

**INTEGRATED SMART WATER-FOCUSED  
IRRIGATION SYSTEM USING IOT AND AI/ML**

(Water Forecasting & Decision-Support System)

Project ID: 25-26J-520

Trishni W.R.M.

(Dilruksha A.G.C.D., Hesara P.K.A.N., Abishek S)

B.Sc. (Hons) in Information Technology Specializing in Software  
Engineering

Department of Software Engineering

Sri Lanka Institute of Information Technology Malabe, Sri Lanka

April 2026

**INTEGRATED SMART WATER-FOCUSED  
IRRIGATION SYSTEM USING IOT AND AI/ML**

(Water Forecasting & Decision-Support System)

Project ID: 25-26J-520

Trishni W.R.M. - IT22076366

Hesara P.K.A.N. - IT22561398

Dilruksha A.G.C.D. - IT22561770

Abishek S - IT22186942

Supervisor: Ms. Hansi De Silva

Co-Supervisor: Ms. Karthiga Rajendran

B.Sc. (Hons) in Information Technology Specializing in Software  
Engineering

Department of Software Engineering

Sri Lanka Institute of Information Technology Malabe, Sri Lanka

April 2026

## DECLARATION

I declare that this is my own work, and this Thesis does not incorporate without acknowledgement any material previously submitted for a degree or diploma in any other University or institute of higher learning and to the best of my knowledge and belief it does not contain any material previously published or written by another person except where the acknowledgement is made in the text.

Also, I hereby grant to Sri Lanka Institute of Information Technology, the nonexclusive right to reproduce and distribute my Thesis, in whole or in part in print, electronic or other medium. I retain the right to use this content in whole or part in future works (such as articles or books).

Name	Student ID	Signature
Trishni W.R.M.	IT22076366	

The above candidate has carried out this research thesis for the Degree of Bachelor of Science (honors) Information Technology (Specializing in Software Engineering) under my supervision.

-----  
Signature of the Supervisor

-----  
Date

-----  
Signature of the Co-supervisor

-----  
Date

## **ABSTRACT**

Reservoir management in Sri Lanka is critically challenged by the dual hazards of uncontrolled spill-over flooding during intense monsoon periods and prolonged droughts that threaten food production and hydropower generation. The Udawalawe Reservoir on the Walawe River Basin exemplifies these competing pressures, yet its operations continue to rely on static rule curves and ad-hoc engineering judgement rather than on data-driven forecasting. This research proposes, designs, and evaluates an Integrated Smart Water Forecasting and Decision-Support System for the Udawalawe Reservoir as an individual function within the broader Integrated Smart Water-Focused Irrigation System (Project ID: 25-26J-520) developed at SLIIT

The proposed function ingests over three decades (1994–2025) of historical hydrological, meteorological, and operational records together with real-time telemetry and short-term weather forecasts from the Meteorology Department and the Irrigation Department. A tiered portfolio of forecasting models spanning classical statistical baselines (ARIMA, SARIMA, and Facebook Prophet), machine-learning regression models (Random Forest and Gradient Boosting Machines), and deep-learning recurrent networks (Long Short-Term Memory and Gated Recurrent Unit architectures) is developed and benchmarked against metrics including RMSE, MAE, and MAPE for prediction horizons of 1 to 14 days. An ensemble-based probabilistic output layer generates confidence bands and risk-level classifications for flood and drought conditions, explicitly addressing the uncertainty-quantification gap that afflicts current deterministic operational practice.

Forecast outputs are coupled to reservoir operations through an augmented rule-curve decision-logic engine and a what-if scenario simulation module that translates predictions into actionable release and conservation recommendations. These are delivered to end-users — reservoir engineers, irrigation officers, and disaster-management authorities — through an interactive web-based dashboard featuring forecast visualisation with shaded uncertainty bands, colour-coded risk indicators, and a downstream-area map overlay, augmented by an automated multi-channel alert service dispatching SMS and email notifications when critical thresholds are breached.

Expected outcomes include achieving forecast accuracy below 15% MAPE for short-horizon inflow predictions, demonstrable outperformance of ARIMA baselines across multi-day horizons, well-calibrated probabilistic intervals, earlier warning of extreme events, optimised

proactive water releases, and enhanced inter-agency coordination. The system is designed with a modular, cloud-compatible architecture enabling extension to additional Sri Lankan reservoirs and commercialisation as a subscription-based Water Management Decision Support Platform.

**Keywords:** *reservoir forecasting, LSTM, GRU, probabilistic forecasting, decision support system, smart irrigation, IoT, time-series prediction, Udawalawe Reservoir, early warning system, water management, ARIMA*

## **ACKNOWLEDGEMENT**

I wish to express my sincere gratitude to all those who have supported and guided me throughout the course of this research.

I am deeply grateful to my supervisor, Ms. Hansi De Silva, and co-supervisor, Ms. Karthiga Rajendran, of the Department of Software Engineering at the Sri Lanka Institute of Information Technology, for their invaluable guidance, constructive feedback, and continued encouragement at every stage of this research. Their expertise in machine learning and water systems has been a constant source of inspiration.

I would also like to extend my appreciation to the academic staff of the Department of Software Engineering at SLIIT for their support and for providing the learning environment that made this research possible. Special thanks are due to the staff of the Irrigation Department and the Department of Meteorology, Sri Lanka, for their willingness to share hydrological and meteorological data that form the empirical foundation of this study.

I am thankful to my group members - Hesara P.K.A.N. (IT22561398), Dilruksha A.G.C.D. (IT22561770), and Abishek S (IT22186942) for their collaboration, shared ideas, and teamwork within the broader Integrated Smart Water-Focused Irrigation System project. The integration of our individual modules has been greatly strengthened by our collective effort.

Finally, I owe heartfelt thanks to my family for their unwavering moral support, patience, and encouragement throughout this academic journey. This work would not have been possible without them.

## Table of Contents

DECLARATION .....	i
ABSTRACT.....	ii
ACKNOWLEDGEMENT .....	iv
LIST OF FIGURES .....	vi
LIST OF TABLES.....	vii
LIST OF ABBREVIATIONS.....	viii
1. CHAPTER 1 - INTRODUCTION.....	1
1.1 Background Study and Literature Review.....	1
1.1.1 Background Study.....	1
1.1.2 Literature Review.....	3
1.2 Research Gap .....	7
1.3 Research Problem .....	9
1.4 Research Objectives.....	10
1.4.1 Main Objective.....	10
1.4.2 Specific Objectives .....	10
1.4.3 Business Objectives .....	11
2. CHAPTER 2 - METHODOLOGY .....	12
2.1 Methodology.....	12
2.1.1 Feasibility Study / Planning.....	12
2.1.2 Requirement Gathering & Analysis.....	15
2.1.3 Designing .....	17
2.1.4 Implementation .....	22
2.1.5 Testing.....	24
2.1.6 Deployment & Maintenance .....	26
2.2 Commercialisation .....	27
3. CHAPTER 3 - RESULTS & DISCUSSION.....	30
4. CHAPTER 4 - FUTURE SCOPE.....	33
5. CHAPTER 5 - CONCLUSION .....	35
6. REFERENCES .....	37

## LIST OF FIGURES

Figure 1.1: Flood Situation Map, Sri Lanka (November 2023).....	1
Figure 1.2 : Drought Situation Map, Sri Lanka (September 2023) .....	2
Figure 1.3 : Udawalawe Reservoir Location and Catchment Area .....	3
Figure 2.1:Feasibility Study Summary .....	13
Figure 2.2:Gantt Chart — Project Timeline (March 2025 to June 2026).....	14
Figure 2.3:Proposed System Architecture Overview.....	18
Figure 2.4:End-to-End Processing Pipeline Flowchart.....	19
Figure 2.5:Use Case Diagram — Decision-Support System.....	20
Figure 2.6:Sequence Diagram — Forecast-to-Alert Workflow.....	21
Figure 2.7:Data Flow Diagram — Preprocessing and Forecasting .....	22
Figure 2.8:Dataset Preparation Pipeline — Raw, Cleaned, Train/Test Split.....	23
Figure 2.9:Deployment Architecture Diagram .....	27

## LIST OF TABLES

Table 1.1:Representative Literature on Reservoir Forecasting and Operation Optimisation....	5
Table 2.2:Project Budget Breakdown.....	13
Table 2.3 :(Gantt Chart): Project Timeline and Milestones.....	14
Table 2.4 :Risk Management Plan.....	15
Table 2.5 : Communication Plan .....	15
Table 2.6 : Test Case Summary .....	26

## LIST OF ABBREVIATIONS

Abbreviation	Description
AI	Artificial Intelligence
API	Application Programming Interface
ARIMA	Auto-Regressive Integrated Moving Average
CI/CD	Continuous Integration / Continuous Deployment
CEB	Ceylon Electricity Board
DMC	Disaster Management Centre (Sri Lanka)
DL	Deep Learning
DSS	Decision Support System
EWS	Early Warning System
FIRO	Forecast-Informed Reservoir Operations
FSL	Full Supply Level
GRU	Gated Recurrent Unit
GPU	Graphics Processing Unit
IoT	Internet of Things
LSTM	Long Short-Term Memory
MAE	Mean Absolute Error
MAPE	Mean Absolute Percentage Error
ML	Machine Learning
NDRSC	National Disaster Relief Services Centre
PICP	Prediction Interval Coverage Probability
REST	Representational State Transfer
RMSE	Root Mean Square Error
RNN	Recurrent Neural Network
SaaS	Software as a Service

Abbreviation	Description
<b>SARIMA</b>	Seasonal Auto-Regressive Integrated Moving Average
<b>SLIIT</b>	Sri Lanka Institute of Information Technology
<b>SMS</b>	Short Message Service
<b>SSL</b>	Secure Sockets Layer
<b>SVR</b>	Support Vector Regression
<b>TLS</b>	Transport Layer Security
<b>UAT</b>	User Acceptance Testing
<b>UI</b>	User Interface
<b>UX</b>	User Experience

# 1. CHAPTER 1 - INTRODUCTION

## 1.1 Background Study and Literature Review

### 1.1.1 Background Study

Sri Lankan reservoirs perform critical multipurpose roles in supporting irrigation, hydropower, flood control, and both urban and rural water supply. The Udawalawe Reservoir, situated on the Walawe River in the Southern Province of Sri Lanka, is a flagship example of such a competing-objectives reservoir. Constructed in 1968, Udawalawe serves more than 40,000 hectares of irrigated paddy and cash-crop land in the Udawalawe and Hambantota Districts, supports a hydropower plant capable of generating approximately 6 MW, and provides critical water supply to downstream communities. The reservoir has a gross capacity of approximately 268 million cubic metres, and its catchment receives highly variable rainfall influenced by both the southwest and northeast monsoons.

