic descriptors (devoid of pest signals), a conventional ESN stripped of the penalty correction (i.e., receiving identical input channels including FAW records but left to internalize pest consequences on its own), and a deep neural architecture fed the same predictor suite. The principal error yardsticks employed comprised Root Mean Square Error (RMSE) and Mean Absolute Error (MAE) for tonnage projection (expressed in t/ha), alongside the coefficient of determination R^2 to assess the fraction of variance accounted for. [Table 2](#) condenses the outcome of these head-to-head comparisons on the evaluation fold.

Table 2. Benchmarking of model behaviour on harvest estimation (test partition). The advocated ESN construct that merges FAW intelligence with a penalty apparatus registers the smallest discrepancy and the highest R^2 .

Model	RMSE (tons/ha)	MAE (tons/ha)	R^2 (–)
Proposed ESN (FAW-aware + penalty)	0.0628	0.0531	0.55
ESN (with FAW input, no penalty)	0.1945	0.1539	–2.12
DNN (with FAW input)	0.3734	0.3432	–10.48
Linear Regression (weather-only)	0.2129	0.1683	–4.14

As evident above, the FAW-cognizant ESN returned the strongest $R^2 = 0.55$, signifying that close to 55 % of the dispersion in observed tonnage was encompassed. This marks a meaningful leap over the weather-restricted linear formulation ($R^2 = -4.14$), which routinely inflated estimates under infestation duress owing to its complete blindness to pest signals. While an R^2 of roughly 0.5 might appear modest when set beside figures reported from tightly constrained laboratory trials, it sits squarely within the band documented by field-scale tonnage estimation research that acknowledges biotic pressures. Contemporary syntheses of agroclimatic and pest-conscious frameworks document R^2 scores spanning from 0.30 to 0.80 across heterogeneous cropping systems, geographies, and modelling strategies [46, 47]. Notably, investigations that overtly encode FAW burden seldom account for more than half of yield dispersion, a consequence of the sizable, intrinsically capricious interplay linking infestation chronology, crop phenophase, and meteorological shocks [48]. Thus, our finding places within the upper quartile of published field investigations and signals that the FAW-sensitive ESN seizes as much of the harvest variation as can realistically be anticipated amidst operational noise and husbandry inconsistency. Even when measured against an ESN that digests pest readings without the overt penalty, the penalized variant registered substantially depressed error (an approximate 68 % shrinkage in RMSE). The deep network model performed marginally below the ESN in our trials (for instance, RMSE of 0.37 versus 0.19 for the bare ESN), plausibly attributable to memorization on the modest corpus combined with its heavier fitting demand. Tellingly, both the linear regression and the deep network faltered in years stamped by acute FAW epidemics. During one season where FAW larval density surged anomalously high, the linear construct overshot tonnage by 0.32 t/ha (failing to sense that roughly one-third of the standing crop would be forfeited), whereas the deep network missed certain non-linear interdependencies and undershot output in a rebound scenario. The unpenalized ESN did absorb some pest influence (its

$R^2 = -2.12$ reveals it internalized a share of the pest–yield coupling), yet it inclined toward underrating output erosion during the most intense epidemic spikes, as it lacked any overt mechanism constraining it to align with documented impairment thresholds. The insertion of the penalty mapping rectified this skew, harmonizing model conduct with agronomic expectations amid epidemic episodes. **Table 3** underneath catalogues the findings obtained.

Table 3. Benchmarking of model behaviour on harvest estimation (test partition). The advocated ESN construct that merges FAW intelligence with a penalty apparatus registers the smallest discrepancy.

Year	CI_upper	CI_lower	Error penalty ESN (%)	Error base ESN (%)	Penalized yield	Predicted yield	Actual yield	FAW level
2014	2.983	2.125	8.10%	8.10%	2.554	2.554	2.78	0
2016	3.003	2.145	3.80%	4.50%	2.574	2.591	2.48	3.1
2017	2.888	2.03	2.40%	2.70%	2.459	2.586	2.518	13.6
2018	2.943	2.085	14.80%	20.00%	2.514	2.627	2.189	24.7
2019	2.91	2.052	4.60%	9.80%	2.481	2.604	2.371	19.4
2020	2.846	1.988	1.70%	4.20%	2.417	2.563	2.459	35.8
2021	2.554	1.696	13.40%	8.50%	2.125	2.247	2.455	43.7
2022	2.57	1.712	10.10%	4.60%	2.141	2.272	2.382	32.1
2023	2.634	1.776	0.80%	6.80%	2.205	2.336	2.188	44.1
2024	2.642	1.784	1.60%	7.40%	2.213	2.339	2.178	22.2

Under restrained FAW circumstances (2016, 2017), both formulations deliver estimates tightly clustered around observed tonnage. This implies that when pest loading stays muted, the penalty mapping enacts negligible correction and the two formulations behave comparably. In elevated FAW contexts, the bare ESN begins to diverge appreciably; for illustration, the parcel subjected to a 44.1 FAW index (2023) yielded an actual harvest of 2.188 t/ha, whereas the baseline ESN returned 2.336 t/ha (a meaningful 6.8% overshoot). Under this scenario, the penalized ESN drew the forecast downward to 2.205 t/ha, shrinking the misfit to a mere 6% through the imposition of a 5.6% yield deduction. This instance showcases how the penalty apparatus sharpens fidelity precisely when pest duress runs high. Although the inferred penalty routinely trims cross-validated misfit, certain excluded years stay inadequately rectified. Three dominant sources account for this. First, our cross-validation splits encompass scarcely more than a dozen cropping cycles, so any single idiosyncratic epidemic signature can dominate the isotonic adjustment, yielding penalties that are either overly steep or too shallow. Second, condensing pest stress into its peak value discards temporal framing. Late-cycle surges inflict less devastation than those striking early in crop ontogeny; hence, identical $P_{\max}$ magnitudes may call for divergent corrections. Third, inaccuracies in the ESN's initial forecast cascade non-linearly. Should the ESN already under-project in high-P seasons, the penalty can counterintuitively inflate the estimate toward the recorded tonnage. Collectively, these constraints illuminate the residual discrepancies where the penalized outlook does not uniformly surpass the raw ESN. The error structure corroborates that fusing agronomic insight via our penalty function yields markedly dependable tonnage projections, most conspicuously under regimes of heightened pest loading.

