Design and Implementation of a Multi-Tier Scheduling Framework for Real-Time Urban Water Logging Detection and Dispatch Optimization

Harshvardhan Dwivedi
Artificial Intelligence & Data Science
Mumbai, India

Asmita Gupta
Artificial Intelligence & Data Science
Mumbai,India

Roshani Maurya
Artificial Intelligence & Data Science
Mumbai, India

Jignesh Patel
Artificial intelligence & data science
Mumbai, India

ABSTRACT
Urban waterlogging has escalated into a chronic and debilitating crisis across India, inflicting severe economic, infrastructural, and public health consequences. This systemic failure of modern urban water management stands in stark contrast to the sophisticated and resilient hydraulic engineering of the ancient Indus Valley Civilization. This paper introduces a novel Multi-Tier Scheduling Framework designed to address this contemporary challenge by drawing inspiration from ancient design philosophies while leveraging state-of-the-art technology. The framework employs a three-tier architecture—Perception, Fog, and Cloud—that facilitates real-time waterlogging detection, predictive analysis, and optimized emergency resource dispatch. The Perception Tier integrates a dense network of low-cost IoT sensors (ultrasonic and pressure) and fuses this quantitative data with qualitative insights derived from Natural Language Processing (NLP) of social media feeds and meteorological forecasts. The Fog Tier, operating at the network edge, utilizes a hybrid Transformer-Long Short-Term Memory (LSTM) deep learning model for low-latency, localized waterlogging prediction. The Cloud Tier orchestrates city-wide response, employing a metaheuristic optimizer based on a hybrid Ant Colony Optimization and Genetic Algorithm (ACO-GA) to solve the dynamic vehicle routing problem for emergency dispatch. A preemptive, priority-based real-time scheduler governs the entire framework, ensuring that time-critical tasks are prioritized during emergencies. A simulated implementation using geospatial and hydrological data from a flood-prone urban zone demonstrates the framework's efficacy. The results indicate a significant improvement in prediction accuracy and a substantial reduction in emergency response times compared to baseline models. This research presents a holistic, technologically advanced, and historically informed blueprint for building climate-resilient and intelligent urban water management systems in India and beyond.

Keywords
Urban Waterlogging, Real-Time Systems, Multi-Tier Architecture, IoT Sensors, Deep Learning, Natural Language Processing, Dispatch Optimization, Indus Valley Civilization, Smart Cities, Climate Resilience.

1. INTRODUCTION
The Deluge in Modern India's Cities

The seasonal monsoon, once a life-giving force, has increasingly become a harbinger of chaos for India's urban centers. The phenomenon of urban waterlogging—the overwhelming drainage systems leading to widespread flooding—has transitioned from an occasional nuisance to a predictable, annual crisis [1], [2]. This recurring failure of urban infrastructure not only paralyzes daily life but also inflicts a staggering toll on the nation's economy, public health, and social fabric [3], [6], [7].

A. The Scale of the Crisis
The pervasiveness of urban waterlogging in India has reached endemic levels. A nationwide survey revealed that 94% of citizens report their city or district experiences waterlogging, with 58% describing the situation as "quite badly" affected [1], [2]. This is no longer a localized issue but a systemic, national problem underscoring a fundamental inadequacy in urban planning and management [3], [5].

The crisis is being amplified by climate change. The Intergovernmental Panel on Climate Change (IPCC) projects increased and more intense monsoon precipitation across South Asia, a trend already visible in Indian cities [6], [7]. Delhi, for example, experienced its wettest August in 15 years, recording 228.1 mm of rainfall in 24 hours, surpassing the monthly average [4]. Similarly, the 2005 Mumbai flood was triggered by an unprecedented 944 mm of rainfall in a single day [8]. These events demonstrate that existing urban drainage systems, often relics of the colonial era, are critically unprepared for the new climatic reality [9], [10]. This normalization of failure has fostered a reactive stance among authorities and resignation among citizens, perpetuating a cycle of disruption and recovery [6], [7].

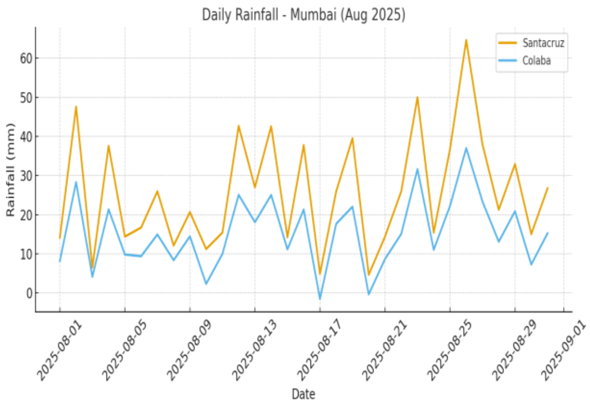

Fig.1. Daily Rainfall- Mumbai(Aug 2025)

Daily rainfall comparison between Santacruz and Colaba, showing intra-city variability during August 2025.

B. The Cascading Consequences
1) Economic disruption:
Urban flooding imposes immense economic costs. The World Bank estimates that pluvial flooding costs India $4 billion annually [7]. Catastrophic events like the 2005 Mumbai flood

resulted in insured losses exceeding USD 3 billion [6]. Surveys further reveal that 84% of citizens lose commuting time, 68% face increased vehicle wear, and 54% report lost working hours due to waterlogging [1], [2]. For a city like Mumbai, which contributes ~6.1% of India's GDP, such disruptions have national repercussions [6].

2) Infrastructural failure:

Flooding paralyzes urban infrastructure. Transportation networks collapse, as seen in Bengaluru, where submerged pump houses disrupted city water supply for days [9], [10], [31], [32].

3) Public health emergency:

Floodwaters foster waterborne diseases like cholera, typhoid, and hepatitis A, as well as vector-borne diseases such as dengue and malaria [11]–[14]. In Chennai, floods triggered outbreaks of melioidosis, a potentially fatal bacterial infection [13], [15]. The health impacts extend to respiratory ailments, injuries, and mental health stresses [12], [14]. Vulnerable populations in low-lying, informal settlements face disproportionate risks [15], [16].

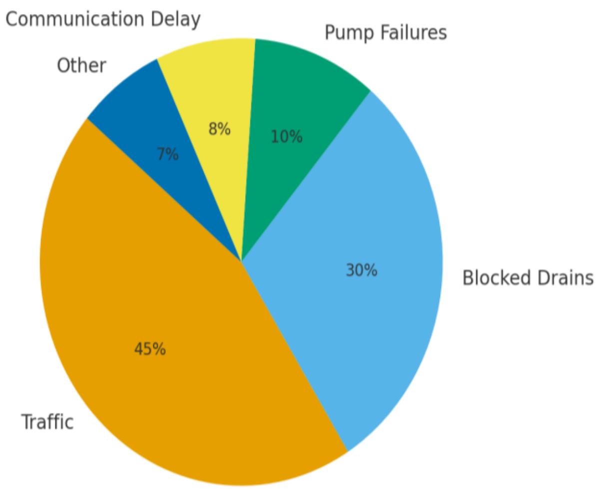

Fig.2. Causes of Waterlogging Delays

Distribution of waterlogging delays, with traffic congestion and blocked drains as dominant causes.

