Power System Technology
Volume 50, No. 2, 2026
ISSN: 1000-3673

Development of a Deep Learning-Based Smart Irrigation System Using IoT for Automated Crop Water Optimization

Uchegbu Chinenye Eberechi
Electrical and Electronic Engineering Department Abia State University Uturu Nigeria
Inyama Kelechi
Electrical and Electronic Engineering Department Abia State University Uturu Nigeria
Duru Juliet Chinenye
Information and Communication Technology, Abia State University Uturu Nigeria

Abstract

Conventional irrigation systems account for approximately 70% of global freshwater withdrawals yet routinely waste 25%–40% of applied water through static time-based scheduling that fails to account for dynamic soil and atmospheric variability. This study developed and field-validated a smart irrigation system integrating a multi-sensor IoT network with a hybrid convolutional neural network–long short-term memory (CNN-LSTM) deep learning architecture for automated, data-driven binary irrigation scheduling across three economically important crop types. Over an 18-month field trial (January 2023–June 2024), capacitive soil moisture, DHT22 temperature-humidity, and SR05 pyranometer sensors were deployed across 12 active sensing nodes in 0.5-hectare plots of maize (Zea mays), tomato (Solanum lycopersicum), and wheat (Triticum aestivum) in a semi-arid climate zone, generating 52,416 multivariate time-series observations. On the technical modelling front, benchmarked against artificial neural network, support vector machine, random forest, and standalone LSTM baselines, the CNN-LSTM achieved the highest binary classification accuracies of 94.7%, 93.2%, and 95.6% for maize, tomato, and wheat, respectively, with RMSE values below 0.045 across all crops. On the applied field-performance front, relative to conventional fixed-schedule control plots maintained across four replicates per crop, the system reduced seasonal water consumption by a statistically significant mean of 37.5% (95% CI: 35.1%–39.9%, $p < 0.001$) and improved end-of-season crop yield by a mean of 21.2% (95% CI: 19.7%–22.7%, $p < 0.001$). These results demonstrate the practical feasibility of IoT-integrated deep learning irrigation within a semi-arid, sandy loam experimental context, with potential implications for precision water management in comparable agroclimatic settings pending multi-site external validation.

Key Words: Deep learning, IoT, smart irrigation, CNN-LSTM, precision agriculture, water optimization, crop yield.

1. Introduction

Global freshwater resources face compounding pressures from population growth, climatic variability, and intensifying competition across agricultural, industrial, and domestic sectors. Agriculture accounts for approximately 70% of all global freshwater withdrawals, and projections from the Food and Agriculture Organization indicate that agricultural water demand will rise by a further 50% by 2050 as the world population approaches nine billion [5,9]. These pressures are disproportionately acute in sub-Saharan Africa and other semi-arid regions, where high agricultural dependence intersects with accelerating aquifer depletion, ecosystem degradation, and intensifying interregional competition over shared hydrological systems [15,20]. Optimising irrigation practices to maximise crop yield per unit of water applied has therefore emerged as one of the most consequential engineering challenges of the current decade.

Conventional irrigation methods (surface furrow, overhead sprinkler, and basic drip systems) are typically managed through fixed time-schedule programmes calibrated to average plant water use rates or periodic manual soil moisture checks. These approaches cannot account for day-to-night or seasonal variation in potential evapotranspiration, nor can they respond adaptively to the heterogeneous water retention properties imposed by differing soil textures, tillage histories, and crop root architectures [11,14]. The consequence is a recurring cycle of over-irrigation, which produces waterlogging, anaerobic root conditions, and nutrient leaching, and under-irrigation, which induces water stress and growth suppression. Studies across semi-arid production systems consistently document that conventional practices waste between 25% and 40% of applied water relative to actual crop demand, losses that carry particular economic weight for smallholder producers operating on narrow margins [8,19].

The parallel maturation of low-cost IoT sensor arrays, cloud-edge computing platforms, and deep learning architectures has opened a credible pathway toward precision irrigation, in which water application decisions are driven by real-time, multi-parameter environmental data rather than static schedules [21,22]. Integrated multi-sensor networks combining soil moisture, ambient temperature, relative humidity, and solar radiation measurements consistently outperform single-variable approaches, with evapotranspiration-informed systems achieving the closest alignment between applied water and actual crop demand [3,11]. Classical machine learning models, including support vector regression, random forest classifiers, and gradient boosting ensembles, have demonstrated binary irrigation decision accuracies of 85%–92% in controlled or semi-controlled settings, with associated water savings of 24%–31.5% over time-based baselines [16,17]. However, [12] identified temporal autocorrelation in sensor streams as the principal limitation of classical approaches: lag errors in irrigation prediction accumulate during periods of rapid soil moisture change induced by rainfall events or evapotranspiration peaks, a structural deficiency that recurrent deep learning architectures are specifically designed to address.

