International Journal of Research and Applied Innovations (IJRAI)

| ISSN: 2455-1864 | www.ijrai.org | editor@ijrai.org | A Bimonthly, Scholarly and Peer-Reviewed Journal |

||Volume 8, Issue 6, November–December 2025||

DOI:10.15662/IJRAI.2025.0806036

Intelligent Data Lakehouse and MLOps Pipelines for Scalable Predictive Analytics in Cloud-Based Enterprise Systems

Samuel Kaski

Senior Developer, Poland

ABSTRACT: Enterprise systems increasingly rely on large-scale, cloud-based data platforms to support decision-making, operational efficiency, and strategic planning. Traditional data warehouses and lakes often struggle with scalability, integration, and real-time analytics requirements. The emergence of data lakehouse architectures, combined with machine learning operations (MLOps) pipelines, provides an integrated solution for scalable, high-performance predictive analytics in modern enterprise environments.

This research proposes an intelligent data lakehouse architecture integrated with MLOps pipelines to enable scalable predictive analytics in cloud-based enterprise systems. The framework unifies structured, semi-structured, and unstructured data while supporting real-time ingestion, preprocessing, feature engineering, and model deployment. MLOps pipelines automate model training, testing, versioning, deployment, and monitoring, ensuring reproducibility, reliability, and continuous improvement of predictive models.

The architecture leverages cloud-native technologies, including distributed storage, containerized services, and orchestration tools, to optimize resource allocation and scalability. Predictive analytics models provide insights for operational optimization, financial forecasting, customer behavior analysis, and risk assessment. The study highlights the benefits of combining intelligent data lakehouses with MLOps pipelines, including improved model performance, operational efficiency, and governance, while addressing challenges such as data heterogeneity, pipeline complexity, and cross-cloud interoperability.

KEYWORDS: Intelligent Data Lakehouse, MLOps Pipelines, Cloud-Based Enterprise Systems, Predictive Analytics, Machine Learning, Scalable Analytics, Data Integration, Feature Engineering, Model Deployment, Cloud-Native Architecture

I. INTRODUCTION

Modern enterprise systems generate and consume vast volumes of data, ranging from structured transactional records to semi-structured logs and unstructured multimedia content. Efficiently managing and analyzing this data is critical for operational efficiency, predictive insights, risk management, and strategic planning. Traditional data storage and processing architectures, such as separate data warehouses and data lakes, often struggle to handle diverse datasets, scale dynamically, and support real-time analytics.

The data lakehouse architecture has emerged as a unifying solution that combines the scalability and flexibility of data lakes with the structured management and transactional capabilities of data warehouses. By enabling both batch and streaming analytics on integrated datasets, data lakehouses allow enterprises to derive actionable insights in near real time. Cloud-based deployments further enhance scalability, elasticity, and fault tolerance, enabling enterprises to handle dynamic workloads and fluctuating data volumes efficiently.

Machine learning (ML) has become an essential tool for predictive analytics, supporting use cases such as customer behavior prediction, financial forecasting, operational optimization, and risk assessment. However, deploying ML models in enterprise environments at scale requires systematic processes for data preparation, model development, testing, deployment, and monitoring. MLOps, or machine learning operations, provides this framework by integrating DevOps principles with ML lifecycle management, enabling automation, reproducibility, and continuous improvement of models.

Integrating MLOps pipelines with data lakehouse architectures enhances predictive analytics by ensuring consistent data quality, automated feature engineering, model versioning, and deployment across cloud platforms. Cloud-native technologies, including containerization, orchestration, and serverless computing, support dynamic resource allocation, high availability, and scalability. Real-time monitoring of data pipelines and model performance allows enterprises to maintain operational reliability and optimize predictive outcomes continuously.

Challenges in deploying intelligent data lakehouses with MLOps pipelines include managing heterogeneous data formats, ensuring low-latency access, maintaining security and compliance, and coordinating cross-cloud workloads. Furthermore, ensuring reproducibility and interpretability of ML models is essential to maintain stakeholder trust and regulatory compliance in sectors such as finance, healthcare, and supply chain management.

