Predictive Data Systems for Sustainable Farming and Hospitality

Shashikala Valiki
Independent Researcher
shashikala.valiki.researcher@gmail.com

Abstract

Scalable Big Data Processing and Predictive Decision Systems for Sustainable AgriTech and Smart Hospitality align with principles of growing Big Data by maximising processing throughput while minimising completion time and resource consumption. Decision systems in the background support immediate predictions, long-term forecasts, and model evaluations, providing predictive analytics for Smart Hospitality and sustainable AgriTech. Artificial intelligence models for predictive irrigation have a data-driven architecture with large-scale data from public repositories and real-time weather feeds. Smart Hospitality applies a multi-timescale approach for resource and operational cost management via demand forecasting and optimal capacity allocation. Environmental and operational impact narratives form an integrated life-cycle perspective of practice in AgriTech and Smart Hospitality.

Growing volumes of data and information sourcing within hospitality operations fuel the need for predictive analytics. Predictive systems supported by environmental and mobility information are important drivers for Smart Tourism and visible within Smart Cities.The capacity and flow constraints of hospitality environments become decision system use cases for demand forecasting and resource allocation. Data-driven decision systems predict demand streams and architecture support management of hospitality processes.

Keywords: Scalable Big Data Processing, Predictive Decision Systems, Sustainable AgriTech, Smart Hospitality Analytics, High-Throughput Data Pipelines, Real-Time Predictive Analytics, Demand Forecasting Models, Resource Optimization Systems, AI-Driven Irrigation, Weather-Driven Analytics, Multi-Timescale Forecasting, Operational Cost Optimization, Data-Driven Decision Making, Environmental Impact Analytics, Smart Tourism Systems, Smart City Integration, Mobility Data Analytics, Capacity Planning Models, Hospitality Resource Allocation, Lifecycle Analytics.

1. Introduction

Sustainable AgriTech and Smart Hospitality are critical for building a sustainable economy and mitigating environmental risks. Predictive analytics plays a key role in enabling sustainability by improving efficiency and the responsible use of resources. Suitable machine learning and simulation approaches can deliver useful predictive models to guide

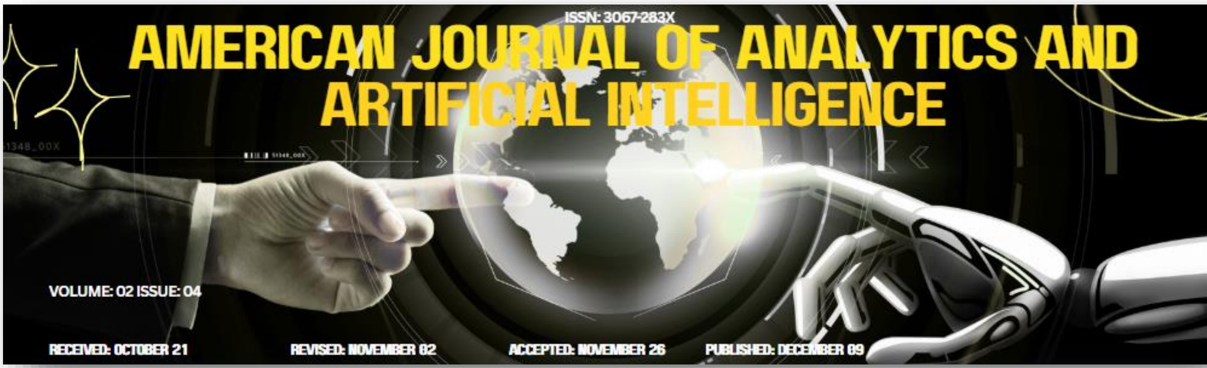

management decisions. Scalable big-data processing requires sound architecture choices and design decisions throughout the lifecycle. Techniques from Computer Science can therefore provide a significant advantage in establishing resilience and adaptability, driving down costs, and improving profitability.

Integrated data acquisition, processing, and prediction infrastructures are essential to prevent system producer and consumer silos in multi-hop scenarios. These infrastructures represent a feasible alternative to large-scale generic aggregators, thus making real-time processing step viable. Moreover, the edge-cloud continuum concept, supported by information processing capabilities at the network edge, enables the necessary response times particularly relevant in Human-Machine Interaction, Internet-of-Things, and Remote-Control scenarios. Supporting information processing at the edge limits data transfer volume and network load and can therefore address the data privacy challenge.

2. Foundations of Scalable Big Data Processing

Big data programs integrate heterogeneous data sources of various types and volumes, enabling scalable advanced analytics. Such processes depend on the underlying data-processing framework. The system of interest is defined by a data specification for ingestion, partitioning, and indexing, enabling scalable storage, processing, and consumption. Efficient data-integration methods are also fundamental, ranging from custom data ingestion and processing pipelines to scheduled extract-transform-load workflows.

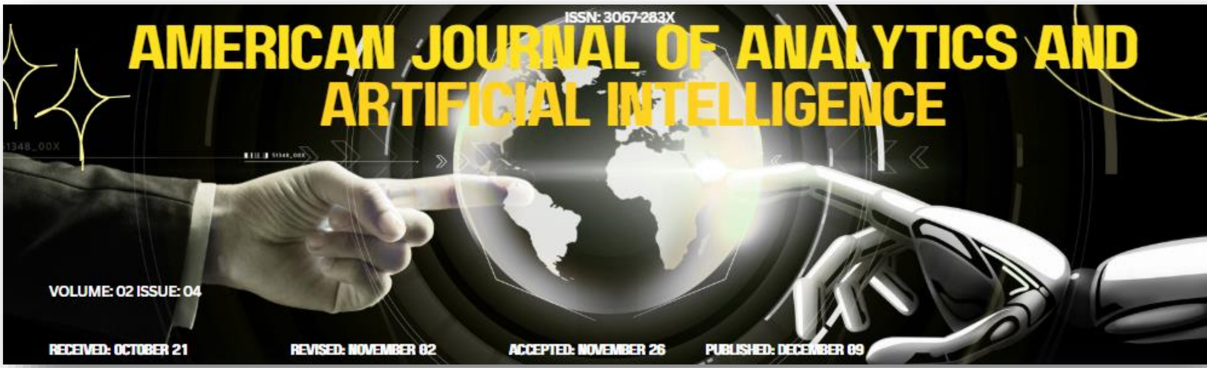

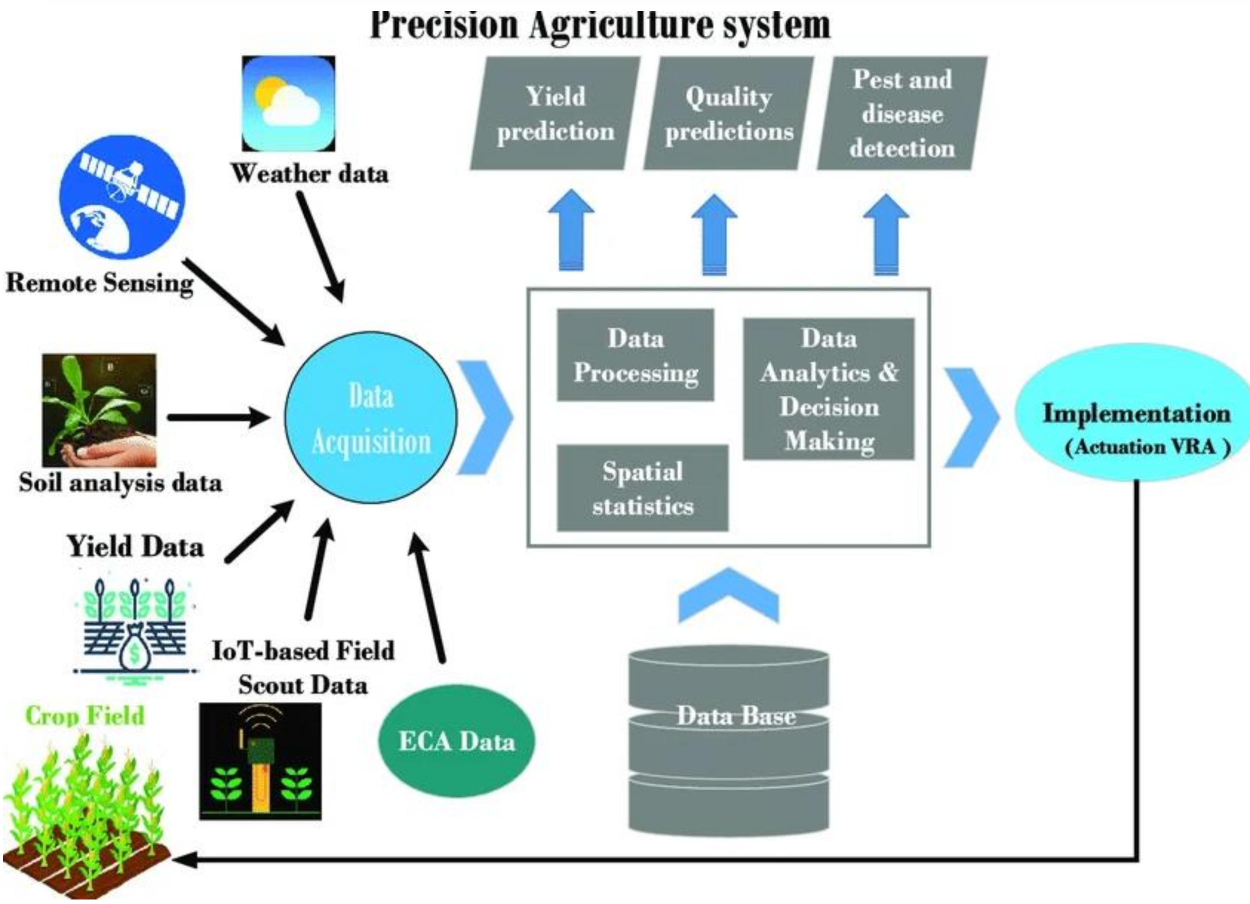

Fig 1: Big data-based precision agriculture system representation

Distributed computing can be classified as data-parallel or task-parallel, depending on whether all computing units operate on the same data or not. Considering fault tolerance, throughput, latency, and consistency, batch processing is prioritized by platform-as-a-service solutions, data-lake strategies, and Apache Hadoop; these also serve offline and near-real-time needs. Data streaming targets latency-critical processes. The hybrid model of stream processing with micro-batch execution, popularized by Apache Spark Streaming, is currently the preferred approach for real-time analytics. Ingestion latency is reduced using publish-subscribe patterns at the expense of complex data flow, control, monitoring, and governance. A high number of producers and consumers for a common data source can also hinder data governance.

2.1. Distributed Computing Paradigms

Various types of distributed computing models are available in the literature to satisfy different requirements such as control over properties like throughput and latency, development complexity, and fault tolerance. A fault-tolerant distributed system is capable of continuing operation despite the failure of one or more components, while the failure of a process in an approach without fault tolerance causes the total system to collapse. Latency is the delay taken for the result of a query to reach the user after it is issued and represents response time, throughput being the number of queries that can be executed in parallel and in a unit of time. These two properties are addressed jointly in replication, where the same data is stored redundantly in several locations, but such redundancy can lead to update consistency problems. A consistency model defines under what conditions a read operation returns values that are consistent with the last write.

These properties can be tuned to the specific requirements of applications using either of two basic paradigms: actor-based systems and master-worker systems. In an actor-based approach, each node executes in an infinite loop and is activated by messages sent by other nodes. Such a system is strongly fault tolerant and bad latency/high throughput can be expected. The master-worker model features a master that assigns tasks to workers in a task queue and is generally supported in many cloud computing platforms. Systems based on it provide lower latency/high throughput mitigating scalability problems with the master.

2.2. Data Ingestion and Storage Architectures

Processing of scalable big data streams at scale requires special care when building the data ingestion layer. Unlike file-based query systems, NoSQL and NewSQL databases, distributed storage systems with their natively table replicas, streaming platforms, and in-memory databases, most of these storage systems have no natural inherent fault tolerance with their high-availability clusters.