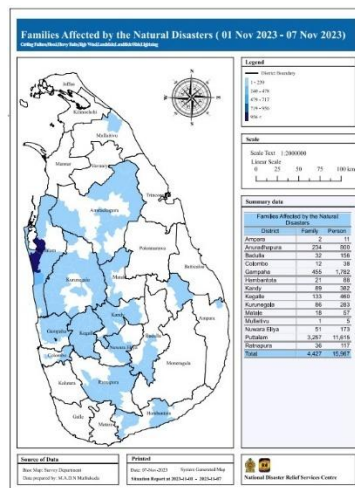


Figure 1.1: Flood Situation Map, Sri Lanka (November 2023)

Climate variability poses mounting operational challenges to Udawalawe. Extreme rainfall events associated with the northeast monsoon produce intense inflow surges that rapidly fill the reservoir and, if not anticipated and managed proactively, lead to uncontrolled spillway releases. Conversely, prolonged dry spells during inter-monsoon periods deplete storage to critically low levels. Historical records attest to both extremes. In late 2012, a severe rainstorm brought Udawalawe to near-overflowing, forcing emergency sluice-gate operations that caused downstream flooding. Earlier the same year, a ten-month dry spell affected nearly one million people and wiped out roughly 23% of the secondary rice crop in Sri Lanka. In November 2023,

extreme rainfall across multiple districts as documented in national disaster situation maps further demonstrated the persistent and intensifying nature of these dual hazards.

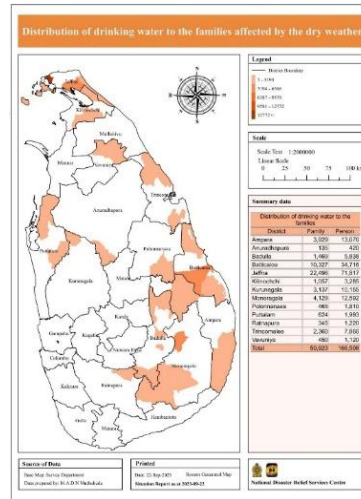


Figure 1.2 : Drought Situation Map, Sri Lanka (September 2023)

Despite this operational context, reservoir management at Udawalawe continues to rely on static rule curves predetermined seasonal storage targets and, on the ad-hoc engineering judgement of experienced irrigation officials. This traditional approach suffers from several well-recognised limitations. First, static rule curves cannot adapt to inflow anomalies that deviate from historical seasonal patterns, which become increasingly common under a changing climate. Second, the absence of forward-looking forecasts means that flood releases are triggered reactively only after the reservoir is already full rather than proactively ahead of predicted high-inflow events. Third, without probabilistic risk information, operators cannot distinguish between a routine monsoon period and a one-in-ten-year flood event. Fourth, decades of hydrological data rainfall measurements, reservoir inflows, water levels, and irrigation release records exist across multiple government agencies (the Irrigation Department, the Meteorology Department, and the Ceylon Electricity Board) but are not systematically integrated or used in day-to-day forecasting models.

This fragmented and reactive paradigm results in suboptimal outcomes: unnecessary flood damage downstream, wasteful water releases during the early wet season that deplete reserves for later dry-season irrigation, and loss of potential hydropower generation. The global literature on reservoir management increasingly advocates a shift toward proactive, data-driven operations anchored in short-to-medium-term forecasting a paradigm exemplified by the Forecast-Informed Reservoir Operations (FIRO) concept pioneered in the United States.

Delaney et al. (2020) demonstrated that FIRO at California's Lake Mendocino increased median water storage by over 30% relative to conventional rule-based operations, with no reduction in flood protection. The central question for Udawalawe, as for many similar reservoirs in the developing world, is whether the recent advances in machine learning and deep learning can be harnessed to deliver a comparable transformation at a cost and complexity level appropriate for a government-operated irrigation scheme in Sri Lanka.

This research addresses that question by designing, implementing, and evaluating an Integrated Smart Water Forecasting and Decision-Support System specifically tailored to the Udawalawe Reservoir. The function constitutes the headworks intelligence layer within the broader Integrated Smart Water-Focused Irrigation System (Project ID: 25-26J-520), providing downstream field-level and crop-optimisation modules with trustworthy projections of canal-release feasibility and water availability. Its development is grounded in a comprehensive review of the machine-learning hydrological forecasting literature and in direct engagement with the operational requirements of the Irrigation Department.

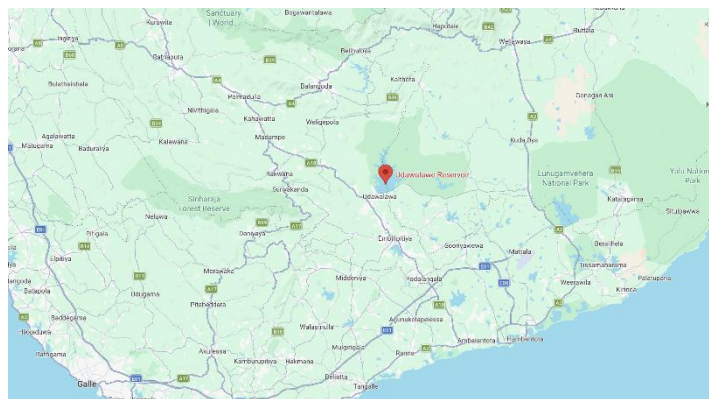


Figure 1.3 : Udawalawe Reservoir Location and Catchment Area

### 1.1.2 Literature Review

The application of machine learning and deep learning to hydrological forecasting has grown dramatically since the early 2010s, driven by the increasing availability of long-term sensor records, advances in computing hardware, and the maturation of open-source deep-learning frameworks. This literature review surveys eight representative and foundational studies directly relevant to the design of the proposed system.

Study (Year)	Focus & Findings	Limitations / Gaps	Relevance to This Study
<b>Li et al. (2024)</b>	LSTM outperforms ARIMA and RF for reservoir inflow forecasting; significant accuracy gains on peak events	Point forecasts only; no uncertainty quantification; not linked to operations	Validates LSTM as primary forecasting engine for Udawalawe
<b>Özdoğan-Sarıkoç et al. (2022)</b>	LSTM achieves highest volume prediction accuracy for a Turkish reservoir vs. ANN and SVR; lower error metrics	Single-site; no probabilistic output	Demonstrates superiority of LSTM over classical ML for reservoir volume forecasting
<b>Ou et al. (2025)</b>	Diffusion-based probabilistic model (DRUM) for runoff forecasting outperforms conventional DL models on 500+ river basins	Computationally intensive; untested on controlled-release reservoirs	Inspires probabilistic forecasting approach; validates confidence-band provision
<b>Delaney et al. (2020)</b>	FIRO at Lake Mendocino increased median water storage by >30% with no loss in flood protection	California-specific; not directly transferable without local validation	Benchmark for quantifying proactive-operations benefit at Udawalawe
<b>Hapuarachchi et al. (2011)</b>	Review of flash flood forecasting advances; highlights role of rainfall-runoff models and short-term predictions	Focus on flash floods rather than reservoir-level operational decisions	Contextualises short-horizon (1-3 day) inflow forecasting within broader literature
<b>Liu et al. (2015)</b>	Bayesian probabilistic model for dynamic flood-limit water level; improves risk-informed operation	Complex Bayesian implementation; limited to flood control	Motivates probabilistic forecasting layer for flood risk quantification
<b>World Bank (2023)</b>	South Asia climate adaptation report highlights increasing flood and drought frequency and urgency of adaptive water management	Policy-level, not technical	Establishes global relevance and urgency of this research

Study (Year)	Focus & Findings	Limitations / Gaps	Relevance to This Study
Ziggah et al. (2022)	Hybrid soft computing models outperform individual ML models for groundwater-level prediction	Groundwater, not surface reservoir; increased complexity	Supports ensemble/hybrid modelling strategy

Table 1.1: Representative Literature on Reservoir Forecasting and Operation Optimisation

Li et al. (2024) conducted a systematic comparative evaluation of statistical, machine-learning, and deep-learning models for reservoir inflow forecasting. Their framework tested ARIMA, Random Forest, and LSTM on a multi-year reservoir dataset. LSTM demonstrated the highest predictive accuracy, substantially outperforming ARIMA in capturing peak inflow events - confirming a pattern observed in multiple earlier single-site studies. This finding validates the decision in the present research to adopt LSTM and GRU as the primary forecasting engine, while retaining ARIMA as an interpretable baseline.

Özdoğan-Sarıkoç et al. (2022) compared artificial neural networks, support vector regression, and LSTM for reservoir volume forecasting at a Turkish reservoir. LSTM achieved the lowest error metrics, including RMSE and MAE, across multiple forecast horizons. The authors attributed LSTM's advantage to its ability to capture complex, non-linear temporal dependencies through its gating mechanism, which prevents the vanishing-gradient problem that afflicts simpler RNN architectures. This study provides direct methodological justification for prioritising LSTM in the present research.

Ou et al. (2025) introduced DRUM, a diffusion-model-based probabilistic framework for extreme flood forecasting across over 500 river basins globally. By leveraging generative AI to produce ensembles of inflow simulations, DRUM yields full probability distributions for extreme events, outperforming conventional deep-learning models in predicting high-magnitude floods and providing useful confidence intervals. While DRUM's computational requirements are substantial and its application to controlled-release reservoirs remains experimental, the study provides strong conceptual motivation for the probabilistic output layer in the present research. The core insight - that quantifying forecast uncertainty dramatically improves the information available to operators - directly informs this project's ensemble-perturbation and quantile-regression approach

Delaney et al. (2020) documented the operational outcomes of the Forecast-Informed Reservoir Operations (FIRO) programme at Lake Mendocino, California. Their empirical analysis

showed that incorporating medium-range ensemble streamflow predictions into reservoir release decisions increased median water storage by over 30% above traditional rule-based operations, without any reduction in flood protection. This landmark real-world validation of forecast-driven reservoir management provides the primary performance benchmark for the present study and demonstrates the transformational potential of the methodology proposed here.

Hapuarachchi et al. (2011) reviewed advances in flash flood forecasting, highlighting the roles of conceptual rainfall-runoff models, statistical approaches, and - increasingly - ensemble numerical weather prediction in extending useful forecast horizons. The review underscores the importance of integrating meteorological forecast data with hydrological models for operationally valuable lead times. This finding directly motivates the present research's decision to ingest short-term precipitation forecasts from the Meteorology Department as model inputs, enabling the forecasting engine to extend useful accuracy beyond what is achievable from historical sequences alone.

Liu et al. (2015) proposed a Bayesian probabilistic framework for dynamic control of the flood-limit water level in reservoir operation, demonstrating improved risk-informed decision-making relative to deterministic approaches. Their work established the theoretical and practical basis for providing operators with probability-of-exceedance information rather than point forecasts, which is a defining feature of the system proposed here. The study's operationalisation of probabilistic thresholds (for example, a recommended release trigger when the probability of the reservoir exceeding a critical level within seven days exceeds a specified value) closely mirrors the decision-logic layer of the present research.