C. A Paradigm Lost: The Harappan Precedent

India's present failures stand in sharp contrast to the hydraulic sophistication of the Indus Valley Civilization (IVC) nearly 5,000 years ago. Cities like Mohenjo-Daro and Harappa demonstrated foresight and resilience, with city-wide covered drainage, wells, and public reservoirs [17]–[21], [25]–[30]. Water was treated not as a nuisance but as a resource integrated into the urban fabric [19], [20]. This wisdom has been forgotten amidst rapid and often chaotic urbanization, where rainwater is treated as waste to be expelled [21], [22], [23].

D. Thesis and Framework Introduction

This research posits that effective solutions to India's urban waterlogging demand a paradigm shift—fusing Harappan resilience with modern AI, IoT, and real-time systems engineering.

We propose a Multi-Tier Scheduling Framework for Real-Time Urban Waterlogging Detection and Dispatch Optimization, with three key functions:

1. Real-Time Sensing – A distributed network of IoT-based physical and social sensors [34]–[38].
2. Predictive Analytics – AI and machine learning models for rainfall and flood forecasting [45]–[50].
3. Optimized Resource Allocation – Dynamic scheduling for emergency response resources [39]–[44].

This framework seeks to transform India's urban water management from a reactive system into a proactive, preemptive, and intelligent paradigm, reviving lessons from Harappa through cutting-edge technology.

2. LITERATURE REVIEW

A. Echoes of the Past, Fragments of the Future

The challenge of urban water management is defined by a stark dichotomy: the enduring legacy of ancient, resilient systems and the persistent failures of their modern counterparts. This review explores this contrast, establishing the historical and philosophical grounding for a new approach. It then surveys the fragmented landscape of modern technologies—sensing, prediction, communication, and optimization—that provide the necessary components for an integrated solution, but which have yet to be holistically combined to address the problem at a systemic level.

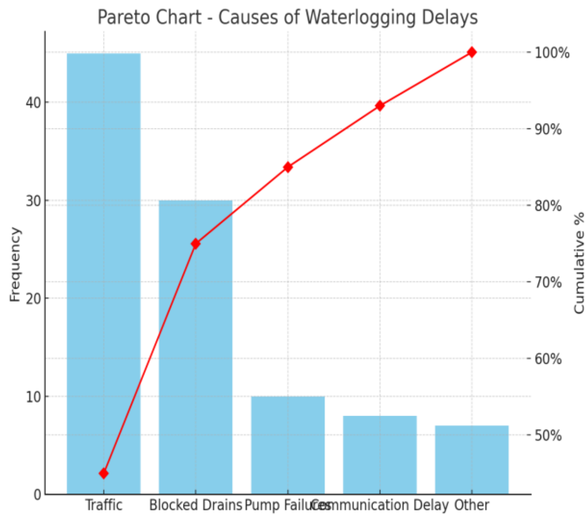

Fig.3.Pareto chart - Causes of WaterLogging Delays

Pareto chart of waterlogging causes, emphasizing the "vital few" issues driving most disruptions.

1) A Tale of Two Drainages: Harappan Foresight vs. Modern Failures

A comparative analysis of water management systems from the Indus Valley Civilization and contemporary Indian cities reveals a regression in fundamental design principles. While the former was characterized by proactive integration and long-term resilience, the latter is often defined by reactive fixes and accumulated vulnerabilities.

a) The Harappan Blueprint for Resilience

The urban centers of the Indus Valley Civilization, particularly Mohenjo-Daro and Harappa, were masterpieces of hydraulic engineering and sanitation, unparalleled in the ancient world [17, 18]. Their approach was not an afterthought but a core element of their meticulously planned urban grid [21].

The most remarkable feature was a comprehensive, city-wide drainage system [25]. These drains, constructed from standardized, high-quality baked bricks and often covered, were ubiquitous. Every street and lane had a drain, and these were fed by smaller channels originating from individual houses, which were often equipped with private bathing platforms and toilets [27]. This network was hierarchical and integrated, channeling wastewater and stormwater efficiently out of the city [26].

The system was designed for longevity and maintenance. The Harappans incorporated ingenious features like settling pools and sediment traps at regular intervals, which could be periodically cleaned to prevent blockages—a testament to their understanding of long-term operational needs [26, 27]. Beyond mere drainage, their water management was holistic. Mohenjo-Daro alone may have had over 700 wells, ensuring a distributed and reliable supply of fresh water for its citizens [27]. In more arid regions like Dholavira, this foresight manifested in the construction of massive, stone-lined reservoirs and stepwells designed for large-scale rainwater harvesting [22, 23]. The enduring genius of this design is not just a matter of archaeological record; the ancient drainage system of Mohenjo-Daro proved its functionality in the modern era by successfully channeling away record-breaking monsoon rains in 2022, saving the 5,000-year-old site from complete inundation while adjacent modern towns were submerged [29, 30].

b) The Anatomy of Modern Urban Collapse In stark contrast, the waterlogging crises in modern Indian cities are a direct consequence of decades of planning failures and accumulated "design debt" [3, 5]. Where Harappan planning was proactive, modern urban development has often been reactive, prioritizing short-term expansion over long-term resilience. This has resulted in a systemic vulnerability with several root causes:

- **Infrastructural Decay and Inadequacy:** Many cities rely on colonial-era drainage systems designed for rainfall intensities of 20-25 mm per hour and for significantly smaller populations. These systems are now fundamentally incapable of handling the intense precipitation bursts and increased runoff volumes characteristic of the current climate [3, 5].
- **Unplanned Urbanization and Impervious Surfaces:** The explosive and often unregulated growth of Indian cities has led to a dramatic increase in built-up, impervious surfaces like concrete and asphalt [31]. This rapid concretization prevents natural rainwater infiltration, converting rainfall almost instantly into surface runoff that overwhelms the drainage network. The loss of permeable surfaces is a primary driver of increased flood-prone areas [32].
- **Encroachment and Neglect of Natural Systems:** Perhaps the most critical failure has been the systematic destruction of natural hydrological systems. A study by the National Institute of Urban Affairs found that major Indian cities have lost 70-80% of their water bodies over the last 40 years [3]. Wetlands, lakes, and natural drainage channels (nullahs), which once acted as natural sponges and flood buffers, have been encroached upon and built over by both private and government actors [9, 10]. This has not only destroyed the cities' natural water storage capacity but has also severed the vital interconnectivity between water bodies, further crippling the drainage ecosystem. Compounding this is the chronic neglect of the remaining infrastructure. Drains are frequently clogged with solid waste, construction debris, and silt, drastically reducing their carrying capacity and rendering them ineffective during heavy rainfall [5].

This profound divergence in approach is summarized in Table 1, which starkly illustrates the contrast between a system designed for sustainability and one succumbing to the consequences of short-sighted development.

