A Simulation-Driven Neuro-Symbolic Explainable AI Framework with Blockchain-Enabled Validation for Trustworthy Precision Agriculture

Dr. Syed Ibad Ali [1], Dr. (Prof.) Shashi Kant Gupta [1]

[1]Lincoln University College, 47301, Petaling Jaya, Selangor Darul Ehsan , Malaysia.

pdf.syedibadali@lincoln.edu.my, raj2008enator@gmail.com

Abstract

Artificial intelligence (AI) is increasingly being used in precision agriculture; however, predictive accuracy alone is insufficient to ensure trustworthy agricultural decision-making. Dynamic field conditions, noisy and incomplete sensor data, agronomic constraints, and weakly connected post-hoc explanations can limit the reliability of AI-based recommendations. To address these challenges, this paper proposes PoISV-NSDT, a Proof-of-Integrity Smart Validation and Neuro-Symbolic Digital Twin framework that integrates multimodal data processing, a physics-informed procedural digital twin, reinforcement learning, constraint-guided neuro-symbolic reasoning, explainable AI, and permissioned blockchain validation. The framework initially performs temporal and contextual integration of soil, weather, crop, multispectral, and telemetry data using z-score normalization, similarity-aware KNN imputation, and stress-aware synthetic data enrichment. A domain-randomized Unity ML-Agents digital twin subsequently supports Proximal Policy Optimization (PPO) under diverse soil, crop, and weather conditions. The learned policy is integrated with a three-layer constraint-guided neuro-symbolic policy network that combines Temporal Fusion Transformer (TFT)-based temporal perception, lightweight temporal convolution for imagery, symbolic agronomic rules, knowledge-graph constraints, and an Arctic Puffin Optimization (APO) action tuner. For interpretability, SHAP-based feature attributions are combined with symbolic rule traces to generate a unified explanation of each decision. A permissioned blockchain further records cryptographic hashes associated with input data, model versions, outputs, and explanations to support auditability and provenance. The reported experimental study uses 1,200 season-level episodes with an 80:20 training-to-testing split. The framework reports an R^2 of 0.919, 2% agronomic rule violations, 92% robustness under sensor noise, 89.7% explanation fidelity under stress, and blockchain throughput of 82 transactions per second with latency below 60 seconds. It also reports an 18% increase in yield, 11% reduction in water consumption, and 7.5% reduction in fertilizer usage. Overall, the proposed framework establishes a system-level approach to trustworthy precision agriculture by integrating prediction, simulation, constraint enforcement, explainability, and verifiable provenance within a unified architecture.

Keywords

Precision agriculture; trustworthy AI; neuro-symbolic AI; digital twin; reinforcement learning; explainable AI; SHAP; blockchain; PPO; agronomic constraints; smart farming.

1. Introduction

Precision agriculture (PA) increasingly combines Internet-of-Things sensing, remote sensing, machine learning and automated decision support to optimize crop production and resource use. However, agricultural environments differ substantially from static benchmark datasets. Soil properties, crop phenology, weather, irrigation response and sensor behavior change over time, creating domain shift and uncertainty. A decision system may therefore be accurate on a historical test set while producing infeasible recommendations when field conditions change. Recent work on agricultural digital twins similarly emphasizes the value of integrating sensing, simulation and analytics, while also identifying challenges in interoperability and deployment [1].

The trust problem in AI-driven PA has at least four dimensions. First, data quality can degrade because of sensor noise, missing observations and drift. Second, purely statistical models do not inherently enforce agronomic requirements such as nutrient-water budgets or management thresholds. Third, post-hoc explainability can provide feature importance without establishing how domain rules contributed to the final recommendation. Fourth, model artifacts, simulations and explanations are difficult to audit when their provenance is not independently verifiable.

This paper addresses these issues through an integrated framework termed PoISV-NSDT. The framework follows the principle that trustworthy agricultural AI should be evaluated as a system rather than as a collection of individually accurate algorithms.

Its five stages are: multimodal data acquisition and robust preprocessing; domain-randomized digital-twin simulation and PPO policy learning; a constraint-guided neuro-symbolic policy network; fused SHAP and symbolic explanations; and permissioned blockchain-based proof-of-integrity validation.