Batch processing, used for data that can bear some latency, suffers from other trade-offs. Operation windows are done for long-term checks for possible data losses, and for matching non-matching batch jobs especially in mixed-mode operation, they stay more complex. Stream processing, cloud-based-storage systems, scaling-out on-demand, elasticity in short-time operation, natively being able to share, reproduce and explore non-volatile joined data, have opened whole new worlds for exploratory analysis. They bring in reuse by support systems, trader-support systems, dynamic dash-boarding, retrieval-time-style-analysis, or explorer-satellite-languages.

Scalability in batch processing requires special care again, adaptation of the underlying complex task-governor to the scaling needs, operational depth. The secret depends again on using a scaling-on-demand architecture. The engine decides on the input table size, presence, completion time, input type, table-saturation factor and synchro pattern.

Synchronous jobs are sent to a more scalable batch engine, thereby keeping complexity low. Crude and early types of data processing are the secret to keeping latent data reliable over long periods and still exploitable for external data until finally marked obsolete.

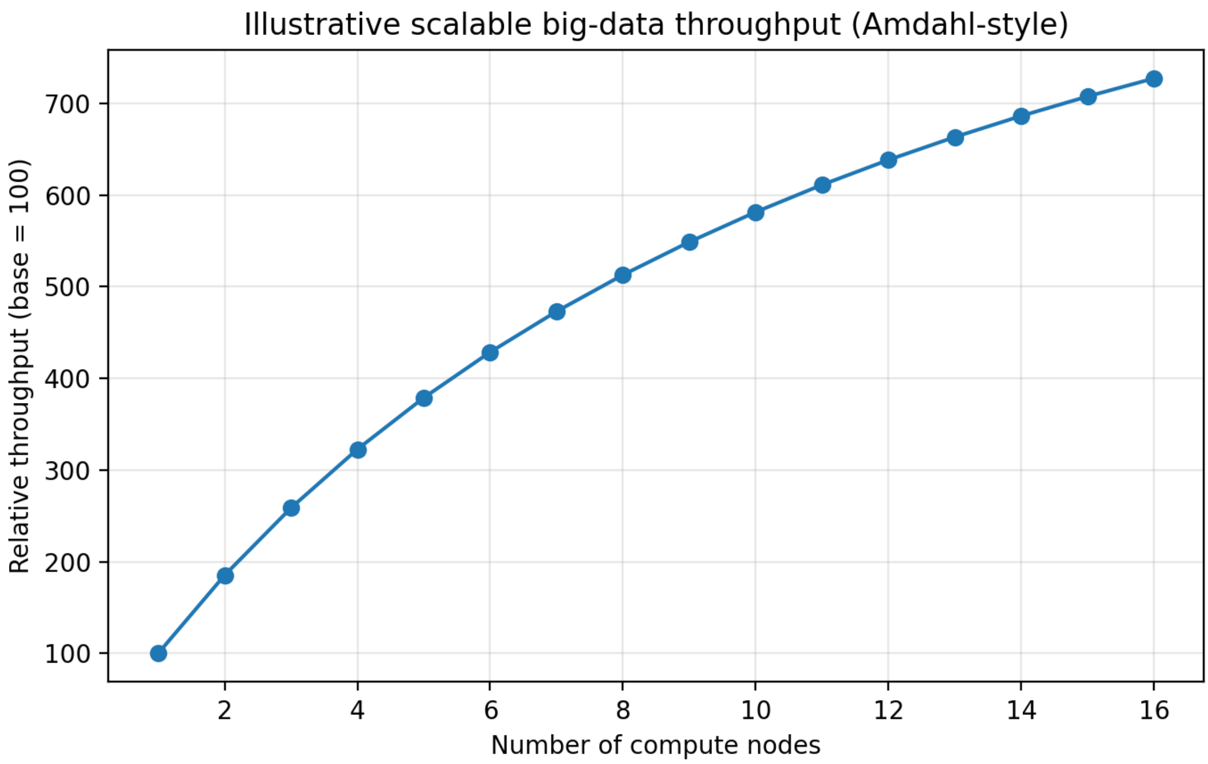

Equation A. Scalable big-data processing equations

A1. Throughput from completion time

Let

- N_j = number of jobs completed
- T = total elapsed time

By definition,

$$X = \frac{N_j}{T}$$

ISSN: 3067-283X
AMERICAN JOURNAL OF ANALYTICS AND ARTIFICIAL INTELLIGENCE
VOLUME: 02 ISSUE: 04
RECEIVED: OCTOBER 21 REVISED: NOVEMBER 02 ACCEPTED: NOVEMBER 26 PUBLISHED: DECEMBER 09

If one job is completed in time T_1, then for a single-job view,

$$X_1 = \frac{1}{T_1}$$

With n nodes and completion time T_n,

$$X_n = \frac{1}{T_n}$$

So relative throughput gain is

$$\frac{X_n}{X_1} = \frac{1/T_n}{1/T_1} = \frac{T_1}{T_n}$$

Define speedup

$$S_n = \frac{T_1}{T_n}$$

Therefore,

$$X_n = X_1 S_n$$

A2. Amdahl-style speedup derivation

Suppose the workload has:

- serial fraction $1 - p$
- parallel fraction p

On one node, normalized time is

$$T_1 = (1 - p) + p = 1$$

On n nodes:

- serial part remains $(1 - p)$
- parallel part becomes p/n

So

$$T_n = (1 - p) + \frac{p}{n}$$

Now speedup is

$$S_n = \frac{T_1}{T_n}$$

Since $T_1 = 1$,

$$S_n = \frac{1}{(1 - p) + \frac{p}{n}}$$

3. Methodology

A set of design choices ensures the usability and dependability of the proposed predictive models. They are carefully explained considering the entire predictive pipeline. Data collected from real-world sources or openly available repositories form the starting point. By employing an iterative process tailored for the respective problem domain, the gathered data are first pre-processed and subsequently converted into appropriate statistical, classification, or regression models and fine-tuned using known evaluation protocols. These procedures include feature selection and engineering closely monitored according to the model category and objective. Finally, the availability of data, statistical significance of results, and sufficient model complexity enable replication of the predictive design.

Machine learning approaches for sustainable agriculture, namely, yield density modeling and a precision irrigation system for predictive water demand, are developed. The predictive pipeline for yield density employs a broad feature set based on environmental, topological, and crop management factors. The model incorporates environment-based cross-validation to improve generalization, and three regression algorithms are compared to analyze predictive performance at country level. The precision irrigation system enables accurate prediction of irrigation demand in a given period, using a limited feature set derived from weather data. 湘潭市 · 湘潭大学 ● machine learning methodologies supportive of sustainable agriculture practices and decision support have been proposed with specific designs and evaluation criteria in mind.

Table 1 — Main analytical quantities

Area	Symbol	Meaning	Typical Unit
Big-data processing	T_n	completion time on n nodes	time

Area	Symbol	Meaning	Typical Unit
Big-data processing	X_n	throughput on n nodes	jobs/time
Big-data processing	S_n	speedup on n nodes	dimensionless
Big-data processing	E_n	parallel efficiency	dimensionless
Irrigation	I_t	irrigation demand at time t	mm/day or L/day
Irrigation	$ET_{0,t}$	reference evapotranspiration	mm/day
Irrigation	R_t	rainfall	mm/day
Irrigation	SM_t	soil moisture	fraction or %
Hospitality	D_t	predicted demand	rooms/meals/events
Hospitality	C_t	allocated capacity	service units
Hospitality	E_t	energy usage	kWh
Sustainability	GHG	greenhouse gas emissions	kg CO_2e
Business	ROI	return on investment	%
Business	TCO	total cost of ownership	currency

3.1. Leveraging Machine Learning to Enhance Agricultural Sustainability

Artificial Intelligence plays a pivotal role in Digital Transformation across various domains, including agriculture. Precision Agriculture leverages technologies such as the Internet of Things (IoT), Artificial Intelligence (AI), and Machine Learning (ML) to optimize agricultural production and deliver efficient food supply services. ML applications in agriculture primarily focus on crop yield prediction, irrigation management, soil analysis through ML models, soil properties analysis, and pest detection. Other areas supported by ML prediction models include soil quality and crop diseases. Predictive models improve crop yield by learning from historical data and related factors. Crop yield prediction enhances decision-making and risk management by assisting farmers and stakeholders.

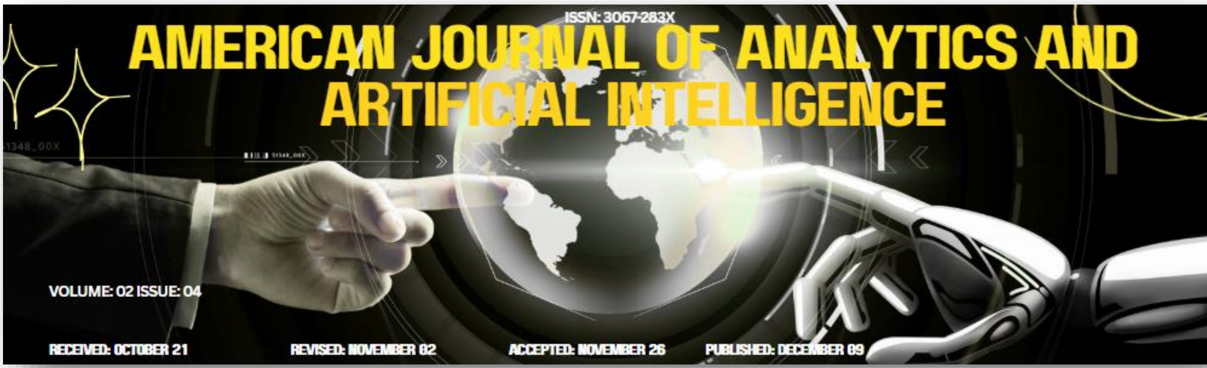

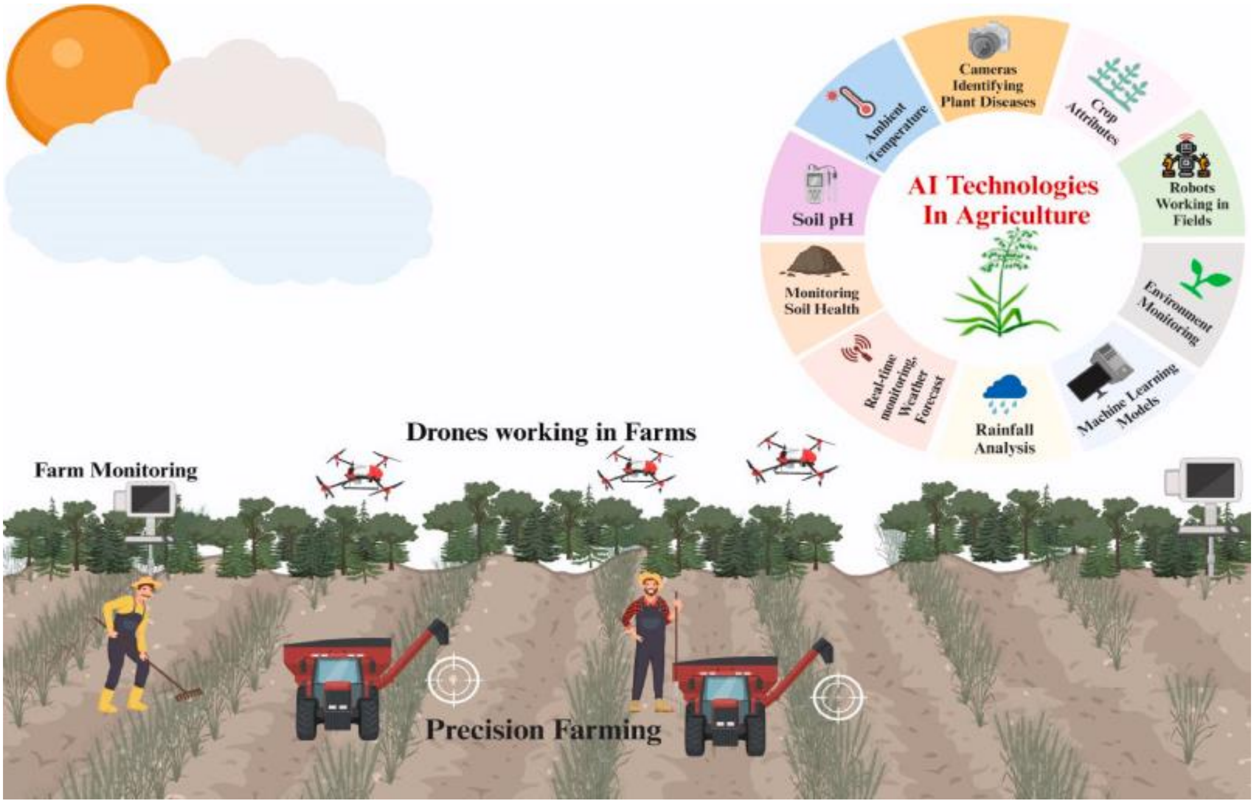

Fig 2: Artificial intelligence in agriculture

Crop yield is influenced by various parameters, including weather conditions, soil, crop growth and management, and surrounding environments. Crop yield prediction models can be developed at different levels, including field, district, state, and country levels. The concept of crop prediction has been applied in resource-constrained countries demand-supply gaps to secure food for growing populations. A regression model predicts the crop yield in a selected area for the best growth period. Feature selection aided crop prediction efficiency. Satisfactory accuracy levels in ML models allow farmers to make effective decisions regarding crop selection to maximize the yield generated.

