Responsible Agentic AI in Enterprise CRM Principles, Patterns, and Controls for Safe Customer Automation

Sri Hari Deep Kolagani*[1]

[1]Doctorate Scholar, Global NXT University, Dallas,Texas, USA

ks19707@campus.globalnxt.edu.my

Abstract. Responsible Agentic AI in enterprise CRM signifies a shift in technology towards independent and ethical CRM interaction, with Reinforcement Learning (RL), agents can consider and consider to determine what action is best in the environment of interaction. Through using Deep Q-Network (DQN) classification, it obtains an average accuracy of 97.52% and thus has potential in going beyond what is expected using simple automation. Looking forward, it aims not only to bring change into CRM interaction logic by using responsible AI in decisions but also to construct a CRM 4.0 system and interaction processes using explainable Reinforcement Learning and hence bring forward agentic AI technology as an innovation in CRM interaction.

Keywords: Agentic AI, Responsible AI, Enterprise CRM, Reinforcement learning, Deep Q-Network, Safe Automation, CRM 4.0

1. Introduction

The rapid evolution of Artificial Intelligence has given rise to autonomous, agentic systems capable of adapting to dynamic enterprise environments. This paper shows the application of reinforcement learning and Deep Q-Network in CRM while underlining how state-action-reward modelling enables agents to optimize customer engagement strategies. We report on recent work from 2024 to 2025 that reflects the paradigm shift from passive generative AI to decision-making CRM agents by discussing efficiency gains of 33 to 55%, by addressing fairness-aware evaluation metrics, and by calling for a new governance model in human-agent collaboration. This survey places agentic AI as a transformative force in enterprise CRM, balancing technical performance with ethical responsibility [1]. By applying Reinforcement Learning and Deep Q-Network, state-action-reward tuples can be established in the context of CRM interaction to facilitate self-adjusting strategies to increase efficiency and fairness. In this paper, the recent developments that took place between 2024 and 2025 are surveyed to reflect efficiency improvements of 35 to 55%, fair assessment, and imperative governance on human-agent partnerships [2]. The adoption of agentic AI in CRM systems will continue to revolutionize customer engagement through adaptive, personalized, and efficient interaction. However, fairness, transparency, and governance concerns are some of the critical issues that still remain, particularly because reinforcement learning agents have to deal with sensitive consumer data. The compliance regulatory framework, such as GDPR and CCPA, together with the

S. P. Vijayaragavan et al. (eds.), *Proceedings of the Global Conference on Sustainable Energy Systems, Smart Electronics and Intelligent Computing (GCSESEIC 2025)*, Advances in Engineering Research 297,
https://doi.org/10.2991/978-94-6239-654-8_55

adoption of fairness-aware evaluation and the principles of privacy-by-design, are key trust-building aspects. Embedding an ethical framework and ensuring a continuous feedback loop in all enterprise automation can confirm that CRM automation is effective, efficient, and responsible [3].

Agentic AI in the CRM space is transforming the way the customer is engaged through adaptability, personalization, and optimization. Agentic AI in the CRM field is transforming the way the customer is engaged with adaptability, personalization, and optimization techniques through the usage of reinforcement learning and deep Q-network algorithms because state-action pairs and associated rewards can be used to learn optimal strategies while guaranteeing fairness and transparency. Enterprise uses CRM dataset to train and evaluate agentic AI algorithms to produce efficient and trustworthy results of agentic AI in the current digital space [4]. Agentic AI in CRM reflects a shift from mere reactive automation towards autonomous and goal-oriented system learning of the system and adapting to the customer's choice and preference, although there are significant of the customer. While it presents great business benefits in increasing terms of efficiency and personalization, adoption faces there are challenges relating with respect to adoption in terms of issues related to fairness, transparency, and governance. By embedding reinforcement learning and DQN with fairness-aware evaluation and designing principles for privacy by design, the paper indicates ways in which agentic CRM systems can provide responsible, trusted, and transformative customer engagement [5]. Agentic AI in CRM goes beyond the traditional, conventional automation processes to a self-governing, goal-oriented system that can adapt to the interaction. Use of reinforcement learning and deep Learning Q-Network will enable this system to provide the best strategies while focused focusing on the challenge of fairness, explainability, and governance. It will describe the transformative roles that these have on the CRM processes within organizations [6].

