Anomaly Detection in Crop Price: A Review of Machine Learning and Deep Learning, Versus the Emerging Need for Neuro-Symbolic Approaches

Grace Ojochenemi Emmanuelanorue[1] , Sadiya Mohammed Dala[2] , ADEWOYE Munir Adeniyi[3] , JIMOH Abdulrasheed[4]

[1]Computer Science Department, Federal College of Education (Tech.) Gombe

[2]MIS Directorate, Federal College of Education (Tech.) Gombe

[3,4]Computer Science Department, Federal College of Education (Tech.) Gombe

DOI: https://doi.org/10.47772/IJRISS.2026.100900675

Received: 06 September 2026; Accepted: 11 September 2026; Published: 20 October 2026

ABSTRACT

Systems which support the detection of sudden shifts in crop prices in a deep learning fashion often fail to provide explanations for the alerts they provide, which is important for the farmers, extension officers and policy makers to whom they must be directed to act. This paper is a narrative review and a conceptual proposal. The review combines 27 studies that use ML and DL for agricultural commodity price forecasting and anomaly detection (2022-2025) with 24 studies that apply NeSy to anomaly detection in cybersecurity, healthcare, manufacturing and related areas (but not in agriculture). Based on the second body of evidence, it suggests a NeSy architecture for crop price anomaly detection that integrates a domain knowledge graph (DKG) that contains information on both supply and demand events and neural time-series model that can detect and explain the anomaly in terms more easily understood by a non-specialist. The figures are presented without any comparisons between the studies and as published by the original authors. Knowledge-graph design, rule base and data sources and validation plan are described in sufficient detail to be implemented, but the architecture is not built or tested. We believe the marriage of pattern recognition with symbolic reasoning is a viable path to accurate, explainable and usable price anomaly systems and we outline the empirical studies yet to be carried out to demonstrate this.

keywords — Neuro-symbolic AI, price anomaly detection, explainable AI, agricultural economics, interpretable machine learning, real-time monitoring, crop economics, Nigeria, food security, market volatility

INTRODUCTION

Agricultural commodity markets should be stable to ensure food security and overall economic resilience. These markets, however, are extremely volatile. Various factors influence the price dynamics, including weather conditions, political factors, transportation expenses, and the complexity of the supply chain from farm to consumer (Choruma et al., 2024; Theofilou et al., 2025). It is therefore essential to first recognise price trends that are out of line with historical trends in order to take appropriate action in time. When prices move away from what they usually do, it can be a sign of trouble ahead, whether it's because of a shortage, a scam, or simply a sign of economic instability.If prices are moving unusually then it's a sign of trouble ahead and may be a sign of a shortage, scam or just economic stress, and this early warning can help you respond properly.

As machines have become more intelligent, so have the analytical tools that detect such price changes. Early approaches relied on classical statistical methods such as ARIMA and seasonal-trend decomposition (Paul et al., 2025; Ranaweera et al., 2023). These methods are intuitive and easy to interpret, but they cannot process

current agricultural data, which is nonlinear, high-dimensional, and multi-source (Manogna et al., 2025). These approaches were enhanced by machine learning techniques that learned nonparametric patterns of anomalies directly from the data, such as Isolation Forest, Local Outlier Factor, and support vector machines (Sari et al., 2024). Recently, deep-learning architectures such as LSTM networks and Transformers have been developed to handle complex tasks of forecasting and anomaly detection in multivariate time-series, leveraging long-range dependencies (Sediyono et al., 2025; Wen et al., 2022). Some simpler models cannot capture subtle, context-dependent anomalies, which is another advantage offered by hybrid LSTM-autoencoder models (Lee et al., 2024).

As great as these advances have been, there is still a major constraint. Most machine learning and deep learning approaches work like a black box, they can detect statistical deviations, but can not give any explanation on their causes (Garcez & Lamb, 2023; Kautz, 2022). That is, an LSTM model can be used to recognize a spike in the data stream caused by a data entry mistake, by a shortage of data, or when a new trade policy is implemented, but a simple sub-symbolic model might not be able to do so. This lack of interpretability is one of the key challenges for decision makers who might need more than just alerts; they require defensible explanations as well (Atzmueller et al., 2024).

In this paper, the authors argue that a hybrid approach combining explicit symbolic reasoning with data-driven learning, called neuro-symbolic AI (NeSy), could be a viable solution to this interpretability challenge. Neuro-symbolic systems (NSS) integrate the pattern-recognition capabilities of neural networks with the transparent, rule-based reasoning of symbolic AI (W. Wang et al., 2025; Liang et al., 2025; Garcez & Lamb, 2023). These systems aim to detect domain knowledge, including agronomic principles, market rules, and causal relationships, while also providing accurate, explainable results (Ali & Tatheer Fatima, 2025; De Arriba-Pérez et al., 2024). For instance, a price rise may be considered anomalous only if it contradicts what is known about seasonal production, a judgment that requires a system that is both neuro and symbolic.

The review aspect aims to synthesize published ML and DL research on crop price anomaly detection and neuro-symbolic studies from other domains. The perspective component suggests a neuro-symbolic approach to agricultural price monitoring but does not put it into practice. Published evidence is cited to support claims, and the author's reasoning is explicitly identified for each claim.

Scope and Contribution

Some earlier reviews have broadly examined the topic of machine learning applications in agricultural forecasting, with the crop yield prediction review by Jabed & Azmi Murad (2024) and the systematic review of staple crop price prediction by Theofilou et al. (2025). None of them, however, differentiates between anomaly detection and forecasting, nor does any explore the neuro-symbolic approach, which, as shown in this survey, were not yet used in the agricultural market. This paper has two main contributions. It starts by analysing the literature on ML/DL related to crop price anomaly detection, and categorising it to clarify which papers are comparable. Second, it brings in evidence from other areas where NeSyAI has been applied and discusses the rationale and feasibility of using this paradigm for monitoring agricultural prices, especially in markets with specific characteristics in Nigeria and other markets in Sub Saharan Africa, such as volatile exchange rates, limited historical price data, fuel subsidy shocks and informal trading.

REVIEW METHODOLOGY

This review combines a structured literature search with narrative synthesis. It is not a PRISMA-style systematic review: there was no registered protocol, no independent dual screening or formal risk-of-bias assessment, and coverage should not be assumed exhaustive.

Search Strategy

The literature search was conducted across the following databases using combinations of the following keywords: crop price, commodity price, anomaly detection, price forecasting, machine learning, deep

learning, neuro-symbolic, neurosymbolic, and knowledge graph, with "explainable" included for the NeSy strand. The publication period was restricted to 2022–2025 to capture the latest research, and prior approaches (ARIMA, GARCH, classical outlier detection) were included only when used as a baseline within the selected studies. The retrieved reviews were then carefully checked in their reference lists for additional relevant studies.

Screening

Records were screened in two stages by the authors, first by title and abstract and then by full text, with duplicates removed at the first stage. Each full-text candidate was assigned to the core price set, the prior-review set, or the NeSy set, or excluded. Disagreements were resolved by discussion.

Inclusion and Exclusion.

Studies that applied a computational method to predict, or find anomalies in, agricultural commodity prices, and that resulted in a method and outcome were included in the core set (Table 1). Studies related to other agricultural tasks (crop yield, disease risk, produce classification, anomaly detection in sensor data) were not included where there was no price component, as data, tasks and metrics are not commensurable with price work. No previous reviews are included as context. Table 2 does not contain any papers that apply NeSy to agricultural prices (only the NeSy strand), and we included papers from any domain that used a NeSy architecture for anomaly detection, diagnosis or a closely related task, but not purely theoretical NeSy papers, except where relevant in the narrative.

Data Extraction

For each included study we recorded the citation, the algorithm or tool used, the commodity or application domain, the reported performance metric(s) as stated by the authors, and the research gap the authors themselves identified. No metric was recomputed or converted.

Quality and Comparability

The methodological quality of individual studies was not formally assessed, and reported figures were not verified against source data. The F1, precision, recall, and accuracy values in Tables 1 and 2 are presented as published, based on each study's dataset, commodity, and train/test split. Table 3 is limited to qualitative trade-offs.

