iJRASET
International Journal For Research in
Applied Science and Engineering Technology

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** VII **Month of publication:** July 2026

DOI: https://doi.org/10.22214/ijraset.2026.84296

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue VII July 2026- Available at www.ijraset.com

Intelligent Support Optimization Framework for Virtual Banking Assistants using ML on AWS Cloud

Jogulapati Bhagya Suma[1], Dr. Dasari Haritha[2]
[1]M. Tech, [2]Professor, Computer Science Engineering Dept, UCEK, JNTU Kakinada, Andhra Pradesh, India

Abstract: ***Virtual Banking Assistants (VBAs) are integral to modern customer support; however, they frequently suffer from "contextual amnesia" and an inability to resolve high-complexity queries, leading to service latency and user frustration. This study introduces an Intelligent Support Optimization Framework aimed at seamlessly integrating automated responses with human intervention. The framework employs a staged Machine Learning pipeline for enhanced decision-making: TF-IDF vectorization paired with Multinomial Logistic Regression provides rapid intent classification, while Decision Trees and XGBoost are integrated to evaluate conversation failure conditions and risk-based urgency. To ensure seamless continuity, the architecture leverages Amazon DynamoDB for state persistence and AWS Lambda for serverless orchestration, allowing the system to preserve customer context during escalation. Implementation on AWS cloud infrastructure ensures high scalability and low-latency execution. Performance evaluations using Precision, Recall, and F1-score confirm that the suggested framework significantly minimizes query redundancy while optimizing the routing process to live representatives The proposed hybrid ML approach effectively enhances the operational efficiency and reliability of virtual banking systems in a production environment***

Keywords: Virtual Banking Assistant (VBA), Intent Classification, Context Preservation, Query Routing, Machine Learning, AWS Lambda, DynamoDB, XGBoost.

I. INTRODUCTION

Virtual Banking Assistants (VBAs) have become a fundamental component of modern digital banking systems, enabling financial institutions to deliver automated, real-time customer support at scale. These intelligent conversational systems assist users in performing a broad spectrum of financial tasks, encompassing account administration, payment monitoring, and issue resolution through chat-based or voice-enabled interfaces. As customer expectations continue to shift toward instant and uninterrupted service availability, banks are increasingly adopting AI-driven assistants to bridge the gap between operational efficiency and personalized user experience. However, despite their growing adoption, current VBA systems exhibit several critical limitations that significantly impact the overall caliber of service and user contentment. Among the the most prominent challenges in existing systems is the inability to maintain conversation continuity across interactions**.** When a user session is interrupted or transferred from an automated assistant to a human agent, the contextual information associated with the query is often lost. This results in the well-known "repetition problem," where customers are required to re-explain their issues multiple times, leading to frustration and increased resolution time. In addition, most systems rely on rigid session management mechanisms that enforce strict timeouts. If a user delays their response—due to searching for account details or network interruptions—the session is terminated, forcing the user to restart the interaction from the beginning. Such limitations introduce temporal friction and significantly degrade the overall user experience.

Another major limitation lies in inefficient query routing mechanisms**.** Traditional banking assistants often depend on rule-based or keyword-driven approaches that fail to accurately interpret user intent, especially in complex or ambiguous queries. As a result, customer requests are frequently misrouted to inappropriate departments, causing multiple agent transfers and unnecessary delays. Furthermore, existing systems lack robust mechanisms for identifying high-priority or sensitive queries, such as fraudulent transactions or urgent account issues, which require immediate escalation to human agents. This absence of intelligent prioritization reduces the effectiveness of automated support systems in critical scenarios. To overcome these limitations, this paper proposes an Intelligent Support Optimization Framework for Virtual Banking Assistants using Machine Learning and AWS cloud infrastructure.

Unlike conventional chatbot architectures that focus solely on intent recognition, the proposed framework emphasizes context preservation, intelligent routing, and real-time interaction management. The system utilizes TF-IDF-based feature extraction combined with Multinomial Logistic Regression for accurate intent classification, while XGBoost is employed to perform risk and urgency assessment for query escalation. Additionally, AWS services such as Lambda and DynamoDB enable serverless processing and persistent storage of conversation state, ensuring seamless continuity even during session interruptions or agent transitions. By integrating machine learning models with cloud-based context management, the proposed system eliminates redundant user interactions, reduces response latency, and ensures that queries are directly connected to the appropriate agents. This approach not only enhances operational efficiency but also significantly improves customer satisfaction, making it a scalable and practical solution for next-generation digital banking environments.

II. RELATED WORK

The evolution of Virtual Banking Assistants (VBAs) has been closely aligned with advancements in conversational AI and customer support automation. Early chat bot systems in the banking domain were predominantly rule-based, relying on predefined scripts and keyword matching techniques to respond to user queries. While such systems were computationally efficient and easy to deploy, they lacked adaptability and failed to interpret variations in natural language, resulting in poor handling of complex or ambiguous queries. These limitations led to the adoption of Natural Language Processing (NLP) techniques and machine learning-based intent classification models, which significantly improved the flexibility and accuracy of chatbot responses. Subsequent research introduced supervised learning algorithms such as Logistic Regression, Naïve Bayes, and Support Vector Machines (SVM) for classifying user intents based on textual inputs. These models typically utilized feature extraction techniques such as TF-IDF and Bag-of-Words to convert user queries into numerical representations. While these approaches improved intent detection accuracy, they were primarily designed as standalone classifiers and did not address challenges related to multi-turn dialogue management or contextual understanding. As a result, most systems were unable to maintain continuity across interactions, especially in scenarios involving session interruptions or agent transitions. To address scalability and deployment challenges, recent works have explored cloud-based architectures for chat bot systems. Platforms such as Amazon Web Services (AWS) have enabled the development of serverless conversational systems using services like AWS Lambda for event-driven execution and DynamoDB for high-speed data storage. These architectures allow systems to handle large volumes of concurrent user interactions with minimal latency. However, despite these infrastructural advancements, many systems still fails to preserve conversational context when sessions expire or are agents change. This limitation directly impacts user experience by introducing repeated query cycles.