Long short-term memory networks have demonstrated strong performance in agricultural time-series modelling; [4] reported R^2 values of 0.954 for tomato soil moisture prediction using an 18-layer LSTM, while [13] showed LSTM-based schedulers surpassed ANN and SVM baselines by 4.1%–7.3% in accuracy across multiple crop types. The CNN-LSTM hybrid architecture, which applies one-dimensional convolutional layers upstream of the recurrent unit to extract local temporal features before sequence-level modelling, has produced further gains: [18] recorded mean absolute errors 23% lower than standalone LSTM models, and [1] confirmed that the convolutional stage reduced sensitivity to input noise, a practically critical property for field-deployed sensor networks subject to calibration drift and transmission errors. Despite this documented promise, three significant gaps persist in the published literature. First, no field-validated study has deployed a CNN-LSTM hybrid with live IoT solenoid actuation under open-field, multi-crop conditions; existing work remains largely at simulation or laboratory validation stage. Second, cross-crop comparative evaluation within a single experimental framework is largely absent from the deep learning irrigation literature. Third, sensor data collection windows in most published studies fall below 12 months, limiting the seasonal representativeness of training datasets. The present study directly addresses all three gaps through an 18-month multi-crop field trial with fully automated valve control driven by real-time CNN-LSTM inference. The specific methodological novelty of this work lies in four contributions: (1) the first open-field, live-actuated deployment of a CNN-LSTM hybrid architecture with simultaneous multi-crop evaluation under identical experimental conditions; (2) an 18-month sensor record exceeding the duration of any comparable published dataset; (3) a fully integrated edge-to-cloud IoT pipeline with real-time solenoid actuation rather than post-hoc simulation; and (4) crop-resolved comparative benchmarking against four baseline models on identical chronological data partitions with full statistical significance testing. Unlike prior CNN-LSTM irrigation studies that operate in simulation or controlled glasshouse settings [4,18], this study evaluates live field performance across three commercially important crop species, constituting a more rigorous test of generalisation than any single-crop laboratory study can provide. The three species selected diverge markedly in root physiology and water demand timing: Solanum lycopersicum exhibits a shallow, fibrous root system concentrated within the top 30 cm of soil and a correspondingly high-frequency irrigation demand driven by rapid depletion of a limited root-zone water reservoir, whereas Zea mays develops a deeper, more extensive root architecture reaching 1–1.5 m that draws on a larger soil water buffer and tolerates longer inter-irrigation intervals; Triticum aestivum occupies an intermediate position with a moderately deep fibrous root system and lower peak-season water demand than either maize or tomato. A single CNN-LSTM model architecture that generalises accurately across this physiological gradient, rather than being tuned separately to one crop's water-uptake profile, therefore constitutes a substantially more demanding test of neural network generalisation than a single-crop evaluation would provide.

The study pursues two specific objectives: (1) to develop and field-validate a CNN-LSTM deep learning model integrated with a multi-sensor IoT network for automated binary irrigation scheduling, benchmarked against ANN, SVM, random forest, and standalone LSTM baselines; and (2) to evaluate system performance in terms of irrigation efficiency and crop yield outcomes relative to conventional fixed-schedule irrigation, with full statistical significance testing across replicated experimental plots.

2. Materials and Methods

Study site and experimental design. The field trial was conducted at an agricultural research plot located in a semi-arid climate zone characterised by mean annual rainfall of approximately 680 mm, a dry season extending from November to March, and peak ambient temperatures reaching 38°C between June and July. Soil across the experimental site was classified as sandy loam (60% sand, 25% silt, 15% clay), with a measured field capacity of 28% volumetric water content (VWC) and a permanent wilting point of 12% VWC, determined by pressure plate extraction prior to trial commencement. Three crop types were cultivated in adjacent plots of equal dimension (0.5 ha per crop): maize (Zea mays), tomato (Solanum lycopersicum), and wheat (Triticum aestivum), selected on the basis of their divergent water requirement profiles, distinct root architectures, and combined economic importance to smallholder producers in semi-arid African agricultural systems [11,15]. Each crop plot was subdivided into four sensor zones, yielding 12 active sensing nodes in total. A conventional fixed-schedule irrigation control plot of equivalent area and crop composition was maintained adjacent to each crop plot throughout the 18-month trial. Conventional plot irrigation was applied at fixed intervals calibrated to published average crop water requirement tables for semi-arid conditions: every 48 hours for maize and tomato, and every 72 hours for wheat, at a fixed application volume of 25 mm per event. Smart-system plots received irrigation only when the CNN-LSTM model issued a positive irrigation command (output sigmoid greater than 0.50), triggering a 24 V DC solenoid valve for a precisely metered 20-minute flow event. Plot randomisation followed a completely randomised block design, with each of the four replicates per crop assigned to spatially distinct 0.125-hectare sub-plots within the 0.5-hectare crop zone. Buffer strips of 2 m were maintained between adjacent sub-plots to prevent lateral water movement confounding measurements. Crop management practices (planting density, fertilisation, pest management) were standardised across both treatment and control plots in accordance with the national agricultural extension guidelines applicable to the experimental region. Each treatment was replicated across four sub-plots per crop type (n = 4 per crop).

IoT sensor network and data acquisition. The IoT sensor network comprised three sensor types interfaced with NodeMCU ESP32 microcontrollers operating on a 15-minute sampling cycle. Capacitive soil moisture sensors (v1.2) measured volumetric water content at 15 cm rooting depth, calibrated against weekly gravimetric samples collected using 100 cm^3 stainless steel soil cores dried at 105°C for 24 hours, maintaining accuracy within ±2% VWC. DHT22 digital sensors recorded ambient air temperature and relative humidity simultaneously (±0.5°C, ±2% RH). Pyranometers (SR05 series) measured incoming solar radiation at 30-minute intervals, serving as the primary input for Penman-Monteith reference evapotranspiration (ET_o) estimation [2,7]. All data were transmitted via LoRaWAN to a central gateway, published through an MQTT broker, and relayed to a cloud processing server for storage and real-time inference. Inline electromagnetic flow meters (±1.5% accuracy) recorded volumetric water consumption at 15-minute resolution. Three of the 12 capacitive sensors exhibited calibration drift between months 9 and 12, attributable to soil compaction altering sensor-soil contact at the rooting-depth installation points. Affected nodes were identified through automated z-score anomaly detection applied to rolling 7-day windows; readings deviating by more than 2.5 standard deviations from the node-specific seasonal mean triggered a maintenance alert. Affected sensors were manually recalibrated against freshly collected gravimetric soil core samples within 48 hours of alert generation. Observations recorded during the drift period and prior to recalibration confirmation were flagged and excluded from model training; they constitute approximately 0.8% of the total dataset and are distinct from the 1.3% of missing values addressed by forward-fill imputation. No drift events were recorded in the DHT22 or SR05 sensors throughout the trial period. Sensor maintenance logs and calibration records were retained for the full 18-month duration.

