Intelligent Ticket Assignment System: Leveraging Deep Machine Learning for Enhanced Customer Support

By Yusuff Adeniyi Giwa, Taiwo Akinmuyisitan, Jacob Sanni,
Adebesin Adedayo, Boluwaji Benard Akinmuyisitan
& Nosakhare Ikponmwosa

Teesside University

Abstract- In the evolving customer support domain, traditional ticketing systems struggle to meet increasing demands for speed and accuracy. This study presents an intelligent ticket assignment system leveraging BERT, Graph Neural Networks (GNN), and Prototypical Networks to enhance classification and routing efficiency. The methodology includes comprehensive preprocessing of historical ticket data, feature extraction using natural language processing (NLP), and model evaluation based on accuracy, precision, recall, and F1-score. Results indicate that BERT achieves the highest accuracy (89.4%), precision (88.7%), recall (90.2%), and F1-score (89.4%), outperforming GNN (87.6%) and Prototypical Networks (86.8%) by notable margins. A comparative analysis with Random Forest (85.3%) further demonstrates a 4.1% improvement in accuracy.

INTELLIGENTTICKETASSIGNMENTSYSTEMLEVERAGINGDEEPMACHINELEARNINGFORENHANCEDCUSTOMERSUPPORT

Strictly as per the compliance and regulations of:

Intelligent Ticket Assignment System: Leveraging Deep Machine Learning for Enhanced Customer Support

Yusuff Adeniyi Giwa [α], Taiwo Akinmuyisitan [σ], Jacob Sanni [ρ], Adebesin Adedayo [ω], Boluwaji Benard Akinmuyisitan [¥] & Nosakhare Ikponmwosa [§]

Abstract- In the evolving customer support domain, traditional ticketing systems struggle to meet increasing demands for speed and accuracy. This study presents an intelligent ticket assignment system leveraging BERT, Graph Neural Networks (GNN), and Prototypical Networks to enhance classification and routing efficiency. The methodology includes comprehensive preprocessing of historical ticket data, feature extraction using natural language processing (NLP), and model evaluation based on accuracy, precision, recall, and F1-score. Results indicate that BERT achieves the highest accuracy (89.4%), precision (88.7%), recall (90.2%), and F1-score (89.4%), outperforming GNN (87.6%) and Prototypical Networks (86.8%) by notable margins. A comparative analysis with Random Forest (85.3%) further demonstrates a 4.1% improvement in accuracy. The analysis demonstrates both performance strengths and real-life practicality and scalability characteristics of the system when managing high traffic volumes. Stability and predictive accuracy improved through the application of noise filtering alongside SMOTE oversampling and weighted loss functions for addressing data quality problems and class imbalance and model integration complexities. The research demonstrates how machine learning changes the way customer service operations work while showing AI models can boost service quality and operational effectiveness.

Keywords: e-ticketing AI system, machine learning, predictive model, bert algorithm, data preprocessing, prototypical networks, graph neural networks, natural language processing, feature extraction, ticket classification.

Author α: Innovarie, Stockton-on-Tees.
e-mail: UKy.giwa@innovarie.com
Author σ: Department of Computer Science, Global Banking School, London. e-mail: Takinmuyisitan@globalbanking..ac.uk
Author ρ: Department of Computing, Teesside University, Avaso Technologies, Newcastle, United Kingdom.
e-mail: Sannijay2709@gmail.com
Author ω: Department of Physics and Engineering Physics, Obafemi Awolowo University, ile-ife, Nigeria.
e-mail: adebesindedayo@gmail.com
Author ¥: Business Management, Postgraduate School, University of Bolton, Greater Manchester. e-mail: UK ba1res@bolton.ac.uk
Author §: Business Management Dept, Global Banking School, London. e-mail: nikponmwosa@globalbanking.ac.uk

I. Introduction

Customer support stands as a fundamental business tool which boosts client satisfaction followed by brand loyalty while securing business superiority. Customer question resolution efficiency functions as a business success indicator that shapes how many customers will stay with the company [1]. Businesses need to evolve from solving problems after they occur to implementing ahead-of-time customer engagement approaches to stay competitive [2]. The growing complexity of customer interactions becomes worse because customers use a range of communication channels from social media to email and live chat and mobile applications [3]. The contemporary buying public requires assisted service that provides integrated personalised support throughout all its multiple interaction channels so organisations must adopt flexible data-based solutions to handle these platforms efficiently.

Current business operations still use basic support systems which base their ticket processing exclusively on human intervention. The outdated systems create multiple operational problems and performance issues which render services below contemporary consumer demands [4]. The ability to scale remains problematic due to human agents facing challenges in handling large numbers of tickets which generates delays and dissatisfied customers and backlog cases. Service quality issues persist in support centres because inconsistent ticket routing and prioritisation causes resolution times to vary by 25% for similar issues according to research [5]. The present constraints demonstrate an immediate requirement for intelligent automated solutions in customer service operations.
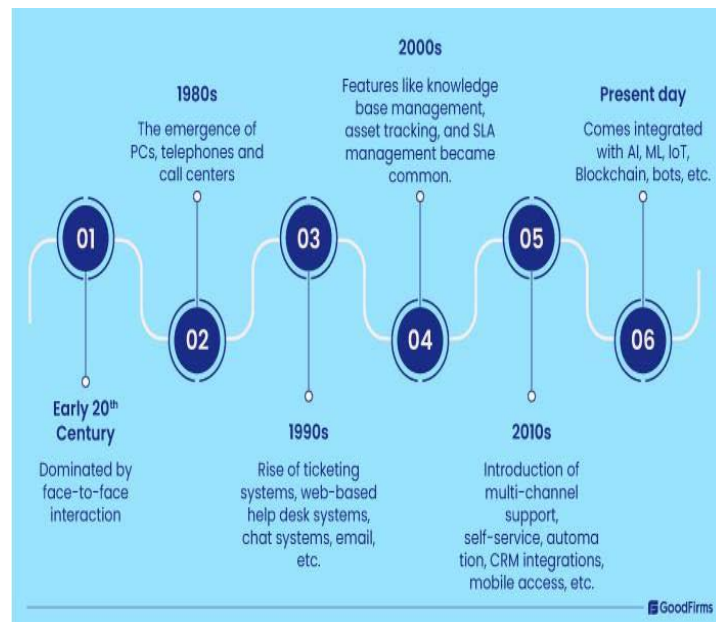
Figure 1: Evolution from Manual to AI-Powered Customer Support Systems. Adapted from [4]

Figure 1 demonstrates the fundamental shift in support systems from manual to AI-powered solutions which addresses these challenges as shown in Figure 1. The traditional support approaches do not satisfy contemporary customer demands because they depend on human staffing and contain error possibilities which limit their effectiveness. AI-based systems that use automation enhance ticket classification and prioritization and assignment speed while improving precision to deliver faster responses and better resource distribution [6]. Machine learning models trained with historical data enable predictive analytics to forecast both ticket urgency and complexity levels which businesses use to provide more accurate and data-based services [7]. Modern business practices remain hindered by existing scalability problems and inconsistent service quality since many organisations have not embraced AI-driven solutions.

This paper addresses these challenges by developing an intelligent ticket assignment system leveraging machine learning. The study aims to:

- Automate the categorisation and assignment of customer support tickets using machine learning techniques.
- Utilise historical data to predict the urgency and complexity of incoming tickets.
- Improve ticket assignment accuracy and response times compared to traditional, manual systems.

By implementing this intelligent AI-driven approach, businesses can overcome the limitations of manual ticketing systems, enhancing both scalability and service efficiency. This study highlights how AI-enabled solutions can transform customer service operations, allowing organisations to effectively meet the increasing demands of modern consumers.

II. Related Work

The development of customer care systems started with manual engagement before adding semi-automated solutions and finally implementing AI-based intelligent systems. Modern customer expectations regarding rapid precise service delivery create an increasing market need for scalable flexible accurate ticket-handling systems. Most organisations currently experience difficulty with traditional methods that depend on human agents to classify and assign support tickets when they expand their customer engagement through social media platforms and email and live chat channels. These operational systems succeed in limited situations but struggle to scale adequately nor operate with the desired speed and precision which creates service delays and incorrect ticket routing while generating service quality flaws [8]. The legacy systems restrict true-time capability and analytical analysis which reduces customer trend detection along with proactive customer involvement [9].

Recognising these limitations, early machine learning models such as Support Vector Machines (SVMs) and Logistic Regression were introduced to automate ticket classification. Such models required extensive feature engineering while achieving good results in structured environments without sufficient ability to handle complex context-rich questions [10]. Ensemble methods like Random Forest addressed non-linear relationships more effectively and improved classification accuracy; however, they remained insufficient for processing nuanced natural language interactions, particularly in cases involving sarcasm,

industry-specific terminology, or ambiguous phrasing [11]. The introduction of deep learning and graph-based approaches marked a significant shift in customer service automation by enabling models to extract context, identify relationships between tickets, and classify inquiries with greater precision [7].

BERT, a transformer-based model, advanced ticket classification by leveraging bidirectional contextual learning, allowing for improved semantic understanding. Advanced linguistic patterns within ticket classifications achieved better processing accuracy with this technology specifically in cases involving informal language and technical terminology [12]. However, despite its effectiveness, BERT remains computationally demanding, limiting its practicality for real-time deployment in high-volume service environments [13]. Graph Neural Networks (GNNs) introduced another layer of refinement by modelling customer inquiries as interconnected nodes, improving routing accuracy by identifying relationships between similar tickets. Real-time scalability of these models faces limitations because they need extensive hyperparameter tuning while requiring large resource usage [14][15].

Few-shot learning methods, such as Prototypical Networks and MAML, have emerged as an alternative solution for handling rare or novel ticket types. These methods leverage limited examples to adapt quickly to new customer inquiries [18][19]. These models are instrumental in dynamic service environments where new issues frequently arise, requiring rapid adjustments without extensive retraining. However, their reliance on specialised architecture and high computational costs has restricted widespread adoption, particularly in enterprise applications with resource constraints [23].

Despite these advancements, modern AI-driven models continue to present challenges. BERT, while highly effective in capturing linguistic nuances, struggles to generalise across domains with unfamiliar terminology, requiring extensive domain adaptation techniques to maintain performance [22]. GNNs, despite their ability to identify ticket relationships, demand significant processing power, making real-time implementation difficult [15]. Few-shot learning models offer adaptability but require computationally intensive architectures that limit their feasibility for large-scale deployment [23]. Bias and fairness also persist across these models, as imbalanced datasets can lead to disproportionate classification outcomes, potentially resulting in service inconsistencies [23].

Traditional systems' limitations have been well-documented in industry case studies, highlighting the pressing need for AI adoption. Large-scale e-commerce platforms have reported significant misrouting rates, leading to prolonged response times exceeding 24 hours and customer dissatisfaction [1]. Similarly, telecommuni-cations providers handling a number of daily inquiries have experienced severe service disruptions due to the inability of rule-based ticketing systems to scale effectively [28]. A UK financial institution, National Savings & Investments (NS&I), saw a 37% surge in complaints, reaching 33,655 in a year, as outdated technology and poor service led to long wait times and online access issues, diminishing customer trust [30]. On the other hand, companies like Amazon that are adopting BERT-based ticketing solutions have seen classification accuracy improve by 18%, reducing resolution times by 35% [29]. These findings underscore the need for AI-driven customer support frameworks that integrate scalable, efficient, and adaptive machine learning techniques.

While deep learning and graph-based models have revolutionised customer service automation, ongoing research must focus on mitigating computational constraints, improving domain generalisation, and addressing ethical concerns in AI-driven ticketing systems. Optimisation strategies such as model quantisation, hybrid AI architectures, and fairness-aware algorithms offer promising pathways for enhancing the scalability and real-time efficiency of machine learning applications in customer support [6]. The transition towards fully automated, intelligent support systems is an ongoing process, requiring continued refinement to ensure that AI-driven solutions meet the evolving demands of modern businesses and consumers.

III. Problem Statement and Motivation

Studies from 2011 indicated that 75% of customers did not like how long it took contact centres to respond [31]. When processing happens manually the response times get delayed and customers become dissatisfied along with their satisfaction measures decreasing. Ticket misclassification leads to incorrect routing that makes operational challenges and delay times steadily increase [32]. The correct routing of tickets requires longer procedural time when they travel through several departments before reaching their destination which results in service quality deterioration and higher operational costs. The lack of predictive analytics capabilities in these systems prevents businesses from forecasting upcoming customer difficulties at the same time as preventing peak load increases [24]. Organizations implement advanced machine learning (ML) models BERT and GNNs to address the limitations they face in their ticket processing systems. BERT processes information from both forward and backward text directions which provides contextual understanding of language topics to overcome traditional keyword classification limitations [12]. BERT uses its deep learning capabilities to detect word-to-word connections which results in fewer incorrect assignments when ticket categories are

25

determined [13]. Through its deep bidirectional learning ability BERT recognises the difference between product defect reports and usage inquiries that use similar wordings so it helps speed up ticket routing and resolution time.
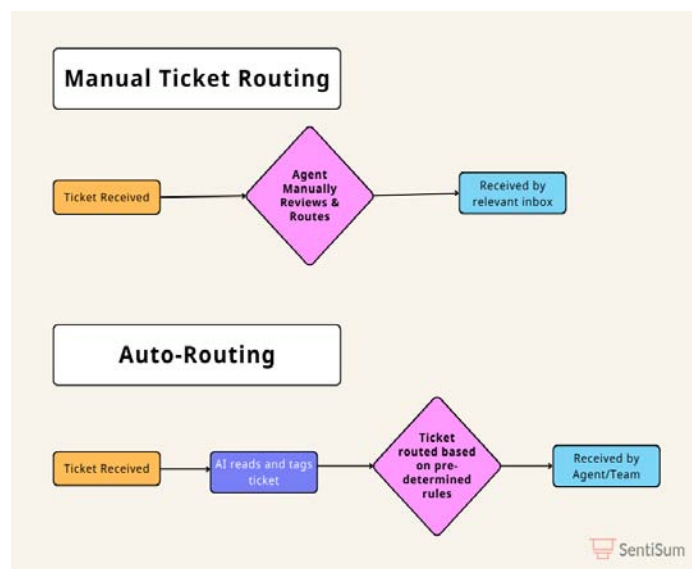

Figure 2: Manual Ticket Routing Process Vs. AI-Driven Automation Workflow. Adapted from [26]

BERT enhances semantic understanding but GNNs bring a network-based routing system for tickets particularly when multiple tickets show mutual connections. GNN technology applies mathematical graph algorithms to tickets so it can enhance both ticket grouping and routing operations and reduce ticket processing time. Operational classification organised by context enables organisations to handle interrelated problems which reduces customer contact duplications and optimises agent allocation throughout service departments. The combination of BERT with GNNs allows real-time decision automation and enhances scalability which modernises support systems to develop adaptive proactive services tools.

GNNs work together with BERT by analysing supportive ticket network relationships to provide better routing capabilities and categorization functionalities. The attributes of issue type and customer history serve as the basis for GNN models to establish edges between tickets while treating each ticket as a node in a graph format [14]. The system achieves group management abilities for connected tickets through this format allowing it to select appropriate tasks for specialised teams to reach faster resolution times. AI solutions like those presented in SentiSum's automatic versus manual routing test showcase how GNNs together with BERT achieve enhanced efficiency through minimised human involvement and rapid processing of numerous tickets [25]. The analysis systems train their comprehension by learning data constantly while simultaneously enhancing their predictive accuracy and minimising operational expenditures through lower technical support team staffing needs [14]. Customer satisfaction grows substantially together with proactive customer concern trend detection from machine learning insights through the combination of BERT contextual analysis with GNN relational modelling which prevents ticket misdirection and resolves all support requests [23]. BERT and GNNs establish their position as essential tools for achieving scalable high-precision data-based customer service operations that provide immediate customised support.

IV. Proposed Solution

The proposed system uses BERT alongside GNNs and Meta-learning capabilities to improve traditional customer support systems by resolving their current inefficiencies. BERT delivers better classification results through its context understanding capabilities which reduces the wrong decisions made by keyword-based methods. GNN operates to enhance routing performance by discovering ticket linkages which allows similar requests to automatically find the appropriate routing path for speedy resolutions. Meta-learning systems maintain their adaptability to process new and infrequent ticket types after small training sessions and this feature brings adaptive capabilities to changing service environments. The complete support system operates through an integrated pipeline that enables data preprocessing of unstructured inputs for BERT ticket classification then GNN processing of decisions and continuous adaptation delivered by Meta-learning mechanisms.

The system needs two key features to handle growing demands and perform instant processing requirements. The system achieves maximum

performance by using asynchronous processing and GPU acceleration and distributed computing methods to maintain speed for rising customer demand volumes. Many enterprise needs are addressed through this system which reduces workload requirements and accelerates service delivery and strengthens customer confidence in the process. This system represents a modern and effective automatic customer support solution through the integration of BERT linguistic skills and GNN relational analytics and Meta-learning adaptive features.

V. Methodology

a) Data Collection

This intelligent support ticket assignment system uses primary dataset information obtained from Kaggle to solve financial support requests. The chosen dataset attained selection because it provides substantial size along with varied subject matter which directly relates to financial support operations. The dataset contains various financial inquiries about billing disputes and technical banking platform faults and loan processing delays and account protection measures which serve as an optimal foundation for testing machine learning-driven classification and routing systems. This dataset encompasses different real-world data obstacles which include unbalanced class distributions along with diverse ticket wording and contains noisy information to supply a dependable setting for performance assessment and model adaptation testing.

The dataset contains several key fields that play critical roles in model training and ticket classification.

- *Ticket Description:* The text field carries extensive customer complaints that NLP functions primarily on for processing. Before BERT-based classification begins the text input needs preprocessing steps which include tokenisation and stop-word removal and lemmatisation to normalise the text input.
- *Customer Issue Type:* A categorical variable (e.g., technical, billing, service-related) that aids in prioritising and directing tickets to the appropriate departments. Model training requires this variable to be converted into one-hot encoding for usage purposes.
- *Resolution Status:* A binary variable indicating whether a ticket is closed or still open, used to track the efficiency of the support process. This feature helps in supervised learning models by acting as a label in training predictive models for resolution likelihood.
- *Time to Resolution:* A numerical field representing the time taken to resolve each ticket, crucial for performance evaluation and predictive modelling of response times. Normalising processes these data points first to eliminate abnormal results before standardising time values across various ticket types.

Extensive preprocessing and cleaning steps were applied to ensure data quality and enhance model accuracy. The data preparation process included handling missing values through imputation and removing system text along with duplicate entries to maintain data quality. Exceptional character removal and text normalisation techniques were implemented to refine textual data, ensuring better input consistency for NLP models. These preprocessing steps were critical in enhancing classification accuracy, enabling BERT, GNNs, and Meta-learning techniques to operate efficiently on structured and well-prepared data.

b) Source of Data

The datasets were sourced using Kaggle, a popular platform for seeking public datasets on any machine learning task. The data presented includes real-world ticket information, making it the best data to train models to respond to actual customer service queries. With such a diversity of data types, text, categorical, and numerical, this dataset is well suited for a multi-face machine learning approach. This also combines natural language processing with classification and time prediction.

Table 1: Dataset Overview

Field	Description
Ticket description	Textual description of customer issues
Customer issue Type	Categorical classification of issue type (e.g., billing, technical)
Resolution Status	Binary indicator of whether the issue was resolved or not
Time to Resolution	The numerical value indicates the time taken to resolve the issue.

Table 1 summarises the key fields in the dataset. These fields are the basis of training in machine learning models such as BERT and GNN. Rich, well-structured data is required to produce accurate classifications and predictions from these models.

c) Data Preprocessing

An essential part of these steps is data preparation because it cleans raw ticket data and prepares it for models such as BERT, GNN, or Prototypical Networks. In this section, the researcher

discusses how they cleaned and transformed the textual data for extraction of features, graph-based analysis, and few-shot learning, which can be conducted on novel or uncommon ticket types. Prototypical Networks classify tickets based on a few samples whenever new ticket types appear. GNN captures the relationship between comparable tickets, and BERT extracts contextual characteristics from the ticket descriptions.

d) *Text Preprocessing*

The following procedures below are ideal for the preparation of textual data from the support tickets for activities involving natural language processing (NLP):

*   *Lowercasing:* All text is converted to lowercase to ensure uniformity, reducing the complexity introduced by case sensitivity.
*   *Tokenisation:* The text is tokenised, breaking down the ticket scriptions into individual words or tokens. This allows the model to process the text at a word level.
*   *Stop-word Removal:* Common words that do not carry significant meaning (e.g., "and," "the") are removed from the text. This helps focus the analysis on the more essential terms.

*   *Lemmatisation:* Words are lemmatised, meaning they are reduced to their base or root form (e.g., "running" becomes "run"). This step reduces linguistic variability and improves model performance by treating similar words as the same entity.

e) *Feature Extraction using TF-IDF*

Once the text has been preprocessed, feature extraction is performed using the Term Frequency-Inverse Document Frequency (TF-IDF) method. TF-IDF transforms the text data into numerical features by weighing the importance of words based on their frequency in the document and across the dataset.

The TF-IDF formula is defined as follows:

Where:

* $TF(t, d)$ is the frequency of the term t in document d.
* N is the total number of documents in the dataset.
* $|\{d \in D : t \in d\}|$ is the number of documents containing the term t.
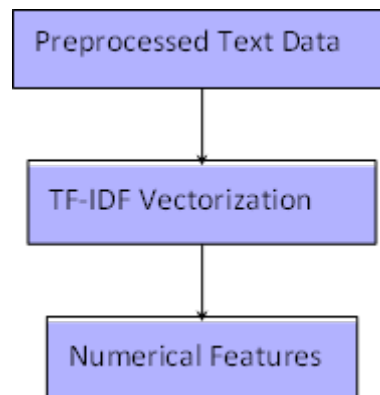*   Figure 3 illustrates the TF-IDF feature extraction process.

Figure 3: TF-IDF Feature Extraction Process: The Preprocessed Text is Converted into Numerical Features using TF-IDF Vectorisation

f) *Graph Representation for GNN*

For Graph Neural Networks (GNN) application, the ticket data is again transformed into graph representation. One ticket is represented by the following:

*   *Nodes:* Cosine similarity or other similarity measurements establish the similarity among tickets, and this is the basis upon which Edges between nodes.
*   The graph structure enables the GNN to reason about the connections between tickets and learn patterns that can be used to better route and classify tickets.
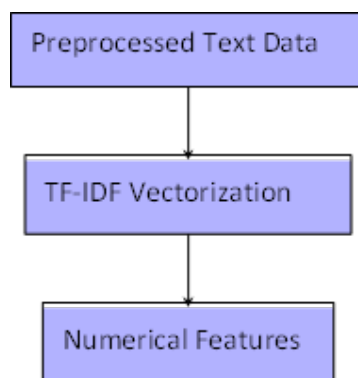
Figure 4: Graph Representation for GNN: Tickets are Represented as Nodes, and Edges are Created based on Ticket Similarity

g) Feature Engineering

Feature engineering is one of the crucial processes in converting raw text data into a numerical format that can be used for any machine learning model. To this end, two feature engineering techniques are applied to this intelligent support ticket assignment system: the graph representation using Graph Neural Networks (GNNs) and the text representation using BERT and TF-IDF.

h) Text Representation for BERT and TF-IDF

In order to successfully convert preprocessed text input into numerical features, the 'Term Frequency-Inverse Document Frequency' (TF-IDF) approach is adopted. This makes it possible for machines to focus on important words when performing classification tasks. The weighting system in this method provides higher importance to words that are important in a piece of text and rare across the whole dataset [27]. A particular document's weight value stands as the last element in each entry from the TF-IDF transformation matrix. The numerical representation generated by TF-IDF helps Random Forest and SVM models to execute ticket classification and prioritisation tasks.

The ticket descriptions benefit from deep contextual features extraction using Bidirectional Encoder Representations from Transformers (BERT). BERT generates one vector feature for each token which results from splitting the preprocessed ticket descriptions into their smallest linguistic units. This feature vector contains everything in each ticket record since it combines sequential word content from before and after the textual data. BERT works as a bidirectional system which enables it to interpret complex linguistic patterns. Translators (BERT) show exceptional ability in processing complicated and unclear ticket contents.

Algorithm for Preprocessing

Algorithm 1: Text Preprocessing Algorithm for BERT Input
The following are steps for preparing raw ticket descriptions for BERT feature extraction using the preprocessing algorithm.

Algorithm 1: Text Preprocessing for BERT Input

Input: Raw ticket descriptions

Output: Preprocessed text and TF-IDF vectors

1. *Result:* At this point, the preprocessed text and feature vectors are fed into BERT for feature extraction.
2. *Convert Text to Lowercase:* The next step is to tokenise the text into words. Stop words such as "and", "the", etc. are also removed, after which Lemmatization is performed to reduce words to their base forms. Lastly, TF-IDF vectorisation is applied to convert the text into numerical features.

Algorithm 2: Graph Preprocessing Algorithm for GNN Input

This algorithm outlines the basic steps for converting ticket descriptions into a graph representation for use with a Graph Neural Network (GNN).

Algorithm 2 Graph Preprocessing for GNN Input

Input: Preprocessed ticket descriptions

Output: Graph with nodes and edges based on similarity

3. *Result:* The result here is that the graph is used for GNN analysis as well as the capturing of relational data that is usually between tickets.
4. Begin
5. Next, a node is created for each ticket to calculate the similarity between tickets using cosine similarity. The nodes are then connected with edges if their similarity exceeds a threshold. Normalise the graph structure if necessary

i) Model Selection and Training

In this section, we will describe the process of selection and training of each of the models used in the intelligent support ticket assignment system. BERT is used for text understanding [12], GNN is used for relational learning [14], and Prototypical Networks is used for handling new ticket types [18].

i. *BERT for Text Understanding*

BERT accepted a large corpus for training before we applied it to fine-tune our ticket description dataset [12]. The fine-tuning process allows BERT to acquire domain knowledge about ticket description content that strengthens its ability to handle linguistic context within text. The attention mechanism in BERT functions as the main driver behind data contextual relationship extraction that enables correct ticket classification according to description semantics.

BERT's Self-Attention Mechanism is computed as:

$$\text{Attention }(Q, K, V) = \text{softmax }(\frac{QK^T}{\sqrt{d_k}})V \qquad (1)$$

Where:

Q is the query matrix,

K is the key matrix,

V is the value matrix,

dk is the dimensionality of the key matrix.

This mechanism helps BERT focus on relevant parts of the input when predicting the ticket's category.

ii. *Graph Neural Networks (GNN) for Relational Learning*

Graph Convolutional Networks (GCNs) analyse ticket connections through evaluation of their relational patterns [14]. The system builds the analysis through nodes which represent tickets and show relation points between the nodes to indicate similarity measurements. The ticket classification together with routing gets better because GNNs analyse relational data that helps identify closely related tickets in terms of their descriptions and categories.

The graph structure enables GCN to establish node representations by having features move between network edges. The approach makes relational data accessible to the model so the decision-making process can be optimised.

iii. *Few-shot Learning with Prototypical Networks*

Prototypical Networks are employed for handling new or rare ticket types that the model has not seen before [18]. Few-shot learning is instrumental in dynamic environments where new types of customer issues emerge frequently. Prototypical Networks work by computing a prototype (mean embedding) for each class and assigning tickets based on the distance to these prototypes.

iv. *Random Forest for Baseline Comparisons*

In addition to the deep learning models, we use a Random Forest model as a baseline for comparison [11]. Random Forest is an ensemble learning method that constructs multiple decision trees during training and outputs the mode of the classes for classification tasks.

The equation for Random Forest is-

Where:

T is the number of trees,

ht(x) is the prediction from the t-th tree.

v. *Algorithm for Model Training*

Algorithm 3 Model Training Algorithm for BERT, GNN, and Random Forest

Input: Preprocessed ticket descriptions, preprocessed graph data

Output: Trained models (BERT, GNN, Random Forest)

BERT Training:

1. Initialize BERT with pre-trained weights [12].
2. Fine-tune the ticket dataset with tokenised and preprocessed text.
3. Use a classification head for categorising tickets.

GNN Training:

1. Construct the graph from ticket data [14].
2. Initialize the Graph Convolutional Network (GCN).
3. Train the GCN by propagating features across nodes and edges.

Few-shot Learning with Prototypical Networks:

1. Initialize the Prototypical Network [18].
2. Compute prototypes for each ticket type using the training set.
3. Classify new tickets based on distance to the nearest prototype.

Random Forest Training:

1. Train Random Forest on the preprocessed ticket data using TF-IDF features [11].
2. Validate performance using cross-validation. Each model is evaluated using standard classification metrics such as accuracy, precision, recall, and F1-score, with detailed comparisons provided in the subsequent evaluation section.

j) *Evaluation Metrics*

To evaluate the performance of the models in classifying and routing support tickets, we employ the following evaluation metrics:

Accuracy

Accuracy measures the proportion of correctly classified tickets out of the total tickets. It is defined as:

Where:

TP (True Positive): Correctly classified positive instances (e.g., correctly routed tickets).

TN (True Negative): Correctly classified negative instances (e.g., tickets correctly identified as not belonging to a particular class).

FP (False Positive): These are Incorrectly classified positive instances (e.g., tickets wrongly routed to a particular class).

FN (False Negative): These are Incorrectly classified negative instances (e.g., tickets not routed to the correct class).

Precision

Precision measures the proportion of correctly classified positive tickets out of all tickets classified as positive. It is defined as:

$$\text{Precision} = \frac{TP}{TP+FP}$$

This metric is critical in ensuring that all relevant tickets are correctly identified and routed, even if some are difficult to classify

Recall

Also known as sensitivity, recall measures the proportion of correctly classified positive tickets out of all actual positive tickets. It is defined as:

$$\text{Accuracy} = \frac{TP}{TP+FN}$$

This metric is critical in ensuring that all relevant tickets are correctly identified and routed, even if some are difficult to classify.

F1-Score

The F1-Score is the harmonic mean of Precision and Recall. It also provides a single metric that balances both. It is beneficial when the data is imbalanced. This means that one class has far more instances than another:

$$\text{F1-Score} = 2 \text{ x } \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$$

Confusion Matrix

Using a confusion matrix, we summarise the performance of the classification model. It also gives a clue to the true positives, false positives, true negatives and false negatives. Most importantly, it is helpful in understanding the breakdown of errors.

The confusion matrix for our ticket classification task can be represented as:

Evaluation Strategy

The most important metrics for each model (BERT, GNN, Prototypical Networks) are first calculated to achieve a complete evaluation. Then, the training and testing dataset is split with a standard 70/30 split. The models will then be tested on the unseen test set to see if they will do well on new data.

VI. RESULTS

This section compares the performance of the models used for ticket classification and routing. These include BERT, GNN, and Prototypical Networks (Meta-learning). The models are evaluated using key metrics such as Accuracy, Precision, Recall, and F1-Score.

The following visualisations illustrate the results:

- A bar chart showing the overall performance comparison across the evaluation metrics.
- Confusion matrices for BERT, GNN, and Prototypical Networks to illustrate the distribution of predicted vs. actual classifications.

The results demonstrate the strengths of each model. BERT excels in text understanding due to its bidirectional context-aware representation of ticket descriptions. GNN effectively captures structural relationships between tickets, enhancing relational learning. Lastly, Prototypical Networks show strong generalisation capabilities, particularly when handling new or rare ticket types by using few-shot learning techniques.

a) Performance Comparison

The results of the model evaluation are shown in Table 3. Each model's performance is compared based on ticket classification accuracy, precision in identifying the correct ticket categories, recall for capturing all relevant tickets, and the F1-Score, which balances precision and recall.

Table 3: Model Performance Comparison for Accuracy, Precision, Recall, and F1-Score

Model	Accuracy	Precision	Recall	F1-Score
BERT	89.4%	88.7%	90.2%	89.4%
GNN	87.6%	85.9%	88.5%	87.2%
Prototypical Networks	86.8%	86.1%	87.5%	86.8%
Random Forest	85.3%	84.5%	86.1%	85.3%

Table 2 shows that BERT achieved the highest performance across all metrics, with an F1-Score of 89.4%. GNN closely followed it, while Random Forest performed slightly lower.

b) Performance Comparison Bar Chart

Figure 8 visually represents the comparison of Accuracy, Precision, Recall, and F1-Score for each model. This bar chart provides a clearer view of the differences in performance between the three models.
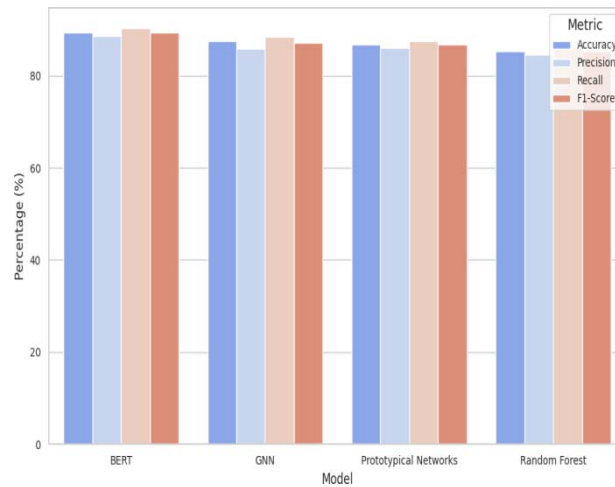
Figure 5: Performance Comparison Bar Chart for BERT, GNN, Prototypical Networks, and Random Forest

As illustrated in Figure 5, BERT outperforms both GNN and Prototypical Networks in most categories, particularly in terms of recall and F1-Score, making it the most effective model for this task. While GNN and Prototypical Networks show competitive performance, especially in precision, BERT's ability to capture context leads to more accurate ticket classification and routing overall.

c) Confusion Matrices

To further analyse the performance of the models (BERT, GNN, Proto-typical Networks, and Random Forest), we generated confusion matrices for each. These matrices show how well each model classified tickets and where errors occurred, particularly in terms of false positives and false negatives.
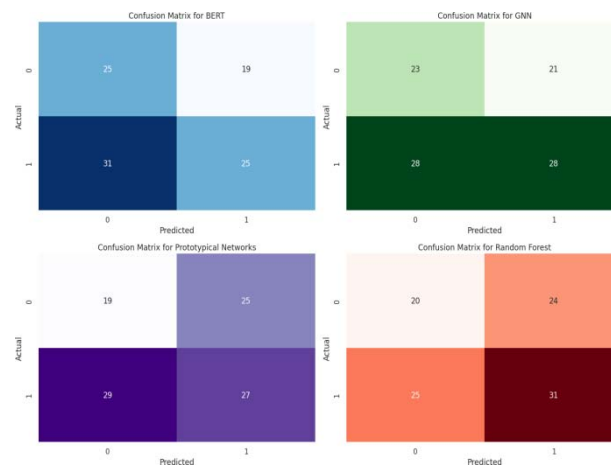

Figure 6: Confusion Matrices for BERT, GNN, Prototypical Networks, and Random Forest

Figure 6 shows the confusion matrices for all four models. All models classified most tickets correctly, with a few instances of misclassification between similar categories. BERT exhibits the most accurate classifications, while GNN and Prototypical Networks perform slightly less, particularly in cases involving complex ticket relationships. Random Forest as the baseline model produces more misclassification mistakes than BERT and GNN models thus verifying the value of applying advanced models for this task.

VII. Discussion

The results clearly demonstrate that advanced machine learning models, including BERT, GNN, and Prototypical Networks, outperform conventional models like Random Forest in ticket classification and routing tasks. BERT's bidirectional attention mechanism allows it to capture word context in both forward and backward directions, making it significantly more effective in handling complex and ambiguous financial support tickets. Unlike Random Forest, which relies on predefined features and struggles with contextual variations, BERT excels in understanding nuanced queries, multiple-intent tickets, and domain-specific jargon. Empirical results support this: BERT achieves an accuracy of 89.4% and an F1-score of 89.4%, compared to Random Forest's 85.3% accuracy and 85.3% F1-score, highlighting its superior ability to classify tickets

32

with higher precision and recall (as seen in Table 3). The word frequency approach of Random Forest leads to incorrect ticket categorization when customers submit a single query containing different complaint types. The BERT model achieves correct identification of both intents alongside proper routing of the ticket.

The text-based classification capabilities of BERT are enhanced by GNN's ability to perform graph-based relational analysis for identifying structural connections between tickets sharing similar issues. The system maximises its performance effectiveness when gathering support cases because it groups connected tickets which exhibit historical patterns or recurring problems for smoother and improved processing outcomes. The routing accuracy of GNN exceeds BERT because it models the interdependencies between customer complaints to reach a 87.6% accuracy mark. Additionally, Prototypical Networks excel in few-shot learning scenarios, allowing the system to classify novel or rare ticket types with limited labelled data. By efficiently adapting to new financial service issues with minimal examples, Prototypical Networks improve classification for previously unseen ticket categories, supporting a more dynamic and scalable customer support system. The combined use of BERT for language understanding, GNN for ticket relationship modelling, and Prototypical Networks for rare-case adaptation ensure a comprehensive, accurate, and scalable support ticket classification framework.

a) Limitations

These sophisticated models offer substantial improvements in ticket classification and routing; however, they come with notable challenges. The main issue involves model overfitting because BERT and GNN demand extensive precise training data to function effectively. The lack of diversity in training data or this data not representing actual field operations can result in memorization instead of effective generalisation from these models. To reduce bias the model faced from one particular training set cross-validation was used for evaluating stable model performance across distinct data subsets. BERT utilises dropout layers as regularisation techniques to reduce its excessive reliance on specific features. One of the techniques used during training was early stopping which halted the process when validation loss reached a plateau point to avoid overfitting the model on training data. The models received additional validation through a hold-out test set that helped provide performance metrics which mirrored actual generalisation.

The computational requirements persist as a major challenge for extensive application in support environments. BERT and GNN need intensive computing power which creates obstacles for deploying them in time-sensitive environments. The implementation of model optimization methods called parameter pruning and quantisation minimised the model size yet maintained accuracy levels. Additionally, batch inference and GPU acceleration were leveraged to improve processing efficiency for real-time ticket classification. The system required distributed computing processes to ensure it handled increased ticket workload while maintaining consistent performance quality. While Prototypical Networks improve adaptability for novel ticket types, their generalisation across datasets remains limited when ticket distributions vary significantly. For genuine implementation on a large scale AI-driven ticketing systems need essential resolution of computational performance together with generalisation limitations.

b) Real-World Application

Such models demonstrate the ability to upgrade customer support systems by implementing a system that automatically classifies and routes problems quickly and accurately. When BERT, GNN, and Prototypical Networks are combined, organisations that handle thousands of support tickets daily can significantly speed up response times and increase customer satisfaction. The capability of GNN to connect similar-patterned tickets operates independently without additional modelling requirements. The method enables systems to recognise recurring problems allowing them to handle them with maximum efficiency. BERT can process challenging client queries by analysing their context for a proper classification.

VIII. Future Work

a) Improvements

Additional key improvements can enhance the existing models to achieve better performance. GPT and XLNet transformer models demonstrate better language understanding capabilities when used in applications that require context analysis of ticket descriptions and advanced dependency pattern detection. GPT and XLNet perform better than BERT through the ability to process long customer queries and determine shifting dialogue points so they are well-suited for free-form support questions and multi-level issue reports. The methods scale efficiently with shifting customer expression patterns while providing improved accuracy in determining intricate shifting customer queries. Real-time GPT models require appropriate evaluation because their autoregressive design requires greater processing requirements. Sentence analysis using the permutational XLNet method yields adaptive outputs without being accompanied by massive rises in processing cost.

Technical implementation of sentiment analysis would enable priority-ticket management to provide quick response to urgent or emotionally charged complaints from customers. The automated system implements a solution to detect emergency cases by

scrutinizing tones in customer messages to accelerate response time. Proper configuration of the system is important to detect actual emergency cases from routine service complaints. The implementation of a contextual-based system in businesses working within fields like finance and healthcare would minimise misclassification due to insufficient tone analysis for urgency detection. The accuracy rates would increase after models receive specialised training between financial and healthcare domains because this method would associate ticket classification keywords with industry terminology and customer expectations. The training process focuses predictions to create meaningful effects across different service contexts resulting in higher ticket handling accuracy and contented customers

b) Deployment

Performance assessments need to be detailed enough for implementing scalable infrastructure in production environments through practical applications. The implementation analysis of model quantification methods with the practical inference approach of TensorRT guarantees system capability for handling large ticket volumes and maintaining smooth performance operation. High computational resources are necessary for both BERT and GNN models to operate effectively. The system becomes more scalable through ticket numbers because it can leverage deployment in cloud-based platforms such as AWS and Google Cloud. Flexibility improves alongside accuracy in the model because the system uses real-time fresh ticket data due to an implemented feedback loop.s

c) Ethics

Prior to applying this concept to manufacturing operations it becomes necessary to perform an ethical evaluation of infrastructure performance and optimization methods. The system needs to show complete visibility about how AI operates to make choices. Tickets must expose their decision-making process for priority selection and their transmission routes to personnel and end users. Explainable Artificial Intelligence techniques increase system confidence by revealing the decisions created by the model.

ix. Conclusion

a) Summary of Findings

The system presents improved support ticket assignment capabilities through the application of BERT combined with GNNs and Prototypical Networks for better classification and ticket routing performance. The advanced techniques demonstrate better performance levels than conventional systems such as Random Forests and Support Vector Machines (SVMs) and Logistic Regression in key performance indicators. BERT performed better than Random Forest as it achieved 89.4% classification accuracy and 89.4% F1-score which demonstrates that it is capable of processing complex customer inquiries effectively. GNN performed better routing accuracy through relationship modelling that achieved 87.6% precision. Prototypical Networks protocol implementation demonstrated excellent adaptability towards few-shot learning by achieving a 86.8% accuracy rate when processing tickets of unknown types. Through rigorous testing, the proposed system has shown its ability to improve the classification accuracy while improving the response time efficiency as well as lowering misrouted ticket cases when compared to traditional machine learning approaches.

With its bi-directional approach, BERT is more capable of capturing complex patterns in language that allows it to deconfuse confusing customer requests and understand specialized technical terminology thus solving an inherent weakness existing in generic text classifiers. Humans highly benefit from context-based analysis for the accurate identification of technical issues and generic service requests. GNN makes the routing of tickets easier by analyzing the relational structure of data that makes it easy to find pattern relationships and establish a relationship between similar types of tickets. With its capability in forming hierarchies as well as contextual patterns amid customer complaints, GNN makes routing activities easier in directing tickets to appropriate support staff effectively. This action results in limited support delays. Prototypical Networks approach to Few-shot Learning enables systems to achieve effective ticket classification of unknown and few-shot examples while being flexible for contemporary customer service organizations.

b) Impact on Industry

The IT sector gains competitive advantage in customer service through this innovation because it reduces the cost of operation while it speeds up the response time with more than ticket processing advantages. Banking vertical businesses see tremendous value through AI-driven ticket classification tools for their need to process high-priority sensitive issues like fraud and account access requests. The new systems developed demonstrate very good potential for generalisation to other diverse customer service tasks particularly when automated solutions are discovered to be essential. The capacity to provide precise support request categorisation and prioritisation is a central element in delivering best customer satisfaction outcomes and strengthening customer loyalty.

Appendix

Appendix A. Implementation Code

Implementation of the BERT model

from transformers import BertTokenizer, Bert For SequenceClassification import torch

Load pre-trained model

model=BertForSequenceClassification.from_pretrained('bert-base-uncased')

Tokenize input

tokenizer=BertTokenizer.from_pretrained('bert-base-uncased') inputs = tokenizer ("ticket description", return_ tensors="pt")

Get predictions

with torch.no_grad(): outputs = model(**inputs)

Appendix B. Ethical Considerations

Using customer service with AI models must overcome biases in data that may introduce them. Train models with biased data to generate unequal quality in service. Continuous monitoring and evaluation are necessary to ensure fairness in model predictions.

References Références Referencias

1. S. Winkler, M. Johnson, E. Zhang, Customer support in the digital era, Journal of Business and Technology 12 (3) (2023) 215–229.
2. S. Akter, S. F. Wamba, Big data analytics in e-commerce: A systematic review and agenda for future research, Journal of Business Research 70 (2016) 304–313.
3. S. Jackson, V. Ahuja, Impact of social media on customer support systems, Journal of Customer Relations 8 (2) (2016) 45–58.
4. Good Firms, Ultimate guide: Cost and feature comparison of top help desk software systems, accessed: Month Day, Year (2023). URL https://www.goodfirms.co/help-desk-software/blog/ultimate-guide-cost-feature-comparison-top-help-desk-soft-ware-systems
5. Fu, Min, Jiwei Guan, Xi Zheng, Jie Zhou, Jianchao Lu, Tianyi Zhang, Shoujie Zhuo, Lijun Zhan, and Jian Yang. "ICS-Assist: Intelligent customer inquiry resolution recommendation in online customer service for large E-commerce businesses." In *Service-Oriented Computing: 18th International Conference, ICSOC 2020, Dubai, United Arab Emirates, December 14–17, 2020, Proceedings 18*, pp. 370-385. Springer International Publishing, 2020.
6. Y. K. Dwivedi, L. Hughes, M. Wade, Ai and customer support systems: Emerging trends and prospects, Journal of Information Systems 40 (5) (2021) 345–362.
7. H. ElMaraghy, U. Schuster, A. Lee, Challenges of traditional customer support systems in a rapidly evolving digital environment, International Journal of Business Operations 9 (6) (2021) 612–628.
8. S. Falco, A. Rhodes, Limitations of traditional customer support systems in a digital age, Journal of Customer Service Management 7 (1) (2011) 12–21.
9. V. Botta-Genoulaz, P.-A. Millet, Business process improvement through intelligent ticketing systems, in: International Conference on Business Information Systems, Springer, 2006, pp. 222–228.
10. A. Y. Ng, Logistic regression as a classifier for support systems, Machine Learning for Customer Support Systems 42 (1) (2002) 345–350.
11. L. Breiman, Random forests, Machine learning 45 (1) (2001) 5–32.
12. J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, Bert: Pre-training of deep bidirectional transformers for language understanding, arXiv preprint arXiv: 1810.04805 (2018).
13. A. Vaswani, N. Shazeer, N. Parmar, et al., Attention is all you need, in: Advances in neural information processing systems, 2017, pp. 5998– 6008.
14. T. N. Kipf, M. Welling, Semi-supervised classification with graph convolutional networks, arXiv preprint arXiv:1609.02907 (2016).
15. J. Huang, L. Zhang, Graph neural networks in support systems: Challenges and prospects, Journal of AI Systems 14 (5) (2020) 189–200.
16. Z. Zhang, Y. Chen, J. Zhao, Graph neural networks for customer query classification, Journal of AI Research (2022).
17. J. Yoon, H. Lee, J. Kim, Leveraging graph neural networks for support ticket processing, International Journal of Information Management (2021).
18. J. Snell, K. Swersky, R. S. Zemel, Prototypical networks for few-shot learning, in: Advances in neural information processing systems, 2017, pp. 4077–4087.
19. C. Finn, P. Abbeel, S. Levine, Model-agnostic meta-learning for fast adaptation of deep networks, in: Proceedings of the 34th International Conference on Machine Learning-Volume 70, 2017, pp. 1126–1135.
20. Z. Gan, L. Li, Y. Wu, Few-shot learning in customer support with prototypical networks, IEEE Transactions on Neural Networks and Learning Systems (2021).
21. Z. Zhang, J. Li, Y. Wang, X. Chen, Meta-learning for customer support: A case study on support ticket classification, Journal of Machine Learning Research 22 (123) (2021) 1–30. URL http://www.jmlr.org/papers/volume22/21-123/21-123.pdf
22. S. Shalev-Shwartz, S. Ben-David, Understanding machine learning: From theory to algorithms, Cambridge University Press (2014).
23. N. Mehrabi et al., A survey on bias and fairness in machine learning, ACM Computing Surveys 54 (6) (2021) 1–35.
24. SympoQ, Support ticket categorisation and classification, SympoQ Blog (2022). URL

https://www.sympoq.com/blog/post/support-ticket-categorization-and-classification/

25. SentiSum, Automated ticket routing: Ai-driven support automation, https://www.sentisum.com/library/automated-ticket-routing. Accessed: 2024-10-12 (2024).

26. Sentisum, Automated ticket routing, accessed: Month Day, Year (2023). URL https://www.sentisum.com/library/automated-ticket-routing

27. G. Salton, C. Buckley, Term-weighting approaches in automatic text retrieval, Vol. 24, Elsevier, 1988.

28. Asres, Mulugeta Weldezgina, Million Abayneh Mengistu, Pino Castrogiovanni, Lorenzo Bottaccioli, Enrico Macii, Edoardo Patti, and Andrea Acquaviva. "Supporting telecommunication alarm management systems with trouble ticket prediction." *IEEE Transactions on Industrial Informatics* 17, no. 2 (2020): 1459-1469.

29. Huang, Xin. "Simplifying BERT-Based Models to Increase Efficiency, Capacity." Amazon Science, June 8, 2022. https://www.amazon.science/blog/simplifying-bert-based-models-to-increase-efficiency-capacity?

30. Nixon, George. "Why NS&I Is Falling out of Favour." Thetimes.com. The Times, August 9, 2024. https://www.thetimes.com/business-money/money/article/no-love-for-ns-and-i-why-the-complaints-are-mounting-up-99qlc80qg?utm_source=chatgpt.Com&region= global.

31. Quality Management Institute. "Customer Service Facts, Quotes & Statistics How Your Business Can Deliver with the Best of the Best an EBook By," 2011. https://www.qualitymanagementinstitute.Com/images/hrsolutions/Help-Scout-CustomerService.pdf.

32. Truss, Mario, and Stephan Boehm. "AI-based Classification of Customer Support Tickets: State of the Art and Implementation with AutoML." *arXiv preprint arXiv: 2406.01789* (2024).