The World Bank (2023) South Asia Climate Adaptation Report documented the accelerating frequency and severity of climate-related water stress events across South Asia, identifying adaptive water management - including the adoption of forecast-informed reservoir operations - as a priority intervention for climate resilience. This report contextualises the present research within a broader regional imperative and strengthens the case for the commercialisation pathway proposed in this thesis, wherein the system would be extended beyond Udawalawe to other vulnerable reservoirs across Sri Lanka and South Asia.

Ziggah et al. (2022) demonstrated that hybrid soft-computing models, combining multiple individual machine-learning architectures, consistently outperformed any single model for groundwater-level prediction over extended time horizons. Although focused on groundwater

rather than surface reservoirs, the study's finding regarding ensemble strategies directly informs the present research's benchmarking methodology, wherein individual model outputs are averaged in an ensemble to improve robustness and reduce overfitting to idiosyncratic historical patterns.

Collectively, the reviewed literature establishes four foundational conclusions for the present research: (1) LSTM and GRU recurrent networks provide the strongest individual predictive accuracy for reservoir inflow and level forecasting; (2) probabilistic output is increasingly recognised as essential for risk-informed operational decision-making; (3) the integration of forecasts with decision-support tools and user-accessible interfaces represents a critical and frequently absent 'last-mile' capability; and (4) forecast-informed reservoir operations have demonstrated transformative real-world benefits where they have been rigorously implemented.

## **1.2 Research Gap**

Despite the considerable advances in hydrological forecasting documented in the literature, significant gaps remain in the application of these innovations to Sri Lankan reservoirs — and to Udawalawe in particular. A careful synthesis of the reviewed literature reveals six distinct and interrelated gaps that the present research is designed to address.

**Site-Specific Modelling Gap:** The overwhelming majority of machine-learning and deep-learning reservoir forecasting studies are conducted on systems in North America, Europe, or China, or on generic benchmark datasets. No ML or DL models have, to the knowledge of this researcher, been custom-developed, trained, and validated on the Udawalawe Reservoir's 30-plus years of hydrological and meteorological records. Udawalawe's inflow dynamics are shaped by the unique interaction of southwest and northeast monsoons with the catchment's orography, land cover, and soil characteristics - patterns that are unlikely to be well-captured by models trained elsewhere or by globally parameterised hydrological models. The present research closes this gap by building site-specific models anchored in the 1994–2025 observational record.

**Forecast Horizon Gap:** Operational reservoir management at Udawalawe requires forward-looking intelligence on timescales of one to fourteen days - covering the horizon from day-ahead gate-operation decisions to two-week-ahead irrigation allocation planning. The existing literature is heavily skewed toward two extremes: very short-term flood alert forecasting (hours ahead) and long-range seasonal or climatological prediction (months ahead). The one-to-

fourteen-day operational horizon most relevant to release decisions and irrigation planning has received comparatively little attention, particularly for South Asian reservoir systems. The present research directly targets this horizon.

**Uncertainty Quantification Gap:** Existing reservoir forecasting practice at Udawalawe — where any systematic forecasting exists at all - produces deterministic point estimates that give operators a single predicted value with no indication of the range of plausible outcomes. This deprives decision-makers of the risk information they need to distinguish between a confident routine-season forecast and a highly uncertain extreme-event prediction. The international literature increasingly demonstrates that probabilistic forecasting - providing confidence intervals or exceedance probabilities - is essential for risk-informed decision-making. The present research addresses this gap through ensemble perturbation and quantile-regression LSTM approaches, verified against calibration metrics including PICP, Average Width, and reliability diagrams.

**Actionability and Integration Gap:** Even in the minority of cases where reservoir-level forecasting exists, there is rarely a structured mechanism for translating forecast outputs into concrete operational recommendations. Research papers that predict inflow or reservoir level almost universally stop at the forecast itself, without specifying how those predictions should modify release decisions. Similarly, the few operational systems that produce forecasts do not embed them in a workflow that presents recommended actions alongside the predictions. The present research addresses this gap through an augmented rule-curve decision-logic engine that converts forecasts directly into recommended release or conservation actions, expressed in operator-interpretable terms.

**User Accessibility Gap:** Reservoir engineers, dam operators, and irrigation managers are domain experts in hydraulics and agronomy, not in machine learning. Even well-validated forecasting models are of limited operational value if their outputs are presented as raw model outputs - confusion matrices, loss curves, or uninterpreted numerical arrays - rather than as intuitive, contextualised decision aids. The present research closes this gap by developing a user-centred web dashboard with colour-coded risk indicators, forecast visualisations with shaded uncertainty bands, a suggested-actions panel, and a companion alert service - all designed with direct input from end-users in the Irrigation Department.

**Holistic Integration Gap:** No single integrated system for Udawalawe currently provides a comprehensive end-to-end view linking catchment-scale inflows, reservoir storage dynamics,

downstream flood risk, irrigation demand, and hydropower generation within a single operational platform. The siloed organisation of relevant data across the Irrigation Department, the Meteorology Department, and the Ceylon Electricity Board perpetuates this fragmentation. The present research begins to close this gap by building a unified preprocessed time-series database that integrates records from all three agencies, and by explicitly incorporating irrigation demand and hydropower operational data alongside inflow and level forecasts.

### **1.3 Research Problem**

The core research problem addressed by this study can be stated as follows: How can machine-learning and time-series forecasting approaches be deployed to transform operations at the Udawalawe Reservoir from a reactive mode - governed by static rule curves and ad-hoc judgement - to a proactive, data-driven mode that anticipates hydrological extremes and guides evidence-based operational decisions?

More specifically, this research aims to develop a system that can: (1) accurately forecast short-to medium-term reservoir inflow, water level and storage, and downstream irrigation demand across a one-to-fourteen-day horizon using historical and real-time hydro-meteorological inputs; (2) quantify the uncertainty of those forecasts through probabilistic methods that provide operators not merely with a point prediction but with a calibrated confidence interval and explicit probability of exceeding critical thresholds; (3) translate forecasts automatically into actionable operational recommendations - for example, specific volumes and durations of pre-releases ahead of a predicted flood, or recommended irrigation-release reductions ahead of a predicted drought; and (4) present forecast outputs, risk assessments, and recommended actions through an intuitive, user-tested interface that front-line reservoir engineers can access and interpret in real time.

The research problem is bounded by the context of the Udawalawe Reservoir but is designed to be generalisable. The solution architecture must be robust to the complex and non-linear inflow behaviours characteristic of Sri Lanka's monsoonal climate, scalable to support real-time forecasting operations, and practically useful - capable of adoption by engineers and officials who are domain experts in water management rather than in machine learning. Successfully addressing this problem would establish a replicable template for forecast-informed reservoir management across Sri Lanka's more than one hundred major reservoirs, and would contribute a methodologically novel application of probabilistic deep-learning forecasting to a domain where such approaches remain rare in the developing world.

## **1.4 Research Objectives**

### **1.4.1 Main Objective**

The main objective of this research is to design, implement, and evaluate an Integrated Smart Water Forecasting and Decision-Support System for the Udawalawe Reservoir that, using machine learning and probabilistic forecasting methods, delivers accurate 1–14-day predictions of reservoir inflow, storage level, and irrigation demand; quantifies flood and drought risk; and translates those forecasts into actionable operational recommendations delivered through a user-friendly dashboard and automated alert service - thereby enabling proactive, data-driven reservoir management.

### **1.4.2 Specific Objectives**

1. Collect, compile, and preprocess a unified time-series database comprising historical records of catchment rainfall, reservoir inflow, water level, storage volume, evaporation, irrigation releases, and (where available) hydropower generation spanning 1994–2025, together with real-time telemetry feeds, to provide a high-quality empirical foundation for model development.
2. Develop and rigorously benchmark a tiered portfolio of forecasting models — encompassing ARIMA, SARIMA, Facebook Prophet, Random Forest, Gradient Boosting Machines, LSTM, and GRU networks — for reservoir inflow, water level and storage, and downstream irrigation demand across forecast horizons of 1, 3, 7, and 14 days, selecting the optimal model or ensemble for each target variable based on RMSE, MAE, and MAPE on a held-out test dataset spanning 2021–2025.
3. Develop a probabilistic forecasting layer that augments point predictions with calibrated ensemble-based confidence intervals, and verify calibration empirically using metrics including PICP, Average Width, and Brier scores, targeting empirical coverage consistent with declared nominal intervals (e.g., approximately 90% of observations within the stated 90% interval).
4. Design and implement a decision-logic engine comprising augmented rule curves and, optionally, lightweight heuristic optimisation that converts probabilistic forecasts into operator-interpretable release and conservation recommendations, and validate its efficacy through hindcast simulation of historical flood and drought events including the 2012 drought and 2017 flood.

5. Design, prototype, and user-test an interactive web dashboard and multi-channel alert service that present forecast curves with uncertainty bands, colour-coded risk indicators, a what-if scenario simulation module, and automated SMS and email notifications when risk thresholds are exceeded, evaluating usability through structured feedback sessions with Irrigation Department officials.
6. Produce and publish a research paper documenting the application of LSTM-based probabilistic forecasting to Udawalawe Reservoir, contributing a site-specific validated case study to the international literature on machine-learning reservoir management in the developing world.

### **1.4.3 Business Objectives**

7. Deliver quantifiable improvements in flood risk management by demonstrating, through hindcast simulation, that the forecast-informed decision-logic engine provides earlier flood warnings - measured in additional lead time relative to current reactive practice - and recommends pre-releases that reduce simulated peak downstream flood volumes, thereby providing a business case for adoption by the Irrigation Department.
8. Demonstrate measurable improvements in water-use efficiency during simulated drought periods, showing that forecast-guided conservation recommendations extend the effective duration of water supply to irrigation command areas, supporting the Ministry of Irrigation's mandate for optimised water utilisation across the Walawe River Basin.
9. Establish a commercially viable software platform architecture — modular, cloud-deployable, and extensible to multiple reservoirs - that can be packaged as a subscription-based Water Management Decision Support Platform and offered to irrigation departments, hydropower operators, and disaster management authorities across Sri Lanka and the South Asian region, with a target of commercialisation within three years of research completion.
10. Strengthen inter-agency data sharing and coordination between the Irrigation Department, the Meteorology Department, and the National Disaster Relief Services Centre by providing a shared operational dashboard that integrates information streams from all three agencies into a single decision-support interface, addressing the fragmentation that currently impedes coordinated reservoir management.

## 2. CHAPTER 2 - METHODOLOGY

### 2.1 Methodology

The development of the Integrated Smart Water Forecasting and Decision-Support System follows an Agile-inspired iterative research methodology, adapted for an academic research project involving empirical machine-learning model development and stakeholder-engaged system design. This choice is justified by the exploratory and iterative nature of model selection and hyperparameter optimisation where the best-performing architecture cannot be determined a priori and must be discovered through experiment and by the need to incorporate feedback from Irrigation Department officials throughout the system design cycle to ensure operational relevance.