The main contributions are threefold. First, a robust data and simulation pipeline is defined to improve resilience to missing data, outliers, sensor drift and climate variation. Second, a constraint-guided neuro-symbolic decision core combines temporal and image representations with explicit agronomic rules and multi-objective action optimization. Third, the framework links explanations and decision artifacts to cryptographic hashes on a permissioned ledger, enabling selective auditability.

2. Related Work and Research Gap

2.1 Hybrid AI and Deep Learning for Precision Agriculture

Modern PA systems increasingly use convolutional, recurrent, attention-based and reinforcement-learning architectures to model crop state and management actions. The source study considers representative approaches including CAViT, CNN-LSTM, DFE-RGI, SCAML, ABR and Transformer-GRU with reinforcement learning. These approaches can improve predictive performance, but the stated limitation is sensitivity to domain shift, irregular sampling and the lack of explicit agronomic constraint enforcement.

The proposed framework uses TFT as its temporal perception component. TFT was designed for multi-horizon forecasting with feature selection, gating and interpretable attention, making it suitable for heterogeneous temporal covariates [2]. PPO is used for policy learning in the digital twin because it provides a practical clipped policy-gradient mechanism for repeated optimization of a surrogate objective [3].

2.2 Explainability and Constraint-Aware Decision Making

Explainable AI methods can expose relationships between inputs and predictions, but a feature attribution by itself does not necessarily encode the agronomic rule that makes a recommendation acceptable or unacceptable. SHAP provides a unified framework for additive feature attribution and is used here as the neural explanation component [4]. The proposed architecture complements SHAP with explicit temporal-logic rules, agronomic thresholds and nutrient-water budget knowledge. Constraint violations are handled through localized repair instead of accepting an infeasible action or retraining the entire model.

Constraint-aware reinforcement learning has been studied as a way to incorporate safety requirements directly into policy optimization [5]. The present framework extends this idea operationally by placing symbolic validation between neural inference and action tuning: the optimizer can modify only actions that have passed symbolic checks.

2.3 Blockchain and Trust in Agricultural Systems

Blockchain has been widely investigated for agricultural traceability and provenance. Reviews identify transparency, traceability and tamper resistance as major motivations, while also noting challenges involving standardization, participation and interoperability [6], [7]. In contrast to supply-chain-only use, PoISV applies the ledger to AI decision provenance. Hashes bind inputs, model versions, neural outputs and explanations, while large artifacts are stored off-chain and referenced through their hashes.

2.4 Research Gap

The research gap addressed in this study is the absence of a unified workflow that simultaneously couples robust preprocessing, digital-twin simulation, constraint-guided neuro-symbolic learning, integrated explanation and on-chain integrity validation. The objective is not merely to increase prediction accuracy, but to create a decision pipeline in which feasibility, interpretability and auditability are first-class properties.

3. Proposed PoISV-NSDT Framework

PoISV-NSDT is organized as a closed decision-support loop. Physical and digital agricultural observations are transformed into conditioned state representations. The digital twin generates stress-tested trajectories and supports reinforcement-learning policy development. The learned policy is then interpreted and constrained by a neuro-symbolic core before actions are optimized. Finally, the decision and its supporting artifacts are converted into a verifiable provenance record.

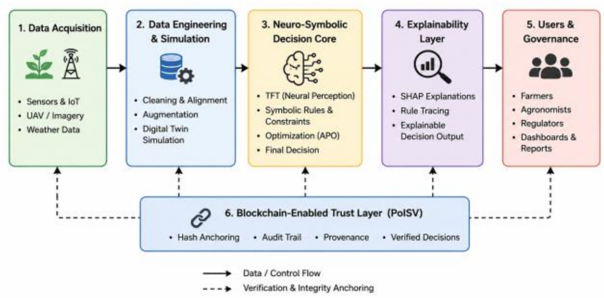

Fig. 1. Integrated PoISV-NSDT workflow reproduced from the supplied research presentation.

3.1 Multimodal Data Acquisition and Robust Preprocessing

The data layer integrates soil moisture, pH, electrical conductivity (EC), temperature, nitrogen-phosphorus-potassium (N-P-K), multispectral imagery, weather observations and telemetry. Temporal and geospatial alignment is applied before model ingestion so that measurements from different modalities refer to a consistent field state.

Z-score normalization is selected to reduce the effect of scale differences and to provide a more stable representation under outliers and sensor drift than direct min-max scaling. Missing observations are reconstructed using similarity-aware KNN imputation, with the objective of preserving local temporal trends rather than replacing missing values with a global mean or a simple forward-fill strategy.

