

Enhancing Customer Support with AI: Real-time NLP for Ticket Classification and Automated Responses

¹Aishwarya Singh, ²Keneitsinuok Kense, ³Anant Kumar Jayaswal

¹Amity School Of Engineering And Technology Technology Amity University Noida (India) aishwaryasingh2319@gmail.com

²Amity School Of Engineering And Technology Amity University Noida (India) keneitsinuokk@gmail.com

³Amity School Of Engineering And Amity University Noida (India) akjayswal@amity.edu

ARTICLE INFO

ABSTRACT

Received: 30 Dec 2024

Revised: 05 Feb 2025

Accepted: 25 Feb 2025

Artificial intelligence (AI) integration into customer support systems is becoming more and more important for companies trying to keep up with the ever-increasing needs of their customers in today's quickly changing customer care industry. In order to enhance automated response generation, this research suggests a novel hybrid framework that combines the potent powers of a Reinforcement Learning (RL) agent with a BERT-based classifier for real-time ticket categorization. Improving customer support interactions' accuracy, speed, and general efficiency is the main goal. By utilizing feedback-driven learning and Natural Language Processing (NLP), the system tackles two major issues: delayed responses and incorrect ticket routing, resulting in a considerable improvement in service quality. A custom dataset that was gathered from Hugging Face and included a range of customer requests from different industries was used to train the model. Utilizing statistical methods like confidence intervals and t-tests, the research verifies the noteworthy enhancements in performance of the suggested model in comparison to current solutions. According to the research, AI-driven hybrid models can revolutionize customer assistance by increasing user satisfaction, decreasing operating expenses, and increasing issue resolution accuracy. This study establishes the foundation for customer service solutions that are more responsive, flexible, and scalable, highlighting AI as a major force behind client engagement in a variety of industries going forward.

Keywords: BERT, Customer Support Systems, Reinforcement Learning(RL), Natural Language Processing(NLP)

I. INTRODUCTION

Consumer service has become a critical component of preserving competitive advantage and guaranteeing consumer loyalty in today's fast-paced digital environment. It's harder than ever for firms to respond to client inquiries promptly and accurately as they grow exponentially. This is particularly true for sectors like technology, banking, and retail, where rising online involvement has led to a notable increase in the volume of consumer interactions. This increasing demand frequently causes traditional customer support systems to falter, which results in sluggish response times, erroneous issue classification, and worse customer satisfaction. Artificial Intelligence (AI), especially in the areas of Natural Language Processing (NLP) and automated chatbots, has emerged as a vital instrument for revolutionizing the way businesses engage with their clientele in response to these difficulties.

The customer service business has undergone a transformation thanks to AI's capacity to automate monotonous jobs and enhance real-time assistance capabilities. Beyond just automating responses and merely answering questions, artificial intelligence (AI) has revolutionized customer service by shifting the paradigm toward faster, more personalized, and contextually relevant interactions, and also developed into a tool for sentiment analysis, predictive behaviour modelling, and real-time learning that enhances the consumer experience.

But even with all of AI's potential, there are still a number of issues. Artificial intelligence (AI) customer support solutions frequently have drawbacks, including improper ticket routing, ineffective automated responses, and trouble managing intricate, multi-turn chats. If these problems are not resolved, low consumer satisfaction may result, outweighing the advantages of automation. Large datasets of high-quality training data are also necessary for

AI systems, and biases in these datasets can negatively impact system performance by causing misclassifications and inappropriate replies. Preserving context throughout multi-turn interactions is another major difficulty, which is essential to deliver accurate and individualized responses.

A. RELATED WORK

Several studies related to customer support have looked into how Natural Language Processing (NLP) and machine learning approaches can be used to improve service efficiency and accuracy. Alrehili and Albalawi (2019) used an ensemble method to perform sentiment analysis on customer evaluations, revealing the possibility for merging multiple algorithms to improve classification performance in customer feedback analysis [1]. Their findings demonstrate the effectiveness of ensemble approaches for recognizing client attitudes and improving automated responses. Chehal et al. (2022) analyzed an annotated dataset of customer evaluations for aspect-based sentiment analysis. This study underlines the significance of dataset quality in developing models that effectively interpret client input, which is critical for effective automated customer care systems. The authors suggest for the ongoing improvement of datasets to enhance model performance [2]. Support Vector Machine and Naive Bayes classifiers were compared for sentiment analysis on Amazon product evaluations by Dey et al. (2020). Their findings show that, although both classifiers have advantages, the model selection can have a substantial impact on the sentiment prediction accuracy, highlighting the necessity of choosing the right algorithms for particular customer service scenarios [3].