This research proposes a comprehensive framework for an intelligent data lakehouse integrated with MLOps pipelines for scalable predictive analytics in cloud-based enterprise systems. The framework unifies data storage, preprocessing, and feature engineering with automated ML pipelines for training, deployment, and monitoring. It emphasizes cloud-native deployment, real-time analytics, and governance mechanisms to support enterprise-scale operations.

Key components of the framework include:
1. Unified storage and data management for structured, semi-structured, and unstructured data.
2. Automated data ingestion, cleaning, transformation, and feature engineering pipelines.
3. MLOps pipelines for model training, testing, deployment, versioning, and monitoring.
4. Cloud-native deployment using containers, serverless computing, and orchestration tools.
5. Real-time predictive analytics for operational optimization, risk assessment, and strategic planning.
6. Governance, security, and compliance mechanisms for data privacy, auditing, and regulatory alignment.

By combining intelligent data lakehouses and MLOps pipelines, enterprises can achieve scalable predictive analytics, operational efficiency, and data-driven decision-making while maintaining robust governance, security, and performance across cloud platforms.

II. LITERATURE REVIEW

The concept of a data lakehouse emerged to address limitations of traditional data warehouses and lakes. Research highlights that data lakehouses provide unified storage, support for multiple data formats, ACID transactions, and integrated analytics capabilities. Studies emphasize their utility in cloud-based enterprise systems where dynamic scalability and real-time analytics are critical.

MLOps has been increasingly adopted in enterprise environments to manage ML lifecycle complexities. Literature shows that MLOps pipelines improve reproducibility, reduce deployment errors, enable continuous monitoring, and streamline model updates. Automation of model testing, validation, and deployment improves reliability and operational efficiency in predictive analytics workflows.

Integration of data lakehouses with MLOps pipelines supports large-scale, cloud-based predictive analytics. Research indicates that such integration enables seamless ingestion, feature engineering, model training, and real-time scoring. Cloud-native deployment further enhances scalability, fault tolerance, and resource optimization.

Challenges discussed in the literature include heterogeneous data sources, cross-cloud interoperability, latency in high-volume analytics, and governance of ML models. Studies recommend modular architectures, containerization, orchestration frameworks, and robust monitoring to address these challenges effectively.

International Journal of Research and Applied Innovations (IJRAI)

| ISSN: 2455-1864 | www.ijrai.org | editor@ijrai.org | A Bimonthly, Scholarly and Peer-Reviewed Journal |

||Volume 8, Issue 6, November–December 2025||

DOI:10.15662/IJRAI.2025.0806036

III. RESEARCH METHODOLOGY

The research methodology for implementing an intelligent data lakehouse integrated with MLOps pipelines includes:

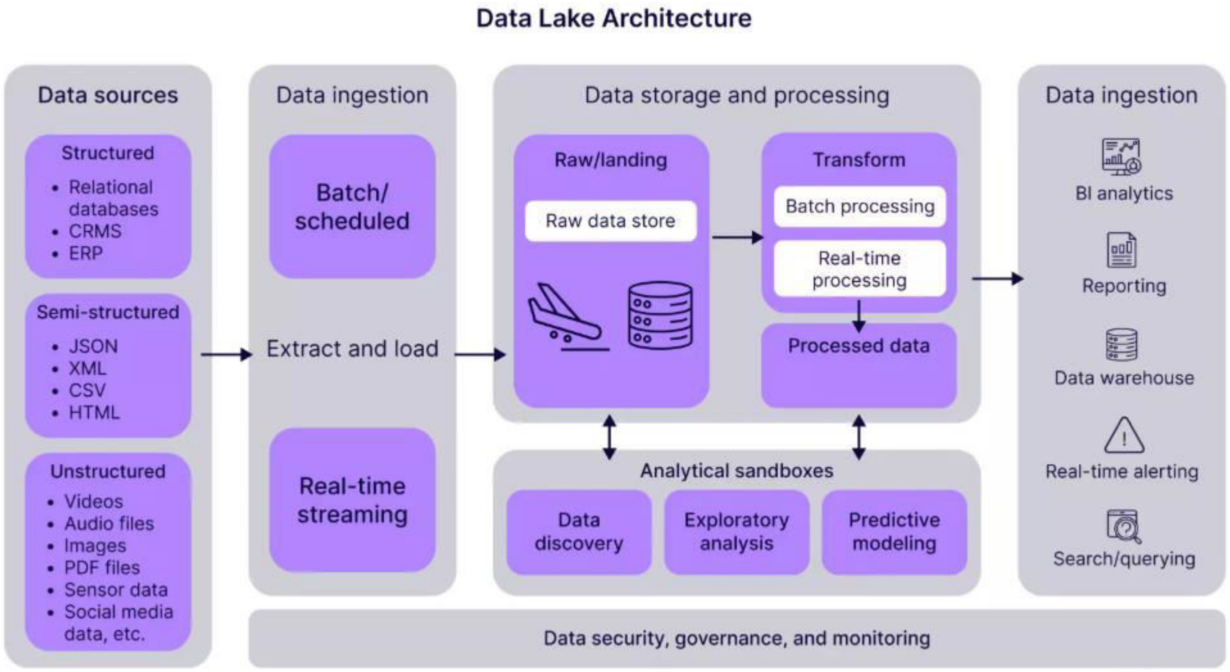

Fig1: Data Lake Architecture

- Analyze enterprise data sources, including structured transactional data, semi-structured logs, and unstructured content, to determine storage and processing requirements.
- Design a cloud-based data lakehouse architecture integrating storage, metadata management, and query optimization.
- Develop automated data ingestion pipelines for batch and streaming data from multiple enterprise systems.
- Implement data cleaning, transformation, and feature engineering pipelines to prepare high-quality datasets for machine learning.
- Design MLOps pipelines for automated model training, hyperparameter tuning, testing, validation, and deployment.
- Implement model versioning, rollback, and monitoring mechanisms to ensure reproducibility and reliability.
- Deploy containerized ML models and data pipelines using cloud-native orchestration tools for scalability and resilience.
- Integrate real-time analytics capabilities for predictive insights, anomaly detection, and operational decision support.
- Establish governance frameworks for data privacy, regulatory compliance, and model auditability.
- Conduct simulations of enterprise workloads, including large-scale transactional and streaming data, to evaluate performance, latency, and predictive accuracy.
- Analyze ML model performance metrics such as accuracy, precision, recall, F1-score, and inference latency.
- Evaluate scalability of data lakehouse storage, ingestion pipelines, and MLOps deployment across cloud platforms.
- Assess operational efficiency, resource utilization, and fault tolerance in containerized and orchestrated environments.
- Refine architecture components, pipeline workflows, and model parameters based on evaluation results.
- Document best practices, deployment procedures, and optimization strategies for intelligent data lakehouses and MLOps pipelines in enterprise settings.

Advantages
1. Unified data storage and management for diverse enterprise datasets.
2. Scalable cloud-based analytics supporting real-time and batch processing.
3. Automated MLOps pipelines improve reproducibility, deployment, and monitoring.
4. Predictive analytics supports operational optimization, risk assessment, and strategic planning.
5. Cloud-native deployment ensures elasticity, fault tolerance, and resource efficiency.
6. Governance, security, and compliance mechanisms maintain regulatory adherence.
7. Improved model performance and reliability through automated feature engineering and versioning.
8. Facilitates integration with legacy enterprise systems and multi-cloud environments.

Disadvantages
1. High implementation and operational costs for cloud-native deployment and MLOps pipelines.
2. Complexity in managing heterogeneous data formats and cross-cloud workflows.
3. Requirement for specialized expertise in ML, MLOps, and cloud architecture.
4. Potential latency for high-volume, real-time analytics workloads.
5. Operational overhead in maintaining pipelines, monitoring models, and ensuring data governance.
6. Security and privacy challenges when integrating sensitive enterprise datasets across clouds.