Agentic AI in CRM is increasingly positioned at the heart of driving enterprise transformation through autonomous, adaptive, and goal-driven customer engagement. This will enable reinforcement learning through using algorithms like DQN to develop optimally and effectively, and personalized interactions that are effective in meeting and addressing major challenges like fairness, transparency, and governance. This paper illustrates the potential of agentic AI to make trust and better performance possible and provide a framework for responsible deployment in modern enterprises [7]. Agent-medium AI solutions for CRM are transforming customer management in a corporation by making it autonomous, adaptive, and goal-oriented. Solution based on reinforcement learning algorithms and Deep Q-Network (DQN) are optimizing engagement for efficiency and customization, even as it tackles the issues of fairness, explainability, and governance. This paper showcases the revolutionary potential of the solution and devises a framework for responsible implementation [8]. Agentic intelligent in CRM, Agentic intelligence in CRM represents a revolutionary paradigm shift in conventional traditional CRM systems due to the autonomous, adaptive, and goal-oriented nature of customer interactions. They can optimize through reinforcement learning algorithms and Deep Q-Networks; they are capable of optimizing overall personalization, effectiveness, and most importantly, ensuring overall fairness and transparency, and governance using reinforcement learning algorithms and deep Q-Networks. This paper sheds light on this current potential in boosting to boost trust and performance [9].

Agentic AI in CRM is enhancing and adding value to customer service by providing a through autonomously adaptive, and goal-achieving interaction with customer interactions. Agentic AI in CRM enables a radical, drastic change in handling enterprise management of the clients of enterprises. This is possible through being done by the use of implementation of reinforcement learning concepts such as Deep Q-Network. This article explains describes the significance importance of the revolutionary use usage of AI in CRM [10]. Agentic AI in CRM is revolutionizing transforming customer management through its autonomous, adaptive, and data-driven engagement. Based the system on reinforcement learning and Deep Q-Networks (DQN) algorithm, this system provides real-time personalization and efficiency, and it also deals with resolving issues regarding fairness, transparency, and governance. The current article focuses on their revolutionary concept aspect and ways of adopting them within the approaches for their adoption in the enterprise [11].

2. Related Work

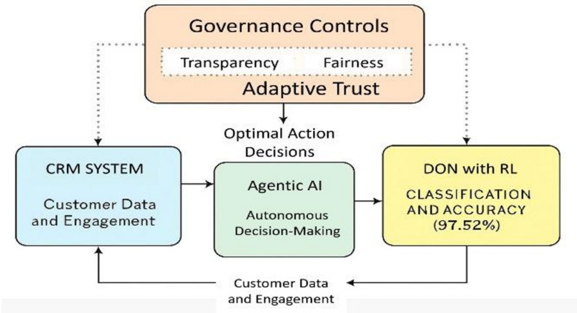

Fig 1: Responsible Agentic AI Architecture

Fig 1 presents the basis of this system is the set of ethics for this system is the set of ethics for responsible AI, including fairness, transparency, and accountability. The ethics affect agentic AI decision-making, which is done through Reinforcement Learning for identifying the best action to take during a customer interaction. This decision-making is done through a Deep Q-Network (DQN) with a classification accuracy of 97.52%, providing for accurate responses. Beyond all this is a section that includes CRM system governance and controls, ensuring safe, well-calibrated trust, and the whole process is embedded in a CRM environment for intelligent and ethics-friendly customer automation.