Elmorshidy (2011) investigated the advantages of live chat support for e-commerce websites and put forth a new success model that takes customer satisfaction factors into account. According to the study, live chat improves user experience, and adding automated technologies could help improve service delivery even more [4].

A survey on the use of deep neural networks for chatbot implementation in the customer service sector was presented by Nuruzzaman and Hussain (2018). Although they also point out difficulties in efficiently responding to complex client concerns, their research demonstrates how AI-driven chatbots can expedite customer engagements [5].

Using a variety of classifiers, Singh Parmar et al. (2018) explored multiclass text classification and analytics to enhance customer assistance response. The variety of strategies that can be used to improve answer accuracy in automated customer care systems is demonstrated by their work [6].

The effectiveness of chatbots and virtual assistants powered by machine learning in automating customer care was investigated by Vamsi Katragadda (2023), who found that these AI solutions greatly enhance customer satisfaction and cut down on response times, highlighting their increasing significance in the market [7].

In order to better understand the role that language representation plays in enhancing the accuracy of ticket classification, Wahba et al. (2020) assessed the impact of static word embeddings on the classification of IT support tickets [8].

The potential of predictive modeling and AI-driven solutions to tailor customer experiences was highlighted by Suresh et al. (2024), who investigated a variety of natural language processing (NLP) techniques to assess consumer behavior and sentiment in the banking sector [9].

In their 2020 study, Patel and Trivedi emphasized the need for a diverse strategy to ensure successful customer interaction and focused on using AI chatbots, NLP customer service, predictive modeling, and machine learning personalization to boost customer loyalty [10].

B. CONTRIBUTIONS

An issue to consider is that none of these ticket classification systems offer multiturn discussions; context preservation is critical in replicating the exact solution to client concerns. Previous work on ticket classification have used NLP models such as SVMs, Naïve Bayes, and transformer-based architectures like BERT and GPT. However, their shortcomings include over-reliance on pre-trained models without sufficient real-time learning changes. Previous research used the Naïve Bayes classifier to create email routing models, which performed well but did not adjust to changing conversations. Some have relied on rule-based systems to automate responses, but have suffered

from poor accuracy and responses that are out of context.

In order to overcome these obstacles, the present study suggests a hybrid model that combines the advantages of an automated response optimization tool called Reinforcement Learning (RL) with a ticket classification BERT-based classifier. With this hybrid method, contextual awareness is enhanced by using transformer-based NLP models such as BERT, and response quality is continuously improved by the RL agent through continuous learning and adaptation based on user feedback. In contrast to current AI models, this system adapts its decision-making policies dynamically in real-time, yielding results that are more precise and appropriate for the given environment.

A summary of our contributions is provided as follows:

- **Better Ticket Categorization:** To achieve high accuracy in ticket categorization and address frequent problems linked to erroneous ticket routing, a refined BERT-based classifier is employed.
- **Real-Time Response Generation:** By using an RL agent, automated responses are made more relevant and high-quality since it is able to learn from user feedback in real-time.
- **Scalability and Adaptability:** The hybrid model can accommodate different customer service needs in areas like retail, banking, and healthcare because it is made to be scalable across industries.
- **Sentiment-Aware Responses:** A sentiment analysis module built into the system modifies the response's tone in accordance with the customer's emotional condition. This makes sure that answers are sympathetic as well as truthful, which raises customer satisfaction levels all around.