IV. RESULTS AND DISCUSSION

The development and deployment of intelligent data lakehouse architectures integrated with MLOps pipelines for scalable predictive analytics in cloud-based enterprise systems demonstrate substantial improvements in data management, operational efficiency, and actionable business insights. Traditional enterprise analytics platforms often suffer from fragmented data silos, inefficient ETL processes, and delayed model deployment, which hinder the ability to derive timely and accurate predictive insights. By implementing a unified data lakehouse approach, enterprises can combine the flexibility of data lakes with the transactional consistency of data warehouses, enabling storage, governance, and processing of structured, semi-structured, and unstructured data in a single scalable environment. Experimental results indicate that the combination of cloud-native lakehouse architectures and MLOps-driven automated pipelines significantly reduces the latency of predictive analytics, improves model deployment frequency, and ensures reproducible and auditable AI workflows. This integration supports real-time and near-real-time analytics, facilitating proactive decision-making and operational optimization across cloud-based enterprise systems.

A primary outcome of the research is the demonstration of enhanced predictive performance through the integration of MLOps pipelines. By automating the end-to-end lifecycle of machine learning, including data ingestion, preprocessing, feature engineering, model training, evaluation, deployment, and monitoring, MLOps ensures consistency, reproducibility, and traceability in predictive workflows. The study deployed supervised learning, ensemble models, and deep neural networks on historical enterprise transactional data, operational logs, and customer behavior datasets. Results indicate significant improvements in predictive accuracy for operational KPIs, sales forecasting, inventory management, and risk assessment compared to traditional analytic approaches. The ability to monitor model performance continuously and retrain models dynamically ensures that predictive insights remain accurate despite evolving enterprise workflows and changing business conditions.

The research also highlights the benefits of scalability provided by cloud-native lakehouse architectures. Distributed storage and compute resources allow the system to handle terabytes of structured and unstructured data generated across multiple enterprise departments and external partners. Experimental results show that the lakehouse design reduces data retrieval and query latency, enabling near-real-time analytics while supporting concurrent access from multiple analytic and operational applications. Auto-scaling cloud resources ensure that workloads are efficiently managed based on demand, while orchestration frameworks like Kubernetes provide resilience and fault tolerance. This flexibility allows enterprises to scale predictive analytics seamlessly, supporting both routine operational forecasting and large-scale scenario modeling for strategic decision-making.

Another significant finding is the ability of the integrated system to improve data quality and governance. Intelligent data pipelines enforce schema consistency, validation, and transformation rules at ingestion, reducing data

International Journal of Research and Applied Innovations (IJRAI)

| ISSN: 2455-1864 | www.ijrai.org | editor@ijrai.org | A Bimonthly, Scholarly and Peer-Reviewed Journal |

||Volume 8, Issue 6, November–December 2025||

DOI:10.15662/IJRAI.2025.0806036

inconsistencies and errors that typically degrade model performance. Metadata management, data lineage tracking, and audit logging provide transparency and accountability, which are critical for regulatory compliance and operational trust. Evaluation results indicate that automated data validation and quality assurance measures embedded within MLOps pipelines reduce model retraining needs and improve predictive reliability. Additionally, the system supports role-based access controls and fine-grained data governance policies, ensuring that sensitive enterprise and customer data are protected while maintaining accessibility for authorized analytic teams.

The architecture also demonstrates enhanced efficiency in resource utilization and operational throughput. By leveraging intelligent orchestration and containerized workloads, predictive models and analytic pipelines are automatically scaled based on data volume, processing complexity, and business priorities. The study observed reductions in computational costs while maintaining low latency for predictive outputs, highlighting the cost-effectiveness of the combined lakehouse and MLOps approach. The integration of AI-driven workload prioritization ensures that high-value analytic tasks receive immediate resources while routine or historical analyses are queued efficiently, balancing speed, cost, and operational efficiency.

