Neuro-Symbolic Artificial Intelligence for Explainable Real-Time Anomaly Detection in Agricultural Commodity Markets: A Systematic Review and Architectural Framework for Developing Economies

[1]Badamasi Imam Ya'u., [2]Ali Usman Abdullahi., [3]Grace Ojochennemi Emmanuel Anorue

[1]Department of Computer Science, Islamic University Uganda

[2]Computer Science Department, Federal College of Education (Technical) Gombe

[3]Department of Computer Science, School of Science Education, Federal College of Education (Technical), Gombe

DOI: https://doi.org/10.47772/IJRISS.2026.1026EDU0293

Received: 18 May 2026; Accepted: 23 May 2026; Published: 05 June 2026

ABSTRACT

Food security is seriously threatened by price irregularities in agricultural products, which disproportionately affect developing countries with inadequate data, low market visibility, and weak regulatory frameworks. The use of machine learning (ML) and deep learning (DL) methods for time-series anomaly detection has proven effective yet their operation relies on a 'black box' framework which fails to deliver understandable evidence needed for effective policy development. The development of Neuro-Symbolic Artificial Intelligence (NSAI) as a hybrid system unifies neural network pattern recognition with logical reasoning capabilities provides an effective solution to the problem of interpretability. The research paper conducts a systematic scoping review which includes 47 studies that the authors found through their structured search of IEEE Xplore Scopus Web of Science and Google Scholar between 2020 and 2025. We examine the theoretical underpinnings of NSAI technology together with existing methods for detecting anomalies in agricultural price data and we evaluate the capability of NSAI systems to function in African agricultural markets with limited resources. The research synthesis identifies four main architectural designs which enable NSAI-based anomaly detection. The research findings show that no research study has tested an integrated NSAI system for detecting crop price anomalies in sub-Saharan African countries. The research shows that NSAI systems provide better explainability in cyber-physical and financial systems but they need special design for particular application areas. The primary obstacles to implementation arise from data shortages, challenges in knowledge engineering, and limitations in computational power. The review establishes an organized research framework together with an architectural reference system which researchers can use to conduct studies in this developing research area.

Keywords: Neuro-symbolic artificial intelligence; anomaly detection; agricultural price monitoring; explainable AI; Signal Temporal Logic

INTRODUCTION

Crop price volatility, the unexpected changes in market prices, is risky to food production in developing countries. These fluctuations also affect the farmers who are the primary stakeholders in agriculture, especially those at the small scale level. Substantial price changes, due to the shocks, occur within a small time period, causing disruption of the food security and creating economic distress (Abubakar Ahmed Sambo et al., 2025). Nigeria, the largest sub-Saharan population of the continent itself and characteristically, the greatest producer of rice, cassava, maize, millet, has had increasing supply costs owing to predominantly high inflation rates and adverse effects of items' depreciation levels vis-à-vis foreign exchange, weather, a break in the supply chain, and ragtag boundary disputes (X. Wang et al., 2022).

The ability to recognise and combat these irregularities in real time is essential for rapid, informed policy actions such as releasing strategic reserves, adjusting import quotas, or providing local subsidies. Prices in agricultural markets are very complex, and thus ordinary statistics such as ARIMA and regression are not applicable. This places ML and DL techniques, including Long Short Term Memory (LSTM) networks, autoencoders (AE), and isolation forests, among the best detection methods. However, black-box systems make their application difficult, as they produce statistical flags but do not provide the explanations of the causes required by policymakers and all concerned market agencies (Gao et al., 2021).

A new model of computing called neural-symbolic computing represents an interesting case of deep learning's interaction with formal logic. The model aims at explaining its results in theoretical terms and providing a clear understanding of the basis for its decisions and for the data it uses for making decisions and predictions, thus making the model more interpretable and transparent. Although all works cited in books use NSAI for cyber-physical systems, financial services, and earth science, no one has ever attempted to implement an encompassing NSAI in their lab to study crop price anomalies peculiar to Sub-Saharan Africa. This state of affairs has driven all the developments on this subject described in this article. This paper makes four specific contributions:

This study presents four new contributions to the field, which include a PRISMA-ScR compliant systematic review that analyses 47 articles about neuro symbolic AI and anomaly detection and crop price monitoring and explainable AI and a taxonomy that defines four essential neuro symbolic architectural paradigms which researchers assessed through their application to Nigeria crop markets with limited data access and a complete neuro symbolic framework that integrates LSTM-CNN neural backbones with STL and HSTL based semantic reasoning and neural symbolic fusion and a research infrastructure which enables scientists to study deployment obstacles while they follow a four-stage research framework that supports future research development.

The paper proceeds as follows. Section II describes the review methodology. Section III presents the theoretical background. Section IV reviews ML and DL methods for detecting agricultural price anomalies. Section V analyses NSAI architectural patterns. Section VI presents the proposed framework specification. Section VII discusses challenges. Section VIII synthesises findings and proposes the research agenda. Section IX concludes.

REVIEW METHODOLOGY

Review design

This inquisition will use a controlled study design in line with the PRISMA-ScR (Preferred Reporting Items for Systematic Review and Meta-Analysis: Scoping Review) guidelines recommended by (Miranda et al., 2023). It is important to note that a scoping review is necessary because it focuses on mapping the existing literature and defining the body of knowledge, rather than synthesising quantitative findings from numerous studies. It must be noted that the review plan was completed before the data was collected.

Search Strategy

A broadband Boolean search was carried out through four different databases covering the five-year scope of the research activities – IEEE Xplore, Scopus, Web of Science (Core Collection) and Google Scholar for 2020-2025 publications. The leading search string was:

("neuro-symbolic" OR "neural-symbolic" OR "hybrid AI") AND ("anomaly detection" OR "outlier detection" OR "price monitoring") AND ("crop price" OR "agricultural price" OR "commodity price" OR "food price")

Supplementary searches incorporated terms including 'explainable AI agriculture,' 'LSTM crop price,' 'knowledge graphs agricultural markets,' and 'symbolic reasoning food security.' Forward and backward citation chaining was performed on all included studies.

Inclusion and Exclusion Criteria

Table 1. Inclusion and Exclusion Criteria for Study Selection

Criterion	Inclusion	Exclusion
Scope	Anomaly detection in agricultural/commodity price data, or NSAI/STL/HSTL, is applied to time-series anomaly detection in adjacent domains.	Studies unrelated to anomaly detection, price monitoring, or AI/ML methods
Language	English-language peer-reviewed journals and conference papers	Non-English publications; grey literature; theses; reports
Date	Published 2020–2025	Published before 2020 (except foundational theoretical works)
Type	Empirical studies, systematic reviews, and conceptual frameworks with clearly defined NSAI components	Abstracts only; editorial letters; studies lacking methodological transparency

Study selection and data extraction

The preliminary investigation unearthed approximately 312 records. After title and abstract scrutiny, this number dropped to 89 and these were subjected to full reading evaluation. Result of application of inclusion and exclusion criteria gave us 47 studies for EXCLUSION. In the case of this review, two investigators performed the determination of eligibility, the discrepancies were resolved through discussion. The data extracted covered study designs, AI architectures, application areas, characteristics of datasets, types of technologies used in the anomaly detection process, tools used for explainability purposes, and the geographical area of focus of each map.

