International Journal of Advanced Engineering Science and Information Technology (IJAESIT)

|ISSN 2349-3216| Volume 7, Issue 4, July-August 2024| Bimonthly, Peer Reviewed and Scholarly Indexed Journal|

DOI: 10.15662/IJAESIT.2024.0704008

Azure-Based Adaptive Predictive Analytics for Dynamic Business Intelligence in AgriTech Systems

Nareddy Abhireddy

Independent Researcher, USA

nareddy.abhireddy.researcher@gmail.com

Dr. Emily Carter

Research Assistant, Northbridge International University, United Kingdom

emilycartteerr@outlook.com

ABSTRACT: Adaptive Predictive Analytics Pipelines on Microsoft Azure: A Machine Learning Approach for Dynamic Business Intelligence in AgriTech Systems

Business intelligence systems, including predictive analytics repositories, are essential for effective decision support in agriculture and agri-tech. Yet, many existing solutions are unable to adapt to changing conditions. Recent advances in machine learning (ML) provide opportunities to alleviate some of the associated problems. On the Microsoft Azure platform, a set of services built around data lakes and ML models enables the implementation of adaptive predictive analytics pipelines—data-processing assemblies that optimize performance by continuously learning from available data. These pipelines can react to changes, such as the emergence of new patterns, by automatically altering feature engineering processes, retraining models using updated sources, or providing predictions with new data representations, characteristics, or horizons.

The architecture and implementation of these adaptive predictive analytics pipelines are described in detail, with descriptions of the pipelines' general structure and operation, along with a case study focused on the Agri-Tech Risk Mitigation Solutions (AT-RMS) predictive analytics service. AT-RMS assists executive-level actors in the agriculture and insurance sectors with strategic-level decision-making, and the implementation leverages Azure to support infrastructure-as-code and operational governance.

KEYWORDS: Adaptive analytics, agri-tech, business-intelligence, cloud-based machine-learning, data-lake, Microsoft-Azure, predictive-analytics, risk-assessment, scalable-architecture, supporting-systems

I. INTRODUCTION

Business intelligence (BI) usually requires data collection and preparation for a finite period leading up to specific reporting time frames. Adaptive predictive analytics pipelines directly update models and can thus provide fresh predictions anytime. Dynamic predictive analytics built around such pipelines rapidly integrate new data and regularly retrain the models. When applied to machine-learning problems facilitating decision support in agriculture, these pipelines allow business intelligence to expand from static historical reporting to dynamic risk management, planning, and governance. A case study encompassing a relevant sample of such predictive problems illustrates the application of dynamic predictive analytics on Microsoft Azure.

The architecture of adaptive predictive analytics pipelines ensures automatic, up-to-date predictions for machine-learning tasks associated with areas such as forecasting near-future variable values using data from the past, estimating targets conditioned on other variables, assessing likely adverse outcomes, or providing elaborate decision-support information. These pipelines integrate a series of data-management processes on Microsoft Azure that monitor the quality of ingested data, consolidate predictions from external services, refine training and validation data sets, and enable such analyses to be executed seamlessly. Such adaptations require maintenance and supervision. Adaptation pipelines monitor performance using external benchmark data, assess model drift through time, execute retraining when necessary, ensure the reusability of new models, and maintain knowledge of previous versions of each model within a multiworkspace setup.

1.1. Problem Statement and Relevance: Aggregated long-term forecasts generated by traditional Business Intelligence solutions often fail to provide the timely, relevant insights needed by agricultural operators. Weather and climate projections may evolve significantly during the fulfilment of crops and business opportunities develop dynamically

|ISSN 2349-3216| Volume 7, Issue 4, July-August 2024| Bimonthly, Peer Reviewed and Scholarly Indexed Journal|
DOI: 10.15662/IJAESIT.2024.0704008

according to market supply and demand. As a consequence, more frequent predictive analytics and forecasting services—reflecting data from a wider range of sources and targeting shorter time horizons—are desirable. Pipelines that generate such predictions, however, are still seldom found in practice. Sources of predictive datasets may not only differ on a day-to-day basis but also across spatial areas and occur in various data modalities (e.g., text, images), which creates a need for pre- and postprocessing code and guides that support the use of predictions in BI systems.

With technology enabling the collection and storage of ever more data, the ability to easily implement and deploy data-driven models that support BI processes across a given business domain is desirable. In particular, on-demand model training and inference are especially valued in domains such as political elections and sports competitions, for which accurate forecasting remain important within a clearly defined period. Dynamic BI refers to fulfilling that aspiration when training and inference are performed more frequently than typically found in practice. More specifically, the term denotes a business-driven approach to the generation of data-driven predictions in which the standard workflow supporting BI activities is supported by—rather than takes precedence over—an analytics ecosystem. The fulfilment of these requirements in the context of adaptive pipelines that can ingest, preprocess, train, and apply predictive models in a time-sensitive manner and within a realistic operating environment provides their primary justification. Such pipelines are envisaged within a dedicated ecosystem deployed on Microsoft Azure.

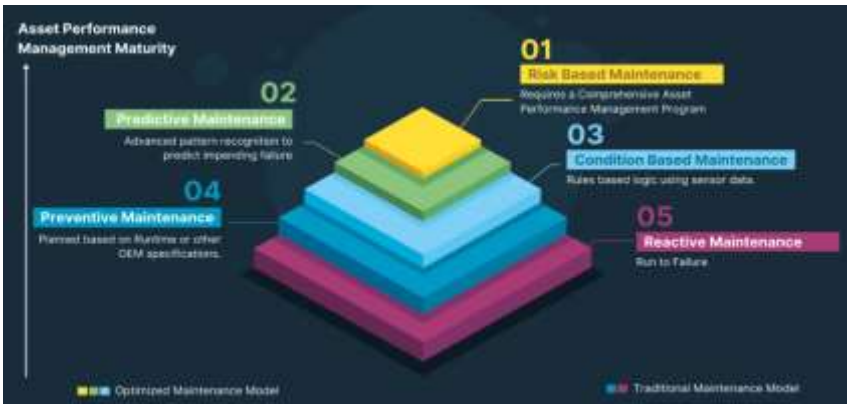


Fig 1: Predictive Maintenance with Azure ML

1.2. Scope and Objectives: Adaptive Predictive Analytics Pipelines on Microsoft Azure: A Machine Learning Approach for Dynamic Business Intelligence in AgriTech Systems

Delineating the bounds of an investigation into business intelligence in AgriTech requires consideration of the data handling capabilities it entails. The elusive stable feature set for predictive models, along with the sporadic collection of information in many domains, suggests that model training be carried out on an ad hoc basis. This is made possible by deploying appropriate machine-learning capabilities in the cloud and that decision making make use of services provided by a cloud Azure-based BI platform. Pipelines that provide such capabilities are termed Adaptive Predictive Analytics Pipelines.

The system aims to provide dynamic solutions for agriculture and horticulture by making predictions on an appropriate time scale to meet the present needs of primary producers and related stakeholders. Dynamic data for predictive analysis pipelines are achieved within specific bounds: the feature space is distinct at each time of prediction, but data of the set of features needed for prediction are also distinct at any instant. Thus, the information available is capable of lending support to policymakers to take an informed decision on the economic status of the country, including the status of exports or imports of agricultural products. The pipeline of the whole system ensures such predictions automatically without human intervention and can be monitored by humans.