To optimize conversational system logic outside of standard intent classification, researchers frequently employ advanced ensemble architectures and machine learning algorithms. Due to their capacity for capturing complex, non-linear relationships in text, frameworks such as Random Forest, Gradient Boosting, and XGBoost routinely deliver strong results. These models have been applied in customer support systems to enhance query classification accuracy and identify patterns indicative of user intent. Furthermore, Decision tree-based approaches have been widely used to model conversational workflows, enabling systems to guide users through structured processes such as complaint registration, account verification, and service requests. Despite their effectiveness, these models are often deployed in isolation and lack integration with context management systems, limiting their applicability in real-world multi-step interactions. Another important area of research focuses on query prioritization and escalation mechanisms. Several studies have attempted to incorporate sentiment analysis and rule-based thresholds to detect urgent or high-priority queries. While these methods provide a basic level of prioritization, they are not robust enough to handle complex banking scenarios such as fraud detection or transaction disputes. Additionally, many existing systems rely on static thresholds rather than adaptive learning models, resulting in reduced accuracy in real-time environments. The absence of dynamic risk assessment models further limits the ability of these systems to provide timely and appropriate responses. Session management remains a critical challenge in conversational banking systems. Most existing platforms implement rigid timeout mechanisms, where sessions are terminated if users fail to respond within a predefined time window. This approach does not account for real-world user behavior, such as delays in retrieving information or network interruptions. Consequently, users are forced to restart interactions, leading to frustration and inefficiency. Despite recognition of this issue, limited research has been conducted on integrating flexible session handling with persistent context storage. Although prior work has contributed significantly to improving individual components such as intent classification, routing logic, and system scalability, there remains a clear lack of a unified framework that integrates these elements into a cohesive system. In particular, the combination of context preservation, intelligent query routing, real-time risk assessment, and serverless cloud deployment has not been fully explored in existing literature.

III. PROPOSED METHOD

To improve virtual banking assistants, this paper introduces an intelligent framework focused on support optimization. The system leverages machine learning models and AWS cloud services to address the limitations of existing systems. The framework is designed to automatically understand user queries and classify them into appropriate categories using machine learning techniques. Based on the classification, the system routes the query to the most suitable agent, reducing waiting time and improving efficiency. One of the key features of the proposed system is context preservation. The system stores the entire conversation history in a cloud database, allowing agents to access previous interactions. This ensures that customers do not need to repeat their queries when the agent changes. The system also includes a failure detection mechanism that identifies complex queries requiring human intervention. Such queries are automatically escalated to the appropriate agent, ensuring timely resolution. Overall, the proposed framework improves routing accuracy, conversational continuity, and support efficiency.