Machine Learning and Deep Learning Approaches To Crop Price Anomaly Detection

When using crop price data, there is no universal best ML approach; it is mainly determined by the quantity of labeled data. When labeled histories are abundant, supervised methods like random forest, gradient boosting, and support vector machines are preferred; when market conditions change more rapidly than labels can be acquired, clustering and one-class methods, such as k-means, DBSCAN, Isolation Forest, and One-Class SVM, are preferred (Sari et al., 2024; Sediyono et al., 2025). However, if the number of anomaly records is limited and computational resources are constrained, it is practical to use a smaller labeled set and a larger unlabeled set, i.e., semi-supervised designs. Yet all of these practices have one problem: In agriculture, price data is often incomplete, fragmented, and poorly balanced, which degrades the reliability of the tools just when they are most needed.

In addition to these efforts, deep learning models can model longer and more complex temporal dependencies. When anomalies are typically unlabeled, autoencoders learn what the normal price behaviour is, and deviations from it are detected via reconstruction error, while attention-based and Transformer models have been employed to predict prices and score anomalies in the same pipeline (Sediyono et al., 2025). There are two practical constraints. As with classical ML, deep models provide no explanation for deviations, which makes them less useful for policy decisions that must be explained, and they tend to require larger training sets than are available for any single commodity or region.

What the Crop Price Evidence Shows

Collectively, the studies in Table 1 demonstrate a progression from statistical to deep learning methods, but they were not designed for direct comparison, and are not interpreted as such in this review. Almost all are forecasting studies; only one (Sediyono et al., 2025) treats anomaly detection as an explicit objective, combining a Transformer with an attention-boosted LSTM-VAE and reporting an F1 score of 0.92, with precision and recall of 0.91 and 0.93, for multi-commodity forecasting and anomaly detection. It is the only study in the core set to report the full precision/recall/F1 metrics, which is itself notable: the recent crop-price literature remains focused on forecasting, with anomaly detection typically considered a secondary objective.

Among the forecasting studies, the highest reported accuracies are associated with decomposition-plus-deep-learning hybrids (F. Li & Tang, 2025; Sun et al., 2023; Choudhary et al., 2025), metaheuristic-tuned Transformers (Nayak et al., 2025), and tuned ensembles applied to a single commodity in a single market (Mao & Soonthornphisaj, 2024). Studies incorporating exogenous inputs such as weather or cross-market prices (Nayak et al., 2024; Jaiswal et al., 2025a) also report improvements over price-only models, consistent with the broader finding that richer context enhances performance. However, none of the studies in Table 1 provides an account of why a given prediction or flag was produced, which is the gap addressed in the remainder of this paper.

Table 1 lists the 27 core studies. Metrics are reported as published and are not comparable across studies, because commodity, region, dataset and evaluation protocol all differ. Where a study reported only error measures or a metric outside the table, the entry reads “not reported in comparable form”.

Table 1: Crop and agricultural commodity price forecasting and anomaly-detection studies (core scope, 2022–2025)

No.	**Study**	**Commodity / setting**	**Method**	**Reported metric (as published)**	Gap noted by the authors
1	An integrated framework for multi-commodity agricultural price forecasting and anomaly detection using attention-boosted models (Sediyono et al., 2025)	Multi-commodity; MSME markets	Transformer + attention-boosted LSTM-VAE	F1 0.92; Precision 0.91; Recall 0.93; Accuracy 0.94	Anomaly detection under-served in price forecasting; little MSME focus
2	Frequency-aware transformer for potato price forecasting in agricultural management (Zhang et al., 2025)	Potato	PatchTST + Mixture-of-Experts	Not reported in comparable form	Frequency-domain analysis under-used; price shocks poorly handled
3	Deep learning approaches for potato price forecasting: comparative analysis of LSTM, Bi-LSTM and AM-LSTM (Praveenkumar et al., 2025)	Potato	AM-LSTM, Bi-LSTM, LSTM	Accuracy 0.85	Attention mechanisms under-explored; limited interpretability
4	An enhanced LSTM recurrent neural network deep learning model for potato price prediction (Alzakari et al., 2024)	Potato	Enhanced LSTM	Accuracy 0.87	Few enhancement techniques; no cross-regional validation
5	NARX model for potato price prediction utilising	Potato; multiple	NARX	Accuracy 0.88	Single crop; cross-market

	multimarket information (Jaiswal et al., 2025a)	markets			information hard to integrate
6	Predicting potato prices in Agra, UP, India: an H2O AutoML approach (Kumari et al., 2024)	Potato; Agra, India	H2O AutoML	Accuracy 0.92	AutoML is a black box; limited hyperparameter control
7	Deep learning technique for forecasting the price of cauliflower (Paul et al., 2023)	Cauliflower; India	Deep learning (unspecified)	Accuracy 0.84	High-volatility vegetables under-studied; weak seasonal adjustment
8	Machine learning techniques for forecasting agricultural prices: a case of brinjal in Odisha, India (Paul et al., 2022)	Brinjal; Odisha, India	GRNN, SVR, RF, GBM	Accuracy 0.86	Vegetable-specific challenges; single-market validation
9	Recurrent neural network architecture for forecasting banana prices in Gujarat, India (Kumari et al., 2023)	Banana; Gujarat, India	RNN	Accuracy 0.85	Single crop; no comparison with newer architectures
10	Prediction of corn futures prices with decomposition and hybrid deep learning models (F. Li & Tang, 2025)	Corn futures	CEEMDAN + CNN-LSTM-Attention + PSO	Accuracy 0.94	Hyperparameter tuning; limited feature extraction from drivers
11	Thailand's maize prices forecasting using ensemble technique (Mao & Soonthornphisaj, 2024)	Maize; Thailand	Bagging-SVR with feature importance	Accuracy 0.96	Feature importance not tested on other crops; limited generalisability
12	Enhancing agricultural commodity price forecasting with deep learning (Manogna et al., 2025)	Multi-commodity	GRU Nataliia, F. (2026). Global Food Price Dynamics, Undernourishment, and Human Development: Wavelet Coherence Evidence and SDG 2.1 Resilience Scenarios up to 2030. Sustainability, 18(8), 3724., LSTM, ARIMA, ESN	Not reported in comparable form	Traditional models struggle with volatility; little multi-commodity analysis
13	Meta-transformer: leveraging metaheuristic algorithms for agricultural commodity price forecasting (Nayak et al., 2025)	Multi-commodity	Transformer + PSO/GWO/WOA	Accuracy 0.93	Hyperparameter tuning; little optimisation work in agriculture
14	Exogenous variable driven deep learning models for improved price forecasting of TOP crops in India (Nayak et al., 2024)	Tomato, onion, potato; India	TransformerX, NBEATSX	Accuracy 0.92	Weather variables under-used as exogenous inputs
15	Can deep learning models enhance the accuracy of agricultural price forecasting?	Multi-commodity; India	GRU, LSTM, RNN vs ARIMA, ANN, SVR	Accuracy 0.89	High-volatility crops poorly predicted; limited regional diversity