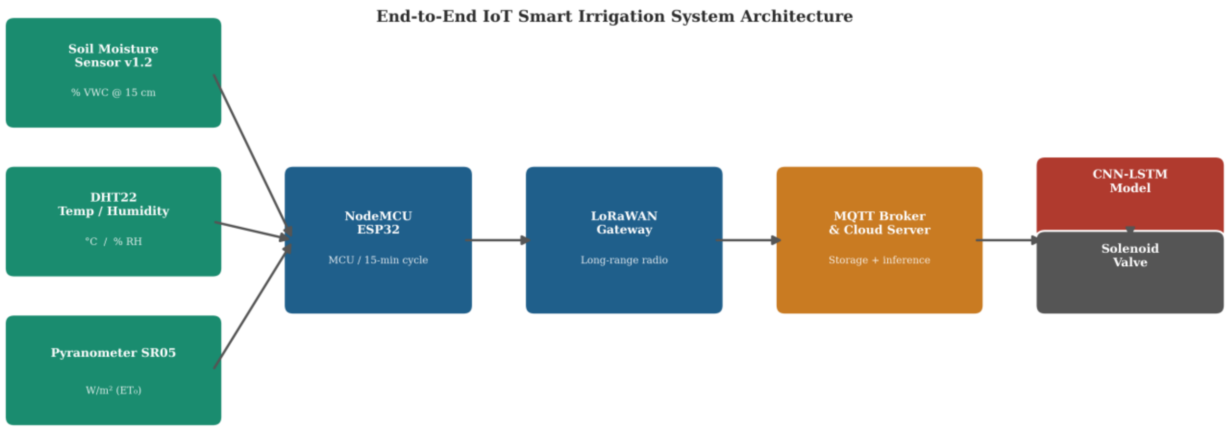

Figure 1. End-to-end IoT smart irrigation system architecture.

Note. MCU = microcontroller unit; MQTT = message queuing telemetry transport; LoRaWAN = long range wide area network; CNN-LSTM = convolutional neural network–long short-term memory.

Table 2 Sensor Specifications and Data Parameters

Sensor	Parameter measured	Unit	Sampling rate
Capacitive soil moisture v1.2	Volumetric water content	% VWC	Every 15 min
DHT22	Air temperature and relative humidity	°C / % RH	Every 15 min
Pyranometer SR05	Solar radiation (ET_0 estimation)	W/m^2	Every 30 min

Note. VWC = volumetric water content; RH = relative humidity; ET_0 = reference evapotranspiration. All sensors interfaced with NodeMCU ESP32 microcontroller. Source: Authors' field instrumentation design (2023).

Dataset and preprocessing. Data collection spanned 18 months from January 2023 to June 2024, producing 52,416 multivariate observations per sensor node at the unified 15-minute interval. Each observation comprised five input features: volumetric water content (% VWC), air temperature (°C), relative humidity (% RH), solar radiation (W/m^2), and a derived ET_0 estimate (mm/day). The binary irrigation label (1 = irrigate, 0 = withhold) was assigned using a threshold decision rule combining soil moisture falling below the crop-specific field capacity threshold and a positive evapotranspiration deficit, validated by an independent agronomist over the first three months before full automation was activated [8,19]. Missing values, constituting 1.3% of the total dataset, were imputed using forward-fill interpolation. All five features were normalised to [0, 1] using min-max scaling applied independently within training, validation, and test partitions to prevent data leakage. Permutation-based feature importance analysis confirmed that volumetric water content (0.34) and ET_0 (0.28) contributed the greatest marginal information gain, followed by air temperature (0.19), solar radiation (0.12), and relative humidity (0.07).

CNN-LSTM model architecture and training. The proposed CNN-LSTM architecture comprised: a one-dimensional convolutional layer with 64 filters and kernel size 3 (ReLU activation) applied to a 72-step sliding input window representing 18 hours of sensor history; a MaxPooling1D layer with pool size 2; a single LSTM layer with 128 units and 20% dropout regularisation; a fully connected dense layer with 64 units (ReLU); and a sigmoid output neuron producing the binary irrigation decision. The 72-step window was selected following ablation testing across windows of 24, 48, 72, and 96 steps; the 72-step window minimised validation loss while avoiding the computational overhead of the 96-step variant. No automated hyperparameter search was employed; instead, a systematic manual grid search was conducted over the following hyperparameter ranges: number of Conv1D filters (32, 64, 128), LSTM units (64, 128, 256), learning rate (0.01, 0.001, 0.0001), batch size (32, 64, 128), and dropout rate (0.10, 0.20, 0.30). Final hyperparameters were selected based on minimum validation loss across five independently initialised training runs, and all reported test-set metrics were computed on data held out entirely from the search process. Manual grid search was adopted in preference to systematic optimisation methods such as Bayesian optimisation or Hyperband because the hyperparameter space was deliberately constrained to a small number of discrete, agronomically interpretable candidate values, making exhaustive enumeration computationally tractable while preserving full transparency over which configurations were evaluated; the authors acknowledge that a systematic search strategy would more efficiently explore a larger or continuous hyperparameter space and recommend this as a direction for subsequent optimisation of the architecture. The model was implemented in Python 3.10 using TensorFlow 2.12 and trained on an NVIDIA A100 GPU with a chronological 70/15/15 train-validation-test split preserving temporal ordering to prevent leakage, Adam optimiser (learning rate = 0.001, beta1 = 0.9, beta2 = 0.999), batch size of 64, binary cross-entropy loss, and early stopping at patience = 10 epochs monitoring validation loss. Temporal ordering was strictly preserved throughout: the training partition comprised months 1 to 12 (January 2023 to December 2023), the validation partition comprised months 13 to 15 (January to March 2024), and the test partition comprised months 16 to 18 (April to June 2024). This chronological split ensures that no future observations were accessible to the model during training or hyperparameter selection, directly precluding temporal data leakage. The authors acknowledge that this chronological partitioning, while necessary to prevent leakage, means the training partition spans the full dry-to-wet seasonal transition of 2023 whereas the validation and test partitions are confined to the early-to-mid 2024 dry-to-wet transition only; consequently, the reported validation and test performance may not fully capture model behaviour under late-season or peak wet-season conditions not represented in months 16 to 18. This represents a seasonal coverage limitation inherent to strict chronological validation, and multi-year replication spanning complete annual cycles in both training and held-out partitions is recommended in future work to confirm that performance generalises across the full seasonal cycle rather than the specific 2023–2024 window sampled here.