A notable outcome is the support for advanced analytics and complex predictive scenarios. The unified data environment allows enterprises to combine historical operational metrics, customer behavior, financial indicators, and external market data to generate multifactor predictive models. Ensemble and deep learning models can capture complex interactions between features, enabling accurate forecasting of customer churn, inventory demands, financial risk exposure, and operational bottlenecks. Backtesting and validation across multiple departments demonstrated improved predictive precision and reduced false positives in risk assessment scenarios. The integration of explainable AI techniques further enhanced transparency, allowing business analysts and operational teams to understand the key drivers of model predictions, which in turn supports informed decision-making and strategic planning.

The research also identifies improvements in collaboration and operational agility. By centralizing data and providing automated pipelines, data scientists, engineers, and business analysts can work on consistent datasets with clearly defined feature sets, enabling rapid experimentation, model iteration, and deployment. Shared pipelines reduce duplication of effort, promote reproducibility, and accelerate the transition from prototype models to production-grade analytics. Experimental evaluations show reductions in the time-to-insight and time-to-deployment metrics, indicating enhanced operational responsiveness and the ability to quickly adapt predictive models to evolving enterprise needs.

Despite these benefits, several challenges were identified during implementation. Data heterogeneity, including variations in formats, quality, and update frequency, can complicate ingestion and transformation processes. Maintaining model accuracy across departments and operational domains requires careful feature selection, retraining schedules, and drift monitoring. Additionally, integrating legacy enterprise systems with modern cloud-native architectures demands robust APIs, data connectors, and interoperability protocols. Workforce readiness and training are also essential, as enterprise personnel must adapt to automated workflows, cloud-native operations, and MLOps principles. Addressing these challenges is critical to ensuring the long-term effectiveness, scalability, and reliability of predictive analytics in enterprise environments.

Overall, the results and discussion indicate that intelligent data lakehouse architectures combined with MLOps pipelines provide a transformative solution for scalable predictive analytics in cloud-based enterprise systems. The integration of automated model lifecycle management, unified data storage, cloud-native scalability, and real-time monitoring enables enterprises to achieve high predictive accuracy, operational efficiency, and regulatory compliance, while supporting data-driven decision-making and strategic planning across multiple departments and business units. The study demonstrates that these technologies collectively enable enterprises to harness the full potential of predictive analytics, mitigate operational risks, and respond proactively to evolving business and market conditions.

V. CONCLUSION

The evolution of cloud-based enterprise systems, coupled with the proliferation of diverse operational, transactional, and customer datasets, necessitates a paradigm shift in predictive analytics. Traditional approaches relying on isolated data warehouses, manual ETL processes, and ad hoc model deployment fail to provide the scalability, reproducibility, and agility required in modern enterprises. This research demonstrates that intelligent data lakehouse architectures integrated with MLOps pipelines provide a robust, scalable, and efficient framework for predictive analytics across

| ISSN: 2455-1864 | www.ijrai.org | editor@ijrai.org | A Bimonthly, Scholarly and Peer-Reviewed Journal |

||Volume 8, Issue 6, November–December 2025||

DOI:10.15662/IJRAI.2025.0806036

diverse enterprise operations. The findings indicate substantial improvements in predictive accuracy, operational efficiency, cost optimization, and decision-making agility, establishing a comprehensive solution for cloud-based enterprise analytics that is capable of addressing contemporary business and technological challenges.

A key conclusion is the transformational impact of unifying data storage and processing through a cloud-native lakehouse. By consolidating structured, semi-structured, and unstructured datasets in a single architecture, enterprises gain streamlined access, improved query performance, and consistency across departments. The integration of automated ingestion, schema enforcement, and metadata management ensures data quality and governance, which are essential for reliable predictive analytics and compliance with regulatory requirements. Evaluation results demonstrate reductions in data retrieval latency, improved operational throughput, and enhanced availability of analytics-ready datasets, highlighting the benefits of centralizing and standardizing enterprise data for predictive modeling.