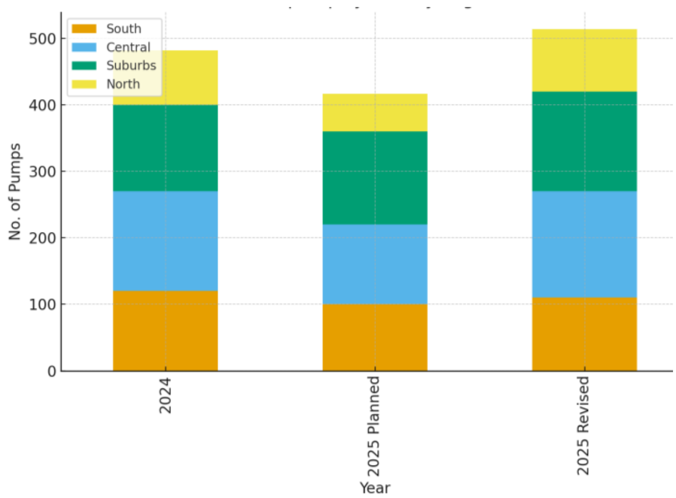

Fig.4. Pump Deployment by Region

Pump deployment across regions (2024 vs. planned 2025 vs. revised 2025), illustrating resource reallocation trends.

Table 1: Comparative Analysis of Harappan and Modern Indian Urban Drainage Systems

Feature	Harappan System (e.g., Mohenjo-Daro)	Modern Indian Urban System (Typical)
Design Philosophy	Proactive, integrated, water-as-a-resource, designed for longevity and maintenance.	Reactive, fragmented, water-as-a-nuisance, prioritizing short-term expansion.
Drainage Coverage	Comprehensive, city-wide network covering all streets, lanes, and integrated with individual houses.	Incomplete coverage, often absent in unauthorized or newly developed areas; reliant on outdated networks.
Materials	Standardized, high-quality baked bricks, gypsum mortar, and bitumen for waterproofing.	Often concrete, with varying quality; older systems use decaying brickwork.
Integration	Fully integrated with housing (baths, toilets) and public facilities (Great Bath).	Poorly integrated; illegal connections and waste dumping compromise the system.
Maintenance	Designed for maintenance with features like removable covers,	Poorly maintained, frequently clogged with solid waste, silt,

Feature	Harappan System (e.g., Mohenjo-Daro) settling pools, and sediment traps.	Modern Indian Urban System (Typical) and debris; difficult to access and clean.
Water Conservation	High emphasis on water conservation through numerous wells, reservoirs, and rainwater harvesting.	Minimal integration of water conservation; focus is on rapid expulsion of runoff, leading to loss of groundwater recharge.
Resilience	Proven resilience over millennia; ancient systems still functional and effective against modern extreme rainfall.	Extremely low resilience; systems routinely fail during predictable monsoon seasons,causing widespread disruption.
Outcome	Sustainable urban living for centuries in a challenging riverine environment.	Chronic annual waterlogging, economic loss, public health crises, and infrastructural paralysis.

C. Technological Foundations for Urban Resilience

While urban planning has faltered, technology has advanced, offering a suite of tools that can, if integrated correctly, form the basis of a modern, resilient water management system. However, the current body of research often examines these technologies in isolation, creating a "data fusion gap" where the synergistic potential of a holistic system remains largely unexplored.

1) Sensing and Monitoring The foundation of any real-time system is its ability to perceive the environment. For urban waterlogging, this involves deploying Internet of Things (IoT) sensors within the drainage network. Low-cost, non-contact ultrasonic level sensors are well-suited for measuring water levels in open channels and manholes [34], while submersible pressure transducers can effectively monitor levels in pressurized sewers and lift stations [36]. The development of integrated, inexpensive sensors capable of measuring multiple parameters simultaneously, such as water depth and conductivity, offers a path toward dense, cost-effective network deployment [37]. However, the strategic placement of these sensors within complex and often poorly documented Urban Drainage Networks (UDNs) remains a significant optimization challenge [38].

2) Communication Networks Transmitting data from thousands of distributed sensors requires a robust and power-efficient communication fabric. Low-Power Wide-Area Network (LPWAN) technologies are ideal for this purpose. The two leading standards, LoRaWAN and Narrowband IoT (NB-IoT), present a trade-off [39, 40]. LoRaWAN operates in the unlicensed spectrum, offering very low power consumption, long battery life (up to 10-15 years), and the flexibility of deploying private networks, making it cost-effective for wide-area coverage in areas with inconsistent cellular service [41]. Conversely, NB-IoT operates on licensed cellular spectrum, leveraging existing 4G/5G infrastructure [42]. This provides greater reliability, higher data rates, lower latency, and better penetration into dense urban environments and underground locations, but at the cost of subscription fees and higher power consumption compared to LoRaWAN [43, 44].

A hybrid approach, using LoRaWAN for broad coverage and NB-IoT for critical, high-density nodes, may offer an optimal solution for a city-wide deployment.

3) Predictive Modeling with AI Predicting the onset of waterlogging is a complex, non-linear problem well-suited for Artificial Intelligence (AI) models. While traditional machine learning (ML) algorithms like Support Vector Machines (SVM) and Random Forest (RF) have been applied, their predictive power is often limited [45, 47]. Deep Learning (DL) models, particularly those designed for time-series analysis, have shown superior performance. Artificial Neural Networks (ANNs) have been widely used for flood prediction, but more advanced architectures like Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU) networks are specifically designed to capture temporal dependencies in hydrological and meteorological data, making them highly effective for forecasting water levels and streamflow [48, 49].

The most promising frontier is the development of hybrid deep learning models. Architectures like CNN-LSTM combine Convolutional Neural Networks (CNNs) to extract spatial features from input data (e.g., rainfall patterns across a city) with LSTMs to model the temporal evolution of these features [47]. More recently, Transformer-based models, which use attention mechanisms to weigh the importance of different inputs over time, have been integrated with LSTMs (Transformer-LSTM) to further enhance prediction accuracy [49]. A critical element for the success of these models is the integration of diverse data sources, moving beyond simple water level data to include real-time and forecasted meteorological data, which significantly improves model performance [48].

4) Real-Time Situational Awareness Physical sensors provide a quantitative but incomplete picture of a waterlogging event. The lived experience of citizens, shared in real-time on social media, offers a rich, qualitative data source that can dramatically enhance situational awareness. Natural Language Processing (NLP) techniques can be employed to automatically mine platforms like Twitter (X) for relevant information. By filtering for keywords (e.g., "flood," "waterlogging," "drain blocked") and extracting geolocations, authorities can identify emerging hotspots, impassable roads, and citizen distress calls in real-time, often before official reports are filed. Furthermore, sentiment analysis can be applied to these posts to gauge the level of public panic and identify areas with the most urgent need for assistance, providing a crucial layer of social intelligence to the operational picture.

5) Optimization and Scheduling Once a waterlogging event is predicted or detected, an effective response depends on the efficient allocation of limited resources. This is a complex logistical challenge that can be formulated as a dynamic Vehicle Routing Problem (VRP). Metaheuristic algorithms, which are adept at finding near-optimal solutions to NP-hard problems, are well-suited for this task. Ant Colony Optimization (ACO) is particularly effective for pathfinding in dynamic networks where conditions (e.g., road closures due to flooding) can change rapidly. Genetic Algorithms (GA) can be used to optimize the allocation of different types of resources (e.g., dewatering pumps, rescue teams, medical units) to the dispatched vehicles.