The residual examination (**Figure 3**) further substantiates the merit of our methodology. Inspecting forecast discrepancies (residuals) plotted against recorded yields uncovers instructive signatures of model conduct. The penalized ESN variant registers considerably shrunken residuals spanning the entire tonnage spectrum; the baseline ESN persistently overestimates output (positive residuals stretching from 0.008 to 0.256 t/ha), whereas the penalized counterpart's residuals stay substantially tighter (0.006 to 0.081 t/ha), signalling superior alignment. When residuals are charted versus

observed harvests, the bare ESN manifests heteroscedasticity—error spread that swells alongside FAW indices, a pattern starkly apparent in the 2022 observation where it overshoot tonnage by 0.256 t/ha (10.7%). By comparison, the penalized ESN upholds more uniform error magnitudes throughout the FAW gradient.




Figure 3. Residual analysis.

The linear regression construct, fed exclusively by climatic descriptors, displays residuals that are larger still and inflate dramatically with mounting FAW duress, consistently projecting tonnage above 2.5 t/ha even as empirical yields receded substantially under pest aggression, since it possesses no pathway for encoding FAW herbivory. A paired t-test on the absolute forecast discrepancies verifies that gains from the penalty function attain statistical significance ($P < 0.05$) specifically within high FAW settings ($FAW > 30$), whereas for subdued FAW conditions (2024 with $FAW = 22.2$), the margin separating the two formulations proved negligible and lacked statistical distinction. Pooled across all test seasons, the penalized ESN shrank absolute misfits by a mean of 64.3% relative to the baseline ESN, with the most sizable enhancement surfacing in 2022 (a high FAW year), where error contracted from 10.7% to 3.4%—a 68.2% decrease. Expressed in tangible agronomic equivalents, slashing the mean absolute error from 0.256 t/ha to 0.081 t/ha (a 68.2% contraction) corresponds to approximately 7.5 tonnes of grain preserved on a prototypical 100-hectare holding. At the prevailing regional farm-gate valuation of USD 240 per tonne, this equates to an averted revenue erosion of USD 4,200 per cropping cycle (maize pricing intelligence is accessible through the FAO Food Price Monitoring Tool [49]). Because cultivators armed with advanced-alert projections can forgo blanket prophylactic spraying, the framework additionally unlocks a commensurate curtailment of insecticide outlays. Drawing on Zambia's mean FAW suppression expenditure of USD 22 per hectare for two treatment rounds, a pinpointed single application confined exclusively to the 30 hectares of model-flagged "critical zones" would shrink agrochemical deployment by roughly 70% while securing an equivalent yield result [50]. This residual dissection ratifies that our penalty function design effectively confronts the structural

overprediction bias plaguing harvest estimation under pest load. Embedding discipline-specific awareness of FAW herbivory into the predictive engine produced a more trustworthy forecasting apparatus that conserves accuracy across a spectrum of infestation milieus. The formulation now appropriately encodes output forfeiture during epidemic intervals without forfeiting dependability in seasons of lighter colonization. This pattern mirrors precisely the behaviour sought for sound pest governance decision assistance. The structure embedded in the residuals plainly illustrates that the penalized ESN furnishes more credible projections, capable of better equipping growers with knowledge about prospective harvest shrinkage attributable to FAW colonization. This, in turn, permits sharper logistical planning and more surgical pest suppression actions. The findings drawn from this investigation signal that wedding algorithmic learning methodologies with agronomic understanding can meaningfully elevate the fidelity of crop output estimation, most especially when biotic stressors such as Fall Armyworm incursions are present. A penalty mapping was embedded within an Echo State Network. This mapping calibrates anticipated output diminution arising from FAW. Through this tactic, estimation fidelity was heightened, above all under acute pest intensity. The proposed technique outstrips foundational machine learning formulations that lean exclusively on climatic predictors [51-54].

Component removal investigation

To rigorously dissect the incremental value of each module within the proposed architecture, we carried out a systematic ablation experiment, sequentially stripping away core elements and documenting the resulting erosion of predictive power. **Table 4** consolidates these observations.

Table 4. Ablation experiment outcomes revealing the marginal contribution of separate components to overall model fidelity.

Feature removed	MAPE	MAE	RMSE	R2
None	4.058598	0.103644	0.153996	0.553264
NIR	8.005574	0.197649	0.230053	0.02534
GPR	12.76491	0.315219	0.375581	−1.597799
pH	24.475578	0.614224	0.668658	−7.233917
nutrients	31.11154	0.768365	0.931807	−14.990073
tmin	17.094873	0.419772	0.531408	−4.200601
tmax	23.012303	0.57043	0.720883	−8.570335
humidity	56.732728	1.382546	2.13682	−83.087954
pressure	18.736607	0.451232	0.681116	−7.543598
rain	15.133409	0.358884	0.558558	−4.74558
cumulative_rain	17.432963	0.423911	0.563527	−4.848263
faw_infestation	13.738203	0.348733	0.42074	−2.260064

The feature-dropping probe brought to light several non-obvious performance levers. Eliminating relative humidity sent R² plunging to −83.09 while error magnitudes ballooned, hinting that atmospheric moisture not only mediates maize transpiration physiology but also decisively shapes FAW egg survival rates [55]. Soil nutritional status (NPK levels) ranked as the second most indispensable block: discarding these measurements drove R² down to −14.99, reaffirming the thoroughly documented role that accessible nutrients play in both crop robustness and pest endurance.

Thermal variables likewise proved critical, with daily maximum temperature standing out. Singling out t_{max} for removal depressed R^2 to -8.57 , a pattern that mirrors FAW's elevated consumption rates during sunlit warmth and the organism's overall heat sensitivity. In marked contrast, withdrawing FAW trap catches and larval density readings produced a comparatively subdued impact ($R^2 = -2.26$). This pattern implies that the network can reconstruct a substantial share of the pest-driven signal secondarily from ambient environmental cues, even before direct infestation measurements are presented.

Taken collectively, these results illuminate the dense ecological web that determines realized productivity under biotic pressure. Humidity and soil nourishment jointly regulate both constitutive plant defences and pest demographic trajectories, while temperature modulates FAW activity rhythms. Appreciating these broader contextual variables may enhance future generations of forecast engines more than simply adding pest count streams.

This investigation introduces an original fusion of pest population behaviour into a harvest forecasting pipeline built around an Echo State Network (ESN). Through the incorporation of Fall Armyworm (FAW) demographic data along with a pest-triggered yield deduction mechanism, the framework attained heightened prediction fidelity during infestation events.