Stress-aware synthetic enrichment combines conditional generative modeling with density-aware oversampling. Conditional GANs allow generated samples to be conditioned on contextual variables [8], while SMOTE provides the conceptual basis for synthetic minority generation [9]. In the proposed workflow, these methods are used to increase representation of growth-stage and stress conditions without deliberately amplifying isolated outliers.

Dataset provenance is maintained through version identifiers and cryptographic hashes. The hash becomes the reference point for subsequent blockchain validation, allowing a later auditor to determine whether the data used by a decision system corresponds to the registered dataset artifact.

3.2 Digital Twin Simulation and Policy Learning

The simulation layer is implemented as a physics-informed procedural crop-soil-weather environment using Unity ML-Agents. Unity-based simulation environments are suitable for configurable, interactive learning scenarios [10]. The digital twin represents field state through soil moisture, crop growth, weather variability and sensor observations, while the domain is randomized across soil parameters, crop phenology and weather conditions.

PPO is trained using lifecycle-controlled episodes. At initialization, field parameters are sampled from the domain space. At the beginning of each episode, crop, soil and climate states are reset. The observation vector contains multimodal state information together with short-term temporal context. Actions correspond to irrigation rate, fertilizer adjustment and management timing. The reward combines dense growth feedback with sparse yield and rule-related objectives. PPO clipped updates are then used to improve policy stability under randomized conditions.

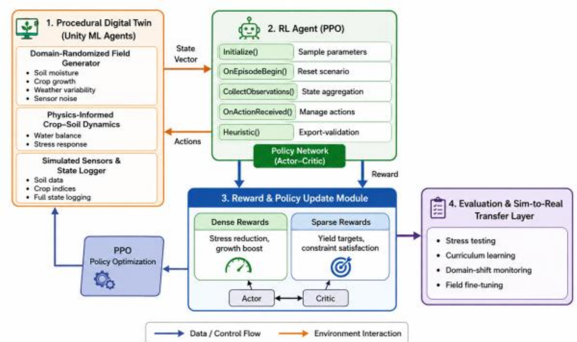

Fig. 2. Digital-twin-driven reinforcement-learning pipeline reproduced from the supplied research presentation.

3.3 Constraint-Guided Neuro-Symbolic Policy Network

The decision core follows a three-layer learn-reason-optimize architecture. The neural perception layer uses TFT for multivariate temporal information and a lightweight temporal convolution component, denoted LTCE in the source study, for imagery. The two representations are fused using confidence-weighted sigmoid gating so that the relative contribution of temporal and image evidence can change with the reliability of the observations.

The symbolic reasoning layer encodes temporal-logic rules, agronomic thresholds and nutrient-water budget knowledge. A candidate action is checked against these constraints before being accepted. When a violation is detected, the system performs localized constraint-guided repair instead of blindly accepting the action or retraining the entire neural model.

After symbolic validation, the Arctic Puffin Optimization (APO) action tuner searches the feasible action space. The source framework describes APO as combining Lévy-flight aerial exploration with underwater local refinement and elitist selection. Its multi-objective fitness is expressed conceptually as:

$$F(a) = YieldGain(a) - \lambda_1 Irrigation(a) - \lambda_2 Fertilizer(a) - \lambda_3 Stress(a) - \lambda_4 RuleViolation(a)$$

where a denotes a candidate management action and λ terms control the relative penalties. Because only symbolically validated actions are passed to the optimizer, the search process is restricted to a feasible region rather than optimizing an unconstrained objective.

3.4 Explainability Layer

The explanation layer combines SHAP attributions with symbolic rule traces. SHAP produces feature-level contributions for tabular inputs and image embeddings, while the rule engine records which agronomic constraints were satisfied, violated or used for repair. The two evidence streams are fused into a single decision narrative. This design is intended to reduce the disconnect between a post-hoc model explanation and the operational reason an action was selected.

3.5 Proof-of-Integrity Smart Validation

The PoISV blockchain component uses six role-segmented lanes: administrator, authorized user, blockchain client, smart contract, ledger and external auditor. A smart contract recomputes hashes, enforces validation conditions and can issue correction feedback. The append-only ledger records cryptographic references to inputs, neural outputs, explanations and model versions. Large artifacts are maintained off-chain in IPFS and referenced from the ledger by the hashes.