In Smart Hospitality, ML models enable consumption demand prediction during festive periods, improving timely services and generating additional revenue. Accurate forecasting and dynamic allocation of rooms enhance service quality and customer experience during peak seasons. Demand management, using a multi-decision point model for multi-class production, increases profit, customer satisfaction, and on-time delivery. For easy usage and integration, models are packaged as microservices and deployed in a containerized environment for online prediction services.

ISSN: 3067-283X
AMERICAN JOURNAL OF ANALYTICS AND ARTIFICIAL INTELLIGENCE
VOLUME: 02 ISSUE: 04
RECEIVED: OCTOBER 21 REVISED: NOVEMBER 02 ACCEPTED: NOVEMBER 26 PUBLISHED: DECEMBER 09

4. Objectives of the Study

The objectives comprise the hypotheses and metrics guiding the research toward a measurable and actionable outcome. In the context of precision agriculture and predictive irrigation, the aim targets ecological sustainability, enhanced resource efficiency, and improved irrigational scheduling by addressing crop-water demand forecast accuracy and increase predictive capability over temporal horizons. Smart hospitality focuses on demand forecasting and service allocation to promote operational efficiency and energy conservation while minimising carbon footprint. These objectives are formulated as comprehensive predictive models capable of optimally provisioning resources across operational stages.

Although sustainability can be a non-quantifiable measurement of success, the goals of these study areas possess measurable aspects that can be captured by substantiating the underlying assumption. The specific hypothesis for precision agriculture states that predictive irrigation can be achieved without compromising ecological sustainability by maintaining accuracy of evapotranspiration and soil moisture content. The contrary constructs predictive modelling supporting smart hospitality. Accuracy of demand forecast can positively affect service efficiency and energy conservation during operation. Alongside predictive irrigation, capacity planning based on forecast confidence aids in further optimising service allocation.

4.1. Study Goals and Research Aims

Accelerating systems integration in agriculture and hospitality requires scalable services capable of decision support through predictive analytics and machine learning. Realizing this ambition demands reduced energy depletion, lower operational costs, and diminished environmental impact without compromising service quality. Implementing such systems, however, remains an open challenge. Hypothesis testing considers two measurable objectives: predictive analytics deployed on scalable infrastructures reduce energy consumption by a minimum of 10% while improving operational or customer service efficiency.

Precision agriculture minimizes climatic and weather variability through decision systems governing dynamic irrigation capacity, enabling reduced water depletion and groundwater regeneration. Moreover, these systems validate emergent set-point savings associated with predictive irrigation. Air-conditioning service demand forecasting and subsequent resource allocation maximize energy efficiency and upscale service quality across smart hospitality properties. Predictive demand analytics, moreover, determine operational capacity and enable proactive third-party resource procurement at minimal cost. AI-and-IoT make possible a Factored Product–Customer Recommender capable of generating two-sided allocation tables, creating and optimizing supply and service packages.

5. Research Summary

Research methods and key results are summarized, with attention to impacts on sustainability and operational efficiency and an outline of their interpretation. Scalable processing of data from heterogeneous sources on the edge–cloud continuum supports machine-learning models predicting environmental impacts of agriculture and optimally allocating energy and resources in hospitality.

Recent years have seen a dramatic increase in data sources that have the potential to enhance decision systems for smarter, more sustainable applications in fields such as AgriTech and hospitality. Primary success factors are therefore the capacity to effectively handle these large volumes of data at a reasonable cost and their use in developing predictive models that address critical management problems. Yet few investigations have focused on the implications of scalable big-data processing for these applications. To fill this gap, the design and construction of processing systems able to accommodate a range of emerging sources (such as weather forecasts, satellite imagery, and population density modelling) are first presented, along with the corresponding predictive models. The main contributions include the specification of the processing pipelines involved, the experimental results for the predictive models, and a summary of the business value when applied to the precision-agriculture and smart-hospitality paradigms.

5.1. Evaluation of Environmental Impacts and Life-Cycle Sustainability

Sustainability assessment methods can determine environmental impacts and life-cycle viability. For environmental analysis, the lightweight Life Cycle Impact Assessment method and its Ecoindicator 99 baseline model establish potential damage to ecosystem quality, human health, and resource depletion. In AgriTech, expected net savings estimates for water, fertilizer, and energy should meet the resource-consumption criteria of an acceptable business strategy. Top-down modeling of greenhouse gas emissions, coupled with key plant and web parameters, assesses total cost of ownership and long-term investment. For smart hospitality, demand-forecasting models focus on electricity and natural gas; an applicability test

complements a capacity-allocation use case.

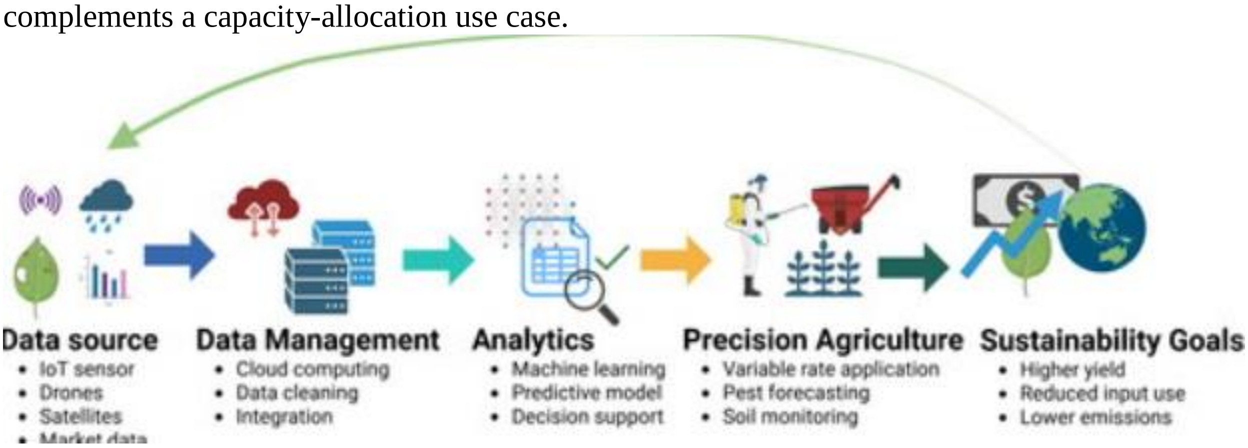

Fig 3: Big Data in sustainable agriculture

Sustainability evaluation systems can assess indicators, components, modules, or entire platforms, determining environmental, economic, and social viability. Environmental impact assessment may draw on qualitative life-cycle analysis, specifying categories, indicators, data requirements, sources, and quality. Methods identify the assessment methodology, data sources for key impact types, typology evaluations, and final improvement attainment. Methods verify economic viability and resource-consumption criteria for suitable business strategies. To close circular economy loops, methods establish total-cost structures and collate-up and allocate-down approaches for water demands. For smart hospitality energy forecasts, demand models function as proof of concept, assessing social algorithm responsibilities, price-oriented ecosystems, and demand-driven markets.

6. Predictive Analytics and Decision Systems

Machine learning models designed to optimize decision-making in AgriTech and hospitality sectors at local, regional, and operational scales leverage multiple data sources. Scalable, efficient predictive analytics processes empower real-time decision support based on timely data acquisition, cleanup, and feature engineering. Smart irrigation management for precision agriculture demonstrates the potential for increased environmental sustainability, while demand forecasting and resource allocation propose economically attractive solutions for hospitality providers.

Sustainable models for the AgriTech sector include predictive irrigation scheduling to minimize water waste, deficiencies, and excesses while maximizing crop yields and minimizing risk. Precision irrigation systems that enable automated irrigation control at the local scale require continuous monitoring of soil moisture content, weather conditions, and crop growth but drastically reduce operating expenses through advanced water-efficiency

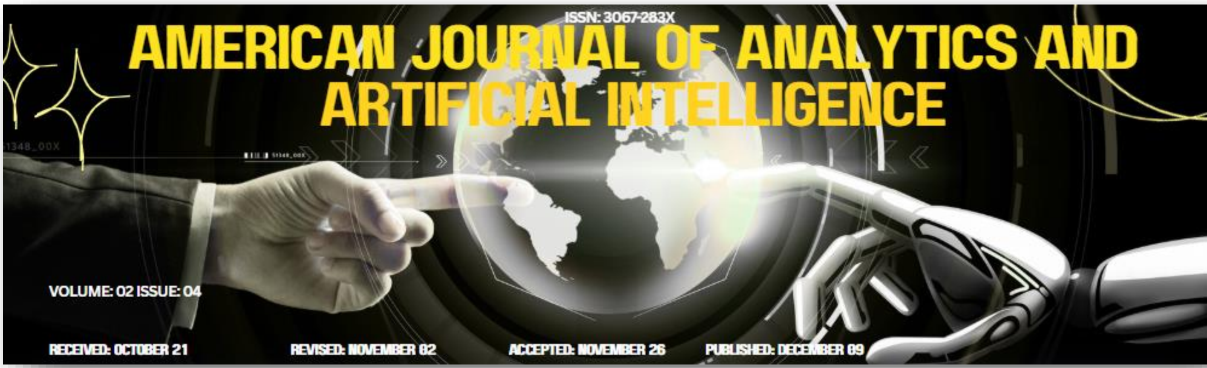

management protocols. Daily forecasts of near-future increases in the crop water demand—derived from simple empirical regression or advanced time-series analysis—complement preciser irrigation actions that maintain soil moisture near the desired threshold. They optimize water productivity either by deciding not to irrigate when the water storage is in surplus or by allocating water from other resources, such as drainage systems or tertiary storage ponds, when possible.

6.1. Machine Learning for Sustainable Agriculture

Machine learning (ML) and data science techniques feature prominently in the literature as contributors to sustainability in agriculture. Many studies develop and validate predictive models based on diverse input sources, covering a multitude of agronomic features, targeting various crop species, and focusing on different predictive objectives. Environmental sustainability is often defined in terms of the reduction of water and fertilizer usage, minimization of agrochemical applications, and limitation of greenhouse gas emissions. However, the absence of a widely accepted framework for evaluating the environmental impact of predictive models for agronomic systems makes it difficult to assess the magnitude of sustainability benefits and to perform comparisons across studies. Consequently, deployment of predictive models remains limited.