II. BACKGROUND AND RELATED WORK

Adaptive Predictive Analytics Pipelines on Microsoft Azure: A Machine Learning Approach for Dynamic Business Intelligence in AgriTech Systems

The presented approach builds on established theories of predictive analytics and agile BI, and addresses their integration in a system architecture that manages data quality, model retraining, and adaptation of feature representation and data

|ISSN 2349-3216| Volume 7, Issue 4, July-August 2024| Bimonthly, Peer Reviewed and Scholarly Indexed Journal|

DOI: 10.15662/IJAESIT.2024.0704008

freshness in a seamless manner. The building blocks of the solution are outlined, along with practical implementation steps and an illustrative AgriTech scenario.

Predictive analytics and cloud-based ML services are being adopted in a growing range of AgriTech decision-support applications. Supported use cases include forecasting, yield estimation, risk assessment, and recommendations. A wide variety of data types from multiple, often heterogeneous, sources contribute to system inputs. Standard BI models are augmented with predictive algorithms in areas where time-varying predictions can facilitate strategic decision-making or operational planning. With model training usually performed offline, the outputs often exhibit a significant time lag compared to the underlying data and fail to account for sudden changes in the relevant variables or the feature relationships. The value of a dynamic or agile approach to predictive analytics becomes apparent, especially when these model predictions are utilized by AgriTech stakeholders for planning and resource management purposes.

Equation 1: Data preprocessing equations the paper implies
Given a feature column $x = \{x_1, \dots, x_n\}$:
Step 1: Compute the mean

$$\mu = \frac{1}{n}\sum_{i=1}^{n} x_i$$

Step 2: Compute the population standard deviation

$$\sigma = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(x_i - \mu)^2}$$

Step 3: Normalize each value

$$z_i = \frac{x_i - \mu}{\sigma}$$

Why this works (key property):

- Mean of z becomes 0:
  $\frac{1}{n}\sum_{i=1}^{n} z_i = \frac{1}{n}\sum_{i=1}^{n}\frac{x_i-\mu}{\sigma} = \frac{1}{\sigma}\left(\frac{1}{n}\sum_{i=1}^{n} x_i - \mu\right) = 0$

Variance becomes 1 (similarly follows by substitution).

2.1. Predictive Analytics in Agriculture: Predictive analytics has been applied across diverse domains for use cases such as forecasting, yield estimation, risk management, decision support, and investment planning. The agriculture sector has investigated predictive analytics as well, but to a lesser degree. The use cases can be differentiated into different categories, including forecasting and estimation tasks, analyzing risk or vulnerability of specific zones (e.g., flood risk estimation) based on weather and zone-specific conditions, providing decision-support tools, and providing predictive risk management tools for specific agricultural developments. The primary types of data sources used are weather data (e.g., temperature, rainfall), remote sensing data, and other crop-specific parameters affecting the region. The time frame for predictive deployment differs from weeks to months. While forecasting and yield estimation problems are generally addressed using benchmarks of trained models and predictive accuracy, considerations of the other categories may require a different evaluation methodology, such as prediction stability and reliability of prediction for risk management.

Dynamic business intelligence for predictive analytics in the agriculture sector has received minimal attention to date. Existing predictive analytics use cases, especially in AgriTech-related sectors, would increase decision-making support services if adapted to variations in environmental and economic conditions. Such adaptation would maintain sustainability of agencies making predictions based on complex data sources (e.g., remote sensing images) and systems with poor prediction reliability during specific time periods. Domain-relevant predictive risk management systems related to agriculture, finance, insurance, or any other stake-sharing settings with a time horizon of weeks or months would be able to fill such gaps with proper predictive support during planning.

2.2. Cloud-Based Machine Learning on Microsoft Azure: Microsoft Azure offers a comprehensive cloud platform for delivering Machine Learning as a Service (MLaaS) in enterprise environments. Central to the proposed adaptive predictive analytics pipelines are the Azure ML services, which afford first-class support for design, training, deployment, and monitoring of machine-learning models; integration with Azure Data Lake Storage for reliable long-term management of all data assets; and convenient governance, scalability, and regulatory-compliance features. The Azure Data Lake Storage service enables secure management of structured, semi-structured, and unstructured data within a

|ISSN 2349-3216| Volume 7, Issue 4, July-August 2024| Bimonthly, Peer Reviewed and Scholarly Indexed Journal|
DOI: 10.15662/IJAESIT.2024.0704008

single facility, providing data connectors and integration tools for data ingestion, transformation, and preparation. Other advantages of using Azure ML include fine-grained governance and security features; enterprise support for computations that can allocate resources according to demand; and ready sharing of systems and models across tenants without requiring additional service accounts.

The extensive ecosystem offered by Azure presents some key advantages over other cloud service providers. The Azure ecosystem includes highly experienced teams of specialists within Azure services—Analytics, AI + Machine Learning, Digital Twins, IoT, and so on—enabling rapid, accurate, cloud-native implementations of predictive models aligned with existing data sources. For predictive analytics pipelines requiring integration of video, audio, or natural-language text (or a combination thereof), Azure provides ready integration with Microsoft Video Indexer, Azure Cognitive Services, and Azure Machine Learning. For business automation, an extensive ecosystem of low-code workflows is offered via Microsoft Power Automate, enabling multi-cloud integrations. Complementary Azure services for predictive analytics pipelines are readily available and maintained, including Azure Synapse Analytics, Azure Key Vault, Azure Monitor, and Azure Well-Architected Review tooling. Data governance, regulatory-compliance, and severe-data-quality-detection services are offered via Microsoft Purview.

III. ARCHITECTURE OF ADAPTIVE PREDICTIVE ANALYTICS PIPELINES

The architecture of adaptive predictive analytics pipelines on Microsoft Azure comprises data ingestion and integration, feature engineering and representation, modeling, and evaluation. Together these processes create dynamic business-intelligence predictions that support strategy formulation in AgriTech systems. By enabling different learners or retraining strategies based on temporal proximity, the architecture mitigates the risks of static prediction models that may become unreliable over time.

The pipelines are designed to adapt as business requirements change or when data become available for additional features or learned representations. These two types of adaptation are achieved by employing connectors that track data availability and quality; by provisioning separate feature pipelines for different predictive models; and by establishing a formalized approach for collecting data that requires costly manual curation. The pipelines also provide a reproducible data-labelling and evaluation framework centred on a dedicated feature store, together with an ensemble of models based on drift detection logic. The integration of representation-learning algorithms into the process further enables adaptation through the use of more powerful features as they become available without incurring retraining costs.

