

# An Integrated AI- and IoT-Based Smart Water Management Framework for Reservoir-Fed Irrigation Agriculture in Sri Lanka

DILRUKSHA A.G.C.D.  
Department of Computer  
Science and Software  
Engineering  
Sri Lanka Institute of  
Information Technology  
Malabe, Sri Lanka  
[it222561770@my.sliit.lk](mailto:it222561770@my.sliit.lk)

HESARA P.K.A.N  
Department of Computer  
Science and Software  
Engineering  
Sri Lanka Institute of  
Information Technology  
Malabe, Sri Lanka  
[it22561398@my.sliit.lk](mailto:it22561398@my.sliit.lk)

TRISHNI W.R.M.  
Department of Computer  
Science and Software  
Engineering  
Sri Lanka Institute of  
Information Technology  
Malabe, Sri Lanka  
[it22076366@my.sliit.lk](mailto:it22076366@my.sliit.lk)

ABISHEK S.  
Department of Computer  
Science and Software  
Engineering  
Sri Lanka Institute of  
Information Technology  
Malabe, Sri Lanka  
[it22186942@my.sliit.lk](mailto:it22186942@my.sliit.lk)

**Abstract** - In Sri Lanka, reservoir-fed irrigation schemes distribute water through administrative quotas, yet the everyday operational choices when to release water, what crops to plant, when to open field valves, and how to respond to crop stress fall to separate actors who rarely coordinate their data. While reservoir forecasting, IoT-based field actuation, satellite-derived crop-health monitoring, and seasonal crop-area planning have each attracted independent research attention, no prior work has woven all four into a shared decision loop tailored to a Sri Lankan reservoir-fed context. We describe ASICOP, the Adaptive Smart Irrigation and Crop Optimization Platform, which ties these four streams together through a common API gateway and event bus for the Udawalawe scheme. The platform includes a forecasting service that targets 1- to 14-day water-level and rainfall horizons, though full multi-year validation is still pending. An IoT service governs field-valve actions using ESP32 sensor telemetry and postpones irrigation when rainfall is anticipated. A hybrid crop-health service generates per-zone stress maps from Sentinel-2 spectral indices alongside an image-based disease classifier. A crop-planning service applies Fuzzy-TOPSIS ranking and allocates cultivable areas against the reservoir forecast. In testing, the F1 release regressor achieved  $R^2 = 0.7949$  with  $RMSE = 0.7108$  MCM/day across 1,095 held-out days drawn from a 32-year hydrological record, and the F2 disease classifier reached 95.43% accuracy on the 38-class PlantVillage benchmark. Both the F3 forecasting module and the F4 crop-planning module remain at prototype stage, and no end-to-end deployment evaluation has been carried out. What we contribute is a connected decision-support framework, not four standalone tools.

**Keywords:** *adaptive irrigation, crop health monitoring, decision-support systems, Fuzzy-TOPSIS, Internet of Things (IoT), machine-learning forecasting, remote sensing, reservoir operation, smart water management, Sri Lanka.*

## I. INTRODUCTION

Freshwater withdrawal for agriculture accounts for the single largest share of global water use, and demographic growth, intensifying yields, and shifting rainfall under climate change continue to tighten that resource [9]. Sri Lanka channels an estimated 87% of its freshwater into farming, most of it routed through dry-zone reservoir schemes that store seasonal rain and parcel it out via canal networks under administrative quotas [10], [11]. The Udawalawe scheme, fed by the Walawe Ganga reservoir in the southern dry zone, is representative of this class. Within it, reservoir operators set daily release volumes by rule of

thumb, farmers settle on crop choices without knowing the aggregate water demand across the command area, field-valve schedules run on fixed calendars regardless of whether rain is imminent, and crop-health problems surface only when an extension officer happens to walk through the affected zone.

Recent advances have opened a path toward linking these fragmented decisions. Gradient-boosted tree ensembles [3] and long short-term memory networks [4] can learn multi-decade reservoir patterns well enough to produce useful daily release forecasts. Sentinel-2 satellite imagery, revisiting the same ground every five days at 10 m resolution, picks up drops in canopy vigour or water content that point to stressed zones. ESP32-class microcontrollers paired with capacitive soil-moisture probes and MQTT messaging can relay field conditions to a central event bus at minimal cost. Multi-criteria ranking approaches like Fuzzy-TOPSIS [6] can weigh crop options against water supply, soil fit, and market price simultaneously. Each of these capabilities has been demonstrated on its own, but pulling all four into a single operational loop for a quota-governed Sri Lankan scheme has not been attempted before. We built such a loop and call it the Adaptive Smart Irrigation and Crop Optimization Platform, or ASICOP. This paper puts forward four specific claims about it: first, that it provides a working prototype of a reservoir-to-field-to-feedback framework with clearly defined cross-module event flows; second, that its release-prediction model, trained on 32 years of Udawalawe hydrology, performs well on a chronological holdout; third, that its hybrid satellite-plus-image crop-health pipeline delivers useful zone-level triage alongside leaf-level disease screening; and fourth, that its crop and area planning stack can combine market signals, multi-criteria suitability scores, and quantile water forecasts into a single allocation.

## II. LITERATURE REVIEW

### A. Reservoir Forecasting and Water-Level Prediction

The practice of reservoir operation has gradually moved away from purely rule-based and physically parameterized approaches SWAT [12], HEC-ResSim [13] and their relatives toward data-driven and hybrid strategies. Long short-term memory networks [4], [14] and gradient-boosted ensembles

[3] have matched or bettered traditional statistical seasonal models wherever sufficiently long training records are available. Probabilistic forecasting has gained particular traction in operational settings: the P10/P50/P90 spread of a forecast often proves more useful to an operator than a single best-guess number [15]. Two gaps in this body of work are relevant here. First, the bulk of published ML reservoir-forecasting research draws on large, well-instrumented schemes in China, the United States, and the Nordic countries, leaving small- and medium-scale Sri Lankan schemes comparatively understudied [1]. Second, the forecasts produced by these studies are almost always delivered as self-contained outputs a PDF bulletin or a web dashboard rather than machine-readable signals that feed into downstream planning or actuation components.