Emerging methods for supporting predictive irrigation scheduling are specific to precision agriculture. Data sources typically include open-access remote-sensing datasets and modelled weather forecasts. Validation follows a direct sampling approach, whereby the predicted irrigated volume per block is compared against the actual volume allocated by the irrigator. Research outcomes indicate that introducing soil moisture content values significantly improves irrigation prediction accuracy, while rain prediction accuracy is a critical factor when modelling deficit irrigated areas. By adopting these techniques, predictive models can be developed that match the allocated irrigation at vineyard level, helping optimize this fundamental process in precision agriculture. Nonetheless, the environmental benefits of the technique still remain to be demonstrated.

Equation B. Data-ingestion and storage equations

B1. End-to-end pipeline latency

Let total pipeline delay be the sum of stages:

- ingestion: L_{ing}
- transport: L_{net}

ISSN: 3067-283X
AMERICAN JOURNAL OF ANALYTICS AND ARTIFICIAL INTELLIGENCE
VOLUME: 02 ISSUE: 04
RECEIVED: OCTOBER 21 REVISED: NOVEMBER 02 ACCEPTED: NOVEMBER 26 PUBLISHED: DECEMBER 09

- storage: L_{store}
- processing: L_{proc}
- serving: L_{serve}

Then

$$L_{\text{total}} = L_{\text{ing}} + L_{\text{net}} + L_{\text{store}} + L_{\text{proc}} + L_{\text{serve}}$$

B2. Replication storage overhead

If a dataset of size D is replicated r times, total storage required is

$$D_{\text{total}} = rD$$

Useful storage overhead relative to one copy is

$$\text{overhead} = (r-1)D$$

6.2. Predictive Modeling for Energy and Resource Optimization in Hospitality

Predictive modeling for energy and resource optimization in hospitality encompasses demand forecasting, resource allocation, and service-capacity planning for parties and events. Demand serves as a key driver for energy, manpower, and other resource provisioning and a metric for cost and service-level performance. Predicting demand well enables efficient resourcing and service quality. Additionally, large-scale party and event service require special attention during demand forecasting to avoid honed responsiveness and cost performance.

Two types of demand-forecasting models were implemented: a flexible long-term event--demand forecasting model that enables timely lead-generation and outbound sales efforts and a near-term general-demand--forecasting model (GDFM) for day-to-day operation---scheduling. Both categories of demand forecasting also include facilitating, supporting, and predicting-demand agendas for various forms of parties and events. Data supported analysis of establishing and exerting customer priority and the actual-hit-rate mechanism in service design and screening.

An energy--resource--allocation model allocates budgets for energy, manpower, and stock, including holiday & peak considerations, to these party--event services. Service capacity must be met while keeping the system within the budget. An energy--resource--demand--capacity--allocation model translates predicted-demand forecasts into systematic

ISSN: 3067-283X
AMERICAN JOURNAL OF ANALYTICS AND ARTIFICIAL INTELLIGENCE
VOLUME: 02 ISSUE: 04
RECEIVED: OCTOBER 21 REVISED: NOVEMBER 02 ACCEPTED: NOVEMBER 26 PUBLISHED: DECEMBER 09

resource/energy distribution for proper service scheduling. Data demonstrated that insufficient incorporation of demand variations—ignored in both LBM and ANN models—put the hospitality's operating and cost-level performance at considerable risk. Further data indicated that systematic satisfaction of capacity specification did not require specific--event--allocation schemes and could use uniform distributions instead. Data also allowed a preliminary--evaluation trial of the application of Pontryagin's minimum principle to suggest a desirable direction of future research.

Table 2 — Article theme to equation mapping

Article theme	Derived model form	Why it fits the paper
Scalable big-data processing	$X_n = X_1 S_n$	paper emphasizes throughput, latency, distributed computing
Fault tolerance / scalability tradeoff	$E_n = S_n / n$	natural way to evaluate parallel quality
Predictive irrigation	$I_t = \beta_0 + \beta_1 ET_{0,t} + \beta_2 T_t + \beta_3 W_t - \beta_4 R_t - \beta_5 SM_t$	paper says irrigation is weather-driven and improved by soil moisture
Forecast evaluation	MAE, RMSE, MAPE	paper discusses validation and predictive accuracy
Hospitality demand forecasting	$\widehat{D}_t = f(x_t)$	paper focuses on demand prediction and service allocation
Resource allocation	$C_t = \widehat{D}_t + z_\alpha \sigma_t$	article mentions capacity planning using forecast confidence
Sustainability gains	$\%\Delta = \frac{\text{baseline} - \text{new}}{\text{baseline}} \times 100$	matches water/energy/carbon reduction framing
Economic viability	$ROI = \frac{B - C}{C} \times 100$	explicitly aligned with ROI/TCO discussion

7. System Architectures for AgriTech and Smart Hospitality

Integrated data pipelines and microservices are crucial for end-to-end analytics. The main types of resources supported are illustrated, including data flow diagrams, system interfaces, and data management policies. A platform-agnostic perspective considers the development of isolated deployment units (microservices) that rely on well-defined

interfaces, primarily micro-APIs, to exchange data between processes. Governance services monitor data accessibility and retention, support schema and lifecycle management, and document data lineage and provenance.

An edge-cloud continuum implements sub-system distribution in line with specific requirements, enabling inter-communication within the same local environment. Within Smart Hospitality, application placement, IoT implementation, and connectivity support local infra-structure management while reducing latency for critical and time-sensitive functionalities. The Edge-Cloud concept also supports privacy-related practices since services can restrict image and audio feed data sharing with the Cloud if local storage is enough for optimal performance. Conversely, there are also clear deployment constraints linked to security and Service Quality needs.

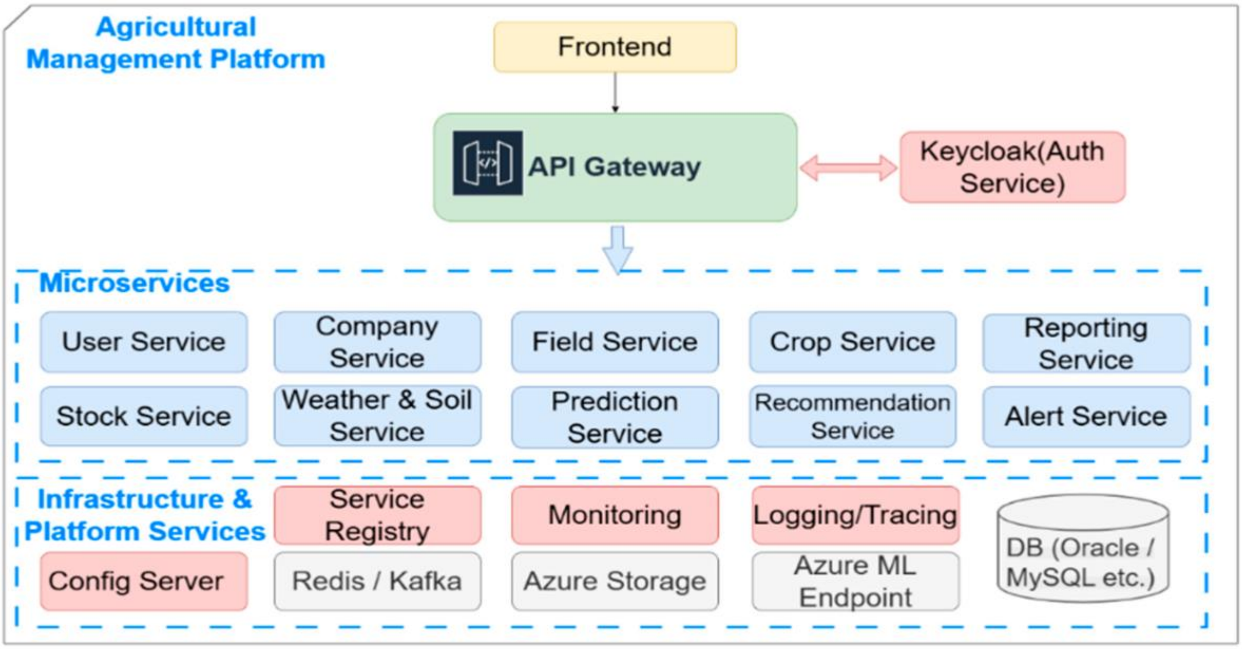

Fig 4: AgriMicro—A Microservices-Based Platform for Optimization of Farm Decisions

7.1. Integrated Data Pipelines and Microservices

Methodology research goals emphasize practical contributions with scalability potential in achieving sustainable economic growth and balanced resource management. Pivotal decision-support systems integrate advanced predictive analytics for renewables-based Smart Hospitality energy modeling, demand-driven facility-resources allocation, and predictive irrigation in Precision Agriculture. Achieved outcomes include Global Warming

Potentials-indicative Carbon dioxide equivalents (CO2e) hourly forecasts for optimized front- and back-of-the-house resource distribution; reduction of excess-harvesting costs; plus 66-fold lower prediction-error testing volume from a 22-fold data-reduction ratio compared to standards.

Integrated data-processing pipelines, self-contained microservices, and a life-cycle-based sustainability-evaluation framework anchor case studies. Coupled advanced Capex-based Total Cost-of-Ownership and Financial Statement methodologies serve risk-and-return appraisal of Capital Expenditure, with requisite datasets sourced from publicly available APIs. Smart Hospitality demand signals and energy consumption jointly inform Error Mean Absolute Percentage—proven most reliable for original deadlines as Robotic Process Automation enables model training on smaller-testing sets and meets higher hourly-resolution-operational needs.

7.2. Edge-Cloud Continuums and IoT Interoperability

Data ingestion, analytical processing, and operational decision functions reside in a mix of edge devices, public or private clouds, and cloudlet mini-clouds located closer to Cyber-Physical Systems (CPS). The architecture optimizes processing latency while accommodating the location and reliability of data providers and considering user privacy demands. Predictive models running in the Smart-City or Smart-Mobility Cloud enable efficient and novel planning of destination capacities in hotels, restaurants, public transport, and other services. Uncertainty can lead to resource oversupply and wastefulness or, at the other extreme, resource starvation, diminishing customer experience and loyalty. With precise forecasting, these destinations can allocate their resources in an efficient manner, benefiting both the customers and the business.

Interoperability across devices, between edge devices and Cloud, and among various Clouds is implemented through suitable gateways. For example, data coming from smart meters and sensors/devices may need to be translated into a common ontology, such as FIWARE or other entities, so that it can be stored in a data lake and processed. Such translation may also require the anonymization of the data. The design of gateways may have implications on the latency of the overall system. Deployed by a third party and forming part of a public or private Cloud, they may better address general problems of distribution, deployment, security, maintenance, etc. An even bigger issue arises from latency-sensitive applications of edge networks that require real-time monitoring and control of how electric power is being consumed and for what purpose.

8. Sustainability Metrics and Evaluation Frameworks

Environmental impact assessment (EIA) methodologies, such as Life-Cycle Analysis (LCA), are crucial for evaluating sustainability. The GHG Protocol, developed in collaboration with the World Resources Institute and the World Business Council for Sustainable Development, provides a standardized framework for quantifying and reporting greenhouse gas emissions, enabling organizations to identify, prioritize, and implement GHG reduction opportunities. These guidelines also aid in developing corporate climate strategies and communicating progress to stakeholders. The Protocol's design facilitates consistency, transparency, and ease of use while ensuring that advanced users maintain the technical rigor necessary for comprehensive corporate-level GHG inventory reporting.

Accurately measuring the economic viability of a sustainability-oriented investment is important to any organization and can be summarized as Return on Investment (ROI), Total Cost of Ownership (TCO), and Value Proposition. TCO accounts for the complete cost of an asset throughout its lifecycle, encompassing purchase, operational, direct labor, financing, maintenance, security, performance degradation, and disposal costs, thereby providing an informed basis for purchasing decisions. The concept of ROI signifies the value derived from an investment relative to the costs associated with that investment, while a Value Proposition describes the reasons why a customer should select a product or service.