Across our experimental evaluations, the FAW-sensitive ESN reached an R^2 approaching 0.55 on withheld test observations, exceeding simpler benchmark formulations. This gain buttresses the understanding that meteorological extremes, together with biological forces such as pest colonization, constitute principal drivers of yield dispersion and therefore demand explicit representation within forward-looking models [56, 57]. We noted that a purely data-fitted construct (lacking overt pest descriptors) showed a systematic tendency to underrepresent harvest forfeiture whenever infestation severity ventured beyond the calibration envelope—a deficit that our penalty correction directly addresses. Through the fusion of heterogeneous sensor streams (climatic, edaphic, crop condition, and pest metrics), the ESN accounted for substantially more yield dispersion than frameworks resting on any single predictor category. These outcomes corroborate that a broad-spectrum, multi-input modelling stance can materially elevate both harvest estimation and pest surge anticipation. Furthermore, the capacity of the framework to refresh its outputs contemporaneously as fresh pest intelligence flows in means it can operate as a continuously updating early-alert system for impending crop contraction, warning operators without delay when productivity projections dip (for instance, "expected yield has contracted by 34%") as soon as a FAW incursion registers. Such temporally precise outlooks hold immense tactical value, enabling anticipatory countermeasures.

Situating the approach within prevailing pest surveillance paradigms

To locate our contribution amid established practice, a side-by-side comparison with current on-farm pest observation conventions proves illuminating. A broadly adopted foundation for operational decision-making is the Food and Agriculture Organization's Agroecosystem Analysis (AESA) framework, a community-engaged protocol deployed within Farmer Field Schools for systematic crop and pest observation; it is rooted in Agroecology principles [58]. AESA foregrounds structured field appraisals that weave together pest scouting, ecological state evaluation, and agronomic indicators, empowering cultivators and outreach agents to reach evidence-informed husbandry choices. This integrative philosophy fosters recognition of the intricate linkages within the farming landscape, contemplating how herbivores, beneficial organisms, soil condition, weather patterns, and cultivation practices collectively govern harvest trajectories. AESA's potency lies in equipping agriculturalists with an organized yet intuitive scaffold for decoding complex ecosystem behaviour and in nurturing knowledge circulation and group alignment at the communal level. Regularly scheduled AESA convenings support prompt flagging of crop stress and pest challenges, enabling growers to react pre-emptively. The methodology also champions ecologically sound practices by pursuing balanced ecosystem management over indiscriminate pest elimination; farmers are guided toward regulating pest densities with minimal chemical footprints, a philosophy that inherently serves to moderate excessive pesticide reliance [59].

Virtues of AESA

The bedrock strengths of the AESA philosophy encompass its participatory and blending character. Producers conduct routine assessments of plant vitality and developmental tempo (customarily on a weekly rotation), a cadence that aids prompt recognition of stress agents. Pest communities are surveilled in parallel with their naturally occurring opponents, furnishing a layered appreciation of pest dynamics nested within the ecosystem. Agriculturalists are encouraged to perceive cross-linkages—for instance, how below-ground conditions or weather sequences might predispose fields to outbreaks or alter the efficacy of indigenous predators. This expansive, systems-sensitive lens carries particular value for small-scale farming households throughout Sub-Saharan Africa, as it confers the insight necessary for independent field evaluation.

Constraints of AESA

Notwithstanding its benefits, AESA confronts tangible operational hurdles, especially when confronting highly dispersive and explosively breeding invasive organisms like FAW across broad terrain. One substantive bottleneck is the method's drain on labour and clock time. Undertaking meticulous in-field inspections coupled with group deliberation on a weekly rhythm can impose heavy demands on the schedules of growers and advisory personnel alike. During peak agronomic windows, reserving one or two hours per week for AESA exercises may prove burdensome, potentially leading to diminishing session frequency or depth. This labour appetite constrains the feasibility of extending AESA monitoring across wide geographies or of preserving temporal consistency. A related drawback is the technique's inherently bounded spatial resolution. In practice, farmers applying AESA might scrutinize only a sliver of the plant population within a given field to extrapolate pest presence. Such sparse canvassing can miss incipient infestation pockets or localized hotspots in unexamined field sectors, most notably within sizeable or internally variable plots. As a result, colonizations risk being underrated or entirely overlooked until they have escalated to damaging proportions. Moreover, the calibre of AESA intelligence hinges intimately on the skill and steadiness of the observer. Determinations of pest identity, tissue removal magnitude (for example, gauging percent leaf defoliation), and enumerations of beneficial fauna are inherently vulnerable to subjectivity. Less seasoned growers may misdiagnose pest life stages (such as early FAW instars or egg masses) or miscalculate the breadth of crop lesioning. Even among veteran observers, personal drift and exhaustion can inject scatter and bias. Such variability can yield inconsistent records, complicating the task of dependably activating management responses at the ideal moment.

Supplementing AESA through detector-driven technologies

To remedy these shortcomings, weaving modern sensor platforms into pest surveillance represents a promising evolutionary step. Mechanization through Internet of Things (IoT) nodes, uncrewed aerial vehicles, and analogous devices can markedly relieve the labour encumbrance of field observation while simultaneously amplifying its recurrence frequency and spatial sweep. For instance, self-operating lure traps can sustain uninterrupted logs of FAW moth interceptions, obviating the requirement for farmers to manually visit and tally insects on a weekly basis. Similarly, drone-borne imaging facilitates rapid reconnaissance of extensive fields, capturing high-resolution proxies of crop condition (such as the Normalized Difference Vegetation Index (NDVI) and canopy closure) that covary with herbivory stress. These technological aids can flag pest-instigated strain (for example, zones of nascent canopy flagging or tissue removal) ahead of when such damage becomes visually conspicuous to ground-level inspection, thereby serving as a sky-borne advance-warning apparatus. Beyond this, detector webs can substantially elevate the temporal density and spatial granularity of surveillance. Instead of intermittent weekly spot-checks at scattered stations, distributed IoT nodes (for example, a network of mechanized traps or in-field cameras arrayed across the property) can generate continuous, real-time data torrents enveloping the entire holding. This fine-meshed information fabric amplifies the odds of capturing a FAW incursion at its earliest genesis, even if it originates within a confined patch of the landscape. It also ensures that no segment of the field is routinely bypassed, shrinking surveillance voids that could shelter cryptic pest aggregations. A further decisive merit of sensor-centred strategies is the impartiality and uniformity of the captured record.