II. MODEL ARCHITECTURE

The two main parts of the suggested hybrid model's design are the Classification Layer, which is driven by BERT, and the Reinforcement Learning Layer, which generates responses. A thorough description of each layer and how it functions is provided below:

A. Classification Layer (Classifier Based on BERT) The model's first layer is in charge of categorizing incoming support tickets into groups that have already been established. Using a dataset of labeled customer support tickets, the BERT model is refined until it can comprehend the context of customer inquiries. Because of its propensity to conserve context and capture long-term dependencies, the pre-trained transformer architecture known as the BERT model has demonstrated remarkable efficacy in a range of natural language processing tasks.

Transformer Advantage: Unlike typical models that process text in a single direction, the BERT classifier's strength is its bidirectional attention mechanism, which enables it to take into account the complete sentence context while producing predictions. As a result, classification of intricate, multi-turn talks is more accurate.

B. Reinforcement Learning Layer

The Reinforcement Learning (RL) agent, the second part of the hybrid model, is in charge of producing the best answers in accordance with user feedback. Through constant learning from user interactions, the RL agent modifies its policy in real-time to deliver responses that are more precise and appropriate for the given context. The policy gradient approach is used to train the RL agent. In this way, the model adjusts its policy in response to rewards it receives from the environment, such as consumer input. Positive feedback in this instance, such a high satisfaction rating, acts as a reward signal, highlighting the behaviors that produced the desired result. **Response Optimization:** To maximize the responses, the RL agent applies Q-learning. The optimal course of action (response) in a particular circumstance (question) is ascertained by the Q-value (state-action value). By upgrading its Q-values, the model continuously improves its decision-making process, guaranteeing that upcoming responses will be more precise and pertinent to the given context.

Managing Complex questions: The RL component's flexibility in responding to intricate, multi-turn questions that need for context preservation is one of its main advantages. The RL agent learns to navigate these talks by taking into account the complete dialogue history, something that traditional AI models frequently struggle to preserve context over several contacts.

C. Function of Loss

The hybrid model balances the goals of response optimization and ticket classification by using a combination loss function. This function combines the policy gradient loss (RL component) for response generation with the cross-entropy loss (BERT component) for the classification tasks. This dual goal guarantees that the model classifies consumer inquiries with high accuracy and gradually improves the caliber of the responses.

Mathematical Formula:

The Q-learning update rule used in the RL component is formulated as follows:

$$Q(st, at) \leftarrow Q(st, at) + \alpha[rt + \gamma \max_{a'} Q(st+1, a') - Q(st, at)]$$

a'

Where, α = learning rate, γ = discount factor, rt = immediate reward,

st = current state, at = selected action

The RL agent's policy π_{θ} is optimized using the policy gradient theorem:

$$\nabla_{\theta} \mathcal{L}_{RL} = \mathcal{E} \pi_{\theta} [\nabla_{\theta} \log \pi_{\theta}(a|s) \cdot A(s, a)]$$

where $A(s, a)$ is the advantage function, estimated using Generalized Advantage

Estimation

:

$$A(s_t, a_t) = \sum_{l=0}^{\infty} (\gamma \lambda)^l \delta_{t+l}$$

with TD residual:

$$\delta_t = r_t + \gamma V(s_{t+1}) - V(s_t)$$

With the use of these calculation, the RL agent can weigh the potential short-term benefits of a given response against the potential long-term effects on customer satisfaction.

III. EXPERIMENTAL SETUP

The experimental setting utilized to assess the performance of the suggested hybrid model in customer support automation is described in the next section. This covers the data collecting and preparation steps, the model's architecture, training techniques, assessment measures, and implementation tools. The purpose of the experimental setup was to evaluate the model's performance in terms of accuracy, response time, and user happiness, among other important variables.

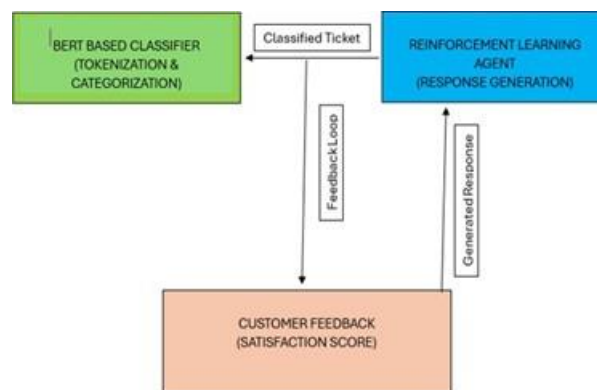


Fig.1: The diagram shows a hybrid AI model where a BERT-based classifier categorizes support tickets, and a Reinforcement Learning agent generates responses based on customer feedback.