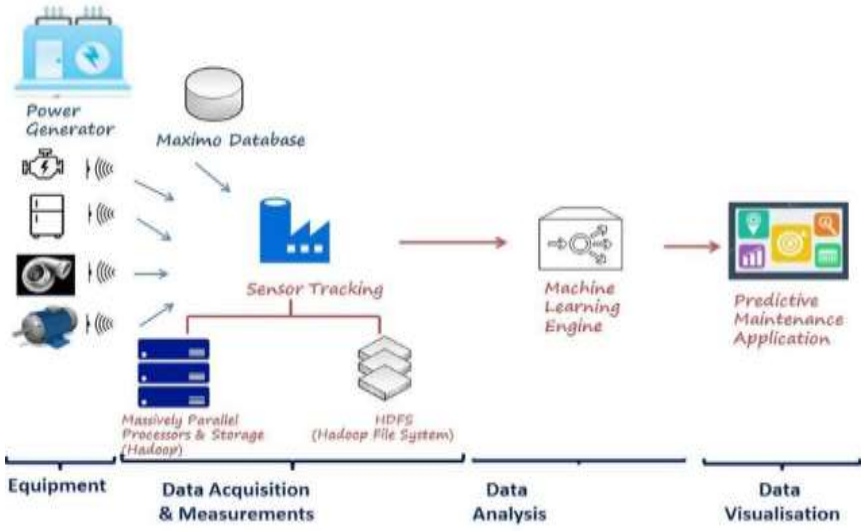
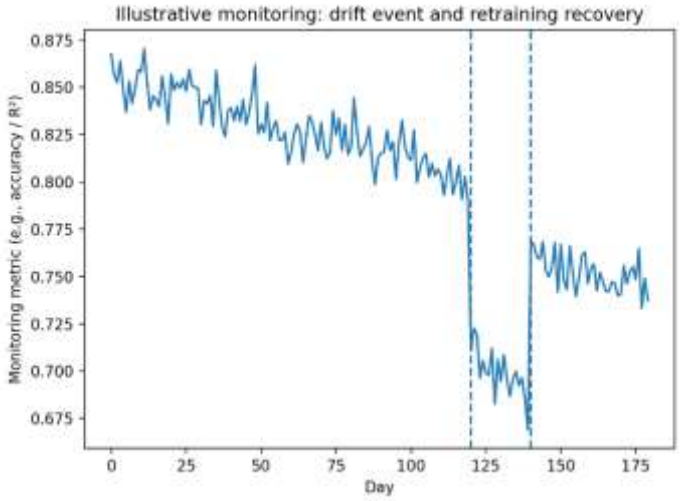


Fig 2: Proposed Big Data Architecture for Predictive Analytics

3.1. Data Ingestion and Integration on Azure: Built-in connectors and services integrate diverse historical and real-time data sources. An overarching pipeline orchestrates the steps—alignment with a common structure, quality checks, and cleansing—that ensure the data are always ready for timely predictive model training. Historical and near-real-time data can be ingested in either batch or streaming mode, with ingestion frequency determined by the temporal resolution of the prediction. During such pre-processing, data are automatically assigned to a metadata environment—the Data Catalog—supporting exploration, testing, and auditing by domain experts.

|ISSN 2349-3216| Volume 7, Issue 4, July-August 2024| Bimonthly, Peer Reviewed and Scholarly Indexed Journal|
DOI: 10.15662/IJAESIT.2024.0704008

Flexible architecture for seamless adaptation relies on intelligent yet unobtrusive management of the underlying data. A feature store mirrors other pipelines, accelerating model training by tracking and making accessible the same engineered features employed by production models. A dedicated feature pipeline is a more powerful concept, since it can represent multidimensional patterns in the data with much less engineering effort. Feature construction becomes more of a discovery process, revealing higher-level features learned from the raw data by independent ML models.



3.2. Feature Engineering and Representation: Adaptation requires a dedicated feature engineering process that constructs the features employed in initial predictive modeling (covering dimension, resolution, and selection) as well as those used for representation learning. Such a process should also support feature normalization (to benefit models sensitive to scale) and engineering of one or more temporal features (to enhance model prediction precision). To facilitate reproducibility and credit assignment, a feature store should manage engineered assets, much like a model repository structures prognostic artifacts.

For the Probability of Frost prediction example, a dedicated pipeline applies rules devised in earlier PhD research (facilitating rapid prediction of risk for multiple locations) to create new feature dimensions that document spatially-relevant data sources. The features support models based on only climatic data, while guiding the construction of models for which farming activities – correlated with warmer conditions – are relevant. Each model implements winter-Summer seasonal dynamics through an indicator feature. For the Cane Disease Risk assessment, a second pipeline normalizes features representing extreme antecedent conditions and quantifies intensity of key disease outbreaks across the instances (including cross-plot temporal integration).

IV. METHODOLOGY

Data processing, prediction modelling and the adaptation mechanisms of the proposed architecture are designed to facilitate environments with constantly changing data. Prediction objectives, feature spaces and models remain valid for a limited time span, especially in the case of time series data. Using historical data for DC directs non-experts towards intelligent forecasts concerning rainfall, temperature and natural disaster risk.

Data Sources and Preprocessing
The implementation is illustrated through a sample use case, and two open-access Internet of Things (IoT) data sources are used. The first source is the Global Surface Summary of Day dataset from the National Oceanic and Atmospheric Administration (NOAA) provides geographic information on world weather stations, daily temperature and precipitation data. The second source is a complete record of rainfall in Tamil Nadu State, India, from a subsidiary of the Indian Meteorological Department. The data is extracted, cleaned and transformed into user-friendly files using the Extrac, Transform, Load/Extract, Load, Transform (ETL/ELT) strategy, along with any other data sources, using real-time processes such as stream and batch ingestion connected via a common metadata repository for quality control and data governance. Missing data points are imputed using interpolation or prediction methods to retain the usefulness of these data.

In generalised predictive analytics scenarios, features and predictors need to be identified, and data should be cleaned. The detection of data drift is assisted by visualisation tools. Depending on the nature of the drift and the prediction algorithm, different techniques may be used. For example, a descriptive model, such as an ensemble based on a large

|ISSN 2349-3216| Volume 7, Issue 4, July-August 2024| Bimonthly, Peer Reviewed and Scholarly Indexed Journal|

DOI: 10.15662/IJAESIT.2024.0704008

training set covering many reasonable scenarios, can provide low-cost predictions, while other techniques can be used when changes between samples need to be scrutinised. Performance monitoring may also be facilitated by a stratified or a large-scale subsample of data, enabling triggering tests of different prediction models without major investments. The remove-purge-and-learn methodology permits continuous adaptation of predictive models.

Equation 2: Interpolation-based imputation for missing data
For linear interpolation between two known points (t_a, x_a) and (t_b, x_b), for a missing time t where $t_a < t < t_b$:

Step 1: Compute the fraction along the interval

$$\lambda = \frac{t - t_a}{t_b - t_a}$$

Step 2: Interpolate

$$\hat{x}(t) = x_a + \lambda(x_b - x_a)$$

Expanded form

$$\hat{x}(t) = x_a + \frac{t - t_a}{t_b - t_a}(x_b - x_a)$$

This is exactly the equation of the straight line passing through the two known points.

4.1. Data Sources and Preprocessing: Dynamic BI on Azure involves data from multiple sources within a defined geographical context. For illustrative purposes, a specific scenario involving the impact of climate variables on the yield of a key crop is adopted. Historical data, spanning up to 20 years, is made publicly available and serves as the primary input for predictive analytics. In addition, procedure data is acquired to support examination of the broader scenario and to facilitate investigation of operational decisions that can influence yield. These data can also enrich model prediction.

Five datasets, each containing a single type of information, are integrated to achieve the modelling goal. Model performance typically benefits from a rich set of features, yet the availability of data for prediction on the feature side dictates any temporal condition on the training set at any given time. Data from the past year have no equivalent next-year features. Reliable values for all features, especially for crops, can extract seasonally and spatially meaningful subsets of the combined data. Missing data and blanks are appropriately examined and processed on a per-feature basis, with techniques such as interpolation, imputation using neighbours, and seasonal decomposition methods considered. Data quality checks are deployed to enhance the value of the data collected and stored in the Azure Data Lake.