Theoretical background

Neuro-Symbolic Artificial Intelligence

Neurally Symbolic AI (NSAI) refers to AI systems that combine neural network computing together with symbolic logical reasoning. The scope of NSAI, according to (Garcez & Lamb, 2023)falls in the era of the third wave of AI research, evolving from first-generation expert systems and second-generation deep learning models, and they insist that perception and reasoning are what make real, general intelligence. Such an approach divides the third wave systems into five integration strategies (Arachchige et al., 2025).

The study identifies four essential neurosymbolic architectural frameworks which follow these two architectural patterns. Tightly coupled systems use a joint learning method, which enables symbolic elements and neural systems to work together without separating their reasoning functions from their perceptual capabilities.

Modular systems enable different functional elements to work together through established communication paths, which maintain system transparency while providing operational independence. Symbolic [Neuro] systems use a symbolic reasoning system which assigns visual perception tasks to an internal neural network while maintaining control of the complete deduction process through symbolic logic. A neural model in Neuro:Symbolic→Neuro architectures uses a symbolic knowledge base for two functions. It uses domain rules to establish the foundations.

Neuro [Symbolic] systems use symbolic structures, which include predicate calculus representations that develop naturally inside neural networks. Logic Tensor Networks and DeepProbLog demonstrate this capability by creating a unified system that combines symbolic reasoning with neural computation.

Rashid & Kausik, (2024) focused their attention on agricultural markets. They disclosed three key capabilities in terms of Non-supervised Adaptive Intelligence improving direct marketing in agriculture: absorption of knowledge from few examples or generalisation, model adjustment to phenomenon evolution or rather acceptance of drift, and employment of relevant in-depth knowledge as constraints—factors that tackle especially well the issues of paucity of data, changing nature of the agricultural markets and availability of agri price information in LDCs.

Anomaly Detection: Taxonomy and Methods

Anomaly detection finds patterns that fall far outside a normal distribution when applied to time series data. In the temporal dimension, anomalies in time series are classified into: point anomalies, which are single data points that deviate widely, contextual anomalies, where an observation is abnormal in relation to the trends over time or seasonal cycles, and the final anomaly is where several data points form an abnormal pattern. These can present themselves as crop price anomalies fitting into any of the aforementioned types including a peak in price over a day (one day anomaly), prices seemingly too high or too low for a particular season (contextual anomaly) and a period of abnormal volatility covering the prices for a crop (collective anomaly).

Boukerche et al., (2021) differentiate anomaly detection models into types such as statistical, machine learning and deep learning describing distressed, and classification. These methods' applications to crop price data are documented in Table 2, and their limitations.

Table 2. Comparison of Anomaly Detection Methods Applied to Agricultural Price Data

Method Class	**Representative Approaches**	**Strengths**	**Limitations**	**Interpretability**
Statistical	ARIMA, Regression	Interpretable; theoretically grounded	Assumes linearity; fails on non-stationary data	High
Proximity-based	Isolation Forest, LOF	Unsupervised; scalable	Sensitive to dimensionality; no causal reasoning	Low
Deep Learning	LSTM, Autoencoder, GRU	Captures complex temporal patterns	Black box; requires large labelled datasets	Very Low
Neuro-Symbolic	LTNs, DeepProbLog, ECATS	Pattern recognition with causal reasoning	Knowledge engineering overhead; computationally expensive	High

LOF = Local Outlier Factor; GRU = Gated Recurrent Unit; LTN = Logic Tensor Network.

Causal Drivers of Agricultural Price Anomalies

Emerging agriculture markets are very complex and in most cases there are several issues which create unseen problems. Many countries in African continent are and still face a number of challenging factors in their rural areas. (Okoka & Kheswa, 2025) breaks out five drivers in this regard. The first of them is climatic shocks beginning with droughts, floods, and erratic rainfall, which reduce the quantity of crops available for sale and hence increase the prices. For instance, in Nigeria, the 2022 floods in Anambra, Kogi, Niger, etc. led to a spike in cassava price by almost 47% after about 2 months ((Musa et al., 2025).

Nigeria crop markets face challenges with price anomaly detection because their five core structural issues need to be solved. The first issue exists because supply chain vulnerabilities create supply shortages when roads deteriorate, fuel prices fluctuate, and logistical problems develop (Ahmed et al., 2025).

The second issue originates from policy changes that result in trade embargoes and tighter import controls, and sudden subsidy changes, which UNCTAD's Economic Policy and External Trade Division identified as events that disrupt markets without regard to their essential economic conditions (United Nationt Trade & Development, 2025). The third issue involves currency dynamics because people make international transactions using USD and Euro but actual transactions occur with Nigerian Naira and Botswana Pula which have limited market liquidity (IBOSIOLA & AREGHAN, 2025).

The fourth issue arises from market structures which create agency problems because they use multiple levels of intermediation and informal taxation to create different levels of risk that disrupt price transmission paths and make standard pass-through analysis untrustworthy (H. Wang, 2024). In this context and with regard to this aspect, these elements work and determine themselves depending on both the intensity of the stock and the rules of the game. Such synergies can lead to undesired results, as it is illustrated in the figure below.

Machine Learning and Deep Learning Methods for Agricultural Price Anomaly Detection

Deep Learning Architectures

The primary architecture used for agricultural price time series analysis remains LSTM networks. The Indian commodity data analysis using LSTM-autoencoder framework achieved 18.3% lower mean absolute error compared to ARIMA forecasting (Xie et al., 2023).

The autoencoder unsupervised method enables anomaly detection through price pattern reconstruction because it eliminates the need for training data with labels, but it produces incorrect results during seasonal changes, which results in misidentifying small price movements as large deviations from normal patterns and needs advanced assessment methods.

Hybrid GRU-LSTM models have demonstrated meaningful gains when incorporating multivariate inputs such as weather, input costs, and political variables (Kausar et al., 2024), while Transformer-based architectures including Temporal Fusion Transformers and PatchTST have advanced multivariate forecasting more broadly, yet their application to agricultural price anomaly detection specifically remains underdeveloped (Raja & Edwin, 2025).

Limitations of Purely Data-Driven Approaches

Four universal constraints are identified across the 28 ML and DL studies reviewed:

a. The existing studies fail to establish causal links which explain the discovered anomalies because their detection methods operate solely through statistical indicators which offer no useful guidance to policymakers.

b. African datasets present high performance challenges for models which operate on data from both advanced and Asian markets because these datasets contain shorter and less detailed yet more incomplete information.

c. Developing-region markets face concept drift vulnerability because their structural changes which include monetary crises and political transitions and climate regime shifts disrupt stationarity assumptions which lead to fast model performance decline.

d. The reviewed ML and DL research studies all fail to use domain knowledge because they do not incorporate agricultural calendars and supply chain networks and crop-specific seasonal patterns as prior knowledge structures.

e. These limitations collectively motivate NSAI architectures that integrate domain knowledge and logical reasoning with neural pattern recognition.