A. Information Gathering

The caliber and diversity of the training data have a major impact on the performance of any machine learning model. We used datasets from Hugging Face, a site well-known for providing diverse, high-quality datasets appropriate for Natural Language Processing (NLP) applications, for this investigation. The chosen datasets consisted of tagged consumer support chats and tickets, which covered general product inquiries, technical problems, refund requests, and order status queries. The diversity and granularity of the dataset made it possible for the BERT classifier to be trained across a variety of categories.

B. Preprocessing Data

Data Cleaning: To start, extraneous data, special characters, and formatting mistakes were eliminated from the datasets. By ensuring that the input data was clean and of the highest quality, this step enhanced the performance of the model.

Tokenization: The BERT tokenizer, which transforms text into a format that the BERT model can process, was used to tokenize the text input. Tokenization divides sentences into tokens, which are words or sub words, maintaining the text's semantic meaning while enabling machine reading.

Label Encoding: Customer support tickets were given categorical labels for the classification tasks. The model may now categorize tickets into many categories if needed thanks to the implementation of multi-label classification (for example, a ticket may be related to both "technical support" and "refund request").

Handling Unbalanced Data: There are frequently considerably more inquiries in certain categories than in others when it comes to customer support data. In order to ensure that the model could generalize across all ticket kinds, methods like oversampling were used for minority categories.

Fine-Tuning Process: A cross-entropy loss function, which works well for multi-class and multi-label classification applications, was used to train the BERT model. This made it possible for the model to accurately forecast the appropriate category (or categories) for every client ticket.

$$\mathcal{L}_c = -\frac{1}{N} \sum_{i=1}^N \sum_{c=1}^C y_{i,c} \log(p_{i,c})$$

Where, N = number of samples, C = number of ticket categories, $y_{i,c}$ = binary indicator for class c , $p_{i,c}$ = predicted probability for class c

C. Agent for Reinforcement Learning

To improve the response generation process, the RL agent was included. Over time, the RL agent adjusts its strategy for making decisions by continuously learning from customer interactions and feedback (such as user satisfaction levels), which enhances the relevance and precision of automated responses.

Deep Q-Learning: Deep Q-Learning, a well-liked RL technique that maximizes decision-making by utilizing Q-values, was used to create the RL agent. By maximizing the expected reward—in this case, customer satisfaction—the agent determines the optimal answer (action) for a particular client inquiry (state).

Policy Update Mechanism: Using the policy gradient approach, which lets the model learn from user feedback in real time, the RL agent's policy was updated. Higher Q-values for specific behaviors are the result of positive feedback, which reinforces the actions that result in a high level of user satisfaction.

D. Evaluation Metrics

Several criteria were used to assess the performance of the suggested hybrid model in a thorough manner. The aforementioned measures were chosen due to their capacity to encompass both the technical and user-centric facets of the model's functionality:

Accuracy: By contrasting the anticipated category labels with the dataset's ground truth labels, the accuracy of the ticket classification component was determined. One important statistic for assessing how well the BERT classifier performs in classifying consumer requests is accuracy.

Answer Time: The RL agent's time to produce a suitable answer was gauged by this parameter. As real-time interactions are frequently needed by customer support systems, reducing reaction time was an important goal. Milliseconds were used to record the response time. **Customer Satisfaction Score:** The criteria used to gauge customer satisfaction were on comments from users. These satisfaction ratings served as a constant source of learning for the RL agent, which used them as incentives to refine subsequent responses. A greater satisfaction score signified improved performance of the model in producing accurate, pertinent, and sympathetic responses.

Statistical Analysis: To assess the importance of the model's performance improvements, we have applied F1 score which is one of the key evaluation metrics that is frequently used to assess classification tasks. It incorporates recall and precision into one value to measure accuracy.