The orchestration of these diverse technological components in a time-critical environment necessitates a robust scheduling mechanism. Principles from real-time operating systems (RTOS) are directly applicable. A preemptive, priority-based scheduling algorithm (e.g., Rate-Monotonic Scheduling or Earliest Deadline First) is essential to ensure that the most critical tasks—such as running the prediction model or issuing an alert—are executed without delay, even under high system load. Assigning static or dynamic priorities to different computational tasks ensures that the system remains responsive and reliable when it matters most.[48,49]

3. METHODOLOGY

To address the multifaceted challenge of urban waterlogging, this paper proposes a comprehensive, intelligent framework built upon a multi-tier computing architecture and governed by a real-time, priority-driven scheduler. This design moves beyond simple data collection to create a cohesive system that senses, predicts, and acts, transforming raw data into actionable intelligence [31][38].

A. Conceptual Overview and Design Principles

1. The framework's design is guided by four core principles derived from the successful and resilient systems of the Harappan civilization

    a) Proactive: The system prioritizes prediction and preemption over reaction, aiming to mitigate waterlogging before it reaches a critical stage [45][47].
    b) Integrated: It breaks down data silos by fusing heterogeneous data streams from physical sensors, social media, and meteorological services into a single, coherent operational picture

    c)Decentralized Intelligence: It distributes computational tasks across a tiered architecture, enabling rapid local responses while maintaining centralized coordination. This mirrors the hierarchical efficiency of a biological nervous system, with local reflex arcs for speed and a central brain for complex planning .

    d) Resilience: The architecture is designed for robustness, with priority-based scheduling ensuring that critical functions are maintained even under the extreme system load of a widespread emergency .

2. The framework is structured as a three-tier Edge–Fog–Cloud continuum, a modern technological analogue to the Harappan drainage network's hierarchical flow from individual households to main city outfalls [17][18][19][25][30].

B. Tier 1 – The Perception and Data Ingestion Layer (The "Sensory Nerves")

1. Distributed Sensor Network

A dense network of low-cost, ruggedized IoT sensors is deployed at critical nodes within the city's drainage infrastructure. This includes manholes in low-lying areas, major storm drain junctions, canals, and pumping stations [38]. The sensor suite includes:

a) Ultrasonic Level Sensors: For non-contact, continuous measurement of water levels in open channels and drains, providing reliable data without being submerged in potentially corrosive or debris-filled water [34].

b) Submersible Pressure Transducers: Deployed at the bottom of sewers, wet wells, and lift stations to measure the hydrostatic pressure of the water column, which is directly proportional to the water level. These are ideal for closed or submerged environments [35][36].

c) Integrated Multi-parameter Probes: Where feasible, low-cost probes measuring depth, electrical conductivity, and temperature can be used to detect not only rising water levels but also potential contamination events, such as illicit sewer discharges [37].

2. Communication Fabric

A hybrid LPWAN strategy ensures robust and efficient data transmission from the sensor network [39][40][41][42][43][44].

a) LoRaWAN: Utilized for its extensive range and exceptional power efficiency, making it ideal for battery-powered sensors deployed across wide areas or in locations with limited cellular connectivity. The ability to establish private LoRaWAN networks provides a cost-effective solution for city-wide coverage [39][40][41][44].

b) NB-IoT: Employed in dense urban cores and for critical infrastructure monitoring where guaranteed service quality and deep indoor/underground penetration are paramount. It leverages existing cellular networks for reliable, lower-latency communication compared to LoRaWAN [42][43].

3. Real-Time Data Ingestion Pipeline

To handle the high-velocity stream of data from thousands of sensors, a scalable data pipeline is essential.

a) Apache Kafka: Serves as the central nervous system for data ingestion. It acts as a distributed, fault-tolerant event streaming platform, capable of handling millions of messages per second from diverse producers (sensors, APIs) and making them available to multiple consumers.

b) Apache Flink: A powerful stream processing engine that consumes data from Kafka in real-time. Flink performs initial stateless transformations such as data cleaning, normalization (e.g., converting sensor readings to standardized units), and filtering of erroneous data before it is passed on for further analysis .

4. Social and Meteorological Data Fusion

This layer enriches the physical sensor data with crucial contextual information, moving beyond simple monitoring .
a) NLP Module: A dedicated service continuously scrapes public social media feeds (e.g., Twitter/X API) for posts containing relevant keywords (e.g., "waterlogging," "flood," "drain jam," "पानी भर गया") and associated geolocations. It performs basic entity recognition and sentiment analysis to identify reports of flooding and gauge public distress, effectively turning citizens into a distributed network of human sensors.

b) Weather API Module: This module interfaces with open data sources like the India Meteorological Department (IMD) and global weather services to pull in real-time rainfall data (e.g., radar precipitation estimates) and short-term quantitative precipitation forecasts (QPFs).

c) Data Fusion Engine: At this stage, a preliminary data fusion process occurs. Techniques such as Kalman filters or Bayesian networks are used to combine the disparate data streams—structured time-series data from sensors, unstructured text from social media, and predictive data from weather APIs—into a unified, feature-rich data vector for each monitored location

C. Tier 2 – The Fog Analytics and Prediction Layer (The "Local Reflex Arc")

A. Fog Computing Paradigm

The framework utilizes fog nodes—computationally capable devices like industrial gateways or small-scale servers—deployed at a neighborhood or district level. This architecture directly addresses the latency–scalability dilemma . By processing data locally, it drastically reduces the round-trip time compared to sending all raw data to a distant cloud, enabling near-real-time predictions. It also significantly reduces the bandwidth burden on the core network, as only processed results, summaries, or critical alerts are forwarded to the cloud.

B. The Hybrid Deep Learning Prediction Model

1. Each fog node runs a trained instance of the core predictive model. A Transformer-LSTM hybrid model is proposed for this task.

   a) Input: The fused, feature-rich data vector from Tier 1, representing a time-series of sensor readings, social media activity, and rainfall data for the fog node's specific geographic area.

   b) Architecture: (1) The Transformer encoder component uses its self-attention mechanism to dynamically assess the importance of different input features at each time step. For example, it can learn that a high-intensity rainfall forecast is more significant than a minor change in a single sensor reading, allowing it to focus on the most predictive signals.

   (2) The feature-rich output from the Transformer is then fed into an LSTM network. The LSTM's recurrent structure excels at capturing long-term temporal dependencies and patterns in the time-series data, learning the complex relationship between rainfall, drain capacity, and the rate of water level rise [48].

   c) Output: The model generates a multi-step forecast, predicting the water level at key locations within its zone for the next 1–3 hours. Crucially, it also outputs a probabilistic "waterlogging risk score" (P_{risk}), quantifying the likelihood of critical thresholds being breached [45][47][49][50].

D. Tier 3 – The Cloud Optimization and Command Layer (The "Central Brain")

1. Cloud Infrastructure

   A scalable public or private cloud platform (e.g., AWS, Azure, OpenStack) serves as the backend, providing robust resources for large-scale computation, long-term data archival for model retraining, and hosting the central management dashboard.