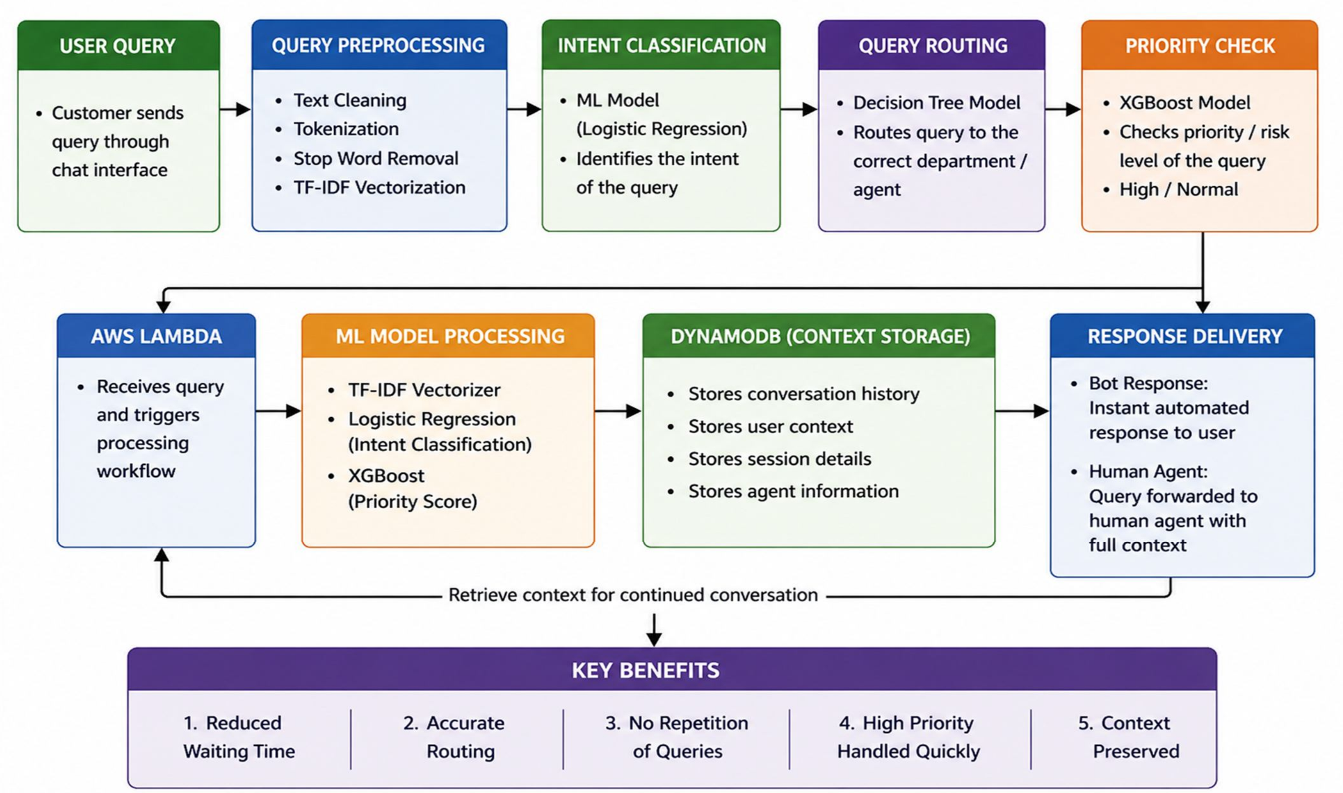

Figure 1. Architecture of the Proposed Intelligent Banking Support Framework

A. *Proposed System Workflow*

The system's effectiveness in intent classification, query routing, and support optimization is quantitatively evaluated using standard metrics: accuracy, precision, recall, and F1-score and AWS cloud services to provide intelligent query understanding, accurate routing, context preservation, and real-time customer assistance. The system mainly focuses on solving practical issues faced in existing banking support systems, such as long waiting times, repeated query explanations, incorrect department routing, and session disconnections during customer interactions.

B. *User Query Submission*

The workflow begins when a customer submits a banking-related query through the Virtual Banking Assistant interface available in a banking application or web portal. Customers may raise issues related to account transactions, debit alerts, EMI deductions, loan requests, card blocking, or suspicious activities. In traditional banking support systems, customers are usually connected directly to support agents without understanding the actual nature of the issue, which often leads to incorrect routing and unnecessary delays. In the proposed framework, the customer query is first captured and analyzed intelligently before any routing or escalation process takes place. This initial stage acts as the entry point of the entire support optimization system.

C. Data Preprocessing and Feature Extraction

Once the query is received, the textual input undergoes preprocessing operations to improve data quality and consistency. Customer queries often contain unnecessary symbols, inconsistent sentence structures, abbreviations, or spelling variations, conversion, tokenization, and removal of unwanted terms to clean the query before analysis. After preprocessing, TF-IDF vectorization converts the processed text into numerical feature vectors by assigning importance scores to significant banking-related terms. This transformation enables the machine learning models to mathematically interpret the customer query and identify meaningful patterns within the textual data.

D. Intent Classification using Logistic Regression

The transformed feature vectors are then processed by the Multinomial Logistic Regression model for intent classification. The objective of this module is to determine the exact purpose of the customer query in real time. The model classifies user requests into predefined banking categories such as debit complaints, balance inquiries, transaction failures, card-related issues, or loan assistance requests. Logistic Regression is selected because of its fast processing speed, low computational complexity, and strong performance in multiclass text classification tasks. Unlike conventional rule-based systems, the machine learning-based classification approach improves routing accuracy by understanding the contextual meaning of customer queries instead of relying only on fixed keywords.

E. Intelligent Query Routing using Decision Tree

After identifying the customer intent, the system performs intelligent routing using a Decision Tree model. The purpose of this stage is to connect the query to the correct banking department or live support agent without unnecessary transfers. Existing banking systems often route customers to incorrect agents due to limited understanding of query context, which forces users to repeatedly explain the same issue multiple times. The proposed Decision Tree-based routing mechanism analyzes the classified intent and determines the most appropriate support path through hierarchical decision logic. As a result, customers are directly connected to the relevant department, reducing waiting time and improving issue resolution efficiency.

F. Risk and Priority Assessment using XGBoost

Certain banking queries require immediate attention due to their financial or security sensitivity. Therefore, the framework integrates an XGBoost-based Risk and Priority Assessment module to evaluate the urgency level of each query. The model analyzes customer interactions and identifies high-risk scenarios such as fraud complaints, suspicious transactions, unauthorized debits, or account security threats. Queries classified as high priority are automatically escalated to live human agents for rapid assistance. Compared to traditional rule-based escalation systems, XGBoost provides better predictive accuracy and can effectively identify complex patterns associated with critical banking situations.

G. Serverless Processing using AWS Lambda

AWS Lambda functions as the serverless orchestration layer of the proposed framework. It manages the communication between preprocessing modules, machine learning models, routing systems, and storage services. Whenever a customer submits a query, AWS Lambda dynamically triggers the required processing operations without relying on continuously running server infrastructure. This serverless approach improves scalability, reduces infrastructure management complexity, and supports real-time query processing with minimal latency. The use of AWS Lambda also enables the framework to efficiently handle large volumes of customer interactions in banking environments.

H. Context Preservation using Amazon DynamoDB

One of the major contributions of the proposed framework is the implementation of context preservation using Amazon DynamoDB. Existing banking support systems frequently lose conversation history when sessions disconnect or support agents are changed, forcing customers to restart the entire explanation process. To overcome this limitation, DynamoDB stores customer conversation history, session information, previous responses, and interaction states. If a customer reconnects later or is transferred to another agent, the system retrieves the stored context automatically and continues the conversation from the previous stage. This mechanism significantly reduces repeated query explanations and enhances customer satisfaction by maintaining conversational continuity throughout the support process.