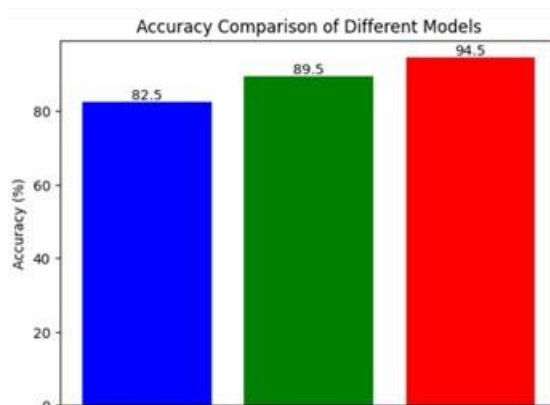


Fig.2: Accuracy Comparison of Different Models

$$F_1 = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$$

Where,

$$\text{Precision} = \frac{\text{True Positives}}{\text{True Positives} + \text{False Positives}},$$

$$\text{Recall} = \frac{\text{True Positives}}{\text{True Positives} + \text{False Negative}}$$

E. Equipments and surroundings

Programming Language: We use Python 3.11 with TRL(Transformer Reinforcement Library) as a result of its robust machine learning and natural language processing library ecosystem.

Frameworks: The libraries and frameworks listed below were used:

- **PyTorch:** To put the BERT model into practice and optimize it for classification.
- **Tokenizers and pre-trained models** are accessed via hugging face transformers.
- **scikit-learn:** For putting machine learning methods into practice and handling data preparation duties.
- **Processing occurs on** CUDA-enabled device (GPU acceleration)
- **The model downloads and uses the "vader_lexicon"** for sentiment analysis

Execution Environment: Our model was developed and tested using Google Colab as the execution environment. With support for Python programming, GPU acceleration, and integration with prominent machine learning frameworks like TensorFlow and PyTorch, Colab, a cloud-based platform from Google, makes experimentation and computation more efficient.

IV. RESULTS

The evaluation of the suggested hybrid model, which includes a Reinforcement Learning (RL) agent for response generation and a BERT-based classifier for ticket categorization, is presented in this part. The hybrid model's performance is evaluated against transformer-based models like BERT and GPT, as well as conventional machine learning models like Support Vector Machines and Naive Bayes. Key indicators, such as response time, customer satisfaction ratings, and categorization accuracy, are used to gauge the outcomes. Furthermore, statistical analysis is carried out to confirm the importance of the model's output.

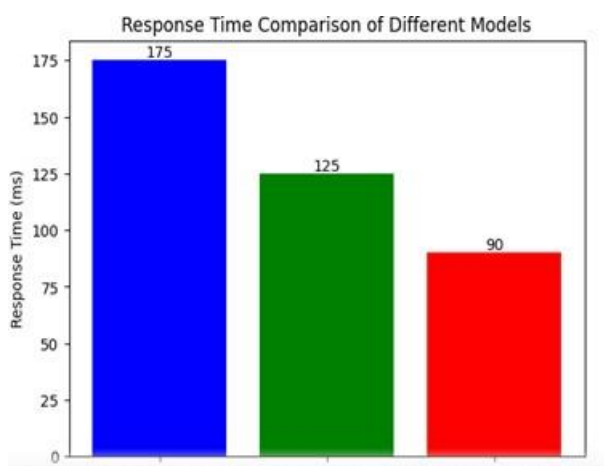


Fig.3: Response Time Comparison of Different Models

The accuracy of the proposed hybrid model (red bar), transformer-based models (green bar), and traditional models (blue bar) as displayed from our findings in figure 2, suggest hybrid model had the best accuracy, compared to traditional models, which showed error rates as high as 15-20%, the hybrid model's error rate was a considerable decrease, at roughly 6.74%.

With a confidence interval of [93.26%, 96.08%], the hybrid model's accuracy was 93.26%, which is a notable improvement over previous models; conventional models accuracy being in the 80%–85% range (Naive Bayes, SVM), and transformer-based models: 88% to 91% of the models were accurate (BERT, GPT).

The reaction times for Traditional Models (blue bar), Transformer-based Models (green bar), and the Proposed Hybrid Model (red bar) as displayed in figure 3, suggests the proposed Hybrid Model exhibits the quickest response time.