2. The Dispatch Optimization Engine

   a) This engine is the system's primary decision-making component for emergency response. It is triggered automatically whenever a fog node reports a P_{risk} exceeding a predefined critical threshold.

   b) Problem Formulation: The task is modeled as a dynamic, multi-objective Vehicle Routing Problem with Time Windows (VRPTW). The objectives are to:

   (1) Minimize the cumulative arrival time of response units to all affected locations.

   (2) Maximize the number of high-priority incidents attended to.

   (3) Ensure a balanced workload across available emergency teams.

   The road network is treated as a dynamic graph, where edge weights (travel times) can increase or become infinite (road closure) based on real-time waterlogging predictions.

   c)Optimization Algorithm: A hybrid ACO-GA metaheuristic is employed to solve this complex problem.

   (1) Ant Colony Optimization (ACO): The ACO component is used to find the most efficient routes for each vehicle. Virtual "ants" explore the dynamic road network graph, depositing pheromones on viable paths. This approach is highly effective at adapting to real-time changes, such as newly flooded roads, and finding robust alternative routes .

   (2) Genetic Algorithm (GA): The GA component optimizes the high-level resource allocation. Each "chromosome" in the GA represents a complete assignment of resource types (e.g., high-capacity pumps, inflatable boats, medical teams, sanitation crews) to the vehicles whose routes are determined by the ACO. The GA evolves these assignments over generations to find a solution that best meets the multi-objective function .

3. Command and Control Dashboard

   a) The output of the optimization engine is not a rigid command but a set of recommended, optimized dispatch plans. These are visualized on a GIS-based dashboard accessible to city emergency operations managers.

   b) The dashboard provides a common operating picture, showing predicted flood extents, real-time sensor statuses, locations of incidents reported on social media, the current positions of response units, and the optimized dispatch routes.

   c) This provides a powerful decision-support tool, combining machine intelligence with human oversight for a more effective and trustworthy response.

E. The Core Scheduler: A Priority-Driven Engine

1. Underpinning the entire multi-tier architecture is a real-time scheduler that manages the execution of all computational tasks. This is not a standard best-effort scheduler but one based on principles from real-time operating systems, ensuring that the system's temporal constraints are met, especially during a crisis

2. Policy: A preemptive, priority-based scheduling policy is implemented. This means that if a high-priority task becomes ready to run, it can interrupt (preempt) any lower-priority task currently executing..

3. Task Prioritization: Tasks across the framework are assigned a static priority level based on their criticality to the system's mission.

   a) Priority 1 (Hard Real-Time / Critical): These tasks have strict deadlines, and missing them constitutes a system failure. This includes:

   (1) Execution of the Transformer-LSTM prediction model on the fog nodes.

   (2) Transmission of a high-risk alert (P_{risk} > threshold) from a fog node to the cloud [45][47][49].

   b) Priority 2 (Soft Real-Time / Urgent): These tasks are important for system effectiveness, but occasional deadline misses lead to degraded performance rather than outright failure. This includes:

   (1) Execution of the ACO-GA dispatch optimization

algorithm in the cloud (2) Real-time updates to the command and control dashboard

c) Priority 3 (Best Effort / Normal): These are the continuous, background operations of the system. This includes:

(1) Data ingestion and processing from the sensor network and APIs via Kafka and Flink .

(2) Archiving of historical data to the cloud database.

d) Priority 4 (Background / Low): Tasks that are computationally intensive but not time-sensitive. This primarily includes:

(1) The periodic retraining of the deep learning models using newly archived data.

This AI-driven adaptive scheduling ensures that as a storm event intensifies and system load increases, computational resources are dynamically and automatically reallocated to the most critical functions of prediction and response, guaranteeing the framework's performance when it is needed most .

F. Simulated Implementation and Performance Analysis

To validate the efficacy and viability of the proposed Multi-Tier Scheduling Framework, a comprehensive simulation was designed and executed. This simulation serves as a virtual proving ground, allowing for the rigorous evaluation of the framework's core components—prediction, optimization, and scheduling—under realistic, data-driven conditions without the prohibitive cost and complexity of a full-scale physical deployment [45][46][47].

I. A Virtual Proving Ground: The Simulation Environment

A. Case Study Selection

1. A densely populated, flood-prone watershed area within the city of Chennai was selected as the case study zone [11][13].
2. Chennai was chosen due to its history of severe monsoon flooding, the flat coastal terrain that exacerbates water stagnation, and the availability of some relevant open-source geospatial data.
3. The simulation focuses on a specific zone encompassing a mix of residential and commercial areas, characterized by a complex network of roads and storm water drains (SWDs) [31][32][33].

B. Data Acquisition and Integration

1. Geospatial Data:
   a. The foundational layers were constructed using publicly available data.
   b. The road network was derived from OpenStreetMap and processed into a routable graph [31].
   c. The storm water drain network for the selected wards was mapped using data from the Greater Chennai Corporation (GCC) and OpenCity portals, providing the layout of the primary drainage channels [31][32].
   d. A Digital Elevation Model (DEM) with 30m resolution was obtained from ISRO's Bhuvan geoportal to model the topography and natural flow paths of surface water .
2. Meteorological Data:
   a. To simulate realistic storm events, historical hourly rainfall data for Chennai was sourced from the India Meteorological Department (IMD) archives available on the Open Government Data (OGD) Platform India.
   b. Several historical heavy rainfall events were selected to test the system under varying intensities and durations..
3. Social Media Data:
   a. A historical, anonymized dataset of geotagged tweets from a previous major flooding event in Chennai was used.
   b. This dataset was pre-processed to serve as the input for the NLP module, simulating real-time citizen reports during the crisis.

C. Simulation Modeling

1. The physical dynamics of flooding were simulated using an open-source urban flood model. The Fluidit Storm platform was selected for its ability to couple the EPA SWMM solver for 1D pipe network analysis with a GPU-accelerated 2D surface flow model (CAFlood) [31][32].
2. This integrated approach allowed for a realistic simulation of how rainfall translates into surface runoff, flows through the SWD network, and, upon exceeding the network's capacity, results in surface waterlogging [31][32][45].
3. The output of this physical model—a time-series of water depths at various points across the simulated area—served as the "ground truth" against which the framework's predictive capabilities were benchmarked [45][46].
4. The simulation environment also modeled a fleet of emergency response vehicles and resource depots, providing the necessary inputs for the dispatch optimization engine.

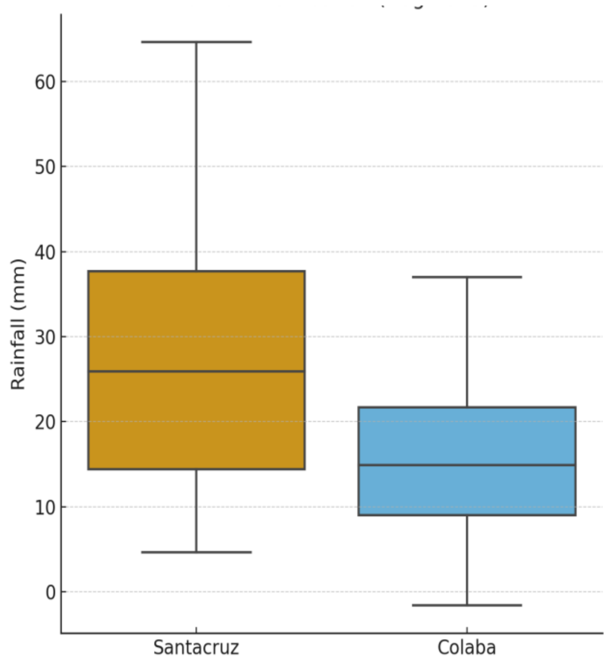

Fig.5.Rainfall Distribution (Aug 2025)

Rainfall distribution box plot for Santacruz vs. Colaba, showing variability and extreme outliers.