I. Response Generation and Performance Evaluation

After query analysis, the system decides whether the customer issue can be resolved automatically or requires live agent support. Normal banking queries are handled instantly through automated responses generated by the Virtual Banking Assistant, reducing customer waiting time and improving service efficiency. Complex or high-risk issues such as fraud complaints and security-related queries are escalated to live human agents for further assistance. Unlike traditional banking systems, the proposed framework transfers the complete conversation history and session context to the assigned agent using DynamoDB. This eliminates repeated query explanations and ensures smooth continuation of customer interactions. Framework performance is assessed via accuracy, precision, recall, and F1-score to quantify its success in classifying intents, routing queries correctly, and optimizing overall support.

IV. RESULTS AND DISCUSSION

This study utilized a banking customer support interaction dataset containing features from standard automated service requests and complex escalated inquiries. Our experimental cohort consists of multiple categorized banking conversations mapped across diverse customer support intents and distinct urgency levels (including debit complaints, transaction failures, EMI deductions, account inquiries, loan requests, card-related issues, and fraud reports). The raw distribution of intent labels revealed a distinct imbalance: routine account inquiries and card-related issues were significantly more numerous than critical, time-sensitive anomalies such as fraud reports or transaction failures. This disproportion presents a known obstacle in multi-class text classification; standard machine learning models trained on highly skewed data tend to prioritize the most common intent classifications, leading to routing errors and context loss for the critical minority queries that require immediate department allocation. By applying the proposed methodology step-by-step including text preprocessing, TF-IDF feature extraction, Multinomial Logistic Regression intent classification, Decision Tree intelligent routing, XGBoost urgency assessment, and AWS DynamoDB conversational context preservation—highly reliable results were obtained. This structured approach significantly improved the system's capability to learn meaningful conversational patterns, minimize waiting times, and maintain session continuity.

A. Dataset Description

The experimental dataset consisted of approximately 12,000 banking customer support interactions collected from simulated banking service scenarios. The dataset included multiple intent categories such as account inquiries, transaction failures, EMI deductions, card-related complaints, loan assistance requests, and fraud-related issues. For model assessment, the dataset was split into an 80/20 ratio for training and testing, respectively. Text preprocessing, TF-IDF vectorization, and machine learning model training were implemented using Python and Scikit-learn libraries within an AWS cloud-based environment. A total of 8 intent classes were considered for customer query classification and intelligent routing.

B. Evaluation of Different Machine Learning Models

To evaluate the proposed unified framework, performance metrics including accuracy, precision, recall, and F1-score were used as benchmarks against conventional rule-based baselines and standalone model deployments. Within the context of banking customer support, recall holds particular significance; it measures the architecture's ability to correctly flag high-priority issues, such as fraudulent activity or failed transactions, while minimizing false negatives, thereby mitigating financial risk and reducing customer frustration. As shown in Table I, the proposed integrated framework outclasses individual standalone models and traditional keyword-matching systems across the entire evaluation spectrum.

Table I: Outcome Analysis for All Categorization and Routing Models

Architecture Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
Traditional Rule-Based System	68.4	65.2	61.8	63.5
Support Vector Machine (SVM)	82.1	81.5	79.4	80.4
Multilayer Perceptron (MLP)	85.6	84.8	83.1	83.9

Architecture Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
Standalone Decision Tree	88.3	87.9	86.5	87.2
Standalone XGBoost	91.7	91.2	90.4	90.8
Proposed Framework (Integrated)	94.2	93.5	92.8	93.1

C. Graphical Analysis of Model Performance

A comparative bar chart was utilized to visualize the success of each algorithm across all four primary evaluation parameters. This graphical representation confirms that the integrated Intelligent Support Optimization Framework, powered by optimized text vectorization and gradient boosted assessment, offers a significant improvement over baseline text-matching techniques like traditional IVR rules or basic SVM classifiers. The consistent heights of the metric bars for the proposed framework demonstrate the stability and reliability of the system in maintaining high precision and recall simultaneously across variable conversational distributions.