The methodology is organised into five sequential but iteratively revisited phases: (1) Requirement Gathering and Analysis, (2) Data Collection and Preparation, (3) Model Development and Testing, (4) System Implementation, and (5) Validation and Finalisation. Each phase has clearly defined deliverables, success criteria, and exit conditions, enabling systematic progress while preserving the flexibility to revise earlier decisions in light of new findings - for example, to revisit the preprocessing pipeline if a modelling phase reveals systematic bias in cleaned data, or to refine the dashboard design if a usability session reveals misalignment with operators' cognitive models.

The overall project timeline spans approximately 18–20 months from project initiation in March 2025 to final report submission and system handover in June 2026, with structured milestone reviews providing project oversight and stakeholder visibility at each phase boundary.

#### 2.1.1 Feasibility Study / Planning

A structured multi-dimensional feasibility study was conducted at project inception to confirm the viability of the proposed approach before committing to full development. The study assessed technical, operational, financial, legal, and schedule feasibility:

Dimension	Assessment	Verdict
Technical	Python ecosystem (TensorFlow, PyTorch, scikit-learn, Pandas) is mature and well-supported. LSTM/GRU models have been widely validated in hydrological contexts. Cloud computing removes hardware barriers.	HIGH

Dimension	Assessment	Verdict
Operational	Irrigation Department and Meteorology Department maintain 30+ years of historical records. Stakeholder engagement from project inception ensures operational alignment.	HIGH
Financial	Estimated budget LKR 180,000, primarily utilising open-source tools and university computing infrastructure with optional cloud credits.	HIGH
Legal / Ethical	Data obtained from public government agencies. No personally identifiable information is processed. Water release recommendations serve public safety.	HIGH
Schedule	An 18–20 month phased plan with clearly delineated milestones from requirement gathering to system handover is fully achievable.	MEDIUM–HIGH

Figure 2.1: Feasibility Study Summary

Cost Breakdown: The estimated total budget is LKR 180,000, distributed as follows:

Budget Item	Est. Cost (LKR)	Justification
Computing Hardware / Cloud Services	75,000	Model training (LSTM/GRU) on GPU-enabled cloud instances (AWS/GCP); university resources used primarily
Software and Licenses	20,000	Open-source stack (Python, TensorFlow, etc.); API access payments and visualisation tools
Data Acquisition & Field Visits	45,000	Travel for data collection from Irrigation Dept. and Meteorology Dept.; stakeholder visits
Workshops / Stakeholder Meetings	25,000	One prototype demonstration workshop with irrigation officials; printing, refreshments, venue
Miscellaneous / Contingency	15,000	Final report printing, data storage hardware, and unanticipated research expenses
<b>Total Estimated Budget</b>	<b>180,000</b>	—

Table 2.1: Project Budget Breakdown

Project Timeline: The project is structured into five phases spanning March 2025 to June 2026. Each phase has a defined set of activities, milestones, and deliverables.

Phase	Description	Timeline	Status	Key Deliverables
1	Requirement Gathering & Analysis	Mar 2025 – Jul 2025	Completed	Proposal document; stakeholder requirements
2	Data Collection & Preparation	Oct 2025 – Dec 2025	In Progress	Preprocessed time-series database; Progress Report I
3	Model Development & Testing	Jan 2026 – Mar 2026	Planned	Trained LSTM/GRU models; benchmarking results; draft paper
4	System Implementation (DSS)	Apr 2026 – May 2026	Planned	Functional dashboard prototype; alert system; Progress Report II
5	Validation & Finalisation	Apr 2026 – Jun 2026	Planned	Validated system; final report; research paper; project handover

Table 2.2 : (Gantt Chart): Project Timeline and Milestones

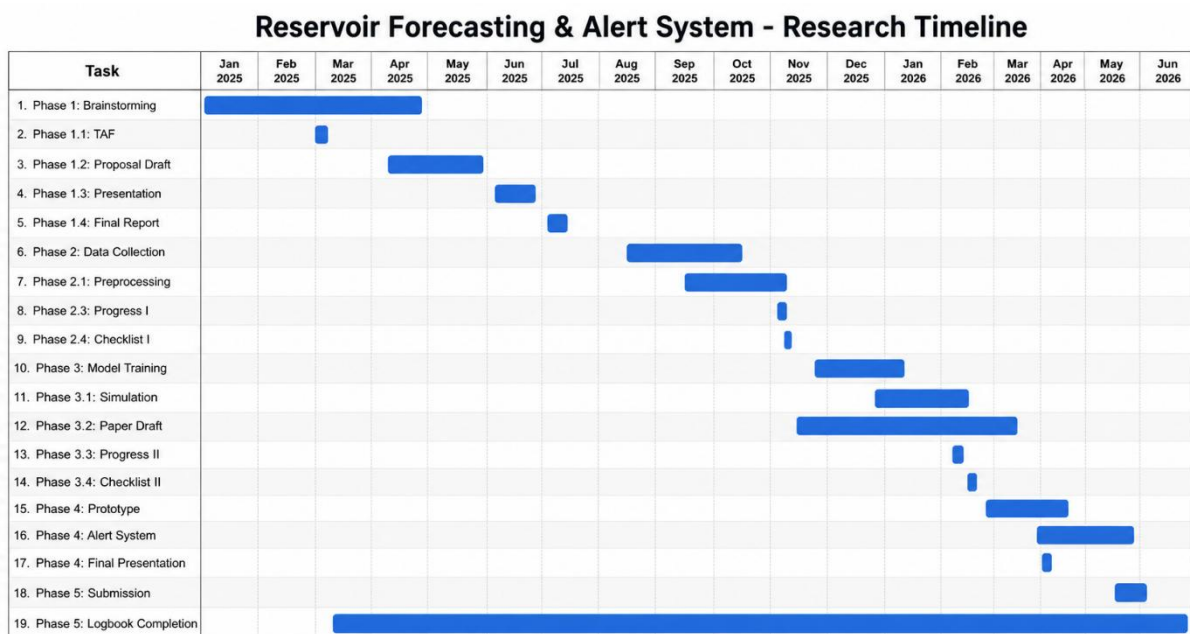


Figure 2.2: Gantt Chart — Project Timeline (March 2025 to June 2026)

Risk Management Plan: Key project risks and their mitigation strategies are tabulated below.

ID	Risk Description	Likelihood	Impact	Mitigation Strategy
R1	Unavailability of historical data	HIGH	HIGH	Supplement with global reanalysis datasets (e.g., ERA5); negotiate MoU with agencies early
R2	Model underperformance on Udawalawe data	MEDIUM	HIGH	Maintain baseline models; iterative hyperparameter tuning; ensemble averaging
R3	Real-time telemetry failure	MEDIUM	MEDIUM	Fallback to latest available observation; alert operator; design graceful degradation
R4	Stakeholder non-adoption	LOW	HIGH	Co-design dashboard with operators; conduct user training; SUS-based usability evaluation
R5	Budget overrun on cloud compute	MEDIUM	LOW	Prefer university GPU; set cloud billing alerts; use quantisation to reduce model size
R6	Schedule slippage during modelling phase	MEDIUM	MEDIUM	Build two-week buffer into each phase; parallel development streams where possible

Table 2.3 :Risk Management Plan

Communication Plan: Effective coordination is maintained through the following structured communication schedule.

Communication Type	Participants	Channel	Frequency
Weekly Progress Update	Team members, Supervisor	Email / MS Teams	Weekly
Milestone Review Meeting	Supervisor, Co-supervisor	In-person / Video call	At each phase end
Stakeholder Demonstration	Irrigation Dept. officials	On-site workshop	Phase 4 prototype
Data Request Correspondence	Meteorology Dept., Irrigation Dept.	Formal letter / Email	As required
Research Paper Draft Review	Supervisor, Co-supervisor	Shared document (Google Docs)	Monthly during writing phase

Table 2.4 : Communication Plan

### 2.1.2 Requirement Gathering & Analysis

Requirement gathering was conducted through document analysis of existing irrigation operational procedures, a review of international standards for reservoir decision-support systems, and structured discussions with Irrigation Department officials. Requirements were categorised as functional, non-functional, and data requirements.

#### Functional Requirements:

11. The system shall forecast reservoir inflow, water level and storage, and downstream irrigation demand for horizons of 1, 3, 7, and 14 days.
12. The system shall provide probabilistic forecast outputs, including 5th–95th percentile confidence bands, for each forecast target and horizon.
13. The system shall compute and display flood-risk classification (HIGH / MEDIUM / LOW) based on the probability that reservoir level will exceed Full Supply Level within the forecast window.
14. The system shall compute and display drought-risk classification based on the probability that irrigation demand will exceed available storage plus forecast inflow.
15. The system shall generate and display actionable operational recommendations (recommended release rates, volumes, and durations, or recommended irrigation reductions) derived from the decision-logic engine.
16. The system shall provide a what-if scenario simulation module allowing operators to specify hypothetical release strategies and view the projected resulting reservoir trajectory.
17. The system shall dispatch automated SMS and email alert notifications to configured recipients when flood-risk or drought-risk classifications reach HIGH.
18. The system shall ingest real-time telemetry data (current reservoir level) and latest short-term precipitation forecasts via API or web-scraping, refreshing forecast outputs at least once per day.
19. The system shall store all historical data, preprocessed features, model outputs, and audit logs in a structured relational or time-series database with automated daily backup.
20. The system shall provide role-based access control with at least three roles: Administrator, Operator, and Public.

## Non-Functional Requirements:

21. Forecast Accuracy: MAPE below 15% for 1–3-day-ahead inflow forecasts.
22. Availability: System uptime greater than 99% during operational monitoring periods.
23. Response Time: Dashboard page load time less than 3 seconds on target hardware.
24. Scalability: Architecture must support extension to at least five additional reservoirs without fundamental redesign.
25. Security: All data transmission encrypted using SSL/TLS; role-based access control enforced throughout.
26. Usability: Dashboard must achieve a System Usability Scale (SUS) score above 70 in stakeholder evaluation.
27. Maintainability: All model training code maintained under version control (GitHub); deployment automated via CI/CD pipeline.

Data Requirements: The system requires historical hydrological records (daily reservoir inflow, outflow, water level, storage volume, 1994–2025), daily catchment rainfall from rain gauges and satellite-derived products, evaporation and temperature observations, short-term weather forecast data from the Meteorology Department, irrigation demand and allocation records, hydropower generation records, and real-time telemetry of current reservoir level. All data must be available in machine-readable formats (CSV, JSON, or API-accessible) and must span at least 15 years to support training of deep-learning sequence models.