Model contribution and originality

From a methodological standpoint, the tailored ESN equipped with a pest-contingent penalty function introduces a novel template for merging process-based agronomic understanding with algorithmic learning. Conventional machine learning

constructs for yield estimation operate as opaque "black boxes," deducing associations exclusively from observational records. By contrast, our formulation attains a degree of semi-interpretability: domain knowledge concerning pest herbivory is woven directly into the architecture through the penalty term. Consequently, the model's conduct during an infestation episode can be expressed in agronomic language (for instance, the model trimming projected output by roughly 20% once FAW trap enumerations cross a predefined boundary corresponds closely with specialist expectation for that colonization intensity). Such lucidity proves vital for cultivating confidence in AI-generated guidance among growers and field advisors [60]. When the model's internal logic echoes established agricultural precepts, end-users become more disposed to accept and act upon its alerts. To illustrate, when the system autonomously depresses a tonnage outlook because FAW lure interceptions surpassed a known danger mark, an agronomist can readily interpret this as the model effectively invoking an economic threshold criterion—analogue to what an extension professional would apply. This interpretability, along with consonance with entrenched Integrated Pest Management (IPM) frameworks, renders the model's advisories more persuasive and easier to corroborate under real farming conditions.

Our observations also resonate with a growing body of scholarship indicating that conventional time-series formalisms (such as ARIMA or stationary neural architectures) regularly falter when confronting the non-linear, multi-causal character of sensor-derived agricultural data [61]. More sophisticated designs like ESNs, most notably when augmented with discipline-specific calibrations or blended strategies, prove better matched to these complexities. Notably, a nascent literature is emerging around fusing machine learning with specialist insight for pest projection [62, 63]. Certain recent contributions, for example, deploy learning algorithms alongside economic damage thresholds or multi-sensor amalgamation to anticipate infestation outbreaks [64, 65]. Our investigation stretches this conceptual thread by tethering pest forecasts directly to output estimates, furnishing an instrument that spans the full chain from environmental sensing through to anticipated crop forfeiture. In this respect, it supplies a design pattern for how artificial intelligence may be harnessed not merely to project pest occurrence itself, but to estimate the agronomic penalty exacted by that pest—the metric of ultimate concern to farmers.

Ramifications for precision farming

The engineered AI-driven pipeline exemplifies the manner in which precision agriculture stands to gain from the confluence of IoT and machine learning. By delivering faithful advance estimates of pest-attributable output erosion, both producers and governance bodies can mount anticipatory actions in a temporally appropriate fashion. Should the framework, for instance, project a 25% tonnage contraction in a designated zone owing to a mounting FAW surge, a grower could opt to step in early with site-directed biorational sprays or release natural predators at the model-flagged concentration points, plausibly heading off the brunt of the impairment. Conversely, where the model signals only a negligible harvest repercussion, the cultivator might elect to forgo treatment altogether, side-stepping gratuitous outlay and ecological burden. This selective intervention doctrine lies at the heart of IPM—engaging countermeasures solely when surveillance confirms that the economic calculus warrants doing so.

Constraints

Notwithstanding the encouraging outlook of the proposed methodology, we recognize several bounds circumscribing the current investigation and its accompanying implementation. The penalty mapping encoding output forfeiture from FAW feeding was constructed heuristically, drawing upon published economic damage benchmarks and a confined set of localized field investigations. Pest-inflicted yield depression can assume strongly non-linear contours and hinges on dimensions such as colonization phenology (early vegetative onset versus late reproductive phase) and the intrinsic tolerance of the crop genotype. If the penalty specification proves misaligned, the framework risks either over-attenuating or under-correcting for pest consequences. We anchored it through the accessible corpus on FAW-linked output loss [35, 41, 43]; nonetheless, broader on-farm experimentation would serve to sharpen this mapping. A more empirically guided avenue could be pursued in subsequent work, wherein the model internalizes the pest-to-tonnage linkage directly by conditioning on paired datasets drawn from colonized and protected plots—essentially permitting the network to deduce the yield penalty from manipulative trials. Such a route would presumably encode subtleties more faithfully than the present hand-assembled penalty specification.

The current instantiation focuses on a single pest organism (FAW) affecting a single host crop (maize). In operational reality, cultivated fields are routinely beset by multiple coincident stresses—additional herbivores (stalk borers, locusts, and similar), pathogens, or nutrient shortfalls—each capable of independently modulating final output. In its existing form, the ESN remains insensitive to these concurrent pressures. A vital expansion will involve accommodating multi-pest or multi-stressor settings. One feasible avenue toward this goal could entail incorporating supplementary input conduits for other pests and pathogens and formulating an array of penalty mappings (or else a composite damage aggregator) that jointly rescale the output forecast.

A further substantial obstacle plaguing investigations into the linkage between pest colonization and crop output is the dearth of thorough datasets. Despite the considerable economic toll exacted by pest-driven yield contraction, the scientific record remains sparse in quantitative couplings relating infestation intensity to realized harvest depression. This lacuna is especially acute for transboundary invaders like the fall armyworm across Africa and Asia, where extended monitoring archives are only now beginning to materialize.

The present inquiry into echo state network (ESN)-based harvest estimation brings this data poverty into sharp relief, having been compelled to draw upon synthesized and imputed figures owing to the insufficiency of locally sourced measurements. The preponderance of extant scholarship concentrates on either pest demography or output endpoints in isolation, seldom erecting robust quantitative bridges between the two. This situation underscores the pressing requirement for controlled manipulative field experimentation designed to gauge yield responses along a deliberate gradient of pest loading, employing harmonized protocols that would enable meta-analytical synthesis and cross-site model portability.

Forward agenda and refinements

We are actively exploring composite architectures that wed the ESN's capacity for temporal motif distillation with supplementary machine learning paradigms. One avenue under consideration, for example, involves channelling the reservoir activation vector into an auxiliary regression tier (such as an Elastic Net or another regularized sparse linear formulation) jointly with static covariates like soil taxonomy or husbandry regime. This arrangement would effectively unite the non-linear sequential modelling prowess of the ESN with a transparent feature-weighting apparatus, potentially yielding gains in both interpretability and predictor selection. Such an assembled framework could illuminate which input variables or sensor modalities hold the greatest predictive leverage over both tonnage and pest hazard, thereby delivering auxiliary diagnostic value—for instance, quantifying whether precipitation volatility or pest abundance exerted the larger sway on yield projections within a given cropping cycle.

Deep reservoir computing

A further avenue involves investigating deep or hierarchically arranged reservoir computing (frequently termed Deep ESN) configurations. Agronomic chronosequences commonly exhibit tiered structure—brief, within-week oscillations layered beneath protracted seasonal rhythms. A multi-tiered ESN (cascading several reservoirs or deploying a "reservoir of reservoirs") could concurrently apprehend patterns operating at distinct temporal scales. Earlier contributions have demonstrated that deep ESNs can render intricate spatio-temporal data with greater fidelity. Within our setting, a deep ESN might, for example, encode intra-seasonal dynamics in one tier while absorbing inter-annual shifts (such as the footprint of an El Niño episode) in another. This could further sharpen precision, particularly when projecting across extended horizons or adapting to year-on-year realignments in pest population baselines.