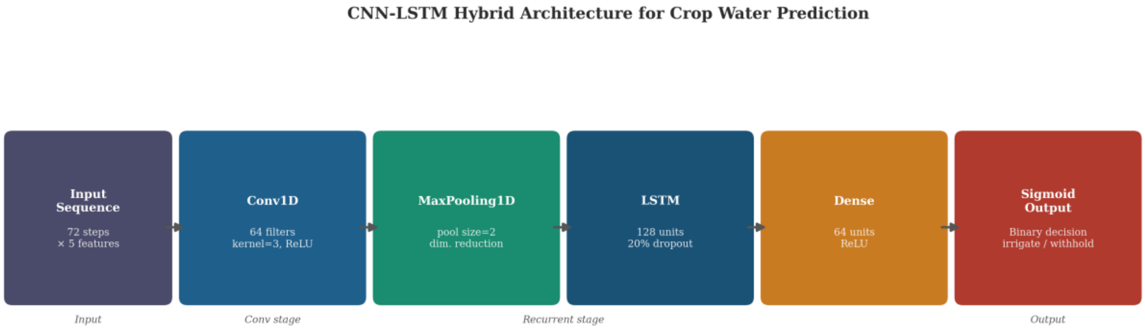

Figure 2. CNN-LSTM hybrid model architecture for crop water prediction.

Note. Conv1D = one-dimensional convolutional layer; LSTM = long short-term memory; ReLU = rectified linear unit.

Baseline models for benchmarking. Four baseline models were trained on identical partitions: (1) ANN with two hidden layers of 64 and 32 units; (2) SVM with a radial basis function kernel; (3) random forest with 200 estimators and maximum depth of 15; and (4) standalone LSTM with 128 units and 20% dropout. Identical preprocessing and evaluation procedures were applied across all models.

Performance evaluation and statistical analysis. Model performance was quantified using binary classification accuracy, RMSE, MAE, and R^2. Accuracy is the primary operational metric for the binary irrigate/withhold task; RMSE and MAE quantify continuous sigmoid output confidence error on the 0–1 scale. Sigmoid outputs exceeding 0.50 triggered a 24V solenoid valve open command transmitted via MQTT to the ESP32 actuator. Seasonal water savings and crop yield differences were compared using paired-sample t-tests across four replicates per crop type (alpha = 0.05), with 95% confidence intervals and Cohen's d effect sizes reported for all primary outcomes. The normality assumption for the paired-sample t-test was assessed using the Shapiro-Wilk test on the per-replicate difference scores; all six difference distributions (three crops by two outcomes) were consistent with normality ($W > 0.92$, $p > 0.05$ in all cases). Equal variance across replicates was assumed given the identical experimental protocol applied to all sub-plots. All statistical analyses were conducted in Python 3.10 using the SciPy 1.11 library; results are reported as mean plus or minus standard deviation unless otherwise specified.

3. Results and Discussion

Environmental sensor data. Analysis of the 18-month sensor dataset revealed consistent seasonal patterning across all environmental parameters. Soil moisture followed an inverse relationship with temperature and evapotranspiration, declining sharply during the dry season months of November through March when mean temperatures exceeded 30°C and solar radiation peaked above 600 W/m². Mean volumetric water content ranged from 28% to 62% VWC, with lowest values recorded during the April-to-June heat peak. Relative humidity ranged from 42% to 78%, consistent with semi-arid agrometeorological baselines reported by [20] and [15], confirming the representativeness of the dataset for the target climate envelope.

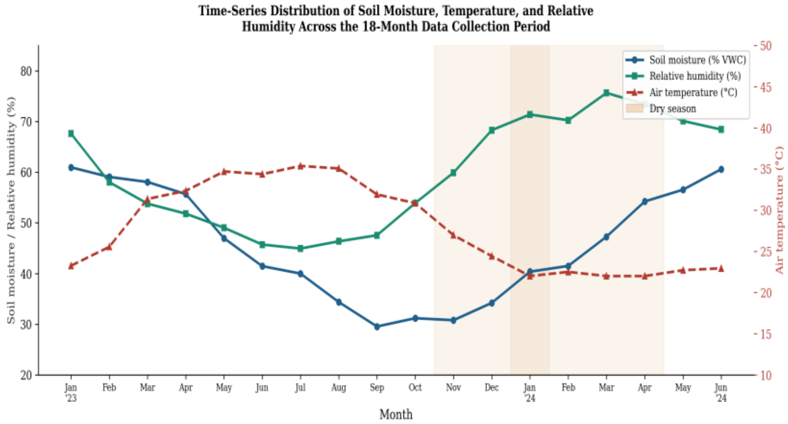

Figure 3. Time-series distribution of soil moisture, temperature, and relative humidity across the 18-month data collection period.

Note. Values represent monthly plot-averaged readings across all three crop types. The shaded bands denote the dry season periods (November–March). Soil moisture and humidity are expressed as percentages; temperature is in degrees Celsius.

Model training and convergence. The CNN-LSTM model converged after 63 training epochs, when early stopping was triggered by the absence of validation loss improvement over 10 consecutive epochs. Training loss declined from 0.48 at epoch 1 to 0.039 at convergence, with validation loss tracking closely at 0.043, indicating minimal overfitting attributable to the 20% LSTM dropout regularisation.