	Insights from India (Paul et al., 2025)				
16	Agricultural commodity price prediction based on secondary decomposition and LSTM (Sun et al., 2023)	Multi-commodity	VMD + EEMD + LSTM	Accuracy 0.91	Noise in price series; decomposition effectiveness limited
17	Hidden Markov guided deep learning models for forecasting highly volatile agricultural commodity prices (Avinash et al., 2024)	High-volatility commodities	HMM + deep learning	Accuracy 0.88	Volatility and regime change poorly modelled
18	A genetic algorithm optimised hybrid model for agricultural price forecasting based on VMD and LSTM (Choudhary et al., 2025)	Agricultural price series	GA-VMD-LSTM	Accuracy 0.92	Complex architecture; high compute; limited real-time use
19	Forecasting agricultural commodities prices using deep learning-based models: basic LSTM, Bi-LSTM, stacked LSTM, CNN-LSTM and convolutional LSTM (Murugesan et al., 2022)	Multi-commodity	LSTM variants	Accuracy 0.88	Limited comparative analysis; little hyperparameter optimisation
20	STL-LSTM hybrid model for forecasting seasonal agricultural price series (Jaiswal et al., 2025b)	Seasonal price series	STL-LSTM	Accuracy 0.91	Decomposition effectiveness varies; non-seasonal components
21	Commodity price forecasting via neural networks for coffee, corn, cotton, oats, soybeans, soybean oil, sugar and wheat (Xu & Zhang, 2022)	Eight commodities	Neural networks	Accuracy 0.87	Cross-commodity relationships under-explored
22	An ingenious method for estimating future crop prices that emphasises machine learning and deep learning models (Banerjee & Mondal, 2023)	Multi-crop	ML + DL hybrid	Accuracy 0.90	Integration methodology unclear; few validation datasets
23	Hybrid forecasting of agricultural commodity prices: integrating machine learning, time series and stochastic simulation models (Celik & Celik, 2025)	Multi-commodity	ML + time series + stochastic simulation	Not reported in comparable form	Stochastic elements rarely integrated with deep learning
24	Various optimised machine learning techniques to predict agricultural commodity	Multi-commodity	RF, SVM, XGBoost (optimised)	Accuracy 0.90	Few optimisation frameworks; limited comparative

	prices (Sari et al., 2024)				analysis
25	Forecasting commodity prices in Brazil through hybrid SSA-complex seasonality models (Palazzi et al., 2023)	Commodities; Brazil	SVR + LSTM; SVR + AdaBoost	Accuracy 0.91	Ensembles little explored; regional constraints
26	Artificial bee colony-based combination approach to forecasting agricultural commodity prices (J. Wang et al., 2022)	Multi-commodity	ABC + SSA/EMD/VMD + ARIMA/SVR/LSTM	Accuracy 0.93	Complex ensemble; interpretability; high compute
27	Forecastability of agricultural commodity futures realised volatility with daily infectious disease-related uncertainty (Shiba et al., 2022)	Agricultural futures volatility	GARCH models	Not reported in comparable form	No ML integration

Neuro Symbolic AI regarding anomaly detection: evidence from other domains

Table 2 summarises 24 studies that apply NeSyAI to anomaly detection, diagnosis or a closely related task in cybersecurity, healthcare, manufacturing, transport and other domains. None concerns agricultural prices; to our knowledge, no published study yet does. We review them to establish what the paradigm has delivered and where it has struggled before proposing it for a new domain. They are not comparable with Table 1 and we do not compare them.

Table 2: Neuro-symbolic AI applications in anomaly detection across domains (2022–2025)

No.	Study	Approach	Domain	Reported performance	Challenges and limitations
1	A neuro-symbolic classifier with optimised satisfiability for monitoring security alerts in network traffic (Onchis et al., 2022)	Logic Tensor Networks	Network security monitoring	Effective alert monitoring; optimises satisfiability of logical constraints	Computational complexity; needs well-defined logical constraints
2	A neuro-symbolic approach for anomaly detection and complex fault diagnosis in the automotive domain (Bohne et al., 2023)	Hybrid neuro-symbolic framework	Automotive fault diagnosis	High diagnostic accuracy; knowledge-based methods complement ML	Implementation complexity; domain expertise needed for knowledge engineering
3	Neuro-symbolic approaches for context-aware human activity recognition (Arrotta et al., 2023)	Semantic loss with context constraints	Human activity recognition	Outperforms purely data-driven models; no symbolic module needed at inference	May not generalise to all scenarios; depends on context quality

4	Neuro-symbolic empowered denoising diffusion probabilistic models for real-time anomaly detection in Industry 4.0 (Capogrosso et al., 2023)	Diffusion models + industrial ontologies	Industry 4.0 real-time monitoring	Real-time prediction; deployable on embedded systems	Implementation complexity; needs accurate ontologies
5	Knowledge-enhanced neurosymbolic AI for cybersecurity and privacy (Piplai et al., 2023)	Knowledge-enhanced NeSy	Cybersecurity and privacy	Improved threat detection and interpretability	Needs high-quality knowledge bases; real-time complexity
6	ECATS: explainable-by-design concept-based anomaly detection for time series (Ferfoglia et al., 2024)	Concept-based NeSy architecture	Cyber-physical time series	High classification performance with local interpretability on CPS datasets	Preliminary experiments; needs validation on diverse real-world data
7	A neuro-symbolic explainer for rare events: predictive maintenance (Gama et al., 2024)	LSTM autoencoder + rule learning	Predictive maintenance (public transport)	Local and global explanations; effective on rare events	Online-learning overhead; limited evaluation across datasets
8	Neurosymbolic learning and domain knowledge-driven explainable AI for IoT network attack detection (Kalutharage et al., 2025)	Knowledge-infused AI with domain knowledge graphs	IoT attack detection	Detection accuracy 0.97; 100% threat-intelligence accuracy using MITRE ATT&CK	Depends on knowledge-graph completeness; continuous updates required
9	Improved graph-based model for real-time anomaly detection in healthcare using hybrid CNN-LSTM and federated learning (Nagamani & Kumar, 2024)	Hybrid CNN-LSTM + federated learning	Healthcare anomaly detection	Accuracy 94%; false-positive rate 2% (10,000-patient dataset)	Multimodal integration complexity; heavy compute for federated learning
10	Data-driven model for predictive maintenance of CT equipment (H. Zhou et al., 2024)	SAX-HCBOP dictionary-based model	Healthcare predictive maintenance	Accuracy 0.904; Recall 0.747; Precision 0.417; F1 0.535	Low precision; extensive feature engineering; domain-specific

11	Explainable diagnosis prediction through neuro-symbolic integration (Zang et al., 2024)	NeSy integration for diagnosis	Healthcare diagnosis	Improved diagnostic accuracy with logical reasoning	Needs structured medical knowledge bases; real-time deployment complexity
12	Anomaly-based threat detection in smart health using machine learning (Tabassum et al., 2024)	Transformer-based anomaly detection	Smart-health threat detection	F1 0.92 (multimodal)	High compute; dataset quality and availability
13	KnowGraph: knowledge-enabled anomaly detection via logical reasoning on graph data (A. Zhou et al., 2024)	KnowGraph	Fraud and intrusion detection on graphs	Over 50% improvement in generalisation; handles distribution shift	Knowledge-engineering complexity; needs well-defined rules
14	On the use of neurosymbolic AI for defending against cyber attacks (Grov et al., 2024)	NeSy cyber-defence system	Cyber-attack detection and response	Enhanced detection; explainable responses	Scalability on large networks; expert-defined rules
15	Designing NORIA: a knowledge graph-based platform for anomaly detection and incident management in ICT systems (Tailhardat et al., 2023)	NORIA (ML + ontology + reasoning)	Network infrastructure monitoring	Multi-domain incident correlation; lower operator cognitive load	Comprehensive ontology engineering; heterogeneous environments
16	Neuro-symbolic temporal point processes (Y. Yang et al., 2024)	NeSy temporal point processes	Temporal event prediction and explanation	Efficient discovery of compact temporal logic rules; improved interpretability	Needs sufficient training data; rules hard for non-experts to read
17	NeuSyRE: neuro-symbolic visual understanding and reasoning based on scene graph enrichment (Khan et al., 2024)	NeuSyRE	Visual scene understanding	Beats state of the art on Recall@K and mean Recall@K; better SPICE and CIDEr	Dataset imbalance; needs high-quality knowledge graphs
18	Towards neuro-symbolic video understanding (M. Choi et al., 2024)	Temporal-logic video understanding	Autonomous-driving video analysis	F1 improvement of 9–15% over GPT-4 on Waymo and NuScenes	Needs labelled video; temporal-logic specifications are hard to write
19	Anomaly detection in cloud computing using	Knowledge-graph embedding +	Cloud anomaly detection	Effective real-time detection; preserves infrastructure graph	Scalability on very large infrastructures; continuous graph updates