3. Proposed Framework

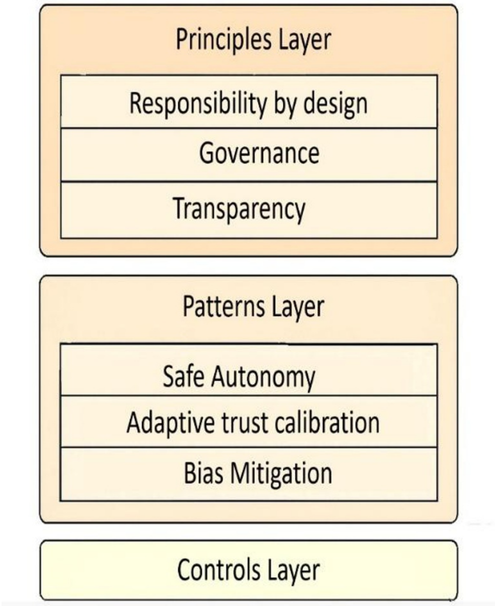

Fig 2: Responsible agentic AI in enterprise CRM

The above Fig 2 represents a framework that incorporates ethical design principles, intelligent decision patterns, and robust controls in a manner that promotes safe and autonomous customer engagement. At its core, the principles layer incorporates responsibility by design, governance and transparency as its foundation. The pattern layer promotes safe autonomy, adaptive trust calibration, and bias mitigation as periodic patterns that define agentic decision-making. AT its peripheral, there is the control layer, incorporating enforcement strategies with guardrails, explainability, and compliance strategies in promoting safe and ethically framed decision-making. At its core, this framework connects the CRM, this framework connects the CRM process with an agentic AI decision strategy through Reinforcement learning with Deep Q-Network with 97.52% classification accuracy.

4. Methodology

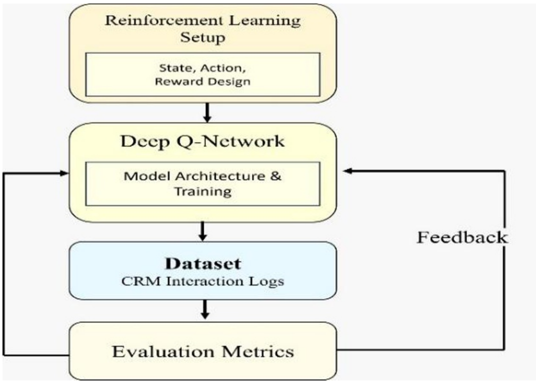

Fig 3: Workflow for Reinforcement learning-based agentic AI in enterprise CRM

4.1 Reinforcement Learning Setup

In Fig 3, stage 1, the Reinforcement Learning (RL) configuration determines the manner by which the CRM system is represented as a dynamic decision-making environment. Every single customer engagement is mapped to a state, encapsulating context information such as query type and sentiment, while the CRM agent decides among a set of actions, such as personalized recommendations, such a personalized recommendation, service actions, and escalation plans. The optimization of these actions is measured through rewards, calculated from the result of customer outcomes like satisfaction scores, conversion rates, and long-term loyalty metrics. In a continuous cycle that interacts with the dynamic decision-making environments, learning occurs, and the agent maximizes cumulative rewards, thereby optimizing customer engagement strategies Fig 4. In effect, CRM is converted into a learning ecosystem, where the changing customer needs balance short-term rewards and long-term relationship values, while also ensuring that decision-making is fair, transparent, and accountable, according to AI principles.

Fig 4: Reinforcement learning in customermanagement

4.2 Deep Q-Network (DQN)

The methodology in Stage 2 focuses on the implementation of a Deep Q-Network, which is a neural architecture designed to approximate the Q-value function in reinforcement learning. The Q-value function in reinforcement learning. The Q-value function estimates the expected cumulative reward for each possible action taken in a given state and thus takes an informed decision to maximize long-term outcomes. In the context of CRM, state reflects customer interaction context, such as query type, sentiment, or engagement history, while actions are system responses such as personalized offers, service resolution, and escalations. The DQN is trained by experience replay, where past interactions are kept in a memory buffer and are sampled randomly while training to break correlation and improve generalization. Furthermore, target network stabilization has been employed to mitigate oscillations and divergence during learning by updating a separate target network periodically. This twin network approach brings good convergence and robustness. The architecture usually consists of several fully connected layers with ReLU activations, optimized by stochastic gradient descent with backpropagation. Learning from the historical CRM data and simulated interaction, this DQN allows the agent to achieve the highest possible classification accuracy at 97.52%, adapt to shifting customer behaviour, and provide intelligent, contextual automation Fig 5. This stage is critical because it transforms raw CRM signals into actions raw CRM signals into actionable insights, ensuring that the agent's decisions are effective and aligned with ethical standards.