### B. Smart Field Irrigation and IoT Actuation

Research on sensor-driven field irrigation has matured considerably over the past decade. Early systems coupled fixed soil-moisture thresholds to relay-switched valves; subsequent designs introduced fuzzy rule bases and, more recently, supervised classifiers trained on paired sensor-and-yield histories [16]. The arrival of ESP32 microcontrollers, low-cost capacitive probes, and lightweight MQTT messaging has brought the hardware cost down to a level that even smallholder deployments can consider. Yet most of this work treats the irrigated field as an isolated control problem: the scheduling logic does not draw on any scheme-wide reservoir forecast, nor does it adjust its priorities when satellite imagery flags stress in a neighbouring zone. That disconnect is the gap ASICOP seeks to bridge on the field-actuation side.

### C. Satellite and Image-Based Crop Health Monitoring

Sentinel-2 vegetation indices NDVI for canopy greenness and NDWI or NDMI variants for canopy water content processed at 10 m spatial resolution can highlight stressed zones within hours of a satellite overpass [17], [18]. On the image-classification front, transfer-learning models such as MobileNet and ResNet fine-tuned on the PlantVillage dataset routinely achieve high validation accuracy under controlled single-leaf conditions, but their performance drops noticeably in the field, where leaf occlusion, mixed canopies, and variable lighting are the norm [19]. ASICOP’s crop-health module accordingly treats both its satellite and image outputs as advisory triage signals, not as definitive agronomic diagnoses a design choice that acknowledges both the proxy-label nature of the satellite track and the controlled-imaging origin of the disease classifier.

### D. Crop Recommendation, Suitability, and Area Optimization

Deciding which crops to grow and how much land to give each one, subject to water and soil constraints, is a well-studied operations-research problem [20]. Multi-criteria decision methods AHP, TOPSIS, and their fuzzy extensions have been applied to blend agronomic suitability, market outlook, and risk tolerance into a unified ranking [6], [22]. Machine-learning crop recommenders, typically random forests or gradient-boosted classifiers, trained on soil-and-climate feature vectors have also become common in precision-agriculture platforms [21]. Two points bear on

ASICOP’s design. First, most crop-planning studies take a single expected water-availability figure as input; very few consume quantile reservoir forecasts (P10, P50, P90) that express the uncertainty around that figure. Second, the Hector Kobbekaduwa retail price dataset that our F4 module draws on was originally assembled as a market-monitoring record, not as a cultivation-decision dataset, so it lacks the soil and nutrient features found in purpose-built crop-recommendation benchmarks.

### E. Comparative Summary

TABLE I. COMPARISON OF REPRESENTATIVE PRIOR SYSTEMS WITH ASICOP

System	Rsrvr	IoT	Sat.	Img.	Crop	Quota	SL
LSTM rainfall–runoff [14]	✓	—	—	—	—	—	—
ESP32/IoT smart-irrig. [16]	—	✓	—	—	—	—	—
Sentinel-2 stress [18]	—	—	✓	—	—	—	—
MobileNet/PlantVillage [19]	—	—	—	✓	—	—	—
IoT + remote-sensing [21]	—	✓	✓	—	—	—	—
LP/MIP crop-area [20]	✓	—	—	—	✓	—	—
Multi-track AI/IoT [1]	—	✓	✓	✓	—	—	—
<b>ASICOP (this paper)</b>	<b>✓</b>	<b>✓</b>	<b>✓</b>	<b>✓</b>	<b>✓</b>	<b>✓</b>	<b>✓</b>

## III. RESEARCH GAP, PROBLEM, AND CONTRIBUTIONS

The literature reviewed above reveals three interconnected gaps. The four technical streams reservoir forecasting, IoT irrigation, satellite health monitoring, and crop optimisation each have well-developed stand-alone research base, yet attempts to unify them under one decision-support umbrella are rare. Where integration has been tried, the resulting platforms typically pair two of the four streams and leave the others untouched. At the same time, the geographic and governance context that matters for ASICOP is underserved: the published work overwhelmingly targets large, data-rich schemes, with Sri Lankan reservoirs operating under seasonal administrative quotas largely absent from the machine-learning literature [1]. Even in the few cases where researchers have attempted cross-stream integration, the connection mechanism tends to be shallow a shared dashboard or a manual hand-off rather than a structured data contract between services.

Translated into a concrete operational problem: at the Udawalawe scheme, four tightly coupled decisions are made by different people at different cadences with almost no shared data flowing between them. The reservoir operator decides how much water to release each day in light of the next one-to-fourteen-day forecast. The farming community collectively decides what to plant and on how much land each season, given expected prices and the scheme-wide water budget. Individual field-valve schedules respond to current soil moisture, anticipated near-term rainfall, and any active crop-stress signal. And zone-level crop-health anomalies need to be caught before they translate into yield loss.

**Contributions.** We address these gaps with four contributions. (1) A prototype reservoir-to-field-to-feedback decision framework for Sri Lankan quota-based schemes, wired with MQTT event flows, Redis-cached cross-service calls, and graceful-degradation fallbacks, verified through service-log inspection; system-level water-saving and yield outcomes remain to be established in a deployment study. (2) A release-prediction model trained on 32 years of Udawalawe daily hydrology (HistGradientBoostingRegressor;  $R^2 = 0.7949$ , RMSE = 0.7108 MCM/day on a 1,095-day chronological holdout), which outperforms both Ridge and Random Forest baselines on RMSE and  $R^2$ . (3) A two-track crop-health pipeline pairing a Random Forest zonal classifier [2] trained on Sentinel-2 NDVI and a vegetation water-content index evaluated against proxy-derived labels with a MobileNetV2 disease classifier [8] fine-tuned on PlantVillage to 95.43% validation accuracy across 38 classes. (4) A crop and area optimisation stack that chains a Random Forest market-signal recommender [2] achieving 97.0% test accuracy on five vegetable classes, a neural-network price predictor ( $R^2 = 0.167$ ), a Fuzzy-TOPSIS suitability scorer [6], and a greedy area allocator that reads F3 P10/P50/P90 quantile forecasts as water constraints. We also documented eight specific limitations and paired each with a concrete remediation pathway (Section VII).