	knowledge graph embedding and ML (Mitropoulou et al., 2023)	GraphSAGE		structure	
20	Towards knowledge graph reasoning for supply chain risk management using GNNs (Kosasih et al., 2022)	NeSy reasoning for supply chains	Supply-chain risk management	Uncovers hidden risks; predicts multiple relationship types	Needs comprehensive supply-chain data; ontology engineering
21	Surveying neuro-symbolic approaches for reliable artificial intelligence of things (Lu et al., 2024)	NeSy survey for AIoT	Reliable IoT / AIoT	Addresses testability, verifiability and interpretability limits of DL	Data quality; coverage testing; scalability for mission-critical systems
22	Robust novel defect detection with neurosymbolic AI (Theodoropoulos et al., 2024)	NeSy defect detection	Manufacturing defect detection	More robust to novel defects; better generalisation	Expert rule definition; limited evaluation across scenarios
23	Knowledge-augmented explainable and interpretable learning for anomaly detection and diagnosis (Atzmueller et al., 2024)	Knowledge-augmented explainable learning	General anomaly detection and diagnosis	Better understandability and transparency across domains	Knowledge-engineering complexity; interpretability–performance trade-off
24	Neuro-symbolic AI for cloud intrusion detection: a hybrid intelligence approach (Gajjar, 2025)	Hybrid DL + symbolic reasoning	Cloud intrusion detection	Improved accuracy, adaptability and explainability	Integration of diverse techniques; high computational overhead

What the NeSy evidence shows (Table 2)

The earliest studies in the set, from 2022 and 2023, made the basic case that combining symbolic inference with neural networks can make anomaly detection more transparent without giving up much accuracy. Logic Tensor Networks produced self-explaining intrusion detection (Onchis et al., 2022), and a hybrid NeSy framework gave naturally interpretable fault diagnoses in the automotive domain (Bohne et al., 2023). Over 2023 and 2024 the reported benefits widened to include speed, with several real-time deployments in manufacturing and healthcare (Bougzime et al., 2025; Nawaz et al., 2025). Diffusion models combined with industrial ontologies were run on embedded hardware for Industry 4.0 monitoring (Capogrosso et al., 2023), the ECATS architecture used Signal Temporal Logic formulae as interpretable concepts for time-series anomaly detection (Ferfoglia et al., 2024), and a hybrid CNN-LSTM healthcare system reported 94% accuracy with a 2% false-positive rate (Nagamani & Kumar, 2024). The 2024 work leans more heavily on knowledge graphs: KnowGraph reported a generalisation improvement of more than 50% under distribution shift (A. Zhou et al., 2024), NORIA paired Semantic Web reasoning with deep learning for network incident management (Tailhardat et al., 2023), and graph neural networks over supply-chain knowledge graphs were used to surface hidden risk relationships (Kosasih et al., 2022).

The limitations are just as consistent as the gains, and they deserve equivalent weight. Almost every study in Table 2 reports some combination of high computing demand, the effort of building and upkeeping the knowledge structures the symbolic layer depends on, difficulty scaling to large or heterogeneous environments, and sensitivity to the quality and completeness of that knowledge (Atzmueller et al., 2024; Lu et al., 2024; Nawaz et al., 2025). Several are explicitly preliminary, evaluated on one dataset or site. Where knowledge graphs are incomplete or stale, performance degrades (Kalutharage et al., 2025; Lu et al., 2024), and hand-written rules inherit their authors' blind spots. The field's own agenda (automated knowledge-graph construction, faster reasoning, shared benchmarks) is a fair list of what remains unsolved, and every item would carry over to agriculture.

A Proposed NeSy Framework for Crop Price Anomaly Detection

Status of This Section

Everything in this section is a design proposal. The framework has not been implemented and no performance figures exist for it. We describe it at the level of detail a research team would need to build and test it, distinguishing design choices backed by the Table 2 evidence from those that rest on our own judgement about agricultural price data.

Architecture Overview

The framework has four parts. (1) A neural detection layer scores each new price observation for abnormality. (2) A domain knowledge graph holds what is known about the market: commodities, regions, seasons, inputs, policies and their relationships. (3) A symbolic reasoning layer applies explicit rules over the knowledge graph to interpret each neural score, confirming, suppressing or re-ranking the alert and attaching an explanation. (4) A reporting layer renders the result for a non-technical user. The coupling between (1) and (3) is loose: the neural component produces a score and the symbolic component decides what to do with it. This is the same two-stage pattern used by Gama et al. (2024) and Bohne et al. (2023), and we prefer it to tighter integration (for example, encoding rules directly into the loss function, as in Arrotta et al., 2023) for a practical reason: rules that live outside the network can be inspected and edited by an agricultural economist who has never trained a model.

Neural Detection Layer

The candidate approaches are those that performed best in Table 1 and that work with little or no labelled data: an LSTM autoencoder scoring anomalies by reconstruction error, and a Transformer-driven forecaster scoring them by forecast residual (Sediyono et al., 2025; Wen et al., 2022). Both are unsupervised at training time, which is essential because labelled anomaly events in Nigerian price data are rare. Inputs would be weekly or monthly price series per commodity and market, decomposed into trend, seasonal and residual components before modelling (as in Jaiswal et al., 2025b), together with exogenous covariates: rainfall, fuel price, the naïra–dollar exchange rate and transport cost indices. The output of this layer is a continuous anomaly score and, where the model permits, an attribution of which inputs drove it.

Knowledge Graph Construction

The knowledge graph is the part of the framework that carries the most risk and the most value, so we describe it concretely. Its entity types would be: commodity (maize, rice, sorghum, yam, tomato and so on), market (a named physical market with a state as well as geopolitical zone), season (planting, growing, harvest and lean periods per commodity and zone), input (fuel, fertiliser, transport), policy event (subsidy change, border closure, import ban, exchange-rate regime change), climate event (drought, flood, late rains) and security event (conflict-related disruption of a supply route). Relationships would include harvested_in (commodity→season→zone), sourced_from (market→production zone), substitutes_for (commodity→commodity), input_to (input→commodity, with a typical lag), affects_route (security event→market pair) and in_effect_during (policy event→date range). Each relation would carry a typical lag and an expected direction of effect where the literature or expert opinion supports one.

Populating the graph would draw on three sources. First, structured public data: crop calendars from FAO GIEWS and FEWS NET, agro-ecological zone boundaries, and administrative market lists from the National Bureau of Statistics. Second, event records: dated policy announcements from the Central Bank of Nigeria and federal ministries, rainfall and flood records from the Nigerian Meteorological Agency, and security-incident data from public conflict trackers. Third, expert elicitation: systematic interviews with agricultural economists, extension officers and market traders to capture relationships that are not written down anywhere, such as which markets actually supply which, and how long a fuel price rise usually takes to show up in a tomato price in a given city. Each elicited relation would be recorded with its source alongside a confidence level, so that the graph is auditable and so that low-confidence relations can be tested rather than trusted.

Symbolic Rule Development

Rules would be written in a simple, readable form (if–then over graph relations and price statistics) so that domain experts can review them. We anticipate four rule families. Seasonal expectation rules suppress alerts that coincide with a known seasonal pattern: for instance, a 20% maize price rise in the lean season in a northern zone is expected, not anomalous. Input-propagation rules link an upstream shock to a downstream price movement through the input_to lag: a fuel price rise in week t followed by a transport-cost rise and then a crop price rise within the expected lag is explained, and the explanation names the chain. Policy and event rules do the same for dated policy, climate and security events. Data-integrity rules catch things that are almost certainly recording errors, such as a single-week price that departs from both neighbours by an implausible multiple, or a price reported in the wrong unit. A neural anomaly that survives all four families is reported as an unexplained anomaly, which is the case that most warrants human attention.

Rules would be drafted from the literature and the elicitation interviews, then reviewed by a panel of domain experts before adoption. The rule base would start small (tens of rules, not hundreds), grow with use, and be versioned with every change logged, because a wrong rule will systematically suppress or over-report a whole class of irregularities.