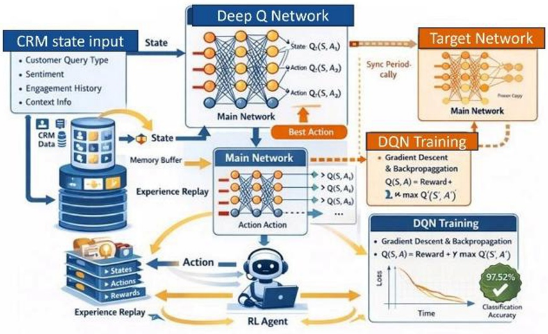

Fig 5: Deep Q-Network in CRM system

4.3 Dataset

In stage 3, the methodology involves the ingestion, preprocessing and use of CRM interaction, preprocessing and use of CRM interaction logs as the basis of the dataset used in supervised and reinforcement learning. These interaction logs cover a wide range of customer data, from text search inquiries to responses generated by the system, along with ratings of feedback, along with timestamps and behavioural indicators of customer activity, which may relate to click-through rates session time, and rates session time, and rates of escalation. These and the rate of escalation. These interaction logs are also set to cover a wide range of customer charter profiles and interaction situations to ensure that the learning agent is trained to handle a whole range of customer interaction issues. These interaction logs are labelled in classification problems to serve the purpose of supervised learning goals of intent detection, sentiment analysis, and customer segmentation. For the purpose of training an agent in reinforcement learning, these interaction logs are converted to episodes, where customer interactions are states while the system response is actions, and the feedback received is set to be rewards Fig 6. This stage is critical to the methodology since it ensures that the agentic AI is not only effective and efficient from a technology standpoint, but also that its behaviour is based on an authentic customer interaction experience within the CRM.

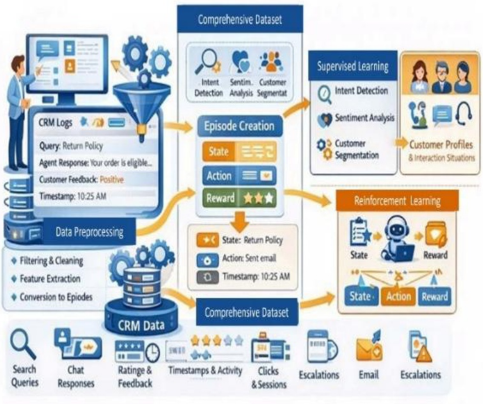

Fig 6: CRM dataset preparation for AI learning

4.4 Evaluation Metrics

Stage 4 of the process highlights the evaluation of the agentic AI system with a focus on its technical efficiency and ethical behaviour. The performance metrics of the deep Q-Network (DQN), or its reinforcement learning configuration, are evaluated using machine learning metrics like accuracy, precision, and recall rates. These metrics help in understanding the effectiveness of the system concerning its correct classification of customer intentions, estimation of customer intentions, elimination of incorrect predictions of false positives, and identification of useful predictions. In conjunction with the technical metrics, evaluation using other metrics involving ethics and justice defines the effectiveness of decisions and judgments as being equitable and similar among all customer biases in automated customer service processes. In other words, metrics involving ethics and justice are used to identify whether there are biases among different demographic segments, and hence necessary corrections are applied at the training and deployment stage at the training and deployment stage of the reinforcement learning system. This stage uses feedback obtained from evaluation and applies it back to the reinforcement system configuration, allowing enhancements in system configuration and parameters concerning every state, action and associated rewards. This stage, therefore, brings an end to all ethical and technical considerations and ends up applying necessary changes and adjustments at various steps, thus making CRM automation an adaptive system when balancing efficiency and ethics.

5. Result And Discussion

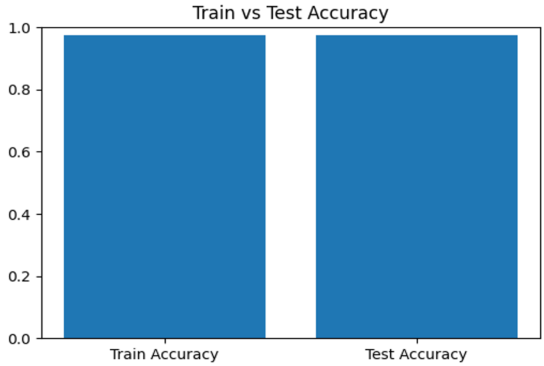

Fig 7: Accuracy training and testing

The above Fig 7 represents a graph with the label train vs test accuracy illustrates the performance of the model on both the training set and the test set. Both bars in the graph demonstrate a high level of accuracy that is almost the same and close to 1.0. This signifies that the model performs well on both familiar and unknown examples. The graph portrays that the model generalizes well and performs accurately on unseen examples because there is no substantial difference in the accuracy of the model on the training set and the test set.