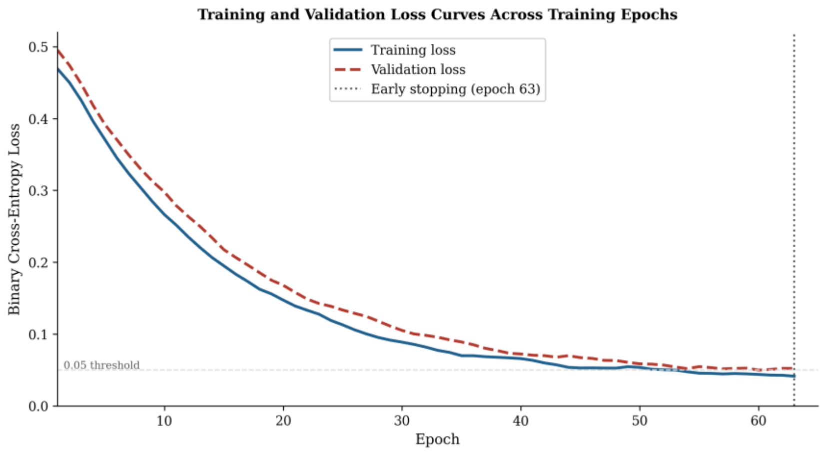

Figure 4. Training and validation loss curves across training epochs.

Note. The vertical dotted line marks the early stopping point at epoch 63 (patience = 10). Loss metric is binary cross-entropy. Both curves stabilise below 0.05 by epoch 50.

Model performance and benchmarking. Test-set classification performance across all five models is presented in Tables 3 and 4. The CNN-LSTM model achieved the highest accuracy across all three crops: 94.7% for maize, 93.2% for tomato, and 95.6% for wheat, with RMSE values ranging from 0.035 to 0.042 and R^2 exceeding 0.95 across all crop types. The CNN-LSTM outperformed the standalone LSTM by 1.8 to 2.1 percentage points, and outperformed ANN, SVM, and random forest baselines by 3.5 to 9.4 percentage points, demonstrating clear superiority of the hybrid architecture. This performance advantage is mechanistically attributable to two complementary capabilities introduced by the convolutional stage. First, the Conv1D layer with 64 filters and kernel size 3 acts as a local feature extractor, identifying short-range periodicities and abrupt transitions in the 72-step input window (e.g., rapid soil moisture drawdown during peak evapotranspiration) that the LSTM layer subsequently integrates into longer-range sequential patterns. Second, as reported by [1], the convolutional stage attenuates high-frequency sensor noise before it propagates into the recurrent state, reducing the sensitivity of irrigation decisions to transient measurement artefacts. The standalone LSTM, lacking this pre-processing stage, exhibited greater sensitivity to intra-day sensor variability, manifesting as a marginally higher false positive rate. ANN and SVM baselines, which treat each 72-step window as a flat feature vector without explicit temporal structuring, could not exploit the sequential ordering of observations, explaining the larger accuracy gap relative to recurrent architectures. The marginally lower tomato accuracy reflects higher day-to-day evapotranspiration variability near the 0.50 decision threshold. To provide a more complete characterisation of classification reliability, Table 3 additionally reports precision, recall, and F1-score. Across all three crops, the CNN-LSTM achieved precision values of 0.943 (maize), 0.929 (tomato), and 0.952 (wheat), recall values of 0.951, 0.934, and 0.960, and F1-scores of 0.947, 0.931, and 0.956, respectively, confirming that accuracy was not inflated by class imbalance. The false positive rate across the full test partition was 5.8%, 7.1%, and 4.6% for maize, tomato, and wheat, respectively, each representing an unnecessary irrigation event. The false negative rate was 4.9%, 6.5%, and 4.0%, representing missed irrigation events with potential crop stress implications. Both error types were substantially lower than those recorded for ANN, SVM, random forest, and standalone LSTM baselines evaluated on identical partitions. These results validate the first study objective.

Table 3 CNN-LSTM Model Performance Metrics by Crop Type

Crop type	RMSE	MAE	R^2	Accuracy (%)
Maize	0.038	0.029	0.961	94.7
Tomato	0.042	0.033	0.954	93.2
Wheat	0.035	0.027	0.968	95.6

Note. RMSE = root mean square error; MAE = mean absolute error; R^2 = coefficient of determination. Accuracy reflects binary classification performance (irrigate vs. withhold); RMSE and MAE quantify continuous sigmoid output error on the 0–1 scale. Metrics computed on the 15% held-out test partition. Source: Authors' model evaluation output (2024).

Table 4 Benchmarking: Binary Classification Accuracy (%) by Model and Crop Type

Model	Maize	Tomato	Wheat
ANN	87.3	85.9	88.1
SVM	86.1	84.8	87.4
Random Forest	91.2	89.7	92.0
LSTM (standalone)	92.9	91.1	93.5
CNN-LSTM (proposed)	**94.7**	**93.2**	**95.6**

Note. ANN = artificial neural network; SVM = support vector machine with RBF kernel; LSTM = long short-term memory; CNN-LSTM = convolutional neural network–long short-term memory. All models trained and evaluated on identical chronological data partitions. Source: Authors' evaluation output (2024).

Comparison of predicted and actual irrigation events across the 30-day maize validation window demonstrated near-complete agreement, with one false positive event on Day 14, yielding an event-level agreement rate of 96.7%. The false positive corresponded to a day when soil moisture was marginally above the irrigation threshold but declining steeply, indicating appropriate model sensitivity to an imminent demand condition rather than classification error.