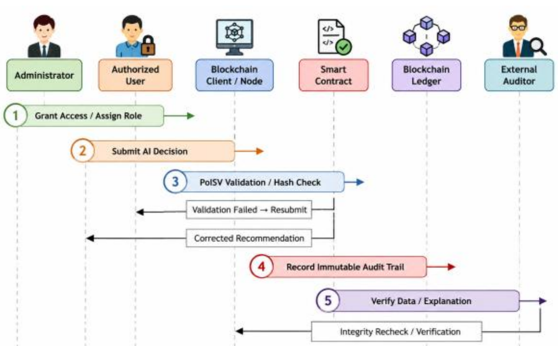

eir

Fig. 3. Role-segmented PoISV validation workflow reproduced from the supplied research presentation.

4. Experimental Setup

The reported implementation uses Python, TensorFlow, Unity ML-Agents, SHAP and a permissioned blockchain. The experimental dataset is represented as 1,200 season-level episodes, divided into 80% training and 20% testing subsets. The stated PPO configuration uses a learning rate of 3×10^{-4}, discount factor $\gamma=0.99$ and clipping parameter 0.2. TFT uses a hidden dimension of 128 and four attention heads, while LTCE is configured with kernel setting k=3 and 64 channels. APO uses a population size of 30 and 50 iterations.

The comparison set listed in the study comprises CAViT, CNN-LSTM, DFE-RGI, SCAML, ABR and QYieldOpt. An ablation study is reported to assess the contribution of the proposed components. Stress testing is performed under sensor noise and climate variation, with optional fine-tuning on limited real data.

5. Results and Discussion

Table I summarizes the quantitative results reported in the supplied presentation. These values are presented as study results and should be retained only after the underlying experiments, datasets, random seeds and statistical tests have been independently reproduced or verified for the final conference submission.

Metric	Reported value	Interpretation
Regression performance	$R^2 = 0.919$; RMSE = 1.3; NRMSE = 0.027	Strong reported predictive fit
Convergence	78 epochs	Earlier convergence than stated baselines (>100 epochs)
Agronomic rule violations	2% vs. ~10%	Lower reported infeasible-decision rate
Resource efficiency / WUE	+15%	Reported improvement in resource-use efficiency
Robustness to sensor noise	92%	Reported stress robustness
Explanation fidelity under stress	89.7%	Reported explanation consistency
Blockchain throughput	82 TPS; <60 s latency	Reported validation performance

On-chain storage reduction	78%	Reduced ledger storage burden
Yield	+18%	Reported improvement
Water use	−11%	Reported saving
Fertilizer use	−7.5%	Reported reduction

5.1 Data and Robustness Performance

The preprocessing stage reports 92.8% missing-data recovery, outlier detection precision/recall of 0.88/0.85, a KL divergence value of 0.21 for the stated drift assessment, approximately 40% lower drift, and a post-augmentation class balance of 1.0. These values indicate that the preprocessing pipeline is intended to stabilize downstream learning under imperfect sensing.

5.2 Policy Learning and Agronomic Feasibility

The reported R^2 of 0.919 and NRMSE of 0.027 indicate strong predictive performance within the stated experimental setting. More importantly for decision support, the agronomic rule violation rate is reported as 2%, compared with approximately 10% for the comparison setting. This suggests that explicit symbolic validation contributes a distinct feasibility objective that is not captured by regression accuracy alone.

5.3 Explainability and Auditability

The reported explanation fidelity under stress is 89.7%. The architecture attributes this to the combination of SHAP and symbolic traces rather than to SHAP alone. On the provenance side, the system reports 82 TPS with less than 60 seconds latency and a 78% reduction in on-chain storage. The use of hashes and off-chain artifact storage is intended to balance auditability with the practical cost of storing large agricultural data products.

5.4 Resource and Yield Outcomes

The study reports a 15% resource-efficiency/WUE gain, an 18% yield improvement, 11% water savings and a 7.5% fertilizer reduction. Taken together, these measures support the intended multi-objective formulation: improve yield while reducing water, fertilizer and stress costs subject to agronomic constraints.

5.5 Ablation and System-Level Interpretation

The source presentation states that ablation confirms the contribution of individual components. The conceptual interpretation is that the modules address different failure modes: preprocessing addresses data uncertainty, digital-twin randomization addresses domain shift, symbolic reasoning addresses feasibility, SHAP and rule traces address interpretability, and PoISV addresses provenance. The framework therefore treats trust as an emergent system property produced by the interaction of these modules.