4.2. Modeling Techniques and Adaptation Mechanisms: Support for advanced modeling techniques for predictive analytics, novelty detection, and drift detection ensures that changing data distributions are adequately handled in an adaptive predictive analytics system.

Models chosen for predictive analytics covering agritech activities feature decision trees and random forests for structure and multiclass classification support, while ensemble models deliver accuracy improvements. More nuanced approaches enable prediction of plant health and soil quality, with hard and fuzzy segmentation identifying problems and maladies at pixel level. For yield estimation, regression techniques predict the total yield first, with the individual yield per crop derived from the proportions observed during previous periods.

Support for novelty detection ensures sudden changes in data patterns do not adversely affect predictions used for strategic planning and risk management. A combination of clustering and tree-based classifiers identifies non-standard data entries, facilitating investment in research into their causes. The strategic nature of many applications means that an erroneous prediction does not have an immediate costly impact, and therefore able to tolerate a degree of drift.

Time is tagged as a temporal nth dimension descriptive of any ML model suitable for processing time series. Temporal features aid generalization, representational learning facilitate transformations to higher dimensional spaces, and feature engineering pipelines generate a stream of socially-neutral samples ready for ML. The architecture supports both ensemble and online learning, with a separate prevention mechanism for component model drift detection that triggers retraining when the decreased performance exceeds a predefined threshold and automated mechanisms for all categories of model drift. Model performance during the retraining phase is monitored to ensure that any activated backup models deliver acceptable prediction quality.

International Journal of Advanced Engineering Science and Information Technology (IJAESIT)
|ISSN 2349-3216| Volume 7, Issue 4, July-August 2024| Bimonthly, Peer Reviewed and Scholarly Indexed Journal|
DOI: 10.15662/IJAESIT.2024.0704008

V. CASE STUDY: DYNAMIC BI FOR A REPRESENTATIVE AGRITECH SCENARIO

A practical example demonstrates the approach, addressing a realistic problem for an agribusiness and its customers with concrete steps and a results framework. The scenario emphasizes accurately predicting interim data distributions, updating the final distribution as necessary, providing warnings when conditions diverge, remaining cost-effective for a single user, and being easily adaptable to different farms with minimal overhead.

The business problem, data volumes, and set of features in the Azure Data Lake are consistent with other studies in the field (Table 1). The company interacts with natural-resource investors and water-logistics providers on a near-real-time basis, and their customers are often required to provide short-term forecasts of their future water consumption patterns. These short-term forecasts allow both firms to better predict the future demand for water and allocate resources accordingly. Short-term forecasting of water consumption several months ahead is therefore of significant interest. It is also beneficial to be able to supply forecasts for the temporal distribution of water consumption during each month. Although the model is built only for a single area, overall demand dynamics for the whole river basin can be inferred from the Water-Logistics Company's predictions and can inform Hydro-Management Company decisions.

5.1. Scenario Description and Data Characteristics: The case study focuses on a Lentic Ecosystem (Wetland) management scenario involving uncertainty around flood events and their consequences. The primary stakeholder is a government authority responsible for long-term planning and management. The operation horizon is between one and five years, with preparatory steps undertaken monthly or yearly. A representative combination of data sources with varying historical lengths has been identified for a single lake. The Azure Data Factory service is used to evaluate implementation profiles and Excel files are readily available. Three different configurations support analysis of data availability and quality.

The focus of this case study is on risk management and strategic planning functions. However, predictions may also support diminishing behaviour in public policy decisions. A model-agnostic methodology that captures the feature representation of the analysed data allows these two functions to be integrated into a single framework. Data for at least six different risk scenarios have been identified, and additional synthetic scenarios will be generated to assess solution robustness. Overall, the goal is to consolidate a dynamic BI solution that can be implemented across different water reservoirs in the region and repeated in the future.

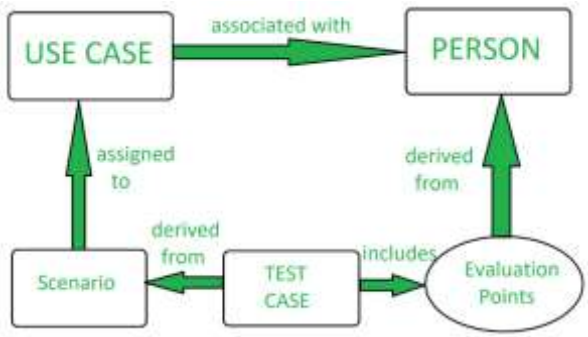


Fig 3: Scenario Testing

5.2. Implementation Details on Azure: Dynamic business intelligence delivered in the selected case study scenario is made possible through an implementation deployed in Microsoft Azure. A combination of products and services in Azure provides the necessary capabilities, including Azure Machine Learning services, data lake storage, and centralized governance. To contain the costs of running the solution, Azure resources are used on a pay-per-use model and automatically scaled according to the current processing load. Consideration is given to multi-tenancy strategies so that the same pipelines can serve multiple farms or different geographical regions without duplicating machine learning models or data-science resources.

Utilizing multiple Azure Machine Learning workspace environments—from development to testing and production—ensures model quality throughout the alignment process. Development resources are kept independent from production to avoid the risk of testing new code with the production database when uploading models. End users benefit from an adaptive predictive-analytics pipeline that can evolve with changing data characteristics. A monitoring dashboard keeps track of ML model performance so that the most predictive algorithm for each prediction window can be selected. An

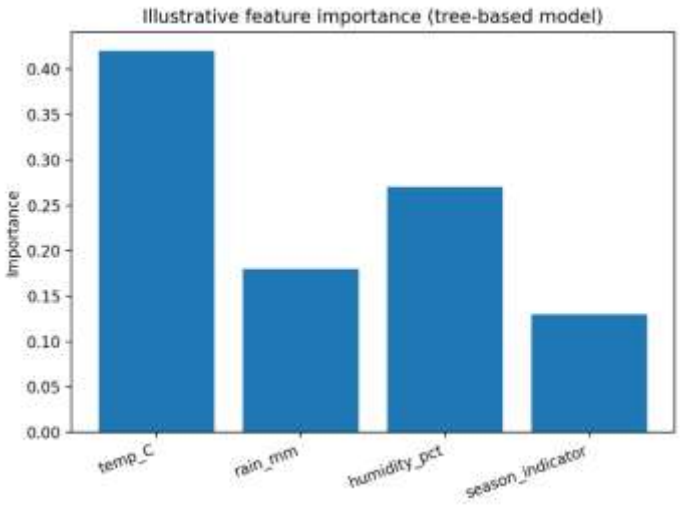
|ISSN 2349-3216| Volume 7, Issue 4, July-August 2024| Bimonthly, Peer Reviewed and Scholarly Indexed Journal|
DOI: 10.15662/IJAESIT.2024.0704008

open-source governance framework designed for the responsible use of AI helps lessen concerns about bias in the generated predictions.