Enriched remote sensing assimilation

As remote observation products become accessible at finer spatial granularity and higher revisit cadences, these can be woven into the framework to refresh estimates more frequently and with improved exactness. By way of illustration, daily orbital captures or even quasi-real-time drone surveillance could stream into the ESN, enabling it to monitor shifts in crop condition throughout the growing cycle rather than only at terminal harvest. Such a capability would permit within-season

tracking of both yield trajectory and pest burden. Should a moisture deficit episode or pest-driven defoliation event materialize, the model could dynamically revise productivity expectations. High-resolution feeds would likewise assist in generating location-specific advisories (down to sub-parcel zones), capturing the very essence of precision farming.

Transferability to alternative crops and pest complexes

We intend to probe the portability of our methodology to additional host species (such as sorghum or rice) and distinct pest organisms. FAW has by now extended its range into Asia and may shortly pose a threat to southern Europe; accordingly, applying the construct across these newly colonized territories will represent a crucial examination of its generalizability. Every fresh context is liable to introduce divergent interactions—contrasting climatic regimes, the presence of indigenous natural antagonists, or dissimilar cropping architectures. The framework, most notably the penalty mapping together with any embedded presuppositions about pest population behaviour, may demand recalibration for such scenarios. Broadening the model to entertain multiple co-occurring herbivore species attacking the crop (as flagged among the constraints) likewise figures on the development agenda. The ultimate aspiration is a resilient system wherein appending a novel pest data conduit and its established impact parameters would equip the model to jointly estimate yield forfeiture arising from the entire pest assemblage present.

Partnership with disciplinary specialists

Subsequent investigations stand to gain substantially from tighter engagement with agronomists and entomologists. Their domain acumen can steer the formulation of richer representations of pest herbivory processes inside the model. One illustration is the embedding of pest phenology simulators or crop developmental stage vulnerability functions. In a parallel manner, as manipulative field experimentation progresses, specialist feedback can assist in diagnosing model discrepancies and proposing refinements. Conceivably, the construct may have overlooked a late-cycle crop rebound, or may have exaggerated impairment under particular circumstances—insights that disciplinary experts are uniquely positioned to supply.

Conclusion

The fusion of IoT sensing with AI-driven modelling showcased in this undertaking epitomizes a plausible trajectory toward more intelligent and ecologically sound agriculture. In place of adhering to calendar-driven or uniformly broadcast pest suppression, cultivators could harness statistically derived intelligence to govern pests in an environmentally considerate, spatially exact fashion. The FAW case stands as emblematic of a twenty-first-century agricultural predicament: an invasive organism that defies uniform management prescriptions and demands perpetual vigilance coupled with adaptivity. Our ESN-grounded predictive architecture, endowed with the capacity to absorb heterogeneous data conduits and impose pest-sensitive recalibrations, illustrates the promise that technology holds for confronting such biological realities at the farm level.

For operational deployment, we conceive of this model being housed within a handheld application accessible to growers through smartphones, issuing automated cautions whenever pest metrics breach critical demarcations. The requisite physical infrastructure would encompass a web of inexpensive pheromone lure stations fitted with IoT connectivity, positioned strategically at a density of one trap per five hectares, supplemented by intermittent drone sorties for visual corroboration of colonization patterns. Monthly orbital imagery combined with readings from localized meteorological stations would deliver the ancillary environmental inputs demanded by the framework. Compact computing boards such as the Raspberry Pi, built on a standard Von Neumann architecture, possess the processing headroom to execute the model and ingest multivariate sensor feeds [66]. These edge computation units can be augmented by more substantial processing capacity provisioned through cloud infrastructure.

Nonetheless, several scalability hurdles would require resolution before widespread uptake becomes feasible. Computational loads, while restrained for the ESN proper (which can execute on edge hardware), would mount when handling satellite imagery at landscape extents. Sensor interoperability and harmonization across disparate agricultural

regions would necessitate coordination among equipment fabricators and agricultural advisory bodies. Sparse telecommunications coverage across many rural settings constitutes an additional obstacle, potentially demanding off-grid operability with intermittent synchronization features. The expanding footprint of low-Earth orbit (LEO) satellite communication constellations may, however, deliver pervasive connectivity within a handful of years. Since 2019, SpaceX has deployed Starlink [67], OneWeb operates its own LEO constellation, and Amazon similarly envisions a worldwide satellite network [68]. These developments could suppress connectivity expenditure, rendering it economically viable to field IoT sensor webs together with edge AI and ML capabilities. Additionally, the constructed model would require retuning for distinct maize genotypes and local environmental envelopes beyond the boundaries of our study domain.

The present investigation drew upon a concise multi-season archive and an isotonic regression-fitted penalty function that encodes Fall Armyworm herbivory exclusively. Performance may consequently erode under seasons, geographies, or cropping species that deviate substantially from our calibration corpus. Forthcoming efforts will broaden the repository with supplementary annual cycles and satellite-derived indices, infer the pest-damage relationship directly from manipulative field experimentation, and stretch the architecture to accommodate multiple pests along with within-field spatial heterogeneity. These measures will extend the framework's validity envelope and facilitate more expansive adoption, moving us nearer to functional, deployment-ready instantiations capable of tangibly shaping sustainable pest husbandry.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 30 Jul 2022

Revised: 08 Sep 2022

Accepted: 31 Oct 2022

Published online: 15 March 2023

Rights and permissions

Open Access The author(s) retain copyright. This article is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License. It may be shared and adapted for non-commercial purposes with appropriate attribution, an indication of changes, and distribution of adaptations under the same license. Third-party material may be subject to separate terms identified in its credit line. View the license at <https://creativecommons.org/licenses/by-nc-sa/4.0/>.

References

Domingues T, Brandão T, Ferreira JC. Machine learning for detection and prediction of crop diseases and pests: a comprehensive survey. *Agriculture*. 2022;12(9):1350.
<https://doi.org/10.3390/agriculture12091350>.

Day R, Abrahams P, Bateman M, Beale T, Clotey V, Cock M, et al. Fall Armyworm: impacts and implications for Africa. *Outlooks Pest Manag*. 2017;28(5):196-201.
https://doi.org/10.1564/v28_oct_02.

Balla A, Bhaskar M, Bagade P, Rawal N. Yield losses in maize (*Zea mays*) due to fall armyworm infestation and potential IoT-based interventions for its control. *J Entomol Zool Stud*. 2019;7(5):920-7.