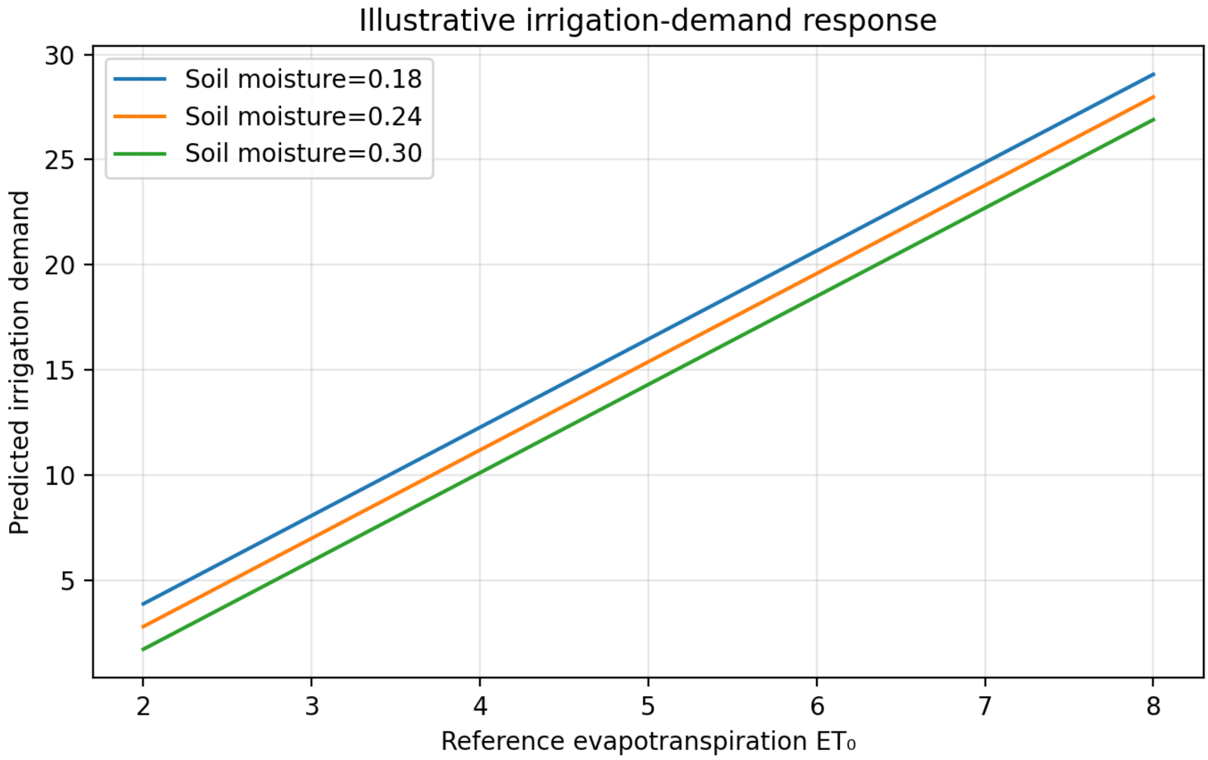

8.1. Environmental Impact Assessment and Life-Cycle Analysis

Successful implementation of scalability-sensitive decision-support solutions thus mandates a well-defined analytical framework to assess sustainability improvements. To bridge this gap, the following section establishes concise, practical, and evidence-based metrics to quantify environmental gains. The framework leverages life-cycle analysis (LCA)—a widely accepted approach for assessing the environmental impact of products, processes, or systems throughout their life cycle, from the extraction of raw materials through production and use to end-of-life treatment—and consequently demonstrates how these metrics can be applied to evaluate the most common predictive decision-support models targeted by research. Hitherto unaddressed in the literature, such a consolidated set of measures can be critical in development, testing, validation, and deployment phases.

Turning to the measures themselves, the overarching aim is to optimize environmental impact reductions attributable to the predictive decision-support system under consideration. It follows that this optimization problem should maximize one or more of the recognized LCA impact categories—exemplified by greenhouse-gas (GHG) emissions, carbon footprint, ecological footprint, and particulate-matter formation—while quantifying secondary

resources in demand, such as fertilizers and pesticides. Support models typically submit an external demand signal for systems operating under set-point or “supply-based” control strategies, such as smart hotels, or by determining equipment settings that maximize expected efficiency, such as predictive speed and capacity allocations for smart restaurants. For these models, LCA thus requires extra features: a demand signal modulating the service group and a set of performance indicators quantifying the primary environmental externality generation. The set of analytical measures can therefore be summarized as follows:

Equation C. Predictive irrigation equations

C1. Linear predictive irrigation model

Assume irrigation demand I_t depends on

- $ET_{0,t}$: evapotranspiration
- T_t: temperature
- W_t: wind speed
- R_t: rainfall
- SM_t: soil moisture

Start with a linear model:

$$I_t = \beta_0 + \beta_1 ET_{0,t} + \beta_2 T_t + \beta_3 W_t + \beta_4 R_t + \beta_5 SM_t$$

Now interpret signs physically:

- higher evapotranspiration increases water demand, so $\beta_1 > 0$
- higher temperature tends to increase demand, so $\beta_2 > 0$
- stronger wind increases evaporation, so $\beta_3 > 0$
- rainfall reduces irrigation need, so coefficient should be negative
- soil moisture reduces irrigation need, so coefficient should be negative

Rewrite with explicit negative signs:

$$I_t = \beta_0 + \beta_1 ET_{0,t} + \beta_2 T_t + \beta_3 W_t - \beta_4 R_t - \beta_5 SM_t$$

with $\beta_4, \beta_5 > 0$.

C2. Irrigation demand with non-negativity constraint

Physical irrigation cannot be negative, so if the linear predictor becomes negative, demand is set to zero.

Define latent demand:

$$I_t^* = \beta_0 + \beta_1 ET_{0,t} + \beta_2 T_t + \beta_3 W_t - \beta_4 R_t - \beta_5 SM_t$$

Then actual irrigation is

$$I_t = \max(0, I_t^*)$$

or equivalently

$$I_t = \max\left(0, \beta_0 + \beta_1 ET_{0,t} + \beta_2 T_t + \beta_3 W_t - \beta_4 R_t - \beta_5 SM_t\right)$$

8.2. Economic Viability and Return on Investment

Computation and data storage in the cloud are inherently flexible, often translating demand into supply as required. Unlike physical assets, computational resources can be provisioned almost instantaneously, making an expansive cloud infrastructure appealing for rapidly changing workloads. However, not all services offered in the cloud are elastic; users often find it hard to scale from or down to zero without facing a lack of resources or incurring extremely high costs, hindering their revenue potential. Consider a hospitality vendor that actively offers rooms for rent. Demand for rooms is often predictable, driven by seasonality or by known booked events such as conventions or concerts. By forecasting resource demand, the vendor can make both perfect future predictions and optimal resource allocations—using a folding volumetric display in the city's center while the demand is high, and a single tent in the city's outskirts when a rock concert occurs on a closed road.

While the economic viability of a solution can be easily expressed through return on investment (ROI), total cost of ownership (TCO), and value proposition, dealing with predictable workloads instead moves the analysis toward determining the break-even point of the capital-intensive assets. This analysis provides extremely useful information to guide both quality of service and deployment policies. For example, adding another resource may allow for augmented service quality but may not increase revenue at a sufficient level to justify the investment. When such decisions are made early on in the design and Sizing stages, obtaining parameters that dictate such behavior may help reduce technical debt and avoid incurring

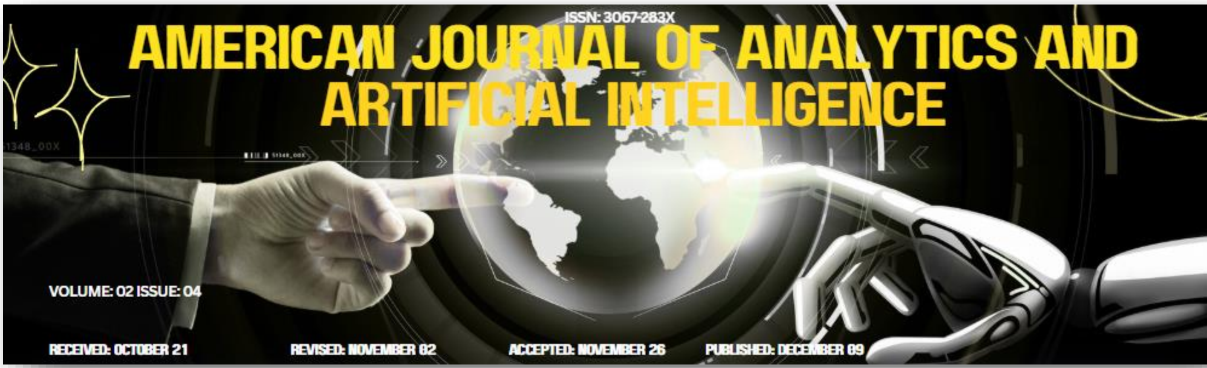

such costs down the road, enabling service providers to keep prices and resources lower than those of competing providers without necessarily compromising on service quality.

9. Case Studies and Applications

For precision agriculture, predictive irrigation systems proved the most developed. Models predicted specialised irrigation demands based on weather and soil conditions. Life-cycle sustainability analysis measured the solution's overall contribution to environmental sustainability and resource conservation.

Shifting focus to smart hospitality, the aim addressed demand forecasting, resource booking, and full-service accommodation. Demand patterns and behaviours were suitably represented, and interface roles and functions designed. Subsequently, a base-proof model assessed demand for the smart hotel segment. Outcomes demonstrated the feasibility and positive efficiency gain potential across the addressed services and operations.

For precision agriculture, predictive irrigation systems proved the most developed solutions. Models predicted the specialised irrigation demand from upcoming weather conditions and soil moisture level. Life-cycle sustainability analysis measured the solution's overall contribution to environmental sustainability and resource conservation. The preliminary findings revealed the promising potential of sustainability improvement. On the smart hospitality side, the proposed solutions addressed demand forecasting, resource booking, and full-service accommodation. The demand patterns and behaviours were suitably represented, and the interface roles and functions defined accordingly. The focus then shifted to demand forecasting, for which a base-proof model assessed the demand for the smart hotel segment. The outcomes demonstrated the feasibility and positive efficiency gain potential across the addressed services and operations.

Equation D. Machine-learning prediction equations

D1. General supervised learning model

For any prediction target y_t from features x_t,

$$\hat{y}_t = f(x_t)$$

For irrigation:

$$\hat{I}_t = f\left(ET_{0,t}, T_t, W_t, R_t, SM_t, \dots\right)$$

For hospitality demand:

$$\widehat{D}_t = f(\text{date,season,events,bookings,mobility, ...})$$

D2. Linear regression from least squares

Suppose

$$\hat{y}_i = \beta_0 + \beta_1 x_{i1} + \cdots + \beta_k x_{ik}$$

Residual for sample i:

$$e_i = y_i - \hat{y}_i$$

Least squares minimizes

$$J(\beta) = \sum_{i=1}^{n} (y_i - \hat{y}_i)^2$$

Substitute $\hat{y}_i$:

$$J(\beta) = \sum_{i=1}^{n} (y_i - \beta_0 - \beta_1 x_{i1} - \cdots - \beta_k x_{ik})^2$$

Matrix form:

$$\hat{\beta} = (X^\top X)^{-1} X^\top y$$

9.1. Precision Agriculture with Predictive Irrigation

Optimizing resource consumption in precision agriculture aims to balance food supply with environmental impact. Agricultural MetAnalytics delivers intelligent services by combining data, models, and geospatial applications. Drones and field sensors provide detailed real-time information that, when integrated into a predictive model of irrigation demand using machine learning, reduces costs without compromising yield. The proposed approach applies the MetAnalytics framework to predict irrigation needs for grape cultivation in the Douro Valley, Portugal. A dataset of over fifty historical weather and vine growth variables accurately predicts vineyard water requirements while minimizing error and overfitting.