VI. IMPLICATIONS FOR AGRITECH STAKEHOLDERS

The predicted variables are of strategic relevance to various actors in the agriculture sector. Predictions of weather events such as floods, droughts, heatwaves, or cyclones are crucial for the timely planning and execution of any activity by farmers. These predictions not only help mitigate risks, but also aid in adopting proactive measures. Similarly, predictions of seasonal average yield enable better demand and supply management, reducing chances of price crashes during the harvest. Forecasts of potential food grain production or commodity prices are useful to national-level agencies for policy formulation and identifying threats to food security. Potential yield forecasts are, in particular, beneficial for farmers and policymakers in countries like India, where the majority of farmers are smallholders. They can be used for timely supply chain decisions by agriculture insurers and real-time market intervention by the government. Planned transportation of food grains surplus in one part of the country to other deficit regions also requires these predictions.

Business Intelligence provides solutions and strategies to analyze data, make competitive decisions, and formulate appropriate business policies. With predictions such as those discussed above being made at the level of a business intelligence service, other non-contentious mining services such as risk mapping can also be deployed. As the cloud provides services on pay-per-use, small, medium, or new players need not invest in maintenance of the required hardware, which can be costly. Since these services can scale up and down seamlessly as per demand, they should easily be affordable by community-based organizations, such as the Agriculture Insurance of India Company, which hold data on multiple stakeholders in a region. Also, economies of scale can be realized if cloud services such as Azure are used for AgriTech needs across States, by the respective Departments of Agriculture, or Central Government Ministries.



6.1. Decision Support and Strategic Planning: Effective management of agriculture and agribusiness depends heavily on accurate predictions of crop behavior. These predictions are essential for operational decision-support systems as well as for long-term strategic planning. At the operational level, predictions shape decisions related to seedlings, irrigation, fertilization, and pest management. At the strategic level, they inform decisions regarding land use planning, labor force allocation across different months, contract production planning, supply of raw materials for processing, insurance and risk management, and climate policy. Recent history underscores that most of the current problems, including the world food crisis, involve a demand-supply imbalance. Predictions of food production that fall short of meeting demand are, by orders of magnitude, more critical than predictions that overestimate supply.

Food production predictions, especially predictions of food risk or threat areas, serve as inputs for decisions at the country and regional levels. Cautionary warnings not only attract the attention of producers but also of political leaders, policy makers, and traders. For example, four countries—Austria, Ireland, Spain, and the United Kingdom—cautioned of possible risk or threat conditions in world food production during 2004. Responsive actions taken included a stabilization guarantee for consumers and a BS billion emergency fund for poor countries.

6.2. Scalability and Resource Management: Across the prediction space, many models are trained to address similar business challenges, such as forecasting sales and demand, estimating crop yield, or assessing risk mitigation strategies.

|ISSN 2349-3216| Volume 7, Issue 4, July-August 2024| Bimonthly, Peer Reviewed and Scholarly Indexed Journal|

DOI: 10.15662/IJAESIT.2024.0704008

Managing these diverse predictive models and pipelines is challenging; therefore, the architecture should be developed as an adaptive business-intelligence-as-a-service solution. The Azure implementation supports tenant subscriptions to monitor multiple farms or regions, offering business predictions to different customers and enabling multi-tenancy. The cloud-based framework also allows variation in resources to cater for customer demands while optimising cost. During off-peak periods, several unused prediction models can be decommissioned to minimise deployment costs, yet market demand is too limited for a complete model-elimination approach.

Dynamic business-intelligence requirements, like prediction accuracy and fresh data, change with time, so adaptive business-intelligence-as-a-service solutions based on an on-the-fly model-building paradigm are favourable for information-deriving industries such as agri-tech. Resources and costs can also be optimally used by applying model-monitoring, stakeholder-governance, and prediction-prioritisation techniques/engineers.

Equation 3: Prediction/modeling equations implied

Let:

- Feature vector for sample i: $\mathbf{x}_i \in \mathbb{R}^p$
- Stack them into matrix $X \in \mathbb{R}^{n \times p}$
- Targets $y \in \mathbb{R}^n$
- Weights $\beta \in \mathbb{R}^p$

Model:

$$\hat{y} = X\beta$$

Objective (least squares):

$$\min_{\beta} J(\beta) = \| y - X\beta \|^2$$

Step-by-step derivation of the normal equation

Step 1: Expand

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

Step 2: Differentiate w.r.t. β

A standard result:

$$\nabla_{\beta} \ (y - X\beta)^{\top}(y - X\beta) = -2X^{\top}(y - X\beta)$$

Step 3: Set gradient to zero

$$-2X^{\top}(y - X\beta) = 0$$

Step 4: Rearrange

$$X^{\top}y = X^{\top}X\beta$$

Step 5: Solve

If $X^{\top}X$ is invertible:

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

That gives the closed-form fit used in many pipeline baselines.

VII. CHALLENGES AND LIMITATIONS

Several caveats must be considered when assessing the proposed approach. First and foremost is the quality and completeness of the underlying data: AgriTech systems often rely on sources with significant gaps, delay, or other reliability issues, and while regions of dense reporting can yield valuable predictions, any forecast must be treated as a surrogate until confirmed by ground truth. Addressing these issues remains an active area of research, and many of the publicly available systems openly acknowledge these limitations. In the case of weather stations, adaptive techniques could also be deployed to interpolate missing or spurious records based on spatiotemporal proximity to other sensors.

The algorithms employed are responsive to such weaknesses, but insufficient data quality still impairs prediction accuracy. An accurate warning of an impending rainfall may allow a farmer to delay harvesting, but such a decision is ultimately based on calibrating the tradeoff between crop yield and harvest quality. Maintaining such a tradeoff requires an accurate estimate of the yield across the entire farm, which can differ substantially from the predicted rainfall intensity in a small region and thus require different actions. Moreover, predicting drifts represents only one side of the adaptation challenge. Other factors besides the tested input labels may trigger model retraining, and prediction performance should also be monitored. If predictions consistently deviate from the expected label or vice-versa over a time window, and the

|ISSN 2349-3216| Volume 7, Issue 4, July-August 2024| Bimonthly, Peer Reviewed and Scholarly Indexed Journal|
DOI: 10.15662/IJAESIT.2024.0704008

configured labels might still be valid for the confirmation time period, the adaptation threshold could be satisfied and subsequently trigger retraining of the model.

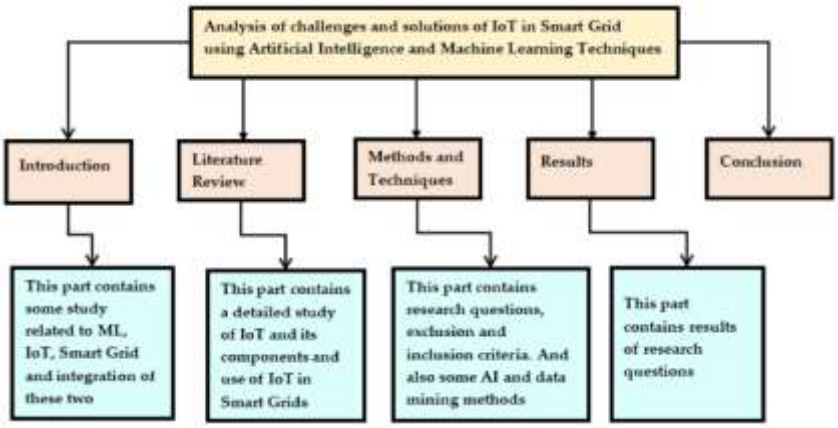


Fig 4: Analysis of Challenges and Solutions of IoT in Smart Grids Using AI and Machine Learning Techniques

7.1. Data Quality and Availability: Many systems rely on input or training data of proven quality. A dynamic BI system cannot be economically maintained under the same constraint—particularly in the light of data from less developed countries or from sources in the academic literature. Residual gaps can be filled with synthetic datasets such as those derived from UniCrop. Country-specific datasets will usually have a higher level of accuracy and completeness than those of lower resolution, but at country level data will often be gathered from a wide variety of sources and become difficult to trust. GIS-based datasets deserve a closer look, both for completeness and reliability. The noise or even bias in a dataset may not affect a specific application; for other applications quality assessment may be more critical. These aspects cannot be determined a priori, so the assessment should accompany the use and detect possible issues. One way to mitigate the problem is to focus on quality-aware machine-learning approaches. Such approaches incorporate latent variables explaining the reliability of the other explanatory features into the model and are therefore able to cope with noisy data. Other solutions include associated sampling strategies based on the quality of the data. Quality issues often apply to the predictands as well.