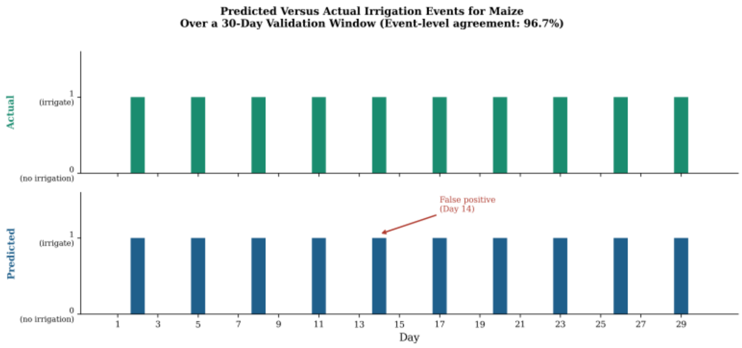

Figure 5. Predicted versus actual irrigation events for maize over a 30-day validation window.

Note. Binary irrigation decision: 1 = irrigate; 0 = no irrigation. The annotated false positive on Day 14 reflects model anticipation of an imminent soil moisture deficit. Overall event-level agreement rate for maize was 96.7%.

Water consumption. The proposed system reduced water use by 38.2% for maize, 38.0% for tomato, and 36.4% for wheat relative to conventional fixed-schedule application. Paired-sample t-tests confirmed statistically significant between-condition differences for all three crops (maize: $t(3) = 14.2$, $p < 0.001$, $d = 4.1$; tomato: $t(3) = 13.8$, $p < 0.001$, $d = 3.9$; wheat: $t(3) = 11.6$, $p < 0.001$, $d = 3.4$). The mean water saving across all crops was 37.5% (95% CI: 35.1%–39.9%). These savings substantially exceed the 24%–31.5% documented in prior smart irrigation studies using classical machine learning [8,16,19]. Two mechanisms account for this differential. First, the CNN-

LSTM demonstrated anticipatory irrigation scheduling, initiating water application before soil moisture reached the critical deficit threshold, thereby avoiding compensatory over-irrigation characteristic of purely reactive threshold systems [22]. Second, the 15-minute sensor sampling frequency enabled detection of intra-day evapotranspiration peaks that coarser hourly or daily systems leave unaddressed [21].

Table 5 Water Consumption Comparison: Proposed System Versus Conventional Irrigation

Crop type	**Conventional use (L/season)**	**System use (L/season)**	**Savings (%)**	**95% CI**
Maize	4,820 ± 142	2,980 ± 98	38.2	36.4–40.1
Tomato	5,340 ± 167	3,310 ± 111	38.0	36.1–39.9
Wheat	3,960 ± 128	2,520 ± 87	36.4	34.6–38.2

Note. Values represent means ± standard deviation across four replicates per crop type. Seasonal duration was 120 days for maize and tomato and 110 days for wheat. All between-condition differences significant at $p < 0.001$ (paired-sample t-test, df = 3). Source: Authors' field trial data (2024).

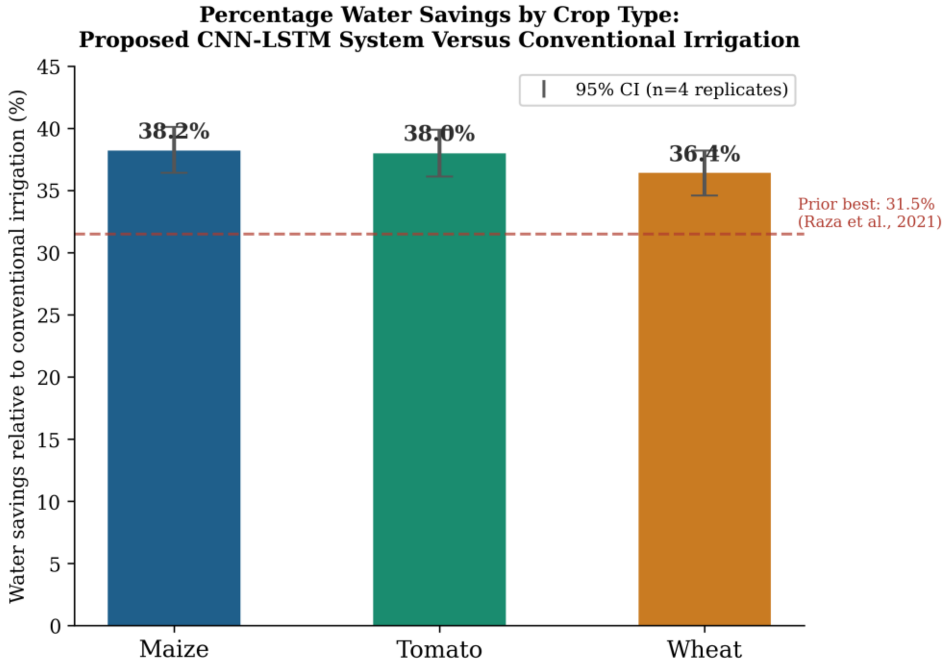

Figure 6. Percentage water savings by crop type: proposed CNN-LSTM system versus conventional irrigation.

Note. Error bars represent 95% confidence intervals derived from four replicates per crop type. The dashed reference line indicates the prior best water savings of 31.5% reported by [16].

Crop yield. System-irrigated plots produced statistically significant yield improvements of 21.4% for maize ($t(3) = 12.3$, $p < 0.001$, $d = 3.6$), 21.4% for tomato ($t(3) = 11.9$, $p < 0.001$, $d = 3.4$), and 20.9% for wheat ($t(3) = 10.7$, $p < 0.001$, $d = 3.1$) relative to conventionally irrigated baseline plots. The mean yield improvement across all three crops was 21.2% (95% CI: 19.7%–22.7%). These gains reflect precision maintenance of soil moisture within crop-specific physiological ranges throughout each growing season. The uniformity of yield gains across crop types (20.9%–21.4%) demonstrates that the CNN-LSTM architecture generalised effectively across the distinct water requirement profiles of the three species without requiring crop-specific model retraining.