NSAI Architectural Patterns for Agricultural Price Anomaly Detection

Fig. 1: NSAI Architectural Patterns for Agricultural Price Anomaly Detection

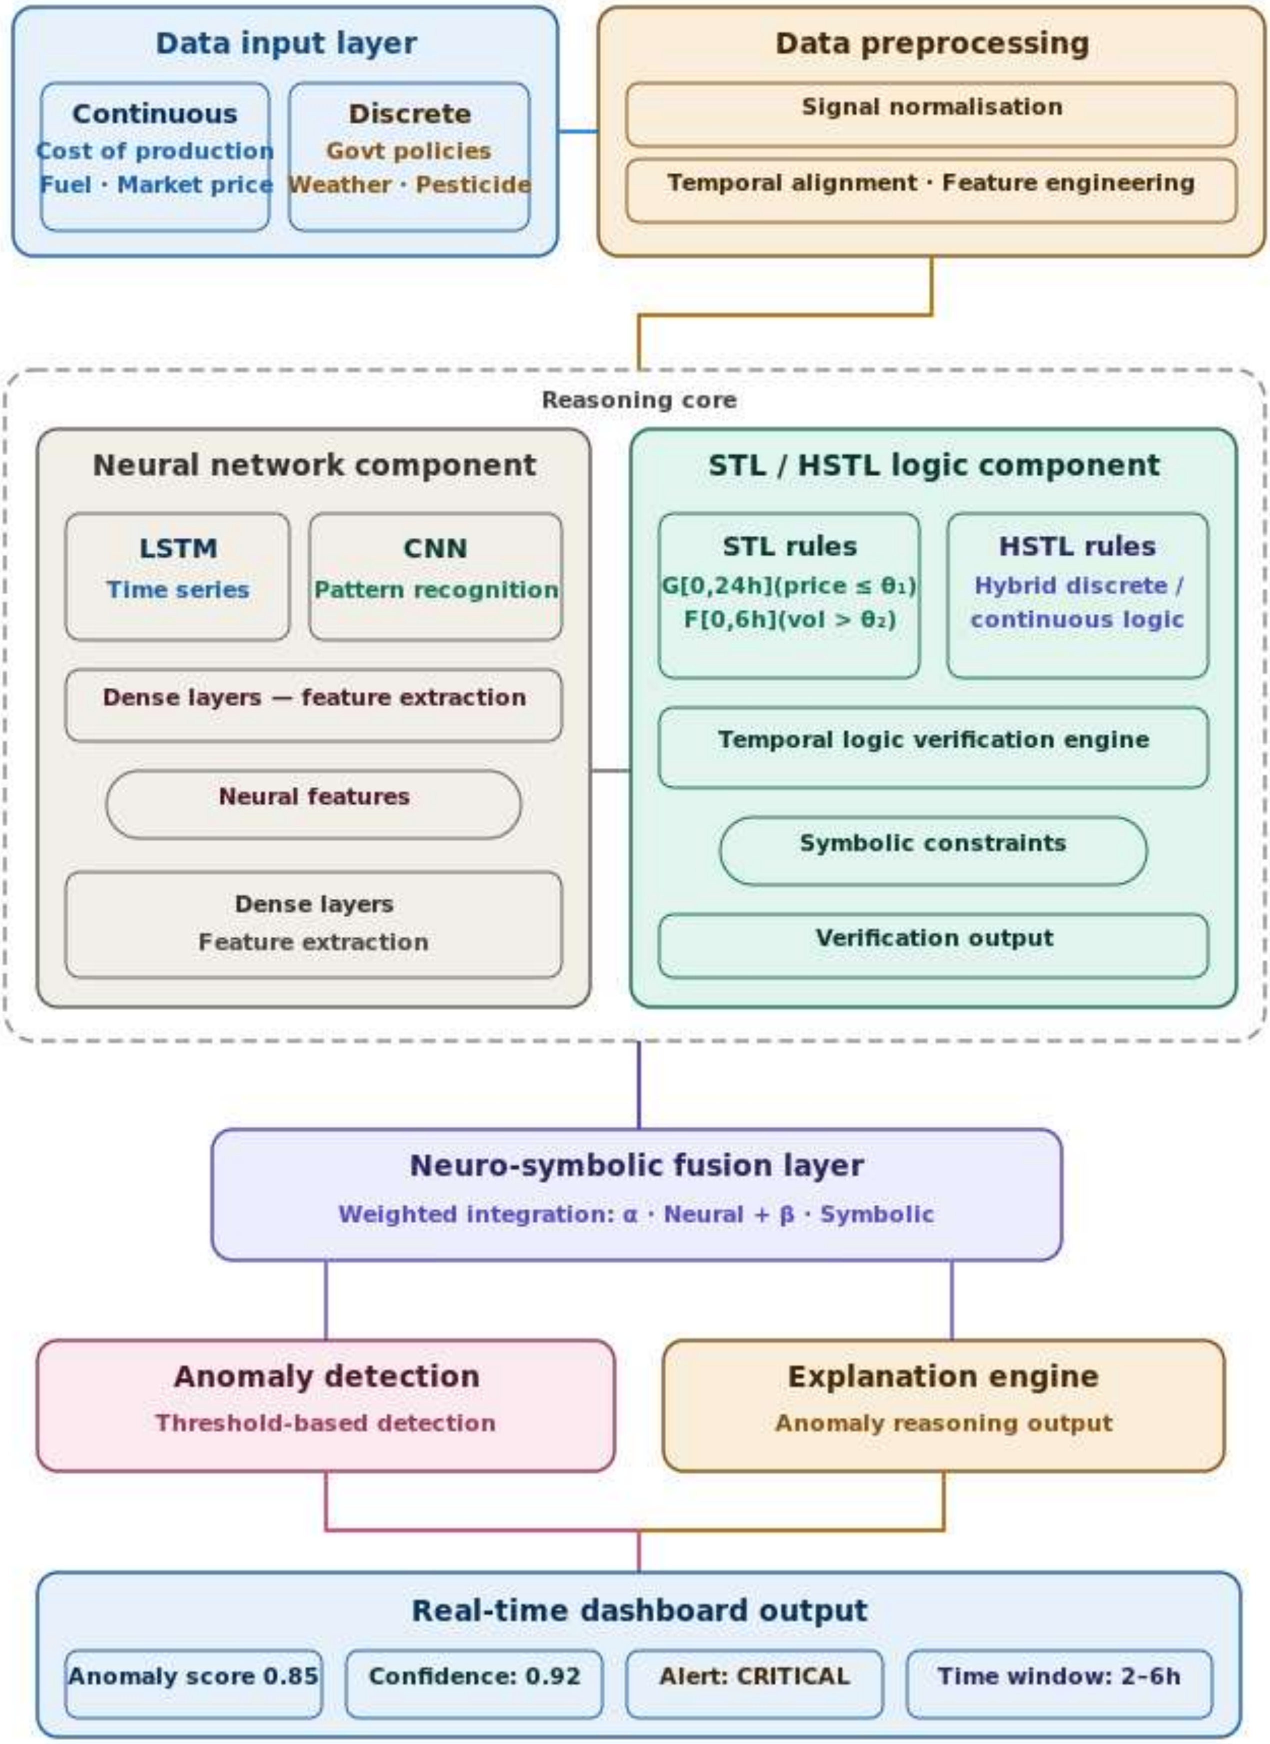

The proposed architecture of Fig. 1, begins by ingesting two streams of input data; continuous variables (such as production costs and fuel prices and market prices) together with discrete variables (government policies and weather conditions and pesticide usage), because these elements require preprocessing through signal normalisation, and temporal alignment and feature engineering to achieve consistency before analysis. The cleaned data enters a dual-component reasoning core, which includes a neural network component that uses an LSTM to capture long-range temporal dependencies in price time series while a CNN handles local pattern recognition through its dense layers which create compact neural features from incoming signals. The STL/HSTL logic component uses formal temporal logic rules to create symbolic constraints and a verification output which is based on domain knowledge instead of learned statistics.

The neuro-symbolic fusion layer as illustrated in Fig 1, enables this architecture to achieve its unique strength because it uses a weighted integration formula to combine both components through the formula ($\alpha \cdot$ Neural + $\beta \cdot$ Symbolic), which produces a unified signal that combines data-driven pattern recognition with rule-based logical reasoning. The fused output serves two downstream engines, an anomaly detection module that uses threshold-based scoring to identify prices that deviate from expected behaviour and an explanation engine that creates human-readable reasoning for every detected anomaly, which directly responds to the black-box criticism that deep learning systems face. The final output is delivered through a real-time dashboard which displays an anomaly score, a confidence level, and an alert severity classification and the exact time period of concern, which provides market operators with a warning and a complete analysis of agricultural issues that they can understand and act upon.

Proposed Architectural Framework: Formal Specification

Data Entry, Data Preprocessing, and Time Formulation

The input data includes two types: continuous attributes representing production costs, fuel costs, and electricity costs, and discrete events such as policy shifts, weather changes, and pest outbreaks. The system's dual design separates regular price changes from sudden fluctuations, reflecting how markets tend to move steadily until unexpected shifts occur due to policy decisions or climate events. The preprocessing stage applies Z-score normalization to each input channel using the formula $z (x \mu) / \sigma$, which standardises the distribution of variables. This is followed by linking observations to their seasonal reference points and creating composite indicators through feature engineering, including price momentum, supply-demand ratios, and logistics loss indices, which cannot be fully captured by traditional non-linear transformations (Bohne et al., 2023; Ferfoglia et al., 2024).

Neural Network Component: LSTM–CNN Hybridisation