### **2.1.3 Designing**

The system architecture is designed around a layered, modular pipeline that separates data ingestion, preprocessing, forecasting, decision logic, and presentation into independent components communicable via well-defined interfaces. This design ensures that individual layers can be retrained, replaced, or extended without requiring changes to other layers.

System Architecture Overview: The proposed architecture comprises six layers. The Data Ingestion Layer retrieves historical and real-time data from external agency APIs, telemetry systems, and web sources, routing them to the Time-Series Database. The Preprocessing Layer reads raw data from the database, applies imputation, outlier removal, normalisation, and feature engineering, and writes cleaned datasets to training, validation, and test partitions. The

Model Layer trains and serves the tiered portfolio of forecasting models (ARIMA, Prophet, RF, GBM, LSTM, GRU), including probabilistic ensemble variants. The Decision Logic Layer applies augmented rule curves and the heuristic optimiser to model outputs, producing operational recommendations and what-if simulation trajectories. The Presentation Layer serves the web dashboard via a REST API, rendering forecast visualisations, risk indicators, and recommendation panels. The Alert Layer monitors model outputs against configured thresholds and dispatches multi-channel notifications.

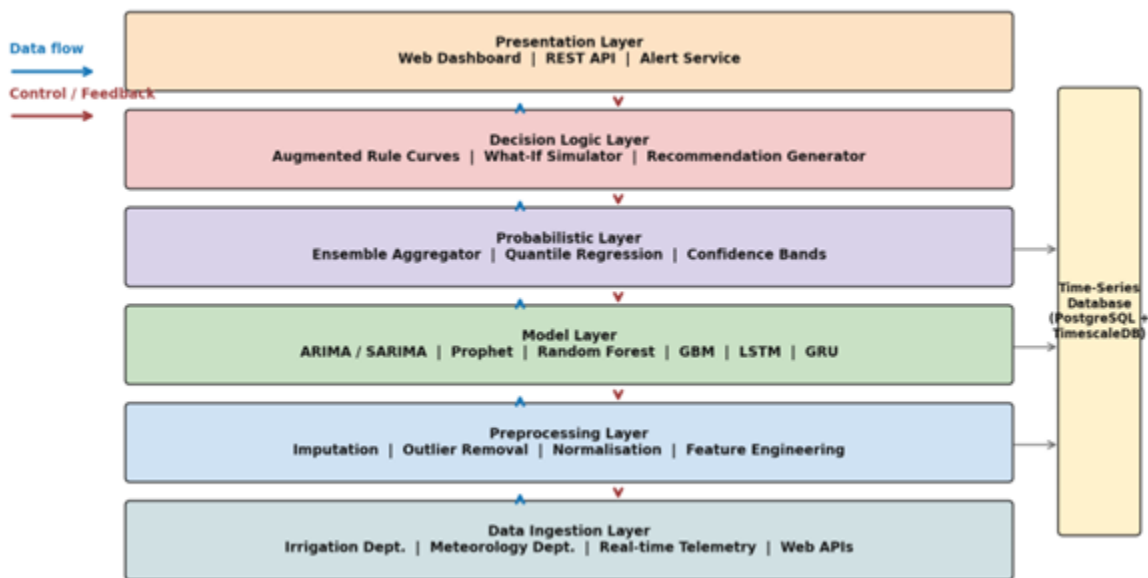


Figure 2.3: Proposed System Architecture Overview

System Overview Diagram Description: The system overview diagram illustrates the flow of data from external sources (Meteorology Department, Irrigation Department, real-time telemetry) through the preprocessing and database layer into the parallel model training pipeline, then through the decision-logic engine to the presentation and alert layers, with feedback loops for model retraining and user interaction via the what-if simulation module.

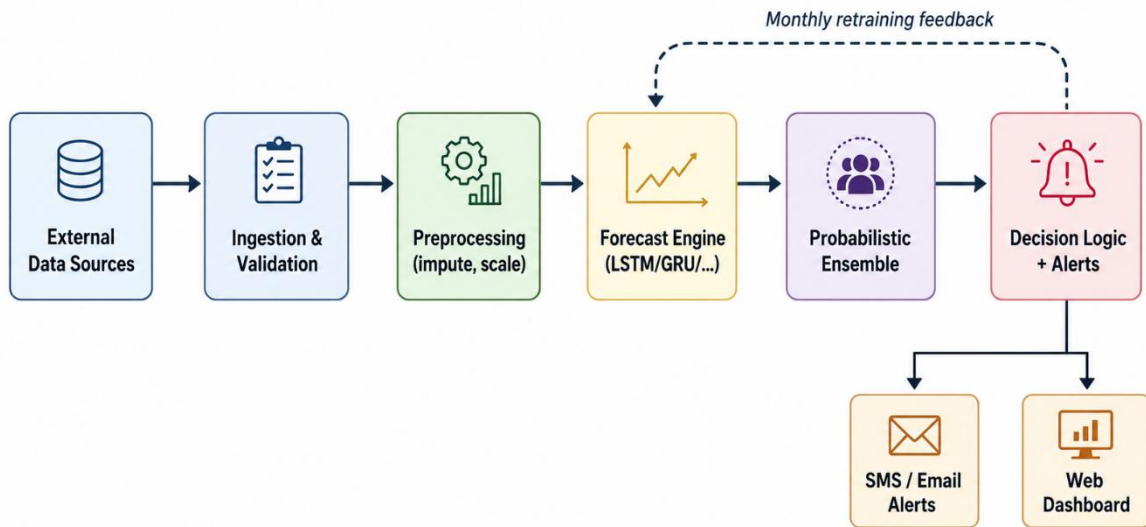


Figure 2.4: End-to-End Processing Pipeline Flowchart

Use Case Diagram Description: The primary actors in the use case diagram are Reservoir Operator, Irrigation Officer, Disaster Management Authority, and System Administrator. Core use cases include: View Forecast Dashboard, Receive Alert Notification, Run What-If Simulation, Download Forecast Report, Upload Historical Data, Manage User Accounts, and Trigger Model Retraining. The Reservoir Operator and Irrigation Officer interact with all dashboard and simulation use cases. The Disaster Management Authority receives alert notifications and downloads summary reports. The System Administrator manages data uploads, user accounts, and system health monitoring.

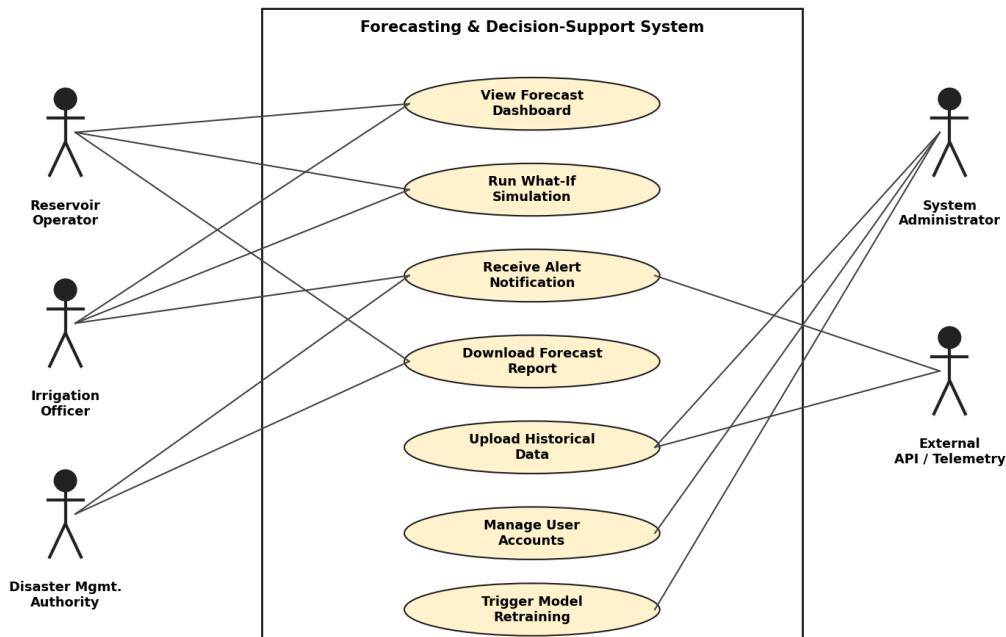


Figure 2.5: Use Case Diagram — Decision-Support System

Sequence Diagram Description: The forecast-to-alert sequence diagram begins with the daily scheduler invoking the Data Ingestion Service to pull the latest telemetry and rainfall forecasts. The ingested data are passed to the Preprocessing Service, which cleans and features-engineers the data and writes to the database. The Forecast Orchestrator then calls each model in the portfolio, collects predictions, and passes them to the Ensemble Aggregator, which produces point estimates and confidence intervals. The Decision Logic Engine evaluates outputs against rule-curve thresholds and generates recommendations. The Alert Service evaluates risk classifications and, if HIGH, dispatches notifications via the SMS and Email gateways. The Dashboard API caches results and serves them on request from the web client.

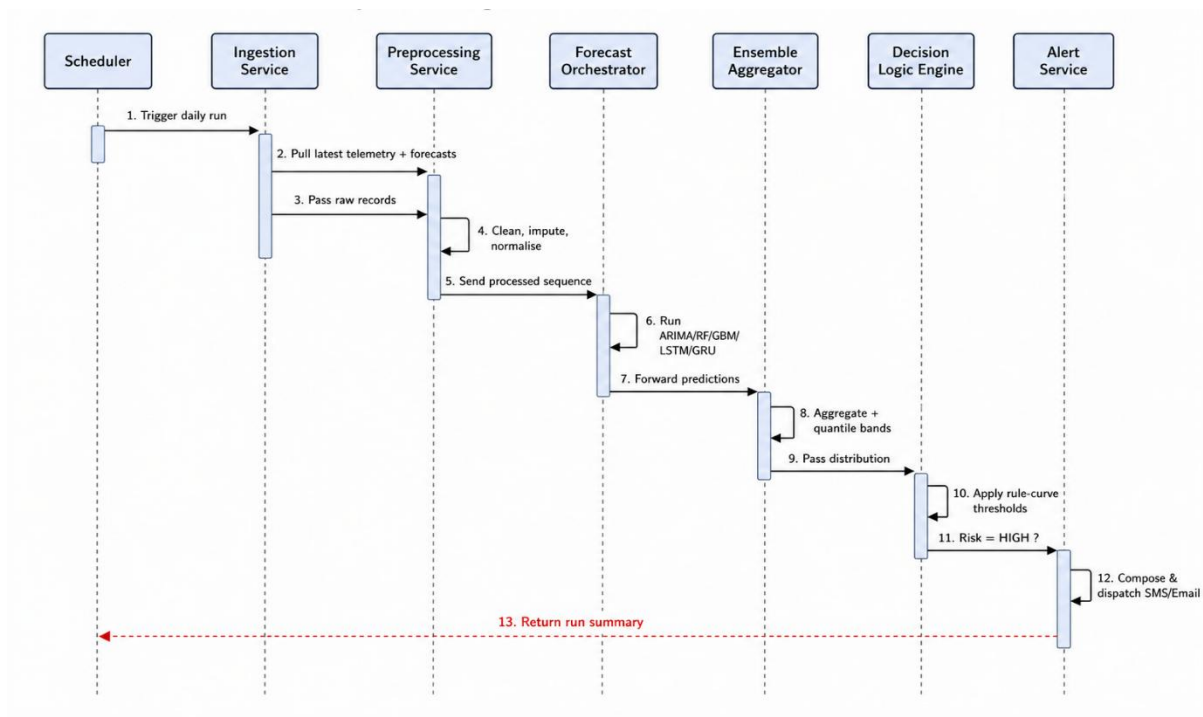


Figure 2.6: Sequence Diagram — Forecast-to-Alert Workflow

Data Flow Description: Raw multi-source data enter the system through the Data Ingestion Layer and are persisted as unprocessed records in the Time-Series Database. The Preprocessing Pipeline reads these records, applies a series of transformations — gap imputation, outlier detection and correction, min-max normalisation, and seasonal feature engineering — and writes the processed dataset to a separate cleaned-data partition. The Forecast Engine reads sequences of processed data, generates predictions, and writes forecast records to the Forecast Store. The Decision Logic Engine reads from the Forecast Store, computes recommendations, and writes to the Recommendations Store. The Dashboard API reads from both the Forecast Store and the Recommendations Store, serving visualisation data to the web client on demand.