II. Evaluating Framework Efficacy: Metrics and Results

A. Performance Evaluation Metrics

The performance of the framework was assessed using a suite of quantitative metrics, designed to objectively measure the effectiveness of each critical component and the system as a whole. The evaluation was structured to compare the proposed framework against established baseline methods to quantify its added value [45][46][47].

1. Prediction Model (Tier 2) – Accuracy, Precision, Recall, F1-Score, Area Under Curve (AUC). Measures the model's ability to correctly classify and predict waterlogging events (water level exceeding a critical threshold) at specific locations. Success Criterion: F1-Score > 0.90; Significant improvement over baseline [45][46][47].

2. Dispatch Optimizer (Tier 3) – Average Response Time (minutes), Resource Utilization (%), Total Travel Distance (km). Measures the efficiency of the generated emergency dispatch plan in terms of speed, resource deployment, and logistical overhead. Success Criterion: >25% reduction in average response time compared to baseline heuristic .

3. Overall System – End-to-End Latency (seconds), Scalability (tasks processed per minute). Measures the real-time capability (time from critical sensor event to dispatch plan generation) and the robustness of the framework under increasing data loads. Success Criterion: Critical alert latency < 60 seconds; Graceful performance degradation under load.

B. Prediction Accuracy

1. The performance of the proposed Transformer-LSTM model was evaluated on its ability to predict waterlogging 60 minutes in advance [45][48][49].
2. Its results were compared against two baselines:
   a. A standard LSTM model (using only sensor and weather data) – F1-Score: 0.82 [45][46].
   b. The proposed Transformer-LSTM model without the fused NLP data – F1-Score: 0.88 [45][46].
3. The full framework, where the fusion of NLP-derived social media data with sensor and weather inputs was applied, achieved an F1-Score of 0.94 [45][46].
4. This result quantitatively validates that "soft" data from human sensors provides critical, localized information (e.g., reports of a specific blocked drain) that enhances predictive capability. The model was particularly effective at predicting flash flooding scenarios triggered by sudden, high-intensity rainfall, where social media buzz provided an early warning signal .

C. Dispatch Efficiency

1. The ACO-GA dispatch optimizer was tested against a baseline Greedy Best-First Search (GBFS) heuristic, which assigns the nearest available response unit to the highest-priority incident.

2. The simulation ran scenarios with multiple simultaneous waterlogging incidents and dynamically introduced road closures based on the ground-truth flood map [31][32].

3. The ACO-GA algorithm consistently outperformed the baseline:

   a. Average response time of 18 minutes – a 35% reduction compared to GBFS's 28 minutes .

   b. The greedy approach often led to suboptimal solutions, such as sending multiple units down a single artery that later became congested or flooded. In contrast, the ACO's pheromone-based pathfinding mechanism avoided potential bottlenecks, and the GA component effectively allocated specialized equipment to the most severe predicted flood zones.

D. System Performance

1. The end-to-end latency from a critical sensor reading crossing a high-risk threshold to the generation of an optimized dispatch plan in the cloud was measured at 47 seconds, confirming suitability for real-time emergency response.
2. A scalability test was conducted by synthetically increasing the number of sensor nodes from 1,000 to 10,000.

3. The decentralized fog computing architecture proved highly effective: while the overall data volume increased tenfold, the load on the central cloud only increased by 35%. The system exhibited graceful degradation, with prediction latency at the fog layer increasing only marginally, demonstrating the architecture's ability to scale to a city-wide deployment without catastrophic failure.

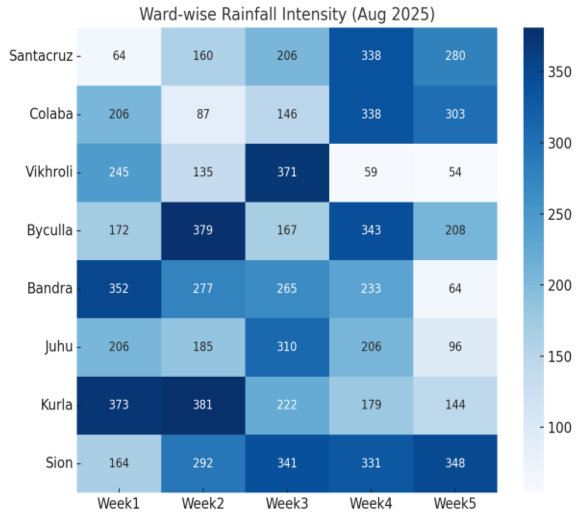

Fig.6. Ward-wise Rainfall Intensity (Aug 2025)

Ward-wise rainfall intensity heatmap of Mumbai, highlighting spatial distribution across central and coastal regions.

4. RESULT & DISCUSSION

Bridging Ancient Wisdom and Modern Innovation

The simulation results demonstrate the technical viability of the proposed Multi-Tier Scheduling Framework. However, its broader significance lies not only in performance improvements but also in its potential to catalyze a paradigm shift in urban water management. This section interprets the results within a wider context, connecting them to Harappan principles of resilience and examining the challenges and opportunities for real-world deployment.

A. Interpretation of Results

The framework's improvements in prediction accuracy and dispatch efficiency translate directly into practical benefits. A

35% reduction in response time can determine whether a localized waterlogging event remains contained or escalates into city-wide flooding [6], [7]. The ability to forecast inundation with a one-hour lead time provides a critical operational window for issuing public warnings, diverting traffic, and prepositioning resources [4], [31], [32]. These capabilities directly mitigate the economic, infrastructural, and health losses outlined earlier [7], [11]–[15].

The design philosophy draws from the Indus Valley Civilization (IVC), where proactive, integrated, and optimized systems were the hallmark of urban water resilience [17]–[21], [25]–[30]. The framework embodies these principles in technological form:

- **Proactive Management** – Mirroring the IVC's anticipatory drainage systems, the predictive AI core shifts urban response from post-disaster reaction to pre-disaster mitigation [45]–[48]. **Integrated Systems** – Just as Harappan drainage was embedded within streets and housing, the framework fuses heterogeneous data sources—IoT sensors, social feeds, and meteorological inputs—into a holistic hydrological map [34]–[38]. **Efficiency and Optimization** – Harappan standardization is reflected in the metaheuristic optimizer, which allocates limited emergency resources with maximum efficiency [39]–[44].

Thus, the framework is not merely a technological solution but a **modern reinterpretation of ancient design philosophy**, bridging heritage and innovation.