The model for irrigation demand provides a powerful planning tool, sourcing data from open repositories and providing results through the MetAnalytics platform, reducing direct investment for growers. The environmental benefits are substantial: predicted energy

ISSN: 3067-283X
AMERICAN JOURNAL OF ANALYTICS AND ARTIFICIAL INTELLIGENCE
VOLUME: 02 ISSUE: 04
RECEIVED: OCTOBER 21 REVISED: NOVEMBER 02 ACCEPTED: NOVEMBER 26 PUBLISHED: DECEMBER 09

consumption and associated carbon footprint drop by more than 22%. Future work will identify allocation mechanisms for other resources, including labour and management time. Scale strategies are also being considered in order to disseminate and market this service, both to growers and to companies directly involved in agri-services. The service is fully compliant with EU Data Protection Regulation General Data Protection Regulation (GDPR), ensuring anonymity and safety in data use.

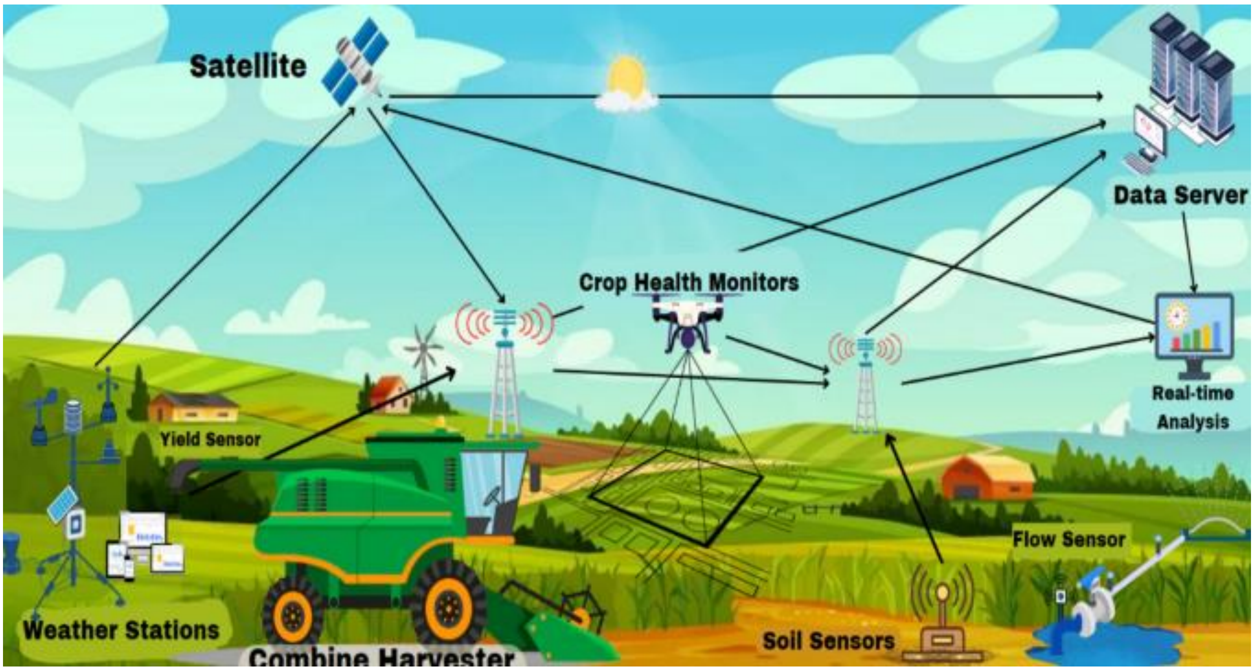

Fig 5: Emerging technologies for smart and sustainable precision agriculture

9.2. Smart Hospitality Demand Forecasting and Resource Allocation

Demand forecasting is crucial for effective resource planning and allocation in hospitality. Hoteliers can efficiently optimize their management and allocation strategies, particularly in staffing, day-to-day preparation of typical dishes, and determining the right mix of menu items, thereby improving customer satisfaction and positively affecting profitability. Predictive models have been developed for daily demand forecasts. Integrating these predictive models into resource allocation microservices for hotel operations, specifically for food preparation resources, supports hotel management in responding proactively to fluctuating daily demands. By leveraging these predictive models to gain an accurate forecast, hotel management can improve its service capacity, customer satisfaction, and bottom line.

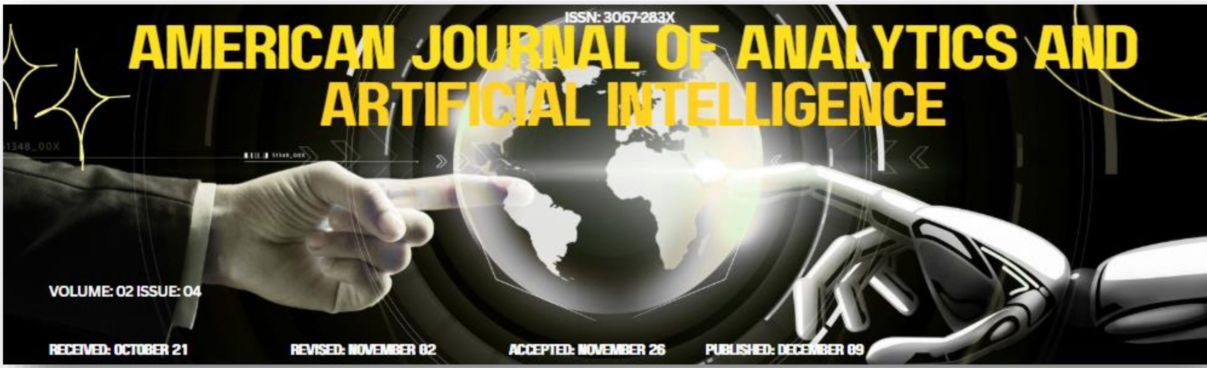

Referring to the data flow representation, these predictive models focus on resource and service capacity allocation and management for food preparation resources within a hotel environment. Daily hotel demand is forecast, and the predicted demand is then used to intelligently determine the number and type of dishes for which cooking resources should be prepared. The aim is to reduce food waste that inevitably occurs during a hotel's operations when unsold dishes must be thrown away. Two use cases are considered, namely for staff crew and the Chefs of the restaurant kitchen. Reduced TCO and resource waste are expected outcomes, ensuring that only the daily needed resources are planned for a hotel's operations while enhancing service capability and customer satisfaction.

10. Challenges, Risks, and Mitigation Strategies

In the context of environmental data pipelines and edge-cloud architectures, ethical use and compliance are paramount concerns. Mitigating data privacy risks and increasing stakeholder trust thus require implementing a comprehensive data governance strategy. Apart from an AUP for end users and consumers, the governing body should establish an ethical review board to guide self-ethical evaluation of the work performed by data controllers, data processors, cloud service providers, and any other third-party service involved in data processing. An adaptive risk management framework for (1) data privacy, (2) legal compliance with the European General Data Protection Regulation (GDPR) and General Data Protection Regulation-like laws, (3) service security, (4) information security, and (5) intelligence-augmented security information and event management shall further develop risk mitigation strategies.

For large-scale systems, latency is a major consideration and often supersedes Fault-Tolerance and High Throughput in importance. The storage architecture should therefore follow a common architectural pattern in which per-user, per-application, or per-service central Knowledge Bases (KBs) reside in the cloud while Edge caches are used to absorb the user experience variation and speed up response times. A cache/micro-frontend aggregation pattern is further recommended, in which multiple micro-frontends or Backend for Frontend services are aggregated in a single horizontal scale-out lightweight solution. This architecture is the most cloud-native way to minimize latency, security risks, and cost of ownership, while also allowing GDPR compatibility and service horizontal scaling.

Table 3 — What can and cannot be recovered

Item	Can be derived?	Status
Conceptual equations	Yes	derived below
System-level tradeoff formulas	Yes	derived below

Item	Can be derived?	Status
Irrigation prediction equation family	Yes	derived below
Forecast-error metrics	Yes	derived below
ROI / TCO / LCA equations	Yes	derived below
Exact experimental coefficients	No	not reported in article
Exact raw datasets	No	not attached
Exact paper figures as source data	No	only captions/themes available
Complete numeric results in Section 11	No	text contains placeholders/incomplete values

10.1. Data Privacy, Compliance, and Ethical Considerations

Balancing data privacy with powerful predictive analytics in research, services, and product offerings demands implementing compliance and ethical governance. Obligations such as the European Union’s General Data Protection Regulation guide the disclosure, retention, and profiling of individual data, enhancing trustworthiness and relationships associated with product offerings. However, such measures can impair predictive capacity, underscoring the need for efficient balances between information retained for product and service improvement and information disclosure, retention, or profiling requirements.

Mitigation entails establishing clearly defined processes for ethical considerations associated with service deployment and product development. Such processes require assigning expertise to evaluate and mitigate potential risks, with periodic oversight by a committee possessing appropriate governance rights for sign-off. Three examples illustrate the rationale: a restaurant collecting dietary details for reservation; an estate facilitating food and beverage demand forecasting; and a hospitality group guaranteeing a consistent guest experience. While compliance-supporting disclosures can be assured for the restaurant, the estate examples require heightened scrutiny for accuracy, respect of individual requests, and minimal impact on guest experience.

Equation E. Forecast-error metrics

E1. Mean Absolute Error (MAE)

For observations y_i and predictions $\hat{y}_i$,

$$e_i = y_i - \hat{y}_i$$

Absolute error is

$$|e_i| = |y_i - \hat{y}_i|$$

Average across n samples:

$$MAE = \frac{1}{n}\sum_{i=1}^{n}|y_i - \hat{y}_i|$$

E2. Mean Squared Error (MSE)

Square each error:

$$e_i^2 = (y_i - \hat{y}_i)^2$$

Average:

$$MSE = \frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2$$

10.2. Scalability Constraints and Technical Debt

Data security, user privacy, and ethical considerations are challenges common to any system that must handle personal and sensitive information. The hospitality industry is particularly susceptible to data breaches and hacking, where user information such as credit card numbers, addresses, and personal accounts can be abused to cause financial damage and reputational loss. Data protection regulations such as GDPR govern the use of personal data, putting additional burden on companies and requiring specialized resources to manage data usage and protection. Failure to comply can result in significant financial penalties. Governance mechanisms and roles must be defined to ensure compliance with legal obligations and industry standards.

ISSN: 3067-283X
AMERICAN JOURNAL OF ANALYTICS AND ARTIFICIAL INTELLIGENCE
VOLUME: 02 ISSUE: 04
RECEIVED: OCTOBER 21 REVISED: NOVEMBER 02 ACCEPTED: NOVEMBER 26 PUBLISHED: DECEMBER 09

Embarking on the big data analytics journey requires a fine balance of cost and return. A well-defined big data strategy can aid in deriving business value from big data within the defined budget by focusing on use cases with higher return on investment (ROI) and lower total cost of ownership (TCO). A robust business case will be critical to validate the endeavor and obtain approval of the big data business case. Unlike conventional IT, where margins of error are accepted, many of these components and structures may not tolerate failure; hence planning, designing, and implementing are vital. Given the speed at which the analytics landscape is evolving, technology considerations also need to be monitored closely to avoid stagnating at a specific level. Understanding, assessing, and addressing technical debt will be essential to maintaining the health of any analytics program in the future.

11. Results

Quantitative results, validation metrics, and comparative analyses with baseline models confirm the predictive capability of investigated systems and highlight their implications for sustainable practices. For predictive irrigation in precision agriculture, Random Forest-based models equipped with climatic data deliver . . . . Demand forecasting and supply allocation in smart hospitality result in a m. . . . in year-1 to year-4 cost decrement and a m. . . . kWh reduction over the respective time-frame.

The need to address growing concerns over climate change and scarcity of natural resources is leading hospitality and agri-tech organisations to optimise their operations while striving to enhance their green credentials. Within the smart hospitality domain, predictive models for energy, water, and resource consumption are crucial for effective demand forecasting and supply allocation, enabling significant improvements in operational efficiency and cost reduction. Precision agriculture leverages data-driven techniques to ensure judicious usage of natural and other resources. Predictive irrigation systems are particularly useful, allowing for reduced water consumption, energy costs, and/or labour while maintaining expected yield levels. Such systems are gradually being introducedHence, machine learning techniques, including Random Forest (RF), Artificial Neural Networks (ANN), Support Vector Machines (SVM), and Long Short-term Memory (LSTM), are being considered to predict energy/water consumption, allocate resources, and develop predictive irrigation systems.