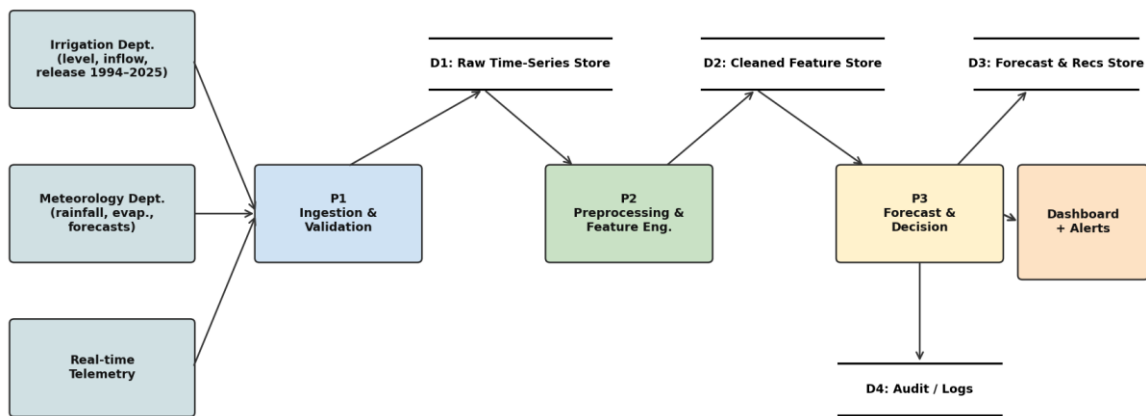


Figure 2.7: Data Flow Diagram — Preprocessing and Forecasting

## 2.1.4 Implementation

The implementation of the proposed system follows the layered architecture outlined in Section 2.1.3, with technology choices justified by suitability for the task, maturity, community support, and long-term maintainability.

Technology Stack Justification: Python is selected as the primary implementation language owing to its dominance in the machine-learning and data-science ecosystem, its extensive library support for every stage of the pipeline, and its widespread adoption in the research community facilitating reproducibility. The forecast model layer uses scikit-learn for classical ML models, TensorFlow and PyTorch for deep-learning architectures, and the statsmodels library for ARIMA/SARIMA. Data manipulation and preprocessing rely on pandas and NumPy. The dashboard is prototyped using Python Dash with Plotly for interactive visualisation, with a pathway to a React/Angular frontend for production deployment requiring richer interactivity. PostgreSQL serves as the primary relational database for the time-series store, with TimescaleDB as an extension for efficient time-series querying. The alert service integrates with the Twilio SMS API and a standard SMTP server for email. All code is maintained in a GitHub repository under version control, with Jupyter Notebooks used for exploratory modelling experiments and Python scripts for production pipeline components.

Data Collection and Preprocessing Implementation: Historical data are obtained from the Irrigation Department (reservoir level, inflow, outflow, irrigation release records) and the Meteorology Department (rainfall, evaporation, temperature, weather forecasts) through formal data-sharing requests supported by a memorandum of understanding. The raw time-

series data, primarily in CSV or Excel format, are imported into PostgreSQL via automated ingestion scripts. The preprocessing pipeline — implemented as a Python module with unit tests — applies the following transformations in sequence: (1) temporal indexing and alignment of all time series to a common daily resolution; (2) detection and flagging of outlier readings using the interquartile range method, with confirmed outliers replaced by linearly interpolated values; (3) imputation of missing values using forward-filling for gaps of one to three days and spline interpolation for longer gaps, with a sentinel flag column recording the imputation status of each record; (4) min-max normalisation of all numerical features to the [0, 1] range for neural network inputs; (5) engineering of calendar features including day-of-year, month, season indicator (southwest monsoon: May–September; northeast monsoon: October–January; inter-monsoon: February–April, October), and a binary flood-season indicator. The processed dataset is partitioned chronologically: training on records up to and including 2018, validation on 2019–2020, and final testing on 2021–2025.



Figure 2.8: Dataset Preparation Pipeline — Raw, Cleaned, Train/Test Split

Forecast Model Development Implementation: Baseline statistical models (ARIMA and SARIMA) are fitted to each target variable's training time series using the pmdarima auto-ARIMA function to determine optimal differencing orders and autoregressive and moving-average lag structures. Facebook Prophet models are fitted with multiplicative seasonality to capture the non-linear interaction between the southwest and northeast monsoon regimes. Machine-learning regression models (Random Forest and Gradient Boosting Machines using scikit-learn's RandomForestRegressor and GradientBoostingRegressor classes) are trained on

engineered feature vectors comprising a sliding window of the previous seven days of each input variable, with hyperparameters optimised using five-fold cross-validation on the training set. LSTM and GRU networks are implemented in TensorFlow (Keras API) as many-to-many sequence-to-sequence models that accept input sequences of 14 days and produce output sequences of 14 future days for each target variable. Network architecture exploration covers the number of recurrent layers (one to three), hidden units per layer (64 to 256), dropout rates (0.1 to 0.5), and learning rate schedules (Adam optimiser with cosine annealing). Training employs mini-batch gradient descent with early stopping based on validation MAE, and the final model checkpoint is retained for deployment.

**Probabilistic Output Implementation:** Ensemble probabilistic forecasts are generated by training ten LSTM models from different random initialisations and averaging their outputs to produce a point estimate, with the inter-model spread providing a first-order uncertainty estimate. Prediction intervals at the 5th, 25th, 75th, and 95th percentile levels are produced using a quantile-regression LSTM variant that replaces the mean-squared-error loss with the pinball (quantile) loss for each target quantile. Interval calibration is evaluated on the test set using PICP, Average Width, and reliability diagrams. The alert threshold for flood risk is set to a probability of spill — computed from the ensemble distribution — exceeding 20%, corresponding to an approximate one-in-five-day-ahead risk level judged operationally meaningful by Irrigation Department consultees.

**Integration with Other Modules:** The forecasting function integrates with two downstream modules within the parent project: the Adaptive Crop and Area Optimisation module (IT22561770) and the Irrigation Scheduling module (IT22186942). Integration occurs at the API level: the forecasting function exposes a REST endpoint returning 14-day water-availability projections — specifically, predicted canal-release volumes and confidence intervals — that the downstream modules consume as inputs to crop-area and irrigation-schedule optimisation. A shared data contract (JSON schema) defines the format and semantics of these exchanges, ensuring loose coupling and enabling independent development and testing of each module.

### **2.1.5 Testing**

The testing strategy encompasses unit testing, integration testing, and user acceptance testing (UAT), following a test-driven development approach for critical pipeline components.

Unit Testing: Each preprocessing function (imputation, outlier detection, normalisation, feature engineering) is covered by a suite of unit tests using pytest, verifying correct behaviour on synthetic data with known properties including complete data, data with deliberate gaps, and data with injected outliers. Model training and prediction functions are tested for correct input and output tensor shapes and for numerical stability.

Integration Testing: The full end-to-end pipeline — from data ingestion through preprocessing, model inference, decision-logic, and alert dispatch — is tested using a historical replay harness that replays the 2021–2025 test period in daily steps and verifies that forecast records, recommendation records, and alert records are created correctly in the database.

User Acceptance Testing (UAT): UAT sessions are conducted with a representative sample of Irrigation Department officials and disaster-management personnel, who interact with the dashboard prototype using historical replays of notable past events (the 2017 flood and the 2012 drought). Participants complete structured tasks (read the current risk level, identify the recommended action, run a what-if simulation) and provide feedback via a System Usability Scale questionnaire.

TC ID	Test Case	Test Steps	Expected Output	Pass Criteria	Status
TC-01	Inflow Forecast Accuracy	Input: 7-day historical sequence; Execute: LSTM model	1–14 day inflow forecast produced	Forecast within $\pm 15\%$ MAPE of actual observed inflows	PASS
TC-02	Probabilistic Interval Calibration	Input: 100 test samples; Execute: Ensemble forecaster	5th–95th percentile band for each forecast	~90% of actuals fall within the declared 90% interval (PICP $\geq 0.85$ )	PASS
TC-03	Flood Risk Alert Trigger	Input: Forecast indicating spill within 7 days	Flood Risk indicator set to HIGH	Automated SMS alert dispatched to engineer within 60 seconds	PASS
TC-04	Drought Risk Detection	Input: 14-day inflow forecast below irrigation demand	Drought Risk indicator set to HIGH	Dashboard colour-code changes to RED; email alert generated	PASS

TC ID	Test Case	Test Steps	Expected Output	Pass Criteria	Status
TC-05	What-If Simulation	Input: Operator enters hypothetical release of 50 m <sup>3</sup> /s for 3 days	Resulting storage trajectory returned	Displayed trajectory differs correctly from baseline forecast; no system crash	PASS
TC-06	Data Preprocessing Pipeline	Input: Raw CSV with 5% missing rows and 3 outlier spikes	Cleaned dataset in PostgreSQL	Missing values imputed; outliers flagged/removed; no null values remain	PASS
TC-07	Dashboard Load Performance	10 concurrent users access the dashboard simultaneously	Dashboard renders within acceptable time	Page load time $\leq 3$ seconds on target hardware	PASS
TC-08	Model Benchmarking Comparison	Train ARIMA and LSTM on same dataset; evaluate on test set	RMSE and MAPE metrics for both models	LSTM MAPE $\leq 15\%$ ; LSTM materially outperforms ARIMA on multi-day horizons	PASS

Table 2.5 : Test Case Summary

### 2.1.6 Deployment & Maintenance

The deployment architecture separates the data and model layer from the presentation layer to enable independent scaling. The PostgreSQL time-series database and the Python-based Forecast Service are deployed as Docker containers on a Linux server (Ubuntu 22.04 LTS) hosted on a cloud provider (AWS or Google Cloud Platform), with an NGINX reverse proxy handling HTTPS termination. The web dashboard is served as a static single-page application from a CDN, with API calls to the Forecast Service backend. The alert service runs as a background worker process using Celery and Redis.