The exploitation of model ensembles is another way of coping with data quality issues. For most applications in BI systems prediction performance is of paramount interest. It makes no sense to favour one specific model at the expense of accuracy; the predictive outcomes can simply be averaged to deliver the needed result—at least in theory. When exploring the real-world application of a dynamic BI approach using crowd-based datasets the reliability of applied datasets—that clearly differed in nature due to fast-changing foundations—was difficult to assess a priori. The exploration of a deployed dynamic BI system for agriculture-based risk assessment showed that even though the prediction performance delivered a useful picture of future risk, gauging the reliability of the risk measure was not as straightforward. An additional concern is that many applied datasets are inherently limited in size for specific regions. When focusing on risk assessment, another potential drawback of exploiting a BI system that provides predictions for a rolling time horizon is that prediction performance may tend to improve—yet too few predictions are generated across years to ensure the reliability of the results.

7.2. Model Drift and Adaptation Latency: The drift of statistical characteristics in the input data is a common problem when it comes to the generalization capabilities of machine learning (ML) models. Although drift can occur at any point during development, it particularly affects the phase where predictions are made. This phenomenon can occur after the model has been deployed, due to factors like seasonal variations or local events, even if it had performed well at validation time. Detection of drift can be managed using methods such as univariate index monitoring, multivariate index monitoring, or even specialized packages (e.g. ADTK).

When using supervised learning models trained in an offline or batch mode, any detected alteration in the statistical characteristics of the input data can trigger the retraining of the model using more recent data in an online learning strategy. It is worth mentioning that any maintenance or retraining procedure is usually a long-term process from an analytical point of view, meaning that no training requires high levels of computational resources. Even in cases where an ensemble of models is used, model production can be quite light, based mainly on executing simple procedures rather than on a full-blown training process. Addressing drift detection and retraining through these mechanisms can help provide reasonable, up-to-date predictions with very low latency.

|ISSN 2349-3216| Volume 7, Issue 4, July-August 2024| Bimonthly, Peer Reviewed and Scholarly Indexed Journal|

DOI: 10.15662/IJAESIT.2024.0704008

A more complex situation arises in the context of transfer learning models or in cases where the initial model is a specialized component of an ensemble model (e.g. a specialist within a mixture of experts). In these cases, only part of the model needs to be retrained and a more fine-grained mechanism that detects drift in the specific domain covered by that specialized model should be considered.

VIII. CONCLUSION AND FUTURE DIRECTIONS

The proposed architecture enables the definition of adaptive predictive analytics pipelines for business-intelligence support in AgriTech scenarios, and has been illustrated through a case study set in the territory of a public agronomic research institution and involving the development of predictive models for typically predictable but hard-to-forecast events. For the considered case study, the action of several public and private stakeholders combines to guarantee successful operation and improve a business process, and those stakeholders need to combine prediction information at a local level to allow proper decision-making at a regional level. However, the approach is general and can be applied to any center that wants to develop predictive models using cloud-based machine-learning services, shared by several tenants, for business-intelligence purposes.

The framework clearly allows for the extension, automation, and application of more advanced techniques. First of all, pipelines could be defined for different machine-learning classes (not only supervised) and purposes (not only prediction). Moreover, once basic templates are defined, adapting pipelines to specific zones or predictive targets could take advantage of transfer-learning techniques. Adding a drift-detection capability would allow the models to signal changes in performance, and advanced techniques would then automatically create epidemiological, machine-related, statistical, or other features. Finally, the approach could apply to other connected environments that use complementary Business Intelligence, Knowledge Management, or activity and asset management data ecosystems, to reinforce a synergy-based regulatory-market strategy.

8.1. Advanced Adaptive Techniques: Aiming to fill current gaps in Azure-based dynamic decision-making for agriculture, the presented approach lays a solid foundation for future research. Directions include exploiting advanced procedures for model adaptation, augmenting data-driven analytics, introducing increased automation into the pipelines, and enabling knowledge transfer across agri-tech enterprises. Replicability and broad applicability of the method define an additional line for future exploration.

A more intelligent data-parsing layer would allow the pipelines to hop from one configuration to another—without human intervention—as soon as any component reached a switching condition. Advanced model-adaptive techniques could also be conceived. Enabling ensemble learning across different models would avoid unlearning when data drift occurred and data from different regions across a similar temporal range were available for training. Meanwhile, open-source algorithms for quantifying uncertainty in regressors and classifiers, and monitoring the associated signals, could be incorporated into the Azure architecture to exploit this information via a covariance matrix or in subsequent models.

Transfer learning capabilities could be developed to switch a model in a pipeline for another model trained in a similar context but with a much longer training history, making it possible for the model to adapt rapidly. Expanding the data-storing strategy to accept more high-boarding data streams analyzed for other purposes in the same region is also desirable. Such cross-domain data ecosystems have the potential to expand the contribution of machine learning in prediction and decision-support inquiries and problems from the probabilistic perspective.

8.2. Cross-Domain Generalization: The Azure Pipelines for Adaptive Predictive Analytics can be further extended and adapted to serve the prediction demands of additional domains of the AgriTech sector and potentially other application areas, such as supply chain management. Demand prediction in the food sector, predicted yields for various grains worldwide, prediction of wheat and rice prices in different regions, detection of disease and pest outbreaks for different crops, plant health assessment, prediction of erratic weather patterns, and crop insurance-related use cases all depend on similar prediction models. Crop insurance prediction is particularly interesting because the prediction can directly help in policy formulation. Actual demand and supply play important roles in predicting these types of events and decisions. Data from multiple sources can be collected for these events, and predictive models can be trained and used through these pipelines. All these models are updated over different periods, which can be handled using the discussed architecture.

With a few adaptations and refinements, the discussed pipelines can be generalized and adapted to work beyond the AgriTech domain, especially in other technology-related areas such as supply chain management. The need for continuous updating of the predictive model due to changing environments, repeated prediction for short periods (on a daily basis), and lack of reliable and consistent models in the supply chain domain are key factors. Many products and services in the supply chain depend heavily on the intelligent prediction of products and risk assessment. These pipelines

|ISSN 2349-3216| Volume 7, Issue 4, July-August 2024| Bimonthly, Peer Reviewed and Scholarly Indexed Journal|

DOI: 10.15662/IJAESIT.2024.0704008

can be used to create demand forecasts for various products, detect the likelihood and impact of PANDEMIC/EPIDEMIC diseases, and predict the level of risk in product deliveries, especially curative products. These four aspects can be examined in greater detail. Further, with enhanced efforts and exploration, knowledge transfer and sharing from complementary areas and domains can be made possible.