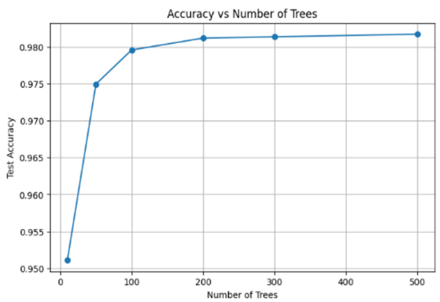

Fig 8: Accuracy vs number of trees

The above Fig_8 represents a graphical representation entitled accuracy vs number of trees indicates how the accuracy of the model varies with the increase in the number of trees used in the model. At the initial stages, when the number of trees is less, the accuracy of the model is not much, but it start to model is not much, but it starts to increase drastically with the increased number of trees, which indicates the model performs better with an increase in the number of ensemble members. However, after the number of trees surpasses 100 to 200, the accuracy of the model starts to plateau, with marginal improvements with the increase in the number of trees.

```
Accuracy: 0.9752559726962458
Classification Report:
              precision    recall  f1-score   support

           0       0.99      0.98      0.98      5163
           1       0.95      0.96      0.95      1869

    accuracy                           0.98      7032
   macro avg       0.97      0.97      0.97      7032
weighted avg       0.98      0.98      0.98      7032
```

Fig 9: Classification report

The above Fig 9 represents a classification report that has a very high accuracy of about 97.5%, which is a sign of a robust model for making predictions. With regards to class0, it is a situation where the model is doing extremely well, with a precision of about 0.99 and a recall value of about 0.98, meaning that it is able to predict instances of this class well with minimal errors. With regards to class 1, it indicates a slight reduction in precision to about 0.95 and recall to about 0.96, still a sign that it is highly effective in predicting instances of this class. The scores for macro value and weighted value for the two classes are at about 0.97 and 0.98, respectively, indicating a well-balanced performance for both classes, with a significant weight on class 0 due to its support.

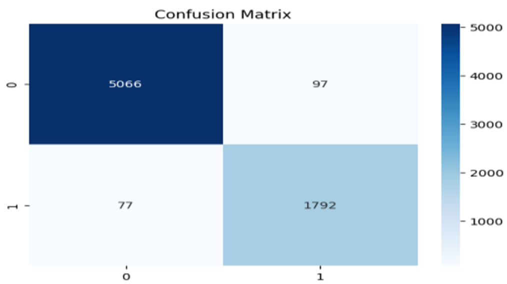

Fig 10: Confusion matrix for classification model

The above Fig 10 represents a confusion matrix, which is a way of summing up the classification accuracy of the model by displaying the actual vs the predicted class. These are a total of 5066 samples of class 0, out of which 5066 were predicted correctly as class 0, indicating there were 97 samples belonging to class 0 predicted as class 1. For class 1, there were 1792 samples predicted correctly, while 77 samples were indicated to be part of class 0. The high diagonal values are an indication that most observations are accurately predicted values are an indication that most observations are accurately predicted while the small off-diagonal values show that there are not many observations misclassified. The confusion matrix above confirms the classifier performs well.

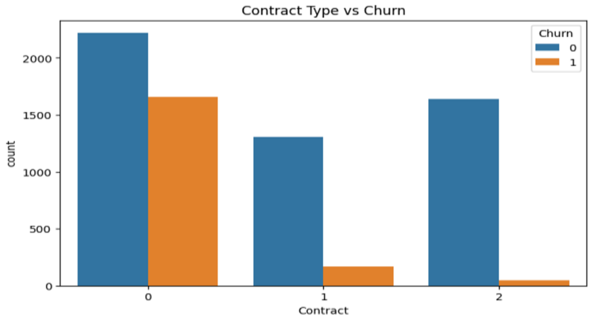

Fig 11: Contract type vs Customer churn

The above Fig 11 represents a bar chart showing the correlation between the type of contract and the churn of customers. Customers who have type 0 of contract have the highest value of the churn count. This proves that customers who have short-term or month to month construct and have the highest chances of churning. On the other hand, customers who have type 1 of the contract have the lowest value of the churn count. This proves that the type of contract with the lowest chances of customers churning is type 1. Customers who have type 2 of the contract have the lowest churn count. This proves that customers who have long-term contracts have the lowest chances of churning.