Data Sources

Price data would be sourced from the National Bureau of Statistics' monthly food price releases, supplemented, where available, by FEWS NET and WFP market monitoring series, as well as NGO and cooperative records for markets not covered nationally. Covariate series would be obtained from the Nigerian Meteorological Agency (rainfall), the Central Bank of Nigeria (exchange rates), and published fuel-price series. It is important to acknowledge at the outset that these series contain gaps, changes in collection methods, and inconsistent market coverage; addressing these issues is an integral part of the research process.

Validation Strategy

Validation would proceed in four stages, and none of these stages has yet been conducted. (1) Retrospective event validation: the system would be applied to historical data, and its outputs compared to a curated list of known market shocks, such as the 2016 exchange-rate liberalisation, the 2019 land-border closure, and the 2023 fuel-subsidy removal, to assess whether the system correctly identifies and attributes these events. (2) Expert-labelled evaluation: a panel of agricultural economists would independently label a sample of price series for anomalous periods, providing a ground truth for calculating precision, recall, and F1 scores, with cross-validation across commodities and zones to assess robustness and control the false-positive rate. (3) Ablation: the evaluation would be repeated with the symbolic layer deactivated, allowing measurement of the knowledge graph and rules' contribution. (4) Comprehension study: system-generated explanations would be presented to extension officers and cooperative members, and their understanding of the alert and intended actions would be measured, compared to the same alert presented without an explanation. Stage 4 is considered the most critical test for the central claim of this paper, and it is the area least addressed in the existing NeSy literature.

Anticipated Implementation Challenges

Implementing this framework will present significant challenges, so it is important to specify the sources of difficulty. Knowledge engineering represents the most substantial cost. Constructing and validating the knowledge graph and rule base requires sustained engagement with domain experts, whose availability is limited. The resulting system will be tailored to Nigeria, and adapting it to other contexts, such as Ghana or Kenya, would require considerable redevelopment. Data availability remains a challenge, as agricultural price series are often sparse and inconsistent; a knowledge-rich system reduces but does not eliminate the need for reliable data. Computational demands are also considerable: running both a neural detector and a reasoning engine requires more resources than either component alone (Bougzime et al., 2025; Lu et al., 2024), which may constrain near-real-time deployment on limited infrastructure. Rule brittleness is another risk: rules based on outdated market structures may misclassify current anomalies, necessitating ongoing maintenance and oversight of the rule base. Evaluation presents further complexity, as determining what constitutes a "true" price anomaly often requires expert judgment, which may not always be unanimous. These challenges underscore the need for caution in interpreting the current proposal, as no empirical results have yet been demonstrated.

Illustrative Output for a Non-Technical Stakeholder

Figure 1 provides a mock-up illustration of how the reporting layer could present results to a user monitoring three related indicators (fuel price, production cost, and crop price) over time. This **is a conceptual design, not output from an implemented or evaluated system.** In this concept, anomalies are marked on each series, and the symbolic layer links related anomalies across indicators: a fuel-price rise, followed by a rise in production cost, and then a rise in crop price, each occurring within the lag expected by the knowledge graph. The user is presented with the causal chain, rather than only the final alert. Generating such output from real data would require the knowledge graph and rule base described above to be constructed and validated.

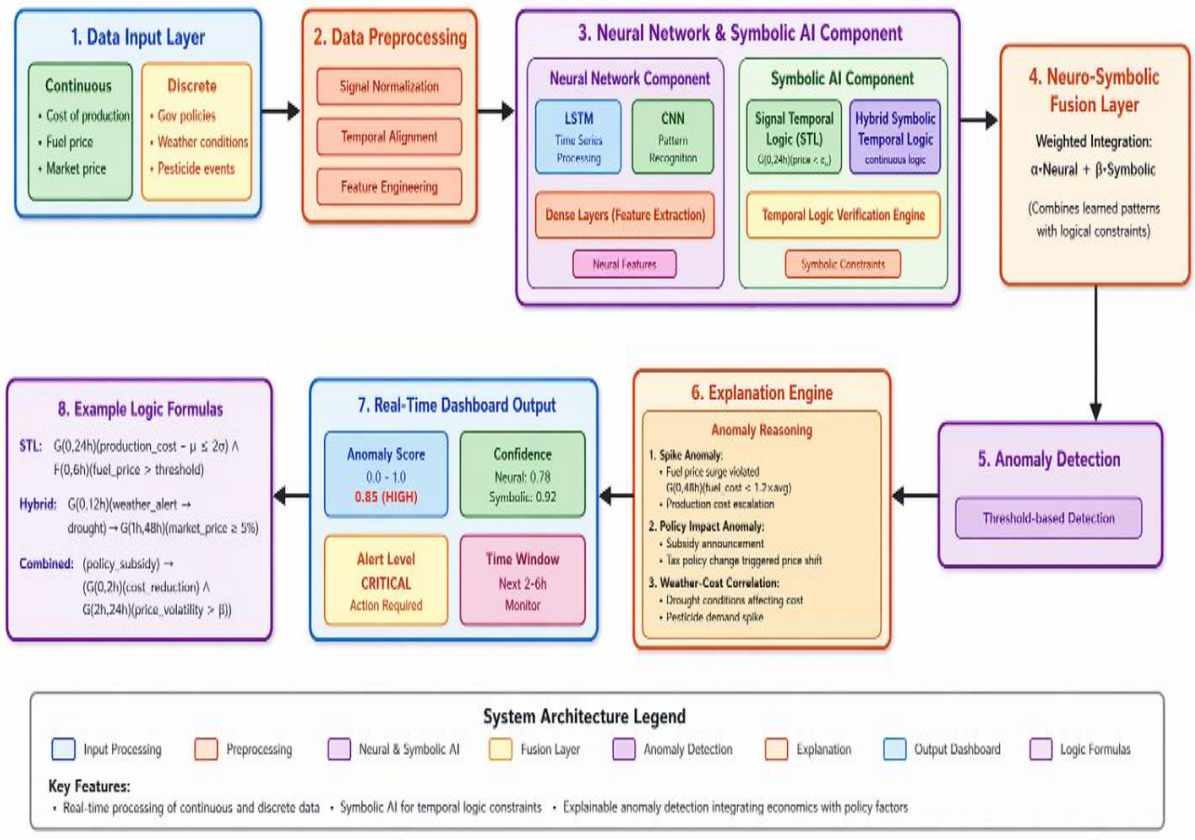


Figure 1: Illustrative design mock-up of the proposed reporting interface. Anomalies detected by the neural

layer (circled) are linked across the three indicators by the symbolic layer, which reports the inferred causal chain, its confidence and a suggested action. All series are synthetic; the figure is a design concept, not output from an implemented system.

Table 3: Qualitative comparison of ML/DL and NeSyapproaches for crop price anomaly detection

Note. Numerical performance ranges are deliberately not compared across the two columns, because the two literatures address different tasks, datasets and metrics. The table is confined to the qualitative properties that motivate this paper's proposal, and the right-hand column states the basis for each entry.

Evidence base in agricultural price data	Substantial: 27 studies, mostly forecasting, one with explicit anomaly detection	None: no published application to agricultural prices	Tables 1 and 2
Interpretability of an alert	Low; output is a score or class with no stated reason	High where symbolic rules or knowledge graphs are exposed; explanations are part of the output	Reported across Table 2 (e.g., Bohne et al., 2023; Gama et al., 2024)
Labelled-data requirement	High for supervised methods; unsupervised DL still needs long histories	Reported to be lower where domain knowledge substitutes for labels, though this is not tested on price data	Table 2 and Kautz (2022); untested in agriculture
Behaviour under distribution shift	Regional and temporal limits noted by several authors	Improved generalisation reported in some studies (e.g., KnowGraph)	Table 1 gap column; A. Zhou et al. (2024)
Real-time operation	Rarely documented in crop-price studies	Demonstrated in several industrial and healthcare settings	Table 2 (Capogrosso et al., 2023; Nagamani & Kumar, 2024)
Computational cost	Moderate to high, rising with model size and ensembles	Typically higher than either component alone; reasoning adds overhead	Table 2 limitations column; Bougzime et al. (2025)
Knowledge-engineering cost	Low; features are learned from data	High; knowledge graphs and rules must be built and maintained by experts	Table 2 limitations column
Risk of confident error	Unexplained false positives and negatives	Confidently wrong explanations if knowledge is incomplete or stale	Lu et al. (2024); authors' assessment
Maturity	Established in agricultural forecasting	Emerging; many Table 2 studies are preliminary or single-site	Table 2

DISCUSSION

Two findings emerge, and they pull in different directions. First, in the domains where NeSy systems have been tried, they report performance broadly comparable to purely neural methods while adding an explanation and, in several cases, better behaviour under distribution shift (A. Zhou et al., 2024; Kalutharage et al., 2025; Tailhardat et al., 2023). All of that evidence comes from outside agriculture.