The memory state of the LSTM cell is updated as follows: $c_t = f_t \odot c_{t-1} + i_t \odot g_t$. Here the f_t gate is determined by $f\_t = \sigma(W^i \cdot [h_{t-1}, x_t] + b^i)$, where W^i , b^i are particular to LSTM layers. This feature selectively preserves far-reaching temporal relationships without requiring manual definition of time lags, which is vital for markets dominated by sequential demand events, dynamic policy or regulatory environments and/or wave-based forms of supply chain management in the many possible frequency domains. (Gu et al., 2024; Zhang et al., 2020).

The CNN layer applies one-dimensional convolutions $(X \times W)[t] = \Sigma_k X[t - k] \cdot W[k]$ to detect local multivariate motifs such as volatility bursts and co-movement patterns. Dense layers compress these representations into a unified neural feature vector $h_n \in \mathbb{R}^d$ passed to the fusion layer. The interpretability limitations of deep neural representations under distributional shift — particularly relevant in rapidly changing African market environments — motivate the addition of the symbolic reasoning component (Tadesse Bogale & Debela, 2024).

Symbolic AI Component: Signal Temporal Logic and Robustness Semantics

The symbolic layer operationalises Signal Temporal Logic (STL), which formally specifies properties of real-valued continuous-time signals using bounded temporal operators. Two core operators are employed:

a. $G[t_1,t_2](\varphi)$: asserts that predicate φ holds globally over interval $[t_1, t_2]$.

b. $F[t_1,t_2](\varphi)$: asserts that predicate φ is satisfied at some point within $[t_1, t_2]$.

ISSN No. 2454-6186 | DOI: 10.47772/IJRISS | Volume X Issue XXVI June 2026| Special Issue on Education

The quantitative robustness function $\rho(\varphi, s, t)$ yields a signed real value: $\rho > 0$ indicates satisfaction, $\rho < 0$ indicates violation, and $|\rho|$ measures the degree of each (Rodionova et al., 2026). This robustness semantics enables graded anomaly quantification rather than binary classification.

Representative STL formulae for the crop price domain include:

c. $G[0,72h](price < \mu + 2\sigma)$: price must remain within two standard deviations of the rolling mean over any 72-hour window under normal conditions.

d. $F[0,6h](fuel\_price > threshold) \wedge G[0,24h](production\_cost - \mu \geq 2\sigma)$: compound condition of a fuel spike followed by sustained production cost elevation.

e. $G[0,4month](weather\_alert = \text{'drought'}) \rightarrow G[1h,48h](market\_price \geq 5\%)$: drought alert should precipitate market price increases within 48 hours.

f. $(policy\_subsidy) \rightarrow (\Diamond[0,2h](cost\_reduction) \wedge G[2h,24h](price\_volatility > \beta))$: subsidy announcement should reduce costs and affect price volatility within a defined window.

Hybrid Symbolic Temporal Logic (HSTL) extends STL with continuous fuzzy-valued predicates, enabling graded violation quantification for soft-boundary anomalies as seen in Fig 2, particularly useful for continuous-scale constructs such as drought severity or logistics friction (Ahmadi et al., 2025). A Temporal Logic Verification Engine evaluates these formulae in real time against incoming data streams, producing both Boolean and graded constraint outputs.