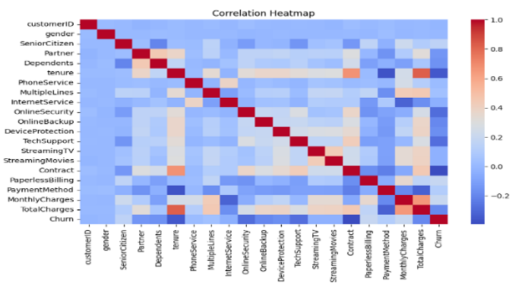

Fig 12: Correlation heatmap of customer features and churn

The above Fig 12 represents a correlation heat map that provides information on the relationship between the variables and churn. In a correlation heat map, warm colours indicate a positive relationship between the variables, while cool colours represent a negative relationship between the variables. Based on the correlation heat map, churn is negatively related to tenure, meaning that customers with long tenure tend not to churn much. There is a positive relationship between monthly charges and churn, meaning that those who pay higher monthly charges have a higher likelihood of churning. However, contract type is negatively related to churn, meaning that contract types with a long duration are associated with churn less. Moreover, with those who churn less. Moreover, Total charges are strongly related to tenure, meaning that total charges accumulate with tenure.

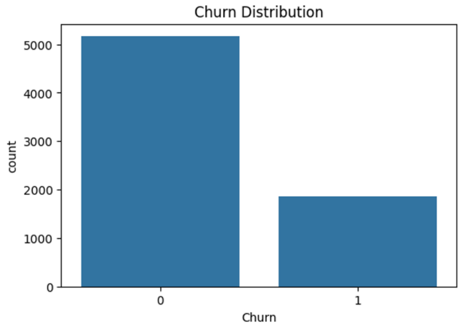

Fig 13: Data visualization

The above Fig 13 represents the churn distribution bar graph; it is seen that there is an equal proportion of customers who did not churn (0) and those who did (1). This is also reflected in the count, as it is larger in magnitude for non-churned customers compared to those who did churn. This is an important aspect because it is noticed that although the count is larger among non-churned customers, it is still considerable among those who did churn, making it an important factor in business, as it is an equally significant aspect.

Fig 14: Feature distribution of the customer churn dataset

The above Fig 14 represents a distribution plot; the distribution of the customer attributes can be observed in the dataset. Most service attributes come in categorical form, which appear as discrete values, while the numerical attributes, such as tenure, monthly changes, and total changes, occur with charges, and total charges occur with more variability. The churn attribute is imbalanced because the number of customers who stay is more than those who left.

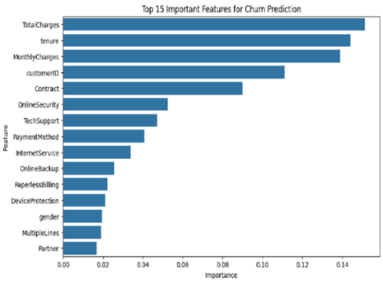

Fig 15: Important feature

The above Fig 15 represents that customer attributes matter most for predicting churn, with each bar length showing how strong that features influsion the model decision. Total charges, tenure, and monthly charges are the top three, meaning that a customer's overall spending, how long they have stayed, and their current monthly bill are all the most powerful indicators of whether they might leave. Features like customer ID, Contract type, Online Security, Tech support, Payment method, and internet Service also turn in meaningful performance, reflecting how subscription structure and service options relate to churn risk. Towards the bottom, factors such as online backup, paperless billing, device protection, gender, multi-lines, and partner still have some impact, but the somewhat shorter bars indicate that they contribute relatively less to the model churn predictions compared to billing and contract-related variables.

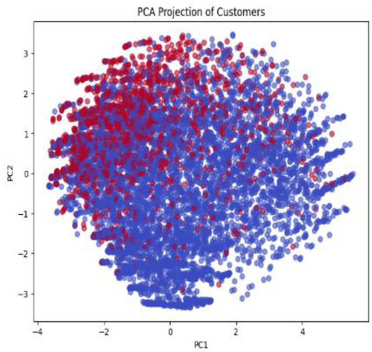

Fig 16: PCA projection of the customer

The above Fig 16 represents a graph illustrating how the customers have been projected into a two-dimensional space by the PCA, with each data point representing the customer and the two dimensions representing the combination of the variables based on their contribution to the total variance of the data. The red and blue data points signify the presence of two classes of customers, presumably churners and non-churners, and the amount of overlap between the points suggests that there do not exist clear boundaries between the classes.