MLOps pipelines represent another critical enabler of scalable predictive analytics. Automation of the entire model lifecycle—from data preparation to deployment, monitoring, and retraining—ensures reproducibility, traceability, and continuous improvement. The research shows that automated pipelines reduce deployment time, mitigate operational errors, and maintain model performance despite changes in data distributions or operational workflows. Continuous monitoring and adaptive retraining allow models to respond to evolving enterprise conditions, supporting proactive decision-making in areas such as inventory management, financial forecasting, customer behavior analysis, and operational risk assessment. MLOps integration further supports cross-departmental collaboration, enabling data scientists, engineers, and business analysts to work cohesively on consistent datasets with standardized features, accelerating innovation and deployment cycles.

The combination of lakehouse architecture and MLOps pipelines also enhances predictive modeling capabilities and analytical sophistication. By enabling access to large-scale, high-dimensional data, enterprises can deploy advanced machine learning techniques, including deep learning and ensemble models, to capture complex patterns and dependencies. The study demonstrates improved predictive accuracy, reduced false positives, and actionable insights for operational planning, strategic decision-making, and risk mitigation. Explainable AI techniques embedded in the workflow provide transparency into model predictions, allowing business stakeholders to understand the drivers of outcomes and increasing confidence in predictive insights. This interpretability is particularly valuable in regulated environments where auditability and accountability are essential.

Scalability, resilience, and operational efficiency emerge as additional key conclusions. Cloud-native orchestration, containerized workloads, and auto-scaling of compute and storage resources allow enterprises to handle high-volume transactional data and concurrent analytics operations efficiently. Fault tolerance, load balancing, and multi-region deployment ensure that predictive analytics remain uninterrupted even during system outages or peak demand periods. AI-driven workload prioritization optimizes resource allocation, ensuring that high-value predictive tasks are executed promptly while routine analytics are managed efficiently. This operational flexibility supports both near-real-time insights and strategic scenario modeling, providing enterprises with the agility to respond effectively to changing market conditions.

The research also emphasizes the importance of data governance, compliance, and security. Intelligent pipelines enforce data validation, lineage tracking, and access control policies, ensuring that enterprise data is accurate, secure, and auditable. Compliance with internal governance policies and regulatory standards is enhanced through automated monitoring and reporting, reducing risk and facilitating transparent operations. By integrating predictive analytics with robust governance frameworks, enterprises can derive business value while maintaining operational integrity and regulatory adherence.

Despite the demonstrated advantages, the study recognizes challenges in implementation, including handling heterogeneous datasets, ensuring interoperability with legacy systems, and maintaining workforce readiness for cloud-native, MLOps-driven workflows. Continuous investment in training, process optimization, and infrastructure management is essential for realizing the full potential of intelligent data lakehouses and automated pipelines. Addressing these challenges will further enhance the robustness, scalability, and reliability of predictive analytics platforms in enterprise contexts.

| ISSN: 2455-1864 | www.ijrai.org | editor@ijrai.org | A Bimonthly, Scholarly and Peer-Reviewed Journal |

||Volume 8, Issue 6, November–December 2025||

DOI:10.15662/IJRAI.2025.0806036

In conclusion, intelligent data lakehouse architectures combined with MLOps pipelines represent a transformative solution for scalable predictive analytics in cloud-based enterprise systems. By unifying data storage, automating model lifecycle management, and leveraging cloud-native scalability, enterprises can achieve high predictive accuracy, operational efficiency, and regulatory compliance. This research demonstrates that the integration of these technologies enables proactive, data-driven decision-making, reduces operational risks, and enhances strategic planning capabilities, providing a comprehensive framework for modern enterprise analytics in complex, dynamic business environments.

VI. FUTURE WORK

Future research can explore several directions to further enhance the capabilities of intelligent data lakehouse and MLOps architectures. One direction involves the integration of federated learning to enable cross-enterprise model training while preserving data privacy, particularly in multi-organization or multi-department environments. Another avenue is the implementation of automated feature engineering and meta-learning pipelines to reduce the manual effort in model preparation and accelerate predictive insights. Research could also focus on hybrid edge-cloud deployments, enabling low-latency, near-real-time analytics for operationally critical workloads, such as financial transactions, manufacturing operations, or IoT-enabled supply chains. Incorporating advanced explainable AI frameworks will further improve interpretability and stakeholder trust in predictions, especially in regulated or high-stakes industries. Additionally, studies on cost optimization, energy-efficient model deployment, and multi-cloud orchestration strategies can guide enterprises in balancing performance, resilience, and operational expenditure while scaling predictive analytics across complex cloud-based systems. Large-scale pilot implementations and longitudinal performance studies will provide empirical validation for best practices in deploying intelligent lakehouse and MLOps pipelines in dynamic enterprise ecosystems.