Second, in agricultural price forecasting, which has a long history, ML and DL still hold a modest edge in raw accuracy: decomposition hybrids and tuned Transformers in Table 1 report 0.93 to 0.94 (F. Li & Tang, 2025; J. Wang et al., 2022; Nayak et al., 2025), and a tuned ensemble on a single national maize series reports 0.96 (Mao & Soonthornphisaj, 2024). We do not expect a NeSy system to beat those figures on the same data. What we expect is a trade-off, a small cost in peak accuracy on some tasks in exchange for an output a policymaker can interrogate. Whether that trade is worth making is an empirical question this paper cannot settle.

Research Gap

The gap this paper addresses is therefore specific. NeSy has not been applied to agricultural markets, and no existing NeSy architecture accounts for fuel-subsidy shocks, exchange-rate volatility, transport-cost swings, supply-chain insecurity, rainfall-driven production cycles, and informal trading that shape price formation in Nigeria. No published knowledge graph encodes Nigerian agricultural policy, regional climate, or cross-border trade, and no NeSy method has been designed for the sparse, inconsistently recorded price histories typical of Sub-Saharan Africa. Results reported elsewhere suggest the approach is worth trying; they do not show that it works here.

LIMITATIONS

Three limitations shape how this paper should be read. The literature search, as noted, was not systematic, and some selection bias is possible.

The performance figures in Tables 1 and 2 are taken as published, based on each study's own data and protocol; we did not reproduce or verify them, and some studies report a single accuracy value for what is essentially a regression task without specifying how it was computed. We have therefore avoided numerical comparisons across paradigms and confined Table 3 to qualitative trade-offs.

The proposed framework is conceptual and has not been evaluated using agricultural price data anywhere. Building it depends on the knowledge engineering, data access, computing infrastructure, and ground-truth validation described above, none of which has been completed. Any performance figures will have to come from future field testing, ideally in partnership with Nigerian agricultural agencies or cooperatives that can provide both data and expert judgment on what the system flags.

CONCLUSION

This paper reviewed recent ML and DL evidence on crop price anomaly detection and, separately, NeSy evidence from other fields, and used the latter to argue that combining the two is a promising but untested direction for agricultural markets. The proposed framework would let anomalies be flagged and explained in causal terms, distinguishing genuine disruption from seasonal movement or tracing a price anomaly to an upstream fuel or weather shock, in a form a stakeholder with limited technical expertise can follow. It is offered as a case for doing that work, not as a report of a system that has been built.

REFERENCES

1. Ali, H., & Fatima, T. (2025). Integrating Neural Networks and Symbolic Reasoning: A Neurosymbolic AI Approach for Decision-Making Systems. [Unpublished manuscript]. https://doi.org/10.13140/RG.2.2.35901.99046
2. Alzakari, S. A., Alhussan, A. A., Qenawy, A.-S. T., Elshewey, A. M., & Eed, M. (2024). An Enhanced Long Short-Term Memory Recurrent Neural Network Deep Learning Model for Potato Price Prediction. Potato Research. https://doi.org/10.1007/s11540-024-09744-x
3. Arrotta, L., Civitarese, G., & Bettini, C. (2023). Neuro-Symbolic Approaches for Context-Aware Human Activity Recognition (Version 1). arXiv. https://doi.org/10.48550/ARXIV.2306.05058
4. Atzmueller, M., Bohne, T., & Windler, P. (2024). Knowledge-Augmented Explainable and Interpretable Learning for Anomaly Detection and Diagnosis (Version 1). arXiv. https://doi.org/10.48550/ARXIV.2412.00146
5. Avinash, G., Ramasubramanian, V., Ray, M., Paul, R. K., Godara, S., Nayak, G. H. H., Kumar, R. R., Manjunatha, B., Dahiya, S., & Iquebal, M. A. (2024). Hidden Markov guided Deep Learning models for forecasting highly volatile agricultural commodity prices. Applied Soft Computing, 158, 111557. https://doi.org/10.1016/j.asoc.2024.111557

6. Banerjee, S., & Mondal, A. C. (2023). An ingenious method for estimating future crop prices that emphasises machine learning and deep learning models. International Journal of Information Technology, 15(8), 4291–4313. https://doi.org/10.1007/s41870-023-01474-6
7. Besold, T. R., d'Avila Garcez, A., Jimenez-Ruiz, E., Confalonieri, R., Madhyastha, P., & Wagner, B. (Eds.). (2024). Neural-Symbolic Learning and Reasoning: 18th International Conference, NeSy 2024, Barcelona, Spain, September 9–12, 2024, Proceedings, Part I (Vol. 14979). Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-71167-1
8. Bohne, T., Windler, A.-K. P., & Atzmueller, M. (2023). A Neuro-Symbolic Approach for Anomaly Detection and Complex Fault Diagnosis Exemplified in the Automotive Domain. Proceedings of the 12th Knowledge Capture Conference 2023, 35–43. https://doi.org/10.1145/3587259.3627546
9. Bougzime, O., Jabbar, S., Cruz, C., & Demoly, F. (2025). Unlocking the Potential of Generative AI through Neuro-Symbolic Architectures – Benefits and Limitations. arXiv. https://doi.org/10.48550/arXiv.2502.11269
10. Capogrosso, L., Mascolini, A., Girella, F., Skenderi, G., Gaiardelli, S., Dall'Ora, N., Ponzio, F., Fraccaroli, E., Di Cataldo, S., Vinco, S., Macii, E., Fummi, F., & Cristani, M. (2023). Neuro-symbolic Empowered Denoising Diffusion Probabilistic Models for Real-time Anomaly Detection in Industry 4.0 (Version 2). arXiv. https://doi.org/10.48550/ARXIV.2307.06975
11. Celik, B. A., & Celik, S. (2025). Hybrid forecasting of agricultural commodity prices: Integrating machine learning, time series, and stochastic simulation models. Borsa Istanbul Review, 25(6), 1440–1462. https://doi.org/10.1016/j.bir.2025.10.004
12. Choi, H., Kim, D., Kim, J., Kim, J., & Kang, P. (2022). Explainable anomaly detection framework for predictive maintenance in manufacturing systems. Applied Soft Computing, 125, 109147. https://doi.org/10.1016/j.asoc.2022.109147
13. Choi, M., Goel, H., Omama, M., Yang, Y., Shah, S., & Chinchali, S. (2024). Towards Neuro-Symbolic Video Understanding. https://doi.org/10.48550/ARXIV.2403.11021
14. Choruma, D. J., Dirwai, T. L., Mutenje, M. J., Mustafa, M., Chimonyo, V. G. P., Jacobs-Mata, I., & Mabhaudhi, T. (2024). Digitalisation in agriculture: A scoping review of technologies in practice, challenges, and opportunities for smallholder farmers in sub-saharan africa. Journal of Agriculture and Food Research, 18, 101286. https://doi.org/10.1016/j.jafr.2024.101286
15. Choudhary, K., Jha, G. K., Jaiswal, R., & Kumar, R. R. (2025). A genetic algorithm optimized hybrid model for agricultural price forecasting based on VMD and LSTM network. Scientific Reports, 15(1), 9932. https://doi.org/10.1038/s41598-025-94173-0
16. De Arriba-Pérez, F., García-Méndez, S., Leal, F., Malheiro, B., & Burguillo, J. C. (2024). Exposing and explaining fake news on-the-fly. Machine Learning, 113(7), 4615–4637. https://doi.org/10.1007/s10994-024-06527-w
17. Ferfoglia, I., Saveri, G., Nenzi, L., & Bortolussi, L. (2024). ECATS: Explainable-by-Design Concept-Based Anomaly Detection for Time Series. In T. R. Besold, A. d'Avila Garcez, E. Jimenez-Ruiz, R. Confalonieri, P. Madhyastha, & B. Wagner (Eds.), Neural-Symbolic Learning and Reasoning (Vol. 14980, pp. 175–191). Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-71170-1_16
18. Gajjar, S. R. (2025). Neuro-symbolic AI for cloud intrusion detection: A hybrid intelligence approach. GSC Advanced Research and Reviews, 22(2), 142–144. https://doi.org/10.30574/gscarr.2025.22.2.0049
19. Gama, J., Ribeiro, R. P., Mastelini, S., Davarid, N., & Veloso, B. (2024). A Neuro-Symbolic Explainer for Rare Events: A Case Study on Predictive Maintenance (Version 1). arXiv. https://doi.org/10.48550/ARXIV.2404.14455
20. Garcez, A. d'Avila, & Lamb, L. C. (2023). Neurosymbolic AI: The 3rd wave. Artificial Intelligence Review, 56(11), 12387–12406. https://doi.org/10.1007/s10462-023-10448-w
21. Grov, G., Halvorsen, J., Eckhoff, M. W., Hansen, B. J., Eian, M., & Mavroeidis, V. (2024). On the Use of Neurosymbolic AI for Defending Against Cyber Attacks. In T. R. Besold, A. d'Avila Garcez, E. Jimenez-Ruiz, R. Confalonieri, P. Madhyastha, & B. Wagner (Eds.), Neural-Symbolic Learning and Reasoning (Vol. 14979, pp. 119–140). Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-71167-1_7