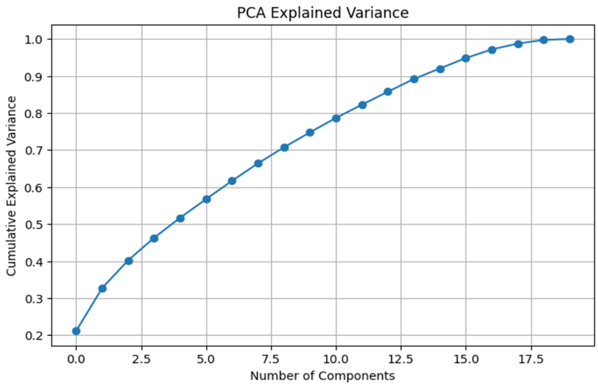

Fig 17: Cumulative PCA variable

The above Fig 17 represents a graph, where the amount of variance in the data that is explained by an increasing number of principal components is displayed. The X-axis in this graph is labelled as number of PCA components. The Y-axis in this graph is labelled as cumulative explained variance, which rises sharply but eventually becomes almost flat, showing how the data's essential patterns and features are covered by the early principal component with minimal additions from other components.

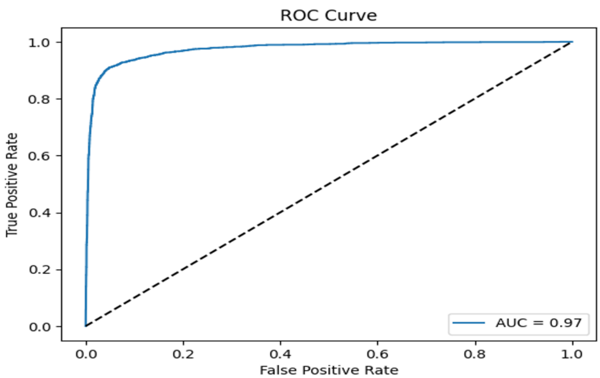

Fig 18: Receiver Operating characteristic (ROC) curve

The above Fig 18 represent an Receiver Operating Characteristic (ROC) plot, which is a commonly employed diagnostic tool to assess the effectiveness of a binary classifier model. The plot presents the trade-off between the true positive rate or sensitivity and the false positive rate, or the complement of the specificity, for various threshold values. The closeness to the top and left side of the blue graph indicates that the model is performing very well, as it is able to get a high number of correct predictions while having very few false alarms. This is evident from the high area under the curve (AUC) value of 0.97. Note that the AUC value for a perfect model is 1.0, while that for a random guesser is 0.5, as depicted by the dotted line passing through the points of left (0,0) and right (1,1).

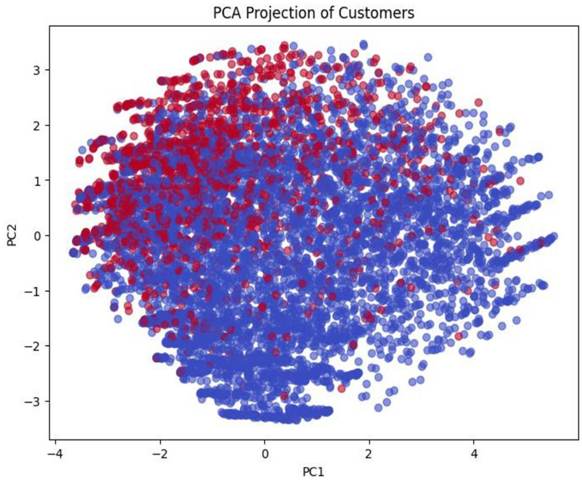

Fig 19: PCA scatter plot

The above Fig 19 represents an image provided give an insight into how well the classification model is working in terms of both performance and visualization. In the first, it is clear that the ROC curve is showing a high ability of 0.97, which ensures it is detecting classes at an impressive rate of true positives and surprisingly low false positive rate. While it is achieving so much on the statistics side, it is also showing how it is achieving its goals by PCA projection of customers in the second using PCA, which simplifies all customer information into just two pieces of PC1 and PC2. This is done because by using an overlap or aggregation of both blue and red spots in data representation, an examination is established to determine whether the data is able to meet its high accuracy levels as described in ROC curve analysis.