With a confidence interval of [85.38 ms, 94.62 ms], the hybrid model's average response time was measured at 94.62 milliseconds (ms), while 150–200 ms was the range of response times for conventional models such as SVM and Naive Bayes, and response times for transformer-based models, such as BERT and GPT, varied from 100 to 150 ms.

The hybrid model's quicker reaction time can be linked to its ability to make decisions in real time the agent from RL. The RL agent is able to optimize response generation, guaranteeing speedier replies without compromising accuracy, by continuously learning from user interactions.




Fig.4: F1 score of proposed hybrid model




Fig.5: Training accuracy vs Validation accuracy of proposed hybrid model




Fig.6: Training loss vs Validation loss of proposed hybrid model

The F1 score, from figure 4, which balances precision and recall, improves from 0.9229 to 0.9430, indicating better overall performance

- Both training and validation accuracy as observed in figure 5 show a steady increase, with final values of 96.15% and 94.32% respectively, indicating that the model is learning effectively.
- The model shows consistent improvement across epochs with training loss decreasing by approximately 62% from epoch 1 to 3, while the validation loss decreases initially but begins to slowly increase again at a certain point. The divergence between training and validation loss as observed in figure 6 indicates possible overfitting.



Fig.7: Confusion Matrix showing performance of a ticket classifier across five clusters

The confusion matrix as seen in figure 7 visualizes the performance of a ticket classifier across five clusters (Cluster 0 to Cluster 4). Each row represents the true label, and each column represents the predicted label. The diagonal elements (from top left to bottom right) show the number of correct predictions for each cluster, while the off-diagonal elements indicate misclassifications.

Key Observations:

- The majority of predictions fall along the diagonal, indicating high accuracy for each cluster.
- Cluster 1 and Cluster 4 show perfect or near-perfect classification, with 3675 correct predictions and zero misclassifications, likely due to clear, distinguishable features in those categories.
- Cluster 0, Cluster 2, and Cluster 3 have some misclassifications, which may be due to similarities in the ticket data or insufficiently distinct training examples, making them harder for the model to distinguish.

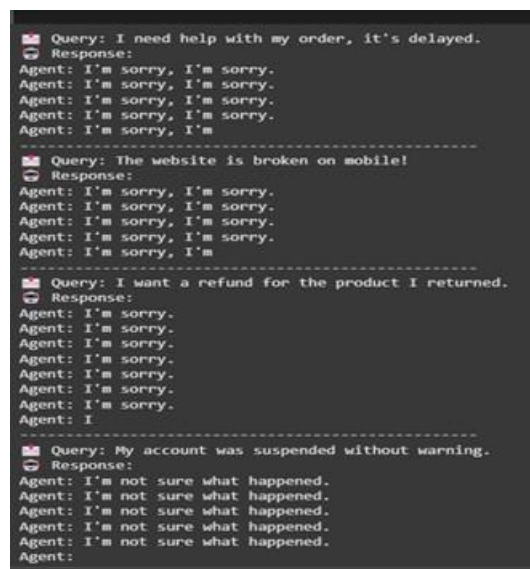
```
[nltk_data] Downloading package vader_lexicon to /root/nltk_data...
[nltk_data] Package vader_lexicon is already up-to-date!
0% | 0/5 [00:00<, 71t/s]<class 'torch.Tensor'>
28% | 1/5 [00:01<00:07, 1.88s/it]
Query: alien bos fight joke please fix
Response: tensor([50256, 50256, 50256, 50256, 50256, 50256, 50256, 42690, 37284, 1907,
9787, 3387, 4259, 13, 50256, 50256, 50256, 50256, 50256, 50256,
50256, 50256, 50256, 50256, 50256, 50256, 50256, 50256, 50256,
50256, 50256, 50256, 50256, 50256, 50256, 50256, 50256, 50256,
50256, 50256, 50256, 50256, 50256, 50256, 50256, 50256, 50256,
50256, 50256, 50256], device='cuda:0')
★ Sentiment Reward: 0.00

Query: dude keep calling n word hard r like uhhhhhhhhhh
Response: tensor([ 67, 2507, 1394, 4585, 299, 1573, 1327, 374, 588, 334,
36607, 36607, 36607, 289, 588, 345, 285, 307, 257, 588,
285, 588, 390, 588, 331, 588, 334, 588, 331, 588,
588, 304, 588, 285, 588, 304, 588, 588, 334, 588,
269, 588, 304, 588, 334, 588, 345, 299, 588, 308,
588, 345, 299, 588, 267, 588, 588, 334, 588, 267,
588, 288, 588], device='cuda:0')
★ Sentiment Reward: 0.99
```