22. Jabed, Md. A., & Azmi Murad, M. A. (2024). Crop yield prediction in agriculture: A comprehensive review of machine learning and deep learning approaches, with insights for future research and sustainability. Heliyon, 10(24), e40836. https://doi.org/10.1016/j.heliyon.2024.e40836
23. Jaiswal, R., Jha, G. K., Kumar, R. R., & Choudhary, K. (2025a). NARX Model for Potato Price Prediction Utilising Multimarket Information. Potato Research, 68(3), 2609–2630. https://doi.org/10.1007/s11540-025-09846-0
24. Jaiswal, R., Jha, G. K., Kumar, R. R., & Choudhary, K. (2025b). STL-LSTM Hybrid Model for Forecasting Seasonal Agricultural Price Series. Annals of Data Science. https://doi.org/10.1007/s40745-025-00590-3
25. Kalutharage, C. S., Liu, X., & Chrysoulas, C. (2025). Neurosymbolic learning and domain knowledge-driven explainable AI for enhanced IoT network attack detection and response. Computers & Security, 151, 104318. https://doi.org/10.1016/j.cose.2025.104318
26. Kautz, H. A. (2022). The third AI summer: AAAI Robert S. Engelmore Memorial Lecture. AI Magazine, 43(1), 105–125. https://doi.org/10.1002/aaai.12036
27. Khan, M. J., G. Breslin, J., & Curry, E. (2024). NeuSyRE: Neuro-symbolic visual understanding and reasoning framework based on scene graph enrichment. Semantic Web, 15(4), 1389–1413. https://doi.org/10.3233/SW-233510
28. Kosasih, E. E., Margaroli, F., Gelli, S., Aziz, A., Wildgoose, N., & Brintrup, A. (2022). Towards knowledge graph reasoning for supply chain risk management using graph neural networks. International Journal of Production Research. Advance online publication. https://doi.org/10.1080/00207543.2022.2100841
29. Kumari, P., Goswami, V., N., H., & Pundir, R. S. (2023). Recurrent neural network architecture for forecasting banana prices in Gujarat, India. PLOS ONE, 18(6), e0275702. https://doi.org/10.1371/journal.pone.0275702
30. Kumari, P., M, S. K., Vekariya, P., Kujur, S. N., Macwan, J., & Mishra, P. (2024). Predicting Potato Prices in Agra, UP, India: An H2O AutoML Approach. Potato Research. https://doi.org/10.1007/s11540-024-09726-z
31. Lee, M., Lin, J., & Katsikas, S. (2024). Exploring the effects of RNNs and deep learning frameworks on real-time, lightweight, adaptive time series anomaly detection. Concurrency and Computation: Practice and Experience, 36(28), e8288. https://doi.org/10.1002/cpe.8288
32. Li, F., & Tang, M. (2025). Prediction of corn futures prices with decomposition and hybrid deep learning models. Knowledge and Information Systems, 67(3), 2427–2452. https://doi.org/10.1007/s10115-024-02301-5
33. Liang, B., Wang, Y., & Tong, C. (2025). AI Reasoning in Deep Learning Era: From Symbolic AI to Neural–Symbolic AI. Mathematics, 13(11), 1707. https://doi.org/10.3390/math13111707
34. Lu, Z., Afridi, I., Kang, H. J., Ruchkin, I., & Zheng, X. (2024). Surveying neuro-symbolic approaches for reliable artificial intelligence of things. Journal of Reliable Intelligent Environments, 10(3), 257–279. https://doi.org/10.1007/s40860-024-00231-1
35. Manogna, R. L., Dharmaji, V., & Sarang, S. (2025). Enhancing agricultural commodity price forecasting with deep learning. Scientific Reports, 15(1), 20903. https://doi.org/10.1038/s41598-025-05103-z
36. Mao, S., & Soonthornphisaj, N. (2024). Thailand's Maize Prices Forecasting using Ensemble Technique. ASEAN Journal of Scientific and Technological Reports, 27(4), e252279. https://doi.org/10.55164/ajstr.v27i4.252279
37. Mitropoulou, K., Kokkinos, P., Varvarigos, M., & Soumplis, P. (2023). Anomaly Detection in Cloud Computing using Knowledge Graph Embedding and Machine Learning Mechanisms. Journal of Grid Computing, 22(1), 6. https://doi.org/10.1007/s10723-023-09727-1
38. Murugesan, R., Mishra, E., & Krishnan, A. H. (2022). Forecasting agricultural commodities prices using deep learning-based models: Basic LSTM, bi-LSTM, stacked LSTM, CNN LSTM, and convolutional LSTM. International Journal of Sustainable Agricultural Management and Informatics, 8(3), 242. https://doi.org/10.1504/IJSAMI.2022.125757
39. Nagamani, G. M., & Kumar, C. K. (2024). Design of an improved graph-based model for real-time anomaly detection in healthcare using hybrid CNN-LSTM and federated learning. Heliyon, 10(24), e41071. https://doi.org/10.1016/j.heliyon.2024.e41071