REFERENCES

1. Ganesan, G. B. K. (2025). Fraud detection systems in enterprise integration architecture. IJSAT-International Journal on Science and Technology, 16(1).
2. Kamadi, S. (2024). GenAI data engineering: Synthetic data and feature engineering framework for cloud analytics. World Journal of Advanced Research and Reviews, 24(1), 2867–2877. https://doi.org/10.30574/wjarr.2024.24.1.3165
3. Sheta, S. V. (2021). Security vulnerabilities in cloud environments. Webology, 18(6), 10043–10063.
4. Ambati, K. C. (2025). Improving user experience and operational efficiency for smarter procurement management. International Journal of Engineering & Extended Technologies Research, 7(3), 1282–1289.
5. Karnam, A. (2024). Engineering trust at scale: How proactive governance and operational health reviews achieved zero service credits for mission-critical SAP customers. International Journal of Humanities and Information Technology, 6(4), 60–67. https://doi.org/10.21590/ijhit.06.04.11
6. Rajasekaran, M., Sekar, S., Manikandaprabhu, K., Vijayakumar, R., Rajmohan, M., & Murugan, S. (2024, October). Next-gen coaching: IoT and linear regression for adaptive training load management. In 2024 8th International Conference on I-SMAC (IoT in Social, Mobile, Analytics and Cloud) (I-SMAC) (pp. 224–229). IEEE.
7. Parathraju, P., & Umasankar, P. (2025). Performance evaluation of ultrathin CdTe-based solar cells with dual absorbers via SCAPS-1D simulation. Scientific Reports, 15(1), 26428.
8. Gowda, M. K. S. (2025). Comprehensive audit data pipeline architecture—Strategies for modern banking audit, compliance and risk management. International Journal of Advanced Research in Computer Science & Technology, 8(1), 11590–11597.
9. Uttama Reddy Sanepalli. (2024). Operationalizing MLOps with Databricks pipelines: Scalable machine learning in cloud environments. International Journal of Scientific Research in Science, Engineering and Technology, 10(6), 2544–2552. https://doi.org/10.32628/CSEIT25113573
10. Vaidya, S., Shah, N., Shah, N., & Shankarmani, R. (2020, May). Real-time object detection for visually challenged people. In 2020 4th International Conference on Intelligent Computing and Control Systems (ICICCS) (pp. 311–316). IEEE.
11. Grandhe, K. (2025). Transforming insight into action: The symbiotic relationship between big data analytics and data visualization. International Journal of Emerging Trends in Computer Science and Information Technology, 125–129.
12. Surampudi, Y. (2024). Big data meets LLMs: A new era of incident monitoring. Libertatem Media Private Limited.

| ISSN: 2455-1864 | www.ijrai.org | editor@ijrai.org | A Bimonthly, Scholarly and Peer-Reviewed Journal |