6. Limitations and Deployment Considerations

Several limitations should be made explicit before conference submission. First, the paper relies on the experimental results and configuration reported in the supplied presentation; detailed dataset provenance, crop type, geographic region, sensor specifications, random seeds, confidence intervals and statistical significance tests are not provided in the source material. Second, the physics-informed digital twin is described procedurally, but the exact crop-soil equations and calibration procedure are not specified. Third, APO is presented as the action optimizer, but its formal mathematical derivation and convergence analysis are outside the supplied material. Fourth, blockchain throughput and latency depend strongly on the permissioned network configuration and should be reported with hardware, consensus and transaction-size details.

These limitations do not invalidate the architectural contribution, but they affect reproducibility. A final conference version should therefore provide dataset identifiers, complete hyperparameters, implementation details, baseline configurations, statistical testing, ablation tables and a precise definition of each reported metric. If the numerical results are preliminary or simulation-only, this should be stated explicitly rather than presenting them as field-validated outcomes.

7. Conclusion

This paper presented PoISV-NSDT, a simulation-driven neuro-symbolic explainable AI framework with blockchain-enabled validation for trustworthy precision agriculture. The framework integrates robust multimodal preprocessing, a domain-randomized digital twin, PPO policy learning, constraint-guided neuro-symbolic reasoning, SHAP-based explanation and permissioned blockchain provenance. The central design principle is that trustworthy PA requires more than predictive accuracy: decisions must remain feasible under agronomic constraints, interpretable in terms of both data and rules, and auditable through verifiable provenance.

The reported experiments indicate improvements in prediction, convergence, rule compliance, robustness, explanation fidelity and resource efficiency. The framework also reports positive yield and input-use outcomes. Future work should focus on adaptive digital-twin refinement using field feedback, continuously evolving agronomic knowledge bases, and governance-aware blockchain interoperability for decentralized deployment. These directions can move the framework from a controlled simulation-oriented architecture toward validated, field-scale decision support.

References

[1] S. Pang, S. W. Teng, M. Murshed, C. V. Bui, P. Karmakar, Y. Li, and H. Lin, “A survey on evaluation of blockchain-based agricultural traceability,” Computers and Electronics in Agriculture, vol. 227, 2024, Art. no. 109548.

[2] B. Lim, S. Ö. Arık, N. Loeff, and T. Pfister, “Temporal Fusion Transformers for interpretable multi-horizon time series forecasting,” International Journal of Forecasting, vol. 37, no. 4, pp. 1748–1764, 2021.

[3] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, “Proximal Policy Optimization Algorithms,” arXiv:1707.06347, 2017.

[4] S. M. Lundberg and S.-I. Lee, “A Unified Approach to Interpreting Model Predictions,” in Advances in Neural Information Processing Systems 30, 2017.

[5] J. Achiam, D. Held, A. Tamar, and P. Abbeel, “Constrained Policy Optimization,” arXiv:1705.10528, 2017.

[6] M. Kamilaris, A. Fonts, and F. X. Prenafeta-Boldú, “The rise of blockchain technology in agriculture and food supply chains,” Trends in Food Science & Technology, vol. 91, pp. 640–652, 2019.

[7] A. T. Kamilaris, F. X. Prenafeta-Boldú, and others, “Blockchain in Agriculture Traceability Systems: A Review,” Applied Sciences, vol. 10, no. 12, 4113, 2020.

[8] M. Mirza and S. Osindero, “Conditional Generative Adversarial Nets,” arXiv:1411.1784, 2014.

[9] N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer, “SMOTE: Synthetic Minority Over-sampling Technique,” Journal of Artificial Intelligence Research, vol. 16, pp. 321–357, 2002.

[10] A. Juliani, V.-P. Berges, E. Teng, A. Cohen, J. Harper, C. Goy, Y. Gao, H. Henry, M. Mattar, and D. Lange, “Unity: A General Platform for Intelligent Agents,” arXiv:1809.02627, 2018.

[11] A. De Antoni, M. Rucco, A. M. Cattaneo, E. Gezer, G. Sulis, P. Draicchio, G. Iacca, A. Pugliese, and M. V. Mancini, “PesTwin: A biology-informed Digital Twin for enabling precision farming,” arXiv:2603.12294, 2026.