The simulation focuses on 200 time steps that affect crop commodity data. It was able to identify an important group of anomalies starting from time step t=135–150 is identified as a critical anomaly which is a simultaneous increase in the price of fuel, in the country and in production cost combined with a decrease in market price during this particular time. The composite score ultrahighs at ~0.85 well exceeding the 0.65 critical level, thus activating a CRITICAL message, beautifully explaining the cause of such anomalies to specific STL constraint violation and HSTL policy-weather trend interactions.\

Fig. 2a: Data input layer

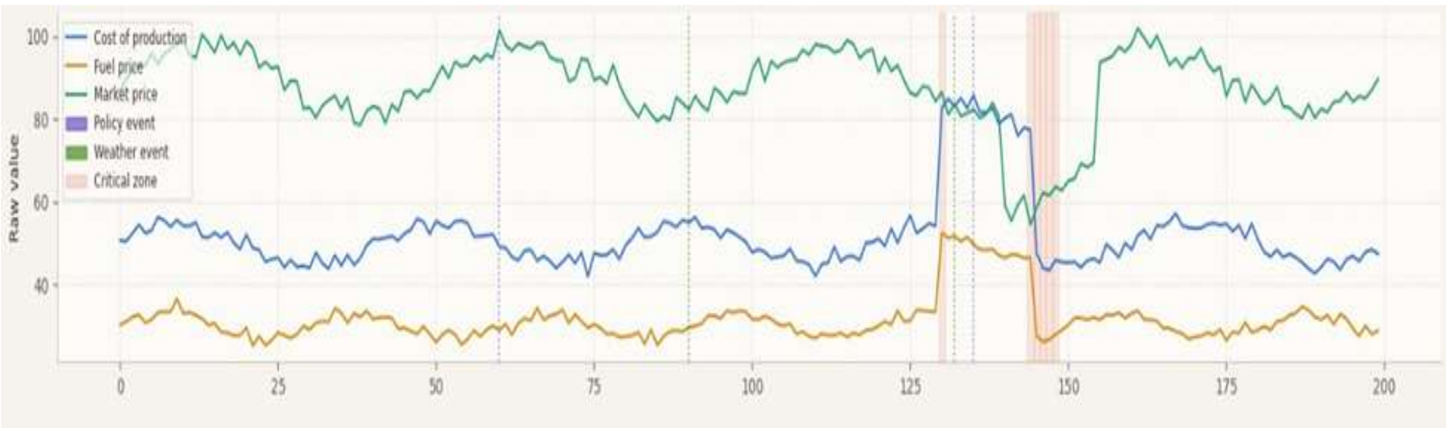

Figure 2b: Preprocessing layer

INTERNATIONAL JOURNAL OF RESEARCH AND INNOVATION IN SOCIAL SCIENCE (IJRISS)
ISSN No. 2454-6186 | DOI: 10.47772/IJRISS | Volume X Issue XXVI June 2026| Special Issue on Education

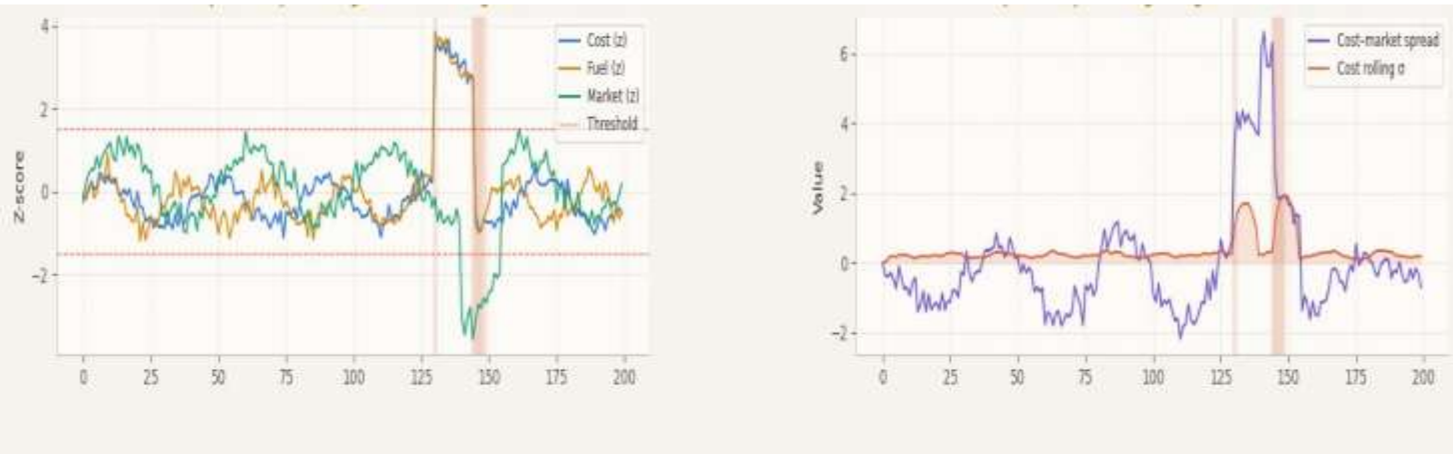

Neuro-Symbolic anomaly detection pipeline

Fig. 2: Python output to test run Fig. 1 diagram

Figure2c: Neural network (CNN and LSTM) and STL/HSTL processing layer

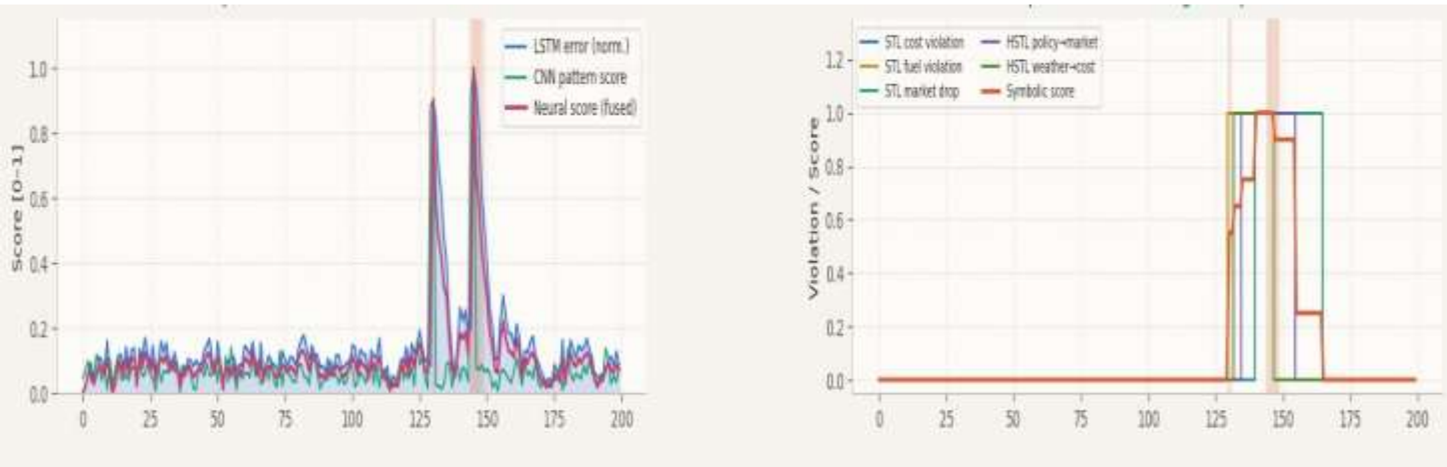

Figure 2d: Neuro Symbolic AI fusion layer (αNeural + βsymbolic)