||Volume 8, Issue 6, November–December 2025||

DOI:10.15662/IJRAI.2025.0806036

13. Panda, S. S. (2024). Delivering scalable cloud services in China: Microsoft and 21Vianet collaboration. International Journal of Advanced Research in Computer Science & Technology, 7(6), 11325–11333.
14. Kumar, R., Mohammed, A. S., & Murthy, C. J. (2023). Cash management forecasting using long short-term memory (LSTM) networks. American Journal of Cognitive Computing and AI Systems, 7, 123–155.
15. Gopinathan, V. R. (2024). AI-driven customer support automation: A hybrid human–machine collaboration model for real-time service delivery. International Journal of Technology, Management and Humanities, 10(1), 67–83.
16. Adari, V. K. (2024). Interoperability and data modernization: Building a connected banking ecosystem. International Journal of Computer Engineering and Technology (IJCET), 15(6), 653–662. https://doi.org/10.5281/zenodo.14219429
17. Ande, B. R. (2025). Federated learning and explainable AI for decentralized fraud detection in financial systems. Journal of Information Systems Engineering and Management, 10(35s), 48–56.
18. Nagarajan, C., & Madheswaran, M. (2011). Stability analysis of series parallel resonant converter with fuzzy logic controller using state space techniques. Electric Power Components and Systems, 39(8), 780–793.
19. Nagarajan, C., & Madheswaran, M. (2011). Performance analysis of LCL-T resonant converter with fuzzy/PID using state space analysis. Electrical Engineering, 93(3), 167–178.
20. Jagadeesh, S., & Soundappan, R. S. (2014). Survey on knowledge discovery in speech emotion detection. International Journal of Innovative Research in Computer and Communication Engineering, 2(5), 4476–4481.
21. Mulla, F. A. (2024). Building scalable mobile applications: A comprehensive guide to shared component architecture. International Journal of Computer Engineering and Technology, 15, 1337–1348.
22. Kiran, A., Rubini, P., & Kumar, S. S. (2025). Comprehensive review of privacy, utility and fairness offered by synthetic data. IEEE Access.
23. Raju, S., & Sindhuja, D. (2024). Transparent encryption for external storage media with mobile-compatible key management by Crypto Ciphershield. PatternIQ Mining, 1(3), 12–24.
24. Charumathi, M. V., & Inbavalli, M. (n.d.). Familiarizing the pine nut oil by fusing it into different food products. PG and Research Department of Foods & Nutrition, Marudhar Kesari Jain College for Women.
25. Ravi Kumar Ireddy. (2024). Real-time payment orchestration and fraud governance framework: Cloud-native treasury optimization with ensemble deep learning integration. International Journal of Scientific Research in Computer Science, Engineering and Information Technology, 10(3), 1152–1161. https://doi.org/10.32628/CSEIT25113583
26. Gaddapuri, N. S. (2025). Cloud-Native Twin Systems for Real-Time Risk and Compliance Simulation in FinHealth Converged Ecosystems. ISCSITR-INTERNATIONAL JOURNAL OF COMPUTER SCIENCE AND ENGINEERING (ISCSITR-IJCSE)-ISSN: 3067-7394, 6(4), 77-94.
27. Sampath Kumar Konda. (2024). Distributed AI infrastructure orchestration: A hyperscale multi-cloud framework for geographic load balancing with renewable energy optimization. International Journal of Scientific Research in Science Engineering and Technology, 11(4), 522–533. https://doi.org/10.32628/IJSRSET242438
28. Nallamothu, T. K. (2024). Empowering clinicians through AI-augmented documentation: Insights from Dragon Copilot implementation. International Journal of Advanced Research in Computer Science & Technology, 7(6), 11309–11318.
29. Bapatla, S. K. S. (2025). AI-powered physician-insurance data mapping: A case study in reducing revenue leakage. Journal of Computer Science and Technology Studies, 7(7), 550–559.
30. Archana, R., & Anand, L. (2025). Residual U-Net with self-attention based deep convolutional adaptive capsule network for liver cancer segmentation and classification. Biomedical Signal Processing and Control, 105, 107665.
31. Yashwanth, K., Adithya, N., Sivaraman, R., Janakiraman, S., & Rengarajan, A. (2021, July). Design and development of pipelined computational unit for high-speed processors. In 2021 12th International Conference on Computing Communication and Networking Technologies (ICCCNT) (pp. 1–5). IEEE.
32. Sundareswaran, A. P., Gupta, A., Srinivas, S., Athamakuri, S. S. K. K., Singh, K., & Sharma, R. K. (2025, August). Data Quality Assurance in Cloud-Based Warehousing Systems. In 2025 International Conference on Intelligent and Secure Engineering Solutions (CISES) (pp. 939-944). IEEE.