40. Nawaz, U., Anees-ur-Rahaman, M., & Saeed, Z. (2025). A review of neuro-symbolic AI integrating reasoning and learning for advanced cognitive systems. Intelligent Systems with Applications, 26, 200541. https://doi.org/10.1016/j.iswa.2025.200541
41. Nayak, G. H. H., Alam, M. W., Singh, K. N., Avinash, G., Kumar, R. R., Ray, M., & Deb, C. K. (2024). Exogenous variable driven deep learning models for improved price forecasting of TOP crops in India. Scientific Reports, 14(1), 17203. https://doi.org/10.1038/s41598-024-68040-3
42. Nayak, G. H. H., Alam, Md. W., Naik, B. S., Varshini, B. S., Avinash, G., Kumar, R. R., Ray, M., & Singh, K. N. (2025). Meta-transformer: Leveraging metaheuristic algorithms for agricultural commodity price forecasting. Journal of Big Data, 12(1), 138. https://doi.org/10.1186/s40537-025-01196-5
43. Onchis, D., Istin, C., & Hogea, E. (2022). A Neuro-Symbolic Classifier with Optimized Satisfiability for Monitoring Security Alerts in Network Traffic. Applied Sciences, 12(22), 11502. https://doi.org/10.3390/app122211502
44. Palazzi, R. B., Maçaira, P., Meira, E., & Klotzle, M. C. (2023). Forecasting commodity prices in Brazil through hybrid SSA-complex seasonality models. Production, 33, e20220025. https://doi.org/10.1590/0103-6513.20220025
45. Paul, R. K., Yeasin, M., Kumar, P., Paul, A. K., & Roy, H. S. (2023). Deep Learning Technique for Forecasting the Price of Cauliflower. Current Science, 124(9), 1065. https://doi.org/10.18520/cs/v124/i9/1065-1073
46. Paul, R. K., Yeasin, M., Tamilselvi, C., Paul, A. K., Sharma, P., & Birthal, P. S. (2025). Can Deep Learning Models Enhance the Accuracy of Agricultural Price Forecasting? Insights From India. Intelligent Systems in Accounting, Finance and Management, 32(1), e70002. https://doi.org/10.1002/isaf.70002
47. Paul, R. K., Yeasin, Md., Kumar, P., Kumar, P., Balasubramanian, M., Roy, H. S., Paul, A. K., & Gupta, A. (2022). Machine learning techniques for forecasting agricultural prices: A case of brinjal in Odisha, India. PLOS ONE, 17(7), e0270553. https://doi.org/10.1371/journal.pone.0270553
48. Piplai, A., Kotal, A., Mohseni, S., Gaur, M., Mittal, S., & Joshi, A. (2023). Knowledge-enhanced Neuro-Symbolic AI for Cybersecurity and Privacy (Version 1). arXiv. https://doi.org/10.48550/ARXIV.2308.02031
49. Praveenkumar, A., Jha, G. K., Madival, S. D., Lama, A., & Kumar, R. R. (2025). Deep Learning Approaches for Potato Price Forecasting: Comparative Analysis of LSTM, Bi-LSTM, and AM-LSTM Models. Potato Research, 68(2), 1941–1963. https://doi.org/10.1007/s11540-024-09823-z
50. Ranaweera, H. M. B. P., Rathnayake, R. M. G. H. N., & Ananda, A. S. G. J. K. (2023). Crop Price Prediction Using Machine Learning Approaches: Reference to the Sri Lankan Vegetable Market. Journal of Management Matters, 10(1), 19–34. https://doi.org/10.4038/jmm.v10i1.45
51. Sari, M., Duran, S., Kutlu, H., Guloglu, B., & Atik, Z. (2024). Various optimized machine learning techniques to predict agricultural commodity prices. Neural Computing and Applications, 36(19), 11439–11459. https://doi.org/10.1007/s00521-024-09679-x
52. Sediyono, E., Hartomo, K. D., Arthur, C., Utami, I., Prabowo, R., & Chiong, R. (2025). An integrated framework for multi-commodity agricultural price forecasting and anomaly detection using attention-boosted models. Journal of Agriculture and Food Research, 22, 102021. https://doi.org/10.1016/j.jafr.2025.102021
53. Shiba, S., Aye, G. C., Gupta, R., & Goswami, S. (2022). Forecastability of Agricultural Commodity Futures Realised Volatility with Daily Infectious Disease-Related Uncertainty. Journal of Risk and Financial Management, 15(11), 525. https://doi.org/10.3390/jrfm15110525
54. Sun, C., Pei, M., Cao, B., Chang, S., & Si, H. (2023). A Study on Agricultural Commodity Price Prediction Model Based on Secondary Decomposition and Long Short-Term Memory Network. Agriculture, 14(1), 60. https://doi.org/10.3390/agriculture14010060
55. Tabassum, M., Mahmood, S., Bukhari, A., Alshemaimri, B., Daud, A., & Khalique, F. (2024). Anomaly-based threat detection in smart health using machine learning. BMC Medical Informatics and Decision Making, 24(1), 347. https://doi.org/10.1186/s12911-024-02760-4
56. Tailhardat, L., Chabot, Y., & Troncy, R. (2023). Designing NORIA: a Knowledge Graph-based Platform for Anomaly Detection and Incident Management in ICT Systems. In Proceedings of the

4th International Workshop on Knowledge Graph Construction (KGCW 2023), co-located with ESWC 2023. CEUR Workshop Proceedings.

57. Theodoropoulos, S., Makridis, G., Kyriazis, D., & Tsanakas, P. (2024). Robust Novel Defect Detection with Neurosymbolic AI. In M. Thürer, R. Riedel, G. Von Cieminski, & D. Romero (Eds.), Advances in Production Management Systems. Production Management Systems for Volatile, Uncertain, Complex, and Ambiguous Environments (Vol. 732, pp. 381–396). Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-71637-9_26
58. Theofilou, A., Nastis, S. A., Michailidis, A., Bournaris, T., & Mattas, K. (2025). Predicting Prices of Staple Crops Using Machine Learning: A Systematic Review of Studies on Wheat, Corn, and Rice. Sustainability, 17(12), 5456. https://doi.org/10.3390/su17125456
59. Van Der Heever, W., Satapathy, R., Cambria, E., & Goh, R. S. M. (2024). Neurosymbolic AI for Mining Key Aspects of Socially Responsible Investing. 2024 IEEE International Conference on Data Mining Workshops (ICDMW), 503–510. https://doi.org/10.1109/ICDMW65004.2024.00070
60. Wang, J., Wang, Z., Li, X., & Zhou, H. (2022). Artificial bee colony-based combination approach to forecasting agricultural commodity prices. International Journal of Forecasting, 38(1), 21–34. https://doi.org/10.1016/j.ijforecast.2019.08.006
61. Wang, W., Yang, Y., & Wu, F. (2025). Towards Data-And Knowledge-Driven AI: A Survey on Neuro-Symbolic Computing. IEEE Transactions on Pattern Analysis and Machine Intelligence, 47(2), 878–899. https://doi.org/10.1109/TPAMI.2024.3483273
62. Wen, Q., Zhou, T., Zhang, C., Chen, W., Ma, Z., Yan, J., & Sun, L. (2022). Transformers in Time Series: A Survey. https://doi.org/10.48550/ARXIV.2202.07125
63. Xu, X., & Zhang, Y. (2022). Commodity price forecasting via neural networks for coffee, corn, cotton, oats, soybeans, soybean oil, sugar, and wheat. Intelligent Systems in Accounting, Finance and Management, 29(3), 169–181. https://doi.org/10.1002/isaf.1519
64. Yang, Y., Yang, C., Li, B., Fu, Y., & Li, S. (2024). Neuro-Symbolic Temporal Point Processes (Version 1). arXiv. https://doi.org/10.48550/ARXIV.2406.03914
65. Zang, H., Hu, A., Xu, X., Ren, H., & Xu, L. (2024). Development of machine learning models to predict perioperative blood transfusion in hip surgery. BMC Medical Informatics and Decision Making, 24(1), 158. https://doi.org/10.1186/s12911-024-02555-7
66. Zhang, A., Chen, S., Wu, C., Li, G., & Liu, L. (2025). PatchTST-DME: Frequency-Aware Transformer for Potato Price Forecasting in Agricultural Management. Potato Research. https://doi.org/10.1007/s11540-025-09907-4
67. Zhou, A., Xu, X., Raghunathan, R., Lal, A., Guan, X., Yu, B., & Li, B. (2024). KnowGraph: Knowledge-Enabled Anomaly Detection via Logical Reasoning on Graph Data. Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security, 168–182. https://doi.org/10.1145/3658644.3690354
68. Zhou, H., Liu, Q., Liu, H., Chen, Z., Li, Z., Zhuo, Y., Li, K., Wang, C., & Huang, J. (2024). Healthcare facilities management: A novel data-driven model for predictive maintenance of computed tomography equipment. Artificial Intelligence in Medicine, 149, 102807. https://doi.org/10.1016/j.artmed.2024.102807