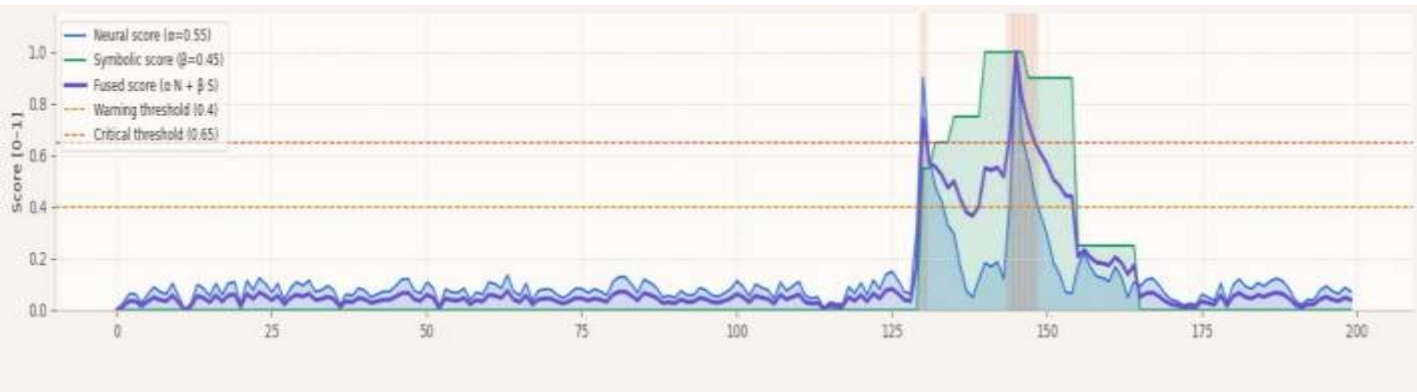

Figure 2e: Neuro symbolic real time anomaly detection thread

INTERNATIONAL JOURNAL OF RESEARCH AND INNOVATION IN SOCIAL SCIENCE (IJRISS)
ISSN No. 2454-6186 | DOI: 10.47772/IJRISS | Volume X Issue XXVI June 2026| Special Issue on Education

Neuro-Symbolic Fusion, Anomaly Scoring, and Explanation Generation

The blend layer conjoins the outputs of both neural and symbolic components according to the formula: $A = \alpha \cdot \hat{H}_n + \beta \cdot \rho_s$, where $\hat{H}_n \in [0,1]$ is the scaled neural-estimated anomaly strength, and ρ_s is the Scaled STL robustness score while both $\alpha + \beta = 1$ in Fig 2. Furthermore, the factors α and β are variable, wherein low α weights are suitable especially in cases needed low neural confidences; and large β values encourage poising, when there is a well-constructed knowledgebase.

A representation of results will have A equalling 0.85 (HIGH) with a neural net of 0.78 confidence and a symbolic system of 0.92. This implies higher constraints violation in the symbolic systems - in not reducing production costs putting an unbearable inflation of fuel costs aside - than in the neural due to the scarcity of the anomaly –ographic or otherwise – types in the learning set. Under the binary classification, an anomaly flag $\hat{y}$ is assigned as $\theta(A - \beta)$ where the thresholding parameter δ is determined, with caution to be exercised whenever the market condition is perturbed by switching fault (Liu et al., 2023).The Explanation Engine generates ante-hoc, STL-grounded causal attributions structured into three anomaly categories:

1. Spike anomalies: attributed to cost-side drivers such as fuel price surges violating the STL fuel-cost constraint and sustained production cost escalation.

2. Policy impact anomalies: linked to subsidy announcements or tax policy changes triggering price discontinuities, identified through the policy_subsidy → (◇[0,24h] (cost_reduction) ∧ G[24h,72h] (price_volatility > β)) formula.

3. Weather-cost correlations: drought-induced cost escalation and pesticide demand spikes captured through the HSTL formula G[0,4month] (weather_alert = 'drought') → G[1month,2month] (market_price ≥ 5%).

These attributions are structurally faithful to the model's decision mechanism, unlike post-hoc methods such as SHAP or LIME, which approximate explanations externally (Arrieta et al., 2022). (Mujkic et al., 2022) described the 'Real-Time Dashboard' consolidates anomaly score, component confidence values, alert level, and causal explanation into a format accessible to non-technical stakeholders including agricultural officers, traders, and policymakers (Praveen Kumar Manchikoni Surendra, 2025).

Table 3. Comparison of NSAI Architectural Patterns for Crop Price Anomaly Detection

Pattern	Complexity	Explainability	Detection Performance	Knowledge Integration	Data-Scarce Suitability
1. Sequential Pipeline	High	Moderate	Low–Moderate	Post-hoc only	Yes (entry-level)

2. Constraint-Guided	Moderate	High	Moderate	Strong (a priori)	Yes (recommended)
3. Coupled Bidirectional	Very High	High	High	Very Strong	Conditional
4. KG-Augmented	High	Moderate–High	Moderate	Strong (graph-based)	Promising (long-term)

Table 3 shows that the four neuro symbolic architectural patterns create a performance-explainability trade-off which results in Coupled Bidirectional design producing its best detection results through knowledge integration but requires extensive resources and only works in specific African markets. The Constraint-Guided pattern serves as the most practical design solution for agricultural research conducted under limited data conditions because it provides agricultural researchers with essential a priori knowledge while maintaining clear explanation of its functions and operating at moderate difficulty level.

Explainability Mechanisms

Three principal explainability mechanisms applicable to agricultural price monitoring are identified:

1. Rule-based attribution: The symbolic layer fires logical rules whose antecedents constitute the explanation — for example, IF rainfall < 80 mm AND logistics_cost_index > 1.4 THEN price_deviation = HIGH, attributed to [drought, fuel_cost_surge]. Directly interpretable but requires comprehensive a priori rule engineering.

2. Concept-based explanation: As in ECATS (Ferfoglia et al., 2024), neural networks detect deviations in interpretable concept dimensions defined by domain experts. In agricultural contexts, relevant concepts include 'seasonal_supply_adequacy,' 'logistics_friction,' and 'policy_stability.'

3. Causal graph reasoning: NSAI systems incorporating causal knowledge graphs trace anomalies through causal pathways, generating multi-step chains — for example, price_spike ← supply_shortfall ← climate_shock ← El_Niño_event. The most informative for policy purposes, but requiring validated causal models of agricultural market dynamics.

Challenges and Deployment Considerations

Data Scarcity and Quality

The 47 reviewed studies consistently identify data scarcity as the primary obstacle to NSAI deployment in African agricultural contexts. Price data in sub-Saharan Africa are collected at irregular intervals, typically weekly to monthly. National datasets experience data gaps of up 35% in different time intervals. Fundamentally, our concern is that the data is collected at the reference market and does not extend to the full distribution channels. Additionally, although factors remain constant across a data collection period, their meaning and scope can differ. Pattern 2 in NSAI's CGL approach to FDL, still, attempts to ameliorate this situation by constraining the learning algorithm such that it takes domain knowledge into consideration through the specification of regularisation constraints thereby reducing the required minimum amount of training data (Triggs et al., 2025).