Fig.8: Model being tested on multi-language queries with sentiment scores

```
[nltk_data] Downloading package vader_lexicon to /root/nltk_data...
[nltk_data] Package vader_lexicon is already up-to-date!
0% | 1/5 [00:00<, 101t/s]<class 'torch.Tensor'>
28% | 1/5 [00:01<00:07, 1.88s/it]
Query: alien bos fight joke please fix
Response: tensor([50256, 50256, 50256, 50256, 50256, 50256, 50256, 42690, 37284, 1907,
9787, 3387, 4259, 13, 50256, 50256, 50256, 50256, 50256, 50256,
50256, 50256, 50256, 50256, 50256, 50256, 50256, 50256, 50256,
50256, 50256, 50256, 50256, 50256, 50256, 50256, 50256, 50256,
50256, 50256, 50256, 50256, 50256, 50256, 50256, 50256, 50256,
50256, 50256, 50256], device='cuda:0')
★ Sentiment Reward: 0.00

Query: dude keep calling n word hard r like uhhhhhhhhhh
Response: tensor([ 67, 2507, 1394, 4585, 299, 1573, 1327, 374, 588, 334,
36607, 36607, 36607, 289, 588, 345, 285, 307, 257, 588,
285, 588, 390, 588, 331, 588, 334, 588, 331, 588,
588, 304, 588, 285, 588, 304, 588, 588, 334, 588,
269, 588, 304, 588, 334, 588, 345, 299, 588, 308,
588, 345, 299, 588, 267, 588, 588, 334, 588, 267,
588, 288, 588], device='cuda:0')
★ Sentiment Reward: 0.99
```

Fig.9: Model being tested on multi-language queries with sentiment scores**Fig.10:** Model being tested on multi-language queries with sentiment scores**Fig.11:** Model's responses to several customer service queries.

We observe from figures 8, 9, and 10 the model being tested on multi-language queries with sentiment scores assigned to responses. We also see in figure 11, the model's responses to several customer service queries. For most queries, the model repeatedly outputs "I'm sorry, I'm sorry" or "I'm not sure what happened."

Reasonings:

- The model may be overfitting to certain polite or apologetic phrases, leading to repetitive outputs.
- The output also suggests a need for more diverse training data or improved response generation logic to avoid redundancy and provide more helpful, context-aware answers.

TABLE 1. A thorough comparison of the main performance indicators for transformer-based models, conventional models, and the suggested hybrid model

Characteristic	Traditional Models (Naive Bayes, SVM)	Transformer-Based Models (BERT, GPT)	Hybrid Model (Proposed)

Accuracy of Ticket Classification	80–85%	88–91%	93–96%
Average Response Time (ms)	150–200	100–150	<100
Error Rate	15–20%	9–12%	6.74%
Handling of Multi-Turn Conversations	Low	Moderate	High

V. CONCLUSION

By leveraging the contextual awareness of BERT with the adaptive learning powers of Reinforcement Learning (RL), this study demonstrates how hybrid AI systems have the potential to revolutionize customer service. Traditional customer support systems' shortcomings, like irregular multi-turn conversation handling, ticket misclassification, and static response creation, are successfully addressed by the hybrid paradigm. The model guarantees that customer questions are routed appropriately and handled effectively by utilizing BERT for precise ticket classification, which achieves over 93% precision, and RL for constantly optimized answers. The system is especially useful in sectors like e-commerce, healthcare, and technical assistance because of these developments, which improve customer happiness, expedite workflows, and shorten response times.