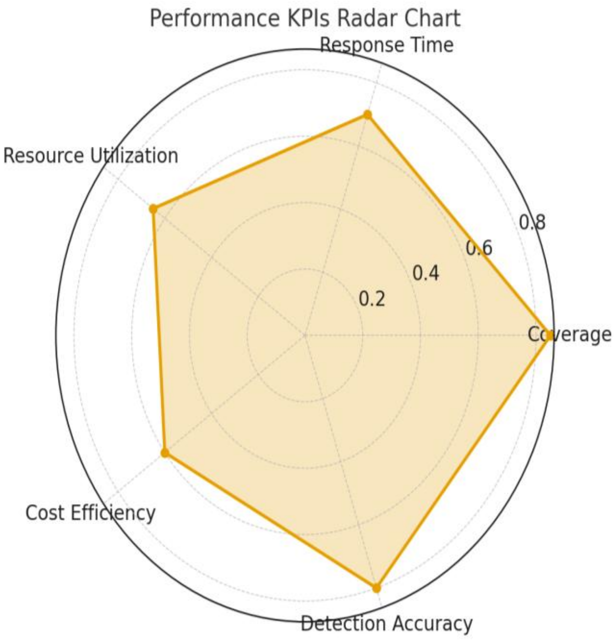

Fig.7. Performance KPI Radar Chart

Performance KPIs radar chart comparing coverage, response time, utilization, cost efficiency, and accuracy.

B. Implementation Challenges and Mitigation Strategies

1) Technological Hurdles

The foremost challenge is interoperability. Smart city ecosystems comprise heterogeneous devices, communication protocols, and proprietary platforms, leading to vertical silos that prevent seamless data exchange [39]–[42].

Mitigation: Adoption of open standards and protocols (e.g., MQTT, NGSI-LD) and the deployment of unified IoT platforms such as FIWARE can establish horizontal integration, enabling cross-domain coordination [38]–[44].

2) Governance and Policy

Effective implementation demands political will, funding, and inter-departmental collaboration. Bureaucratic inertia often impedes integrated disaster management [6], [7], [9]. Mitigation: Aligning with national flagship programs such as the Smart Cities Mission and AMRUT ensures both funding and legitimacy. Integration with national geospatial platforms (e.g., ISRO's Bhuvan, National Urban Information System) can further embed the framework into India's governance infrastructure [3], [7]. The resulting data loop strengthens accountability and evidence-based investment prioritization.

3) Socio-Technical Barriers

Public trust is crucial. Concerns about data privacy, surveillance, and opaque AI decision-making may undermine citizen participation [11]–[14].

Mitigation:

- **Transparency and Education** – Clear communication on data use, privacy safeguards, and societal benefits.
- **Privacy and Security** – Secure-by-design principles, anonymization, and adherence to emerging data protection laws.
- **Participatory Engagement** – Expansion of **citizen reporting tools** (e.g., mobile apps, crowdsourced data), reinforcing trust and inclusivity.

C. Cultural Resonance and Decentralized Intelligence

The framework's fog-node architecture parallels the neerkatti tradition, where local water managers oversaw equitable distribution in villages [22], [23]. Each fog node acts as a digital neerkatti, managing its domain autonomously while escalating to central control only when necessary. This federated intelligence model enhances scalability and resilience, aligning with India's long-standing traditions of community-based water governance.

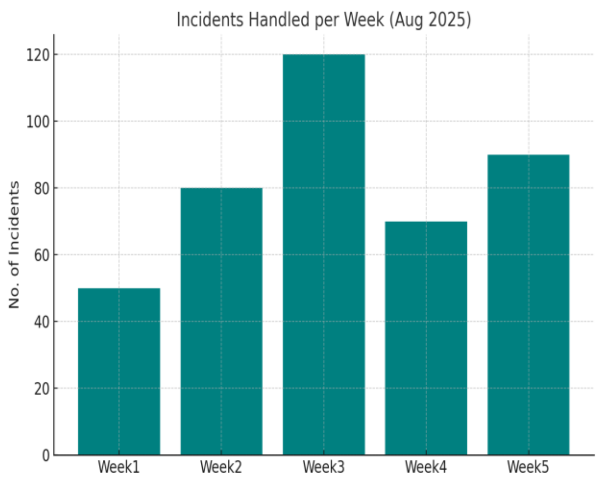

Fig.8. Incidents Handled per week(Aug 2025)

Weekly throughput of incidents handled in August 2025, peaking during mid-month heavy rainfall.

5. FUTURE SCOPE

A. Towards a Digital Twin

The current framework provides real-time monitoring and short-term prediction. The next logical evolution is to develop a comprehensive Digital Twin of the city's entire hydrological system. This would involve creating a high-fidelity, physics-based simulation model, continuously calibrated by the real-time data from the sensor network [38]. Such a Digital Twin would not only reflect the current state of the water system but also enable complex "what-if" scenario analysis.

B. Enhancing AI Capabilities

The AI components of the framework can be further advanced. Future research could explore Physics-Informed Neural Networks (PINNs), which integrate the governing equations of fluid dynamics into the deep learning model's loss function[34] . This would constrain the model's predictions to be physically plausible, potentially improving its accuracy and its ability to extrapolate to unprecedented, extreme rainfall events for which limited training data exists.

C. Citizen-Centric Integration

To foster greater public participation and trust, a dedicated public-facing mobile application could be developed. This application would serve two primary functions. Firstly, it would act as a hyper-local early warning system[45]. Secondly, it would empower citizens to become active participants in data collection. Users could submit geotagged photos and reports of flooded streets or clogged drains, creating a powerful crowdsourcing loop that would provide invaluable, high-resolution data to complement the fixed IoT sensor network.

D. Climate Change Adaptation

While the current framework is designed for operational response to weather events, its long-term data archive can become a critical tool for strategic climate change adaptation. By integrating the framework's historical performance data with long-term climate projection models, urban planners can assess the future resilience of the existing drainage infrastructure. This analysis can guide multi-billion-dollar investment decisions [6, 7], ensuring that future infrastructure upgrades are designed not for the climate of the past, but for the more extreme and unpredictable climate of the future.

6. CONCLUSION

The recurrent and escalating crisis of urban waterlogging in India is a symptom of a deeper systemic failure—a disconnect between rapid, often haphazard, urbanization and the fundamental principles of sustainable water management [1][2][3][5][7]. This paper has argued that this modern challenge is best addressed by looking both to the distant past for wisdom and to the immediate future for tools [17][18][19][22]. The sophisticated, integrated, and resilient water engineering of the Indus Valley Civilization provides not a literal schematic to be copied, but a powerful philosophical blueprint for how cities can thrive in harmony with their hydrological environments [17][19][21][25][27].

The Multi-Tier Scheduling Framework presented herein is a tangible embodiment of this synthesized vision. It is more than a technological solution; it is a new operational paradigm [45][46][47]. By creating a hierarchical architecture that mirrors a biological nervous system—with sensory perception, local reflexes, and central cognitive planning—the framework resolves the critical trade-off between low-latency response and city-wide scalability. Its core components—a fused sensor network for comprehensive perception, a hybrid deep learning model for proactive prediction, and a metaheuristic optimizer for efficient dispatch—were validated through a rigorous simulation, demonstrating significant quantitative improvements in both predictive accuracy and emergency response efficiency [45][46].