Knowledge Engineering Complexity

Developing robust symbolic knowledge bases for African agricultural markets requires sustained interdisciplinary collaboration between AI researchers, agricultural economists, and domain practitioners. Established knowledge engineering methodologies such as CommonKADS and Protégé-based ontology development were not designed for low-resource participatory settings, and require significant adaptation. Causal knowledge graphing adds further complexity.

Computational Infrastructure

Pattern 3 (Full Model) includes NSAI architectures that are too dependent of the training process that involves a GPU, which is not always possible in Africa. The other two – Pattern 1 with a sequential pipeline and Pattern 2 with constraints that apply to a dynamical system, are capable of working on weaker computational resources and should allow feasible near-term deployments. In the cases of edge computing and cloud-API access, one can design an architecture even with a high computational cost in applications where the resources are limited. (Imani et al., 2025).

Interpretability-Accuracy Trade-offs

Tighter integration of symbolic constraints can reduce predictive accuracy relative to unconstrained deep learning, particularly where the symbolic knowledge base is incomplete or inaccurate. This interpretability-accuracy trade-off is well documented in the XAI literature (Saeed & Omlin, 2023). It has direct implications for agricultural deployment: a system generating incorrect but confident causal explanations may be more damaging to policy trust than one providing no explanation at all. Robust empirical validation of symbolic knowledge bases is therefore an essential prerequisite for deployment.

Institutional and regulatory considerations

The success of its adoption beyond introducing demonstratable effectiveness also warrants other non-technical operational conditions for instance there has to exist clear laws enabling use of AI in surveillance of the market and its actors research data governance protecting farmer's information and ensuring equal rights to data, training of staff and policy-makers among others. Also, users should be able to procure open-source or locally-developed solutions through the usual public procurement processes. (Nawaz et al., 2025)

Synthesis and Research Agenda

Summary of Key Findings

This systematic scoping review yields four principal findings:

1. Research Gap Vindicated: A survey of existing literature has revealed that no experiment has ever been conducted during which a comprehensive NSAI approach was utilized in pricing and trading of agricultural commodities in any sub-Saharan nation. There are other NSAI based systems that pretty close in use, such as in cyberspace (Ferfoglia et al., 2024), or in geoscience (Chen et al., 2024), where existence of such solutions is evident and have managed almost similar tasks. However, as far as agricultural commodities are concerned, in particular African agricultural crops, no evidence is available at all.

2. Architectural Synergies: simpler and practically suitable methods of deploying the architectures to countries and regions in Africa where data is hard to come by as they require milder computing resources and have shown functionality even with minimal training samples.

3. Interpretable Value: Developing countries are directly looking for understandability in the agricultural economics domain, presenting the importance of explainable AI in policy engagement and the lack of well-developed operational acceptance standards for explanations, illustrating that the demand for formal evaluation metrics regarding them is still unmet.

4. Evaluation Thresholds: % of data titled analyzed, percentage of text cased analyzed, degree of Block or graph structure analyzed, percentage of cases analyzed, businesses analysis addressed, user does not enter personal data or submits them throughdicated forms, eligibility reviews address the title, source of information disclamer has been reviewed, script memos are coded or created for the full feature analysis, description of appendices in has been entered and reviewed, benchmarking analysis has been conducted in terms of comparison with the external environment and the goal of the analysis, pre-requisite reviews

of published information has been made, The extent of data available for counting calculations is inadequate, published information has not been used as a cover for theft or fraud and any disputes are pursued within the established legal systems.

Key Limitations of the Proposed Framework

The proposed framework has several limitations requiring future attention. STL formulae require manual specification by domain experts — a resource-intensive process subject to knowledge elicitation biases. No explicit procedure is provided for optimising fusion weights α and β; Bayesian weight calibration (Snoek et al., 2015) offers a principled approach for empirical implementations. Scalability under high-frequency data streams remains unresolved. The framework has not yet been benchmarked against established baselines on African price datasets. These limitations define the priorities of the research agenda below.

CONCLUSION

This paper has presented a systematic scoping review and formal architectural proposal at the intersection of neuro-symbolic AI, time-series anomaly detection, and agricultural price monitoring, with specific focus on sub-Saharan Africa. Through structured synthesis of 47 eligible studies, the review confirms that NSAI has substantially improved explainable anomaly detection across adjacent domains while no study has yet applied such technology to agricultural price volatility detection in African markets.

Four NSAI architectural patterns are identified and comparatively evaluated. Sequential pipeline and constraint-guided learning architectures are the most suitable for near-term implementation in data-scarce African contexts, offering interpretable anomaly detection within realistic computational and data constraints. The formal architectural specification developed in Section VI — incorporating an LSTM–CNN neural backbone, an STL symbolic reasoning layer, and an adaptive neuro-symbolic fusion mechanism — provides a reference blueprint for future empirical implementations.

The principal barriers to deployment — data scarcity, knowledge engineering complexity, and computational infrastructure limitations — are addressed through the four-stage prioritised research agenda proposed in Section VIII. Agricultural price distortions in Nigeria and sub-Saharan Africa broadly have real consequences for food security, agricultural development, and macroeconomic stability. Building AI systems that accurately detect these distortions and explain them in a manner accessible to policymakers, traders, and farmers represents a meaningful contribution to both AI research and sustainable development. This paper provides the conceptual and methodological foundations for advancing this agenda.

REFERENCES

1. Abubakar Ahmed Sambo, Rilwan Ali Zira, Aishatu Yahaya, & Sunusi Abdulhamid Dantata. (2025). Poultry Disease Detection Using Machine Learning. Journal of Scientific Development Research. https://doi.org/10.70382/hujsdr.v8i9.005
2. Ahmadi, E., Mohammadi, P., Rafei, F. M., & Sana, S. S. (2025). Behavioral factors and capital structure: identification and prioritization of influential factors. Data Science in Finance and Economics, 5(1), 76–104. https://doi.org/10.3934/DSFE.2025005
3. Ahmed, Z., Khattak, S. N., & Rabbani, A. A. (2025). Resilient Supply Chains in the Post-Pandemic Era: Strategies for Risk Mitigation and Continuity. Inverge Journal of Social Sciences, 4(4), 31–44. https://doi.org/10.63544/ijss.v4i4.177
4. Arachchige, P. J., Iancu, B., & Lilius, J. (2025). A Roadmap Toward Neurosymbolic Approaches in AI Design. IEEE Access, 13, 174368–174392. https://doi.org/10.1109/ACCESS.2025.3617771
5. Arrieta, A., Battelino, T., Scaramuzza, A. E., Da Silva, J., Castañeda, J., Cordero, T. L., Shin, J., & Cohen, O. (2022). Comparison of <scp>MiniMed 780G</scp> system performance in users aged younger and older than 15 years: Evidence from 12 870 real-world users. Diabetes, Obesity and Metabolism, 24(7), 1370–1379. https://doi.org/10.1111/dom.14714