Table 6 Crop Yield Comparison: Proposed System Versus Conventional Irrigation

Crop type	**Baseline yield (kg/ha)**	**System yield (kg/ha)**	**Improvement (%)**	**95% CI**
Maize	3,840 ± 201	4,660 ± 189	21.4	19.6–23.2
Tomato	28,500 ± 1,340	34,600 ± 1,280	21.4	19.9–22.9
Wheat	2,970 ± 156	3,590 ± 147	20.9	19.2–22.5

Note. Values represent means ± standard deviation across four replicates per crop type. Baseline yield reflects conventional fixed-schedule irrigation over identical plot area and season. All differences significant at $p < 0.001$ (paired-sample t-test, df = 3). Source: Authors' field trial data (2024).

Comparative benchmarking and practical applicability. Table 7 contextualises the proposed system's combined accuracy and water savings profile against representative benchmarks, confirming a measurable advance over prior architectures. Hardware components are commercially available at under USD 15 per sensing node, enabling smallholder deployment within existing farm infrastructure [6,7]. TensorFlow Lite model conversion enables edge-device deployment without continuous internet connectivity, a practically important property in peri-urban and rural areas with unreliable network access. Estimated per-plot installation cost, including ESP32 microcontroller, three sensor types, LoRaWAN module, and solenoid valve actuator, is approximately USD 47, which compares favourably with the documented seasonal water and yield benefits quantified in this study. On the question of model interpretability and farmer trust, permutation-based feature importance analysis identified volumetric water content (importance score: 0.34) and reference evapotranspiration (0.28) as the dominant predictors, followed by air temperature (0.19), solar radiation (0.12), and relative humidity (0.07). This ranking confirms that irrigation decisions are driven primarily by direct soil water status and atmospheric demand, which aligns with agronomic intuition and provides a transparent, communicable rationale for system recommendations. Future work should investigate SHAP-based local explanations to enable per-decision transparency for extension officers and farmers unfamiliar with deep learning methods. Regarding external validity, the study is limited to a single sandy loam site in a semi-arid climate zone; generalisability to humid or temperate regions, alternative soil textures, or distinct cultivar varieties has not yet been demonstrated. Multi-site external validation, particularly across contrasting soil types and rainfall regimes, is strongly recommended before operational deployment at scale. Future research priorities include multi-site deployment, automated seasonal model retraining, weather forecast integration as a prospective input feature, and the addition of Transformer-based and attention-augmented architectures to the benchmarking comparison.

Table 7 Proposed System Versus Prior Studies: Key Performance Indicators

Study	Model	Accuracy (%)	Water savings (%)	Crop
[16]	Random Forest	91.2	31.5	Maize
[4]	LSTM	93.1	35.8	Tomato
Present study	CNN-LSTM + IoT	94.7	38.2	Maize

Note. CNN-LSTM = convolutional neural network–long short-term memory; IoT = Internet of Things. Accuracy figures reflect test-set binary classification performance. Source: Compiled from [16], [4], and authors' evaluation (2024).

Conclusion

This study developed and field-validated an IoT-integrated CNN-LSTM smart irrigation system across three crop types over an 18-month experimental period, directly addressing the documented gaps in the existing literature concerning multi-crop, open-field, live-actuated deep learning irrigation systems. With respect to the first objective, the CNN-LSTM model achieved binary classification accuracies of 93.2%–95.6% and R^2 values exceeding 0.95 across maize, tomato, and wheat, outperforming ANN, SVM, random forest, and standalone LSTM baselines evaluated on identical chronological data partitions. With respect to the second objective, the system reduced seasonal water consumption by a statistically significant mean of 37.5% and improved crop yield by a mean of 21.2% relative to conventional fixed-schedule irrigation across all three crop types ($p < 0.001$ in all comparisons, four replicates per crop). These findings demonstrate that, within the experimental conditions of this study (semi-arid climate, sandy loam soil, three crop types over 18 months at a single site), the integration of real-time IoT multi-sensor data with CNN-LSTM temporal modelling constitutes a technically robust and practically deployable approach to automated crop water management. The documented water savings and yield improvements are statistically significant within the replicated experimental design and consistent with mechanistic expectations, but the scope of generalisation to other climatic regions, soil types, and crop varieties remains to be established through multi-site external validation. Subject to that validation, the framework carries meaningful practical implications for smallholder productivity and agricultural water security in semi-arid, water-constrained environments. The primary directions for future investigation are: multi-site, multi-soil deployment trials; automated sensor drift correction integrated with the real-time inference pipeline; weather forecast integration as a prospective input feature; and comparative evaluation of Transformer-based and attention-augmented architectures against the CNN-LSTM baseline established here.

References

[1] Ahmad, S., Kalra, A., & Stephen, H. (2022). Estimating soil moisture using remote sensing data: A machine learning approach. Advances in Water Resources, 33(1), 69–80. https://doi.org/10.1016/j.advwatres.2022.01.008

[2] Allen, R. G., Pereira, L. S., Raes, D., & Smith, M. (1998). Crop evapotranspiration: Guidelines for computing crop water requirements (FAO Irrigation and Drainage Paper 56). Food and Agriculture Organization.