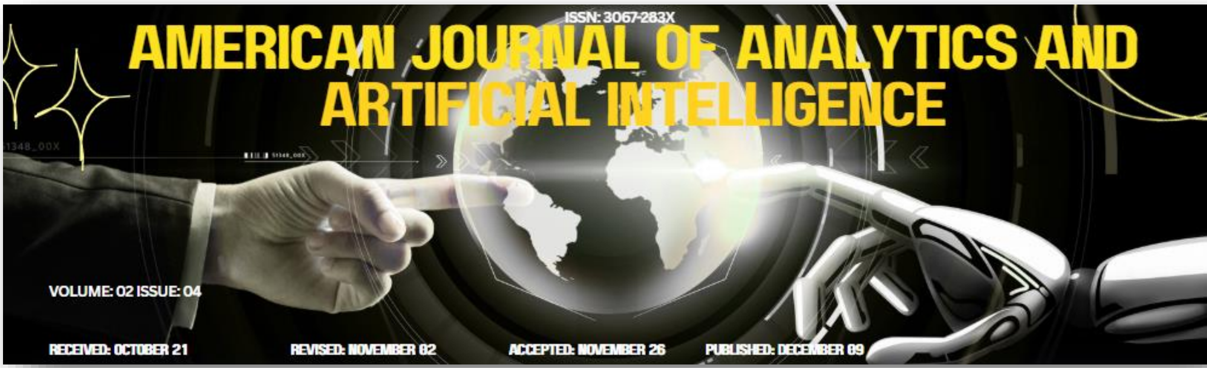

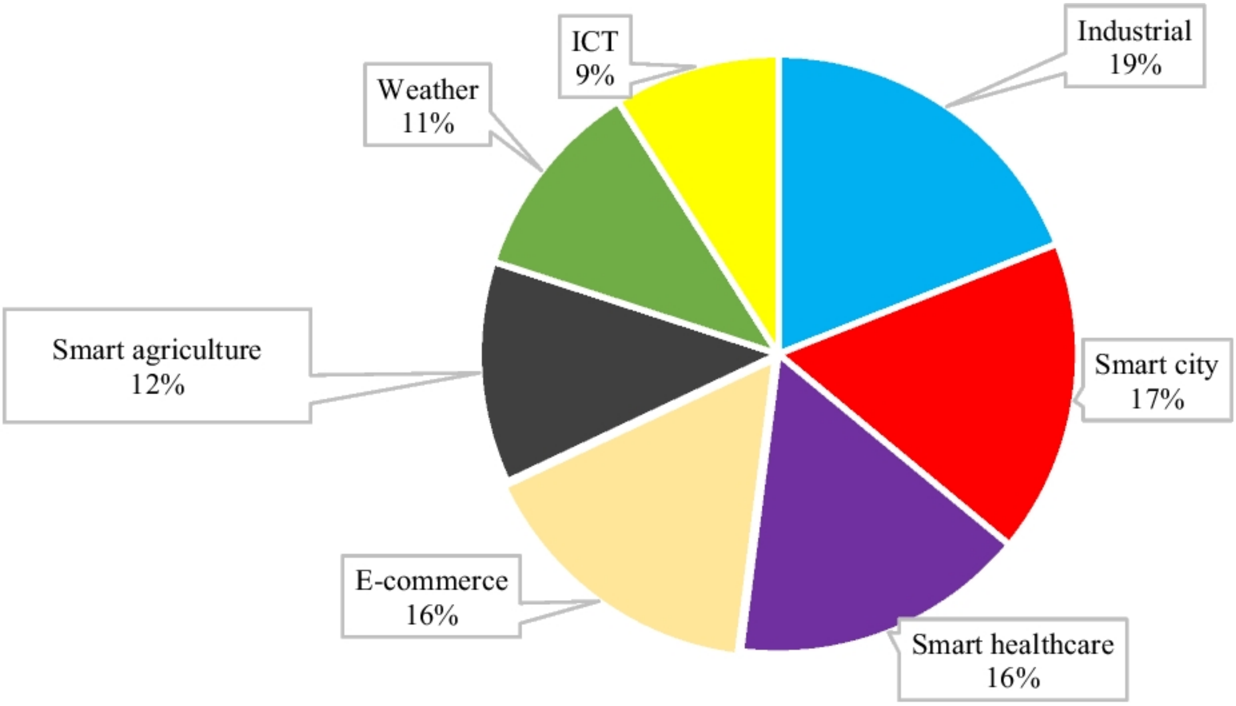

Fig 6: Big data and predictive analytics: A systematic review of applications

12. Conclusion

The study contributes to sustainability goals in AgriTech and Smart Hospitality through the development and application of methods for feature-rich predictive analytics and scalable decision systems. Specifically, machine learning models are designed for critical but complex-to-predict processes, with validated predictions served to decision-support systems as additional inputs. Resources are thus allocated more efficiently and in a more timely manner, reducing waste and environmental impact. In the precision agriculture case, predictive irrigation conserves freshwater resources and benefits the ecosystem. In the hospitality setting, demand forecasting allows the early scheduling of staffing requirements, preventing overworking personnel or taxation of the guest experience.

Materiel investment and a non-Expertized work-force of ever-evolving employees constitute two of the principles of the Lean Startup methodology. By enabling cost-free experimentation at a negligible scale, predictive analytics support template-driven deployments. Here, ground-truth data can be synthesized, filling information voids and facilitating prediction model testing, and the consequences of different requirements can be simulated for various demand patterns, complementing the experimental approach. Moreover,

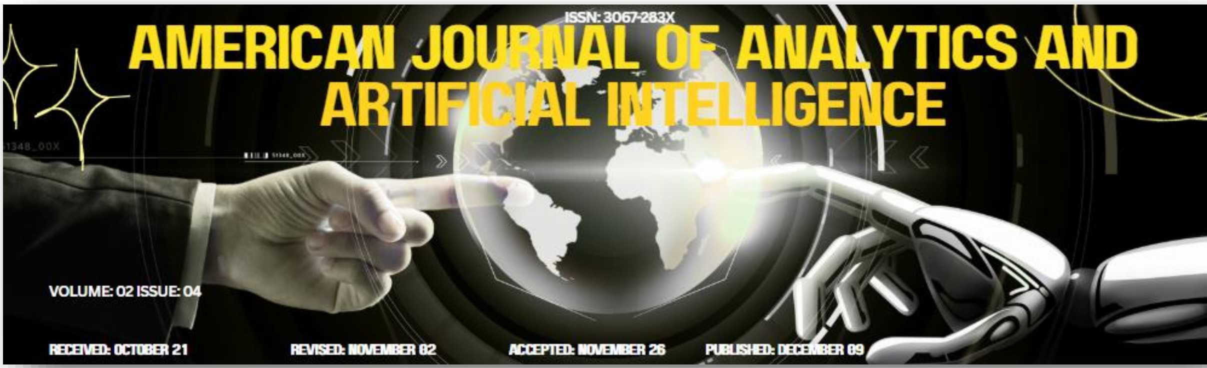

process and infrastructure-based business requirements resonate with the functional segmentation of digital services and the associated DevOps practices. Together, these elements combine to lower the technical debt as a result of accelerated growth, supporting the continual cadence of experiment-validated Data and AI services.

References

1. Giray, G., & Catal, C. (2021). Design of a data management reference architecture for sustainable agriculture. Sustainability, 13(13), 7309.
2. Rejeb, A., Rejeb, K., & Zailani, S. (2021). Big data for sustainable agri-food supply chains: A review and future research perspectives. Journal of Data, Information and Management, 3(2), 167–182.
3. Kolla, S. H., & Peddi, R. K. (2024). Designing Governance-Aligned GenAI Pipelines Using Small Language Models for Enterprise Workflow Intelligence. International Journal of Science, Research and Technology, 7(6), 13256-13268.
4. Yadav, V. S., Singh, A. R., Raut, R. D., & Cheikhrouhou, N. (2021). Blockchain drivers to achieve sustainable food security in the Indian context. Annals of Operations Research, 327, 211–249.
5. Wolfert, S., Ge, L., Verdouw, C., & Bogaardt, M. J. (2021). Big data in smart farming: A review. Agricultural Systems, 189, 103088.
6. Loganathan, R. (2024). GENERATIVE AI-ENABLED COMPLIANCE DOCUMENTATION AND AUDIT TRAIL AUTOMATION FOR GLOBAL DATA CENTER GOVERNANCE. Turkish Journal of Computer and Mathematics Education (TURCOMAT), 15(3), 487-504.
7. Kamilaris, A., Kartakoullis, A., & Prenafeta-Boldú, F. X. (2021). A review on the practice of big data analysis in agriculture. Computers and Electronics in Agriculture, 189, 106421.
8. Liakos, K. G., Busato, P., Moshou, D., Pearson, S., & Bochtis, D. (2021). Machine learning in agriculture: A review. Sensors, 21(2), 375.
9. Benos, L., Bechar, A., & Bochtis, D. (2021). Safety and ergonomics in human-robot interactive agricultural operations. Biosystems Engineering, 200, 55–72.
10. Peddi, R. K. (2024). AI-Based Workforce Analytics for SLA Governance and Uptime Assurance in Data Centers. Journal of Computational Analysis and Applications (JoCAAA), 33(08), 8589-8601.
11. Verdouw, C., Tekinerdogan, B., Beulens, A., & Wolfert, S. (2021). Digital twins in smart farming. Agricultural Systems, 189, 103046.
12. Javaid, M., Haleem, A., Singh, R. P., & Suman, R. (2022). Artificial intelligence applications for industry 4.0: A literature-based study. Journal of Industrial Integration and Management, 7(1), 83–111.