Timilsena BP, Niassy S, Kimathi E, Abdel-Rahman EM, Seidl-Adams I, Wamalwa M, et al. Potential distribution of fall armyworm in Africa and beyond, considering climate change and irrigation patterns. *Sci Rep*. 2022;12(1):539.
<https://doi.org/10.1038/s41598-021-04369-3>.

Kumar RM, Gadratagi BG, Paramesh V, Kumar P, Madivalar Y, Naveena N, et al. Sustainable management of invasive fall armyworm, *Spodoptera frugiperda*. *Agronomy*. 2022;12(9):2150.
<https://doi.org/10.3390/agronomy12092150>.

Kansiime MK, Rwomushana I, Mugambi I. Fall armyworm invasion in Sub-Saharan Africa and impacts on community sustainability in the wake of Coronavirus Disease 2019: reviewing the evidence. *Curr Opin Environ Sustain*. 2023;62:101279.
<https://doi.org/10.1016/j.cosust.2023.101279>.

Aziz D, Sajid Z, Zia MA, Shahid M, Ashraf M, Ali A, et al. Remote sensing and artificial intelligence: revolutionizing pest management in agriculture. *Front Sustain Food Syst*. 2025;9:1551460.
<https://doi.org/10.3389/fsufs.2025.1551460>.

Buchailot ML, Cairns J, Hamadziripi E, Wilson K, Hughes D, Chelal J, et al. Regional monitoring of fall armyworm (FAW) using early warning systems. *Remote Sens*. 2022;14(19):5003.
<https://doi.org/10.3390/rs14195003>.

Wadhwa D, Malik K. A generalizable and interpretable model for early warning of pest-induced crop diseases using environmental data. *Comput Electron Agric*. 2024;227:109472.
<https://doi.org/10.1016/j.compag.2024.109472>.

Brown ME, Mugo S, Petersen S, Klauser D. Designing a pest and disease outbreak warning system for farmers, agronomists and agricultural input distributors in East Africa. *Insects*. 2022;13(3):232.
<https://doi.org/10.3390/insects13030232>.

Fernandes de Oliveira R, Cardoso IM, Carole Muggler C, de Jesus Pereira A, do Carmo DL. Agroecological pest and disease control: the result of action research in agrarian reform settlement. *Agroecol Sustain Food Syst*. 2022;46(2):165-80.
<https://doi.org/10.1080/21683565.2021.1983741>.

Shaheb MR, Sarker A, Shearer SA. Precision Agriculture for Sustainable Soil and Crop Management. In: *Soil Science - Emerging Technologies, Global Perspectives and Applications*. IntechOpen; 2022.
<https://doi.org/10.5772/intechopen.101759>.

Chisonga C, Chipabika G, Sohathi PH, Harrison RD. Understanding the impact of fall armyworm (*Spodoptera frugiperda* J. E. Smith) leaf damage on maize yields. *PLoS One*. 2023;18(6):e0279138.
<https://doi.org/10.1371/journal.pone.0279138>.

Togola A, Meseka S, Menkir A, Badu-Apraku B, Boukar O, Tamò M, et al. Fall armyworm (*Spodoptera frugiperda*) in Africa: insights into biology, ecology and impact on staple crops, food systems and management approaches. *Front Agron*. 2025;7:1538198.
<https://doi.org/10.3389/fagro.2025.1538198>.

De Lange ES, Xiao Y, Beukeboom LW. Biology of fall armyworm – an introduction. *Entomol Exp Appl.* 2023;171(7):488-91.
<https://doi.org/10.1111/eea.13324>.

Kiki MPAF, Ahouandjinou SAR, Assogba KM. Towards a smart farming for early detection of fall armyworm pest. In: *Proceedings of the 2019 International Conference on Smart Applications, Communications and Networking (SmartNets)*. Sharm El Sheikh, Egypt: IEEE; 2019. p. 1-5.
<https://doi.org/10.1109/SmartNets48225.2019.9069756>.

Kiki MPAF, Ahouandjinou SAR, Assogba KM, Sama YN. An E-nose using metal oxide semiconductor sensors array to recognize the odors of fall armyworm pest for its early detection in the farm. In: Mambo AD, Gueye A, Bassioni G, editors. *Innovations and Interdisciplinary Solutions for Underserved Areas*. Vol. 449. Cham: Springer Nature Switzerland; 2022. p. 69-81. (Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering).
https://doi.org/10.1007/978-3-031-23116-2_5.

Krishnamurthy V, Rojas CR. Slow Convergence of Interacting Kalman Filters in Word-of-Mouth Social Learning. Urbana, IL, USA: IEEE; 2024. p. 1-6.
<https://doi.org/10.1109/Allerton63246.2024.10735278>.

Shahi S, Fenton FH, Cherry EM. Prediction of chaotic time series using recurrent neural networks and reservoir computing techniques: a comparative study. *Mach Learn Appl.* 2022;8:100300.
<https://doi.org/10.1016/j.mlwa.2022.100300>.

McDermott PL, Wikle CK. Deep echo state networks with uncertainty quantification for spatio-temporal forecasting. *Environmetrics.* 2019;30(3):e2553.
<https://doi.org/10.1002/env.2553>.

Steiner P, Jalalvand A, Birkholz P. Exploring unsupervised pre-training for echo state networks. *Neural Comput Appl.* 2023;35(34):24225-42.
<https://doi.org/10.1007/s00521-023-08988-x>.

Joshi A, Pradhan B, Chakraborty S, Varatharajoo R, Gite S, Alamri A. Deep-transfer-learning strategies for crop yield prediction using climate records and satellite image time-series data. *Remote Sens.* 2024;16(24):4804.
<https://doi.org/10.3390/rs16244804>.

Farmonov N, Esmaeili M, Abbasi-Moghadam D, Sharifi A, Amankulova K, Mucsi L. HypsLiDNet: 3-D–2-D CNN model and spatial–spectral morphological attention for crop classification with DESIS and LiDAR data. *IEEE J Sel Top Appl Earth Obs Remote Sens.* 2024;17:11969-96.
<https://doi.org/10.1109/JSTARS.2024.3418854>.

Esmaeili M, Abbasi-Moghadam D, Sharifi A, Tariq A, Li Q. ResMorCNN model: hyperspectral images classification using residual-injection morphological features and 3DCNN layers. *IEEE J Sel Top Appl Earth Obs Remote Sens.* 2024;17:219-43.
<https://doi.org/10.1109/JSTARS.2023.3328389>.