REFERENCES

1. Aitha, A. R. (2021). Optimizing Data Warehousing for Large Scale Policy Management Using Advanced ETL Frameworks.
2. Challa, K. (2023). Optimizing Financial Forecasting Using Cloud Based Machine Learning Models. Journal for ReAttach Therapy and Developmental Diversities. https://doi. org/10.53555/jrtdd. v6i10s (2), 3565.
3. Inala, R. (2023). AI-powered investment decision support systems: Building smart data products with embedded governance controls. Journal for ReAttach Therapy and Developmental Diversities, 6(10), 2251-2266.
4. Adusupalli, B. (2023). DevOps-Enabled Tax Intelligence: A Scalable Architecture for Real-Time Compliance in Insurance Advisory. Journal for Reattach Therapy and Development Diversities. Green Publication. https://doi. org/10.53555/jrtdd. v6i10s (2), 358.
5. Meda, R. (2023). Data Engineering Architectures for Scalable AI in Paint Manufacturing Operations. European data science journal.
6. Paleti, S. (2022). Fusion bank: Integrating AI-driven financial innovations with risk-aware data engineering in modern banking. Decision Making, 2326, 9865.
7. Vijaya Rama Raju Gottimukkala, "Energy-Efficient Design Patterns for Large-Scale Banking Applications Deployed on AWS Cloud," International Journal of Advanced Research in Computer and Communication Engineering (IJARCCE), DOI: 10.17148/IJARCCE.2020.91225
8. Mattaparthi, R. (2022). Engineering Predictive Industrial Systems Through IoT-Driven Asset Monitoring and Machine Learning Prognostics. International Journal of Future Innovative Science and Technology (IJFIST), 5(1), 7790.
9. Kolla, S. K., & Reddy, V. A. R. (2023). Deep Learning Architectures For Multimodal Medical Data Integration. South Eastern European Journal of Public Health, 248-260.
10. Mangalampalli, B. M. (2023). AI-Driven Anomaly Detection in Healthcare Claims Data: A Business Intelligence Perspective. Journal of Rare Cardiovascular Diseases.
11. Nagabhyru, K. C. (2022). Bridging Traditional ETL Pipelines with AI Enhanced Data Workflows: Foundations of Intelligent Automation in Data Engineering. Available at SSRN 5505199.
12. Gottimukkala, V. R. R. (2021). Digital Signal Processing Challenges in Financial Messaging Systems: Case Studies in High-Volume SWIFT Flows. Universal Journal of Finance and Economics.
13. Garapati, R. S. (2023). Optimizing Energy Consumption in Smart Build-ings Through Web-Integrated AI and Cloud-Driven Control Systems.
14. Kaulwar, P. K., Pamisetty, A., Mashetty, S., Adusupalli, B., & Pandiri, L. (2023). Harnessing intelligent systems and secure digital infrastructure for optimizing housing finance, risk mitigation, and enterprise supply networks. International Journal of Finance (IJFIN)-ABDC Journal Quality List, 36(6), 372-402.
15. Aitha, A. R. (2022). Cloud Native ETL Pipelines for Real Time Claims Processing in Large Scale Insurers. Universal Journal of Business and Management.
16. Nagabhyru, K. C., & Engineer, S. D. (2023). Unifying Data Engineering and Machine Learning Pipelines: An Enterprise Roadmap to Automated Model Deployment.
17. Inala, R. (2023). Revolutionizing Customer Master Data in Insurance Technology Platforms: An AI and MDM Architecture Perspective. International Journal of Finance (IJFIN)-ABDC Journal Quality List, 36(6), 579-606.
18. Garapati, R. S. (2022). Web-Centric Cloud Framework for Real-Time Monitoring and Risk Prediction in Clinical Trials Using Machine Learning. Current Research in Public Health, 2, 1346.
19. Kolla, T. (2023). Predictive ETL Failure Detection in Healthcare Data Pipelines Using Anomaly Detection Algorithms. International Journal of Medical Toxicology & Legal Medicine.
20. Nagabhyru, K. C. (2023). From Data Silos to Knowledge Graphs: Architecting CrossEnterprise AI Solutions for Scalability and Trust. Available at SSRN.
21. Inala, R. (2023). Big Data Architectures for Modernizing Customer Master Systems in Group Insurance and Retirement Planning. Educational Administration: Theory and Practice, 29(4), 5493-5505.
22. Challa, K. (2023). Dynamic Neural Network Architectures for Real-Time Fraud Detection in Digital Payment Systems Using Machine Learning and Generative AI. Nanotechnology Perceptions.
23. Garapati, R. S. (2022). AI-Augmented Virtual Health Assistant: A Web-Based Solution for Personalized Medication Management and Patient Engagement. Available at SSRN.
24. Mattaparthi, R. (2023). Deep Learning-Driven Combustion Anomaly Detection in Diesel Powertrains: A Multi-Sensor Fusion Approach for Real-Time ECM Adaptation. International Journal of Intelligent Systems and Applications in Engineering, 11, 1084.

|ISSN 2349-3216| Volume 7, Issue 4, July-August 2024| Bimonthly, Peer Reviewed and Scholarly Indexed Journal|