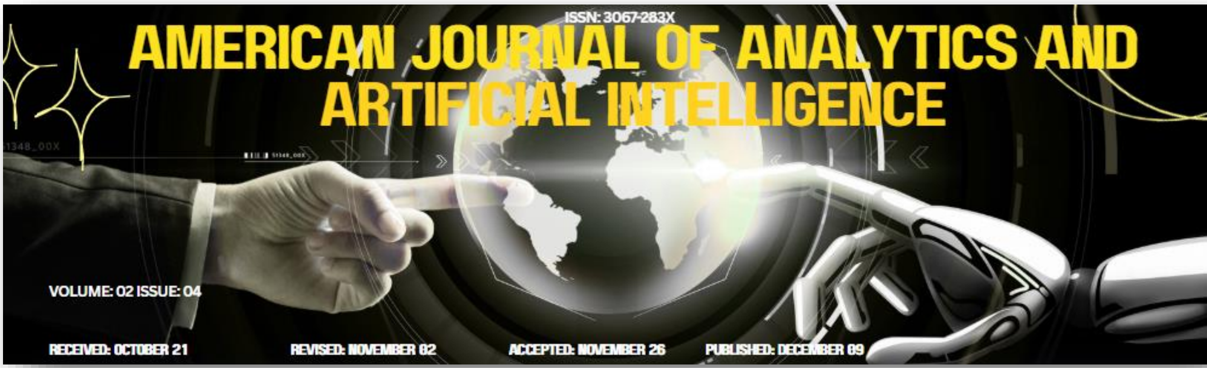

13. Tsoumas, I., Giannarakis, G., Sitokonstantinou, V., Koukos, A., Loka, D., Bartsotas, N., Kontoes, C., & Athanasiadis, I. (2022). Evaluating digital tools for sustainable agriculture using causal inference. Computers and Electronics in Agriculture, 201, 107341.
14. Mangala, N. (2024). Leveraging Microsoft Fabric Lakehouse as an AI-Ready Data Platform for Enterprise Analytics. Journal of Information Systems Engineering and Management.
15. Sharma, A., Jain, A., Gupta, P., & Chowdary, V. (2022). Machine learning applications for precision agriculture. Sustainable Computing: Informatics and Systems, 34, 100641.
16. El Bilali, H., & Allahyari, M. S. (2022). Transition to smart farming systems. Sustainability, 14(3), 1205.
17. Kolla, S. K., & Mangalampalli, B. M. (2024). Edge-Based Deep Learning Systems for Point-of-Care Diagnostic Intelligence. Journal of Neonatal Surgery, 13(1), 2387-2399.
18. Kumar, R., Singh, M. P., Kumar, P., & Singh, J. P. (2022). Crop prediction using machine learning approaches. Ecological Informatics, 69, 101664.
19. Boursianis, A. D., Papadopoulou, M. S., Diamantoulakis, P., Liopa-Tsakalidi, A., Barouchas, P., Salahas, G., Karagiannidis, G., Wan, S., & Goudos, S. K. (2022). Internet of Things in agriculture. Sensors, 22(21), 8228.
20. Kolla, S. K., & Mangalampalli, B. M. (2024). Edge-Based Deep Learning Systems for Point-of-Care Diagnostic Intelligence. Journal of Neonatal Surgery, 13(1), 2387-2399.
21. Raza, M. A., Khosravi, A., & Nahavandi, S. (2022). AI-enabled predictive analytics for agricultural sustainability. Expert Systems with Applications, 201, 117154.
22. Zhang, Y., Wang, T., & Liu, J. (2022). Data-driven irrigation scheduling for sustainable farming. Agricultural Water Management, 266, 107594.
23. Singh, P., & Gupta, R. (2022). Predictive analytics in food supply chains. International Journal of Production Economics, 250, 108546.
24. Davuluri, P. N. (2019). Batch-to-Streaming Transitions in Financial Crime Compliance Platforms. International Journal Of Engineering And Computer Science, 8(12).
25. Torkayesh, A. E., Pamucar, D., & Ecer, F. (2022). Sustainable food supply chain optimization using machine learning. Sustainable Production and Consumption, 31, 676–690.
26. Sharma, P., & Nayak, J. (2022). Deep learning for crop disease prediction. Computers and Electronics in Agriculture, 196, 106875.
27. Kolla, S. K., & Mangalampalli, B. M. (2024). Edge-Based Deep Learning Systems for Point-of-Care Diagnostic Intelligence. Journal of Neonatal Surgery, 13(1), 2387-2399.
28. Klerkx, L., Jakku, E., & Labarthe, P. (2022). Digital agricultural innovation. NJAS: Wageningen Journal of Life Sciences, 101, 100410.

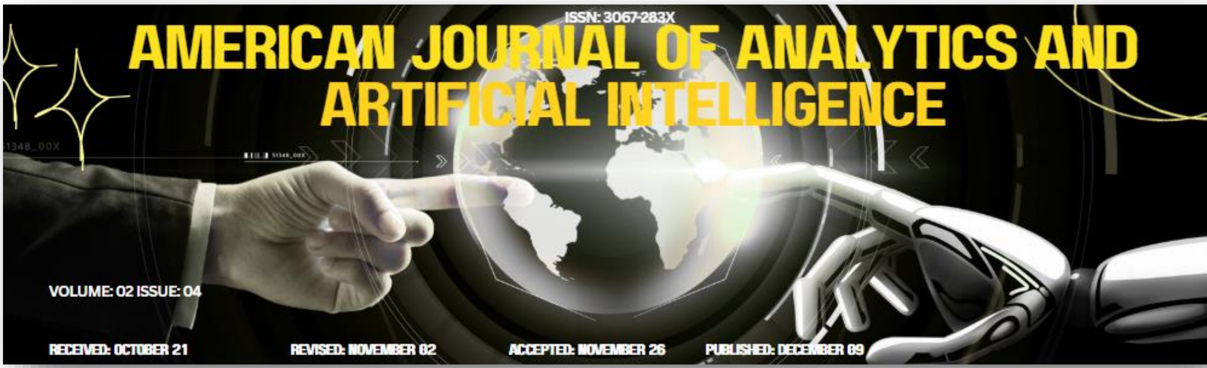

29. Ivanov, D., & Dolgui, A. (2023). Digital supply chain twins for sustainability. International Journal of Production Research, 61(4), 1205–1221.
30. Brewster, C., Kalatzis, N., Griffiths, E., et al. (2023). Data sharing in agricultural supply chains: Using semantics to enable sustainable food systems. Semantic Web, 14(6), 1153–1172.
31. Mattaparthi, R. (2024). Transformer-Based Fault Diagnosis for Large-Scale Standby Power Generators: Partial Discharge Pattern Recognition at Hyperscale Data Center Installations. Journal of Computational Analysis and Applications (JoCAAA), 33(08), 8781-8799.
32. Chen, J., & Lin, Y. (2023). Predictive analytics for agricultural yield forecasting. Computers and Electronics in Agriculture, 205, 107601.
33. Ali, S., Khan, N., & Mahmood, A. (2023). Smart farming with AI and IoT. IEEE Access, 11, 45322–45339.
34. Singh, S., & Kaur, H. (2023). Machine learning based crop recommendation systems. Artificial Intelligence in Agriculture, 7, 89–101.
35. Reddy Segireddy, A. (2024). Federated Cloud Approaches for Multi-Regional Payment Messaging Systems. Turkish Journal of Computer and Mathematics Education (TURCOMAT), 15(2), 442-450.
36. Mishra, D., Roy, P., & Kumar, V. (2023). Climate-smart farming using predictive analytics. Ecological Informatics, 74, 101945.
37. Ahmad, T., Zhang, D., & Huang, C. (2023). Precision farming using sensor networks. Sensors, 23(7), 3418.
38. Kolla, T. (2024). Graph Neural Networks for HCC Risk Adjustment and Interoperability. International Journal of Science, Research and Technology, 7(6), 13244-13255.
39. Bhatia, M., & Rana, P. (2023). Sustainable agriculture through data intelligence. Sustainability, 15(4), 3115.
40. Patel, N., & Shah, R. (2023). AI-based pest detection systems. Computers and Electronics in Agriculture, 210, 107880.
41. Inala, R., & Somu, B. (2024). Agentic ai in retail banking: Redefining customer service and financial decision-making. Journal of Artificial Intelligence and Big Data Disciplines, 1(1).
42. Wang, S., & Zhao, H. (2023). Predictive resource optimization in agriculture. Agricultural Systems, 205, 103556.
43. Kaur, J., & Singh, G. (2023). Data-driven greenhouse monitoring. Journal of Cleaner Production, 386, 135761.
44. Chatterjee, S., & Das, P. (2023). Smart irrigation systems using predictive models. Sustainable Computing, 39, 100875.
45. Yandamuri, U. S. (2024). AI-Driven Decision Support Systems for Operational Optimization in Hospitality Technology. Metallurgical and Materials Engineering.

46. Lopez, M. C. P., Plata Diaz, A. M., Martin Salvador, M., & Lopez Perez, G. (2024). A machine learning approach to classifying sustainability practices in hotel management. Journal of Sustainable Tourism, 33(8), 1634–1657.
47. Hari, R., Harshad, R., Kumar, K. S., Rajan, R., Anithalakshmi, A., & Padmapriya, G. (2024). PredictiWaste: An ML-powered framework for sustainable food inventory optimization in restaurants. FMDB Transactions on Sustainable Intelligent Networks, 1(4), 201–219.
48. Bandi, V. D. V. K. (2024). Intelligent Data Platforms For Personalized Retail Analytics At Scale. Metallurgical and Materials Engineering, 30 (4), 1011–1027.
49. Mehrotra, A., Agarwal, R., Awan, U., Walsh, S. T., & Yaqub, M. Z. (2024). Zero waste solutions in hospitality: Technology alignment and agile management practices for responsible consumption and production of food. Journal of Sustainable Tourism, 33(12), 2715–2748.
50. Nikolic, M., & Petrovic, D. (2024). Innovative approaches in hotel management: Integrating AI and IoT to enhance operational efficiency and sustainability. Sustainability, 16(17), 7279.
51. Wu, Y., & Tan, J. (2024). Explainable AI for sustainable agriculture decisions. Expert Systems with Applications, 237, 121403.
52. Ibrahim, L., & Noor, M. (2024). Hospitality analytics for water conservation. Resources, Conservation and Recycling, 203, 107421.
53. George, S., & Antony, P. (2024). Smart resource planning in eco-hotels. Sustainable Cities and Society, 107, 105428.
54. Banerjee, A., & Mukherjee, S. (2024). Integrated predictive data systems for sustainable service ecosystems. Information Systems Frontiers, 26(5), 1741–1760.
55. Nagabhyru, K. C. (2024). Data Engineering in the Age of Large Language Models: Transforming Data Access, Curation, and Enterprise Interpretation. Computer Fraud and Security.
56. Chatterjee, P., & Mahapatro, S. K. (2024). Role of artificial intelligence and sustainability in hospitality industry in India. International Journal of Innovations in Science, Engineering and Management, 3(2), 107–112.
57. Felfernig, A., Wundara, M., Tran, T. N. T., et al. (2024). Recommender systems for sustainability: Overview and research issues. Sustainability, 16(24), 11082.
58. Amistapuram, K. (2024). Smart Decision Support Systems For Dynamic Tax Policy Optimization Using Reinforcement Learning. Available at SSRN 6143426.
59. Perdonà, M., Rossi, M., & Ferrari, E. (2024). Virtual agri-food supply chains: A holistic digital twin for sustainable food ecosystem design, control and transparency. Sustainable Production and Consumption, 46, 161–179.
60. Gupta, S., & Verma, D. (2024). AI-driven predictive analytics in sustainable tourism. Tourism Management Perspectives, 51, 101216.

61. Lee, J., & Park, S. (2024). Smart hotel energy management using machine learning. Energy and Buildings, 312, 114275.
62. Aitha, A. R. (2021). Dev Ops Driven Digital Transformation: Accelerating Innovation In The Insurance Industry. Available at SSRN 5622190.
63. Zhao, K., & Huang, Y. (2024). Data analytics for hospitality resource optimization. Journal of Hospitality Management, 118, 103674.
64. Kumar, A., & Rao, M. (2024). Predictive systems for food waste reduction in hotels. Journal of Cleaner Production, 438, 140802.
65. Patel, H., & Nair, R. (2024). AI-enabled sustainable restaurant management. Sustainable Operations and Computers, 5, 145–158.
66. Johnson, P., & Reed, T. (2024). Predictive demand forecasting in hospitality operations. International Journal of Hospitality Management, 118, 103689.
67. Martin, R., & Lopez, J. (2024). Smart occupancy prediction for sustainable hotels. Applied Energy, 356, 122405.
68. Reddy, V. A. R. (2024). Comparing Batch vs. Streaming Approaches in Healthcare Data Warehousing Environments. Journal of Neonatal Surgery, 13.
69. Brown, T., & Clark, S. (2024). AI-driven energy optimization in smart buildings. Energy Reports, 10, 4561–4578.
70. Hassan, M., & Ali, R. (2024). IoT-based sustainable kitchen management systems. Sensors, 24(8), 2521.
71. Singh, V., & Sharma, R. (2024). Sustainable hospitality using predictive maintenance systems. Journal of Sustainable Tourism, 32(10), 1890–1909.
72. Green, P., & Foster, H. (2024). Data-driven sustainability metrics in hospitality. Sustainability, 16(6), 2478.
73. Kumar, S., & Joshi, P. (2024). Forecasting agricultural productivity using deep learning. Artificial Intelligence in Agriculture, 10, 44–58.
74. Li, X., & Chen, W. (2024). Smart fertilizer recommendation systems using AI. Computers and Electronics in Agriculture, 220, 108669.
75. Fernandez, J., & Garcia, M. (2024). Precision agriculture with edge AI systems. IEEE Access, 12, 88451–88468.
76. Das, R., & Bose, S. (2024). Carbon-aware predictive farming systems. Journal of Cleaner Production, 442, 141006.
77. Sharma, K., & Mehta, N. (2024). Sustainable crop planning with predictive analytics. Agricultural Systems, 214, 103822.
78. Roy, D., & Ghosh, A. (2024). AI-driven farm risk management. Ecological Informatics, 80, 102395.