#### IV. PROPOSED SYSTEM ARCHITECTURE

ASICOP follows a microservice design deployed behind a shared API gateway, split across three planes. On the client side, a Progressive Web Application serves reservoir operators and agricultural officers with forecast dashboards and crop-plan outputs; a simpler mobile interface gives farmers field-level valve status and irrigation advisories; and ESP32 microcontrollers in the field push sensor telemetry soil moisture, air temperature, local rainfall, field-level water height over the MQTT event bus while receiving valve-actuation commands in return. The service plane seats seven backend services behind an API gateway on port 8000: JWT authentication, F1 IoT irrigation, F2 crop health, F3 forecasting, F4 optimisation, and an IoT broker that routes MQTT traffic between field devices and the irrigation service. The data plane pairs PostgreSQL for time-series and user data persistence with Redis for cross-service result caching at a ten-minute TTL, Mosquitto for event-bus brokering, and a Prometheus-plus-Grafana stack for operational monitoring.

The cross-module event that distinguishes ASICOP from a loose collection of tools are summarized in Table II.

TABLE II. CROSS-MODULE DATA FLOWS

Producer	Consumer	Data Passed	Mechanism	Effect
F3	F1	Rainfall >5 mm flag	MQTT forecast.update.v1	Suppresses valve-open
F2	F1	stress_penalty_factor	Redis-cached REST	Elevates irrig. priority
F2	F4	stress_penalty_factor	Redis-cached REST	Penalises TOPSIS suit.

F3	F4	P10/P50/P90 release	Redis-cached REST	Bounds area allocation
----	----	---------------------	-------------------	------------------------

When an upstream service goes down, the consuming service falls back to its Redis-cached last-known-good output. If even that cache has expired, the downstream module logs a DEGRADED\_MODE warning and carries on with whatever plan it last committed, so that a single service outage cannot cascade across the platform.

#### V. METHODOLOGY

All four modules are described below using a common structure inputs, features, model, outputs, and integration so that the reader can compare them side by side.

##### A. F1 - IoT Smart Water Management and Field Irrigation Control

**Inputs.** The primary data source is the Udawalawe Hydrological Dataset maintained by the Irrigation Department of Sri Lanka. It arrives as an Excel workbook whose 32 annual sheets span 1 January 1994 through 31 December 2025, recording daily reservoir water level in metres above mean sea level (mMSL), rainfall in millimetres, inflow in million cubic metres (MCM), main-canal and branch-canal releases in MCM, and evaporation estimates. A secondary stream consists of field-level IoT telemetry soil moisture (%), ambient temperature ( $^{\circ}\text{C}$ ), local rainfall (mm), and water height (cm) pushed by ESP32 microcontrollers over MQTT. In the current prototype this telemetry is represented by 1,000 synthetic samples drawn with a fixed random seed (seed = 42) to approximate realistic field-sensor distributions.

**Features.** After iterating and concatenating all 32 sheets, we obtained 11,687 raw daily rows. Cleaning removal of non-numeric entries, rows missing more than one hydrological signal, and instrument-failure outliers left 10,945 usable records. Remaining gaps were forward-filled to maintain the time-series continuity that lagged feature construction depends on. The dataset was then split chronologically without shuffling: the first 9,850 rows (1994-2022) went to training and the final 1,095 rows (2023-2025) were held out for testing. From the ten base hydrological signals the feature-engineering pipeline builds 46 inputs: twenty lag features covering lag-1, lag-2, lag-3, and lag-7 days across five key signals; nine rolling means at 3-, 7-, and 14-day windows over three signals; and three calendar variables (month, day-of-year, binary season indicator).

**Model.** We trained and compared three regressors on the same chronological split. Ridge Regression served as the linear baseline, providing an interpretable reference. Random Forest [2] with 100 estimators was the ensemble tree baseline. HistGradientBoostingRegressor (max\_depth = 6, learning\_rate = 0.07, max\_iter = 250, random\_state = 42) was our selected model; its native handling of missing values and histogram-based splitting suited the multi-year record, where missingness patterns vary across earlier years. Hyperparameters were set through a small grid search over learning\_rate  $\in \{0.05, 0.07, 0.10\}$ , max\_depth  $\in \{4, 6, 8\}$ , and max\_iter  $\in \{150, 250, 400\}$ . Separately, a RandomForestClassifier [2] (n\_estimators = 10, random\_state = 42) was trained on the 1,000 synthetic field-

sensor samples to generate valve-actuation decisions. Its labelling rule reads: EMERGENCY\_RELEASE if reservoir level  $L \geq 95$  mMSL; OPEN if soil moisture  $m < \theta_{\text{crop}}$  AND forecast release  $\hat{r} > r_{\text{min}}$  AND  $L \geq 80$  mMSL; CLOSE otherwise [23]. Because the very same rule produced the training labels, any high accuracy the classifier shows merely confirms that it has learnt the rule, not that it has learnt actual field irrigation behaviour. This caveat is formally recorded in Section VII.




Fig 1: Water Release Predicting

**Outputs and integration.** The HGB regressor outputs a scalar next-day canal-release forecast in MCM/day, broadcast as a reservoir.release.v1 MQTT event and cached in Redis. The RF valve classifier dispatches per-field valve decisions to the IoT broker. When the F3 service signals rainfall above 5 mm within the next 24 hours, F1 defers any pending valve-open commands for that scheduling cycle. When the F2 service raises a crop.stress.v1 event, F1 retrieves the stress penalty factor for the affected zone through a cached service call and bumps irrigation priority proportionally.

### B. F2 - Hybrid Satellite and Image-Based Crop Health Monitoring

**Inputs.** Track A draws on Sentinel-2 Level-2A multispectral imagery covering the Udawalawe scheme perimeter, pulled through Google Earth Engine (project fyp-gee-srilanka, COPERNICUS/S2\_SR catalogue). A 2023 median composite was generated within a 15 km buffer centred on 6.42°N, 80.89°E, yielding a cloud-minimised surface-reflectance mosaic at 10 m spatial resolution. In the production deployment, this composite is simulated from cached band reflectances rather than live GEE API calls, owing to authentication constraints. Track B ingests the PlantVillage dataset [7], which contains roughly 54,306 RGB leaf images distributed across 38 disease-class labels spanning 14 crop species.