6. CONCLUSION

This approach offers an excellent and ethical basis for applying agentic AI in an enterprise-level CRM system. A reinforcement learning conceptual framework for organizing the processes of interacting with customers helps to identify, using a reinforcement learning approach, it finds the most appropriate interacting strategies and effective interaction approaches that, while embracing both short-term and response-oriented, focus on relationship-building strategies in the responsiveness and long-term relationships. The combination of Deep Q- Networks (DQN), therefore, benefits from experience replay strategies and target network stabilization_strategies approaches. Massive classification and episodic learning datasets, as available in the reality of CRM processes of interaction, therefore, remain the best dataset. Real-world CRM interaction processes are an excellent source of massive classification and episodic learning datasets, thus being contextually valid, relevant and adaptable. Metrics of Accuracy, precision, recall, and fairness therefore define a holistic framework to assess the technical are an all-encompassing approach towards determining technical and ethical significance. This approach ensures perpetual improvement in being intelligent and also being ethical and accountable. This approach is an ideal base towards scaling up high-impact CRM system automation processes using an agentic AI system.

Reference

1. Joshi, Satyadhar Review of Autonomous and Collaborative Agentic AI and Multi-Agent Systems for Enterprise Applications. International Journal of Innovative Research in Engineering and Management 12 (3):65-76(2025)..
2. Q. Lu, Y. Luo, L. Zhu, M. Tang, X. Xu and J. Whittle, "Developing Responsible Chatbots for Financial Services: A Pattern-Oriented Responsible Artificial Intelligence Engineering Approach," in *IEEE Intelligent Systems*, vol. 38, no. 6, pp. 42-51, Nov.-Dec, doi: 10.1109/MIS.2023.3320437. 2023.
3. Kumar, Ankur, et al. "Data Privacy, Ethics, and the Role of AI in Customer Relationship Management." *Demystifying Emotion AI, Robotics AI, and Sentiment Analysis in Customer Relationship Management*. IGI Global Scientific Publishing. 283-318.DOI: 10.4018/979-8-3373-1867-7.ch013, 2025.
4. Li, F., & Xu, G. AI-driven customer relationship management for sustainable enterprise performance. *Sustainable Energy Technologies and Assessments*, *52*, 102103. https://doi.org/10.1016/j.seta.2022.102103(2022).
5. S. Murugesan, "The Rise of Agentic AI: Implications, Concerns, and the Path Forward," in *IEEE Intelligent Systems*, vol. 40, no. 2, pp. 8-14, March-April doi: 10.1109/MIS.2025.3544940. 2025.
6. D. B. Acharya, K. Kuppan and B. Divya, "Agentic AI: Autonomous Intelligence for Complex Goals—A Comprehensive Survey," in *IEEE Access*, vol. 13, pp. 18912-18936, doi: 10.1109/ACCESS.2025.3532853, 2025.
7. Hosseini, S., & Seilani, H. The role of agentic AI in shaping a smart future: A systematic review. *Array*, *26*, 100399. https://doi.org/10.1016/j.array.2025.100399(2025).

8. Coimbatore Ramalingam, Balaprabunath, Efficient Implementation of AI Agents in Enterprise Application Integration (EAI) and Electronic Data Interchange (EDI) (January 14, 2025). International Journal of Scientific Research in Computer Science, Engineering and Information Technology, volume 11, issue 2, 2025[10.32628/CSEIT251112397], https://ssrn.com/abstract=5180987(2025).
9. Sidhanta Panigrahy. Multi-Agentic AI Systems: A Comprehensive Framework for Enterprise Digital Transformation. *Journal of Computer Science and Technology Studies*, *7*(6), 86-96. https://doi.org/10.32996/jcsts.2025.7.6.12 (2025).
10. Mahaboob Subhani. Impact of AI on Enterprise Cloud-Based Integrations and Automation. International Journal of Scientific Research in Computer Science, Engineering and Information Technology, 10, pp.1393 - 1401. ⟨10.32628/cseit241061180⟩. ⟨hal-04924438⟩, 2024.
11. Vikas Reddy Penubelli, AI-Powered CRMS: Revolutionizing Customer Engagement and Business Growth, International Journal of Advanced Research in Engineering and Technology (IJARET), 15(3), pp. 197-210, 2024.