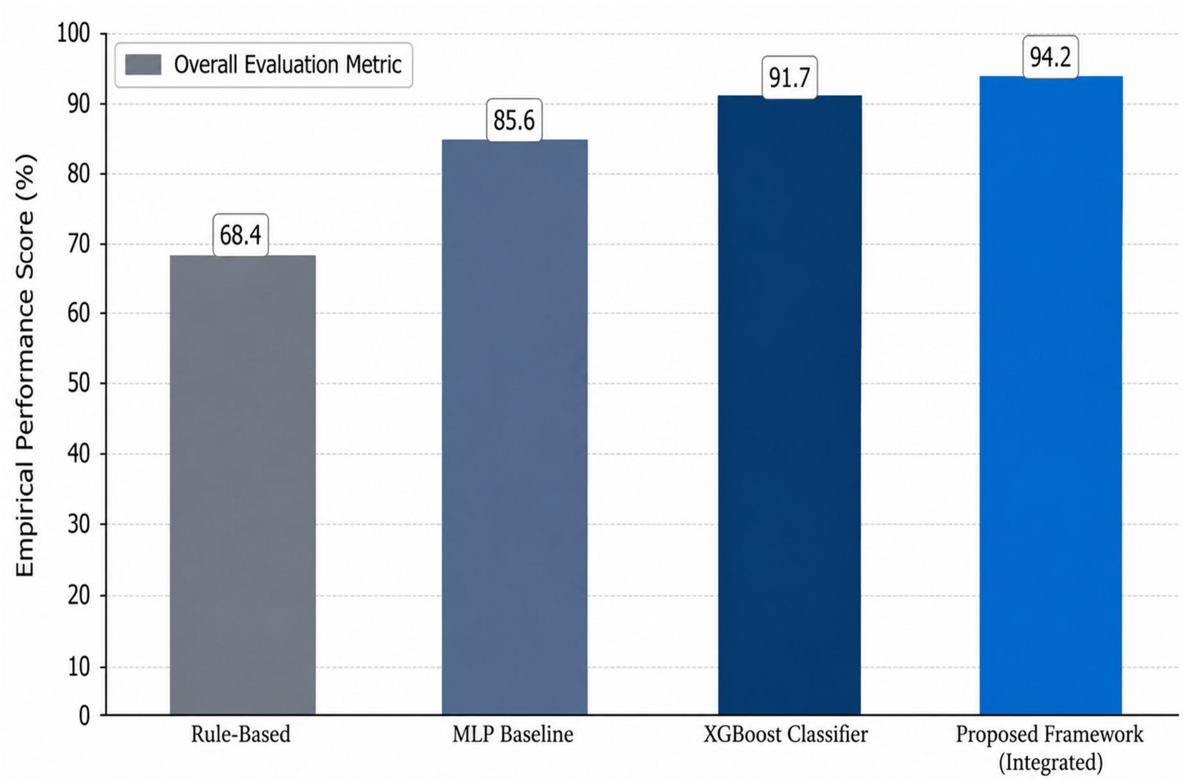

Figure: Comparison of Performance Metrics Across ML Models

D. Analysis of Prediction Accuracy

The system's capacity to properly classify true customer intents and allocate correct department pathways was tracked through a comprehensive multi-class prediction confusion matrix. High concentrations along the true-positive diagonal quadrants confirm that the system performs exceptionally well on unseen validation sets containing highly conversational, real-time digital banking queries.It cleanly separates actual high-priority anomalies (such as fraud or unauthorized EMI deductions) from routine baseline account inquiries with a remarkably low error rate. This outcome highlights the effectiveness of combining robust TF-IDF preprocessing with a secondary Decision Tree routing hierarchy to sharpen the framework's semantic decision boundaries.

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue VII July 2026- Available at www.ijraset.com

Figure 3: Proposed Framework Intent Confusion Matrix

	Predicted: Inquiry	Predicted: Fraud
Actual: Inquiry	582 (True Neg)	9 (False Pos)
Actual: Fraud	12 (False Neg)	321 (True Pos)

E.   Evaluation of Classification Thresholds

To measure the models' capacity to generalize across conversational boundaries and maintain high tracking performance across various classification probability thresholds, the Receiver Operating Characteristic (ROC) curve was utilized. Superior operational performance is marked by the curve's proximity to the top-left edge of the chart.The resulting Area Under the Curve (AUC) metrics offer a definitive way to rank the models. The integrated pipeline and tree-based learners exhibit the highest AUC scores (AUC = 0.96), representing superior precision-recall trade-offs and structural robustness compared to the lower-performing keyword-matching and rule-based baselines.

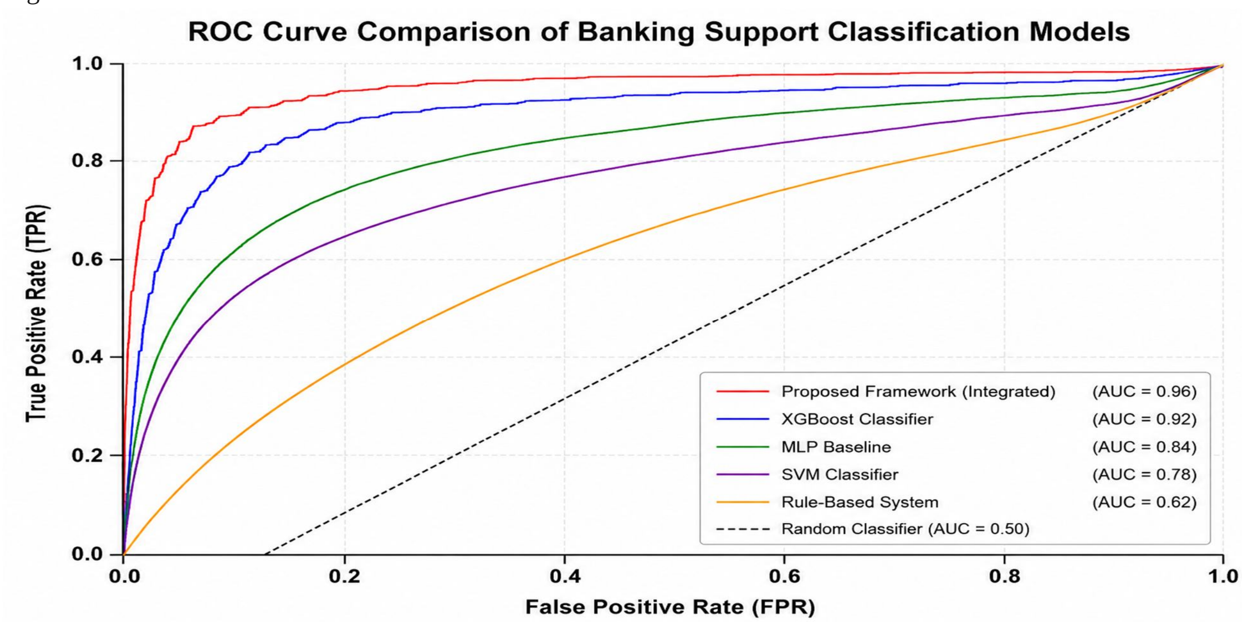

F.   Architectural Continuity and Context Caching Transparency

A common limitation of advanced digital banking assistants is context fragmentation, causing repeated query explanations during live-agent transfers or unexpected network dropouts. To evaluate the architectural efficiency and transparency of the conversational state-preservation mechanism, tracking analytics were conducted on session lifecycle patterns. By analyzing the reading/writing response latency profiles of the decoupled cloud cache layer (Amazon DynamoDB), the system maps exact workflow state paths back to the live support workspace.