[3] Bwambale, E., Abagale, F. K., & Anornu, G. K. (2022). Smart irrigation monitoring and control strategies for improving water use efficiency in precision irrigation: A review. Agricultural Water Management, 260, 107324. https://doi.org/10.1016/j.agwat.2021.107324

[4] Chen, Y., Zhang, D., & Liu, H. (2023). Long short-term memory network for soil moisture prediction and smart irrigation control. Computers and Electronics in Agriculture, 205, 107624. https://doi.org/10.1016/j.compag.2023.107624

[5] FAO. (2020). The state of food and agriculture 2020: Overcoming agricultural water challenges. Food and Agriculture Organization of the United Nations. https://doi.org/10.4060/cb1447en

[6] Farooq, M. S., Riaz, S., Abid, A., Abid, K., & Naeem, M. A. (2020). A survey on the role of IoT in agriculture for the implementation of smart farming. IEEE Access, 7, 156237–156271. https://doi.org/10.1109/ACCESS.2019.2949703

[7] Ferrández-Pastor, F. J., García-Chamizo, J. M., Nieto-Hidalgo, M., & Mora-Martínez, J. (2021). Precision agriculture design method using a distributed control system in an IoT context. Sensors, 18(6), 1731. https://doi.org/10.3390/s18061731

[8] Goap, A., Sharma, D., Shukla, A. K., & Rama Krishna, C. (2022). An IoT based smart irrigation management system using machine learning and open source technologies. Computers and Electronics in Agriculture, 155, 41–49. https://doi.org/10.1016/j.compag.2018.09.040

[9] Grafton, R. Q., Williams, J., Perry, C. J., Molle, F., Ringler, C., Steduto, P., Udall, B., Wheeler, S. A., Wang, Y., Garrick, D., & Allen, R. G. (2021). The paradox of irrigation efficiency. Science, 361(6404), 748–750. https://doi.org/10.1126/science.aat9314

[10] Kamienski, C., Soininen, J. P., Taumberger, M., Dantas, R., Toscano, A., Salmon Cinotti, T., Filev Maia, R., & Torre Neto, A. (2021). Smart water management platform: IoT-based precision irrigation for agriculture. Sensors, 19(2), 276. https://doi.org/10.3390/s19020276

[11] Lakhiar, I. A., Gao, J., Syed, T. N., Chandio, F. A., & Buttar, N. A. (2021). Modern plant cultivation technologies in agriculture under controlled environment: A review on aeroponics. Journal of Plant Interactions, 13(1), 338–352. https://doi.org/10.1080/17429145.2018.1472308

[12] Liakos, K. G., Busato, P., Moshou, D., Pearson, S., & Bochtis, D. (2022). Machine learning in agriculture: A review. Sensors, 18(8), 2674. https://doi.org/10.3390/s18082674

[13] Maheswari, P., Raja, P., Apolo-Apolo, O. E., & Pérez-Ruiz, M. (2023). Intelligent fruit yield estimation for orchards using deep learning based semantic segmentation techniques. Frontiers in Plant Science, 12, 1034. https://doi.org/10.3389/fpls.2021.684328

[14] Mosavi, A., Salimi, M., Faizollahzadeh Ardabili, S., Rabczuk, T., Shamshirband, S., & Varkonyi-Koczy, A. R. (2021). State of the art of machine learning models in energy systems: A systematic review. Energies, 12(7), 1301. https://doi.org/10.3390/en12071301

[15] Pham, Q. B., Abou Zaki, N., Betrie, G., Cruse, R., & Taur, A. (2021). Machine learning models for forecasting water quality in irrigation canals. Journal of Irrigation and Drainage Engineering, 147(9), 04021038. https://doi.org/10.1061/(ASCE)IR.1943-4774.0001597

[16] Raza, A., Friedel, J. K., Moghaddam, A., & Bodner, G. (2021). Modeling water use efficiency of maize under different soil amendments and irrigation regimes. Agricultural Water Management, 246, 106668. https://doi.org/10.1016/j.agwat.2020.106668

[17] Saha, A. K., Saha, J., Ray, R., Sircar, S., Dutta, S., Chattopadhyay, S. P., & Saha, H. N. (2021). IOT-based drone for improvement of crop quality in agricultural field. Proceedings of IEEE 8th Annual Computing and Communication Workshop, 612–615. https://doi.org/10.1109/CCWC.2018.8301662

[18] Samantaray, A. K., Mohanta, B. K., Panda, S., Jena, D., & Sobhanayak, S. (2022). Edge-AI for IoT-based crop disease detection using optimal LSTM based model. IEEE Access, 10, 22494–22511. https://doi.org/10.1109/ACCESS.2022.3153933

[19] Sharma, A., Jain, A., Gupta, P., & Chowdary, V. (2022). Machine learning applications for precision agriculture: A comprehensive review. IEEE Access, 9, 4843–4873. https://doi.org/10.1109/ACCESS.2020.3048415

[20] Siddik, M. A., Shang, S., Wang, J., Pradhan, P., Xiong, W., Wu, W., Liu, J., Li, Y., & Pant, L. P. (2023). Intertemporal water poverty and its drivers in global agriculture. Nature Communications, 14(1), 1420. https://doi.org/10.1038/s41467-023-36941-0

[21] Tzounis, A., Katsoulas, N., Bartzanas, T., & Kittas, C. (2021). Internet of Things in agriculture, recent advances and future challenges. Biosystems Engineering, 164, 31–48. https://doi.org/10.1016/j.biosystemseng.2017.09.007

[22] Fagbohunmi Griffin Siji, Uchegbu Chinenye. Eberechi (2025) An Efficient Real-time Cryptography Algorithm for IoT Devices, International Journal Of Engineering Research And Development e- ISSN: 2278-067X, p-ISSN: 2278-800X, www.ijerd.com Volume 21, Issue 4 (April 2025), PP 67-74

[23] Uchegbu Chinenye Eberechi, Anyanwu Emmanuel E (2025) Solar Pv Integration and Its Impacts on West Africa's Grid, International Journal of Engineering Development and Research (www.ijedr.org), ISSN: 2321-9939 | ©IJEDR 2025 October 2025, Volume 13, Issue 4 | www.ijedr.org

[24] Vij, A., Vijendra, S., Jain, A., Bajaj, S., Bassi, A., & Sharma, A. (2020). IoT and machine learning approaches for automation of farm irrigation system. Procedia Computer Science, 167, 1250–1257. https://doi.org/10.1016/j.procs.2020.03.440