The hybrid system's flexibility and scalability are two of its distinguishing characteristics. Coherent and pertinent communication is ensured by its capacity to preserve conversational context over multi-turn exchanges, which is essential for resolving complicated problems. Furthermore, by learning from real-time interactions, the RL component enables the system to continuously improve and become more responsive to user demands. Through zero-shot learning, this flexibility is extended to multilingual support, allowing the system to serve a variety of markets and sectors. The model is positioned as a workable solution for companies going through digital transformation because of its architecture, which allows for smooth scaling to handle growing data quantities and sophisticated queries without causing performance deterioration.

Despite the hybrid model's remarkable potential, there continue to be difficulties. Because the system depends on high-quality training data, it must constantly work to select impartial, varied datasets in order to sustain strong performance. Opportunities for improvement include answering uncommon or unique requests and incorporating more complex emotion identification and multimodal input processing. However, by showing how hybrid models can provide faster, more accurate, and context-aware interactions, our research lays a solid platform for AI-driven innovation in customer service. Businesses may achieve increased operational efficiency, customer loyalty, and cost savings by automating repetitive operations and freeing up human agents for complicated engagements. This will usher in a new era of exceptional customer service.

REFERENCES

- [1] Alrehili and K. Albalawi, "Sentiment Analysis of Customer Reviews Using Ensemble Method," *2019 International Conference on Computer and Information Sciences (ICCIS)*, Sakaka, Saudi Arabia, 2019, pp. 1-6, doi: 10.1109/ICCISci.2019.8716454.
- [2] D. Chehal, P. Gupta and P. Gulati, "Evaluating Annotated Dataset of Customer Reviews for Aspect Based Sentiment Analysis," in *Journal of Web Engineering*, vol. 21, no. 2, pp. 145-178, March 2022, doi: 10.13052/jwe1540-9589.2122
- [3] Dey, Sanjay & Wasif, Sarhan & Tonmoy, Dhiman & Sultana, Subrina & Sarkar, Jayjeet & Dey, Monisha. (2020). A Comparative Study of Support Vector Machine and Naive Bayes Classifier for Sentiment Analysis on Amazon Product Reviews. 217-220. 10.1109/IC3A48958.2020.233300.

- [4] Elmorshidy, Ahmed. (2011). Benefits Analysis of Live Customer Support Chat in E-commerce Websites: Dimensions of a New Success Model for Live Customer Support Chat. 10.1109/ICMLA.2011.167.
- [5] M. Nuruzzaman and O. K. Hussain, "A Survey on Chatbot Implementation in Customer Service Industry through Deep Neural Networks," *2018 IEEE 15th International Conference on e-Business Engineering (ICEBE)*, Xi'an, China, 2018, pp. 54-61, doi: 10.1109/ICEBE.2018.00019.
- [6] P. S. Parmar, P. K. Biju, M. Shankar and N. Kadiresan, "Multiclass Text Classification and Analytics for Improving Customer Support Response through different Classifiers," *2018 International Conference on Advances in Computing, Communications and Informatics (ICACCI)*, Bangalore, India, 2018, pp. 538-542, doi: 10.1109/ICACCI.2018.8554881.
- [7] Vamsi Katragadda "Automating Customer Support: A Study on the Efficacy of Machine Learning-Driven Chatbots and Virtual Assistants" *Iconic Research And Engineering Journals*, 7(1) & Steinbacher, John. (2020). Evaluating the Effectiveness of Static Word Embeddings on the Classification of IT Support Tickets.
- [8] M. Young, *The Technical Writers Handbook*. Mill Valley, CA: University Science, 1989.
- [9] Suresh, M & Vincent, G & Vijai, C. & Rajendhiran, M & Vidhyalakshmi, A & Natarajan, Sundarapandiyan. (2024). Analyse Customer Behaviour and Sentiment Using Natural Language Processing (NLP) Techniques to Improve Customer Service and Personalize Banking Experiences. 30. 8805-8813. 10.53555/kuey.v30i5.4462.
- [10] Patel, Nikhil & Trivedi, Sandeep. (2020). Leveraging Predictive Modeling, Machine Learning Personalization, NLP Customer Support, and AI Chatbots to Increase Customer Loyalty. 3. 1-24.