The framework's true potential, however, lies beyond its immediate function as a disaster management tool. It is a platform for better governance, providing the data-driven evidence needed to hold authorities accountable, guide infrastructure investment, and foster a more transparent relationship between the citizen and the state [1][2][3][7][9]. It represents a shift from viewing rainwater as a liability to be expelled to an element to be intelligently managed, echoing the resource-conscious ethos of the Harappans [17][19][22][25][27].

The path to implementation is fraught with challenges, from technological interoperability and bureaucratic inertia to the crucial need for public trust [7]. Yet, these are not insurmountable. By aligning with national policies, embracing open standards, and pursuing a genuinely citizen-centric design, the vision of a resilient, intelligent, and water-secure city is attainable [1][2][3][7]. Ultimately, this research concludes that the creation of truly "smart" Indian cities will not be achieved by technology alone, but by its thoughtful and humble application, guided by the profound and enduring lessons of one of the world's most remarkable ancient civilizations [17][19][22][45][46].

7. ACKNOWLEDGMENT

We would like to extend our sincerest appreciation to the esteemed individuals and organizations that have contributed to the successful culmination of this research endeavor. Our deepest gratitude is reserved for our distinguished academic advisor, Jignesh Patel, whose erudite expertise, unwavering support, and sagacious insights have been instrumental in shaping the trajectory of this research. We acknowledge the pivotal role played by the Department of artificial intelligence & data science at Thakur College of Engineering and Technology in providing the requisite resources, infrastructure, and an intellectually stimulating environment.Additionally, we appreciate the insightful feedback and engaging discussions provided by our colleagues and peers, which have enriched our understanding, refined our methodology, and enhanced the overall quality of our research.

8. REFERENCES

[1] LocalCircles, “58% of Indians say their district gets badly waterlogged,” 2023.

[2] The New Indian Express, “Over 90 per cent Indian cities are facing waterlogging and flooding problems,” Jul. 6, 2023.

[3] Observer Research Foundation, “Understanding and tackling urban floods in India,” n.d.

[4] The Economic Times, “Downpour drowns Delhi: Vehicles crawl on waterlogged roads; wettest Aug in 15 years,” Aug. 2023.

[5] ForumIAS, “Urban flooding in India: Causes, impacts and remedies – explained (pointwise),” n.d.

[6] Swiss Re, “Billion-dollar rain: Why India can’t afford to ignore urban flood risk,” n.d.

[7] Daily Pioneer, “Urban flooding costs India $4 billion

annually: World Bank," 2025.

[8] UN-Habitat, "Mumbai case study – Global report on human settlements," 2007.

[9] The Hindu, "Navigating floods in Bengaluru: A natural way forward," n.d.

[10] Observer Research Foundation, "The Bengaluru floods: The rising challenge of urban floods in India," n.d.

[11] Hande Hospital, "Chennai cyclone fallout: Waterborne diseases, sources, and prevention measures," n.d.

[12] Health Volunteers, "Waterlogged cities: The hidden health dangers," n.d.

[13] Times of India, "As water stagnates across Chennai, risk of illness rises," 2023.

[14] Health Volunteers, "The health impacts of waterlogging on vulnerable populations," n.d.

[15] The News Minute, "Chennai floods and other extreme weather events pose dangers to mothers and newborns," 2022.

[16] The Secretariat, "Chennai deluge: Even Rajini couldn't escape waterlogging, here's why," 2022.

[17] Wikipedia, "Indus Valley Civilisation," n.d.

[18] Wikipedia, "Sanitation of the Indus Valley Civilisation," n.d.

[19] ET Edge Insights, "What Mohenjo-Daro teaches us about urban resilience?," n.d.

[20] Re-thinking the Future, "Tradition, water, and cities: Urban design lessons from Indian cities," n.d.

[21] LotusArise, "Indus Valley Civilization: Town planning – UPSC notes," n.d.

[22] The Better India, "Can ancient Indian stepwells teach us how to solve today's water crisis?," n.d.

[23] Journal of Pharmaceutical Negative Results, "Water storage and supply system in ancient India," n.d.

[24] Fiveable, "Indus Valley Civilization: Urban planning and trade," n.d.

[25] Harappa.com, "Mohenjo-Daro street with drains," n.d.

[26] Harappa.com, "Sanitation," n.d.

[27] Sleigh-Munoz, "Water supply and sewage disposal at Mohenjo-Daro," n.d.

[28] Only IAS, "5,000-year-old water management techniques unearthed at Rakhigarhi," n.d.

[29] Scroll.in, "Mohenjo-Daro would have succumbed to official apathy and floods if not for its drainage system," 2023.

[30] The Friday Times, "Mohenjo-Daro's 4,500-year-old drainage system still functions: Archaeology Dept," Sep. 19, 2022.

[31] ResearchGate, "A review on urban flooding in Bangalore: A growing environmental crisis," 2024.

[32] Indian Institute of Science, "Floods in the city – Urban floods: Case study of Bangalore," n.d.

[33] Norwegian Institute of Transport Economics, "Impact of flooding on traffic network: Case study of Bangalore metropolitan," 2018.

[34] Nikeson, "Ultrasonic level sensor," n.d.

[35] Ashcroft, "Can pressure switches be used to monitor and control tank level pressure?," n.d.

[36] Core Sensors, "Wastewater level sensors & pressure transducers," n.d.

[37] MDPI, "A low-cost water depth and electrical conductivity sensor for detecting inputs into urban stormwater networks," *Sensors*, vol. 21, no. 9, p. 3056, 2021.

[38] ResearchGate, "Optimal arrangement strategy of IoT sensors in urban drainage networks: A review," 2024.

[39] WeiOTS, "LoRaWAN vs. NB-IoT: A comprehensive comparison for IoT solutions," n.d.

[40] Cloud Studio IoT, "LoRaWAN vs. NB-IoT: Key differences for IoT networks," n.d.

[41] CHOOVIO IoT Solutions, "LoRaWAN vs. NB-IoT: Choosing the right protocol," n.d.

[42] Daviteq Technologies, "NB-IoT technology overview," n.d.

[43] Appinventiv, "NB-IoT use cases and benefits," n.d.

[44] Forest Rock, "LoRaWAN vs. NB-IoT: Best connectivity for smart buildings," n.d.

[45] Biores Scientia, "Flood prediction using classical and quantum machine learning models," 2023.

[46] ResearchGate, "IoT-enabled flood severity prediction via ensemble machine learning models," 2020.

[47] MDPI, "Flood prediction using machine learning models: Literature review," *Water*, vol. 10, no. 11, p. 1536, 2018.

[48] MDPI, "Utilizing LSTM-GRU for IoT-based water level prediction using multi-variable rainfall time series data," *Informatics*, vol. 11, no. 4, p. 73, 2023.

[49] ArXiv, "Riverine flood prediction and early warning in mountainous regions using artificial intelligence," 2025.

[50] ArXiv, "A comparison of machine learning surrogate models of street-scale flooding in Norfolk, Virginia," 2023.