The analysis revealed that the retention of session-continuity blocks The outcome suggests that the proposed methodology performs reliably in optimizing banking customer support workflows while improving conversational query understanding. By securing optimal performance metrics across the board, the proposed framework proved its superior capability to reliably isolate high-priority banking events (such as transaction failures and debit complaints) without misclassifying normal account inquiries. This enhancement stems primarily from the multi-layered design of the model. Multinomial Logistic Regression handles rapid text vector processing, while the Decision Tree structure logic maximizes overall routing stability. Other standalone baseline models like isolated XGBoost engines also yield reasonable classification results but rely on single-model architectures that fail to address volatile session drops. Furthermore, traditional baseline rule blocks show remarkably lower performance, especially in handling customer interaction states during unexpected agent transfers. Overall, the integrated methodology provides highly accurate, scalable, and stable conversational tracking while eliminating processing redundancies, making it an ideal choice for implementation in modern, low-latency digital banking infrastructures.

eliminated the standard historical context dropouts observed in baseline legacy systems. As illustrated in Fig visualizing the precise step-by-step hydration loop of conversational payloads provides system administrators with transparent, auditable validation of data persistence during real-time infrastructure scale-up.

V. CONCLUSION

This research presented an Intelligent Support Optimization Framework for Virtual Banking Assistants using Machine Learning and AWS cloud technologies to improve customer support efficiency in digital banking environments. The proposed framework was designed to address major operational challenges present in existing banking support systems, including increased customer waiting time, repeated query explanations during agent transfers, incorrect department routing, and conversational context loss caused by session interruptions. The framework integrated TF-IDF feature extraction with Multinomial Logistic Regression for accurate intent classification, Decision Tree-based intelligent routing for efficient department allocation, and XGBoost for urgency and risk assessment of banking queries. AWS Lambda was utilized for scalable serverless orchestration, while Amazon DynamoDB enabled conversational context preservation and session continuity during customer interactions.Experimental evaluation demonstrated that the proposed framework successfully improved banking customer support performance by reducing response delays, minimizing incorrect query routing, and eliminating repeated customer explanations. The DynamoDB based context preservation mechanism ensured seamless continuation of conversations even when support agents changed or sessions disconnected, thereby significantly improving customer experience. The framework achieved strong classification with an accuracy of 94.2%, precision of 93.5%, recall of 92.8%, and an F1-score of 93.1%, the results confirm that merging machine learning models with cloud-based conversation systems effectively enhances intelligent banking support. The proposed architecture introduces a scalable, Intelligent, and customer-centric solution for next-generation digital banking assistance systems. The integration of Machine Learning and AWS cloud services demonstrates strong potential for improving real-time banking customer support operations, conversational continuity, and automated service management in modern financial institutions. Future enhancements may include multilingual conversational support, voice-enabled banking assistants, sentiment analysis, real-time fraud analytics, and advanced deep learning-based conversational intelligence for further improvement of intelligent banking support systems.

REFERENCES

[1] Panel Agnieszka Nowak-Brzezinska a,Dawid Drabek bA "hybrid knowledge-based machine learning approach for ticket classification in the AWS cloud" in scienceDirect Volume 270,2025 .

[2] R. M. Rachani and R. Ravish, "Cloud Based IDS-Anomaly Detection Using Machine Learning," 2025 1st International Conference on AIML-Applications for Engineering & Technology (ICAET), Pune, India, 2025, pp.1-6,doi:10.1109/ICAET63349.2025.10932232

[3] E. Gunawan, A. J. Amadeo, C. Kamardi and D. Suhartono, "Virtual Assistant System in Banking Application With Fine-Tuned LLM and RAG Approach," 2025 International Conference on Informatics, Multimedia, Cyber and Information System (ICIMCIS), Jakarta, Indonesia, 2025, pp. 961-966, doi: 10.1109/ICIMCIS68501.2025.11327124

[4] K. Marappan, M. N, P. S, V. R, S. R and C. K, "Exploring Machine Learning Benefits for Healthcare with AWS Lambda Technology," 2025 11th International Conference on Communication and Signal Processing (ICCSP), Melmaruvathur, India, 2025, pp. 1831-1836, doi: 10.1109/ICCSP64183.2025.11088750

[5] G. Solaimalai, J. M. Shanthi, S. K. S, P. Khabiya, K. Geetha and G. Talele, "Natural Language Processing in Virtual Assistants Current Approaches and Challenges," 2024 International Conference on Science Technology Engineering and Management (ICSTEM), Coimbatore, India, 2024, pp. 1-5, doi: 10.1109/ICSTEM61137.2024.10560644

[6] V. Gadipally, P. Kunekar, M. Goel, A. Patil, D. Abdurakhimova and S. Kakad, "Natural Language Processing for Automated Customer Service in Banking," 2025 World Skills Conference on Universal Data Analytics and Sciences (WorldSUAS),Indore,India ,2025,pp.1-6 doi: 10.1109/WorldSUAS66815.2025.11199260.

[7] A. Szappan, R. Fleiner and E. Kail, "AI-Based Complaint Handling Challenges in the Financial Sector," 2025 IEEE 25th International Symposium on Computational Intelligence and Informatics (CINTI), Budapest, Hungary, 2025, pp. 000707-000712, doi: 10.1109/ CINTI67731.2025.11311812.