Qiao M, He X, Cheng X, Li P, Luo H, Zhang L, et al. Crop yield prediction from multi-spectral, multi-temporal remotely sensed imagery using recurrent 3D convolutional neural networks. *Int J Appl Earth Obs Geoinf.* 2021;102:102436.
<https://doi.org/10.1016/j.jag.2021.102436>.

Marzvan S, Moravej K, Felegari S, Sharifi A, Askari MS. Risk assessment of alien *Azolla filiculoides* Lam in Anzali Lagoon using remote sensing imagery. *J Indian Soc Remote Sens.* 2021;49(8):1801-9.
<https://doi.org/10.1007/s12524-021-01362-1>.

Chakravarty S, Paikaray BK, Mishra R, Dash S. Hyperspectral image classification using spectral angle mapper. In: Proceedings of the 2021 IEEE International Women in Engineering (WIE) Conference on Electrical and Computer Engineering (WIECON-ECE). Dhaka, Bangladesh: IEEE; 2021. p. 87-90.

<https://doi.org/10.1109/WIECON-ECE54711.2021.9829585>.

Felegari S, Sharifi A, Khosravi M, Sabanov S. Using experimental models and multitemporal Landsat-9 images for cadmium concentration mapping. *IEEE Geosci Remote Sens Lett.* 2023;20:1-4.

<https://doi.org/10.1109/LGRS.2023.3291019>.

Mirhoseini Nejad SM, Abbasi-Moghadam D, Sharifi A. ConvLSTM-ViT: a deep neural network for crop yield prediction using earth observations and remotely sensed data. *IEEE J Sel Top Appl Earth Obs Remote Sens.* 2024;17:17489-502.

<https://doi.org/10.1109/JSTARS.2024.3464411>.

Khaki S, Wang L, Archontoulis SV. A CNN-RNN framework for crop yield prediction. *Front Plant Sci.* 2020;10:1750.

<https://doi.org/10.3389/fpls.2019.01750>.

Khaki S, Pham H, Wang L. Simultaneous corn and soybean yield prediction from remote sensing data using deep transfer learning. *Sci Rep.* 2021;11(1):11132.

<https://doi.org/10.1038/s41598-021-89779-z>.

Mahdipour H, Sharifi A, Sookhak M, Medrano CR. Ultrafusion: optimal fuzzy fusion in land-cover segmentation using multiple panchromatic satellite images. *IEEE J Sel Top Appl Earth Obs Remote Sens.* 2024;17:5721-33.

<https://doi.org/10.1109/JSTARS.2024.3360648>.

Sharifi A, Safari MM. Enhancing the spatial resolution of Sentinel-2 images through super-resolution using transformer-based deep-learning models. *IEEE J Sel Top Appl Earth Obs Remote Sens.* 2025;18:4805-20.

<https://doi.org/10.1109/JSTARS.2025.3526260>.

García-Sánchez F, Galvez-Sola L, Martínez-Nicolás JJ, Muelas-Domingo R, Nieves M. Using near-infrared spectroscopy in agricultural systems. In: Kyprianidis KG, Skvaril J, editors. *Developments in Near-Infrared Spectroscopy*. InTech; 2017.

<https://doi.org/10.5772/67236>.

Buchailot ML, Cairns J, Hamadziripi E, Wilson K, Hughes D, Chelal J, et al. Multi-Scale Remote Sensing for Fall Armyworm Monitoring and Early Warning Systems. In: *IGARSS 2020. USA: IEEE; 2020. p. 4886-9.*

<https://doi.org/10.1109/IGARSS39084.2020.9323181>.

Buchailot ML, Cairns J, Hamadziripi E, Wilson K, Hughes D, Chelal J, et al. Regional monitoring of fall armyworm (FAW) using early warning systems. *Remote Sens.* 2022;14(19):5003.

Sun C, Song M, Cai D, Zhang B, Hong S, Li H. A systematic review of echo state networks from design to application. *IEEE Trans Artif Intell.* 2024;5(1):23-37.

<https://doi.org/10.1109/TAI.2022.3225780>.

Feng GL, Yang J, Wang Q, Zhang W, Sun D, Zhao J, et al. Improved prediction model for flood-season rainfall based on a nonlinear dynamics-statistic combined method. *Chaos Solit Fractals.* 2020;140:110160.

<https://doi.org/10.1016/j.chaos.2020.110160>.

Bürkner PC, Gabry J, Vehtari A. Efficient leave-one-out cross-validation for Bayesian non-factorized normal and Student-t models. *Comput Stat.* 2021;36(2):1243-61.

<https://doi.org/10.1007/s00180-020-01045-4>.

Midega C, Hadi B, McGuire S, Zaidi Z, Rajagopalan P. Fall Armyworm: Measuring Damage and Loss Caused by a Novel Invasive Pest as a Guide to Sustainable Management. *Food and Agriculture Organization.* p. 23-38.

Yaméogo IS, Ouattara D, Ouattara S, Traoré F, Sanou A, Ouédraogo I, et al. Level of damage and yield losses associated with the fall armyworm (*Spodoptera frugiperda*) on maize (*Zea mays*), millet (*Pennisetum glaucum*) and sorghum (*Sorghum bicolor*) on station in Burkina Faso. *Crop Prot.* 2024;182:106743.

<https://doi.org/10.1016/j.cropro.2024.106743>.

Picanço Filho MC, da Silva RS, de Oliveira IR, Pereira EJG, de Oliveira Neto SS, de Oliveira AC, et al. Economic injury levels and economic thresholds for *Leucoptera coffeella* as a function of insecticide application technology in organic and conventional coffee (*Coffea arabica*) farms. *Plants.* 2024;13(5):585.

<https://doi.org/10.3390/plants13050585>.

Usman MS, Adamu RS, Onu I, Kogi E, Musa N, Ahmed IA. Economic injury levels of fall armyworm (*Spodoptera frugiperda* Smith) at early-whorl vegetative growth stage of quality protein maize variety. *J Agric Environ.* 2024;20(1):253-64.

Pereira CJ, Hellman JL. Economic injury levels for *Spodoptera frugiperda* (Lepidoptera: Noctuidae) on silage corn in Maryland. *J Econ Entomol.* 1993;86(4):1266-70.

Makgoba MC, Tshikhudo PP, Nnzeru LR, Makhado RA. Impact of fall armyworm (*Spodoptera frugiperda*) (J.E. Smith) on small-scale maize farmers and its control strategies in the Limpopo province, South Africa. *Jamba J Disaster Risk Stud.* 2021;13(1).

<https://doi.org/10.4102/jamba.v13i1.1016>.