**Features.** Track A computes two spectral indices per pixel:  $NDVI = (B8 - B4) / (B8 + B4)$ , which captures canopy greenness via the contrast between red absorption and near-infrared reflectance, and a vegetation water-content index  $NDWI_{\text{veg}} = (B8 - B11) / (B8 + B11)$ , which targets leaf water status through the shortwave-infrared response. Raw band reflectances serve as supplementary signals for the Random Forest. Before any classification, pixels pass through a five-stage validation pipeline: cloud masking, water-body exclusion, urban-area exclusion, a minimum-pixel coverage check per zone, and per-zone index computation. A stratified random sample of 5,000 pixels is then drawn and each pixel assigned a proxy label healthy, mild stress, or severe stress based on published NDVI/NDWI threshold guidelines [17]. Track B features are the intermediate activations of the frozen MobileNetV2 backbone, squeezed to a 1,280-dimensional

embedding by a GlobalAveragePooling2D layer before the trainable classification head. Fig. 6 shows the validation pipeline.




Fig 2: NDVI and NDWI class-boundary logic used to derive zone-level crop-health labels

**Model.** For Track A, a Random Forest classifier [2] was fitted on the 5,000-pixel sample and returned roughly 99% accuracy on its held-out pixel subset. That number reflects the model faithfully reproducing the NDVI/NDWI threshold rule that generated the labels; it is not evidence of agreement with any independent agronomic ground truth, and we flag this explicitly in Section VII. For Track B, we adopted a MobileNetV2 backbone [8] pre-trained on ImageNet, froze its convolutional layers, and attached a classification head comprising GlobalAveragePooling2D  $\rightarrow$  Dense(128, ReLU)  $\rightarrow$  Dense(38, Softmax), giving approximately 49,190 trainable parameters. Training used the Adam optimiser at a learning rate of 0.001 with categorical cross-entropy loss, a batch size of 32, and ran for 10 epochs across about 1,358 batches per epoch, totalling around 5.8 hours of compute. The highest validation accuracy of 95.43% appeared at epoch 9; the lowest validation loss of 0.1452 appeared at epoch 4.

**Outputs and integration.** Track A delivers a zone-level health classification map (healthy / mild stress / severe stress). Track B returns a 38-class disease label with a confidence score for a submitted leaf image. The F2 service fuses these into a scalar stress\_penalty\_factor via the formula  $\text{stress\_index} = \min(1.0, (\text{mild\_ratio} \times 0.5) + \text{severe\_ratio})$ , where the half-weight on mild stress encodes the judgement that mild stress should raise irrigation priority but should not, on its own, penalise crop suitability as harshly as confirmed severe stress. When the zone stress index crosses the alert threshold, F2 publishes a crop.stress.v1 event. F1 then raises irrigation priority for the affected zone, and F4 applies  $\text{effective\_suitability} = \text{base\_suitability} \times (1 - \text{penalty\_factor})$  before Fuzzy-TOPSIS ranking.

### C. F3 - ML Time-Series Forecasting and Alerting

**Inputs.** The F3 module works from the same Udawalawe Hydrological Dataset as F1, supplemented by live weather data from the Open-Meteo API (temperature, precipitation, humidity, wind speed for the Udawalawe catchment coordinates). The production service iterates all 32 sheets programmatically, exactly as F1 does. The notebook prototype, however, calls `pd.read_excel(DATA_PATH, header=2)` without iterating sheets, inadvertently loading only the first sheet year 1994 and restricting the training set to just 286 rows after cleaning from the original 365. This loading error is the direct cause of every negative  $R^2$  value reported for the notebook experiments, and we document it as the module's primary limitation in Section VII.

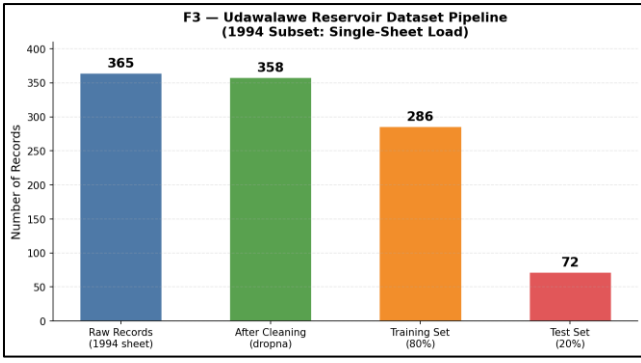


Fig 3: Dataset loading and preparation pipeline for the F3 forecasting module

**Features.** The supervised-model feature set comprises lagged water-level and rainfall values at lag-1, lag-2, lag-3, and lag-7 days, rolling means at 7-day and 30-day windows, and calendar variables (month, season indicator). Feature-importance analysis on the best notebook Gradient Boosting model placed `water_level_lag_1` at the top with 49.3% of total importance, consistent with the strong day-to-day autocorrelation one would expect for reservoir levels.

**Model.** The production service organises its forecasting logic into five layers: (1) a Ridge Regression baseline that supplies a linear trend extrapolation; (2) seasonal and non-seasonal ARIMA models fitted by AIC grid search to capture the bimodal Yala–Maha cycle; (3) three advanced supervised models Random Forest [2], Gradient Boosting [3], and a stacked LSTM [4] whose architecture runs LSTM(64, return\_sequences=True) → Dropout(0.2) → LSTM(32) → Dropout(0.2) → Dense(16) → Dense(1) for a total of 32,673 trainable parameters, trained with MSE loss, Adam, and early stopping on validation loss; (4) an ensemble aggregator whose combination weights are fitted on the validation partition; and (5) an anomaly-detection layer where Isolation Forest, Local Outlier Factor, DBSCAN, and Z-score detectors vote by majority on whether an incoming reading is anomalous before it enters the pipeline. In parallel, three quantile Gradient Boosting models [3] are trained at  $\alpha = 0.10, 0.50, \text{ and } 0.90$ , each minimising the pinball loss, with post-processing that sorts the three predictions at every step to enforce  $P10 \leq P50 \leq P90$  and prevent quantile crossing.

**Outputs and integration.** The module publishes daily point forecasts of reservoir water level (mMSL) and canal release (MCM/day), P10/P50/P90 quantile bands over the configured horizon, and drought or flood alert flags whenever the P50 level breaches the 80 mMSL minimum-safe or 95 mMSL emergency threshold. Alerts flow to the operator dashboard and to downstream services via `forecast.update.v1` events. F1 subscribes to suppress or accelerate valve actions; F4 reads the current P10/P50/P90 daily release forecasts through a Redis-cached REST call before computing the water-constrained area allocation, taking P50 as the primary bound and P10 as the conservative floor.

#### D. F4 - Adaptive Crop and Area Optimisation

**Inputs.** The F4 module consumes four datasets. Its primary source is the Hector Kobbekaduwa Agrarian Research and Training Institute Government Retail Price dataset [25], which contains 71,737 weekly retail price records for five vegetable crops tomatoes, carrot, green beans, long beans,

and leeks across 37 Sri Lankan districts from 2015 to 2024. Three supplementary datasets widen the feature space: a Sri Lanka Climate dataset (Kaggle, rasulmah/sri-lanka-weather-data; 314,000 daily records) for weather covariates, a Paddy Cultivation dataset (Kaggle; 1,039 seasonal records) for rice production context, and a Rice Time Series dataset (Kaggle; 324 monthly records) for paddy price modelling. Paddy does not appear in the Hector dataset a gap we discuss under limitation L6 in Section VII.

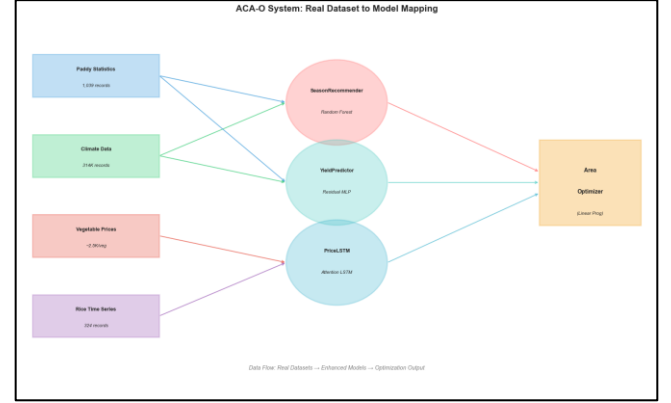


Fig 4: Dataset-to-model mapping for the crop recommendation, price prediction, suitability scoring, and area-allocation components

**Model.** The F4 stack has four components. The *Crop Recommender* is a RandomForestClassifier [2] with 100 estimators, trained on the Hector-derived feature set after applying leakage controls: the target crop label and all direct crop-identity encodings are excluded from the input matrix, preprocessing is fitted on the training partition only, and duplicate district–season–crop tuples are removed before splitting. On the held-out test partition the classifier reaches 97.0% accuracy (training accuracy 98.5%) with a macro-averaged F1 of 0.97, which is 77 percentage points above the 20% random baseline for five classes. We interpret this as strong separability within the available market-feature space; it is not, on its own, a validation of agronomic suitability for a real planting decision. The *Price Predictor* is a dual-embedding neural network (PricePredictorNN) implemented in PyTorch:  $\text{Embedding}(n_{\text{crops}}, 8) \oplus \text{Embedding}(n_{\text{districts}}, 8) \oplus x_{\text{cont}} \rightarrow \text{Linear}(d_{\text{in}}, 32) \rightarrow \text{ReLU} \rightarrow \text{Dropout}(0.2) \rightarrow \text{Linear}(32, 1)$ , totalling 5,409 parameters, trained with MSE loss and Adam at  $\text{lr} = 0.001$ . It reaches  $R^2 = 0.167$ ,  $\text{MAE} = \text{Rs. } 115.66/\text{kg}$ , and  $\text{RMSE} = \text{Rs. } 175.81/\text{kg}$  on the test partition. Five-fold cross-validation gives a mean  $R^2$  of  $0.155 \pm 0.009$  with a train/validation  $R^2$  ratio of 1.01, confirming stable generalisation with limited overfitting. A Random Forest [2] trained in log-price space reaches  $R^2 = 0.486$ , outperforming the neural network; the intended production predictor is LightGBM [5] with 24 features, but that model has not yet been run. The *Fuzzy-TOPSIS Suitability Scorer* evaluates each crop–zone combination on five criteria soil suitability (weight 0.25), water coverage ratio (0.25), historical yield in tonnes per hectare (0.20), water sensitivity (0.15), and growth duration in days (0.15). Criterion scores are mapped to trapezoidal fuzzy numbers and processed through the standard six-step Fuzzy-TOPSIS algorithm [6]; F2 stress penalties are applied before the final ranking as  $\text{effective\_suitability} =$

$\text{base\_suitability} \times (1 - \text{penalty\_factor})$ . The *Area Allocator* currently applies a greedy heuristic, ranking crops by expected revenue per millimetre of water (Rs/mm) and assigning quota in descending order until exhausted. A PuLP mixed-integer programming formulation that would support rotation constraints and multi-season planning is coded in the service but is not yet wired as the active solver (see limitation L5).

**Outputs and integration.** The module delivers a ranked crop-suitability table per field zone (Fuzzy-TOPSIS scores after stress adjustment), a quota-aware area allocation plan giving hectares per crop per zone calibrated to P50 water availability, and expected-revenue estimates under P10, P50, and P90 water scenarios. F4 reads the current quantile forecasts from F3 and zone-level stress summaries from F2 before every allocation cycle. Seasonal crop plans are published to the API gateway for the operator dashboard and passed to F1 as expected crop-water requirements for the upcoming season.

## VI. EXPERIMENTAL SETUP AND RESULTS

Table III collects the dataset, record count, provenance, and train/test split for every module.

TABLE III. MODULE DATASETS AND PARTITIONS

Module	Dataset	Records	Source	Split
F1, Release reg.	Udawalawe 1994–2025 (32 sheets)	11,687→10,945	Irrig. Dept. SL	Chrono: 9,850/1,095
F1, Valve cls.	Synthetic (seed=42)	1,000	Generated	80/20 random
F2, RF sat.	Sentinel-2 2023 composite	5,000 px, 3 cls.	Copernicus/GE E	80/20 stratified
F2, MobileNet	PlantVillage (38 cls.)	~54,306 img.	Mohanty et al. [7]	80/20 cls-strat.
F3, Notebook	Udawalawe 1994 only	365→358	Irrig. Dept. SL	286/72
F3, Prod. svc.	Udawalawe 1994–2025 (not run)	10,945 (proj.)	Irrig. Dept. SL	Proj. chrono.
F4, Crop rec.	Hector Prices 2015–2024	71,737, 5 crops	Hector ARTI	Holdout
F4, Price pred.	Hector + Climate	71,737+314k	Hector; Kaggle	5-fold CV

**F1 results.** Table IV shows how the three regression models performed on the chronological 2023–2025 holdout ( $n = 1,095$  daily records).

TABLE IV. F1 RELEASE REGRESSOR COMPARISON ( $N = 1,095$ )

Model	MAE (MCM/d)	RMSE (MCM/d)	R <sup>2</sup>
Ridge Regression (baseline)	0.5191	0.7497	0.7718
Random Forest (baseline) [2]	0.4366	0.8514	0.7058
HistGradBoostingReg. (selected)	0.4412	0.7108	0.7949

The HistGradientBoostingRegressor was selected on the basis of RMSE and R<sup>2</sup>. Random Forest returned a slightly lower MAE, but we prioritised RMSE because underestimating a large release event carries a heavier operational penalty potential over-draw and downstream flood risk than minor day-to-day forecast noise. We frame this model as advisory-grade release prediction suitable for driving the rain-suppression and gate-timing logic, not as a validated signal for autonomous gate operation.

**F2 results.** On the 38-class PlantVillage validation partition the MobileNetV2 classifier [8] peaked at 95.43% accuracy (epoch 9) with a best validation loss of 0.1452 (epoch 4). This falls within the range commonly reported for ImageNet-pretrained transfer learners on this dataset [19]. With only about 49,190 trainable parameters, the model is compact, though the gap between its best-accuracy and best-loss checkpoints points to mild overfitting in the final training epochs. We stress that these numbers represent benchmark validation on PlantVillage, not field validation in the Udawalawe command area. PlantVillage images are photographed under controlled single-leaf conditions [7], so accuracy on real smallholder field photographs will almost certainly be lower, and we have not measured by how much. Paddy, the dominant crop grown under the Udawalawe scheme, is entirely absent from PlantVillage's species list. The Sentinel-2 Random Forest track, which returned roughly 99% accuracy on its proxy-labelled subset, is best understood as a test of how well the model learnt the labelling rule, not as an independent measure of stress-detection capability.

**F3 results.** Table V reports the notebook results on the 1994-only test partition ( $n = 72$ ). Every metric is coloured by the single-sheet loading bug.

TABLE V. F3 NOTEBOOK COMPARISON

Model	RMSE (mMSL)	MAE (mMSL)	R <sup>2</sup>
Gradient Boosting [3]	2.8763	2.7027	-2.2912
Random Forest [2]	3.1813	3.0068	-3.0260
LSTM [4] (32,673 params)	3.6707	3.2944	-17.5414
P80 quantile coverage	-	-	8.33% ( $\geq 70\%$ )

A negative R<sup>2</sup> tells us that every one of these models predicted worse than a simple mean on the 72-record test set. That outcome is entirely expected when training data is drawn from just one hydrological year: 286 rows from a single monsoon season cannot teach a model the inter-annual variability that shapes the reservoir's multi-decade trajectory. These results do not indict gradient boosting [3] or LSTM [4] as methods; they indict the inadvertent reduction of a 32-year dataset to a single-year subset.

**F4 results.** On the held-out partition of the five-class crop-recommendation task the Random Forest classifier achieved 97.0% accuracy with a macro-averaged F1 of 0.97. Per-class F1 scores ranged from 0.96 (green beans) to 0.98 (carrot); the highest precision was 0.98 for leeks. We take this as evidence of strong class separability within the market-feature space, not as evidence that the classifier by itself validates agronomic crop suitability it still has no access to soil-type, nutrient, or crop-choice labels. The price predictor (PricePredictorNN) returned  $R^2 = 0.167$  (MAE Rs.

115.66/kg, RMSE Rs. 175.81/kg); the Random Forest log-scale baseline outperformed it at  $R^2 = 0.486$ . Five-fold cross-validation confirmed the neural network's result is stable rather than a single-split artefact (mean  $R^2 = 0.155 \pm 0.009$ ). We therefore describe the neural network as an exploratory model and note that inside ASICOP the area allocator treats predicted prices as advisory inputs, relying primarily on suitability and water-efficiency scores to drive the allocation.

**Cross-module integration.** All four data flows listed in Table III have been exercised end-to-end at the service-log level: F3 rainfall events trigger the F1 valve-suppression branch, F2 stress events raise F1 irrigation priority for the affected zone, F2 stress penalties propagate into the F4 Fuzzy-TOPSIS scoring step, and F3 P10/P50/P90 quantile values are read by the F4 area allocator before each cycle. What we have verified is that the wiring works correctly. What we have not yet demonstrated is that this integrated platform makes better operational decisions than the status quo of uncoordinated, single-stream tools.

## VII. LIMITATIONS

We record eight specific limitations. Each is stated with its root cause and its bearing on the results claimed above.

**L1 - F1 synthetic-only valve training.** The RF valve classifier was trained exclusively on 1,000 synthetic sensor samples generated from a two-condition rule (OPEN if moisture < 30% AND temperature > 25°C). That rule captures none of the variation in crop growth stage, inter-field soil differences, real sensor noise patterns, or micro-climate effects that shape actual irrigation need. The classifier is therefore a rule-informed prototype, not a validated field-level smart-irrigation controller. Real-world retraining on ESP32 telemetry is required before any operational reliance.

**L2 - F2 MobileNetV2: no paddy class.** Paddy is the dominant crop in the Udawalawe command area, yet it is absent from PlantVillage's roster of 14 species and 38 disease classes [7]. Whether the classifier generalises usefully to rice blast, bacterial leaf blight, or sheath blight is unknown and untested. The 95.43% PlantVillage accuracy says nothing about paddy-field performance.

**L3 - F2 RF satellite: proxy-derived labels.** The RF zonal classifier was trained on 5,000 pixels whose labels were derived by applying NDVI/NDWI threshold rules to the very same Sentinel-2 median composite the model is asked to classify. The resulting ~99% accuracy shows that the model can reproduce the labelling rule; it does not demonstrate agreement with independent agronomic ground-truth measurements of crop water stress.

**L4 - F3 single-sheet loading.** The F3 notebook calls `pd.read_excel` without iterating sheets, loading only the first year-sheet (1994) and restricting training to 286 rows from a single hydrological year. Every negative  $R^2$  value reported in the F3 results Gradient Boosting at -2.29, Random Forest at -3.03, LSTM at -17.54 traces directly to this loading error, not to any inherent inadequacy of the models.

**L5 - F4 greedy area allocator.** The active area allocator ranks crops by expected revenue per millimetre of water and fills quota in descending order. This greedy approach cannot enforce inter-crop rotation constraints, multi-season planning

horizons, or field-specific agronomic compatibility rules. A PuLP mixed-integer programming formulation that handles these constraints is coded in the F4 service but has not been wired as the active solver.

**L6 - Hector dataset excludes paddy.** The Hector Kobbekaduwa Government Retail Price dataset covers five vegetable crops but not paddy. Paddy price modelling within F4 relies on a separate 324-record monthly Rice Time Series dataset far smaller and at a coarser temporal resolution than the 71,737-record weekly Hector data.

**L7 - Quota enforcement is advisory.** ASICOP offers decision support within the existing administrative quota structure. It can recommend reservoir releases, crop plans, and valve actions, but it cannot physically compel farmers or field operators to follow those recommendations. Actual irrigation outcomes depend on human adherence.

**L8 - Udawalawe-specific calibration.** Every numerical threshold in ASICOP the 80 mMSL minimum-safe reservoir level, the 95 mMSL emergency release trigger, the 30% valve moisture threshold, the Fuzzy-TOPSIS criterion weights of [0.25, 0.25, 0.20, 0.15, 0.15], and the Rs/mm efficiency ranking in the allocator is calibrated to the Udawalawe scheme's particular hydrology, cropping calendar, and governance rules. Deploying the platform to a different Sri Lankan scheme would require re-parameterising all of these against that scheme's own hydrological record and agronomic data.

## VIII. CONCLUSION AND FUTURE WORK

This paper has described ASICOP, a prototype AI- and IoT-based decision-support framework that links reservoir forecasting, field-level irrigation control, satellite and image-based crop-health monitoring, and adaptive crop-area planning for the Udawalawe quota-based irrigation scheme in Sri Lanka. The four modules communicate through MQTT events, Redis-cached service calls, and graceful-degradation fallbacks. On a 32-year chronological holdout the F1 release regressor ( $R^2 = 0.7949$ , RMSE = 0.7108 MCM/day) provides advisory-grade daily release prediction. The F2 disease classifier (95.43% on PlantVillage) supplies a point-of-care leaf-screening tool, while the Sentinel-2 zonal track sorts the command area into zones needing attention versus zones that do not. The F4 crop and area planning stack brings together a market-signal recommender, a Fuzzy-TOPSIS suitability ranker, and a water-aware area allocator fed by F3 quantile forecasts. The F3 notebook, hampered by a single-sheet data-loading error, yields negative  $R^2$  on all three supervised models and is best read as an architecture proposal awaiting a corrected data run. No end-to-end field measurement of water savings, flood-damage reduction, or yield improvement has been carried out, so the integration contribution remains conceptual rather than empirically validated at this stage.

Six concrete follow-up actions emerge from the eight limitations above. (R1) Correcting the F3 `pd.read_excel` call to iterate all 32 sheets and re-running the Gradient Boosting, Random Forest, and LSTM models on the full 10,945-row dataset is the single highest-impact step. (R2) Deploying ESP32 sensors in instrumented Udawalawe fields for six to twelve months will generate the real telemetry needed to retrain the F1 valve classifier beyond its current synthetic-

rule foundation. (R3) Collecting a paddy-specific leaf-image dataset covering rice blast, bacterial leaf blight, and sheath blight and fine-tuning a second MobileNetV2 or EfficientNet-B0 classifier would close the dominant-crop gap in F2. (R4) Wiring the PuLP MIP formulation as the active F4 area-allocation solver, once the rotation-constraint matrix is defined and validated against the greedy baseline, replaces the current heuristic with a globally optimal planner. (R5) Adding paddy as a sixth class to both the recommendation and price-prediction models, backed by the LightGBM [5] production predictor with its 24-feature design, will align ASICOP with the actual dominant land use of the Udawalawe command area. (R6) Running the platform across a set of instrumented farmer fields for one full Maha or Yala cropping season measuring water use, crop yield, disease incidence, and farmer adherence against matched control fields is the experiment that will determine whether the integrated design delivers measurable operational improvements over the current fragmented status quo.

## IX. ACKNOWLEDGMENTS

The authors wish to thank the Irrigation Department of Sri Lanka for providing the 32-year Udawalawe daily hydrological record; the European Space Agency and the EU Copernicus Programme for Sentinel-2 imagery, which was accessed through the Google Earth Engine platform [24]; Mohanty, Hughes, and Salathé for making the PlantVillage dataset publicly available [7]; the Hector Kobbekaduwa Agrarian Research and Training Institute, Government of Sri Lanka, for the Government Retail Price dataset [25]; the Kaggle community contributors behind the Sri Lanka Weather Data (rasulmah/sri-lanka-weather-data), Paddy Cultivation SL, and Rice Time Series SL datasets; and our academic supervisors Ms. Hansi de Silva and Ms. Karthiga Rajendran at the Department of Software Engineering, Sri Lanka Institute of Information Technology (SLIIT), for their guidance throughout this research.

## X. REFERENCES

- [1] D. Herath, L. Weerakoon, D. Dissanayake, and S. Seneviratne, "Advancing food sustainability: A case study on improving rice yield prediction in Sri Lanka using weather-based, feature-engineered machine learning models," *Discover Applied Sciences*, vol. 6, no. 11, Art. no. 580, 2024. doi: 10.1007/s42452-024-06300-7.
- [2] L. Breiman, "Random forests," *Machine Learning*, vol. 45, no. 1, pp. 5–32, Oct. 2001.
- [3] J. H. Friedman, "Greedy function approximation: A gradient boosting machine," *The Annals of Statistics*, vol. 29, no. 5, pp. 1189–1232, Oct. 2001.
- [4] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, Nov. 1997.
- [5] G. Ke et al., "LightGBM: A highly efficient gradient boosting decision tree," in *Proc. NeurIPS*, 2017, pp. 3146–3154.
- [6] C.-T. Chen, "Extensions of the TOPSIS for group decision-making under fuzzy environment," *Fuzzy Sets and Systems*, vol. 114, no. 1, pp. 1–9, Aug. 2000.
- [7] S. P. Mohanty, D. P. Hughes, and M. Salathé, "Using deep learning for image-based plant disease detection," *Frontiers in Plant Science*, vol. 7, p. 1419, Sep. 2016.
- [8] M. Sandler et al., "MobileNetV2: Inverted residuals and linear bottlenecks," in *Proc. IEEE/CVF CVPR*, 2018, pp. 4510–4520.
- [9] Food and Agriculture Organization of the United Nations (FAO), "AQUASTAT Country Profile – Sri Lanka," FAO, Rome, Italy, 2012. [Online]. Available: <https://www.fao.org/aquastat>.
- [10] International Water Management Institute (IWMI), "Water management in the Walawe River Basin: Prospects and problems," IWMI Research Report, Colombo, Sri Lanka, 2006.
- [11] Department of Agriculture, Sri Lanka, "Agricultural Statistics 2021/2022," Ministry of Agriculture, Colombo, Sri Lanka, 2022.
- [12] G. Arnold et al., "Large area hydrologic modeling, Part I," *J. AWRA*, vol. 34, no. 1, pp. 73–89, 1998.
- [13] U.S. Army Corps of Engineers, "HEC-ResSim User's Manual v3.1.," Davis, CA, 2013.
- [14] F. Kratzert et al., "Rainfall–runoff modelling using LSTM networks," *HESS*, vol. 22, pp. 6005–6022, 2018.
- [15] G. Papacharalampous, H. Tyralis, A. Langousis, A. W. Jayawardena, B. Sivakumar, N. Mamassis, A. Montanari, and D. Koutsoyiannis, "Probabilistic hydrological post-processing at scale: Why and how to apply machine-learning quantile regression algorithms," *Water*, vol. 11, no. 10, Art. no. 2126, Oct. 2019. doi: 10.3390/w11102126
- [16] G. P. Pereira, M. Z. Chaari, and F. Daroge, "IoT-enabled smart drip irrigation system using ESP32," *IoT*, vol. 4, no. 3, pp. 221–243, Jul. 2023. doi: 10.3390/iot4030012
- [17] European Space Agency (ESA), "Sentinel-2 User Handbook," ESA Standard Document, Issue 1, Rev. 2, Jul. 2015. [Online]. Available: <https://sentinels.copernicus.eu/documents/247904/685211/Sentinel-2\ User\ Handbook>
- [18] R. Sonobe, Y. Yamaya, H. Tani, X. Wang, N. Kobayashi, and K. Mochizuki, "Crop classification from Sentinel-2-derived vegetation indices using ensemble learning," *J. Applied Remote Sensing*, vol. 12, no. 2, Art. no. 026019, May 2018. doi: 10.1117/1.JRS.12.026019
- [19] E. C. Too, L. Yujian, S. Njuki, and L. Yingchun, "A comparative study of fine-tuning deep learning models for plant disease identification," *Computers and Electronics in Agriculture*, vol. 161, pp. 272–279, Jun. 2019. doi: 10.1016/j.compag.2018.03.032
- [20] L. M. A. Santos, A. M. Costa, M. N. Arenales, and R. H. S. Santos, "Mixed integer linear programming models for optimal crop selection," *Computers & Operations Research*, vol. 81, pp. 26–39, May 2017. doi: 10.1016/j.cor.2016.12.010
- [21] S. Pudumalar, E. Ramanujam, R. H. Rajashree, C. Kavya, T. Kiruthika, and J. Nisha, "Crop recommendation system for precision agriculture," in *Proc. IEEE 8th Int. Conf. Advanced Computing (ICoAC)*, Chennai, India, Jan. 2017, pp. 32–36. doi: 10.1109/ICoAC.2017.7951740
- [22] M. R. N. Qureshi, R. K. Singh, and M. A. Hasan, "Decision support model to select crop pattern for sustainable agricultural practices using fuzzy MCDM," *Environment, Development and Sustainability*, vol. 20, no. 2, pp. 641–659, Apr. 2018. doi: 10.1007/s10668-016-9903-7
- [23] Irrigation Department of Sri Lanka, "Operational guidelines for Udawalawe Reservoir," Internal Technical Document, Ministry of Irrigation, Colombo, Sri Lanka (unpublished).
- [25] Hector Kobbekaduwa Agrarian Research and Training Institute (HARTI), "Weekly food commodities bulletin," Government of Sri Lanka, Colombo, 2024. [Online]. Available: <https://www.harti.gov.lk/index.php/en/market-information/weekly-food-commodities-bulletin>
- [26] P. Zippenfenig, "Open-Meteo.com Weather API," Zenodo, 2023. doi: 10.5281/zenodo.7970649. [Online]. Available: <https://open-meteo.com>.