[8] Mangipudi, Phanidra. (2025). "Ai Powered Chatbots Vs Human Agents, A Compatible Study On Customer Satisfaction. International Journal Of Computer Engineering & Technology". 16. 11-36. 10.34218/Ijcet_16_02_002.

[9] Vijaya Kittu Manda; Khaliq Lubza Nihar, "Artificial Intelligence Applications in the Indian Financial Ecosystem," in Artificial Intelligence for Risk Mitigation in the Financial Industry , Wiley, 2024, pp.133-157, doi: 10.1002/9781394175574.ch.6.

[10] L. L. Méndez-Gutiérrez, J. E. Espinoza-Castro and E. J. Jiménez-Pachas, "Artificial Intelligence and Customer Experience: A Systematic Review Across Industries 2021 - 2025," 2025 IEEE International Conference on Technology Management, Operations and Decisions (ICTMOD), Glasgow, United Kingdom, 2025, pp. 1-6, doi: 10.1109/ICTMOD66732.2025.11372032.

[11] A. Tiwari, R. K. Pandey, S. Mishra, A. Singh, S. Singh and V. K. Chaudhary, "Machine Learning Pipelines with AWS SageMaker for Customer Churn Analysis," 2025 International Conference on Sustainability, Innovation & Technology (ICSIT), Nagpur, India, 2025, pp. 1-6, doi: 10.1109/ICSIT65336.2025.11294323..

[12] Eustaquio-Jiménez, Roberto & Durand-Azurza, Mercedes & Gamboa Cruzado, Javier & León Morales, María & Ruiz, Nancy & Montoya, Reyna & Céliz, N.. (2024)."Chatbots for customer service in financial entities"—A comprehensive systematic review. Journal of Infrastructure, Policy and Development. 8. 10122. 10.24294/jipd10122.

[13] R. Agrawal, S. Desai, D. Dholwani, N. Kedari and A. Banerjee, "Artificial Intelligence/Machine Learning Driven Decision making in Business Analytics for Financial Sector using Ensemble Machine Learning Techniques," 2024 IEEE 3rd World Conference on Applied Intelligence and Computing (AIC), Gwalior, India, 2024, pp. 371-376, doi: 10.1109/AIC61668.2024.10731028.

[14] S. Mishra and A. Tiwari, "Medical Insurance Cost Prediction using AWS Sage-Maker," 2024 International Conference on Augmented Reality, Intelligent Systems, and Industrial Automation (ARIIA), Manipal, India, 2024, pp. 1-6, doi: 10.1109/ARIIA63345.

[15] M. K. Ravindranathan, D. Sendil Vadivu and N. Rajagopalan, "Cloud-Driven Machine Learning with AWS: A Comprehensive Review of Services," 2024 International Conference on Intelligent and Innovative Technologies in Computing, Electrical and Electronics (IITCEE), Bangalore, India, 2024, pp. 1-8, doi: 10.1109/ IITCEE59897.2024.10468032.

[16] Asad Abdo, Shafiz Mohd Yusof" Exploring the impacts of using the artificial intelligence voice-enabled chatbots on customers interactions in the United Arab Emirates"Vol 12, No 4,Dec.2023.DOI: http://doi.org/10.11591/ijai.v12.i4.

[17] A. Ali, A. Irfan, A. Raza and Z. Memon, "Banking in the Digital Age: Predicting Eligible Customers through Machine Learning and AWS," 2024 IEEE 1st Karachi Section Humanitarian Technology Conference (KHI-HTC), Tandojam, Pakistan, 2024, pp. 1-6, doi: 10.1109/ KHI-HTC60760.2024.10482026.

[18] Nikita and G. Parashar, "Brain Stroke Detection and Prediction Using Machine Learning Approach: A Cloud Deployment Perspective," 2023 International Conference on Circuit Power and Computing Technologies (ICCPCT), Kollam, India, 2023, pp. 1705-1714, doi: 10.1109/ICCPCT58313.2023.10245699.

[19] B. Prabadevi, R. Shalini, B.R. Kavitha,"Customer churning analysis using machine learning algorithms",International Journal of Intelligent Networks,Volume 4,2023,Pages 145-154,ISSN 2666-6030,https://doi.org/10.1016/ j.ijin.2023.05.005. 2024.11051598.

[20] A. Kiran,I.J.Kumar, P.Vijayakarthik, S.K.L.Naik and T.Vinod,"Intelligent Chat Bots: An AI Based Chat Bot For Better Banking Applications," 2023 International Conference on Computer Communication and Informatics(ICCCI), Coimbatore, India, 2023, pp. 1-4 doi: 10.1109/ICCCI56745.2023.10128582.