Halder M, Datta A, Siam MKH, Mahmud S, Sarkar MS, Rana MM. A systematic review on crop yield prediction using machine learning. In: Nguyen TDL, Verdú E, Le AN, Ganzha M, editors. *Intelligent Systems and Networks*. Vol. 752. Singapore: Springer Nature Singapore; 2023. p. 658-67. (Lecture Notes in Networks and Systems).

https://doi.org/10.1007/978-981-99-4725-6_77.

Cedric LS, Adoni WYH, Aworka R, Zoueu JT, Mutombo FK, Krichen M, et al. Crops yield prediction based on machine learning models: case of West African countries. *Smart Agric Technol.* 2022;2:100049.

<https://doi.org/10.1016/j.atech.2022.100049>.

Chisonga C, Chipabika G, Sohathi PH, Harrison RD. Understanding the impact of fall armyworm (*Spodoptera frugiperda* J. E. Smith) leaf damage on maize yields. *Ecology.* 2022.

<https://doi.org/10.1101/2022.12.01.518792>.

Consoli S, Egas Yerovi JJ, Machiorlatti M, Morales Opazo C. Real-time monitoring of food price policy interventions during the first two years of COVID-19. *Food Policy.* 2023;115:102405.

<https://doi.org/10.1016/j.foodpol.2023.102405>.

Chiwamba SH, Phiri J, Nkunika POY, Sikasote C, Kabemba MM, Moonga MN. Automated fall armyworm (*Spodoptera frugiperda*, J.E. Smith) pheromone trap based on machine learning. *J Comput Sci.* 2019;15(12):1759-79.

<https://doi.org/10.3844/jcssp.2019.1759.1779>.

Shahhosseini M, Hu G, Huber I, Archontoulis SV. Coupling machine learning and crop modeling improves crop yield prediction in the US Corn Belt. *Sci Rep.* 2021;11(1):1606.

<https://doi.org/10.1038/s41598-020-80820-1>.

Pham HT, Awange J, Kuhn M, Nguyen BV, Bui LK. Enhancing crop yield prediction utilizing machine learning on satellite-based vegetation health indices. *Sensors.* 2022;22(3):719.

<https://doi.org/10.3390/s22030719>.

Morales A, Villalobos FJ. Using machine learning for crop yield prediction in the past or the future. *Front Plant Sci.* 2023;14:1128388.

<https://doi.org/10.3389/fpls.2023.1128388>.

Aslan MF, Sabanci K, Aslan B. Artificial Intelligence Techniques in Crop Yield Estimation Based on Sentinel-2 Data: A Comprehensive Survey. *Sustainability.* 2024;16(18):8277.

<https://doi.org/10.3390/su16188277>.

He L, Zhao S, Ali A, Ge S, Wu K. Ambient humidity affects development, survival, and reproduction of the invasive fall armyworm, *Spodoptera frugiperda* (Lepidoptera: Noctuidae), in China. *J Econ Entomol*. 2021;114(3):1145-58.

Ray DK, Gerber JS, MacDonald GK, West PC. Climate variation explains a third of global crop yield variability. *Nat Commun*. 2015;6(1):5989.
<https://doi.org/10.1038/ncomms6989>.

Savary S, Willocquet L, Pethybridge SJ, Esker P, McRoberts N, Nelson A. The global burden of pathogens and pests on major food crops. *Nat Ecol Evol*. 2019;3(3):430-9.
<https://doi.org/10.1038/s41559-018-0793-y>.

Diane R, Francis C, editors. *Agroecosystems analysis*. In: *Agronomy Monographs*. Madison, WI, USA: American Society of Agronomy, Crop Science Society of America, Soil Science Society of America; 2004.
<https://doi.org/10.2134/agronmonogr43>.

Altieri MA, Nicholls CI, Dinelli G, Negri L. Towards an agroecological approach to crop health: reducing pest incidence through synergies between plant diversity and soil microbial ecology. *npj Sustain Agric*. 2024;2(1):6.
<https://doi.org/10.1038/s44264-024-00016-2>.

Gardezi M, Joshi S, Rizzo DM, Ryan M, Prutzer E, Brugler S, et al. Artificial intelligence in farming: challenges and opportunities for building trust. *Agron J*. 2024;116(3):1217-28.
<https://doi.org/10.1002/agj2.21353>.

Ma Q, Shen L, Chen E, Tian S, Wang J, Cottrell GW. WALKING WALKing walking: action recognition from Action Echoes. In: *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence*. Melbourne, Australia: International Joint Conferences on Artificial Intelligence Organization; 2017. p. 2457-63.

Von Bloh M, Lobell D, Asseng S. Knowledge informed hybrid machine learning in agricultural yield prediction. *Comput Electron Agric*. 2024;227:109606.
<https://doi.org/10.1016/j.compag.2024.109606>.

Schramowski P, Stammer W, Teso S, Brugger A, Herbert F, Shao X, et al. Making deep neural networks right for the right scientific reasons by interacting with their explanations. *Nat Mach Intell*. 2020;2(8):476-86.
<https://doi.org/10.1038/s42256-020-0212-3>.

Madakkattel I, Zhou A, McDonnell MD, Hyppönen E. Combining machine learning and conventional statistical approaches for risk factor discovery in a large cohort study. *Sci Rep*. 2021;11(1):22997.
<https://doi.org/10.1038/s41598-021-02476-9>.

Upadhyay A, Patel A, Patel A, Chandel NS, Chakraborty SK, Bhalekar DG. Leveraging AI and ML in precision farming for pest and disease management: benefits, challenges, and future prospects. In: Jatav HS, Raiput VD, Minkina T, editors. *Ecologically Mediated Development Sustainable Development and Biodiversity*. Vol. 41. Singapore: Springer Nature Singapore; 2025. p. 511-28.
https://doi.org/10.1007/978-981-96-2413-3_23.

Johnston S, Cox S. The Raspberry Pi: a technology disrupter, and the enabler of dreams. *Electronics*. 2017;6(3):51.
<https://doi.org/10.3390/electronics6030051>.

Mohan N, Jaiswal A, Kumar A, Sharma P, Singh R, Gupta A, et al. A multifaceted look at Starlink performance. In: *Proceedings of the ACM Web Conference 2024*. Singapore: ACM; 2024. p. 2723-34.
<https://doi.org/10.1145/3589334.3645328>.

Osoro OB, Oughton EJ. A techno-economic framework for satellite networks applied to low earth orbit constellations: assessing Starlink, OneWeb and Kuiper. *IEEE Access*. 2021;9:141611-25.

<https://doi.org/10.1109/ACCESS.2021.3119634>.