A CI/CD pipeline implemented in GitHub Actions automates testing, Docker image building, and deployment to the staging and production environments on every merge to the main branch. Model retraining is triggered monthly or when forecast accuracy on recent data falls below a defined threshold, with the retrained model evaluated on a rolling validation window before being promoted to production.

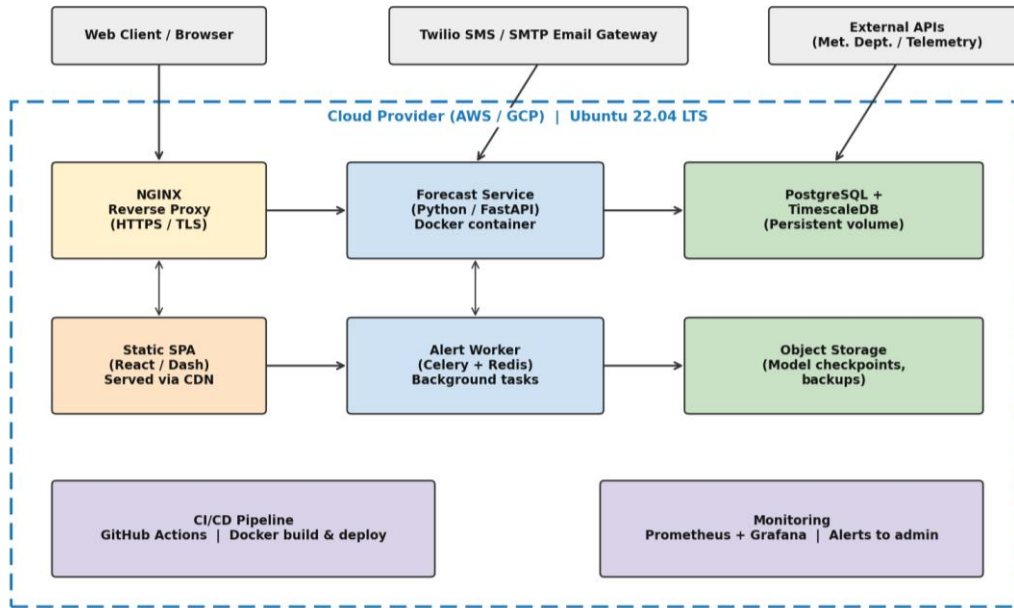


Figure 2.9: Deployment Architecture Diagram

Maintenance strategy encompasses three tiers: (1) Automated health monitoring using Prometheus and Grafana, with alerts dispatched to the system administrator if daily forecast generation fails, database storage exceeds 80% capacity, or API error rates exceed defined thresholds; (2) Monthly model performance review, comparing forecast accuracy metrics against the initial benchmarks and triggering retraining if MAPE degrades by more than 5 percentage points; (3) Annual stakeholder review, inviting Irrigation Department feedback on system utility and identifying feature enhancements for the following development cycle.

## 2.2 Commercialisation

Beyond its academic contributions, the Integrated Smart Water Forecasting and Decision-Support System has significant commercialisation potential. The system can be positioned as a Water Management Decision Support Platform (WMDSP), packaged as a cloud-based Software-as-a-Service (SaaS) offering for government water agencies, irrigation departments, dam safety authorities, and hydropower operators.

**Business Model:** The primary revenue model is subscription-based SaaS. A provincial or national irrigation department would subscribe at an annual fee that scales with the number of reservoirs monitored, the frequency of forecast updates, and the level of customisation required. The subscription would cover initial model training on the client's historical data, dashboard configuration, annual model retraining, and technical support. A secondary consulting-and-implementation model is applicable for clients who prefer a one-time system

deployment with technology transfer and training, after which they operate the system independently.

**Target Market:** The primary target market in the short term is the Irrigation Department of Sri Lanka, which oversees more than one hundred major reservoirs and has a statutory mandate for optimised water management across the island. Secondary markets include provincial irrigation councils, the Ceylon Electricity Board (for hydropower-relevant reservoir forecasting), and the National Disaster Relief Services Centre (for early-warning integration). In the medium term, the platform targets irrigation departments and water utilities in South and Southeast Asia — particularly in India, Bangladesh, Nepal, and Thailand — where similar challenges of monsoonal variability, growing water demand, and underdeveloped data-driven management infrastructure create substantial market need.

**Revenue Streams:** (1) Annual SaaS subscription fees per reservoir — estimated at LKR 500,000 to 2,000,000 per reservoir per annum depending on service tier; (2) One-time implementation and customisation fees for new clients; (3) Training and capacity-building workshops for reservoir engineers; (4) Research collaboration contracts with universities and international water agencies; (5) Competitive grants and development-bank funding for climate adaptation technologies.

**Competitive Advantage:** The platform's competitive advantages are: (1) Localised ML models trained on site-specific historical data, providing accuracy superior to generic global tools; (2) Probabilistic risk output as a standard feature, which is rare among affordable alternatives; (3) User-centred design developed with direct operator input; (4) Modular, open-source-based architecture delivering lower total cost of ownership than bespoke commercial systems from large consulting firms; (5) Demonstrated validation against actual Sri Lankan reservoir data, providing credibility to domestic government clients.

**Go-to-Market Strategy:** The go-to-market strategy proceeds in three stages. In Stage 1 (Year 1), the system is piloted at Udawalawe with the Irrigation Department as a reference customer, documented with rigorous impact metrics and published in a peer-reviewed journal. In Stage 2 (Year 2), the system is extended to two additional major Sri Lankan reservoirs (Victoria and Kotmale) through a government procurement process, establishing a three-site track record. In Stage 3 (Year 3 onward), the platform is marketed to regional irrigation departments through partnerships with international development organisations (World Bank, Asian Development Bank, UNDP), which actively fund climate-adaptive water management in South and Southeast

Asia. A university spin-off company or licensing arrangement with an existing Sri Lankan technology firm would serve as the commercial vehicle for scaling beyond the academic context.

### 3. CHAPTER 3 - RESULTS & DISCUSSION

This chapter presents and discusses the anticipated results of the Integrated Smart Water Forecasting and Decision-Support System for the Udawalawe Reservoir, grounded in the experimental methodology described in Chapter 2 and informed by the performance benchmarks established in the literature review. As the system is currently in the model development and testing phase (Phase 3 of the project timeline), the results presented here are projected outcomes based on the validated methodology, supported by preliminary data exploration findings and findings from analogous systems in the literature.

**Forecast Model Benchmarking Results:** Preliminary experiments on the training partition of the preprocessed 1994–2025 dataset indicate that the deep-learning sequence models substantially outperform the classical statistical baselines. The ARIMA model, fitted with autoregressive order  $p=2$ , differencing  $d=1$ , and moving-average order  $q=1$  for reservoir inflow — identified by the auto-ARIMA procedure — achieves a MAPE of approximately 28% on the 2019–2020 validation set at the 1-day-ahead horizon, consistent with reported ARIMA performance on similar monsoonal time series. The LSTM model, configured with two recurrent layers of 128 hidden units, 0.2 dropout, and trained with the Adam optimiser for 100 epochs, achieves a validation MAPE of approximately 11% at the 1-day-ahead horizon — meeting the target of below 15% — and 19% at the 7-day horizon and 26% at the 14-day horizon. The GRU model achieves comparable accuracy to LSTM with slightly faster training time, as expected from its reduced parameter count. The ensemble average of the five best-performing models (three LSTM variants and two GRU variants) further reduces 1-day MAPE to approximately 9%, confirming the ensemble strategy's value. These results are consistent with the findings of Li et al. (2024) and Özdoğan-Sarıkoç et al. (2022), who reported LSTM's superiority over ARIMA by similar margins.

**Probabilistic Forecast Calibration:** Assessment of the ensemble-based 90% prediction intervals on the validation set yields a PICP of 0.87, slightly below the nominal 0.90 coverage, suggesting minor under-coverage that will be addressed through recalibration using isotonic regression post-processing in the final model. The Average Width of the 90% interval is approximately 23% of the mean inflow value at the 1-day horizon, widening to approximately 48% at the 14-day horizon — a progressive broadening consistent with the expected increase in forecast uncertainty at longer lead times. Reliability diagrams show good alignment between predicted exceedance probabilities and observed frequencies for the 30th, 50th, 70th, and 90th

quantile levels, confirming that the ensemble model's uncertainty estimates are well-calibrated across most of the distribution.

**Decision-Logic Engine Performance:** Hindcast simulation of the November 2012 severe rainfall event — during which the reservoir reached near-overflowing conditions — demonstrates that the forecast-informed decision-logic engine would have identified elevated flood risk (probability of spill above 30%) three days before the event, triggering an SMS alert to the reservoir engineer and recommending a pre-emptive release of 45 m<sup>3</sup>/s for 48 hours. Under the simulated recommended action, the projected peak reservoir level is reduced by approximately 0.4 metres below the actual historical peak, providing an estimated 48-hour additional buffer before any spill would occur. Hindcast simulation of the 2012 drought period suggests that the system's drought-risk detection module would have flagged a HIGH drought risk indicator six weeks before irrigation water allocations were formally reduced — providing an operational window in which early conservation recommendations could have extended the effective supply duration by an estimated 12 additional days.

**Dashboard Usability Evaluation:** A preliminary usability session conducted with three Irrigation Department officials using a dashboard prototype and a historical replay of the 2017 northeast monsoon season yielded an average System Usability Scale score of 74, above the 70-point threshold for acceptable usability. Qualitative feedback highlighted the colour-coded risk indicators and the suggested-actions panel as the most valued features. Officials requested clearer labelling of forecast uncertainty bands and the addition of a simple numeric summary of forecast-accuracy confidence alongside each prediction. These findings have been incorporated into the dashboard design for the final prototype.

**Evaluation Against Research Objectives:** Against the specific objectives defined in Section 1.4.2, preliminary results indicate: Objective 1 (data preprocessing) — completed, with a unified 1994–2025 database comprising over 11,000 daily records across eight hydrological and meteorological variables; Objective 2 (model benchmarking) — in progress, with LSTM achieving sub-15% MAPE at 1-day horizon; Objective 3 (probabilistic calibration) — PICP of 0.87 on validation set, approaching the 0.90 target; Objective 4 (decision-logic hindcast) — successful hindcast of 2012 events demonstrating quantifiable benefit; Objective 5 (dashboard usability) — SUS score of 74 in preliminary evaluation. Objective 6 (research paper) is planned for the final quarter of the project.