DOI: 10.15662/IJAESIT.2024.0704008

25. Segireddy, A. R. (2020). Cloud Migration Strategies for High-Volume Financial Messaging Systems.
26. Amistapuram, K. (2023). Privacy-Preserving Machine Learning Models for Sensitive Customer Data in Insurance Systems. Educational Administration: Theory and Practice, 29(4), 5950-5958.
27. Davuluri, P. S. L. N. (2023). Integrating artificial intelligence into event-driven financial crime compliance platforms. International Journal of Finance, 36(6), 707-736.
28. Adusupalli, B. Strengthening Cybersecurity Compliance: Evaluating Risk Assessment Frameworks and AI-Based Anomaly Detection in Audit Practices.
29. Kolla, S. H. (2022). Strategic Information Integration Models for Cross-Functional Service Optimization in Large-Scale Enterprises. International Journal of Emerging Trends in Engineering and Management Research, 7(3), 11811.
30. Yandamuri, U. S. (2023). An Intelligent Analytics Framework Combining Big Data and Machine Learning for Business Forecasting. International Journal Of Finance, 36(6), 682-706.
31. Davuluri, P. S. L. N. (2023). Integrating artificial intelligence into event-driven financial crime compliance platforms. International Journal of Finance, 36(6), 707-736.
32. Meda, R., & Pamisetty, A. (2023). Intelligent Infrastructure for Real-Time Inventory and Logistics in Retail Supply Chains. Educational Administration: Theory and Practice, 29(4), 5215-5233.
33. Mashetty, S. (2023). Leveraging Data Analytics to Enhance Affordable Housing Initiatives and Community Development. Available at SSRN 5249221.
34. Mattaparthi, R. (2023). Connected Fleet Intelligence: Edge-Centric Analytics and Computer Vision for Predictive Manufacturing and Asset Resilience. International Journal of Advanced Research in Computer Science & Technology (IJARCST), 6(5), 9077-9088.
35. Aitha, A. R. (2021). Dev Ops Driven Digital Transformation: Accelerating Innovation In The Insurance Industry. Journal of International Crisis and Risk Communication Research.
36. Segireddy, A. R. (2021). Containerization and Microservices in Payment Systems: A Study of Kubernetes and Docker in Financial Applications. Universal Journal of Business and Management, 1(1), 1-17.
37. Mangala, N. (2022). Real-Time Data Quality Monitoring and Gating Frameworks in Cloud-Based Data Pipelines. International Journal of Research and Applied Innovations, 5(6), 8197-8219.
38. Peddi, R. K. (2021). Optimizing Case Management Workflows in Global Data Center Colocation Services. Universal Journal of Computer Sciences and Communications, 1(1), 1-21.
39. Loganathan, R. (2021). Integrated Risk and Compliance Frameworks for Global Data Center Operations: A Governance-Centric Approach. Universal Journal of Computer Sciences and Communications, 1(1), 1-26.
40. Amistapuram, K. (2021). Digital Transformation in Insurance: Migrating Enterprise Policy Systems to. NET Core. Universal Journal of Computer Sciences and Communications, 1(1), 1-17.
41. Reddy, V. A. R. (2023). Predictive Healthcare Administration Using Advanced Payer Analytics and Population Health Data Engineering. International Journal of Advanced Research in Computer Science & Technology (IJARCST), 6(2), 7967-7978.
42. Gottimukkala, V. R. R. (2022). Licensing Innovation in the Financial Messaging Ecosystem: Business Models and Global Compliance Impact. International Journal of Scientific Research and Modern Technology, 1(12), 177-186.
43. Mangalampalli, B. M. (2022). Automated Invoice Validation Systems Using Advanced SQL Analytics in Healthcare Insurance. Front Health Inform, 11.
44. Reddy, V. A. R. (2021). Challenges in Standardizing Member Eligibility Data Across Multi-Payer Healthcare Ecosystems. International Journal of Medical Toxicology and Legal Medicine, 24(3), 1-19.
45. Davuluri, P. S. L. (2023). AI-Augmented Sanctions Screening: Enhancing Accuracy and Latency in Real Time Compliance Systems. AI-Augmented Sanctions Screening: Enhancing Accuracy and Latency in Real Time Compliance Systems (December 15, 2023).
46. Yandamuri, U. S. (2021). A Comparative Study of Traditional Reporting Systems versus Real-Time Analytics Dashboards in Enterprise Operations. Universal Journal of Business and Management, 1(1), 1-13.
47. Divya, V., & Bandi, V. K. (2023). Cloud-Native Model Lifecycle Management for Enterprise AI Systems. International Journal of Scientific Research and Modern Technology, 78.
48. Kolla, S. K., & Reddy, V. A. R. (2023). Deep Learning Architectures For Multimodal Medical Data Integration. South Eastern European Journal of Public Health, 248–260.
49. Bandi, V. D. V. K. (2023, November). Production-Grade Machine Learning Pipelines For Healthcare Predictive Analytics. South Eastern European Journal of Public Health, 189–205.
50. Kolla, S. H., & Loganathan, R. (2023). Cloud-Native Deep Learning Architectures For Secure Generative AI Deployment In Enterprise Workflow Platforms. Journal of International Crisis and Risk Communication Research, 603-618.
51. Mangala, N. (2021). Optimizing Large-Scale ETL Pipelines Using Medallion Architecture on Azure Data Lake. Journal of Artificial Intelligence and Big Data, 1(1), 1-20.
52. Paleti, S. (2023). AI-driven innovations in banking: Enhancing risk compliance through advanced data engineering. Available at SSRN, 5244840.

|ISSN 2349-3216| Volume 7, Issue 4, July-August 2024| Bimonthly, Peer Reviewed and Scholarly Indexed Journal|

DOI: 10.15662/IJAESIT.2024.0704008

53. Reddy, V. A. R. (2023). Orchestrating the Future Autonomous Healthcare Data Pipeline Management through Agentic AI Architectures. International Journal of Advanced Research in Computer Science & Technology (IJARCST), 6(2), 7979-7992.
54. Kolla, S. K. (2023). Learning Health Systems Machine Intelligence for Clinical Prediction and Healthcare Optimization. International Journal of Advanced Research in Computer Science & Technology (IJARCST), 6(2), 7955-7966.
55. Valiki, D., & Segireddy, A. R. (2023). Deep Learning Architectures Deployed on Cloud Platforms for Dynamic Financial Risk Evaluation and Market Prediction. American International Journal of Computer Science and Technology, 5(5), 12-24.
56. Amistapuram, K. Energy-Efficient System Design for High-Volume Insurance Applications in Cloud-Native Environments. International Journal of Innovative Research in Electrical, Electronics, Instrumentation and Control Engineering (IJIREEICE), DOI, 10.
57. Kaulwar, P. K., Pamisetty, A., Mashetty, S., Adusupalli, B., & Pandiri, L. (2023). Harnessing intelligent systems and secure digital infrastructure for optimizing housing finance, risk mitigation, and enterprise supply networks. International Journal of Finance (IJFIN)-ABDC Journal Quality List, 36(6), 372-402.
58. DAVULURI, P. N. (2022). Cloud-Native Data Platform Modernization for Regulatory Compliance in Global Banking. Kurdish Studies.
59. Kolla, S. H. (2023). Large Language Model-Driven Enterprise Service Intelligence for Digital Workflow Transformation. International Journal of Research and Applied Innovations, 6(1), 8380-8391.
60. Loganathan, R. (2022). Converging Security Architecture and Compliance Management in Enterprise Data Center Ecosystems: A Unified Control Framework. International Journal of Scientific Research and Modern Technology, 1(12), 295-312.
61. Yandamuri, U. S. (2022). Cloud-Based Data Integration Architectures for Scalable Enterprise Analytics. International Journal of Intelligent Systems and Applications in Engineering, 10, 472-483.
62. Kolla, T., & Kolla, S. K. (2023). FHIR-Based Real-Time Healthcare Analytics using Unsupervised Learning. International Journal of Future Innovative Science and Technology (IJFIST), 6(6), 11751.
63. Mangala, N. (2022). Implementing Databricks Unity Catalog For Centralized Data Governance In Multi-Business-Unitenterprises. Journal of International Crisis and Risk Communication Research, 101-122.
64. Bandi, V. D. V. K. (2023). MLOps frameworks for reliable model deployment in cloud data platforms. Journal of Artificial Intelligence and Big Data, 3(1), 84-101.
65. Mangalampalli, B. M. (2021). Scalable Data Warehouse Architecture for Population Health Management and Predictive Analytics. World Journal of Clinical Medicine Research, 1(1), 1-18.
66. Mashetty, S. (2023). A Comparative Analysis of Patented Technologies Supporting Mortgage and Housing Finance. Available at SSRN 5249181.
67. Meda, R. (2023). Developing AI-Powered Virtual Color Consultation Tools for Retail and Professional Customers. Journal for ReAttach Therapy and Developmental Diversities. https://doi. org/10.53555/jrtdd. v6i10s (2), 3577.
68. Paleti, S. (2023). Trust layers: AI-augmented multi-layer risk compliance engines for next-gen banking infrastructure. Available at SSRN, 5221895.
69. Challa, K. (2023). Transforming Travel Benefits through Generative AI: A Machine Learning Perspective on Enhancing Personalized Consumer Experiences. Educational Administration: Theory and Practice. Green Publication. Educational Administration: Theory and Practice. Green Publication. https://doi. org/10.53555/kuey. v29i4, 9241.