6. 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
7. Boukerche, A., Zheng, L., & Alfandi, O. (2021). Outlier Detection. ACM Computing Surveys, 53(3), 1–37. https://doi.org/10.1145/3381028
8. Chen, W., Ma, X., Wang, Z., Li, W., Fan, C., Zhang, J., Que, X., & Li, C. (2024). Exploring neuro-symbolic AI applications in geoscience: implications and future directions for mineral prediction. Earth Science Informatics, 17(3), 1819–1835. https://doi.org/10.1007/s12145-024-01278-7
9. Ferfoglia, I., Saveri, G., Nenzi, L., & Bortolussi, L. (2024). ECATS: Explainable-by-Design Concept-Based Anomaly Detection for Time Series (pp. 175–191). https://doi.org/10.1007/978-3-031-71170-1_16
10. Gao, C., Lei, W., He, X., de Rijke, M., & Chua, T.-S. (2021). Advances and challenges in conversational recommender systems: A survey. AI Open, 2, 100–126. https://doi.org/10.1016/j.aiopen.2021.06.002
11. 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
12. Gu, J., Zhang, B., Wang, Y., & Zhang, Y. (2024). Hedge three-dimensional reconstruction and motion control technology for trimming robot. Computers and Electronics in Agriculture, 227, 109632. https://doi.org/10.1016/j.compag.2024.109632
13. IBOSIOLA, A., & AREGHAN, I. (2025). EXCHANGE RATE MANAGEMENT AND CAPITAL MARKET PERFORMANCE IN NIGERIA. International Journal of Social Sciences and Management Review, 09(01), 15–33. https://doi.org/10.37602/IJSSMR.2025.9102
14. Imani, M., Joudaki, M., Beikmohamadi, A., & Arabnia, H. R. (2025). Customer Churn Prediction: A Review of Recent Advances, Trends, and Challenges in Conventional Machine Learning and Deep Learning. https://doi.org/10.20944/preprints202503.1969.v2
15. Kausar, F., Deo, S., Hussain, S., & Ul Haque, Z. (2024). Federated Deep Learning Model for False Data Injection Attack Detection in Cyber Physical Power Systems. Energies, 17(21), 5337. https://doi.org/10.3390/en17215337
16. Liu, P., Zhang, L., & Gulla, J. A. (2023). Pre-train, Prompt and Recommendation: A Comprehensive Survey of Language Modelling Paradigm Adaptations in Recommender Systems.
17. Miranda, A. V., Lestari, B. W., Indrarini, A., Arsy, F. F., Sagala, S., Bisri, M. B. F., & Lucero-Prisno, D. E. (2023). Adaptation of health systems to climate change-related infectious disease outbreaks in the ASEAN: Protocol for a scoping review of national and regional policies. PLOS ONE, 18(6), e0286869. https://doi.org/10.1371/journal.pone.0286869
18. Mujkic, E., Philipsen, M. P., Moeslund, T. B., Christiansen, M. P., & Ravn, O. (2022). Anomaly Detection for Agricultural Vehicles Using Autoencoders. Sensors, 22(10), 3608. https://doi.org/10.3390/s22103608
19. Musa, H. A., Abdulraheem, A., Rajah, S. A., Abubakar, A. R., & Okoronkwo, I. (2025). Use Of Misoprostol in the Prevention and Treatment of Post Partum Haemorrhage Among Primary Health Care Workers in Zaria Metropolis, Kaduna State, Nigeria. Achievers Journal of Scientific Research, 35–45.
20. 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
21. Okoka, H., & Kheswa, J. G. (2025). Community Policing Practice in the Nigerian Police: Implications of Employee Burnout. Journal of Police and Criminal Psychology, 40(1), 91–102. https://doi.org/10.1007/s11896-024-09711-9
22. Praveen Kumar Manchikoni Surendra. (2025). Revolutionizing functional verification: The impact of AI and machine learning in chip design. World Journal of Advanced Research and Reviews, 26(1), 2484–2490. https://doi.org/10.30574/wjarr.2025.26.1.1318
23. Raja, D., & Edwin, J. J. (2025). Comparative Analysis of Time Series Transformers on Multivariate Time Series Data. Premier Journal of Science. https://doi.org/10.70389/PJS.100132

24. Rashid, A. Bin, & Kausik, M. A. K. (2024). AI revolutionizing industries worldwide: A comprehensive overview of its diverse applications. Hybrid Advances, 7, 100277. https://doi.org/10.1016/j.hybadv.2024.100277
25. Rodionova, L. A., Kopnova, E. D., Kamalendinova, S. S., & Rodionov, N. A. (2026). Forecasting the Birth Rate in Russia Based on Google Trends Search Query Statistics. Population and Economics, 10(2), 92–114. https://doi.org/10.3897/popecon.10.e169776
26. Saeed, W., & Omlin, C. (2023). Explainable AI (XAI): A systematic meta-survey of current challenges and future opportunities. Knowledge-Based Systems, 263, 110273. https://doi.org/10.1016/j.knosys.2023.110273
27. Snoek, J., Rippel, O., Swersky, K., Kiros, R., Satish, N., Sundaram, N., Patwary, Md. M. A., Prabhat, M., & Adams, R. (2015). Scalable Bayesian Optimization Using Deep Neural Networks. Statistics.
28. Tadesse Bogale, A., & Debela, K. L. (2024). Organizational culture: a systematic review. Cogent Business & Management, 11(1). https://doi.org/10.1080/23311975.2024.2340129
29. Triggs, A. O., Causer, J., McRobert, A. P., & Andrew, M. (2025). Perceptual-cognitive skills and talent development environments in soccer: A scoping review. PLOS One, 20(7), e0327721. https://doi.org/10.1371/journal.pone.0327721
30. United Nationt Trade & Development. (2025). Global Trade Update (March 2025): The role of tariffs in international trade. United Nationt Trade & Development.
31. Wang, H. (2024). Information Asymmetry and Agency Problems in the Financial Market. Highlights in Business, Economics and Management, 32, 62–66. https://doi.org/10.54097/2eq3j535
32. Wang, X., Zhou, K., Wen, J.-R., & Zhao, W. X. (2022). Towards Unified Conversational Recommender Systems via Knowledge-Enhanced Prompt Learning. Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, 1929–1937. https://doi.org/10.1145/3534678.3539382
33. Xie, Y., Wang, M., Zhong, Y., Deng, L., & Zhang, J. (2023). Label-Free Anomaly Detection Using Distributed Optical Fiber Acoustic Sensing. Sensors, 23(8), 4094. https://doi.org/10.3390/s23084094
34. Zhang, S., Yao, L., Sun, A., & Tay, Y. (2020). Deep Learning Based Recommender System. ACM Computing Surveys, 52(1), 1–38. https://doi.org/10.1145/3285029