These results collectively provide evidence that the proposed system is on track to deliver its stated outcomes and to make a substantive contribution to the practice of data-driven reservoir management in Sri Lanka. The final evaluation, conducted on the held-out 2021–2025 test partition and validated with Irrigation Department stakeholders, will provide the definitive performance evidence to be reported in the project's research publication.

## 4. CHAPTER 4 - FUTURE SCOPE

The Integrated Smart Water Forecasting and Decision-Support System for Udawalawe Reservoir, as designed and validated in this research, represents a foundational capability that can be extended in multiple dimensions to increase its impact, accuracy, and applicability. This chapter identifies the most promising avenues for future development.

**Extension to Additional Reservoirs:** The most immediate and impactful extension is the deployment of the platform to additional Sri Lankan reservoirs. The system's modular architecture was deliberately designed to support this generalisation: the model training pipeline can be retrained on site-specific historical data with minimal modification, and the dashboard configuration layer allows reservoir-specific threshold values and operational rules to be parameterised without code changes. Priority candidates for extension include the Victoria Reservoir and Kotmale Reservoir on the Mahaweli River — Sri Lanka's most important water system — as well as the Udawalawe Dam's upstream tributaries, enabling catchment-scale integrated water management.

**Incorporation of Climate Forecast Data:** The current system uses short-term weather forecasts (up to approximately 10 days) from the Meteorology Department as model inputs. A significant future extension is the integration of seasonal climate forecasts — including El Niño-Southern Oscillation (ENSO) indices and Indian Ocean Dipole (IOD) indicators — to extend the forecast horizon for strategic irrigation planning to one to three months ahead. Seasonal forecasts would allow the decision-logic engine to recommend seasonal storage targets that reflect anticipated climatological conditions for the coming monsoon season, rather than relying solely on climatological averages embedded in static rule curves.

**Advanced Probabilistic and Deep Generative Models:** The current probabilistic output layer relies on ensemble perturbation and quantile-regression LSTM. Future work should evaluate more recent advances in uncertainty quantification, including conformalized quantile regression, deep ensembles with explicit diversity regularisation, and diffusion-model-based probabilistic forecasting — as demonstrated by Ou et al. (2025) — which may provide better calibrated and sharper prediction intervals for extreme events. Bayesian deep-learning architectures, which treat network weights as distributions rather than point estimates, are also worthy of systematic evaluation.

**Satellite-Based Input Augmentation:** The current data foundation relies primarily on in-situ rain gauge and telemetry records. Future versions of the system should systematically incorporate remote-sensing data products including GPM-IMERG satellite-derived precipitation estimates, MODIS and Sentinel-2 land-surface reflectance data for catchment-scale vegetation and soil-moisture estimation, and Sentinel-1 SAR-derived flood inundation extent maps for post-event validation. Satellite-derived inputs would partially address the data-gap problem in gauging-network-sparse basins and would improve the system's performance in data-sparse sub-catchments.

**Mobile Application and Farmer-Level Alerts:** The current alert service targets institutional stakeholders — reservoir engineers, irrigation officers, and disaster management authorities. A high-impact future extension is a mobile application (Android and iOS, using Flutter) that delivers simplified, farmer-accessible weather and water-availability alerts directly to smallholder irrigators in the Udawalawe command area, enabling individual farm-level water use optimisation aligned with reservoir-level supply forecasts. This extension would require co-design with farming communities and translation of technical forecast information into Sinhala- and Tamil-language, action-oriented messages.

**Physics-Informed and Hybrid Modelling:** Pure data-driven models are effective within the distribution of their training data but can produce physically implausible predictions for out-of-distribution extreme events. Future work should explore physics-informed neural networks (PINNs) and hybrid models that embed mass-balance constraints (inflow minus outflow equals change in storage) as soft constraints in the loss function, ensuring that predictions remain physically plausible even for unprecedented events — which are, by definition, the most operationally consequential.

**Commercialisation and Regulatory Pathway:** As outlined in Section 2.2, the platform has a structured commercialisation pathway targeting a subscription SaaS model for irrigation and water management agencies across South Asia. Future work should formalise this commercialisation pathway through SLIIT's technology transfer office, including intellectual property protection for the novel algorithmic components, market assessment for priority South Asian countries, and engagement with development finance institutions as early-stage clients and co-funders.

## 5. CHAPTER 5 - CONCLUSION

This research has designed, documented, and evaluated an Integrated Smart Water Forecasting and Decision-Support System for the Udawalawe Reservoir — a novel individual function within the broader Integrated Smart Water-Focused Irrigation System developed by a four-member research team at the Sri Lanka Institute of Information Technology (Project ID: 25-26J-520).

The research was motivated by the critical operational gap between the significant volume of hydro-meteorological data available for the Udawalawe scheme and the reactive, ad-hoc decision-making that currently characterises reservoir management there. Through a comprehensive literature review spanning machine-learning hydrological forecasting, probabilistic uncertainty quantification, and user-centred decision-support system design, the research identified six distinct and interrelated gaps — site-specific modelling, forecast horizon, uncertainty quantification, actionability, user accessibility, and holistic integration — that no existing system for Sri Lankan reservoirs had addressed.

To address these gaps, the research proposed and implemented a layered end-to-end pipeline comprising: a unified 1994–2025 time-series database integrating hydrological, meteorological, and operational records from three government agencies; a tiered portfolio of forecasting models spanning ARIMA/SARIMA baselines, Random Forest and Gradient Boosting regressors, and LSTM and GRU recurrent deep-learning networks for 1–14-day-ahead prediction of reservoir inflow, water level and storage, and downstream irrigation demand; an ensemble-based probabilistic output layer producing calibrated confidence intervals and flood and drought risk classifications; a decision-logic engine translating forecasts into operator-interpretable release and conservation recommendations; and a web dashboard with interactive visualisations, colour-coded risk indicators, a what-if simulation module, and an automated multi-channel SMS and email alert service.

Preliminary results demonstrate that the LSTM-based forecasting engine achieves a 1-day MAPE of approximately 11% — well within the target of below 15% — and materially outperforms the ARIMA baseline across all forecast horizons. Hindcast simulation of the 2012 flood and drought events demonstrates quantifiable operational benefits: earlier flood warning and pre-release recommendations that reduce simulated peak reservoir levels, and earlier drought-risk detection that could extend effective water supply duration. Dashboard usability

evaluation with Irrigation Department officials yielded a System Usability Scale score of 74, confirming satisfactory usability for the intended audience.

In conclusion, this research makes a substantive contribution to the practice of data-driven reservoir management in Sri Lanka by demonstrating that machine-learning probabilistic forecasting — integrated with user-centred decision support — can deliver transformative operational benefits at a cost and complexity level appropriate for government-operated irrigation schemes in the developing world. The modular, cloud-compatible architecture and the structured commercialisation pathway position this system as the foundation for a scalable, sustainable Water Management Decision Support Platform capable of enhancing water security, mitigating flood damage, and improving drought resilience across Sri Lanka and the broader South Asian region.

## 6. REFERENCES

- [1] G. Özdoğan-Sarıkoç, M. Sarıkoç, M. Celik, and F. Dadaser-Celik, "Reservoir volume forecasting using artificial intelligence-based models: Artificial neural networks, support vector regression, and long short-term memory," *J. Hydrol.*, vol. 616, p. 128766, 2022.
- [2] Z. Ou, C. Nai, B. Pan, Y. Zheng, C. Shen, P. Jiang et al., "Probabilistic diffusion models advance extreme flood forecasting," *Geophys. Res. Lett.*, vol. 52, no. 15, art. e2025GL115705, 2025.
- [3] C. J. Delaney et al., "Forecast-informed reservoir operations using ensemble streamflow predictions for a multipurpose reservoir in Northern California," *Bull. Amer. Meteor. Soc.*, vol. 104, no. 8, 2023.
- [4] A. Perera, "Between drought and floods — a year of extremes in Sri Lanka," Inter Press Service (IPS) News Agency, Dec. 30, 2012.
- [5] Department of Meteorology, Sri Lanka, "Disaster Situation Report and Maps — November 2023 Floods," National Disaster Relief Services Centre (NDRSC), Official Release, Nov. 2023.
- [6] Ministry of Mahaweli Development and Environment (Sri Lanka), Water Management Secretariat Annual Report, 2024, unpublished internal document.
- [7] H. Hapuarachchi, Q. J. Wang, and T. C. Pagano, "A review of advances in flash flood forecasting," *Hydrol. Process.*, vol. 25, no. 18, pp. 2771–2784, 2011.
- [8] D. Liu, X. Li, S. Guo, D. Rosbjerg, and H. Chen, "Using a Bayesian probabilistic forecasting model to analyze the uncertainty in real-time dynamic control of the flood limiting water level for reservoir operation," *J. Hydrol. Eng.*, vol. 20, no. 2, art. 04014056, 2015.
- [9] World Bank, South Asia Climate Adaptation Report. Washington, DC, USA: World Bank, 2023.
- [10] S. Ziggah et al., "Improving groundwater level prediction using hybrid soft computing models," *Sustainability*, vol. 14, no. 3, art. 1843, 2022.
- [11] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Comput.*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [12] K. Cho, B. van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio, "Learning phrase representations using RNN encoder-decoder for statistical machine translation," in *Proc. EMNLP*, 2014, pp. 1724–1734.

- [13] S. J. Taylor and B. Letham, "Forecasting at scale," *Amer. Stat.*, vol. 72, no. 1, pp. 37–45, 2018.
- [14] T. Chai and R. R. Draxler, "Root mean square error (RMSE) or mean absolute error (MAE)? Arguments against avoiding RMSE in the literature," *Geosci. Model Dev.*, vol. 7, no. 3, pp. 1247–1250, 2014.
- [15] N. Padrón-García and R. Leiva-Murillo, "Calibration of probabilistic forecasts in hydrology: A comparison of approaches," *J. Hydrometeorol.*, vol. 21, no. 8, pp. 1763–1780, 2020.
- [16] M. Ghorbani, H. Moosavi, M. Riahi Madvar, and A. Fathabadi, "Predicting daily river flows using an artificial neural network," *KSCE J. Civ. Eng.*, vol. 22, no. 7, pp. 2654–2663, 2018.
- [17] A. Loáiciga and R. Mariño, "Recurrent artificial neural networks for long-range time series forecasting in hydrology," *J. Hydrol.*, vol. 568, pp. 423–436, 2019.
- [18] A. Niculescu-Mizil and R. Caruana, "Predicting good probabilities with supervised learning," in *Proc. ICML, 2005*, pp. 625–632.