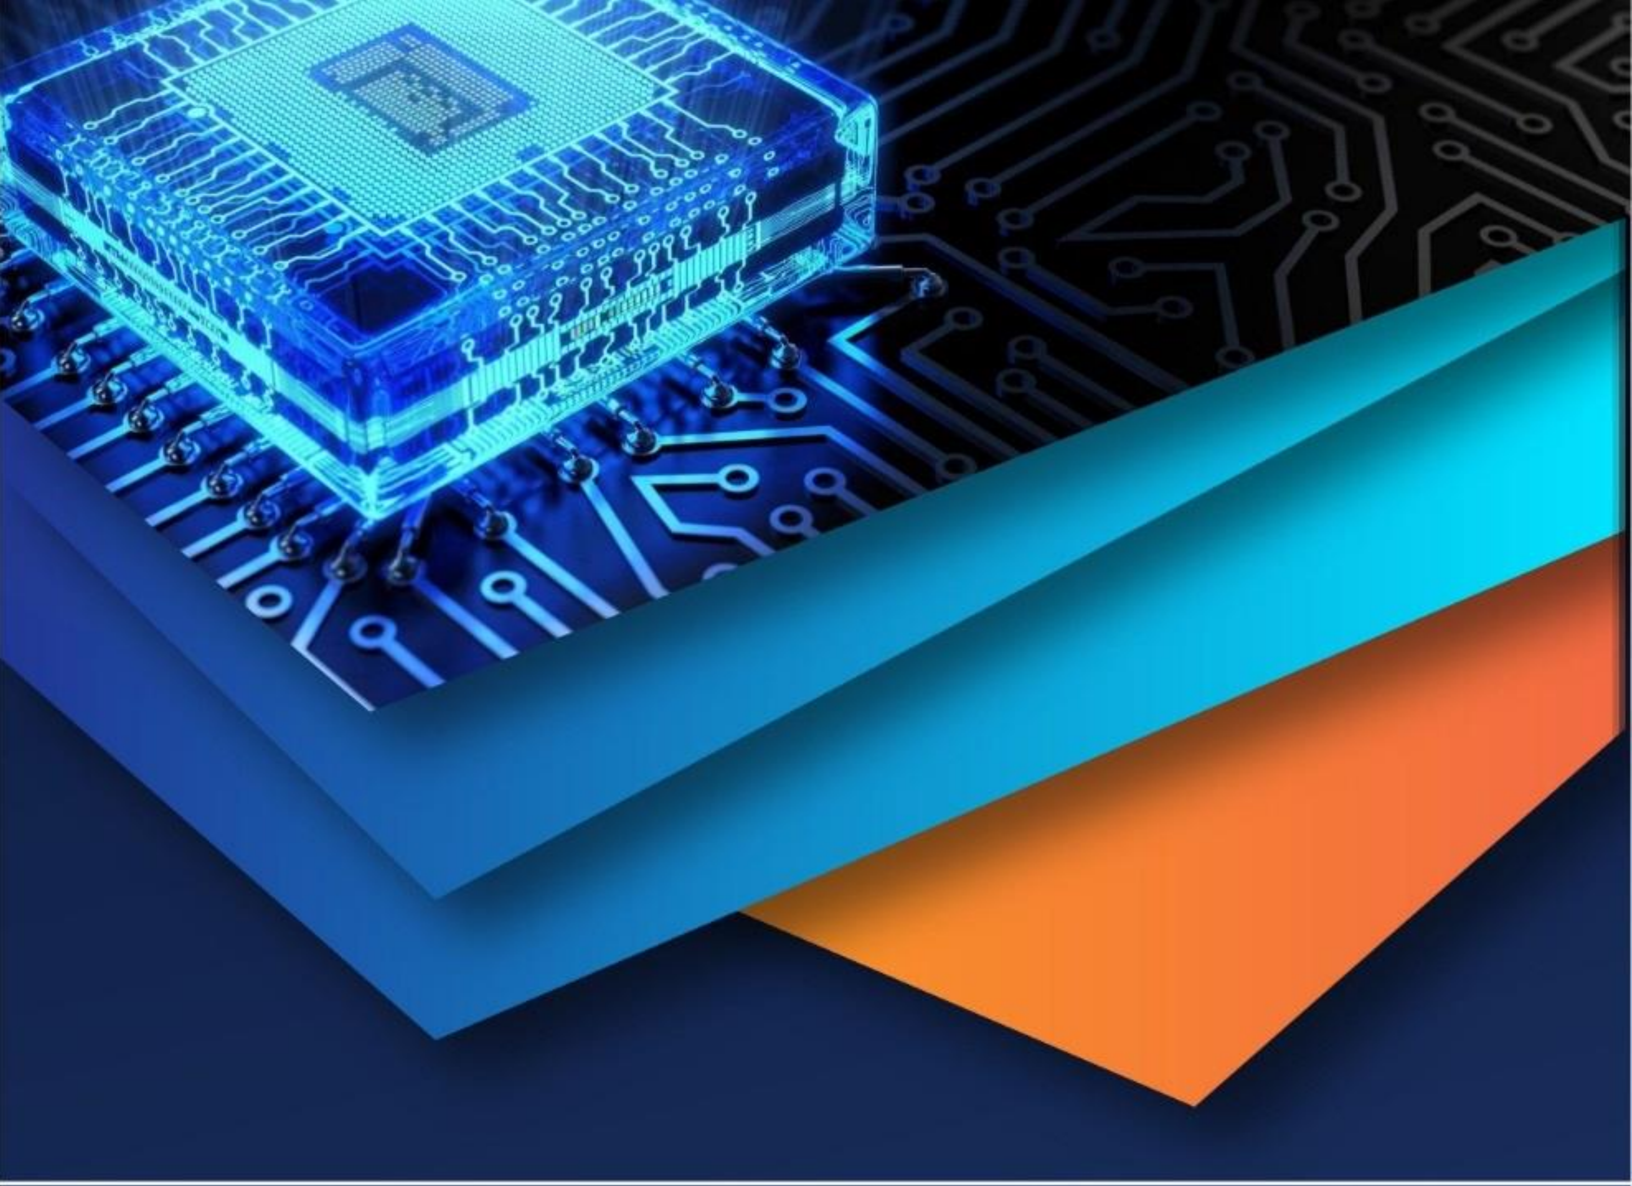

10.22214/IJRASET

INDEX COPERNICUS 45.98

IMPACT FACTOR: 7.129

IMPACT FACTOR: 7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY